

| Committee Draft ISO/IEC 3rd CD 27040 | | Reference number: ISO/IEC JTC 1/SC 27 N12681 | |
|---|--|---|--|
| Date: 2013-06-11 | | Supersedes document SC 27 N11987 | |
| THIS DOCUMENT IS STILL UNDER STUDY AND SUBJECT TO CHANGE. IT SHOULD NOT BE USED FOR REFERENCE PURPOSES. | | | |
| ISO/IEC JTC 1/SC 27 Information technology - Security techniques Secretariat: Germany (DIN) | Circulated to P- and O-members, and to technical committees and organizations in liaison for comments by: 2013-09-12 Please submit your votes/comments via the online balloting application by the due date indicated. | | |
| ISO/IEC 3rd CD 27040 Title: Information technology -- Security techniques – Storage security Project: 1.27.90 (27040) | | | |
| Explanatory Report | | | |
| Status | SC 27 Decision | Reference documents | |
| | | Input | Output |
| For details regarding previous development stages please see 2 nd page. | | | |
| NP 27040 1st WD 27040 | 10 th WG 4 meeting, April 2011 resolutions 1, 6 (N9942). | SoV (N9670); SoContr (N9680); TC 46/SC 11 com. (N9686). | Liaison to: TC 46/SC 11 (N9982); DoC (N9977); Text 1 st WD (N9978). |
| 2nd WD 27040 | 11 th WG 4 meeting, Oct. 2011, resolutions 1, 9 (N10152). | SoCom. (N10137). US supplement comm. (N10688). | WG 4 report (N10167); DoC (N10180); Text f. 2 nd WD (N10181). |
| 1st CD 27040 | 12 th WG 4 meeting, May 2012, resolutions 1, 12, (N10993) and 24 th SC 27 Plenary, May 2012, resolution 4, 12, (Deleg. Of Authority f. DIS), 20 (N11330). | SoCom. (N10881); US supplement comm. (N11365). | Liaison to: JTC 1/SC 25/WG 4 (N11037); WG 4 report (N11000); DoC (N11023); Text f. 1 st CD (N11024). |
| 2nd CD 27040 | 13 th WG 4 meeting, Oct. 2012, resolution 12 (N11941) | SoV (N11532); Draft DoC (N11666); Draft revised text (N11667) | WG 4 report (N11943); DoC (N11986); Text f. 2 nd CD (N11987). |
| 3rd CD 27040 | 14 th WG 4 meeting, April 2013, resolutions 18, 26 (N12740); 25 th SC 27 Plenary, April 2013, Resolution 14 Delegation of Auth. f. DIS (N12739). | SoV (N12238); Draft DoC (N12338); Draft revised text (N12339) US supplement comm. (N12756).. | Liaisons to CSA (N12508); EuroCloud (N12639); ITU-T JCA-Cloud (N12519); ITU-T SG 17 (N12518) WG 4 report (N126334) DoC (N12680); Text f. 3 rd CD (N12681). |
| 3rd CD Consideration In accordance with resolution 26 (see SC 27 N12740) of the 14 th SC 27/WG 4 Plenary meeting in Sophia Antipolis, 22 nd – 26 th April 2013 the attached document is hereby circulated for a 3rd CD letter ballot closing by <div style="text-align: center; font-size: 1.5em; font-weight: bold; margin: 10px 0;">2013-09-12</div> Medium: http://isotc.iso.org/livelink/livelink/open/jtc1sc27 No. of pages: 2 + 116 | | | |

| Explanatory Report (2nd page) | | | |
|--|--|---------------------|---|
| Status | SC 27 Decision | Reference documents | |
| | | Input | Output |
| Study Period Storage security | 7 th WG 4 meeting, Nov. 2009, resolution 8 (N6017). | CA contr. (N6268) | Call f. contr. (N7924). |
| | 8 th WG 4 meeting, April 2010, resolution 9 (N8617) | SoContr. (N8548) | Call f. contr. (N8656). |
| NWIP | 9 th WG 4 meeting, Oct. 2010 resolutions 1, 2, 4, 11, 12, 14, 15 (N084).; 22 nd SC 27 Plenary, Oct. 2010, resolution 21 (N9460). | | Report (N9112); Prelim. draft (N9117); Call f. contr/co-editor (N9445); NWIP (N9444). |

ISO/IEC JTC 1/SC 27 **N12681**

Date: 2013-06-05

ISO/IEC CD 27040.3

ISO/IEC JTC 1/SC 27/WG 4

Secretariat: DIN

Information technology — Security techniques — Storage security

Warning

This document is not an ISO International Standard. It is distributed for review and comment. It is subject to change without notice and may not be referred to as an International Standard.

Recipients of this draft are invited to submit, with their comments, notification of any relevant patent rights of which they are aware and to provide supporting documentation.

Copyright notice

This ISO document is a working draft or committee draft and is copyright-protected by ISO. While the reproduction of working drafts or committee drafts in any form for use by participants in the ISO standards development process is permitted without prior permission from ISO, neither this document nor any extract from it may be reproduced, stored or transmitted in any form for any other purpose without prior written permission from ISO.

Requests for permission to reproduce this document for the purpose of selling it should be addressed as shown below or to ISO's member body in the country of the requester:

Secretariat of ISO/IEC JTC 1/SC 27
DIN German Institute for Standardization
DE-10787 Berlin

Tel. + 49 30 2601 2652
Fax + 49 30 2601 1723
E-mail krystyna.passia@din.de
Web <http://www.jtc1sc27.din.de/en> (public web site)
<http://isotc.iso.org/isotcportal/index.html> (SC 27 documents)

Reproduction for sales purposes may be subject to royalty payments or a licensing agreement.

Violators may be prosecuted.

| | | Page |
|----|--|-----------|
| 1 | Contents | |
| 2 | Foreword | v |
| 3 | Introduction | vi |
| 4 | 1 Scope | 1 |
| 5 | 2 Normative references | 1 |
| 6 | 3 Terms and definitions | 2 |
| 7 | 4 Symbols and abbreviated terms | 7 |
| 8 | 5 Overview and concepts | 12 |
| 9 | 5.1 General | 12 |
| 10 | 5.2 Storage concepts | 12 |
| 11 | 5.3 Introduction to storage security | 13 |
| 12 | 5.4 Storage security risks | 15 |
| 13 | 5.4.1 Background | 15 |
| 14 | 5.4.2 Data breaches | 15 |
| 15 | 5.4.3 Data corruption or destruction | 17 |
| 16 | 5.4.4 Temporary or permanent loss of access/availability | 17 |
| 17 | 5.4.5 Failure to meet statutory, regulatory, or legal requirements | 17 |
| 18 | 6 Supporting controls | 18 |
| 19 | 6.1 General | 18 |
| 20 | 6.2 Direct Attached Storage (DAS) | 18 |
| 21 | 6.3 Storage networking | 19 |
| 22 | 6.3.1 Background | 19 |
| 23 | 6.3.2 Storage Area Networks (SAN) | 19 |
| 24 | 6.3.3 Network Attached Storage (NAS) | 23 |
| 25 | 6.4 Storage management | 25 |
| 26 | 6.4.1 Background | 25 |
| 27 | 6.4.2 Authentication and authorization | 25 |
| 28 | 6.4.3 Secure the management interfaces | 26 |
| 29 | 6.4.4 Security auditing, accounting, and monitoring | 28 |
| 30 | 6.4.5 System hardening | 30 |
| 31 | 6.5 Block-based storage | 31 |
| 32 | 6.5.1 Fibre Channel (FC) storage | 31 |
| 33 | 6.5.2 IP storage | 31 |
| 34 | 6.6 File-based storage | 31 |
| 35 | 6.6.1 NFS-based NAS | 31 |
| 36 | 6.6.2 SMB/CIFS-based NAS | 32 |
| 37 | 6.6.3 Parallel NFS-based NAS | 33 |
| 38 | 6.7 Object-based storage | 34 |
| 39 | 6.7.1 Cloud storage | 34 |
| 40 | 6.7.2 Object-based Storage Device (OSD) | 34 |
| 41 | 6.7.3 Content Addressable Storage (CAS) | 35 |
| 42 | 6.8 Storage security services | 36 |
| 43 | 6.8.1 Data sanitization | 36 |
| 44 | 6.8.2 Data confidentiality | 39 |
| 45 | 6.8.3 Data reductions | 41 |

| | | | |
|----|--------------------|--|-----|
| 1 | 7 | Guidelines for the design and implementation of storage security | 42 |
| 2 | 7.1 | General..... | 42 |
| 3 | 7.2 | Storage security design principles | 42 |
| 4 | 7.2.1 | Defence in depth..... | 42 |
| 5 | 7.2.2 | Security domains | 43 |
| 6 | 7.2.3 | Design resilience | 44 |
| 7 | 7.2.4 | Secure initialization | 44 |
| 8 | 7.3 | Data reliability, availability, and resilience..... | 44 |
| 9 | 7.3.1 | Reliability | 44 |
| 10 | 7.3.2 | Availability | 45 |
| 11 | 7.3.3 | Backups and replication | 46 |
| 12 | 7.3.4 | Disaster recovery and business continuity | 46 |
| 13 | 7.3.5 | Resilience | 47 |
| 14 | 7.4 | Data retention..... | 47 |
| 15 | 7.4.1 | Long-term retention..... | 47 |
| 16 | 7.4.2 | Short to medium-term retention..... | 48 |
| 17 | 7.5 | Data confidentiality and integrity | 49 |
| 18 | 7.6 | Virtualization | 52 |
| 19 | 7.6.1 | Storage virtualization | 52 |
| 20 | 7.6.2 | Storage for virtualized systems | 52 |
| 21 | 7.7 | Design and implementation considerations | 53 |
| 22 | 7.7.1 | Encryption and key management issues | 53 |
| 23 | 7.7.2 | Align storage and policy | 54 |
| 24 | 7.7.3 | Compliance..... | 55 |
| 25 | 7.7.4 | Secure multi-tenancy | 56 |
| 26 | 7.7.5 | Secure autonomous data movement..... | 57 |
| 27 | Annex A | (normative) Media sanitization | 59 |
| 28 | A.1 | Methods used to sanitize media..... | 59 |
| 29 | A.2 | Sanitization for different types of media | 59 |
| 30 | A.3 | Cryptographic erase device guidelines..... | 73 |
| 31 | Annex B | (informative) Selecting appropriate storage security controls | 77 |
| 32 | B.1 | Criteria for selecting controls..... | 77 |
| 33 | B.1.1 | Overview | 77 |
| 34 | B.1.2 | Data sensitivity classes | 77 |
| 35 | B.1.3 | Security priority codes | 79 |
| 36 | B.2 | Summary of storage security controls | 79 |
| 37 | B.2.1 | Supporting controls for storage security..... | 79 |
| 38 | B.2.2 | Storage security design and implementation guidance | 88 |
| 39 | Annex C | (informative) Important security concepts | 97 |
| 40 | C.1 | Authentication..... | 97 |
| 41 | C.2 | Authorization and access control..... | 98 |
| 42 | C.3 | Self-Encrypting Drives (SED) | 100 |
| 43 | C.4 | Sanitization..... | 100 |
| 44 | C.5 | Logging..... | 101 |
| 45 | C.6 | N_Port ID Virtualization (NPIV) | 102 |
| 46 | Bibliography | | 103 |
| 47 | Index..... | | 107 |
| 48 | | | |

1 Foreword

2 ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission)
3 form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC
4 participate in the development of International Standards through technical committees established by the
5 respective organization to deal with particular fields of technical activity. ISO and IEC technical committees
6 collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in
7 liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have
8 established a joint technical committee, ISO/IEC JTC 1.

9 International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

10 The main task of the joint technical committee is to prepare International Standards. Draft International Standards
11 adopted by the joint technical committee are circulated to national bodies for voting. Publication as an
12 International Standard requires approval by at least 75 % of the national bodies casting a vote.

13 Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights.
14 ISO and IEC shall not be held responsible for identifying any or all such patent rights.

15 ISO/IEC 27040 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*,
16 Subcommittee SC 27, *Security techniques*.

1 Introduction

Many organizations face the challenge of implementing data protection and security measures to meet a wide range of requirements, including statutory and regulatory compliance. Too often the security associated with storage systems and infrastructure has been missed because of misconceptions and limited familiarity with the storage technology, or in the case of storage managers and administrators, a limited understanding of the inherent risks or basic security concepts. The net result of this situation is that digital assets are needlessly placed at risk of compromise due to data breaches, intentional corruption, being held hostage, or other malicious events.

Data storage has matured in an environment where security has been a secondary concern due to its historical reliance on isolated connectivity, specialized technologies, and the physical security of data centres. Even as storage connectivity evolved to use technologies such as storage protocols over TCP/IP, few users took advantage of either the inherent security mechanisms or the recommended security measures.

This International Standard provides guidelines for Storage Security in an organization, supporting in particular the requirements of an Information Security Management System (ISMS) according to ISO/IEC 27001. This standard recommends the information security risk management approach as defined in ISO/IEC 27005. It is up to the organization to define their approach to risk management, depending for example on the scope of the ISMS, context of risk management, or industry sector. A number of existing methodologies can be used under the framework described in this International Standard to implement the requirements of an ISMS.

This International Standard is relevant to managers and staff concerned with information security risk management within an organization and, where appropriate, external parties supporting such activities.

The objectives for this International Standard are to:

- help draw attention to the risks,
- assist organizations in better securing their data when stored,
- provide a basis for auditing, designing and reviewing storage security controls.

It is emphasized that ISO/IEC 27040 provides further detailed implementation guidance on the storage security controls that are described at a basic standardized level in ISO/IEC 27002.

It should be noted that this International Standard is not a reference or normative document for regulatory and legislative security requirements. Although it emphasizes the importance of these influences, it cannot state them specifically, since they are dependent on the country, the type of business, etc.

Information technology — Security techniques — Storage security

1 Scope

This International Standard provides detailed technical guidance on how organizations may define an appropriate level of risk mitigation by employing a well-proven and consistent approach to the planning, design, documentation and implementation of data storage security. Storage security applies to the protection (security) of information where it is stored and to the security of the information being transferred across the communication links associated with storage. Storage security includes the security of devices and media, the security of management activities related to the devices and media, the security of applications and services, and security relevant to end-users during the lifetime of devices and media and after end of use.

Storage security is relevant to anyone involved in owning, operating or using data storage devices, media and networks. This includes senior managers, acquirers of storage product and service, and other non-technical managers or users, in addition to managers and administrators who have specific responsibilities for information security and/or storage security, storage operation, or who are responsible for an organization's overall security program and security policy development. It is also relevant to anyone involved in the planning, design and implementation of the architectural aspects of storage network security.

This standard provides an overview of storage security concepts and related definitions. It includes guidance on the threat, design and control aspects associated with typical storage scenarios and storage technology areas. In addition, it provides references to other international standards and technical reports that address existing practices and techniques that can be applied to storage security.

2 Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 17788 (to be published), *Information technology — Distributed application platforms and services — Cloud computing — Vocabulary*

ISO/IEC 27000:2012, *Information technology — Security techniques — Information security management systems — Overview and vocabulary*

ISO/IEC 27001:2005, *Information technology — Security techniques — Information security management systems — Requirements*

ISO/IEC 27002:2005, *Information technology — Security techniques — Code of practice for information security management*

ISO/IEC 27005:2008, *Information technology — Security techniques — Information security risk management*

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/IEC 27000, ISO/IEC 27005 and the following apply.

3.1

block

unit in which data is *stored* (3.43) and retrieved on disk and tape *devices*(3.13)

3.2

clear

method of sanitization in which software or hardware products are used to overwrite storage space on the media with non-sensitive data

3.3

compression

process of removing redundancies in digital data to reduce the amount that should be *stored* (3.43) or transmitted

[SOURCE: ISO/TR 12033:2009, 3.1.]

Note 1 to entry: For *storage* (3.38), lossless compression (i.e., compression using a technique that preserves the entire content of the original data, and from which the original data can be reconstructed exactly) is required.

3.4

cryptographic erase

method of sanitization in which the Media Encryption Key (MEK) for the encrypted target data is *sanitized* (3.33), making recovery of the decrypted target data infeasible

3.5

cryptoperiod

defined period of time during which a specific cryptographic key is authorized for use, or during which time the cryptographic keys in a given system can remain in effect

[SOURCE: ISO 16609:2004, 3.9.]

3.6

data at rest

data *stored* (3.43) on stable, persistent media

3.7

data breach

compromise of security that leads to the accidental or unlawful *destruction* (3.12), loss, alteration, unauthorized disclosure of, or access to protected data transmitted, *stored* (3.43) or otherwise processed

3.8

data in motion

data being transferred from one location to another

3.9

data integrity

property that data has not been altered or destroyed in an unauthorized manner

[SOURCE: ISO 7498-2:1989, 3.3.21.]

3.10**deduplication**

method of reducing *storage* (3.38) needs by eliminating redundant data, which is replaced with a pointer to the unique data copy

Note 1 to entry: Deduplication is sometimes considered a form of *compression* (3.3).

3.11**degauss**

render data unreadable by applying a strong magnetic field to the media

3.12**destruction**

result of actions taken to ensure that media cannot be reused as originally intended and that information is virtually impossible or prohibitively expensive to recover

3.13**device**

mechanical, electrical, or electronic contrivance with a specific purpose

[SOURCE: ISO/IEC 14776-372:2011, 3.1.10.]

3.14**disintegration**

act of physical *destruction* (3.12) that reduces an item to components, fragments, or particles

3.15**Electronically Stored Information****ESI**

data or information of any kind and from any source, whose temporal existence is evidenced by being *stored* (3.43) in or on any electronic medium

Note 1 to entry: ESI includes traditional e-mail, memos, letters, spreadsheets, databases, office documents, presentations and other electronic formats commonly found on a computer. ESI also includes system, application and file-associated *metadata* (3.24) such as timestamps, revision history, file type, etc.

Note 2 to entry: Electronic medium can take the form of, but is not limited to, *storage devices* (3.33) and *storage elements* (3.34).

3.16**Fibre Channel****FC**

serial I/O interconnect capable of supporting multiple protocols, including access to open system *storage* (3.38), access to mainframe *storage* (3.38), and networking

Note 1 to entry: Fibre Channel supports point to point, arbitrated loop, and switched topologies with a variety of copper and optical links running at speeds from 1 gigabit per second to over 10 gigabits per second.

3.17**Fibre Channel Protocol****FCP**

serial SCSI transport protocol used on *Fibre Channel* (3.16) interconnects

3.18**gateway**

device (3.13) that converts a protocol to another protocol

3.19

host

computer system that provides end users with services such as computation and *storage* (3.38) access

3.20

in-band

communications or transmissions that occur within a previously established communication method or channel

Note 1 to entry: The communications or transmissions often take the form of a separate protocol, such as a management protocol over the same medium as the primary data protocol.

3.21

malware

malicious software designed specifically to damage or disrupt a system, attacking confidentiality, integrity and/or availability

Note 1 to entry: Viruses and Trojan horses are examples of malware.

[SOURCE: ISO/IEC 27033-1:2009, 3.22.]

3.22

Mean Time Between Failures

MTBF

the expected time between consecutive failures in a system or component

[SOURCE: ISO/IEC/IEEE 24765:2010, 3.1713, modified — The term was capitalized.]

3.23

Mean Time To Repair

MTTR

expected or observed duration to return a malfunctioning system or component to normal operations

[SOURCE: ISO/IEC/IEEE 24765:2010, 3.1714, modified — The term was capitalized.]

3.24

metadata

data that defines and describes other data

[SOURCE: ISO/IEC 11179-1:2004, 3.2.16.]

3.25

multi-factor authentication

authentication using two or more of the following factors:

- knowledge factor, “something an individual knows”;
- possession factor, “something an individual has”;
- biometric factor, “something an individual is or is able to do”.

[SOURCE: ISO 19092:2008, 4.42.]

3.26**multi-tenancy**

allocation of physical and virtual resources such that multiple tenants and their computations and data are isolated from and inaccessible to one another

[SOURCE: ISO/IEC 17788, 3.2.25.]

3.27**Network Attached Storage****NAS**

storage devices (3.40) or systems that connect to a network and provide file access services to computer systems

3.28**non-volatile storage**

storage (3.38) that retains its contents even after power is removed

3.29**out-of-band**

communications or transmissions that occur outside of a previously established communication method or channel

3.30**point of encryption**

location within the ICT infrastructure where data is encrypted on its way to *storage* (3.38) and conversely where data is decrypted when accessed from *storage* (3.38)

3.31**purge**

render *sanitized* (3.33) data unrecoverable by laboratory attack methods

3.32**reliability**

the ability of a system or component to perform its required functions under stated conditions for a specified period of time

[SOURCE: ISO/IEC/IEEE 24765:2010, 3.2467, modified — The second definition from ISO/IEC 9126-1:2001 and the cf. entry were not included.]

3.33**sanitize**

process to remove information from media such that data recovery is not possible at a given level of effort

3.34**secure multi-tenancy**

type of *multi-tenancy* (3.26) that employs security controls to explicitly guard against *data breaches* (3.7) and provides validation of these controls for proper governance

Note 1 to entry: Secure multi-tenancy exists when the risk profile of an individual tenant is no greater than it would be in a dedicated, single-tenant environment.

Note 2 to entry: In very secure environments even the identity of the tenants is kept secret.

3.35**security strength**

number associated with the amount of work that is required to break a cryptographic algorithm or system

3.36

shred

act of physical *destruction* (3.12) that reduces an item to regular size fragments or particles

3.37

single point of failure

element or component of a system, a path in a system, or a system that if it fails the whole system or an array of systems are unable to perform their primary functions

Note 1 to entry: A single point of failure is often considered a design flaw associated with a critical element.

3.38

storage

device (3.13) into which data may be entered, and from which data may be retrieved

3.39

Storage Area Network

SAN

network whose primary purpose is the transfer of data between computer systems and *storage devices* (3.40) and among *storage devices* (3.40)

Note 1 to entry: A SAN consists of a communication infrastructure, which provides physical connections, and a management layer, which organizes the connections, *storage elements* (3.41), and computer systems so that data transfer is secure and robust.

3.40

storage device

any *storage element* (3.41) or aggregation of elements, designed and built primarily for the purpose of persistent data *storage* (3.38) and delivery

3.41

storage element

device (3.13) that is used to build *storage devices* (3.40) and which contributes to persistent data *storage* (3.38) and delivery

Note 1 to entry: Common examples of a storage element include a disk or tape drive.

3.42

storage medium (storage media)

material on which *Electronically Stored Information* (3.15) or digital data are or can be persistently *stored* (3.43)

3.43

stored

process that results in data being recorded on *volatile storage* (3.45) or *non-volatile storage* (3.28)

3.44

strong authentication

authentication by means of cryptographically derived credentials

[SOURCE: ISO/TS 22600-1:2006, 2.23.]

3.45

volatile storage

storage (3.37) that fails to retain its contents after power is removed

1 **3.46**2 **weak key**

3 key that interacts with some aspect of a particular cipher's definition in such a way that it weakens the security
 4 strength of the cipher

5 **4 Symbols and abbreviated terms**

| | | |
|----|---------|--|
| 6 | ACE | Access Control Entry |
| 7 | ACL | Access Control List |
| 8 | AES | Advanced Encryption Standard |
| 9 | ATA | Advanced Technology Attachment |
| 10 | BC | Business Continuity |
| 11 | BCM | Business Continuity Management |
| 12 | CAS | Content Addressable Storage |
| 13 | CCM | Counter with Cipher block chaining Message authentication code |
| 14 | CDMI | Cloud Data Management Interface |
| 15 | CDP | Continuous Data Protection |
| 16 | CHAP | Challenge Handshake Authentication Protocol |
| 17 | CIFS | Common Internet File System |
| 18 | CLI | Command Line Interface |
| 19 | DAC | Discretionary Access Control |
| 20 | DAS | Direct Attached Storage |
| 21 | DDoS | Distributed Denial of Service |
| 22 | DH-CHAP | Diffie Hellman – Challenge Handshake Authentication Protocol |
| 23 | DLM | Data Lifecycle Management |
| 24 | DMZ | De-Militarized Zone |
| 25 | DNS | Domain Name System |
| 26 | DoS | Denial of Service |
| 27 | DR | Disaster Recovery |
| 28 | DRP | Disaster Recovery Planning |

| | | |
|----|-------|---|
| 1 | EHR | Electronic Healthcare Record |
| 2 | ESI | Electronically Stored Information |
| 3 | ESP | Encapsulating Security Payload |
| 4 | FC | Fibre Channel |
| 5 | FC-SP | Fibre Channel – Security Protocol |
| 6 | FCIP | Fibre Channel over TCP/IP |
| 7 | FCoE | Fibre Channel over Ethernet |
| 8 | FCP | Fibre Channel Protocol |
| 9 | FCS | Fixed Content Storage |
| 10 | FDE | Full Disk Encryption |
| 11 | GCM | Galois/Counter Mode |
| 12 | GUI | Graphical User Interface |
| 13 | HAMR | Heat Assisted Magnetic Recording |
| 14 | HBA | Host Bus Adapter |
| 15 | HDD | Hard Disk Drive |
| 16 | HTTPS | Hypertext Transfer Protocol Secure |
| 17 | ICT | Information and Communications Technology |
| 18 | IDS | Intrusion Detection System |
| 19 | IEEE | Institute of Electrical and Electronics Engineers |
| 20 | IETF | Internet Engineering Task Force |
| 21 | IKE | Internet Key Exchange |
| 22 | ILM | Information Lifecycle Management |
| 23 | I/O | Input/Output |
| 24 | IP | Internet Protocol |
| 25 | IPS | Intrusion Prevention System |
| 26 | iFCP | Internet Fibre Channel Protocol |
| 27 | IPOCM | Incident Preparedness and Operational Continuity Management |

| | | |
|----|-------|--|
| 1 | IPsec | Internet Protocol Security |
| 2 | IRBC | ICT Readiness for Business Continuity |
| 3 | iSCSI | Internet Small Computer Systems Interface |
| 4 | ISMS | Information Security Management System |
| 5 | iSNS | Internet Storage Name Service |
| 6 | KMIP | Key Management Interoperability Protocol |
| 7 | LAN | Local Area Network |
| 8 | LBA | Logical Block Address |
| 9 | LDAP | Lightweight Directory Access Protocol |
| 10 | LUN | Logical Unit Number |
| 11 | MAC | Mandatory Access Control |
| 12 | MD5 | Message-Digest algorithm 5 |
| 13 | MEK | Media Encryption Key |
| 14 | MTBF | Mean Time Between Failure |
| 15 | MTTF | Mean Time To Failure |
| 16 | MTTR | Mean Time To Repair |
| 17 | NAS | Network Attached Storage |
| 18 | NAT | Network Address Translation |
| 19 | NFS | Network File System |
| 20 | NIC | Network Interface Card |
| 21 | NIS | Network Information Service |
| 22 | NPIV | N_Port ID Virtualization |
| 23 | NTLM | NT LAN Manager |
| 24 | NTP | Network Time Protocol |
| 25 | NVM | Non-Volatile Memory |
| 26 | OASIS | Organization for the Advancement of Structured Information Standards |
| 27 | OID | Object Identifier |

| | | |
|----|--------|---|
| 1 | OS | Operating System |
| 2 | OSD | Object-based Storage Device |
| 3 | PCIe | Peripheral Component Interconnect Express |
| 4 | PII | Personally Identifiable Information |
| 5 | PKI | Public Key Infrastructure |
| 6 | pNFS | Parallel Network File System |
| 7 | RADIUS | Remote Authentication Dial In User Service |
| 8 | RAID | Redundant Array of Independent Disks |
| 9 | RBAC | Role-based Access Control |
| 10 | REST | Representational State Transfer |
| 11 | RNG | Random Number Generator |
| 12 | RPC | Remote Procedure Call |
| 13 | SAN | Storage Area Network |
| 14 | SAS | Serial Attached SCSI |
| 15 | SCSI | Small Computer System Interface |
| 16 | SED | Self-Encrypting Drive |
| 17 | SHA | Secure Hash Algorithm |
| 18 | SIEM | Security Information and Event Management |
| 19 | SLP | Service Locator Protocol |
| 20 | SMB | Server Message Block |
| 21 | SMI-S | Storage Management Initiative – Specification |
| 22 | SNIA | Storage Networking Industry Association |
| 23 | SNMP | Simple Network Management Protocol |
| 24 | SOHO | Small Office/Home Office |
| 25 | SSC | Security Subsystem Class |
| 26 | SSD | Solid State Drive |
| 27 | SSH | Secure Shell |

| | | |
|----|------|--|
| 1 | SSHD | Solid State Hard Drive |
| 2 | TCG | Trusted Computing Group |
| 3 | TCP | Transmission Control Protocol |
| 4 | TLS | Transport Layer Security |
| 5 | USB | Universal Serial Bus |
| 6 | VLAN | Virtual Local Area Network |
| 7 | VM | Virtual Machine |
| 8 | VSAN | Virtual Storage Area Network |
| 9 | VPN | Virtual Private Network |
| 10 | WAN | Wide Area Network |
| 11 | WORM | Write Once Read Many |
| 12 | WWN | World Wide Name |
| 13 | XEX | Xor-Encrypt-Xor |
| 14 | XTS | XEX-based Tweaked-codebook mode with ciphertext Stealing |
| 15 | | |

5 Overview and concepts

5.1 General

Computer data storage or information storage, often called storage, refers to computer components, storage elements, storage devices, and storage media that retain Electronically Stored Information (ESI) or digital data in durable form (i.e., non-volatile). Storage is a core function and fundamental component of computers.

To secure storage infrastructure, a clear understanding of the storage technologies and concepts are necessary. In addition, the types of security controls and insight into how they impact and interact with the storage technologies are also important. Finally, the threats to this infrastructure and the major risks arising from these threats are factored into any efforts to secure storage infrastructure or individual storage systems.

5.2 Storage concepts

In the past, storage was simply seen as Hard Disk Drives (HDD) and tape drives attached to a computer to store data. This approach, commonly called Direct Attached Storage (DAS), is still in use within enterprise data centres as well as Small Office/Home Office (SOHO) environments. Alternate approaches based on the use of networking technology for storage have emerged as highly sophisticated technologies became available to provide solutions for managing, connecting, protecting, securing, sharing, and optimizing the storage of data. These solutions become more feasible and cost effective as storage technology evolved from non-intelligent internal and external DAS to intelligent networked storage. The use of networking in these solutions increases the attack surface of these solutions and requires additional attention be paid to their security.

Contemporary storage solutions include some or all of the following elements:

- Storage arrays with storage network interfaces
- Network Attached Storage (NAS)
- Content Addressable Storage (CAS)
- Object-based Storage Devices (OSD)
- Backup/recovery systems, Continuous Data Protection (CDP), etc. (i.e., data protection systems)

Storage has become a prominent and independent layer of ICT infrastructure in enterprise class and midrange computing environments. The requirements for these environments frequently exceed simple data storage capabilities. Examples of applications and functions driving the emergence of new storage technology include:

- Sharing of vast storage resources (measured in petabytes and exabytes) between multiple systems via networks
- Backups that don't require use of a Local Area Network (LAN)
- Remote, disaster tolerant, on-line mirroring of mission critical data
- Clustering of fault tolerant applications and related systems around a single copy of data
- Long-term retention of sensitive and/or high-value business information
- Distributed database and file systems

1 — Support for regulatory and legal compliance requirements

2 — Support for centralized backup and archiving

3 **5.3 Introduction to storage security**

4 Storage security is concerned with the physical, technical and administrative controls as well as the preventive,
5 detective and corrective controls associated with storage systems and infrastructure as outlined in 5.2. Storage
6 security can also force the introduction of specialized technologies like:

7 — Media sanitization

8 — Virtualization security

9 — Self-encrypting storage devices like HDD, Solid State Drives (SSD), and Solid State Hard Drive (SSHD)

10 — Key management services

11 — Data authenticity and integrity services

12 — Data in motion protections (encryption and data reduction)

13 — Directories and other user management systems

14 To better understand the security issues and implications for storage, one should know both how and why the
15 storage technologies are used. As a starting point, consider the following:

16 — The storage systems can function as nodes within storage networks, which can be based on, but not limited
17 to technologies like TCP/IP, Fibre Channel, Fibre Channel over Ethernet (FCoE), and InfiniBand. The
18 potential threats can vary significantly based on the networking technology as well as the topologies used.

19 — When stored, the data is typically represented and accessed as either block data or as files/objects; there are
20 significant differences between these two types of storage methods. Likewise, the security measures
21 associated with each can have radical differences, especially with access controls, encryption, and data
22 integrity.

23 — Storage management is both an element of the storage infrastructure as well as an operation performed on
24 many of the systems. It is common to have privileged users applying configuration changes, provisioning
25 storage, tuning, monitoring, etc. this infrastructure. Some of the management can be performed remotely as
26 well as involve third parties such as vendor support personnel.

27 — Data availability and integrity are key factors in an organization's storage architecture, so it is important for
28 security to be complementary rather than a trade-off and that it not negate high-availability measures by
29 introducing choke-points and additional single points of failure.

30 — Many organizations implement elaborate data resiliency strategies, which are integral to their Disaster
31 Recovery (DR) and Business Continuity (BC) plans. Security mechanisms like data at rest encryption have to
32 be implemented carefully to ensure that resiliency strategies are not impacted.

33 — Virtualization within storage can take many forms and be implemented at different points with the storage
34 infrastructure. This virtualization can mask the physical details associated with the presentation of storage
35 (e.g., a logical unit or filesystem to a server), mask the true capacity of a device, perform policy-driven
36 autonomous data movement (like tiered storage), or completely abstract the storage infrastructure (like cloud

storage). Balancing the security and virtualization to coexist requires careful planning and selection of the right technologies.

— Data growth rates in some organizations are driving increased use of data storage technologies. As an alternative to acquiring additional storage, organizations are employing data reduction technologies such as compression and deduplication. However, these data reduction technologies can be impacted by data at rest encryption mechanisms, and they in turn, can introduce data integrity problems during DR and BC operations.

— As part of the normal data protection strategy, many copies of data end up getting created (e.g., replicated between systems and sites, backups, snapshots, etc.). These copies need to be protected appropriately while they are in use, and then properly sanitized when their usefulness has ended.

— Sensitive and high-value data often need to be protected when transmitted between systems, using mechanisms like the Internet Protocol Security (IPsec). IPsec can have detrimental impacts to the use of certain technologies like NAT, IDS, IPS, or other systems that look deeper into network traffic frames. Whether to rely on IPsec or other in motion protection protocols can hinge on the trade-offs of potentially neutralizing the value of other technologies or where in the network it is to be deployed.

— Many organizations are implementing data at rest encryption to protect sensitive and high-value data. The specific cryptographic mechanisms and the point of encryption are important factors in the actual data protection as well as meeting compliance requirements.

— Successful use of encryption is often predicated on proper management of keying material throughout its lifecycle. This includes correct generation of keys, secure storage and transmittal of key material, replicating keys as part of the normal strategy to ensure availability of the data, and proper disposal of the keying material when it is no longer needed. The sensitivity and importance of the data to be protected may also factor into the key management approach.

Ensuring adequate confidentiality, integrity, and availability of data stored and accessed on current and emerging storage technologies requires a concerted effort within this layer of ICT. Many of these security efforts will focus on:

— Protecting storage management (operations and interfaces)

— Ensuring adequate credential and trust management

— Protecting data backup and recovery resources

— Data in motion protection

— Data at rest protection

— Data availability protection

— Disaster Recovery and Business Continuity support

— Proper sanitization and disposal

— Secure autonomous data movement

— Secure multi-tenancy

5.4 Storage security risks

5.4.1 Background

Storage security risk is created by an organization's use of specific storage systems or infrastructures. Storage security risk arises from:

- a) threats targeting the information handled by the storage systems and infrastructure;
- b) vulnerabilities (both technical and non-technical); and
- c) impact of successful exploitation of vulnerabilities by threats.

Risk management is a key concept in information security. According to ISO/IEC 27005, "the information security risk management process can be applied to the organization as a whole, any discrete part of the organization (e.g. a department, a physical location, a service), any information system, existing or planned or particular aspects of control (e.g. business continuity planning)." The information security risk management process presented in ISO/IEC 27005 consists of context establishment, risk assessment, risk treatment, risk acceptance, risk communication, and risk monitoring and review.

Threats for storage systems and infrastructure include, but are not limited to:

- Unauthorized usage
- Unauthorized access
- Liability due to regulatory non-compliance
- Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks on storage
- Corruption/modification and destruction of data
- Data leakage/breaches
- Theft or accidental loss of media
- Malware attack or introduction
- Improper treatment or sanitization after end-of-use

These threats can give rise to a wide assortment of risks. However, for storage systems and infrastructure the risks associated with data breaches, data corruption or destruction, temporary or permanent loss of access/availability, and failure to meet statutory, regulatory, or legal requirements are the major concerns.

5.4.2 Data breaches

A data breach can be one of the results of a security compromise and it can take many forms. Unauthorized access or disclosure of protected information are two commonly recognized forms of data breaches, but it is important to understand that lesser known forms can include accidental or unlawful destruction, loss, or alteration of data.

Depending on the volume and type of information involved (e.g., personally identifiable information, protected health information, etc.) and the applicable laws and regulations, a data breach can expose the organization to significant risk arising from costs involved in investigating the data breach, making requisite notifications to

affected individuals, litigation expenses, regulatory fines and other legal penalties as well as brand damage accruing from the public disclosure of the data breach.

There are economic and security risks to the entity that has lost their or others' secured information, in that the loss of the information could include things such as:

- secrets or confidential information (e.g., passwords, encryption keys, etc.),
- intellectual property or other sensitive business information,
- Personally Identifiable Information (PII),
- financial account or record information,
- personally identifiable health record information.

Untrusted or unauthorized entities seeking this leaked or spilled information can be of a broad range of sources, be well funded and have diverse motivations.

Table 1 summarizes the storage-based security compromises that are more likely to occur and lists the forms of data breaches that can result from these compromises.

Table 1 — Storage-oriented data breaches

| Security compromise | Potential forms of data breach |
|--|---|
| Theft of storage element or media | Unlawful access, unlawful disclosure, unlawful data loss, unlawful data destruction |
| Loss of storage element or media | Unauthorized access, unauthorized disclosure, accidental data loss, accidental data destruction |
| Accidental configuration changes (e.g., storage management, storage/network resources, incorrect patch management, etc.) by authorized personnel | Unauthorized access, unauthorized disclosure, accidental data destruction, accidental data alteration |
| Malicious configuration changes (storage management, storage/network resources, application tampering, etc.) by external or internal adversaries | Unlawful access, unlawful disclosure, unlawful data destruction, unlawful data alteration |
| Privileged user abuses by authorized users (e.g., inappropriate data snooping) | Unlawful/unauthorized access or disclosure |
| Malicious data tampering by external or internal adversaries | Unlawful data destruction or alteration |
| Denial of service attacks | Unlawful data destruction, loss, or alteration |
| Malicious monitoring of network traffic | Unlawful/unauthorized disclosure |

5.4.3 Data corruption or destruction

Data corruption is the deterioration or damage of computer data (i.e., unintended changes to the original data) caused by human, hardware and software error. It can occur during writing, reading, storage, transmission, or processing. Data corruption may only affect a small, but important portion of data or metadata, which can allow for recovery under the right conditions; or it could also result in permanent data loss if the root cause is allowed to persist. Data destruction on the other hand results in data loss, which can be permanent if data protection mechanisms like backups have not been employed. Both data corruption and destruction can be the result of unintentional or intentional events, and in the latter case, they can be further categorized as malicious or non-malicious.

Events such as fire, flood, power outages, programming bugs and user errors are all examples of general, unintentional sources of data corruption and destruction. Background radiation, head crashes, and aging or wear of the storage hardware are additional sources of problems that are more storage centric. Hardware-based data corruption can generally be detected by the use of checksums, and can often be corrected by the use of error correcting codes, but these "silent corrections" can lead to other problems if storage is not managed well (i.e., temporary correctable errors can turn into permanent ones as the storage device or media deteriorates).

Intentional attacks/events of a malicious nature can be perpetrated by external parties and/or insiders with the purpose of making some or all of the affected data unusable or destroyed. In this context, unusable could mean that unauthorized modifications have been applied, modifications are suspected, or the data can be encrypted with an unknown key or mechanism. Non-malicious attacks typically result from carelessness, lack of knowledge, or intentional circumvention of security for such reasons as "getting the job done," but the impact on the data can be as devastating as malicious attacks.

Employing appropriate mechanisms to detect and remedy data corruption is an important way of maintaining data integrity. Likewise, detecting data loss and recovering this data using data protection mechanisms can guard against the total loss of data.

5.4.4 Temporary or permanent loss of access/availability

Availability is concerned with assuring that there is no or limited denial of authorized access¹⁾ to storage elements, storage network elements, stored information, information flows, services, and applications. In general, data availability is achieved through redundancy involving where the data is stored and how it can be reached.

Loss of availability can often be attributed to a problem with one or more of the following:

— Reliability

— Accessibility

— Timeliness

NOTE 5.4.3: Data corruption or destruction can have the effect of rendering data temporarily or permanently inaccessible, see 5.4.3:

5.4.5 Failure to meet statutory, regulatory, or legal requirements

Organizations can incur significant liabilities and penalties for non-compliance with statutory, regulatory, or legal requirements. These requirements can vary significantly in different jurisdictions, and for multi-national organizations, country specific legislation has a great influence on information security requirements.

1) Within this context, "limited denial of authorized access" means that data continues to be available at a required level of performance within a specified timeframe.

Common compliance issues include:

- breach of country specific privacy requirements
- breach of confidentiality
- non-conformance with an organization's policies (e.g., sanitization)
- inadequate data retention and protection
- insufficient evidence of security (e.g., audit logs, proof of encryption/sanitization, etc.)
- unlawful transfer of data (i.e., moving restricted data out of particular jurisdiction)

These non-compliance issues can result in costly sanctions and remediation (e.g., breach notifications).

6 Supporting controls

6.1 General

Clause 6 provides the controls that *support* storage security technical architectures, their related technical controls, and other controls (technical and non-technical) that are applicable not just to storage. Information on many of these types of controls can be found in ISO/IEC 27002. The controls that are especially important with regard to the use of storage are expanded upon in subclauses 6.2 to 6.8 below, which address securing direct attached storage, storage networks, the management of storage network security, technical controls for different types of storage (block, file, and object-based storage), and storage security services. Data sensitivity, criticality, and value can also be an important consideration in selecting and using controls, so Annex B and specifically B.1.2 should also be consulted.

6.2 Direct Attached Storage (DAS)

A DAS device is a storage element (e.g., HDD, tape, etc.) that is directly connected to a host computer without a storage network in between (i.e., no network device like a hub, router, or switch between the two). DAS devices can take the form of internal storage (i.e., an integral part of the computer system) or external storage (i.e., auxiliary storage). In addition, they are typically dedicated to the system to which they are attached; a DAS device can be shared between multiple computers, as long as it provides multiple interfaces (ports) that allow concurrent and direct access.

These storage elements have limited data access and management interfaces (the latter is usually in-band). As such, the options for securing DAS tend to be limited and include:

- DAS tends to be small in physically size and located in office environment where they can be subjected to malicious attacks (e.g., stolen, destruction, unauthorized access, etc.), so DAS should be physically secured
- To avoid unauthorized access of data on DAS, some form of encryption should be used to protect the at rest data
 - Storage elements with integrated encryption and access control capabilities, also known as Self-Encrypting Drives (SEDs)
 - Host-based or application-based encryption, including Full Disk Encryption (FDE)

- 1 — Media sanitization should be used on all DAS involved with sensitive and high value data (see 6.8.1.2 and
2 Annex A for additional information)
- 3 — If possible, authentication (e.g., FC-SP Authentication) should be used to prevent unauthorized access to
4 data
- 5 — To guard against accidental or intentional data loss or corruption, backups of the DAS contents should be
6 made on a regular basis

7 **6.3 Storage networking**

8 **6.3.1 Background**

9 With the possible exception of DAS, networking plays an important role in storage infrastructures and these
10 networks can include both common networking technologies (e.g., LAN and WAN) as well as storage-specific
11 technologies (e.g., Fibre Channel). In the case of the former, the security guidance offered in ISO/IEC 27033 is
12 instrumental in protecting storage resources that utilize these technologies. The storage-specific technologies are
13 addressed in this International Standard.

14 Storage systems use networking for three primary purposes: 1) storage and retrieval of data, 2) protection of data,
15 and 3) management of storage systems. None of the uses mandate a particular networking technology or
16 approach. For example, some storage management can be performed over the same Fibre Channel interface (i.e.,
17 in-band) used by a server to access data and simultaneously over a TCP/IP connection to the management
18 interface of the storage system.

19 **6.3.2 Storage Area Networks (SAN)**

20 **6.3.2.1 General**

21 A Storage Area Network (SAN) is a specialized, high-speed network that provides block-level network access to
22 storage. SANs are typically composed of servers, switches, storage elements, and storage devices that are
23 interconnected using a variety of technologies, topologies, and protocols. SANs can also span multiple sites (see
24 Figure 1).

25 SANs are often used to improve application availability (e.g., multiple data paths), enhance application
26 performance (e.g., off-load storage functions, use separate networks, etc.), increase storage utilization and
27 effectiveness (e.g., consolidate storage resources, tiered storage, etc.), and improve data protection and security.
28 In addition, SANs typically play an important role in an organization's DR and BC activities.

29 A SAN presents storage devices (such as disk arrays, tape libraries and optical jukeboxes) to a server operating
30 system (OS) such that to the server, the storage appears to be locally attached. This simplified presentation of
31 storage to a server is accomplished through the use of different types of virtualization.

32 SANs are commonly based on Fibre Channel (FC) technology that utilizes the Fibre Channel Protocol (FCP). In
33 addition, the use of FCP over Ethernet, known as Fibre Channel over Ethernet (FCoE), makes it possible to move
34 FC traffic across existing high-speed Ethernet infrastructure and converges storage and IP protocols onto a single
35 cable transport and interface. Other technologies like Internet Small Computing System Interface (iSCSI),
36 commonly used in small and medium sized organization as a less expensive alternative to FC, and InfiniBand,
37 which is commonly used in high performance computing environments, can also be used. Interconnects like Serial

Attached SCSI (SAS) and Peripheral Component Interconnect Express (PCIe), which leverage extenders and switches, are beginning to take on characteristics of a SAN as well.²⁾

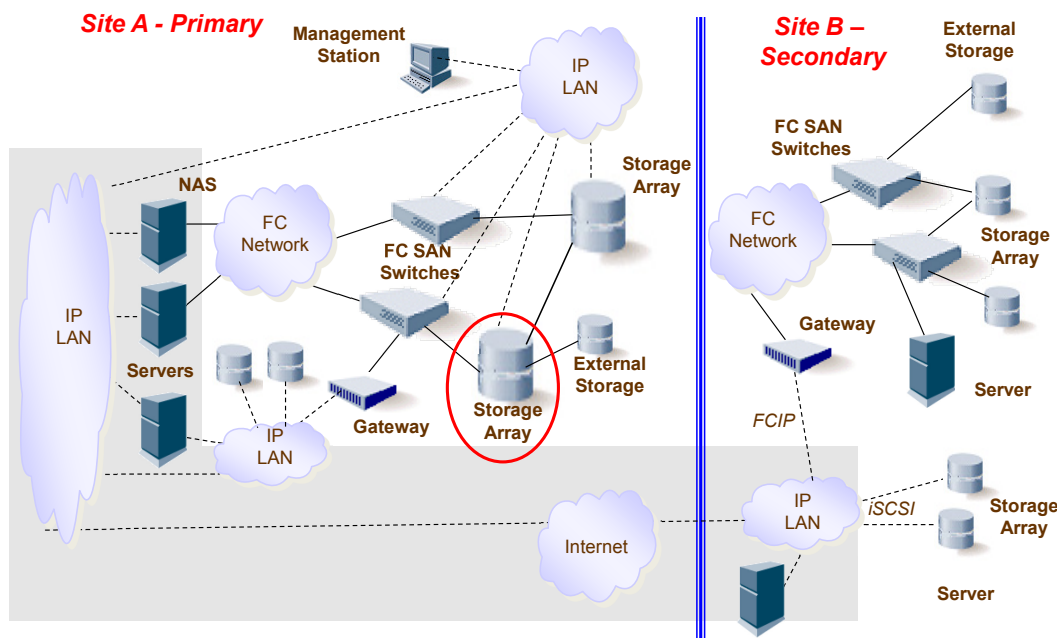


Figure 1 — Storage Area Network (SAN) Example

Security controls relevant to a SAN are grouped into the following categories:

- **Access Control:** Access control on a SAN is implemented through application of zoning, LUN masking, and port binding mechanisms:
 - **Port Binding:** Globally unique identifiers known as World Wide Names (WWN) are used for identification in a SAN. Port binding is a SAN security mechanism that associates a physical port ID and the WWN of the connected device. This association can mitigate snooping attempts by a potential adversary and should be used when possible.
 - **Zoning:** A SAN fabric can be segmented into separate zones to restrict the visibility of portions of a SAN to specific servers and storage devices. Hard zoning uses physical port numbers on a switch for zone membership and is a more secure zoning method since it is not susceptible to spoofing. Zoning should be used in SAN fabrics with a preference for hard zoning.
 - **LUN masking:** A storage device can be divided into different logical units that are identified by numbers (LUNs). Typically not all LUNs are intended to be presented to all servers in the SAN therefore a means to control access to LUNs, known as LUN masking, is necessary. LUN masking can be implemented within the FC switch or at a storage node. LUN masking should be implemented on storage devices but if this is not feasible implement LUN masking on FC switches.

²⁾ This International Standard does not provide guidance specific to InfiniBand, SAS, or PCIe, but the general guidance is applicable.

— **Authentication:** For FC environments it is important for a FC switch to verify the identity of other FC switches in the SAN with whom it communicates. If switch authentication is not implemented a rogue switch could join a SAN and potentially compromise SAN data. For FC environments, FC Switches should mutually authenticate with each other. For iSCSI environments, bi-directional CHAP authentication between iSCSI initiators (clients) and the storage device should be used. CHAP challenges should be random (i.e. not repeated)

— **Encryption:** There are two major components of data confidentiality on a SAN: 1) data in motion and 2) data at rest. Information should be cryptographically protected in SANs when it is in motion as well as when it is at rest on a storage device. For Fibre Channel environments Encapsulating Security Payload (ESP) should be implemented between FC devices to protect user data in motion. For iSCSI environments, IPsec ESP should be implemented between iSCSI devices to provide data confidentiality. To protect data at rest, user data should be encrypted before arriving at the storage device or media. This can require the use of special purpose hardware that can encrypt the data that is being sent to a storage device. Refer to 6.8.2.2 for additional guidance on protection of data in motion and 6.8.2.3 for guidance on protection of data at rest.

A defence in depth strategy (see 7.2.1) helps to mitigate the risk associated with failure of one security control (possible single point of failure) compromising the assets under protection.

Physical and logical isolation of storage elements within a SAN can also play an important role, and can take the form of:

— Physical isolation

— Segregate production from other system classes (e.g., Quality Assurance, Development)

— Where possible, avoid network connections between classes (e.g., a production server connected to both the production and development networks)

— Segregate networks and storage by class where appropriate

— Physically separate systems in each class

— Isolate storage devices from other data centre devices, if practical

— Logical isolation

— Segregate storage traffic from normal server traffic

— Use available network controls to create independent logical domains on common physical infrastructure

— Use trust and access controls to manage membership in the logical domains

— Segregate management traffic from all other traffic

— Carefully review configuration of network gateways

6.3.2.2 Fibre Channel SAN

Fibre Channel is a multi-gigabit-speed network technology used for block-based storage. There are three major Fibre Channel topologies, describing how a number of ports are connected together: point-to-point (two devices are directly connected), arbitrated loop, and switched fabric. Switched fabric topologies along with the Fibre

1 Channel Protocol (FCP), which is the interface protocol used to transmit SCSI traffic on this network technology,
2 are the more interesting from a security perspective.

3 Fabric administrators should take steps to

4 — Control FCP node access

5 — Restrict host access on the switches (e.g., ACLs, binding lists, FC-SP policy³⁾)

6 — Use NPIV (N_Port ID Virtualization) to assign individual N_Port IDs to virtual hosts

7 — Implement switch-based controls

8 — Restrict switch interconnections (e.g., ACLs, binding lists, FC-SP policy)

9 — Carefully consider the adequacy of basic zoning as a security measure

10 — Disable unused ports

11 — Carefully use default zones and zone sets (assume a least privilege posture)

12 — Interconnect storage networks securely by configuring switches, extenders, routers, and gateways (e.g.,
13 FCIP) necessary to meet requirements

14 Subclause 6.5.1 provides guidance on block-based Fibre Channel storage.

15 **6.3.2.3 IP SAN**

16 Internet SCSI or iSCSI, which is described in IETF RFC 3720, is a connection-oriented command/response
17 protocol that runs over TCP, and it is used to access disk, tape and other devices.

18 — Control iSCSI network access and protocols

19 — Avoid connecting iSCSI interfaces to general purpose LANs; segregate for security and performance

20 — Carefully use VLANs when the use of physically isolated LANs is not an option

21 Fibre Channel over TCP/IP (FCIP), defined in IETF RFC 3821, is a pure Fibre Channel encapsulation protocol. It
22 allows the interconnection of islands of Fibre Channel Storage Area Networks through IP-based networks to form
23 a unified Storage Area Network.

24 — Control FCIP network access and protocols

25 — Carefully set up the peer-to-peer relationship between FCIP entities, recognizing that the security
26 policies will be applied uniformly

27 — Consider using a private IP network used exclusively by the FCIP entities

28 — Implement FCIP security measures

3) FC-SP is specified in ANSI INCITS 496-2012, *Information Technology - Fibre Channel - Security Protocols - 2 (FC-SP-2)*

- Use IPsec to secure the communications between FCIP entities.
- Perform cryptographic authentication and data integrity at a minimum
- Protect sensitive data by appropriate confidentiality measures

Subclause 6.5.2 provides guidance on block-based IP storage.

6.3.2.4 FCoE SAN

Fibre Channel over Ethernet (FCoE) is a protocol specification⁴⁾ to encapsulate Fibre Channel frames in Ethernet packets. The Ethernet network that supports FCoE is required to be a lossless Ethernet network,⁵⁾ with switching devices that have internal architectures designed to offer a no-drop packet capability and network flow control mechanisms to enable lossless transmission of packets across the Ethernet infrastructure.

- Leverage the FCP-based security mechanisms (e.g., FC-SP)
- Protect against Ethernet broadcast storms (e.g., allocation of adequate input buffering), which can cause throughput and timeout issues
- ACLs should be used to control network access (e.g., denying specific hosts from unnecessary or unwanted traffic)
- Carefully use VLANs when the use of physically isolated LANs is not an option

6.3.3 Network Attached Storage (NAS)

6.3.3.1 General

Network Attached Storage (NAS) is a data storage technology that provides file-level access to heterogeneous clients over a network. NAS enables a file system physically residing on one server or device to be accessed by remote client computers, appearing to users as a local file system. NAS systems are typically designed and build specifically for NAS purposes, but general purpose server computers can also be used.

NAS systems can be implemented as individual storage servers or as a clustered collection of storage servers that dynamically distributes client connections by slicing and/or striping data and metadata across the clustered storage servers; parallel NFS (pNFS) systems are examples of clustered NAS systems.

Common file system implementations include the Network File System (NFS) and SMB/CIFS, but other technologies like Object-based Storage Device (OSD) and cloud storage exist as well.

Security controls relevant to NAS are grouped into the following categories:

- Authorization controls, such as ACLs, that restrict users' access to file and folder resources provided by the NAS device
- Encryption of data, both in motion and at rest

4) FCoE is specified in ANSI INCITS 462-2010, *Information Technology - Fibre Channel - Backbone - 5 (FC-BB-5)*

5) These networks are full duplex, IEEE 802.3 networks that support 802.3x PAUSE and Jumbo frames to encapsulate the 2KB FC frames.

- 1 — Authentication controls, such as Kerberos, for verifying the identity of users attempting to access NAS data
- 2 Refer to 6.6 for further implementation guidance on NAS and file-based storage.

3 **6.3.3.2 Network File System (NFS)**

4 NFS is a client/server application, communicating with a remote procedure call (RPC)-based protocol. Multiple
5 versions of NFS are specified and in use, including NFS version 3 (specified in IETF RFC 1813), NFS version 4
6 (specified in IETF RFC 3530), and NFS version 4.1 (specified in IETF RFC 5661). From a security perspective,
7 NFS version 3 (NFSv3) is considered less secure and its use should be carefully considered when used with
8 sensitive or high-value data.

9 The following networking guidance is applicable to NFS-based NAS:

- 10 — Control NFS network access and protocols
 - 11 — Enable NFS only if needed. This will eliminate it as a possible attack vector available to an intruder.
 - 12 — Use NFSv4 (or later versions) whenever possible and limit NFSv3 usage
 - 13 — Filter client and management access by IP address for additional security
- 14 — Encrypt client data access (e.g., IPsec) when necessary

15 **6.3.3.3 SMB/CIFS**

16 Server Message Block (SMB) 3.0—the successor to CIFS (Common Internet File System), itself the successor to
17 SMB 1.0—is a protocol intended to provide an open cross-platform mechanism for client systems to request file
18 services from server systems over a network. It is based on the standard SMB protocol widely in use by personal
19 computers and workstations running a wide variety of operating systems.

- 20 — Use later versions of the SMB protocol
- 21 — Turn off low-security session negotiation protocols, such as NTLM v1, LanMan and plaintext; use NTLM v2 or
22 Kerberos instead
- 23 — Maintain up-to-date patch levels
- 24 — Use SMB signing
- 25 — Maintain Active Directory (AD) services securely
- 26 — Use one-way trusts, from leaf domains to parent domains, when possible
- 27 — Control SMB/CIFS network access and protocols
 - 28 — Enable SMB/CIFS only if needed. This will eliminate it as a possible attack vector available to an intruder
 - 29 — Encrypt client data access when necessary

6.4 Storage management

6.4.1 Background

Storage networks and infrastructure elements are complex architectures, which can impose stringent management demands on administrators. To address these demands, organizations implement storage infrastructure management tools and processes to ensure availability and performance of all storage elements, greater data protection and security, centralized auditing, and meeting compliance requirements.

Like the network management description in ISO/IEC 27033-2, storage management also refers to the activities, methods, procedures, and tools that pertain to the operation, administration, maintenance, provisioning and sanitization of storage systems.

- Operation deals with keeping the storage (and the services that the storage infrastructure provides) up and running smoothly. It includes monitoring the storage to spot problems as soon as possible, ideally before users are affected.
- Administration deals with keeping track of resources in the storage infrastructure and how they are assigned. It includes all the "housekeeping" that is necessary to keep the storage under control.
- Maintenance is concerned with performing repairs and upgrades - for example, when equipment has to be replaced, when a storage array needs a microcode update, when a new switch is added to a storage network. Maintenance also involves corrective and preventive measures to make the storage run "better," such as adjusting device configuration parameters.
- Provisioning deals with initializing and equipping a system to prepare it to provide services.
- Sanitization deals with preserving the confidentiality of information remaining on media when it is removed from service or re-purposed by rendering the data unreadable (e.g., by overwriting it with random data, destroying the encryption keys for encrypted data or physically destroying the device).

Performing these storage management activities securely, requires controls associated with authentication and authorization (6.4.2), protecting the storage management interfaces (6.4.3), maintaining accountability and traceability of systems and users (6.4.4), and ensuring the underlying systems used for storage management are adequately hardened (6.4.5). Guidance for each of these topics is provided in the subclauses within this clause 6.4.

6.4.2 Authentication and authorization

6.4.2.1 Authentication

The individuals managing storage systems and infrastructure are generally privileged users. Inappropriate use of system administration privileges (any feature or facility of an information system that enables the user to override system or application controls) can be a major contributory factor to the failures or breaches of systems. To help mitigate these threats, additional authentication measures (see C.1) may be necessary, including but not limited to:

- All users should have a unique identifier (user ID) for their personal use only
- A suitable authentication technique should be chosen to substantiate the claimed identity of a user.
- Use strong passwords (increased minimum number of characters, increased complexity, etc.) with a reduced period of use.

- 1 — Use strong authentication (e.g., a challenge response protocol)
- 2 — Use multi-factor authentication, such as biometric data (e.g. finger-print verification, signature
- 3 verification) and use of hardware tokens (e.g., smart cards)
- 4 — For all remote access, use strong authentication or multi-factor authentication along with secure channels
- 5 — When possible, use a centralized authentication solution (e.g., RADIUS, Single Sign-on, etc.) for improved
- 6 monitoring and control
- 7 — Use multi-factor authentication, complimented by an auto identity provisioned system, when managing
- 8 sensitive and high-value data
- 9 — Disable login to the root account. Remotely log all "sudo" operations

10 In addition to user authentication, storage systems sometimes employ entity authentication, which is the process
 11 by which an agent in a distributed system gains confidence in the identity of a communication partner. This entity
 12 authentication can take place in Transport Layer Security (TLS) and IPsec connections as well as within storage
 13 protocols (e.g., CHAP with iSCSI, DH-CHAP within FCP, etc.). When possible, these entity authentication
 14 mechanisms should be used. Additional guidance is provided in the sections that address these mechanisms.

15 6.4.2.2 Authorization and access control

16 Within market sectors like financial services and healthcare there are trends to align authorization and access
 17 control (see C.2) to a least-privilege model that leverages specific roles. The Storage Networking Industry
 18 Association (SNIA) Storage Management Initiative – Specification (SMI-S) Version 1.5 identifies the following
 19 general roles, which should be implemented and used within storage implementations:

- 20 — *Security Administrator* - This role has view and modify rights to establish and manage accounts, to create and
- 21 associate roles/permissions, for audit logging configurations and contents (audit log event entries can never
- 22 be changed), to establish trust relationships with IT infrastructure (e.g., shared secrets for RADIUS), to
- 23 manage certificate and key stores, to manage encryption and key management, and to set access controls.
- 24 — *Storage Administrator* - This role has "view" and "modify" rights for all aspects of the storage system. No
- 25 access is granted to security-related elements or data.
- 26 — *Security Auditor* - This role has view rights that allow entitlement reviews, verification of security parameters
- 27 and configurations, and inspections of audit logs. No access is granted to the storage, configuration, or data.
- 28 — *Storage Auditor* - This operator-like role has view rights that allow for the verification of storage parameters
- 29 and configurations and inspections of health/fault logs. No access is granted to security-related elements or
- 30 data.

31 Each storage management transaction should be associated with a "security" or "storage" role. These roles can
 32 be important controls to ensure separation of duties with respect to management capabilities.

33 6.4.3 Secure the management interfaces

34 Protecting the management interfaces from unauthorized access and reconnaissance is of paramount importance.
 35 Unauthorized access to management interfaces, occurring due to failure to implement appropriate controls, could
 36 result in data destruction, corruption, and denial of access.

37 Management interfaces for storage systems can take on several physical forms including serial ports (e.g., RS-
 38 232, DB9, DB25, etc.), local area networks, modems, and even the technologies used for the data path (e.g.,

Fibre Channel). Hybrid interfaces (e.g., serial ports plugged into a console concentrator that provides an interface on a LAN) are also relatively common. To protect these physical interfaces organization should:

- Restrict physical access to management interfaces
- Disable and disconnect serial management ports when not in use
- Segregate LAN interfaces used for management from other LAN traffic; physical isolation is preferred, but logical isolation (such as VLANs) should be used at a minimum
- Disable modem ports when not needed

In addition to the physical interfaces, storage systems employ a variety of software and firmware to enable management of storage system. These software interfaces can include simple Command Line Interfaces (CLI), Web-based graphical user interfaces (GUI), support for the Simple Network Management Protocol (SNMP), and host-based proxies that in-band management (i.e., over the data path). To secure these software/firmware interfaces, organizations should:

- Use firewalls and TCP wrappers to restrict access to management networks to authorized hosts and protocols
- Use entity authentication to establish trust relationships between storage systems and the management systems (e.g., FC-SP for in-band management over Fibre Channel)
- Leverage intrusion detection and prevention mechanisms to identify anomalous behaviours and guard against it
- Use ICT infrastructure (DNS, SLP, NTP) with appropriate security controls to avoid indirect attacks
- Employ appropriate privileged user controls, including authentication (see 6.4.2.1), authorization (see 6.4.2.2), and secure auditing/monitoring (see 6.4.4)
- Ensure that operating systems and applications are current and sufficiently hardened against attacks (see 6.4.5)

When storage systems are managed remotely, the following additional security measures should be used:

- Use secure channels for all remote access (VPN, TLS, SSH Secure Shell, HTTPS)
- Employ strong authentication or multi-factor authentication
- Restrict privileges to the minimum needed (i.e., least privilege)

The organization should devise organizational and technical controls to restrict the management interface used for remote (non-local) vendor maintenance sessions. Remote vendor maintenance operations conducted by individuals communicating through an external network such as the Internet impose significant risks, not only regarding availability, but also integrity and confidentiality.

Technical controls should restrict communication traffic (i.e. hosts, ports, and protocols) to the minimum required for remote vendor maintenance operations. After the accessing party is authenticated, additional controls at the access point should be devised to authorize the vendor maintenance session. These include accepting, asking for approval, or denying the requested session. Appropriate logs containing audit records of vendor actions should be generated.

The organization should restrict dial-up access lines to authorized accessing parties. This includes enforcing a modem callback protocol and disabling connection establishment until vendor requests a maintenance session and the request is authorized by the organization.

6.4.4 Security auditing, accounting, and monitoring

Compliance regulations and contractual clauses often include monitoring and reporting requirements. Event logging and systems accounting are key capabilities (see Annex C.5) to help address these requirements. Of these two, event logging is probably the more useful from a storage security perspective because it can be used both real-time and as part of an incident investigation. As such, storage systems and infrastructure need to participate in the organization's event logging program.

- Include storage in the logging policy
 - With regard to storage systems and devices, the following elements of policy should be addressed:
 - Storage systems and devices should participate in audit logging
 - All *significant* storage management events should be collected
 - Log data is preserved
 - Log data is archived and retained according to log data retention policy
 - The device time is synchronized with a reliable, external source
 - The logging policy should include evidentiary expectations (authenticity, chain of custody, etc.)
- Employ external event logging to a trusted⁶⁾ remote source
 - Implement centralized⁷⁾ audit logging to collect events from all sources in a single repository
 - Establish and use a common, accurate time source across the environment to assure that event records from different sources can be correlated
 - For anything other than system health monitoring and debugging, device resident logs are not recommended because they are more easily subjected to tampering or destruction, there is limited storage space available for logs, and they preclude the use of centralized automated analysis, alerting, and archiving.
 - Storage devices should natively log events to one, and preferably multiple,⁸⁾ external log servers
 - Use standard logging protocols like syslog⁹⁾ that support reliable delivery and secure transports (e.g., TLS)

6) A trusted external event logging source is an IT security management product located in a dedicated security zone or domain and is assumed to enforce its security functions correctly.

7) Centralization in this context should not be interpreted as meaning that all audit logging within an organization has to use a common infrastructure. It is more important to have storage systems/ecosystems within a single security domain use a common audit logging infrastructure.

8) Some logging protocols use unreliable network protocols such as UDP and therefore log messages may be lost due to network or server performance. Sending messages to multiple log destinations reduces the risk of inadvertent loss.

- 1 — Audit logging for which compliance, accountability, and/or security serve as the primary drivers should
2 have devices configured to log events as they occur (i.e., no buffering).
- 3 — Implement an analysis protocol to correlate audit log records across event sources to identify significant
4 security events that provide indication of security incidents
- 5 — Ensure that the storage logging is factored into SIEM¹⁰⁾ solutions, when such technology is deployed
- 6 — Ensure complete event logging
 - 7 — Once the types of events to be logged have been determined, then all occurrences of these events
8 should be logged (whether in-band and out-of-band)
 - 9 — The following kinds of events should be logged (a minimum set of security events):
 - 10 — Failed and successful logon attempts
 - 11 — Failed file and object access attempts for sensitive and high-value data
 - 12 — Account and group profile additions, changes, and deletions
 - 13 — Changes to system security configurations (e.g., audit logging, network filtering, zoning changes)
 - 14 — Changes to security server usage (e.g., syslog, network time protocol or NTP, domain name system
15 or DNS, authentication)
 - 16 — System shutdown and restarts
 - 17 — Privileged operations (i.e., administrator initiated changes)
 - 18 — Use of sensitive utilities (e.g., Unix *sudo* or *cron* commands)
 - 19 — Access to critical data files
 - 20 — Movement of virtual servers between physical hosts
 - 21 — Each log entry should include:
 - 22 — a timestamp (date and time),
 - 23 — a severity level,
 - 24 — the source of the log entry (distinguishing name, IP address, etc.),
 - 25 — an event ID as well as a textual description (necessary to enable localization/ internationalization of
26 events – where the event ID remains the same but the textual description could be translated to
27 different languages), and

9) Syslog is defined in IETF RFC 5424 with additional details contained in IETF RFC 3195, IETF RFC 5425, IETF RFC 5426, IETF RFC 5848, IETF RFC 6012, and IETF RFC 6587.

10) ISO/IEC 27044, *Information technology – Security techniques – Guidelines for security information and event management (SIEM)* provides guidelines to assist organizations in preparing to deploy SIEM processes and systems.

— a description of the event.

— Use care when filtering on fields like “severity” as the enterprise logging policy should serve as the guide for determining what kind of filtering is appropriate and what level of information requires long term storage.

— Implement appropriate retention and protection

— Make sure the event log data are handled and retained correctly

— Implement appropriate measures to preserve log integrity and prevent their modification or destruction (either maliciously or accidentally)

— Depending on the importance of the data and/or the type of log entries, protective measures may be required to ensure the confidentiality¹¹⁾ and integrity of the event log data.

— Use special purpose log servers to handle unique and/or sensitive data requirements

— Leverage log relays and log filtering to minimize the impact of specialized storage requirements (e.g., WORM)

6.4.5 System hardening

All operating systems, hypervisors, and applications should be hardened relative to the use of the storage system. There are many existing best practices for various operating systems that should be referenced based on operating system that is being used. Some of the best practices that should be used for any operating system include:

— removal of un-needed/un-used software

— removal of unnecessary accounts

— changes (e.g., rename, disable, change any default password, etc.) to any predefined or default accounts

— closure of all unused ports

— installation of latest patches from a trusted source

— update firmware from a trusted source

— install and maintain malware protection

When elements of the storage infrastructure receive an update (microcode for example) or patches, there should be some assurance that the software to be applied is from a trusted source. Otherwise, attackers can write their own “update” that instead contains malicious code of their choosing, such as a rootkit, botnet, or other malware.

¹¹⁾ Some log entries may expose things like passwords (e.g., when a user types a password instead of the userid), but more subtle problems may exist as well (e.g., search commands that expose specific names and health issues).

6.5 Block-based storage

6.5.1 Fibre Channel (FC) storage

Fibre Channel storage systems use specialized networking (see 6.3.2.2) to present block-based storage resources to computers. These resources usually take the form of logical units and tape devices (including virtual tape).

- Restrict access to storage with WWN filtering (i.e., LUN masking) and other access control mechanisms
- Implement FCP security measures
 - Mutual authentication (per FC-SP-2) should be used with all servers and switches; leverage centralized authentication services when possible
 - If possible, encrypt Fibre Channel connections (e.g., ESP_Header¹²) that leave the protected area (e.g., confines of a physically controlled data centre)

6.5.2 IP storage

Unlike FC storage, IP storage uses TCP/IP networking (see 6.3.2.3), specifically iSCSI, to present block-based storage resources to computers.

- Control iSCSI initiator access by filtering based on source IP addresses and protocols
- Implement iSCSI security measures
 - Use CHAP authentication for both initiators and targets in all iSCSI implementations
 - Consider using IPsec to secure the communication channel when sensitive data could be exposed
 - Use iSNS, SLP, DNS infrastructure with appropriate security controls to avoid indirect attacks

6.6 File-based storage

6.6.1 NFS-based NAS

This type of storage is basically a LAN-attached file server that serves files using the network protocol, NFS (see 6.3.3.2). It consists of an engine that implements the file services and one or more storage devices, on which data is stored. A NAS system can also be SAN-attached, in which case the NAS system is treated just like any other server on the SAN (e.g., provided access to storage, LAN-free backups, etc.). NFS-based NAS systems can take many different forms (e.g., simple NAS filers to highly scalable clusters), and they tend to be highly optimized to handle large numbers of simultaneous file accesses.

For NFS-based NAS systems, the following should be considered:

- Apply access controls to NFS exported filesystems
- Employ user-level authentication whenever possible (e.g., NFSv4 with Kerberos V5)

¹² ESP_Header Authentication and Confidentiality optional headers are defined in *ANSI INCITS 424–2007 Fibre Channel – Framing and Signaling-2 (FC-FS-2)*.

- 1 — Configure the NFS server to export file systems explicitly for the authorized users
- 2 — Configure the NFS server to export file systems with minimum required privileges
- 3 — Avoid granting “root” or “administrator” access to files on network filesystems
- 4 — Make sure NFSv4 ACLs (access control lists) are assigned correctly
- 5 — Use Kerberos authentication for NFSv3
- 6 — Consider using Kerberos Safe and Private modes to sign and encrypt NFS traffic
- 7 — Restrict NFS client behaviours
- 8 — Filter client access to NFS shares whenever possible
- 9 — Do not allow NFS clients to run *suid* and *sgid* programs on exported file systems
- 10 — Secure data on NFS filer
- 11 — Exported file systems should be in their own partitions to prevent system degradation by an attacker
- 12 writing to an exported file system until it is full
- 13 — Encrypt data at rest when necessary
- 14 — Do not allow NFS exports of administrative file systems (e.g., */etc*)
- 15 — Guard against malware (e.g., viruses, worms, rootkits, etc.)
- 16 — Continually monitor content placed in NFS shares and relevant access controls

17 **6.6.2 SMB/CIFS-based NAS**

18 Like NFS-based NAS (see 6.6.1), SMB/CIFS-based NAS is a LAN-attached file server that serves files, but it
19 differs in its use of the network protocols, SMB/CIFS (see 6.3.3.3).

20 For SMB/CIFS-based NAS systems, the following should be considered:

- 21 — Apply access controls to SMB/CIFS exported filesystems
- 22 — Disable unauthenticated access to CIFS shares and NAS devices (i.e. restrict *Anonymous*)
- 23 — Disable “Guest” and “Everyone” access to all CIFS shares
- 24 — Implement authentication and access control via a centralized mechanism (RADIUS, LDAP)
- 25 — Restrict SMB/CIFS client behaviours
- 26 — Enable SMB signing for clients and the NAS device
- 27 — Secure data on SMB/CIFS filers
- 28 — Enable CIFS auditing whenever possible

- Continually review content placed in CIFS shares and relevant access controls
- Encrypt data at rest when necessary
- Guard against malware (e.g., viruses, worms, rootkits, etc.)
- Implement CIFS with strong authentication (NTLMv2, Kerberos)

6.6.3 Parallel NFS-based NAS

As mentioned in 6.3.3.1, NAS devices can be implemented as individual storage servers or as a clustered collection of storage servers. These clusters come in two varieties, symmetric and asymmetric, and the two can be combined. Symmetric clusters allow all of the file servers to be full file servers, with redirection or similar techniques used to select the appropriate server based on the client and what files the client wants to access. A common technique is to partition the filesystem namespace with different servers being responsible for different portions of that namespace - in such a structure, filename resolution can result in the client traversing a namespace path that involves multiple file servers. Asymmetric clusters split functionality across servers - parallel NFS uses at least one primary file server and multiple secondary storage servers that are slaved to the primary server (client has to contact the primary file server in order to understand what data is stored on the secondary storage servers and how to access it).

For a symmetric cluster, including clustering of the primary file servers for pNFS, the primary guidance is consistent application of controls and control mechanisms (e.g., authentication and authorization) across the clustered servers so that the security assurance properties don't depend on which file server the client happens to access.

For asymmetric clusters, that consistent application of controls and control mechanisms is important, but the different roles of the servers can place responsibilities on the client that are not present in a symmetric cluster. A specific complication for pNFS is that the secondary storage servers may not use the same protocol (NFS) as the primary file server(s), requiring control implementation in a consistent fashion across both protocols. Another important example is that the pNFS block/volume layout requires trusting the client to respect the layout information obtained from the primary server and not access block storage for which it does not have a layout - this should be somehow captured in a control that enforces the recommendation to not use the pNFS block/volume layout when clients cannot be relied upon to do this - see the security considerations (Section 4) in RFC 5663.

In both cases, controls should not depend on path traversal of the filesystem namespace across servers - direct client access to servers that the client isn't supposed to "start from" is an important consideration in effective application of controls, as some servers can export a partial filesystem namespace (no "root" that the client is expected to start from) or no filesystem namespace at all. A specific example of the latter is that there should be limited or no control dependence on the fact that some servers cannot export a filesystem namespace (e.g., as is the case for pNFS storage servers). Another example is that a directory ACL that blocks namespace path traversal may be an insufficient control for a namespace path that crosses into another file server - an effective control has to deal with direct client access to that latter file server, as such access would bypass the directory ACL.

- Controls and control mechanisms should be applied consistently across clusters (both symmetric and asymmetric)
- Security assurance properties should not be dependent on the client accessing a specific file server
- For asymmetric clusters, controls should be implemented such that they are consistent across different protocols

— Security controls should not be dependent on path traversal of the filesystem namespace across servers

6.7 Object-based storage

6.7.1 Cloud storage

Both proprietary and standards-based, cloud storage offerings, based on RESTful HTTP technology, are in use. These object-based implementations often have a dependency on HTTPS (HTTP over TLS) to secure the underlying communications. Additional security features may be specified, but there can be significant difference in terms of what is implemented versus what ultimately gets used.

Cloud storage, based on the ISO/IEC 17826:2012 *Cloud data management interface (CDMI)* specification, has security elements that are sufficiently described that specific guidance can be provided. Security measures within CDMI can be summarized as transport security, user and entity authentication, authorization and access controls, data integrity, data and media sanitization, data retention, protections against malware, data at rest encryption, and security capability queries. With the exception of both the transport security and the security capability queries (mechanism to determine what is supported), which are mandatory to implement (use is always optional), the security measures can vary significantly from implementation to implementation.

When using CDMI, the following should be used by clients:

- ensure that Transport Layer Security (TLS) is used for all transactions (see 6.8.2.2)
- check the security capabilities of the cloud service provider's CDMI implementation and make a risk-based decision on whether the offered security is adequate
- authenticate CDMI entities (certificates for servers and HTTP basic authentication for clients)
- use CDMI Domains to provide a place for authentication mappings to external authentication providers
- when possible, enable CDMI security logging and retrieve the event data in a regular and timely fashion
- align the automatic deletion capability (CDMI Deletion) with the organization's data retention policy
- prior to using CDMI Holds, understand the process and mechanism for lifting the CDMI Hold
- for cryptographic functionality, always verify that the implementation has used a requested CDMI Capability (supported operation), and not something different
- use the provided sanitization facilities to clear sensitive data from the cloud service provider's storage

6.7.2 Object-based Storage Device (OSD)

An Object-based Storage Device (OSD) is a computer storage device, similar to disk storage but working at a higher level (i.e., the physical storage locations are hidden under the object interface and managed by the storage device itself). Instead of providing a block-oriented interface that reads and writes fixed sized blocks of data, an OSD organizes data into flexible-sized data containers, called objects. Each object has both data (a linear sequence of bytes) and metadata (an extensible set of attributes describing the object), which is accessed by specifying the Object Identifier (OID) and an (offset, length) tuple. The command interface to the OSD includes commands¹³⁾ to create and delete objects, write bytes and read bytes to and from individual objects, and to set

¹³⁾ The initial OSD standard, ANSI/INCITS 400-204, *Information technology – SCSI Object-based Storage Device Commands (OSD)*, was approved in 2004.

and get attributes on objects. The OSD is responsible for managing the storage of objects and their metadata. The OSD implements a security mechanism that provides per-object and per-command access control.

To ensure secure access to storage, every command is accompanied by a cryptographically secure capability that identifies a specific object and the list of operations that can be performed against a specific object. Capabilities not only provide the per-device security that is lacking in typical block-based storage, but they also facilitate fine-grained access to individual objects. This enables storage-device sharing among diverse applications with unique security requirements.

OSD uses a credential-based access control system composed of three active entities: the object store (the OSD), a security manager, and a client. As a capability-based access control system, all requests to the object store are accompanied by a capability, which encodes a set of rights the holder has on an object, and is cryptographically secured.

To use OSD securely, the following should be followed:

- IPsec should be used for all transactions involving sensitive data on insecure networks
- the object store should verify the authenticity of the capability prior to performing an operation
- clock synchronization between the OSD and the security manager should be implemented using a secure protocol
- capability expiration times should have limits that minimize the amount of time a compromised capability can be used
- working keys (used to generate capability keys) should be refreshed frequently

6.7.3 Content Addressable Storage (CAS)

Content Addressable Storage (CAS), sometimes called Fixed Content Storage (FCS), technology is intended to store data that does not change (is fixed) in time. CAS typically exposes a digest generated by a cryptographic hash function (such as MD5 or SHA-1) from the document it refers to. The main advantages of CAS technology are that the location of the actual data and the number of copies is unknown to the user.

CAS supports retrieval of documents given their content digests, and provides an assurance that the retrieved document is identical to the one originally stored. (If the documents were different, their content addresses would differ.) In addition, since data is stored into a CAS system by what it contains, there is never a situation where more than one copy of an identical document exists in storage. By definition, two identical documents have the same content address.

If the hash function used by the CAS system is weak, this method could be subject to collisions in an adversarial environment (different documents generating the same hash). Therefore, it is important for the CAS system to use a robust hashing mechanism.

Users and applications should be authenticated and authorized before access is granted to the CAS system. This prevents unauthorized users from storing data or retrieving data. Additionally the CAS system should ensure that content will be readable and accessible over its entire life-cycle. Finally, the CAS system should employ a robust hashing mechanism.

CAS is a particularly useful technology when addressing needs for short and medium-term retention requirements (see 7.4.2).

6.8 Storage security services

6.8.1 Data sanitization

6.8.1.1 General

Sanitization refers to the general process of removing data from storage media, such that there is reasonable assurance that the data cannot be easily retrieved or reconstructed (see C.4).

To effectively use this standard for all media types, organizations and individuals should categorize the information, assess the nature of the medium on which it is recorded, assess the risk to confidentiality, and determine the future plans for the media (for example, reuse). Then decide on the appropriate type of sanitization. The selected type should be assessed as to cost, environmental impact, etc., and a decision made that best mitigates the risk to confidentiality and best satisfies other constraints imposed on the process.

Sanitization operations can be costly and time consuming, but they are necessary for security reasons. The level of sanitization operations should be carefully balanced against the risks. Particular attention should be paid to Personally Identifiable Information (PII) and Electronic Healthcare Records (EHR) as well as business or mission critical data (e.g., trade secrets, intellectual property, etc.).

6.8.1.2 Media-based sanitization

When storage media are transferred, become obsolete, or are no longer usable or required by an information system, it is important to ensure that residual magnetic, optical, electrical, or other representation of data that has been deleted is not easily recoverable.

Once a decision is made and after applying relevant organizational environmental factors, then Annex A should be used to determine recommended sanitization of specific media. Although the use of Annex A is recommended here, other methods exist to satisfy the intent of clear, purge (still relevant in some cases), and destroy, and methods not specified in Annex A may be suitable as long as they are vetted and found satisfactory by the organization. Not all types of available media are specified in this International Standard, and for those media not included, organizations should identify and use processes that will fulfil the intent to clear, purge, or destroy their media.

Sanitization of media at end-of-life situations is recommended, even when using encryption methods.

6.8.1.3 Logical sanitization

Many storage devices virtualize the underlying storage media and present it as logical storage. A well-known example is the logical unit on a storage array, which can have a size that far exceeds the capacity of a single storage medium as well as using a large quantity of storage media (for example, a logical unit built from a RAID 5 stripe across 100+ disk drives with spare drives). The situation can be further complicated when logical storage is replicated (i.e., multiple copies of the data exist) to support server virtualization and Disaster Recovery. For these types of situations, it is almost impossible to identify all of the underlying storage media. Further, it may not be appropriate to sanitize all of the physical media because multiple logical storage instances can coexist on shared physical media.

If the logical storage is writeable, then sanitization may be possible using an overwrite technique or for encrypted content (e.g., files), the encryption keying materials can be destroyed. This approach will also sanitize the on-line replications of this logical storage. The overwrite technique will not, however, sanitize off-line storage, which includes backups and CDP storage.

6.8.1.4 Proof of sanitization

When sanitization is being performed as part of a compliance activity, it is important to review the specific requirements. These requirements can take the form of specific overwrite techniques, proof of sanitization, etc.

Organizations should maintain a record of sanitization activities to document what media were sanitized, when, how they were sanitized, and the final disposition of the media. Often when an organization is suspected of losing control of its information, it is because of inadequate record keeping of media sanitization.

Proof of sanitization takes on at least two forms: 1) an audit log trail and 2) a certificate of sanitization. These sanitization records are the evidence that organizations should retain for compliance/legal purposes or they run the risk of sanctions and/or costly data breach notifications. The importance of this proof along with the chain of custody requirements associated with the evidence serve as the primary drivers for placing sanitization under the control of security personnel.

The certificate of sanitization should include the following information at a minimum:

- Manufacturer
- Model
- Serial number
- Media type (e.g., magnetic, flash, hybrid, etc.)
- Media source (i.e., user or system the media came from)
- Sanitization description (i.e., clear, purge, damage, destroy)
- Sanitization method used (e.g., degauss, overwrite, block erase, cryptographic erase, etc.)
- Tool used (including version)
- Verification method (e.g., full, quick sampling, etc.)
- For both sanitization and validation:
 - Name of person
 - Position/title of person
 - Date and Time (completion)
 - Location
 - Contact information (e.g., telephone number, email address, etc.)
 - Field for the signature of the person performing sanitization

In addition to the details associated with the certificate of sanitization, the audit trail should capture time stamped transactions and progress associated with sanitization. For example, the initiation and conclusion of the sanitization operation as well as intermediate overwrite and verification progress should be reflected.

In case devices are in an inoperable state, proof of sanitization should be achieved by additional physical destruction of the media.

6.8.1.5 Verification of sanitized media

The goal of sanitization verification is to ensure that the target data was effectively sanitized. When supported by the device interface (such as an ATA or SCSI HDD or SSD), the highest level of assurance of effective sanitization (outside of a laboratory) is typically achieved by a full reading of all accessible areas to verify that the expected sanitized value is in all addressable locations. A full verification should be performed if time and external factors permit. This manner of verification typically only applies where the device is in an operational state following sanitization so that data can be read and written through the native interface.

If an organization chooses representative sampling then there are three main goals applied to electronic media sanitization verification:

- a) Select pseudorandom locations on the media each time the analysis tool is applied. This reduces the likelihood that a sanitization tool that only sanitizes a subset of the media will result in verification success in a situation where sensitive data still remains.
- b) Select locations across the addressable space. For instance, conceptually break the media up into equally sized subsections. Select a large enough number of subsections so that the media is well-covered. The number of practical subsections depends on the device and addressing scheme. The suggested minimum number of subsections for HDD leveraging LBA addressing is one thousand. Select at least two non-overlapping pseudorandom locations from within each subsection. For example, if one thousand conceptual subsections are chosen, at least two pseudorandom locations in the first thousandth of the media addressing space would be read and verified, at least two pseudorandom locations in the second thousandth of the media addressing space would be read and verified, and so on. In addition to the locations already identified, include the first and last addressable location on the storage device.
- c) Each consecutive sample location (except the ones for the first and last addressable location) should cover at least 5% of the subsection and not overlap the other sample in the subsection. Given two non-overlapping samples, the resulting verification should cover at least 10% of the media once all subsections have had two samples taken.

Editor's Note: The following text describing the process of verification of a cryptographic erase sanitization operation was identified during the 2ndCD ballot as being a problem. Replacement text is still needed.

Cryptographic erase has different verification considerations than procedures such as rewriting or block erasing, because the contents following cryptographic erase may not be known and therefore cannot be compared to a given value. When cryptographic erase is leveraged, there are multiple options for verification, and each uses a quick review of a subset of the media. Each involves a selection of pseudorandom locations to be sampled from across the media.

The first option is to read the pseudorandom locations prior to cryptographic erase, and then again following cryptographic erase to compare the results. This is likely the most effective verification technique. However, this technique cannot always be available because, for example, the person performing the sanitization may not have access to the cryptographic key (e.g., MEK) needed to decrypt the data stored on the drive. Alternatives include searching for strings across the media or looking for files that are in known locations, such as operating system files likely to be stored in a specific area.

The number of locations and size of each sample should take into consideration the risks in transferring the target data to the storage media of the machine hosting the sanitization application. As a result, the proportion of the media covered by verification for the cryptographic erase technique can be relatively small (or at least lower than the above guidance of 10% for verification of non-cryptographic sanitization techniques), but should still be applied across a wide range of the addressable area.

As part of the sanitization process, in addition to the verification performed on each piece of media following the sanitization operation, a subset of media items should be selected at random for secondary verification using a separate validation tool. The secondary validation tool should be from a separate developer. For the secondary validation, a full validation should be performed. At least 20% of sanitized media (by number of media items Sanitized) should be verified. The secondary validation provides assurance that the primary operation is working as expected.

6.8.2 Data confidentiality

6.8.2.1 General

Within storage infrastructures, data confidentiality is typically maintained using some method of encryption. These methods are most often associated with protecting data while it is transferred (sometime referred to as in flight or in motion) within the storage infrastructure or as it is stored (or at rest) within a device or on storage media.

The process of encryption is a matter of applying an encryption algorithm (or cipher) to plaintext data yielding encrypted data (or ciphertext). Conversely, a decryption transforms ciphertext back into its original plaintext. The definition and specification of many important ciphers relevant to storage can be found in: ISO/IEC 18033:2005, NIST FIPS 197, NIST Special Publication 800-67, and IEEE 1619.2-2010.

For some types of ciphers (e.g., n-bit block ciphers) there are multiple ways (called modes of operation) in which the cipher can be used to encrypt plaintext. The definition and specification of common modes of operation can be found in: ISO/IEC 10116:2006, NIST Special Publication 800-38A, NIST Special Publication 800-38C, NIST Special Publication 800-38D, NIST Special Publication 800-38E, and IEEE 1619-2007.

Ciphers work in association with a key and possibly other keying material (e.g., initialization vectors). In a symmetric cipher, the same key is used with both the encryption and decryption algorithms. In an asymmetric cipher, different but related keys are used for encryption and decryption. The management and protection of keys (known as key management) is critically important in maintaining data confidentiality.

The purpose of key management is to provide procedures for handling the cryptographic keying material used with symmetric or asymmetric cryptographic mechanisms. The definition and specification of different aspects of key management can be found in: ISO/IEC 11770 (Part 1 & 2) and NIST Special Publication 800-57 (Part 1 & 2).

6.8.2.2 Encrypting transferred data

Within storage infrastructures, data confidentiality and/or integrity (digital signature or authentication code) of the information being transferred between two points can be of interest, especially for data that leaves the confines of a physically controlled data centre.

Protocols such as FC-SP ESP Header, IPsec,¹⁴⁾ TLS,¹⁵⁾ or even host based encryption techniques can provide additional protection to the information as it is transferred. These methods are most often associated with protecting data while it is in motion (sometimes referred to as in flight or in transit).

In flight data protection is generally a temporary protection of the data, which may exist only while it is being moved. For in motion encryption the sender applies an encryption algorithm and sends the ciphertext. It can also apply an integrity algorithm and send the integrity value. Conversely, a receiver applies a decryption algorithm which transforms ciphertext back into its original plaintext and the receiver performs a check of the integrity value.

14) IETF RFC 6071 *IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap* provides an excellent snapshot of IPsec- and IKE-related RFCs.

15) The most recent version of TLS is specified in IETF RFC 5246 *The Transport Layer Security (TLS) Protocol Version 1.2*.

There are various standard specifications including the FC-SP standard, IPsec RFC's, and TLS RFC's that details alternatives for securing data in motion.

For some protocols there are multiple modes or options of operation in the standards. And in addition there are multiple cipher modes or digital signature (integrity) algorithms. The definition and specification of the modes of operation can be found in: ISO/IEC 10116:2006.

Protection of data in motion works in association with a key establishment or key agreement process or protocol. The management and protection of the initial authentication keys is critically important in maintaining data confidentiality and integrity of data in motion. The previously cited standards detail additional information of critical security parameters that have to be protected when using in motion data protection methods.

— When protection of data in motion is needed, it should provide end-to-end protection

— Encryption of data in motion can impose significant computational burdens on the communicating entities, so appropriate compensations should be implemented to minimize the impacts

— For IPsec, version 3 and IKE version 2 (or later versions) should be used

— For TLS, version 1.2 (or later) should be used

6.8.2.3 Encrypting data at rest

With increasing amounts of sensitive and regulated data being stored, organizations are taking steps to ensure this data is stored in encrypted forms. Although encrypting data as close as possible to its origin and use is the ideal situation, encryption of at rest data within the storage infrastructure does provide a basic level of protection against breaches stemming from the loss of control of media, especially tape. Consequently, organizations are using encryption mechanisms within storage devices (Self-Encrypting Drives as well as controller-based technologies), switches, specialized appliances, host bus adapters, etc.

Implementing data encryption requires much more than just purchasing a device with encryption features and connecting it to an existing storage infrastructure. The positioning of the encryption mechanism (the point of encryption) in the infrastructure needs to be carefully chosen, and arrangements made to provision that location with keying material. The data to be processed needs to be identified, and in some cases its location needs to be changed. In addition, adequate proof of encryption, which is likely to take the form of logs, needs to be created and integrated into the audit log infrastructure. See clause 7.5 for additional information.

The use of all types of encryption for storage relies on the management of cryptographic keys. Poor key management can easily compromise data no matter how strong the encryption is. Ultimately, the security of information protected by cryptography directly depends on the strength of the keys, the effectiveness of mechanisms and protocols associated with keys, and the protection afforded to the keys. All keys need to be protected against modification, and secret (for symmetric encryption) and private (for asymmetric or public key encryption) keys need to be protected against unauthorized disclosure. Key management provides the foundation for the secure generation, storage, distribution and destruction of keys. Overall frameworks for key management are given in ISO/IEC 11770. In addition, the OASIS Key Management Interoperability Protocol (KMIP) specification and profiles¹⁶⁾ are the dominant mechanism for key management within storage infrastructures.

For data at rest encryption on storage, the following should be followed:

16) OASIS KMIP is specified in a pair of documents: *OASIS Key Management Interoperability Protocol Specification Version 1.1* and *OASIS Key Management Interoperability Protocol Profiles Version 1.1*.

- 1 — Encryption algorithms and modes of operations designed specifically for storage technology should be used;
2 for example, XTS-AES (specified in IEEE Std 1619-2007) for HDD and CCM or GCM (as described in IEEE
3 Std 1619.1-2007) for tape
- 4 — Limit the amount of time a key is in plaintext form and prevent humans from viewing plaintext keys
- 5 — Cryptographic keys should only be used for one purpose, specifically, do not use key-encrypting keys (also
6 known as key wrapping keys) to encrypt data or use data encrypting keys to encrypt other keys
- 7 — Randomly choose keys from the entire keyspace by using a cryptographically strong Random Number
8 Generator (RNG)
- 9 — Check for and avoid use of known weak keys
- 10 — Data encryption keys should be limited to a finite cryptoperiod (typically no more than 2 years) or to a
11 maximum amount of data processed
- 12 — When possible, storage systems and infrastructure should use KMIP-compliant key management
13 infrastructure to facilitate centralized key management; compliance should be measured against the "Storage
14 Client Conformance Clauses" in OASIS KMIP Profiles v1.1

15 6.8.3 Data reductions

16 As a routine course of business, organizations may attempt to reduce the amount of data they store and transmit
17 in an effort to reduce costs. Two of the more common approaches are data compression and data deduplication.
18 Data compression seeks to reduce the amount of data by encoding it with a known algorithm¹⁷⁾ to produce a
19 representation of the data that uses fewer bits of storage than the unencoded representation. Data deduplication,
20 on the other hand, attempts to replace multiple copies of data with references to a shared copy. These two
21 techniques can be used together to maximize data reduction.

22 Data compression is commonly used in conjunction with tape storage to reduce the number of tapes required for
23 things like backups. In addition, compression can be an integral part of the network gateways used in remote
24 replication to reduce the bandwidth requirements for Disaster Recovery and business continuity support. Data
25 compression is typically performed in hardware so some care is required to ensure the encoded data can be
26 decoded later (for example, when a tape is read by a different tape drive or when the compressed data is received
27 by a network gateway).

28 Data deduplication can take place at a variety of different points within the storage infrastructure, including at the
29 file system level, in-line to the storage network, and the storage device.

30 In and of themselves, data reduction technologies do not represent security mechanism. However, their presence
31 can be impacted by storage security activities.

32 — When encryption is used along with compression, the compression should be applied before the encryption
33 because ciphertext does not effectively compress; the reverse order should be used on the other end (i.e.,
34 decryption followed by expansion)

35 — When encryption is used along with deduplication, the deduplication should be applied before the encryption
36 because deduplication is not effective on ciphertext; the reverse order should be used when the data is to be
37 decrypted

17) Compression algorithms include lossey (a portion of the original information is lost) and lossless (preserves the entire content of the original data) approaches, but in the storage industry only the lossless algorithms are used. Applying a particular compression algorithm to multiple instances of data will result in identical encoded/compressed data.

- When both compression and deduplication are used along with encryption, the order of use should be deduplication and compression or compression and deduplication, and then encryption; the reverse order should be used when the data is to be decrypted
- Compression and/or deduplication can impact DR and BC implementations, so they should be factored into the design, documentation, and testing of DR and BC solutions

7 Guidelines for the design and implementation of storage security

7.1 General

Despite the increased power of personal computers and departmental workstations, there continues to be a dependency on centralized data centres due to the need for data integration, data consistency, and data quality. With the enormous growth of critical data volumes, many organizations have adopted storage-centric architectures for their ICT infrastructure. Consequently, storage security plays an important role in securing this data, and in many instances, it serves as the last line of defence.

Designing and implementing storage security solutions requires adherence to core security design principles. In addition, the controls and guidance described in Clause 6 have to be integrated into the design and implementation of storage security solutions to counter storage security threats. Data sensitivity, criticality, and value can also be an important consideration in designs (see Annex B and specifically B.1.2).

Common risk areas associated with storage security architectures are design failures due to poor design and/or the lack of appropriate consideration of business continuity planning or the design does not correspond to the current or expected threat level. A design should consider all relevant threats and vulnerabilities in the storage system as described in 5.4.

Information on assessing security risks and associated threats can also be found in ISO/IEC 27001, ISO/IEC 27002 and ISO/IEC 27005. Clause 7 identifies general design issues to consider as part of storage security architecture.

7.2 Storage security design principles

7.2.1 Defence in depth

Organizations need to look at security not just from one perspective, but as a pervasive layered approach that is comprehensive across all applications, systems, networks, storage, and devices. Adopting such a layered approach is considered to be defence in depth especially when it combines policy, design, management and technology. The degree to which defence in depth is pursued is different for each organization, depending on factors like data value and sensitivity, compliance requirements, adversarial capabilities and activities, etc.

An important defence in depth principle leverages the use of multiple security controls or security techniques to help mitigate the risk of one component of the defence being compromised or circumvented. An example could be anti-virus software installed on individual workstations when there is already virus protection on the firewalls and servers within the same environment.

Specific guidance includes:

- Ensure a balanced focus on the three primary elements: people, technology, and operations
- Follow through with effective information assurance policies and procedures, assignment of roles and responsibilities, commitment of resources, training of critical personnel, and personal accountability

- 1 — Deploy protection mechanisms at multiple locations to resist all classes of attacks
- 2 — Deploy multiple defence mechanisms (layered) between potential adversaries and targets
- 3 — Include both detection and protection mechanisms
- 4 — Deploy robust key management and Public Key Infrastructure (PKI) frameworks that support all information
5 assurance technologies and that are highly resistant to attack
- 6 — Maintain visible and up to date system security policies
- 7 — Actively manage the security posture of the storage technology and protection mechanisms (e.g., install
8 security patches and anti-virus updates, maintain ACLs, etc.)
- 9 — Perform regular security threat assessments to determine the continued security readiness
- 10 — Monitor and react to current threats
- 11 Security solutions based on the layered approach are flexible and scalable as well as being adaptable to the
12 security needs of the organization.
- 13 For storage, a layered approach means that security controls are deployed and used throughout the storage
14 infrastructure, including the HBA/NIC in host computers, storage network switches/routers, storage appliance,
15 storage elements, and storage devices.

16 7.2.2 Security domains

- 17 Security domains use the concept that system resources of different sensitivity levels (i.e., different risk tolerance
18 values and threat susceptibility) should be differently located. This creates a way to have the systems make
19 available only such data that is necessary for conducting the tasks for that particular domain. As a design principle,
20 the architecture should enforce domain separation to ensure that resources to which an entity has access cannot
21 be accessed or affected by another entity.
- 22 For storage infrastructure, a security domain will typically be represented as a SAN, especially when sensitive
23 data is being stored and processed within the storage elements. In situations where the data sensitivity is low,
24 zoning and VLANs can be considered acceptable, but it is important to note that this generic capability (as
25 opposed to FC-SP zoning) is not a security mechanism.
- 26 Building on the compartmentalization principle described in ISO/IEC 27033-2, the following storage security
27 design rules should be considered:
- 28 — Factor data sensitivity into the use of security domains
- 29 — Storage and storage networks of different sensitivity levels should be located in different security
30 domains
- 31 — Devices and computer systems providing services for external networks (e.g., the Internet) should be
32 located in different domains (De-Militarized Zone or DMZ) than internal network devices and computer
33 systems
- 34 — Strategic assets should be located in dedicated security domains
- 35 — Devices and computer systems of low trust level should be located in dedicated security domains with
36 limited or no access to storage assets

- 1 — Factor purpose in the use of security domains
- 2 — Storage and storage networks used for different purposes (e.g., development, production, management,
- 3 etc.) and using different technologies (e.g., CIFS/NFS, iSCSI, CDMI, etc.) should be located in separate
- 4 security domains
- 5 — Storage networks should be in different security domains than regular networks (e.g., corporate LANs)
- 6 — Storage device and storage network management systems should be located in dedicated security
- 7 domains
- 8 — Systems in development stage should be located in different domains than production systems
- 9 — Storage devices that may be permitted to reside with a single security domain, but used for multiple purposes
- 10 or hold multiple levels of sensitive data, should be further isolated (using zoning, VLANs, and VSANs) to
- 11 minimize possible interactions

12 7.2.3 Design resilience

13 Storage security design should incorporate several layers of redundancy to eliminate single points of failure and to
 14 maximize the availability of the storage infrastructure. This includes the use of redundant interfaces, backup
 15 modules, standby devices, and topologically redundant paths. In addition, the designs should also use a wide set
 16 of features destined to make the storage more resilient to attacks and network failures.

17 7.2.4 Secure initialization


18 As a design principle, the architecture should support a secure initialization sequence to ensure the transition from
 19 a “down” state after a power-on or reset is applied. During the initialization phase externally accessible processes
 20 and network interfaces should not be available or at a minimum deny access until the subjects are authenticated.
 21 Software and OS load processes should start from a known state with secure values specified by the system
 22 administrator when the system was last operational.

23 7.3 Data reliability, availability, and resilience

24 7.3.1 Reliability

25 At a basic level, reliability is the probability that a device will perform its required function under stated conditions
 26 for a specific period of time. Reliability is quantified as:

- 27 — MTBF (Mean Time Between Failures) for repairable product, which is the average time available for a system
- 28 or component to perform its normal operations between failures (see Figure 2)
- 29 — MTTR (Mean Time To Repair) for repairable product, which is the average time to repair a failed component
- 30 — MTTF (Mean Time To Failure) for non-repairable product, which is the average time available for a system or
- 31 component to perform its normal operations until it fails



$$\text{Time Between Failures} = \{ \text{down time} - \text{up time} \}$$

Figure 2 — Quantification of Reliability

Within the context of storage, system compromises and attacks can have negative impacts on MTBF, MTTR, and MTTF. In addition, the inclusion of security features (like malware protection), the application of system or application patches, or other system hardening measures like those described in 6.4.5 can also have impacts, especially if they are not approved by the vendor.

- The reliability of the storage system and infrastructure should not be adversely impacted by the inclusion of security features
- Vulnerabilities should be proactively managed to minimize their impacts on system reliability
- Controls should be assessed to determine whether they are appropriate to ensure the reliability and security of data

7.3.2 Availability

In the context of storage, data availability typically refers to how accessible data is when stored in some form, usually in reference to remote storage of data through a network or external storage media. This term is often used to refer to several different concepts, primarily how reliable the data is with regard to people trying to access it, in terms of “uptime,” and how quickly someone can access the data.

Availability is usually measured as a probability that something will be there when it is needed (i.e., the proportion of time a system is in a functioning condition), and it can be calculated as the ratio of (a) the total time system is capable of being used during a given interval to (b) the length of the interval. For example, a storage array that had approximately 5 minutes of downtime in a year, assuming 24x7 operations, would have an availability of 0.99999 (99.999%).

To achieve high availability of data, significant amounts of hardware and software redundancy (e.g., automated I/O path failover, redundant components, RAID protection, global hot spares and mirrored data cache with battery back-up) are implemented within contemporary storage systems as well as the storage infrastructures. In addition, data redundancy mechanisms (e.g., mirroring and replication) as well as data protection mechanisms (e.g., backups and CDP) are often used to ensure fast data recoveries in the event of a failure.

- Because of the importance of availability, storage security designs and implementations should strive to minimize impacts to availability (e.g., minimize single points of failure)
- Data encryption keys should be carefully managed to avoid data availability problems when keys are unavailable or inadvertently destroyed

- Data protection mechanisms (like backups, replication, etc.) should be part of availability designs to guard against major outages due to system failures

7.3.3 Backups and replication

Because of the increased dependency on data availability and integrity, many organizations employ a range of data protection mechanisms like backups and replications for increased data resiliency. Unfortunately, the focus is often on the creation of the backups and replicated data sets rather than the ability to use them to recover from problems. All of the data protection solutions should be viewed as data recovery mechanisms.

The data protection mechanisms themselves also need a measure of security, including but not limited to:

- Data protection mechanisms (like backups, replication, etc.) should be designed with quick recoveries in mind, rather than just preservation of data
- Backup security
 - Ensure that the backup approach, especially for business/mission critical data, is aligned with its associated restore strategy
 - Ensure that the backup approach provides adequate protections against unauthorized access (e.g., encryption)
 - Establish a chain of trusted individuals (and vendors) who handle the storage media
 - Implement backup validations to show “proof” that restore requirements are being met
- Replication security
 - Ensure that the replication approach, especially for business/mission critical data, is aligned with its associated reliability, fault-tolerance, or performance requirements.
 - Ensure that the replication approach provides adequate protections against unauthorized access (e.g., encryption in motion).
- CDP security
 - Ensure that the CDP approach (e.g., continuous, near continuous, fixed interval, etc.), especially for business/mission critical data, is aligned with its associated restore strategy
 - In high network bandwidth scenarios (e.g., multimedia files), employ throttling techniques which prioritize network traffic in order to reduce the impact of CDP on day-to-day operation
 - Ensure that the CDP approach provides adequate protections against unauthorized access (e.g., in motion and at rest encryption)

7.3.4 Disaster recovery and business continuity

ISO/PAS 22399:2007 summarizes the Business Continuity Management (BCM) approach to preventing, reacting and recovering from incidents. Activities involved in BCM include Incident Preparedness, Operational Continuity Management (IPOCM), Disaster Recovery Planning (DRP) and risk mitigation which focus on increasing the resilience of the organization and by preparing it to react effectively to incidents and recover within pre-determined timescales.

ISO/IEC 27031:2011 describes the concepts and principles of ICT readiness for business continuity, and provides a framework of methods and processes to identify and specify all aspects (such as performance criteria, design, and implementation) for improving an organization's ICT readiness to ensure business continuity. It applies to any organization (private, governmental, and non-governmental, irrespective of size) developing its ICT Readiness for Business Continuity (IRBC) program, and requiring its ICT services/infrastructures to be ready to support business operations in the event of emerging events and incidents, and related disruptions, that could affect continuity (including security) of critical business functions. It also enables an organization to measure performance parameters that correlate to its IRBC in a consistent and recognized manner.

Storage is typically a critical element of an organization's IRBC program or informal DR/BC activities.

- Ensure the storage ecosystem is factored into the DR/BC planning and implementation

- Prepare for limited disruption events (system failures, adversarial attacks, operator errors)

- Identify and document the unique staffing and facility requirements associated with the storage ecosystem

- Perform on-going planning and regular testing of assumption, which are critical to successful DR/BC; results of DR/BC testing should be fed back into on-going maintenance of the DR/BC plan

7.3.5 Resilience

Resilience is the ability to provide and maintain an acceptable level of service, typically associated with preserving data integrity and availability, in the face of faults (system failures) and challenges (such as attacks, accidents or large-scale natural disasters) to normal operation. This ability is frequently a significant consideration in the deployment of storage systems and infrastructure because of its impact on the overall availability of data.

When considering resilience, failure of individual components may be acceptable, but constitute an incident, as long as the service is still being delivered and the integrity of that service is still there. In practical terms, resilience is a design strategy that aims to reduce vulnerabilities, often by shortening supply lines, improving redundancy in critical areas, bolstering local capacity, and solving for a deeper pattern of dependence and disability.

- Security should be an integral part of the resilience strategy; plan for unit failures and compromises of both the storage and security technologies

- Redundancy should be exploited to the extent possible

- Diverse components that are easily repairable should be used whenever possible

- Security features and functionality (e.g., encryption, centralized authentication, etc.) should be implemented in such a way as to cause no adverse impact to the resilience of the storage system or infrastructure

7.4 Data retention

7.4.1 Long-term retention

Due to the rather short lifetime and limited reliability of traditional storage components, data begins to degrade as soon as it is placed on media. This issue is relatively well understood by those who are involved in the long-term retention of data (e.g., managing data archives) and it is addressed in the following standards, which are applicable to storage infrastructure:

- ISO/TR 10255:2009, *Document management applications – Optical disk storage technology, management and standards*

- 1 — ISO/TR 18492:2005, *Long-term preservation of electronic document-based information*
 - 2 — ISO 16175-1:2010, *Information and documentation – Principles and functional requirements for records in*
3 *electronic office environments – Part 1: Overview and statement of principles*
 - 4 — ISO 16175-2:2011, *Information and documentation – Principles and functional requirements for records in*
5 *electronic office environments – Part 2: Guidelines and functional requirements for digital records*
6 *management systems*
 - 7 — ISO 16175-3:2010, *Information and documentation – Principles and functional requirements for records in*
8 *electronic office environments – Part 3: Guidelines and functional requirements for records in business*
9 *systems*
- 10 Long-term archival storage systems introduce integrity, authentication and privacy threats that do not generally
11 exist in non-archival storage systems. In addition, the long lifetime of data gives attackers a much larger window
12 within which they can attempt to compromise a security system; with archival storage an assailant might have
13 several decades of time to conduct an attack (slow attack).
- 14 — Archival storage assumes a write-once, read-maybe access pattern, thus the integrity of the data in the
15 system should be actively checked at regular intervals rather than waiting to when it is read
 - 16 — Leverage data migrations (e.g., migrating data from outdated components) to effect greater security
 - 17 — Since the data in a long-term archive can out-live the data owners, a secure, archival storage system should
18 be able to authenticate new users and establish their relationship to resources attached to existing users
 - 19 — Secrecy mechanisms (e.g., encryption, secret-sharing, etc.) should function in the complete absence of the
20 user that wrote the data (e.g., a new user who is given rights to read data should also be given the ability to
21 decrypt the data)
 - 22 — Security logging should be sufficiently complete and long-lived (measured in decades) that it assists in
23 detecting slow attacks and maintains an attack history that can be used to make decisions to adjust the data
24 protections
 - 25 — The system should either immediately deal with any compromise or maintain a history of compromises in
26 order to intelligently schedule corrective action
 - 27 — The use of data reduction technologies (e.g., compression and deduplication) should be used carefully
28 because they can compromise data integrity if they are not factored into copies

29 **7.4.2 Short to medium-term retention**

30 A large number of organizations are forced to retain data for periods of time that are shorter than traditional
31 archives (less than 10 years). Often, the retention drivers are based on legal, regulatory, and/or statutory
32 requirements that also include security provisions. Failure to meet the requirements can result in significant
33 liabilities for the organization.

34 To assure successful retention of digital information over short to mid-term retention periods, requires utilization of
35 data protection, Disaster Recovery, and digital preservation and curation practices commensurate with the value
36 of the information being retained, the risk of loss from all factors, and the acceptable amount of loss over the
37 retention period. From a storage perspective, these short and medium-term data retention scenarios usually span
38 one or more generations of technology and require the capture and retention of associated metadata. The
39 following should be considered for short and medium-term retention:

- 1 — Multiple physical and/or logical replicas of the data should be created and preserved;¹⁸⁾ the replicas need to
2 be organized to be as independent as possible (e.g., geographic, administrative/management, and
3 platform/OS), and their number chosen according to the data's value and tolerance of risk
- 4 — On a defined schedule, audit for both obvious and latent faults (e.g., integrity checks), and the damage they
5 cause; repair the corrupted data using the good data from other replicas before that damage spreads
- 6 — Match the access control scheme to the legal and regulatory requirements for the information being
7 preserved
- 8 — Ensure that accountability and traceability measures are adequate and functional; all data accesses may
9 require audit log entries
- 10 — Implement mechanism to demonstrate data authenticity, provenance, and chain of custody, especially for
11 data of an evidentiary nature
- 12 — If encryption is used, archive/escrow the keys and keying material; rekey the data within recommended
13 cryptoperiods or when the underlying cryptographic algorithm needs to be replaced

14 **7.5 Data confidentiality and integrity**

15 There are multiple considerations that need to be made when evaluating the deployment of a storage-based
16 encryption solution, including, but not limited to:

- 17 — Encryption has the potential of impacting other security aspects (e.g., inspection of data, anti-virus, etc.)
- 18 — Although necessary, encryption carries the risk of making data unavailable should anything go wrong with
19 data handling, data transformations, key management, or the actual encryption
- 20 — Encryption can impose significant overhead cost/impacts on hosts and storage elements
- 21 — Centralized key management may be necessary especially when encryption is used in conjunction with out of
22 region replication for DR and BC purposes
- 23 — Encryption can diminish or negate the benefit of data reduction technologies (e.g., compression and
24 deduplication)
- 25 — The quality of the cryptography (security strength, vetted, etc.) can impact that actual protection offered

26 Not all data is worth encrypting. A risk assessment can help identify sensitive and high-value data that warrant the
27 use of encryption as well as assist with the cost benefit analysis (i.e., is the risk reduction worth the cost). It is
28 important to note that there are other vehicles to safeguard the confidentiality of information when the data is
29 considered a critical asset.

30 As mentioned in 6.8.2.3, the point of encryption is important because it represents the location within the IT
31 infrastructure in which the data has to traverse before it is decrypted and usable. A common security perspective
32 is to encrypt as close to the source as possible, as this tends to maximize the protection provided, but there may
33 be many options in selecting a point of encryption (see Figure 3), including:

18) It is not at all about how many copies, rather about the quality and characteristics of the digital archive process.

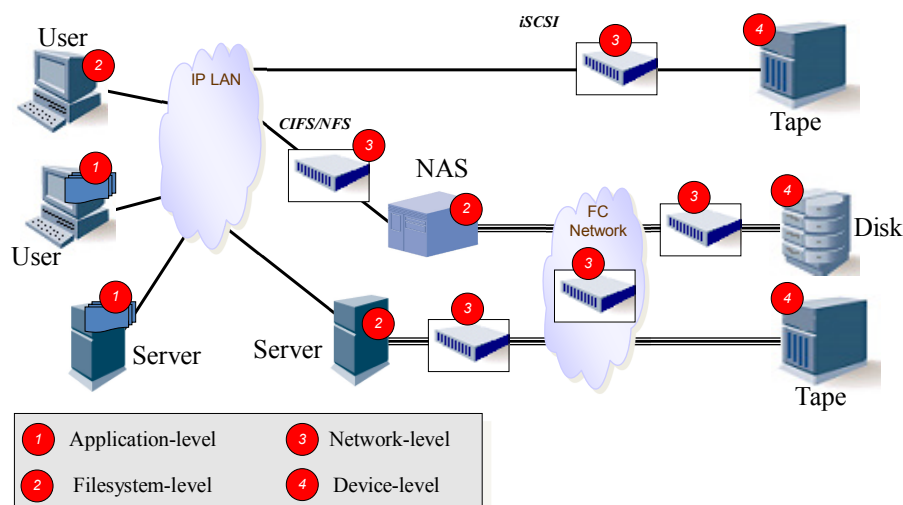


Figure 3 — Sample Points of Encryption

- **Application-level** – under the control of a 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
 - **HBA-, Array Controller-, or Switch-level** – under the control of the network
 - 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”¹⁹⁾
 - **Device-level** – under the control of the end-device (e.g., tape drive, disk array, disk drive, etc.); control at the media level (and possibly at the logical volume level) with limited insights in the “community of users”
- When encryption is deemed necessary, consider the following guidance
- Storage-based encryption should not be the primary form of encryption for sensitive data²⁰⁾
 - Selection of a point of encryption should be influenced by DR and BC (see 7.3.4), data reduction (see 6.8.3), and data protection (see 7.3.3) considerations
 - Data retention (see 7.4) needs should be considered when selecting and deploying encryption
 - The security strength of the encryption solution should be at least 112 bits

¹⁹⁾ The specific user community is unknown, as are their individual access rights. The community is defined by the hosts/servers that have access to the individual logical volumes.

²⁰⁾ The storage encryption is active only while the data is resident on the storage system or media (i.e., it is plaintext once it passes through the point of encryption, which occurs any time the data is accessed).

- 1 — Cryptographic modules used to protect sensitive and/or regulated data should be validated using recognized
2 criteria (e.g., ISO/IEC 19790, ISO/IEC 15408, NIST FIPS 140-2, etc.)
- 3 — Multiple encryption steps can be used, as when data encrypted for privacy purposes is further encrypted by a
4 Self-Encrypting Drive for security purposes
- 5 As with sanitization, it is important that an organization maintain records of its data at rest encryption to document
6 what media were protected as well as when and how they were encrypted. When an organization is suspected of
7 losing control of its storage media, which contain sensitive information, these records or proof of encryption can
8 be instrumental in demonstrating that no data breach occurred, thereby avoiding costly data breach notifications
9 and other liabilities. The following should be considered for proof of encryption:
- 10 — Ensure the encryption mechanisms create appropriate audit log entries (activation, verification, integrity
11 checks, re-keying, etc.)
- 12 — Agree in advance on what audit log material demonstrates (to the satisfaction of the legal department) that
13 encryption was properly performed
- 14 — Perform regular and audited checks that encryption was properly performed and consider outside
15 accreditation
- 16 Successful use of cryptography is dependent on adhering to basic principles associated with keying material as
17 well as implementing key management. As storage systems and devices integrate encryption for data at rest, key
18 management becomes important and should address the following:
- 19 — Fully automate key management whenever possible
- 20 — Sparsely use keys with a long life (i.e., approaches the maximum recommended cryptoperiod, which is
21 typically no more than 1-2 years, depending on the key type)
- 22 — Enforce strict access controls to limit user capabilities and separation of duties constraints (e.g., a security
23 role) for key generation, change and distribution
- 24 — For sensitive and/or high-value data, the encryption should be end-to-end (i.e., in motion and at rest)
- 25 Data integrity is a significant design criterion for most storage systems and infrastructure and it is only rivalled by
26 data availability in its importance to storage personnel. To address data integrity issues, a wide range of
27 technologies are typically deployed in storage infrastructure, including but not limited to, RAID, backups,
28 replications, and CDP. Although important, these data protection technologies are not typically considered part of
29 the storage security controls.
- 30 Data retention and compliance requirements often include provisions for storing data in a manner that blocks
31 record deletion or alteration (i.e., immutable) along with integrity verification (e.g., hashing) and explicit retention
32 periods (e.g., legal holds) that need to be honoured. Several forms of WORM-based storage can be used to meet
33 the immutability (non-editable) requirements. In addition, many CAS (see 6.7.3) implementations combine WORM
34 with metadata that can be used to perform explicit integrity checks as well as enforce data expirations.
- 35 — Malware is a common threat to the integrity of data, applications, and operating system; storage systems
36 should include sufficient malware protections to guard against attacks on data (e.g., corruption, destruction,
37 etc.)
- 38 — WORM-based storage should be used to help meet immutability requirements

7.6 Virtualization

7.6.1 Storage virtualization

Storage virtualization disconnects the logical storage abstractions used by servers and applications from the physical storage systems, devices and/or media on which the information is stored in a fashion that enables that logical to physical relationship to change over time and can mask the details of the physical entities. For example, a logical volume manager in a server or storage array can present portions of multiple physical disk drives as a single mirrored logical volume and be capable of rebuilding the mirrored volume to use another disk drive after a failure of one of the original drives. Another example is that automatic tiering functionality in a storage array can change the drives on which information is stored in response to changed access patterns (e.g., move more frequently accessed information to higher performance drives).

The presence of storage virtualization is an important consideration in control design and application. Controls can be applied to logical and/or physical storage entities. Controls on logical storage entities are unaffected by physical relocation of the information, but controls on physical entities should be applied to the entire domain of physical entities (e.g., storage systems, devices, media) on which information subject to the control may be stored in order to avoid relocation of that information causing the control to be bypassed.

When storage virtualization can store or relocate information across a domain of distributed entities (e.g., information stored on one of multiple storage systems and relocated over time) and storage networking is in use, the appropriate storage networking controls (see 6.3) should be applied to that entire domain, as application of such a control to a subset of the domain can cause the control to be bypassed when information is relocated or new information subject to the control is stored on an entity to which the control has not been applied.

If storage virtualization exposes the physical storage entities that are virtualized (e.g., external storage virtualized by a storage array as shown in Figure 1 in 6.3.2.1) controls should be applied to limit or prevent direct access to the un-virtualized physical elements, as such access is not equivalent to accessing the virtualized storage.

When storage is virtualized, both data sanitization controls (see 6.8.1) and data at rest encryption controls (see 6.8.2) on the physical storage entities should assume that the controlled storage entities (e.g., systems, devices and media) can contain the most sensitive information that may be stored on them. For example, if encryption is used to control the confidentiality of data stored on a disk drive that is removed from a storage array (e.g., because the drive has failed) and that storage array implements storage virtualization, then the encryption algorithm should be appropriate for protection of the most sensitive data that can be stored by the storage array.

Additional virtualization considerations include:

- Ensure appropriate service level objectives for virtual storage
- Match the availability objective for the storage infrastructure to the application requirements
- Match the confidentiality and privacy requirements for the storage infrastructure to the types of information stored
- Address multi-tenancy concerns, as appropriate (see 7.7.4)

7.6.2 Storage for virtualized systems

Server virtualization extends the shared access to resources of typical operating systems to a model in which the virtualization software instead provides the illusion of more than one computer, HDD, printer, etc. The physical server typically runs a hypervisor which is tasked with creating, releasing, and managing the resources of "guest" operating systems, or Virtual Machines (VM). These guest operating systems are allocated a share of resources

of the physical server, typically in a manner in which the guest is not aware of any other physical resources save for those allocated to it by the hypervisor.

When storage systems and infrastructure are used to support virtualized servers, additional care is often necessary to ensure data is available, but not unduly exposed to potential data breaches.

- VM access to storage networks should be controlled via use of access controls in the server virtualization (hypervisor) software

- N_Port ID Virtualization (NPIV) should be leveraged appropriately to limit VM access to storage targets (see C.6 for additional information on NPIV)

- Configure FC SAN zones and present LUNs using the VM-specific WWPNs, so that the LUNs will only be visible to that virtual host and not to any other virtual host

- Avoid scaling problems due to resource limitations (e.g., state related information in servers, network fabrics, and storage) by restricting use of NPIV to creating only the N_Port_IDs that are necessary to provide isolation among larger domains (e.g., the set of VMs for a single organization or a single tenant of a service provider)

- VM migration/movement between physical hosts in an infrastructure should be carefully controlled to avoid having unintended security consequences

- Moving a VM from a lower-risk (more trusted) to a higher-risk (less trusted) domain can expose the sensitive information the server contains or allowed to process unless its configuration is hardened appropriately

- Conversely when a VM is moved from a higher-risk (less trusted) domain to a lower-risk (more trusted) domain, its hardened configuration can interfere with normal operation unless it is matched to that appropriate for the lower-security domain

- VM could move to a compromised virtualized servers thereby putting the data at risk

7.7 Design and implementation considerations

7.7.1 Encryption and key management issues

The use of cryptographic technology introduces certain challenges that cannot be ignored. These challenges can include strict regulations governing the import/export of the technology as well as causing catastrophic losses under certain failure conditions.

- Comply with Import/Export Controls

- Understand and obey government import regulations associated with encryption and key management

- Understand and obey government export regulations associated with encryption and key management

- Comply with corporate and/or government key escrow requirements

- Understand and obey any corporate or government requirements for making encryption keys available to corporate officials, law enforcement authorities, etc. to enable access to and recovery of encrypted data.

- Plan for problems

- 1 — Have a recovery plan in the event of a key compromise
- 2 — Have a key backup²¹⁾ plan²²⁾ in place to ensure continued access to encrypted business/mission critical
- 3 information²³⁾
- 4 — Other problem areas
- 5 — Securely distribute key material among storage devices that process/access the same data. For example
- 6 data is encrypted at one node but decrypted at a second node
- 7 — The effect of encryption on deduplication and compression techniques should be understood and
- 8 factored in designs and implementations
- 9 — The inability to apply security techniques like virus scanning, etc. on encrypted data should be
- 10 understood and mitigated with other mechanisms

11 **7.7.2 Align storage and policy**

12 The presence or absence of policy plays a major role in assuring both security and compliance.

- 13 — Incorporate storage in policies
- 14 — Identify most sensitive (Personally Identifiable Information, intellectual property, trade secrets, etc.) and
- 15 business/mission critical data categories as well as protection requirements
- 16 — Integrate storage-specific policies with other policies (i.e., avoid creating a separate policy document for
- 17 the storage ecosystem)
- 18 — Address data retention and protection (e.g., write-once-read-many or WORM, authenticity, access
- 19 controls, etc.)
- 20 — Address data destruction and media sanitization
- 21 — Conformance with policies
- 22 — Ensure that all elements of the storage ecosystem comply with policy (e.g., ISO/IEC 27001/27002/)
- 23 — Give most sensitive/most critical data a priority

21) Key backup is different from key escrow. Key backup is normally implemented in the context of a specific encryption/key management solution and is focused on providing the solution users (human or machine) access to the keys used to encrypt data within the solution. Key escrow can be implemented separate from an encryption/key management solution and is focused on providing third-party access (e.g. an entity who is not a user of the solution) to the keys used to encrypt data within the solution.

22) Key backup plans can take a number of forms from a simple physical copy of key material to sophisticated key management infrastructures which are designed with high availability and Disaster Recovery in mind.

23) The loss of an encryption key with no key recovery capability (backups, escrow, etc.) renders all of the corresponding ciphertext (i.e., data encrypted under the lost key) unusable. This situation and risk will persist for as long as the data is stored as ciphertext.

7.7.3 Compliance

Complying with legal and regulatory requirements has become an important issue world-wide and this compliance is driving a significant portion of the security agenda and strategy of many organizations. The following elements are key compliance aspects of storage systems and infrastructure that are of concern to an information systems (IS) auditor.

— Accountability

— Ensure that users, especially privileged users, have unique userids (i.e., no shared accounts)

— When possible, grant rights and privileges based on roles

— Log all attempted (successful and unsuccessful) management events and transactions

— Traceability

— Ensure logged event/transaction data contains sufficient application and/or system detail to clearly identify the source

— Ensure that the user information can be traced to a specific individual

— When appropriate, treat log records as evidence²⁴⁾ (chain of custody, non-repudiation, authenticity, etc.)

— Detect, Monitor, and Evaluate

— Ensure that the storage layer participates in the external audit logging measures

— Monitor the audit logging events and issue the appropriate alerts

— Information retention and sanitization

— Implement appropriate data retention measures

— Implement appropriate data integrity and authenticity measures

— Correctly sanitize data upon deletion, repurposing or decommissioning of hardware

— Correctly sanitize virtual server images, and their copies, at end of life

— Privacy

— Implement appropriate data access control measures to control access to data and metadata (e.g., search results); assume a least privilege posture whenever possible

— Implement appropriate data confidentiality measures to prevent unauthorized disclosure

²⁴⁾ ISO/IEC 27037:2012, *Information technology — Security techniques — Guidelines for identification, collection, acquisition, and preservation of digital evidence* provides information and guidance that may be relevant when log records may serve in an evidentiary role.

- 1 — Legal
- 2 — Ensure that the use of data deduplication does not conflict with data authenticity requirements
- 3 — Ensure data and media sanitization mechanisms do not violate preservation orders
- 4 — Ensure proper chain of custody procedures are followed when evidentiary data (e.g., audit logs,
5 metadata, mirror images, point-in time copies, etc.) is handled

6 NOTE Annex B can be a useful resource when auditing storage systems and infrastructure.

7 **7.7.4 Secure multi-tenancy**

8 Multi-tenancy, as defined by ISO/IEC 17788, focuses on the "allocation of physical and virtual resources such that
9 multiple tenants and their computations and data are isolated from and inaccessible to one another." Secure
10 multi-tenancy builds on this concept by adding security controls to explicitly guard against data breaches as well
11 as to allow for verification of the state of these controls (e.g., they are active) and validation of the controls (i.e.
12 assurance that they work).

13 When considering secure multi-tenancy, it is important to include the perspective of the tenants (including their
14 administrators). As such, a secure multi-tenant solution needs the capability to provide secure isolation while still
15 delivering the management and flexibility benefits of shared resources that assures:

- 16 — no tenant can determine the existence or identity of any other tenant
- 17 — no tenant can access the data in motion (network) of any other tenant
- 18 — no tenant can access the data at rest (storage) of any other tenant
- 19 — no tenant can perform an operation that affects an operation performed by another tenant
- 20 — no tenant can perform an operation that might deny service to another tenant
- 21 — each tenant can have a configuration that is independent of other tenant's existence and configuration (For
22 example in naming or addressing.)
- 23 — when a resource (compute, storage or network) is decommissioned from a tenant the resources should be
24 sanitized of all data and configuration information
- 25 — accountability and traceability measures are available at the tenant level

26 Within storage systems and infrastructure that are used in part or in whole for secure multi-tenancy solutions, the
27 following additional security measures should be used:

- 28 — Encrypted storage that is aligned with the tenants' usage of resources
- 29 — Strong at rest encryption (minimum of 112-bits of security strength)
- 30 — Secure and rapid de-provisioning (see Annex A for media sanitization, including cryptographic erase)
- 31 — Trusted third-party data storage management (e.g., SNMPv3, SMI-S with TLS, etc.)
- 32 — Automated key management providing tenant-controlled key management (leverages KMIP v1.1 compliant
33 servers)

- 1 — Secure data replication (e.g., data in motion and at rest encryption)
- 2 — Protect data from administrators (e.g., enforce a least privileges access model, administrators do not have
- 3 access to the keying materials, etc.)
- 4 — Highly available data fabrics (multi-path and diverse path)
- 5 — Centralized and secure audit logging (e.g., syslog over TLS)
- 6 — Validation and certification (e.g., Common Criteria) of cryptographic modules and other security measures
- 7 (e.g., media sanitization, access control, etc.)

8 **7.7.5 Secure autonomous data movement**

9 Many storage systems and infrastructure have the ability to move data between different storage devices and
 10 storage elements (e.g., tiered storage), between data centres (e.g., synchronous and asynchronous data
 11 replication), to data archiving facilities, to data protection systems (e.g., backups on tape robots or virtual tape),
 12 etc. More complex scenarios exist within Information Lifecycle Management (ILM) and Data Lifecycle
 13 Management (DLM) solutions. However, all of these scenarios assume:

- 14 — data movement is policy driven
- 15 — intervention of operators or host computers is not required to initiate or intervene throughout the process.
- 16 Because autonomous data movement takes many forms, the security needs can vary significantly; they can
- 17 include some or all of the following:
- 18 — Accountability and traceability
 - 19 — Configuring policies for data movement should be restricted to authenticated and authorized privileged
 - 20 users
 - 21 — The individual establishing the configurations should be conversant with the security attributes of both
 - 22 source and destination
 - 23 — Configuration changes to implement or terminate autonomous data movement should be reflected in the
 - 24 audit log
 - 25 — All autonomous data movement transactions should be reflected in the audit log of the system
 - 26 conducting the data movement
- 27 — Integrity, authenticity, and immutability
 - 28 — As part of autonomous data movement transactions, the integrity of the moved data should be verified
 - 29 (preferably with a cryptographic hash)
 - 30 — Autonomous data movement transactions should not impact the authenticity of the data (e.g., original
 - 31 system metadata like creation date, last accessed, etc. are correctly represented in the moved data)
 - 32 — Autonomous data movement transactions should not negate the immutability or other data preservation
 - 33 controls (e.g., supporting legal holds)

34

- 1 — Confidentiality
- 2 — Autonomous data movement transactions should not eliminate or weaken encryption controls associated
- 3 with the data
- 4 — Autonomous data movement transactions that span systems should include data in motion encryption for
- 5 sensitive and high value data
- 6 — Sanitization
- 7 — As part of autonomous data movement transactions, the source data or storage media should be
- 8 appropriately sanitized (see 6.8.1.2 and 6.8.1.3) before it is released for re-use
- 9 — Sanitization performed in conjunction with autonomous data movement should also include verification
- 10 (see 6.8.1.5) and some form of proof of sanitization (see 6.8.1.4)
- 11 — Trustworthiness and physical security
- 12 — Autonomous data movement transactions should not cause data to cross security domains (e.g.,
- 13 production to development environments)
- 14 — Autonomous data movement transactions should not cause data to move to systems with inadequate
- 15 certifications and accreditations
- 16 — Autonomous data movement transactions should not cause data to move to systems with inadequate
- 17 physical security

Annex A (normative)

Media sanitization

A.1 Methods used to sanitize media

Several different methods can be used to sanitize media with the three most common being:

- **Clear** - One method to sanitize media is to use software or hardware products to overwrite storage space on the media with non-sensitive data. This process can include overwriting not only the logical storage location of a file(s) (e.g., file allocation table) but also can include all addressable locations. The security goal of the overwriting process is to replace written data with random data. Overwriting cannot be used for media that are damaged or not rewriteable. The media type and size can also influence whether overwriting is a suitable sanitization method

- **Purge** - Degaussing, cryptographic erase (see A.3), and executing the appropriate ATA/SCSI firmware commands are acceptable methods for purging. Degaussing is not applicable to non-magnetic media (e.g. SSD or SSHD).

Degaussing is exposing the magnetic media to a strong magnetic field in order to disrupt the recorded magnetic domains. A degausser is a device that generates a magnetic field used to sanitize magnetic media. Degaussers are rated based on the type (i.e., low energy or high energy) of magnetic media they can purge. Degaussers operate using either a strong permanent magnet or an electromagnetic coil. Degaussing can be an effective method for purging damaged or inoperative media, for purging media with exceptionally large storage capacities, or for quickly purging diskettes.

- **Destroy** - There are many different types, techniques, and procedures for media destruction. If destruction is decided on because of the high security categorization of the information, then after the destruction, the media should be able to withstand a laboratory attack.

- *Disintegration, Pulverization, Melting, and Incineration.* These sanitization methods are designed to completely destroy the media.

- *Shredding.* Paper shredders can be used to destroy flexible media such as diskettes once the media are physically removed from their outer containers. The shred size of the refuse should be small enough that there is reasonable assurance in proportion to the data confidentiality that the data cannot be reconstructed.

A.2 Sanitization for different types of media

There are two primary types of media in common use:

- **Hard Copy.** Hard copy media is physical representations of information, most often associated with paper printouts. However, printer and facsimile ribbons, drums, and platens are all examples of hard copy media. The supplies associated with producing paper printouts are often the most uncontrolled. Hard copy materials containing sensitive data that leave an organization without effective sanitization expose a significant

vulnerability to “dumpster divers” and overcurious employees, risking accidental disclosures. Table A.1 provides guidance for this type of media.

— **Electronic (or soft copy).** Electronic media are the devices containing bits and bytes such as HDD, random access memory (RAM), read-only memory (ROM), disks, memory devices, phones, mobile computing devices, networking devices, office equipment, and many other types. Tables A.2, A.3, A.4, A.5, A.6, A.7, A.8, and A.9 provide guidance for common forms of electronic media.

Table A.1 — Hard Copy Storage Sanitization

| Hard Copy Storage | |
|--------------------------|--|
| Paper and microforms | |
| Clear/ Purge: | N/A, see Destroy. |
| Destroy: | <p>Destruct paper using cross cut shredders which produce particles that are 1 x 5 millimetres in size (or smaller), or pulverize/disintegrate paper materials using disintegrator devices equipped with 3/32 inch security screen.</p> <p>Destruct microforms (microfilm, microfiche, or other reduced image photo negatives) by burning.</p> |
| Notes: | When material is burned, the residue is reduced to white ash. |

Table A.2 — Networking Device Sanitization

| Networking Devices | |
|--|--|
| Routers and Switches (home, home office, enterprise) | |
| Clear: | Perform a full manufacturer's reset to reset the router or switch back to its factory default settings. |
| Purge: | See Destroy. Most routers and switches only offer capabilities to Clear (and not Purge) the data contents. A router or switch may offer Purge capabilities, but these capabilities are specific to the hardware and firmware of the device and should be applied with caution. Refer to the device manufacturer to identify whether the device has a Purge capability that applies media-dependent techniques (such as rewriting or block erasing) to ensure that data recovery is infeasible, and that the device does not simply remove the file pointers. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>For both Clear and (if applicable) Purge, refer to the manufacturer for additional information on the proper Sanitization procedure.</p> <p>Network Devices may contain removable storage. The removable media should be removed and sanitized using media-specific techniques.</p> |

Table A.3 — Mobile Device Sanitization

| Mobile Devices ²⁵⁾ | |
|---------------------------------------|---|
| Apple iPhone and iPad | |
| Clear/ Purge: | Select the full sanitize option (typically in the 'Settings > General > Reset > Erase All Content and Settings' menu). The sanitization operation may take only minutes if cryptographic erase is supported, or may take as long as several hours if media-dependent non-cryptographic sanitization techniques that leverage overwriting are applied by the device (depending on the media size). |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>Following the Clear/Purge operation, manually navigate to multiple areas of the device (such as browser history, files, photos, etc.) to verify that no personal information has been retained on the device.</p> <p>Refer to the manufacturer for proper sanitization procedure, and for details about implementation differences between device versions and OS versions. Proper initial configuration using guides helps ensure that the level of data protection and sanitization assurance is as robust as possible.</p> |
| Blackberry | |
| Clear/ Purge: | Select the full sanitize option (typically in either the ' <i>Options > Security Options > General Settings > [menu button] > Wipe Handheld</i> ' OR in ' <i>Options > Security Options > Security Wipe</i> ' menu), making sure to select all subcategories of data types for sanitization. The sanitization operation may take as long as several hours depending on the media size. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>Following the Clear/Purge operation, manually navigate to multiple areas of the device (such as browser history, files, photos, etc.) to verify that no personal information has been retained on the device.</p> <p>Refer to the manufacturer for additional information on the proper sanitization procedure, and for details about implementation differences between device versions and OS versions. Proper initial configuration using guides helps ensure that the level of data protection and sanitization assurance is as robust as possible. If the device contains removable storage media, ensure that the media is sanitized using appropriate media-dependent procedures.</p> |
| Devices running the Google Android OS | |
| Clear: | Select the full sanitize option (typically in the ' <i>Menu > Settings > [Privacy OR SD and Phone Storage] > Factory data reset</i> ' menu). |

²⁵⁾ Disassembly of battery and display may be required.

| | |
|--|--|
| Purge: | Android settings and capabilities may be modified by device vendors or service providers, and therefore no assumptions should be made about the level of assurance provided by performing a factory data reset. Some versions of Android support encryption, and may support cryptographic erase. Refer to the device manufacturer (and potentially the service provider as well, if applicable) to identify whether the device has a Purge capability that applies media-dependent techniques (such as rewriting or block erasing) or cryptographic erase to ensure that data recovery is infeasible, and that the device does not simply remove the file pointers. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>Proper initial configuration helps ensure that the level of data protection and sanitization assurance is as robust as possible. Following the Clear or (if applicable) Purge operation, manually navigate to multiple areas of the device (such as browser history, files, photos, etc.) to verify that no personal information has been retained on the device.</p> <p>For both Clear and (if applicable) Purge, refer to the manufacturer for additional information on the proper sanitization procedure.</p> |
| All other mobile devices <i>This includes cell phones, smart phones, PDAs, tablets, and other devices not covered in the preceding mobile categories.</i> | |
| Clear: | Manually delete all information, then perform a full manufacturer's reset to reset the mobile device to factory state. |
| Purge: | See Destroy. Many mobile devices only offer capabilities to Clear (and not Purge) the data contents. A mobile device may offer Purge capabilities, but these capabilities are specific to the hardware and software of the device and should be applied with caution. The device manufacturer should be referred to in order to identify whether the device has a Purge capability that applies media-dependent techniques (such as rewriting or block erasing) or cryptographic erase to ensure that data recovery is infeasible, and that the device does not simply remove the file pointers. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>Following the Clear or (if applicable) Purge operation, manually navigate to multiple areas of the device (such as call history, browser history, files, photos, etc.) to verify that no personal information has been retained on the device.</p> <p>For both Clear and (if applicable) Purge, refer to the manufacturer for proper sanitization procedure.</p> |

Table A.4 — Equipment Sanitization

| Equipment | |
|---|---|
| Office Equipment <i>This includes copy, print, fax, and multifunction machines</i> | |
| Clear: | Perform a full manufacturer's reset to reset the office equipment to its factory default settings. |
| Purge: | See Destroy. Most office equipment only offers capabilities to Clear (and not Purge) the data contents. Office equipment may offer Purge capabilities, but these capabilities are specific to the hardware and firmware of the device and should be applied with caution. Refer to the device manufacturer to identify whether the device has a Purge capability that applies media-dependent techniques (such as |

| | |
|-----------------|--|
| | rewriting or block erasing) or cryptographic erase to ensure that data recovery is infeasible, and that the device does not simply remove the file pointers. Office equipment may have removable storage media, and if so, media-dependent sanitization techniques may be applied to the associated storage device. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>For both Clear and (if applicable) Purge, manually navigate to multiple areas of the device (such as stored fax numbers, network configuration information, etc.) to verify that no personal information has been retained on the device.</p> <p>For both Clearing and (if applicable) Purge, the ink, toner, and associated supplies (drum, fuser, etc.) should be removed and destroyed or disposed of in accordance with applicable law, environmental, and health considerations. Some of these supplies may retain impressions of data printed by the machine and therefore could pose a risk of data exposure, and should be handled accordingly. If the device is functional, one way to reduce the associated risk is to print a blank page, then an all-black page, then another blank page. For devices with dedicated colour components (such as cyan, magenta, and yellow toners and related supplies), one page of each colour should also be printed between blank pages. The resulting sheets should be handled at the confidentiality of the Office Equipment (prior to sanitization). Note that these procedures do not apply to supplies such as ink/toner on a one-time use roll, as they are typically not used again and therefore will not be addressed by sending additional pages through the equipment. Office Equipment supplies may also pose health risks, and should be handled using appropriate procedures to minimize exposure to the print components and toner.</p> <p>For both Clear and (if applicable) Purge, refer to the manufacturer for additional information on the proper sanitization procedure.</p> |

1

2

Table A.5 — Magnetic Media Sanitization

| Magnetic Media | |
|---|--|
| Floppies | |
| Clear: | Overwrite media by using organizationally approved software and validate the overwritten data. The Clear procedure should consist of at least a single pass of writes with a fixed data value, such as all zeros. Multiple passes or more complex values may optionally be used. |
| Purge: | Degauss in an organizationally approved degausser. |
| Destroy: | Incinerate floppy disks and diskettes by burning in a licensed incinerator or Shred. |
| Removable Flexible or Rigid Magnetic Disks <i>This includes Zip, Floptical, Jaz, SyQuest, LS-120, etc.</i> | |
| Clear: | Overwrite media by using organizationally approved software and validate the overwritten data. The Clear procedure should consist of at least a single pass of writes with a fixed data value, such as all zeros. Multiple passes or more complex values may optionally be used. |
| Purge: | Degauss in an organizationally approved degausser. |
| Destroy: | Incinerate disks and diskettes by burning in a licensed incinerator or Shred. |
| Notes: | Degaussing disks typically renders the disk permanently unusable. |

3

| Reel and Cassette Format Magnetic Tapes <i>This also includes 8mm, DDS DAT, DLT, QIC, etc.</i> | |
|---|--|
| Clear: | Re-record (overwrite) all data on the tape using an organizationally approved pattern, using a system with similar characteristics to the one that originally recorded the data. For example, overwrite previously recorded sensitive VHS format video signals on a comparable VHS format recorder. All portions of the magnetic tape should be overwritten one time with known non-sensitive signals. Clearing a magnetic tape by re-recording (overwriting) may be impractical for most applications since the process occupies the tape transport for excessive time periods. |
| Purge: | Degauss the magnetic tape in an organizationally approved degausser. |
| Destroy: | Incinerate by burning the tapes in a licensed incinerator or Shred. |
| Notes: | Preparatory steps for Destruct, such as removing the tape from the reel or cassette prior to Destruction, are unnecessary. However, segregation of components (tape and reels or cassettes) may be necessary to comply with the requirements of a Destruction facility or for recycling measures. |
| ATA HDD <i>This includes PATA, SATA, eSATA, etc.</i> | |
| Clear: | Overwrite media by using organizationally approved and validated overwriting technologies/methods/tools. The Clear procedure should consist of at least a single pass of writes with a fixed data value, such as all zeros. Multiple passes or more complex values may optionally be used. |
| Purge: | <p>Four options are available:</p> <ol style="list-style-type: none"> a) Use one of the ATA Sanitize Device feature set commands, if supported, to perform a Sanitize operation. One or both of the following options may be available: <ol style="list-style-type: none"> 1) The OVERWRITE EXT command. Apply one pass of a fixed pattern across the media surface. Some examples of fixed patterns include all 0s or a pseudorandom pattern. <i>Optionally:</i> Instead of one pass, use three total passes of a pseudorandom pattern, leveraging the invert option so that the second pass is the inverted version of the pattern specified. 2) If the device supports encryption and the technical specifications described in this International Standard have been satisfied, the CRYPTO SCRAMBLE EXT command may be used. <i>Optionally:</i> After cryptographic erase is successfully applied to a device, use the overwrite command (if supported) to write one pass of zeros or a pseudorandom pattern across the media. If the overwrite command is not supported, the ATA Security feature set SECURITY ERASE UNIT command or the Clear procedure could alternatively be applied following cryptographic erase. b) Use the SECURITY ERASE UNIT command in Enhanced Erase mode, if supported. The ATA Sanitize Device feature set commands are preferred over the SECURITY ERASE UNIT command when supported by the ATA device. c) Cryptographic erase through the Trusted Computing Group (TCG) Opal Security Subsystem Class (SSC) or Enterprise SSC interface by issuing commands as necessary to cause all MEKs to be changed (if the technical |

| | |
|-----------------|--|
| | <p>specifications described in this International Standard have been satisfied). Refer to the TCG and device manufacturers for more information. <i>Optionally:</i> After cryptographic erase is successfully applied to a device, use the overwrite command (if supported) to write one pass of zeros or a pseudorandom pattern across the media. If the overwrite command is not supported, the SECURITY ERASE UNIT command or the Clear procedure could alternatively be applied following cryptographic erase.</p> <p>d) Degauss in an organizationally approved automatic degausser or disassemble the hard disk drive and Purge the enclosed platters with an organizationally approved degaussing wand.</p> |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>Performing verification is necessary for each technique within Clear and Purge, except degaussing. The assurance provided by degaussing depends on selecting an effective degausser, applying it appropriately and periodically spot checking the results to ensure it is working as expected.</p> <p>When using the OVERWRITE EXT command with the invert option and an odd number of passes (e.g., three passes), the verification process would simply search for the original pattern (which would have been written again during the third pass).</p> <p>The storage device may support configuration capabilities that artificially restrict the ability to access portions of the media as defined in the ATA standard, such as a Host Protected Area (HPA), Device Configuration Overlay (DCO), or Accessible Max Address. Even when a dedicated sanitization command addresses these areas, their presence may affect the ability to reliably verify the effectiveness of the sanitization procedure if left in place. Any configuration options limiting the ability to access the entire addressable area of the storage media should be reset prior to applying the sanitization technique. Recovery data, such as an OEM-provided restoration image may have been stored in this manner, and sanitization may therefore impact the ability to recover the system unless installation media is also available.</p> <p>When cryptographic erase is applied, performing verification is necessary prior to additional sanitization techniques (if applicable), such as a Clear or Purge technique applied following cryptographic erase, to ensure that the cryptographic operation completed successfully. A quick sampling verification as described in 6.8.1.5 should also be performed after any additional techniques are applied following cryptographic erase.</p> <p>Not all implementations of encryption are necessarily suitable for reliance upon cryptographic erase as a Purge mechanism. The decision regarding whether to use cryptographic erase depends upon verification of attributes previously identified in this guidance and in A.3.</p> <p>Given the variability in implementation of the SECURITY ERASE UNIT command, use of this command is not recommended without first consulting with the manufacturer to confirm that the storage device's model-specific implementation meets the needs of the organization.</p> <p>This guidance applies to magnetic media only, and it is critical to verify the media type prior to sanitization. Note that emerging media types, such as HAMR media or hybrid drives may not be easily identifiable by the label. Refer to the manufacturer for details about the media type in a storage device.</p> <p>Degaussing the media in a storage device typically renders the device unusable.</p> |

| SCSI HDD <i>This includes Parallel SCSI, Serial Attached SCSI (SAS), Fibre Channel, USB Attached Storage, SCSI Express, etc.</i> | |
|---|---|
| Clear: | Overwrite media by using organizationally approved and validated overwriting technologies/methods/tools. The Clear procedure should consist of at least a single pass of writes with a fixed data value, such as all zeros. Multiple passes or more complex values may optionally be used. |
| Purge: | <p>Four options are available:</p> <ol style="list-style-type: none"> Apply the SANITIZE command, if supported. One or both of the following options may be available: <ol style="list-style-type: none"> The OVERWRITE service action. Use three total passes of a pseudorandom pattern, leveraging the invert option so that the second pass is the inverted version of the pattern specified. If the device supports encryption, the CRYPTOGRAPHIC ERASE service action may be used. <i>Optionally:</i> After cryptographic erase is successfully applied to a device, use the OVERWRITE service action (if supported) to write one pass of zeros or a pseudorandom pattern across the media. If the OVERWRITE service action is not supported, the Clear procedure could alternatively be applied. Cryptographic erase through the TCG Opal SSC or Enterprise SSC interface by issuing commands as necessary to cause all MEKs to be changed (partial sanitization is not supported). Refer to the TCG and vendors shipping TCG Opal or Enterprise storage devices for more information. <i>Optionally:</i> After cryptographic erase is successfully applied to a device, use the OVERWRITE service action (if supported) to write one pass of zeros or a pseudorandom pattern across the media. If the OVERWRITE service action is not supported, the Clear procedure could alternatively be applied. If neither of the first two options is supported, use the native read and write interface to write least a single pass with a fixed data value, such as all zeros. Multiple passes or more complex values may alternatively be used. Degauss in an organizationally approved automatic degausser or disassemble the hard disk drive and Purge the enclosed platters with an organizationally approved degaussing wand. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>Performing verification is necessary for each technique within Clear and Purge as described in 6.8.1.5, except degaussing. The assurance provided by degaussing depends on selecting an effective degausser, applying it appropriately and periodically spot checking the results to ensure it is working as expected.</p> <p>When using the SANITIZE command with OVERWRITE service action with three passes and the invert (also known as complement) option, the verification process would simply search for the original pattern (which would have been written again during the third pass). While it is widely accepted that one pass of overwriting should be sufficient for Purging the data, the availability of a dedicated command that incorporates the ability to invert the data pattern allows an efficient and effective approach that mitigates any residual risk associated with variations in implementations</p> |

of magnetic recording features across device manufacturers.

The storage device may support configuration capabilities that artificially restrict the ability to access portions of the media, such as the SCSI mode parameter block descriptor's NUMBER OF LOGICAL BLOCKS field (accessible with the MODE SENSE and MODE SELECT commands). Even when a dedicated sanitization command addresses these areas, their presence may affect the ability to reliably verify the effectiveness of the sanitization procedure if left in place. Any configuration options limiting the ability to access the entire addressable area of the storage media should be reset prior to applying the sanitization technique.

When cryptographic erase is applied, performing verification is necessary prior to additional sanitization techniques (if applicable), such as a Clear or Purge technique applied following cryptographic erase, to ensure that the cryptographic operation completed successfully. A quick sampling verification as described in 6.8.1.5 should also be performed after any additional techniques are applied following cryptographic erase.

Not all implementations of encryption are necessarily suitable for reliance upon cryptographic erase as a purge mechanism. The decision regarding whether to use cryptographic erase depends upon verification of attributes previously identified in this guidance and in A.3.

This guidance applies to magnetic media only, and it is critical to verify the media type prior to sanitization. Note that emerging media types, such as HAMR media or hybrid drives may not be easily identifiable by the label. Refer to the manufacturer for details about the media type in a storage device.

Degaussing the media in a storage device typically renders the device unusable.

Editor's Note: For Table A.6, external USB and Firewire drives use SCSI commands; eSATA uses ATA commands. All the basic SCSI and ATA rules thus apply. Consideration should be given to merging any unique material into table A.5 (HDDs) and A.8 (SSDs).

Table A.6 — Peripherally Attached Storage Sanitization

| Peripherally Attached Storage | |
|---|---|
| External Locally Attached HDD <i>This includes, USB, Firewire, etc. (Treat eSATA as ATA HDD.)</i> | |
| Clear: | Overwrite media by using organizationally approved and validated overwriting technologies/methods/tools. The Clear procedure should consist of at least a single pass of writes with a fixed data value, such as all zeros. Multiple passes or more complex values may optionally be used. |
| Purge: | <p>See Destroy. The implementation of External Locally Attached HDD varies sufficiently across models and vendors that the issuance of any specific command to the device may not reasonably and consistently assure the desired sanitization result.</p> <p>When the external drive bay contains an ATA or SCSI HDD, if the commands can be delivered natively to the device the device may be sanitized based on the associated media-specific guidance. However, the drive could be configured in a vendor-specific manner that precludes sanitization when removed from the enclosure. Additionally, if sanitization techniques are applied, the HDD may not work as expected when reinstalled in the enclosure.</p> <p>Refer to the device manufacturer to identify whether the device has a Purge capability that applies media-dependent techniques (such as rewriting, block erasing,</p> |

| | |
|-----------------|---|
| | cryptographic erase, etc.) to ensure that data recovery is infeasible, and that the device does not simply remove the file pointers. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>Verification as described in 6.8.1.5 should be performed for each technique within Clear and Purge.</p> <p>Some external locally attached HDD, especially those featuring security or encryption features, may also have hidden storage areas that might not be addressed even when the drive is removed from the enclosure. The device vendor may leverage proprietary commands to interact with the security subsystem. Please refer to the manufacturer to identify whether any reserved areas exist on the media and whether any tools are available to remove or sanitize them, if present.</p> |

Table A.7 — Optical Media Sanitization

| Optical Media | |
|--------------------------|---|
| CD, DVD, BD | |
| Clear/ Purge: | N/A, see Destroy. |
| Destroy: | <p>Destroy in order of recommendations:</p> <ul style="list-style-type: none"> a) Removing the information-bearing layers of CD media using a commercial optical disk grinding device. Note that this applies only to CD and not to DVD or BD media b) Incinerate optical disk media (reduce to ash) using a licensed facility. c) Use optical disk media shredders or disintegrator devices to reduce to particles that have a nominal edge dimensions of point five millimetres (.5 mm) and surface area of point two five square millimetres (.25 mm²) or smaller. |

Table A.8 — Flash-Based Storage Device Sanitization

| Flash-Based Storage Devices | |
|--|---|
| ATA SSD <i>This includes PATA, SATA, eSATA, etc.</i> | |
| Clear: | <ul style="list-style-type: none"> a) Overwrite media by using organizationally approved and validated overwriting technologies/methods/tools. The Clear procedure should consist of at least a single pass of writes with a fixed data value, such as all zeros. Multiple passes or more complex values may optionally be used. <p>It is important to note that overwrite on flash-based media may significantly reduce the effective lifetime of the media and it may not get rid of data in unmapped physical media (i.e., the old data may well remain).</p> <ul style="list-style-type: none"> b) Use the SECURITY ERASE UNIT command, if supported. |

| | |
|-----------------|--|
| Purge: | <p>Three options are available:</p> <ul style="list-style-type: none"> a) Use one of the ATA Sanitize Device feature set commands, if supported, to perform a Sanitize operation. One or both of the following options may be available: <ul style="list-style-type: none"> 1) BLOCK ERASE EXT command. <i>Optionally:</i> After the BLOCK ERASE EXT is successfully applied to a device, write binary 1s across the user addressable area of the storage media and then perform a second BLOCK ERASE EXT. 2) If the device supports encryption, the CRYPTO SCRAMBLE EXT command may be used. <i>Optionally:</i> After cryptographic erase is successfully applied to a device, use the BLOCK ERASE EXT command (if supported) to block erase the media. If the BLOCK ERASE EXT command is not supported, the ATA Security feature set SECURITY ERASE UNIT command or the Clear procedure could alternatively be applied. b) Use the SECURITY ERASE UNIT command in Enhanced Erase mode, if supported. The ATA Sanitize Device feature set commands are preferred over the SECURITY ERASE UNIT command. c) Cryptographic erase through the TCG Opal SSC or Enterprise SSC interface by issuing commands as necessary to cause all MEKs to be changed. Refer to the TCG and vendors shipping TCG Opal or Enterprise storage devices for more information. <i>Optionally:</i> After cryptographic erase is successfully applied to a device, use the block erase command (if supported) to block erase the media. If the BLOCK ERASE EXT is not supported, the Clear procedure could alternatively be applied. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>Performing verification is necessary for each technique within Clear and Purge as described in 6.8.1.5.</p> <p>When cryptographic erase is applied, performing verification is necessary prior to additional sanitization techniques (if applicable), such as a Clear or Purge technique applied following cryptographic erase, to ensure that the cryptographic operation completed successfully. A quick sampling verification as described in 6.8.1.5 should also be performed after any additional techniques are applied following cryptographic erase.</p> <p>The storage device may support configuration capabilities that artificially restrict the ability to access portions of the media as defined in the ATA standard, such as a Host Protected Area (HPA), Device Configuration Overlay (DCO), or Accessible Max Address. Even when a dedicated sanitization command addresses these areas, their presence may affect the ability to reliably verify the effectiveness of the sanitization procedure if left in place. Any configuration options limiting the ability to access the entire addressable area of the storage media should be reset prior to applying the sanitization technique. Recovery data, such as an OEM-provided restoration image may have been stored in this manner, and sanitization may therefore impact the ability to recover the system unless reinstallation media is also available.</p> <p>Not all implementations of encryption are necessarily suitable for reliance upon cryptographic erase as a Purge mechanism. The decision regarding whether to use</p> |

| | |
|---|--|
| | <p>cryptographic erase depends upon verification of attributes previously identified in this guidance and in A.3.</p> <p>Given the variability in implementation of the SECURITY ERASE UNIT command, use of this command is not recommended without first consulting with the manufacturer to confirm that the storage device's model-specific implementation meets the needs of the organization.</p> <p>Whereas the SECURITY ERASE UNIT command is a Purge mechanism for magnetic media, it is only a Clear mechanism for flash due to variability in implementation and the possibility that sensitive data may remain in areas such as spare cells that have been rotated out of use.</p> <p>Do not rely solely upon Degaussing as a sanitization technique on flash-based storage devices or on hybrid devices that contain non-volatile flash storage media. Degaussing may be used when non-volatile flash media is present if the flash components are sanitized using media-dependent techniques.</p> |
| SCSI SSD <i>This includes Parallel SCSI, Serial Attached SCSI (SAS), Fibre Channel, USB Attached Storage, SCSI Express, etc.</i> | |
| Clear: | <p>Overwrite media by using organizationally approved and validated overwriting technologies/methods/tools. The Clear procedure should consist of at least a single pass of writes with a fixed data value, such as all zeros. Multiple passes or more complex values may optionally be used.</p> <p>It is important to note that overwrite on flash-based media may significantly reduce the effective lifetime of the media and it may not get rid of data in unmapped physical media (i.e., the old data may well remain).</p> |
| Purge: | <p>Two options are available:</p> <ol style="list-style-type: none"> Apply the SANITIZE command, if supported. One or both of the following options may be available: <ol style="list-style-type: none"> The BLOCK ERASE service action. If the device supports encryption, the CRYPTOGRAPHIC ERASE service action may be used. <i>Optionally:</i> After cryptographic erase is successfully applied to a device, use the BLOCK ERASE service action (if supported) to block erase the media. If the BLOCK ERASE service action is not supported, the Clear procedure could alternatively be applied. Cryptographic erase through the TCG Opal SSC or Enterprise SSC interface by issuing commands as necessary to cause all MEKs to be changed. Refer to the TCG and vendors shipping TCG Opal or Enterprise storage devices for more information. <i>Optionally:</i> After cryptographic erase is successfully applied to a device, use the BLOCK ERASE service action (if supported) to block erase the media. If the BLOCK ERASE service action is not supported, the Clear procedure is an acceptable alternative. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>Performing verification is necessary for each technique within Clear and Purge as described in 6.8.1.5.</p> <p>The storage device may support configuration capabilities that artificially restrict the</p> |

| | |
|-------------------------|---|
| | <p>ability to access portions of the media, such as the SCSI mode parameter block descriptor's NUMBER OF LOGICAL BLOCKS field (accessible with the MODE SENSE and MODE SELECT commands). Even when a dedicated sanitization command addresses these areas, their presence may affect the ability to reliably verify the effectiveness of the sanitization procedure if left in place. Any configuration options limiting the ability to access the entire addressable area of the storage media should be reset prior to applying the sanitization technique.</p> <p>When cryptographic erase is applied, performing verification is necessary prior to additional sanitization techniques (if applicable), such as a Clear or Purge technique applied following cryptographic erase, to ensure that the cryptographic operation completed successfully. A quick sampling verification as described in 6.8.1.5 should also be performed after any additional techniques are applied following cryptographic erase.</p> <p>Not all implementations of encryption are necessarily suitable for reliance upon cryptographic erase as a Purge mechanism. The decision regarding whether to use cryptographic erase depends upon verification of attributes previously identified in this guidance and in A.3.</p> <p>Do not rely solely upon Degaussing as a sanitization technique on flash-based storage devices.</p> |
| NVM Express SSDs | |
| Clear: | Overwrite media by using organizationally approved and validated overwriting technologies/methods/tools. The Clear procedure should consist of at least a single pass of writes with a fixed data value, such as all zeros. Multiple passes or more complex values may optionally be used. |
| Purge: | <p>Two options are available:</p> <ol style="list-style-type: none"> Apply the NVM Express Format command, if supported. One or both of the following options may be available: <ol style="list-style-type: none"> The User Data Erase command. If the device supports encryption, the cryptographic erase command. <i>Optionally:</i> After cryptographic erase is successfully applied to a device, use the User Data Erase command (if supported) to erase the media. If the User Data Erase command is not supported, the Clear procedure could alternatively be applied. Cryptographic erase through the TCG Opal SSC or Enterprise SSC interface by issuing commands as necessary to cause all MEKs to be changed. Refer to the TCG and vendors shipping TCG Opal or Enterprise storage devices for more information. <i>Optionally:</i> After cryptographic erase is successfully applied to a device, use the User Data Erase command (if supported) to erase the media. If the User Data Erase command is not supported, the Clear procedure is an acceptable alternative. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | <p>Performing verification is necessary for each technique within Clear and Purge.</p> <p>When cryptographic erase is applied, performing verification is necessary prior to additional sanitization techniques (if applicable), such as a Clear or Purge technique</p> |

| | |
|--|--|
| | <p>applied following cryptographic erase, to ensure that the cryptographic operation completed successfully. A quick sampling verification as described in 6.8.1.5 should also be performed after any additional techniques are applied following cryptographic erase.</p> <p>Not all implementations of encryption are necessarily suitable for reliance upon cryptographic erase as a Purge mechanism. The decision regarding whether to use cryptographic erase depends upon verification of attributes previously identified in this guidance.</p> <p>Do not rely solely upon Degaussing as a sanitization technique on flash-based storage devices.</p> |
| USB Removable Media <i>This includes Pen Drives, Thumb Drives, Flash Drives, Memory Sticks, etc.</i> | |
| Clear: | Overwrite media by using organizationally approved and validated overwriting technologies/methods/tools. The Clear procedure should consist of at least two passes of writes, to include a pattern in the first pass and its complement in the second pass. Additional passes may be used. |
| Purge: | Most USB removable media do not support sanitize commands, or if supported, the interfaces are not supported in a standardized way across these devices. Refer to the manufacturer for details about the availability and functionality of any available sanitization features and commands. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | For most cases where Purging is desired, USB removable media should be Destroyed. |
| Memory Cards <i>This includes SD, SDHC, MMC, Compact Flash, Microdrive, MemoryStick, etc.</i> | |
| Clear: | Overwrite media by using organizationally approved and validated overwriting technologies/methods/tools. The Clear procedure should consist of at least two passes of writes, to include a pattern in the first pass and its complement in the second pass. Additional passes may be used. |
| Purge: | N/A, See Destroy. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | None. |
| Embedded Flash on Boards and Devices <i>This includes motherboards and peripheral cards such as network adapters or any other adapter containing non-volatile flash memory.</i> | |
| Clear: | If supported by the device, reset the state to original factory settings. |
| Purge: | <p>N/A, See Destroy.</p> <p>If the flash can be easily identified and removed from the board, the flash may be Destroyed independently from the disposal of the board that contained the flash. Otherwise, the whole board should be Destroyed.</p> |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | While Embedded flash has traditionally not been specifically addressed in media sanitization guidelines, the increasing complexity of systems and associated use of flash has complementarily increased the likelihood that sensitive data may be present. For example, remote management capabilities integrated into a modern motherboard |

| | |
|--|--|
| | <p>may necessitate storing IP addresses, hostnames, usernames and passwords, certificates, or other data that may be considered sensitive. As a result, for Clearing, it may be necessary to interact with multiple interfaces to fully reset the device state. When this concept is applied to the example, this might include the BIOS/UEFI interface as well as the remote management interface.</p> <p>As with other types of media, the choice of sanitization technique is based on environment-specific considerations. While the choice might be made to neither Clear nor Purge embedded flash, it is important to recognize and accept the potential risk and continue to re-evaluate the risk as the environment changes.</p> |
|--|--|

1

2

Table A.9 — RAM and ROM-Based Storage Device Sanitization

| RAM and ROM-Based Storage Devices | |
|---|---|
| Dynamic Random Access Memory (DRAM) | |
| Clear/ Purge: | Power off device containing DRAM, remove from the power source, and remove the battery (if battery backed). Alternatively, remove the DRAM from the device. |
| Destroy: | Shred, Disintegrate, or Pulverize. |
| Notes: | In either case, the DRAM should remain without power for a period of at least five minutes. |
| Electronically Alterable PROM (EAPROM) | |
| Clear/ Purge: | Perform a full chip Purge as per manufacturer's data sheets. |
| Destroy: | Shred, Disintegrate, or Pulverize. |
| Notes: | None. |
| Electronically Erasable PROM (EEPROM) | |
| Clear/ Purge: | Overwrite media by using organizationally approved and validated overwriting technologies/methods/tools. |
| Destroy: | Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator. |
| Notes: | None. |

3 In the future, organizations will be using media types not specifically addressed by this standard. The processes
4 described in this International Standard should guide media sanitization decision making regardless of the type of
5 media in use.

6 **A.3 Cryptographic erase device guidelines**

7 Cryptographic erase leverages the encryption of target data by enabling sanitization of the target data's
8 encryption key. This leaves only the ciphertext remaining on the media, effectively sanitizing the data.

9 Without the encryption key used to encrypt the target data, the data is unrecoverable. The level of effort needed to
10 decrypt this information without the encryption key then is the lesser of:

11 — The strength of the cryptographic algorithm used to encrypt the data (including mode of operation)

- 1 — The level of entropy of the target data's encryption
- 2 As a result, sanitization of the data is reduced to sanitization of the encryption key(s) used to encrypt the data.
- 3 With cryptographic erase, sanitization may be performed with high assurance much faster than with other
- 4 sanitization techniques. The encryption itself acts to sanitize the data.
- 5 Typically, cryptographic erase can be executed in seconds. This is especially important as storage devices get
- 6 larger resulting in other sanitization methods take more time. Cryptographic erase can also be used as a
- 7 supplement or in addition to other sanitization approaches.
- 8 **Do not use cryptographic erase when devices other than SEDs may also support validated encryption modules**
- 9 **and could support cryptographic erase.** Reliance upon cryptographic erase to purge the media on devices should
- 10 not occur if:
 - 11 — The encryption was enabled after sensitive data was stored on the device without having been sanitized first,
 - 12 or
 - 13 — If it is unknown whether sensitive data was stored on the device without being sanitized prior to encryption,
 - 14 then cryptographic erase should not be relied upon as a Purge mechanism.
- 15 Consider using cryptographic erase if for all devices supporting encryption where cryptographic erase is intended
- 16 for use to purge the media (including SEDs, mobile devices, and other devices), the level of assurance depends
- 17 on the following:
 - 18 — Encryption of all data intended for cryptographic erase prior to storage on the device (including the data, as
 - 19 well as virtualized copies).
 - 20 — Locations on the media where the data encryption key is stored (be it the target data's encryption key or an
 - 21 associated wrapping key) are directly accessible for sanitization (ensuring the actual location on media where
 - 22 the key is stored is addressed) using the appropriate media-specific sanitization technique.
 - 23 — All copies of the encryption keys used to encrypt the target data are sanitized.
 - 24 — If the target data's encryption keys are, themselves, encrypted with one or more wrapping keys, it is
 - 25 acceptable to perform cryptographic erase by Sanitizing a corresponding wrapping key.
 - 26 — And, the ability of a user to clearly identify the commands provided by the device to perform the cryptographic
 - 27 erase operation.
- 28 Other cryptographic erase considerations:
 - 29 — If the encryption key (or any key at or below the level of key sanitized during cryptographic erase) exists
 - 30 outside of the storage device (typically due to escrow or injection), there is a possibility that the key could be
 - 31 used in the future to recover data stored on the encrypted media.
 - 32 — Sanitization using cryptographic erase should not be trusted on devices that have escrowed or injected the
 - 33 key(s) unless the organization has a high level of confidence about how and where the keys were stored and
 - 34 managed outside the device. Such back-up or escrowed copies of data, credentials, or keys should be the
 - 35 subject of a separate device sanitization policy. That policy should address backups or escrowed copies
 - 36 within the scope of the devices on which they are actually stored.
- 37 The choice regarding whether to leverage cryptographic erase on a given device depends upon organizational
- 38 requirements for sanitization, as well as potentially the end user's ability to determine whether the implementation

offers sufficient assurance against future recovery of the data. The level of assurance depends in large part on the factors described in Table A.10.

Table A.10 — Cryptographic erase considerations

| Area | Consideration(s) |
|-------------------------------|--|
| Key Generation | The level of entropy of the random number sources and quality of whitening procedures applied to the random data. This applies to the cryptographic keys, and potentially to wrapping keys affected by the cryptographic erase operation. |
| Media Encryption | The security strength and validity of implementation of the encryption algorithm/mode used for protection of the target data. |
| Key Level and Wrapping | The key being sanitized might not be the Media Encryption Key, but instead a key used to wrap (that is, encrypt) the MEK or another key. In this case, the security strength and level of assurance of the wrapping techniques used should be commensurate with the level of strength of the cryptographic erase operation |

Users seeking to leverage cryptographic erase should identify the mechanisms the storage device implements to address these areas before relying upon cryptographic erase for media sanitization.

— **Make/Model/Version/Media Type:** The product and versions the statement applies to, and the type of media the device uses (i.e., magnetic, SSD, hybrid, other).

— **Key Generation:** Identify whether a deterministic random bit generator such as one of those listed in SP800-90 was used, and how it has been validated.

— **Media Encryption:** Identify the algorithm, key strength, mode of operation, and any applicable validation(s).

— **Key Level and Wrapping:** Identify if the MEK (either wrapped with another value or not wrapped) is directly sanitized, or if a key that wraps the MEK (a key encryption key, or KEK) is sanitized. A description of the wrapping techniques only applies where a KEK (and not the MEK) is sanitized. Wrapping details, when provided, should include the algorithm used, strength, and (if applicable) mode of operation.

— **Data Areas Addressed:** Describe which areas are encrypted and which areas are not encrypted. For any unencrypted areas, describe how sanitization is performed.

— **Key Life Cycle Management:** The key(s) on a device may have multiple wrapping activities (wrapping, unwrapping, and rewrapping) throughout the device's lifecycle. Identify how the key(s) being sanitized are handled during wrapping activities that are not directly part of the cryptographic erase operation. For example, a user may have received an SED that was always encrypting, and simply turned on the authentication interface. Identify how the previous instance of the MEK was sanitized when it was wrapped with the user's authentication credentials.

— **Key Sanitization Technique:** Describe the media-dependent sanitization method for the key being sanitized. Some examples might include three inverted overwrite passes if the media is magnetic, a block erase for an SSD, or other media-specific techniques for other types of media.

— **Key Escrow or Injection:** Identify whether the device supports key escrow or injection at or below the level of cryptographic erase. Identify whether the device supports discovery of whether any key(s) at or below the

- 1 level of the key escrowed has/have ever been escrowed from or injected into the device. If the MEK
2 encryption key is directly sanitized and only a KEK can be escrowed, clearly identify that fact.
- 3 — **Error Condition Handling:** Identify how the device handles error conditions that prevent the cryptographic
4 erase operation from fully completing, such as if a defect is encountered where an instance of the key to be
5 sanitized is stored. For example, if the location where the key was stored cannot be sanitized, does the
6 cryptographic erase operation report success or failure to the user?
- 7 — **Interface Clarity:** Identify which interface commands support the features described in the statement. If the
8 device supports the use of multiple MEKs, identify whether all MEKs are changed using the interface
9 commands available and any additional commands or actions necessary to ensure all MEKs are changed.
- 10 Implementers who choose to apply cryptographic erase should seek either independent validation of these
11 assurance areas or ask the vendor to identify which mechanisms are used to ensure that these concern areas
12 have been addressed. Generally accepted and (where applicable) standardized mechanisms should be used. For
13 example, security requirements for cryptographic requirements are specified in ISO/IEC 19790:2006 and test
14 requirements for cryptographic modules are specified in ISO/IEC 24759:2008. These requirements and tests
15 cover some (but not all) of the concern areas.
- 16 The decision regarding whether to rely upon cryptographic erase should also consider whether the Media
17 Encryption Key has been escrowed or injected, and if so, how the key was protected outside of the storage device.
18 If the Media Encryption Key (or any key at or below the level of key sanitized during CE) exists outside of the
19 storage device, there is a possibility that the key could be used in the future to recover data stored on the
20 encrypted media.

Annex B (informative)

Selecting appropriate storage security controls

B.1 Criteria for selecting controls

B.1.1 Overview

As presented in this International Standard, the storage security guidance may appear to be a collection of controls of equal importance or that need to be implemented in their entirety. Neither of these is true, and organizations can benefit significantly from the adoption of a subset of these controls that are most relevant to their specific environments and needs. The actual set of controls selected can vary from these guidelines via both addition and removal of controls for reasons that could include regulatory requirements, known threats and vulnerabilities, organizational policies, industry or regional guidelines and applicable standards.

This informative Annex B provides a summary of all the storage security controls (see B.2) from the normative clauses²⁶⁾ along with information that can serve as selection criteria based on:

- Data sensitivity classes: provides a data centric focus that leverages three classes, which can be used by organization that have performed basic data classifications
- Security priority codes: provides a security focus that leverages the confidentiality, integrity, and availability aspects of security

Organizations should consider these criteria and guidelines as starting points for storage security control selection, see also ISO/IEC 27002. They can also help organization implement a phased approach to implementing storage security controls.

It is inappropriate to mandate the use of a set of storage security controls listed for either a specific priority or a data sensitivity level in this annex without performing security control selection as part of information security management system planning and implementation, see ISO/IEC 27001. In addition, the security priority and data sensitivity criteria in this Annex B should not be used as the basis for rating or scoring the security of storage systems or infrastructure.

B.1.2 Data sensitivity classes

B.1.2.1 General

Organizations that have performed basic data classifications, based on data sensitivity or criticality, can leverage these classifications to help identify storage security controls that are most relevant to their environments. To assist with such an effort, three generic data sensitivity classes or levels are defined: Low (see B.1.2.2), Moderate (see B.1.2.3), and High (see B.1.2.4).

As a starting point, organizations need to map their specific data classifications to one of the three data sensitivity classes defined in this Annex B. The summary of controls listed in the tables in B.2 can then be consulted to

²⁶⁾ All of the storage security controls in this annex were extracted from clauses 6 and 7 of this standard, and they are provided here in summary form. When additional information or clarification is needed, consult the appropriate source clauses.

1 identify the relevant storage security controls; within these tables, a data sensitivity of "L" corresponds to "Low,"
2 "M" corresponds to "Moderate," and "H" corresponds to "High."

3 **B.1.2.2 Low data sensitivity**

4 Data of this nature are typically easily accessible and determined for internal use within larger groups or
5 organizations (e.g., business entities, government agencies, etc.). In addition, the data is considered less
6 sensitive (e.g., no mandated confidentiality or privacy requirements), have limited value, and are not considered
7 business/mission critical.

8 Minimum protective controls are still needed because unauthorized disclosure or circulation could:

- 9 — have limited negative effects to business, but not breach contractual and/or legal agreements or laws
- 10 — have limited negative effects to government
- 11 — affect individuals in their social and economic circumstances

12 **B.1.2.3 Moderate data sensitivity**

13 Data of this nature are typically restricted to limited groups of people or for internal use within organizational units
14 (e.g., business units, government departments, etc.). In addition, the data is considered sensitive (e.g., have
15 mandated confidentiality or privacy requirements), have significant value, or are considered business/mission
16 critical.

17 Protective controls are needed because unauthorized disclosure or circulation could:

- 18 — have significant negative effects to business and could breach contractual and or/legal agreements or laws
- 19 — constitute a breach of government security that exposes confidential data
- 20 — affect individuals substantially in their social and economic circumstances

21 **B.1.2.4 High data sensitivity**

22 Data of this nature are typically restricted to sole persons or small, namely known groups of people or highly
23 secured organizational units (e.g., business groups/projects, government departments/groups, etc.). In addition,
24 the data is considered very sensitive (e.g., have mandated confidentiality or privacy requirements), have very high
25 value, and/or are considered business/mission critical (e.g., trade secrets).

26 NOTE The difference between moderate and high data sensitivity can be due to an increase in sensitivity, value, or
27 criticality as well as a combination of factors (e.g., both sensitivity and business/mission criticality).

28 Protective controls are absolutely needed because unauthorized disclosure or circulation could:

- 29 — have significant, existence-threatening negative effects to business and would breach contractual and/or
30 legal agreements or laws.
- 31 — constitute a major breach of government security that exposes highly confidential or secret data
- 32 — endanger individuals in their health, life or personal liberty

B.1.3 Security priority codes

Organizations that tend to focus on the confidentiality, integrity, and availability aspects of security can leverage these aspects to help identify storage security controls that are most relevant to their environments. To assist with such an effort, the summary of controls listed in the tables in B.2 include priority codes for each of the security aspects as well as an indicator for system-wide (i.e., the priority codes are identical for all three security aspects); within these tables, a priority indicator of "C" corresponds to "Confidentiality," "I" corresponds to "Integrity," "A" corresponds to "Availability," and "S" corresponds to "System-wide." The priority codes used in the tables within B.2 are numeric and in the range of 0 to 5, with 5 representing the highest priority.

Organizations can leverage the priority code data by first selecting a particular security aspect (confidentiality, integrity, or availability) or system-wide designator. Next, the storage security controls having the highest priority codes associated with the selected designator represent the controls that are most likely to be applicable. Using such an approach, an organization could start with the controls having a priority code of 5 and then proceed in a phased approach where the next set to be addressed correspond to a priority code of 4, followed by a priority code of 3, etc.

B.2 Summary of storage security controls

B.2.1 Supporting controls for storage security

Tables B.1, B.2, B.3, B.4, B.5, B.6, and B.7 summarize the security controls and guidance contained in clause 6 as well as showing how they are relevant to different data sensitivity categories/level (see B.1.2) and priority codes (see B.1.3).

Table B.1 — Direct Attached Storage (Subclause 6.2)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| DAS should be physically secured | | 5 | 3 | 5 | X | X | X |
| Some form of encryption (SED, FDE, host-based, or application-based) should be used to protect the at rest data | | 5 | 3 | 0 | | X | X |
| Media sanitization should be used on all DAS involved with sensitive and high value data | | 5 | 1 | 0 | X | X | X |
| To guard against accidental or intentional data loss or corruption, backups of the DAS contents should be made on a regular basis | | 0 | 5 | 0 | X | X | X |

Table B.2 — Storage networking (Subclause 6.3)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Storage Area Networks (Subclause 6.3.2) | | | | | | | |
| Where possible, avoid network connections between classes (e.g., production or development) | | 5 | 3 | 0 | | X | X |
| Physically isolate storage devices from other data centre device | X | 2 | 2 | 2 | | X | X |
| Logically segregate storage traffic from normal server traffic | X | 4 | 4 | 4 | | X | X |
| Segregate management traffic from all other traffic | | 2 | 3 | 1 | | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Carefully review configuration of network gateways | | 3 | 3 | 4 | X | X | X |
| Restrict host access on the switches (e.g., ACLs, binding lists, FC-SP policy) | | 4 | 3 | 3 | | X | X |
| For FC, use NPIV to assign individual N_Port IDs to virtual hosts | | 5 | 3 | 2 | | X | X |
| For FC, restrict switch interconnections (e.g., ACLs, binding lists, FC-SP policy) | | 4 | 3 | 2 | | X | X |
| For FC, carefully consider the adequacy of basic zoning as a security measure | | 5 | 3 | 3 | X | X | X |
| For FC, disable unused switch ports | | 3 | 4 | 1 | X | X | X |
| For FC, carefully use default zones and zone sets (assume a least privilege posture) | | 3 | 3 | 1 | X | X | X |
| For FC, configure switches, extenders, routers, and gateways with the least amount of access | | 4 | 4 | 2 | | X | X |
| Avoid connecting iSCSI interfaces to general purpose LANs; segregate for security and performance | X | 5 | 5 | 5 | | X | X |
| For iSCSI, carefully use VLANs when the use of physically isolated LANs is not an option | X | 5 | 5 | 5 | X | X | X |
| Carefully set up the peer-to-peer relationship between FCIP entities | | 5 | 3 | 5 | X | X | X |
| Consider using a private IP network used exclusively by the FCIP entities | | 5 | 3 | 5 | | X | X |
| For FCIP, consider using IPsec to secure the communication channel when sensitive data could be exposed | | 5 | 4 | 3 | | X | X |
| For FCoE, leverage the FCP-based security mechanisms (e.g., FC-SP) for FCoE | | 4 | 3 | 2 | X | X | X |
| For FCoE, protect against Ethernet broadcast storms (e.g., allocation of adequate input buffering), which can cause throughput and timeout issues | | 0 | 1 | 3 | X | X | X |
| ACLs should be used to control FCoE network access (e.g., denying specific hosts from unnecessary or unwanted traffic) | | 5 | 4 | 1 | X | X | X |
| For FCoE, carefully use VLANs when the use of physically isolated LANs is not an option | X | 5 | 5 | 5 | X | X | X |
| Network Attached Storage (Subclause 6.3.3) | | | | | | | |
| Enable NFS only if it is needed | | 3 | 3 | 1 | X | X | X |
| Use NFSv4 (or later versions) whenever possible and limit NFSv3 usage | | 3 | 3 | 1 | X | X | X |
| For NFS, filter client and management access by IP address for additional security | | 2 | 3 | 4 | X | X | X |
| For NFS, encrypt client data access (e.g., IPsec) when necessary | | 5 | 5 | 2 | | X | X |
| Use later versions of the SMB protocol | | 3 | 4 | 5 | | | |
| Turn off low-security session negotiation protocols, such as NTLM v1, LanMan and plaintext; use NTLM v2 or Kerberos instead | X | 5 | 5 | 5 | X | X | X |
| Maintain up-to-date patch levels | X | 4 | 4 | 4 | X | X | X |
| Use SMB signing | | 5 | 5 | 0 | X | X | X |
| Maintain Active Directory (AD) services securely | | 3 | 3 | 5 | X | X | X |
| Use one-way trusts, from leaf domains to parent domains, when possible | | 5 | 5 | 2 | X | X | X |
| Enable SMB/CIFS only if it is needed | | 3 | 3 | 1 | X | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| For SMB/CIFS, encrypt client data access (e.g., IPsec) when necessary | | 4 | 3 | 0 | | X | X |

1

2

Table B.3 — Storage management (Subclause 6.4)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Authentication and authorization (Subclause 6.4.2) | | | | | | | |
| All users should have unique ID for their personal use | | 5 | 5 | 0 | X | X | X |
| A suitable authentication technique (strong passwords, strong authentication, or multi-factor authentication) should be chosen to substantiate the claimed identity of a user | | 5 | 5 | 0 | X | X | X |
| For all remote access, use strong authentication or multi-factor authentication along with secure channels | | 5 | 5 | 0 | X | X | X |
| When possible, use a centralized authentication solution for improved monitoring and control | | 4 | 4 | 0 | | X | X |
| Use multi-factor authentication, complimented by an auto identity provisioned system, when managing sensitive and high-value data | | 4 | 4 | 0 | | | X |
| Disable login to the root account. Remotely log all "sudo" operations | | 3 | 3 | 0 | X | X | X |
| When possible, use entity authentication in TLS and IPsec connections as well as within storage protocols | | 3 | 3 | 0 | | X | X |
| Implement and use general roles like Security Administrator, Storage Administrator, Security Auditor, and Storage Auditor within storage | | 3 | 3 | 0 | X | X | X |
| Secure the management interfaces (Subclause 6.4.3) | | | | | | | |
| Restrict physical access to management interfaces | X | 5 | 5 | 5 | X | X | X |
| Disable and disconnect serial management ports when not in use | | 2 | 2 | 1 | X | X | X |
| Segregate LAN interfaces used for management from other LAN traffic; physical isolation is preferred, but logical isolation (such as VLANs) should be used at a minimum | | 3 | 2 | 1 | X | X | X |
| Disable modem ports when not needed | | 2 | 2 | 1 | X | X | X |
| Use firewalls and TCP wrappers to restrict access to management networks to authorized hosts and protocols | | 4 | 4 | 5 | X | X | X |
| Use entity authentication to establish trust relationships between storage systems and the management systems (e.g., FC-SP for in-band management over Fibre Channel) | | 3 | 3 | 4 | | X | X |
| Leverage intrusion detection and prevention mechanisms to identify anomalous behaviours and guard against it | X | 4 | 4 | 4 | | X | X |
| Use ICT infrastructure (DNS, SLP, NTP) with appropriate security controls to avoid indirect attacks | X | 3 | 3 | 3 | | X | X |
| Employ appropriate privileged user controls, including authentication, authorization, and secure auditing/monitoring | X | 5 | 5 | 5 | | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| For storage management, ensure that operating systems and applications are current and sufficiently hardened against attacks | | 5 | 5 | 3 | | X | X |
| For remote storage management, use secure channels for all remote access (VPN, TLS, SSH Secure Shell, HTTPS) | X | 5 | 5 | 5 | | X | X |
| For remote storage management, employ strong authentication or multi-factor authentication | X | 5 | 5 | 5 | | X | X |
| For remote storage management, restrict privileges to the minimum needed (i.e., least privilege) | | 4 | 4 | 2 | | X | X |
| Devise organizational and technical controls to restrict the management interface used for remote (non-local) vendor maintenance sessions | | 3 | 2 | 2 | | X | X |
| Technical controls should restrict communication traffic (i.e. hosts, ports, and protocols) to the minimum required for remote vendor maintenance operations | | 3 | 3 | 5 | X | X | X |
| After the accessing party (vendor maintenance personnel) is authenticated, additional controls at the access point should be devised to authorize the vendor maintenance session, including accepting, asking for approval, or denying the requested session | | 3 | 3 | 5 | | X | X |
| Appropriate logs containing audit records of vendor actions should be generated. | X | 4 | 4 | 4 | | X | X |
| The organization should restrict dial-up access lines to authorized accessing parties, enforcing a modem callback protocol and disabling connection establishment until vendor requests a maintenance session and the request is authorized by the organization | | 2 | 3 | 2 | | X | X |
| Security auditing, accounting, and monitoring (Subclause 6.4.4) | | | | | | | |
| Include storage systems and infrastructure in the logging policy (what is collected, retention/preservation, time synchronization, etc.) | X | 4 | 4 | 4 | X | X | X |
| In the policies, identify and address the evidentiary expectations for storage logs | X | 5 | 5 | 5 | X | X | X |
| Employ external and centralized event logging to a trusted remote source | X | 5 | 5 | 5 | X | X | X |
| Establish and use a common, accurate time source across the storage systems and infrastructure | X | 5 | 5 | 5 | X | X | X |
| Natively log events to one, and preferably multiple, external log servers | X | 4 | 4 | 4 | X | X | X |
| For compliance, accountability, and/or security purposes, events should be logged as they occur (no buffering) | X | 4 | 4 | 4 | X | X | X |
| Ensure that the storage logging is factored into SIEM solutions, when such technology is deployed | X | 3 | 3 | 3 | | X | X |
| Log all occurrences of the minimum set of security events with the necessary data | X | 5 | 5 | 5 | X | X | X |
| Ensure the event log data are handled and retained correctly | X | 5 | 5 | 5 | X | X | X |
| Implement appropriate measures to preserve log integrity and prevent their modification or destruction | X | 5 | 5 | 5 | | X | X |
| Protective measures may be required to ensure the confidentiality and integrity of the event log data | X | 5 | 5 | 5 | | X | X |
| Use special purpose log servers to handle unique and/or sensitive | X | 4 | 4 | 4 | | | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| data requirements | | | | | | | |
| Leverage log relays and log filtering to minimize the impact of specialized storage requirements (WORM) | X | 3 | 3 | 3 | | | X |
| System hardening (Subclause 6.4.5) | | | | | | | |
| Remove un-needed/un-used software | | 2 | 3 | 3 | X | X | X |
| Remove unnecessary accounts | X | 3 | 3 | 3 | X | X | X |
| Eliminate, disable, or change passwords on predefined or default accounts | X | 4 | 4 | 4 | X | X | X |
| Close or disable all unused networking ports | | 1 | 1 | 3 | X | X | X |
| Install latest patches from a trusted source | X | 4 | 4 | 4 | X | X | X |
| Update firmware from a trusted source | X | 4 | 4 | 4 | X | X | X |
| Install and maintain malware protection | X | 5 | 5 | 5 | X | X | X |

1

2

Table B.4 — Block-based storage (Subclause 6.5)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Fibre Channel storage (Subclause 6.5.1) | | | | | | | |
| Restrict access to storage with WWN filtering (i.e., LUN masking) and other access control mechanisms | | 4 | 4 | 1 | X | X | X |
| Mutual authentication (per FC-SP) should be used with all servers and switches; leverage centralized authentication services when possible | | 2 | 2 | 5 | X | X | X |
| If possible, encrypt Fibre Channel connections (e.g., ESP_Header) that leave the protected area | | 4 | 3 | 1 | | X | X |
| IP storage (Subclause 6.5.2) | | | | | | | |
| Control iSCSI initiator access by filtering based on source IP addresses and protocols | | 5 | 3 | 5 | X | X | X |
| Use CHAP authentication for both initiators and targets in all iSCSI implementations | X | 5 | 5 | 5 | X | X | X |
| Consider using IPsec to secure the communication channel when sensitive data could be exposed | | 5 | 3 | 2 | | X | X |
| Use iSNS, SLP, DNS infrastructure with appropriate security controls to avoid indirect attacks | X | 3 | 3 | 3 | X | X | X |

3

4

Table B.5 — File-based storage (Subclause 6.6)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| NFS-based NAS (Subclause 6.6.1) | | | | | | | |
| Employ user-level authentication whenever possible (e.g., NFSv4 with Kerberos V5) | X | 5 | 5 | 5 | X | X | X |
| Configure the NFS server to export file systems explicitly for the authorized users | | 3 | 2 | 1 | X | X | X |
| Configure the NFS server to export file systems with minimum required privileges | | 3 | 2 | 1 | X | X | X |
| Avoid granting “root” or “administrator” access to files on network filesystems | | 5 | 5 | 3 | X | X | X |
| Make sure NFSv4 ACLs are assigned correctly | | 4 | 4 | 2 | X | X | X |
| Use Kerberos authentication for NFSv3 | X | 3 | 3 | 3 | X | X | X |
| Consider using Kerberos Safe and Private modes to sign and encrypt NFS traffic | | 4 | 4 | 2 | X | X | X |
| Filter client access to NFS shares whenever possible | X | 3 | 3 | 3 | X | X | X |
| Do not allow NFS clients to run <i>suid</i> and <i>sgid</i> programs on exported file systems | | 3 | 3 | 5 | X | X | X |
| Exported file systems should be in their own partitions to prevent system degradation by an attacker writing to an exported file system until it is full | X | 4 | 4 | 4 | X | X | X |
| Encrypt data at rest when necessary | | 3 | 3 | 1 | | X | X |
| Do not allow NFS exports of administrative file systems (e.g., /etc) | X | 3 | 3 | 3 | | X | X |
| Guard against malware (e.g., viruses, worms, rootkits, etc.) | X | 5 | 5 | 5 | | X | X |
| Continually monitor content placed in NFS shares and relevant access controls | | 4 | 1 | 2 | | X | X |
| SMB/CIFS-based NAS (Subclause 6.6.2) | | | | | | | |
| Disable unauthenticated access to CIFS shares and NAS devices (i.e. restrict <i>Anonymous</i>) | | 5 | 3 | 4 | X | X | X |
| Disable “Guest” and “Everyone” access to all CIFS shares | | 4 | 4 | 2 | X | X | X |
| Implement authentication and access control via a centralized mechanism (RADIUS, LDAP) | X | 5 | 5 | 5 | | X | X |
| Enable SMB signing for clients and the NAS device | | 3 | 5 | 3 | | X | X |
| Enable CIFS auditing whenever possible | | 3 | 3 | 1 | | X | X |
| Continually review content placed in CIFS shares and relevant access controls | | 4 | 1 | 2 | | X | X |
| Encrypt data at rest when necessary | | 3 | 3 | 1 | | X | X |
| Guard against malware (e.g., viruses, worms, rootkits, etc.) | X | 5 | 5 | 5 | | X | X |
| Implement CIFS with strong authentication (NTLMv2, Kerberos) | | 4 | 4 | 1 | | | |
| Parallel NFS-based NAS (Subclause 6.6.3) | | | | | | | |
| Controls and control mechanisms should be applied consistently across clusters (both symmetric and asymmetric) | | 3 | 5 | 3 | X | X | X |
| Security assurance properties should not be dependent on the client accessing a specific fileserver | X | 3 | 3 | 3 | X | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| For asymmetric clusters, controls should be implemented such that they are consistent across different protocols | X | 4 | 4 | 4 | X | X | X |
| Security controls should not be dependent on path traversal of the filesystem namespace across servers | | 2 | 2 | 4 | X | X | X |

1

2

Table B.6 — Object-based storage (Subclause 6.7)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Cloud storage (Subclause 6.7.1) | | | | | | | |
| When using CDMI, ensure that Transport Layer Security (TLS) is used for all transactions | X | 4 | 4 | 4 | X | X | X |
| Check the security capabilities of the cloud service provider's CDMI implementation and make a risk-based decision on whether the offered security is adequate | X | 5 | 5 | 5 | X | X | X |
| Authenticate CDMI entities (certificates for servers and HTTP basic authentication for clients) | X | 5 | 5 | 5 | X | X | X |
| Use CDMI Domains to provide a place for authentication mappings to external authentication providers | | 4 | 4 | 1 | | X | X |
| When possible, enable CDMI security logging and retrieve the event data in a regular and timely fashion | | 3 | 3 | 1 | | X | X |
| Align the automatic deletion capability (CDMI Deletion) with the organization's data retention policy | X | 3 | 3 | 3 | | X | X |
| Prior to using CDMI Holds, understand the process and mechanism for lifting the CDMI Hold | X | 4 | 4 | 4 | | X | X |
| For cryptographic functionality, always verify that the implementation has used a requested CDMI Capability (supported operation), and not something different | | 4 | 1 | 1 | | X | X |
| Use the CDMI sanitization functionality to clear sensitive data from the cloud service provider's storage | X | 3 | 3 | 3 | | X | X |
| Object-based Storage Device (Subclause 6.7.2) | | | | | | | |
| For OSD, IPsec should be used for all transactions involving sensitive data on insecure networks | | 5 | 5 | 0 | | X | X |
| For OSD, the object store should verify the authenticity of the capability prior to performing an operation | X | 5 | 5 | 5 | X | X | X |
| Clock synchronization between the OSD and the security manager should be implemented using a secure protocol | X | 4 | 4 | 4 | X | X | X |
| For OSD, capability expiration times should have limits that minimize the amount of time a compromised capability can be used | X | 3 | 3 | 3 | X | X | X |
| For OSD, working keys (used to generate capability keys) should be refreshed frequently | X | 3 | 3 | 3 | | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Content Addressable Storage (Subclause 6.7.3) | | | | | | | |
| Users and applications should be authenticated and authorized before access is granted to the CAS system. | X | 5 | 5 | 5 | X | X | X |
| The CAS system should ensure that content will be readable and accessible over its entire life-cycle. | | 0 | 5 | 5 | X | X | X |
| The CAS system should employ a robust hashing mechanism | | 0 | 4 | 4 | X | X | X |

Table B.7 — Storage security services (Subclause 6.8)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Data sanitization (Subclause 6.8.1) | | | | | | | |
| Organizations and individuals should categorize the information, assess the nature of the medium on which it is recorded, assess the risk to confidentiality, and determine the future plans for the media (for example, reuse) | | 5 | 5 | 1 | X | X | X |
| The selected type of sanitization should be assessed as to cost, environmental impact, etc., and a decision made that best mitigates the risk to confidentiality and best satisfies other constraints imposed on the process | | 5 | 3 | 1 | X | X | X |
| The level of sanitization operations should be carefully balanced against the risks, paying particular attention to PII and EHR as well as business or mission critical data (e.g., trade secrets, intellectual property, etc.). | | 5 | 3 | 1 | X | X | X |
| When storage media are transferred, become obsolete, or are no longer usable or required by an information system, it is important to ensure that residual magnetic, optical, electrical, or other representation of data that has been deleted is not easily recoverable. | | 5 | 0 | 1 | X | X | X |
| Once a decision is made and after applying relevant organizational environmental factors, then Annex A should be used to determine recommended sanitization of specific media. | | 5 | 0 | 1 | X | X | X |
| Not all types of available media are specified in this standard, and for those media not included, organizations should identify and use processes that will fulfil the intent to clear, purge, or destroy their media. | | 5 | 0 | 1 | X | X | X |
| Sanitization of media at end-of-life situations is recommended, even when using encryption methods. | | 3 | 0 | 1 | X | X | X |
| If the logical storage is writeable, then sanitization may be possible using an overwrite technique or for encrypted content (e.g., files), the encryption keying materials can be destroyed. | | 2 | 0 | 1 | X | X | X |
| Organizations should maintain a record of sanitization activities to document what media were sanitized, when, how they were sanitized, and the final disposition of the media. | X | 4 | 4 | 4 | X | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| A certificate of sanitization should be produced and contain the appropriate details | X | 4 | 4 | 4 | X | X | X |
| The audit trail associated with sanitization should capture time stamped transactions and progress. | X | 4 | 4 | 4 | X | X | X |
| A full verification of the sanitization should be performed if time and external factors permit. | X | 4 | 4 | 4 | X | X | X |
| If cryptographic erasure is used for sanitization, appropriate verification should be performed. | X | 4 | 4 | 4 | X | X | X |
| Data confidentiality (Subclause 6.8.2) | | | | | | | |
| When data in motion encryption is need, it should provide end-to-end protection | | 3 | 3 | 1 | | X | X |
| Encryption of data in motion can impose significant computational burdens on the communicating entities, so appropriate compensations should be implemented to minimize the impacts | | 4 | 4 | 5 | | X | X |
| For IPsec, version 3 and IKE version 2 (or later versions) should be used | X | 5 | 5 | 5 | | X | X |
| For TLS, version 1.2 (or later) should be used | X | 5 | 5 | 5 | | X | X |
| For at rest encryption, algorithms and modes of operations designed specifically for storage technology should be used | X | 5 | 5 | 5 | | X | X |
| Limit the amount of time a key is in plaintext form and prevent humans from viewing plaintext keys | | 5 | 3 | 3 | | X | X |
| Cryptographic keys should only be used for one purpose, specifically, do not use key-encrypting keys to encrypt data or use data encrypting keys to encrypt other keys | X | 5 | 5 | 5 | | X | X |
| Randomly choose keys from the entire keyspace by using a cryptographically strong Random Number Generator | | 4 | 3 | 3 | | X | X |
| Check for and avoid use of known weak keys | X | 3 | 3 | 3 | | X | X |
| Data encryption keys should be limited to a finite cryptoperiod (typically no more than 2 years) or to a maximum amount of data processed | | 4 | 3 | 3 | | X | X |
| When possible, storage systems and infrastructure should use KMIP-compliant key management infrastructure to facilitate centralized key management | X | 3 | 3 | 3 | | X | X |
| Data reductions (Subclause 6.8.3) | | | | | | | |
| When encryption is used along with compression, the compression should be applied before the encryption | X | 4 | 4 | 4 | | X | X |
| When encryption is used along with deduplication, the deduplication should be applied before the encryption | X | 4 | 4 | 4 | | X | X |
| When both compression and deduplication are used along with encryption, the order of use should be deduplication and compression or compression and deduplication, and then encryption | X | 4 | 4 | 4 | | X | X |
| Compression and/or deduplication can impact DR and BC implementations, so they should be factored into the design, documentation, and testing of DR and BC solutions | X | 5 | 5 | 5 | | X | X |

B.2.2 Storage security design and implementation guidance

Tables B.8, B.9, B.10, B.11, B.12, and B.13 summarize the security controls and guidance contained in clause 7 as well as showing how they are relevant to different data sensitivity categories/level (see B.1.2) and priority codes (see B.1.3).

Table B.8 — Storage security design principles (Subclause 7.2)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Defence in depth (Subclause 7.2.1) | | | | | | | |
| Ensure a balanced focus on the three primary elements: people, technology, and operations | X | 5 | 5 | 5 | X | X | X |
| Follow through with effective information assurance policies and procedures, assignment of roles and responsibilities, commitment of resources, training of critical personnel, and personal accountability | X | 4 | 4 | 4 | X | X | X |
| Deploy protection mechanisms at multiple locations to resist all classes of attacks | X | 3 | 3 | 3 | X | X | X |
| Deploy multiple defence mechanisms (layered) between potential adversaries and targets | X | 3 | 3 | 3 | X | X | X |
| Include both detection and protection mechanism | X | 3 | 3 | 3 | X | X | X |
| Deploy robust key management and PKI that support all information assurance technologies and that are highly resistant to attack | X | 4 | 4 | 4 | | X | X |
| Maintain visible and up to date system security policies | X | 3 | 3 | 3 | X | X | X |
| Actively manage the security posture of the storage technology and protection mechanisms (e.g., install security patches and virus updates, maintain ACLs, etc.) | X | 3 | 3 | 3 | X | X | X |
| Perform regular security threat assessments to determine the continued security readiness | X | 3 | 3 | 3 | X | X | X |
| Monitor and react to current threats | X | 4 | 4 | 4 | X | X | X |
| Security domains (Subclause 7.2.2) | | | | | | | |
| Storage and storage networks of different sensitivity levels should be located in different security domains | | 4 | 4 | 3 | | X | X |
| Devices and computer systems providing services for external networks should be located in different domains than internal network devices and computer systems | | 3 | 3 | 2 | X | X | X |
| Strategic assets should be located in dedicated security domains | | 5 | 3 | 2 | | X | X |
| Devices and computer systems of low trust level should be located in dedicated security domains with limited or no access to storage assets | X | 5 | 5 | 5 | X | | |
| Storage and storage networks used for different purposes (e.g., development, production, management, etc.) and using different technologies (e.g., CIFS/NFS, iSCSI, CDMI, etc.) should be located in separate security domains | | 3 | 2 | 1 | | X | X |
| Storage networks should be in different security domains than regular networks (e.g., corporate LANs) | | 2 | 4 | 2 | | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Storage device and storage network management systems should be located in dedicated security domains | | 1 | 3 | 1 | X | X | X |
| Systems in development stage should be located in different domains than production systems | | 5 | 3 | 3 | | X | X |
| Storage devices that may be permitted to reside with a single security domain, but used for multiple purposes or hold multiple levels of sensitive data, should be further isolated (using zoning, VLANs, and VSANs) to minimize possible interactions | | 4 | 4 | 1 | | X | X |
| Design resilience (Subclause 7.2.3) | | | | | | | |
| Storage security design should incorporate several layers of redundancy to eliminate single points of failure and to maximize the availability of the storage infrastructure. | X | 5 | 5 | 5 | X | X | X |
| The designs should also use a wide set of features destined to make the storage more resilient to attacks and network failures | X | 4 | 4 | 4 | | X | X |
| Secure initialization (Subclause 7.2.4) | | | | | | | |
| As a design principle, the architecture should support a secure initialization sequence to ensure the transition from a “down” state after a power-on or reset is applied. | X | 4 | 4 | 4 | | X | X |
| During the initialization phase externally accessible processes and network interfaces should not be available or at a minimum deny access until the subjects are authenticated. | X | 4 | 4 | 4 | | X | X |
| Software and OS load processes should start from a known state with secure values specified by the system administrator when the system was last operational. | X | 3 | 3 | 3 | | X | X |

1

2

Table B.9 — Data reliability, availability, and resilience (Subclause 7.3)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Reliability (Subclause 7.3.1) | | | | | | | |
| The reliability of the storage system and infrastructure should not be adversely impacted by the inclusion of security features | | 1 | 4 | 4 | X | X | X |
| Vulnerabilities should be proactively managed to minimize their impacts on system reliability | | 1 | 4 | 4 | X | X | X |
| Controls should be assessed to determine whether they are appropriate to ensure the reliability and security of data | X | 3 | 3 | 3 | X | X | X |
| Availability (Subclause 7.3.2) | | | | | | | |
| Because of the importance of availability, storage security designs and implementations should strive to minimize impacts to availability | | 1 | 1 | 5 | X | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Data encryption keys should be carefully managed to avoid data availability problems when keys are unavailable or inadvertently destroyed | | 1 | 1 | 5 | | X | X |
| Data protection mechanism should be part of availability designs to guard against major outages due to system failures | | 1 | 3 | 4 | X | X | X |
| Backups and replication (Subclause 7.3.3) | | | | | | | |
| Data protection mechanisms (like backups, replication, etc.) should be designed with quick recoveries in mind, rather than just preservation of data | | 1 | 4 | 5 | X | X | X |
| Ensure that the backup approach, especially for business/mission critical data, is aligned with its associated restore strategy | | 1 | 4 | 5 | | X | X |
| Ensure that the backup approach provides adequate protections against unauthorized access (e.g., encryption) | | 4 | 4 | 1 | | X | X |
| Establish a chain of trusted individuals (and vendors) who handle the storage media | X | 5 | 5 | 5 | | X | X |
| Implement backup validations to show “proof” that restore requirements are being met | | 0 | 5 | 5 | | X | X |
| Ensure that the replication approach, especially for business/mission critical data, is aligned with its associated reliability, fault-tolerance, or performance requirements | X | 5 | 5 | 5 | | X | X |
| Ensure that the replication approach provides adequate protections against unauthorized access (e.g., encryption in motion) | | 5 | 0 | 3 | | X | X |
| Ensure that the CDP approach (e.g., continuous, near continuous, fixed interval, etc.), especially for business/mission critical data, is aligned with its associated restore strategy | | 0 | 3 | 3 | | X | X |
| In high network bandwidth scenarios (e.g., multimedia files), employ throttling techniques which prioritize network traffic in order to reduce the impact of CDP on day-to-day operation | | 0 | 2 | 2 | | X | X |
| Ensure that the CDP approach provides adequate protections against unauthorized access (e.g., in motion and at rest encryption) | | 4 | 3 | 2 | | X | X |
| Disaster Recovery and Business Continuity (Subclause 7.3.4) | | | | | | | |
| Ensure the storage ecosystem is factored into the DR/BC planning and implementation | X | 5 | 5 | 5 | | X | X |
| Prepare for limited disruption events (system failures, adversarial attacks, operator errors) | X | 4 | 4 | 4 | X | X | X |
| Identify and document the unique staffing and facility requirements associated with the storage ecosystem | X | 3 | 3 | 3 | | X | X |
| Perform on-going planning and regular testing of assumption, which are critical to successful DR/BC; results of DR/BC testing should be fed back into on-going maintenance of the DR/BC plan | X | 3 | 3 | 3 | | X | X |
| Resilience (Subclause 7.3.5) | | | | | | | |
| Security should be an integral part of the resilience strategy; plan for unit failures and compromises of both the storage and security | X | 5 | 5 | 5 | X | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| technologies | | | | | | | |
| Redundancy should be exploited to the extent possible | | 0 | 5 | 5 | X | X | X |
| Diverse components that are easily repairable should be used whenever possible | | 2 | 2 | 4 | X | X | X |
| Security features and functionality should be implemented in such a way as to cause no adverse impact to the resilience of the storage system or infrastructure | X | 4 | 4 | 4 | X | X | X |

1

2

Table B.10 — Data retention (Subclause 7.4)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Long-term retention (Subclause 7.4.1) | | | | | | | |
| Actively check the integrity of the data in the system at regular intervals rather than waiting to when it is read | | 1 | 4 | 4 | X | X | X |
| Leverage data migrations to effect greater security | X | 2 | 2 | 2 | | X | X |
| Ensure that the archival storage system is able to authenticate new users and establish their relationship to resources attached to existing users | | 4 | 1 | 1 | | X | X |
| Ensure that the secrecy mechanisms function in the complete absence of the user that wrote the data | | 5 | 0 | 0 | | X | X |
| Ensure that security logging is sufficiently complete and long-lived that it assists in detecting slow attacks and maintains an attack history that can be used to make decisions to adjust the data protections | X | 4 | 4 | 4 | | X | X |
| The system should either immediately deal with any compromise or maintain a history of compromises in order to intelligently schedule corrective action | X | 5 | 5 | 5 | | X | X |
| The use of data reduction technologies (e.g., compression and deduplication) should be used carefully because they can compromise data integrity if they are not factored into copies | | 1 | 4 | 1 | | X | X |
| Short to medium-term retention (Subclause 7.4.2) | | | | | | | |
| Multiple physical and/or logical replicas of the data should be created and preserved; the replicas need to be organized to be as independent as possible (e.g., geographic, administrative/management, and platform/OS), and their number chosen according to the data's value and tolerance of risk | | 0 | 5 | 5 | | X | X |
| On a defined schedule, audit for both obvious and latent faults (e.g., integrity checks), and the damage they cause; repair the corrupted data using the good data from other replicas before that damage spreads | | 0 | 4 | 2 | | X | X |
| Match the access control scheme to the legal and regulatory requirements for the information being preserved | X | 5 | 5 | 5 | | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Ensure that accountability and traceability measures are adequate and functional; all data accesses may require audit log entries | X | 4 | 4 | 4 | | X | X |
| Implement mechanism to demonstrate data authenticity, provenance, and chain of custody, especially for data of an evidentiary nature | | 4 | 5 | 0 | | X | X |
| If encryption is used, archive/escrow the keys and keying material; rekey the data within recommended cryptoperiods or when the underlying cryptographic algorithm needs to be replaced | | 5 | 5 | 3 | | X | X |

Table B.11 — Data confidentiality and integrity (Subclause 7.5)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Storage-based encryption should not be the primary form of encryption for sensitive data | | 4 | 3 | 2 | | X | X |
| Selection of a point of encryption should be influenced by DR and BC, data reduction, and data protection considerations | X | 4 | 4 | 4 | | X | X |
| Data retention needs should be considered when selecting and deploying encryption | X | 3 | 3 | 3 | | X | X |
| The security strength of the encryption solution should be at least 112 bits | | 4 | 4 | 1 | | X | X |
| Cryptographic modules used to protect sensitive and/or regulated data should be validated using recognized criteria | X | 5 | 5 | 5 | | X | X |
| Multiple encryption steps may be used, as when data encrypted for privacy purposes is further encrypted by a Self-Encrypting Drive for security purposes | X | 3 | 3 | 3 | | X | X |
| Ensure the encryption mechanisms create appropriate audit log entries (activation, verification, integrity checks, re-keying, etc.) | X | 4 | 4 | 4 | | X | X |
| Agree in advance on what audit log material demonstrates (to the satisfaction of the legal department) that encryption was properly performed | X | 4 | 4 | 4 | | X | X |
| Perform regular and audited checks that encryption was properly performed and consider outside accreditation | X | 3 | 3 | 3 | | X | X |
| Fully automate key management whenever possible | X | 3 | 3 | 3 | | X | X |
| Sparsely use keys with a long life (i.e., approaches the maximum recommended cryptoperiod, which is typically no more than 1-2 years, depending on the key type) | X | 4 | 4 | 4 | | X | X |
| Enforce strict access controls to limit user capabilities and separation of duties constraints (e.g., a security role) for key generation, change and distribution | X | 5 | 5 | 5 | | X | X |
| For sensitive and/or high-value data, the encryption should be end-to-end (i.e., in motion and at rest) | X | 3 | 3 | 3 | | X | X |
| To protect data integrity, storage systems should include sufficient malware protections to guard against attacks on data | X | 4 | 4 | 4 | X | X | X |
| WORM-based storage should be used to help meet immutability requirements | | 2 | 4 | 2 | X | X | X |

1

Table B.12 — Virtualization (Subclause 7.6)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Storage virtualization (Subclause 7.6.1) | | | | | | | |
| Ensure appropriate service level objectives for virtual storage | X | 4 | 4 | 4 | X | X | X |
| Match the availability objective for the storage infrastructure to the application requirements | | 0 | 0 | 5 | X | X | X |
| Match the confidentiality and privacy requirements for the storage infrastructure to the types of information stored | | 5 | 0 | 0 | | X | X |
| Address multi-tenancy concerns, as appropriate | | 3 | 3 | 2 | X | X | X |
| Storage for virtualized systems (Subclause 7.6.2) | | | | | | | |
| VM access to storage networks should be controlled via use of access controls in the server virtualization (hypervisor) software | | 0 | 3 | 3 | X | X | X |
| NPIV should be leveraged appropriately to limit VM access to storage targets | X | 3 | 3 | 3 | X | X | X |
| VM migration/movement between physical hosts in an infrastructure should be carefully controlled to avoid having unintended security consequences | X | 3 | 3 | 3 | X | X | X |

2

3

Table B.13 — Design and implementation considerations (Subclause 7.7)

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Encryption and key management issues (Subclause 7.7.1) | | | | | | | |
| Understand and obey government import regulations associated with encryption and key management | | 3 | 0 | 0 | | X | X |
| Understand and obey government export regulations associated with encryption and key management | | 4 | 0 | 0 | | X | X |
| Comply with corporate and/or government key escrow requirements | | 5 | 3 | 2 | | X | X |
| Have a recovery plan in the event of a key compromise | X | 5 | 5 | 5 | | X | X |
| Have a key backup plan in place to ensure continued access to encrypted business/mission critical information | X | 5 | 5 | 5 | | X | X |
| Securely distribute key material among storage devices that process/access the same data | | 5 | 5 | 3 | | X | X |
| The effect of encryption on deduplication and compression techniques should be understood and factored in designs and implementations | X | 4 | 4 | 4 | | X | X |
| The inability to apply security techniques like virus scanning, etc. on encrypted data should be understood and mitigated with other mechanisms | X | 4 | 4 | 4 | | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|--|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Align storage and policy (Subclause 7.7.2) | | | | | | | |
| Identify most sensitive (Personally Identifiable Information, intellectual property, trade secrets, etc.) and business/mission critical data categories as well as protection requirements | | 5 | 2 | 1 | | X | X |
| Integrate storage-specific policies with other policies (i.e., avoid creating a separate policy document for the storage ecosystem) | X | 2 | 2 | 2 | X | X | X |
| Address data retention and protection (e.g., write-once-read-many or WORM, authenticity, access controls, etc.) | | 3 | 4 | 2 | | X | X |
| Address data destruction and media sanitization | | 4 | 1 | 1 | | X | X |
| Ensure that all elements of the storage ecosystem comply with policy (e.g., ISO/IEC 27001/27002) | X | 3 | 3 | 3 | X | X | X |
| Give most sensitive/most critical data a priority | X | 4 | 4 | 4 | | X | X |
| Compliance (Subclause 7.7.3) | | | | | | | |
| Ensure that users, especially privileged users, have unique userids (i.e., no shared accounts) | | 5 | 0 | 0 | X | X | X |
| When possible, grant rights and privileges based on roles | X | 3 | 3 | 3 | X | X | X |
| Log all attempted (successful and unsuccessful) management events and transactions | X | 4 | 4 | 4 | X | X | X |
| Ensure logged event/transaction data contains sufficient application and/or system detail to clearly identify the source | X | 3 | 3 | 3 | X | X | X |
| Ensure that the user information can be traced to a specific individual | X | 3 | 3 | 3 | X | X | X |
| When appropriate, treat log records as evidence (chain of custody, non-repudiation, authenticity, etc.) | X | 4 | 4 | 4 | X | X | X |
| Ensure that the storage layer participates in the external audit logging measures | X | 4 | 4 | 4 | X | X | X |
| Monitor the audit logging events and issue the appropriate alerts | X | 5 | 5 | 5 | X | X | X |
| Implement appropriate data retention measures | | 0 | 4 | 0 | X | X | X |
| Implement appropriate data integrity and authenticity measures | | 4 | 5 | 0 | X | X | X |
| Correctly sanitize data upon deletion, repurposing or decommissioning of hardware | | 4 | 0 | 0 | | X | X |
| Correctly sanitize virtual server images, and their copies, at end of life | X | 1 | 1 | 1 | | X | X |
| Implement appropriate data access control measures to control access to data and metadata (e.g., search results); assume a least privilege posture whenever possible | | 3 | 0 | 1 | X | X | X |
| Implement appropriate data confidentiality measures to prevent unauthorized disclosure | | 5 | 0 | 0 | | X | X |
| Ensure that the use of data deduplication does not conflict with data authenticity requirements | | 3 | 4 | 0 | | X | X |
| Ensure data and media sanitization mechanisms do not violate preservation orders | X | 5 | 5 | 5 | | X | X |
| Ensure proper chain of custody procedures are followed when evidentiary data (e.g., audit logs, metadata, mirror images, point-in time copies, etc.) is handled | | 5 | 5 | 0 | | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| Secure multi-tenancy (Subclause 7.7.4) | | | | | | | |
| Use encrypted storage that is aligned with the tenants' usage of resources | | 4 | 4 | 0 | X | X | X |
| Use strong at rest encryption (minimum of 112-bits of security strength) | | 5 | 3 | 0 | X | X | X |
| Use secure and rapid de-provisioning (media sanitization, including cryptographic erase) | | 5 | 0 | 0 | X | X | X |
| Use trusted third-party data storage management (e.g., SNMPv3, SMI-S with TLS, etc.) | X | 4 | 4 | 4 | X | X | X |
| Use automated key management providing tenant-controlled key management (leverages KMIP v1.1 compliant servers) | X | 3 | 3 | 3 | X | X | X |
| Use secure data replication (e.g., data in motion and at rest encryption) | | 5 | 0 | 0 | X | X | X |
| Protect data from administrators (e.g., enforce a least privileges access model, administrators do not have access to the keying materials, etc.) | | 4 | 0 | 0 | X | X | X |
| Use highly available data fabrics (multi-path and diverse path) | | 0 | 0 | 4 | X | X | X |
| Use centralized and secure audit logging (e.g., syslog over TLS) | | 4 | 4 | 0 | X | X | X |
| Validation and certification (e.g., Common Criteria) of cryptographic modules and other security measures (e.g., media sanitization, access control, etc.) | X | 1 | 1 | 1 | X | X | X |
| Secure autonomous data movement (Subclause 7.7.5) | | | | | | | |
| Configuring policies for data movement should be restricted to authenticated and authorized privileged users | X | 2 | 2 | 2 | X | X | X |
| The individual establishing the configurations should be conversant with the security attributes of both source and destination | X | 3 | 3 | 3 | X | X | X |
| Configuration changes to implement or terminate autonomous data movement should be reflected in the audit log | X | 2 | 2 | 2 | X | X | X |
| All autonomous data movement transactions should be reflected in the audit log of the system conducting the data movement | X | 2 | 2 | 2 | X | X | X |
| As part of autonomous data movement transactions, the integrity of the moved data should be verified (preferably with a cryptographic hash) | | 0 | 3 | 0 | X | X | X |
| Autonomous data movement transactions should not impact the authenticity of the data (e.g., original system metadata like creation date, last accessed, etc. are correctly represented in the moved data) | | 0 | 4 | 2 | X | X | X |
| Autonomous data movement transactions should not negate the immutability or other data preservation controls (e.g., supporting legal holds) | | 0 | 4 | 0 | | X | X |
| Autonomous data movement transactions should not eliminate or weaken encryption controls associated with the data | | 4 | 0 | 0 | | X | X |
| Autonomous data movement transactions that span systems should include data in motion encryption for sensitive and high value data | | 4 | 0 | 0 | | X | X |
| As part of autonomous data movement transactions, the source | | 4 | 0 | 0 | | X | X |

| Controls | Priorities (5 is highest) | | | | Data Sensitivity | | |
|---|------------------------------|---|---|---|---------------------|---|---|
| | S | C | I | A | L | M | H |
| data or storage media should be appropriately sanitized before it is released for re-use | | | | | | | |
| Sanitization performed in conjunction with autonomous data movement should also include verification and some form of proof of sanitization | | 4 | 0 | 0 | | X | X |
| Autonomous data movement transactions should not cause data to cross security domains (e.g., production to development environments) | X | 3 | 3 | 3 | X | X | X |
| Autonomous data movement transactions should not cause data to move to systems with inadequate certifications and accreditations | X | 1 | 1 | 1 | | X | X |
| Autonomous data movement transactions should not cause data to move to systems with inadequate physical security | X | 3 | 3 | 3 | X | X | X |

1

Annex C (informative)

Important security concepts

C.1 Authentication

At a basic level, authentication is the process used to verify the asserted identity of a user or entity through the verification of supplied information (i.e., verify a declared identity). Authentication is a common security issue for the design of systems and protocols, and a wide variety of authentication technologies are available. A common problem is knowing which technology to choose or which of a variety of essentially similar implementations of a given technique to choose.

Authentication can be unidirectional or mutual (bidirectional). For one unidirectional authentication, only the client (user or entity) is typically authenticated to a system (or server). For mutual authentication, both communicating parties authenticate each other. Authentication can also involve more than two parties. In three-party authentication, a trusted third party is the middle man through whom both communicating parties carry on the authentication process.

An authentication factor is a piece of information used to authenticate or verify a user or an entity's identity for security purposes. Three commonly recognized factors are:

- a) Something the authenticating party knows, such as a secret or a password.
- a) Something the authenticating party has, such as a physical hardware token or a key card.
- b) Something the authenticating party is (e.g., a fingerprint, a retinal pattern), does (e.g., a signature), or someplace the party is located at.

A system is said to leverage two-factor authentication (or multi-factor authentication) when it requires at least two of the authentication factors mentioned above. This contrasts with traditional password authentication, which requires only one authentication factor (such as knowledge of a password) in order to gain access to a system. The best authentication mechanisms combine two or more of these factors.

ISO/IEC 27002 (§13.3.1) provides applicable guidance on the use of secret authentication information, which can be summarized as:

- keep secret authentication information confidential; do not share it
- avoid keeping a record of secret authentication information
- change secret authentication information whenever there is any indication of their possible compromise
- select quality passwords when they are used as secret authentication information (minimum length, composition, not easily guessed, do not consist of words included in dictionaries, free of identical consecutive characters, etc.)

Authentication implementations can take on many forms, including:

- 1 — *Local Authentication* - The system needing the authentication service is also the authenticator (i.e., entity
2 making the authentication decision). There is no easy way to synchronize the credential database used for
3 verification, so its usability is limited within larger organizations.
 - 4 — *External Authentication* - The authenticator resides outside of the control and influence of the system needing
5 an authentication decision; further, the authenticator is a trusted, authoritative source.
 - 6 — *Centralized Authentication* - This form of external authentication is designed to support many systems (often
7 heterogeneous) and it often includes redundancy, use of standard protocols, and provides additional useful
8 information (e.g., role identifiers). There is no attempt to make subsequent authentications transparent (i.e.,
9 multiple authentication are often required).
 - 10 — *Single Sign-on (SSO)* - This form of centralized authentication employs a single set of credentials, which are
11 then used transparently to perform subsequent authentications on behalf of the users. In addition, there is
12 typically a close alignment with a centralized authorization system to ensure consistent privileges.
- 13 Many enterprises have centralized their identity management (directory services, NIS, NIS+) and authentication
14 services (e.g., RADIUS, PKI, Kerberos, LDAP, etc.), so there is a natural desire to leverage this infrastructure and
15 the investments made in populating the identity data to help address authentication and authorization.

16 **C.2 Authorization and access control**

17 Authorization is the process by which one determines whether an authenticated party has permission to access a
18 particular resource or service. Although tightly bound, authentication and authorization are two separate
19 mechanisms. Perhaps because of this tight coupling, authentication is sometimes mistakenly thought to imply
20 authorization. Authentication simply validates the identity of a party; authorization defines whether they can
21 perform a certain action. The means to ensure that access to assets and services are granted or restricted
22 appropriately are often implemented as part of an access control mechanism.

23 Numerous access control models and systems (e.g., Bell-LaPadula, Cark-Wilson, etc.) have been developed
24 since the early 1970s. Almost all of these access control models can be formally stated using the following notions
25 and their relationships:

- 26 — *user* - people who interface with the system; the focus is on the human and not the credentials
- 27 — *subjects* - a computer process acting on behalf of a user; they can initiate requests to perform an operation or
28 series of operations on objects
- 29 — *objects* - any resource accessible on a computer system; passive entities that contain or receive information
- 30 — *operations* - an active process invoked by a subject
- 31 — *permissions (or privileges)* - authorizations to perform some action on the system; it typically refers to some
32 combination of object and operation

33 These concepts have been incorporated into a variety of access control policies (rules) and mechanisms,
34 including the following:

- 35 — *Discretionary Access Control (DAC)* - policy permits the granting and revocation of access permissions to be
36 left to the discretion of the individual users
- 37 — *Mandatory Access Control (MAC)* - policy is centrally controlled by a security policy administrator; users do
38 not have the ability to override the policy

— *Role-based Access Control (RBAC)* - policy that assigns permissions to specific roles and roles in turn are assigned users; management of individual user rights becomes a matter of simply assigning the appropriate roles to the user

Implementations of authorization and access control mechanisms can take many forms and have different levels of complexity. A more simplistic implementation is likely to impose few if any controls on an authenticated user (i.e., granting unrestricted access to the system's resources). More sophisticated implementations are likely to impose controls on users based on their membership in groups or holding a particular role. This latter approach is often implemented using Role-based Access Control (RBAC) mechanisms and it is a recommend technique for implementations.

An Access Control List (ACL) is one way of implementing an access control matrix that specifies the operations users or subjects are allowed to perform on an object. In a typical ACL, each entry in the list specifies a subject and an operation; as shown in Table 1, the entry (Alice, Delete) on the ACL for file XYZ gives Alice permission to delete file XYZ.

Table C.1 — ACL for File "XYZ"

| User/Subject | Operations |
|--------------|--------------------|
| Alice | Delete |
| Joe | Read, Write |
| Jan | Execute |

In an ACL-based security model, the system first checks the list for an applicable entry in order to decide whether or not to proceed with the operation a user (subject) requested. The list is often a data structure, usually a table, containing entries that specify individual user or group rights to specific system objects, such as a program, a process, or a file. These entries are sometimes called Access Control Entries (ACE) Each accessible object contains an identifier to its ACL. The privileges or permissions determine specific access rights, such as whether a user can read from, write to, or execute an object. In some implementations an ACE can control whether or not a user, or group of users, can alter the ACL on an object.

It is also possible for the users (subjects) to be grouped so that the ACL would contain the name of the group rather than individual users. This makes the management of ACLs much easier as revoking a user's permissions would involve removing them from membership in the group rather than modifying the ACL itself.

Protection bit mechanisms are similar to ACLs; however, bits are associated with an object rather associating users and operations entries. Protection bit mechanisms are commonly implemented in UNIX operating systems and are used to divide users into different categories, typically user (self), group, and other. The access control system regulates access to a file by associating read (r), write (w), or execute (x) operations with each of the categories of users.

As an access control mechanism, protection bit mechanisms have an assortment of issues, including:

- The user who created a file is the owner, by default.
- The owner of a file is typically the only one (besides the *superuser* or *administrator*) who can modify the protection bits.
- There is only one group available for each file
- The system administrator controls group membership; as membership within groups changes, so will the capabilities of users to access files.

— The system cannot grant access to an object on an individual basis.

Access control decisions are often determined by the roles individual users take on as members of an organization. This includes the specification of duties, responsibilities, and qualifications. For example, the roles an individual associated with a hospital can assume include doctor, nurse, clinician, and pharmacist. Roles in a bank include teller, loan officer, and accountant. Roles can also apply to military systems; for example, target analyst, situation analyst, and traffic analyst are common roles in tactical systems.

C.3 Self-Encrypting Drives (SED)

Many storage manufacturers have released storage elements with integrated encryption and access control capabilities, also known as Self-Encrypting Drives (SEDs). SEDs that feature always-on encryption substantially reduce the likelihood that unencrypted data is inadvertently retained on the device. The end user cannot turn off the encryption capabilities (so all previous data in the designated areas is encrypted). SEDs typically encrypt most or all of the user-addressable area, with the potential exception of clearly identified areas dedicated to the storage of pre-boot applications and associated data.

A significant additional benefit of SEDs is the opportunity to tightly couple the controller and storage media so that the device can directly address the location where any keys are stored, whereas solutions that depend only on the abstracted user access interface through software may not be able to directly address those areas.

SEDs are also well equipped to perform a special form of sanitization, known as cryptographic erase (see A.3), which can be performed with high assurance much faster than with other sanitization techniques. Typically, cryptographic erase is executed in seconds (as compared to hours or days for some alternative media sanitization operations). This is especially important as storage devices get larger and sanitization becomes more cumbersome and time consuming when non-cryptographic approaches are applied.

C.4 Sanitization

ICT systems capture, process, and store information using a wide variety of media. This information is not only located on the intended storage media but also on devices used to create, process, or transmit this information. These media can require special disposition in order to mitigate the risk of unauthorized disclosure of information and to ensure its confidentiality. Efficient and effective management of information that is created, processed, and stored by an ICT system throughout its life, from inception through disposition, is a primary concern of an ICT system owner and the custodian of the data.

With the use of increasingly sophisticated encryption, an attacker wishing to gain access to an organization's sensitive information is forced to look outside the system itself for that information. One avenue of attack is the recovery of supposedly deleted data from media. These residual data may allow unauthorized individuals to reconstruct data and thereby gain access to sensitive information. Sanitization can be used to thwart this attack by ensuring that deleted data cannot be easily recovered.

When storage media are transferred, become obsolete, or are no longer usable or required by an information system, it is important to ensure that residual magnetic, optical, electrical, or other representation of data that has been deleted is not easily recoverable. Sanitization refers to the general process of removing data from storage media, such that there is reasonable assurance that the data cannot be easily retrieved and reconstructed.

Cryptographic erase features (see A.3), a technique that can provide significant benefits in both timeliness and assurance, will likely be widely available in the near future based on support from nearly all major storage vendors. Cryptographic erase could provide substantial value (if well defined) by:

— facilitating rapid eradication of sensitive data (in seconds versus hours or days)

- 1 — reducing the wear on the storage device (therefore potentially extending the life of the device)
- 2 — reducing the amount of man-hours expended performing sanitization
- 3 — addressing media types that may be impractical to address using legacy degaussing and destruction
- 4 techniques
- 5 For all devices supporting encryption where cryptographic erase is intended for use to purge the media (including
- 6 SEDs, mobile devices, and other devices), the level of assurance depends (in large part) on the following:
- 7 a) The assurance that all target data is either encrypted or (for any target data not encrypted) able to be
- 8 effectively sanitized using media-specific techniques.
- 9 b) The level of entropy of the MEK.
- 10 c) If the key sanitized during cryptographic erase is a key that wraps the MEK (and not the MEK itself), the
- 11 strength of the wrapping mechanism(s) and entropy of the wrapping key(s) to be sanitized.
- 12 d) The strength of the encryption algorithm used to encrypt the data, including mode of operation and assurance
- 13 of correct implementation.
- 14 e) The level of difficulty in retrieving the MEK after sanitization, plus any effort to unwrap the key (if it was stored
- 15 wrapped with another value).
- 16 NOTE In some cases, the Media Encryption Key may have been stored in the clear because the sole purpose of media
- 17 encryption may have been to support cryptographic erase.
- 18 Mobile devices (and devices other than SEDs) can also support strong encryption capabilities. The decision
- 19 regarding whether to rely upon cryptographic erase to purge the media on those devices depends, in part, on
- 20 whether all sensitive data is encrypted on the device. If encryption was enabled after sensitive data was stored on
- 21 the device, or if it is unknown whether sensitive data was stored on the device prior to encryption, cryptographic
- 22 erase should not be relied upon as a Purge mechanism (see A.1).

23 C.5 Logging

24 Within storage systems and infrastructure, there are a wide range of transactions or events that can result in the
 25 generation of event log entries (messages). In addition, these event log entries have to be recorded in some
 26 manner for event logging. From a security or compliance perspective, it is important to capture those event log
 27 entries necessary to demonstrate proof of operations (e.g., encryption and retention), enforcement of
 28 accountability and traceability, meeting evidentiary requirements, and adequate monitoring of systems. This
 29 subset of general event logging is commonly called audit logging.

30 Not all event log entries are created equal, as some may only be useful for debugging purposes, provide system
 31 health status, warn of minor configuration problems, etc. From an audit logging perspective the management
 32 events (i.e., what a human did) are always of interest, the data access events are usually of limited interest
 33 (except in situations where critical files and directories need to be tightly monitored), and control events are
 34 typically of the least interest (they can provide useful information during root-cause analysis after an incident).

35 In addition, audit logging often requires the event entries of interest to be handled differently and separately from
 36 most other event log entries generated by a device. This special handling can be accomplished by having the
 37 devices send the audit log entries to special log infrastructure or they can be culled out of the general log stream,
 38 using a log filtering mechanism (a more challenging approach because it requires all the event entries of interest
 39 to be known a priori). Another aspect of this special handling is that an organization often has to demonstrate that

it is monitoring (e.g., generating alerts for anomalous events) and reporting; these actions usually require some form of centralized logging infrastructure beyond simple collectors.

C.6 N_Port ID Virtualization (NPIV)

As a starting point, a Fibre Channel (FC) port is a hardware pathway into and out of a node that performs data communications over an FC link (sometimes called a channel). FC defines many different types of ports, but the following are relevant to NPIV:

- **N_Port:** A network or node port used to connect a node to a FC switch. This could be an HBA (Host Bus Adapter) in a server or a target port on a storage array.
- **F_Port:** A switch port used to connect the FC fabric to a node (N_Port), which is usually a server's HBA or a storage array's target port.
- **E_Port:** An extender port used to connect (cascade) FC switches together; the connection between two E_Ports forms an Inter-Switch Link (ISL).

Normally, an N_Port would have a single N_Port_ID associated with it. This N_Port_ID is a 24-bit address assigned by the Fibre Channel switch during the Fabric Login (FLOGI) process, which occurs after a link is operational.

NOTE The N_Port_ID is not the same as the World Wide Port Name (WWPN), although there is typically a one-to-one relationship between WWPN and N_Port_ID. Thus, for any given physical N_Port, there would be exactly one WWPN and one N_Port_ID associated with it.

NPIV enables a single physical N_Port to have multiple WWPNs, and therefore multiple N_Port_IDs, associated with it. After the normal FLOGI process, an NPIV-enabled physical N_Port can subsequently issue additional commands to register more WWPNs and receive more N_Port_IDs (one for each WWPN). The Fibre Channel switch must also support NPIV, as the F_Port on the other end of the link would "see" multiple WWPNs and multiple N_Port_IDs coming from the host and must know how to handle this behaviour.

Once all the applicable WWPNs have been registered, each of these WWPNs can be used for SAN zoning or LUN presentation. There is no distinction between the physical WWPN and the virtual WWPNs; they all behave in exactly the same manner and they can be used in exactly the same ways.

Each N_Port_ID created by NPIV consumes resources in the server, network fabric and storage for state related to that N_Port_ID. In environments with a large number of virtual hosts, creating N_Port_IDs for every virtual host can cause scaling problems due to limitations on these resources. For such environments, it may be possible to limit use of NPIV to creating only the N_Port_IDs that are necessary to provide isolation among larger domains (e.g., the set of virtual hosts for a single organization or a single tenant of a service provider).

Bibliography

- 2 [01] ISO Guide 73:2009, Risk management - Vocabulary
- 3 [02] ISO 7498-2:1989, *Information technology — Open Systems Interconnection — Basic Reference Model —*
4 *Part 2: Security Architecture*
- 5 [03] ISO 16609:2004, *Banking — Requirements for message authentication using symmetric techniques*
- 6 [04] ISO/PAS 22399:2007, *Societal security — Guideline for incident preparedness and operational continuity*
7 *management*
- 8 [05] ISO/IEC 10116:2006, *Information technology — Security techniques — Modes of operation for an n-bit*
9 *block cipher*
- 10 [06] ISO/TR 10255:2009, *Document management applications — Optical disk storage technology,*
11 *management and standards*
- 12 [07] ISO/TR 18492:2005, *Long-term preservation of electronic document-based information*
- 13 [08] ISO 16175-1:2010, *Information and documentation – Principles and functional requirements for records in*
14 *electronic office environments – Part 1: Overview and statement of principles*
- 15 [09] ISO 16175-2:2011, *Information and documentation – Principles and functional requirements for records in*
16 *electronic office environments – Part 2: Guidelines and functional requirements for digital records*
17 *management systems*
- 18 [10] ISO 16175-3:2010, *Information and documentation – Principles and functional requirements for records in*
19 *electronic office environments – Part 3: Guidelines and functional requirements for records in business*
20 *systems*
- 21 [11] ISO/IEC 11770 (all parts), *Information technology — Security techniques — Key management*
- 22 [12] ISO/IEC 2ndCD 17788, *Information technology — Distributed application platforms and services – Cloud*
23 *computing – Overview and vocabulary*
- 24 [13] ISO/IEC 17826:2012, *Information technology — Cloud Data Management Interface (CDMI)*
- 25 [14] ISO/IEC 19790:2006, *Information technology — Security techniques — Security requirements for*
26 *cryptographic modules*
- 27 [15] ISO/IEC 24759:2008, *Information technology — Security techniques — Test requirements for*
28 *cryptographic modules*
- 29 [16] ISO/IEC 24775:2007, *Information technology — Storage management*
- 30 [17] ISO/IEC 27003:2010, *Information technology — Security techniques — Information security management*
31 *systems implementation guidance*
- 32 [18] ISO/IEC 27031:2011, *Information technology — Security techniques — Guidelines for information and*
33 *communication technology readiness for business continuity*

- 1 [19] ISO/IEC 27033-1:2009, *Information technology — Security techniques — Network security — Part 1:*
2 *Overview and concepts*
- 3 [20] ISO/IEC 27033-2, *Information technology — Security techniques — Network security — Part 2:*
4 *Guidelines for the design and implementation of network security*
- 5 [21] ISO/IEC 27033-3:2010, *Information technology — Security techniques — Network security — Part 3:*
6 *Reference networking scenarios — Threats, design techniques and control issues*
- 7 [22] ISO/IEC 27037:2012, *Information technology — Security techniques — Guidelines for identification,*
8 *collection, acquisition, and preservation of digital evidence*
- 9 [23] ISO/IEC WD 27044, *Information technology – Security techniques – Guidelines for security information*
10 *and event management (SIEM)*
- 11 [24] IEEE/ISO/IEC 24765-2010, *Systems and software engineering -- Vocabulary*
- 12 [25] IEEE 1619-2007, *IEEE Standard for Wide-Block Encryption for Shared Storage Media*
- 13 [26] IEEE 1619.1-2007, *IEEE Standard for Authenticated Encryption with Length Expansion for Storage*
14 *Devices*
- 15 [27] IEEE 1619.2-2010, *IEEE Standard for Cryptographic Protection of Data on Block-Oriented Storage*
16 *Devices*
- 17 [28] IETF RFC 1813 *NFS Version 3 Protocol Specification*
- 18 [29] IETF RFC 3530 *Network File System (NFS) version 4 Protocol*
- 19 [30] IETF RFC 3723 *Securing Block Storage Protocols over IP*
- 20 [31] IETF RFC 3720 *Internet Small Computer Systems Interface (iSCSI)*
- 21 [32] IETF RFC 3821 *Fibre Channel Over TCP/IP (FCIP)*
- 22 [33] IETF RFC 4172 *iFCP - A Protocol for Internet Fibre Channel Storage Networking*
- 23 [34] IETF RFC 5246 *The Transport Layer Security (TLS) Protocol Version 1.2*
- 24 [35] IETF RFC 5661 *Network File System (NFS) Version 4 Minor Version 1 Protocol*
- 25 [36] IETF RFC 5663 *Parallel NFS (pNFS) Block/Volume Layout*
- 26 [37] IETF RFC 6071 *IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap*
- 27 [38] IETF RFC 3195 *Reliable Delivery for syslog*
- 28 [39] IETF RFC 5424 *The Syslog Protocol*
- 29 [40] IETF RFC 5425 *TLS Transport Mapping for Syslog*
- 30 [41] IETF RFC 5426 *Transmission of Syslog Messages over UDP*
- 31 [42] IETF RFC 5427 *Textual Conventions for Syslog Management*

- 1 [43] IETF RFC 5848 *Signed Syslog Messages*
- 2 [44] IETF RFC 6012 *Datagram Transport Layer Security (DTLS) Transport Mapping for Syslog*
- 3 [45] IETF RFC 6587 *Transmission of Syslog Messages over TCP*
- 4 [46] ANSI INCITS 388 (all parts)-2011, *Information Technology – Storage Management*
- 5 [47] ANSI INCITS 424–2007, *Fibre Channel – Framing and Signaling-2 (FC-FS-2)*
- 6 [48] ANSI INCITS 400-204, *Information technology – SCSI Object-based Storage Device Commands (OSD)*
- 7 [49] ANSI INCITS 458-2011, *Information technology – SCSI Object-Based Storage Device Commands – 2*
- 8 (OSD-2)
- 9 [50] ANSI INCITS 462-2010, *Information Technology – Fibre Channel - Backbone – 5 (FC-BB-5)*
- 10 [51] ANSI INCITS 482-2012, *Information Technology – ATA/ATAPI Command Set – 2 (ACS-2)*
- 11 [52] ANSI INCITS 496-2012, *Information Technology – Fibre Channel - Security Protocols – 2 (FC-SP-2)*
- 12 [53] ANSI INCITS 1799-D, *Information Technology – SCSI Block Commands – 3 (SBC-3)*
- 13 [54] NIST FIPS 140-2, *Security Requirements for Cryptographic Modules*
- 14 [55] NIST FIPS 197, *Advanced Encryption Standard*
- 15 [56] NIST Special Publication 800-38A, *Recommendation for Block Cipher Modes of Operation: Three*
- 16 *Variants of Ciphertext Stealing for CBC Mode*
- 17 [57] NIST Special Publication 800-38C, *Recommendation for Block Cipher Modes of Operation: the CCM*
- 18 *Mode for Authentication and Confidentiality*
- 19 [58] NIST Special Publication 800-38D, *Recommendation for Block Cipher Modes of Operation:*
- 20 *Galois/Counter Mode (GCM) and GMAC*
- 21 [59] NIST Special Publication 800-38E, *Recommendation for Block Cipher Modes of Operation: The XTS-AES*
- 22 *Mode for Confidentiality on Storage Devices*
- 23 [60] NIST Special Publication 800-57 Part 1, *Recommendation for Key Management: Part 1: General*
- 24 *(Revision 3)*
- 25 [61] NIST Special Publication 800-57 Part 2, *Recommendation for Key Management: Part 2: Best Practices*
- 26 *for Key Management Organization*
- 27 [62] NIST Special Publication 800-67, *Recommendation for the Triple Data Encryption Algorithm (TDEA)*
- 28 *Block Cipher*
- 29 [63] NIST Special Publication 800-88 Revision 1 (draft), *Media Sanitization*
- 30 [64] Storage Networking Industry Association (SNIA), *Storage Management Initiative – Specification (SMI-S),*
- 31 *Version 1.5, Architecture Book*, http://www.snia.org/tech_activities/standards/curr_standards/smi
- 32 [65] Trusted Computing Group, *Enterprise Security Subsystem Class (SSC)*

- 1 [66] Trusted Computing Group, *Opal Security Subsystem Class (SSC)*
- 2 [67] OASIS, *Key Management Interoperability Protocol Specification Version 1.1*
- 3 [68] OASIS, *Key Management Interoperability Protocol Profiles Version 1.1*

4

5

Index

- 2 3 access control, 13, 31, 54, 94, 100
- 4 Access Control Entry, 7
- 5 Access Control List, 7, 22, 23, 32, 33, 43, 80, 84, 88, 99
- 6 accessibility, 17
- 7 accidents, 47
- 8 accountability, 29, 42, 49, 55, 57, 82, 88, 92, 101
- 9 accreditation, 51, 92
- 10 ACE. *See* Access Control Entry
- 11 ACL. *See* Access Control List
- 12 Advanced Encryption Standard, 7
- 13 Advanced Technology Attachment, 7, 38, 65, 67
- 14 AES. *See* Advanced Encryption Standard
- 15 arbitrated loop, 21
- 16 ATA. *See* Advanced Technology Attachment
- 17 attack
 - 18 history, 48, 91
 - 19 surface, 12
 - 20 vector, 24
- 21 attacks, 27, 43, 44, 45, 47, 51, 82, 88, 89, 92, 100
- 22 adversarial, 47, 90
- 23 denial of service, 15
- 24 indirect, 27, 31, 81, 83
- 25 intentional, 17
- 26 laboratory, 5, 59
- 27 malicious, 17
- 28 malware, 15
- 29 non-malicious, 17
- 30 slow, 48, 91
- 31 audit logging, 26, 28, 29, 55, 57, 94, 95, 101
- 32 audit trail, 37, 87
- 33 authentication, 21, 23, 48, 97
 - 34 centralized, 26, 31, 47, 81, 83, 98
 - 35 external, 34, 85, 98
 - 36 factor, 97
 - 37 multi-factor, 4, 26, 27, 81, 82
 - 38 strong, 6, 26, 27, 33, 81, 82, 84
- 39 authenticity, 28, 35, 54, 55, 85, 94
- 40 autonomous data movement, 13, 14, 57, 58, 95, 96
- 41 availability, 4, 13, 14, 15, 17, 25, 27, 44, 45, 47, 89
 - 42 application, 19
 - 43 data, 45, 46, 90
 - 44 designs, 46, 90
- 45 backups, 13, 14, 17, 19, 31, 36, 41, 45, 46, 51, 57, 74, 79, 90
- 46 BC. *See* Business Continuity
- 47 breach notifications, 18
- 48 Business Continuity, 7, 13, 14, 19, 42, 47, 49, 50, 87, 90, 92
- 49 Business Continuity Management, 46
- 50 CAS. *See* Content Addressable Storage
- 51 CDMI. *See* Cloud Data Management Interface
- 52 CDP. *See* Continuous Data Protection
- 53 certificate of sanitization, 37, 87
- 54 chain of custody, 28, 37, 49, 55, 56, 92, 94
- 55 Challenge Handshake Authentication Protocol, 7
- 56 CHAP. *See* Challenge Handshake Authentication Protocol
- 57 CIFS. *See* Common Internet File System
- 58 Cloud Data Management Interface, 7, 34, 44, 85, 88, 103
- 59 cloud storage, 14, 23, 34, 85
- 60 Common Internet File System, 7, 24
- 61 **compression**, 2, 3, 14, 41, 42, 48, 49, 54, 87, 91, 93
- 62 confidentiality, 27
- 63 Content Addressable Storage, 7, 12, 35, 86
- 64 Continuous Data Protection, 7, 12, 36, 45, 46, 51
- 65 **cryptographic erase**, 2, 37, 38, 56, 61, 62, 64, 65, 66, 67, 69, 70,
 - 66 71, 75, 76, 87, 95, 100, 101
- 67 cryptographic hash, 35
- 68 **cryptoperiod**, 2, 41, 49, 51, 87, 92
- 69 curation, 48
- 70 DAC. *See* Discretionary Access Control
- 71 DAS. *See* Direct Attached Storage
- 72 data at rest, 2, 13, 14, 21, 32, 33, 34, 40, 51, 52, 84
- 73 data authenticity, 13, 49, 56, 92, 94
- 74 data breach, vi, 2, 5, 15, 16, 37, 51, 53, 56
- 75 data confidentiality, 39
- 76 data corruption, 17
- 77 data destruction, 17
- 78 **data in motion**, 2, 13, 14, 21, 40, 87
- 79 **data integrity**, 2, 23
- 80 Data Lifecycle Management, 57
- 81 data protection, vi, 14, 19, 25, 39
 - 82 mechanisms, 17
 - 83 methods, 40
 - 84 strategy, 14
 - 85 systems, 12
- 86 data reduction, 13, 41, 50, 87, 91, 92
- 87 technologies, 14, 41, 48, 49
- 88 data sensitivity, 43
- 89 DDoS. *See* Distributed Denial of Service
- 90 **deduplication**, 3, 14, 41, 42, 48, 49, 54, 56, 87, 91, 93, 94
- 91 defence in depth, 42
- 92 **degauss**, 3
- 93 Denial of Service, 7, 15, 16
- 94 DH-CHAP. *See* Diffie Hellman – Challenge Handshake
 - 95 Authentication Protocol
- 96 Diffie Hellman – Challenge Handshake Authentication Protocol,
 - 97 7
- 98 digital signature, 39, 40
- 99 Direct Attached Storage, 7, 12, 18, 19, 79
- 100 Disaster Recovery, 7, 13, 14, 19, 36, 41, 42, 47, 48, 49, 50, 54,
 - 101 87, 90, 92
- 102 Disaster Recovery Planning, 46
- 103 Discretionary Access Control, 7
- 104 Distributed Denial of Service, 7, 15

- 1 DNS. *See* Domain Name System
- 2 Domain Name System, 7, 27, 29, 81, 83
- 3 DoS. *See* Denial of Service
- 4 DR. *See* Disaster Recovery
- 5 EHR. *See* Electronic Healthcare Record
- 6 Electronic Healthcare Record, 8, 36, 86
- 7 **Electronically Stored Information**, 3, 6, 8, 12
- 8 Encapsulating Security Payload, 8, 21
- 9 encryption, 13, 14, 21, 26, 36, 39, 40, 41, 47, 48, 49, 50, 51, 53,
10 54, 86, 87, 92, 93, 100
11 at rest, 13, 14, 23, 32, 33, 34, 40, 46, 51, 56, 57, 84, 87, 92,
12 95
13 CDP, 46
14 data at rest, 52
15 in motion, 13, 23, 39, 40, 46, 51, 57, 87, 90, 92, 95
16 key, 16, 25, 41, 45, 53, 54, 76, 87, 90
17 proof of, 18, 40, 51, 101
- 18 entity authentication, 26
- 19 ESI. *See* Electronically Stored Information
- 20 ESP. *See* Encapsulating Security Payload
- 21 Ethernet, 23, 80
- 22 event logging, 28, 29, 82, 101
- 23 evidentiary data, 56
- 24 fault-tolerance, 46, 90
- 25 FC. *See* Fibre Channel
- 26 FCIP. *See* Fibre Channel over TCP/IP
- 27 FCoE. *See* Fibre Channel over Ethernet
- 28 FCP. *See* Fibre Channel Protocol
- 29 FCS. *See* Fixed Content Storage
- 30 FC-SP. *See* Fibre Channel – Security Protocol
- 31 FDE. *See* Full Disk Encryption
- 32 **Fibre Channel**, 3, 8, 19, 21, 22, 23, 27, 31, 81, 83
- 33 Fibre Channel – Security Protocol, 8, 19, 22, 23, 27, 31, 39, 40,
34 43, 80, 81, 83, 105
- 35 Fibre Channel over Ethernet, 8, 13, 23
- 36 Fibre Channel over TCP/IP, 8, 22, 23, 80, 104
- 37 **Fibre Channel Protocol**, 3, 8, 19, 22, 23, 26, 31, 80
- 38 Fixed Content Storage, 8
- 39 Full Disk Encryption, 8, 18, 79
- 40 Galois/Counter Mode, 8
- 41 gateways, 21
- 42 GCM. *See* Galois/Counter Mode
- 43 Hard Disk Drive, 8, 12, 13, 18, 38, 41, 52, 64, 66, 67, 68
- 44 HBA. *See* Host Bus Adapter, *See* Heat Assisted Magnetic
45 Recording
- 46 HDD. *See* Hard Disk Drive
- 47 Heat Assisted Magnetic Recording, 8, 65, 66
- 48 Host Bus Adapter, 8
- 49 hypervisor, 52, 53, 93
- 50 ICT Readiness for Business Continuity, 47
- 51 IKE. *See* Internet Key Exchange
- 52 ILM. *See* Information Lifecycle Management
- 53 immutability, 51, 57, 92
- 54 InfiniBand, 13, 19, 20
- 55 Information Lifecycle Management, 8, 57
- 56 Information Security Management System, vi
57 integrity, 27, 46, 48, 55, 94
58 preserving, 47
59 Internet Key Exchange, 8, 39, 40, 87, 104
60 Internet Protocol Security, 9, 14, 21, 23, 24, 26, 31, 35, 39, 40,
61 80, 81, 83, 85, 87, 104
62 Internet Small Computer Systems Interface, 9, 19, 22, 26, 31, 44,
63 80, 83, 88
64 devices, 21
65 environments, 21
66 initiator, 31, 83
67 initiators, 21
68 interfaces, 22, 80
69 network, 22
70 security, 31
71 Internet Storage Name Service, 9, 31, 83
72 IPsec. *See* Internet Protocol Security
73 iSCSI. *See* Internet Small Computer Systems Interface
74 iSNS. *See* Internet Storage Name Service
75 ISO 16175-1, 48
76 ISO 16175-2, 48
77 ISO 16175-3, 48
78 ISO 16609, 2
79 ISO 7498-2, 2
80 ISO/IEC 10116, 39, 40, 103
81 ISO/IEC 11770, 39, 40, 103
82 ISO/IEC 14776-372, 3
83 ISO/IEC 15408, 51
84 ISO/IEC 17788, 1
85 ISO/IEC 17826, 34, 103
86 ISO/IEC 19790, 51, 76, 103
87 ISO/IEC 24759, 76, 103
88 ISO/IEC 24775, 103
89 ISO/IEC 27000, 1, 2
90 ISO/IEC 27001, vi, 1, 42, 54, 94
91 ISO/IEC 27002, vi, 1, 18, 42, 54, 94, 97
92 ISO/IEC 27003, 103
93 ISO/IEC 27005, vi, 1, 2, 15, 42
94 ISO/IEC 27031, 47, 103
95 ISO/IEC 27033, 19
96 ISO/IEC 27033-1, 4, 104
97 ISO/IEC 27033-2, 25, 43, 104
98 ISO/IEC 27033-3, 104
99 ISO/IEC 27037, 104
100 ISO/IEC/IEEE 24765, 5
101 ISO/PAS 22399, 46
102 ISO/TR 10255, 47
103 ISO/TR 12033, 2
104 ISO/TR 18492, 48
105 isolation
106 logical, 21
107 physical, 21
108 Kerberos, 24, 31, 32, 33, 80, 84, 98
109 key escrow, 53, 93
110 key management, 13, 14, 26, 39, 40, 41, 43, 49, 51, 53, 54, 56,
111 87, 88, 92, 93, 95, 103, 105

- 1 Key Management Interoperability Protocol, 9, 40, 41, 56, 87, 95,
- 2 106
- 3 KMIP. *See* Key Management Interoperability Protocol
- 4 LAN. *See* Local Area Network
- 5 LBA. *See* Logical Block Address
- 6 LDAP. *See* Lightweight Directory Access Protocol
- 7 least privilege, 22, 27, 55, 57, 80, 82, 94, 95
- 8 Lightweight Directory Access Protocol, 9, 32, 84, 98
- 9 Local Area Network, 9, 12, 19, 22, 23, 27, 31, 32, 44, 80, 81, 88
- 10 logging, 48, 91
- 11 CDMI, 34, 85
- 12 policy, 30, 82
- 13 logging protocols, 28
- 14 Logical Block Address, 9
- 15 logical unit, 13, 31, 36
- 16 Logical Unit Number, 9, 20
- 17 loss of media, 15
- 18 LUN. *See* Logical Unit Number
- 19 LUN masking, 20, 31, 83
- 20 **malware**, 4, 15, 30, 32, 33, 51, 84
- 21 malware protection, 30, 34, 45, 51, 83, 92
- 22 **Mean Time Between Failures**, 4, 9, 44, 45
- 23 Mean Time To Failure, 9, 44
- 24 Mean Time To Repair, 4, 9, 44, 45
- 25 Media Encryption Key, 2, 9, 75, 76, 101
- 26 media sanitization, 13, 19, 34, 37, 38, 54, 56, 57, 72, 73, 75, 79,
- 27 94, 95, 100
- 28 MEK. *See* Media Encryption Key
- 29 *metadata*, 3, 4, 17, 23, 34, 35, 48, 51, 55, 56, 94
- 30 modem ports, 27, 81
- 31 modes of operations, 41, 87
- 32 MTBF. *See* Mean Time Between Failure
- 33 MTTF. *See* Mean Time to Failure
- 34 MTTR. *See* Mean Time to Repair
- 35 multi-factor authentication, 4, 26, 27, 81, 82
- 36 multiple data paths, 19
- 37 **multi-tenancy**, 5, 52, 56, 93
- 38 N_Port ID Virtualization, 9, 22, 53, 93
- 39 NAS. *See* Network Attached Storage
- 40 natural disasters, 47
- 41 **Network Attached Storage**, 5, 9, 12, 23
- 42 Network File System, 9
- 43 Network Time Protocol, 9, 27, 29, 81
- 44 NFS. *See* Network File System
- 45 non-repudiation, 55, 94
- 46 NPIV. *See* N_Port ID Virtualization
- 47 NTP. *See* Network Time Protocol
- 48 Object-based Storage Device, 10, 12, 23, 34, 35, 85, 105
- 49 OSD. *See* Object-based Storage Device
- 50 Parallel Network File System, 10, 23, 33, 84, 104
- 51 passwords, 16, 72, 83, 97
- 52 strong, 25, 81
- 53 patches, 30, 83
- 54 path failover, 45
- 55 PCIe. *See* Peripheral Component Interconnect Express
- 56 Peripheral Component Interconnect Express, 10, 20
- 57 Personally Identifiable Information, 10, 16, 36, 54, 86, 94
- 58 physical isolation, 23
- 59 PII. *See* Personally Identifiable Information
- 60 PKI. *See* Public Key Infrastructure
- 61 pNFS. *See* Parallel Network File System
- 62 **point of encryption**, 5, 14, 40, 49, 50, 92
- 63 point-in time copies, 56
- 64 point-to-point, 21
- 65 policy, 54
- 66 logging, 28, 30, 82
- 67 port binding, 20
- 68 privacy, 18, 48, 52, 93
- 69 proof, 46, 90
- 70 of encryption, 18, 40, 51, 101
- 71 of sanitization, 18, 37, 38, 58
- 72 provenance, 49, 92
- 73 Public Key Infrastructure, 10, 43, 88, 98
- 74 RADIUS. *See* Remote Authentication Dial In User Service
- 75 RAID. *See* Redundant Array of Independent Disks
- 76 Random Number Generator, 10, 41, 87
- 77 RBAC. *See* Role-based Access Control
- 78 recovery plan, 54, 93
- 79 redundancy, 44, 47, 89, 91
- 80 Redundant Array of Independent Disks, 10, 36, 45, 51
- 81 redundant components, 45
- 82 regulatory requirements, 49, 55, 91
- 83 **reliability**, 5, 17, 44, 45, 46, 47, 89, 90
- 84 Remote Authentication Dial In User Service, 10, 26, 32, 84, 98
- 85 replication, 45, 46, 51, 57, 90, 95
- 86 on-line, 36
- 87 out of region, 49
- 88 remote, 41
- 89 security, 46
- 90 Representational State Transfer, 10
- 91 resilience, 47, 91
- 92 resiliency, 13, 46
- 93 REST. *See* Representational State Transfer
- 94 retention
- 95 cloud data, 34
- 96 data, 18, 50, 51, 54, 55, 92, 94
- 97 drivers, 48
- 98 event log data, 30
- 99 long-term, 12, 47, 91
- 100 medium-term, 48, 91
- 101 periods, 48
- 102 policy, 28, 34, 82, 85
- 103 short-term, 48, 91
- 104 risk management, vi, 15
- 105 RNG. *See* Random Number Generator
- 106 Role-based Access Control, 10, 99
- 107 SAN. *See* Storage Area Network
- 108 sanitization, 14, 15, 25, 36, 52, 58, 85, 100, 101
- 109 proof of, 18, 37, 38, 58
- 110 verification, 38, 87
- 111 sanitize, 2, 5, 14, 36, 37, 38, 39, 55, 59, 60, 61, 67, 68, 69, 75, 76,
- 112 86, 94, 101

- 1 SAS. *See* Serial Attached SCSI
- 2 SCSI. *See* Small Computer System Interface
- 3 secret-sharing, 48
- 4 secure initialization, 44, 89
- 5 **secure multi-tenancy**, 5, 14, 56
- 6 Secure Shell, 10
- 7 security domain, 28, 43, 44, 88, 89
- 8 Security Information and Event Management, 10, 29, 82, 104
- 9 **security strength**, 5, 7, 49, 50, 56, 75, 92, 95
- 10 SED. *See* Self-Encrypting Drives
- 11 Self-Encrypting Drives, 10, 18, 40, 51, 75, 79, 92, 100, 101
- 12 separation of duties, 51, 92
- 13 Serial Attached SCSI, 10, 20
- 14 Server Message Block, 10, 24
- 15 Service Locator Protocol, 10, 27, 31, 81, 83
- 16 SIEM. *See* Security Information and Event Management
- 17 silent corrections, 17
- 18 Simple Network Management Protocol, 10, 27, 56, 95
- 19 **single point of failure**, 6, 13, 21, 44, 45, 89
- 20 SLP. *See* Service Locator Protocol
- 21 Small Computer System Interface, 10
- 22 SMB. *See* Server Message Block
- 23 SMI-S. *See* Storage Management Initiative – Specification
- 24 snapshots, 14
- 25 SNMP. *See* Simple Network Management Protocol
- 26 Solid State Drive, 10, 13, 38, 59, 68, 70, 75
- 27 Solid State Hard Drive, 11, 13, 59
- 28 SSD. *See* Solid State Drive
- 29 SSH. *See* Secure Shell
- 30 SSHD. *See* Solid State Hard Drive
- 31 storage
 - 32 management, 13, 14, 16, 19, 25, 26, 28, 56, 81, 82, 95, 103,
 - 33 105
 - 34 **Storage Area Network**, 6, 10, 19, 20, 21, 31, 43
 - 35 storage device, 3, 5, 6, 12, 13, 17, 19, 20, 21, 28, 31, 34, 38, 40,
 - 36 41, 43, 44, 54, 57, 62, 65, 66, 69, 70, 71, 74, 75, 76, 79, 89, 93,
 - 37 100, 101
 - 38 storage element, 3, 6, 12, 17, 18, 19, 21, 25, 43, 49, 57
 - 39 Storage Management Initiative – Specification, 10, 26, 56, 95,
 - 40 105
 - 41 storage security, vi, 1, 28, 41, 42
 - 42 architecture, 42
 - 43 concepts, 1
 - 44 controls, 18, 51
 - 45 design, 44
 - 46 design and implementation, 42, 45
 - 47 design rules, 43
 - 48 introduction, 13
 - 49 relevance, 1
 - 50 risk, 15
 - 51 scope, 1
 - 52 services, 36
 - 53 threats, 42
 - 54 switched fabric, 21
 - 55 syslog, 28, 29, 57, 95, 104, 105
 - 56 threats, 1, 12, 13, 15, 25, 42, 43, 48, 51, 77, 88
 - 57 storage security, 42
 - 58 timeliness, 17
 - 59 timestamp, 29
 - 60 TLS. *See* Transport Layer Security
 - 61 traceability, 49, 55, 57, 92
 - 62 Transport Layer Security, 11, 26, 27, 28, 34, 39, 40, 56, 57, 81,
 - 63 82, 85, 87, 95, 104
 - 64 unauthorized modifications, 17
 - 65 Universal Serial Bus, 11
 - 66 USB. *See* Universal Serial Bus
 - 67 Virtual Local Area Network, 11, 22, 23, 27, 43, 44, 80, 81, 89
 - 68 Virtual Machine, 11, 52, 53, 93
 - 69 Virtual Private Network, 11, 27, 82
 - 70 virtual storage, 52, 93
 - 71 Virtual Storage Area Network, 11, 44, 89
 - 72 virtualization, 13, 14, 19, 52
 - 73 security, 13
 - 74 server, 36, 52, 53, 93
 - 75 storage, 52, 93
 - 76 VLAN. *See* Virtual Local Area Network
 - 77 VM. *See* Virtual Machine
 - 78 VPN. *See* Virtual Private Network
 - 79 VSAN. *See* Virtual Storage Area Network
 - 80 **weak key**, 7, 41, 87
 - 81 World Wide Name, 11, 20, 31, 83
 - 82 WORM. *See* Write Once Read Many
 - 83 Write Once Read Many, 11, 30, 51, 54, 83, 92, 94
 - 84 WWN. *See* World Wide Name
 - 85 XTS-AES, 41
 - 86 zoning, 20, 22, 29, 43, 44, 80, 89

87

88