

INTERNATIONAL
STANDARD

ISO/IEC
27040

First edition
2015-01-15

**Information technology — Security
techniques — Storage security**

*Technologie de l'information — Techniques de sécurité — Sécurité de
stockage*

Reference number
ISO/IEC 27040:2015(E)



© ISO/IEC 2015



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2015

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Contents

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

7	Guidelines for the design and implementation of storage security	43
7.1	General.....	43
7.2	Storage security design principles.....	43
7.2.1	Defence in depth.....	43
7.2.2	Security domains.....	44
7.2.3	Design resilience.....	45
7.2.4	Secure initialization.....	45
7.3	Data reliability, availability, and resilience.....	45
7.3.1	Reliability.....	45
7.3.2	Availability.....	46
7.3.3	Backups and replication.....	46
7.3.4	Disaster Recovery and Business Continuity.....	47
7.3.5	Resilience.....	48
7.4	Data retention.....	48
7.4.1	Long-term retention.....	48
7.4.2	Short to medium-term retention.....	49
7.5	Data confidentiality and integrity.....	50
7.6	Virtualization.....	52
7.6.1	Storage virtualization.....	52
7.6.2	Storage for virtualized systems.....	53
7.7	Design and implementation considerations.....	54
7.7.1	Encryption and key management issues.....	54
7.7.2	Align storage and policy.....	55
7.7.3	Compliance.....	55
7.7.4	Secure multi-tenancy.....	56
7.7.5	Secure autonomous data movement.....	57
	Annex A (normative) Media sanitization	60
	Annex B (informative) Selecting appropriate storage security controls	75
	Annex C (informative) Important security concepts	96
	Bibliography	109

Foreword

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

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the WTO principles in the Technical Barriers to Trade (TBT) see the following URL: [Foreword - Supplementary information](#)

The committee responsible for this document is ISO/IEC JTC 1, *Information technology, SC 27, Security techniques*.

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 Transmission Control Protocol/Internet Protocol (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 International 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 the following:

- 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 can 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 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 International 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 documents, in whole or in part, are normatively referenced in this document and are indispensable for its application. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ITU-TY.3500 | ISO/IEC 17788:2014, *Information technology — Cloud computing — Overview and vocabulary*

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

ISO/IEC 27001:2013, *Information technology — Security techniques — Information security management systems — Requirements*

ISO/IEC 27005, *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* (3.50) and retrieved on disk and tape *devices* (3.14)

**3.2
clear**

sanitize (3.38) using logical techniques on data in all user-addressable storage locations for protection against simple non-invasive data recovery techniques using the same interface available to the user

**3.3
compression**

process of removing redundancies in digital data to reduce the amount that should be *stored* (3.50) or transmitted

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

Note 1 to entry: For *storage* (3.43), 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* (3.37) in which the encryption key for the encrypted *target data* (3.52) is *sanitized* (3.38), making recovery of the decrypted *target data* (3.52) 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* (3.50) on stable *non-volatile storage* (3.30)

**3.7
data breach**

compromise of security that leads to the accidental or unlawful *destruction* (3.13), loss, alteration, unauthorized disclosure of, or access to protected data transmitted, *stored* (3.50), or otherwise processed

**3.8
data in motion**

data being transferred from one location to another

Note 1 to entry: These transfers typically involve interfaces that are accessible and do not include internal transfers (i.e., never exposed to outside of an interface, chip, or device).

**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* (3.43) 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**destruct**

sanitize (3.38) using physical techniques that make recovery infeasible using state of the art laboratory techniques and results in the subsequent inability to use the media for *storage* (3.43) of data

Note 1 to entry: *Disintegrate* (3.15), *incinerate* (3.21), *melt* (3.25), *pulverize* (3.34), and *shred* (3.41) are destruct forms of *sanitization* (3.37).

3.13**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

3.14**device**

mechanical, electrical, or electronic contrivance with a specific purpose

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

3.15**disintegrate**

destruct (3.12) by separating media into its component parts

3.16**Electronically Stored Information**

data or information of any kind and from any source, whose temporal existence is evidenced by being *stored* (3.50) in, or on, any electronic medium

Note 1 to entry: Electronically Stored Information (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* (3.26) 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* (3.45) and *storage elements* (3.47).

3.17**Fibre Channel**

serial I/O interconnect capable of supporting multiple protocols, including access to open system *storage* (3.43), access to mainframe *storage* (3.43), 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.

3.18**Fibre Channel Protocol**

serial Small Computer System Interface (SCSI) transport protocol used on *Fibre Channel* (3.17) interconnects

3.19**gateway**

device (3.14) that converts a protocol to another protocol

3.20**in-band**

communication or transmission that occurs 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

incinerate

destruct (3.12) by burning media completely to ashes

3.22

malware

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

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

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

3.23

Mean Time Between Failures

expected time between consecutive failures in a system or component

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

3.24

Mean Time To Repair

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.]

3.25

melt

destruct (3.12) by changing media from a solid to a liquid state generally by the application of heat

3.26

metadata

data that define and describe other data

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

3.27

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]

3.28

multi-tenancy

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

[SOURCE: Recommendation ITU-T Y.3500 | ISO/IEC 17788:2014, 3.2.27]

3.29

Network Attached Storage

storage device (3.45) or system that connects to a network and provide file access services to computer systems

3.30

non-volatile storage

storage (3.43) that retains its contents even after power is removed

3.31**out-of-band**

communication or transmission that occurs outside of a previously established communication method or channel

3.32**over provisioning**

technique used by *storage elements* (3.47) and *storage devices* (3.45) in which a subset of the available media is exposed through the interface

Note 1 to entry: *Storage media* (3.48) is used internally and independently by the *storage element* (3.47) to improve performance, endurance, or reliability.

3.33**point of encryption**

location within the Information and Communications Technology (ICT) infrastructure where data are encrypted on its way to *storage* (3.43) and, conversely, where data are decrypted when accessed from *storage* (3.43)

Note 1 to entry: The point of encryption is only applicable for *data at rest* (3.6).

3.34**pulverize**

destruct (3.12) by grinding media to a powder or dust

3.35**purge**

sanitize (3.38) using physical techniques that make recovery infeasible using state of the art laboratory techniques, but which preserves the *storage media* (3.48) in a potentially reusable state

3.36**reliability**

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.]

3.37**sanitization**

process or method to *sanitize* (3.38)

3.38**sanitize**

render access to *target data* (3.52) on *storage media* (3.48) infeasible for a given level of effort

Note 1 to entry: *Clear* (3.2), *purge* (3.35), and *destruct* (3.12) are actions that can be taken to *sanitize* (3.38) *storage media* (3.48).

3.39**secure multi-tenancy**

type of *multi-tenancy* (3.28) 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.

3.40**security strength**

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

3.41

shred

destruct (3.12) by cutting or tearing media into small particles

3.42

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.

3.43

storage

device (3.14), function, or service supporting data entry and retrieval

3.44

Storage Area Network

network whose primary purpose is the transfer of data between computer systems and *storage devices* (3.45) and among *storage devices* (3.45)

Note 1 to entry: A SAN consists of a communication infrastructure, which provides physical connections, and a management layer, which organizes the connections, *storage devices* (3.45), and computer systems so that data transfer is secure and robust.

3.45

storage device

any *storage element* (3.48) or aggregation of *storage elements* (3.47), designed and built primarily for the purpose of data *storage* (3.43) and delivery

3.46

storage ecosystem

complex system of interdependent components that work together to enable *storage* (3.43) services and capabilities

Note 1 to entry: The components often include *storage devices* (3.45), storage elements (3.47), storage networks, storage management, and other Information and Communications Technology (ICT) infrastructure.

3.47

storage element

component that is used to build *storage devices* (3.45) and which contributes to data *storage* (3.43) and delivery

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

3.48

storage medium

storage media

material on which *Electronically Stored Information* (3.16) or digital data are or can be recorded

3.49

storage security

application of physical, technical, and administrative controls to protect storage systems and infrastructure as well as the data *stored* (3.50) within them

Note 1 to entry: Storage security is focused on protecting data (and its storage infrastructure) against unauthorized disclosure, modification, or destruction while assuring its availability to authorized users.

Note 2 to entry: These controls may be preventive, detective, corrective, deterrent, recovery, or compensatory in nature.

**3.50
store**

record data on *volatile storage* (3.53) or *non-volatile storage* (3.30)

**3.51
strong authentication**

authentication by means of cryptographically derived credentials

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

**3.52
target data**

information subject to a given process, typically including most or all information on a piece of *storage media* (3.48)

**3.53
volatile storage**

storage (3.43) that fails to retain its contents after power is removed

**3.54
weak key**

key that interacts with some aspect of a particular cipher's definition in such a way that it weakens the *security strength* (3.40) of the cipher

4 Symbols and abbreviated terms

ACE	Access Control Entry
ACL	Access Control List
AD	Active Directory
AES	Advanced Encryption Standard
ATA	Advanced Technology Attachment
BC	Business Continuity
BCM	Business Continuity Management
CAS	Content Addressable Storage
CBC	Cipher Block Chaining
CCM	Counter with Cipher block chaining Message authentication code
CDMI	Cloud Data Management Interface
CDP	Continuous Data Protection
CHAP	Challenge Handshake Authentication Protocol
CIFS	Common Internet File System
CLI	Command Line Interface
CNA	Converged Network Adaptor
DAC	Discretionary Access Control

ISO/IEC 27040:2015(E)

DAS	Direct Attached Storage
DDoS	Distributed Denial of Service
DH-CHAP	Diffie Hellman – Challenge Handshake Authentication Protocol
DES	Data Encryption Standard
DLM	Data Lifecycle Management
DMZ	De-Militarized Zone
DNS	Domain Name System
DoS	Denial of Service
DR	Disaster Recovery
DRP	Disaster Recovery Planning
EHR	Electronic Healthcare Record
ESI	Electronically Stored Information
ESP	Encapsulating Security Payload
FC	Fibre Channel
FC-SP	Fibre Channel – Security Protocol
FCAP	Fibre Channel Certificate Authentication Protocol
FCEAP	Fibre Channel Extensible Authentication Protocol
FCIP	Fibre Channel over TCP/IP
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FCPAP	Fibre Channel Password Authentication Protocol
FCS	Fixed Content Storage
FDE	Full Disk Encryption
GCM	Galois/Counter Mode
GUI	Graphical User Interface
HAMR	Heat Assisted Magnetic Recording
HBA	Host Bus Adapter
HDD	Hard Disk Drive
HTTPS	Hypertext Transfer Protocol Secure
ICT	Information and Communications Technology
ID	Identifier

IDS	Intrusion Detection System
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IKE	Internet Key Exchange
ILM	Information Lifecycle Management
I/O	Input/Output
IP	Internet Protocol
IPS	Intrusion Prevention System
IPOCM	Incident Preparedness and Operational Continuity Management
IPsec	Internet Protocol Security
IRBC	ICT Readiness for Business Continuity
iSCSI	Internet Small Computer Systems Interface
ISL	Inter-Switch Link
ISMS	Information Security Management System
iSNS	Internet Storage Name Service
KEK	Key Encryption Key
KMIP	Key Management Interoperability Protocol
LAN	Local Area Network
LBA	Logical Block Address
LDAP	Lightweight Directory Access Protocol
LUN	Logical UNit
MAC	Mandatory Access Control
MD5	Message-Digest algorithm 5
MEK	Media Encryption Key
MTBF	Mean Time Between Failure
MTTF	Mean Time To Failure
MTTR	Mean Time To Repair
NAS	Network Attached Storage
NAT	Network Address Translation
NFS	Network File System
NIC	Network Interface Card

ISO/IEC 27040:2015(E)

NIS	Network Information Service
NPIV	N_Port_ID Virtualization
NTLM	NT LAN Manager
NTP	Network Time Protocol
NVM	Non-Volatile Memory
OASIS	Organization for the Advancement of Structured Information Standards
OID	Object Identifier
OSD	Object-based Storage Device
PCIe	Peripheral Component Interconnect express
PII	Personally Identifiable Information
PKI	Public Key Infrastructure
pNFS	parallel Network File System
PRNG	Pseudo-Random Number Generator
RADIUS	Remote Authentication Dial In User Service
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
RBAC	Role-Based Access Control
REST	REpresentational State Transfer
RNG	Random Number Generator
ROM	Read-Only Memory
RPC	Remote Procedure Call
SAN	Storage Area Network
SAS	Serial Attached SCSI
SCSI	Small Computer System Interface
SED	Self-Encrypting Drive
SHA	Secure Hash Algorithm
SIEM	Security Information and Event Management
SLP	Service Locator Protocol
SMB	Server Message Block
SMI-S	Storage Management Initiative – Specification
SNIA	Storage Networking Industry Association

SNMP	Simple Network Management Protocol
SOHO	Small Office/Home Office
SSC	Security Subsystem Class
SSD	Solid State Drive
SSH	Secure SHell
SSHD	Solid State Hard Drive
SSO	Single Sign-On
TCG	Trusted Computing Group
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol
USB	Universal Serial Bus
VLAN	Virtual Local Area Network
VM	Virtual Machine
VSAN	Virtual Storage Area Network
VPN	Virtual Private Network
WAN	Wide Area Network
WORM	Write Once Read Many
WWN	World Wide Name
WWPN	World Wide Port Name
XEX	Xor-Encrypt-Xor
XTS	XEX-based Tweaked-codebook mode with ciphertext Stealing

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. While both volatile and non-volatile forms of storage exist, this International Standards is concerned primarily with non-volatile storage. 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 Information and Communications Technology (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 or high-value business information;
- distributed database and file systems;
- support for regulatory and legal compliance requirements;
- support for centralized data repositories for rapid recovery (e.g., backups) and archiving.

5.3 Introduction to storage security

Storage security is concerned with the physical, technical and administrative controls as well as the preventive, detective and corrective controls associated with storage systems and infrastructure as outlined in [5.2](#). Storage security can also force the introduction of specialized technologies such as:

- media sanitization;
- virtualization security;
- self-encrypting storage devices (see [C.3](#)) like HDD, Solid State Drives (SSD), and Solid State Hard Drive (SSHD);
- key management services;
- data authenticity and integrity services;

- data in motion protections (encryption and data reduction);
- directory services and other user management systems.

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

- The storage systems can function as nodes within storage networks, which can be based on, but not limited to technologies like Transmission Control Protocol/Internet Protocol (TCP/IP), Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), and InfiniBand. The potential threats can vary significantly based on the networking technology as well as the topologies used.
- When stored, the data is typically represented and accessed as either block data or as files/objects; there are significant differences between these two types of storage methods. Likewise, the security measures associated with each can have radical differences, especially with access controls, encryption, and data integrity.
- As part of normal storage operations, many storage device types have more internal media capability than is exposed through the interface. SSDs and SSHDs often are typically over provisioned, where data may be moved internally between physical media areas to improve write latencies. HDDs often contain spare areas where data may be moved internally between physical media areas when there is a temporary access problem. User data may remain on such areas even after the device is written through its interface. Such devices may not be cleared by overwrites through the interface.
- Storage management is both an element of the storage infrastructure as well as an operation performed on many of the systems. It is common to have privileged users applying configuration changes, provisioning storage, tuning, monitoring, etc. this infrastructure. Some of the management can be performed remotely as well as involve third parties such as vendor support personnel.
- Data availability and integrity are key factors in an organization's storage architecture, so it is important for security to be complementary rather than a trade-off and that it not negate high-availability measures by introducing choke-points and additional single points of failure.
- Many organizations implement elaborate data resiliency strategies, which are integral to their Disaster Recovery (DR) and Business Continuity (BC) plans. Security mechanisms like data at rest encryption have to be implemented carefully to ensure that resiliency strategies are not impacted.
- Virtualization within storage can take many forms and be implemented at different points within the storage infrastructure. This virtualization can mask the physical details associated with the presentation of storage (e.g., a Logical Unit or filesystem to a server), mask the true capacity of a device, perform policy-driven autonomous data movement (like tiered storage), or completely abstract the storage infrastructure (like cloud computing storage). Balancing the security and virtualization so they can 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 Network Address Translation (NAT), Intrusion Detection System (IDS), Intrusion Prevention System (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 threats 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 threats	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
Loss of data	Unlawful, unauthorized, or accidental data destruction or corruption
Accidental configuration changes (e.g., storage management, storage/network resources, incorrect patch management, etc.) by authorized personnel	Accidental access, accidental disclosure, accidental data destruction, accidental data alteration

Table 1 (continued)

Security threats	Potential forms of data breach
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	Unauthorized 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 that can allow for recovery under the right conditions; 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 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 of 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;

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.

- timeliness.

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. Requirements can vary significantly in different jurisdictions, and for multi-national organizations, country specific legislation has a greater influence on information security requirements.

Common compliance issues include:

- breach of country specific privacy requirements;
- unlawful transfer of data (i.e., moving restricted data out of particular jurisdiction);
- 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.).

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 clauses [6.2](#) to [6.8](#) below. They 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 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 may be 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 sensitive and high value data on DAS, some form of encryption should be used to protect the data at rest, including:
 - Storage elements with integrated encryption and access control capabilities, also known as Self-Encrypting Drives (SEDs);

- Computer-based or application-based encryption, including Full Disk Encryption (FDE).
- Media sanitization (see [6.8.1.2](#) and [Annex A](#) for additional information) should be used on all DAS involved with sensitive or high value data with either of the following:
 - Use the integrated sanitization functionality in the storage elements;
 - Use computer-based or application-based sanitization.
- If possible, authentication such as Fibre Channel – Security Protocol – 2²) (FC-SP-2) AUTH-AA Authentication (see [C.7.2](#)) should be used to prevent unauthorized access to sensitive and high value data;
- To guard against accidental or intentional data loss or corruption, backups of the DAS contents should be made on a regular basis.

6.3 Storage networking

6.3.1 Background

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

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

6.3.2 Storage Area Networks (SAN)

6.3.2.1 General

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

SANs are often used to improve application availability (e.g., multiple data paths), enhance application performance (e.g., off-load storage functions, use separate networks, etc.), increase storage utilization and effectiveness (e.g., consolidate storage resources, tiered storage, etc.), and improve data protection and security. In addition, SANs typically play an important role in an organization's DR and BC activities (see [Figure 1](#)).

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

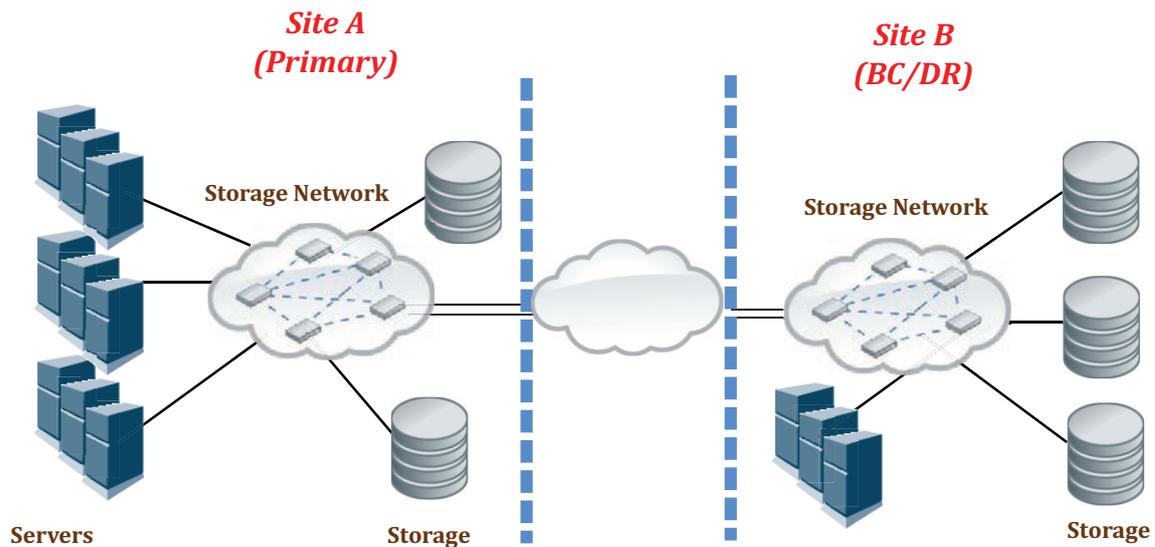


Figure 1 — Storage Area Network (SAN) example

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

SANs are commonly based on Fibre Channel (FC) technology that utilizes the Fibre Channel Protocol (FCP) for open systems and proprietary variants for mainframes. In addition, the use of Fibre Channel over Ethernet (FCoE) makes it possible to move FC traffic across existing high-speed Ethernet infrastructure and converges storage and IP protocols onto a single cable transport and interface. Other technologies like Internet Small Computing System Interface (iSCSI), commonly used in small and medium sized organization as a less expensive alternative to FC, and InfiniBand, commonly used in high performance computing environments, can also be used. Interconnects like Serial Attached SCSI (SAS) and Peripheral Component Interconnect express (PCIe) that leverage extenders and switches are beginning to take on characteristics of a SAN as well.³⁾ In addition, it is possible to move data between different SAN technologies through the use of gateways (see [Figure 2](#)).

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, Logical Unit (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. Soft zoning is based on limiting SAN fabric nameserver responses to queries based on the assumption that servers will not contact storage devices that are not discovered via the nameserver. Hard zoning uses physical port numbers on SAN switches to restrict traffic forwarding and is a more secure zoning method because it does not rely on correct server behaviour and in particular is not vulnerable to spoofing of server identity.

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

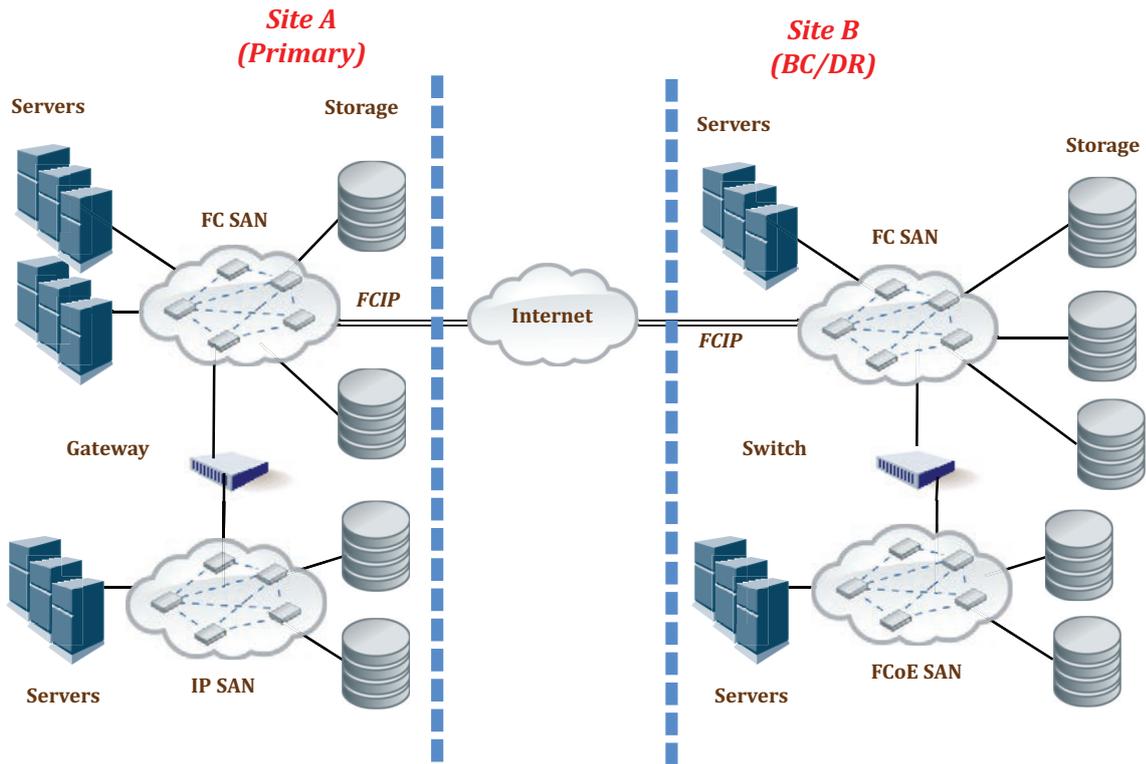


Figure 2 — Storage Area Network (SAN) example

- **LUN masking and mapping:** A storage device can be divided into different logical units (LUNs) that are identified by logical unit numbers. LUN mapping refers to the assignment of a number to a LUN, and it typically takes place in a storage array, but it can also occur as part of a redirection (initial address to a new address) in the switch, Host Bus Adapter (HBA) or Converged Network Adaptors (CNA), and virtualization layer. LUN masking refers to making a LUN visible to some servers and not visible to others.
- **Authentication:** For SANs it is important for a switch to verify the identity of other 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. Likewise, the nodes in a SAN (e.g., storage devices and servers) need to employ authentication to guard against data breaches.
- **Encryption:** There are two major components of data confidentiality on a SAN: 1) data in motion and 2) data at rest. Sensitive and high-value data needs to be cryptographically protected in SANs when it is in motion as well as when it is at rest on a storage device. 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 including:
 - 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 including:
 - segregate storage traffic from normal server traffic using:
 - available network controls to create independent logical domains on common physical infrastructure;
 - trust and access controls to manage membership in the logical domains.
 - segregate storage management traffic from all other traffic;
 - ensure that configurations of network gateways maintain appropriate network segregation.

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

Fabric administrators should take steps to:

- Control FCP node access including:
 - restrict server access on the switches using techniques such as Access Control Lists (ACLs), binding lists, and FC-SP-2 fabric policies (see [C.7.1](#));
 - use NPIV (N_Port_ID Virtualization) enabled HBAs to assign individual N_Port_IDs to virtual servers.
- Implement switch-based controls including:
 - restrict switch interconnections using techniques such as ACLs, binding lists, and FC-SP-2 fabric policies (see [C.7.1](#));
 - zoning should be used in FC SAN fabrics with a preference for hard zoning;
 - determine whether basic zoning is a strong enough security measure for the target environment, and if it is not, use stronger techniques like FC-SP Zoning (see [C.7.5](#)) where supported by the vendor;
 - disable unused ports;
 - carefully use default zones and zone sets (assume a least privilege posture).
- Interconnect storage networks securely by configuring switches, extenders, routers, and gateways necessary to meet requirements.

Subclause [6.5.1](#) provides guidance on block-based Fibre Channel storage.

6.3.2.3 IP SAN

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

Control iSCSI network access and protocols by:

- avoiding connecting iSCSI interfaces to general purpose LANs; segregate for security and performance;
- using Virtual Local Area Networks (VLANs) when the use of physically isolated LANs is not an option.

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

FCIP network access and protocols controls should be controlled by:

- setting up the peer-to-peer relationship between FCIP entities, recognizing that the security policies will be applied uniformly;
- using a private IP network exclusively by the FCIP entities whenever possible.

Implement IPsec security measures in conjunction with FCIP by:

- performing cryptographic authentication and data integrity at a minimum;
- protecting sensitive data by appropriate confidentiality measures.

IETF RFC 3723, *Securing Block Storage Protocols over IP* provides additional useful information on both iSCSI and FCIP. Subclause [6.5.2](#) provides guidance on block-based IP storage. In addition, IETF RFC 7146, *Securing Block Storage Protocols over IP: RFC 3723 Requirements Update for IPsec v3* provides important security updates to IETF RFCs 3720, 3723, and 3821.

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.

FCoE SANs should be protected by:

- leveraging the Fibre Channel security mechanisms (see [C.7](#));
- protecting against Ethernet broadcast storms (e.g., allocation of adequate input buffering) that can cause throughput and timeout issues;
- using ACLs to control network access (e.g., denying specific computers from unnecessary or unwanted traffic);
- using FCoE VLANs when the use of physically isolated LANs is not an option.

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 2 kilobyte FC frames.

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 or striping data and metadata across the clustered storage servers (see [Figure 3](#)); parallel NFS (pNFS) systems are examples of clustered NAS systems.

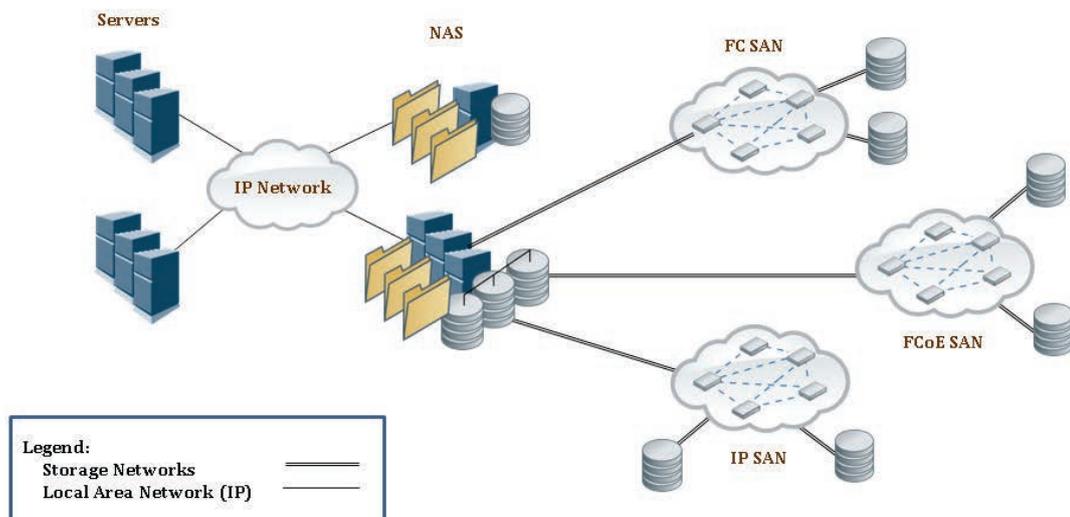


Figure 3 — Network Attached Storage (NAS) Example

Common file system implementations include the Network File System (NFS) and Server Message Block (SMB)/ Common Internet File System (CIFS), but other technologies like Object-based Storage Device (OSD) and cloud computing 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; and
- authentication controls, such as Kerberos, for verifying the identity of users attempting to access NAS data.

Refer to [6.6](#) for further implementation guidance on NAS and file-based storage.

6.3.3.2 Network File System (NFS)

NFS is a client/server application, communicating with a Remote Procedure Call (RPC)-based protocol. Multiple versions of NFS are specified and in use, including NFS version 3 (specified in IETF RFC 1813), NFS version 4 (specified in IETF RFC 3530), and NFS version 4.1 (specified in IETF RFC 5661). From a security perspective, NFS version 3 (NFSv3) is considered less secure and extra care should be exercised when it is used with sensitive or high-value data.

The following networking guidance is applicable to NFS-based NAS and should be used:

- control NFS network access and protocols by:
 - enabling NFS only if needed. This will eliminate it as a possible attack vector available to an intruder;
 - using NFSv4 (or later versions) whenever possible and limit NFSv3 usage;
 - filtering client and management access by IP address for additional security.
- encrypt client data access (e.g., IPsec) when necessary.

6.3.3.3 SMB/CIFS

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

The following networking guidance is applicable to SMB/CIFS-based NAS and should be used:

- use later versions of the SMB protocol;
- turn off low-security session negotiation protocols, such as NT LAN Manager (NTLM) v1, LanMan and plaintext, and use NTLM v2 or Kerberos instead;
- maintain up-to-date patch levels;
- use SMB signing;
- maintain Active Directory (AD) services securely;
- use one-way trusts, from leaf domains to parent domains, when possible;
- control SMB/CIFS network access and protocols by:
 - enabling SMB/CIFS only if needed. This will eliminate it as a possible attack vector available to an intruder;
 - encrypting client data access when necessary.

6.4 Storage management

6.4.1 Background

Storage networks and infrastructure elements are complex architectures that 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 (see [Figure 4](#)).

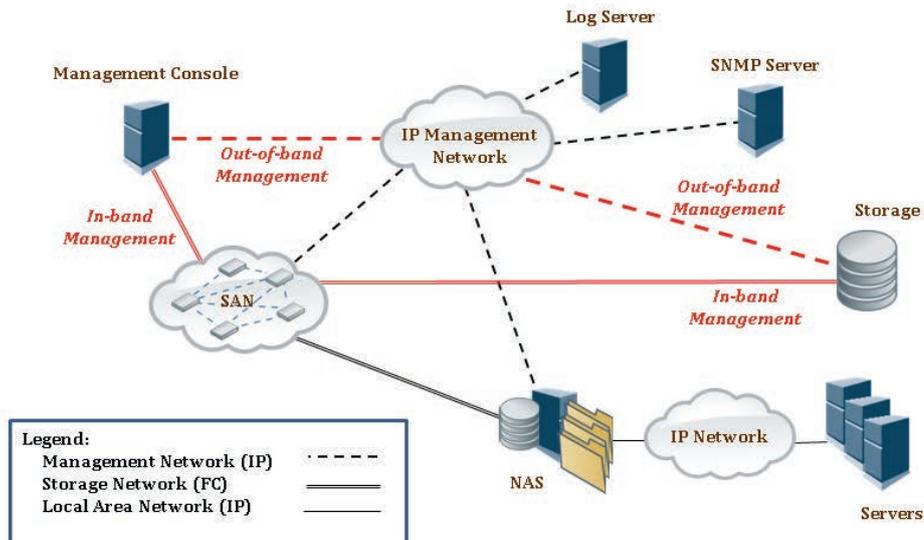


Figure 4 — Storage Management Example

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, or 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. According to ISO/IEC 27002:2013, 9.2.3 the allocation and use of privileged access rights should be restricted and controlled. 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, secure log-on procedures as described in ISO/IEC 27002:2013, 9.4.2 with 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 by using:
 - strong passwords (increased minimum number of characters, increased complexity, etc.) with a reduced period of use;
 - strong authentication (e.g., a challenge response protocol); or
 - multi-factor authentication, such as biometric data (e.g. finger-print verification, signature verification) and use of hardware tokens (e.g., smart cards).
- for all remote access, use strong authentication or multi-factor authentication along with secure channels;
- when possible, use a centralized authentication solution (e.g., Remote Authentication Dial In User Service or RADIUS, Single Sign-on or SSO, etc.) for improved monitoring and control;
- use multi-factor authentication when managing sensitive and high-value data;
- disable login to the root account. Remotely log all privilege escalation operations.

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

6.4.2.2 Authorization and access control

Within market sectors like financial services and healthcare there are trends to align authorization and access control (see [C.2](#)) to a least-privilege model that leverages specific roles. The following roles should be implemented and used within storage technologies:

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

- *Storage Auditor* - This operator-like role has view rights that allow for the verification of storage parameters and configurations and inspections of health/fault logs. No access is granted to security-related elements or data.

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

6.4.3 Secure the management interfaces

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

Management interfaces for storage systems can take on several physical forms including serial ports (e.g., RS-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, noting that physical isolation is preferred, but logical isolation (such as VLANs) should be used at a minimum.

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 server-based proxies that handle 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 systems and protocols;
- use entity authentication to establish trust relationships between storage systems and the management systems (e.g., using FC-SP-2 AUTH-A to authenticate the entities performing in-band management as described in [C.7.1](#));
- leverage IDS and IPS mechanisms to identify anomalous behaviours and guard against them;
- use ICT infrastructure (Domain Name System or DNS, Service Locator Protocol or SLP, Network Time Protocol or 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 (Virtual Private Network or VPN, TLS, Secure Shell or SSH, Hypertext Transfer Protocol Secure or 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. systems, 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 call-back 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 [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 (see also ISO/IEC 27002:2013, 12.4 for relevant guidance).

The following logging guidance is applicable to storage systems and should be used:

- include storage in the logging policy such that:
 - 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; and
 - the device time is synchronized with a reliable, external source.
 - the logging policy (see also [7.7.2](#)) should include evidentiary expectations (authenticity, chain of custody, etc.).
- employ external or centralized event logging to a trusted⁶⁾ remote source by:
 - implementing centralized⁷⁾ audit logging to collect events from all sources in a single repository;
 - establishing and using a common, accurate time source across the environment to assure that event records from different sources can be correlated;
 - avoiding use of device resident logs for anything other than system health monitoring and debugging 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;
 - natively logging events to one, and preferably multiple,⁸⁾ external log servers;

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 User Datagram Protocol (UDP) and therefore log messages may be lost due to network or server performance. Sending messages to multiple log

- using standard logging protocols like syslog⁹⁾ that support reliable delivery and secure transports (e.g., TLS);
 - having devices configured to log events as they occur (i.e., no buffering) when the primary drivers for audit logging are compliance, accountability, or security;
 - implementing an analysis protocol to correlate audit log records across event sources to identify significant security events that provide indication of security incidents;
 - ensuring that the storage logging is factored into Security Information and Event Management (SIEM) solutions, when such technology is deployed.
- ensure complete event logging
- once the types of events to be logged have been determined, then all occurrences of these events should be logged (whether in-band or out-of-band);
 - the following kinds of events should be logged (a minimum set of security events):
 - failed and successful logon attempts;
 - failed file and object access attempts for sensitive and high-value data;
 - account and group profile additions, changes, and deletions;
 - changes to system security configurations (e.g., audit logging, network filtering, zoning changes);
 - changes to security server usage (e.g., syslog, Network Time Protocol or NTP, Domain Name System or DNS, authentication);
 - system shutdown and restarts;
 - privileged operations (i.e., administrator initiated changes);
 - use of sensitive utilities (e.g., privilege escalation commands);
 - access to critical data files;
 - movement of virtual servers between physical servers.
 - each log entry should include:
 - a timestamp (date and time);
 - a severity level;
 - the source of the log entry (distinguishing name, IP address, etc.);
 - an event ID as well as a textual description (necessary to enable localization/internationalization of events – where the event ID remains the same but the textual description could be translated to different languages); and

destinations reduces the risk of inadvertent loss.

9) Syslog is defined in IETF RFC 5424 with additional details contained in IETF RFC 3195, IETF RFC 5425, IETF RFC 5426, IETF RFC 5427, IETF RFC 5848, IETF RFC 6012, and IETF RFC 6587.

- 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 such that:
 - audit log data that may have evidentiary value should be handled correctly (e.g., maintain chain of custody, verifiable integrity and authenticity, etc.);
 - audit log data with specific retention requirements (e.g., for regulatory compliance) should be preserved with the organization’s data retention solution (see [7.4](#));
 - implement appropriate measures to preserve log integrity and prevent their modification or destruction (either maliciously or accidentally);
 - when audit log entries contain sensitive information, the audit log data should be protected with appropriate confidentiality mechanisms¹⁰;
 - for unique audit logging requirements (e.g., high volume, special preservation, event signing, etc.) dedicated and specially hardened and configured systems should be used;
 - Leverage log relays and log filtering to minimize the impact of specialized storage requirements (e.g., Write Once Read Only or WORM).

6.4.5 System hardening

All operating systems, hypervisors, and applications should be hardened relative to the use of the storage system. In addition to the technical vulnerability management guidance in ISO/IEC 27002:2013, 12.6, there are many existing best practices for various operating systems that should be referenced based on the 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;
- only open up network ports that are needed;
- installation of latest patches from a trusted source;
- update firmware from a trusted source;
- install and maintain malware protection (see also ISO/IEC 27002:2013, 12.2).

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.

Vendors should perform the actions described in [6.4.5](#) for elements under their control.

¹⁰ 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 (LUNs) and tape devices (including virtual tape).

For Fibre Channel systems, the following should be considered:

- LUN masking and mapping (WWN filtering) and other access control mechanisms should be used to restrict access to storage
- implement FCP security measures including:
 - mutual authentication using FC-SP-2 AUTH-A (see [C.7](#)) should be used with all servers and switches; leverage centralized authentication services when possible;
 - if possible, Fibre Channel connections that leave the protected area (e.g., confines of a physically controlled data centre) should be encrypted using ESP_Header¹¹⁾ (see [6.8.2.2](#) and [C.7.3](#)).
- implement data at rest encryption measures (see [6.8.2.3](#)) including:
 - sensitive and high-value data should be encrypted while on the storage device or media¹²⁾;
 - encryption should be implemented in storage devices that may come in contact with sensitive or regulated data as well as to facilitate rapid sanitization (see [A.3](#)).
- implement sanitization measures (see [6.8.1](#) and [Annex A](#)) including:
 - media-aligned sanitization (see [6.8.1.2](#)) should be used for sensitive and regulated data;
 - logical sanitization (see [6.8.1.3](#)) should be used to clear virtualized storage (see [7.6.1](#)), especially when the actual storage devices and media cannot be determined.

Vendors should implement the functionality for access control, authentication, sanitization and encryption described in [6.5.1](#) within their products.

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.

For IP storage systems, the following should be considered:

- control iSCSI initiator access by filtering based on source IP addresses and protocols;
- implement iSCSI security measures including:
 - bidirectional Challenge Handshake Authentication Protocol (CHAP) authentication, using random challenges (i.e., not repeated), should be used for both initiators and targets in all iSCSI implementations;
 - IPsec should be used to secure the communication channel when sensitive or high-value data could be exposed (see [6.8.2.2](#));

11) Fibre Channel frame integrity or confidentiality can be provided with ESP_Header optional headers, which are defined in ANSI INCITS 470–2011, *Fibre Channel – Framing and Signaling-3 (FC-FS-3)*.

12) Encryption within FC storage ecosystems provides media-level protection and can be a safety net for data that typically is encrypted by a server, application, etc. as the primary form of protection.

- Internet Storage Name Service (iSNS), SLP, DNS infrastructure should be used with appropriate security controls to avoid indirect attacks.
- implement data at rest encryption measures (see [6.8.2.3](#))
 - sensitive and high-value data should be encrypted while on the storage device or media¹³;
 - encryption should be implemented in storage devices that may come in contact with sensitive or regulated data as well as to facilitate rapid sanitization (see [A.3](#)).
- implement sanitization measures (see [6.8.1](#) and [Annex A](#))
 - media-aligned sanitization (see [6.8.1.2](#)) should be used for sensitive and regulated data;
 - logical sanitization (see [6.8.1.3](#)) should be used to clear virtualized storage, especially when the actual storage devices and media cannot be determined.

Vendors should implement the functionality for access control, authentication, sanitization and encryption described in [6.5.2](#) within their products.

6.6 File-based storage

6.6.1 NFS-based NAS

This type of storage is a LAN-attached file server that presents 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 servers 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 including:
 - employ user-level authentication whenever possible (e.g., NFSv4 with Kerberos V5);
 - configure the NFS server to export file systems explicitly for the authorized users;
 - configure the NFS server to export file systems with minimum required privileges;
 - avoid granting “root” or “administrator” access to files on network filesystems;
 - make sure NFSv4 ACLs (Access Control Lists) are assigned correctly;
 - use Kerberos authentication for NFSv3;
 - consider using Kerberos Safe and Private modes to sign and encrypt NFS traffic.
- restrict NFS client behaviours
 - filter client access to NFS shares whenever possible;
 - do not allow NFS clients to run *suid* and *sgid* programs on exported file systems.
- secure data on NFS server
 - 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;

¹³) Encryption within IP storage ecosystems provides media-level protection and can be a safety net for data that typically is encrypted by a server, application, etc. as the primary form of protection.

- encrypt data at rest when necessary;
- do not allow NFS exports of administrative file systems (e.g., */etc*);
- guard against malware (e.g., viruses, worms, rootkits, etc.);
- continually monitor content placed in NFS shares and relevant access controls.

Vendors should implement the functionality for access control, authentication, and encryption described in [6.6.1](#) within their products.

6.6.2 SMB/CIFS-based NAS

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

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

- apply access controls to SMB/CIFS exported filesystems
 - disable unauthenticated access to CIFS shares and NAS devices (i.e. restrict *Anonymous*);
 - disable “Guest” and “Everyone” access to all CIFS shares;
 - implement authentication and access control via a centralized mechanism (RADIUS, Lightweight Directory Access Protocol or LDAP).
- restrict SMB/CIFS client behaviours by enabling SMB signing for clients and the NAS device;
- secure data on SMB/CIFS servers:
 - 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).

Vendors should implement the functionality for access control, authentication, and encryption described in [6.6.2](#) within their products.

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, the 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 fileserver(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 IETF 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 fileserver - an effective control has to deal with direct client access to that latter fileserver, as such access would bypass the directory ACL.

For pNFS systems, the following should be considered:

- 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 fileserver;
- 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 computing storage

6.7.1.1 Securing cloud computing storage

Both proprietary and standards-based, cloud computing storage offerings are in use and they commonly provide copy capabilities (e.g., mirror some or all of the storage on a system), backups and recovery capabilities, long-term retention capabilities (e.g., archives), and multi-system synchronization capabilities (e.g., allows a user to synchronize data on multiple and potentially different types of devices). However, individuals and organizations hesitate to entrust their data to cloud computing storage unless they have assurances that the relevant security threats and challenges have been addressed.¹⁴⁾

Some of these cloud computing implementations are object-based and 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.

Secure use of cloud computing storage should involve some or all of the following:

- ensure that transport security such as IPsec or Transport Layer Security (TLS) is used for all transactions (see [6.8.2.2](#));
- when sensitive data is stored in a third party cloud environment, data at rest encryption (and appropriate key management processes) should be used to prevent access by unauthorized parties (e.g., cloud service provider personnel, other tenants, adversaries, etc.);

¹⁴⁾ Recommendation ITU-T X.1601, *Security framework for cloud computing* provides a useful summary of cloud computing security threats and challenges.

- secure user registrations and use strong password authentication to protect access to data;
- employ access controls that guard against unauthorized access from other tenants while providing appropriate access privileges to users permitted to access the data;
- use the provided sanitization capabilities to clear sensitive data from the cloud computing storage.

NOTE In some jurisdictions, privacy requirements like the “right of erasure” or the “right to be forgotten” may necessitate additional security controls.

Cloud computing implementations often leverage different forms of virtualization, so the guidance in [7.6](#) may also be relevant.

Vendors should implement the appropriate cloud computing storage functionality for access control, authentication, encryption, logging, sanitization, etc. described in [6.7.1.1](#) within their products.

6.7.1.2 CDMI security

Cloud computing storage, based on the ISO/IEC 17826:2012 *Cloud data management interface (CDMI)* specification, is an object-based storage technology that uses a RESTful HTTP interface. It 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.

CDMI clients should:

- ensure that Transport Layer Security (TLS) is used for all transactions (see [6.8.2.2](#));
- query 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;
- enable CDMI security logging and retrieve the security event data out of the appropriate logging queue on a regular and timely basis;
- 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;
- use data at rest encryption measures to protect sensitive and high-value data;
- 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.

Vendors should implement the appropriate CDMI functionality for access control, authentication, encryption, logging, sanitization, etc. described in [6.7.1.2](#) within their products.

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) that is accessed by specifying the Object Identifier (OID) and an (offset, length) tuple. The OSD interface includes commands¹⁵⁾ to create and delete objects, write bytes and read bytes to and from individual objects, and to set 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 that encodes a set of rights the holder has on an object, and is cryptographically secured.

To use OSD securely:

- 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.

Vendors should implement the appropriate OSD functionality for access control, authentication, encryption, etc. described in [6.7.2](#) within their products.

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 in time (i.e., it 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 user does not need to know the location of the actual data and the number of copies being stored.

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.

15) The initial OSD standard, ANSI INCITS 400-2004, *Information technology – SCSI Object-based Storage Device Commands (OSD)*, was approved in 2004; a more recent version of the OSD standard is ANSI INCITS 458-2011, *Information technology – SCSI Object-Based Storage Device Commands – 2 (OSD-2)*.

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

Vendors should implement the appropriate CAS functionality for authentication, authorization, availability, hashing etc. described in [6.7.3](#) within their products.

6.8 Storage security services

6.8.1 Data sanitization

6.8.1.1 General

An important element of the media handling guidance, as described in ISO/IEC 27002:2013, 8.3, is to ensure that the contents are made unrecoverable when media are no longer needed. Sanitization refers to the general process of rendering previously written data in the storage media irretrievable, 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 their 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.

Disposal of storage devices or storage elements without sanitization should be considered only if information disclosure would have no impact on organizational mission, would not result in damage to organizational assets, and would not result in financial loss or harm to any individuals.

Clear, purge, and destroy (destruct) are actions that can be taken to sanitize storage. [A.1](#) describes each method and provides additional options where appropriate. [A.2](#) supplements this information by providing specific guidance on sanitizing both hard copy and electronic (soft) copy media.

Sanitization operations can be costly and time consuming, but they are necessary for security reasons. The level of sanitization operations should be 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.).

When sanitization is an element of compliance, the specific requirements and associated specifications should be reviewed to determine whether they mandate particular overwrite techniques, documentary proof of sanitization, etc. These requirements can take the form of specific overwrite techniques, proof of sanitization, etc.

To make appropriate sanitization capabilities available, vendors should implement the functionality described in [6.8.1.2](#), [6.8.1.3](#), [6.8.1.4](#), [6.8.1.5](#), and [Annex A](#) within their products.

6.8.1.2 Media-based sanitization

When storage media are transferred, become obsolete, are no longer usable, or are not needed by an information system, the residual magnetic, optical, electrical, or other representation of data should be sanitized.

[Annex A](#) should be used to determine recommended sanitization of specific media. Although the use of [Annex A](#) is strongly 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-use situations is recommended, even when using encryption methods. Even when cryptographic erase is used to sanitize the data on the device, it is recommended that the media itself be sanitized as well.

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 (LUN) on a storage array that can have a size that far exceeds the capacity of a single storage element. The situation can be further complicated when logical storage is replicated (i.e., multiple copies of the data exist) to support server virtualization (see 7.6.2) and Disaster Recovery (see 7.3.4). 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 logical storage (e.g., Logical Unit, filesystem or object store) is writeable, then sanitization, using an overwrite or cryptographic erase technique, should be used to clear the portions of the underlying storage media used by the logical storage; successful application of cryptographic erase for sanitization (see A.3) is predicated on the encryption being active before data is recorded on the logical storage. Data protection technologies (see 7.3.3), which can include replication, backups and CDP storage, are often used in conjunction with logical storage, so separate sanitization operations should be performed on storage associated with data protection mechanisms.

6.8.1.4 Proof of sanitization

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 or costly data breach notifications. The importance of this proof along with the provenance or chain of custody requirements associated with the evidence documenting sanitization 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, 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 the situation where media are in an inoperable state and physical destruction is necessary, proof of sanitization should be achieved by a certificate of sanitization.

6.8.1.5 Verification of sanitized media

The goal of sanitization verification is to assure that the target data was effectively sanitized. When supported by the device interface (such as an Advanced Technology Attachment or 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 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, using a new seed for the Pseudo-Random Number Generator (PRNG) 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 Logical Block Address (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.

Devices that are protected with access control mechanisms have additional verification considerations. Whether such devices were sanitized by overwrite, block erasing or cryptographic erase (see A.3), such devices need to be accessible both before and after sanitization to enable a verification process.

Cryptographic erase has different verification considerations than other procedures, because the contents following cryptographic erase may not be known and therefore cannot be compared to a given value. When cryptographic erase is leveraged, an attempt should be made to apply simple checks such as reading a storage location with known contents (for example, file system metadata) to verify that the expected data is not returned. If it not possible, for whatever reason (e.g., person executing cryptographic erase does not have read access), then verification can be skipped.

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 (sometimes 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, 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). ISO/IEC 27002:2013, 10.1.2 also provides relevant guidance on key management.

To make appropriate data confidentiality capabilities available, vendors should implement the functionality described in [6.8.2.2](#) and [6.8.2.3](#) within their products.

6.8.2.2 Encrypting transferred data

Within storage infrastructures, data confidentiality 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 ESP_Header (see [C.7.3](#)), IPsec,¹⁶⁾ TLS,¹⁷⁾ or even computer 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).

Data in motion 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 that transforms ciphertext back into its original plaintext and the receiver performs a check of the integrity value. There are various standard specifications including the Fibre Channel security standards (see [C.7.1](#)), IPsec RFC's, and TLS RFC's that detail 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.

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

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

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 Internet Key Exchange (IKE) version 2 (or later versions) should be used.
- For TLS, storage clients should comply with the requirements in the Storage Networking Industry Association (SNIA) Technical Position: *TLS Specification for Storage Systems v1.0* (or the latest version).

6.8.2.3 Encrypting data at rest

With increasing amounts of sensitive and regulated data being stored, organizations need to take 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 data at rest within the storage infrastructure does provide a basic level of protection against breaches stemming from the loss of control of media, especially tape. Consequently, encryption mechanisms within storage devices (Self-Encrypting Drives as well as controller-based technologies), switches, specialized appliances, HBAs, etc. should be used.

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 selected to address the identified risks, 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 [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 Organization for the Advancement of Structured Information Standards (OASIS) Key Management Interoperability Protocol (KMIP) specification and profiles¹⁸⁾ are the dominant mechanism for centralized key management within storage infrastructures.

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

- encryption algorithms and modes of operations designed specifically for storage technology like XTS-AES (specified in IEEE 1619-2007) for HDD and Counter with Cipher block chaining Message authentication code (CCM) or Galois/Counter Mode (GCM) (as described in IEEE 1619.1-2007) for tape should be used; if storage specific modes are not available, suitable AES modes like Cipher Block Chaining (CBC) (specified in ISO/IEC 10116:2006) may be used;
- limit the amount of time a key is in plaintext form and prevent humans from viewing plaintext keys
- cryptographic keys should only be used for one purpose, specifically, do not use key-encrypting keys (also known as key wrapping keys) to encrypt data or use data encrypting keys to encrypt other keys;

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

- randomly choose keys from the entire keyspace¹⁹⁾;
- check for and avoid use of known weak keys;
- data encryption keys should be limited to a finite cryptoperiod (typically no more than 2 years) or to a maximum amount of data processed;
- when possible, storage systems and infrastructure should use interoperable, centralized key management infrastructure (e.g., generate and archive encryption keys);
- storage systems and infrastructure should use OASIS approved, KMIP-compliant clients to access and use key management infrastructure (see [C.8](#)).

6.8.3 Data reductions

As a routine course of business, organizations may attempt to reduce the amount of data they store and transmit in an effort to reduce costs. Two of the more common approaches are data compression and data deduplication. Data compression seeks to reduce the amount of data by encoding it with a known algorithm²⁰⁾ to produce a representation of the data that uses fewer bits of storage than the unencoded representation. Data deduplication, on the other hand, attempts to replace multiple copies of data with references to a shared copy. These two techniques can be used together to maximize data reduction.

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

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

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

- When encryption is used along with compression, the compression should be applied before the encryption because ciphertext does not effectively compress; the reverse order should be used on the other end (i.e., decryption followed by expansion).
- When encryption is used along with deduplication, the deduplication should be applied before the encryption because deduplication is often not effective on ciphertext; the reverse order should be used when the data is to be decrypted.
- 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 or deduplication can impact DR and BC implementations, so they should be factored into the design, documentation, and testing of DR and BC solutions.

19) Best practice recommends a cryptographically secure Pseudo-Random Number Generator (PRNG) with the property that when given full knowledge of the algorithm and a sequence of outputs, neither the numbers preceding the sequence nor the numbers following the sequence can be determined using practical computational means.

20) Compression algorithms include lossy approaches (in which a portion of the original information is lost) and lossless approaches (which preserve the entire content of the original data), 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.

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 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;
- deploy protection mechanisms at multiple locations to resist all classes of attacks;
- deploy multiple defence mechanisms (layered) between potential adversaries and targets;
- include both detection and protection mechanisms;
- deploy robust key management and Public Key Infrastructure (PKI) frameworks that support all information assurance technologies and that are highly resistant to attack;
- maintain visible and up to date system security policies;

- actively manage the security posture of the storage technology and protection mechanisms (e.g., install security patches and anti-virus updates, maintain ACLs, etc.);
- perform regular security threat assessments to determine the continued security readiness;
- monitor and react to current threats.

Security solutions based on the layered approach are flexible and scalable as well as being adaptable to the security needs of the organization.

For storage, a layered approach means that security controls are deployed and used throughout the storage infrastructure, including the HBA/CNA/NIC in computers, storage network switches/routers, storage appliance, storage elements, and storage devices.

7.2.2 Security domains

Security domains are based on the concept that system resources of different sensitivity levels (i.e., different risk tolerance values and threat susceptibility) should be differently located. This creates a way to have the systems make available only such data that is necessary for conducting the tasks for that particular domain. As a design principle, the architecture should enforce domain separation to ensure that resources to which an entity has access cannot be accessed or affected by another entity.

For storage infrastructure, a security domain will typically be represented as a SAN, especially when sensitive data is being stored and processed within the storage systems. In situations where the data sensitivity is low, zoning and VLANs can be considered acceptable, but it is important to note that this generic capability is not a security mechanism such as FC-SP Zoning (see [C.7.5](#)).

Building on the compartmentalization principle described in ISO/IEC 27033-2, the following storage security design rules should be considered:

- Factor data sensitivity into the use of security domains
 - storage and storage networks of different sensitivity levels should be located in different security domains;
 - devices and computer systems providing services for external networks (e.g., the Internet) should be located in different domains (De-Militarized Zone or DMZ) than internal network devices and computer systems;
 - strategic assets should be located in dedicated security domains;
 - untrusted devices and computer systems should have limited or no access to storage assets.
- Factor purpose in the use of security domains
 - 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;
 - storage networks should be in different security domains than regular networks (e.g., corporate LANs);
 - storage device and storage network management systems should be located in dedicated security domains;
 - systems in development stage should be located in different domains than production systems.
- 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 Virtual Storage Area Networks or VSANs) to minimize possible interactions.

7.2.3 Design resilience

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

7.2.4 Secure initialization

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

Vendors should implement the functionality for secure initialization described in [7.2.4](#) within their products.

7.3 Data reliability, availability, and resilience

7.3.1 Reliability

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

- MTBF (Mean Time Between Failures) for repairable product, which is the expected time between consecutive failures in a system or component and sometimes thought of as the average time available for a system or component to perform its normal operations between failures (see [Figure 5](#));
- MTTR (Mean Time To Repair) for repairable product, which is the expected or observed duration to return a malfunctioning system or component to normal operations and sometimes thought of as the average time to repair a failed component;
- MTTF (Mean Time To Failure) for non-repairable product, which is the average time available for a system or component to perform its normal operations until it fails.

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. For example, incorrect application of updates or use of updates from non-approved or untrusted sources can have adverse impacts.

- 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 capable of assuring the reliability and security of data.

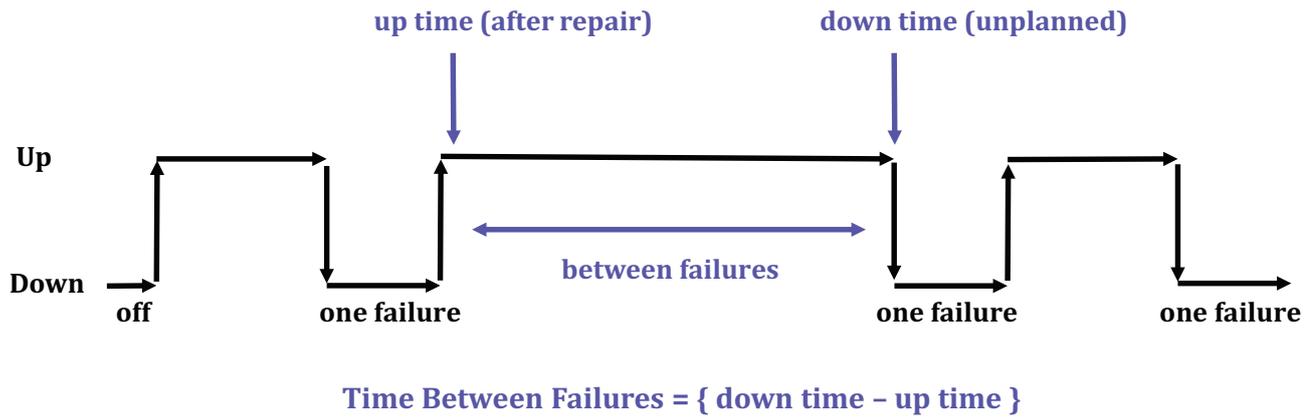


Figure 5 — Quantification of reliability

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 a 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 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 backups²¹⁾ 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;

21) ISO/IEC 27002:2013, 12.3 provides relevant guidance on backups.

- 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 and appropriate protections against unauthorized access (e.g., encryption or user validation);
 - 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., data in motion encryption).
- 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 that 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., data in motion and data 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 (IRBC), 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 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, so it is important to.

- ensure that 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 can become corrupted as the media degrades over time. 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*
- ISO/TR 18492:2005, *Long-term preservation of electronic document-based information*
- ISO 16175-1:2010, *Information and documentation — Principles and functional requirements for records in electronic office environments — Part 1: Overview and statement of principles*
- ISO 16175-2:2011, *Information and documentation — Principles and functional requirements for records in electronic office environments — Part 2: Guidelines and functional requirements for digital records management systems*
- ISO 16175-3:2010, *Information and documentation — Principles and functional requirements for records in electronic office environments — Part 3: Guidelines and functional requirements for records in business systems*

Long-term archival storage systems introduce integrity, authentication and privacy threats that do not generally exist in non-archival storage systems. In addition, the long lifetime of data gives attackers a much larger window within which they can attempt to compromise a security system; with archival storage an assailant might have several decades of time to conduct an attack (slow attack).

- Archival storage assumes a write-once, read-maybe access pattern, thus the integrity of the data in the system should be actively checked at regular intervals rather than waiting to when it is read.
- When migrating archival data to newer storage technologies, introduce available security capabilities that offer enhanced security measures to better secure the data in its new location.

- Since the data in a long-term archive can out-live the data owners, a secure, archival storage system should be able to authenticate new users and establish their relationship to resources attached to existing users.
- Secrecy mechanisms (e.g., encryption, secret-sharing, etc.) should function in the complete absence of the user that wrote the data (e.g., a new user who is given rights to read data should also be given the ability to decrypt the data).
- Security logging should be sufficiently complete and long-lived (measured in decades) that it assists in detecting slow attacks and maintains an attack history that can be used to make decisions to adjust the data protections.
- The system should either immediately deal with any compromise or maintain a history of compromises in order to intelligently schedule corrective action.
- The use of data reduction technologies (e.g., compression and deduplication) should be used in a manner that avoids compromising data integrity (e.g., factored into copies that might not have any association with the data reduction technologies).

7.4.2 Short to medium-term retention

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

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

- Multiple physical or logical replicas of the data should be created and preserved;²²⁾ the replicas need to be organized to be as independent as possible (e.g., geographic, administrative/management, and platform/operating system), and their number chosen according to the data's value and tolerance of risk.
- 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.
- Match the access control scheme to the legal and regulatory requirements for the information being preserved.
- Ensure that accountability and traceability measures are adequate and functional; all data accesses may require audit log entries.
- Implement mechanism to demonstrate data authenticity, provenance, and chain of custody, especially for data of an evidentiary nature.
- 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.

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

7.5 Data confidentiality and integrity

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

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

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

As mentioned in 6.8.2.3, the point of encryption is important because it represents the location within the ICT infrastructure in which the data has to traverse before it is decrypted and usable. A common security perspective is to encrypt as close to the source as possible, as this tends to maximize the protection provided, but there may be many options in selecting a point of encryption (see Figure 6), including:

- **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 operating system or operating system-level application; control at file-level with insights into the users.
- **Network-level** – under the control of the network devices, such as HBA, array controller, or switch
 - 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”²³⁾
- **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.”

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

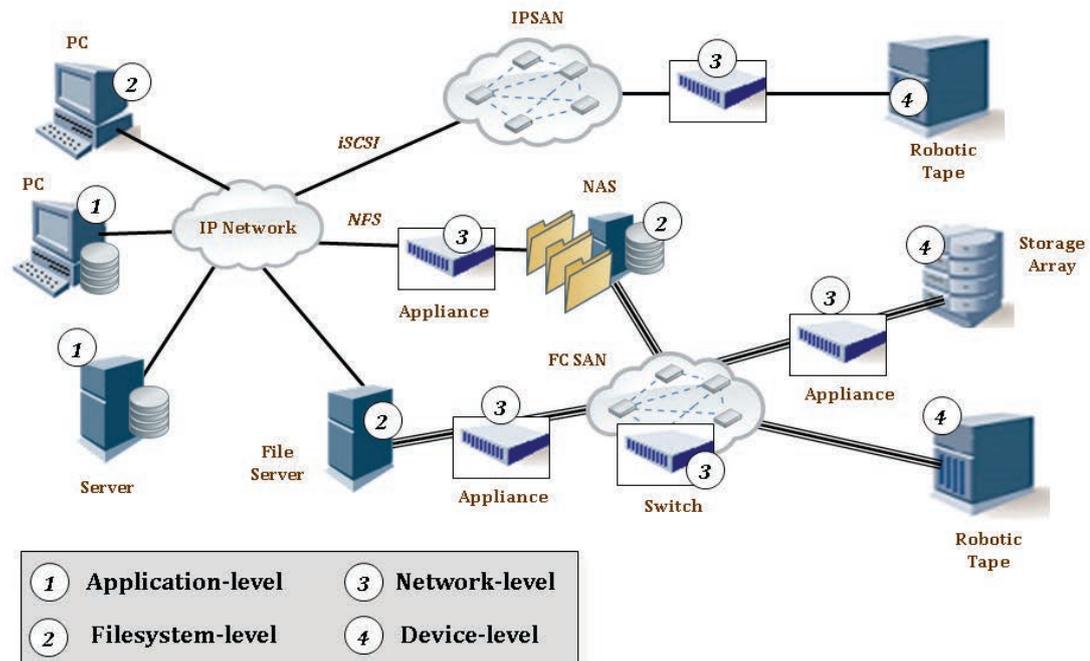


Figure 6 — Sample points of encryption

When encryption is deemed necessary, consider the following guidance:

- storage-based encryption should not be the primary form of encryption for sensitive data²⁴;
- 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 with 128 bits serving as the recommended minimum²⁵;
- cryptographic modules used to protect sensitive or regulated data should be validated using recognized criteria (e.g., ISO/IEC 19790, ISO/IEC 15408, NIST FIPS 140-2, etc.);
- multiple encryption steps can be used, as when data encrypted for privacy purposes is further encrypted by a Self-Encrypting Drive for security purposes.

As with sanitization, it is important that an organization maintain records of its data at rest encryption to document what media were protected as well as when and how they were encrypted. When an organization is suspected of losing control of its storage media, which contain sensitive information, these records or proof of encryption can be instrumental in demonstrating that no data breach

24) 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).

25) Allowing for 112 bits of security strength means that Triple DES is an acceptable option, but not recommended.

occurred, thereby avoiding costly data breach notifications and other liabilities. The following should be considered for proof of encryption:

- ensure that the encryption mechanisms create appropriate audit log entries (activation, verification, integrity checks, re-keying, etc.);
- agree in advance on what audit log material demonstrates (to the satisfaction of the compliance personnel) that encryption was properly performed;
- perform regular and audited checks that encryption was properly performed and consider outside accreditation.

Successful use of cryptography is dependent on adhering to basic principles associated with keying material as well as implementing key management. As storage systems and devices integrate encryption for data at rest, key management becomes important and should address the following:

- leverage centralized key management;
- fully automate key management whenever possible;
- 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);
- enforce strict access controls to limit user capabilities and separation of duties constraints (e.g., a security role) for key generation, change and distribution;
- for sensitive or high-value data, the encryption should be end-to-end (i.e., data in motion and data at rest).

Data integrity is a significant design criterion for most storage systems and infrastructure and it is only rivalled by data availability in its importance to storage personnel. To address data integrity issues, a wide range of technologies are typically deployed in storage infrastructure, including but not limited to, RAID, backups, replications, and CDP. Although important, these data protection technologies are not typically considered part of the storage security controls.

Data retention and compliance requirements often include provisions for storing data in a manner that blocks record deletion or alteration (i.e., immutable) along with integrity verification (e.g., hashing) and explicit retention periods (e.g., legal holds) that need to be honoured. Several forms of WORM-based storage can be used to meet the immutability (non-editable) requirements. In addition, many CAS (see [6.7.3](#)) implementations combine WORM with metadata that can be used to perform explicit integrity checks as well as enforce data expirations.

- Malware is a common threat to the integrity of data, applications, and operating system; storage systems should include sufficient malware protections to guard against attacks on data (e.g., corruption, destruction, etc.)
- WORM-based storage should be used to help meet immutability requirements

Vendors should implement the functionality for encryption, key management, and integrity described in [7.5](#) within their products.

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 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 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) controls should be applied to limit or prevent direct access to the non-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, including:
 - 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 that 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.

The following storage for virtualization guidance is relevant and should be followed:

- 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), including:
 - configure FC SAN zones and present LUNs using the VM-specific World Wide Port Names (WWPNs), so that the LUNs will only be visible to that virtual server and not to any other virtual server;
 - 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 servers in an infrastructure should be controlled to avoid having unintended security consequences, such as:
 - 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 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.

The following encryption and key management guidance is relevant and should be followed:

- comply with import/export controls, including:
 - 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 or government key escrow requirements; and
 - 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:
 - have a recovery plan in the event of a key compromise;

- have a key backup²⁶⁾ plan²⁷⁾ in place to ensure continued access to encrypted business/mission critical information²⁸⁾.
- other problem areas
 - securely distribute key material among storage devices that process/access the same data. For example data is encrypted at one node but decrypted at a second node;
 - the effect of encryption on deduplication and compression techniques should be understood and factored in designs and implementations;
 - the inability to apply security techniques like virus scanning, etc. on encrypted data should be understood and mitigated with other mechanisms.

7.7.2 Align storage and policy

ISO/IEC 27002:2013, 5.1 states that “A set of policies for information security should be defined, approved by management, published and communicated to employees and relevant external parties.” The presence or absence of policy plays a major role in assuring both security and compliance.

- Incorporate storage in policies
 - identify most sensitive (Personally Identifiable Information, intellectual property, trade secrets, etc.) and business/mission critical data categories as well as protection requirements;
 - integrate storage-specific policies with other policies (i.e., avoid creating a separate policy document for the storage ecosystem);
 - address data retention and protection (e.g., write-once-read-many or WORM, authenticity, access controls, etc.);
 - address data destruction and media sanitization.
- Conformance with policies
 - ensure that all elements of the storage ecosystem comply with policy (e.g., ISO/IEC 27001:2013, 5.2 and ISO/IEC 27002:2013, Clause 5);
 - give most sensitive/most critical data a priority.

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. In addition to the relevant compliance guidance in ISO/IEC 27002:2013, Clause 18, the following elements

26) 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.

27) 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.

28) 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.

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 that logged event/transaction data contains sufficient application 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 evidence²⁹⁾ (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.
- Legal
 - ensure that the use of data deduplication does not conflict with data authenticity requirements;
 - ensure that data and media sanitization mechanisms do not violate preservation orders;
 - ensure that proper chain of custody procedures are followed when evidentiary data (e.g., audit logs, metadata, mirror images, point-in time copies, etc.) is handled.

NOTE [Annex B](#) can be a useful resource when auditing storage systems and infrastructure.

7.7.4 Secure multi-tenancy

Multi-tenancy, as defined by Recommendation ITU-T Y.3500 | ISO/IEC 17788:2014, focuses on the “allocation of physical and virtual resources such that multiple tenants and their computations and data are isolated from and inaccessible to one another.” Secure multi-tenancy builds on this concept by adding

²⁹⁾ 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.

security controls to explicitly guard against data breaches as well as to allow for verification of the state of these controls (e.g., they are active) and validation of the controls (i.e. assurance that they work).

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

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

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

- encrypted storage that is aligned with the tenants' usage of resources;
- strong symmetric encryption (i.e., minimum of 128-bits of security strength) to protect data at rest;
- secure and rapid de-provisioning (see [Annex A](#) for media sanitization, including cryptographic erase);
- trusted third-party data storage management (e.g., SNMPv3, SMI-S with TLS³⁰⁾, etc.);
- automated key management providing tenant-controlled key management (leverages KMIP compliant servers);
- secure data replication (e.g., data in motion and at rest encryption);
- protect data from administrators (e.g., enforce a least privileges access model, administrators do not have access to the keying materials, etc.);
- highly available storage networking fabrics (multi-path and diverse path);
- centralized and secure audit logging (e.g., syslog over TLS);
- validation and certification (e.g., Common Criteria) of cryptographic modules and other security measures (e.g., media sanitization, access control, etc.).

Vendors should implement the functionality for secure multi-tenancy described in [7.7.4](#) within their products.

7.7.5 Secure autonomous data movement

Many storage systems and infrastructure have the ability to move data between different storage devices and storage elements (e.g., tiered storage), between data centres (e.g., synchronous and asynchronous data replication), to data archiving facilities, to data protection systems (e.g., backups on tape robots

30) SMI-S v1.5, which is also known as ISO/IEC 24775 *Information technology – Storage management*; is available from the Storage Networking Industry Association (SNIA). More recent versions (e.g., v1.6) are also available from the SNIA.

or virtual tape), etc. More complex scenarios exist within Information Lifecycle Management (ILM) and Data Lifecycle Management (DLM) solutions. However, all of these scenarios assume:

- data movement is policy-driven;
- intervention of operators or computers is not required to initiate or intervene throughout the process.

Because autonomous data movement takes many forms, the security needs can vary significantly; they can include some or all of the following:

- Accountability and traceability
 - configuring policies for data movement should be restricted to authenticated and authorized privileged users;
 - the individual establishing the configurations should be conversant with the security attributes of both source and destination;
 - configuration changes to implement or terminate autonomous data movement should be reflected in the audit log;
 - all autonomous data movement transactions should be reflected in the audit log of the system conducting the data movement;
- Integrity, authenticity, and immutability
 - as part of autonomous data movement transactions, the integrity of the moved data should be verified (preferably with a cryptographic hash);
 - 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);
 - autonomous data movement transactions should not negate immutability or other data preservation controls (e.g., supporting legal holds).
- Confidentiality
 - autonomous data movement transactions should not eliminate or weaken encryption controls associated with the data;
 - autonomous data movement transactions that span systems should include data in motion encryption for sensitive and high value data.
- Sanitization
 - as part of autonomous data movement transactions, the source data or storage media should be appropriately sanitized (see [6.8.1.2](#) and [6.8.1.3](#)) before it is released for re-use;
 - sanitization performed in conjunction with autonomous data movement should also include verification (see [6.8.1.5](#)) and some form of proof of sanitization (see [6.8.1.4](#)).
- Trustworthiness and physical security
 - autonomous data movement transactions should not cause data to cross security domains (e.g., production to development environments);
 - autonomous data movement transactions should not cause data to move to systems with inadequate certifications and accreditations;
 - autonomous data movement transactions should not cause data to move to systems with inadequate physical security.

Vendors should implement the functionality for secure autonomous data movement described in [7.7.5](#) within their products.

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 overwrite through the interface, or using the appropriate ATA/SCSI firmware command to overwrite both logically addressable and logically non-addressable physical media. Overwriting through the interface should include overwriting not only the logical storage location of a file (e.g., file allocation table) but also can include all addressable locations. The security goal of the overwriting process is to replace all previously written data with fixed or 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 to use block erase operations on both logically addressable and logically non-addressable physical media are acceptable methods for purging. Degaussing is not applicable to devices that contain 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.
 - Cryptographic Erase leverages the encryption of target data by enabling sanitization of the target data's encryption key. This leaves only the ciphertext remaining on the media, effectively sanitizing the data.
- **Destruct** - 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.
 - *Disintegrate*. Sanitization method designed to completely destroy the media by breaking or decompose (e.g., acid) it into constituent elements, parts, or small particles.
 - *Incinerate*. Sanitization method designed to completely destroy the media by burning until it is reduced to ashes.
 - *Melt*. Sanitization method designed to completely destroy the media by liquefying it, typically through the application of heat.
 - *Pulverize*. Sanitization method designed to completely destroy the media by grinding it to a powder or dust form.
 - *Shred*. 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

Sanitization Method	Description
Paper and microforms	
Clear/Purge:	N/A, see Destruct.
Destruct:	Destruct paper using cross cut shredders that produce particles that are 1 x 5 millimetres in size (or smaller), or pulverize/disintegrate paper materials using disintegrator devices equipped with 1,5 millimetre security screen. Destruct microforms (microfilm, microfiche, or other reduced image photo negatives) by burning. When material is burned, the residue is reduced to white ash.

Table A.2 — Networking Device Sanitization

Sanitization Method	Description
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 Destruct. 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.
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
For both Clear and (if applicable) Purge, refer to the manufacturer for additional information on the proper sanitization procedure.	
Network Devices may contain removable storage. The removable media should be removed and sanitized using media-specific techniques.	

Table A.3 — Mobile Device Sanitization

Sanitization Method	Description
Apple iPhone and iPad	
Clear/Purge:	Select the full sanitize option. 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).
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
Blackberry	
Clear/Purge:	Select the full sanitize option, 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.
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
Devices running the Google Android operating system	
Clear:	Select the full sanitize option.
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.
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
All other mobile devices <i>This includes cell phones, smart phones, personal digital assistant, 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 Destruct. 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.
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
<p>Disassembly of battery and display may be required.</p> <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> <p>Refer to the manufacturer for additional information on the proper sanitization procedure, and for details about implementation differences between device versions and operating system 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>	

Table A.4 — Equipment Sanitization

Sanitization Method	Description
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 Destruct. 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.
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
<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 Clear 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>	

Table A.5 — Magnetic Media Sanitization

Sanitization Method	Description
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 zeroes. Multiple passes or more complex values may optionally be used.
Purge:	Degauss in an organizationally approved degausser.
Destruct:	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 zeroes. Multiple passes or more complex values may optionally be used.
Purge:	Degauss in an organizationally approved degausser.
Destruct:	Incinerate disks and diskettes by burning in a licensed incinerator or Shred.
Reel and Cassette Format Magnetic Tapes <i>This also includes 8mm, DDS DAT, DLT, QIC, etc.</i>	

Table A.5 (continued)

Sanitization Method	Description
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.
Destruct:	Incinerate by burning the tapes in a licensed incinerator or Shred. 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/SSHD <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 zeroes. Multiple passes or more complex values may optionally be used.
Purge:	<p>Four options are available:</p> <p>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:</p> <p>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.</p> <p>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 zeroes 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.</p> <p>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.</p> <p>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 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 zeroes 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>

Table A.5 (continued)

Sanitization Method	Description
	<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> <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>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>
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
SCSI HDD/SSHD	<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 zeroes. Multiple passes or more complex values may optionally be used.
Purge:	<p>Four options are available:</p> <p>a) Apply the SANITIZE command, if supported. One or both of the following options may be available:</p> <ol style="list-style-type: none"> 1) 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. 2) 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 zeroes or a pseudorandom pattern across the media. If the OVERWRITE service action is not supported, the Clear procedure could alternatively be applied. <p>b) 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 zeroes or a pseudorandom pattern across the media. If the OVERWRITE service action is not supported, the Clear procedure could alternatively be applied.</p> <p>c) 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 zeroes. Multiple passes or more complex values may alternatively be used.</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> <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 of magnetic recording features across device manufacturers.</p>
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.

Table A.5 (continued)

Sanitization Method	Description
	<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>The storage device may support configuration capabilities that artificially restrict the ability to access portions of the media, such as</p> <ul style="list-style-type: none"> — the SCSI mode parameter block descriptor’s NUMBER OF LOGICAL BLOCKS field (accessible with the MODE SENSE and MODE SELECT commands) — a Host Protected Area (HPA), Device Configuration Overlay (DCO), or Accessible Max Address (all defined in the ATA standard) <p>Even when dedicated ATA/SCSI sanitize commands address 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>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> <p>For destruct techniques other than incinerate, the media should be appropriately degaussed before destruction and the resulting pieces should be smaller than a 12 mm grid.</p>

Table A.6 — Peripherally Attached Storage Sanitization

Sanitization Method	Description
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 zeroes. Multiple passes or more complex values may optionally be used.
Purge:	<p>See Destruct. 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, cryptographic erase, etc.) to ensure that data recovery is infeasible, and that the device does not simply remove the file pointers.</p>
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
<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> <p>For destruct techniques other than incinerate, the media should be appropriately degaussed before destruction and the resulting pieces should be smaller than a 12 mm grid.</p>	

Table A.7 — Optical Media Sanitization

Sanitization Method	Description
CD, DVD, BD	
Clear/Purge:	N/A, see Destruct.
Destruct:	<p>Destroy in order of recommendations:</p> <ol 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. <p>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.</p> <p>For destruct techniques other than incinerate on flash-based storage, the resulting pieces should be smaller than a 2 mm grid.</p>

Table A.8 — Flash-Based Storage Device Sanitization

Sanitization Method	Description
ATA SSD <i>This includes PATA, SATA, eSATA, CF, CFast etc.</i>	
Clear:	<p>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 zeroes. 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> <p>b) Use the SECURITY ERASE UNIT command, if supported.</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>
Purge:	<p>Three options are available:</p> <p>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:</p> <ol 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.
	<p>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.</p> <p>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.</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>
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
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 zeroes. 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>

Table A.8 (continued)

Sanitization Method	Description
Purge:	<p>Two options are available:</p> <p>a) Apply the SANITIZE command, if supported. One or both of the following options may be available:</p> <ol style="list-style-type: none"> 1) The BLOCK ERASE service action. 2) 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. <p>b) 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.</p>
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
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 zeroes. Multiple passes or more complex values may optionally be used.
Purge:	<p>Two options are available:</p> <p>a) Apply the NVM Express Format command, if supported. One or both of the following options may be available:</p> <ol style="list-style-type: none"> 1) The User Data Erase command. 2) 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. <p>b) 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.</p>
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
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:	<p>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.</p> <p>For most cases where Purging is desired, USB removable media should be Destructed.</p>
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
Memory Cards <i>This includes SD, SDHC, MMC, Compact Flash, Microdrive, MemoryStick, etc.</i>	

Table A.8 (continued)

Sanitization Method	Description
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 Destruct.
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
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:	N/A, See Destruct. If the flash can be easily identified and removed from the board, the flash may be Destructed independently from the disposal of the board that contained the flash. Otherwise, the whole board should be Destructed.

Sanitization Method	Description
Destruct	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.
<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 ability to access portions of the media, such as</p> <ul style="list-style-type: none"> — the SCSI mode parameter block descriptor's NUMBER OF LOGICAL BLOCKS field (accessible with the MODE SENSE and MODE SELECT commands) — a Host Protected Area (HPA), Device Configuration Overlay (DCO), or Accessible Max Address (all defined in the ATA standard) <p>Even when dedicated ATA/SCSI sanitize commands address 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 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> <p>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 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> <p>For destruct techniques other than incinerate on flash-based storage, the resulting pieces should be smaller than a 2 mm grid.</p>	

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

Sanitization Method	Description
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. In either case, the DRAM should remain without power for a period of at least five minutes.
Destruct:	Shred, Disintegrate, or Pulverize.
Electronically Alterable PROM (EAPROM)	
Clear/ Purge:	Perform a full chip Purge as per manufacturer’s data sheets.
Destruct:	Shred, Disintegrate, or Pulverize.
Electronically Erasable PROM (EEPROM)	
Clear/Purge:	Overwrite media by using organizationally approved and validated overwriting technologies/methods/tools.
Destruct:	Shred, Disintegrate, Pulverize, or Incinerate by burning the device in a licensed incinerator.

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

A.3 Cryptographic erase device guidelines

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

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

- the strength of the cryptographic algorithm used to encrypt the data (including mode of operation);
- the level of entropy of the target data’s encryption.

As a result, sanitization of the data is reduced to sanitization of the encryption key(s) used to encrypt the data. With cryptographic erase, sanitization may be performed with high assurance much faster than with other sanitization techniques. The encryption itself acts to sanitize the data.

Typically, cryptographic erase can be executed in seconds. This is especially important as storage devices get larger resulting in other sanitization methods taking more time. Cryptographic erase can also be used as a supplement or in addition to other sanitization approaches.

Reliance upon cryptographic erase to purge the media on devices should not occur if:

- the encryption was enabled after sensitive data was stored on the device without having been sanitized first; or
- if it is unknown whether sensitive data was stored on the device without being sanitized prior to encryption, then cryptographic erase should not be relied upon as a Purge mechanism.

Where cryptographic erase is intended for use to purge the media (including SEDs, mobile devices, and other devices), the level of assurance depends on the following:

- Encryption of all data intended for cryptographic erase prior to storage on the device (including the data, as well as virtualized copies);
- Locations on the media where the data encryption key is stored (be it the target data’s encryption key or an associated wrapping key) are directly accessible for sanitization (ensuring the actual

location on media where the key is stored is addressed) using the appropriate media-specific sanitization technique;

- All copies of the encryption keys used to encrypt the target data are sanitized;
- If the target data’s encryption keys are, themselves, encrypted with one or more wrapping keys, it is acceptable to perform cryptographic erase by sanitizing a corresponding wrapping key;
- And, the ability of a user to clearly identify the commands provided by the device to perform the cryptographic erase operation.

Other cryptographic erase considerations:

- If the encryption key (or any key at or below the level of key sanitized during cryptographic erase) exists outside of the storage device (typically due to escrow or injection), there is a possibility that the key could be used in the future to recover data stored on the encrypted media.
- Sanitization using cryptographic erase should not be trusted on devices that have escrowed or injected the key(s) unless the organization has a high level of confidence about how and where the keys were stored and managed outside the device. Such back-up or escrowed copies of data, credentials, or keys should be the subject of a separate device sanitization policy. That policy should address backups or escrowed copies within the scope of the devices on which they are actually stored.

The choice regarding whether to leverage cryptographic erase on a given device depends upon organizational 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 level of the key escrowed has/have ever been escrowed from or injected into the device. If the MEK encryption key is directly sanitized and only a KEK can be escrowed, clearly identify that fact.
- **Error Condition Handling:** Identify how the device handles error conditions that prevent the cryptographic erase operation from fully completing, such as if a defect is encountered where an instance of the key to be sanitized is stored. For example, if the location where the key was stored cannot be sanitized, does the cryptographic erase operation report success or failure to the user?
- **Interface Clarity:** Identify which interface commands support the features described in the statement. If the device supports the use of multiple MEKs, identify whether all MEKs are changed using the interface commands available and any additional commands or actions necessary to ensure all MEKs are changed.

Implementers who choose to apply cryptographic erase should seek either independent validation of these assurance areas or ask the vendor to identify which mechanisms are used to ensure that these concern areas have been addressed. Generally accepted and (where applicable) standardized mechanisms should be used. For example, security requirements for cryptographic requirements are specified in ISO/IEC 19790:2006 and test requirements for cryptographic modules are specified in ISO/IEC 24759:2008. These requirements and tests cover some (but not all) of the concern areas.

The decision regarding whether to rely upon cryptographic erase should also consider whether the Media Encryption Key has been escrowed or injected, and if so, how the key was protected outside of the storage device. If the Media Encryption Key (or any key at or below the level of key sanitized during cryptographic erase) exists outside of the storage device, there is a possibility that the key could be used in the future to recover data stored on the 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 clauses³¹⁾ along with information that can serve as selection criteria based on:

- **Data sensitivity classes** – provides a data-centric focus that leverages two classes, which can be used by organization that have performed basic data classifications;
- **Security priority codes** – provides information to assist in making phasing or sequencing decisions for control implementation, based on the relative importance of the confidentiality, integrity, and availability aspects of security.

Organizations should consider these criteria and guidelines as starting points for the selection of storage security controls (see also ISO/IEC 27002:2013,8.2). They can also help organization implement a phased approach to using storage security controls, based on risk assessment (see ISO/IEC 27005).

It is not appropriate 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, two generic data sensitivity classes or levels are defined: Low (see [B.1.2.2](#)) and High (see [B.1.2.3](#)).

As a starting point, organizations need to map their specific data classifications to one of the two data sensitivity classes defined in this [Annex B](#). The summary of controls listed in the tables in [B.2](#) can then be consulted to identify the relevant storage security controls; within these tables, a data sensitivity of “L” corresponds to “Low” and “H” corresponds to “High.”

31) All of the storage security controls in this annex were extracted from [Clauses 6](#) and [7](#), and they are summarized herein. When additional information or clarification is needed, consult the appropriate source clauses.

B.1.2.2 Low data sensitivity

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

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

- have limited negative effects to business, but not breach contractual or legal agreements or laws;
- have limited negative effects to government;
- affect individuals in a limited way in their social and economic circumstances.

B.1.2.3 High data sensitivity

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

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

- have significant, and possibly existence-threatening, negative effects to business and could breach contractual or legal agreements or laws;
- constitute a breach of government security that exposes confidential or possibly secret data;
- affect individuals substantially in their social and economic circumstances; in an extreme situation, endanger individuals in their health, life or personal liberty.

B.1.3 Security priority codes

Organizations can identify and prioritize storage security controls that are most relevant to their environments. This can be accomplished by first conducting an information security risk assessment (see ISO/IEC 27005), which can help identify the data confidentiality, integrity, and availability needs and the relative importance of these security aspects.

If the organization's data confidentiality, integrity, and availability priorities are understood, they can be leveraged to help identify relevant storage security controls, using the tables in [B.2](#). In addition, it may also be possible to address these controls in a phased approach, starting with those controls that mitigate the highest priorities risks identified in the risk assessment.

To assist with such an effort, the storage security controls from [Clauses 6](#) and [7](#) are summarized in the tables in [B.2](#). Each control is shown with a set of priority values for each of the security aspects (confidentiality, integrity, and availability) as well as an indicator for system-wide³²⁾ (i.e., the priority codes are identical for all three security aspects). For the [B.2](#) tables, a priority indicator of "C" corresponds to "Confidentiality," "I" corresponds to "Integrity," "A" corresponds to "Availability," and "S" corresponds to "System-wide," and the priority values used in the tables are in the range of 0 to 5, with 5 representing the highest priority and 0 is the lowest priority; "System-wide" designations are identified with an "X."

To demonstrate how an organizations might use the priority code data, consider a scenario where a risk assessment has identified data confidentiality as the primary problem area. By selecting the confidentiality security aspect as the focus point, the organization is able to review the security controls and identify those controls with the highest values for the "C" priority code. Those controls with a value of 5 are most likely to be applicable. As such, the organization could start implementing the controls

32) A "System-wide" designation for a storage security control is simply a way of showing that there may be a special interdependency between the three security aspects that may warrant special attention.

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 (6.2)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
DAS should be physically secured		5	3	5	X	X
For sensitive and high value data on DAS, some form of encryption (SED, FDE, computer-based, or application-based) should be used to protect the data at rest		5	3	0		X
Media sanitization should be used on all DAS involved with sensitive and high value data		5	1	0		X
If possible, authentication such as FC-SP-2 AUTH-A Authentication should be used to prevent unauthorized access to sensitive and high value data		5	3	0		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

Table B.2 — Storage networking (6.3)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Storage Area Networks (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
Logically segregate storage traffic from normal server traffic	X	4	4	4		X
Segregate storage management traffic from all other traffic		2	3	1	X	X
Ensure that configurations of network gateways maintain appropriate network segregation		3	3	4	X	X
For FC, restrict server access on the switches using techniques such as ACLs, binding lists, FC-SP-2 fabric policies		4	3	3		X
For FC, use NPIV enabled HBAs to assign individual N_Port_IDs to virtual servers		5	3	2		X
For FC, restrict switch interconnections using techniques such as ACLs, binding lists, FC-SP-2 fabric policies		4	3	2		X
Zoning should be used in FC SAN fabrics with a preference for hard zoning		4	3	3	X	X
For FC, determine whether basic zoning is a strong enough security measure for the target environment, and if it is not, use stronger techniques like FC-SP Zoning where supported by the vendor		4	3	3		X

Table B.2 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
For FC, disable unused switch ports		3	4	1	X	X
For FC, carefully use default zones and zone sets (assume a least privilege posture)		3	3	1	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, use VLANs when the use of physically isolated LANs is not an option	X	5	5	5	X	X
Set up the peer-to-peer relationship between FCIP entities		5	3	5	X	X
When possible, a private IP network should be used exclusively by the FCIP entities		5	3	5		X
For FCIP, use IPsec to perform cryptographic authentication and data integrity at a minimum		3	4	3		X
For FCIP, use IPsec to protect sensitive data by appropriate confidentiality measures		5	4	3		X
For FCoE, leverage the Fibre Channel security mechanisms		5	3	2	X	X
For FCoE, protect against Ethernet broadcast storms (e.g., allocation of adequate input buffering) that can cause throughput and timeout issues		0	1	3	X	X
ACLs should be used to control FCoE network access (e.g., denying specific computers from unnecessary or unwanted traffic)		5	4	1	X	X
Use FCoE VLANs when the use of physically isolated LANs is not an option	X	5	5	5	X	X
Network Attached Storage (6.3.3)						
Extra care should be exercised when NFSv3 is used with sensitive or high-value data		4	3	1		X
Enable NFS only if it is needed		3	3	1	X	X
Use NFSv4 (or later versions) whenever possible and limit NFSv3 usage		3	3	1	X	X
For NFS, filter client and management access by IP address for additional security		2	3	4	X	X
For NFS, encrypt client data access (e.g., IPsec) when necessary		5	5	2		X
Use later versions of the SMB protocol		3	4	5	X	X
For SMB/CIFS, turn off low-security session negotiation protocols, such as NTLM v1, LanMan and plaintext, and use NTLM v2 or Kerberos instead	X	5	5	5	X	X
Maintain up-to-date patch levels	X	4	4	4	X	X
Use SMB signing		5	5	0	X	X
Maintain Active Directory (AD) services securely		3	3	5	X	X
Use one-way trusts, from leaf domains to parent domains, when possible		5	5	2	X	X
Enable SMB/CIFS only if it is needed		3	3	1	X	X
For SMB/CIFS, encrypt client data access (e.g., IPsec) when necessary		4	3	0		X

Table B.3 — Storage management (6.4)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Authentication and authorization (6.4.2)						
All users should have a unique ID for their personal use		5	5	0	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
For all remote access, use strong authentication or multi-factor authentication along with secure channels		5	5	0	X	X
When possible, use a centralized authentication solution for improved monitoring and control		4	4	0	X	X
Use multi-factor authentication when managing sensitive and high-value data		4	4	0		X
Disable login to the root account. Remotely log all privilege escalation operations		3	3	0	X	X
When possible, use entity authentication in TLS and IPsec connections as well as within storage protocols		3	3	0		X
Implement and use general roles like Security Administrator, Storage Administrator, Security Auditor, and Storage Auditor within storage systems		3	3	0	X	X
Secure the management interfaces (6.4.3)						
Restrict physical access to management interfaces	X	5	5	5	X	X
Disable and disconnect serial management ports when not in use		2	2	1	X	X
Segregate LAN interfaces used for management from other LAN traffic, noting that physical isolation is preferred, but logical isolation (such as VLANs) should be used at a minimum		3	2	1	X	X
Use firewalls and TCP wrappers to restrict access to management networks to authorized systems and protocols		4	4	5	X	X
Use entity authentication to establish trust relationships between storage systems and the management systems (e.g., using FC-SP-2 AUTH-A to authenticate the entities performing in-band management)		3	3	4		X
Leverage IDS and IPS 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
For storage management, ensure that operating systems and applications are current and sufficiently hardened against attacks		5	5	3		X
For remote storage management, use secure channels for all remote access (VPN, TLS, SSH, 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
Devise organizational and technical controls to restrict the management interface used for remote (non-local) vendor maintenance sessions		3	2	2	X	X

Table B.3 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Technical controls should restrict communication traffic (i.e. systems, ports, and protocols) to the minimum required for remote vendor maintenance operations		3	3	5	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
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 call-back 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 (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
In the policies, identify and address the evidentiary expectations for storage logs	X	5	5	5	X	X
Employ external and centralized event logging to a trusted remote source	X	5	5	5	X	X
Establish and use a common, accurate time source across the storage systems and infrastructure	X	5	5	5	X	X
Natively log events to one, and preferably multiple, external log servers	X	4	4	4	X	X
Use standard logging protocols like syslog that support reliable delivery and secure transports	X	3	3	3	X	X
For compliance, accountability, or security purposes, events should be logged as they occur (no buffering)	X	4	4	4	X	X
Implement an analysis protocol to correlate audit log records across event sources to identify significant security events that provide indication of security incidents	X	3	3	3	X	X
Ensure that the storage logging is factored into SIEM solutions, when such technology is deployed	X	3	3	3		X
Log all occurrences of the minimum set of security events with the necessary data	X	5	5	5	X	X
Audit log data that may have evidentiary value should be handled correctly (e.g., maintain chain of custody, verifiable integrity and authenticity, etc.)	X	5	5	5	X	X
Audit log data with specific retention requirements (e.g., for regulatory compliance) should be preserved with the organization's data retention solution	X	5	5	5	X	X
Implement appropriate measures to preserve log integrity and prevent their modification or destruction	X	5	5	5		X
When audit log entries contain sensitive information, the audit log data should be protected with appropriate confidentiality mechanisms	X	5	5	5		X
For unique audit logging requirements (e.g., high volume, special preservation, event signing, etc.) dedicated and specially hardened and configured systems should be used	X	4	4	4		X

Table B.3 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Leverage log relays and log filtering to minimize the impact of specialized storage requirements (WORM)	X	3	3	3		X
System hardening (6.4.5)						
All operating systems, hypervisors, and applications should be hardened relative to the use of the storage system	X	3	3	3	X	X
Remove un-needed/un-used software		2	3	3	X	X
Remove unnecessary accounts	X	3	3	3	X	X
Eliminate, disable, or change passwords on predefined or default accounts	X	4	4	4	X	X
Only open up network ports that are needed		1	1	3	X	X
Install latest patches from a trusted source	X	4	4	4	X	X
Update firmware from a trusted source	X	4	4	4	X	X
Install and maintain malware protection	X	5	5	5	X	X

Table B.4 — Block-based storage (6.5)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Fibre Channel storage (6.5.1)						
LUN masking and mapping (WWN filtering) and other access control mechanisms should be used to restrict access to storage		4	4	1	X	X
Mutual authentication using FC-SP-2 AUTH-A should be used with all servers and switches; leverage centralized authentication services when possible		2	2	5	X	X
If possible, Fibre Channel connections that leave the protected area should be encrypted using ESP_Header		4	3	1		X
Sensitive and high-value data should be encrypted while on the FC storage device or media		5	3	1		X
Encryption should be implemented in FC storage devices that may come in contact with sensitive or regulated data as well as to facilitate rapid sanitization		5	3	1		X
For FC storage, media-aligned sanitization should be used for sensitive and regulated data		5	3	1		X
Logical sanitization should be used to clear virtualized FC storage, especially when the actual storage devices and media cannot be determined		5	3	1		X
IP storage (6.5.2)						
Control iSCSI initiator access by filtering based on source IP addresses and protocols		5	3	5	X	X
Bidirectional CHAP authentication, using random challenges (i.e., not repeated), should be used for both initiators and targets in all iSCSI implementations	X	5	5	5	X	X
IPsec should be used to secure the communication channel when sensitive or high-value data could be exposed		5	3	2		X

Table B.4 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
iSNS, SLP, DNS infrastructure should be used with appropriate security controls to avoid indirect attacks	X	3	3	3	X	X
Sensitive and high-value data should be encrypted while on the IP storage device or media		5	3	1		X
Encryption should be implemented in IP storage devices that may come in contact with sensitive or regulated data as well as to facilitate rapid sanitization		5	3	1		X
For IP storage, media-aligned sanitization should be used for sensitive and regulated data		5	3	1		X
Logical sanitization should be used to clear virtualized IP storage, especially when the actual storage devices and media cannot be determined		5	3	1		X

Table B.5 — File-based storage (6.6)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
NFS-based NAS (6.6.1)						
Employ user-level authentication whenever possible (e.g., NFSv4 with Kerberos V5)	X	5	5	5	X	X
Configure the NFS server to export file systems explicitly for the authorized users		3	2	1	X	X
Configure the NFS server to export file systems with minimum required privileges		3	2	1	X	X
Avoid granting “root” or “administrator” access to files on network filesystems		5	5	3	X	X
Make sure NFSv4 ACLs are assigned correctly		4	4	2	X	X
Use Kerberos authentication for NFSv3	X	3	3	3	X	X
Consider using Kerberos Safe and Private modes to sign and encrypt NFS traffic		4	4	2	X	X
Filter client access to NFS shares whenever possible	X	3	3	3	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
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
For NFS, encrypt data at rest when necessary		3	3	1		X
Do not allow NFS exports of administrative file systems (e.g., /etc)	X	3	3	3	X	X
For NFS, 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
SMB/CIFS-based NAS (6.6.2)						
Disable unauthenticated access to CIFS shares and NAS devices (i.e. restrict <i>Anonymous</i>)		5	3	4	X	X
Disable “Guest” and “Everyone” access to all CIFS shares		4	4	2	X	X

Table B.5 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
For SMB/CIFS, implement authentication and access control via a centralized mechanism (RADIUS, LDAP)	X	5	5	5		X
Enable SMB signing for clients and the NAS device		3	5	3		X
Enable CIFS auditing whenever possible		3	3	1		X
Continually review content placed in CIFS shares and relevant access controls		4	1	2		X
For SMB/CIFS, encrypt data at rest when necessary		3	3	1		X
For SMB/CIFS, guard against malware (e.g., viruses, worms, rootkits, etc.)	X	5	5	5		X
Implement CIFS with strong authentication (NTLMv2, Kerberos)		4	4	1	X	X
Parallel NFS-based NAS (6.6.3)						
Controls and control mechanisms should be applied consistently across clusters (both symmetric and asymmetric)		3	5	3	X	X
Security assurance properties should not be dependent on the client accessing a specific fileserver	X	3	3	3	X	X
For asymmetric clusters, controls should be implemented such that they are consistent across different protocols	X	4	4	4	X	X
Security controls should not be dependent on path traversal of the filesystem namespace across servers		2	2	4	X	X

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

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Cloud computing storage (6.7.1)						
For cloud storage, ensure that transport security such as IPsec or TLS is used for all transactions		5	4	2	X	X
When sensitive data is stored in a third party cloud environment, data at rest encryption (and appropriate key management processes) should be used to prevent access by the unauthorized parties (e.g., cloud service provider personnel, other tenants, adversaries, etc.)		5	2	2		X
For cloud storage, secure user registrations and use strong password authentication to protect access to data		5	4	3	X	X
For cloud storage, employ access controls that guard against unauthorized access from other tenants while providing appropriate access privileges to users permitted to access the data		4	4	2	X	X
Use the provided sanitization capabilities to clear sensitive data from the cloud computing storage		4	2	2		X
When using CDMI, ensure that TLS is used for all transactions	X	4	4	4	X	X
Query 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
Authenticate CDMI entities (certificates for servers and HTTP basic authentication for clients)	X	5	5	5	X	X

Table B.6 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Use CDMI Domains to provide a place for authentication mappings to external authentication providers		4	4	1		X
Enable CDMI security logging and retrieve the security event data out of the appropriate logging queue on a regular and timely fashion		3	3	1		X
Align the automatic deletion capability (CDMI Deletion) with the organization's data retention policy	X	3	3	3		X
Prior to using CDMI Holds, understand the process and mechanism for lifting the CDMI Hold	X	4	4	4		X
Use data at rest encryption measures to protect sensitive and high-value data		4	1	1		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
Use the CDMI sanitization functionality to clear sensitive data from the cloud service provider's storage	X	3	3	3		X
Object-based Storage Device (6.7.2)						
For OSD, IPsec should be used for all transactions involving sensitive data on insecure networks		5	5	0		X
For OSD, the object store should verify the authenticity of the capability prior to performing an operation	X	5	5	5	X	X
Clock synchronization between the OSD and the security manager should be implemented using a secure protocol	X	4	4	4	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
For OSD, working keys (used to generate capability keys) should be refreshed frequently	X	3	3	3		X
Content Addressable Storage (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
The CAS system should ensure that content will be readable and accessible over its entire life-cycle.		0	5	5	X	X
The CAS system should employ a robust hashing mechanism		0	4	4	X	X

Table B.7 — Storage security services (6.8)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Data sanitization (6.8.1)						
Organizations and individuals should categorize their 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
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
Disposal of storage devices or storage elements without sanitization should be considered only if information disclosure would have no impact on organizational mission, would not result in damage to organizational assets, and would not result in financial loss or harm to any individuals		5	0	1	X	X
When sanitization is an element of compliance, the specific requirements and associated specifications should be reviewed to determine whether they mandate particular overwrite techniques, documentary proof of sanitization, etc.		5	3	1	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
When storage media are transferred, become obsolete, are no longer usable, or are not needed by an information system, the residual magnetic, optical, electrical, or other representation of data should be sanitized.		5	0	1	X	X
Annex A should be used to determine recommended sanitization of specific media.		5	0	1	X	X
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.		5	0	1	X	X
Sanitization of media at end-of-use situations is recommended, even when using encryption methods.		3	0	1	X	X
If the logical storage is writeable, then sanitization, using an overwrite or cryptographic erase technique, should be used to clear the portions of the underlying storage media used by the logical storage; successful application of cryptographic erase for sanitization is predicated on the encryption being active before data is recorded on the logical storage.		2	0	1	X	X
Data protection technologies, which can include replication, backups and CDP storage, are often used in conjunction with logical storage, so separate sanitization operations should be performed on storage associated with data protection mechanisms to eliminate sensitive or high value data.		2	0	1		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
A certificate of sanitization should be produced and contain the appropriate details	X	4	4	4	X	X
The audit trail associated with sanitization should capture time stamped transactions and progress.	X	4	4	4	X	X

Table B.7 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
A full verification of the sanitization should be performed if time and external factors permit.	X	4	4	4	X	X
If cryptographic erasure is used for sanitization, appropriate verification should be performed.	X	4	4	4	X	X
Data confidentiality (6.8.2)						
When data in motion encryption is need, it should provide end-to-end protection		3	3	1		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
For IPsec, version 3 and IKE version 2 (or later versions) should be used	X	5	5	5		X
For TLS, storage clients should comply with the requirements in the SNIA Technical Position: <i>TLS Specification for Storage Systems v1.0</i> (or the latest version)	X	5	5	5		X
To provide a basic level of protection against data breaches associated with loss of control of media, encryption mechanisms within storage devices, switches, specialized appliances, HBAs, etc. should be used		5	2	2		X
For at rest encryption, algorithms and modes of operations designed specifically for storage technology should be used	X	5	5	5		X
Limit the amount of time a key is in plaintext form and prevent humans from viewing plaintext keys		5	3	3		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
Randomly choose keys from the entire keyspace		4	3	3		X
Check for and avoid use of known weak keys	X	3	3	3		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
When possible, storage systems and infrastructure should use interoperable, centralized key management infrastructure (e.g., generate and archive encryption keys)	X	3	3	3		X
Storage systems and infrastructure should use OASIS approved, KMIP-compliant clients to access and use key management infrastructure	X	3	3	3		X
Data reductions (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 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 (7.2)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Defence in depth (7.2.1)						
Ensure a balanced focus on the three primary elements: people, technology, and operations	X	5	5	5	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
Deploy protection mechanisms at multiple locations to resist all classes of attacks	X	3	3	3	X	X
Deploy multiple defence mechanisms (layered) between potential adversaries and targets	X	3	3	3	X	X
Include both detection and protection mechanisms	X	3	3	3	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
Maintain visible and up to date system security policies	X	3	3	3	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
Perform regular security threat assessments to determine the continued security readiness	X	3	3	3	X	X
Monitor and react to current threats	X	4	4	4	X	X
Security domains (7.2.2)						
Storage and storage networks of different sensitivity levels should be located in different security domains		4	4	3		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
Strategic assets should be located in dedicated security domains		5	3	2		X
Untrusted devices and computer systems should have 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
Storage networks should be in different security domains than regular networks (e.g., corporate LANs)		2	4	2	X	X
Storage device and storage network management systems should be located in dedicated security domains		1	3	1	X	X
Systems in development stage should be located in different domains than production systems		5	3	3	X	X

Table B.8 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
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
Design resilience (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
The designs should also use a wide set of approaches destined to make the storage more resilient to attacks and network failures	X	4	4	4	X	X
Secure initialization (7.2.4)						
As a design principle, the architecture of storage systems 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 operating system 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

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

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Reliability (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
Vulnerabilities should be proactively managed to minimize their impacts on system reliability		1	4	4	X	X
Controls should be assessed to determine whether they are capable of assuring the reliability and security of data	X	3	3	3	X	X
Availability (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
Data encryption keys should be managed to avoid data availability problems when keys are unavailable or inadvertently destroyed		1	1	5	X	X
Data protection mechanisms should be part of availability designs to guard against major outages due to system failures		1	3	4	X	X
Backups and replication (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
Ensure that the backup approach, especially for business/mission critical data, is aligned with its associated restore strategy		1	4	5	X	X

Table B.9 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Ensure that the backup approach provides adequate protections against unauthorized access (e.g., encryption)		4	4	1		X
Establish a chain of trusted individuals (and vendors) who handle the storage media	X	5	5	5		X
Implement backup validations to show “proof” that restore requirements are being met		0	5	5		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
Ensure that the replication approach provides adequate protections against unauthorized access (e.g., encryption in motion)		5	0	3		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 that prioritize network traffic in order to reduce the impact of CDP on day-to-day operation		0	2	2		X
Ensure that the CDP approach provides adequate protections against unauthorized access (e.g., data in motion and data at rest encryption)		4	3	2		X
Disaster Recovery and Business Continuity (7.3.4)						
Ensure that 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
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 (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 technologies	X	5	5	5	X	X
Redundancy should be exploited to the extent possible		0	5	5	X	X
Diverse components that are easily repairable should be used whenever possible		2	2	4	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

Table B.10 — Data retention (7.4)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Long-term retention (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
When migrating archival data to newer storage technologies, introduce available security capabilities that offer enhanced security measures to better secure the data in its new location	X	2	2	2		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
Ensure that the secrecy mechanisms function in the complete absence of the user that wrote the data		5	0	0		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
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
The use of data reduction technologies (e.g., compression and deduplication) should be used in a manner that avoids compromising data integrity		1	4	1		X
Short to medium-term retention (7.4.2)						
Multiple physical 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/operating system), 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
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 (7.5)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Storage-based encryption should not be the primary form of encryption for sensitive data		4	3	2		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

Table B.11 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
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 with 128 bits serving as the recommended minimum		4	4	1	X	X
Cryptographic modules used to protect sensitive or regulated data should be validated using recognized criteria	X	5	5	5		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
Ensure that 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 compliance personnel) 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
Leverage centralized key management	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
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
For sensitive or high-value data, the encryption should be end-to-end (i.e., data in motion and data at rest)	X	3	3	3		X
To protect data integrity, storage systems should include sufficient malware protections to guard against attacks on data	X	4	4	4	X	X
WORM-based storage should be used to help meet immutability requirements		2	4	2	X	X

Table B.12 — Virtualization (7.6)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Storage virtualization (7.6.1)						
Controls on physical storage entities should be applied to the entire domain so that the controls are not bypassed when information is relocated.	X	4	4	4	X	X
Appropriate storage networking controls should be applied to the 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	X	4	4	4	X	X
Ensure appropriate service level objectives for virtual storage	X	4	4	4	X	X

Table B.12 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Match the availability objective for the storage infrastructure to the application requirements		0	0	5	X	X
Match the confidentiality and privacy requirements for the storage infrastructure to the types of information stored		5	0	0		X
Address multi-tenancy concerns, as appropriate		3	3	2	X	X
Storage for virtualized systems (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
NPIV should be leveraged appropriately to limit VM access to storage targets	X	3	3	3	X	X
VM migration/movement between physical servers in an infrastructure should be carefully controlled to avoid having unintended security consequences	X	3	3	3	X	X

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

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Encryption and key management issues (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 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
Have a key backup plan in place to ensure continued access to encrypted business/mission critical information	X	5	5	5		X
Securely distribute key material among storage devices that process/access the same data		5	5	3		X
The effect of encryption on deduplication and compression techniques should be understood and factored in designs and implementations	X	4	4	4		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
Align storage and policy (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
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
Address data retention and protection (e.g., write-once-read-many or WORM, authenticity, access controls, etc.)		3	4	2		X
Address data destruction and media sanitization		4	1	1		X

Table B.13 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Ensure that all elements of the storage ecosystem comply with policy (e.g., ISO/IEC 27001/27002)	X	3	3	3	X	X
Give most sensitive/most critical data a priority	X	4	4	4		X
Compliance (7.7.3)						
Ensure that users, especially privileged users, have unique userids (i.e., no shared accounts)		5	0	0	X	X
When possible, grant rights and privileges based on roles	X	3	3	3	X	X
Log all attempted (successful and unsuccessful) management events and transactions	X	4	4	4	X	X
Ensure that logged event/transaction data contains sufficient application or system detail to clearly identify the source	X	3	3	3	X	X
Ensure that the user information can be traced to a specific individual	X	3	3	3	X	X
When appropriate, treat log records as evidence (chain of custody, non-repudiation, authenticity, etc.)	X	4	4	4	X	X
Ensure that the storage layer participates in the external audit logging measures	X	4	4	4	X	X
Monitor the audit logging events and issue the appropriate alerts	X	5	5	5	X	X
Implement appropriate data retention measures		0	4	0	X	X
Implement appropriate data integrity and authenticity measures		4	5	0	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
Implement appropriate data confidentiality measures to prevent unauthorized disclosure		5	0	0		X
Ensure that the use of data deduplication does not conflict with data authenticity requirements		3	4	0		X
Ensure that data and media sanitization mechanisms do not violate preservation orders	X	5	5	5		X
Ensure that 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
Secure multi-tenancy (7.7.4)						
Provide secure isolation while still delivering the management and flexibility benefits of shared resources that assures no tenant can determine the existence or identity of any other tenant, can access the data in motion (network) of any other tenant, can access the data at rest (storage) of any other tenant, can perform an operation that affects an operation performed by another tenant, or can perform an operation that might deny service to another tenant as well as assuring each tenant can have a configuration that is independent of other tenant's existence and configuration	X	4	4	4	X	X

Table B.13 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
When a resource (compute, storage or network) is decommissioned from a tenant the resource should be sanitized of all data and configuration information	X	4	4	4	X	X
Accountability and traceability measures are available at the tenant level	X	4	4	4	X	X
Use encrypted storage that is aligned with the tenants' usage of resources		4	4	0	X	X
Use strong symmetric encryption (i.e., minimum of 128-bits of security strength) to protect data at rest		5	3	0	X	X
Use secure and rapid de-provisioning (media sanitization, including cryptographic erase)		5	0	0	X	X
Use trusted third-party data storage management (e.g., SNMPv3, SMI-S with TLS, etc.)	X	4	4	4	X	X
Use automated key management providing tenant-controlled key management (leverages KMIP compliant servers)	X	3	3	3	X	X
Use secure data replication (e.g., data in motion and at rest encryption)		5	0	0	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
Use highly available storage networking fabrics (multi-path and diverse path)		0	0	4	X	X
Use centralized and secure audit logging (e.g., syslog over TLS)		4	4	0	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
Secure autonomous data movement (7.7.5)						
Configuring policies for data movement should be restricted to authenticated and authorized privileged users	X	2	2	2	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
Configuration changes to implement or terminate autonomous data movement should be reflected in the audit log	X	2	2	2	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
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
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
Autonomous data movement transactions should not negate 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
Autonomous data movement transactions that span systems should include data in motion encryption for sensitive and high value data		4	0	0		X
As part of autonomous data movement transactions, the source data or storage media should be appropriately sanitized before it is released for re-use		4	0	0		X

Table B.13 (continued)

Controls	Priorities (5 is highest)				Data Sensitivity	
	S	C	I	A	L	H
Sanitization performed in conjunction with autonomous data movement should also include verification and some form of proof of sanitization		4	0	0		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
Autonomous data movement transactions should not cause data to move to systems with inadequate certifications and accreditations	X	1	1	1		X
Autonomous data movement transactions should not cause data to move to systems with inadequate physical security	X	3	3	3	X	X

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.
- b) Something the authenticating party has, such as a physical hardware token or a key card.
- c) 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 27002:2013, 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:

- *Local Authentication* - The system needing the authentication service is also the authenticator (i.e., entity making the authentication decision). There is no easy way to synchronize the credential database used for verification, so its usability is limited within larger organizations.
- *External Authentication* - The authenticator resides outside of the control and influence of the system needing an authentication decision; further, the authenticator is a trusted, authoritative source.

- *Centralized Authentication* - This form of external authentication is designed to support many systems (often heterogeneous) and it often includes redundancy, use of standard protocols, and provides additional useful information (e.g., role identifiers). There is no attempt to make subsequent authentications transparent (i.e., multiple authentications are often required).
- *Single Sign-on (SSO)* - This form of centralized authentication employs a single set of credentials that are then used transparently to perform subsequent authentications on behalf of the users. In addition, there is typically a close alignment with a centralized authorization system to ensure consistent privileges.

Many enterprises have centralized their identity management (directory services, NIS, NIS+) and authentication services (e.g., RADIUS, PKI, Kerberos, LDAP, etc.), so there is a natural desire to leverage this infrastructure and the investments made in populating the identity data to help address authentication and authorization.

C.2 Authorization and access control

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

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

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

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

- *Discretionary Access Control (DAC)* - policy permits the granting and revocation of access permissions to be left to the discretion of the individual users
- *Mandatory Access Control (MAC)* - policy is centrally controlled by a security policy administrator; users do 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 recommended 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 C.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 than 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.

For storage devices containing legacy magnetic media, a single overwrite pass with a fixed pattern such as zeroes typically prevents recovery of data even if state of the art laboratory techniques are applied to attempt to retrieve the data. One major drawback of relying solely upon the native read and write interface for performing the overwrite procedure is that areas not currently mapped to active Logical Block Addressing (LBA) addresses (such as defect areas and currently unallocated space) are not addressed. Dedicated sanitize commands support addressing these areas more effectively. The use of such commands results in a trade-off because although they should more thoroughly address all areas of the media, using these commands also requires trust and assurance from the vendor that the commands have been implemented as expected.

Users who have become accustomed to relying upon overwrite techniques on magnetic media and who have continued to apply these techniques as media types evolved (such as to flash-based devices) may be exposing their data to increased risk of unintentional disclosure. Although the host interface (e.g. ATA or SCSI) may be the same (or very similar) across devices with varying underlying media types, it is critical that the sanitization techniques are carefully matched to the media.

Destructive techniques for some media types may become more difficult or impossible to apply in the future. Traditional techniques such as degaussing (for magnetic media) become more complicated as magnetic media evolves, because some emerging variations of magnetic recording technologies incorporate media with higher coercivity (magnetic force). As a result, existing degaussers may not have sufficient force to effectively degauss such media.

Applying destructive techniques to non-magnetic storage media such as flash is also becoming more challenging, as the necessary particle size for commonly applied grinding techniques goes down

proportionally to any increases in flash storage density. Flash chips already present challenges with occasional damage to grinders due to the hardness of the component materials, and this problem will get worse as grinders attempt to grind the chips into even smaller pieces.

Cryptographic erase features (see [A.3](#)), a purge 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);
- reducing the wear on the storage device (therefore potentially extending the life of the device);
- reducing the amount of man-hours expended performing sanitization;
- addressing media types that may be impractical to address using legacy degaussing and destruction techniques.

For all devices supporting encryption where cryptographic erase is intended for use to purge the media (including SEDs, mobile devices, and other devices), the level of assurance depends (in large part) on the following:

- a) The assurance that all target data is either encrypted or (for any target data not encrypted) able to be effectively sanitized using media-specific techniques.
- b) The level of entropy of the MEK.
- c) If the key sanitized during cryptographic erase is a key that wraps the MEK (and not the MEK itself), the strength of the wrapping mechanism(s) and entropy of the wrapping key(s) to be sanitized.
- d) The strength of the encryption algorithm used to encrypt the data, including mode of operation and assurance of correct implementation.
- e) The level of difficulty in retrieving the MEK after sanitization, plus any effort to unwrap the key (if it was stored wrapped with another value).

NOTE In some cases, the Media Encryption Key may have been stored in the clear because the sole purpose of media encryption may have been to support cryptographic erase.

Mobile devices (and devices other than SEDs) can also support strong encryption capabilities. The decision regarding whether to rely upon cryptographic erase to purge the media on those devices depends, in part, on whether all sensitive data is encrypted on the device. If encryption was enabled after sensitive data was stored on the device, or if it is unknown whether sensitive data was stored on the device prior to encryption, cryptographic erase may not be appropriate as a Purge mechanism (see [A.3](#)).

An important issue in cryptographic erasure is ensuring that all copies of the encryption keys are destroyed. SEDs assure this by generating keys internally and never exposing them. Erasure is performed via a key change request operation. It is possible that SEDs may contain multiple partitions, each with a unique key, thus requiring multiple key change requests.

C.5 Logging

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

Not all event log entries are created equal, as some may only be useful for debugging purposes, provide system health status, warn of minor configuration problems, etc. From an audit logging perspective the management events (i.e., what a human did) are always of interest, the data access events are usually of

limited interest (except in situations where critical files and directories need to be tightly monitored), and control events are typically of the least interest (they can provide useful information during root-cause analysis after an incident).

In addition, audit logging often requires the event entries of interest to be handled differently and separately from most other event log entries generated by a device. This special handling can be accomplished by having the devices send the audit log entries to special log infrastructure or they can be culled out of the general log stream, using a log filtering mechanism (a more challenging approach because it requires all the event entries of interest 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 needs to 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 server and need to 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 servers, creating N_Port_IDs for every virtual server 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 servers for a single organization or a single tenant of a service provider).

C.7 Fibre Channel Security

C.7.1 Overview

Fibre channel fabrics may be deployed across multiple, distantly separated sites which make it critical that security services be available to assure consistent configurations and proper access controls.

ANSI INCITS 496-2012, *Information Technology - Fibre Channel - Security Protocols - 2* (FC-SP-2) defines protocols to authenticate Fibre Channel entities, set up session keys, negotiate parameters to ensure frame-by-frame integrity and confidentiality, and define and distribute policies across a Fibre Channel fabric. To claim FC-SP-2 compliance, an implementation is required to support the AUTH-A Compliance Element as described in Annex A (specifically, A.2.1) of the FC-SP-2 standard.

The security architecture defined by FC-SP-2 encompasses the following components:

- a) *Authentication infrastructure* – The fabric security architecture is defined for several authentication infrastructures: secret-based, certificate-based, password-based, and pre-shared key based authentication.
- b) *Authentication* – Authentication protocols are defined to allow entities to assure the identity of communicating entities. Two entities may negotiate whether authentication is required and which authentication protocol may be used. Authentication is defined for switch-to-switch, node-to-switch, and node-to-node, using one of the following protocols:
 - Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP), which is mandatory (see [C.7.2](#));
 - Fibre Channel Certificate Authentication Protocol (FCAP);
 - Fibre Channel Password Authentication Protocol (FCPAP);
 - Fibre Channel Extensible Authentication Protocol (FCEAP);
 - The Security Association Management Protocol (IKEv2-AUTH).
- c) *Security associations* – A subset of the IKEv2 protocol suitable for Fibre Channel (i.e., the Security Association Management protocol) is defined in order to establish Security Associations between entities. Two mechanisms are available to protect specific classes of traffic: the ESP_Header is used to protect Fibre Channel frames, and CT_Authentication is used to protect Common Transport Information Units.
- d) *Cryptographic integrity and confidentiality* – Frame by frame cryptographic integrity and confidentiality, replay protection, and traffic origin authentication is achieved by using the ESP_Header (see [C.7.3](#)). CT_Authentication (see [C.7.4](#)) may be leveraged to provide cryptographic integrity and confidentiality, replay protection, and traffic origin authentication to Common Transport Information Units. ESP_Header processing and CT_Authentication processing are independent.
- e) *Authorization (access control)* – Fabric policies provide basic authorization controls in the form of Access Control Lists (ACLs), and the following two basic types of policies are defined:
 - policies that contain fabric-wide data, distributed to every switch of the fabric; and
 - policies that contain per switch data, sent to an individual switch.

Fabric policies may be used to control which switches are allowed to comprise a fabric and which nodes are allowed to connect to a fabric. Policies may be further used to specify topology restrictions within the fabric environment (e.g., which switches may connect to which other switches or which nodes may connect to which switches).

Fabric policies also provide the mechanism for controlling management access to the fabric and the ability to control authentication choices and to specify security attributes for fabric entities (e.g., nodes and switches). Management access to the fabric may be controlled for Common Transport or IP access.

Figure C.1, which is from clause 4.5 of the FC-SP-2 standard, shows the relationship between the authentication protocols and security associations. The defined authentication protocols are able to perform mutual authentication with optional shared key establishment. The shared key computed at the end of an authentication transaction may be used to establish security associations.

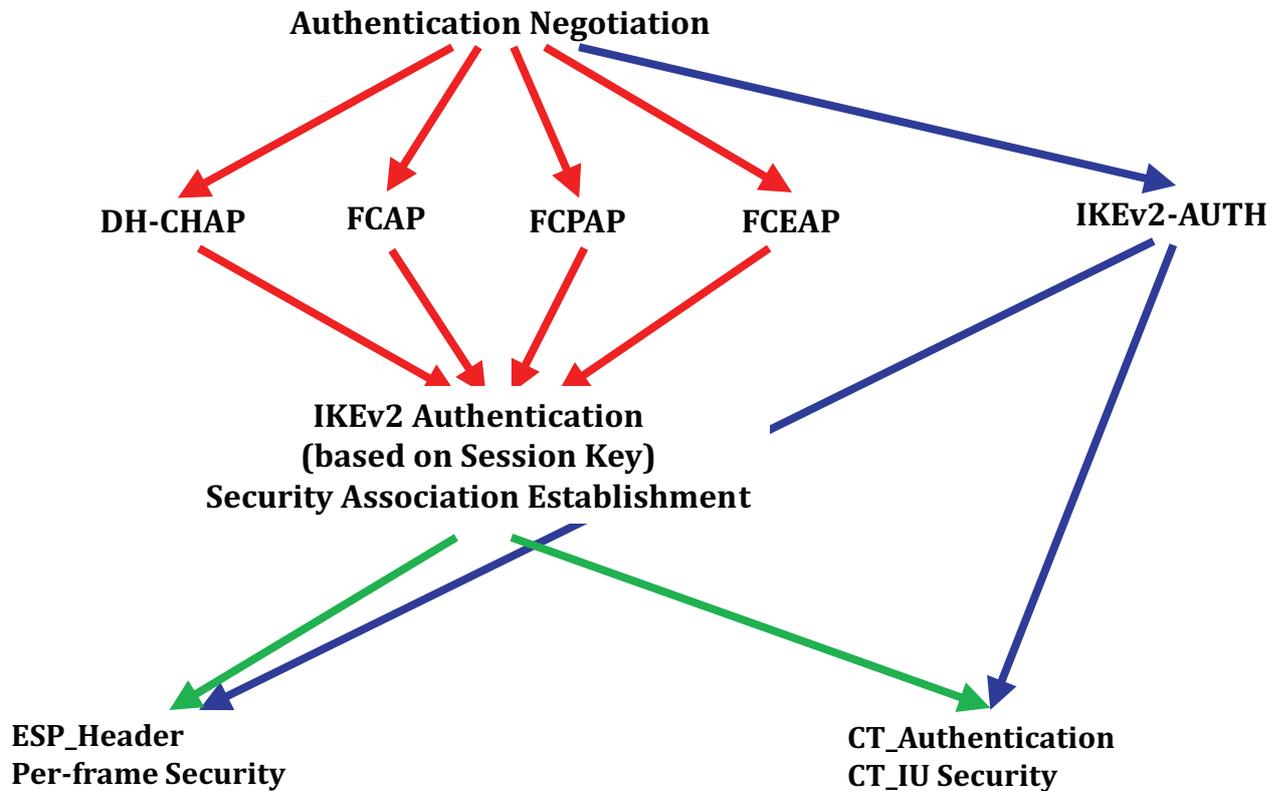


Figure C.1 — Relationship between FC-SP-2 Authentication Protocols and Security Associations

C.7.2 DH-CHAP authentication

DH-CHAP is a secret-based authentication and key management protocol that uses the CHAP algorithm augmented with an optional Diffie-Hellmann algorithm. DH-CHAP provides unidirectional or bidirectional authentication between an *Authentication Initiator* and an *Authentication Responder*. When the Diffie-Hellmann part of the protocol is not used, DH-CHAP reduces its operations to those of the CHAP protocol, and it is referred to as DH-CHAP with a NULL DH algorithm. All FC-SP-2-compliant implementations are required to support DH-CHAP with a NULL DH algorithm. To claim FC-SP-2 compliance, an implementation is required to support the AUTH-A Compliance Element (described in Annex A of FC-SP-2), which includes the 2048-bit DH algorithm.

In addition to identifying the authentication algorithm, FC-SP-2 specifies that authentication is defined for Switch-to-Switch, Device-to-Switch, and Device-to-Device entities, and that the protocols are able to support mutual authentication. Thus, conformant or compliant products are required to also implement each of the following when applicable:

- **Switch-to-Switch**—Products that include authentication between these types of entities must be able to authenticate a switch as well as be authenticated by a switch.
- **Device-to-Switch**—Products that include authentication between these types of entities must be able to authenticate a switch as well as be authenticated by a switch, from a device perspective, or be able to authenticate a device as well as be authenticated by a device, from a switch perspective.
- **Device-to-Device**—Products that include authentication between these types of entities must be able to authenticate a device as well as be authenticated by a device. In addition, each device must complete the appropriate Device-to-Switch authentication prior to performing this authentication.

Conformant products must also implement re-authentication such that the product can be re-authenticated by the other entity at any time.

C.7.3 ESP_Header

ANSI INCITS 470–2011, *Information Technology — Fibre Channel — Framing and Signaling-3* (FC-FS-3) defines optional headers that can be used within Fibre Channel frames. Of these optional headers, the ESP_Header and ESP_Trailer play an important security role because they are the mechanism used to support encryption of frame payloads.

Security Associations for the ESP_Header and CT_Authentication protocols are negotiated by the Fibre Channel Security Association Management Protocol (defined in FC-SP-2). The protocol is a modified subset of the IKEv2 protocol that performs the same core operations, and it uses the Fibre Channel AUTH protocol to transport IKEv2³³⁾ messages. IETF RFC 4595, *Use of IKEv2 in the Fibre Channel Security Association Management Protocol* provides additional information on Fibre Channel use of IKEv2.

NOTE As a result of the separation of Fibre Channel data traffic from control traffic, only one protocol (either ESP_Header or CT_Authentication) is applicable to any FC Security Association.

The Encapsulating Security Payload (ESP), defined in RFC 4303, is a generic mechanism to provide confidentiality, data origin authentication, and anti-replay protection for IP packets. FC-SP-2 defines how to use ESP in Fibre Channel, including any negotiation procedure, additional encryption or authentication algorithms and processing requirements.

FC-FS-3 states that “End-to-end ESP_Header processing shall be applied to FC frames in transport mode (see RFC 4303³⁴⁾), and Link-by-link ESP_Header processing shall be applied to FC frames in tunnel mode (see RFC 4303). The Authentication option shall be used, Confidentiality may be negotiated by the two communicating FC_Ports (see FC-SP-2).”

NOTE An intended application of Link-by-link ESP_Header processing is to secure a link in a Fabric or between Fabrics without requiring use of ESP by every Nx_Port.

C.7.4 CT_Authentication

Fibre Channel defines two security protocols that provide security services for different portions of Fibre Channel traffic: the ESP_Header (see C.7.3) and CT_Authentication defined in ANSI INCITS 463–2010, *Information Technology — Fibre Channel — Generic Services — 6* (FC-GS-6). The CT_Authentication protocol provides origin authentication, integrity, anti-replay protection, and optionally, confidentiality for Common Transport Information Units, which are used to convey control information. As a result of the separation of Fibre Channel data traffic from control traffic, only one protocol (either ESP_Header or CT_Authentication) is applicable to any FC Security Association. Security associations for the ESP_Header and CT_Authentication protocols between two Fibre Channel entities (servers, disks, or switches) are negotiated by the Fibre Channel Security Association Management Protocol (defined FC-FS-3), which is a generic protocol based on the Internet Key Exchange Protocol Version 2 (IKEv2).³⁵⁾

Unlike ESP_Header, which operates at the FC frame level, CT_Authentication operates at the Common Transport (CT) level and provides access to a service (e.g., directory service) with a set of service parameters that facilitates the usage of Fibre Channel constructs. The CT_Authentication protocol is carried in the optional extended CT_IU preamble includes an Authentication Security Association Identifier, a transaction ID, the N_port name of the requesting node, a time stamp used to prevent replay attacks, and an authentication hash block.

33) IETF RFC 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*

34) IETF RFC 4303, IP Encapsulating Security Payload (ESP) describes an updated version of ESP, which is used to provide confidentiality, data origin authentication, connectionless integrity, an anti-replay service (a form of partial sequence integrity), and limited traffic flow confidentiality.

35) IKEv2 is described in IETF RFC 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*. IETF RFC 4595, *Use of IKEv2 in the Fibre Channel Security Association Management Protocol* provides additional useful details that are relevant to CT_Authentication.

C.7.5 FC-SP Zoning

In order to preserve backward compatibility with existing zoning definitions and implementations, FC-SP-2 describes a variant of the Enhanced Zoning model defined in ANSI INCITS 461–2010, *Information Technology — Fibre Channel - Switch Fabric — 5* (FC-SW-5) and ANSI INCITS 463–2010, *Information Technology — Fibre Channel — Generic Services — 6* (FC-GS-6) that follows the general concepts of the policy model, but keeps zoning management and enforcement completely independent from other policies management and enforcement. This variant of zoning is denoted as FC-SP Zoning.

Fabric policies and zoning policies allow an asymmetric distribution of policy information in the fabric with the definition of three types of switches:

- a) Server Switches: Switches that retain all policy objects and all node to node (zoning) information
- b) Autonomous Switches: Switches that retain their own per switch policy objects, all fabric-wide policy objects, and all node to node (zoning) information
- c) Client Switches: Switches that retain their per switch policy objects, all fabric-wide policy objects and the subset of the node to node (zoning) information relevant for their operations, which is pulled from a server switch when needed. In addition, they maintain the zone set database hash and the active zone set hash for the fabric

C.8 OASIS Key Management Interoperability Protocol (KMIP)

Applying protection of information within a storage security environment generally involves the use of cryptographic keys necessitating key lifecycle management (generation, refresh, distribution, tracking of use, life-cycle policies including states, archive, and destruction), key sharing, and long-term availability of cryptographic keys.

The OASIS Key Management Interoperability Protocol (KMIP) is an open, vendor neutral specification that defines a network protocol for communication between an application, device or system that uses cryptographic keys and a system that manages cryptographic keys.

KMIP supports cryptographic key generation by either the client or the server. The key management life-cycle defines the states and transitions permitted between states within KMIP for all managed cryptographic objects.

The key management protocol defined by KMIP encompasses the following components:

- a) Transport Security – Mutually authenticated TLS sessions between the client and the server assure the required security and integrity of the communication. The recommended version of TLS (namely version 1.2) is supported.
- b) Request and Response Messages – KMIP is fundamentally a stateless request/response protocol with the key management client sending request messages to the key management server which replies with response messages.
- c) Authentication – Layered authentication options are provided as a combination of zero or more credentials (in addition to the TLS session based mutual authentication):
 - Attestation Credential
 - Device Credential (with optional Password)
 - Username and Password
- d) Managed Objects – KMIP provides support for the following types of managed cryptographic objects:
 - Certificate
 - Private Key

- Public Key
 - Secret Data
 - Split Key
 - Symmetric Key
- e) Operations – Individual operations that support the full key management life-cycle including:
- Activate
 - Add Attribute
 - Archive
 - Cancel
 - Certify
 - Check
 - Create
 - Create Key Pair
 - Delete Attribute
 - Destroy
 - Discover Versions
 - Get
 - Get Attribute List
 - Get Attributes
 - Get Usage Allocation
 - Locate
 - Modify Attribute
 - Notify
 - Obtain Lease
 - Poll
 - Put
 - Query
 - Re-certify
 - Recover
 - Register
 - Re-key
 - Re-key Key Pair
 - Revoke

- f) Batched Operations – Each Request/Response Message may contain multiple operations within a Batch Item which can be chained to perform a transaction-safe series of key management operations on a managed object.
- g) Attributes – Each Managed Object has a set of associated Attributes that provide details of the managed object to support life-cycle management and include the following:
 - Activation Date
 - Application Specific Information
 - Archive Date
 - Certificate Identifier
 - Certificate Issuer
 - Certificate Length
 - Certificate Subject
 - Certificate Type
 - Compromise Date
 - Compromise Occurrence Date
 - Contact Information
 - Cryptographic Algorithm
 - Cryptographic Domain Parameters
 - Cryptographic Length
 - Cryptographic Parameters
 - Cryptographic Usage Mask
 - Custom Attribute
 - Deactivation Date
 - Destroy Date
 - Digest
 - Digital Signature Algorithm
 - Fresh
 - Initial Date
 - Last Change Date
 - Lease Time
 - Link
 - Name
 - Object Group
 - Object Type

- Operation Policy Name
- Process Start Date
- Protect Stop Date
- Revocation Reason
- State
- Unique Identifier
- Usage Limits
- X.509 Certificate Identifier
- X.509 Certificate Issuer
- X.509 Certificate Subject

In addition to the main specification, KMIP Profiles define functionality sets that address specific scenarios (such as vaulting keys created within a storage environment). KMIP Profiles also define the authentication that need to be used to ensure message confidentiality and integrity. Conformance to these KMIP Profiles is the primary mechanism for ensuring interoperability between KMIP clients and servers.

Bibliography

- [1] ISO Guide 73:2009, *Risk management — Vocabulary*
- [2] ISO 7498-2:1989, *Information processing systems — Open Systems Interconnection — Basic Reference Model — Part 2: Security Architecture*
- [3] ISO 16609:2004, *Banking — Requirements for message authentication using symmetric techniques*
- [4] ISO/PAS 22399:2007, *Societal security - Guideline for incident preparedness and operational continuity management*
- [5] ISO/IEC 10116:2006, *Information technology — Security techniques — Modes of operation for an n-bit block cipher*
- [6] ISO/TR 10255:2009, *Document management applications — Optical disk storage technology, management and standards*
- [7] ISO/TR 18492:2005, *Long-term preservation of electronic document-based information*
- [8] ISO 16175-1:2010, *Information and documentation — Principles and functional requirements for records in electronic office environments — Part 1: Overview and statement of principles*
- [9] ISO 16175-2:2011, *Information and documentation — Principles and functional requirements for records in electronic office environments — Part 2: Guidelines and functional requirements for digital records management systems*
- [10] ISO 16175-3:2010, *Information and documentation — Principles and functional requirements for records in electronic office environments — Part 3: Guidelines and functional requirements for records in business systems*
- [11] ISO/IEC 11770 (all parts), *Information technology — Security techniques — Key management*
- [12] ISO/IEC 17826:2012, *Information technology — Cloud Data Management Interface (CDMI)*
- [13] ISO/IEC 19790:2006, *Information technology — Security techniques — Security requirements for cryptographic modules*
- [14] ISO/IEC 24759:2008, *Information technology — Security techniques — Test requirements for cryptographic modules*
- [15] ISO/IEC 24775, *Information technology — Storage management*
- [16] ISO/IEC 27003:2010, *Information technology — Security techniques — Information security management system implementation guidance*
- [17] ISO/IEC 27031:2011, *Information technology — Security techniques — Guidelines for information and communication technology readiness for business continuity*
- [18] ISO/IEC 27033-1:2009, *Information technology — Security techniques — Network security — Part 1: Overview and concepts*
- [19] ISO/IEC 27033-2, *Information technology — Security techniques — Network security — Part 2: Guidelines for the design and implementation of network security*
- [20] ISO/IEC 27033-3:2010, *Information technology — Security techniques — Network security — Part 3: Reference networking scenarios — Threats, design techniques and control issues*
- [21] ISO/IEC 27037:2012, *Information technology — Security techniques — Guidelines for identification, collection, acquisition and preservation of digital evidence*

- [22] ISO/IEC/IEEE 24765:2010, *Systems and software engineering — Vocabulary*
- [23] IEEE 1619-2007, *IEEE Standard for Wide-Block Encryption for Shared Storage Media*
- [24] IEEE 1619.1-2007, *IEEE Standard for Authenticated Encryption with Length Expansion for Storage Devices*
- [25] IEEE 1619.2-2010, *IEEE Standard for Cryptographic Protection of Data on Block-Oriented Storage Devices*
- [26] IETF RFC 1813 *NFS Version 3 Protocol Specification*
- [27] IETF RFC 3195 *Reliable Delivery for syslog*
- [28] IETF RFC 3530 *Network File System (NFS) version 4 Protocol*
- [29] IETF RFC 3720 *Internet Small Computer Systems Interface (iSCSI)*
- [30] IETF RFC 3723 *Securing Block Storage Protocols over IP*
- [31] IETF RFC 3821 *Fibre Channel Over TCP/IP (FCIP)*
- [32] IETF RFC 4303 *IP Encapsulating Security Payload (ESP)*
- [33] IETF RFC 4595 *Use of IKEv2 in the Fibre Channel Security Association Management Protocol*
- [34] IETF RFC 5246 *The Transport Layer Security (TLS) Protocol Version 1.2*
- [35] IETF RFC 5424 *The Syslog Protocol*
- [36] IETF RFC 5425 *TLS Transport Mapping for Syslog*
- [37] IETF RFC 5426 *Transmission of Syslog Messages over UDP*
- [38] IETF RFC 5427 *Textual Conventions for Syslog Management*
- [39] IETF RFC 5661 *Network File System (NFS) Version 4 Minor Version 1 Protocol*
- [40] IETF RFC 5663 *Parallel NFS (pNFS) Block/Volume Layout*
- [41] IETF RFC 5848 *Signed Syslog Messages*
- [42] IETF RFC 6012 *Datagram Transport Layer Security (DTLS) Transport Mapping for Syslog*
- [43] IETF RFC 6071 *IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap*
- [44] IETF RFC 6587 *Transmission of Syslog Messages over TCP*
- [45] IETF RFC 7146, *Securing Block Storage Protocols over IP: RFC 3723 Requirements Update for IPsec v3*
- [46] ANSI INCITS 400–2004, *Information technology — SCSI Object-based Storage Device Commands (OSD)*
- [47] ANSI INCITS 458–2011, *Information technology — SCSI Object-Based Storage Device Commands – 2 (OSD-2)*
- [48] ANSI INCITS 461–2010, *Fibre Channel — Switch Fabric — 5 (FC-SW-5)*
- [49] ANSI INCITS 462–2010, *Information Technology — Fibre Channel - Backbone — 5 (FC-BB-5)*
- [50] ANSI INCITS 463–2010, *Fibre Channel — Generic Services — 6 (FC-GS-6)*
- [51] ANSI INCITS 470–2011, *Fibre Channel — Framing and Signaling-3 (FC-FS-3)*
- [52] ANSI INCITS 482–2012, *Information Technology — ATA/ATAPI Command Set — 2 (ACS-2)*

- [53] ANSI INCITS 496–2012, *Information Technology — Fibre Channel — Security Protocols — 2 (FC-SP-2)*
- [54] ANSI INCITS 512–2013, *Information Technology — SCSI Block Commands — 3 (SBC-3)*
- [55] NIST FIPS 140–2, *Security Requirements for Cryptographic Modules*
- [56] NIST FIPS 197, *Advanced Encryption Standard*
- [57] NIST Special Publication 800-38A, *Recommendation for Block Cipher Modes of Operation: Three Variants of Ciphertext Stealing for CBC Mode*
- [58] NIST Special Publication 800-38C, *Recommendation for Block Cipher Modes of Operation: the CCM Mode for Authentication and Confidentiality*
- [59] NIST Special Publication 800-38D, *Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC*
- [60] NIST Special Publication 800-38E, *Recommendation for Block Cipher Modes of Operation: The XTS-AES Mode for Confidentiality on Storage Devices*
- [61] NIST Special Publication 800-57 Part 1, *Recommendation for Key Management: Part 1: General (Revision 3)*
- [62] NIST Special Publication 800-57 Part 2, *Recommendation for Key Management: Part 2: Best Practices for Key Management Organization*
- [63] NIST Special Publication 800-67, *Recommendation for the Triple Data Encryption Algorithm (TDEA) Block Cipher*
- [64] NIST Special Publication 800-88 Revision 1 (draft), *Media Sanitization*
- [65] STORAGE NETWORKING INDUSTRY ASSOCIATION (SNIA). *Storage Management Initiative – Specification (SMI-S), Version 1.5, Architecture Book*, http://www.snia.org/tech_activities/standards/curr_standards/smi
- [66] STORAGE NETWORKING INDUSTRY ASSOCIATION (SNIA). *SNIA Technical Position: TLS Specification for Storage Systems v1.0*, http://snia.org/tech_activities/standards/curr_standards/tls
- [67] Trusted Computing Group, *Storage Architecture Core Specification*, Version 2.0, November 2011
- [68] Trusted Computing Group, *Storage Security Subsystem Class: Enterprise*, Version 1.0, January 2011
- [69] Trusted Computing Group, *Storage Security Subsystem Class: Opal*, Version 2.0, February 2012
- [70] OASIS, *Key Management Interoperability Protocol Specification (Version 1.2 or later)*
- [71] OASIS, *Key Management Interoperability Protocol Profiles (Version 1.2 or later)*
- [72] RECOMMENDATION ITU-T X. 1601 (2013), *Security framework for cloud computing*
- [73] ISO/IEC 27002:2013, *Information technology — Security techniques — Code of practice for information security controls*
- [74] ISO/IEC 14776-372:2011, *Information technology — Small Computer System Interface (SCSI) — Part 372: SCSI Enclosure Services–2 (SES-2)*
- [75] ISO/IEC 11179-1:2004, *Information technology — Metadata registries (MDR) — Part 1: Framework*

