International Standard

ISO/IEC 27040

Information technology — Security techniques — Storage security

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

Second edition
2024-01

Reference number
ISO/IEC 27040:2024(en)

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Contents

Page

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.

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 or www.iec.ch/members_experts/refdocs).

ISO and IEC draw attention to the possibility that the implementation of this document may involve the use of (a) patent(s). ISO and IEC take no position concerning the evidence, validity or applicability of any claimed patent rights in respect thereof. As of the date of publication of this document, ISO and IEC had not received notice of (a) patent(s) which may be required to implement this document. However, implementers are cautioned that this may not represent the latest information, which may be obtained from the patent database available at www.iso.org/patents and https://patents.iec.ch. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

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

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see www.iso.org/iso/foreword.html. In the IEC, see www.iec.ch/understanding-standards.

This document was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 27, *Information security, cybersecurity and privacy protection*.

This second edition cancels and replaces the first edition (ISO/IEC 27040:2015), which has been technically revised.

The main changes are as follows:

— the scope has been expanded to cover requirements;

— the clause structure has been more closely aligned with ISO/IEC 27002:2022;

— requirements have been added in Clauses 7, 9, and 10;

— adjustments have been made regarding the storage technologies which are covered;

— a new controls labelling scheme has been added;

— former Annex A, which provided guidance on sanitizing specific types of media, has been removed and text has been added in Clause 10, recommending IEEE 2883 for this purpose;

— former Annex B, which included table to help prioritize the adoption of recommendation, has been replaced with Annex A that summarizes the requirements and guidance contained in this document;

— former Annex C, which provided tutorial oriented material, has been removed and references to appropriate materials have been added in Clause 10.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html and www.iec.ch/national-committees.

Information technology — Security techniques — Storage security

1 Scope

This document provides detailed technical requirements and guidance on how organizations can achieve 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 of data both while stored in information and communications technology (ICT) systems and while in transit across the communication links associated with storage. Storage security includes the security of devices and media, management activities related to the devices and media, applications and services, and controlling or monitoring user activities during the lifetime of devices and media, and after end of use or end of life.

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 products and services, and other non-technical managers or users, in addition to managers and administrators who have specific responsibilities for information or storage security, storage operation, or who are responsible for an organization's overall security programme 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 document provides an overview of storage security concepts and related definitions. It includes requirements and guidance on the threats, 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 are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

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

3 Terms and definitions

3.1 General

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

ISO and IEC maintain terminology databases for use in standardization at the following addresses:

— ISO Online browsing platform: available at https://www.iso.org/obp

— IEC Electropedia: available at https://www.electropedia.org/

3.2 Terms relating to storage technology

3.2.1
block
unit in which data is *stored* (3.2.17) and retrieved on *storage devices* (3.2.14) and *storage media* (3.2.16)

3.2.2
compression
reduction in the number of bits used to represent an item of data

Note 1 to entry: For *storage* (3.2.12), 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 used.

3.2.3
data at rest
data recorded on stable, *non-volatile storage* (3.2.11)

3.2.4
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 outside of an interface, chip, or device).

3.2.5
deduplication
method of reducing *storage* (3.2.12) 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 data reduction.

3.2.6
device
mechanical, electrical, or electronic contrivance with a specific purpose

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

3.2.7
Fibre Channel
serial input/output interconnect capable of supporting multiple protocols, including access to open system *storage* (3.2.12), access to mainframe storage, 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 128 gigabits per second.

3.2.8
Fibre Channel Protocol
Serial Small Computer System Interface (SCSI) transport protocol used on *Fibre Channel* (3.2.7) interconnects

3.2.9
over provision
technique used by *storage devices* (3.2.14) in which a subset of the available *storage medium* (3.2.16) is exposed through the interface

Note 1 to entry: The storage medium is used internally and independently by the storage device to improve performance, endurance, or reliability.

3.2.10
network attached storage
storage device (3.2.14) or system that connects to a network and provides file access services to computer systems

3.2.11
non-volatile storage
storage (3.2.12) that retains its contents after power is removed

3.2.12
storage
device (3.2.6), function, or service supporting data entry or retrieval

3.2.13
storage area network
network whose primary purpose is the transfer of data between computer systems and *storage devices* (3.2.14) and among storage devices

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

3.2.14
storage device
component or aggregation of components made up of one or more *devices* (3.2.6) containing *storage media* (3.2.16), designed and built primarily for the purpose for accessing *non-volatile storage* (3.2.11)

3.2.15
storage ecosystem
system of interdependent components that work together to enable *storage* (3.2.12) services and capabilities

Note 1 to entry: The components often include *storage devices* (3.2.14), storage networks, storage management, and other information and communications technology (ICT) infrastructure.

3.2.16
storage medium
material on which digital data are, or can be, recorded or retrieved

3.2.17
store
record data on *volatile storage* (3.2.20) or *non-volatile storage* (3.2.11)

3.2.18
target data
information subject to a given process, typically including most or all information on *storage* (3.2.12)

3.2.19
virtualized storage
logical storage
abstraction of physical *storage devices* (3.2.14) or *storage media* (3.2.16) that masks the characteristics and boundaries of the physical *storage* (3.2.12)

Note 1 to entry: Virtualized storage can employ multiple levels of virtualization prior to presenting the virtualized storage to a system or an application.

3.2.20
volatile storage
storage (3.2.12) that fails to retain its contents after power is removed

3.3 Terms relating to sanitization

3.3.1
clear
sanitize (3.3.12) using logical techniques on user data on all addressable storage locations for protection against simple non-invasive data recovery techniques using the same host interface available to the user

3.3.2
degauss
render magnetically stored data unreadable by applying a strong magnetic field to the *storage medium* (3.2.16) with an organizationally approved field strength

3.3.3
destruct
sanitize (3.3.12) using physical techniques that make recovery of *target data* (3.2.18) infeasible using state of the art laboratory techniques and results in the subsequent inability to use the *storage medium* (3.2.16) for *storage* (3.2.12)

Note 1 to entry: *Disintegrate* (3.3.5), *incinerate* (3.3.6), *melt* (3.3.7), *pulverize* (3.3.9), and *shred* (3.3.13) are destruct forms of *media sanitization* (3.3.16).

Note 2 to entry: If the storage medium cannot be removed, then the *storage device* (3.2.14) can be subjected to the destruct technique; a storage device can contain multiple storage media.

3.3.4
destruction
result of *destruct* (3.3.3) actions taken to ensure that the *storage medium* (3.2.16) cannot be reused as originally intended and that user data is virtually impossible or prohibitively expensive to recover

3.3.5
disintegrate
destruct (3.3.3) by separating *storage medium* (3.2.16) into its component parts

3.3.6
incinerate
destruct (3.3.3) by burning *storage medium* (3.2.16) completely

3.3.7
melt
destruct (3.3.3) by changing *storage medium* (3.2.16) from a solid to a liquid state generally by the application of heat

3.3.8
cryptographic erase
method of sanitization in which the encryption key for the encrypted *target data* (3.2.18) is *sanitized* (3.3.12), making recovery of the decrypted target data infeasible

3.3.9
pulverize
destruct (3.3.3) by grinding *storage medium* (3.2.16) to a powder or appropriately small particles

3.3.10
purge
sanitize (3.3.12) using physical or logical techniques that make recovery of *target data* (3.2.18) infeasible using state of the art laboratory techniques, but which preserves the *storage media* (3.2.16) and *storage device* (3.2.14) in a potentially reusable state

3.3.11
sanitization
process or method to *sanitize* (3.3.12)

3.3.12
sanitize
render access to *target data* (3.2.18) on *storage* (3.2.12) infeasible for a given level of effort

3.3.13
shred
destruct (3.3.3) by cutting or tearing *storage medium* (3.2.16) into small particles

3.3.14
storage sanitization
logical storage sanitization (3.3.15) or *media sanitization* (3.3.16)

3.3.15
logical storage sanitization
virtual storage sanitization
sanitization (3.3.11) of *virtualized storage* (3.2.19)

Note 1 to entry: *Clear* (3.3.1) and *purge* (3.3.10) are actions that can be taken to *sanitize* (3.3.12) virtualized storage.

Note 2 to entry: Logical storage sanitization is a subset of *storage sanitization* (3.3.14).

3.3.16
media sanitization
sanitization (3.3.11) of *storage media* (3.2.16)

Note 1 to entry: *Clear* (3.3.1), *purge* (3.3.10), and *destruct* (3.3.3) are actions that can be taken to *sanitize* (3.3.12) *storage media* (3.2.16).

Note 2 to entry: Media sanitization is a subset of *storage sanitization* (3.3.14).

3.4 Terms relating to availability

3.4.1
resilience
ability to anticipate and adapt to, resist or quickly recover from a potentially disruptive event, whether natural or man-made

[SOURCE: ISO 15392:2019, 3.21]

3.5 Terms relating to security and cryptography

3.5.1
cryptoperiod
defined period of time during which a specific cryptographic key is authorized for use or during which the cryptographic keys in a given system may remain in effect

[SOURCE: ISO 16609:2022, 3.6]

3.5.2
data breach
compromise of security that leads to the accidental or unlawful *destruction* (3.3.4), loss, alteration, unauthorized disclosure of, or access to protected data transmitted, *stored* (3.2.17), or otherwise processed

3.5.3
data integrity
property that data has not been altered or destroyed in an unauthorized manner

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

3.5.4
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]

3.5.5
malware
malicious software designed specifically to damage or disrupt a system, attacking confidentiality, integrity, and/or availability

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

[SOURCE: ISO/IEC 27033-1:2015, 3.22]

3.5.6
point of encryption
location within the information and communications technology infrastructure where data are encrypted on its way to *storage* (3.2.12) and, conversely, where data are decrypted when accessed from storage

Note 1 to entry: The point of encryption is only applicable for *data at rest* (3.2.3).

3.5.7
security strength
number associated with the amount of work that is required to break a cryptographic algorithm or system

3.5.8
storage security
application of physical, technical, and administrative controls to protect storage systems and infrastructure as well as the data *stored* (3.2.17) 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 can be preventive, detective, corrective, deterrent, recovery, or compensatory in nature.

3.5.9
strong authentication
authentication by means of cryptographically derived multi-factor credentials

[SOURCE: ISO 22600-1:2014, 3.23]

3.6 Terms relating to archives and repositories

3.6.1
archive
<organization> organization or part of an organization responsible for selection, acquisition, *preservation* (3.6.5), and availability of one or more *archives* (3.6.2)

[SOURCE: ISO 5127:2017, 3.2.3.01, modified — Notes to entry 1 to 4 have been omitted; term changed to singular from plural "archives"; domain <organization> has been added.]

3.6.2
archive
<holdings> materials, items, *records* (3.6.6) or documents created or received by a person, family or organization, public or private, in the conduct of their affairs and preserved because of the enduring value contained in them or as evidence of the functions and responsibilities of their creator, especially those materials maintained using the principles of provenance, original order and collective control

[SOURCE: ISO 5127:2017, 3.6.1.03, modified — "items, records or documents" have been included at the start of the definition; Note 1 to entry has been omitted; term changed to singular from plural "archives"; domain <holdings> has been added.]

3.6.3
disposition
range of records processes associated with implementing *records* (3.6.6) *retention* (3.6.9), *records destruction* (3.6.7) or transfer decisions which are documented in disposition authorities or other instruments

[SOURCE: ISO 30300:2020, 3.4.8, modified — "destruction" has been changed to "records destruction".]

3.6.4
evidence
information that can be used either by itself or in conjunction with other information, to establish proof about an event or action

[SOURCE: ISO 30300:2020, 3.2.6, modified — Note 1 to entry has been omitted; "could" has been changed to "can".]

3.6.5
preservation
measures taken to maintain the *usability* (3.6.10), authenticity, reliability and integrity of *records* (3.6.6) over time

Note 1 to entry: Measures include principles, policies, rules, strategies, processes and operations.

[SOURCE: ISO 30300:2020, 3.4.11]

3.6.6
record
information created or received and maintained as *evidence* (3.6.4) and as an asset by an organization, in pursuit of legal obligations or in the course of conducting business

Note 1 to entry: Records are normally used in plural.

Note 2 to entry: In a management system standard (MSS) implementation, the records created to conduct and direct the management system and to document its implementation are called documented information.

[SOURCE: ISO 30300:2020, 3.2.10]

3.6.7
records destruction
eliminating or deleting a *record* (3.6.6), beyond any possible reconstruction

[SOURCE: ISO 30300:2020, 3.4.7, modified — changed the term "destruction" to "records destruction".]

3.6.8
records requirement
requirement for *evidence* (3.6.4) of a business function, activity or transaction and for records processes including how, and how long, *records* (3.6.6) need to be kept

[SOURCE: ISO 30300:2020, 3.3.2]

3.6.9
retention
keeping a *record* (3.6.6) according to *records requirements* (3.6.8)

[SOURCE: ISO 30300:2020, 3.4.14]

3.6.10
usability
property of being able to be located, retrieved, presented and understood

Note 1 to entry: Usability may also refer to the extent to which a system, product, or service can be used by specified users to achieve specified goals with effectiveness, efficiency, and satisfaction in a specified context of use.

[SOURCE: ISO 30300:2020, 3.2.12]

3.7 Miscellaneous terms

3.7.1
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.

3.7.2
metadata
data that defines and describes other data

[SOURCE: ISO/IEC 11179-1:2023, 3.2.26]

3.7.3
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: ISO/IEC 22123-1:2023, 3.4.3]

3.7.4
out-of-band
communication or transmission that occurs outside of a previously established communication method or channel

3.7.5
secure multi-tenancy
type of *multi-tenancy* (3.7.3) that employs security controls to explicitly guard against *data breaches* (3.5.2) 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 in a dedicated, single-tenant environment.

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

4 Symbols and abbreviated terms

ACL	access control list
AES	Advanced Encryption Standard
API	application programming interface
BCM	business continuity management
BMC	baseboard management controller
CCM	counter with cipher block chaining message authentication code
CDMI	Cloud Data Management Interface
CHAP	Challenge Handshake Authentication Protocol
CLI	command line interface
CNA	converged network adaptor
DAS	direct attached storage

DDoS	distributed denial of service
DH-CHAP	Diffie Hellman – Challenge Handshake Authentication Protocol
DNS	domain name system
DoS	denial of service
DRAM	dynamic random access memory
ESP	encapsulating security payload
FC	fibre channel
FC-SP	fibre channel – security protocol
FCIP	fibre channel over TCP/IP
FCP	fibre channel protocol
GCM	Galois/Counter Mode
HBA	host bus adapter
HDD	hard disk drive
HMAC	hash-based message authentication code
HTTPS	hypertext transfer protocol secure
IB	InfiniBand™
ICT	information and communications technology
ID	identifier
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IKE	Internet Key Exchange
IP	Internet Protocol
IPMI	Intelligent Platform Management Interface
IPsec	Internet Protocol Security
IRBC	ICT readiness for business continuity
iSCSI	Internet Small Computer Systems Interface
ISMS	information security management system
iSNS	Internet Storage Name Service
KMIP	key management interoperability protocol
LAN	local area network
LUN	logical unit number

MTBF	mean time between failure
MTTF	mean time to failure
MTTR	mean time to repair
NAS	network attached storage
NFS	network file system
NTP	network time protocol
NVM	non-volatile memory
NVMe®	NVM Express®
NVMe-oF™	NVM Express® over Fabric
NVMe®/FC	NVMe® over Fibre Channel
NVMe®/RDMA	NVMe® over RDMA
NVMe®/TCP	NVMe® over TCP
OASIS	Organization for the Advancement of Structured Information Standards
PCIe®	PCI Express®
PII	personally identifiable information
RADIUS	remote authentication dial in user service
RAID	redundant array of independent disks
RDMA	remote direct memory access
REST	representational state transfer
RMCP	Remote Management Control Protocol
SAN	storage area network
SCSI	Small Computer System Interface
SHA	Secure Hash Algorithm
SLP	Service Locator Protocol
SMB	Server Message Block
SNMP	Simple Network Management Protocol
SSD	solid state drive
SSH	Secure Shell
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol

VLAN	virtual local area network
VM	virtual machine
VPN	virtual private network
WORM	write once read many
WWN	worldwide name
XEX	Xor-Encrypt-Xor
XTS	XEX-based Tweaked-codebook mode with ciphertext Stealing

5 Structure of this document

5.1 General

The basic structure of this document starts with an introduction of storage security in Clause 6, followed by separate clauses that address organization controls (Clause 7), people controls (Clause 8), physical controls (Clause 9), and technological controls (Clause 10) that are relevant to storage systems and ecosystems.

5.2 Controls

The controls of ISO/IEC 27001 and ISO/IEC 27002 generally apply to storage ecosystems and the organizations that procure, manage, and operate them. This document provides additional guidance and requirements that expand upon these standards.

The requirements specified in this document represent a baseline set of storage security controls. To aid in identifying these baseline controls and key guidance, subheadings are used in the text. These subheadings include a control label with a description. The control labels take the form of xx-yyyy-cnn where:

— xx is “OC” for organizational controls, “SC” for people controls, “PC” for physical controls, or “TC” for technological controls;

— yyyy is associated with the topic;

— c is either “R” for a requirement or “G” for guidance;

— nn is a sequence number when multiple controls exist for a single topic; requirements and guidance are numbered independently of each other.

When multiple subheadings are included in a clause or sub-clause, they are presented in an order list to make easier to refer to them. A summary of the controls is listed in Annex A.

Although this document is written from the perspective of consumers of storage technologies, it is important to note that many controls are dependent on vendors implementing certain security features or capabilities as a prerequisite to the use of these controls. The presence or absence of these features and capabilities can be an important consideration when making product selections.

6 Overview and concepts

6.1 General

Data storage or information storage, often called storage, refers to system components, storage devices, and storage media that retain digital data in durable (i.e. non-volatile) form. While both volatile and non-volatile forms of storage exist, this document is concerned primarily with non-volatile storage. Storage is a core function and fundamental component of ICT systems.

To secure storage infrastructure, a clear understanding of the storage technologies and concepts are necessary. In addition, the types of security controls and insights 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 mitigate risks to storage infrastructure or individual storage systems.

6.2 Storage concepts

In its simplest form, storage can be hard disk drives (HDD), solid state drives (SSD), or tape drives attached to an ICT system 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 environments. With the integration of networking technology, storage systems and ecosystems can be highly sophisticated technologies that provide solutions for managing, connecting, securing (i.e. making safe), sharing, and optimizing the storage of data. These solutions have become more feasible and cost effective as storage technology has 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 to be paid to mitigate the associated risks.

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

— Storage devices and media, consisting of materials which are magnetic, solid state, etc., on which data can be recorded or accessed;

— Block-based storage, which is the most basic and fundamental form of persistent digital data storage that is arranged in a sequence of blocks (each consisting of a sequence of bytes or bits) with offsets to represent locations of individual blocks of data;

— File-based storage, which is data storage that organizes data into files and directories on a file system, while abstracting the underlying hardware for local user access or access over a network, typically as network attached storage (NAS). NAS can enable multiple computer systems to have equal access to the stored contents regardless of their native operating system;

— Storage interfaces, which include specific storage direct attachment interfaces as well as storage networking interfaces;

— Object storage, which is a method of storing and subsequently retrieving sets of data as collections of single, uniquely identifiable indivisible items or objects. It applies to any forms of data that can be wrapped up and managed as an object;

— Cloud computing storage, which can be part of a data storage as a service (DSaaS), as defined in ISO/IEC 22123-1 and described in ISO/IEC 22123-2. Cloud storage can include the most basic storage of files or blocks of raw binary data, relational databases, and advanced big data platforms as well as offering automatic backup, disaster recovery, geographical redundancy, advanced failover, and other more sophisticated features;

— Data protection systems that create one or more copies of production data to enable the recovery of a system or its data following a catastrophic event;

— Persistent memory, which is non-volatile memory with densities greater than or equal to dynamic random access memory (DRAM) that resides on the memory bus and has DRAM-like access to data with nearly the same speed and latency of DRAM and the non-volatility of NAND flash.

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

— sharing of vast storage resources (measured in petabytes, exabytes, and yottabytes) between multiple systems via networks;

— need for more speed (lower latency);

— increased storage capacities;

— expanded storage client access (e.g. mobile devices and internet of things);

— support for distributed ICT architectures (e.g. edge computing);

— support for virtualized environments;

— data protection systems that backup data without the use of a local area network (LAN);

— remote, disaster tolerant, online 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 complying with regulatory and legal requirements;

— support for centralized data repositories for rapid recovery (e.g. backups) and archiving;

— resilience to cyber-attacks.

6.3 Introduction to storage security

Storage security is focused on mitigating risk associated with storage systems and infrastructure, as outlined in 6.2, through the use of safeguards and countermeasures (i.e. controls). The controls are categorized according to ISO/IEC 27002:2022, 4.2.

Storage security can also necessitate the introduction of specialized controls to address technologies such as:

— system security hardening;

— storage sanitization;

— virtualization security;

— self-encrypting storage devices and data encryption software;

— key management services;

— data authenticity and integrity services;

— data in motion protections (encryption and data reduction);

— directory services and other user management systems;

— data retention and preservation;

— data protection and recovery.

Knowing both how and why storage technologies are used can lead to a better understand of the security issues and implications for storage. As a starting point, the following points should be considered.

— The storage systems can function as nodes within storage networks, which can be based on, but not limited to, technologies such as Transmission Control Protocol/Internet Protocol (TCP/IP), Fibre Channel (FC), InfiniBand™[1)] (IB), etc. The potential threats can vary significantly based on the networking technology and their topologies.

1) This trademark is provided for reasons of public interest or public safety. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO.

— When stored, the data are typically represented and accessed as either blocks or as files/objects with significant differences between these two types of storage methods. Likewise, the associated security measures 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 capacity than is exposed through the interface. SSDs typically over provision media, where data can be moved internally between physical media areas to improve write latencies and evenly distribute write activity. HDDs often contain spare areas where data can be moved internally between physical media areas when there is an access problem. User data can remain on over provision or spare areas even after subsequent writes to the device occur through its interface. Such data cannot be cleared by overwrites through the interface.

— Storage management is both an element of the storage infrastructure as well as an operation performed on that infrastructure. Commonly this infrastructure can have privileged users applying configuration changes, provisioning storage, tuning, monitoring, etc. Some of the management can be performed remotely and can 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 that security measures are complementary rather than a trade-off and that they do not negate high-availability measures by introducing choke-points and single points of failure.

— Many organizations implement elaborate data resiliency strategies, which are integral to their business continuity management (BCM) plans. Security mechanisms like data at rest encryption can negatively impact these resiliency strategies if not implemented carefully.

— 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 (see NOTE below) or filesystem to a server], mask the true capacity of a device, perform policy-driven autonomous data movement (i.e. tiered storage), or completely abstract the storage infrastructure (i.e. cloud computing storage). Balancing security measures and virtualization to ensure they interoperate requires careful planning and selection of the right technologies.

NOTE In computer storage, a logical unit is a device addressed by the SCSI protocol or protocols which encapsulate SCSI, such as Fibre Channel or iSCSI. A logical unit number (LUN) is a number used to identify a logical unit. The LUN can be used with any device which supports read/write operations, such as a tape drive, but is most often used to refer to a logical storage as created on a SAN. Though not technically correct, the term LUN is often also used to refer to the logical storage itself.

— Data growth rates in some organizations are driving an 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 BCM operations.

— Additional copies of data can be created as a by-product of data protection strategies such as replicating data between systems and sites (see 10.14.3), backups (see 10.14.2), and snapshots (see 10.14.4). Such data copies and residual data have the same sensitivity considerations as the original data.

— Sensitive and high-value data are often transmitted between and within systems (e.g. data in motion), increasing data protection considerations (see 10.5.4).

— Many organizations are implementing data at rest encryption (see 10.5.5) 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 for meeting compliance requirements.

— Successful use of encryption is often predicated on proper management of keying material throughout its lifecycle (see 10.5.2). 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 can 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 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;

— BCM support;

— proper sanitization of storage and disposal of media;

— secure autonomous data movement;

— secure multi-tenancy.

6.4 Storage security risks

6.4.1 Background

Storage security risks are created by an organization's use of specific storage systems or infrastructures. Storage security risks arise from:

— threats targeting the data handled by the storage systems and infrastructure;

— vulnerabilities (both technical and non-technical);

— likelihood of performing a successful exploit of vulnerabilities by threats;

— impact of successful exploitation of vulnerabilities by threats.

Risk management is a key concept in information security. 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 can include:

— unauthorized usage of storage resources;

— unauthorized access;

— liability due to regulatory non-compliance;

— Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks on storage or networks;

— corruption/modification and destruction of data, including backup or recovery copies;

— unauthorized disclosure which can constitute a data breach;

— theft or accidental loss of storage media;

— malware attack or introduction of malicious code (e.g. ransomware);

— improper sanitization or disposal after end-of-use.

These threats can give rise to a wide assortment of risks. For storage systems and infrastructure however, the major concerns are 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.

6.4.2 Data breaches

This document defines a data breach as a compromise of security that leads to the accidental or unlawful destruction, loss, alteration, unauthorized disclosure of, or access to protected data transmitted, stored, or otherwise processed. This definition goes significantly beyond more simplistic data breach definitions that focus on unauthorized access or disclosure of certain types of data.

Depending on the volume and type of data involved 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 loses their or others' secured information. Such information can include:

— 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 information can be well funded and have diverse motivations.

Table 1 summarizes the likely storage-based security threats 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 or loss of storage device or storage media	Unlawful or unauthorized disclosure, data loss, or data destruction
Accidental configuration changes (e.g. storage management, storage/network resources, and incorrect patch management) by authorized personnel	Accidental access, accidental disclosure, accidental data destruction, accidental data alteration, or removal/denial of access
Malicious configuration changes (storage management, storage/network resources, and application tampering) 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	Loss of access by legitimate users and unavailability of storage and data
Malicious monitoring of network traffic	Unlawful/unauthorized disclosure

6.4.3 Data corruption or destruction

Data corruption is the deterioration or damage of data (i.e. undesirable changes to the original data) caused by users, hardware, or software error. It can occur during writing, reading, retention, transmission, or processing. Early detection of data corruption can possibly enable the recovery of data or metadata under the right conditions. If left undetected or corrected, this data corruption can result in permanent data loss if the root cause persists. 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

data 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, and user errors are all examples of general, unintentional sources of data corruption and destruction. Background radiation, HDD head crashes, and aging or wear of the storage media are additional sources of problems that are more storage centric. Data corruption due to storage device or storage media failures can generally be detected by the use of checksums, and can often be corrected by using 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 of a malicious nature can be perpetrated by external parties or insiders with the purpose of making some or all the affected data unusable or destroyed. In this context, unusable can mean that unauthorized modifications have been applied, modifications are suspected, or the data can be encrypted with an unknown key or mechanism (or the key used originally to encrypt the data can be destroyed). Non-malicious incidents typically result from carelessness, lack of knowledge, or intentional circumvention of security measures for such reasons as getting the job done. However, the impact of non-malicious incidents 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 loss of data.

6.4.4 Temporary or permanent loss of access/availability

Availability is concerned with assuring that authorized users and systems have access to data at a required level of performance within a specified time frame. Loss of access, even when temporary, can cause significant harm to an organization. In addition, degraded access (e.g. minimum performance thresholds are not met) can be equally harmful to an organization.

Loss of access or availability can occur due to failures or issues associated with storage devices, storage network elements, stored data, data flows, services, and applications as well as attacks. In general, data availability is achieved through redundancy.

6.4.5 Failure to meet statutory, regulatory, or legal requirements

Organizations can incur significant liabilities and penalties for non-compliance with statutory, regulatory, or legal requirements. For multi-national organizations, country-specific legislation has an important influence on information security requirements.

Common compliance issues include:

— breach of country-specific privacy requirements;

— unlawful transfer of data (e.g. 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 and proof of encryption/sanitization).

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

7 Organizational controls for storage

7.1 General

Users can refer to the extensive set of organizational controls listed in ISO/IEC 27002:2022, Clause 5. Some or all of these controls are relevant to storage systems and storage ecosystems.

7.2 Align storage and policy

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

The following policy guidance and requirements apply.

a) OC-PLCY-G01 Incorporating storage into policies

Storage should be incorporated into policies in order to:

— identify the most sensitive (PII, 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. WORM, authenticity, and access controls);

— address data destruction and storage media sanitization.

b) OC-PLCY-G02 Ensuring storage conforms with policies

Conformance with policies should:

— ensure that all elements of the storage ecosystem comply with policy (e.g. ISO/IEC 27001:2022, 5.2 and ISO/IEC 27002:2022, 5.1);

— give priority to the most sensitive/most critical data.

c) OC-PLCY-R01 Include storage in logging policy

Include storage in the logging policy such that the logging policy shall:

— clearly state that storage systems and devices participate in audit logging;

— identify the significant storage-related events to be collected;

— identify the preservation requirements for storage-related event logs;

— identify the retention and archival requirements for storage-related event logs;

— specify the time synchronization and use requirement for storage-related event logs;

— include evidentiary expectations (authenticity, chain of custody, etc.).

7.3 Business continuity management

ISO 22301 specifies the structure and requirements for implementing and maintaining a business continuity management system that develops business continuity appropriate to the amount and type of impact that the organization accepts following a disruption. The requirements specified in ISO 22301 are generic and intended to be applicable to all organizations, or parts thereof, regardless of type, size and nature of the organization. The extent of application of these requirements depends on the organization's operating environment and complexity. The requirements specified in ISO 22301 address the context of the organization, leadership, planning, support, operations, performance evaluation, and improvement.

ISO/IEC 27002:2022, 5.30 identifies the importance of ICT readiness for business continuity (IRBC) to ensure the availability of the organization's information and other associated assets during disruption. ISO/IEC 27031 describes the concepts and principles of IRBC. It also 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. This framework applies to any organization (private, governmental, and non-governmental, irrespective of size) that is developing its ICT readiness for business continuity programme, and requires its ICT services/infrastructures to be ready to support business operations in the event of emerging events and incidents, and related disruptions, that can 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 programme or informal BCM activities. The controls related to BCM are as follows:

a) OC-IRBC-G01 Factoring storage ecosystem into BCM planning and implementation

Organizations should ensure that the storage ecosystem is factored into the BCM planning and implementation.

b) OC-IRBC-G02 Preparing for limited disruption events

Organizations should prepare for limited disruption events (system failures, adversarial attacks, operator errors).

c) OC-IRBC-G03 Identifying and documenting the unique staffing and facility requirements

Organizations should identify and document the unique staffing and facility requirements associated with the storage ecosystem.

d) OC-IRBC-G04 Performing ongoing planning and testing of BCM

Organizations should perform ongoing planning and regular testing of assumptions, which are critical to successful BCM. Results of BCM testing should be fed back into ongoing maintenance of the BCM plan.

7.4 Compliance

The importance of complying with legal and regulatory requirements drives a significant portion of the security agenda and strategy of many organizations. The following elements are key compliance aspects of storage systems and infrastructure that are of concern to an information systems (IS) auditor.

The controls related to compliance are as follows:

a) OC-CPLC-R01 Ensure storage meets user accountability obligations

Make certain that user accountability obligations are met by ensuring that:

— users, especially privileged users, shall have unique user IDs (i.e. no shared accounts);

— users rights and privileges shall be explicitly granted based on their duties (e.g. roles) in accordance with the principle of least privilege (i.e. users are granted the minimum system resources and authorizations that they need to perform their function);

— all attempted (successful and unsuccessful) management events and transactions shall be logged.

b) OC-CPLC-G01 Ensuring storage meets user traceability obligations

Ensure that storage meets user traceability obligations such that:

— logged event/transaction data should contain sufficient application or system detail to clearly identify the source;

— user information should be traceable to a specific individual;

— when appropriate, log records should be treated as evidence (chain of custody, non-repudiation, authenticity, etc.).

c) OC-CPLC-G02 Ensuring storage meets user monitoring obligations

Ensure that storage meets user monitoring obligations such that:

— the storage layer should participate in the external audit logging measures;

— the audit logging events should be monitored and alert issues when appropriate.

d) OC-CPLC-G03 Ensuring storage meets data retention and sanitization obligations

Ensure that storage meets data retention and sanitization obligations such that:

— appropriate data retention measures should be implemented;

— appropriate data integrity and authenticity measures should be implemented;

— correct data sanitization should be implemented prior to the repurposing or decommissioning of hardware;

— correct sanitization of virtual server images, and their copies, should be implemented at their end of life.

e) OC-CPLC-G04 Ensuring storage meets privacy obligations

Ensure that storage meets privacy (see ISO/IEC 27701) obligations such that:

— appropriate data access control, based on least privilege, should be implemented to control access to data and metadata (e.g. search results);

— appropriate data confidentiality measures should be implemented to prevent unauthorized disclosure.

f) OC-CPLC-G05 Storage taking legal obligations into consideration

Legal obligations can apply to storage such that:

— use of data deduplication should not conflict with data authenticity requirements;

— use of data and storage media sanitization mechanisms should not violate preservation requirements;

— proper chain of custody procedures should be followed when evidentiary data (e.g. audit logs, metadata, mirror images, and point-in-time copies) is handled.

8 People controls for storage

Users can refer to the extensive set of people controls listed in ISO/IEC 27002:2022, Clause 6. Some or all of these controls are relevant to storage systems and storage ecosystems.

ISO/IEC 27002:2022, 6.3 states that an organization should identify, prepare and implement an appropriate training plan for technical teams whose roles require specific skill sets and expertise. For individuals responsible for aspects of storage security, the skill and expertise can be highly specialized as demonstrated by this document.

The controls related to expertise are as follows:

a) SC-XPTS-G01 Ensuring adequate storage protection expertise

Storage teams and management should understand the technologies and practices associated with storage-based data protection technologies (see 10.14). In addition, this understanding should include the applicability of data protection in the organization.

b) SC-XPTS-G02 Ensuring adequate storage security expertise

Storage teams and management should understand the technologies and practices for securing storage devices, systems and ecosystems, as well as leveraging storage-based security controls.

9 Physical controls for storage

9.1 General

Users can refer to the extensive set of physical controls listed in ISO/IEC 27002:2022, Clause 7. Some or all of these controls are relevant to storage systems and storage ecosystems.

9.2 Physically secure storage

Storage systems and ecosystems are particularly susceptible to physical threats. They can be subjected to theft, destruction, and unauthorized access unless protected when not in use or not under observation. The size of storage media, devices, and systems can vary significantly, so the methods to secure them can also vary.

The controls related to physically securing storage are as follows:

a) PC-PHYS-G01 Physically securing storage media

All storage media with recorded data should be stored in a safe, secure environment according to their information classification and which protects them against environmental threats (such as heat, moisture, humidity, electronic field or ageing), in accordance with manufacturers' specifications.

b) PC-PHYS-G02 Physically securing storage devices

Depending on their size and location, storage devices should be physically secured:

— small-size units in an office environment can be tethered to furniture or walls as well as being stored in locked cabinet or safes when not in use; or

— large-size units can be rack-mounted and resident in laboratory or data centres where they benefit from the security measures of the facility (e.g. locking covers or cages).

9.3 Protect physical interfaces to storage

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, can result in data destruction, corruption, or denial of access.

Management interfaces for storage systems can take on several physical forms including serial ports, 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.

The controls related to protecting physical storage interfaces are as follows:

PC-PHYS-R01 Protect physical interfaces

To protect these physical interfaces, the organization shall:

— 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) is considered to be a best practice;

— disable modem ports when not needed.

9.4 Isolation of storage systems

Physical and logical isolation of storage devices (e.g. within a SAN) can also play an important role. The controls related to isolation of storage are as follows:

a) PC-PHYS-G03 Physically isolating storage systems

Physical isolation should be used to:

— segregate production from other system classes (e.g. quality assurance, development) including:
  — where possible, avoiding network connections between classes (e.g. a production server connected to both the production and development networks),
  — segregating networks and storage by class where appropriate,
  — physically separating systems in each class;
— isolate storage devices from other data centre devices, if practical.

b) PC-PHYS-G04 Logically isolating storage systems

Logical isolation should be used to:

— 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;
— configure network gateways for appropriate network segregation.

10 Technological controls for storage

10.1 General

Users can refer to the extensive set of physical controls listed in ISO/IEC 27002:2022, Clause 8. Some or all of these controls are relevant to storage systems and storage ecosystems.

10.2 Design and implementation of storage security

10.2.1 General

With enormous growth in volumes of critical data, many organizations have adopted storage-centric architectures for their ICT infrastructure. Consequently, storage security plays an important role in protecting this data, and in many instances, it serves as the last line of defence. The effectiveness of this storage security is often influenced by design considerations.

The controls related to storage design and implementation are as follows:

a) TC-DSGN-G01 Adhering to core security design principles

Designing and implementing storage security solutions requires adherence to core security design principles. In addition, the controls and guidance described in Clauses 7, 8, and 9 should be integrated into the design and implementation of storage security solutions to counter applicable threats. Data sensitivity, criticality and value can also be an important consideration in designs (see 10.2.2).

b) TC-DSGN-G02 Considering relevant threats in design

Common risk areas associated with storage security architectures are design failures due to poor design or the lack of appropriate consideration for business continuity management planning, or lack of correspondence to the current or expected threat level. A design should consider all relevant threats and vulnerabilities in the storage system as described in 6.4.

Information on assessing security risks and associated threats can also be found in ISO/IEC 27001, ISO/IEC 27002 and ISO/IEC 27005.

10.2.2 Storage security design principles

10.2.2.1 Defence in depth

Increasingly, organizations are considering security from 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 such as data value and sensitivity, compliance requirements, adversarial capabilities and activities.

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 can be anti-virus software installed on individual workstations when there is already virus protection on the firewalls and servers within the same environment.

The controls related to defence in depth are as follows:

TC-DSGN-G03 Deploying defence in depth

Storage systems and solutions should:

— 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 multiple classes of attacks;

— deploy multiple layers of defence mechanisms between potential adversaries and targets;

— include both detective and protective mechanisms;

— leverage robust (e.g. support many use cases and volumes) and highly attack resistant (e.g. zero trust architectures) security infrastructures (e.g. key management, public key infrastructure, and identity management) that are centrally managed and monitored;

— 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, anti-virus updates, and maintain ACLs);

— perform regular security threat assessments to evaluate security readiness;

— monitor and react to current threats.

For storage, multiple layered defence mechanisms mean that security controls are deployed and used throughout the storage infrastructure, including the converged network adaptor (CNA) in host computers, storage network switches/routers, storage appliance, and storage devices.

NOTE A CNA is a single network interface device that provides the functionality of both a FC host bus adapter (HBA) and a TCP/IP Ethernet network interface card.

10.2.2.2 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) are segregated. This creates a way to ensure the systems make available only the data that is necessary for conducting the tasks for that particular domain. As a design principle, the architecture enforces domain separation to ensure that resources to which an entity has access cannot be accessed or affected by another domain.

For storage infrastructure, a security domain is typically represented as a SAN, especially when sensitive data are being stored and processed within the storage systems. In situations where the data sensitivity is low, use of 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-2 Zoning.[61]

Building on the compartmentalization principle described in ISO/IEC 27033-2, the following storage security design recommendations are relevant:

a) TC-DSGN-G04 Factoring data sensitivity into design of security domains

Data sensitivity should be factored into the design of security domains such that:

— 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) 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.

b) TC-DSGN-G05 Factoring purpose into design of security domains

Purpose should be factored into the design of security domains such that:

— storage and storage networks used for different purposes (e.g. development, production, and management) and using different technologies (e.g. SMB, NFS, iSCSI, and CDMI) should be located in separate security domains;

— storage networks should be in different security domains than regular networks using the same technology (e.g. the IP networks carrying iSCSI traffic should be segregated from normal corporate LANs);

— storage device and storage network management systems should be located in dedicated security domains;

— development systems should be in different domains than production systems.

c) TC-DSGN-G06 Using further isolation within a security domain

Storage devices that are permitted to reside within a single security domain, but used for multiple purposes or hold multiple levels of sensitive data, should be further isolated to minimize possible interactions.

10.2.2.3 Design resilience

The controls related to design resilience are as follows:

TC-DSGN-G07 Including resilience in design

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.

10.2.2.4 Secure initialization

The controls related to secure initialization are as follows:

TC-DSGN-G08 Supporting a secure initialization sequence

As a design principle, the architecture should support a secure initialization sequence during the transition from a down state to the operating state (e.g. after a power-on or reset). During the initialization phase, externally accessible processes and network interfaces should be denied access or not be available 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.

10.2.3 Storage system quality attributes

10.2.3.1 Reliability

ISO/IEC 27000 defines reliability as the property of consistent intended behaviour and results. For storage, reliability is often considered the probability that a device performs its required function under stated conditions for a specific period of time and is quantified as:

— MTBF (mean time between failures) for a repairable product, which is the expected time between consecutive failures in a system or component and is sometimes thought of as the average time available for a system or component to perform its normal operations between failures (see Figure 1);

— MTTR (mean time to repair) for a repairable product, which is the expected or observed duration to return a malfunctioning system or component to normal operations and is sometimes thought of as the average time to repair a failed component;

— MTTF (mean time to failure) for non-repairable a product, which is the average time available for a system or component to perform its normal operations until it fails.

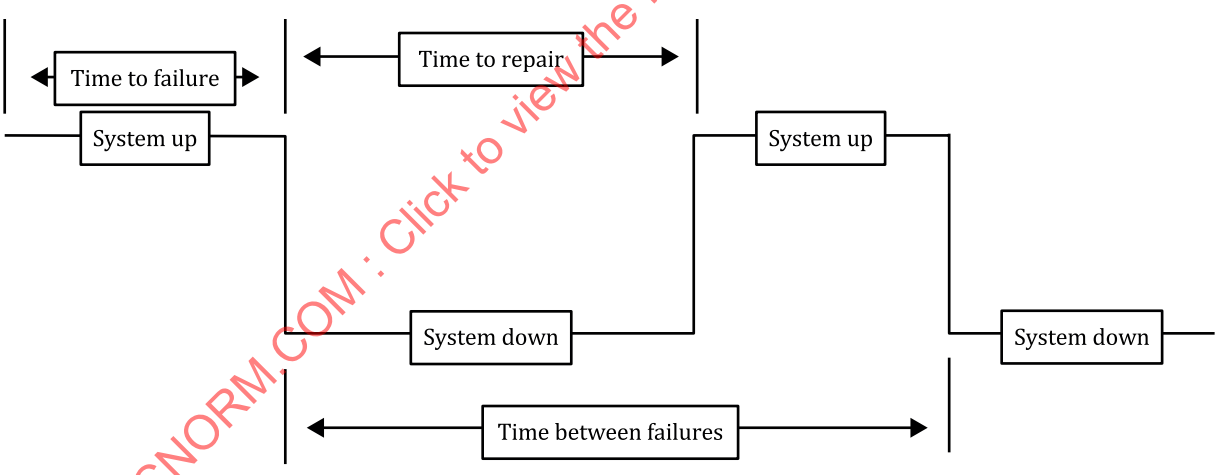

Figure 1 — Quantification of reliability

Within the context of storage, system compromises and attacks (such as DDoS) can have negative impacts on MTBF, MTTR (e.g. MTTR can also involve recovery operations after an attack), and MTTF. In addition, the inclusion of security features, the application of system or application patches, or other system hardening measures like those described in 10.3.1 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 controls related to reliability are as follows:

TC-DSGN-G09 Minimizing impacts on storage reliability

Specific reliability guidance includes:

— the reliability of the storage system and infrastructure should not be adversely impacted by the inclusion of security features;

— storage vulnerabilities should be proactively managed to minimize their impacts on system reliability;

— controls should be assessed to determine whether they can ensure the reliability and security of data.

10.2.3.2 Availability

In the context of storage, data availability typically refers to how accessible data are 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 reliably the data are available for access by authorized users, in terms of uptime, and how quickly someone can access the data.

Availability is measured as the proportion of time something is available for access by authorized users with certain performance assurances (e.g. no more than 5 milliseconds is required to retrieve a piece of information). For example, a storage array that had approximately 5 minutes of downtime in a year, assuming 24-hour per day, 7- days per week operations, has an availability of 0,999 99 (99,999 %).

To achieve high availability of data, significant amounts of hardware and software redundancy (e.g. automated input/output path failover, redundant components, 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) are often used to ensure fast data recoveries in the event of a failure.

The controls related to availability are as follows:

TC-DSGN-G10 Minimizing impacts on data availability

Because of the importance of availability, impacts on data availability should be minimized due to:

— inadequate storage security designs and implementations (e.g. minimize single points of failure);

— insufficient management of data encryption keys which can cause problems when keys are unavailable where they are needed or they are inadvertently destroyed;

— inadequate data recovery mechanisms to guard against major data losses or storage outages.

10.2.3.3 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 to normal operation (such as attacks, accidents or large-scale natural disasters). 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 can be acceptable, if the service is still being delivered and the integrity of that service is still there. However, the resilience of a storage system or storage infrastructure is often determined by the least resilient component in the system, and cost/performance or other factors can limit the extent to which resilience is possible or practical.

The controls related to resilience are as follows:

TC-DSGN-G11 Minimizing impacts on storage resilience

Impacts on resilience should be minimized due to:

— failure to include security as an integral part of the resilience strategy that accounts for unit failures, attacks, and compromises of both the storage and security technologies;

— insufficient redundancy storage components;

— inadequate use of diverse components that are easily repairable;

— inappropriate implementation of security features and functionality (e.g. encryption and centralized authentication) that reduces the overall resiliency of the storage system and its infrastructure.

10.2.3.4 Integrity

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. Data stored on a storage device or transmitted across a network in response to a storage request can be corrupted due to hardware or software malfunction. A malfunction in hardware can also trigger software misbehaviour resulting in serious damage to stored data. Bugs in software, like device drivers, can also result in unexpected modification of data. Unreliable networks can corrupt data that pass through them.

In addition, important information can be modified by malicious programs or malicious users, or faulty system components. For example, virus code can be inserted into binary executables, potentially resulting in the loss of all data stored on a system. Operating systems that allow access to raw disks can inadvertently aid an attacker to bypass security checks in the file system and cause damage to stored data.

User errors can compromise data integrity at the application level. For example, an inadvertent deletion of a critical file (e.g. database schema file) can result in data corruption.

The controls related to integrity are as follows:

TC-DSGN-G12 Minimizing impacts on data integrity

To address data integrity issues, organizations should employ common integrity assurance techniques, which include data replication or mirroring, RAID parity, or check summing. Integrity assurance mechanisms that perform preventive steps so as to avoid specific types of integrity violations (e.g. read-only storage and journaling file systems) or that are capable of recovering from damage once a problem is detected, should also be used.

10.2.4 Retention, preservation, and disposal of data

Retention is focused on keeping records according to records requirements for evidence of a business function, activity or transaction and for records processes including how, and how long records are kept. Preservation, on the other hand, is focused on the measures taken to maintain the usability, authenticity, reliability and integrity of records over time. The terms “retention” and “preservation” are often used interchangeably and incorrectly, resulting in different and conflicting records requirements that govern how the same information is maintained, how long it should be kept, and whether and how it is protected and secured.

The controls related to retention, preservation, and disposal of data are as follows:

a) TC-DSGN-G13 Ensuring data retention and preservation

Organizations that have data retention and preservation obligations should store data in a manner that blocks records destruction or alteration (i.e. immutable) along with integrity verification (e.g. hashing) and enforcement of explicit retention periods (e.g. legal holds). To meet immutability (non-editable) requirements, organizations should use write once read many (WORM)-based storage or object-based storage (see 10.12) implementations that combine WORM with metadata that can be used to perform explicit integrity checks as well as enforce data expirations.

b) TC-DSGN-G14 Ensuring proper data disposal

Within common records and information management frameworks (see NOTE 1 below), disposition is the last stage of a record's lifecycle. Within these frameworks, disposition does not necessarily mean records destruction, but rather transfer to archives (see 10.15). In the latter case, this can simply delay when records destruction occurs for most records (few records outside of government should be retained

indefinitely). When records (data) are no longer required, the destruction of the data becomes a critical and often necessary component of an effective data governance programme. Organizations should use a data destruction process that removes information in a way that renders it unreadable for paper records or irretrievable for digital records (see NOTE 2 below). In the latter case, storage sanitization techniques (see 10.6) are often used.

NOTE 1 ISO 15489-1 is one of many frameworks for planning and implementing a records management program.

NOTE 2 In the digital world, making data irretrievable is caveated to a specified level of effort to retrieve it.

10.3 Storage systems security

10.3.1 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 27002:2022, 8.8, there are many existing best practices for various operating systems that can be referenced based on the operating system that is being used.

In storage systems, operating systems exist in many types of devices (e.g. storage arrays, SAN switches, virtualization appliances, backup and archiving appliances). Other lower-level devices (e.g. HBAs, storage directors, network adaptors) also require periodic software and firmware updates. Certain storage operating system security features are not turned on by default, and routine maintenance and updates can reset hardening operations previously performed by storage administrators.

The controls related to system hardening are as follows:

a) TC-HARD-G01 Performing basic operating system hardening

As part of the security hygiene of storage systems, organizations should:

— remove unneeded/unused software, services, and protocols;

— remove unnecessary accounts;

— rename or disable predefined or default accounts when possible as well as change all default passwords;

— only open up network ports that are needed;

— install the latest patches from a trusted source;

— update firmware from a trusted source;

— install and maintain malware protections (see ISO/IEC 27002:2022, 8.7);

— apply vendor-recommend security configurations.

b) TC-HARD-G02 Using software updates and patches from trusted sources

When elements of the storage infrastructure receive an update (e.g. firmware) 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.

10.3.2 Security auditing, accounting, and monitoring

Compliance regulations and contractual clauses often include monitoring and reporting requirements. Event logging and systems accounting are key capabilities to help address these requirements. Of these two, event logging is probably more useful from a storage security perspective because it can be used both in real-time and as part of an incident investigation.

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) that can be recorded in some manner of event logging. From a security or compliance perspective, it is important to capture those event log entries which are 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 are only useful for debugging purposes, providing system health status, warning of minor configuration problems, etc. From an audit logging perspective, the management events (i.e. what a user did) are always of interest, the data access events are usually of limited interest (except in situations where critical files and directories are tightly monitored), and control events are typically of the least interest (though 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 is often required 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.

The following security auditing, accounting, and monitoring requirements and guidance are applicable to storage systems:

a) TC-HARD-R01 Perform logging on storage

Perform logging on storage such that:

— logging shall be enabled on storage systems and devices where possible;

— storage systems and devices shall use external or centralized event logging and can use local logging;

— a common, accurate time source shall be used across the environment to ensure that event records from different sources can be correlated;

— where possible, storage systems and devices shall natively log events using standard logging protocols such as syslog that support reliable delivery and secure transports (e.g. TLS);

NOTE 1 Syslog is defined in IETF RFC 5424[47] with additional details contained in IETF RFC 3195,[41] IETF RFC 5425,[48] IETF RFC 5426,[49] IETF RFC 5427,[50] IETF RFC 5848,[51] IETF RFC 6012,[52] and IETF RFC 6587.[54]

— external or centralized event logging should be used with a trusted remote source;

NOTE 2 A trusted external event logging source is an ICT security management product located in a dedicated security zone or domain and is assumed to enforce its security functions correctly.

— use of device logs for anything other than system health monitoring and debugging should be avoided because device resident logs are more easily subjected to tampering or destruction, there is limited storage space available for logs, and they preclude the use of centralized automated analysis, alerting, and archiving;

— storage systems and devices should use multiple, external log servers;

NOTE 3 Some logging protocols use unreliable network protocols such as User Datagram Protocol (UDP)[39] and therefore log messages can be lost due to network or server performance. Sending messages to multiple log destinations reduces the risk of inadvertent loss.

— storage systems and devices should be configured to log events as they occur (i.e. no buffering) when the primary drivers for audit logging are compliance, accountability, or security.

b) TC-HARD-G03 Ensuring completeness of storage audit logging

The usefulness of event logging is somewhat dependent on the information captured (e.g. timestamps, sources, types of event, etc.). For storage systems, ensuring completeness of the audit log entries includes:

— once the types of events to be logged have been determined, then all occurrences of these events should be consistently 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 can be translated to different languages);

  — a description of the event.

c) TC-HARD-G04 Implementing appropriate monitoring of storage

Implementing appropriate monitoring such that:

— careful use of filtering (e.g. on fields like severity) that complies with the logging policy should be employed by storage systems and the audit log infrastructure;

— an analysis protocol should be implemented to correlate audit log records across event sources to identify significant security events that provide indication of security incidents;

— the storage logging should be included in the security information and event management solutions, when such technology is deployed;

— the storage systems and devices should be included in the information security continuous monitoring (ISCM) solutions, when such technology is deployed.

d) TC-HARD-G05 Using log retention and protection for storage

Implement appropriate retention and protection such that:

— audit log data that can have evidentiary value should be handled correctly (i.e. maintain chain of custody, and verifiable integrity and authenticity);

— audit log data with specific retention requirements (e.g. for regulatory compliance) should be preserved with the organization's data retention solution;

— appropriate measures to preserve log integrity and prevent their modification or destruction (either maliciously or accidentally) should be implemented;

— when audit log entries contain sensitive information, the audit log data should be protected with appropriate confidentiality mechanisms;

NOTE 4 Some log entries can expose things like passwords (e.g. when a user types a password instead of the user ID), but more subtle problems can exist as well (e.g. search commands that expose specific names and health issues).

— for unique audit logging requirements (e.g. high volume, special preservation, and event signing) dedicated and specially hardened and configured systems should be used;

— log relays and log filtering should be leveraged to minimize the impact of specialized storage requirements (e.g. write once read only or WORM).

10.3.3 Storage vulnerability management

ISO/IEC 27002:2022, 8.8, identifies the importance of managing technical vulnerabilities as well as providing specific guidance. This guidance addresses the identification, evaluation, and remediation of technical vulnerabilities.

The controls related to storage vulnerability management are as follows:

TC-HARD-G06 Including storage in vulnerability management programmes

Due to the specialized nature of storage technologies, storage systems are not always included in an organization's vulnerability management programme. In addition, many of the tools used to identify vulnerabilities do not provide extensive coverage of storage operating systems and applications. Organizations should include storage systems in their vulnerability management programmes.

10.4 Storage management

10.4.1 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 devices, greater data protection and security, centralized auditing, and meeting compliance obligations.

Storage systems can be managed using in-band or out-of-band mechanisms (see Figure 2). In this context, in-band management typically means that the storage system is managed through the path over which the user data are transferred. Storage management commands are sent through an input/output connection using serial-attached SCSI, Fibre Channel, or internet SCSI, PCI Express®[2], etc. Out-of-band management, on the other hand, uses an alternate path, different from the one used by the user data traffic, to gain access to the management interface of the storage system. Out-of-band management sends commands over a network connection or sideband interfaces. The management functionality and security considerations can vary significantly between these two types of storage management, even when both are available on a storage system.

2) This trade name is provided for reasons of public interest or public safety. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO.

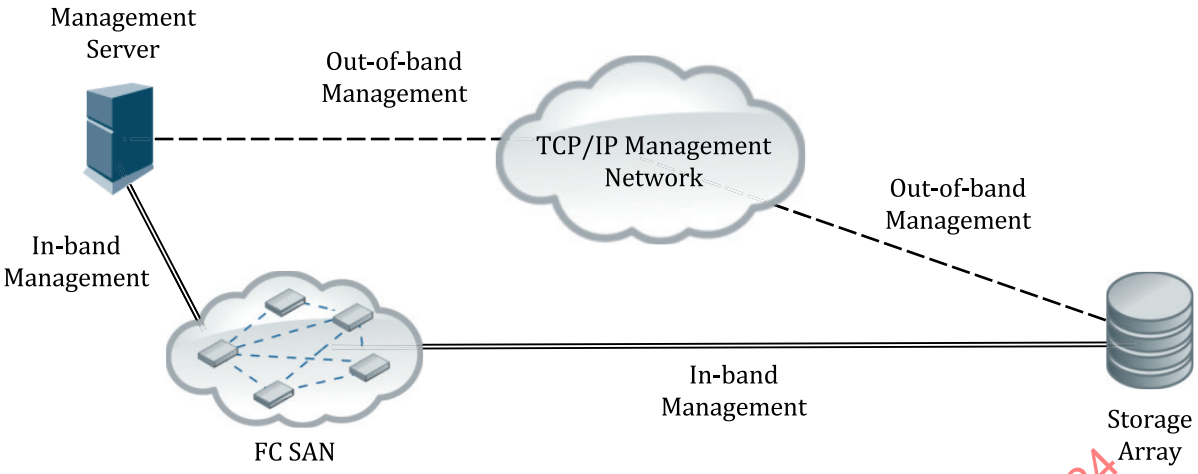

Figure 2 — Storage management example

Like the network management description in ISO/IEC 27033-2, storage management also refers to the activities, methods, procedures, and tools that apply to storage systems, such as:

— 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, e.g. when equipment is replaced, when storage array firmware is updated, 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 the process for defining storage assets to be assigned to a host;

— Sanitization (see 10.6) deals with preserving the confidentiality of data remaining on storage medium when it is removed from service or re-purposed by rendering the data irretrievable.

Performing these storage management activities securely requires controls associated with authentication and authorization (see 10.4.2), protecting the storage management interfaces (see 10.4.3), maintaining accountability and traceability of systems and users (see 10.3.2), and ensuring the underlying systems used for storage management are adequately hardened (see 10.3.1).

10.4.2 Authentication and authorization

10.4.2.1 Authentication

The individuals managing storage systems and infrastructure are generally privileged users. Inappropriate use of system administration privileges (any feature or facility of an information system that enables the user to override system or application controls) can be a major contributory factor to the failure or compromise of systems. To help mitigate these risks, secure log-on procedures as described in ISO/IEC 27002:2022, 8.5 with additional authentication measures as necessary can be used.

The controls related to authentication are as follows:

a) TC-MGMT-R01 Minimum user authentication measures

Storage management typically involves privileged operations that are performed by users who have a minimum level of authentication that includes:

— all users shall have a unique identifier (user ID) for their personal use only;

— to substantiate the claimed identity of a user, one of the following suitable authentication techniques shall be used:

  — password of sufficient complexity and secrecy that is impractical for an attacker to guess or otherwise discover;

  — strong authentication (e.g. challenge-response protocol); or

  — multi-factor authentication, such as biometric data (e.g. finger-print verification or retina scan) and use of hardware tokens (e.g. smart cards).

— all remote access shall use strong authentication or multi-factor authentication along with secure channels.

b) TC-MGMT-G01 Using centralized authentication solutions

A centralized authentication solution, such as Remote Authentication Dial-in User Service (RADIUS), single sign-on, Open Authorization (OAuth), Security Assertion Markup Language (SAML), etc. should be used to improve monitoring and control.

c) TC-MGMT-G02 Using multi-factor authentication

Multi-factor authentication should be used when managing sensitive and high-value data.

d) TC-MGMT-G03 Disabling login to the root or admin account

Logins to the root or admin account should be disabled.

e) TC-MGMT-G04 Remotely logging all privilege escalation operations

All privilege escalation operations should be remotely logged.

f) TC-MGMT-G05 Using entity authentication mechanisms

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 such as Challenge Handshake Authentication Protocol (CHAP) with iSCSI, Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP) within FCP, etc. When possible, these entity authentication mechanisms should be used.

10.4.2.2 Authorization and access control

Within market sectors such as financial services and healthcare, there are trends to align authorization and access control to the principle of least privilege by leveraging specific roles.

The controls related to authorization and access control are as follows:

TC-MGMT-G06 Separating security and non-security roles

The following roles should be implemented and used within storage technologies:

— Security Administrator. This role has read-only/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 except for security-related elements or data (i.e. those covered by the security administrator);

— 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 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.

10.4.3 Secure the management interfaces

In addition to the physical interfaces (see 9.3), storage systems employ a variety of software and firmware to enable management of the storage system. These software interfaces can include simple command line interfaces (CLI), Web-based interfaces, including graphical user interfaces or REST application programming interfaces (API), support for the Simple Network Management Protocol (SNMP), and server-based proxies that handle in-band management (i.e. over the data path).

The controls related to securing management interfaces are as follows:

a) TC-MGMT-G07 Securing the network interfaces to management software/firmware

Storage management software/firmware interfaces are typically secured with the following:

— firewalls and TCP wrappers should be used to restrict access to management networks to authorized systems and protocols;

— entity authentication should be used to establish trust relationships between storage systems and the management systems (e.g. using FC-SP-2 AUTH-A[61] to authenticate the entities performing in-band management);

— intrusion detection system and intrusion prevention system mechanisms should be leveraged to identify anomalous behaviours and guard against them;

— ICT infrastructure such as Domain Name System (DNS), Service Location Protocol (SLP), or Network Time Protocol (NTP) should be used with appropriate security controls to avoid indirect attacks;

— appropriate privileged user controls, including authentication (see 10.4.2.1), authorization (see 10.4.2.2), and secure auditing/monitoring (see 10.3.2) should be employed;

— operating systems and applications should be current and sufficiently hardened against attacks (see 10.3.1).

b) TC-MGMT-R02 Secure the remote management

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

— use secure channels such as virtual private network (VPN), TLS, Secure Shell (SSH), or Hypertext Transfer Protocol Secure (HTTPS) for all remote access;

— employ strong authentication or multi-factor authentication;

— restrict privileges to the minimum necessary (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 to availability, integrity and confidentiality.

c) TC-MGMT-R03 Restrict vendor remote management

Technical controls shall 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 shall be generated.

d) TC-MGMT-R04 Restrict dial-up access use

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

e) TC-MGMT-R05 Secure IPMI

Some storage systems include hardware-based platform management systems making it possible to control and monitor systems centrally, including the ability to manage servers in remote physical locations regardless of the installed operating system. This software does not require permission from the storage system's operating system as it runs on separate hardware attached to a motherboard or server.

In some implementations, a baseboard management controller (BMC) sits between on-board sensors and interfaces and out-of-band communication interfaces such as Ethernet. The Intelligent Platform Management Interface (IPMI) addresses a range of interfaces, including BMCs, that provide low-level access to a system that can override operating system controls. IPMI messages can be transmitted to and from the BMC in Remote Management Control Protocol (RMCP) UDP datagrams (UDP target port 623 for asf-rmcp) encapsulated. This functionality is also called "IPMI-over-LAN". IPMI also defines LAN-specific configuration settings, somewhat like those for IP addresses. An additional packet format (RCMP+) has also been defined in IPMI 2.0. RMCP+ supports encrypted data transmission in addition to various extensions for authentication.

When IPMI is used, the following additional security measures are necessary:

— IPMI shall be disabled as the default configuration setting, and IPMI should only be enabled on a temporary basis when needed;

— IPMI traffic (usually UDP port 623) shall be restricted to trusted internal networks such as a management VLAN segment with strong network controls;

— Strong, unique passwords shall be set and used for the IPMI service on devices running IPMI;

— Encryption should be enabled on IPMI (e.g. with RMCP or RMCP+ protocols), if possible;

— "Cipher 0" (enabled by default on many IPMI enabled devices) and anonymous logins shall be disabled to prevent attackers from bypassing authentication and sending arbitrary IPMI commands;

— Stored passwords shall be eliminated at the system's end of life.

10.5 Data confidentiality

10.5.1 General

Within storage infrastructures, data confidentiality is typically maintained using some method of encryption. These methods are most often associated with protecting data while it is transferred (sometime referred to as in flight or in motion) within the storage infrastructure or as it is stored (or at rest) within a device or on storage media (see Reference [74] for a useful summary of encryption and key management for storage systems and ecosystems).

The process of encryption is a matter of applying an encryption algorithm (or cipher) to plaintext data yielding encrypted data (or ciphertext). Conversely, a decryption transforms ciphertext back into its original

plaintext. The definition and specification of many important ciphers relevant to storage can be found in the ISO/IEC 18033 series, NIST FIPS 197, and IEEE 1619.2-2021.

For some types of ciphers (e.g. n-bit block ciphers) there are multiple ways (called modes of operations) 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, NIST Special Publication 800-38A, NIST Special Publication 800-38C,[66] NIST Special Publication 800-38D,[67] NIST Special Publication 800-38E,[68] and IEEE 1619-2018.

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 the ISO/IEC 11770 series[and NIST Special Publication 800-57 Part 1[69] and Part 2.[70] ISO/IEC 27002:2022, 8.24 also provides relevant guidance on key management.

10.5.2 Encryption and key management issues

The use of cryptographic technology introduces certain issues that cannot be ignored. Failure to address these issues can expose an organization to regulatory penalties as well as causing catastrophic losses under certain conditions.

The controls related to encryption and key management are as follows:

a) TC-CNFD-G01 Complying with import/export regulation for cryptography

Cryptographic technologies can have strict regulations governing the import/export of the technology. To avoid issues, organizations are expected to:

— understand and obey government import regulations associated with encryption and key management;

— understand and obey government export regulations associated with encryption and key management.

b) TC-CNFD-G02 Complying with key escrow and disclosure requirements

Some organizations use key escrow services to manage third party access, including mandatory key disclosures, to certain parts of their systems. When key escrow services are used, it is presupposed that an organization will:

— comply with corporate or government key escrow requirements;

— understand and follow 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.

c) TC-CNFD-G03 Planning for key failures

The loss or corruption of encryption keys can render data unusable, so organizations often make provisions to protect their encryption keys (e.g. a key backup of a key management server). Key backup is normally implemented in the context of a specific encryption/key management solution and is focused on providing the solution access to the keys used to encrypt data within the solution.

To ensure there is adequate protection against key loss or compromise, organizations should have:

— a recovery plan in the event of a key loss or compromise;

— a key backup plan in place to ensure continued access to encrypted business/mission critical data.

d) TC-CNFD-G04 Limiting operational impacts of cryptography

The use of cryptography, encryption in particular, can have operational impacts that can impact the effectiveness or efficiency of ICT infrastructure. Organizations should minimize these impacts, including:

— diminished effectiveness of network-based data loss prevention technologies due to the use of end-to-end encryption on communications;

— expanded network or storage utilization due to the inability to apply data reduction technologies (deduplication and compression techniques) on ciphertext;

— inability to apply centralized malware scanning on encrypted data.

10.5.3 Encryption of storage

There are multiple considerations to be addressed 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 and anti-virus);

— encryption carries the risk of making data unavailable if anything goes wrong with data handling, data transformations, key management, or the actual encryption;

— encryption can require non-trivial computational resources;

— encryption can necessitate centralized key management especially when encryption is used in conjunction with out of region replication for BCM purposes;

— encryption can diminish or negate the benefit of data reduction technologies (e.g. compression and deduplication) because encrypted data are not easily compressed;

— the quality of the cryptography (security strength and industry-tested and accepted algorithms) can impact the actual protection offered.

Not all data are 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 mechanisms to safeguard the confidentiality of data when the data are considered a critical asset.

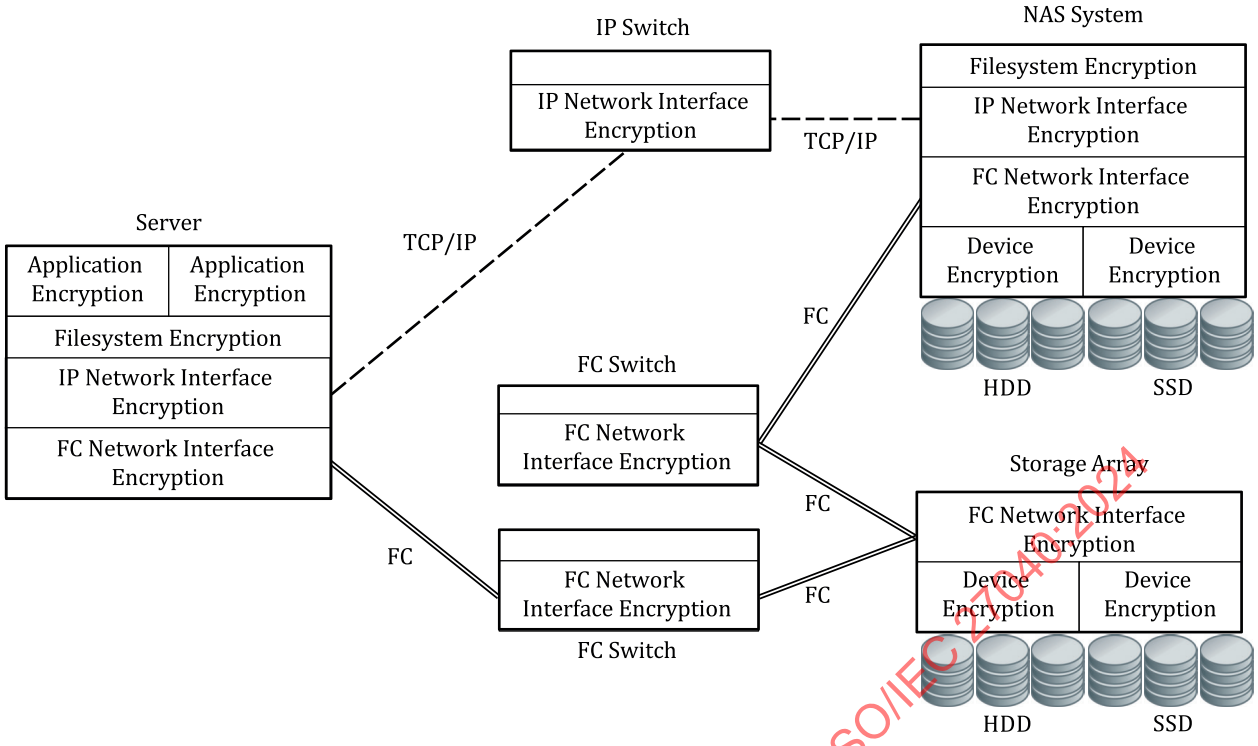

Figure 3 — Example points of encryption

Protecting data in flight typically involves two or more communicating entities establishing an encrypted channel which is used to transfer data (see 10.5.4). This connection is usually ephemeral and the security is usually negotiated whenever it is needed.

Protecting data at rest (see 10.5.5) typically involves a single point in the data path (point of encryption) that is used for encrypting/decrypting the data. The point of encryption is important because it represents the location within the data path where the data are encrypted (ciphertext) and where it is usable (plaintext form). A common security perspective is to encrypt as close to the source or use as possible, as this tends to maximize the protection provided, but there can be many options in selecting a point of encryption (see Figure 3), including:

— Application-level, e.g. a specific application or database, provides finest granularity of control and maximum insight into the data (type, users, sensitivity).

— Filesystem-level, e.g. an operating system or operating system-level application, provides control at file-level with insights into the users.

— Network-level, e.g. an HBA, array controller, or switch, where:

  — file-based (NAS) provides control at the share/filesystem-level (possibly file-level) with moderate insights into the users;

  — block-based provides control at the logical volume level with limited insights in the community of users.

NOTE 1 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.

— Device-level, e.g. tape drive, storage array, and disk drive, provides control at the storage media level (and possibly at the logical volume level) with limited insights in the community of users.

Without careful design and on-going monitoring of changes to the storage ecosystem, there is a possibility that multiple points of encryption can end up being implemented. Such a situation can have negative impacts, especially when the ICT staff are unaware of the data at rest encryption deployments.

The following can be used to protect data on storage using data at rest encryption (see also 10.5.5):

a) TC-CNFD-R01 Use cryptography with at least 128 bits of security strength

Cryptography with at least 128 bits of security strength shall be used throughout the encryption solution.

b) TC-CNFD-G05 Avoiding use of storage encryption as primary protection for sensitive data

Storage-based encryption should not be the primary confidentiality protection for sensitive data.

NOTE 2 The storage encryption is typically active only while the data are 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 are accessed).

c) TC-CNFD-G06 Selecting an appropriate point of encryption

Selection and implementation of a point of encryption should be compatible with BCM (see 7.3), data reduction (see 10.13), data protection (see 10.14), and data confidentiality (see 10.5) requirements.

d) TC-CNFD-G07 Using appropriate encryption and key management compatible with data retention and preservation requirements

Encryption and key management solutions should be compatible with data retention and preservation requirements.

e) TC-CNFD-G08 Using validated cryptographic modules for sensitive or regulated data

Cryptographic modules used to protect sensitive or regulated data should be validated using recognized criteria (e.g. ISO/IEC 19790, ISO/IEC 15408 series and NIST FIPS 140-3[63]).

f) TC-CNFD-G09 Producing and retaining storage encryption records

As with sanitization, it is important that an organization maintain records of its data at rest encryption to document the storage media that 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 data, these records or proof of encryption can be instrumental in demonstrating that no data breach occurred, thereby avoiding costly data breach notifications and other liabilities. The following should be included for proof of encryption:

— ensure that the encryption mechanisms create appropriate audit log entries (activation, verification, integrity checks, re-keying, etc.);

— perform regular and audited checks that encryption was properly performed and consider outside accreditation.

g) TC-CNFD-G10 Following basic key management principles

Successful use of cryptography is dependent on adhering to basic principles associated with keying material as well as key management. Storage systems and devices that integrate data at rest encryption and key management can further improve security with the following:

— centralized key management should be used for key lifecycle management;

— key management should be automated whenever possible;

— keys with a long life (i.e. approaches the maximum recommended cryptoperiod, which is typically no more than 1 to 2 years, depending on the key type) should be used sparsely;

— strict access controls should be used to limit user capabilities and separation of duties constraints (e.g. a security role) for key generation, change, and distribution.

10.5.4 Encrypting transferred data

10.5.4.1 General

Within storage infrastructures, data confidentiality or integrity (digital signature or authentication code) of the data being transferred between two points can be of interest, especially for data that leaves the confines of a physically controlled data centre. In addition, the transfer of data within storage systems can be of concern.

Protocols such as FC ESP_Header,[60] IPsec (see 10.5.4.3), TLS (see 10.5.4.2), or even computer-based encryption techniques can provide additional protection to the data as it is transferred. These methods are most often associated with protecting data while it is in motion (also known as in flight or in transit).

Data in motion protection is generally a temporary protection of the data, which exist only while data are being moved. For data 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, IPsec RFCs, and TLS RFCs that details alternatives for securing data in motion.

For some protocols there are multiple modes or options of operation in the standards. Additionally, 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.

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 should be protected when using data in motion protection methods.

The controls related to encrypting transferred data are as follows:

a) TC-CNFD-G11 Providing end-to-end security protections for data in motion

When protection of data in motion is required, it should provide end-to-end protection (i.e. a method of securing communication that prevents third parties from accessing data while it is transferred from one end system or device to another).

b) TC-CNFD-G12 Compensating for computational impacts of data in motion encryption

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

10.5.4.2 Transport Layer Security (TLS)

TLS is a security protocol designed to facilitate privacy and data security for communications over the internet. TLS protocol version 1.3 is specified in IETF RFC 8446.[58]

The controls related to TLS are as follows:

TC-CNFD-R02 TLS minimum requirements

When TLS is used for data in motion protection, the implementation:

— shall use TLS for storage management such that the storage clients and servers comply with an organizationally approved TLS profile (e.g. ISO/IEC 20648 or Reference [71]);

— should use TLS version 1.3[58] or later for data access.

10.5.4.3 IP Security (IPsec)

There are several IETF RFCs that are associated with the specification of Internet Protocol Security (IPsec). IETF RFC 6071[53] provides an excellent overview of IPsec- and IKE-related RFCs.

IPsec can have detrimental impacts on the use of certain technologies like network address translation, intrusion detection system, intrusion prevention system, or other systems that look deeper into network traffic frames. Whether to rely on IPsec or other data in motion protection protocols can hinge on the trade-offs of potentially neutralizing the value of other technologies.

The controls related to IPsec are as follows:

TC-CNFD-R03 IPsec minimum requirements

When IPsec is used for data in motion protection, the implementation shall:

— support IPsec, version 3;

— support either tunnel mode or transparent mode of IPsec;

— support at least one Security Policy Database (SPD) (e.g. IETF RFC 4301[44]);

— support IPsec ESP using cryptographic algorithms (e.g. RFC 4303[45]);

— support Internet Key Exchange (IKE) version 2 (or later versions) key exchange algorithms;

— support encrypted payload using IKEv2;

— ensure that the IKE protocols perform peer authentication using RSA/ECDSA algorithm that use X.509v3 certificates that conform to the pre-shared key method (e.g. IETF RFC 4945[46]).

10.5.5 Encrypting data at rest

With increasing amounts of sensitive and regulated data being stored, organizations are taking steps to ensure this data are 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 storage media, especially tape. Consequently, encryption mechanisms within storage devices (self-encrypting drives as well as controller-based technologies), switches, specialized appliances, HBAs, etc. can provide useful protections.

The controls related to encrypting data at rest are as follows:

a) TC-CNFD-G13 Using an appropriate point of encryption

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 should be selected to address the identified risks, and arrangements made to provision that location with keying material. The data to be processed should be identified, and in some cases its location should be changed.

b) TC-CNFD-G14 Creating appropriate proof of encryption

In addition, adequate proof of encryption, which is likely to take the form of logs, should be created and integrated into the audit log infrastructure. See 10.5.3 for additional information.

c) TC-CNFD-R04 Protect keys used to encrypt storage

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 data 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 shall be protected against modification. Secret (for symmetric encryption) and private (for asymmetric or public key encryption) keys shall 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 the ISO/IEC 11770 series.

d) TC-CNFD-R05 Use appropriate encryption algorithms and modes of operation for storage

Encryption algorithms and modes of operations appropriate for storage technology shall be used and can include:

— for HDDs and SSDs, AES with the XTS mode as described in IEEE 1619-2018;[36]

— for tape, AES with the counter with cipher block chaining message authentication code (CCM) or Galois/Counter Mode (GCM) modes as described in IEEE 1619.1-2018.[37]

e) TC-CNFD-G15 Limiting plaintext exposure of plaintext keys

The amount of time a key is in plaintext form should be limited and users should be prevented from viewing plaintext keys.

f) TC-CNFD-R06 Use cryptographic keys for one purpose

Cryptographic keys shall only be used for one purpose. Do not use key-encrypting keys (also known as key wrapping keys) to encrypt data or use data encrypting keys to encrypt other keys.

g) TC-CNFD-R07 Randomly generate keys using the entire keyspace

Keys shall be randomly generated from the entire keyspace.

Best practice recommends a cryptographically secure pseudo-random number generator which ensures, 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.

h) TC-CNFD-R08 Key use limited to finite cryptoperiod or maximum amount of data processed

If appropriate, use of data encryption keys shall be limited to a finite cryptoperiod (typically no more than 2 years) or to a maximum amount of data processed.

NOTE 1 Not all keys can be replaced (e.g. a static endorsement key in a trusted platform module).

i) TC-CNFD-G16 Using centralized key management infrastructure

When possible, storage systems and infrastructure should use interoperable, centralized key management infrastructure (e.g. generate and archive encryption keys).

j) TC-CNFD-G17 Using OASIS KMIP to access and use centralized key management infrastructure

The Organization for the Advancement of Structured Information Standards (OASIS) Key Management Interoperability Protocol (KMIP) specification and profiles define the dominant mechanism for accessing centralized key management within storage infrastructures.

NOTE 2 OASIS KMIP is specified in References [75] and [76].

Storage systems and infrastructure should use clients that are OASIS KMIP Version 2.0 (or later) compliant to access and use key management infrastructure.

10.6 Storage sanitization

10.6.1 General

ISO/IEC 27002:2022, 7.10 and 7.14 provides guidance on storage media disposal and re-use to prevent leakage of information. ISO/IEC 27002:2022, 8.10 also provides related guidance on information deletion. In general, sensitive data recorded on storage media should be eliminated prior to re-use or disposal of the storage media.

To address key elements of the guidance contained in ISO/IEC 27002 on disposal and re-use, it is possible to employ a storage sanitization programme that is aligned with the data classification scheme adopted by the organization. Such a programme often includes:

— Specification of the minimum acceptable sanitization method (see 10.6.2);

— Verification steps needed to determine the adequacy of the sanitization performed (see 10.6.6);

— Identification of the records or evidence necessary to meet compliance obligations (see 10.6.7).

Storage sanitization refers to the general process of rendering previously recorded data in the storage irretrievable, such that there is reasonable assurance that the data cannot be easily retrieved or reconstructed. Storage sanitization that is performed within storage can take the form of logical sanitization (see 10.6.4) or media-based sanitization (see 10.6.3). Due to the importance of storage sanitization, there can be additional obligations to validate the outcomes of the sanitization operations as well as to produce records (evidence) of sanitization operations (see 10.6.7).

The controls related to storage sanitization are as follows:

a) TC-SNTZ-G01 Including storage sanitization as part of data governance

Storage sanitization should be an element of the organization's data governance process. The decision to use storage sanitization should be based on the organization's data classification, focusing on the data that are classified as sensitive. Common examples of sensitive data include personal data, PII, and electronic healthcare records as well as certain business data (e.g. trade secrets, intellectual property, customer records, and financial records) or mission critical data (e.g. national security). Failure to sanitize storage or to keep adequate records of sanitization operations can trigger data breach notifications when an organization loses control of storage devices or storage media.

NOTE Sensitive data in this context is data for which disclosure can have an impact on organizational mission, result in damage to organizational assets, or result in financial loss or harm to the organization or individuals.

b) TC-SNTZ-R01 Sanitize storage prior to disposal

Logical storage (see 10.6.4) or media-based storage (see 10.6.3) that has been used to record sensitive data shall be sanitized prior to disposal or transferred to a party outside of the organization (i.e. the organization loses control of the storage). Sanitization shall be used prior to internal transfers of storage within the same organization (e.g. from human resources to engineering) when the sensitivity of the data warrants confidentiality protections (e.g. compliance obligations or organizational policies).

c) TC-SNTZ-R02 Verify storage sanitization outcomes

Verification of the storage sanitization outcome shall be performed, when such sanitization is mandated, prior to the disposal or transfer of storage to ensure the organizational risks have been adequately addressed.

10.6.2 Selection of sanitization methods

Multiple sanitization methods can be used, depending on the storage (logical or media-based), and they take the form of:

— Clear: Involves the use of software or hardware to overwrite the target data with non-sensitive data.

— Purge: Involves the use of physical or logical techniques that make recovery infeasible, using state of the art laboratory techniques, while preserving the storage in a potentially reusable state.

— Destruct: Involves the use of physical techniques (e.g. disintegrate, incinerate, melt, pulverize, and shred) to destroy the storage. This sanitization method is not applicable to logical storage.

Each of the above sanitization methods provide different assurances that the data cannot be easily retrieved or reconstructed, and they are based on the level of effort an adversary expends to defeat the protections.

From this perspective, the clear method provides the least assurances and the destruct method offers the most assurances when the sanitization method is performed correctly.

Not all sanitization methods are applicable for all types of logical storage (see 10.6.4) or media-based storage (see 10.6.3).

The controls related to selecting storage sanitization methods are as follows:

a) TC-SNTZ-R03 Select minimum acceptable storage sanitization method

The selected storage sanitization method shall be specified as the minimum acceptable. A sanitization method that provides a stronger assurance level shall be permitted (e.g. when destruct is used, but clear was the minimum required). A sanitization method that provides a weaker assurance level shall not be used.

b) TC-SNTZ-G02 Considering cost and environmental impacts of storage sanitization

The selected type of storage sanitization should be assessed in terms of factors such as cost and environmental impact, and a decision should be made that best mitigates the risk to confidentiality and best satisfies other constraints imposed on the process.

10.6.3 Media-based sanitization

Data storage technologies change and evolve at a rapid pace, as do the attacks against storage. Thus, the specific sanitization techniques associated with the three sanitization methods (see 10.6.2) for the various type of storage devices or storage media are updated on a frequent basis. This document does not provide the details on how specific types of media can be sanitized, but instead, defers to other resources which provide such information.

The controls related to media-based sanitization are as follows:

a) TC-SNTZ-R04 Sanitize media in conformance with acceptable standards

When the requirement to sanitize storage devices or storage media has been established, then the sanitization shall be performed based on the selected sanitization method (clear, purge, or destruct) and in a manner that conforms with a standard which is identified as acceptable by organizational policy (e.g. IEEE 2883, which provides additional information about selecting appropriate sanitization methods for use, as well as technology-specific sanitization techniques).

When the purge method of cryptographic erase is used, see 10.6.5 for additional considerations or requirements.

b) TC-SNTZ-R05 Verify adequacy of media sanitization outcomes

Verification (see 10.6.6) of the adequacy of the sanitization outcomes shall be performed on media-based storage.

10.6.4 Logical sanitization

Many storage devices virtualize the underlying storage media and present it as logical storage (see 10.16.1). A well-known example is the logical unit on a storage array that can have a size that far exceeds the capacity of a single storage device; cloud computing storage (see 10.11) can take this virtualization to even higher levels of abstraction. The situation can be further complicated when logical storage is replicated (i.e. multiple copies of the data exist) to support server virtualization (see 10.16.2) and BCM (see 7.3). For these types of situations, it is almost impossible to identify all the underlying storage media on which sensitive data are recorded. Further, sanitizing all the physical media is often not appropriate or permissible because multiple logical storage instances can coexist on shared physical media.

The controls related to logical sanitization are as follows:

a) TC-SNTZ-R06 Sanitize logical storage in conformance with acceptable standards

If the logical storage (e.g. logical unit, filesystem, object store, or cloud computing storage) is writeable, then sanitization can be performed either using the clear method or, when encryption has been used appropriately, the purge method of cryptographic erase (see 10.6.5). When the requirement to sanitize logical storage has been established, then the sanitization shall be performed by any of the following, as appropriate:

— Clear by overwriting (replacing) all the addressable logical storage space, through provided interfaces, with known, non-sensitive data (typically zeros); or

— Purge using cryptographic erase as described in 10.6.5.

b) TC-SNTZ-R07 Verify adequacy of logical storage sanitization outcomes

Verify (see 10.6.6) the adequacy of the sanitization outcomes shall be performed on logical storage.

c) TC-SNTZ-G03 Considering additional storage sanitization for data protection mechanisms

Data protection technologies (see 10.14), which can include replication and backups, are often used in conjunction with logical storage, so separate sanitization operations should be performed on storage associated with data protection mechanisms.

10.6.5 Cryptographic erase

At a fundamental level, cryptographic erase leverages the encryption of target data by enabling sanitization of the encryption key used to encrypt the target data. This leaves only the ciphertext remaining on the storage, effectively sanitizing the data.

The controls related to cryptographic erase are as follows:

a) TC-SNTZ-R08 Use cryptographic erase for purge under correct conditions

To use cryptographic erase as a purge method, the following conditions shall be met, at a minimum:

— all data intended for cryptographic erase shall be encrypted prior to recording on the storage;

— the strength of the cryptographic algorithm (including mode of operation) used to encrypt the target data shall be at least 128 bits;

— the bits of entropy shall be at least the number of bits used by the encryption key which is used to encrypt the target data;

— all copies of the encryption keys used to encrypt the target data shall be 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.

NOTE While it can be tempting to combine cryptographic erase with another sanitization method (e.g. clear), such an approach does not improve security, but can significantly slow down the sanitization operation as well as potentially impede the ability to verify the cryptographic erase. Justifications for such an approach often include efforts to reduce the attack surface by preventing access to the ciphertext, but this simply highlights that cryptographic erase is probably not appropriate for the sensitivity level of the data.

b) TC-SNTZ-G04 Seeking assurances on the quality of cryptography used for cryptographic erase

Those who choose to apply cryptographic erase should seek either independent validation of the following assurance areas or ask the vendor to identify which mechanisms are used to ensure that these concern areas have been addressed:

— 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: sanitizing a key used to wrap (that is, encrypt) the media encryption key or another key can increase the importance of the security strength and level of assurance of the wrapping techniques used (e.g. commensurate with the level of strength of the cryptographic erase operation).

Generally accepted and (where applicable) standardized mechanisms should be used. For example, cryptographic requirements are specified in ISO/IEC 19790 and test requirements for cryptographic modules are specified in ISO/IEC 24759. These test requirements and tests cover some (but not all) of the concern areas.

c) TC-SNTZ-G05 Determining whether destroyed encryption keys are recoverable

When deciding whether to rely upon cryptographic erase, it should also be considered whether the encryption keys can be recovered either internally or externally (e.g. injected from a key management server or from a key escrow service). If the encryption key (or any key at or below the level of key sanitized during cryptographic erase) exists outside of the storage, there is a possibility that the key can be used in the future to recover data stored on the encrypted storage.

10.6.6 Verification of storage sanitization

Verification of the sanitization outcomes is an important element of a storage sanitization programme when it is necessary to determine the adequacy or effectiveness of the storage sanitization. This verification differs depending on the sanitization method. For clear or purge, the device interface is used to check the results of the sanitization operation. For destruct, physical inspection is used to check the sanitization outcomes. This verification is important because there can be errors or anomalies that necessitate additional actions to complete the sanitization or a decision on the part of the organization to accept any residual risk.

To demonstrate the importance of verification, a hypothetical scenario can be considered where the sanitization method of destruct, using a shred technique, is performed on an optical disc (PII is recorded). In addition, the resulting pieces from the shred shall be no larger than 3 millimetres by 3 millimetres in size. However, the shred of the optical disc produced pieces that are 5 millimetres by 5 millimetres. Given the possible existence of sensitive PII data on the optical disc, the organization is confronted with a decision to either accept the outcome (i.e. the organization accepts the risks with the large sized pieces) or to not accept the outcome and then to use an alternate sanitization method of destruct (e.g. incinerate, melt, or pulverize) on the pieces. In this scenario, the organization accepts the shred outcome, but documents the deviations from 3 millimetres by 3 millimetres maximum shred size.

The controls related to verification of storage sanitization are as follows:

a) TC-SNTZ-G06 Verifying clear sanitization method outcomes

For the clear sanitization method, a representative sampling for sanitization verification of the storage medium should be performed. IEEE 2883 identifies two representative sampling options, including random sampling a percentage of the user addressable space and an approach that divides the user addressable space into a predefined number of bands, which are then randomly sampled.

b) TC-SNTZ-G07 Verifying purge sanitization method outcomes

For the purge sanitization method, a full verification of the storage medium should be performed. If cryptographic erase was used to perform the sanitization, it is possible the verification cannot be performed because the storage medium is not accessible or the remaining random bit pattern on the storage medium does not provide a basis for comparison.

Storage devices that are protected with access control mechanisms have additional verification considerations; storage devices should be accessible both before and after the sanitization to enable a verification process. This can be an issue because certain purge operations can result in the target data on the storage devices not being accessible.

c) TC-SNTZ-R09 Verification of destruct sanitization method outcomes

Physical inspection is the only option when destruct is the sanitization method because the storage is unusable (by definition). When verification is required for destruct-based sanitization, the outcomes shall

be inspected and compared to a standard identified as acceptable by organizational policy to determine the adequacy (e.g. IEEE 2883). If after reviewing the verification findings associated with the destruct outcomes, a determination is made that the sanitization outcomes are not adequate, destruct-based sanitization shall be repeated with consideration given to using an alternate form of destruct.

10.6.7 Proof of sanitization

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

The controls related to proof of sanitization are as follows:

a) TC-SNTZ-G08 Producing and retaining storage sanitization records

Proof of sanitization takes on at least two forms: 1) an audit log trail and 2) a certificate of sanitization documenting the sanitization. These sanitization records are the evidence that organizations should retain for compliance/legal purposes, to prevent 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.

b) TC-SNTZ-G09 Recording minimum information for certification of sanitization

To generate a certificate of sanitization, certain information shall be gathered, before performing the sanitization if possible (e.g. destruct of an HDD). The minimum information recorded for a certificate of sanitization should include:

— manufacturer;

— model;

— serial number;

— storage media type (e.g. magnetic, flash, and hybrid);

— storage media source (i.e. user or system the storage media came from);

— sanitization method used (i.e. clear, purge, or destruct);

— description of the sanitization technique used (e.g. degauss, overwrite, block erase, and cryptographic erase);

— description of the outcome of the sanitization (e.g. success/failure and errors/anomalies);

— for sanitization verification for clear or purge:

  — tool used (including version);

  — verification method (e.g. full, and quick sampling).

— for both sanitization and validation:

  — name of person;

  — position/title of person;

  — date and time of completion;

  — location;

  — contact information (e.g. telephone number and email address);

— 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.

10.7 Direct attached storage

A direct attached storage (DAS) device is a storage device (e.g. HDD, SSD, and tape) that is directly connected to a host computer without an intervening storage network (i.e. no network device like a hub, router, or switch). 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), which can include expanders that increase the number of drives that can be connected. DAS can also include storage devices with removable media (e.g. optical disc, flash disks, Secure Digital cards) such that the storage media can be added or removed without powering down the storage device or the system. Although they are typically dedicated to the system to which they are attached, a DAS device can be shared between multiple computers, if it provides multiple interfaces (ports) that allow concurrent and direct access.

These storage devices have limited data access and management interfaces (the latter is usually in-band).

The controls related to DAS are as follows:

a) TC-DASS-G01 Protecting DAS against unauthorized access

To avoid unauthorized access of sensitive or high value data on DAS, some form of authenticated access control (drive locking) and storage device encryption should be used to protect the data at rest.

b) TC-DASS-G02 Sanitizing DAS prior to repurposing or disposing

Prior to repurposing or disposing, DAS used with sensitive or high value data should be sanitized as follows:

— use the integrated storage sanitization functionality in the storage devices (see 10.6.2);

— use computer-based or application-based sanitization.

c) TC-DASS-G03 Backing up DAS to ensure recovery after data loss

To guard against accidental or intentional data loss or corruption, backups (see 10.14.2) of the DAS contents should be made on a regular basis and properly secured in a different location.

10.8 Storage networking

10.8.1 Background

Networking plays an important role in storage infrastructures and can include common networking technologies (e.g. LAN and wide area network), 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 the ISO/IEC 27033 series is instrumental in protecting storage resources that utilize these technologies. The storage-specific network protocols and technologies are addressed in this document.

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 over a TCP/IP connection to the management interface (out-of band) of the storage system.

10.8.2 Storage area networks

10.8.2.1 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, 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 and use of separate networks), increase storage utilization and effectiveness (e.g. consolidate storage resources and tiered storage), and improve data protection and security. In addition, SANs typically play an important role in the BCM activities (see 7.3) of an organization.

Figure 4 shows an example of a single SAN that is geographically distributed. Such a SAN allows the computers in either site to access the storage resources in both sites (e.g. advantageous for localization, response and failover). In addition, the storage systems can replicate data to other storage systems independently of their location. The interconnection mechanisms between the two sites can introduce additional security considerations.

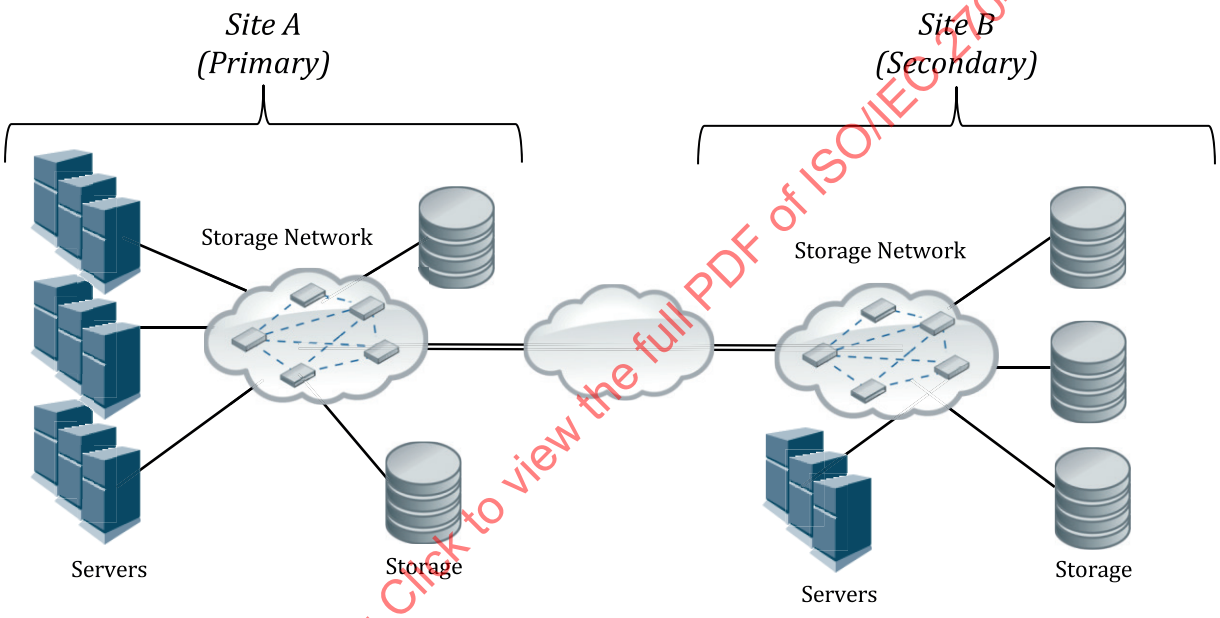

Figure 4 — Example of a storage area network

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 using different types of virtualization.

SANs are commonly based on:

— Fibre Channel (FC) technologies that utilize the Fibre Channel Protocol (FCP) for SCSI for open systems and proprietary variants for mainframes;

— Internet Small Computing System Interface (iSCSI), commonly used in small and medium-sized organization as a less expensive alternative to FC;

— InfiniBand, commonly used in high performance computing environments;

— NVMe®[3] over Fabric (NVMe-oF™[4]),[78] using different fabric transports (see 10.8.25), is used for its multitasking speed at low latency and high throughput;

— Interconnects that leverage extenders and switches take on characteristics of a SAN as well.

In addition, it is possible to move data between different SAN technologies using network gateways (see Figure 5). Such interconnective can be important for business continuity management support.

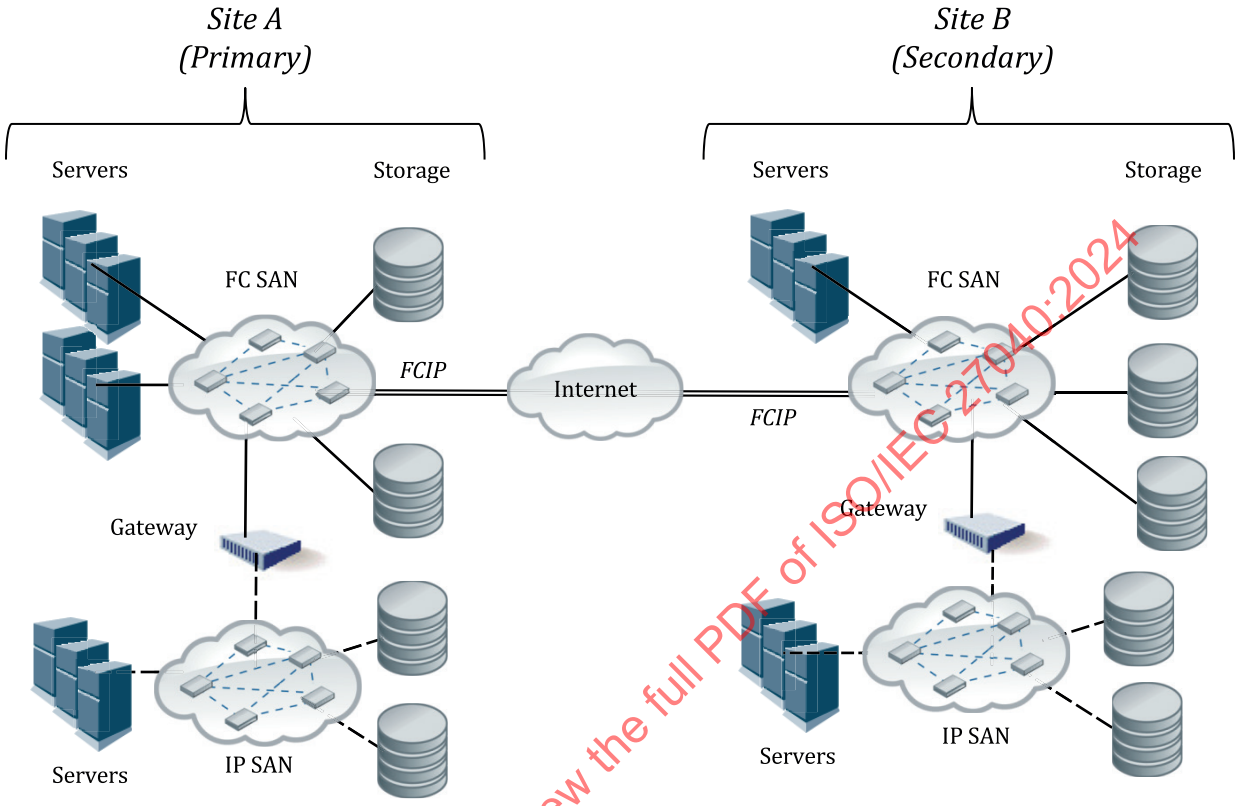

Figure 5 — Example of multi-site and multiple types of SANs

A defence in-depth strategy (see 10.2.2.1) helps to mitigate the risk associated with failure of one security control (possible single point of failure) compromising the assets under protection.

10.8.2.2 Fibre Channel SAN

Fibre Channel storage area network is a multi-gigabit-speed network technology used for block-based storage (see 10.9). There are three major Fibre Channel topologies, describing how multiple 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.

Security controls relevant to a FC SAN can be grouped into access control, authentication, and encryption (see Reference [73] for a useful summary of FC security for storage systems and ecosystems).

Access control on a SAN is implemented through application of zoning, logical unit number (LUN) masking, and port binding mechanisms:

— Port Binding: Globally unique identifiers known as worldwide names (WWN) for FC are used for identification in a SAN. Port binding is a SAN security mechanism that associates a physical port ID

3) This trade name is provided for reasons of public interest or public safety. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO.

4) This trade name is provided for reasons of public interest or public safety. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO.

and the WWN of the connected device. This association can mitigate snooping attempts by a potential adversary.

— 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 do 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.

— LUN masking and mapping: A storage device can be divided into different logical units that are identified by LUNs. 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, HBA or CNA, and virtualization layer. LUN masking refers to making a LUN visible to some servers and not visible to others.

For SANs, it is important for a switch to verify the identity of other switches in the SAN with which it communicates. If switch authentication is not implemented, a rogue switch can join a SAN and potentially compromise SAN data. Likewise, the nodes in a SAN (e.g. storage devices and servers) can employ authentication to restrict access to SAN data.

There are two major components of data confidentiality on a SAN: data in motion, and data at rest. Cryptographic protections can be used for sensitive and high-value data in SANs when data are in motion as well as when at rest on a storage device. This protection can require the use of special purpose hardware that can encrypt the data that are being sent to a storage device. Refer to 10.5.4 for additional guidance on protection of data in motion and 10.5.5 for guidance on protection of data at rest.

The controls related to Fibre Channel SANs are as follows:

a) TC-FCSS-G01 Controlling FCP node access

FCP node access should be controlled by restricting server access on the switches using techniques such as access control lists (ACLs), binding lists, and FC-SP-2[61] fabric policies.

b) TC-FCSS-G02 Using FC switch-based controls

Switch-based controls should be implemented to:

— restrict switch interconnections using techniques such as ACLs, binding lists, and FC-SP-2 fabric policies;

— establish zoning to 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 not, use stronger techniques like FC-SP-2 Zoning where supported by the vendor;

— disable unused ports;

— carefully use default zones and zone sets (assume a least privilege posture).

c) TC-FCSS-G03 Configuring FC device to meet security requirements

When configuring switches, extenders, routers, and gateways (e.g. FCIP) to interconnect storage networks, configurations should meet security requirements.

10.9.1 provides guidance on block-based Fibre Channel storage.

10.8.2.3 IP SAN

Internet SCSI storage area network (iSCSI), which is described in IETF RFC 7143,[55] is a connection-oriented command/response protocol that runs over TCP, and is used for network access to disk, tape, and other devices.

The controls related to IP SANs are as follows:

a) TC-IPSS-G01 Using iSCSI network access and protocols

iSCSI network access and protocols should be controlled by:

— segregating iSCSI interfaces from general purpose LANs to provide security and better performance;

— using virtual local area networks (VLANs) when the use of physically isolated LANs is not an option.

Fibre Channel over IP (FCIP), defined in IETF RFC 3821,[43] 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 SAN.

b) TC-IPSS-G02 Using FCIP network access and protocols

FCIP network access and protocols should be controlled by:

— setting up the peer-to-peer relationship between FCIP entities, recognizing that the security policies are applied uniformly;

— using a private IP network, used exclusively by the FCIP entities, whenever possible.

c) TC-IPSS-G03 Using IPsec to secure FCIP

IPsec security measures (see 10.5.4.3) should be implemented in conjunction with FCIP by:

— performing cryptographic authentication and data integrity at a minimum;

— protecting sensitive data by appropriate confidentiality measures.

IETF RFC 3723[42] provides additional useful information on both iSCSI and FCIP. 10.9.2 provides guidance on block-based IP storage. In addition, IETF RFC 7146[56] provides important security updates to IETF RFC 3723 and IETF RFC 3821.[43]

10.8.2.4 InfiniBand

InfiniBand is a low-latency, high-bandwidth interconnect which requires low processing overhead and is ideal to carry multiple traffic types (clustering, communications, storage, management) over a single connection. InfiniBand is a switch-based point-to-point interconnect architecture that operates both on the printed circuit board as a component-to-component interconnect, as well as a chassis-to-chassis interconnect.

The InfiniBand architecture defines multiple devices for system communication: a channel adapter (see NOTE 1), switch, router, and a subnet manager (see NOTE 2). Within a subnet, InfiniBand requires that there be at least one channel adapter for each end node and a subnet manager to set up and maintain the link. In addition, InfiniBand requires all channel adapters and switches to contain a subnet management agent for handling communication with the subnet manager.

NOTE 1 A channel adapter connects InfiniBand to other devices. There are two types of channel adapters, a host channel adapter and a target channel adapter.

NOTE 2 The subnet manager configures the local subnet and ensures its continued operation. InfiniBand requires that there be at least one subnet manager present in the subnet to manage all switch and router setups and for subnet reconfiguration when a link goes down or a new link comes up.

The controls related to InfiniBand SANs are as follows:

TC-IBSS-G01 protecting InfiniBand SANs

InfiniBand SANs should be protected by:

— keeping all IB hosts attached to IB fabric secure because an IB fabric is only as secure as the least secure IB host attached to it;

— maintaining physical security because an attacker able to connect rogue host to an IB switch can be used to compromise the security of the IB fabric.

10.8.2.5 NVMe over Fabrics

NVM Express (NVMe)[78] is used by a processor to communicate with non-volatile memory across a PCI Express (PCIe) bus,[80] including NVMe-attached solid state drives (SSDs) in many form factors. NVMe is designed to take advantage of the low latency and internal parallelism of faster media, such as SSDs and flash memory-based technologies.

NVMe over Fabric (NVMe-oF™)[78] is a communication protocol for exposing NVMe targets from a remote system to a client via encapsulating NVMe Admin and input/output commands[79] over a variety of transports called fabrics. These fabrics can be memory-based transports, e.g. remote direct memory access (RDMA), and message-based transports, e.g. TCP and FC. NVMe-oF specifies three families of fabrics:

— NVMe over Fibre Channel (NVMe/FC);[62]

— NVMe over TCP (NVMe/TCP);[81]

— NVMe over RDMA (NVMe/RDMA),[82] including:

  — NVMe over RDMA over converged Ethernet (RoCE);

  — NVMe over InfiniBand;

  — NVMe over iWARP (over traditional Ethernet).

The controls related to NVMe-oF are as follows:

a) TC-NVSS-G01 Using NVMe-oF authentication

NVMe-oF[78] supports two high-level types of authentications (fabric secure channel and in-band authentication). One or both authentication mechanisms should be used.

b) TC-NVSS-G02 Using NVMe/FC security controls

Security controls for NVMe/FC can include some of the controls used with FC SAN controls (see 10.8.2.2). For NVMe/FC, the following should be used:

— adequate physical security, including segregated networks;

— FC zoning to partition the FC fabrics;

— LUN masking on the storage to restrict access to specific LUN;

— authentication of Fibre Channel devices, secure key exchange, and secure (i.e. encrypted) communication between Fibre Channel devices.

NOTE The NVMe-oF[78] standard defers NVMe/FC specific security concerns to the FC-SP-2[61] technical standard, which includes protocols to enhance Fibre Channel security in several areas.

c) TC-NVSS-G03 Using NVMe/TCP security controls

Security controls for NVMe/TCP should be used, including:

— authentication only (requires authentication secrets provisioning, one per entity) with DH-HMAC-CHAP;

— secure channel concatenated to an authentication transaction (in which the authentication transaction generates on the fly an ephemeral pre-shared key to be used by the secure channel protocol);

— secure channel alone (requires provisioning a pre-shared key for each pair of entities allowed to communicate);

— appropriate communications security with NVMe-oF solutions such that:

  — TLS is used and no version of the Security Socket Layer (SSL) Protocol is used;

  — NVMe-oF solutions based on NVMe-oF Revision 1.1 do not use TLS protocol versions less than 1.2;

  — NVMe-oF solutions based on the NVMe Base Revision 2.0 specification[78] and the NVMe over TCP Revision 1.0[81] specification do not use TLS protocol versions less than 1.3.

d) TC-NVSS-G04 Using NVMe/RDMA security controls

Security Controls for NVMe/RDMA should use in-band authentication provided by DH-HMAC-CHAP.

10.8.3 Network Attached Storage protocols

10.8.3.1 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 built 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 6). In addition, NAS systems can use one or more SANs to store data locally or out of region (shown as FC SAN-1 and FC SAN-2 in Figure 6).

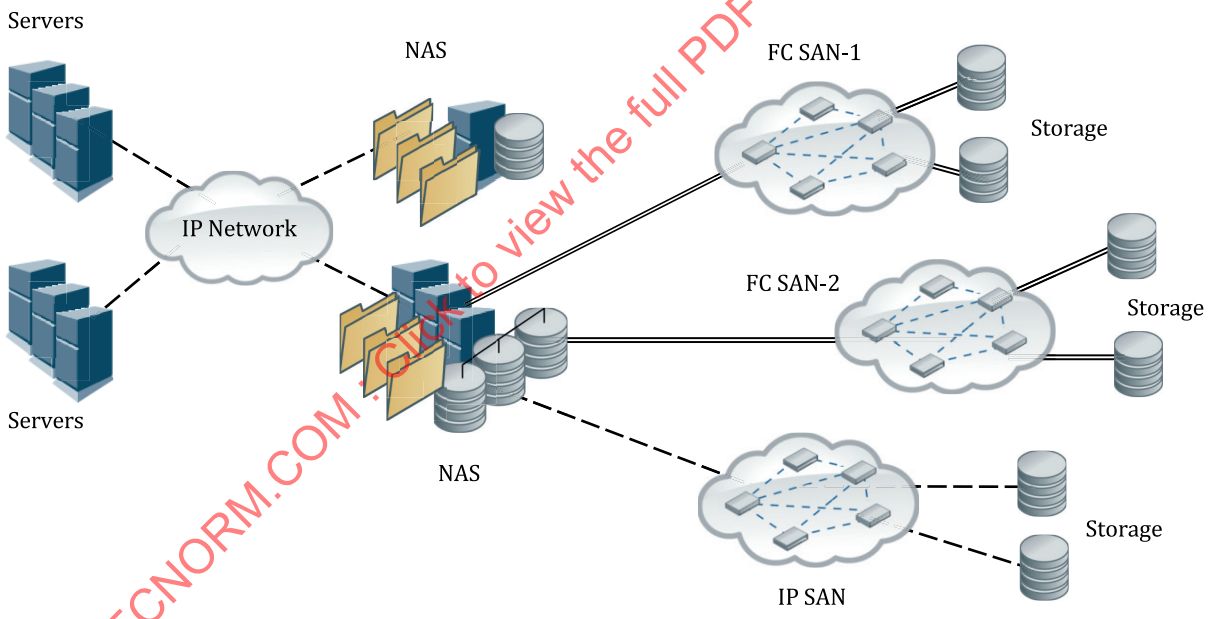

Figure 6 — Example of network attached storage

Common file system implementations include the Network File System (NFS) and Server Message Block (SMB) Common Internet File System, in addition to other technologies such as the Object-based Storage and cloud computing storage. These protocols can be configured to help secure NAS.

Refer to 10.10 for additional controls for NAS and file-based storage.

10.8.3.2 Network File System

Network file system (NFS) is a client/server application, communicating with a remote procedure call-based protocol. Multiple versions of NFS are specified and in use, including NFS version 3 (specified in IETF RFC 1813[40]), NFS version 4 (specified in IETF RFC 7530[57]), and NFS version 4.1 (specified in IETF RFC

8881[59]). From a security perspective, NFS version 3 (NFSv3) is considered less secure and extra care is required when it is used with sensitive or high-value data.

The controls related to NFS are as follows:

a) TC-NASP-G01 Using NFS network access and protocols

NFS network access and protocols should be controlled by:

— enabling NFS only if required to eliminate it as a possible attack vector available to an adversary;

— using NFSv4 (or later versions) whenever possible and limit NFSv3 usage;

— filtering client and management access by IP address for additional security.

b) TC-NASP-G02 Using encryption to secure NFS

NFS client data access should use encryption (e.g. IPsec or TLS).

10.8.3.3 Server Message Block (SMB)

SMB 3 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 controls related to SMB are as follows:

TC-NASP-G03 Using SMB network access and protocols

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

— use later versions of the SMB protocol;

— turn off low-security session negotiation protocols and use Kerberos instead;

— maintain up-to-date patch levels;

— use SMB signing;

— maintain directory services securely;

— use one-way trusts, from leaf domains to parent domains, when possible;

— control SMB network access and protocols by:

  — enabling SMB only if necessary. This eliminates it as a possible attack vector available to an adversary;

  — encrypting client data access when necessary.

10.9 Block-based storage

10.9.1 Fibre Channel (FC) storage

Fibre Channel storage systems use specialized networking (see 10.8.2.2) to present block-based storage resources to computers. These resources usually take the form of logical units (logical storage) and tape devices (including virtual tape).

The controls related to FC storage are as follows:

a) TC-BBFC-G01 Using FC LUN masking and mapping

LUN masking and mapping (worldwide port name filtering) as well as other access control mechanisms should be used to restrict access to storage.

b) TC-BBFC-G02 Using FCP for SCSI security measures

FCP for SCSI security measures should be used, including:

— mutual authentication using FC-SP-2 AUTH-A[61] with all servers and switches, leveraging centralized authentication services when possible;

— encryption of Fibre Channel connections that leave the protected area (e.g. confines of a physically controlled data centre) with ESP_Header.

NOTE 1 Fibre Channel frame integrity or confidentiality can be provided with ESP_Header optional headers, which are defined in Reference [60].

c) TC-BBFC-G03 Using data at rest encryption for FC storage

Data at rest encryption measures (see 10.5.5) should be used to:

— protect the confidentiality of sensitive or high-value data recorded on FC storage devices or media;

NOTE 2 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.

— facilitate rapid elimination of data on FC storage with cryptographic erase (see 10.6.5).

d) TC-BBFC-G04 Using storage sanitization for FC storage

Storage sanitization measures (see 10.6) in the form of:

— media-aligned sanitization (see 10.6.3) should be used for FC storage media and storage devices for sensitive and regulated data;

— logical sanitization (see 10.6.4) should be used for virtualized FC storage (see 10.16.1), especially when the actual storage devices and media cannot be determined.

10.9.2 IP storage

Unlike FC storage, IP storage uses TCP/IP networking (see 10.8.2.3), specifically iSCSI, to present block-based storage resources to computers.

The controls related to IP storage are as follows:

a) TC-BBIP-G01 Filtering iSCSI initiator access

iSCSI initiator access should be controlled by filtering based on source IP addresses and protocols.

b) TC-BBIP-G02 Using iSCSI security measures

iSCSI security measures should be used, including:

— bidirectional CHAP authentication, using random challenges (i.e. not repeated), for both initiators and targets;

— IPsec to secure the communication channel when sensitive or high-value data can be exposed (see 10.5.4.3);

— Internet Storage Name Service (iSNS), SLP, DNS infrastructure with appropriate security controls to avoid indirect attacks.

c) TC-BBIP-G03 Using data at rest encryption for IP storage

Data at rest encryption measures (see 10.5.5) should be used to:

— protect the confidentiality of sensitive or high-value data recorded on IP storage devices or media;

NOTE 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.

— facilitate rapid elimination of data on IP storage with cryptographic erase (see 10.6.5).

d) TC-BBIP-G04 Using storage sanitization for IP storage

Storage sanitization measures (see 10.6) in the form of:

— media-aligned sanitization (see 10.6.3) should be used for IP storage media and storage devices for sensitive and regulated data;

— logical sanitization (see 10.6.4) should be used for virtualized IP storage (see 10.16.1), especially when the actual storage devices and media cannot be determined.

10.10 File-based storage

10.10.1 General

Security controls relevant to file-based storage (typically 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 at rest;

— authentication controls, such as Kerberos, for verifying the identity of users attempting to access NAS data.

10.10.2 NFS-based NAS

This type of storage is basically a LAN-attached file server that presents files using the network protocol, NFS (see 10.8.3.2). It consists of an engine that implements the file services and one or more storage devices, on which data are stored. A NAS system can also be node-attached, in which case the NAS system is treated just like any other server on the SAN (e.g. provided access to storage and LAN-free backups). 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.

The controls related to NFS-based NAS are as follows:

a) TC-FBNF-R01 Apply NFS access controls

Access controls shall be applied to NFS exported filesystems to:

— 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;

— when appropriate, use Kerberos Safe and Private modes to sign and encrypt NFS traffic.

b) TC-FBNF-R02 Restrict NFS client behaviours

NFS client behaviours shall be restricted to:

— filter client access to NFS shares whenever possible;

— disallow NFS clients to run programs that provide temporary elevated permissions during execution on exported file systems.

c) TC-FBNF-G01 Securing data on NFS servers

Data on NFS servers should be secured by:

— ensuring exported file systems are in their own partitions to prevent system degradation by an attacker writing to an exported file system until it is full;

— encrypting data at rest when necessary;

— disallowing NFS exports of administrative file systems;

— guarding against malware (e.g. viruses, worms, and rootkits);

— continually monitoring content placed in NFS shares and relevant access controls.

10.10.3 SMB-based NAS

Like NFS-based NAS (see 10.10.2), SMB-based NAS is a LAN-attached file server that serves files, but it differs in its use of the network protocols, SMB (see 10.8.3.3).

The controls related to SMB-based NAS are as follows:

a) TC-FBSM-R01 Minimum acceptable SMB protocol

For SMB-based NAS, SMB version 3 or later shall be used and all previous versions shall be disabled.

b) TC-FBSM-R02 Apply SMB access controls

Access controls shall be applied to SMB exported filesystems to:

— disable unauthenticated access to SMB shares and NAS devices (i.e. restrict Anonymous access);

— disable Guest and Everyone access to all SMB shares;

— implement authentication and access control via a centralized mechanism (RADIUS, Lightweight Directory Access Protocol).

c) TC-FBSM-R03 Restrict SMB client behaviours

SMB client behaviours shall be restricted by enabling SMB signing for clients and the NAS device.

d) TC-FBSM-G01 Securing data on SMB servers

Data on SMB servers should be secured by:

— enabling SMB auditing whenever possible;

— continually reviewing content placed in SMB shares and relevant access controls;

— encrypting data at rest when necessary;

— guarding against malware (e.g. viruses, worms, and rootkits);

— using SMB with strong authentication (Kerberos).

10.11 Cloud computing storage

10.11.1 Securing cloud computing storage

Both proprietary and standards-based, cloud computing storage offerings are in use and commonly provide copy capabilities (e.g. mirroring 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. allowing 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 assurance that the relevant security threats and challenges have been addressed (see Reference [77] for a useful summary of cloud computing security threats and challenges).

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 can be specified, but there can be significant difference in terms of what is implemented versus what ultimately gets used.

The controls related to securing cloud computing storage are as follows:

a) TC-CCSS-G01 Using transport security for cloud transactions

Transport security such as IPsec or Transport Layer Security (TLS) should be used for all transactions (see 10.5.4).

b) TC-CCSS-G02 Using data at rest encryption for cloud storage

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, and adversaries).

c) TC-CCSS-G03 Using strong authentication to gain access to cloud storage

User registrations should be secured and strong authentication should be used to protect access to data.

d) TC-CCSS-G04 Using access controls to protect data on cloud storage

It is recommended to use access controls that guard against unauthorized access from other tenants while providing appropriate access privileges to users who are permitted to access the data.

e) TC-CCSS-G05 Using storage sanitization for cloud storage

The provided storage sanitization capabilities should be used to eliminate sensitive data from the cloud computing storage.

NOTE In some jurisdictions, privacy requirements like the right of erasure or the right to be forgotten can require additional security controls.

Cloud computing implementations often leverage different forms of virtualization, so the guidance in 10.16 can also be relevant.

10.11.2 CDMI security

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

The controls related to CDMI security are as follows:

a) TC-CDMI-R01 Use TLS for all CDMI transactions

Transport Layer Security (TLS) shall be used for all CDMI transactions (see 10.5.4.2).

b) TC-CDMI-R02 Mutually authenticate all CDMI entities

CDMI entities (certificates for servers and HTTP basic authentication for clients) shall be authenticated.

c) TC-CDMI-G01 Using CDMI capability queries to assess security adequacy

The adequacy of the security capabilities of the cloud service provider's CDMI implementation should be determined with CDMI capability queries to identify the offered security functionality.

d) TC-CDMI-G02 Using CDMI domains

CDMI domains should be used to provide a place for authentication mappings to external authentication providers.

e) TC-CDMI-G03 Integrating CDMI logging into audit logs

CDMI security logging should be enabled and the security event data contained in the appropriate logging queues should be retrieved on a regular and timely basis.

f) TC-CDMI-G04 Configuring CDMI retention to align auto-deletions with policy

The automatic deletion capability (CDMI deletion) should be aligned with the organization's data retention policy.

g) TC-CDMI-G05 Verifying the ability to lift CDMI holds prior to using the capability

Prior to using CDMI holds, the process and mechanism for lifting the CDMI hold should be understood.

h) TC-CDMI-G06 Using data at rest encryption for CDMI storage

Data at rest encryption measures should be used to protect sensitive and high-value data.

i) TC-CDMI-G07 Using storage sanitization for CDMI storage

The provided storage sanitization should be used to eliminate sensitive data from the cloud service provider's storage.

10.12 Object-based storage

Object-based storage is a data storage architecture for handling large amounts of unstructured data. Discrete units of data (objects) are stored in a structurally flat data environment. Each object is a simple, self-contained repository that includes the data, metadata (descriptive, customizable information associated with an object), and a unique identifying ID number (instead of a file name and file path). The metadata or ID can be used by an application to locate and access the objects.

Objects (data) in an object-storage system are accessed via APIs. The native API for object storage is typically an HTTP-based RESTful API (also known as a RESTful web service).

The controls related to object-based storage are as follows:

a) TC-OBSS-G01 Using transport security for object-based storage transactions

Transport security such as IPsec or Transport Layer Security (TLS) should be used for all object-based storage transactions (see 10.5.4).

b) TC-OBSS-G02 Using data at rest encryption for object-based storage

Data at rest encryption (at the object-level or tenant-level) should be used to prevent access by unauthorized parties (e.g. service provider personnel and adversaries).

c) TC-OBSS-G03 Enabling data immutability for object-based storage

Data objects should have appropriate data immutability protections enabled to guard against malicious accidental deletions or encryption (e.g. ransomware).

d) TC-OBSS-G04 Aligning security mechanisms with tenants on object-based storage

Security mechanism (e.g. authentication and key management) should be aligned with tenants on the object-based storage.

e) TC-OBSS-G05 Using storage sanitization for object-based storage

The provided storage sanitization should be used to eliminate sensitive data from the object-based storage.

10.13 Data reductions

As a routine course of business, organizations often 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 algorithm 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.

NOTE 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 can result in identical encoded/compressed data if the same starting conditions (e.g. history buffer) exist.

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 BCM support. Data compression is typically performed in hardware so some care is required to ensure the encoded data can be decoded later (e.g. when a tape is read by a different tape drive or when the compressed data are 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 mechanisms. However, their presence can necessitate adjustments to storage security activities.

The controls related to data reductions are as follows:

a) TC-DRDC-G01 Using compression before encryption

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 decompression).

b) TC-DRDC-G02 Using deduplication before encryption

When encryption is used along with deduplication, the deduplication should be applied before the encryption prior to recording data because deduplication is often not effective on ciphertext; when reading data the reverse order should be used.

c) TC-DRDC-G03 Using correct order for multiple data reductions and encryption

When both compression and deduplication are used along with encryption, the order of use should be:

— deduplication, compression, and encryption prior to recording data (with the reverse order used when reading data); or

— compression, deduplication, and encryption prior to recording data (with the reverse order used when reading data).

d) TC-DRDC-G04 Using data reductions compatible with BCM

Compression or deduplication can impact BCM implementations, so they should be factored into the design, documentation, and testing of BCM solutions.

10.14 Data protection and recovery

10.14.1 General

From a storage perspective, data protection is the process of safeguarding important data from corruption or loss. Data recovery is the process of restoring data that has been lost, accidentally deleted, corrupted or made inaccessible. Organizations that employ data protection mechanisms and processes typically implement them in a way that allows for rapid and complete recovery of data (see Reference [72] for a useful summary of data protection with storage systems and ecosystems). Owing to the increased dependency on data availability and integrity, many organizations employ a range of data protection mechanisms like backups (see 10.14.2) and replications (see 10.14.3) 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 data protection solutions can be viewed as data recovery mechanisms.

TC-PROT-G01 Designing data protection mechanisms for quick recoveries

Data protection mechanisms (like backups and replication) should be designed with quick recoveries in mind, rather than just preservation of data. Specific recovery time objectives and recovery point objectives should be included in these designs.

10.14.2 Storage backups

While data loss due to technology failures are less common because of more reliable hardware and software, a greater number of threats have emerged from viruses and ransomware. This risk of data loss has forced organizations to use a technology known as storage backups. These backups typically take the form of a:

— full backup, which copies all files in a directory to storage media regardless of previous backups;

— incremental backup, which only copies those files in a directory that have changed since the last backup (full or incremental), to storage media;

— differential backup, which copies all files in a directory that have changed since the last full backup, to storage media.

A major challenge with traditional backups (full, incremental, or differential) is that they are associated with moving large chunks of data, such as terabytes or petabytes. Moving large data sets is a time-consuming process and can eventually result in exceeding the recovery time required by the organization.

ISO/IEC 27002:2022, 8.13 provides useful on information backups.

The controls related to storage backups are as follows:

a) TC-PROT-G02 Using data backup measures and operations securely

Backup security should:

— 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.

b) TC-PROT-G03 Using cyber-attack recovery backups

While many organizations implement backups for disaster recovery and business continuity (non-malicious) purposes, some organization are also performing special backups to assist with cyber-attack recovery (e.g. ransomware attacks). In addition to the normal backup security measures, cyber-attack recovery backups should:

— not be accessible to regular IT staff, but rather, a very narrow group of specifically authorized personnel;

— be retained for much longer periods of time to allow for investigations and deal with attacks that have remained latent until triggered much later;

— be stored off-site and separate from where production storage backups are stored;

— use independent, full backup copies on a regular basis (i.e. have no dependencies on production baseline data);

— be restored onto an isolated staging (or air-gapped) environment rather than directly onto the target hosts or application;

— use immutable storage.

10.14.3 Storage replication

Replication involves making a copy of the production data ready to be used immediately with no further movement or changes. Replication is a complex and costly process, and it is usable in limited scenarios. Replication has become a leader in disaster recovery solutions since it can offer zero data loss. However, replication does not protect from deletion, corruption, or ransomware event. Hence, despite its efficiency, replication is not a substitute for backup.

The controls related to storage replication are as follows:

TC-PROT-G04 Using data replication measures and operations securely

Replication security should:

— 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).

10.14.4 Storage snapshots

A storage snapshot can be of a storage volume, file or database as they appeared at a given point in time. In the event of a failure, users can restore their data from the most recent snapshot (or any other snapshot previously saved) before the failure. Unlike backups that involve making copies of data, snapshots maintain pointers to the data and record changes to this data (deletions, additions, moves, etc.). Snapshots typically do not take up much storage space individually, but their total volume can grow, especially if there are many deleted blocks/files, so suppliers usually limit the number of snapshots that can be retained.

The controls related to storage snapshots are as follows:

TC-PROT-G05 Using snapshots in conjunction with backups

Snapshots should be used in conjunction with a backup strategy to provide more frequent protection, measured in minutes or hours, while backups are used for daily protection. The interval of the snapshots should be based on the granularity requirements for restoring from a specific point in time. Snapshot retention periods should allow for one or two backups to have taken place in that period (i.e. before a snapshot is deleted).

TC-PROT-G06 Using snapshot security

Snapshot security should:

— ensure that the snapshot approach, especially for business/mission critical data, is aligned with its associated recovery strategy;

— ensure that snapshots are protected from changes (e.g. read only);

— ensure that the snapshot approach provides adequate protections against unauthorized access (e.g. encryption).

10.15 Data archives and repositories

10.15.1 General

While the argument can be made that a data archive is simply a special type of a data repository, there is often a difference of emphasis between a repository and an archive. Both serve access and preservation functions, but repositories normally emphasize access, while archives normally emphasize preservation. From a storage and security perspective, this difference has major implications on the underlying technology and controls needed to protect the data.

10.15.2 Data archives

10.15.2.1 General

ISO 14721 points out that the term “archive” has come to be used to refer to a wide variety of storage and preservation functions and systems. Further to that, traditional archives are understood as facilities or organizations which preserve records, originally generated by or for a government organization, institution, or corporation, for access by public or private communities. The archive accomplishes this task by taking ownership of the records, ensuring that they are understandable to the accessing community, and managing them so as to preserve their information content and authenticity.

An archive is a collection of data objects that represent an official working copy of the data, but is managed separately from more active production data, for such purposes as long-term preservation and better cost economics. Further, archives are often used for storing data sets that have specific compliance obligations, and they are normally used for auditing or analysis rather than for application recovery. In addition, the retention requirements can vary (e.g. short, medium and long-term), but the archive should ensure proper integrity, immutability, authenticity, confidentiality and provenance.

ISO/TR 18492 defines long-term preservation as the period of time, ranging between a few years to hundreds of years, that electronic document-based information is maintained as accessible and authentic evidence. Retention periods are often determined by regulatory compliance, legal requirements and business needs, so they can vary significantly from one organization to another.

The controls related to storage archives are as follows:

a) TC-DARS-G01 Addressing key preservation issues in storage archives

Organizations should consider and address the following key issues for storage archives contained in ISO/TR 18492 when developing a long-term preservation strategy.

— Readable electronic document-based information. The bit stream comprising electronic document-based information should be accessible on the computer system or device that initially created it, currently stores it, currently accesses it, or can be used to store it in the future; media obsolescence and data formatting are also considerations.

— Intelligible electronic document-based information. The intelligibility of electronic document-based information is a function of information concerning what the bit stream actually represents and the processing software’s capacity to take appropriate action based on this information.

— Identifiable electronic document-based information. Document-based information should be organized, classified and described in such a way that users and information systems can distinguish between information objects based upon a unique attribute such as name or ID number. Facilitating search and retrieval is also a consideration.

— Retrievable document-based information. Discrete information objects (or parts of them) can be retrieved and displayed. Retrievability is typically software-dependent in that it requires keys or pointers that link the logical structure of information objects (e.g. data fields or text strings) to their physical storage location.

— Understandable document-based information. Conveying information to both computers and users beyond the document contents, including context of creation and use (i.e. metadata) as well as relationships among other documents.

— Authentic electronic document-based information. Ensure the information is what it purports to be, that is, information which, over time, has not been altered, changed or otherwise corrupted.

b) TC-DARS-G02 Employing security services to address evidentiary aspects of archives

Archives that have potential evidentiary relevance should address data authenticity, provenance and chain of custody aspects, including retaining, protecting, and maintaining significant amounts of metadata. Organizations should use the following security services identified by ISO 14721 for both the data and metadata.

— Identification/authentication service confirms the identities of requesters for use of information system resources. In addition, authentication can apply to providers of data. The authentication service should occur at the initiation of a session or during a session.

— Access control service prevents the unauthorized use of information system resources. This service also prevents the use of a resource in an unauthorized way. This service should be applied to various aspects of access to a resource (e.g. access to communications to the resource, the reading, writing or deletion of an information/data resource, the execution of a processing resource) or to all accesses to a resource.

— Data integrity service ensures that data are not altered or destroyed in an unauthorized manner. This service applies to data in permanent data stores and to data in communications messages.

— Data confidentiality service ensures that data are not made available or disclosed to unauthorized individuals or computer processes. This service can be applied to devices that permit interaction between the user and the information system. In addition, this service can ensure that observation of usage patterns of communications resources is not possible.

— Non-repudiation service ensures that entities engaging in an information exchange cannot deny being involved in it. This service can take one or both of two forms. First, the recipient of data are provided with proof of the origin of the data. This protects against any attempt by the sender to falsely deny sending the data or its contents. Second, the sender of data are provided with proof of delivery of data. This protects against any subsequent attempt by the recipient to falsely deny receiving the data or its contents.

These security services should be applied during the storage and transfer of the data and metadata to and from the archive. Equally important, care should be exercised when the security services/controls are adjusted/replaced to avoid exposing the archived data to attack or disclosure (i.e. risk).

In many of standards and publications, privacy is often not directly addressed in the context of archives. However, with the increase in privacy (protection of PII) regulations around the world, this is an important aspect to address.

Provenance and authenticity are essential elements of most archives, which means that proper metadata handling is required. Chain of custody measures can be necessary as well to address evidentiary requirements, which can complicate the nature of the archive solutions used (e.g. cloud computing-based storage can be unable to provide the needed details).

Many archives are concerned with "proving" data have not been changed (authenticity), using integrity verification approaches. An alternative strategy is to employ immutability measures (e.g. WORM storage) to prevent changes.

10.15.2.2 Short to medium-term archives

A large number of organizations must 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.

Successful retention and preservation of data over short to medium-term retention periods can involve the use of data protection, business continuity management, and digital preservation and curation practices. Specific measures are typically commensurate with the value of the data 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 controls related to short to medium-term archives are as follows:

a) TC-DARS-G03 Making multiple physical or logical replicas of data retained in archives

Multiple physical or logical replicas of the data should be created and preserved; archived data should have no dependency on production or disaster recovery baseline data. The replicas should 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.

NOTE The important aspect to consider is the quality and characteristics of the digital archive process, rather than how many copies there are.

b) TC-DARS-G04 Performing regular integrity audits of data retained in archives

Integrity audits should be performed on a defined schedule, looking for both obvious and latent faults (e.g. integrity checks) and the damage they cause. Corrupted data should be repaired using the good data from other replicas before that damage spreads.

c) TC-DARS-G05 Access controls for data retained in archives matched to legal and regulatory requirements

As part of addressing legal and regulatory obligations (e.g. protecting PII) associated with data access, the access control scheme should be sufficiently robust to guard against inappropriate access even as the obligations change over time.

d) TC-DARS-G06 Using accountability and traceability measures for data access of archives

Accountability and traceability measures should be implemented and periodically checked to determine whether they are adequate and functional. All data accesses of sensitive or high-value data should be recorded in audit log entries.

e) TC-DARS-G07 Using mechanisms to address data authenticity, provenance, and chain of custody for data retained in archives

Mechanisms to address data authenticity, provenance, and chain of custody, especially for data of an evidentiary nature, should be implemented.

f) TC-DARS-G08 Ensuring appropriate key lifecycle management for data retained in archives

If encryption is used, the keys and keying material should be archived or escrowed. The data should be rekeyed within recommended cryptoperiods or when the underlying cryptographic algorithm is being replaced.

10.15.2.3 Long-term archives

Due to the rather short lifetime and limited reliability of traditional storage components, data can become corrupted as the storage 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). It is addressed in the following standards, which are applicable to storage infrastructure:

— ISO/TR 10255;

— ISO/TR 18492;

— ISO 16175-1;

— ISO/TS 16175-2.

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 can have several decades of time to conduct an attack (slow attack).

The controls associated with short to medium-term archives (see 10.15.2.2) apply to long-term archives. Additional controls specific to long-term archive can apply as well.

The controls related to long-term archives are as follows:

a) TC-DARS-G09 Actively checking data integrity of long-term archival storage

Archival storage assumes a write-once, limited read access pattern, thus the integrity of the data in the system should be actively checked at regular intervals rather than waiting until it is read.

b) TC-DARS-G10 Upgrading security during technology refreshes of long-term archival storage

When migrating archival data to newer storage technologies, security capabilities that offer enhanced security measures to better secure the data in its new location should also be upgraded.

c) TC-DARS-G11 Managing users and access to long-term archival storage

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 and data attached to existing users.

d) TC-DARS-G12 Maintaining the data confidentiality measures while retained in the long-term archival storage

Secrecy mechanisms (e.g. encryption and key sharing) should function in the complete absence of the user that stored the data (e.g. a new user who is given rights to read data should also be given the ability to decrypt the data).

e) TC-DARS-G13 Retaining security logs for the long-term archival storage

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.

f) TC-DARS-G14 Detecting and addressing compromises of long-term archival storage

The archive system should either immediately deal with any compromise or maintain a history of compromises in order to intelligently schedule corrective action.

g) TC-DARS-G15 Ensuring that data reduction technologies do not impact data integrity of long-term archival storage