NIST Special Publication 800
NIST SP 800-209r1 ipd

Security Guidelines for Storage Infrastructure

Initial Public Draft

Ramaswamy Chandramouli
Eric A. Hibbard

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-209r1.ipd

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST Special Publication 800
NIST SP 800-209r1 ipd

Security Guidelines for Storage Infrastructure

Initial Public Draft

Ramaswamy Chandramouli
Computer Security Division
Information Technology Laboratory

Eric A. Hibbard
Samsung Semiconductor, Inc.

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-209r1.ipd

July 2026

U.S. Department of Commerce
Howard Lutnick, Secretary of Commerce

National Institute of Standards and Technology
Arvind Raman, NIST Director and Under Secretary of Commerce for Standards and Technology

Certain equipment, instruments, software, or materials, commercial or non-commercial, are identified in this paper in order to specify the experimental procedure adequately. Such identification does not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at https://csrc.nist.gov/publications.

Authority

This publication has been developed by NIST in accordance with its statutory responsibilities under the Federal Information Security Modernization Act (FISMA) of 2014, 44 U.S.C. § 3551 et seq., Public Law (P.L.) 113-283. NIST is responsible for developing information security standards and guidelines, including minimum requirements for federal information systems, but such standards and guidelines shall not apply to national security systems without the express approval of appropriate federal officials exercising policy authority over such systems. This guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130.

Nothing in this publication should be taken to contradict the standards and guidelines made mandatory and binding on federal agencies by the Secretary of Commerce under statutory authority. Nor should these guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or any other federal official. This publication may be used by nongovernmental organizations on a voluntary basis and is not subject to copyright in the United States. Attribution would, however, be appreciated by NIST.

NIST Technical Series Policies

Copyright, Use, and Licensing Statements
NIST Technical Series Publication Identifier Syntax

Publication History

Approved by the NIST Editorial Review Board on YYYY-MM-DD [Will be added to final publication.]
Supersedes NIST Series XXX (Month Year) DOI [Will be added to final publication, if applicable.]

How to Cite this NIST Technical Series Publication

Chandramouli R, Hibbard EA (2026) Security Guidelines for Storage Infrastructure. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-209r1 ipd. https://doi.org/10.6028/NIST.SP.800-209r1.ipd

Author ORCID iDs

Ramaswamy Chandramouli: 0000-0002-7387-5858
Eric Hibbard: 0009-0001-8112-1263

Public Comment Period
July 22, 2026 – September 8, 2026

Submit Comments
sp800-209r1-comments@nist.gov

National Institute of Standards and Technology
Attn: Computer Security Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930

Additional Information
Additional information about this publication is available at https://csrc.nist.gov/pubs/sp/800/209/r1/ipd, including related content, potential updates, and document history.

All comments are subject to release under the Freedom of Information Act (FOIA).

Abstract

Storage technology has evolved in two primary directions: increased media storage capacity and architectural changes that provide a software-based abstraction over all forms of background storage technologies. However, this architectural shift has increased management complexity and the probability of configuration errors and associated security threats. This document provides an overview of the evolution of the storage technology landscape, current security threats, and the resultant risks to provide a comprehensive set of security recommendations in the form of storage security controls. These recommendations span security management areas that are common to an information technology (IT) infrastructure as well as those specific to storage infrastructure.

Keywords

backup; block storage service; cloud storage; data protection; file storage service; hyper-converged storage; network-attached storage; object storage service; replication; software-defined storage; storage area network; storage array; storage virtualization.

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in federal information systems. The Special Publication 800-series reports on ITL's research, guidelines, and outreach efforts in information system security, and its collaborative activities with industry, government, and academic organizations.

Call for Patent Claims

This public review includes a call for information on essential patent claims (claims whose use would be required for compliance with the guidance or requirements in this Information Technology Laboratory (ITL) draft publication). Such guidance and/or requirements may be directly stated in this ITL Publication or by reference to another publication. This call also includes disclosure, where known, of the existence of pending U.S. or foreign patent applications relating to this ITL draft publication and of any relevant unexpired U.S. or foreign patents.

ITL may require from the patent holder, or a party authorized to make assurances on its behalf, in written or electronic form, either:

a) assurance in the form of a general disclaimer to the effect that such party does not hold and does not currently intend holding any essential patent claim(s); or

b) assurance that a license to such essential patent claim(s) will be made available to applicants desiring to utilize the license for the purpose of complying with the guidance or requirements in this ITL draft publication either:

   i. under reasonable terms and conditions that are demonstrably free of any unfair discrimination; or

   ii. without compensation and under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

Such assurance shall indicate that the patent holder (or third party authorized to make assurances on its behalf) will include in any documents transferring ownership of patents subject to the assurance, provisions sufficient to ensure that the commitments in the assurance are binding on the transferee, and that the transferee will similarly include appropriate provisions in the event of future transfers with the goal of binding each successor-in-interest.

The assurance shall also indicate that it is intended to be binding on successors-in-interest regardless of whether such provisions are included in the relevant transfer documents.

Such statements should be addressed to: sp800-209r1-comments@nist.gov

Table of Contents

List of Tables

Acknowledgments

The authors express their sincere thanks to NetApp, Dell, SNIA, IEEE, Infinidat, FCIA, Center for Cybersecurity Standards at NSA and Richard Austin for their extensive, insightful feedback. Last but not the least, the authors would like to express their thanks to Isabel Van Wyk of NIST for her detailed and extensive editorial review.

Executive Summary

Storage, computing, and networking form the three fundamental building blocks of an information technology infrastructure. Like computing and network technologies, storage technology has also evolved over the years, specifically higher capacity media and storage system architecture. These architectural developments have enabled storage services to support new computing use cases but have also introduced storage management complexity and many security challenges.

This document provides an overview of the current storage technology landscape, including traditional storage services (e.g., block, file, and object storage), storage virtualization, storage architectures designed for virtualized server environments, and storage resources hosted in the cloud. It also describes and analyzes various threats and risks to storage resources and infrastructure as well as their impacts.

The primary purpose of this document is to provide a comprehensive set of security recommendations in the form of storage security controls for the current landscape of the storage infrastructure. The security focus areas include those that are common to the entire IT infrastructure, such as physical security, authentication and authorization, change management, configuration control, incident response, and recovery. Within these areas, security controls that are specific to storage technologies, such as network-attached storage (NAS) and storage area networks (SAN), are also covered. In addition, specific security recommendations are provided for the following areas of operation in the storage infrastructure:

- Data protection
- Isolation
- Restoration assurance
- Encryption

1. Introduction

Storage, computing, and networking form the three fundamental building blocks of any information technology (IT) infrastructure. Storage infrastructure has evolved over the years to become feature-rich in areas such as performance and efficiency due to developments on two fronts. One is the area of storage media, which features solid-state drives (SSDs) with high capacity and storage efficiency features (e.g., deduplication, compression) compared to hard disk drives (HDDs). The second is the area of storage system architecture that uses concepts such as storage virtualization. However, development on this second front has also introduced a great deal of management complexity, including the task of providing security assurance.

Briefly tracing the history of storage system architecture shows that the earliest form of digital storage infrastructure is direct-attached storage (DAS), where the storage element or device (e.g., tape, hard disk) is directly attached to the host server without any intervening network. The next evolution of storage infrastructure is one where the storage resources are logically pooled, located across the network, accessed through networking protocols, and accessible by multiple hosts or servers. This type of storage infrastructure is the only way that the data access needs of distributed systems can be supported since application components that need to share data are in different nodes of a network.

In this stage of evolution, the storage infrastructure has taken two forms, depending on the type of networking protocol. In one form, the storage resource is simply a node in a network that uses common networking technology (e.g., local area network [LAN], wide-area network [WAN]). An example is network-attached storage (NAS), which provides file-level access to heterogeneous clients across a network using higher level protocols, such as network file system (NFS) or server message block (SMB). In the other form, there is a dedicated network for communicating with all storage resources. This is exemplified by the storage area network (SAN), an implementation of which uses a specialized, high-speed network (e.g., Fibre Channel) that provides block-level access to storage. Another implementation of a SAN uses the Internet Small Computer Systems Interface (iSCSI) protocol over a possibly shared LAN/WAN infrastructure.

The next wave of storage evolution involves the introduction of cloud storage, which offers a highly scalable and durable set of storage services that are completely software-defined. Cloud storage services often include:

- Block storage services, which expose software-defined block devices that can be presented to virtual hosts running in the cloud
- Object storage services, which can be mapped to hosts, applications, or even other cloud services and allow for addressing discrete, unstructured data elements by ID or metadata
- Scalable shared file systems, which can allow a scalable set of hosts to access the same file system at a high speed
- A variety of replication, caching, mirroring, and point-in-time copy services to all of the above

Another type of storage infrastructure contains interfaces to support the data storage needs of emerging stateful applications that are designed using microservices-based architecture and deployed using containers organized into clusters with container orchestration platforms. These platforms have a standard plug-in mechanism by way of a container storage interface (CSI) that connects the clusters configured by them to different types of persistent storage implementations.

1.1. Scope

This document provides security recommendations for the following storage technologies:

- Traditional enterprise storage technologies classified by storage service interface type (e.g., block, file, object)
- Network-based storage (e.g., NFS, SAN)
- Storage systems that have a layer of software abstraction (e.g., software-defined storage, storage virtualization)
- Storage systems designed exclusively for virtualized server environments (e.g., storage for virtual machines [VMs] and containers, converged and hyperconverged storage systems)
- Storage systems designed with Application Programming Interfaces (APIs) for cloud access

The security recommendations span the following operations:

- Operations that are carried out for other infrastructures (e.g., computing, networking) in which specific tasks are applicable to storage infrastructure, such as physical security, authentication and authorization, audit logging, network configuration, isolation, configuration control, change management, and training
- Operations that are unique to storage infrastructure, such as data protection and restoration assurance

Storage infrastructure for mainframes is out of scope for this document.

1.2. Target Audience

The target audience for the security recommendations discussed in this document includes:

- The Chief Security Officer (CSO) or Chief Technology Officer (CTO) of an IT department in a private enterprise, government agency, or cloud service provider who wishes to formulate enterprise- or data center-wide policies for the entire infrastructure, including storage infrastructure
- System or storage administrators who set up specific deployment configurations for storage, converged, or virtualized systems

1.3. Relationship to Other NIST Documents

This document focuses on a particular type of infrastructure that provides access to all data resources and services, similar to how the computing infrastructure provides access to computing services and the networking infrastructure provides access to communication services. Hence, some security recommendations related to computing and networking are relevant security strategies for the storage infrastructure discussed in this document. These common recommendations are either included here with a brief description or incorporated by reference. The relevant NIST documents containing recommendations that span all infrastructures (i.e., computing, networking, and storage) are:

- Special Publication (SP) 800-52r2 (Revision 2), *Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations* (2019)
- SP 800-53r5, *Security and Privacy Controls for Information Systems and Organizations* (2020)
- SP 800-57pt1r5 (Part 1), *Recommendation for Key Management: Part 1 – General*
- SP 800-63A-4, *Digital Identity Guidelines: Identity Proofing and Enrollment* (2025)
- SP 800-63B-4, *Digital Identity Guidelines: Authentication and Authenticator Management* (2025)
- SP 800-88r2, *Guidelines for Media Sanitization* (2025)
- SP 800-125Ar1, *Security Recommendations for Server-based Hypervisor Platforms* (2018)
- SP 800-125B, *Secure Virtual Network Configuration for Virtual Machine (VM) Protection* (2016)

1.4. Document Organization

The organization of this document is as follows:

- Section 2 provides an overview of traditional enterprise storage technologies, storage access technologies that provide a level of abstraction, storage architectures tailored for virtualized server environments, and storage resources in the cloud. This section also describes certain general principles of storage administration.
- Section 3 explores threats to storage infrastructure and associated risks, such as unauthorized storage configuration changes, media theft, and insecure storage images. The resulting risks to storage infrastructure (e.g., data breach or exposure, unauthorized data alteration and addition, data corruption, data obfuscation and encryption, tampering of storage-related log and audit data, compromising backup and firmware) are also analyzed based on the possibility of the realization of these threats and their impacts.
- Section 4 provides security recommendations in the form of storage security controls for all facets of storage infrastructure management.

- Appendix A provides a list of acronyms used throughout this document.
- Appendix B provides a glossary of important terms used throughout this document.
- Appendix C provides a summary of the controls as well as mapping of the mitigating controls to the storage-oriented threats.
- Appendix D provides a summary of the changes for this version of NIST SP 800-209.

2. Data Storage Technologies Background

Data storage technology encompasses the devices, objects (e.g., storage arrays, network switches, media), and processes (e.g., protocols, interfaces) used to store computer data in non-volatile (i.e., durable) form. Hence, this technology can be viewed from the following two taxonomies:

- **Based on the location of a storage resource:** The storage device is either directly attached to the storage client or host computer (i.e., DAS), or there is a network separating the host computer and the storage device (i.e., networked storage).
- **Based on storage type (access type):** This classification is based on the service interface offered by the storage system that is used by the client software. Examples include block-based storage (i.e., block storage service), file-based storage (i.e., file storage service), and object-based storage (i.e., object storage service).

In DAS, the storage device can be either an integral part of the computer (attached to the bus) or external storage (attached to a computer port, such as serial or USB).

Networked storage is broadly classified based on the type of access. For example, NAS provides file-level access across the network, and SAN's protocols provide block-level access across the network. Furthermore, in SAN, either the entire network stack can be comprised of storage-specific protocols (e.g., Fibre Channel), or it can consist of storage-specific protocols running over or encapsulated within common networking protocols (e.g., iSCSI by design running over Transmission Control Protocol/Internet Protocol [TCP/IP], Fibre Channel over IP [FCIP], cloud block storage). In networked storage, remote storage devices are presented as if they are locally attached to the host system on which the storage client software is running.

The taxonomy based on storage, access, or service type includes block, file, and object types or services. Since the choice of a storage service is dictated by the specific IT system use case (e.g., volume of data, control over data, performance, nature of data representation), this document will present an overview of storage technology in terms of these services (synonymous with storage/access types).

The data storage media landscape is constantly evolving and includes many technologies, such as:

- Magnetic (e.g., spinning disk drives, tapes)
- Optical (e.g., CD-R, DVD-R, Blu-ray) and magneto-optical
- Semiconductor (e.g., SSD, flash drives, persistent memory devices)
- Hybrid (e.g., devices that incorporate both magnetic and semiconductor technologies)

Experimental and less prevalent technologies include molecular memory (e.g., polymer-based), holographic (e.g., crystal-based), and DNA-like.

2.1. Block Storage Service

A block storage service offers an interface that reads and writes fixed-size blocks of data, typically with high-bandwidth and low-latency access to storage devices at the block level. Each storage device in a block-level storage system can be controlled as an individual hard drive, and the blocks are managed by the host operating system (OS).

2.1.1. Storage Area Network (SAN)

The building blocks of SAN-based systems typically include (a) host computers (clients); (b) topology, which involves the distribution of switches; and (c) storage devices/arrays. All three components are interconnected using various network stacks. SAN is a specialized, high-speed network for block-level network access to storage [1], and its variants are the results of different types of network stacks with different protocols in certain layers of the stack. There are several SAN protocols, including:

- Fibre Channel SAN (FC SAN)
- IP SAN
- Protocols of Non-Volatile Memory express (NVMe) over fabrics (NVMe-oF)

An FC SAN is a network stack that uses the Fibre Channel protocol with five layers (i.e., FC-0 through FC-4, unlike the seven-layer Open Systems Interconnection [OSI] stack). The logical storage resource addressed by Fibre Channel SAN is called the logical unit number (LUN). IP SAN is a network stack that uses the IP protocol at the network layer. iSCSI is an example of IP SAN. Although such use is technically inaccurate, a block device in an FC SAN or IP SAN is often referred to as “LUN.” The historical roots of the term stem from SCSI logical unit numbers (LUNs) that were carried along to SANs.

NVMe is the standard host controller interface for systems that use PCI Express (PCIe)-based SSD. The NVMe over Fabrics (NVMe-oF) specification defines a protocol interface and related extensions that enable the NVMe commands to be transmitted over the network. Thus, NVMe-oF extends the NVMe deployment from a local host to a remote host for a scale-out NVMe storage system.

The protocols of NVMe-oF include Remote Direct Memory Access (RDMA) (a family of protocols), Fibre Channel, and TCP. RDMA [4] enables server-to-server data movement directly between application memory without any CPU involvement, thus allowing a local application to read or write data on a remote computer’s memory with minimal demands on memory bus bandwidth and CPU processing overhead while preserving memory protection semantics. Infiniband [5] is one of the industry standard interconnects that enables RDMA for high-performance computing (HPC) environments.

The topology of the nodes together with the various hardware elements in a SAN system are called the SAN fabric. For example, a Fibre Channel SAN fabric consists of a Name Server (NS) that registers and communicates with switches and end points (i.e., host bus adapters [HBAs]). There are two commonly used topologies in FC SAN: point-to-point (two devices are directly

connected) and switched fabric. In switched fabric, there is a set of hardware switches acting as one big logical switch that connects the host computer and the storage resources. Security recommendations in this document cover only this commonly deployed topology.

2.1.2. Other Forms of Networked Block Storage

Other forms of block storage that can be presented to hosts over IP networks include:

- Cloud block storage service, which is offered in all cloud environments (see Sec. 2.6)

2.2. File Storage Service

This type of service presents storage resources in a file system model with files contained in directories within volumes. Files can be replicated by making redundant copies or be encrypted. The different variations of this service and their associated protocols are:

- **NAS with Network File System (NFS) protocol [6]:** A module that is part of the protocol implementation system (NFS client driver) mounts the volumes that are relevant for the client in its environment. The volume can be shared by multiple clients. The files or folders can also be shared from either a dedicated appliance (“NAS device” or “NAS array”) or from any host running the NFS server service.
- **NAS with SMB protocol connection:** This is provided by a LAN-attached file server, like those that provide NFS protocol connections, but with the standard SMB protocol that is found in the network stack of OSs used in personal computers and workstations.
- **NAS with multi-protocol support:** There are file service storage offerings that support multi-protocol exports of a folder or file system (e.g., both NFS and SMB, concurrently). Each of these can have slightly different access control structures (i.e., access control list [ACL]/permission specifications), and some conflicts in access control rights may have to be resolved during access requests.
- **NAS with parallel NFS protocol (pNFS):** This is provided through a clustered collection of storage servers that slices and/or strips data and metadata at the back end while providing dynamic, distributed client connections at the front end across the set of clustered hosts. The pNFS is implemented by either (a) partitioning file system namespace and assigning storage resources (i.e., files) that belong to different namespaces to different servers (symmetric clustering) or (b) splitting functionality across servers (asymmetric clustering) by having a primary fileserver provide the directory information for the location of secondary storage servers, the data contained in them, and the method for accessing them. This service is used for large-scale content repositories (because of its scalability), media stores, and development environments [7].

2.3. Object Storage Service

An object storage service presents data as flexible, discrete buckets or containers that store objects. Unlike the fixed-size blocks offered by a block storage service or the directories and subdirectories of a traditional file system or NAS service, each object can be of arbitrary size and is assigned a unique identifier (an object ID number or OID) that is compiled with other OIDs into a flat index used to access the data in each object. In addition, dynamic metadata can also be attached to an object to facilitate flexible search and addressing.

The primary benefit of object storage is scalability since its flat index is more efficiently searched than a traditional file system. With only an OID needed to search for data (plus an offset into the object itself), lookups are essentially a two-step process versus the multiple steps used to walk the directory tree of a traditional file system. The immediate effect of this is a reduction in metadata handling by the storage system, which means faster file access, especially as the system grows to very large proportions.

2.4. Storage Virtualization

Storage virtualization is the act of abstracting, hiding, or isolating the internal function of a storage (sub) system or service from applications, compute servers, or general network resources for the purpose of enabling independent application and network management of storage or data. Storage virtualization allows multiple storage devices or arrays to be pooled (abstracted) so that they can be managed as one entity. Virtualization can aggregate and manage storage resources as logical storage across a wide range of physical storage devices in large networks (e.g., SAN) or data centers. Virtualization also allows for segregating a storage resource or pool to multiple virtual representations of such resources. This technique provides the flexibility to change the logical to physical relationship over time and mask the details of physical storage resources [10].

The following are example scenarios in which storage virtualization is deployed:

- Portions of multiple physical disk drives can be presented as a single mirrored logical volume using a logical volume manager (LVM) in a host or storage array. The composition of physical disk drives in the mirrored volume can be changed, and devices can be written concurrently on both mirrors (“active-active”).
- Sensing changes to access patterns, the drives on which data is stored can be changed (e.g., store frequently accessed data on high performance drives) to provide an automatic tiering functionality.
- Large-scale data migrations need to be performed transparently to hosts and applications.

The benefits of storage virtualization are scalability, performance, redundancy, and increased storage resource utilization.

2.5. Storage for Virtualized Servers and Containers

A virtualized server is one in which a single physical server runs multiple computing stacks that consist of an OS, storage, network, and applications. These virtual machines (VMs) use software called a hypervisor.

Containers offer a lighter form of packaged compute, network, and storage units. Multiple containers can run on a server, a VM, or specialized clusters that provide orchestration services. Persistent storage for containers [14] is provided by creating volumes through a new file directory that is local to the host on which the container is running or through mapping to an external SAN or NAS device using plugins. These volumes can be created ahead of time or at when the container starts and can be shared by multiple containers. Plugins are provided by storage vendors to facilitate the process of volume creation while conforming to the specification of the container engine/orchestrator. Plugins automate the process of LUN/volume creation as well as mapping to the host and, eventually, the container.

2.6. Cloud-Based Storage Systems

Storage systems in the cloud may be either standards-based or proprietary and may include object-, block-, or file-based services. The technical reasons for enterprises to use storage systems in the cloud are [20]:

- To accommodate new demand for storage resources without building an additional data center
- To respond to changes in demand for storage, such as peaks and valleys
- The need for immediate storage capacity
- Increasing management complexity of on-premises storage infrastructure

Storage systems based in the cloud provide several sophisticated data services [21], including:

- **Collaboration capability:** Includes features such as (a) notifications when files are changed by others, (b) file sharing with the ability to set editing and view-only permissions, (c) simultaneous editing, and (d) change tracking and versioning.
- **Data integration and analytics capability:** Ability to integrate data resident on several cloud sources, perform complex analytics, and either instantly serve the extracted data or store it in a persistent storage for later access by cloud service customers.
- **Advanced data protection services:** Includes replication, mirroring, auditing, and encryption.

Cloud storage services often include:

- Block storage services that expose software-defined block devices that can be presented to virtual hosts running in the cloud
- Object storage services that can be mapped to hosts, applications, or even other cloud services

- Scalable shared file systems that can allow a scalable set of hosts to access the same file system at a high speed
- A variety of replication, caching, mirroring, and point-in-time copy services to all of the above

Additional cloud services (e.g., managed database services, data lakes, memory caches, messages queues) may also be offered, all of which can store stateful and transient data. However, experts are divided over whether to classify them as storage services in and of themselves.

2.7. Storage and Data Management

Storage management refers to all activities geared toward ensuring reliability, resilience, performance, and the security of storage resources through management tools and processes. Non-security control-related activities that are followed as a state of practice are:

- Storage resource configuration and resource management
- Data classification or categorization
- Storage sanitization
- Data retention
- Data protection
- Data reduction

2.7.1. Storage Resource Configuration and Resource Management

Storage resource configuration and resource management involve the complete life cycle management of storage infrastructure and include:

- Management and control of physical storage devices, such as storage arrays, SAN switches, software updates, device configuration, and device onboarding and disposal
- Change orchestration across multiple assets
- Management of storage resources (e.g., block device, filesystems, pools), including assignment to hosts, replication, point-in-time copy management, data migration, and data tiering
- Performance management and optimization
- Capacity management and optimization
- Inventory management
- Event management

2.7.2. Data Classification or Categorization

Enterprise data can be classified along several dimensions, such as:

- Sensitivity (e.g., sensitive versus non-sensitive)
- Frequency (e.g., frequently accessed versus non-frequently accessed)
- Environment (e.g., production versus development versus testing versus staging)

Sensitivity classification is necessary to enable the provisioning of appropriate security controls (e.g., authentication, authorization, encryption, key management, sanitization). Furthermore, the sensitivity category may have sub-categories based on regulations applicable to the data, such as personally identifiable information (PII), Health Insurance Portability and Accountability Act (HIPAA)-related, and the Payment Card Industry Digital Security Standard (PCI-DSS). Frequency classification is necessary to provision the appropriate storage media (e.g., SSD versus HDD). The environment classification may be relevant for both media selection and security controls. Other classification schemes based on project, application, or other factors may exist but are not listed here since they do not necessarily have security control implications.

2.7.3. Storage Sanitization

Storage sanitization is the process of rendering previously written data irretrievable, such that there is reasonable assurance that the data cannot be accessed or reconstructed [10]. There are three methods for sanitization:

- Clear (e.g., overwriting existing data)
- Purge (e.g., using a strong magnetic field for magnetic media degaussing, cryptographic erase for encrypted data)
- Destruct (e.g., incineration, pulverizing)

Factors that determine the appropriate type of sanitization include the category of data on the media, the nature of the media (e.g., solid-state, magnetic, optical), and reuse plans for the media.

Sanitization can be applied to individual storage media or to logical data (e.g., in the cloud). To achieve its goals, sanitization must consider the type and characteristics of media. For example, overwriting data is effective on magnetic disks but not on SSDs that use flash memory because those devices do not overwrite data in place.

2.7.4. Data Retention

There may be situations in which preserving access to particular data is needed for a short/medium-term (i.e., less than 10 years) or long-term duration. This retention may be due to operational, legal, regulatory, or statutory obligations. Short/medium-term retentions are often handled within storage infrastructures using immutability technologies or locking

approaches that prevent changes to data. Long-term retentions often involve specialized services, such as archives or cloud storage.

Immutability involves the ability to prevent data destruction or alteration along with integrity verification (e.g., hashing), and the enforcement of explicit retention periods (e.g., legal holds). To meet immutability (non-editable) requirements, organizations often use “write once, read many” (WORM)-based storage or object-based storage implementations that combine WORM with metadata that can be used to perform explicit integrity checks and enforce data expirations.

2.7.5. Data Protection

Data protection refers to activities that ensure that data is accessible, usable, uncorrupted, and available for all authorized purposes with an acceptable level of performance and is handled in accordance with compliance obligations, including privacy and all physical, administrative, and technical means, to provide assurance against accidental or unauthorized disclosure, modification, or destruction.

Data protection involves activities and mechanisms that span the following [22]:

- Data at rest/at the endpoint: On a server or client device
- Data in transit: Between storage devices, client to server, or server to server
- Data in use: While viewing, modifying, or synchronizing between devices

The range of objectives and associated activities provides a taxonomy for classifying data protection activities under three facets: storage, privacy, and information assurance/security [22]. Activities related to information assurance/security are predominantly technical controls that each require a dedicated section to discuss their details. Hence, this section only discusses storage-related data protection activities and controls. This category of controls includes:

- Data backup and recovery
- Replication and resilience technologies
- Immutability
- Point-in-time copies and snapshots

Backup is an operation wherein data stored in storage devices is accessed by production systems and periodically copied to another set of storage devices (some of which may be offline). Because of the changing nature of data content, a backup taken at an earlier time is often made obsolete by a backup taken at a later time. Backups can either be “file backups,” which back up a select portion of the data in a storage device (often based on logical data structures, such as files, directories, or data under a database schema), or “image backups,” which contain the entire content of a particular device (e.g., an individual LUN).

Data replication is the process of writing the same data to at least two separate locations [20]. Replication is often used as part of the data recovery process and involves copying data from one site to another. Generally, there are two types of replication: synchronous and

asynchronous. Synchronous replication involves the real-time copying of data from site A (e.g., a production platform) to site B (e.g., a specially designated disaster recovery [DR] site).

Asynchronous replication involves time delay and may be performed continuously or using a designated frequency for writing data from site A to site B. The time delay and frequency are dictated by the enterprise's disaster recovery policy and are described in terms of specific recovery time objective (RTO) and recovery point objective (RPO) goals.

Point-in-time copies are usually created by the immediate element storing the source data (e.g., storage array, file system, database) and are designated for fast recovery and a wide variety of other uses, such as cloning production data for testing purposes. Point-in-time copies can also be taken from remote replicas and referred to as "remote point-in-time copies."

A snapshot is a storage-efficient form of a point-in-time copy that stores only the individual portions of the data changed (the delta) from a given point in time in reference to the source data (the base). This allows the current state of the data to be recreated from the base by applying the snapshot. This often means that if the base is unavailable, the snapshots will be unusable.

2.7.6. Data Reduction

Data reduction is the process of reducing the amount of data stored and/or transmitted in an effort to reduce costs and improve efficiency. Two common approaches to data reduction are data deduplication and compression, which can be used together.

Lossless data compression (sometimes performed on hardware) seeks to reduce the amount of data by applying a known algorithm to produce a representation of data that uses less storage than the unencoded representation [10] with the property that the compressed version can be uncompressed to restore the original representation of the data. Data compression is most commonly used as part of tape backups and during remote data replication in network gateways to reduce the bandwidth needed for DR and business continuity (BC) operations. Interoperability is an important aspect of data compression since compression and subsequent decompression may be performed by different entities.

Data deduplication attempts to replace multiple copies of data with references to a shared copy by eliminating identical blocks of storage. For example, if a storage system has 500 identical blocks, the storage array will store just one copy, thereby eliminating the need to store the other 499 copies [20]. This can take place at the storage device level, the transmission stage, or the file system level.

3. Threats and Risks

SP 800-30r1 [21] provides guidelines on conducting risk assessments for federal information systems and organizations and amplifies the recommendations in SP 800-39 [22] and ISO/IEC 27005 [23]. The risk management strategies described in these documents address how organizations intend to assess, respond to, and monitor risk by making risk perceptions explicit and transparent.

This section provides background information regarding storage system security threats and risks.

3.1. Threats

A threat is the potential cause of an unwanted incident, which can result in harm to a system or organization. Storage systems and storage ecosystems are subject to many of the same threats as other elements of ICT. To gain an understanding of these threats, the following sections of SP 800-30r1 [21] should be consulted:

- Appendix D — Potentially relevant threat sources
- Appendix E — Potentially relevant threat events
- Appendix F — Potentially relevant vulnerabilities and predisposing conditions
- Appendix G — Potentially relevant likelihoods of threat event initiation/occurrence
- Appendix H — Potentially relevant adverse impacts to organizational operations and assets
- Appendix I — Potentially relevant adversarial and non-adversarial risk determinations

The following are examples of generic threats that are particularly relevant to storage systems:

- Credential theft or compromise (e.g., SP 800-63)
- Flawed encryption and/or key management implementations (e.g., SP 800-57pt1r5)
- Exposure of sensitive data while in- transit (e.g., no end-to-end encryption)
- Patch management problems (e.g., patches have to come from the vendor, which can be slow to provide updates; regression testing can be complex)
- Inappropriate privileged access controls can result in severe consequences to due to attacks or mistakes and damage data

The following sections provide a brief overview of threats related to storage infrastructure.

3.1.1. Malware Infection and Ransomware

Malware is the general term for any software that is designed to damage, disrupt, or provide means to compromise a system, device or service. Malware compromises exploit a vulnerability (either social or ethical) to gain access, elevate privileges, install other malicious software (e.g.,

a persistent back-door), etc. Results of successful compromise can include data theft, gaining control of a device or system, and harvesting system resources for unauthorized activities. It can be installed on privileged targets such as a storage management host and then abuse that privileged target to cause harm, such as credential theft, privilege escalation, data corruption, loss or alteration, and the compromise of future backups.

Ransomware is a form of malware that encrypts stored data, rendering it unusable. The adversary then typically demands a ransom to restore access to the data upon payment. Increasingly, adversaries recognize that organizations have recovery mechanisms in place to mitigate ransomware attacks, so prior to encrypting data, some adversaries will exfiltrate data with an intention to publish confidential data that was collected from the storage system if the ransom is not paid.

3.1.2. Unpatched Vulnerabilities

As with all ICT, unpatched vulnerabilities can be used directly or indirectly to bypass other security controls. Exploitable vulnerabilities on management interfaces can allow adversaries to reconfigure systems such that an organization temporarily loses access to its data or the data can be destroyed. Exploitable vulnerabilities on data access mechanisms of storage systems can result in authorized access to sensitive data.

Once vulnerabilities are known — and especially if they are discovered in software versions that are still publicly supported — vendors typically issue a software fix in the form of a patch or a new version to remediate the vulnerability. The timely deployment of such fixes is of paramount importance as adversaries tend to exploit the window of vulnerability between when a vulnerability is announced and when the fix is installed. A large portion of successful attacks are based on vulnerabilities for which a fix has already been published. For some storage systems (particularly those based on open-source software), generic patches for vulnerabilities cannot be directly applied because they may cause unpredictable behaviors on the affected storage system (which may include proprietary enhancements to the underlying open-source software) or potentially void warranties. Approved vendor-supplied patches can be slow to arrive, so interim remediations (e.g., disabling a vulnerable service) may be necessary while awaiting the delivery of patches.

3.1.3. Physical Theft or Inappropriate Disposal of Storage Media

All data is ultimately stored in one or more copies on physical media, which is susceptible to theft. Media, whether online or offline, can be removed from its designated stationary location or while being physically transported between locations (e.g., backup media being transported for off-site storage, storage equipment being shipped as part of a data center relocation project). Theft can be opportunistic (e.g., the content of the stolen media is not known in advance or targeted), or specific data may be of interest to perpetrators who have the means to determine which media to seize.

An organization may choose to dispose of storage media by charitable donation, internal or external transfer, or recycling if that storage media is not needed, obsolete or no longer usable.

Without appropriate storage sanitization, sensitive data may be exposed to unauthorized parties.

3.1.4. Platform Security Compromises

Adversaries can attempt to interfere with a storage device's software distribution, update, or installation process in order to introduce incorrect, outdated, or maliciously modified code (e.g., binaries, images, firmware, drivers). Affected storage components can include disk drives, tape drives, libraries, network cards and controllers (e.g., HBAs, network interface cards, NICs), switches, and storage enclosures and arrays. Software update processes can rely on complex delivery chains (e.g., vendor, third party, open-source community), delivery methods (e.g., transmission or download, shipping of installation media, file copy by vendor employee), and local copies kept by an individual organization (e.g., proxies, internal file servers). Each link in the chain can be targeted to introduce tampered software and affect platform security. Issuers, for example, can be infiltrated to infect source-code libraries, obtain access to signing software or equipment, and publish altered signed binaries on their download sites or update servers. A variety of other strategies can be devised to compromise other links in the chain.

3.2. Risks to Storage Infrastructure

The threats described in Sec. 3.1 can give rise to a wide range of risks. For storage systems and infrastructure, the major concerns are the risks associated with data breaches, data corruption or destruction, temporary or permanent loss of access/availability, and the failure to meet statutory, regulatory, or legal requirements [6].

3.2.1. Data Breach and Data Exposure

A data breach[1] is an incident that involves sensitive data being copied, transmitted, viewed, deliberately exposed to the public, or used by unauthorized individuals or entities. Data breaches can be exercised by an external source (e.g., adversary) or an internal one (e.g., malicious insider, disgruntled employee). They can be performed in a covert manner with traces being concealed or entirely removed or in a manner that can be easily identified (e.g., due to a lack of sophistication). The impact of data breaches can span a wide range, from inconvenience to users to the devastating exposure of sensitive or confidential data, resulting in irreparable damage to the reputation and operational health of the organization.

While certain data breaches involve a high degree of sophistication, many are made possible due to simple, inadvertent security settings of data assets. Among the many possible root-causes are weak (or absent) encryption at rest or in transit, software flaws, the loss of custody of removable media, media theft, incorrect or relaxed access limitations, improper or

[1] ISO/IEC 27040 [6] 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 document takes a narrower perspective.

incomplete data sanitization implementation (e.g., deleted objects, retired or repurposed media), sending or transmitting data to the wrong recipient, or data being uploaded to an incorrect location or in an incorrect manner (e.g., uploading protected data to a public data store).

3.2.2. Unauthorized Data Alteration and Addition

Data alteration refers to the process of modifying data before or after it is entered into the system. In this case, the adversary gains access to the data storage infrastructure and modifies the data in a way that will affect future application transactions or other uses of the data. Data alteration and addition can be exploited by an external or internal source in a covert or easily identifiable manner. The perception of unauthorized data alteration or addition can be sufficient to render data suspect and unusable. The impacts of data alteration and addition can range from the loss of funds to permanent damage to reputation and trust.

3.2.3. Data Corruption

Data corruption refers to damage to or errors in data that occur during writing, reading, storage, transmission, or processing and that introduce unintended changes to the original data. An object that contains corrupted data can produce unexpected results (e.g., minor losses, system crash) when accessed by a system or related application. For example, if a document file is corrupted, a user may be unable to open it, or it might open with some or all of the data rendered unintelligible. Some types of malware can intentionally corrupt or destroy files as part of their payloads, usually by overwriting them with meaningless data or otherwise securely erasing their content. While some forms of data corruption will result in storage device, OS, or software errors upon access, others can be designed to affect data without issuing errors.

3.2.4. Compromised Data Resilience/Protection

Backing up or retaining copies (including replicas and snapshots) of data assets is important to enable their recovery if they are damaged or lost. Satisfactory recovery is possible only if the backup copies are generated correctly with an appropriate retention and currency, stored securely, and accessible in a manner that supports timely restoration. Since these prerequisites closely depend on each other, backups are sensitive to multiple failures. For example, incorrect configuration can involve a live database backup performed without applying techniques to ensure consistency or write-order fidelity. Insufficient currency or retention can mean that at least some portion of the data will be unrecoverable. An adversary, therefore, has a high motivation to target not only a primary data asset but also its backup and copies. When existing copies cannot be compromised, another viable attack strategy can be to interfere with the backup process itself, thereby gradually “poisoning” future copies. When enough time has passed, the adversary can return to the original goal of compromising the primary data assets, knowing that the only available copies for recovery are too old.

Another type of "poisoning" strategy is to specifically infect backup copies of compute or application assets, such as OS images, software packages, firmware, or even source code repositories. This way, when an individual component or even an entire environment is rebuilt in an attempt to battle a malware infection, at least some portions of the malware will be included in the restored environment, allowing the adversary to quickly regain control or inflict more damage.

3.2.5. Malicious Data Obfuscation and Encryption

The malicious use of reversible obfuscation or the encryption of data can result in data becoming unavailable to the user or organization unless it is recovered using information (e.g., a key) held only by an adversary. This type of risk is commonly used in ransomware attacks (see Sec. 3.1.1). Although it originally targeted data or files on users' computers or enterprise servers, ransomware has evolved to also include other storage components, such as NAS and backup appliances]. Data obfuscation and encryption typically originate from an external source but can also be inflicted by an internal one. These attacks are often accompanied by a threat and ransom instructions. The impacts of data obfuscation and encryption can range from the loss of funds to permanent damage to reputation and trust.

3.2.6. Data Unavailability and Denial of Service

In a data availability or denial-of-service incident, the data user cannot gain access to some or all of their data due to purposeful or unintended damage to the communication path or access configuration. The damage can be physical (e.g., disconnection along the communication path) or logical (e.g., misconfiguration of an endpoint of network components). Although the damage may be reversible (e.g., by restoring the settings that were altered or deleted), it can cause lengthy disruptions and downtime for the system or service. A DoS attack can also cause disruptions to data availability by flooding the targeted storage devices, management interfaces, clients, or network with superfluous requests in an attempt to overload systems and prevent some or all legitimate requests from being fulfilled. DoS attacks can potentially impact individual data assets and clients as well as the entire fabric.

3.2.7. Compromising Storage OS or Binaries, Firmware, and Images

The compromise of storage software (e.g., storage device OSs, firmware, images) can produce a wide range of undesirable outcomes, such as providing adversaries with a means for remote access to read, copy, alter, or destroy data and data copies; change security settings; expose data; or alter the behavior of the storage infrastructure. Storage behavior alteration is of particular concern as it can be used to introduce a variety of latent, hard-to-detect attack capabilities, including presenting incorrect data to storage clients (even if stored data is intact), providing an incorrect status (e.g., falsely reporting the existence or state of snapshots and security settings), bypassing environmental safeguards (e.g., thermal, power consumption, speed limits), or stripping or altering encryption.

4. Storage Security Controls

The following subsections provide security recommendations and guidelines for storage infrastructure. Each section is dedicated to a specific aspect of storage security with a uniquely identifying naming convention and numbering scheme. Primary unique identifiers take the form 'xxxx-ss-Cyy', where 'xxxx' is a four-letter combination that relates to the subsection's heading; 'ss' is a two-letter combination that relates to the subheading within the subsection; and 'yy' is a sequential numerical identifier preceded by letter 'C' standing for control. Those controls that do not have a subheading will carry the two-letter 'GN' standing for generic control. For example, in Sec. 4.2.1, "Physical Controls for Data Storage" and subsection "Physically Security Storage," primary identifiers are labeled PSDS-PS-C03, PSDS-PS-C04, and so forth. Secondary alphabetical identifiers (i.e., '(a)', '(b)', …) are sometimes used to address an entire compound guideline by its primary identifier (e.g., 'PSDS-PS-C03' to address media security) as well as specific portions thereof (e.g., 'PSDS-PS-C03.a' (adherence to SP 800-53r5, Sec. 3.10). Finally, bulleted lists are sometimes used as part of a specific guideline when there is no need to address each list item individually.

4.1. Data Storage Governance (DSGV)

4.1.1. Policy and Planning (PP)

DSGV-PP-C01 – Comprehensive storage security policy: Create a comprehensive storage security policy, either as a dedicated policy or as part of the organization's security policy. It should include configuration baselines for storage systems and can be based on:

- Recommendations from this publication and cited sources
- Storage-related security standards internal to the organization
- Relevant vendor/industry security best practices

DSGV-PP-C02 – Keep the storage security policy current: The storage security policy should be reviewed and updated periodically (at least annually). The security baseline should be updated with the latest vendor and industry recommendations available for storage systems and/or specific storage devices (at least quarterly).

DSGV-PP-C03 – Compliance with storage security policy: Periodically and proactively assess storage system configurations to verify compliance with the storage security policy, including:

(a) Confirming that the actual configuration meets the storage security baselines and identifying gaps

(b) Tracking the remediation of gaps in a timely manner

(c) Developing key performance indicators (KPIs) to track compliance with storage security baselines based on types of data, their organization function, and their sensitivity

DSGV-PP-C04 – Data protection plan or policy: A data protection plan or policy should be established prior to deployment and should include, at the minimum, the following:

(a) Tiering, frequency, and number of copies to meet the organization's recovery goals. Specification may include more than one tier (e.g., continuous, hourly, daily, weekly) and should include for each tier:

- Frequency and retention (e.g., 48 hourly snapshots, 30 daily backups)
- Type (e.g., full, incremental, continuous, replication, point-in-time copies)

(b) Types of media to be used

(c) Encryption capabilities for data at rest and in transit as well as encryption key retention and rotation

(d) Other protection capabilities, such as digital signing, location, facility security (including fire, explosion, and magnetic interference protection), immutability and locking, minimum number of copies per backup set, and the geographic distribution of such copies

(e) References to applicable regulatory frameworks with appropriate controls

(f) Comprehensive life cycle management, including tracking data copies and backups against protection and retention policies (e.g., affirmative deletion of those that are no longer needed)

(g) Restoration procedures

DSGV-PP-C05 – Completeness of data protection plan or policy: The data protection plan or policy should be comprehensive enough to:

(a) Cover all data assets of the enterprise, irrespective of where they reside (i.e., on-premises, off-premises or in the cloud). It is acceptable to refrain from protecting data assets that have no significance to the organization or that can be recreated from other protected data sources within the desired RTO, but such omissions should be documented.

(b) Be organized by the type of data involved (e.g., Tier 1, Tier 2).

(c) Address data integrity at the application and business process levels (e.g., if two components should be recovered to the same point in time to function properly, then federated consistency mechanisms or equivalent should be planned and implemented).

(d) Address the desired restoration speed to meet business or regulatory obligations so that implementation is based on technology stacks with appropriate characteristics (e.g., disk-based, point-in-time copies versus tape or over-WAN recovery).

DSGV-PP-C06 – Backup plan or policy: In addition to a backup plan or policy, standard operating procedures related to the backup should:

(a) Monitor the execution of backups based on policy and associated notification mechanisms.

(b) Periodically test backups (at least monthly for critical data) to verify their integrity and their ability to be restored. For applications with a strict RTO, an end-to-end test restore should be performed (e.g., a complete recovery of the dataset to a sandbox recovery environment, simulating real-life restoration scenarios).

(c) An up-to-date recovery catalog should be kept to track each copy (e.g., backup, replication, point-in-time copies) and record which anti-malware tools were used and the results of the scans. For sensitive data, it is further recommended to periodically scan at least a subset of past copies with current anti-malware tools to identify "poisoned" copies. See additional cataloging recommendations in DSGV-SR-C02.

(d) Periodically review (at least annually) the backup plan and operations procedures.

(e) Maintain an audit trail that provides the information necessary to determine conformance of the backup operations consistent with the policy.

(f) Employ special controls when necessary (e.g., refreshing old, at risk, or no longer supported media by copying to new ones).

DSGV-PP-C07 – Ensure the completeness of recovery copies: All storage elements that contain components of critical data assets should be protected and backed up to support both DR and cyber-attack recovery. This includes storage volumes, critical file systems, databases, software images, certificates, encryption keys, startup files, catalog info, ACLs, virtualization settings, and configuration files.

DSGV-PP-C08 – Protect all dependent ICT components: Dependent components (e.g., directory services, DNS, external key management systems) should be protected to enable full recovery. If automated build processes are used to configure storage, then source-code repositories, build-environment, and build-procedures should be protected as well.

DSGV-PP-C09 – Document the DR plan, resources, mapping to production, flow, and test procedures: A disaster recovery plan for storage infrastructure should be written, which includes all of the resources, mapping to production, flows, and test procedures.

DSGV-PP-C10 – Storage consideration for policy: Storage should be incorporated into policies in order to:

- Identify the most sensitive (e.g., PII, intellectual property, trade secrets) 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, access controls)
- Address storage media sanitization
- Address the role of storage in business continuity management and disaster recovery
- Unique staffing and facility requirements associated with the storage ecosystem

DSGV-PP-C11 – Include storage infrastructure in the logging and monitoring policy: Include storage in the logging and monitoring policy such that the policy:

- Clearly states that storage systems and devices participate in audit logging
- Identifies the significant storage-related events to be collected
- Identifies the preservation requirements for storage-related event logs
- Identifies the retention requirements for storage-related event logs
- Specifies the time synchronization and use requirement for storage-related event logs
- Includes evidentiary expectations (e.g., authenticity, chain of custody)
- Establishes expectations for monitoring and alerting

4.1.2. Documentation and Testing (DT)

DSGV-DT-C01 – Include storage in BCM planning/testing: Organizations should perform ongoing planning and regular testing of storage-oriented assumptions/dependencies, which are critical to a successful BCM. The results of BCM testing should be fed back into ongoing maintenance of the BCM plan.

DSGV-DT-C02 – Document data retention and sanitization obligations: Organizations should document and confirm that storage infrastructure complies with data retention and storage sanitization obligations, including the implementation of:

- Appropriate data retention measures
- Appropriate data integrity and authenticity measures
- Storage sanitization prior to the repurposing or decommissioning of hardware
- Storage sanitization of virtual server images and their copies at their end of life

DSGV-DT-C03 – Document and verify storage compliance with privacy obligations: Organizations should document and verify that storage meets privacy obligations such that:

- Appropriate data access control based on least privilege is implemented to control access to data and metadata (e.g., search results)
- Appropriate data confidentiality measures are implemented to prevent unauthorized disclosure

DSGV-DT-C04 – Document and verify storage compliance with legal obligations: Organizations should document and verify that storage infrastructure complies with legal obligations, including:

- Use of data deduplication does not conflict with data authenticity requirements
- Data and storage media sanitization mechanisms do not violate preservation requirements

- Proper chain-of-custody procedures are followed when evidentiary data (e.g., audit logs, metadata, mirror images, point-in time copies) is handled

4.1.3. Storage Resource Management (SR)

The purpose of configuration management is to provide visibility and control over settings, behavior, and the physical and logical attributes of storage assets throughout their life cycle. In the context of storage security, this involves:

- Maintaining a comprehensive and current inventory
- Managing change
- Ensuring that the configuration continually meets the organization's security baselines and current industry best practices and that it is free of known risks

To this end, appropriate controls, policies, processes, and tools are needed. Comprehensive guidance for IT configuration management is provided in SP 800-53r5 [10]. The following paragraphs provide recommendations that are applicable to storage infrastructure.

DSGV-SR-C01 – Create a comprehensive inventory of all storage devices: This includes identifying the name, address, location, and software, firmware, or driver versions for all storage components, including:

- Arrays
- Storage virtualization systems
- Management consoles
- Hosts used to monitor remote network connectivity statuses (e.g., witness hosts)
- Hosts installed with storage management software or plugins
- Data protection appliances
- Backup clients and servers
- Storage network switches
- Storage adapters or host bus adapters (HBAs)
- I/O multipathing software
- Pairing of primary and replication destination storage systems
- Designated backup servers for hosts or off-site backup
- Tape libraries and drives
- Disk drives and removable media

DSGV-SR-C02 – Create a comprehensive inventory of all data and configuration assets: This includes identifying logical data components and data access configurations through the following assets:

- Storage pools, LUNs, masking, and zoning
- Initiators and initiator groups
- File shares and ACLs
- Object storage pools and buckets
- Replicas and snapshots
- Backup catalog and access rights
- Backup sets (e.g., on-premises, virtualized in the cloud, on tapes)
- Users, groups, roles, and rights
- Host access configuration to storage assets (e.g., LUNs, file shares, global file systems, object storage)
- Images of storage software and virtual appliances

DSGV-SR-C03 – Network topology documentation: Maintain current storage-related network documentation, including topology drawings (FC and IP).

DSGV-SR-C04 – Detect unauthorized storage security changes: There should be a process for detecting unauthorized changes to storage configuration using logging, comparison of configuration storage assets to past states, or comparison to organization approved baselines.

DSGV-SR-C05 – Develop a response plan for storage component compromise: The following elements in organizational risk analysis, isolation, remediation, restoration, and testing procedures should be addressed:

(a) Compromise of an entire storage array or cloud-based storage asset (e.g., SAN, NAS, object store, elastic file system)

(b) Compromise of a backup system

(c) Compromise of an individual storage element (e.g., share, block device)

(d) Compromise of an FC SAN fabric, including individual switches and SAN services

DSGV-SR-C06 – Develop a change management process for storage: Create a storage change management process as a dedicated process or as part of the organization's general change management process. It should cover:

(a) Planning, reviewing, and approving storage configuration changes

(b) Updating environment documentation and inventory (e.g., infrastructure, data, configuration)

(c) Assessing compliance with relevant security baselines following any change to the sensitive storage environment

DSGV-SR-C07 – Perform periodic isolation reviews: Isolation recommendations (see Sec. 4.2.3) should be checked at least once per year as part of a periodic audit to determine whether there are configuration gaps or drifts that could compromise the isolation of the cyber-attack recovery copy. Sensitive and high-value storage systems should be audited at least quarterly and after every major change, whichever comes first. Audit results should be documented.

DSGV-SR-C08 – Set RPO and RTO: Set an RPO for each data asset, which is the amount of data that can be lost following a failure, expressed in time (i.e., measures backward from the moment of the disruption). Likewise, set an RTO for each data asset, which is the amount of downtime that can be tolerated before systems are expected to be restored and functional. The design and implementation of data protection solutions should support data recovery while meeting these objectives.

DSGV-SR-C09 – Meeting organizational frequency and retention objectives: The data retention and copy frequency objectives for each data asset (see **DSGV-PP-C06** for more details) should be set. The design and implementation of the backup and data copy technologies should support these objectives.

4.1.4. People Controls (PC)

DSGV-PC-C01 – Storage security training: A storage security training program should be defined and incorporated into existing organizational training activities and schedules to accommodate the following audiences:

- **Information security professionals:** To provide them with a fundamental background of storage security
- **Storage administrators:** To familiarize them with storage security principles, organization policies, and security baselines
- **Managers:** To help them understand the fundamentals of data protection

DSGV-PC-C02 – Ensuring adequate storage data protection expertise: Storage teams and management should understand the technologies and practices associated with storage-based data protection technologies (e.g., backups, replications), including the applicability of data protection in the organization.

DSGV-PC-C03 – 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.

4.2. Physical Controls for Data Storage (PSDS)

Physical security is fundamental to the overall safeguarding of any IT infrastructure. Most software-based security controls can be compromised if an adversary gains access to the physical facility and equipment. In many regards, physical security for storage infrastructure is identical to those of other infrastructure elements, like computers and network equipment (e.g., facility security, surveillance, transportation). These are addressed in multiple

publications, including SP 800-53r5 [10] and SP 800-171 [13]. Additional valuable discussion regarding media disposal and destruction is available in ISO 27040 [6] SP 800-88r2 [11].

This section provides focused guidance on physical security aspects that are unique to storage infrastructure or are less emphasized in other publications.

4.2.1. Physically Secure Storage (PS)

PSDS-PS-C01 – Physically securing storage media: All storage media with recorded data should be stored in a safe, secure environment according to their information classifications, protections against environmental threats (e.g., heat, moisture, humidity, electronic field, aging), manufacturer specifications.

PSDS-PS-C02 – Physically securing storage devices: Depending on their size and location, storage devices should be physically secured as follows:

- Small-size units in an office environment can be tethered to furniture or walls or stored in locked cabinet or safes when not in use.
- Large-size units can be rack-mounted and resident in laboratory or data centers, where they benefit from the security measures of the facility (e.g., locking covers, cages).

PSDS-PS-C03 – Use media security measures:

(a) Follow general recommendations in SP 800-53r5, Sec. 3.10, including policies, access, marking, storage, transport, sanitization, cryptography, removable media, confidentiality, and disposal.

(b) Life cycle management should include purchasing media with adequate supply chain protection.

(c) For sensitive data, physical media for backup should be stored in a location that is sufficiently distant from that in which the primary data is stored.

(d) For sensitive data, a comprehensive inventory of storage media (i.e., cataloging) should be kept to track its location, ownership, capacity, and other relevant configuration attributes. Particular attention should be paid to tracking the actual content of media, including:

- Sensitivity level
- Classification (e.g., what type of data it stores, what applications and business services it relates to)
- Encryption level
- Potential impact if compromised or stolen (e.g., compromise of financial or medical information; leakage of passwords, certificates, or encryption keys)
- Mitigation/contingency steps or procedures to employ (e.g., changing passwords, reissuing keys, re-encrypting data, notifying relevant stakeholders)

- Dependencies between the data and other applications

(e) Consider using advanced tracking controls on sensitive removable media, such as radio-frequency identification (RFID) tags, Global Positioning System (GPS) tracking devices, and tamper protection.

(f) For extremely sensitive information, consider using self-activated and/or remotely controlled self-destruction mechanisms. When implemented, carefully address how to protect such capabilities as they present an attack vector that an adversary can use to trigger destruction of the devices or to harm nearby equipment, data, and personnel.

PSDS-PS-C04 – Protect all sensitive administrative equipment: Sensitive workstations, which can be used to obtain administrative access to storage infrastructure, should be managed using organization-approved security controls for access, surveillance, and auditing, including physical security. The security measures taken to protect the management workstations should be at least as strict as those used to protect the data it manages and the systems that use that data. This includes workstations located outside of the facility storing the data and work-from-home environments, when used.

PSDS-PS-C05 – Use of storage sanitization: The sanitization approach should cover storage infrastructure in detail, including non-obvious storage components. Certain elements capable of storing sensitive data beyond the media itself are sometimes overlooked when disposing of storage equipment (see Sec. 4.7), including non-volatile memory and cache objects (often found in storage arrays, SAN switches, and routers), firmware/BIOS settings, and HBA-level settings that can contain organizationally identifiable addresses (e.g., IP and SAN Fabric World Wide Names [WWNs], masking configuration, passwords). Validate that all of those elements are considered as part of the organization data sanitization policy, including policy definition, operations, and audits.

PSDS-PS-C06 – Use of immutable storage: If possible and appropriate, immutable storage should be used to help further isolate and protect recovery data (e.g., retention locking, vault locking, immutability policies).

4.2.2. Protect Physical Interfaces (PI)

PSDS-PI-C01 – Protect physical management interfaces to storage: To protect the management of physical interfaces, the organization should:

- Restrict physical access to management interfaces
- Disable and disconnect serial management ports when not in use
- Segregate LAN interfaces used for management from other LAN traffic, noting that physical isolation is preferred but logical isolation (e.g., VLANs) is acceptable
- Disable modem ports when not needed

PSDS-PI-C02 – Protect physical SAN interfaces: Limit which SAN Fibre Channel physical and logical ports can be used for in-band management on all SAN switches and storage arrays.

4.2.3. Isolation of Storage Systems (IS)

When production data is damaged or lost, organizations should be able to recover it using replicated or backed up data copies. If the damage is the result of a malicious attack, and the adversaries were also able to compromise the backup data copies, the attack on the production environment can have a devastating effect since the organization will not have the ability to recover. To improve the resilience of backup copies, sufficient isolation should be guaranteed between data assets and their recovery copies.

In this context, organizations should distinguish between at least two separate data protection scenarios:

- **Non-malicious recovery:** Having data copies that can be used in the event of a natural disaster, hardware failure, human error, or other event. These can include local copies (e.g., snapshots taken before performing maintenance), DR copies, backups, and off-site copies. The "closer" the copy is to the production environment, the more likely it is to be mapped to compute systems for the purpose of testing and DR.
- **Cyber-attack recovery:** Having data copes that are hardened, locked, and kept in isolation. The design should strive to achieve a state in which these copies cannot be impacted by *anything*, including scenarios wherein production volumes or other types of link copies have been compromised.

Keeping copies for cyber-attack recovery that are completely separate from those used for non-malicious recovery is recommended for critical and sensitive systems but not mandatory. It is important, however, that the data protection scheme in use is suitable to fully support the two scenarios. If existing copy retention mechanisms (e.g., DR or BC copies, off-site copies) are intended to support cyber-attack recovery, their configuration should be reviewed and adjusted to facilitate data isolation. This includes verification that at least a subset of those recovery copies and systems are inaccessible and independent from the production environment and that a compromise of production would not allow adversaries to impact those copies.

The following security recommendations apply to the creation and management of data copies and the associated management system.

PSDS-IS-C01 – Separation of storage systems: Organizations should distinguish between at least two separate data protection scenarios: non-malicious recovery and cyber-attack recovery. The following security recommendations apply to the creation and management of data copies:

(a) Cyber-attack recovery copies should be created on designated, separated storage environments. In private clouds, this implies physically separated storage systems. In public clouds, this implies separate accounts (or equivalent).

(b) Backup systems should be separated from production data storage systems.

PSDS-IS-C02 – Separation of management systems: Storage systems that store cyber-attack recovery copies should be managed from designated management systems that are separated from the production environment and any other system connected to production (e.g., data

protection mechanisms). It should not be possible to access such management systems with regular credentials, including production and regular backups. The system should be hosted on a dedicated environment that is only connected to an isolated network.

PCDS-IS-C03 – Isolation of storage systems: Physical isolation should be used to:

- Segregate production from other system classes (e.g., quality assurance, development), including:
  - Avoiding network connections between classes (e.g., a production server connected to both the production and development networks), when possible
  - Segregating networks and storage by class, when appropriate
  - Physically separating systems in each class
- Isolate storage devices from other data center devices, if practical

PCDS-IS-C04 – 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

4.3. Storage Systems Security (DSSS)

DSSS-GN-C01 – Defense in depth for storage: A storage systems and solutions defense-in-depth approach [25] should:

- Have a balanced focus on the three primary elements: people, technology, and operations
- Establish effective information assurance policies and procedures, the assignment of roles and responsibilities, the commitment of resources, the training of critical personnel, and personal accountability
- Deploy protection mechanisms at multiple locations to resist multiple classes of attacks
- Deploy multiple layers of defense 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, 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, antivirus updates, maintain ACLs)
- Perform regular security threat assessments to evaluate security readiness
- Monitor and react to current threats

DSSS-GN-C02 – Establish security domains for storage: Data sensitivity should be factored into the design of security domains [25] such that:

- Storage and storage networks of different sensitivity levels should be located in different security domains.
- Devices and computer systems that provide services for external networks (e.g., the internet) should be located in different domains (i.e., demilitarized 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.

DSSS-GN-C03 – Align storage security domains with purpose: Purpose should be factored into the design of security domains such that:

- Storage and storage networks used for different purposes (e.g., development, production, management) and using different technologies (e.g., SMB, NFS, iSCSI, 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 general-purpose 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.

4.3.1. Storage Systems Hardening (SH)

DSSS-SH-C01 – Perform 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 and change all default passwords
- Only open network ports that are needed

- Install the latest patches from a trusted source
- Update firmware from a trusted source
- Install and maintain malware protections
- Apply vendor-recommend security configurations

DSSS-SH-C02 – Use software updates and patches from trusted sources: When elements of the storage infrastructure receive an update (e.g., firmware) or patch, there should be some assurance that the software to be applied is from a trusted source. Otherwise, adversaries can write their own update malicious code of their choosing, such as a rootkit, botnet, or other malware.

DSSS-SH-C03 – Limit network access to management ports of SAN switches: Network access to the management ports of SAN switches should be limited to devices and administrators specifically assigned to manage the switches through a mechanism, such as an ACL.

DSSS-SH-C04 – Limit storage administrative capabilities: Control and limit the devices and components that have administrative capabilities to the minimum necessary. This includes CLI servers, management consoles, API gateways, witness hosts, and storage devices with control permissions. In particular:

(a) Actively discover components that have storage administration capabilities and verify that only the components that are authorized have them. Remove unnecessary ones, if found, and debrief.

(b) Remove unnecessary rights and capabilities from authorized devices.

DSSS-SH-C05 – Favor API access control over CLI/shell access: API access, if available, should be used instead of CLI/shell access because of the latter's ability to access the underlying OS and file systems, including configuration files. If CLI is the only option available, then it should be used over secure protocols (e.g., SSH), as stated in **DSSS-SH-C23**.

DSSS-SH-C06 – Restrict management console OS privileges: Access to management consoles should only be provided through designated storage accounts and not as an OS administrative account (see **DSSS-AM-C14**).

DSSS-SH-C07 – Manage web user interfaces: The web service providing access to management consoles should be hardened to meet or exceed the minimum standards of other web-application servers in the organization.

DSSS-SH-C08 – Secure storage management interfaces: Storage management software/firmware interfaces should be secured with the following:

- Firewalls and network ACLs to restrict access to management networks to authorized systems and protocols
- Entity authentication to establish trust relationships between storage systems and the management systems (e.g., authenticate the entities performing in-band management)

- Intrusion detection system and intrusion prevention system mechanisms to identify anomalous behaviors and guard against them
- ICT infrastructure (e.g., DNS, Service Location Protocol [SLP], Network Time Protocol [NTP]) with appropriate security controls to avoid indirect attacks
- Appropriate privileged user controls, including authentication, authorization, and secure auditing/monitoring
- Up-to-date operating systems and applications that are sufficiently hardened against attacks

DSSS-SH-C09 – Secure the remote management: When storage systems are managed remotely, the following additional security measures should be used:

- Secure channels for all remote access, such as VPNs, TLS, Secure Shell (SSH), or HTTPS
- Strong authentication or multi-factor authentication
- Restricted privileges to the minimum necessary (i.e., least privilege)

DSSS-SH-C10 – Restrict vendor maintenance interfaces: The organization should devise organizational and technical controls to restrict the management interface used for remote (i.e., non-local) vendor maintenance sessions. Remote vendor maintenance operations conducted by individuals communicating through an external network (e.g., the internet) impose significant risks to availability, integrity, and confidentiality.

DSSS-SH-C11 – Restrict vendor maintenance communications: Technical controls should restrict communication traffic (i.e., systems, ports, and protocols) to the minimum required for remote vendor maintenance operations. After the accessing party is authenticated, additional controls at the access point should be devised to authorize the vendor maintenance session. These include accepting, asking for approval, or denying the requested session. Appropriate logs containing audit records of vendor actions should be generated.

DSSS-SH-C12 – Restrict vendor dial-up maintenance: The organization should restrict dial-up access lines to authorized accessing parties. This includes enforcing a modem call-back protocol and disabling connection establishment until the vendor requests a maintenance session and the request is authorized by the organization.

DSSS-SH-C13 – Secure Intelligent Platform Management Interface (IPMI): When IPMI is used, the following additional security measures are necessary:

- IPMI should be disabled as the default configuration setting and only be enabled on a temporary basis when needed.
- IPMI traffic (usually UDP port 623) should be restricted to trusted internal networks, such as a management VLAN segment with strong network controls.
- Strong, unique passwords should 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 should be disabled to prevent adversaries from bypassing authentication and sending arbitrary IPMI commands.
- Stored passwords should be sanitized at the system's end of life (see **PSDS-PS-C05**).

DSSS-SH-C14 – Restrict host storage control privileges: In certain shared data compute cluster configurations (e.g., clusters, geo-clusters, scale-sets, storage virtualization infrastructure), hosts are granted administrative access to storage in order to control shared cluster data resource allocation and behavior. When such administrative access is necessary, restrict the scope and privileges granted to hosts to:

(a) Only the particular elements (e.g., LUNs, shares, files, objects) that the hosts need to control

(b) Only the specific actions that the hosts need to perform.

DSSS-SH-C15 – Command device or gatekeeper configuration: Certain storage arrays allow in-band administrative control to hosts with access to special block devices (e.g., devices referred to by certain vendors as "command device" and "gatekeeper"). Commands are transferred using I/O operations on those special devices. When used, the following security guidelines are recommended:

(a) **Limit the use of control devices to the minimum possible:** If feasible, eliminate the use of such devices completely (e.g., using API access instead). If not, map them to specific approved hosts only (e.g., management hosts).

(b) **Scan for control devices** – Perform network scanning to discover control devices, and confirm that they are mapped to the authorized hosts only.

DSSS-SH-C16 – Disable or limit call home or remote access: Storage infrastructure systems can have the ability to send certain telemetry and diagnostic data back to the manufacturer, such as logs. In some cases, they even enable remote connection to the system by the manufacturer with administrative rights. These mechanisms are in place to allow the manufacturer to investigate and resolve technical issues and to perform automated software updates. These capabilities can potentially be exploited by adversaries and should be disabled if they are not used. However, if they are used, they should be limited and controlled by implementing the following settings:

(a) **Change the default credentials:** Modify the default credentials used for the remote connection.

(b) **Limit permissions:** Limit access permissions to only the minimal level needed.

(c) **Enforce encryption:** Secure protocols (e.g., TLS/SSH/IPSEC using FIPS-approved encryption algorithms) should be used.

(d) **Limit access with an "allow list":** Utilize an allow list that limits access by specific IPs and specific users.

(e) **Log remote access:** All remote access should be fully logged for auditing purposes.

(f) **Enable built-in data obfuscation features:** This is applicable for storage devices that allow the obfuscation of sensitive data, such as IP addresses, WWNs, device names, and usernames.

(g) **Limit the scope of data sent** to the minimum necessary.

(h) **Review and approve:** Periodically evaluate the data that is occasionally or automatically being sent to the vendor to confirm that it does not contain sensitive information, such as IP addresses, usernames, or the actual content of storage devices. Review that the connection is performed to valid vendor IP addresses.

(i) **Authorize each connection:** If possible, implement a mechanism that will ask permission before allowing each connection.

(j) **Restrict access to gateway system:** When remote access by the vendor is performed through a gateway device, server, or appliance, take particular care to secure and restrict access to the gateway system.

(k) **Disable software updates over vendor remote access links:** In sensitive environments, downloading and deploying software components and updates (manual or automated) should not be allowed through remote vendor connection links.

DSSS-SH-C17 – Access control for management networks: In addition to separating management from other traffic in sensitive environments, further access control for management networks is recommended using mechanisms such as:

(a) VPNs, IPsec, or one or more “jump servers” or “login proxies,” which are dedicated servers in the management network that are the only ones accessible from outside of the management network and can serve to connect to other servers after proper authentication and authorization

(b) Enhanced logging and tracing, such as session recording

DSSS-SH-C18 – Secure and protect core storage management files and binaries: Storage management software often includes configuration files that present various options to control how the storage system would operate, including undocumented options. Such sensitive directories and files should be kept with appropriate limited permission, ownership, and group membership. This includes:

- **Configuration files:** Outlining users, roles, network settings, consistency groups, device groups, and other storage options. The configuration files that define consistency and device groups are often automatically propagated from central management hosts to other hosts that are attached to the managed storage system. Thus, if compromised, it can affect multiple systems.
- **Scripts:** Storage management services, daemons, and the binaries themselves should be kept in a secure way to control starting, monitoring, and stopping.

The following controls should be applied to configuration files, scripts, and any other important management-related files:

(a) Restrict access and permissions, and control the ownership of key folders and files.

(b) For sensitive environments, monitor for content changes in such files to prevent unauthorized changes or additions.

DSSS-SH-C19 – Use an approved PKI mechanism for management access: Use an organization-approved and certified centralized PKI system for storage device management and storage management consoles rather than device or software self-signed certificates.

DSSS-SH-C20 – Perform regular security audits of data: Regular audits of the security settings for storage data of all types (e.g., files, objects) should be performed to determine whether there are drifts. Audit results should be documented.

DSSS-SH-C21 – Disallow cleartext protocols: Cleartext protocols (e.g., HTTP, Telnet, File Transfer Protocol [FTP], remote shell [RSH]) should not be used. Cleartext protocols are vulnerable to sniffing, interception, and other attacks as they do not encrypt traffic or logon details, making it easy for an eavesdropper to intercept this data. In some implementations, HTTP is only supported for redirecting to HTTPS to address mistyped URLs. This should not be allowed in sensitive storage environments.

DSSS-SH-C22 – Encryption for storage management API sessions: Storage management APIs and CLIs are used for administrative access to storage systems. All API and CLI client sessions should be encrypted, leveraging features such as vendor configuration options within the management software or the API/CLI software component.

DSSS-SH-C23 – Encryption for administrative access sessions: Administrative sessions over HTTP should use TLS (HTTPS). CLI access should be encrypted using SSH rather than Telnet. The authentication during API access should not use cleartext, and the session itself should be encrypted.

DSSS-SH-C24 – Enable FIPS mode for FIPS-based environments: FIPS 140-3 [32] specifies that a cryptographic module should be a set of hardware, software, firmware, or some combination of those that implements cryptographic functions or processes, including cryptographic algorithms and key generation, and is contained within a defined cryptographic boundary. The cryptographic modules used in storage infrastructure to protect sensitive data should be FIPS 140-3-validated, and the FIPS mode should be enabled.

4.3.2. Security Auditing and Monitoring (SA)

Storage infrastructure components generate event log entries for a wide range of transactions or events. From a security or compliance perspective, it is important to capture event log entries that are necessary to demonstrate proof of operations (e.g., encryption, retention), the enforcement of accountability and traceability, meeting evidentiary obligations, and adequate monitoring of systems. This subset of general event logging is commonly called audit logging.

The following audit logging events are relevant for security purposes:

- **Management events,** such as resetting user passwords, account creation/deletion, modification of privileges, role changes, group membership changes, privileged actions, and creating or changing configuration. These events are always of interest.
- **Security-related events,** such as changes to user profiles and security configurations, failed/blocked attempts to storage, and blocked logins. These events are most often of interest, though some of them may overlap with management events.
- **Data access events,** such as access information. These events are of interest for incident response in monitoring sensitive information (e.g., determining what an adversary has touched).

Deficiencies in security logging and analysis allow adversaries to hide their location, malicious software, and activities on victims' machines. Even if the victims know that their systems have been compromised, without protected and complete logging records, they are blind to the details of the attack and to subsequent actions taken by the adversaries. Without adequate audit logs, an attack can go unnoticed indefinitely, and the actual damage done may be irreversible. Sometimes, logging records are the only evidence of a successful attack. Many organizations keep audit records for compliance purposes, but adversaries rely on the fact that such organizations rarely look at the audit logs and do not know that their systems have been compromised. Because of poor or nonexistent log analysis processes, adversaries sometimes control victim machines for months or years without anyone in the target organization knowing, even though evidence of the attack can be obtained in unexamined log files. Based on the criticality of event log data for attack detection and forensic investigation, the following are the security recommendations for implementing audit logging capabilities.

DSSS-SA-C01 – Configure and use logging on storage infrastructure: Storage infrastructure should participate in logging such that:

- Event logging is enabled on storage systems and devices, where possible.
- External or centralized event logging is used and can use local logging.
- A common, accurate time source is used across the environment so that event records from different sources can be correlated.
- Events are natively logged using standard logging protocols (e.g., syslog) that support reliable delivery and secure transports (e.g., TLS), where possible.
- External or centralized event logging is used with a trusted remote source.
- Device logs for anything other than system health monitoring and debugging are avoided because device-resident logs are more easily subjected to tampering or destruction, there is limited storage space available for logs (e.g., logs are overwritten), and they preclude the use of centralized automated analysis, alerting, and retention.
- Multiple, external log servers are used for redundancy.

- Events are logged as they occur (i.e., no buffering) when the primary drivers for audit logging are compliance, accountability, or security.

DSSS-SA-C02 – Ensure completeness of storage logging: The usefulness of event logging is somewhat dependent on the information captured (e.g., timestamps, sources, types of events). For storage systems, ensuring the 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, NTP, DNS, authentication)
    - System shutdown and restart
    - 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 (i.e., date and time)
    - Severity level
    - The source of the log entry (e.g., distinguishing name, IP address)
    - An event ID
    - Textual description of the event to enable localization/internationalization

DSSS-SA-C03 – Monitor storage infrastructure: Storage infrastructure monitoring should be implemented, including:

- Careful filtering (e.g., on fields like severity) that complies with the logging policy
- An analysis protocol to correlate audit log records across event sources to identify significant security events
- Storage logging in security information and event management (SIEM) solutions, when such technology is deployed
- Storage systems and devices in the information security continuous monitoring (ISCM) solutions, when such technology is deployed

DSSS-SA-C04 – Protect and retain storage logs: Implement appropriate log retention and protection such that:

- Audit log data with evidentiary value is handled correctly (e.g., maintain chain of custody, 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 are implemented to preserve log integrity and prevent their modification or destruction (either maliciously or accidentally).
- When audit log entries contain sensitive information, the audit log data should be protected with appropriate confidentiality mechanisms to prevent their becoming another vector for information leakage.
- Dedicated, hardened, and configured systems should be used for unique audit logging requirements (e.g., high volume, special preservation, event signing).
- Log relays and log filtering should be leveraged to minimize the impact of specialized storage requirements (e.g., WORM).

DSSS-SA-C05 – Enable audit logging: Audit logging should be enabled on all storage infrastructure components using reliable delivery and secure protocols.

DSSS-SA-C06 – Reliable, external time synchronization: An NTP service is critical for time synchronization across the infrastructure. If the NTP service is disabled, dependent systems can suffer from inaccurate timestamps on messages, events, alerts; inconsistent time across different devices; and the subsequent failure to perform log analysis, correlation, anomaly detection or forensics. Establishing and using a common, accurate time source across the environment can help correlate event records from different sources. The following are recommendations regarding the deployment and integration of NTP with storage-related devices:

(a) The NTP service should be enabled on all devices, including log servers and storage infrastructure.

(b) All devices should be configured to synchronize time with a time source server, such as an NTP server.

(c) Time synchronization validity should be monitored on all devices (e.g., that the service remains activated, that organization-approved time servers are configured in each device), and alerts for detected anomalies should be handled at a high priority.

(d) Redundancy for time source servers should be provided by deploying at least three synchronized time sources that are geographically distributed.

(e) Certificate-based authentication should be used to authenticate the time source server.

(f) Access control options (e.g., "ntpd" access restrictions) should be leveraged to restrict access to the time source servers.

DSSS-SA-C07 – Collect logs in a centralized fashion: Writing logs to central log servers (e.g., syslog server, cloud logging services) secures them within the internal network and lowers the risk of those logs being lost or altered. The following are recommendations regarding the deployment and integration of central logging with storage-related devices:

(a) Organizational logging standards for storage devices should be defined to specify the logging level. Management events and security-related events should be logged for all types of data. For extremely sensitive data, data access events should also be logged (see **DSSS-SA-C08**).

(b) All devices should be configured to transmit log event data to organization-approved central log servers according to the applicable organizational logging standard.

(c) Central logging configuration validity should be monitored on all devices (e.g., that the logging service remains activated, that logging level configuration matches the organization standards, that organization-approved log servers are configured in each device), and alerts for detected anomalies should be handled at a high priority.

(d) Multiple syslog servers should be deployed to enable continuous logging and prevent a single point of failure.

(e) At least one off-site copy for each log should be maintained.

(f) To prevent the loss of entries if the logging process is stopped and restarted before all entries are written, logging should be configured to be written to disk in real time with no buffers in place and sent over reliable protocols.

DSSS-SA-C08 – Auding logging level: The following events related to the storage protection of all storage-related objects, sites, and accounts should be included in the storage audit logging:

(a) Read-only API calls in sensitive environments

(b) All denied access attempts to services, ports, files, objects, or devices

(c) Cryptographic key management operations spanning the entire key life cycle (e.g., key generation, key deletion, certificate management), especially for events such as key shredding

DSSS-SA-C09 – Audit log retention and protection: The following measures should be adopted for audit log retention and protection:

(a) Retain log data for a sufficiently long period of time, as it often takes a while to notice that a compromise has occurred or is occurring.

(b) Allocate sufficient storage space and proactively monitor free space and unusual growth rates of log data to prevent log destinations from filling up. A known attack pattern involves filling up logs with meaningless entries to disrupt forensics, and appropriate monitoring can help identify such attacks in real time.

(c) Retained log data should be protected from tampering (e.g., using WORM or immutable storage, object locking, MFA approval for delete). If supported, the central log servers should also use such storage options.

(d) Restrict access to log data and servers using designated roles and accounts.

(e) Enable encryption since access to log data can provide adversaries with valuable insight into assets and possible attack vectors.

DSSS-SA-C10 – SIEM integration: If supported, storage infrastructure logs should be integrated with SIEM software for potential threat detection.

4.3.3. Storage Vulnerability Management (VM)

DSSS-VM-C01 – Include storage infrastructure in vulnerability management programs: Due to the specialized nature of storage technologies, storage systems are not always included in an organization's vulnerability management program. 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 programs.

DSSS-VM-C02 – Software updates and patches:

(a) **Apply storage software releases:** There should be a process for periodically updating storage software to the latest stable and secure storage release available. This includes management software, API and CLI packages, array and HBA firmware versions, and OS drivers.

(b) **Apply important security updates and patches:** There should be a process to proactively and frequently install important and urgent storage security fixes and patches.

(c) **Mitigation plan for missing patches:** Storage components with a critical vulnerability for which the vendor has not issued an update or patch should be suspended from use unless an appropriate mitigation plan can be defined and put in place.

DSSS-VM-C03 – Scan files that contain sensitive information with anti-malware tools on access: Every time a file with sensitive information is accessed, it should first be scanned with organization-approved anti-malware tools to determine whether it has been compromised.

4.3.4. Storage Accounting and Management (AM)

4.3.4.1. Authentication Recommendations

DSSS-AM-C01 – Unique identifier for all users: All users (particularly administrators) should have a unique identifier for their personal use. Identifiers assigned to administrators should, at a minimum, meet identity assurance level 3 (IAL 3) (see Sec. 4.2 and 4.5 in SP 800-63A[30]). The only exception is an emergency-use account whose secure usage is outlined in **DSSS-AM-C10**. This principle is important for accountability, audit purposes, and the ability to control access on the individual user level.

DSSS-AM-C02 – Use centralized authentication solution: In a large-scale environment, a centralized authentication solution (e.g., Active Directory, Lightweight Directory Access Protocol [LDAP], single sign-on [SSO], organization-approved cloud authentication services) should be deployed to enable the close monitoring and control of user access to storage resources as well as the uniform enforcement of the organization's authentication policies. The use of built-in authentication and permissions management capabilities should be avoided and preferably disabled.

DSSS-AM-C03 – Configuration of authentication servers:

(a) The designation of servers to perform authentication services should be strictly controlled, and their validity should be periodically checked to detect and prevent the introduction of any rogue or unauthorized authentication servers.

(b) There should be multiple authentication servers to ensure availability and avoid single points of failure.

DSSS-AM-C04 – Secure connection to centralized authentication server: All communication between the centralized authentication server and the authenticating clients should be secured through state-of-the-art protocols, such as Transport Layer Security (TLS) 1.2 or higher.

DSSS-AM-C05 – Use of multi-factor authentication: Access configuration to storage infrastructure components that store mission-critical data should be protected using a minimum of two-factor authentication with authenticator capabilities outlined in Sec. 5.1.9 of SP 800-63B [15]. MFA should be used for access by users assigned to Security Administrator and Storage Administrator roles.

DSSS-AM-C06 – Entity authentication: 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 TLS and IPsec connections as well as storage protocols (e.g., Challenge-Handshake Authentication Protocol [CHAP] with iSCSI, Diffie-Hellman Challenge-Handshake Authentication Protocol [DH-CHAP] within FCP). When possible, these entity authentication mechanisms should be used.

DSSS-AM-C07 – Secure password policies should cover service accounts: Secure password policies should be applied to individual accounts, service accounts (e.g., Simple Network Management Protocol [SNMP], Network Data Management Protocol [NDMP]), and accounts used by automation tools. The passwords should comply with Sec. 5.1.1 of SP 800-63B [15] for memorized secrets.

4.3.4.2. Account Management Recommendations

DSSS-AM-C08 – Use of accounts not associated with system users: Accounts not associated with any system user (e.g., not in an active directory, such as "guest," "anonymous," "nobody") should be disabled. If they need to be used, they should not be mapped to any system user, and all of their default configurations (e.g., password, privileges) should be changed to conform to organization-wide policies.

DSSS-AM-C09 – Account lockout: Users should be locked out after a certain number of unsuccessful login attempts to mitigate brute force password guessing attempts. Implementations of account locking include automatic reset (i.e., account unlock) after a specified period of time or a power cycle. Automatic reset should not be allowed on sensitive storage systems.

DSSS-AM-C10 – A local user account for emergency purposes: A single local user account should be maintained for access to storage resources in order to provide emergency-only access if the centralized authentication system is down. This account should conform to all organizational policies (e.g., password length). In addition, its usage should be allowed from a physically protected location and following well-documented procedures that include the appropriate approvals of relevant stakeholders and notifications of use.

DSSS-AM-C11 – Eliminate or disable default user accounts: Any default user accounts that come with the storage system installations should be removed or disabled immediately. If the accounts cannot be disabled or removed, or there is a justified reason to keep any of those accounts, the privileges assigned to those accounts should be kept to the minimum necessary.

DSSS-AM-C12 – Limit local and default user accounts: As much as possible, eliminate the use of local and default accounts. If this is not possible:

(a) Limit the use of such accounts and the privileges they have.

(b) Password policies should apply to all user, local, and default accounts, including those with administrative rights.

4.3.4.3. Privilege and Session Management Recommendations

DSSS-AM-C13 – Roles and responsibilities configuration: At a minimum, the four roles in the ISO Standard ISO/IEC 27040 [6] should be implemented for all access to storage resources (e.g., security administrator, storage administrator, security auditor, storage auditor) (see **DSSS-AM-C18**). Storage products that offer only one or two levels or privileges should not be used for storing sensitive information unless compensating controls are available to provide equivalent functionality of granular roles. For example, all management traffic should be routed through a management proxy host or “command gateway” with privilege management tools installed to restrict commands available to each role.

DSSS-AM-C14 – Separation of duties: A critical aspect of storage security is to separate administrative control planes. For example, if an adversary gains control over a host or compromises a host administrator role, they should not be able to trivially compromise its data assets, backups, and replicas. At a minimum, this includes:

(a) The privileges for data management (e.g., create and map a volume or share) and data protection (e.g., configure, stop, and delete backups) should be assigned to different roles, and these roles should not be assigned to the same user.

(b) The privileges for data management and host administration (e.g., creating/deleting objects in the storage controller) should be assigned to different roles, and these roles should not be assigned to the same user.

DSSS-AM-C15 – The privileges assigned to any role should adhere to the principle of least privilege: The permissions assigned to a role should be no more than what is necessary to perform the functions designated for that role [25], such as access to storage-specific resources (e.g., block-devices, files, objects).

DSSS-AM-C16 – Secure session management: All sessions between the client and a storage infrastructure system should be managed based on the required authentication assurance level conforming to the requirements in Sec. 7 of [15], including termination and automated logout.

DSSS-AM-C17 – Implement a "message of the day" and "login banner" notice: A "message of the day" or "login banner" notice should appear before every login to any storage infrastructure component or system via user interface (UI), command-line Interface (CLI), or API, if applicable. The message should include a legal notice, a warning that the user is accessing a restricted system with sensitive data, and any additional warnings or meaningful messages according to the organization's security and privacy policies.

4.3.4.4. Administrative Access

Storage infrastructure systems are administered by designated users who use various accounts to access these systems. The administrative users and their management hosts constitute an important attack vector that can be exploited by adversaries. Since the individuals who manage the storage systems and infrastructure are generally privileged users, the allocation and use of privileged access rights should be restricted and controlled. The inappropriate use of system administration privileges can be a major contributing factor to the failures or breaches of storage systems.

DSSS-AM-C18 – Separate security and non-security roles: A least privilege model that leverages specific roles should be implemented. According to ISO Standard ISO/IEC 27040 [6], 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 (e.g., audit log event entries can never be changed), to establish trust relationships with IT infrastructure (e.g., shared secrets for RADIUS), to manage certificate and key stores, to manage encryption and key management, and to set access controls.
- *Storage Administrator.* This role has view and modify rights for all aspects of the storage system. No access is granted to security-related elements or data (i.e., those covered by the security administrator).
- *Security Auditor.* This role has view rights that allow entitlement reviews, the 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 the separation of duties with respect to management capabilities.

DSSS-AM-C19 – Implement the principle of least privilege: Limit the administrative rights of users to the minimum necessary. This includes specifying the minimum actions that the user can carry out and limiting the scope of these permissions to include only the relevant systems or regions. Full administrative rights should only be granted to users who have legitimate needs for these rights.

DSSS-AM-C20 – Limit the access rights of service accounts: Service accounts (e.g., monitoring tools) should be limited to read-only and metadata-only access.

DSSS-AM-C21 – Authenticate and authorize all CLI/API access: CLI/API usage should be subject to authentication and authorization processes. If it is not possible to perform authentication or authorization, secure unauthorized access with additional security measures (e.g., using privilege management tools to restrict control to the minimum necessary commands and objects).

4.3.5. Protocols (PR)

DSSS-PR-C01 – SNMP security:

(a) If SNMP is not in use, it should be disabled.

(b) Change the default, known community strings, even if SNMP is not enabled. The configured strings should meet the organizational password policy.

(c) Use different community strings for devices with different levels of confidentiality.

(d) Use at least SNMP version 3.

(e) Enforce SNMP authentication and encryption (privacy) features.

(f) Do not configure SNMP with read-write access unless it is absolutely needed. In this case, limit and control the use of read-write SNMP.

(g) Use access control lists to control access to devices through SNMP.

(h) Periodically validate that SNMP traps are sent to authorized, intended managers.

(i) Refer to the Cybersecurity and Infrastructure Security Agency (CISA) TA17-156A [18] for additional guidance.

DSSS-PR-C02 – Control IP addresses used for SNMP: When configuring SNMP, all traffic should be directed to valid organization-internal IP addresses as destinations. The validity of the configuration should be periodically reviewed.

DSSS-PR-C03 – Authenticate directory, domain, and similar services: The service configurations in all storage elements (e.g., devices, switches, management workstations, management software) should be actively and periodically reviewed and adjusted (remediating any discrepancies) so that only the approved ones are used.

DSSS-PR-C04 – Considerations for using standard and non-standard TCP/IP or UDP ports: Most applications and services have a default TCP/IP or UDP port that is used to connect to the application or service. However, systems often include the ability to reassign the ports. Port reassignments should be made to improve security (e.g., limit the number of ports used by an application to a small, fixed range) rather than to obfuscate an application or service, which rarely deters adversaries as tools such as port scanners (e.g., nmap) can usually identify these alternate port assignments. Using non-standard ports can make it difficult for security scanning tools to identify suspicious activities since they are designed to expect specific behaviors on standard ports.

DSSS-PR-C05 – Use of NDMP security features: NDMP provides a means for direct transport between storage arrays and backup devices. When used, the following security features should be configured:

(a) Access control over which hosts can initiate NDMP sessions

(b) The challenge-response authentication (do not use the plaintext authentication option)

(c) Log NDMP connection attempts

(d) An NMDP password that meets the organizational password policy (e.g., length, complexity)

(e) Restricted NDMP-related rights

(f) Encrypted NDMP control connections

(g) NDMP throttling per session or per server

DSSS-PR-C06 – Use TLS in LDAP: Use TLS to secure LDAP connections when setting up directory service options for storage systems.

DSSS-PR-C07 – Additional protocols: When additional protocols (e.g., SymAPI, Storage Management Initiative Specification [SMI-S], Global Name Service [GNS]) are used, the following measures should be applied:

(a) Isolate traffic for data access and management from other environments.

(b) Limit TCP and UDP ports.

(c) Enable encryption.

DSSS-PR-C08 – Use of TLS for encrypted communications: TLS protocol should be used to support encrypted communication between storage clients and servers. To prevent the use of insecure or outdated configuration, the selection and configuration of TLS protocol implementations and the choice of hashing and encryption algorithms should be based on the following guidelines (or more current versions when published):

- SP 800-52r2, *Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations* [16]
- SNIA TLS Specification for Storage Systems [17]

4.3.6. Storage Platform Security (PS)

Verifying the integrity of systems or ecosystems is necessary to building trust, particularly in confidential computing and multi-cloud environments that process or handle sensitive workloads and data. Storage is often an important element of these environments.

DSSS-PS-C01 – Support secure initialization: Storage systems and devices should support a secure initialization sequence (secure boot) 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.

DSSS-PS-C02 – Protect system secrets: The system or device should not reveal secrets (e.g., passwords, asymmetric private keys, symmetric keys, any other secret information) even in encrypted form via programmatic methods, logging operations (e.g., debug outputs), and within externally accessible configuration files or settings.

DSSS-PS-C03 – Secure system BIOS: For storage systems/devices that include a Basic Input/Output System (BIOS), the following elements of SP 800-147 [29] guidelines should be followed:

(a) Approved BIOS update mechanisms

(b) BIOS update authentication

(c) Secure local update (optional)

(d) Integrity protection

(e) Non-bypassability

4.3.6.1. Firmware Security

A platform's firmware code and critical data should always be in a state of integrity to ensure that a computing system can be operated free from malware. Broadly speaking, the platform is comprised of the hardware and firmware necessary to boot the system to a point at which software or an operating system can be loaded [24].

DSSS-PS-C04 – Protection and update of firmware code: For all platform devices that protect firmware code and critical data, the following elements of SP 800-193 [24] guidelines should be followed:

(a) Protect and update of mutable code

a. Authenticated update mechanism
b. Integrity protection
c. Non-bypassability

(b) Protection of immutable code
(c) Runtime protection of critical platform firmware
(d) Protection of critical data
(e) Non-bypassability

4.3.6.2. Debugging/Diagnostic Interfaces

DSSS-PS-C05 – Disable debugging ports: All debug ports on systems/devices should be disabled before the system/device leaves the factory. Alternatively, the ports should only be accessible in the field after a successful exchange of a challenge-response mechanism using an asymmetric cryptographic scheme (see SP 800-63). The debugging port state should be reset to inaccessible on any reset or power cycle.

DSSS-PS-C06 – Sensitive information in telemetry/debugging logs: By design, systems/devices should not include user data, passwords, keys, or any secret or sensitive information in any telemetry or debug log pages. Prior to supplying telemetry or debugging logs to third parties (e.g., vendors), the contents should be inspected to confirm that no sensitive data is present. When such a determination cannot be made (e.g., the contents are in a binary form or encrypted), then the telemetry or debugging data should not be provided to a third party until there is sufficient assurance that there is no sensitive data in the content.

4.3.6.3. Attestation

Attestation allows platforms to determine whether a device/component is permitted to operate. Devices that fail authentication or integrity checks can be isolated before they affect system stability or security.

DSSS-PS-C07 – Enhance trustworthiness with attestation: When attestation is an element of a storage system or device's security, it should function as an attester that reports its identity and measurements to the verifier (e.g., host, server), which collects and verifies them. The requester should compare the returned measurements to known values and use the comparison results to make decisions on the trustworthiness of the storage.

DSSS-PS-C08 – Anatomy of a trusted attester: An attester is a collection of hardware, software, firmware, and a root of trust (RoT), which should have the ability to provide reliable evidence of trustworthiness (i.e., measurements) to the verifier. Each attester device must have an RoT to calculate measurements of the security state of the attester device (e.g., firmware digests, boot parameters).

DSSS-PS-C09 – Use SPDM for storage attestation: The Security Protocol and Data Model (SPDM) [28] enables access to low-level security capabilities and operations by specifying a

managed level for device authentication, firmware measurement, and certificate management. Storage systems/devices should use SPDM (version 1.4 or later) to provide hardware identity through a certificate as well as firmware identity (i.e., firmware code and configuration data), which can be used by hosts/servers to assess the trustworthiness of the storage.

4.4. Data Protection and Recovery (DPRC)

Section 2.7.5 discusses the objectives and associated activities of data protection, three facets based on the range of objectives, and primary controls from the point of the view of the storage facet. To reiterate, these controls are:

- Data backup and recovery
- Replication technologies
- Point-in-time copies and snapshots

DPRC-GN-C01 – Data protection configuration management: The data protection configuration management (e.g., backup, point-in-time copies, replication) should be centrally managed and separated from the data consumption plane. In particular, servers and clients should not be allowed to change their own data protection configuration. This should not be interpreted as excluding redundancy mechanisms that protect against a single point of failure.

DPRC-GN-C02 – Design data protection for recovery: Data protection mechanisms (e.g., backups, replication) should be designed with quick recoveries in mind rather than just the preservation of data. Specific RTOs and RPOs should be included in these designs.

4.4.1. Storage Backups (SB)

DPRC-SB-C01 – Availability of all relevant software and hardware components: All of the relevant software and hardware components (e.g., drivers, firmware) used to run the system should be backed up, protected, and available for a restore operation.

DPRC-SB-C02 – Elected backup and data copy technologies and media should match organizational RTO: RTO is a KPI used to define the expected recovery speed. The ability to meet RTOs should be examined holistically, including all dependent and related components (e.g., restoration of data, configuration files, encryption keys) while also balancing the actual recovery speed that is desired with the cost that it would take to align all of the dependent components to enable this expected recovery speed.

DPRC-SB-C03 – Test restore to ensure RTO: Perform a periodic test restore to validate that it is completed successfully and meets the desired time frame.

DPRC-SB-C04 – Validate health of backup copies: Periodically validate that the backup copies are in good health. This includes checking that there are no relevant logged errors and that the backup and data copy media are in a healthy state.

DPRC-SB-C05 – Enable the separate restoration of data and application: The data should be separated from the application so that it can be restored without restoring infected code or software.

DPRC-SB-C06 – Use data backup measures and operations securely: The backup approach (e.g., measures, operations) should provide adequate and appropriate protections against data breaches (e.g., encryption, access controls). Such activities are often handled through a chain of trusted individuals and vendors.

4.4.2. Storage Replication and Mirroring (RM)

DPRC-RM-C01 – Synchronous and asynchronous replication: In both synchronous and asynchronous replication, the same level of data protection (e.g., encryption of data at rest, access restrictions) that is used in the primary storage should be carried over to the secondary storage.

DPRC-RM-C02 – Eliminate unnecessary replication trust between storage devices: When arrays have shared replicated volumes, their privileges with respect to one another should be limited to the volumes they share. When arrays do *not* have shared replicated volumes, disable the replication trust relationship between them.

DPRC-RM-C03 – Secure replication/mirroring traffic: The confidentiality and integrity of data in transit during replication and mirroring should be protected using encryption. This recommendation can be relaxed when appropriate mitigating controls exist (e.g., mirroring over a short distance within the same enclosure or server room) (see **ENKM-ET-C03**).

DPRC-RM-C04 – Automated I/O suspension: Automated I/O suspension should be enabled for synchronous replication if the replica cannot be allowed to fall behind the primary data. Enabling this feature implies that the primary storage device will disallow any write operations on the data it stores if its synchronization with the secondary storage server is lost, and it can only resume processing when synchronization is restored. For this reason, enforcing this feature comes at the cost of opening an additional attack vector. An adversary who is aware of or suspects that it is in use can trigger primary storage denial of service by attacking the replication network paths. This trade-off should be carefully weighed.

DPRC-RM-C05 – Eliminate obsolete replicas: Obsolete replicas should be removed to reduce opportunities for data exposure.

4.4.3. Storage Snapshots (SS)

Point-in-time copies refer to a variety of storage technology internal data copy mechanisms that can be used to create a copy of original data as it appeared at a past point in time. This includes storage arrays, built-in copy-generation tools (e.g., snapshots, clones), file system-level copies, database-level copies, cloud storage pool copies, and more.

DPRC-SS-C01 – Configure point-in-time copies: When point-in-time copies (e.g., snapshots) are used as part of the backup scheme, they should be configured accordingly:

(a) To meet the RPO of the target data sets in the snapshot. For example, if the business or compliance standards are that no more than five minutes of committed data can be lost in recovery, then the snapshot interval should be five minutes or less.

(b) To meet retention obligations. For example, if hourly copies are needed for at least 48 hours, confirm that a sufficient number of hourly snapshots is preserved.

DPRC-SS-C02 – Eliminate obsolete snapshots/clones: Obsolete snapshots or clones should be removed to reduce opportunities for data exposure.

DPRC-SS-C03 – Snapshot security: Snapshot security should:

- Align the snapshot approach with its associated recovery strategy, especially for business/mission-critical data
- Protect snapshots from changes (e.g., read only)
- Provide adequate protections against unauthorized access (e.g., encryption)

4.4.4. Cyber Recovery Systems (CR)

Incident response planning is an important part of cybersecurity. The *NIST Framework for Improving Critical Infrastructure Cybersecurity* [19] provides a comprehensive discussion of the role of incident response in cybersecurity program management as well as guidelines on building a framework for cybersecurity improvement. Storage-related incidents should be handled as an integral part of the organization's incident response process, including isolation, root-cause analysis, defining and managing a response plan, testing, and periodical process review and refresh.

DPRC-CR-C01 – Disable all unnecessary services and protocols: Unnecessary services and protocols should be disabled on cyber-attack recovery storage systems. In environments where API or CLI are sufficient for management, interactive web interfaces should also be disabled.

DPRC-CR-C02 – Recovery assets immutability during incident management: To avoid compromising cyber recovery copies, immutability protections should be used on the cyber recovery copies during incident management.

DPRC-CR-C03 – Validate the hygiene of recovered compute components: Verify that recovered executables, applications, containers, and operating system images are free from infection prior to deploying them in production.

DPRC-CR-C04 – Restrict access to cyber-attack recovery systems: For sensitive information, cyber-attack recovery copies and their systems should not be accessible to regular IT staff but only to a single person (e.g., CISO) or a very narrow group of executives or security managers who use credentials separate from those used for other day-to-day duties. This ensures that if the credentials of an IT administrator are compromised, the adversary cannot use those credentials to access cyber-attack recovery copies. This restricted team can have access to the cyber-attack recovery copies, but an even smaller subset should have administrative rights that include granting permissions to other users.

DPRC-CR-C05 – Off-site storage: Cyber-attack recovery copies should be stored off-site rather than where the production data is stored. If adversaries have physical access to the production site or manage to compromise the physical site, they cannot access or compromise the off-site cyber-attack recovery copies.

DPRC-CR-C06 – Use an independent, full baseline copy: Backup systems often make use of incremental backups that capture changes to the data relative to a baseline copy. These incremental copies cannot be used during recovery without the baseline copy. For certain types of backup schemes (e.g., snapshots), only incremental copies are used (i.e., the baseline copy is the production data itself). To handle a recovery scenario properly, dependencies between copies should be accounted for, and sufficient isolation between different types of copies should be maintained. In particular:

(a) Replicated disaster recovery copies should have no dependency on production baseline data.

(b) Cyber-attack recovery copies should have no dependency on production baseline data. Dependency on disaster recovery baseline data is allowed only if those copies are properly isolated from production baseline data and meet the recommendations in **PSDS-IS-C01**, **PSDS-IS-C02**, and **DPRC-CR-C04**.

DPRC-CR-C07 – Independence from hosts and applications: Cyber-attack recovery copies should not be mounted, exported, or mapped to a host or application, and they should be restored (pushed) onto an isolated staging environment (or air-gapped) rather than directly onto the target hosts or applications. A less secure option is to allow the target hosts or applications limited read-only access (e.g., mapping, mounting) during restore only, and remove such access as soon as restoration is complete.

DPRC-CR-C08 – Consider setting up an air gap: Organizations should set up an air gap around the cyber-attack recovery copies of sensitive data. Strict implementations of air-gapping should provide full physical and network-level separation. Certain storage technologies also introduce less strict isolation technologies that enable shutting down data ports and opening them during a limited time for the periodic sync with the production system. It is important to weigh the effectiveness of each of those techniques based on the value of the data and the capabilities of the adversary. When strict implementations are chosen, attention should be paid to circumventing known vulnerabilities of air-gaped systems, such as:

(a) Preventing visual, audible, and thermal signal transmission between air-gapped systems and other equipment (e.g., maintaining sufficient distance, using dampening and/or sufficient physical distance)

(b) Preventing any potential wireless transmission capabilities in air-gapped equipment

(c) Disabling exposed data ports (e.g., USB, network)

(d) Using power conditioning or separate power circuits

DPRC-CR-C09 – Cyber hygiene of data copies: For mission-critical information, cyber-attack recovery copies should be scanned with various anti-malware scanning tools and vulnerability scanners. Ideally, all copies should be scanned. If that is not possible, scan a subset of the

copies, and keep a record of those copies scanned and secure. Cyber hygiene tools include anti-malware, vulnerability scanning, and security analytics.

DPRC-CR-C10 – Perform periodic audits: The above recommendations should be reviewed as part of a periodic audit to check the completeness of copies, reevaluate dependencies, address software and hardware needs, and consider the suitability of technology to support recovery speed, RPO, retention, health-checking, DR plan, and cyber hygiene. Gaps should be identified, tracked, and remediated. The frequency of auditing should be no less than once per year and match the sensitivity and value of protected data. Sensitive and high-value storage systems should be audited at least a quarterly and after every major change, whichever comes first. Audit results should be documented.

DPRC-CR-C11 – Secure cyber-attack recovery backups: While many organizations implement backups for disaster recovery and business continuity purposes, some organizations also perform 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 executives or managers
- 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

4.5. Storage Networking and Technologies (SNTC)

4.5.1. Block-Based Storage Networking (BK)

SNTC-BK-C01 – Block-device access control: The hosts that can access a set of SAN storage devices should be restricted through zoning (i.e., software or hardware) and masking to limit the access.

SNTC-BK-C02 – Block-device copy and replica access control: The hosts that can access a set of SAN-replicated block-devices, snapshots, and other types of point-in-time copies of such block-devices should be restricted through zoning and masking. In many cases, a host that is granted access to a device should not be allowed to access a copy.

SNTC-BK-C03 – Use zoning is SAN-switched fabrics: The zoning should be implemented in a switched SAN fabric based on sound logic, particularly as it relates to the separation of environments and traffic type, which should be separated to the maximum possible extent, by:

(a) Environment (e.g., development versus test versus production)

(b) Type of traffic (e.g., data access versus management versus replication versus backup)

(c) Type of hosts (e.g., virtualized versus physical)

(d) Storage device type (e.g., tape versus disk)

SNTC-BK-C04 – Software zoning implementation: When software zoning is implemented, hosts should only be allowed to connect to storage devices provided by the simple name server (SNS) by looking it up at the software zoning table and not directly using device discovery.

SNTC-BK-C05 – Controlling devices that can join fabric: The policy specification feature in SAN that enables the creation of an allowable list of switches, arrays, and hosts that can join the fabric should be leveraged where applicable.

4.5.1.1. Fibre Channel

A SAN fabric may be divided into "zones", which are groupings of ports that are allowed to communicate with each other. One level of security is to construct zones to prohibit communication between particular ports.

- **Soft zoning –** Soft zoning uses filtering implemented in Fibre Channel switches to prevent ports from being seen from outside of their assigned zones. The security vulnerability in soft zoning is that the ports can still be accessible if the user in another zone correctly guesses the Fibre Channel address. In this case, the FC switch will place a host WWN in a zone without evaluating the port numbers in the FC switch, which were used for connection. World Wide Port Name (WWPN) identification is considered more secure than port number identification (used in hard zoning) because any device that is physically connected to a port can grant storage access to an unauthorized host. If the SAN spans across facilities with different physical security controls, and if there is a risk that physical ports can be accessed by unauthorized individuals, soft zoning may be preferable.
- **Hard zoning –** Hard zoning uses physical port numbers on SAN switches, thereby physically blocking access to a zone from any device outside of the zone. This type of zoning protects from WWN spoofing attacks as it does not rely on host identity. If the organization's physical access is thoroughly protected (i.e., it is improbable that an intruder will access a physical port), this method may be preferable.

SNTC-BK-C06 – FC host and switch authentication: Every host and storage switch should have a unique identity and should be authenticated before joining the network (e.g., FC-SP-3).

SNTC-BK-C07 – Use an approved FC PKI mechanism: Use an organization-approved and certified centralized PKI system for the management of switch certificates (e.g., Fibre Channel Certificate Authentication Protocol or FCAP) rather than the devices' self-signed certificates.

SNTC-BK-C08 – Use Fibre Channel zoning: Zoning should be used in FC SAN fabrics with a preference for hard zoning. Default zones and zone sets should be used with care (assume a least privilege posture). Determine whether basic zoning is a strong enough security measure

for the target environment, and if not, use stronger techniques like FC-SP-3 zoning, where supported by the vendor**.**

SNTC-BK-C09 – A blended approach to FC zoning: Implement a zoning approach that blends different types of zoning mechanisms. This is preferable to simply zoning using a single type (i.e., host, switch, and storage device):

(a) Host-based zoning mechanisms control what storage resources or devices are visible to an application on a host as well as the devices that it can access. At the lowest level, the masking capability in an HBA's firmware or driver can be used to control whether the host can interact with any storage device. At the next level, OS capabilities can be used to control which devices the host tries to mount as a storage volume. Finally, the host centralized management software for volume management (e.g., LVM), clustering tools, and the file system can be utilized to control device access by applications.

(b) In switch-based zoning, the switches (especially the FC switches) have the capability to specify which devices on which ports can access other devices or ports. Port-based zoning uses hardware to enforce zoning and is therefore also called "hard zoning." In other words, switches should support zone control at the port WWN's level rather than at the switch (node) WWN level.

(c) In storage device-based zoning, the storage array is configured with a list that shows which hosts (even more specifically, which HBA ports) can access which block devices and on which ports. Access requests from unlisted hosts or HBA ports are ignored or rejected.

(d) If the zone set feature is available, it should be leveraged. This will help create multiple zones dedicated to a particular purpose, such as testing, dynamic reconfiguration, backup, and maintenance.

SNTC-BK-C10 – FC masking recommendation: Masking refers to making a block device visible or invisible to hosts. Masking should be as close to the data as possible and as far from the data consumer or client as possible (e.g., favor array over switch masking, core switch over edge switch, and switch over HBA).

SNTC-BK-C11 – A backup of FC switch configuration data: Create a backup of the switch's configuration data, including zone configuration file. The backup should be kept outside of the SAN switches to enable redeployment upon erroneous or malicious corruption or deletion.

SNTC-BK-C12 – Limit FC switch management capabilities to the minimum necessary:

(a) When implementing the SAN fabric, there should be well-defined policies that specify and minimize the set of switches that are authorized to distribute configuration data while providing acceptable redundancy.

(b) Unnecessary configuration management permissions and services (e.g., password distribution) should not be enabled.

SNTC-BK-C13 – Limit communication between FC switches: Communication between SAN switches should be limited based on security policies while ensuring that switches can only communicate with switches that are necessary.

SNTC-BK-C14 – Disable unused FC SAN ports: Unused FC SAN ports should be persistently disabled to prevent the accidental or deliberate connection of unauthorized equipment.

SNTC-BK-C15 – Audit FC SAN security configuration: Over time, some security changes might not be reliably propagated across all switches in the fabric. FC SANs should be periodically reviewed (similar to IP and Ethernet-based networking) to assess their security, identify and prioritize gaps, and define a remediation plan. Security review should be performed at least annually, and in sensitive environments, at least quarterly or after any major change, whichever comes first. Audit results should be documented.

SNTC-BK-C16 – Confidentiality protections for Fibre Channel storage: Data-at-rest encryption at the LUN-level should be used to prevent access by unauthorized parties (e.g., service provider personnel, adversaries). Fibre Channel connections that leave the protected area (e.g., confines of a physically controlled data center) should be encrypted using ESP_Header.

SNTC-BK-C17 – Sanitization for Fibre Channel-based storage: The Fibre Channel storage should include storage sanitization capabilities for LUNs (i.e., logical sanitization) and the underlying storage technologies (i.e., media-aligned sanitization).

4.5.1.2. NVMe Over Fibre Channel (FC-NVMe)

The NVMe-oF standard defers NVMe/FC-specific security concerns to the FC-SP-3 technical standard, which includes protocols to enhance Fibre Channel security in several areas.

SNTC-BK-C18 – Use NVMe-oF authentication: NVMe-oF supports two high-level types of authentications: fabric secure channel and in-band authentication. One or both authentication mechanisms should be used.

SNTC-BK-C19 – Use NVMe/FC security: Security controls for NVMe/FC can include some of the controls used with FC SAN controls (see Sec. 4.5.1.1). For NVMe/FC, the following should be used:

- Adequate physical security, including segregated networks
- FC zoning to partition the FC fabrics
- LU masking on the storage to restrict access to specific LU
- Authentication of Fibre Channel devices, secure key exchange, and secure (i.e., encrypted) communication between Fibre Channel devices

4.5.1.3. IP Storage Networking

IETF RFC 3723 provides additional useful information on both iSCSI and FCIP. In addition, IETF RFC 7146 provides important security updates to IETF RFC 3723 and IETF RFC 3821.

SNTC-BK-C20 – Segregate iSCSI networks: For iSCSI networks, access and protocols should be controlled by:

- Segregating iSCSI interfaces from general-purpose LANs
- Using virtual local area networks (VLANs) when the use of physically isolated LANs is not an option

SNTC-BK-C21 – Secure iSCSI networks: 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
- Control iSCSI initiator access by filtering based on source IP addresses and protocols
- Internet Storage Name Service (iSNS), SLP, and DNS infrastructure with appropriate security controls to avoid indirect attacks

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

SNTC-BK-C22 – Secure FCIP networks: FCIP network access and protocols should be controlled by:

- Setting up the peer-to-peer relationship between FCIP entities while recognizing that the security policies are applied uniformly
- Using a private IP network used exclusively by the FCIP entities, whenever possible
- Performing cryptographic authentication and data integrity at a minimum
- Protecting sensitive data by appropriate confidentiality measures

SNTC-BK-C23 – Secure NVMe/RDMA networks: Security Controls for NVMe/RDMA should use in-band authentication provided by DH-HMAC-CHAP.

SNTC-BK-C24 – Confidentiality protections for IP storage: Data-at-rest encryption measures should be used to protect the confidentiality of sensitive or high-value data recorded on IP storage devices or media.

SNTC-BK-C25 – Sanitization for IP-based storage: For IP-based storage, storage sanitization measures include the following:

- Media-aligned sanitization should be used for IP storage media and storage device for sensitive and regulated data.
- Logical sanitization should be used for virtualized IP storage, especially when the actual storage devices and media cannot be determined.

SNTC-BK-C26 – IP storage network separation: For storage-related communication over IP networks, sound logic should be applied to the separation of environments and traffic type. Sensitive environments should be separated to the maximum possible extent based on:

(a) Type of traffic: Data access protocols versus management versus replication versus backup versus host and application networking.

(b) In sensitive environments, further separate the management traffic of different solutions, vendors, and technologies. For example, if two or more storage solutions are in use (e.g., different array technologies, Server-Based SAN products, switch technologies, storage virtualization, or any combination thereof), and each one has a separate set of management tools, management traffic for each environment should be separated from the others.

(c) Data access protocols (e.g., iSCSI).

(d) Type of servers or hosts accessing data: Virtualized hosts versus physical hosts.

SNTC-BK-C27 – Isolate IP/Ethernet management ports: The IP or Ethernet management ports of SAN switches should reside in an isolated subnet, including separation from subnets used for data access between hosts and storage and for host-to-host communication.

SNTC-BK-C28 – Enable device IP access control: With respect to IP network accessibility, storage device security features that regulate IPs, ports, and protocols should be turned on and configured on all storage devices, where applicable. This includes but is not limited to built-in firewall rules, IP filtering, and access lists in order to:

(a) Control and limit data access between only the hosts or applications and the storage objects they use and

(b) Separately control management IP traffic between management hosts, management applications, and the relevant storage management interfaces they use.

SNTC-BK-C29 – Enable network IP access control: Restrictions should be applied at the network level (e.g., routing, firewall, access lists, Virtual Private Cloud [VPC] security groups, server-based SAN clients) to restrict all traffic types (e.g., data access, management) to allowed IP addresses and TCP/UDP ports and protocols only:

(a) Between hosts or applications and the storage objects they use and

(b) Between management hosts and applications and the relevant storage management interfaces of storage objects they manage.

SNTC-BK-C30 – Limit iSCSI ports: Hosts on the iSCSI network should be prevented from accessing any TCP ports other than those designated for iSCSI on that network.

SNTC-BK-C31 – Use iSCSI authentication: Using one of the supported methods to authenticate iSCSI initiators upon opening a session (e.g., CHAP, server routing protocol, Kerberos, Simple Public-Key GSS-API Mechanism [SPKM] 1/2). When using CHAP, prefer using two-way authentication over one-way authentication while noting that the use of authentication does not provide encryption or integrity protection to the channel.

4.5.1.4. NVMe over TCP (NVMe/TCP)

SNTC-BK-C33 – Use NVMe/TCP Security: 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 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 and the NVMe over TCP Revision 1.0 specification do not use TLS protocol versions less than 1.3

4.5.2. File-Based Storage Networking (FL)

4.5.2.1. Network File System (NFS)

SNTC-FL-C01 – 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 limiting NFSv3 usage
- Filtering client and management access by IP address for additional security

SNTC-FL-C02 – Apply NFS authentication: Authentications should be applied from all NFS access:

- Employ user-level authentication whenever possible (e.g., NFSv4 with Kerberos V5)
- Use Kerberos authentication for NFSv3
- When appropriate, use Kerberos Safe and Private modes to sign and encrypt NFS traffic

SNTC-FL-C03 – Apply NFS access controls: Access controls should be applied to NFS exported file systems to:

- 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
- Grant permissions at a finer level of granularity rather than a coarser one (e.g., file over folder, label over share)
- Avoid granting root or administrator access to files on network file systems
- Assign NFSv4 ACLs correctly

SNTC-FL-C04 – Secure NFS client access: NFS client behaviors should be restricted to:

- Filter client access to NFS shares whenever possible
- Use encrypted communications (e.g., IPsec or TLS) for data access
- Disallow NFS clients to run programs that provide temporary elevated permissions during execution on exported file systems

SNTC-FL-C05 – Secure data on NFS servers: Data on NFS servers should be secured by:

- Ensuring that export file systems are in their own partitions to prevent system degradation by an adversary 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, rootkits)

SNTC-FL-C06 – Secure NFS by restricting root access: This includes using the “nosuid” option and avoiding “no_root_squash” to prevent programs from being executed as a root user on the client and the modification of shared files by remote root users. In general, NFS clients should not be allowed to run “suid” and “sgid” programs on exported file systems.

SNTC-FL-C07 – NFS “read only” mode: For files that are to be used in the read-only mode in NFS, the mount configuration for corresponding NFS shares should always have the “noexec” option.

SNTC-FL-C08 – NFS export of administrative file systems should not be allowed: This includes the ‘/’ file systems and restricted operating system or storage array system folders.

SNTC-FL-C09 – Disable NFS non-authenticated access: Users that need to authenticate should be disabled (e.g., Anonymous, null, guest, “public access”). An exception may be provided to allow organization-critical functions (e.g., network discovery), in which case, this class of users should be mapped to the “nobody” user group and not to “ID 0.”

4.5.2.2. Server Message Block (SMB)

SNTC-FL-C10 – Apply SMB networking controls: The following networking controls should be used for SMB-based NAS:

- Turn off low-security session negotiation protocols, and use Kerberos instead.
- Maintain up-to-date patch levels.

- Use SMB signing for clients and the NAS device.
- Access 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.

SNTC-FL-C11 – Apply SMB access controls: Access controls should be applied to SMB exported file systems 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 (e.g., LDAP)

SNTC-FL-C12 – Secure 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, rootkits)
- Using SMB with strong authentication (Kerberos)

SNTC-FL-C13 – Disable insecure versions of SMB: Outdated, unrecommended, or unsecured protocol versions (e.g., SMB v1 or v2, CIFS) should be blocked and disabled on both the client side and the server side. SMB version v3 or later should be used.

SNTC-FL-C14 – Restrict SMB "full control" permissions: When SMB is used, "Full Control" permissions should not be granted to any user since the recipient can use it to modify the permissions, thus resulting in the leakage of privileges.

4.5.3. Object-Based Storage Networking (OB)

SNTC-OB-C01 – Permission for default zone: The permission for default zone (which may be product-specific) should always be configured as "deny all."

SNTC-OB-C02 – Non-public storage objects: Block any public access to non-public storage objects, particularly from the internet.

SNTC-OB-C03 – Public access: For storage objects that need public access, sufficient controls should be implemented, including:

(a) Minimizing access

(b) Using physically and logically separate storage subnets and, preferably, separate storage devices and pools from those used for non-public storage objects

(c) Addressing protection from DoS attacks

(d) Cached copies (e.g., using content delivery network, replicas, and proxies) retaining at least the same security characteristics as the source data

(e) Addressing regulatory obligations (e.g., confidentiality, storage location restrictions)

(f) Any additional applicable security controls (e.g., encryption, authentication)

SNTC-OB-C04 – Restricting access to object storage data of all types (e.g., files, objects) to the minimum possible: Follow the least privilege principle, including:

(a) Access to object storage data through any protocol (e.g., SMB, NFS, public cloud object storage) should be restricted based on client Ips, relevant subnets, and the ports/protocols being used.

(b) If supported, finer-grained access control mechanisms (e.g., by role, ID, labels, accounts, VPC, VPC endpoints) should also be used.

(c) Access should only be granted to centrally managed users and roles (e.g., those found in enterprise directories or approved commercial services) and not to local users of the specific system.

(d) The default access to any share should be set to “deny all” or equivalent.

(e) The default shares should be disabled or removed. If there is a specific purpose for using them, the access rights should be restricted to the minimum necessary.

(f) The access rights (e.g., read, write, execute, modify, delete, view ACLs, change ACLs), should be individually assigned on a need-to-know basis.

(g) If the feature to define object storage ACLs is available, it should be leveraged in addition to the use of native OS user, group, or admin permission models.

(h) If a policy definition feature to define file-level access patterns is available, it should be leveraged. The feature to detect violations to the pattern, which sends notifications, should also be implemented.

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

SNTC-OB-C05 – Use object protection to prevent unauthorized deletion: For sensitive information, when supported, use advanced controls to prevent unauthorized object deletion (e.g., MFA for object deletion, locking objects against deletion). Data objects should have appropriate data immutability protections enabled to guard against malicious accidental deletions or encryption (e.g., ransomware).

SNTC-OB-C06 – Secure transfer of objects: Transport security (e.g., IPsec, TLS) should be used for all object-based storage transactions.

SNTC-OB-C07 – Confidentiality protections for stored objects: Data-at-rest encryption at the object or tenant level should be used to prevent access by unauthorized parties (e.g., service provider personnel, adversaries).

SNTC-OB-C08 – Sanitization for object-based storage: Object-based storage should include storage sanitization capabilities for tenants and the underlying storage technologies.

4.6. Encryption and Key Management (ENKM)

ISO/IEC 27040 identifies multiple issues related to storage-based encryption solutions that should be addressed, such as:

- Encryption can adversely impact other security mechanisms (e.g., inspection of data, antivirus).
- Data can be permanently lost if anything goes wrong with data handling, data transformations, key management, or the actual encryption.
- Encryption can be computationally intensive, necessitating non-trivial computational resources. Embedded encryption within storage devices (e.g., SSD and HDD) may offset some of the computational impacts.
- Centralized key management may be necessary, especially when encryption is used in conjunction with out-of-region replication for business continuity/disaster recovery purposes.
- Encryption can diminish or negate the benefit of data reduction technologies, such as compression and/or deduplication (when such technologies are used, they should be applied before encryption – see **ENKM-EN-C04**).
- The quality of the cryptography (e.g., security strength, quantum resistance, industry-tested and accepted algorithms) can impact the actual protection offered.

ISO/IEC 27040 further states that not all data are worth encrypting. A risk assessment can help identify sensitive and high-value data that warrant the use of encryption and assist with cost-benefit analysis.

4.6.1. General Storage Encryption and Key Management Controls

ENKM-GN-C01 – Use encryption to protect sensitive, regulated, and high-value data: Encryption and appropriate key management should be used to protect sensitive, regulated, and high-value data in transit, at rest, and in use.

ENKM-GN-C02 – Storage encryption provides limited confidentiality protection: Storage encryption is typically active only while the data are resident on the storage system or media (e.g., it is plaintext once it passes through the point of encryption, which occurs any time the data are accessed). Storage-based encryption should not be the primary confidentiality protection for sensitive data.

ENKM-GN-C03 – Use encryption to protect administrative/management control of storage: Encryption should be used to protect the confidentiality and integrity of sensitive system software (e.g., firmware), secrets, and details (e.g., logs, dumps, telemetry) that are transmitted and/or stored.

ENKM-GN-C04 – Limit the operational impacts of cryptography: The use of cryptography, particularly encryption, can have operational impacts on the effectiveness or efficiency of ICT infrastructure, such as:

- Increased management complexity
- 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 (e.g., deduplication, compression techniques) on ciphertext
- Inability to apply centralized malware scanning on encrypted data

ENKM-GN-C05 – Security strength of cryptography: Cryptography used within storage systems and infrastructure should have at least 128 bits of security strength [20].

ENKM-GN-C06 – Use validated cryptographic modules: Cryptographic modules used to protect sensitive or regulated data should be validated using recognized criteria [32] [31]. When sensitive government data is being protected, this validation should be at least FIPS 140-3 Level 2 [32].

ENKM-GN-C07 – Comply with import/export regulation for cryptography: Cryptographic technologies can have strict regulations governing their import/export. Failure to address these regulations can expose an organization to regulatory penalties or cause catastrophic losses under certain conditions.

4.6.2. Encrypting Data at Rest (EN)

ENKM-EN-C01 – At-rest encryption of sensitive data: At-rest encryption protects against a variety of data-related risks (e.g., unauthorized access, compromise in case of media loss or theft) and should be enabled for sensitive data. Encryption algorithms and modes of operations that are appropriate for storage technology should be used and can include:

- AES with the XTS mode [35] [for HDDs and SSDs
- AES with the counter with cipher block chaining message authentication code (CCM) or Galois/Counter Mode (GCM) modes [36] for tape

ENKM-EN-C02 – Use an appropriate point of encryption: Protecting data at rest typically involves a single point in the data path (i.e., point of encryption) that is used for encrypting and decrypting the data. This point of encryption is important because it is the location within the data path where the data are protected (i.e., ciphertext) or unprotected and usable (i.e., plaintext). The point of encryption in the storage infrastructure should be selected to address the identified confidentiality risks for sensitive data. The type of storage (e.g., block, file, object)

can also impact the options for selecting the point of encryption. In some situations, data migrations may be necessary to take full advantage of this point of encryption.

ENKM-EN-C03 – Create appropriate proof of encryption: Encryption can be an important protection against data breaches and instrumental to demonstrating that no data breach occurred. To fully realize the benefits of data-at-rest encryption:

- The encryption mechanisms should produce appropriate records of the operational state of the encryption (e.g., activation, verification, integrity checks, re-keying).
- The organization should perform regular and audited checks that encryption was properly performed and consider outside accreditation.
- The organization should confirm that these records are retained (e.g., in the audit log infrastructure).

ENKM-EN-C04 – Use of data reduction technologies with encryption: When encryption is used along with compression and/or deduplication, the data reduction technologies should be applied before the encryption because ciphertext does not effectively compress. The reverse order of these techniques should be used when decrypting the data.

4.6.3. Encrypting Data in Transit (ET)

ENKM-ET-C01 – Use of end-to-end encryption: When the protection of data in transit 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).

ENKM-ET-C02 – Data in transit encryption: Appropriate data-in-transit encryption should be used, including:

(a) **Block over Fibre Channel:** Link encryption for FC, which is defined in ANSI/ INCITS 545-2019, Information Technology – Fibre Channel - Framing and Signaling - 5 (FC-FS-5). For sensitive information, use end-to-end (i.e., host to storage) encryption.

(b) **Block over IP:** IP storage traffic is subject to the same security risks as regular IP networks. By default, block-over IP protocols do not provide data confidentially, integrity, or authentication per packet. While specifications for IP storage traffic encryptions exist, current technology does not natively support it. When using block-over IP protocols (e.g., iSCSI, FCIP, proprietary protocols), IPsec tunneling should be used for exposed network segments. For sensitive information, use end-to-end (i.e., host to storage) encryption.

(c) **File and object storage access:** Data encryption in-flight options should be enabled for backup systems and remote replication whenever supported. For file access, sensitive data should be transmitted and encrypted using mechanisms, such as SMB encryption; available NFS encryption options, such as those offered by cloud vendors; or NFS over TLS using tunneling, such as ‘stunnel’. Objects should be accessed through HTTPS with TLS.

(d) **Extended network communication beyond the boundaries of a physically protected domain:** Particular attention should be paid to enable encryption on all connectivity segments that extend network communication beyond the boundaries of a physically protected domain (e.g., an ISL link between two physically separated datacenters, IP traffic over WAN or the internet).

ENKM-ET-C03 – Communication between storage system components should be encrypted: Storage system component interactions should be reviewed, and available encryption options should be utilized. Encryption should be used to protect communication between storage nodes and managers, active-active storage nodes with witness devices, communication with policy servers, and antivirus servers.

4.6.4. Key Management for Storage (KM)

Encryption relies on the availability and management of cryptographic keys. All communication parties should have access to the necessary keys, which need to be generated, distributed, and disposed of. Detailed recommendations for key management are provided in SP 800-57pt1r5 [20].

ENKM-KM-C01 – 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 the mechanisms and protocols associated with the keys, and the protection afforded to the keys. For storage, the following are of particular importance:

(a) All keys should be protected against modification.

(b) Secret (for symmetric encryption) and private (for asymmetric or public key encryption) keys should be protected against unauthorized disclosure.

(c) Foundational key management (e.g., secure generation, storage, distribution and destruction of keys) should conform to overall frameworks for key management [20].

ENKM-KM-C02 – Recommendations for encryption key management: The following recommendations [20] for key management should be followed:

(a) **Limit key use to finite cryptoperiod or maximum amount of data processed:** Symmetric data encryption keys used to protect the confidentiality of sensitive data should be limited to a finite cryptoperiod (typically no more than 2 years) or to a maximum amount of data processed [20]. An example of the latter case is XTS-AES [35], which limits the amount of data protected by a single key to no more than 256 TiB of data.

(b) **Use cryptographic keys for one purpose:** Cryptographic keys should only be used for one purpose. For example, key-encrypting keys (i.e., key-wrapping keys) should not be used to encrypt data.

(c) **Randomly generate keys using the entire keyspace:** Keys should be randomly generated from the entire keyspace. Key generation should follow the SP 800-90A [37] specification.

(d) **Limit 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.

(e) **Plan for key failures:** The loss or corruption of encryption keys can render data unusable, so organizations should protect their encryption keys (e.g., a key backup of a key management server) and have key-recovery plans and mechanisms in place. Key backup is often 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.

ENKM-KM-C03 – KM compatible with data retention and preservation: Encryption and key management solutions should be compatible with data retention and preservation obligations.

ENKM-KM-C04 – Use centralized key management: When possible, storage systems and infrastructure should use interoperable, 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.

4.7. Storage Disposal and Reuse (SDRU)

SDRU-GN-C01 – Storage disposal considerations: When storage systems or media are no longer needed, the organization should identify and document:

- The sensitivity and value of data previously stored on the system or media
- The reuse potential of the system or media
- The proposed disposal approach for reuse (e.g., transfer, sale, donation) or elimination (e.g., recycling of components, discarding, destruction)

SDRU-GN-C02 – Secure disposal of storage: Storage systems/media that contain sensitive or high-value data should be handled as part of a storage sanitization program based on SP 800-88 [11] and using the sanitization techniques specified in IEEE Std 2883 [33] and NSA Policy Manual 9-12 [34]prior to disposal.

SDRU-GN-C03 – Third-party collection and disposal services: Many third-party entities offer collection and disposal services for storage media. Care should be exercised when selecting and using a suitable external party supplier with adequate controls, experience, and certifications. For certain types of data, an organization's loss of control of the storage system/media without sanitization can constitute a data breach.

SDRU-GN-C04 – Accumulation of storage for disposal: The accumulation of non-sanitized storage systems and media for later disposal should be carefully handled because the aggregation effect can cause a large quantity of storage device/media with:

- Non-sensitive information to become sensitive
- Sensitive information to become a high-value target for adversaries

5. Summary and Conclusions

Along with compute (e.g., operating systems, host hardware) and network infrastructures, storage infrastructure is one of the three fundamental pillars of IT. However, compared to its counterparts, it has received relatively limited attention when it comes to security, even though data compromise can have as much of a negative impact on an enterprise as security breaches in compute and network infrastructures.

This document provided an overview of the storage technology landscape and discussed the threats and resulting risks to the safe utilization of storage resources. It also provided detailed security recommendations for the secure deployment, configuration, and operation of storage resources in various security focus areas, including:

- Areas that are common to all IT infrastructures, such as physical security, authentication and authorization, audit logging, network configuration, change management, incidence response and recovery, administrative access, and configuration management
- Areas that are specific to storage infrastructures, such as data protection and confidentiality protection using encryption, isolation, and restoration assurance

Building an effective risk management program for storage infrastructure based on the security controls described in this document and tightly integrating it with existing cybersecurity frameworks [20] could significantly improve an organization's resilience to attacks on data resources.

It must be mentioned that modern storage infrastructures that support AI/ML workloads face the same type of threats and risks as ones supporting traditional programs, though at a higher scale and velocity. Hence the same mitigating controls such as Isolation, Encryption and Restoration Assurance that we have outlined in this document apply for storage infrastructures associated with those types of systems.

References

[1] SNIA Online Dictionary as of April 2026. Available at https://www.snia.org/resources/online-dictionary

[2] Black D, Koning P, (2014) Remote Direct Memory Access Protocol. (Internet Engineering Task Force (IETF) Network Working Group), IETF Request for Comments (RFC) 7146. https://tools.ietf.org/html/rfc7146

[3] Infinibandta.org (2020) *InfiniBand Architecture Specification*. Available at https://www.infinibandta.org/ibta-specification/

[4] Haynes T (2016) Network File System (NFS) Version 4 Minor Version 2 Protocol. (Internet Engineering Task Force (IETF) Network Working Group), IETF Request for Comments (RFC) 7862. https://tools.ietf.org/html/rfc7862

[5] Black D, Glasgow J, Faibish S (2012) Parallel NFS (pNFS) Block Disk Protection. (Internet Engineering Task Force (IETF) Network Working Group), IETF Request for Comments (RFC) 6688. https://tools.ietf.org/html/rfc6688

[6] International Organization for Standardization/International Electrotechnical Commission (2024) ISO/IEC 27040 – Information technology — Security techniques — Storage security (ISO, Geneva, Switzerland). Available at https://www.iso.org/standard/80194.html

[7] Trust Radius (2020) Cloud Storage Systems. Available at https://www.trustradius.com/categories/cloud-storage

[8] Galibus T, Krasnoproshin VV, De Oliveira Albuquerque R, Pignaton de Freitas E (2016) Elements of Cloud Storage Security (Springer Nature, Switzerland AG)

[9] SNIA *Storage Security: Data Protection* Available at https://www.snia.org/sites/default/files/security/SNIA-Data-Protection-TechWhitepaper.pdf

[10] Joint Task Force (2020) Security and Privacy Controls for Information Systems and Organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-53r5. https://doi.org/10.6028/NIST.SP.800-53r5

[11] Chandramouli R, Hibbard E (2025) Guidelines for Media Sanitization. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-88r2. https://doi.org/10.6028/NIST.SP.800-88r2

[12] Kass H (2019) Ransomware Attacks Target Backups, NAS (Network Attached Storage). *MSSPAlert*. Available at https://www.msspalert.com/cybersecurity-news/attacks-target-nas-backups/

[13] Ross R, Pillitteri V (2024) Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-171r3. https://doi.org/10.6028/NIST.SP.800-171r3

[14] Temoshok D, Abruzzi C, Choong Y, Fenton J, Galluzzo R, LaSalle C, Lefkovitz N, Regenscheid A, Vachino M (2025) Digital Identity Guidelines: Identity Proofing and Enrollment. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-63A-4. https://doi.org/10.6028/NIST.SP.800-63A-4

[15] Temoshok D, Fenton J, Choong Y, Lefkovitz N, Regenscheid A, Galluzzo R, Richer J (2025) Digital Identity Guidelines: Authentication and Authenticator Management. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-63B-4. https://doi.org/10.6028/NIST.SP.800-63B-4

[16] McKay K, Cooper D (2019) Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-52r2. https://doi.org/10.6028/NIST.SP.800-52r2

[17] SNIA (2020) TLS Specification for Storage Systems Version 2.2. Available at https://www.snia.org/sites/default/files/technical-work/tls/release/SNIA-TLS-Specification-v2.2.pdf

[18] Department of Homeland Security (2017) Reducing the Risk of SNMP Abuse. Available at https://us-cert.cisa.gov/ncas/alerts/TA17-156A

[19] National Institute of Standards and Technology (2018) Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1. (National Institute of Standards and Technology, Gaithersburg, MD). https://doi.org/10.6028/NIST.CSWP.04162018

[20] Barker E (2020) Recommendation for Key Management: Part 1 – General. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-57pt1r5. https://doi.org/10.6028/NIST.SP.800-57pt1r5

[21] Joint Task Force (2012) Guide for Conducting Risk Assessments. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-30r1. https://doi.org/10.6028/NIST.SP.800-30r1

[22] Joint Task Force (2011) Managing Information Security Risk: Organization, Mission, and Information System View. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-39. https://doi.org/10.6028/NIST.SP.800-39

[23] International Organization for Standardization/International Electrotechnical Commission (2022) ISO/IEC 27005 – Information security, cybersecurity and privacy protection — Guidance on managing information security risks (ISO, Geneva, Switzerland). Available at https://www.iso.org/standard/80585.html

[24] Regenscheid A (2018) Platform Firmware Resiliency Guidelines (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-193. https://doi.org/10.6028/NIST.SP.800-193

[25] Ross R, Winstead M, McEvilley M (2022) Engineering Trustworthy Secure Systems (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-160v1r1. https://doi.org/10.6028/NIST.SP.800-160v1r1

[26] Hild D, McEvilley M, Winstead M (2021) Principles for Trustworthy Design of Cyber Physical Systems. MITRE Technical Report, MTR210263. https://apps.dtic.mil/sti/trecms/pdf/AD1146447.pdf

[27] International Organization for Standardization/International Electrotechnical Commission/Institute of Electrical and Electronics Engineers (2015) ISO/IEC/IEEE 15288:2015 – Systems and software engineering – Systems life cycle processes. https://www.iso.org/standard/63711.html

[28] DMTF DSP0274, Security Protocol and Data Model (SPDM) Specification 1.4, available at https://www.dmtf.org/standards/spdm
[29] Cooper D, Polk W, Regenscheid A, Souppaya M (2011) BIOS Protection Guidelines: Recommendations of the National Institute of Standards and Technology (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-147. https://doi.org/10.6028/NIST.SP.800-147
[30] Temoshok D, Abruzzi C, Choong Y, Fenton J, Galluzzo R, LaSalle C, Lefkovitz N, Regenscheid A, Vachino M (2025) Digital Identity Guidelines: Identity Proofing and Enrollment, (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-63A-4, https://doi.org/10.6028/NIST.SP.800-63A-4
[31] International Organization for Standardization/International Electrotechnical Commission (2026) ISO/IEC 15408-series, Information security, cybersecurity and privacy protection — Evaluation criteria for IT security (ISO, Geneva, Switzerland). Available at https://www.iso.org/standard/15408-1
[32] National Institute of Standards and Technology (2019) Security Requirements for Cryptographic Modules. (Department of Commerce, Washington, DC), Federal Information Processing Standards Publication (FIPS) NIST FIPS 140-3. https://doi.org/10.6028/NIST.FIPS.140-3
[33] IEEE Standards Association, *IEEE Std 2883 – IEEE Standard for Sanitizing Storage* (IEEE Standards Association, Piscataway, New Jersey). Available at https://standards.ieee.org/ieee/2883/10277/
[34] National Security Agency/Central Security Service (2020) NSA/CSS Storage Device Sanitization and Destruction Manual. NSA/CSS Policy Manual 9-12, December 4, 2020. Available at https://www.nsa.gov/Helpful-Links/NSA-FOIA/Declassification-Transparency-Initiatives/NSA-CSS-Policies/
[35] IEEE Standards Association, *IEEE Std 1619-2025 – IEEE Standard for Cryptographic Protection of Data on Block-Oriented Storage Devices* (IEEE Standards Association, Piscataway, New Jersey). Available at https://standards.ieee.org/ieee/1619/11552/
[36] IEEE Standards Association, *IEEE Std 1619.1-2018 – IEEE Standard for Authenticated Encryption with Length Expansion for Storage Devices* (IEEE Standards Association, Piscataway, New Jersey). Available at https://standards.ieee.org/ieee/1619.1/7153/
[37] Barker E, Kelsey J (2015) Recommendation for Random Number Generation Using Deterministic Random Bit Generators. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-90A Revision 1. https://doi.org/10.6028/NIST.SP.800-90Ar1

Appendix A. List of Symbols, Abbreviations, and Acronyms

Selected acronyms and abbreviations used in the document are defined below.

ACL
Access Control List

API
Application Programming Interface

BC
Business Continuity

BCM
Business Continuity Management

BIOS
Basic Input/Output System

CDN
Content Delivery Networks

CHAP
Challenge-Handshake Authentication Protocol

CIFS
Common Internet File System

CISO
Chief Information Security Officer

CLI
Command Line Interface

CPU
Central Processing Unit

CSO
Chief Security Officer (interchangeable with CISO)

CSI
Container Storage Interface

CTO
Chief Technology Officer

DAS
Directly Attached Storage

DDoS
Distributed Denial of Service

DMA
Direct Memory Access

DoS
Denial of Service

DR
Disaster Recovery

FC
Fibre Channel

FCIP
Fibre Channel over IP

FIPS
Federal Information Processing Standards

GNS
Global Name Server

FTP
File Transfer Protocol

GPS
Global Positioning System

HBA
Host Bus Adapter

HDD
Hard Disk Drives

HIPAA
Health Insurance Portability and Accountability Act

HPC
High-Performance Computing

HTTP
HyperText Transfer Protocol

HTTPS
HyperText Transfer Protocol Secure

I/O
input/output

IP
Internet Protocol

IS
Information Security

iSCSI
Internet Small Computer Systems Interface

ISL
Inter Switch Link

iSNS
Internet Storage Name Service

IT
Information Technology

KPI
Key Performance Indicator

LAN
Local Area Network

LDAP
Lightweight Directory Access Protocol

LUN
Logical Unit Number

LVM
Logical Volume Manager

MAC
Message Authentication Codes

MFA
Multi-Factor Authentication

NAS
Network-Attached Storage

NDMP
Network Data Management Protocol

NFS
Network File System

NIC
Network Interface Card

NS
Name Server

NTP
Network Time Protocol

NVM
Non-Volatile Memory

NVMe
Non-Volatile Memory express

NVMe-oF
NVMe over fabrics

OS
Operating System

OSI
Open Systems Interconnection

PBKDF2
Password-Based Key Derivation Function 2

PCI
Payment Card Industry; Peripheral Component Interconnect

PCIe
PCI express

PCI-DSS
Payment Card Industry Digital Security Standard

PII
Personally Identifiable Information

PKI
Public-Key Infrastructure

pNFS
Parallel NFS

QA
Quality Assurance

QoS
Quality of Service

RDMA
Remote Direct Memory Access

RFID
Radio-Frequency Identification

RPO
Recovery Point Objective

RSH
Remote Shell

RTO
Recovery Time Objective

SAN
Storage Area Network

SCSI
Small Computer System Interface

SHA
Secure Hash Algorithm

SIEM
Security Information and Event Management

SMB
Server Message Block

SMI-S
Storage Management Initiative Specification

SNMP
Simple Network Management Protocol

SNS
Simple Name Server

SoC
System on Chip

SPKM
Simple Public-Key GSS-API Mechanism

SRP
Server Routing Protocol

SSD
Solid-State Drive

SSH
Secure SHell

SSO
Single Sign-On

TCP
Transmission Control Protocol

TLS
Transport Layer Security

UDP
User Datagram Protocol

UI
User Interface

USB
Universal Serial Bus

VLAN
Virtual LAN

VM
Virtual Machine

VPC
Virtual Private Cloud

VPN
Virtual Private Network

WAN
Wide-Area Network

WORM
Write Once Read Many

WWN
World Wide Name

WWPN
World Wide Port Name

XML
eXtensible Markup Language

Appendix B. Glossary

block
A unit in which data is stored and retrieved on storage media. [1]

block storage
A method of storing data in blocks. [1]

Fibre Channel
A serial I/O interconnect capable of supporting multiple protocols. [1]

incident
Anomalous or unexpected event, set of events, condition, or situation at any time during the life cycle of a project, product, service, or system. [27]

Input/output
The process of moving data between a computer system's main memory and an external device or interface. [1]

Internet Small Computer Systems Interface
A transport protocol that provides for the SCSI protocol to be carried over a TCP based IP network, standardized by the Internet Engineering Task Force and described in RFCs 791, 1122, 2003, and 3720. [1]

network-attached storage
A storage device that connects to a network and provides file access services to clients. [1]

point in time copy
A fully usable copy of a defined collection of data that contains an image of the data as it appeared at a single instant in time. [1]

A point in time copy is considered to have logically occurred at that point in time, but implementations may perform part or all of the copy at other times (e.g., via database log replay or rollback) as long as the result is a consistent copy of the data as it appeared at that point in time.

point of encryption
Location within the Information and Communications Technology (ICT) infrastructure where data are encrypted on its way to storage (3.43) and, conversely, where data are decrypted when accessed from storage. [6]

recovery
The recreation of a past operational state of an entire application or computing environment. [1]

Recovery point objective
The maximum acceptable time period prior to a failure or disaster during which changes to data may be lost as a consequence of recovery. [1]

Recovery time objective
The maximum acceptable time period required to bring one or more applications and associated data back from an outage to a correct operational state. [1]

security domain
A domain that implements a security policy and is administered by a single authority. [13]

snapshot
A point in time copy of a defined collection of data. [1]

Clones and snapshots are full copies. Depending on the system, snapshots may be of files, LUNs, file systems, or any other type of container supported by the system.

storage area network
A serial I/O interconnect capable of supporting multiple protocols. [1]

storage virtualization
The application of virtualization to storage services or devices. [1]

trust
A belief that an entity meets certain expectations and, therefore, can be relied upon. [26]

Note: The term *belief* implies that trust may be granted to an entity whether the entity is trustworthy or not.

trust relationship
An agreed-upon relationship between two or more system elements that is governed by criteria for secure interaction, behavior, and outcomes relative to the protection of assets. [25]

Note: This refers to trust relationships between system elements implemented by hardware, firmware, and software.

vulnerability
A weakness that can be exploited or triggered to produce an adverse effect.

The inability to withstand adversity. [25]

Note: Vulnerability can exist anywhere throughout the life cycle of a system, such as in the CONOPS, procedures, processes, requirements, design, implementation, utilization, and sustainment of the system.

Appendix C. Controls List and Mapping of Mitigating Controls to Storage-Oriented Threats

This appendix consists of two tables:

- Table 1 provides a comprehensive list of all mitigating controls provided in this document, including their labels and titles.
- Table 2 provides a mapping of the relevant controls in this document to the storage-oriented threats identified in Sec. 3.1.

Not all the controls are listed in the following table as some are related to generic ICT threats.

Table 1. Comprehensive list of control labels and titles

Control Label	Control Description
Data Storage Governance (DSGV) (4.1)	
DSGV-PP-C01	Comprehensive storage security policy
DSGV-PP-C02	Keep the storage security policy current
DSGV-PP-C03	Compliance with storage security policy
DSGV-PP-C04	Data protection plan or policy
DSGV-PP-C05	Completeness of data protection plan or policy
DSGV-PP-C06	Backup plan or policy
DSGV-PP-C07	Ensure the completeness of recovery copies
DSGV-PP-C08	Protect all dependent ICT components
DSGV-PP-C09	Document the DR plan, resources, mapping to production, flow, and test procedures
DSGV-PP-C10	Storage consideration for policy
DSGV-PP-C11	Include storage infrastructure in logging and monitoring policy
DSGV-DT-C01	Include storage in BCM planning/testing
DSGV-DT-C02	Document data retention and sanitization obligations
DSGV-DT-C03	Document and verify storage compliance with privacy obligations
DSGV-DT-C04	Document and verify storage compliance with legal obligations
DSGV-SR-C01	Create a comprehensive inventory of all storage devices
DSGV-SR-C02	Create a comprehensive inventory of all data and configuration assets
DSGV-SR-C03	Network topology documentation
DSGV-SR-C04	Detect unauthorized storage security changes
DSGV-SR-C05	Develop a response plan for storage component compromise
DSGV-SR-C06	Develop a change management process for storage
DSGV-SR-C07	Perform periodic isolation reviews
DSGV-SR-C08	Set RPO and RTO
DSGV-SR-C09	Meeting organizational frequency and retention objectives
DSGV-PC-C01	Storage security training
DSGV-PC-C02	Ensuring adequate storage data protection expertise
DSGV-PC-C03	Ensuring adequate storage security expertise
Physical Controls for Data Storage (PSDS) (4.2)	
PSDS-PS-C01	Physically securing storage media
PSDS-PS-C02	Physically securing storage devices
PSDS-PS-C03	Use media security measures
PSDS-PS-C04	Protect all sensitive administrative equipment
PSDS-PS-C05	Use of storage sanitization
PSDS-PS-C06	Use of immutable storage
PSDS-PI-C01	Protect physical management interfaces to storage
PSDS-PI-C02	Protect physical SAN interfaces

Control Label	Control Description
PSDS-IS-C01	Separation of storage systems
PSDS-IS-C02	Separation of management systems
PSDS-IS-C03	Isolation of storage systems
PSDS-IS-C04	Logically isolating storage systems
Storage Systems Security (DSSS) (4.3)	
DSSS-GN-C01	Defense in depth for storage
DSSS-GN-C02	Establish security domains for storage
DSSS-GN-C03	Align storage security domains with purpose
DSSS-SH-C01	Perform basic operating system hardening
DSSS-SH-C02	Use software updates and patches from trusted sources
DSSS-SH-C03	Limit network access to management ports of SAN switches
DSSS-SH-C04	Limit storage administrative capabilities
DSSS-SH-C05	Favor API access control over CLI/shell access
DSSS-SH-C06	Restrict management consoles OS privileges
DSSS-SH-C07	Management web user interface
DSSS-SH-C08	Secure storage management interfaces
DSSS-SH-C09	Secure the remote management
DSSS-SH-C10	Restrict vendor maintenance interfaces
DSSS-SH-C11	Restrict vendor maintenance communications
DSSS-SH-C12	Restrict vendor dial-up maintenance
DSSS-SH-C13	Secure Intelligent Platform Management Interface (IPMI)
DSSS-SH-C14	Restrict host storage control privileges
DSSS-SH-C15	Command device or gatekeeper configuration
DSSS-SH-C16	Disable or limit call home or remote access
DSSS-SH-C17	Access control for management networks
DSSS-SH-C18	Secure and protect core storage management files and binaries
DSSS-SH-C19	Use an approved PKI mechanism for management access
DSSS-SH-C20	Perform regular security audits of data
DSSS-SH-C21	Disallow cleartext protocols
DSSS-SH-C22	Encryption for storage management API sessions
DSSS-SH-C23	Encryption for administrative access sessions
DSSS-SH-C24	Enable FIPS mode for FIPS-based environments
DSSS-SA-C01	Configure and use logging on storage infrastructure
DSSS-SA-C02	Ensure completeness of storage logging
DSSS-SA-C03	Monitor storage infrastructure
DSSS-SA-C04	Protect and retain storage logs
DSSS-SA-C05	Enable audit logging
DSSS-SA-C06	Reliable, external time synchronization
DSSS-SA-C07	Collect logs in a centralized fashion
DSSS-SA-C08	Auding logging level
DSSS-SA-C09	Audit log retention and protection
DSSS-SA-C10	SIEM integration
DSSS-VM-C01	Include storage infrastructure in vulnerability management programs
DSSS-VM-C02	Software updates and patches
DSSS-VM-C03	Scan files containing sensitive information with anti-malware tools on-access
DSSS-AM-C01	Unique Identifier for all users
DSSS-AM-C02	Use centralized authentication solution
DSSS-AM-C03	Configuration of authentication servers
DSSS-AM-C04	Secure connection to centralized authentication server

Control Label	Control Description
DSSS-AM-C05	Use of multi-factor authentication
DSSS-AM-C06	Entity authentication
DSSS-AM-C07	Secure password policies should cover service accounts
DSSS-AM-C08	Use of accounts not associated with system users
DSSS-AM-C09	Account lockout
DSSS-AM-C10	A local user account for emergency purposes
DSSS-AM-C11	Eliminate or disable default user accounts
DSSS-AM-C12	Limit local and default user accounts
DSSS-AM-C13	Roles and responsibilities configuration
DSSS-AM-C14	Separation of duties
DSSS-AM-C15	The privileges assigned to any role should adhere to the principle of “least privilege”
DSSS-AM-C16	Secure Session Management
DSSS-AM-C17	Implement a “message of the day” and “login banner” notice
DSSS-AM-C18	Separate Security and Non-security Roles
DSSS-AM-C19	Implement the principle of least-privilege
DSSS-AM-C20	Limit the access rights of service accounts
DSSS-AM-C21	Authenticate and authorize all CLI/API access
DSSS-PR-C01	SNMP security
DSSS-PR-C02	Control of IP addresses used for SNMP
DSSS-PR-C03	The authenticity of directory, domain, and similar services (e.g., AD, Domain Name System (DNS), LDAP)
DSSS-PR-C04	Considerations for using standard and non-standard TCP/IP or UDP ports
DSSS-PR-C05	Use of Network Data Management Protocol (NDMP) security features
DSSS-PR-C06	Use TLS in LDAP
DSSS-PR-C07	Additional protocols
DSSS-PR-C08	Use of TLS for encrypted communications
DSSS-PS-C01	Support for secure initialization
DSSS-PS-C02	Protection of system secrets
DSSS-PS-C03	Securing system BIOS
DSSS-PS-C04	Protection and update of firmware code
DSSS-PS-C05	Disable debugging ports
DSSS-PS-C06	Sensitive information in telemetry/debugging logs
DSSS-PS-C07	Enhance trustworthiness with attestation
DSSS-PS-C08	Anatomy of a trusted attester
DSSS-PS-C09	Use SPDM for storage attestation
Data Protection and Recovery (DPRC) (4.4)	
DPRC-GN-C01	Data protection configuration management
DPRC-GN-C02	Design Data Protection for Recovery
DPRC-SB-C01	The availability of all relevant software and hardware components
DPRC-SB-C02	The elected backup and data copies technologies and media should match organizational RTO
DPRC-SB-C03	Test restore to ensure RTO
DPRC-SB-C04	Validate health of backup copies
DPRC-SB-C05	Enable the separate restoration of data and application
DPRC-SB-C06	Use data backup measures and operations securely
DPRC-RM-C01	Synchronous and asynchronous replication
DPRC-RM-C02	Eliminate unnecessary replication trust between storage devices
DPRC-RM-C03	Secure replication/mirroring traffic
DPRC-RM-C04	Automated I/O suspension
DPRC-RM-C05	Eliminate obsolete replicas

Control Label	Control Description
DPRC-SS-C01	Configure point-in-time copies
DPRC-SS-C02	Eliminate obsolete snapshots/clones
DPRC-SS-C03	Snapshot security
DPRC-CR-C01	Disable all unnecessary services and protocols
DPRC-CR-C02	Recovery assets immutability during incident management
DPRC-CR-C03	Validate the hygiene of recovered compute components
DPRC-CR-C04	Access restriction to cyber-attack recovery system
DPRC-CR-C05	Off-site storage
DPRC-CR-C06	Use of an independent, full baseline copy
DPRC-CR-C07	Independence from hosts and applications
DPRC-CR-C08	Consider setting up an air gap
DPRC-CR-C09	Cyber hygiene of data copies
DPRC-CR-C10	Perform periodic audits
DPRC-CR-C11	Secure cyber-attack recovery backups
Storage Networking and Technologies (SNTC) (4.5)	
SNTC-BK-C01	Block-device access control
SNTC-BK-C02	Block-device copy and replica access control
SNTC-BK-C03	Use zoning is SAN-switched fabrics
SNTC-BK-C04	Software zoning implementation
SNTC-BK-C05	Controlling devices that can join fabric
SNTC-BK-C06	FC host and switch authentication
SNTC-BK-C07	The use of an approved FC PKI mechanism
SNTC-BK-C08	Use Fibre Channel zoning
SNTC-BK-C09	A blended approach to FC zoning
SNTC-BK-C10	FC masking Recommendation
SNTC-BK-C11	A backup of FC switch configuration data
SNTC-BK-C12	Limit FC switch management capabilities to the minimum necessary
SNTC-BK-C13	Limit communication between FC switches
SNTC-BK-C14	Disable unused FC SAN ports
SNTC-BK-C15	Audit FC SAN security configuration
SNTC-BK-C16	Confidentiality protections for Fibre Channel storage
SNTC-BK-C17	Sanitization for Fibre Channel-based storage
SNTC-BK-C18	Use NVMe-oF authentication
SNTC-BK-C19	Use NVMe/FC Security
SNTC-BK-C20	Segregate iSCSI networks
SNTC-BK-C21	Secure iSCSI networks
SNTC-BK-C22	Secure FCIP networks
SNTC-BK-C23	Secure NVMe/RDMA networks
SNTC-BK-C24	Confidentiality protections for IP storage
SNTC-BK-C25	Sanitization for IP-based storage
SNTC-BK-C26	IP storage network separation
SNTC-BK-C27	Isolate IP/Ethernet management ports
SNTC-BK-C28	Enable device IP access control
SNTC-BK-C29	Enable network IP access control
SNTC-BK-C30	Limit iSCSI ports
SNTC-BK-C31	Use iSCSI authentication
SNTC-BK-C32	Use NVMe/TCP Security
SNTC-FL-C01	NFS network access and protocols
SNTC-FL-C02	Apply NFS authentication

Control Label	Control Description
SNTC-FL-C03	Apply NFS access controls
SNTC-FL-C04	Secure NFS client access
SNTC-FL-C05	Secure data on NFS servers
SNTC-FL-C06	Secure NFS by restricting root access
SNTC-FL-C07	NFS “read only” mode
SNTC-FL-C08	NFS export of administrative file systems should not be allowed
SNTC-FL-C09	Disable NFS non-authenticated access
SNTC-FL-C10	Apply SMB networking controls
SNTC-FL-C11	Apply SMB access controls
SNTC-FL-C12	Secure data on SMB servers
SNTC-FL-C13	Disable insecure versions of SMB
SNTC-FL-C14	Restrict SMB “full control” permissions
SNTC-OB-C01	Permission for default zone
SNTC-OB-C02	Non-public storage objects
SNTC-OB-C03	Public access
SNTC-OB-C04	Restricting access to object storage data of all types (e.g., Files, Objects) to the minimum possible
SNTC-OB-C05	Use object protection to prevent unauthorized deletion
SNTC-OB-C06	Secure transfer of objects
SNTC-OB-C07	Confidentiality protections for stored objects
SNTC-OB-C08	Sanitization for object-based storage
Encryption and Key Management (ENKM) (4.6)	
ENKM-GN-C01	Use encryption to protect sensitive, regulated, and high-value data
ENKM-GN-C02	Storage encryption provides limited confidentiality protection
ENKM-GN-C03	Use encryption to protect administrative/management control of storage
ENKM-GN-C04	Limiting operational impacts of cryptography
ENKM-GN-C05	Security Strength of Cryptography
ENKM-GN-C06	Use Validated Cryptographic Modules
ENKM-GN-C07	Complying with import/export regulation for cryptography
ENKM-GN-C01	Complying with import/export regulation for cryptography
ENKM-EN-C01	At-rest encryption of sensitive data
ENKM-EN-C02	Using an appropriate point of encryption
ENKM-EN-C03	Creating appropriate proof of encryption
ENKM-EN-C04	Using data reduction technologies with encryption
ENKM-ET-C01	Use of end-to-end encryption
ENKM-ET-C02	Data in transit encryption
ENKM-ET-C03	Communication between storage system components should be encrypted
ENKM-KM-C01	Protect keys used to encrypt storage
ENKM-KM-C02	Recommendations for Encryption key management
ENKM-KM-C03	KM compatible with data retention and preservation
ENKM-KM-C04	Use centralized key management
Storage Disposal and Reuse (SDRU) (4.7)	
SDRU-GN-C01	Storage disposal considerations
SDRU-GN-C02	Secure disposal of storage
SDRU-GN-C03	Third-party collection and disposal services
SDRU-GN-C04	Accumulation of storage for disposal

2925

2926

2927 **Table 2. Mapping of storage-oriented threats to mitigating controls**

Storage-Oriented Threats	Mitigating Controls
Infection of Malware and Ransomware (Sec. 3.1.1)	**(4.1.1) DSGV-PP-C04** Data protection plan or policy **(4.1.1) DSGV-PP-C05** Completeness of data protection plan or policy **(4.1.1) DSGV-PP-C06** Backup plan or policy **(4.1.1) DSGV-PP-C07** Ensure the completeness of recovery copies **(4.1.1) DSGV-PP-C09** Document the DR plan, resources, mapping to production, flow, and test procedures **(4.1.3) DSGV-SR-C04** Detect unauthorized storage security changes **(4.1.3) DSGV-SR-C05** Develop a response plan for storage component compromise **(4.2.3) PSDS-IS-C01** Separation of storage systems **(4.2.3) PCDS-IS-C03** Isolation of storage systems **(4.2.3) PCDS-IS-C04** Logically isolating storage systems **(4.3.2) DSSS-SA-C01** Configure and use logging on storage infrastructure **(4.3.2) DSSS-SA-C03** Monitor storage infrastructure **(4.3.3) DSSS-VM-C03** Scan files containing sensitive information with anti-malware tools on-access **(4.4.1) DPRC-SB-C06** Use data backup measures and operations securely **(4.4.4) DPRC-CR-C08** Consider setting up an air gap **(4.4.4) DPRC-CR-C11** Secure cyber-attack recovery backups
Unpatched Vulnerabilities (Sec. 3.1.2)	**(4.3.1) DSSS-SH-C01** Perform basic operating system hardening **(4.3.1) DSSS-SH-C02** Establish security domains for storage **(4.3.2) DSSS-SA-C01** Ensure completeness of storage logging **(4.3.2) DSSS-SA-C03** Monitor storage infrastructure **(4.3.3) DSSS-VM-C01** Include storage infrastructure in vulnerability management programs **(4.3.3) DSSS-VM-C02** Software updates and patches
Physical Theft or Inappropriate Disposal of Storage Media (Sec. 3.1.3)	**(4.2.1) PSDS-PS-C01** Physically securing storage media **(4.2.1) PSDS-PS-C02** Physically securing storage devices **(4.2.1) PSDS-PS-C03** Use media security measures **(4.2.1) PSDS-PS-C05** Use of storage sanitization **(4.7) SDRU-GN-C01** Storage disposal considerations **(4.7) SDRU-GN-C02** Secure disposal of storage
Platform Security Compromises (Sec. 3.1.4)	**(4.3.2) DSSS-SA-C01** Configure and use logging on storage infrastructure **(4.3.2) DSSS-SA-C03** Monitor storage infrastructure

2928

Appendix D. Change Log

This document revises SP 800-209 (2020) as follows:

Section 2: Data Storage Technologies Background

- Content in Sec. 2 and elsewhere that was associated with obsolete and/or niche technologies (e.g., Fibre Channel over Ethernet, software defined storage, hyper-converged storage, continuous data protection) has been removed.

Section 3: Data Storage Technologies Background

- Section 3 was revised to primarily focus on the threats and risks that have particular relevance to storage systems and ecosystems. References to other NIST publication are provided for more generic ICT threats and risks.
- The phrase “attack surfaces" was removed from the title, and the associated subsection (Sec. 3.3) was removed. The removed attack surface descriptions were generic and duplicated content in other NIST publications.
- Section 3.1.4, “Platform Security Compromises,” has been added to focus on threats associated with the underlying/supporting infrastructure for data storage.
- Section 3.2.4, “Compromised Data Resilience/Protection,” has been added to focus on threats to secondary data assets and the secondary data asset generation process (i.e., backups).

Section 4: Data Storage Technologies Background

- The title has been changed from “Security Guidelines for Storage Deployments” to “Storage Security Controls,” and the recommendations have been reorganized from nine categories into seven formalized control families with standardized control acronyms.
- The guidance in Sec. 4 has been restructured to improve readability.
- Generic security controls (e.g., password-specific controls) were replaced with references to other NIST publications.
- A new label scheme was applied to the controls in Sec. 4, and these new labels (and short descriptions) are leveraged in the new Appendix C.
- Section 4.1, “Data Storage Governance (DSGV),” has been added to address policy, planning, storage resource management, and people controls. It incorporates some of the configuration management concepts discussed in the original document.
- Section 4.2, “Physical Controls for Data Storage (PSDS),” consolidates the two former sections on physical storage security and isolation.
- Section 4.3, “Storage Systems Security (DSSS),” now encompasses storage system hardening, security auditing and monitoring, vulnerability management, and platform security. This effectively merges the former stand-alone sections on authentication and data access control, administrative access, and audit logging.

- Section 4.4, “Data Protection and Recovery (DPRC),” replaces and merges the previous data protection and restoration assurance sections.
- Section 4.4.4, “Cyber Recovery Systems,” is a new subsection that addresses incident response measures under the umbrella of “cybersecurity program management” but focuses on storage infrastructure and data assets.
- Section 4.5, “Storage Networking and Technologies (SNTC),” categorizes security controls grouped by block-based, file-based, and object-based storage networking architectures.
- Section 4.6, “Encryption and Key Management (ENKM),” expands upon the original encryption section by adding controls related to key management and associated infrastructure.
- Section 4.7, “Storage Disposal and Reuse (SDRU),” is entirely dedicated to the secure disposal and sanitization of storage media.
- Appendix C was added to provide a complete list of storage security controls as well as a mapping of mitigating controls for the storage-oriented threats identified in Sec. 3.