The 2025 Security Maturity of Storage & Data Protection Systems

CONTINUITY

Assessing The Security Risks of Enterprise Storage & Backup Systems

Building on the tradition established in 2021, Continuity is proud to release its 2025 research report, offering an in-depth analysis of the security risks affecting enterprise storage and data protection systems.

Among the three core pillars of modern IT – Compute, Network, and Storage – Storage has historically received the least attention from InfoSec teams. This is largely because most cyberattacks have traditionally focused on the Compute and Network layers, leaving storage systems comparatively overlooked.

This has quickly changed in recent years, with exponential growth in data-targeted attacks, such as ransomware, data exfiltration, and data destruction. Threat actors have become increasingly proficient in targeting storage, backup and data protection systems – either as a primary target (e.g., in order to exfiltrate large amounts of data "under the radar"), or a secondary one (e.g., destroy all backup copies prior to activating a ransomware payload).

Modern security frameworks (such as NIST and ISO) have adapted to address this gap*, and international regulation is also maturing and evolving accordingly**.

We compiled anonymized inputs from a large number of storage and data protection risk assessments performed in 2024, using Continuity's flagship solution, StorageGuard. This provided a unique insight into the security maturity of storage and data protection systems.

The analyzed data covers multiple storage and data protection vendors and models – including **Dell Technologies, IBM, Hitachi Vantara, NetApp, Pure, Infinidat (Lenovo), Cohesity (Veritas), Rubrik, Commvault, Cisco, Brocade (Broadcom)**, and others.

StorageGuard check for thousands of possible security misconfigurations and vulnerabilities at the storage, backup, and data protection systems level that pose a security threat to enterprises' data.

In preparation of this report, thousands of discrete security risks that were detected in the risk assessments were reviewed, allowing us to uncover recurring patterns and important security considerations many organizations fail to get right when managing storage and data protection environments.

*See NIST SP 800-209 Security Guidelines for Storage Infrastructure and ISO 27040 Security techniques — Storage security

** DORA, PCI 4.0.x.

Key Findings

6,085 discrete security issues were analyzed.

323 environments assessed, with 11,435 storage & data protection devices, of which 627 were selected for analysis (*)

An enterprise storage & data protection device has on average 10 vulnerabilities

Out of 10 vulnerabilities, 5 are high or critical risk

The 5 most common areas for security risk include:

01 Authentication and Identity management

02 Unaddressed CVEs

03 Network and Protocol Security

04 Encryption and Key Management

05 Access Control and Authorization

- The research scope has slightly increased compared to the previous report with 11,425 storage and data protection devices** across 323 organizational IT environments in North America and EMEA.
- 53% of organizations were from the Banking, Insurance, and Financial Services sectors. The remaining 47% included a broad range of industries, including Technology, Logistics, Transportation, Telecommunication, Construction, and Postal Services.
- **627 enterprise storage & data protection devices were analyzed** (similar to our previous report), and a total of 6,085 discrete security vulnerabilities and misconfigurations were detected, spanning more than 390 security principles that were not adequately followed. Most frequent, and other significant findings are discussed in more detail below.
- On average, an enterprise storage & data protection device **has 10 security risks**, out of which 5 were of high or critical risk rating (i.e., could present significant compromise if exploited). This finding is slightly higher than previous years***
- As with previous reports, there was little correlation between geographic location and security maturity. In other words, the frequency and severity of issues remained consistent across environments – regardless of where they were located.
- **We didn't detect any significant correlation between industry and security maturity.** Although it is commonly accepted that certain industries, like financial services, tend to have more mature security strategies, this report shows that the entire field of storage & data protection security across all industries is still overlooked. While this was similar to previous years' findings, it is still surprising.

* To prevent any bias, device selection was performed by the organizations who participated in the risk assessments (and not Continuity). Each organization was asked to choose a representative sample from each of their environments.

** See "Methodology" section for more details about the types of devices covered in this report

*** As further described in the Methodology section below, a more refined classification of CVE severity was used which explains the increase.

Top five security risk categories found in this year's analysis:

01	Authentication and Identity management
02	Unaddressed CVEs
03	Network and Protocol Security
04	Encryption and Key Management
05	Access Control and Authorization

In addition to the five most common risk categories, we included in this report a few notable ones that are of particularly high risk, including:

06	Incorrect use of ransomware-protection features
07	Insecure session management
08	End of support devices used in production
09	Vulnerabilities In Software Supply-Chain Management
10	Undocumented And Insecure API / CLI

Recommendations

The state of enterprise storage & data protection security is significantly lagging behind that of compute and network security.

This is a significant gap that should be addressed as soon as possible; with growing sophistication of data-centric attacks, and with tightened regulations, the business implications of ineffective security could rapidly increase.

Determine if knowledge gaps exist in terms of storage & data protection security, and build a plan to address them

---

Review your organization security maturity, considering the industry gaps identified in this report, and strengthen your security program as needed

---

Proactively address risks. Introduce automation to frequently assess the security posture of your storage and data protection systems. This could dramatically reduce the exposure.

Observations

Storage vs. Data Protection Systems maturity

We did not see a significant difference in the number of security misconfigurations or unresolved CVEs in data protection (e.g. backup) environments compared to storage systems; both suffer equally from a lack of hardening.

In the past few years, online backups were frequently targeted as part of cyber attacks. The motivation of adversaries in this regard is twofold:

Prevent recovery: deletion of data protection copies, or other manipulation to make them unusable (*), will prevent recovery of compromised data, which is critical to the adversary goal (e.g., in the case of ransomware, to force the victim's hand to pay the ransom)

Exfiltrate data: it is often much harder to exfiltrated data directly from actual production systems than it is from offsite copies (especially cloud). The former are typically protected by advanced access control, DLP solutions, anomaly detection, honeypots, and other security tools, whereas the latter are in many cases far less secured.

Note (*): other than deletion, data protection copies could be rendered unusable through re-encryption, poisoning, and other strategies. For more information discover the five backup lessons learned from the 2024 breach of UnitedHealth's backup environment

Our recommendation is to maintain an immutable, preferably offline copy of backups, to test them regularly to ensure viability, and to make sure they are configured with at least the same level of security as their source production data.

International Standards and Regulations

International standard organizations have recognized the importance of storage and data protection, and recent years have witnessed the publication of NIST SP 800-209 ("Security Guidelines for Storage Infrastructure"), and ISO/IEC 27040:2024 ("Security techniques — Storage security").

Global industry standards such PCI-DSS v4.0.x are evolving as well to put more emphasis on data protection, IT resilience, and testing. Finally, State and Federal laws are similarly forming and evolving to define much stricter expectations and guidelines related to data protection. Most notable is DORA (the EU Digital Operational Resilience Act), but similar publications have also been formed in the US, the UK, and several APAC countries.

The significance of the new regulations and the notably enhanced bar they set to achieve storage and data protection security include:

- Adapting internal security programs to meet enhanced guidelines would significantly reduce the attack surface, and increase organizational resilience to data-targeted attacks
- Audited organizations should put greater emphasis in identifying relevant requirements, and adapting their risk management programs accordingly. Failure to demonstrate compliance could have significant financial implications.
- It could be highly beneficial to introduce automation around evidence gathering, as some frameworks emphasize the need to measure security outcomes.

Organizations are failing to plan for storage and data protection-targeted attacks

As already noted above, modern cyber security frameworks are clear about the need to tighten incident response plans, and introduce more robust testing – in particular as it relates to recovery from cyber incidents.

However, the sad reality is that many organizations fail to plan accordingly. The result is that many of the organizations experiencing a cyberattack in 2024 were NOT able to resume normal operations in a reasonable time. Some of the more notable examples include UnitedHealth which restored core services in a month, but the complete resolution of all issues, particularly those affecting healthcare providers, took several months following the cyberattack.

Of course, the direct and indirect financial implications of such long service disruptions are staggering. Some of the areas that could lead to notable improvement include:

- Better risk modeling, and incident response planning. Clearly define your resilience goals, in particular, the required recovery speed for each business service. Consider events such as compromise of a storage array (SAN or NAS), and a backup appliance.
- Ensure the tools used to protect data and recover applications are fit-to purpose (e.g., if you need a 2-hour recovery, offsite tapes are likely not the right solution…)
- Rigorously test data and application recovery. Document test results, as you may be required to present them following a breach, or during an audit.
- Consider the possibility that your most recent backup copies might be infected or tampered. Devise and implement solutions that will optimize recovery speed, and increase your chances of selecting a clean copy for recovery.

Ransomware Protection

More modern storage and data protection solutions contain mechanisms for ransomware protection. These include detection of anomalous and suspicious I/O patterns (e.g., mass encryption of files, appearance of known file prefixes, etc.), and in response - automated alerting, and reinforced retention of recovery copies. In many surveyed IT environments included in this report, such mechanisms were not configured, or were configured in an ineffective way.

Furthermore, many storage and data protection systems were not patched or updated to include the latest ransomware protection tools.

Immutability

In our analysis, we detected a year-on-year increase in the number of storage and data protection environments that were configured with immutable data copy technologies. This is, of course, a welcome trend. Immutability, if well-defined and configured, can greatly enhance the security posture of organizations. However, incorrect, or partial implementation can lead to a false sense of security, and, unfortunately, we did detect a significant number of misconfiguration issues in many deployments.

When misconfigured, it is possible to delete supposedly immutable data (for example, by manipulating time/date settings on the storage device to bypass retention enforcement mechanisms). Even when configured correctly, an attacker with access to the data source can poison an immutable data store over time, corrupting it such that it becomes useless when needed for recovery.

For more information on Poisoning, go to https://www.continuitysoftware.com/blog/dont-rely-on-immutable-backup-for-protection-against-ransomware/

Global Tension and Conflicts

Ransomware groups linked to state sponsors in conflict regions—such as Russia, Iran, and North Korea—have been implicated in numerous cyberattacks. Notable groups include Sandworm, Conti, Lazarus Group, and APT33.

A Microsoft Digital Security Unit report on the Russia-Ukraine conflict identified no fewer than eight malware families leveraged by Russia-aligned cybercriminals , most of which are designed for data encryption, destruction, or exfiltration. The same report called for heightened vigilance against further state-sponsored attacks against NATO countries, other states supporting Ukraine, and their constituent organizations or businesses.

Experts also indicate that Nation-state tools can often fall to the hands of "regular" criminals, who weaponizes them for large-scale data breaches, and ransom attacks. Finally, while nation-state actors may use novel tools, their avenues of attack often fall along familiar lines like spear-phishing that can be mitigated by standard cyber security protocols laid out by NIST, ISO, and others.

Note that adoption of these protocols is quickly becoming a prerequisite for cyber insurance coverage, discussed in further detail below.

Shared Responsibility Model 1: *Vendor vs. User*

Storage, backup, and data protection vendors ship their systems with a minimal base level of security configuration, one that is often insufficient for use in a production environment, and provide separate guidance for further hardening by the IT department.

It is the IT Infrastructure department's responsibility to develop their own security baseline following current industry best practices, determine how best to implement a configuration adhering to that baseline, and then ensure that configuration is enforced continuously over time – not just at initial deployment.

Shared Responsibility Model 2: *IT Infrastructure vs. Security*

We have noted a repeated pattern of division between IT infrastructure and security teams, whereby security teams develop security policies and procedures that IT infrastructure teams are tasked with implementing, sometimes with minimal direction.

Often, security teams are not aware of cyber resiliency capabilities offered by storage and data protection systems. While IT infrastructure teams are more focused on day-to-day operations and less concerned with reducing the potential for cyberattacks.

This division is underlined by our findings, which show the use of insecure protocols and unpatched CVEs continue to be the top security risks. These issues are among the most basic aspects of a strong data security posture. An opportunity now exists to increase the level of security literacy among the teams who manage data storage and data protection, while improving the storage-specific knowledge and toolsets available to security teams.

[1]Special Report: Ukraine – April 27, 2022, Microsoft Digital Security Unit

Background

Among the three core pillars of IT infrastructure—Compute, Network, and Storage—Storage holds a uniquely critical role from both a business and security standpoint. While a compromise in any of these layers can lead to downtime, an attack on the storage layer carries the highest risk of being **irreversible**, potentially resulting in permanent data loss.

Consider a scenario in which a coordinated attack on a bank succeeds in compromising both current and long-term customer financial records (e.g., attacking both primary storage and its protective copies, such as snapshots, backup, and archived copies). What would be the consequences for customers, for the bank itself, and for the economy?

We argue that the storage layer should be secured and hardened to a similar if not greater extent than that employed for the Compute and Network layers[2]. A comprehensive storage & data protection security practice should cover the entire lifecycle of data[3].

With growing industry and government attention to data storage & data protection security, resources are now available to guide organizations on building a secure storage management practice, including NIST SP-800-209 'Security Guidelines for Storage Infrastructure', ISO 27040, and a series of educational storage security papers by SNIA.

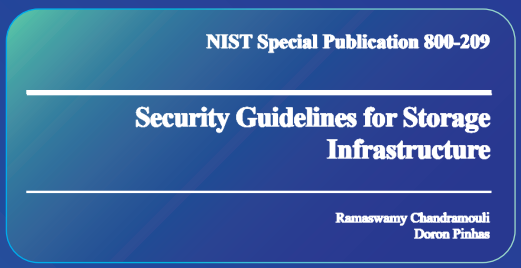
NIST Special Publication 800-209

Security Guidelines for Storage Infrastructure

Ramaswamy Chandramouli
Doron Pinhas

Given the growing evidence that new forms of malware and ransomware are specifically targeting storage and data protection systems, we came to realize it would be valuable to research and compile an industry benchmark for the state of storage & data protection security, to gauge the overall market maturity and to identify if common areas of weakness or oversight exist.

Encouraged by enthusiastic interest in this series of issues published since 2021, we are pleased to provide an updated analysis of the industry security maturity. This year we've expanded our analysis scope, highlighted the major trends in 2025, and explored similarities and differences between this report and previous ones.

It is our hope that these reports could help organizations increase awareness of this important area, help identify gaps in existing plans, and provide insights based on community data.

[2] While many of the principles involved in securing storage and backup are similar in nature to those used for compute and network infrastructure (e.g., authentication and authorization, access control, vulnerability management, etc.), certain aspects are unique to storage and backup. These include proper design, implementation and testing of data protection and recovery, securing storage protocols and storage networking, and data immutability features.

[3] Encompassing secure design, enforcement of security principles during all deployment and maintenance phases, comprehensive testing, and ongoing auditing, vulnerability assessment and anomaly detection.

Detailed Information About Key Security Risks

01 | Authentication & Identity Management

Most storage and data protection platforms support both the use of built-in Authentication and Identity Management ("AIM"), and of centralized solutions. In some implementations, both types of solutions may be used together (e.g., using Microsoft Active Directory authentication for data access and routine management, as well as built-in accounts for break-glass emergency administration).

The most frequently identified AIM-related misconfigurations we encountered can be roughly divided into two types:

- **Infrastructure-related -** such as insecure service configuration (e.g., weak, or no encryption, lack of hardening), lack of utilization of device-specific modern security features, such as MFA, OTP, Dual Control, unapproved or unnecessary administrative user accounts, specialized security accounts (e.g., security officers, command approvers), and restriction of data access to users authenticated only through centralized AIM.

- **Industry best-practice-related -** such as device-, or technology-specific recommendations for disabling inactive users, the use of MFA, expiry and lockout, password rules, etc.

While some of the identified misconfigurations are unique to storage infrastructure, most were not. It is somewhat surprising that in 2025, organizations are still failing to secure storage and data protection systems*. In addition to the questions this raises around process maturity, this could also indicate there's lack of automation, or lack of tools for security posture management of these infrastructure elements.

Business impact

Incorrect and insecure configuration can allow cybercriminals to take full control over storage and data protection systems, and enable them to exfiltrate and destroy the data – and its copies.

Recommendations

Lock and rename or delete factory default users, where possible

Eliminate the use of local user accounts – use centralized authentication mechanisms such as Active Directory or LDAP except for backup systems where local user accounts can be used for isolation purposes. If still needed, harden and restrict access to systems using local accounts (e.g., limit access to isolated network interfaces only, deny data access to local users)

Enable multifactor authentication

Enable and configure Dual Control over destructive workflows and operations (backup deletion, immutability setting changes, backup policy or schedule changes and so on)

Define target system IAM security settings

Isolate your backup systems at identity layer (in addition to the network and domain)

Configure systems for zero trust

Periodically assess the identity management settings of all storage and data protection devices to identify configuration drifts.

Additional recommendations can be found in this article: 'Top 15 Security Controls for Storage & Backup Systems': https://www.continuitysoftware.com/blog/top-15-security-controls-for-storage-backup-systems/

* For example, according to available public information, the notorious attack on United Health services last year, that resulted in weeks of service disruption, could have been avoided or significantly diminished by better enforcing AIM security best practices on its storage and data protection system.

02 | Unaddressed CVEs

Storage and data protection systems involve a surprisingly large number of software components that get routine updates, including:

- Storage arrays, backup appliances, and Fibre-Channel storage switches – all have Operating Systems, which are often proprietary, or highly specialized and restricted versions of commercial or open-source operating systems
- IO Controller (such as HBAs, FCoE, NVMEoF adapters) have dedicated firmware
- Management software suites have multiple components
- API servers (e.g., storage connectors for virtual environments)
- Client-side software (such as OS drivers, and agents) on hosts using storage and data protection services

Vulnerabilities for such devices and components are discovered on an ongoing basis, and Common Vulnerability and Exposure (CVE) records are accordingly published. In most cases, a fix in the form of an upgrade or configuration change is recommended.

Here's a selection of recent News Headlines related to exploited vulnerabilities in storage & backup systems:

Critical Dell Product Vulnerabilities Let Attackers Compromise Affected Systems

The vulnerabilities, identified as CVE-2024-37143 and CVE-2024-37144, impact various versions of Dell PowerFlex appliances, racks, custom nodes, InsightIQ, and Data Lakehouse products.

CVE-2024-37143, the more severe of the two, is an Improper Link Resolution Before File Access vulnerability. This flaw allows an unauthenticated attacker with remote access to execute arbitrary code on affected systems. With a CVSS score of 10.0, this vulnerability poses a critical threat to system security.

The second vulnerability, CVE-2024-37144, involves Insecure Storage of Sensitive Information. While it requires a high-privileged attacker with local access, it can lead to information disclosure. Exploiting this vulnerability may allow attackers to gain unauthorized access to pods within the cluster. This flaw has been assigned a CVSS score of 8.2, indicating its high severity.

NHS England Warns of Critical Veeam Vulnerability Under Active Exploitation

CVE-2024-40711: Critical Veeam Vulnerability Exploited in Frag Ransomware Attacks.

CVE-2022-26500 and CVE-2022-26501: These vulnerabilities allow remote, unauthenticated attackers to execute arbitrary code. They were actively exploited by ransomware groups like Monti and Yanluowang shortly after discovery, emphasizing the importance of timely patching.

CVE-2023-27532: This high-severity vulnerability allows attackers to bypass authentication and access sensitive data. It has been exploited by ransomware actors such as the ransomware operation known as EstateRansomware, showcasing the persistent threat to enterprise environments.

The National Health Service (NHS) noted that enterprise backup applications are valuable targets for cyber threat groups. Veeam noted that unsupported product versions are not tested.

Acronis Warns Of Critical-Severity Vulnerability Being Exploited In Their Storage And Cyber Protection Platform

A critical vulnerability in Acronis Cyber Infrastructure (ACI), tracked as CVE-2023-45249, was highlighted by CISA as being actively exploited by malicious actors.

This vulnerability allows threat actors to execute arbitrary code remotely due to the use of default passwords. Considering ACI is a secure storage solution, this exploited vulnerability has a double effect – it can put mass amount of production data at risk as well as jeopardize backup data – which will hinder cyber recovery.

Despite a patch being available for several months, many organizations are unaware and have not yet applied it, leading to ongoing exploitation in the wild.

The ALPHV Ransomware Operation Exploits Veritas Backup Exec Bugs For Initial Access

U.S. Cybersecurity and Infrastructure Security Agency (CISA) increased its list of security issues that threat actors have used in attacks, three of them in Veritas Backup Exec exploited to deploy ransomware.

CVE-2021-27876: This vulnerability allows unauthorized file access through the Backup Exec Agent.

CVE-2021-27877: This involves improper authentication, potentially allowing attackers to access sensitive information.

CVE-2021-27878: This vulnerability permits command execution, allowing attackers to run arbitrary commands on affected systems.

These vulnerabilities have been actively exploited, highlighting the risks associated with unpatched backup solutions.

Mainstream Vulnerability Management tools used by organizations **do not detect the majority of storage and data protection CVEs** (but rather focus on server OS, traditional network, and software products). The result is in an unacceptably large percentage of storage and data protection devices being exposed.

281 different CVEs were identified in the environments covered in this research (of course, thousands are documented), with an alarming 39% of the devices analyzed being exposed.

Business impact

Each CVE details the possible exposures and outcomes it presents – and these span a wide range. Among the risks identified were the ability to exfiltrate files, initiate denial-of-service attacks, and even take ownership of files and block devices.

Recommendations

Improve proactive CVE identification

use storage-specific tools to scan storage and data protection environments for CVEs, instead of server-specific vulnerability management tools that cannot identify storage and data protection platforms appropriately

Reduce remediation time for important vulnerabilities

identify and patch CVEs with critical and high CVSS scores as quickly as possible, using all relevant tools (in-house scans, vendor security announcements, etc.)

03 | Network & Protocol Security

Storage protocols span both traditional networking[4] (IP over Ethernet and WAN) and dedicated networking (such as Fibre-Channel and Infiniband media & protocols)[5]. It is critical to secure storage and data protection network settings both during session establishment, and while exchanging data. However, in too many cases, and in most storage and data protection environments, it is still common to find configuration gaps such as:

- Not disabling legacy versions of storage protocols, or even worse, defaulting to their use (e.g., SMBv1, NFSv3)
- Failure to configure redundancy for core services, such as DNS, NTP, Fibre-Channel (FC) transport, SNMP
- Failure to use secure versions, and to enforce protocol security options on storage and data protection protocols (such as FC, SAN Fabric Management protocols, NDMP, SNMP, replication)
- Failure to meet vendor network and protocol security best-practices
- And many others (e.g. allowing cleartext HTTP sessions, using unsecure SNMP community strings, enabling unnecessary services, etc.)
- Please note that network and protocol security is also largely influenced by encryption in-transit considerations – a topic covered in the following section, and therefore omitted here (these include the use of unrecommended or obsolete cyphers, failure to protect all data feeds, etc.)

Business impact

Cybercriminals can exploit configuration mistakes to retrieve configuration information and stored data, and in many cases, to also tamper with the data itself (modify, destroy, or lock), including the copies used to protect the data.

Failure to protect storage network management protocols could also result on Denial of Service Attacks on storage and data protection infrastructure.

Recommendations

Close knowledge gaps
refer to resources such as NIST 800-209, ISO 27040:2024, and SNIA to get familiar with storage and data protection network security concepts, risks, and best practices

Define internal requirements
to adapt industry recommendations to business requirements

Identify and remediate gaps
between requirements and actual settings

Build an effective, ongoing process
to continually evaluate the security posture of storage and data protection environments

[4] Mostly used for file and object storage, with a steadily growing use for block-storage

[5] Encompassing FC switches, FC protocols, and FC network management protocols

04 | Encryption & Key Management

Modern storage and data protection systems offer a wide range of encryption capabilities:

- In-transit – many storage protocols support both encrypted and non-encrypted forms, and the default settings are often insecure. Encryption should be considered both for protocols carrying the data itself, as well as for those used for management operations (such as path negotiation, and device configuration). Some data paths are well understood by infosec specialists (such as the ones connecting a host to a storage device, or to a backup tool), while others are less so (such as replication transport)
- At-rest – storage and data protection devices offer encryption for both data and configuration. The encryption can sometimes be implemented at multiple layers (e.g., disk, volume, array)
- In-use – some devices offer additional configuration options to safeguard against memory leaks.

Keying materials and encryption infrastructure should be protected against failure, and backed up in a secure way to facilitate recovery from disasters and cyber-attacks.

While certain improvements were observed when comparing the issues identified in this report to those uncovered in previous years. This is still an area plagued with multiple common misconfigurations, including:

- The use of no-longer recommended cypher suites (e.g., allowing TLS 1.0 and 1.1, not disabling SSL 2.0 and 3.0) – some of which must be disabled to comply with regulatory frameworks (e.g., PCI DSS)
- Not enforcing data encryption for some critical data feeds, such as management transport, replication transport, backup transport.
- Failure to protect data copies with equivalent (or higher) level of encryption to that used for the source data (which leaves the data exposed to exfiltration).
- Failure to meet vendor-specific encryption best-practices

Business impact

Lack of sufficient encryption in transit can expose the data to exfiltration and denial of service.

Lack of encryption at-rest can cause violation of industry standards such as PCI-DSS and HIPPA, and will leave physical media exposed if lost or stolen. It is important to note that media shipping is still a widely-used practice by many large organizations.

Failure to protect keying materials can lead to data loss when restoring data, as there may be no way to decrypt it.

Recommendations

Review voluntary and mandatory requirements applicable to your organization

such as NIST, ISO, CIS, PCI-DSS, HIPPA, DORA, HITRUST, NERC CIP, CRI, SWIFT, NCSC Cyber Essentials, NIS2, FCA/PRA, OSFI, AICPA Trust Services Criteria (TSC), DOD STIG, MAS TRM and others.

Review encryption best practices

published by vendors for each storage and data protection product family used by your organizations

Map dependencies between data sources

in particular between production data and copies kept by storage and data protection systems – such as snapshots, replicas, backup, sets, archived and offsite data – and make sure that the level of encryption used is appropriate, both in-transit and at-rest.

Define and maintain encryption baselines

for storage and data protection systems,

Use automation

to periodically evaluate and enforce encryption security baselines

08 | Access Control & Authorization (Over-Exposure)

Access control to storage and data protection systems includes several different configuration levels:

- Access to storage elements - such as block devices, network shares, or even individual files and objects should be granted only to designated resources (e.g., individual hosts or applications). This is done both at the device level (e.g., share configuration, LUN mapping) and by using network filtering techniques (e.g., IP filters, SAN zoning and Masking)
- Access to the data itself – is often configured for users, hosts, or applications. Different access attributes can be individually set, such as mode (e.g., read, write, and modify permissions), ownership, ACLs., etc.
- Access to advanced storage capabilities - e.g., management, control, replication, snapshot management

In addition, NIST SP 800-209 also recommends that the design and implementation of access control and authorization for Storage and data protection systems should be cognizant of different access-planes. These can include:

- The management plane – the set of access operations required to configure and manage the storage and data protection platforms (e.g., create volumes, define replication, configure backup tasks, execute data protection tasks, etc.)
- The data plane – the set of access operations required to read, write and process the actual content of storage and data protection systems
- The backup plane – the set of access operations required to make copies of data for the purpose of recovery from various data loss scenarios

It is recommended to isolate environments using separate roles, and dedicated access control settings for each: access plane, application and business service, and operational environment role within a business service (e.g., production, development, and testing, etc.)

It is further required to assess the access setting of dependent data objects – such as source data, and its backup copies – and make certain that dependent objects are at least as restricted as their source.

A large number of devices were affected by improper configuration, including unrestricted access to shared storage, unrecommended zoning and masking configuration, ability to reach storage elements from external networks, and more.

Business impact

Incorrect access control and authorization can at best lead to data exposure, and at worst to compromise of the data itself and its copies. In some cases, it can also lead to compromises of the operating systems and applications running on the hosts that use the storage.

Recommendations

Implement appropriate least-privilege access models

both for data access (file/LUN/object etc.), and management. Further separate between planes, applications, and environments

Make sure data copies are at least as secured

and restricted as the source data.

Audit and correct exposures

on a frequent basis. Use automation to detect deviations and enforce proper isolation.

In addition to the five most common risks, other risks that appeared less frequently but were classified as high priority, included:

06 Incorrect use of ransomware-protection features

07 Insecure session management

08 End of support devices used in production

09 Vulnerabilities In Software Supply-Chain Management

10 Undocumented And Insecure API / CLI

06 | Incorrect Use of Ransomware-Protection Features

Modern storage and data protection systems provide advanced ransomware detection and prevention capabilities. These include:

- The ability to recognize I/O patterns that would indicate an ongoing attack (e.g., identification of known file name and extension patterns, abnormal high-rate of file modification and deletion, etc.),
- Capabilities for locking retained copies, protecting critical data from tampering and deletion
- Certain forms of air-gapping

Our research revealed that these features are often overlooked or misunderstood. In many environments, even when properly licensed, they were found to be inactive. Moreover, even when enabled and in use, critical vendor best practices were frequently not followed—for example, retained immutable copies were not locked, time services were not hardened (potentially allowing attackers to manipulate retention expiration), and dual authorization for delete operations was improperly configured or entirely absent.

Business impact

With limited or no use of ransomware protection features, cybercriminals can easily circumvent or disable data protection mechanisms.

Recommendations

Implement and follow vendor best practice for ransomware-defense features

including data immutability capabilities where appropriate (backups, large NAS shares, etc.), hardening of core services, activation of advanced security features, such as security roles, and dual authorization, activation of air-gapping when appropriate.

Be aware of potential impact

while implementations differ, immutability features typically increase overall storage consumption with trickle-down effects on replication and performance

Data older than several weeks is typically less relevant for recovery

keep this in mind when architecting immutability and retention periods

Periodically assess the security

of all storage and data protection devices to identify configuration drifts.

07 | Insecure session management

Session security management is crucial to protecting any system from unauthorized access and data breaches.

A significant portion of the storage and data protection systems reviewed in this report, were not properly secured, leaving them vulnerable to threats such as session hijacking, replay attacks, and compliance violations.

Some of the common misconfigurations and risks include:

- Improper encryption of sessions
- Insecure session tokens:
- Absence of, or improper session timeouts

Business impact

Misconfigurations could lead to the exposure of sensitive data, and allow an attacker to take control of storage and backup systems, and alter, delete or corrupt data, by impersonating users and administrators, or hijacking active sessions. This could lead to massive data loss, as entire storage and backup systems might be impacted, potentially affecting thousands of servers, and preventing successful data recovery.

Recommendations

Follow session security best practices published by your storage and data protection vendors. Default settings do not often provide the required level of protection.

Periodically assess the security of session settings in all storage and data protection devices to identify configuration drifts.

Additional guidance is provided in NIST SP 800 209 and ISO 27040:2024.

08 | End-of-support devices used in production

A notable percentage of the surveyed systems have reached an end-of-support status. This can be explained in part by the fact that storage and backup systems are engineered with significant reliability, redundancy, and resilience, that often far exceed that of other IT systems. Of course, it is only natural that organizations would be tempted to realize the economic saving of keeping them in service for a few more years. However, using end-of-support system can leave organizations vulnerable to massive data breaches, as such systems do not receive critical security updates.

Business impact

Running critical systems such as storage and data protection systems without current security updates could allow attackers to compromise data at a massive scale – and in severe cases, to destroy all backup copies and prevent recovery.

Recommendations

- Make sure to prepare a timely migration plan to more current platforms. Note that with some storage platforms, it may be possible to re-use or trade-in some of the infrastructure (e.g., disk enclosures), as in most cases only the controllers themselves are unsafe for use.
- Maintain current inventory of all storage and data protection systems, and keep track of upcoming end-of-support devices.
- Periodically inspect inventory, and search for undocumented devices

09 | Vulnerabilities In Software Supply-Chain Management

As already discussed, updates to storage & data protection systems are regularly issued. Most storage and data protection platforms can obtain updates, as well as send support information, by establishing connections to designated vendor support environments outside of the organization premises.

Although many organizations have published policies that prohibit such data exchanges, in some environments we have analyzed, active sessions settings were still detected.

Whether downloaded directly to the devices, or to a staging environment, and whether the process is manual or automated, safeguards need to be established when obtaining software updates. These include proper signing and signature validation, end-to-end encryption when obtaining binaries, proper authentication, IP filtering to restrict download sites, etc.

Binaries stored onsite should be validated before each installation.

In several environments, we detected configuration issues that could allow unapproved images to be deployed, or enable cybercriminals to tamper with data transfer and support sessions.

Business impact

Improper control and enforcement of software supply-chain paths could allow cybercriminals to tamper with the storage OS, and thereby gain full control over the devices, the data, and their protective copies.

Recommendations

Implement whitelists disallow access to all but specifically approved and verified sources / IP addresses for software and firmware updates

Use only secure transmission methods only allow downloads or remote support efforts over secured connections from hosts with verified certificates

Verify packages use 'md5sum' or other hash checking to verify a software package prior to every installation

Do not allow externally-initiated support connections – require that any connection to a third party, e.g., remote support and software download, originate within your network

Periodically assess the security of devices, review policies, and audit binaries, to identify configuration drifts or suspicious activity

10 | Undocumented & Insecure API / CLI

There are a surprising number of ways storage and data protection systems can be manipulated and managed, including:

- Device APIs
- Management hosts and API gateways
- In-band – using storage protocols
- Dedicated host agents
- Storage agents (or adapters) on virtual infrastructure

Most of those control methods can be further configured to refine and restrict the access level it will provide (e.g., specify which actions are allowed, such as creation, destruction, mapping, copying, etc.), limiting scope (e.g., "all" vs. named components), and controlling the access path (e.g., applying filtering rules to restrict access to specific IPs, interfaces, or endpoints).

It is vital to approve and document all allowed connections, limit their access level and scope to the minimum, and actively block any other connection.

In some of the storage and data protection environments, undocumented API entry points were found, whose purpose could not be accounted for, and in more than 10% of the environment, approved mechanisms were not properly hardened and limited.

Business impact

Undocumented and insecure API and CLI access paths can provide cybercriminals with a backdoor to control storage devices, exfiltrate data, and tamper with storage content and its backups.

Recommendations

Document all approved control points
disable any other method of access

Allow administrative connections from only designated hosts
jump boxes, management VLANs, etc.

Implement least-privileged access models
where possible, strictly scope allowed API calls by role

Periodically assess the security
of all storage and data protection devices to identify configuration drifts.

Summary & Recommendations

It appears that the state of enterprise storage & data protection security is significantly lagging behind that of compute and network security. This is a significant gap that should be addressed as soon as possible; with growing sophistication of data-centric attacks, and with tightened regulations, the business implications of ineffective resolution could rapidly increase.

On the bright side, awareness of storage and data protection security is growing, and new resources and guidance are available to help organizations build an effective program to address the gap.

It is recommended to evaluate existing internal security processes to determine if they cover storage and data protection infrastructure to a sufficient degree. Some of the questions that could help clarify the level of maturity of storage security planning are:

- Do our security policies cover specific storage, storage networking, and backup risks?
- Are we evaluating the security of our storage, backup, and data protection infrastructure on an ongoing basis?
- Do we have detailed plans and procedures for recovery from a successful attack on a storage or backup system? Do we test such procedures?
- How confident are we that the key findings highlighted in this report, and similar ones do not, and cannot occur in our environment?

If needed, vendors could be consulted or invited to be involved in such an evaluation. Based on the findings, we'd recommend:

- Determining if knowledge gaps exist in terms of storage & data protection security, and building a plan to address them
- Improving security program to address identified gap
- Proactively address risks, by using an automated solution that continually validates the security posture of your storage and data protection systems

Finally, we encourage you to learn more about securing your storage & data protection environments. A good place to start is:

- Read the *NIST SP-800-209 Security Guidelines for Storage Infrastructure* - co-authored by Continuity.
- Read *ISO/IEC 27040:2024 'Security techniques — Storage security'*
- There's also a selection of practical guides on *www.continuitysoftware.com*

Methodology

Continuity has 20 years of expertise in evaluating and validating the configuration of storage and data protection systems. Our product, StorageGuard is a dedicated Security Posture Management solution for storage and data protection systems, scanning these critical systems for security misconfigurations and vulnerabilities, while auto-remediating many of those risks.

For this research, we compiled anonymized inputs from 323 customer environments, providing a unique cross-industry insight into the state of storage and data protection security. Of the 11,435 total storage and data protection devices deployed in these environments, a representative sample of 627 (5.5%) devices were analyzed - similar to the previous years. To avoid bias – the organizations themselves chose which subset of devices would best represent the entire estate.

A total of 6,085 discrete security vulnerabilities and misconfigurations were detected, spanning more than 390 security principles that were not adequately followed. The analysis of these misconfigurations allowed us to uncover recurring patterns, and we hope the analysis and insights in this report support organizations in assessing the maturity of their security programs as they relate to storage and data protection systems.

The graphs below show the demographics of surveyed organizations and devices by industry, vendor and geography:

2025 Demographics -by industry

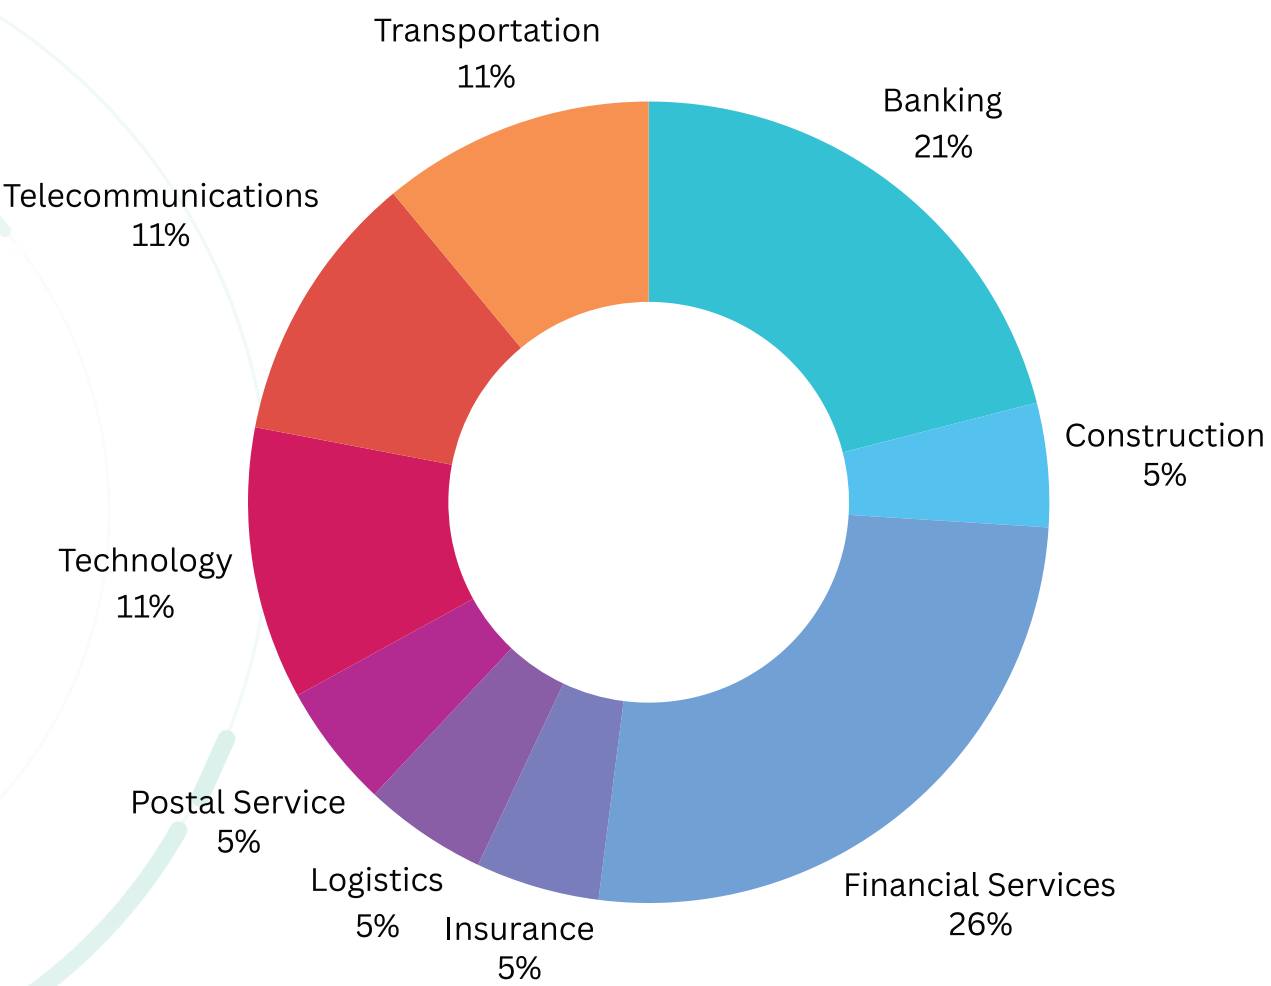

2025 Demographics -by vendor

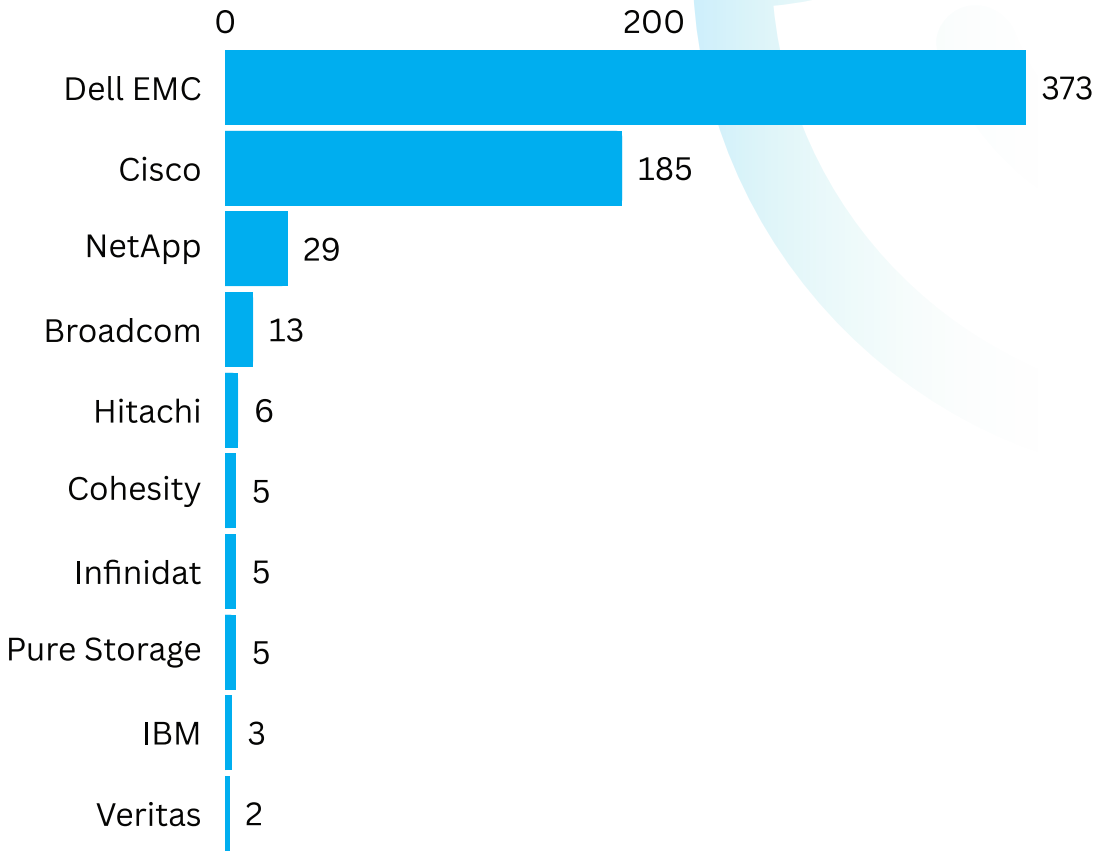

2025 Demographics -by geography

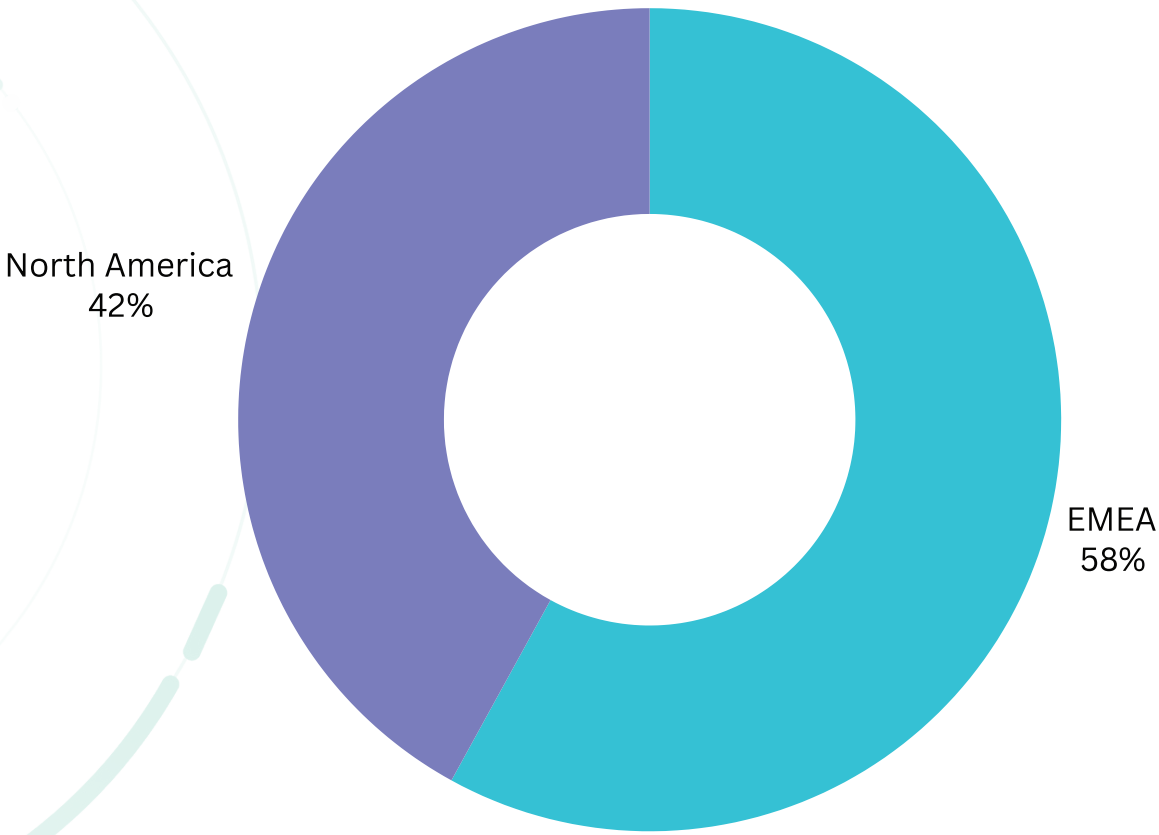

As can be seen in the vendor graph, the data in this report was collected and analyzed from configuration data across multiple storage and data protection vendors and models, including Dell, IBM, Hitachi Vantara, Commvault, Cisco, Brocade (Broadcom), NetApp, Cohesity (Veritas), and others.

The analysis covered in the configuration of block, object and IP storage systems, SAN / NAS, storage management servers, storage appliances, virtual SAN, storage network switches, data protection appliances, storage virtualization systems, backup software, backup appliances, and other storage devices.

Our automated risk detection engines check for thousands of possible security misconfigurations and vulnerabilities at the storage and data protection system level that pose a security threat to enterprises' data. These security risks fit into 4 main categories:

1. Violations of vendor security configuration guidelines
2. Violation of compliance framework requirements (CIS, NIST, PCI DSS and others)
3. Identified Storage Common Vulnerabilities and Exposures (CVEs)
4. Deviation from community-driven best practices (gathered and generalized from dozens of enterprise internal security baselines for storage – representing shared community insights)

Each finding is tagged with a security risk index (1-5), and is tracked with a wide array of tags, that allow for detailed assessment, aggregation, and drill down. These tags include:

- Demographics: Industry, country & region, organization size (# of devices, # of employees, …)
- Device tags: vendor, model, model, capacity, firmware level, …
- Security principle (e.g., authentication, authorization, logging, encryption, least-privileges, and their sub-categories)
- Security frameworks (compliance framework, organization baselines)
- And more

A note about risk scoffing methodology changes vs. previous years

We have refined the way we weigh CVE security risk index to better match modern frameworks such as CVSS 3.0, and the result was a slight increase in the average number of critical and high security risks per device relative to last years – from 3 to 5. Adjusting for this methodology change, the average number would have only slightly increased from 3 to 3.5, which is not statistically significant.

CONTINUITY

www.continuitysoftware.com