IBM Storage
Solutions Brief

Storage for Data Resilience

Minimize business impacts from ransomware, cyber-attacks, hardware failures, natural disasters, and other threats

Highlights
New risks and regulations have made data resilience a top enterprise priority

The NIST framework provides a vendor-neutral blueprint for building data resilience

IBM Storage Defender software provides AI-based end-to-end resilience

IBM FlashSystem delivers highly resilient primary storage

Only IBM can address data resilience from end to end, from data security software to mainframes and tape.

IT leaders today have made data resilience a top priority – for good reason. Their core asset – data – is under constant threat from ransomware and other cyber-attacks, as well as human error, hardware failures, and natural disasters.

Increased expense
For organizations around the globe, failing to implement effective data resilience can be expensive. In 2023, the average cost of a data breach reached an all-time high of USD 4.45 million, according to the 2023 Cost of a Data Breach Report from IBM Security.[1] Worse still, on average it took 204 days for organizations to identify a data breach, and 277 days – *more than nine months* – to identify and contain a data breach.[2]

It's not just the financial risk – the operational costs are also significant. More than 85% of organizations hit by ransomware were unable to recover all their data, according to a November 2023 report from analysts at ESG.[3]

Increased regulation
The massive economic and reputational costs of a cyber-attack should provide organizations with all the incentive they need to deploy and maintain robust data resilience practices. But in case that's not enough, governments at all levels are enacting increasingly severe laws to ensure that companies' data operations are properly protected.

The United States federal government is in the early stages of implementing the Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA), which will eventually require covered entities to report covered cyber incidents and ransomware payments to the federal government.

The European Union has already enacted significantly more stringent regulations, including DORA (the Digital Operational Resilience Act), which comes into effect in January 2025, and NIS2 (an update to the EU's Network and Information Security directive), which comes into effect in April 2025. Both can impose significant financial penalties on non-compliant organizations.

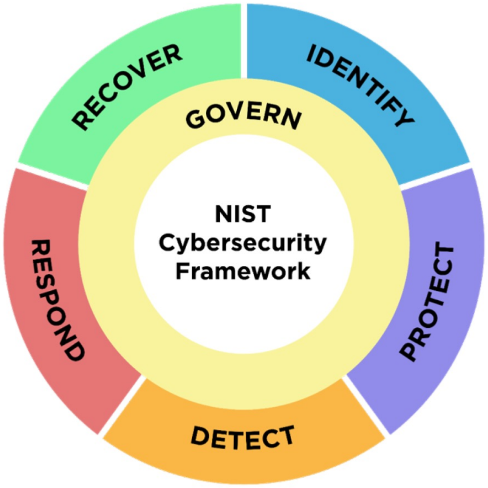

Figure 1. The NIST framework provides organizations with a simple blueprint for maximizing their data resilience.

A Standard Framework for Data Resilience

Many organizations today are designing their data resilience strategies around the cybersecurity framework developed by the US government’s National Institute for Standards and Technology (NIST).

NIST Cybersecurity Framework

The NIST cybersecurity framework focuses on managing risk, and encompasses five key areas:

- **Identify**: Understand and manage cybersecurity risks to all systems, assets, data, and capabilities;
- **Protect**: Implement safeguards to ensure the delivery of critical infrastructure services;
- **Detect**: Deploy activities to identify the occurrence of a cybersecurity event;
- **Respond**: Deploy activities to respond to a detected cybersecurity event;
- **Recover**: Deploy activities to restore services impacted by a cybersecurity event.

The framework has recently added a governance layer at the core, to reflect emerging changes in business practices in response to cyber threats, as well as increased regulatory attention in this area.

Although our primary focus is on data *resilience*, it’s useful to first clarify some key concepts and concerns in data *security*.

Data Security

A good way to think about data security is that it focuses on *preventing* impacts to an organization. Data security is concerned with ensuring the integrity, availability, and confidentiality of information.

Data security is a continuous process, not a product. It doesn’t come from a specific piece of hardware or software; it’s a practice that requires the right people, procedures, and technology, and rigorous adherence to a set of operational principles, such as access controls and monitoring, network security, patch management, vendor security assessments, and employee training.

Zero trust
In recent years, the “zero-trust” security model has become the de facto standard, based on the assumptions that data and resources are inaccessible by default, and that every connection and endpoint is considered a threat. The zero-trust model is about context – organizations start by classifying resources based on risk, defining granular resource boundaries, and separating users according to roles and duties. Only then can an organization create and enforce policies that ensure all users are able to access all the resources they need, and only those resources.

Once effective policies are in place, tools are needed to enforce these policies, along with procedures that outline the targeted actions to take in any potential data breach, such as revoking access for individual users or devices, adjusting network segmentation, quarantining users, wiping devices, creating an incident ticket, or generating compliance reports. The final component is to continuously evaluate and adjust the policies, authorization actions, and remediation tactics to tighten each resource's perimeter.

84%

Percentage of critical infrastructure incidents where initial access vector could have been mitigated.

In most of the critical infrastructure incidents that the IBM X-Force incident team responded to in 2023, the initial access vector could have been mitigated with best practices and security fundamentals, such as asset and patch management, credential hardening and the principle of least privilege.[4]

SOAR and SIEM

Two essential components in enterprise data security are SIEM (security information and event management) and SOAR (security orchestration, automation, and response), which play distinct and complementary roles in an organization's ability to detect, respond to, and mitigate security threats.

SIEM systems are designed to collect, aggregate, and analyze log data and other security event information generated throughout an organization's technology infrastructure, including host systems and applications, network and security devices, and other sources. The goal is to identify and respond to security incidents and events in real-time.

Typical capabilities of SIEM systems include collecting and storing logs and alarms, analyzing the data to identify anomalies that may indicate security threats, generating alerts and reports, and monitoring and reporting on security events to help organizations meet regulatory compliance requirements.

SOAR systems focus on the effectiveness of incident response processes by automating repetitive tasks, orchestrating workflows, and facilitating collaboration among security teams.

Typical capabilities for SOAR systems include running predefined playbooks to automatically respond to incidents, coordinating interactions between multiple security tools and systems when responding to security incidents, investigating and managing cases, and enabling efficient information sharing among security teams to accelerate time-to-recovery.

To put them in context, SIEM focuses on monitoring and analyzing security events, while SOAR enhances incident response by automating and orchestrating actions. The two work together – SIEM systems provide the initial detection and analysis of security events, generating alerts that can be ingested by SOAR platforms. This coordination is crucial in handling the large volume of security events that organizations face daily.

IBM's data security solutions include IBM Security X-Force, IBM Security QRadar Suite, and IBM Security Guardium Data Protection.

- **IBM Security X-Force** is a team of security experts, hackers, responders, researchers, and analysts, working together to help organizations build, manage, and refine their security programs. It includes the X-Force incident response team, which knows where threat actors hide and how to stop an attack.

- **IBM QRadar** is a threat detection and response solution designed to help security teams quickly and effectively manage events through the full incident lifecycle. In addition to SIEM and SOAR capabilities, it offers integrated products for log management and endpoint management, all with a common user interface, shared insights, and connected workflows.

- **IBM Guardium** automatically discovers and classifies sensitive data from across the enterprise, providing real-time data activity monitoring. The software detects behavioral vulnerabilities such as account sharing, excessive administrative logins, and unusual after-hours activity, and identifies threats and security gaps in databases that could be exploited by hackers.

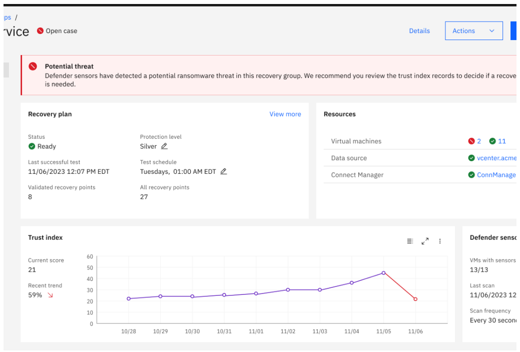

Potential threat
Defender sensors have detected a potential ransomware threat in this recovery group. We recommend you review the trust index records to decide if a recove
is needed.

Recovery plan

Status: Ready
Protection level: Silver
Last successful test: 11/06/2023 12:07 PM EDT
Test schedule: Tuesdays, 01:00 AM EDT
Validated recovery points: 8
All recovery points: 27

Resources

Virtual machines	2 11
Data source	vcenter.acme
Connect Manager	ConnManage

Current score: 21
Recent trend: 59% ↘

VMs with sensors: 13/13
Last scan: 11/06/2023 12
Scan frequency: Every 30 secon

Figure 3 – An IBM Storage Defender dashboard provides quick access to key information about the protection status of all workloads.

Comprehensive Resilience – IBM Storage Defender

Organizations today need a data resilience strategy that encompasses every aspect of their on-premises and cloud environments and that supports all their traditional, hybrid cloud, and virtualized workloads.

IBM Storage Defender software is designed to meet that need. It offers end-to-end data resilience in modern hybrid multi-cloud IT environments that includes virtual machines (VMs), databases, applications, file systems, SaaS workloads, and containers.

IBM Storage Defender features exceptional scalability, multiple layers of cyber resilience, broad application support, and cost-saving data reduction technologies. It uses SLA-based policies to automate the entire data protection process, including backup, replication, and secure data retention on-premises and in the cloud. Cyber resilience is enhanced by support for logical air gap to object storage (WORM technology) as well as the ability to physically air-gap data to tape.

Key capabilities of IBM Storage Defender include:

- Threat detection: IBM Storage Defender is designed to detect threats and anomalies from backup metadata, array snapshots, and other relevant threat indicators. It includes a data resilience service that enhances existing security systems by including storage-specific malware and anomaly detection, as well as providing a trust index to help IT leaders decide where to prioritize the allocation of resources.
- Rapid recovery: Can enable organizations to validate, recover, and restore data more quickly and completely.
- IBM Storage Defender is designed to integrate with other IBM Storage and IBM Security solutions, including IBM QRadar, IBM Guardium, IBM FlashSystem, IBM Storage Scale, IBM Storage Ceph, and IBM Fusion. It also includes copy data management tools to manage and orchestrate application-integrated, hardware snapshots by making copies available when and where users need them for instant data recovery or data reuse, automatically cataloging and managing copy data across hybrid cloud infrastructures.
- Flexible licensing: Licensing is based on resource units, providing a cloud-like, utility-based consumption model for organizations to consume any service within IBM Storage Defender.

IBM Storage Defender 2.0 added new data resilience capabilities across three critical areas:

- Application and storage inventory management: This feature helps organizations assess their landscape of applications and data and determine how they fit into their business continuity plan, using information from FlashSystem and VMware vCenter as well as user-supplied data about custom applications.
- Expanded threat detection: An AI-powered Trust Index provides IT operators with a score to indicate the relative trustworthiness of copied data and VMs by combining signals from existing solutions and new detection methodologies developed by IBM Research.
- Accelerated recovery capabilities: When an attack occurs that results in compromised data, IBM Storage Defender can orchestrate and automate the recovery of VMware applications.

Figure 2 – IBM FlashSystem 9500 storage system.

Resilient Primary Storage – IBM FlashSystem

IBM Storage FlashSystem solutions provide block storage for a wide variety of enterprise workloads, including databases, analytics, IT modernization, application consolidation, virtualization, and data protection.

Data resilience is baked into every FlashSystem solution. IBM FlashSystem products scan all incoming data down to block level granularity without impact to performance as it's being written, using inline data corruption detection software and cloud-based AI to help identify anomalies that might indicate the start of a cyber-attack, thereby enabling the system to detect, respond, and rapidly recover with immutable copies.

The IBM FlashCore Module 4 technology in FlashSystem products is designed to continuously monitor statistics gathered from every single I/O using machine learning models to detect anomalies like ransomware in less than a minute.[5]
The FlashSystem inline data corruption detection capability can be integrated with IBM QRadar SOAR, so that detection of a threat automatically triggers the creation of a new immutable copy.

Policy-based High Availability (HA) provides a zero recovery point objective (RPO) solution for two storage systems in different locations where the storage is synchronously replicated across metro-area distances. In this configuration, FlashSystem solutions enable servers at each data center to access data concurrently with seamless failover for zero recovery time objective (RTO).

To build effective business resilience and disaster recovery (DR) architectures at greater distances, enterprises can configure multiple FlashSystem solutions in a policy-based replication connection and asynchronously replicate data across regions. Data is written to the local FlashSystem and the I/O is completed on the local system before that data is sent to the remote system. This approach can achieve very low RPO times.

IBM Safeguarded Copy and Cyber Vault
IBM Safeguarded Copy creates isolated immutable snapshots of data to help protect against potential data loss. These are typically taken multiple times throughout the day, so that a recent snapshot is always available with an RPO that meets the organization's service level agreements (SLAs) to internal and external clients. Safeguarded Copy snapshots are stored on the same FlashSystem storage as operational data but are logically isolated from the production servers. This architecture helps ensure that recovery occurs more quickly than with systems where backup copies are stored separately.

FlashSystem Cyber Vault is a blueprint implemented by IBM Lab Services or IBM business partners that is designed to help speed cyberattack detection and recovery. The Cyber Vault solution runs continuously and monitors snapshots as they are created by Safeguarded Copy. It leverages a sandbox or clean room environment with logical partitions or VMs to run data validation processes without affecting production workloads.

Using standard database tools and automation software, FlashSystem Cyber Vault checks Safeguarded Copy snapshots for corruption. If FlashSystem Cyber Vault finds such changes, it recognizes that an attack may be occurring. When preparing a response, knowing the last snapshots with no evidence of an attack can speed the determination of which snapshot to use. FlashSystem Cyber Vault is designed to help reduce cyberattack recovery time from days to minutes or hours.

Cyber Recovery Guarantee

Leveraging its breadth of expertise in data security, IBM offers a complete range of storage hardware and software solutions to secure data across primary and secondary storage systems with logical, operational, and physically air-gapped systems. These solutions not only safeguard an organization's data, but also detect and remediate threats as data moves across storage tiers, improving business continuity in the event of an attack.

The IBM FlashSystem Resilience Guarantee is designed to help organizations persist through a ransomware attack, protected by a FlashSystem with Safeguarded Copy immutable snapshots. The guarantee provides clients with the assurance that they'll be able to recover protected data in 60 seconds or less.

In the event of a ransomware or cyber-attack where data is affected, clients can restore the previous snapshot for the volume and make it available to the operating systems or applications in 60 seconds or less on a local array. This enables the client to continue business operations with minimal disruption while preserving critical evidence for incident forensics. IBM guarantees recovery of a SafeGuarded Copy (immutable snapshot) restore point within 60 seconds or less. Should a SafeGuarded Copy not be able to recover within 60 seconds, the Cyber Resilience Guarantee provides 40 hours of IBM Technology Expert Care guidance for remediation and resolution.

FlashSystem and IBM Storage Defender – Better Together

FlashSystem storage systems provide additional resilience capabilities when paired with IBM Storage Defender software.

FlashSystem arrays can be registered with an IBM Storage Defender Data Protect cluster, enabling administrators to create protection groups that include specific volumes and are automatically backed up according to user-defined policies. Data can be restored or recovered to multiple target locations, including new locations when recovering from a cyber-attack. Plus, snapshot copies can be replicated to another Data Protect cluster for an additional layer of protection.

New settings allow administrators to automate the creation of Safeguarded Copy snapshots, cyber-resilient point-in-time copies of volumes that cannot be changed or deleted through user errors, malicious actions, or ransomware attacks. Isolating these backup copies from production data can help recover data quickly in case of a cyberattack.

IBM Storage Sentinel

IBM Storage Sentinel is workload-specific software that detects, diagnoses, and identifies the sources of ransomware attacks and provides automated recovery orchestration for Oracle, SAP HANA, and the Epic healthcare system.

IBM Storage Sentinel automates the creation of immutable backup copies of your data. Then it uses machine learning to detect signs of possible corruption and to generate forensic reports that help you quickly diagnose and identify the source of the attack. Because IBM Storage Sentinel can intelligently isolate infected backups, your organization can identify the most recent verified and validated backup copies, which can accelerate your time to recovery.

Resilient Software-defined Storage

Figure 4 – The IBM Storage Scale System 6000 can store 1.44PB of uncompressed data in a single 4U rack configuration and can scale to yottabytes.

Some organizations now prefer to deploy *software-defined* storage solutions, where the storage software is decoupled from the hardware. Software-defined storage can reduce costs, and provides centralized management, enhanced automation and orchestration, improved resource utilization, data mobility, and the ability to easily scale up or down as business requirements change.

IBM's software-defined storage offerings include IBM Storage Scale, IBM Fusion, and IBM Storage Ceph. IBM hardware options are available for all three solutions.

IBM Storage Scale
IBM Storage Scale provides software-defined file and object storage for AI and data intensive workloads and is used by organizations around the world that need to create a high performance, globally connected, cost-optimized global data platform. Data resilience is fundamental to the IBM Scale architecture, with support for a variety of two-site and three-site configurations. A single Storage Scale cluster can be configured using nodes and storage from two data centers, so that one site always remains active, even if the other site or link fails.

IBM Storage Scale System is a hardware appliance that allows you to deploy IBM Storage Scale software on thousands of nodes with TB/s performance, low latency, and tens of millions of IOPS per node. It's designed for organizations that are looking to accelerate deployment, lower acquisition costs, and simplify storage management when deploying high-performance file and object storage systems.

IBM Fusion
Driven by the need to rapidly develop and deploy innovative applications, many organizations have adopted cloud-native computing based on Kubernetes, which enables agile, consistent, scalable, and portable deployment across different environments – on premises, at the edge, or in the cloud.

IBM Fusion is a container-native data services platform that delivers simplified infrastructure to application developers and data scientists, enabling them to build applications without concerning themselves with underlying layers.

Fusion is available in two forms:

- IBM Fusion software runs anywhere Red Hat OpenShift runs – on-premises, on one or more public clouds, on bare metal, and on virtual machines.
- IBM Fusion HCI provides a fully integrated, turnkey platform for running and maintaining on-premises Red Hat OpenShift applications.

IBM Storage Ceph
IBM Storage Ceph is software-defined block, file, and object storage, designed to decouple data from physical storage hardware to provide unparalleled scaling and data resilience capabilities. Data resilience is central to IBM Storage Ceph, with data automatically replicated across multiple physical devices for redundancy.

Ceph is self-healing, automatically detecting and recovering from hardware failures or corrupted data. Its capabilities include snapshots, object versioning, multi-factor delete, fault tolerance, and rolling upgrades, enabling administrators to upgrade the software and perform maintenance on the cluster without downtime.

IBM Storage Ready Nodes provide a simple building-block approach that combines Ceph software with the optimal hardware and support services from IBM.

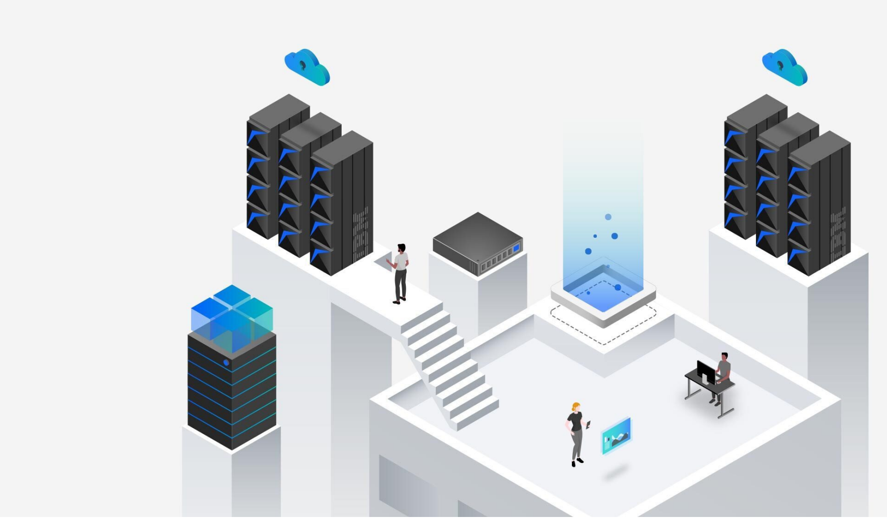

Resilient Mainframe and Tape Storage

Most global organizations today are reliant on modern mainframes to run their hybrid cloud infrastructure. IBM is a major supplier of high-performance computing in environments where security and reliability are paramount, including finance and banking, insurance, government, aviation, healthcare, retail, and manufacturing.

Mainframe resilience – IBM DS8000F
The IBM DS8000F is a family of enterprise storage systems designed for high-performance, business-critical data storage in large enterprises, known for their reliability, scalability, and advanced features. These systems are built for speed, specifically the low latency needed for use cases such as real-time fraud detection, business analytics and visualization, and machine learning-based anomaly detection.

IBM DS8000F storage systems feature resilience capabilities that include hardware snapshot and replication capabilities, encryption, redundant components, advanced error detection and correction, and hot-swappable components to minimize downtime.

Air-gapped resilience – IBM Tape
The market for data storage on tape has grown in recent years, especially among large organizations and hyperscalers. Data resilience is a major reason for this growth, along with capacity, total cost of ownership, endurance, and sustainability. Tape storage offers unparalleled data resilience because it provides a physical air gap between your archived data and the outside world.

IBM is a market leader in tape storage, providing tape drives, autoloaders, libraries, virtual tape systems, and software, with capabilities designed for entry-level, midrange, and enterprise system environments.

Cyber Resilience Assessment

An excellent way for organizations to audit their cyber defenses is the Cyber Resilience Assessment, which IBM and its business partners provide to clients at no charge. The team conducting the assessment uses the NIST cybersecurity framework to evaluate a client's current cyber resilience posture, including their data storage, security, and infrastructure capabilities.

The Cyber Resilience Assessment is a vendor-neutral overview that helps clients identify the strengths and weaknesses of their existing systems and provides recommendations to further strengthen their cyber resilience. To learn more, go to https://www.ibm.com/downloads/cas/W7VJLDPE.

For more information

To learn more about data resilient storage solutions, contact your IBM representative or IBM Business Partner, or visit https://www.ibm.com/storage-data-resilience.

1. IBM Cost of a Data Breach Report 2023, August 2023.
2. IBM Cost of a Data Breach Report 2023, August 2023.
3. Ransomware Preparedness research report, Enterprise Research Group, December 2023.
4. IBM X-Force Threat Intelligence Index 2024, February 2024.
5. Disclaimer: Internal experimentation by IBM Research has demonstrated detection of ransomware within 1 minute of the ransomware starting its encryption process. This experiment was done on a FlashSystem 5200 with 6 FCMs with the 4.1 firmware load. The 5200 had 8.6.3 GA level software loaded. The host connected to the 5200 was running Linux with XFS Filesystem. In this particular case, the IBM ransomware simulator called WannaLaugh was used. Underlying system must be compatible with FCM4.1 and version 8.6.3 GA level software loaded in order to receive results obtained.

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the
United States of America
March 2024

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

IBM products are warranted according to the terms and conditions of the agreements under which they are provided.