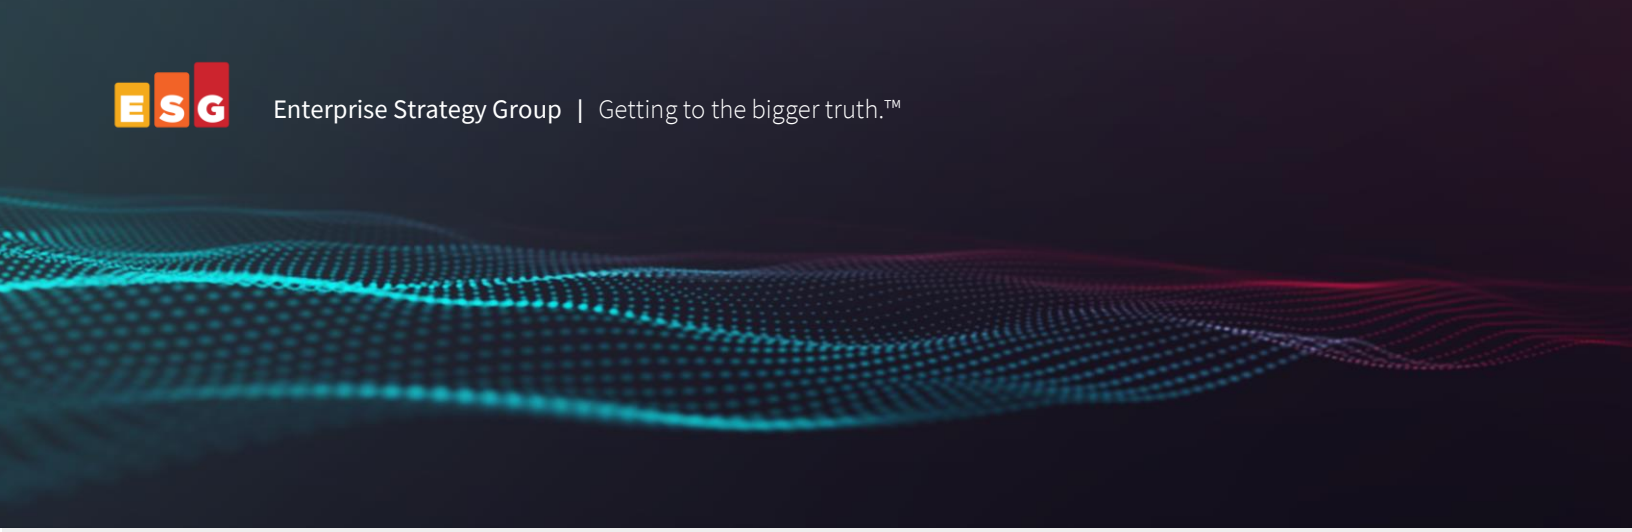

ESG WHITE PAPER

Solve Cyber Resilience Challenges with Storage Solutions

By Jack Poller, ESG Senior Analyst; and Leah Matuson, Research Analyst

August 2020

This ESG White Paper was commissioned by IBM and is distributed under license from ESG.

Contents

Executive Summary

The accelerating adoption of modern technologies—cloud, mobility, the internet of things (IoT), and now work-from-home facilitators—has led to an increase in the complexity of IT infrastructure, as well as a proliferation of data. These factors put organizations and their IT infrastructures at greater risk for malicious attacks, human errors, and negligent behavior.

Unfortunately, legacy strategies cannot adequately ensure continued business operations during and after cyber incidents. Companies can weave together capabilities in an attempt to prevent attacks and data breaches, but functional gaps, poor integration, and management complexity slow implementation and make security objectives difficult to meet.

Changing organizational mindsets from prevention to preparation for a security incident—cyber resilience—and implementing storage solutions with built-in cyber resilience is key to safeguarding critical data assets and quickly responding to and recovering from cyberattacks.

Introduction

IT faces new challenges. Nearly two-thirds (64%) of survey respondents say IT is more complex today than it was two years ago. This can be the result of ongoing digital transformation (as cited by 27% of respondents citing increased complexity) and/or supporting modern enterprises with an increasing mobile workforce (24%). Other top drivers of increased IT complexity include high data volumes (37%), the rapid evolution of the cybersecurity landscape (31%), and new data security and privacy regulations (30%).[1]

Organizations also must address a problematic shortage of critical IT skills: 44% of organizations report a lack of cybersecurity skills, the most often cited response. Additionally, these organizations are more likely to point to additional complicating factors that increase the size and scope of the security perimeter they're tasked with protecting, such as application, device, and remote/mobile worker sprawl.[2]

The Cost of Complexity Is Rising

Organizations face many cybersecurity threats. ESG research highlights that ransomware attacks are a common occurrence for many enterprises, with 60% of respondents experiencing ransomware attacks in the past year and 13% of those organizations experiencing *daily* attacks. Interestingly, organizations that report a cybersecurity skills shortage are also much more likely (67% versus 54%) to report being targeted by ransomware over the past 12 months.[3]

There may be several important reasons driving the increase in cyberattacks, but certainly the financial incentive to attack organizations is increasing. For example, complaints to the FBI's Internet Crime Complaint Center (IC3) increased 62% in the past five years, while reported losses from cybercrimes hit $3.5 Billion in 2019—a 218% increase during the same period.[4] According to Ponemon Institute research sponsored by IBM, on average a data breach costs an organization $8.19M. And costs are long-lived: 67% of the costs came in the first year, 22% in the second year, and 11% continued more than two years after a breach.[5]

[1] Source: ESG Master Survey Results, *2020 Technology Spending Intentions Survey*, January 2020.
[2] ibid
[3] ibid
[4] Source: 2019 FBI Internet Crime Report,
[5] Source: 2019 Cost of a Data Breach Report

Perhaps most revealing, IT complexity amplifies data breach costs. For example, system complexity increases the cost of a data breach by $290K. Third-party involvement in IT infrastructures adds $370K. And a breach during a cloud migration adds $300k.[6]

Solve Cyber Resilience Challenges with Storage Solutions

It's clear that a strong correlation exists between IT complexity and vulnerability to cyberattack. And as IT grows more complex, cyberattacks will increase in frequency and cost. The truth is, given enough time and effort, anything can be breached. This explains why organizations are shifting focus and resources are moving beyond cybersecurity toward a more comprehensive data and business security posture called "cyber resilience."

Cybersecurity encompasses the technology and processes that protect data and the infrastructure from malicious attacks. Cybersecurity reduces organizational risk from attacks and protects data and the organization from deliberate exploitation. But even with the best security, data breaches—via attack or negligence—still occur. And when they do, they aren't always quickly identified. According to Ponemon Institute research, in 2019 it took an average of 206 days to identify a breach and another 73 days to contain the breach, for a total of 279 days, a 4.9% increase over 2018. Also, while cybersecurity-oriented organizations rightly focus on malicious or criminal attacks, which accounted for 51% of breaches in 2019, 25% of data breaches involved system glitches, including both IT and business process failures, and 24% of these breaches were due to negligent employees or contractors (human error).[7]

Organizations can shift from prevention to preparedness. This approach helps organizations continue to operate and deliver intended business outcomes despite adverse cyber events—whether malicious or inadvertent. This strategy is called *cyber resilience*. A cyber-resilient organization stops reacting to cyber incidents and starts anticipating them. It puts in place the infrastructure, tools, and processes to protect its core asset—its data. Such an organization can respond to and recover from attacks and system failures, plus guarantee the continuity of operations before, during, and after any cyber event.

Cyber-resilient strategies take into consideration all data-handling components—hardware, software, people, and processes. When developing cyber resilience, organizations shift from asking "How do I protect?" to asking, "What if?" What if we're hit by ransomware? What if an employee misconfigures cloud storage?

Data storage is a critical element of cyber resilience. But too often organizations do not realize the central role played by data management and storage in a comprehensive cyber resilience strategy.

NIST Cybersecurity Framework

To help guide the conversation and develop a cyber resilience strategy, organizations can look to the cybersecurity framework provided by the National Institute of Standards and Technology (NIST). The *Framework for Improving Critical Infrastructure Cybersecurity* from NIST states: "*The Cybersecurity Framework Core provides a set of activities to achieve specific cybersecurity outcomes…*"[8] The NIST framework provides guidance to organizations on how to plan, develop, and implement solutions to achieve specific business operations outcomes. Typically, the outcome is continuing business operations in the event of cyber incidents. Thus, the five elements of the NIST cyber resilience framework are:

[6] ibid
[7] ibid
[8] Source: NIST Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1, April 2018.

- Identify—Develop an organizational understanding for managing cyber risk to systems, people, assets, data, and capabilities. To ensure continuous operations, organizations must identify the essential elements of business operations and how to protect those elements, then develop and prioritize a protection plan.
- Protect—Develop and implement appropriate safeguards to ensure delivery of critical services. Implementation of various policies and controls can include the principle of least privilege access and access controls, identity management, data security, data protection, application and code protection, and user awareness and training.
- Detect—Develop and implement appropriate activities to identify the occurrence of cyber incidents that put the organization at risk. Early detection and alerting help organizations minimize the impact and maximize time to recover from a cyber event.

- Respond—Develop and implement appropriate activities to take action regarding a detected cyber incident. Put in place the tools and processes to contain the impact of an event.
- Recover—Develop and implement appropriate activities to maintain plans for resilience and to restore any capabilities, data, or services that were impaired due to a cyber incident. Advanced planning and testing of tools and processes help ensure that organizations can restore capabilities and services that were impaired.

Fundamental Capabilities of Cyber-resilient IT Infrastructure

Organizations developing a cyber resilience strategy must require that resilience is a foundational design element of their IT infrastructure. Some of the fundamental capabilities of a cyber-resilient IT infrastructure include:

- Data discovery, categorization, and copy management.
- Encryption.
- Access control.
- Cyber incident detection.
- Immutable storage, including air-gap data protection capabilities.
- Data recovery.

Several of these foundational capabilities of modern IT infrastructure deserve added discussion because their impacts and technical aspects may not be as well known.

Discovery of Sensitive Data and Data Management

In the 21st century, most enterprises can state: *Data is the business and the business requires data*. ESG research confirms that 20% of organizations state flatly that data is their business—their core products and services are information-based—

while another 35% of enterprises offer both tangible and information-based products and services.[9] Thus, protecting the organization means protecting the organization's data. This is why data security ranks at the top of organizations' overall security strategy (30%).[10]

Businesses rarely throw away data because the cost of making a mistake and losing critical data is far greater than the cost of storing that data. Thus, an organization's volume of data is always growing, on-premises, in the cloud, and in backup and archive systems. According to ESG research, more than half of organizations have greater than 250TB of data, and 29% have more than 500TB of data.[11] This volume of data inhibits an organization from manually tracking the location and use of all its data assets.

Enterprises cannot protect their data if they don't know the data exists; where the data is stored; and what tools, applications, and people are using the data. They need help from the IT infrastructure to discover, identify, and classify their data assets. Data discovery ensures that the organization knows about, and can therefore protect, data critical to business operations.

However, data discovery is not enough. Organizations create copies of data for a variety of reasons, including snapshots, backups, archives, regulatory and standards compliance, application development and testing, analytics, and access speed. They must ensure that each copy of data is correctly protected. And when multiple copies exist, they must identify which copy is the source of truth, so that changes made in test/dev or for other reasons don't inadvertently propagate to backups or production data.

Copy data management (CDM), which helps an organization identify and manage its sensitive data repositories, is an important element of data discovery, and a component of the "identify" element of the cyber resilience framework. CDM solutions enable IT administrators to automate and orchestrate otherwise complex data reuse scenarios for multiple business solutions such as automated disaster recovery, test/dev, and business analytics. CDM automation and orchestration enables push-button deployment of resources in fenced/segregated environments, quick promotion to production, simplified post-operation clean-up, and daily validation of proper operation. Repeatability and auditability enable enterprises to decrease the number of data copies, reducing data sprawl and concomitant risk of data breach while helping to lower costs.

Copy data is one of the more susceptible vectors by which an organization can experience cyber events because, too often, each copy of data is not tracked, scrubbed for sensitive information, or protected by stringent levels of security. Deploying an effective CDM solution can help strengthen an organization's copy data security posture by employing robust data masking and by using policy- and role-based filtering to keep sensitive data away from prying eyes.

Encryption

Negative publicity surrounding the many high-profile data breaches over the last few years has made organizations aware of the risks arising from data exposures, whether due to misconfiguration, user error, or malicious activity. Thus, organizations desire to encrypt data both at rest and in transit to ensure that sensitive information is protected along its entire trajectory. According to ESG research, 31% of organizations believe that data encryption is one of the most effective capabilities to protect sensitive data, and 89% of organizations are using or plan to use native data encryption in the next 12-24 months.[12]

[9] Source: ESG Master Survey Results, *The Evolution from Data Backup to Data Intelligence*, January 2020.
[10] Source: ESG Master Survey Results, *Trends in Endpoint Security*, March 2020.
[11] Source: ESG Master Survey Results, *The Evolution from Data Backup to Data Intelligence*, January 2020.
[12] Source: ESG Master Survey Results, *Trends in Cloud Data Security*, January 2019.

But it's important to note that even when data is encrypted, organizations are still at risk of data loss. If a malicious actor accesses an encryption key, the actor can access the associated data. If the key is lost or corrupted, the organization loses access to sensitive and mission-critical data. With a massive amount of data continuing to proliferate across the enterprise, the volume and the variety of keys continue to rise, and organizations may find themselves looking at millions of keys and creating more new keys every day.

Cyber resilience demands always-on encryption, protecting all data all the time. However, data is not only stored, it is utilized by different applications and users and is always moving throughout the organization. Thus, data must be encrypted both at rest in the storage system and in transit in the network.

It is possible to add encryption to existing IT infrastructure using "bump-on-the-wire" technology that encrypts data passing through individual systems or devices. However, this type of bolt-on solution is inefficient, hard to manage, and costly to implement because the encryption solution must be added to every IT device in the organization, including laptops, tablets, and mobile phones. And there is always the risk of missing a component.

Complete native encryption, where each data storage device and network connection use built-in encryption capabilities, keeps all data encrypted and protected all the time, in flight and at rest. Such native solutions are easier to integrate with centralized key management solutions, removing another threat and implementation hurdle.

Access Control

Along with complete native encryption, controlling access to data is a key component of cyber resilience strategies. Modern organizations are implementing data security solutions based on the concept of least privilege access. Applying this principle—giving users the bare minimum permissions they need to perform their work—is critical in helping to prevent inadvertent or malicious unauthorized access to data.

One tool used to implement solutions based on the least privilege access principle is role-based access control (RBAC), which involves restricting access based on the roles of individual users within an enterprise. RBAC lets users have access rights only to the information they need to do their jobs and prevents them from accessing information that doesn't pertain to them. Limiting access is especially important for organizations that have many workers, employ contractors, or permit access to third parties like customers and vendors. Companies that depend on RBAC can better secure their sensitive data and critical applications. RBAC solutions also bring a number of ancillary benefits, such as improved operational efficiency, enhanced compliance, increased visibility into security mechanisms, reduced costs, and fewer data breaches; plus, RBAC enables self-service data reuse to improve the speed and efficacy of analytics, application development, and testing.

Two-factor authentication, where users provide two different authentication factors, increases the accuracy and reliability of user authentication by requiring the user to provide something they know (the first factor, usually a password) and something they have (the second factor). Common second authentication factors include biometrics (fingerprints and facial scans), security tokens, and mobile devices. 2FA has long been used to control access to sensitive systems and data; now, most organizations are using 2FA to protect user credentials from being used by hackers who have access to a database of leaked or stolen passwords, or from guessing common and weak passwords. 2FA solutions continue to evolve. For example, the dual control security authentication now offered on certain IBM enterprise and mainframe-oriented storage solutions is implemented with a "maker" and "checker" approach to proactively reduce the risk of human error for accidental deletion and malicious damage.

Immutable Storage

Buggy software, failing hardware, rogue users, and cyber criminals and their ransomware can make unauthorized modifications to sensitive data. These changes, whether from test/dev or via malicious intent, can propagate to primary storage and backup solutions. For example, many organizations have suffered from undetected ransomware attacks that corrupted both primary storage and backups, hampering or preventing recovery. In December 2019, The Heritage Company ceased operations because it could not recover from a ransomware attack, despite paying ransom. In September 2019, Wood Ranch Medical ceased operations when it could not recover patient records due to an August ransomware attack.

One method to protect against unauthorized modifications to data is by using immutable storage. This technology ensures data can never be changed or deleted for a specified retention period, regardless of the reason. Data that can be written once and read many times (WORM) is immutable and protected from malicious attacks, inadvertent modifications, and even the classic human failure—mistakenly deleting the wrong file. The immutability of the storage system can be a physical characteristic of the device, such as a layer that is physically altered during writing (optical disks, punch cards, and paper tape), or the software and hardware can enforce immutability, preventing storage from being written more than once and preventing data from being deleted for a specified retention period.

An object storage system can be configured to create new objects rather than modifying existing objects. Thus, the object storage system can both enforce immutability and automatically maintain an audit trail and history of data modifications, along with metadata capturing the who, why, how, and when. Changes can be rolled back to known valid states using the audit trail and previous versions of data, helping recovery from cyber incidents.

A particular approach to immutable storage is called “air gapping.” Air-gapped systems are physically isolated. They have no network connection and data can only be transferred using removable media. Virtual air-gapped systems are isolated from unsecured networks—i.e., the system is not directly connected to the internet, nor is it connected to any other system that is connected to the internet. Air gapping can protect sensitive data from malicious attacks such as ransomware. Because the ransomware cannot connect to the system over the network, it cannot encrypt the data.

Backup, archive, and data protection solutions can be physically and logically air gapped. Tape and other removable media create a physical air gap. Once data is stored on the removable media and the media itself is removed from the system, the data is no longer accessible, nor is it susceptible to being overwritten by ransomware. Immutable object storage and external cloud services that are separated from a company’s network are examples of logically air-gapped data. Air gapping both protects data from unauthorized access and aids in recovering from a cyber incident.

Data Recovery

The 3-2-1 backup rule is a proven approach to ensuring data recovery. The rule is: keep at least three copies of your data, and store two backup copies on different storage media, with one of them located offsite. Following this rule gives companies multiple recovery options, and companies can easily implement the 3-2-1 rule by utilizing end-to-end SLA-based polices that automate the date protection process, including operational data backups, data replication, and long-term data retention.

Data recovery performance is another important consideration. While tape excels at isolating backup data; disk, hybrid, and all-flash storage enable faster data recovery. Moreover, modern data protection solutions offer global search, space-efficient snapshots, and support for native data formats to enable near-instant data recovery.

Other capabilities that assist with data recovery are the ability to create temporary data copies for disaster recovery testing, the ability to restore data from application-consistent snapshots, and the ability to recover data in a fenced off environment to find valid recovery points before restoring data to a production system.

Shifting from Cybersecurity to Cyber Resilience with IBM

Fortunately, IBM has developed an extensive suite of storage and data protection solutions to help proactive organizations shift from cybersecurity to developing and implementing a comprehensive cyber resilience strategy. The IBM solutions incorporate all facets of a complete cyber resilience storage strategy—data discovery, encryption, immutable storage, air gapping, and having multiple recovery options to enable organizations of any size to increase their cyber resilience and improve their ability to identify, protect against, detect, respond to, and recover from cyber incidents (see Table 1).

Table 1. IBM Solutions

IBM Solution	Description	Cyber Resiliency Capabilities	Identify	Protect	Detect	Respond	Recover
IBM FlashSystem Storage	All-flash and hybrid storage with IBM FlashCopy snapshots for space-efficient immutable copies of data. Enables quick restores and accelerates recovery from unauthorized data modification.	Encryption Immutable		✓		✓	✓
IBM DS8900F Storage	High-capacity, high-performance all-flash storage supporting continuous operations, immutable storage, and recovery using many immutable recovery copies across multiple volumes or storage systems, with the ability to prevent ransomware attempts to delete or modify data. In addition, DS8900F provides secure authentication and encryption in flight for IBM Z environments.	Encryption Immutable		✓	✓	✓	✓
IBM Cloud Object Storage and IBM Spectrum Discover	Cloud-based object storage for archiving and data protection, providing immutable storage using WORM technology with the ability to specify legal holds and retention periods at an object level. The IBM Spectrum Discover description is below.	Encryption Immutable	✓	✓		✓	✓
IBM Tape Storage	Tape cartridges physically air gap data because they are offline. Virtual tape libraries can provide logical WORM storage.	Encryption Immutable Air gap		✓		✓	✓
IBM Spectrum Archive	Features the IBM Linear Tape File System (LTFS) format standard to provide direct, intuitive, and graphical access to data stored on tape cartridges.	Encryption Immutable Air gap		✓		✓	✓

IBM Solution	Description	Cyber Resiliency Capabilities	Identify	Protect	Detect	Respond	Recover
IBM Spectrum Copy Data Management	Identifies and manages multiple data copies, provides data protection for multi-cloud environments, and can recover data in an isolated network.	Discovery Data management Encryption	✓	✓		✓	✓
IBM Spectrum Protect Suite	Leveraging IBM's long history with data protection, including inventing tape and other removable media solutions, IBM Spectrum Protect Suite supports physical systems, VMs, containers, applications, and multiple cloud services. The software-defined storage solution can store data on flash, disk, object storage, and physical and virtual tape, and can detect malware and ransomware activity by identifying large deviations from normal access patterns.	Detection Access control Backups Data management Encryption Immutable Air gap	✓	✓	✓	✓	✓
IBM Spectrum Scale with IBM QRadar, IBM Spectrum Discover, and IBM Spectrum Archive	Data recovery using snapshots and synchronous and asynchronous replication. IBM Spectrum Scale works in combination with IBM QRadar to detect potential threats using AI-enhanced capabilities. The IBM Spectrum Discover description is below, and the IBM Spectrum Archive description is above.	Encryption Immutable Threat detection	✓	✓	✓	✓	✓
IBM Transparent Cloud Tiering (TCT)	Enables hybrid clouds as an additional storage tier and provides logical air gapping for data protection and recovery.	Encryption Air gap		✓		✓	✓
IBM Spectrum Discover	Enables companies to rapidly identify and categorize large-scale, heterogeneous data repositories using metadata management.		✓				
IBM Cyber Resiliency Services	Helps organizations assess their needs and develop and implement cyber resilience strategies and integrated solutions across storage and security.		✓	✓	✓	✓	✓

The Bigger Truth

IT infrastructures are continuing to grow more complex, increasing the opportunity for human error, system failures, or negligence. Simultaneously, malicious actors—both inside and outside the organization—are relentless in their attacks against IT infrastructure, searching for and exploiting weak links.

Therefore, there is little doubt that security incidents will happen. This understanding drives a change in mindset from reactive to proactive; from attempting to prevent attack to preparing for and responding to failures. This is the transformation from cybersecurity to cyber resilience.

Many organizations are modeling their cyber resilience strategies after the guidance provided by the NIST Cybersecurity Framework, which recommends that organizations identify critical resources, protect those resources, detect failures and breaches, and plan for response and recovery from cyber incidents. Leading organizations are paying special attention to IT infrastructure capabilities that can enhance their cyber resilience through capabilities such as data discovery, copy management, encryption, access control, immutable storage, and having multiple data recovery options.

For IT and business leaders, cyber resilience is all about making the right technology decisions and the right business decisions with the goal of keeping the business operational.

Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

www.esg-global.com contact@esg-global.com 508.482.0188