IDC

Tri-Delta Resources

SPOTLIGHT

Sponsored by: IBM

Integrating AI/ML-driven capabilities into enterprise storage systems holds out the potential for faster cyberthreat detection and recovery.

Business Requirements for Cyber, Operational, and Data Resiliency Are Driving AI-Driven Capabilities in Flash Storage Systems

February 2025

Written by: Carol Sliwa, Research Director, Storage and Converged Systems, Worldwide Infrastructure Research

Introduction

Strengthening cyber, operational, and data resilience is a critical priority for organizations pursuing new digital business initiatives. The growing use of AI applications is heightening concerns about the potential exposure or misuse of sensitive data. The ever-present threat of ransomware and other cyberattacks also drives demand for greater security and data protection capabilities in enterprise infrastructure. Many organizations must confront regulatory requirements, including the European Union's Digital Operational Resilience Act, AI Act, Cyber Resilience Act, and Directive 2022/2025 for protecting network and information systems, also known as NIS2. Hardware failures, system outages, natural disasters, and other catastrophic events pose additional threats that can lead to costly business disruptions.

According to IDC's August 2024 *Future Enterprise Resiliency and Spending (FERS) Survey, Wave 8*, enhancing cyber-resiliency and -recovery is the top area where organizations expect to increase spending in 2025. IDC's November 2024 *FERS Survey, Wave 11,* showed the primary drivers for the increased spending plans are the insufficiency of current cyber-recovery/cyber-resiliency solutions in preventing or mitigating losses from cyberincidents and their failure to meet requirements for compliance or cyberinsurance. Looking deeper into the impact of ransomware attacks, IDC's June 2024 *FERS Survey, Wave 6,* found that 22% of respondents had a business disruption of a day or less from their most recent incident, 45% experienced a few days of disruption, and 32% had the impact stretch for a week or more. Among the affected organizations that paid a ransom even though they had data backups, the top reason was that they could not determine the breach's extent in a timely manner and considered the payment necessary to ensure complete recovery.

AT A GLANCE

KEY STATS

According to IDC survey data:

- Top drivers for increased spending on cyber-recovery/cyber-resiliency include the insufficiency of current solutions to prevent or mitigate losses from cyberattacks and their failure to meet requirements for compliance.
- Cyberincidents disrupted business for a few days at 45% of the polled organizations and for a week or more at 32% of the polled organizations.

KEY TAKEAWAYS

- Enterprise storage vendors are adding AI/ML capabilities designed to detect cyberthreats with greater speed and accuracy.
- Early detection offers the potential to help contain the cyberattack's impact and accelerate recovery.

Cyber-resiliency and cybersecurity are now key focus areas for many providers of enterprise storage infrastructure. Some are bolstering their systems with AI and ML capabilities designed to identify abnormal behavior and activities that could signal the onset of a cyberattack or an infrastructure problem. Earlier detection holds out the possibility of a faster response to limit the impact of the attack or issue, speed recovery and, in some cases, help organizations take proactive steps to improve security or prevent a system outage.

Definitions

Resiliency, or the ability to bounce back from an adverse event, can take many forms in enterprise environments. Cyber-resiliency generally refers to security incidents and attacks. Data resiliency applies to data loss or breaches. Operational resiliency relates to recovery from business disruptions.

NAND flash is a type of nonvolatile memory that does not require power to retain data. Flash-based solid state drives (SSDs) and custom-built modules have gained widespread adoption as a higher-performance alternative to spinning hard disk drives (HDDs). NAND flash stores data in individual memory cells. The flash types available today are:

- Single-level cell (SLC), storing 1 bit per cell
- Multilevel cell (MLC), storing 2 bits per cell
- Triple-level cell (TLC), storing 3 bits per cell
- Quad-level cell (QLC), storing 4 bits per cell

Among the flash types, SLC flash supports the highest levels of cell endurance and data write speeds but costs the most on a price-per-gigabyte basis. Increasing flash density lowers performance, endurance, and storage manufacturing costs.

Benefits

Enterprise storage systems with AI and ML capabilities offer the potential to detect irregular system behavior and unusual activity with greater speed and accuracy to help organizations respond to cyberattacks and IT system issues. Advanced AI/ML models can analyze massive amounts of data to distinguish normal behavior patterns from those that could signal a malicious attack or ransomware encryption, and those capabilities are taking on added importance as bad actors also exploit AI to mount increasingly sophisticated threats. Storage may be the proverbial last line of defense, but early detection holds out the prospect of a quicker or even an automated response to contain the impact of the cyberattack and accelerate the recovery process. AI can also analyze historical data to predict potential threat vectors and vulnerabilities and provide helpful insight that organizations can use to strengthen security and curb the potential for systemwide outages.

Considering IBM FlashSystem

IBM introduced its FlashSystem storage array in 2013, leveraging technology from the company's 2012 acquisition of Texas Memory Systems. IBM redesigned FlashSystem the following year and, over time, added a wide range of enterprise storage features, including data compression and deduplication, dynamic tiering, thin provisioning, snapshots, replication, and support for latency-lowering nonvolatile memory express (NVMe) technologies. According to IDC's Enterprise Storage Systems Tracker, IBM FlashSystem is the company's leading storage array with more than $1 billion in annual revenue and 3EB of shipped capacity each year since 2022. IDC data shows that IBM ranks among the top 5

vendors worldwide in all-flash arrays (AFAs), and the company's AFA revenue and shipped capacity each grew by more than 15% on a year-over-year basis during the first three quarters of 2024.

A major distinguishing characteristic of IBM FlashSystem is the custom FlashCore Module (FCM) that IBM designed for high-performance data storage instead of using "off-the-shelf SSDs" from third-party suppliers. IBM's NAND flash modules have evolved considerably over the past decade, from single-level and multilevel cell to more economical triple- and quad-level cell, and from two-dimensional (2D) to three-dimensional (3D) lithography, to enable greater storage density in a smaller datacenter footprint. To address the reduced endurance and performance inherent with the shift to denser flash, manufacturers such as IBM have developed technologies and techniques to better manage flash, optimize data placement, and make QLC suitable for enterprise use. IBM introduced QLC in 2020 with its second-generation FCM2, at a maximum raw capacity of 38.4TB and "effective" capacity of up to 88TB after data compression. Improvements to the hardware-implemented, always-on inline compression in 2022 with the FCM3 update raised the effective capacity to 115TB, without impact to performance, according to IBM. The FCM3 update also showcased IBM's design efforts to enable SLC pages for high performance and QLC flash pages for high capacity within the same module, using custom software to ensure the most active data stays in the fastest flash. IBM claimed the FCM3 improvements would enable QLC to perform like TLC and allow IBM FlashSystem to target enterprise workloads without restrictions, regardless of their performance or capacity requirements.

One of the main new additions with the FCM4 update in 2024 was AI-driven threat detection capabilities designed to work in tandem with IBM system and management software to speed the identification of a potential cyberattack, reduce the spread and impact, and accelerate recovery. FCM4 collects information in real time on every input/output (I/O) operation, analyzes more than 40 types of data statistics, and transmits a summary to IBM Storage Virtualize. The Storage Virtualize software supports inline data corruption detection and runs an AI inference engine on every IBM FlashSystem. The capabilities can be helpful to protect against not only cyberattacks but also system outages caused by issues such as corrupted data, misapplied updates, and operator mistakes. The AI inference engine analyzes normal system behavior patterns and uses IBM Research–developed ML models trained on real-world ransomware to identify unusual activity that could signal the start of a ransomware attack or the presence of malware or a potential system issue. IBM FlashSystem can then trigger alerts to IBM Storage Insights Pro to initiate a response, potentially to IBM support for customers that enable the "phone home" feature. The system can also share the threat detection information with IBM Storage Defender security software, which is available for purchase separately.

The 2024 Storage Defender release added AI-driven sensors that IBM Research developed to improve the accuracy of ransomware and general threat detection. Defender can alert security tools to try to reduce the blast radius of the security breach and assist enterprises in recovery from the attack. New Storage Defender workload and inventory management capabilities give organizations the option to develop a business continuity plan. Defender also supports the ability to orchestrate and automate the recovery of VMware applications. Using IBM FlashSystem and Storage Defender together enables potential capabilities such as building protection groups of specific volumes based on user-defined policies to help organizations speed the recovery of critical data. IBM FlashSystem can automate the creation of immutable "safeguarded copy" snapshots, and users have the option to recover the immutable copies at multiple target locations and replicate them to another Storage Defender cluster for added protection. IBM offers a FlashSystem cyber-recovery guarantee for the restoration of the immutable snapshots from the safeguarded copy within 60 seconds, once configured by IBM Expert Labs.

Challenges

While AI/ML-driven technologies can help speed the detection of potential cyberattacks, they can also produce false positives or errant results that could potentially disrupt storage systems, applications and, in turn, business operations. IBM has taken steps to improve the accuracy of its threat detection capabilities with the latest IBM FlashSystem and FCM4 releases, but tuning models and algorithms remains a challenge for vendors offering AI/ML-enabled technologies, since normal behavior patterns can sometimes mimic suspicious activity.

Although IBM FlashSystem's and Storage Defender's anomaly detection capabilities may help with the identification of and recovery from cyberattacks, organizations need to take a defense-in-depth approach well beyond the scope of storage systems to guard against the increasingly sophisticated ransomware, malware, and other threats to their applications and datacenter infrastructure. Key components also include endpoint detection and response (EDR), cloud security gateways, security information and event management (SIEM), and user and entity behavior analytics (UEBA).

Conclusion

Strengthening cyber-resilience and operational resilience is a critical priority for organizations that must confront the ongoing threat of ransomware and other cyberattacks, growing concerns over data privacy and security with the use of AI applications, and expanding regulatory requirements for data protection and business continuity. IDC research shows that enhancing cyber-resiliency and -recovery is the top spending priority for 2025, and many enterprise storage vendors are adding AI- and ML-enabled capabilities designed to recognize potential cyberattacks with greater speed and accuracy. In 2024, IBM FlashSystem, with its custom-built FlashCore Module, introduced new AI/ML-driven anomaly detection capabilities that IBM claims can quickly identify ransomware and other cyberthreats and accelerate response and recovery in tandem with IBM Storage Insights and, optionally, Storage Defender software. Although AI/ML technologies can produce false positives, IDC believes the technology will continue to improve and become more commonplace in enterprise storage systems. AI/ML-based enhancements such as the ones in IBM FlashSystem and FCM4 will not obviate the need for a comprehensive defense-in-depth strategy beyond storage systems to combat increasingly sophisticated cyberthreats, including AI-driven attacks. However, to the extent that IBM can confront the challenges outlined in this paper, the company has a significant opportunity to help enterprises address their needs for enhanced cyber, operational, and data resiliency.

About the Analyst

Carol Sliwa, *Research Director, Storage and Converged Systems, Worldwide Infrastructure Research*

Carol Sliwa is a research director in the Storage and Converged Systems practice of IDC's Worldwide Infrastructure Research organization. She is the lead analyst for IDC's storage infrastructure coverage, and her core research spans block, file, and object storage, with a special focus on unstructured data. Carol's coverage also includes flash storage media and scale-out storage systems for performance-intensive computing environments. With more than 25 years of experience as a technology journalist, including 13 years covering enterprise storage, Carol gained extensive insight into the ways in which the IT industry has developed technologies, platforms, and systems over time to address the evolving needs of IT organizations.

MESSAGE FROM THE SPONSOR

IBM FlashSystem is designed to enhance your operational resilience strategy with its advanced cyber resilience and storage management capabilities.

Assess your business's preparedness against cyber-attacks today with a free cyber resilience assessment at: **www.tri-delta.com**

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
blogs.idc.com
www.idc.com