IBM'S ENHANCED FLASHSYSTEM MALWARE DETECTION

STEVE MCDOWELL, CHIEF ANALYST
FEBRUARY 27, 2024

CONTEXT

IBM's new FlashSystem ransomware detection capabilities are driven by advanced AI technologies, new fourth-generation FlashCore Module (FCM) technology, and IBM Storage Defender software enhancements.

These innovations aim to bolster organizational defenses against the increasing threat of ransomware and other cyber attacks by providing earlier and more accurate detection, enabling faster response and recovery.

BACKGROUND: IBM FLASHCORE MODULE (FCM)

The IBM FlashCore Module is an advanced computational storage device that significantly enhances the capabilities of standard solid-state drives (SSDs). Unlike ordinary SSDs that primarily serve data across an NVMe interface, the FlashCore Module incorporates a complex blend of technologies and computational power to improve performance, reliability, and storage efficiency, particularly for QLC (quad-level cell) flash memory.

This device marks a significant evolution in storage technology, offering benefits not only to IBM's portfolio but also to its storage customers on a broader scale.

Here's a detailed look at its features and functionalities:

- **Computational Storage Device:** The FlashCore Module combines NAND flash memory with DRAM and MRAM for caching alongside substantial computational resources. This setup allows it to perform tasks typically handled by the storage array, such as data compression, directly on the drive. This architecture makes the FlashCore Module far more efficient and flexible than traditional SSDs.

- **Enhanced QLC Flash Performance:** QLC flash is known for its higher density and lower cost than TLC (triple-level cell) flash, albeit at the expense of speed and endurance. The FlashCore Module's sophisticated computational logic and management algorithms allow it to mitigate these drawbacks, enabling QLC flash to deliver performance and reliability on par with, or even superior to, TLC-based solutions.

- **High Capacity and Endurance:** IBM has engineered the FlashCore Module to offer exceptional storage capacities, ranging from 22 TB to 116 TB, all within a compact 2.5-inch U.2 NVMe form factor. Moreover, it boasts twice the endurance of standard NVMe flash drives, making it an ideal solution for enterprise storage environments that demand high durability and longevity.

- **Customizable and Adaptive:** The module's computational capabilities are powered by ARM processor cores embedded in a reprogrammable FPGA (Field-Programmable Gate Array). This design allows flexibility and adaptability, enabling IBM to tailor the FlashCore Module's functionalities to meet specific storage requirements and challenges.

- **Advanced Data Management:** The FlashCore Module addresses complex data management tasks beyond typical storage functions. Initially focused on data compression to improve efficiency and reduce costs, IBM has expanded the module's capabilities to include advanced functions like real-time analytics, data filtering, and unstructured data management directly at the storage level.

NEW FLASHSYSTEM MALWARE DETECTION CAPABILITIES

IBM's new FlashSystem ransomware detection capabilities operate on a multifaceted approach that uses AI/ML technologies within its storage system to intelligently detect malware in the I/O path, triggering the appropriate notifications to the system administrator.

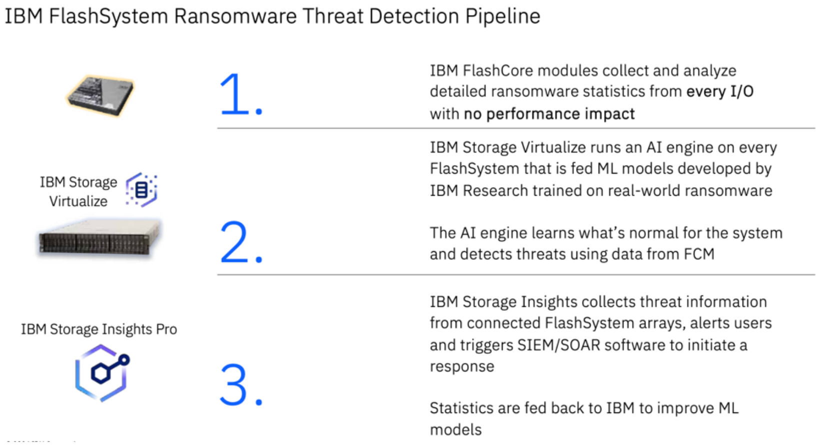

IBM FlashSystem Ransomware Threat Detection Pipeline

1. IBM FlashCore modules collect and analyze detailed ransomware statistics from **every I/O** with **no performance impact**

2. IBM Storage Virtualize runs an AI engine on every FlashSystem that is fed ML models developed by IBM Research trained on real-world ransomware

The AI engine learns what's normal for the system and detects threats using data from FCM

3. IBM Storage Insights collects threat information from connected FlashSystem arrays, alerts users and triggers SIEM/SOAR software to initiate a response

Statistics are fed back to IBM to improve ML models

Here's a breakdown of how this detection works:

1. Inline Data Scanning at Block Level:

IBM FlashSystem products scan all incoming data with block-level granularity as it's written. This process involves inline data corruption detection software and cloud-based AI to identify anomalies that might signal the onset of a cyber-attack, including ransomware. This early detection mechanism allows for immediate response actions to mitigate the attack.

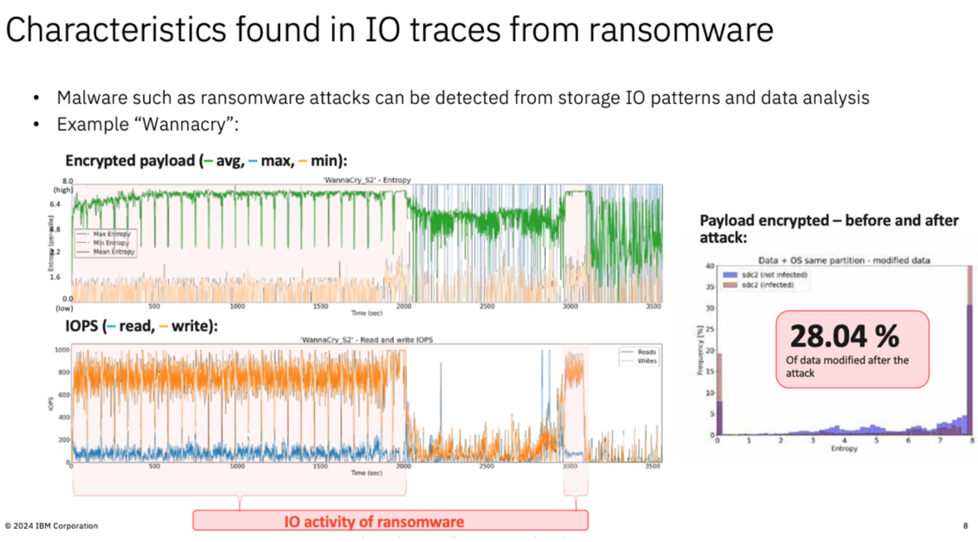

Characteristics found in IO traces from ransomware

- Malware such as ransomware attacks can be detected from storage IO patterns and data analysis
- Example "Wannacry":

Encrypted payload (– avg, – max, – min):

IOPS (– read, – write):

Payload encrypted – before and after attack:

28.04 %

Of data modified after the attack

IO activity of ransomware

© 2024 IBM Corporation

8

2. AI-Powered Anomaly Detection:

The fourth generation of FlashCore Module (FCM) technology enables AI capabilities within the IBM Storage FlashSystem family. Using machine learning models, it continuously monitors statistics gathered from every input/output (I/O) operation. IBM Research trained these modesl to detect anomalies that resemble ransomware behavior, such as unusual encryption patterns, in less than a minute. This rapid detection allows organizations to respond quickly to threats.

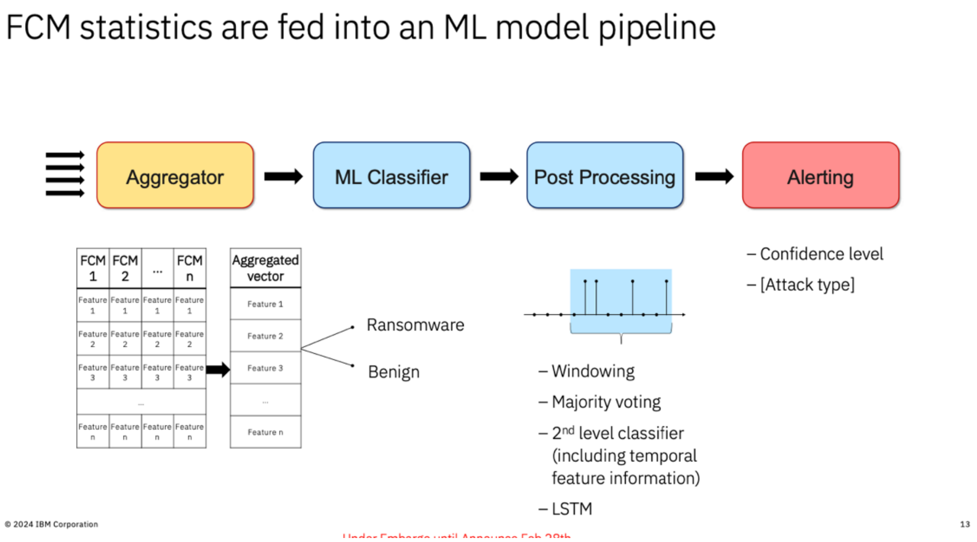

FCM statistics are fed into an ML model pipeline

– Windowing
– Majority voting
– 2nd level classifier (including temporal feature information)
– LSTM

– Confidence level
– [Attack type]

© 2024 IBM Corporation

13

Under Embargo until Announce Feb 28th

3. Monitoring Data Characteristics:

IBM FlashSystem products analyze parameters such as the data's compressibility and randomness (entropy). This information is passed to IBM Storage Insights software, which can alert operators when a workload anomaly is detected. For instance, if ransomware begins to encrypt an application's data, the system can detect this unusual behavior based on the change in data characteristics.

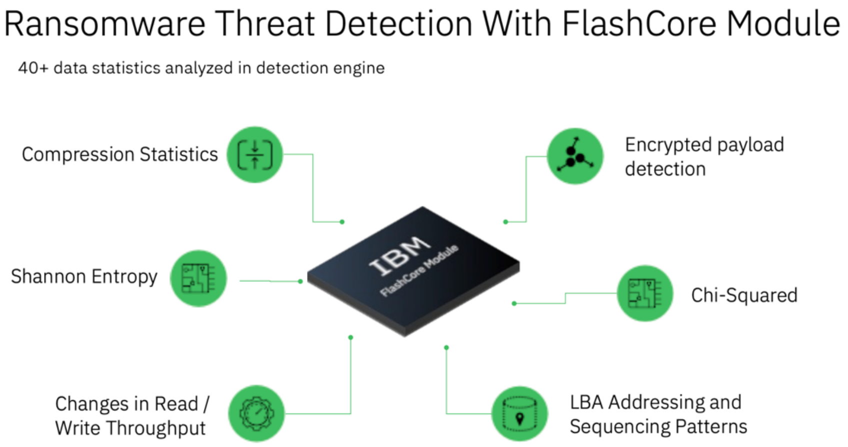

Ransomware Threat Detection With FlashCore Module

40+ data statistics analyzed in detection engine

Compression Statistics
Encrypted payload detection
IBM
FlashCore Module
Shannon Entropy
Chi-Squared
Changes in Read / Write Throughput
LBA Addressing and Sequencing Patterns

4. Real-Time Statistics and Machine Learning Models:

The FCM technology captures and summarizes detailed statistics about every I/O in real-time. By leveraging machine learning models, the system can distinguish between normal behavior and patterns indicative of ransomware or malware, allowing immediate corrective actions to protect the data.

5. Integration with IBM Storage Defender Software:

The updated IBM Storage Defender software extends these capabilities across various IT environments. It includes AI-powered hardware and software sensors that provide an industry-leading index of the relative trustworthiness of data copies, whether they are backup or primary snapshots. This helps in identifying potentially compromised data more accurately.

Holistic Data Resiliency with IBM Storage Defender
Unified Management for Data Resiliency across all storage

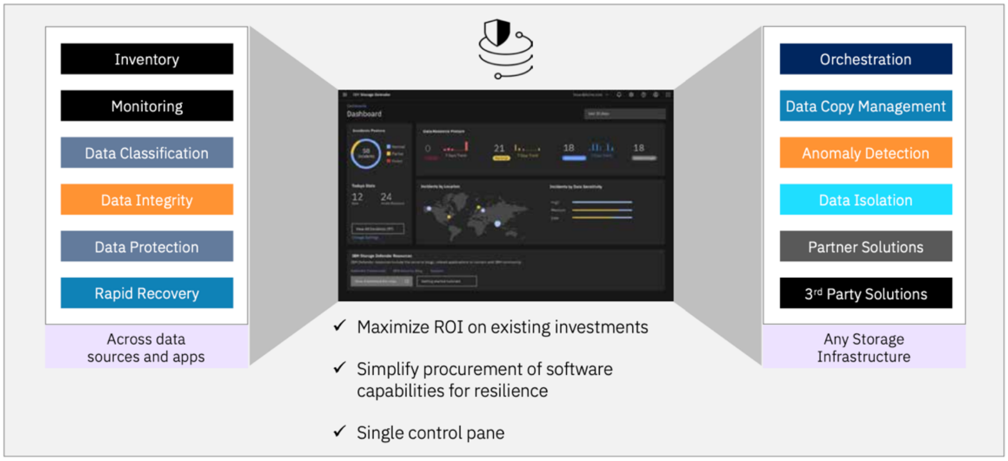

- ✓ Maximize ROI on existing investments
- ✓ Simplify procurement of software capabilities for resilience
- ✓ Single control pane

UPDATED: IBM STORAGE DEFENDER

Central to IBM's data protection is its upgraded IBM Storage Defender software. The new release uses AI-powered sensors to integrate FlashSystem detection with broader threat detection capabilities across various IT environments, including VMs, databases, applications, and containers.

The new version also introduces workload and storage inventory management features to aid organizations in incorporating their assets into a business continuity plan for post-attack recovery. Integration with other IBM Storage and Security solutions and third-party data platforms provides end-to-end data resilience.

<is there a graphic?>

Key innovations include the ability for storage administrators to create protection groups for automatic backup and to restore or recover immutable data copies free from threat signatures to multiple locations. IBM has also integrated settings for creating Safeguarded Copy snapshots, which are cyber-resilient and isolated from production data, enabling quicker recovery after data loss events.

ANALYSIS

IBM's new ransomware detection capabilities are a significant advancement in cyber defense, leveraging the power of AI to enhance the detection, response, and recovery processes. By integrating these technologies into its storage solutions, IBM aims to provide enterprises with a more robust defense against the evolving landscape of cyber threats, ensuring greater data resilience and faster recovery from incidents.

IBM’s sophisticated approach enables its storage solutions to provide early warnings of cyber threats, allowing enterprises to recover faster and maintain the confidentiality and security of their data against ransomware and other cyber-attacks.

Data protection and cyber-resiliency features are quickly becoming standard offerings within enterprise storage solutions. Immutable snapshots, for example, are now part of solutions from nearly every top-tier storage vendor. Scanning snapshots using entropy calculations like those used by IBM for anomalies is also becoming popular.

The challenge with the approaches taken by most storage solutions is that detection often happens after the fact. Once a snapshot is found to be corrupt, it may be too late. You have to work backward in time to find good data to restore.

IBM's approach closes that gap, alerting the user nearly instantly when an anomaly is detected. This happens long before the corrupted data is written to a snapshot. That's precisely the kind of protection you want in your enterprise.

IBM isn’t the only storage provider building this level of real-time anomaly detection directly into its storage arrays. NetApp’s Autonomous Ransomware Detection feature in its ONTAP products also looks at entropy to detect anomalies and raise alerts.

Unlike IBM, NetApp does the computation on the storage controllers, not the underlying drives. NetApp's approach also doesn't detect ransomware signatures; instead, it relies on entropy, file extensions, and the amount of file write activity. IBM is currently alone in delivering this level of storage-based threat detection.

While most of the storage industry spends R&D money trying to coax enterprise-class storage out of commodity server-class hardware, IBM is one of only a small handful of technology providers continuing to innovate storage with a full-stack view of the underlying platform. IBM’s storage solutions remain among the most technologically advanced on the market, making them ideal to protect critical enterprise data. The new ransomware detection features only solidify that position.

The information contained within this document was believed by NAND Research to be reliable and is provided for informational purposes only. The content may contain technical inaccuracies, omissions, or typographical errors. This document reflects the opinions of NAND Research, which is subject to change. NAND Research does not warranty or otherwise guarantee the accuracy of the information contained within.

NAND Research is a technology-focused industry analyst firm providing research, customer content, market and competitive intelligence, and custom deliverables to technology vendors, investors, and end-customer IT organizations.

Contact NAND Research via email at info@nand-research.com or visit our website at nand-research.com.