Ransomware Detection
Using Storage-Embedded AI

Andy Walls – Consultant
Owner Great Walls of
Storage LLC

12/10/2024

IEEE-CNSV
Consultants' Network
of Silicon Valley

My Story

- **Retired IBM Fellow and CTO Flash Storage**
- **43 years with IBM**
- **Started Great Walls of Storage**
- **Renown expert in SSD design, AFAs and IT.**
- **I Consult on all aspects of Storage Systems development and deployment.**

LOCK FILE

ALL YOUR IMPORTANT FILES ARE ENCRYPTED!

Any attempts to restore your files with the thrid-party software will be fatal for your files!

Restore you data posible only buying private key from us.

There is only one way to get your files back:

01.

contact us

UTox Email

qTox ID:
- B2F873769EB6B508EBC2103DDEB7366CEFB7B09AB8314DAD0C4346169072

https://tox.chat/download.html

- Email: contact@contipauper.com

02.

Through a Tor Browser - recommended

- Download Tor Browser - https://www.torproject.org/ and install it.
- Open link in Tor Browser - http://zqaflhty5hyziovsxgqvj2mrz5e5rs6oqxzb54zolccfnvtn5w2johad.onion This link only works in Tor Browser!
- Follow the instructions on this page

ATTENTION!

- Do not try to recover files yourself. this process can damage your data and recovery will become impossible
- Do not rename encrypted files.
- Do not waste time trying to find the solution on the Internet. The longer you wait, the higher will become the decryption key price
- Decryption of your files with the help of third parties may cause increased price (they add their fee to our).
- Tor Browser may be blocked in your country or corporate network. Use https://bridges.torproject.org or use Tor Browser over VPN.
- Thanks to the warning wallpaper provided by lockbit, it's easy to use

Secure & Resilient

Data is Under threat

85%

not able to fully restore data from back up after an attack

66%

of breaches were not identified by the organization's internal security teams and tools

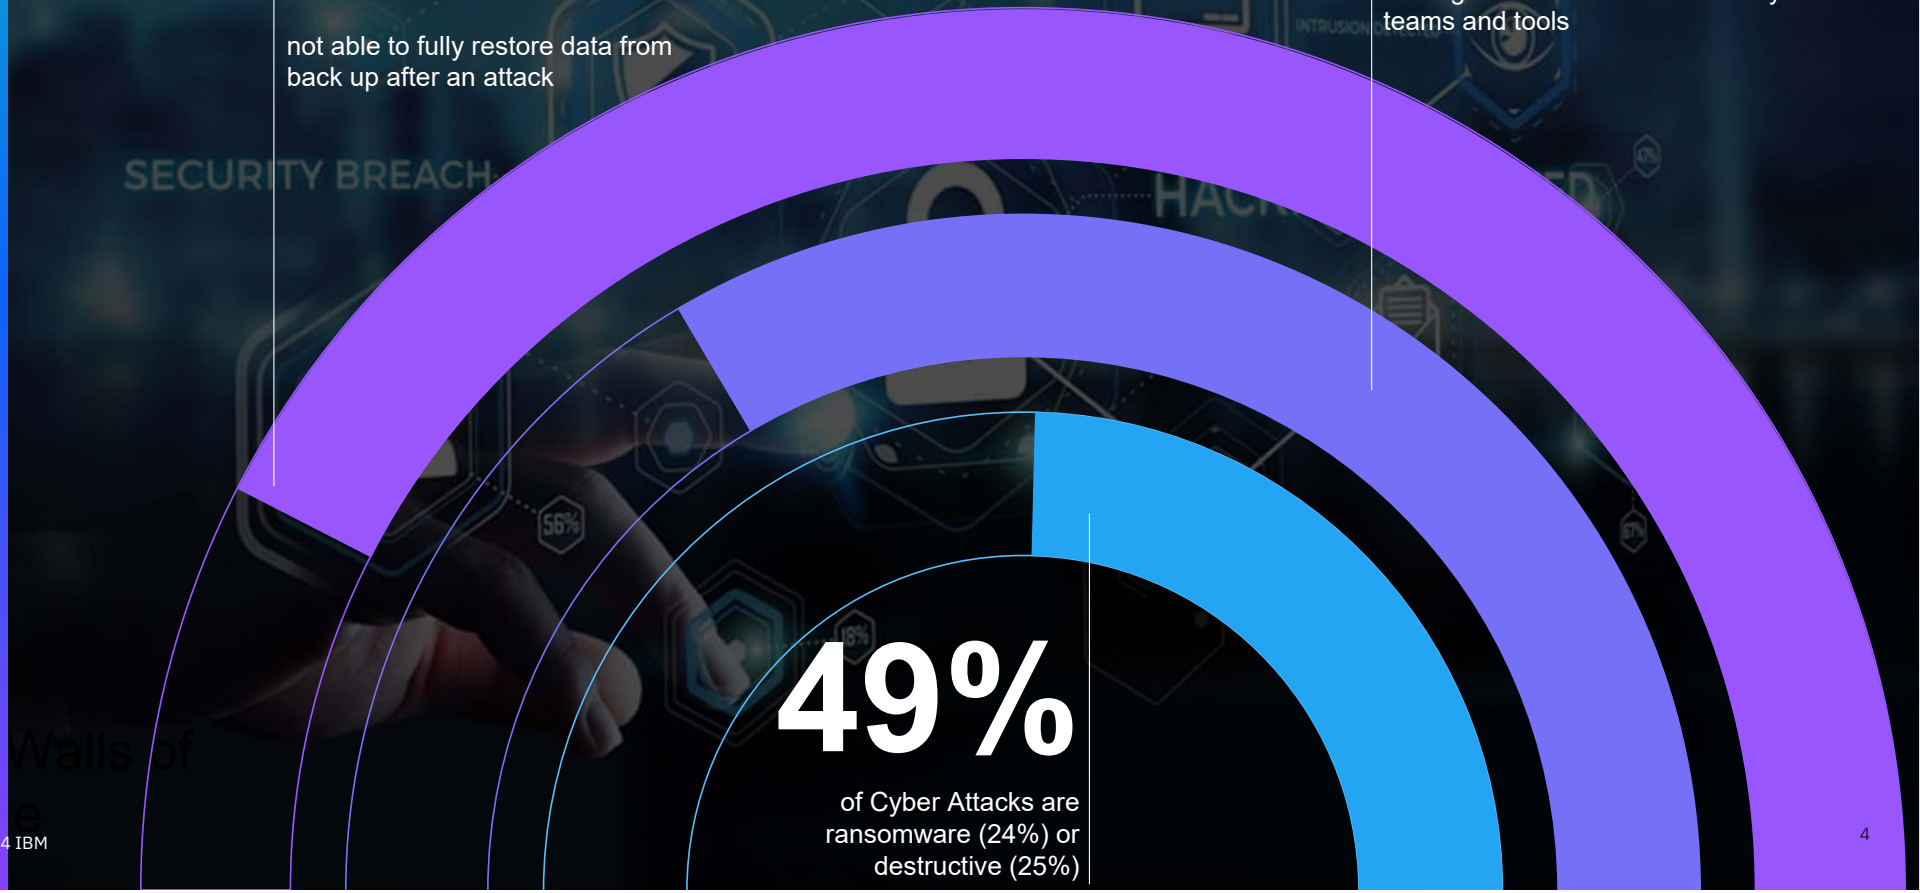

49%

of Cyber Attacks are ransomware (24%) or destructive (25%)

Walls of
e

Bad actors are getting faster!

It takes less than 5 mins to encrypt 100GB

Too late if we detect AFTER encryption

60+ days
2019 ransomware deployment time

9.5 days
2020 ransomware deployment time

3.85 days
2021 ransomware deployment time

Bad Actors are moving Faster
Less than 4 days to a Ransomware attack

Data Extortion on the rise – and it is all over the world.

23
days, average recovery after a ransomware attack

Takes over 287 days for full recovery

Cyber Attacks are on the Rise, getting more sophisticated

51%
of Cyber Attacks are ransomware (24%) or exfiltration (27%)

26%
clients who paid the ransom still could not recover the data

108
days faster identification and containment of a breach with extensive security AI & automation

2X
Cyber Attacks YTY 2022 vs 2021, 2023 YTY 2.5x so far!

23
days, average recovery after a ransomware attack

66%
of breaches were not identified by the organization's internal security teams and tools

Cost of Data Breach

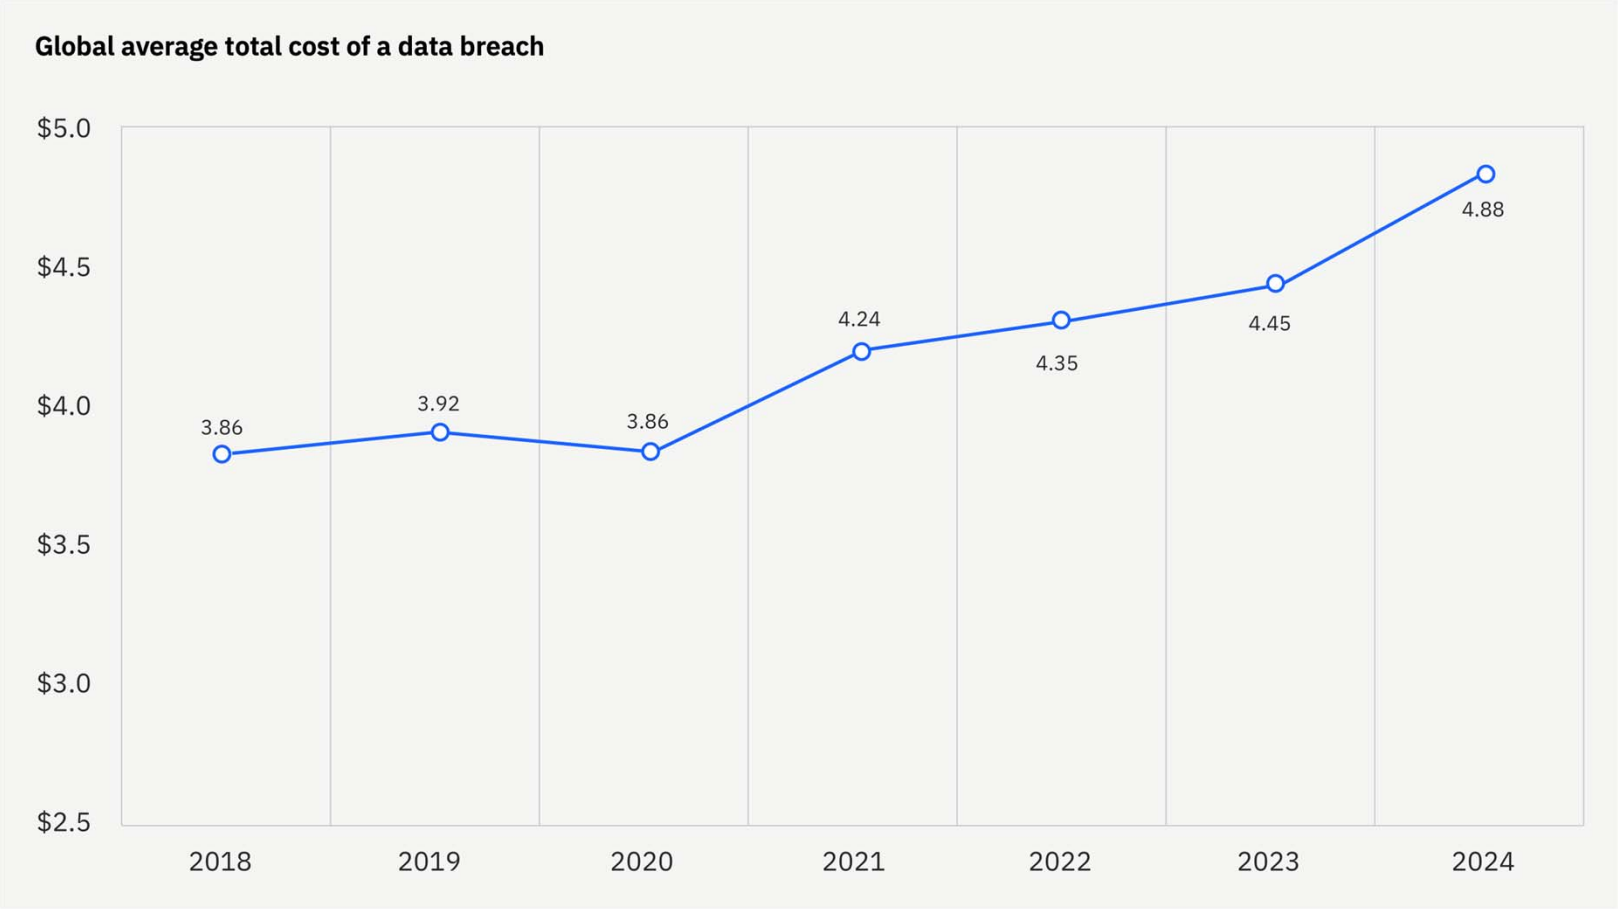

Figure 1. Measured in USD millions

Data provided by IBM – Cost of Data Breach Report 2024

Ransomware attack steps

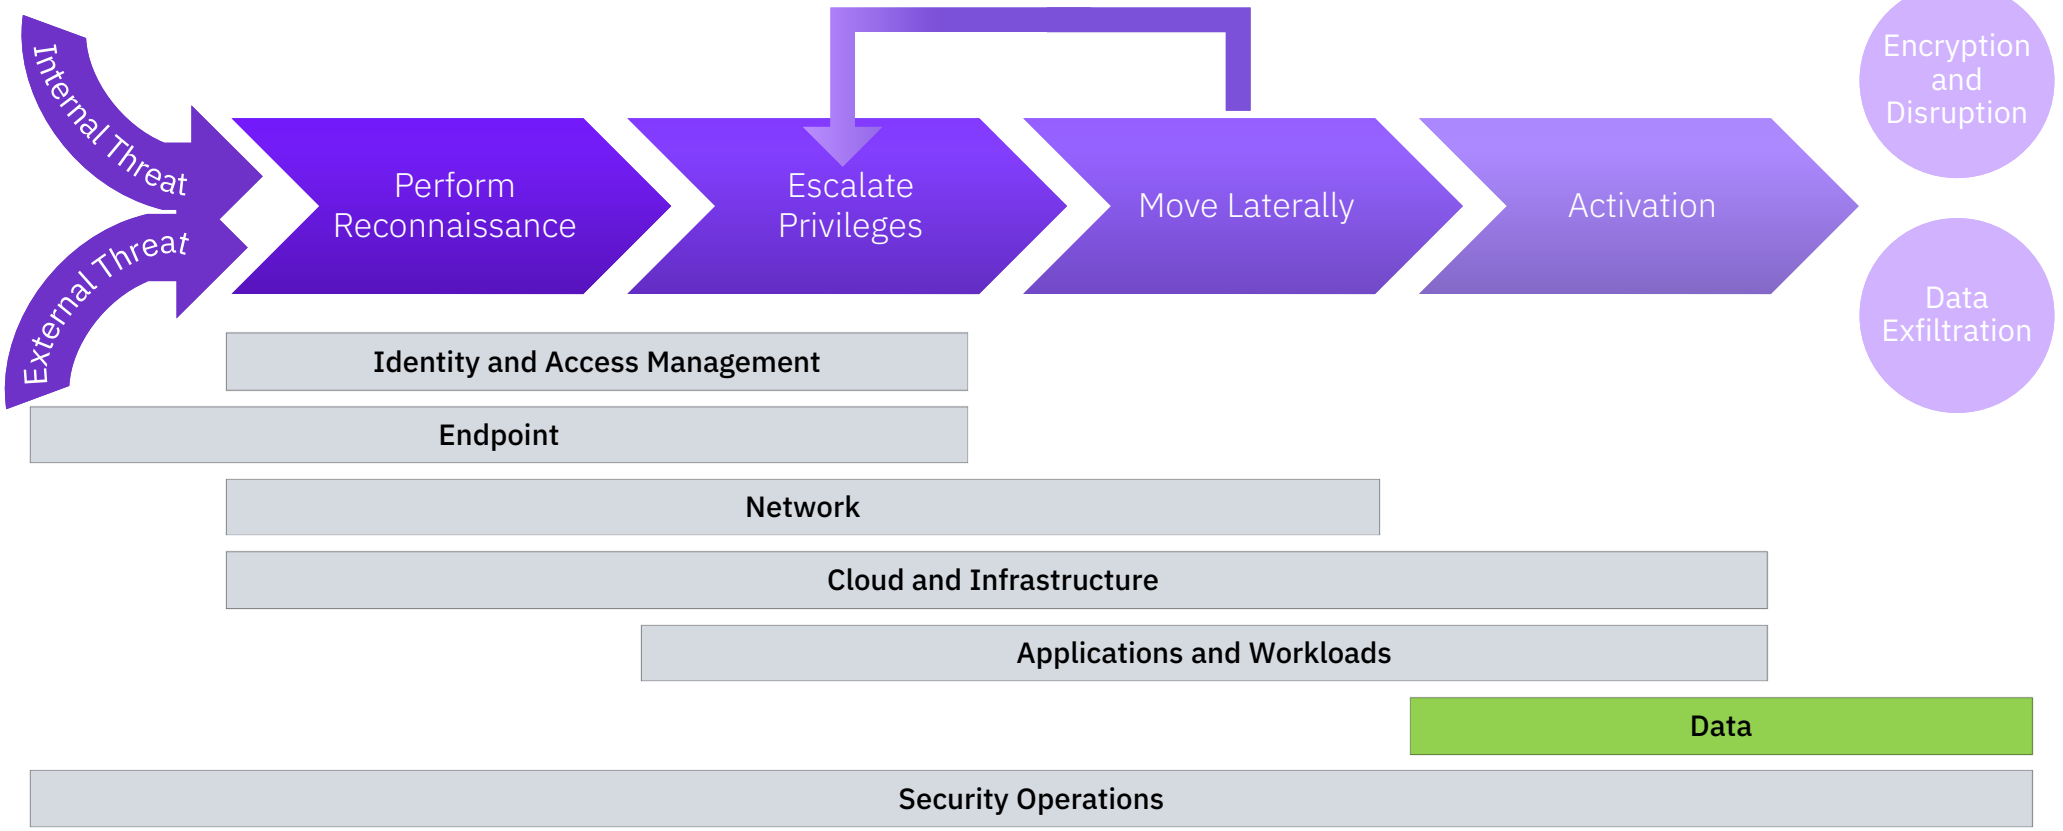

Attacks in the News

1. NHS London: **Qilin ransom gang unleashed an attack that compromised the data of almost 1 million National Heathcare System patients in London hospitals**. The attackers published personal information about patients with sensitive medical conditions like cancer and sexually transmitted diseases.Oct 16, 2024
2. In an update on January 22, LoanDepot confirmed that around 16.6 million of its customers had their sensitive personal information stolen in the incident, including Social Security numbers and financial account numbers.
3. In February, reports emerged that US healthcare payment provider, Change Healthcare, had been hit by a ransomware attack.
4. A ransomware attack on Australian medical prescriptions provider MediSecure in May led to 12.9 million individuals' personal and health data being compromised.
5. The City of Columbus, Ohio, revealed it had been hit by a ransomware attack in July, resulting in outages to some resident-facing IT services.
6. An August cyber-attack on the Port of Seattle, a local government agency overseeing the seaport of Seattle and Seattle–Tacoma International Airport (SEA), heavily disrupted travel to and from the state ahead of the US Labor Day holiday.

But How Do You Detect Ransomware

Detection By

Threat Signature	Sample Hash Comparison
Data Behavior Signals	Monitoring for Anomalies
Network Signals	Network-Level Monitoring for Anomalies

But How Do You Detect Ransomware

Detection By

Threat Signature	Sample Hash Comparison
Data Behavior Signals	Block Level Monitoring for Anomalies
Network Signals	Network-Level Monitoring for Anomalies

A Realization:

Block Storage is missing some context other parts of the system have

BUT: It can generate data needed for determining Ransomware attacks with less performance impact then any other part of the system

IBM FlashSystem excels in ingesting large amounts of data fast.

If the storage could **analyze** the data **as it is stored** we can generate critical **insights** more efficiently than external backup scanning applications.

The impressive history of FlashCore Technology

2014

MicroLatency Module

Proprietary interface, single-layer cell (SLC) flash, followed up with multi-layer cell (MLC) flash, and in both cases the data path is in hardware

Multiple protection features, including ECC error correction, variable stripe RAID data protection, overprovisioning, and three-dimensional (AE3 flash modules) or two-dimensional (AE2 flash modules) flash RAID

2018

FCM1

NVMe interface, re-implemented into a standard 2.5” form factor, triple-layer cell (TLC) flash with inline 2-to-1 data compression and encryption with no performance penalty

2020

FCM2

NVMe interface, quad-layer cell (QLC) flash with better than TLC performance, inline 2-to-1 data compression and encryption with no performance penalty

2022

FCM3

NVMe interface, quad-layer cell (QLC) flash with SLC abilities, optimized with a “Hinting Architecture” to optimize data placement, with up to 3-to-1 inline data compression, encryption with no performance penalty, L and XL modules based on PCIe G4,

2024

FCM4

NVMe interface, quad-layer cell (QLC) flash with SLC abilities, optimized with a “Hinting Architecture” to optimize data placement, with up to 3-to-1 inline data compression and, encryption, with no performance penalty, all modules based on PCIe G4, and inline intrusion detection

A breakthrough in cyber security

FCM-4

FLASHCORE MODULE 4

Built in Data Resilience to accelerate recovery from Ransomware attacks

IBM Safeguarded Copy prevents point-in-time copies of data from being modified or deleted by user errors, malicious destruction, or ransomware attacks – Logical Air Gap of Data

Separation of duties

Additional security capabilities to prevent non-privileged users from compromising production data

Protected copies of data

Capabilities to regularly create secure, immutable point in time copies – Up to 15,000 copies

Speed of recovery

Functionality that enables different use cases to restore corrupted data **in minutes or hours vs days or weeks**

IBM has developed technology not just to recover from attacks

But to detect them early!

Two types of protection implemented

- Entropy and compressibility statistics sent back to a cloud based product that looks for trends and sends alerts if an anomalous behavior is discovered.
- AI Based inline Inferencing using a trained model looking for ransomware attacks

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):

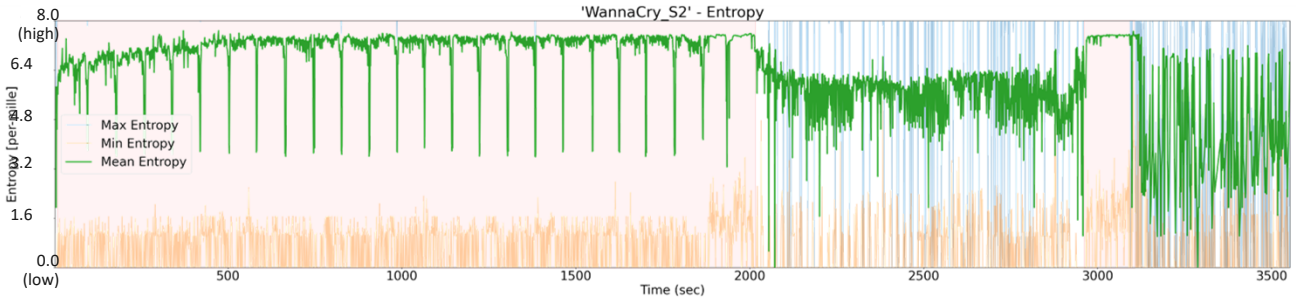

IOPS (– read, – write):

IO activity of ransomware

Payload encrypted – before and after attack:

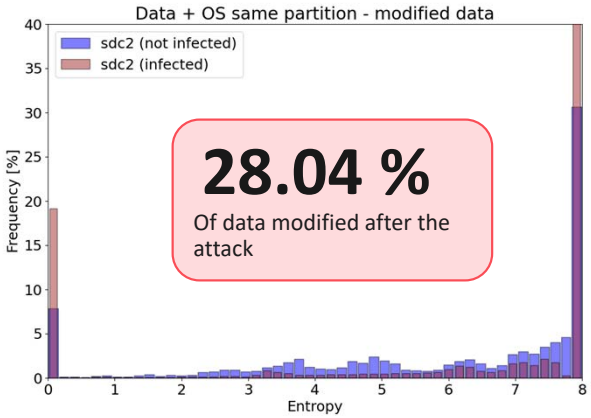

LBA access analysis – WannaCry - 1 hour

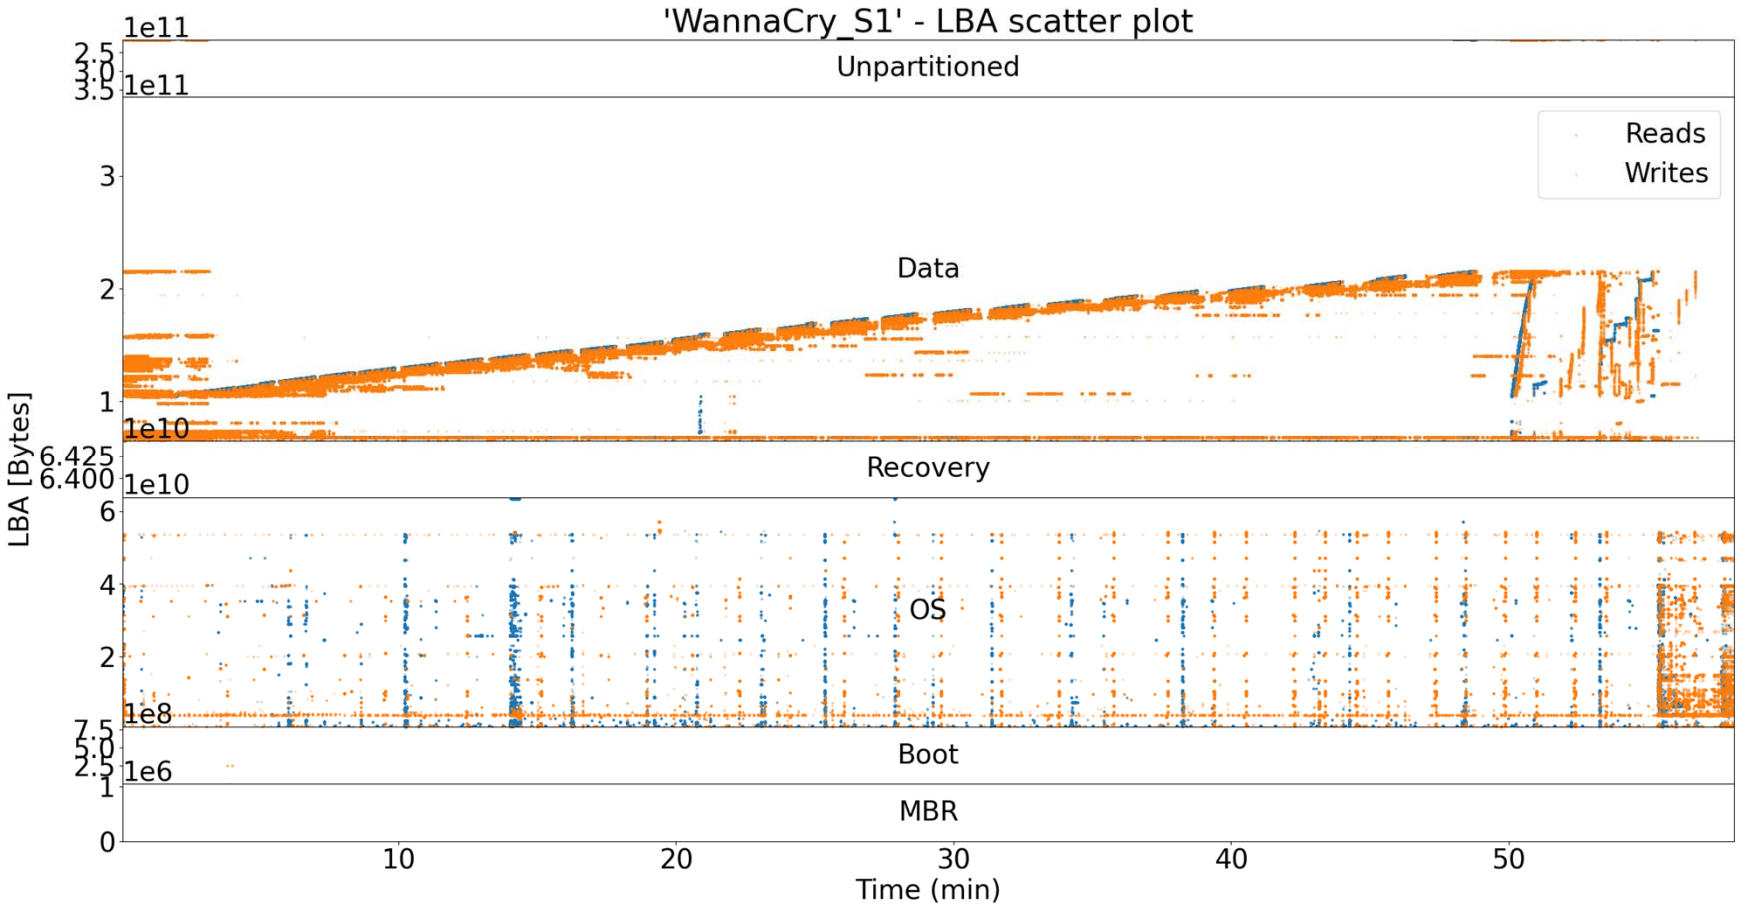

Characteristics found in IO traces from ransomware

- Malware such as ransomware attacks can be detected from storage IO patterns and data analysis
- Example “Conti”:

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

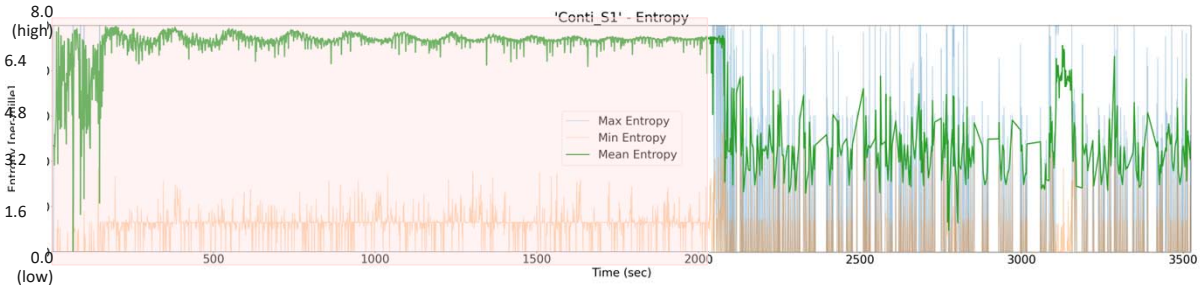

IOPS (– read, – write):

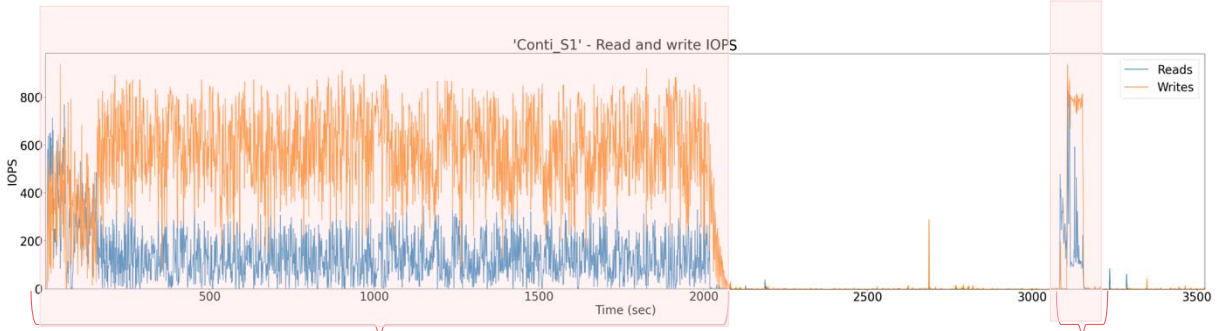

IO activity of ransomware

Payload encrypted – before and after attack:

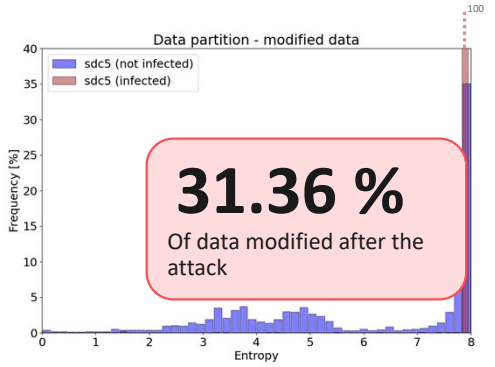

31.36 %

Of data modified after the attack

LBA access analysis – Conti - 1 hour

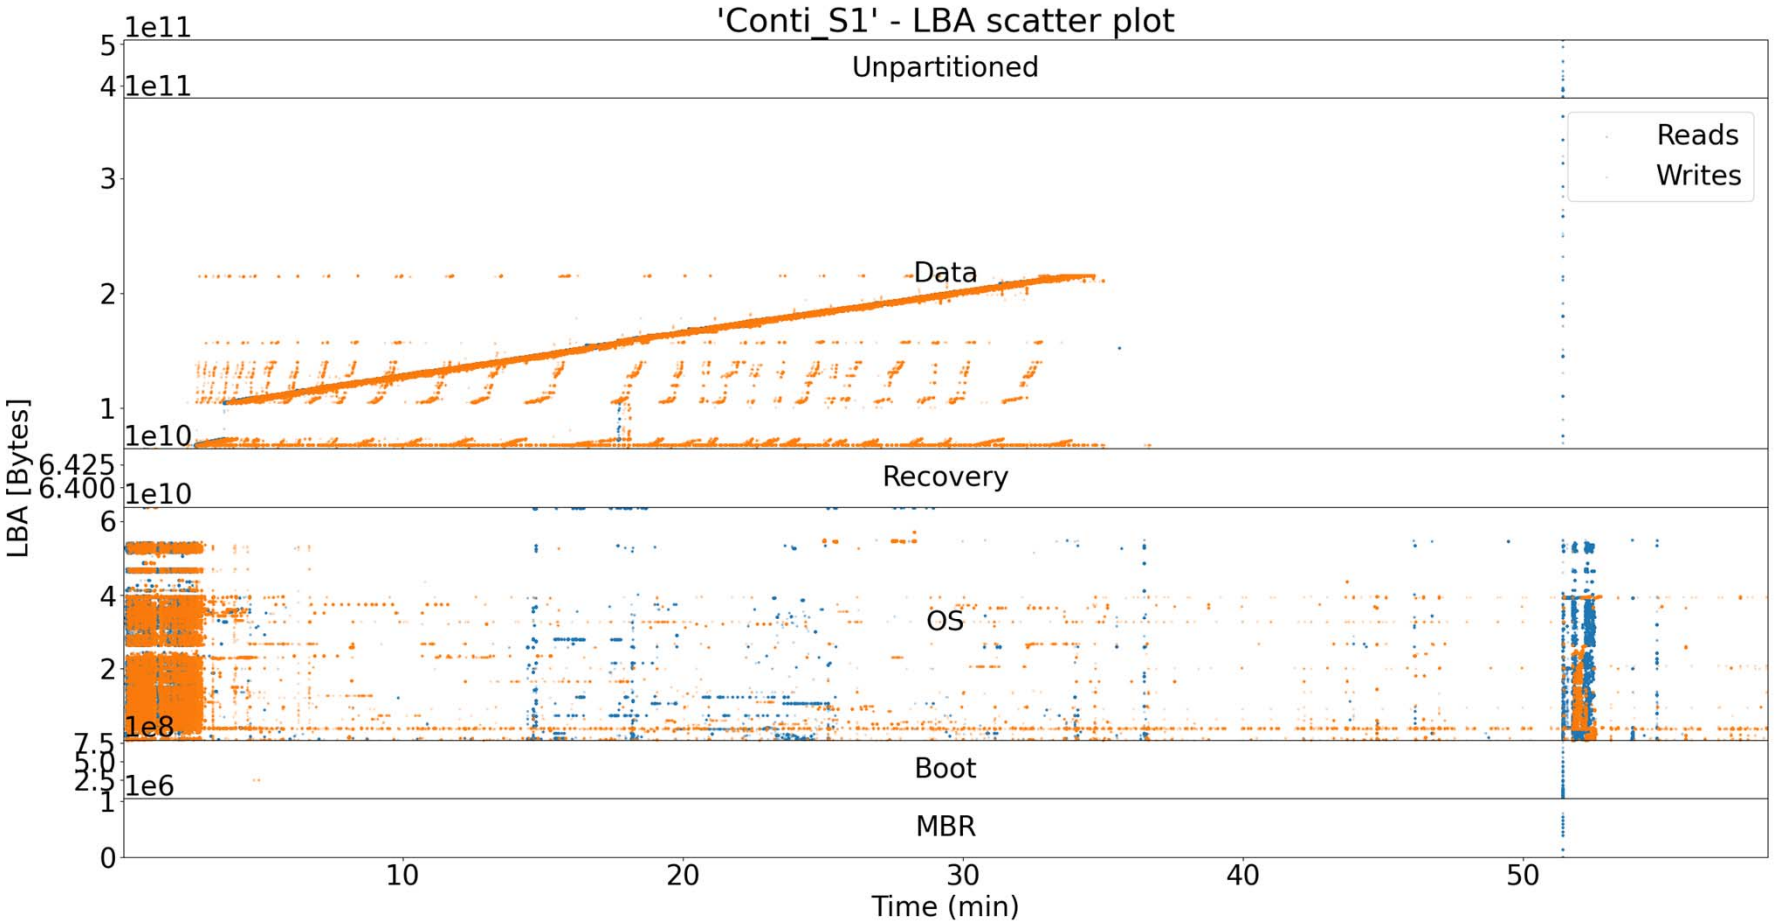

Cyber Resiliency with IBM Storage FlashSystem

Address cyber threats with a holistic approach

Early Detection and Prevention

AI powered ransomware threat detection at the drive, controller and fleet level, with FlashCore Modules, Storage Virtualize and Storage Insights

System-level resilience with Multi-Factor Authentication and dual user controls

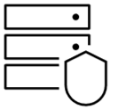

Safe Recovery

Use immutable Safeguarded Copies to create known clean copies of data and recover within 60 seconds guaranteed

Integrations with leading data protection software such as IBM Defender, Veeam, Commvault, Veritas for full business recovery

SecOps Integration

Integrations with SIEM and SOAR software such as Predatar, Qradar and Splunk for coordinated response

Orchestration with RedHat Ansible for automated recovery playbooks

Ransomware Threat Detection With FlashCore Module

40+ data statistics analyzed in detection engine

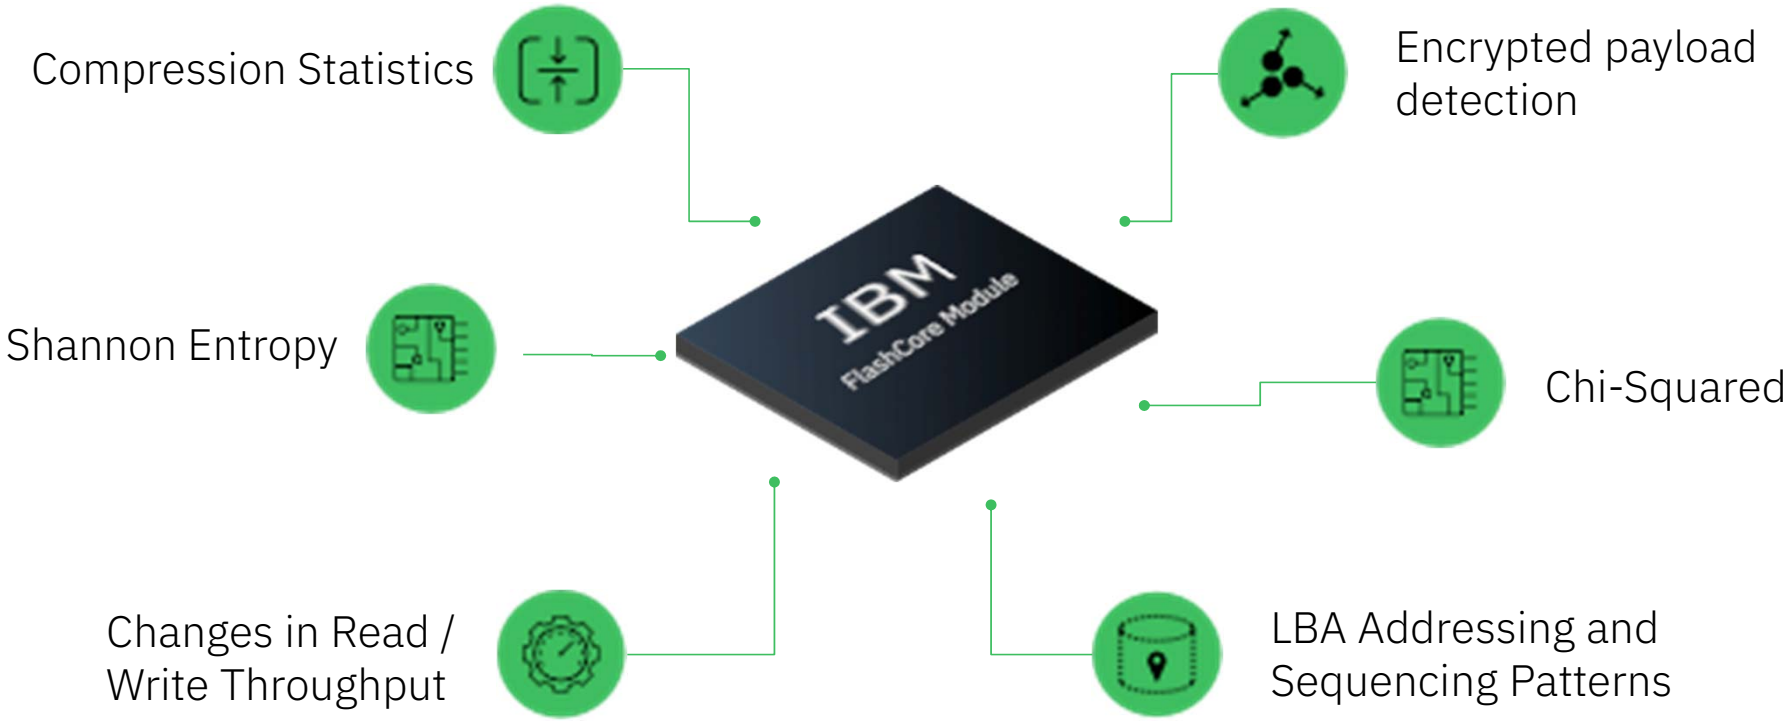

Processed on **EVERY** write with ZERO performance impact!

FCM4 and Ransomware Detection

- FCM4 calculates entropy (estimate of randomness) and change in compression on every IOP
- FCM4 keeps statistics on each IOP like block size, LBA , Rd
- FCM 4 has 2 small RISC cores process all this information
- All this information is statistically summarized into a relatively small amount of information per volume
- These summaries are passed every 2 seconds to an inference engine in Storage Virtualize.

FlashSystem ransomware detection conceptual model

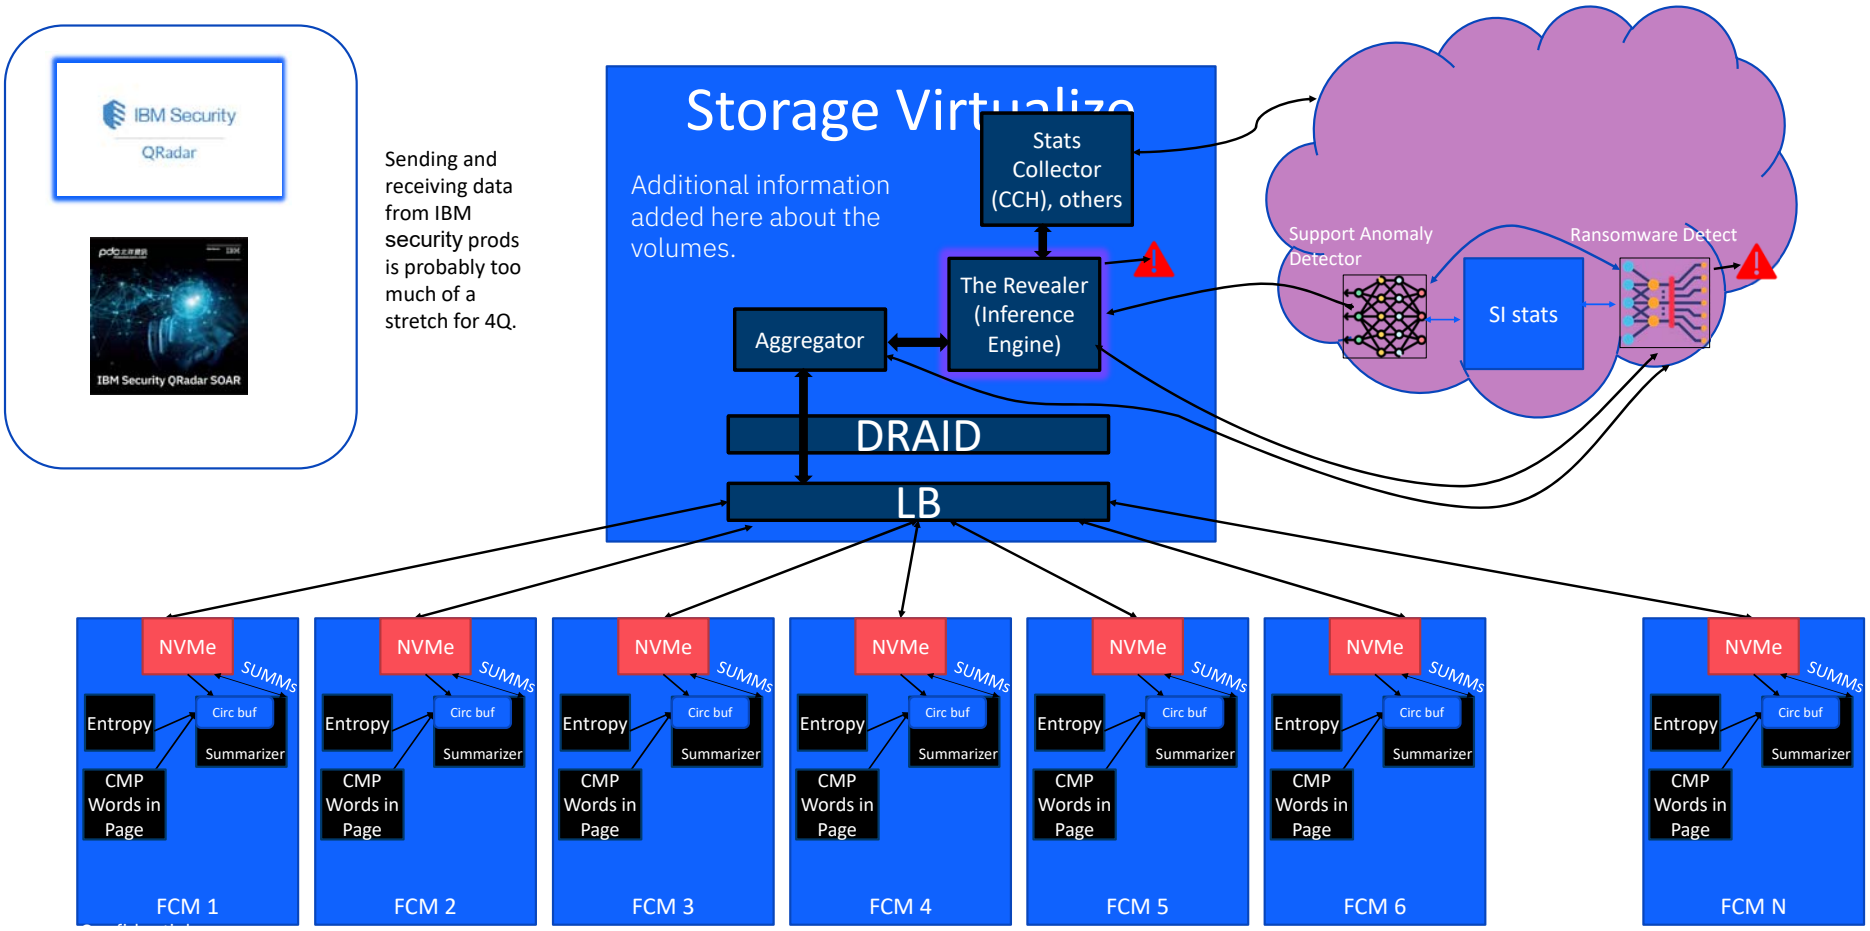

Ransomware Threat Detection

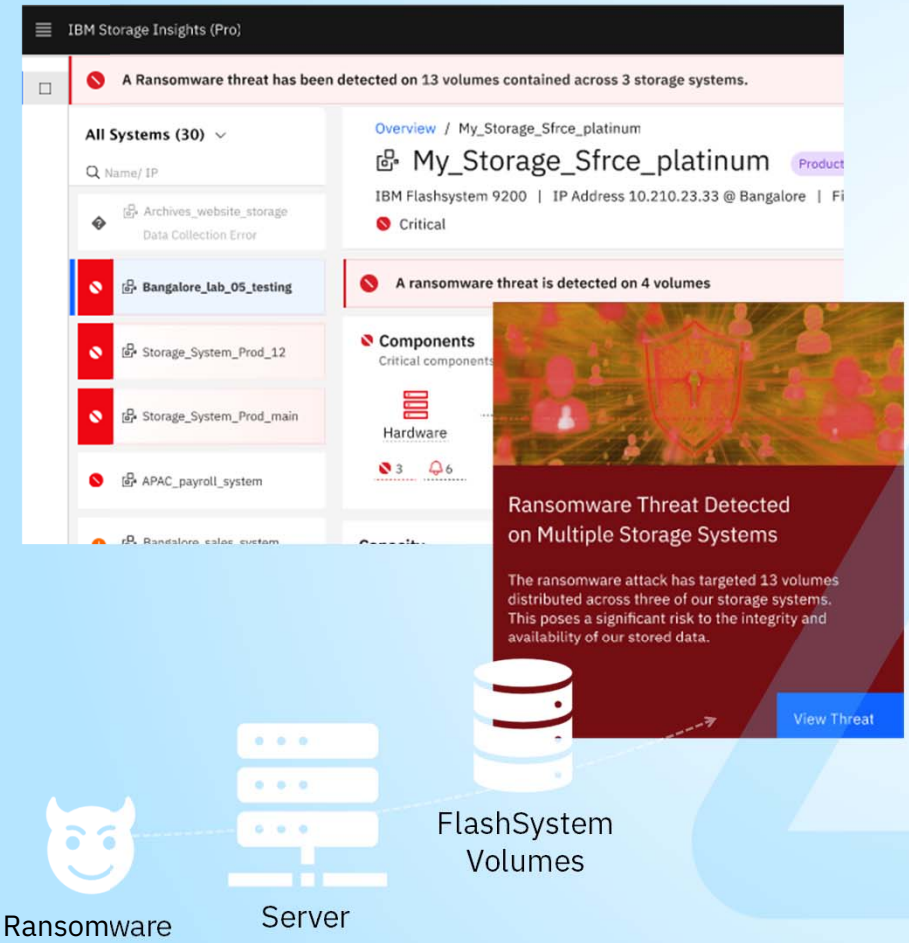

2Q24: Alert on Ransomware with FCM4

Using hardware in the FCM4 drives, each I/O is analyzed for multiple vectors of statistics, with signals passed to Storage Insights Pro

Users are warned on a per-volume basis, marking volumes and copies as compromised while the user determines whether the threat is real.

Generate an actionable alert or pass it on to other SIEM software via the open webhook API

Create a feedback loop to help with continuous detection improvements

Management: Ongoing Detection Improvements

3Q24: Using user feedback to tune detection of FCM4 based alerts

If a user determines an alert is a false positive, they can report it
Data collected related to the alert will be sent back to IBM to investigate
Changes to the algorithm can be upgraded onto Virtualize devices

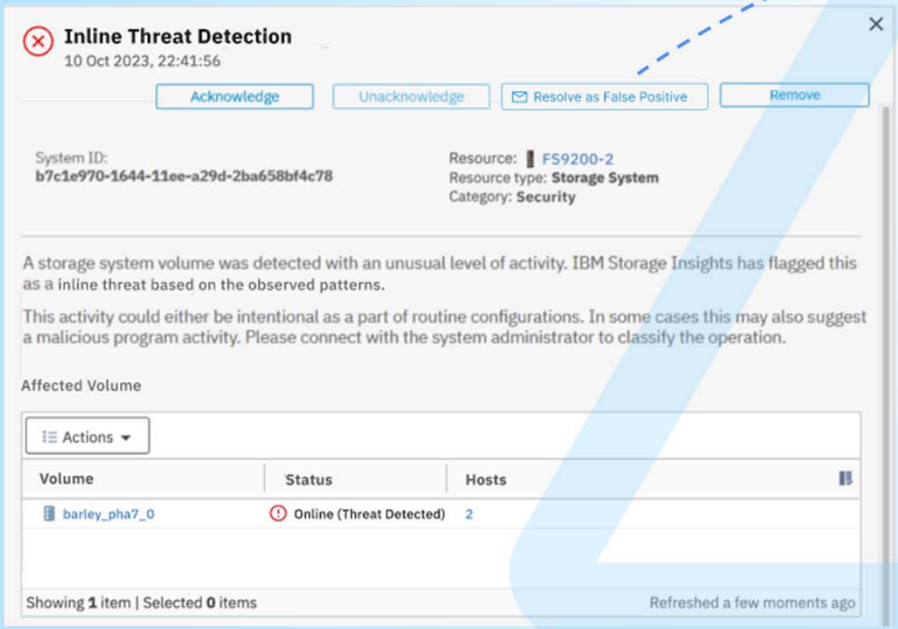

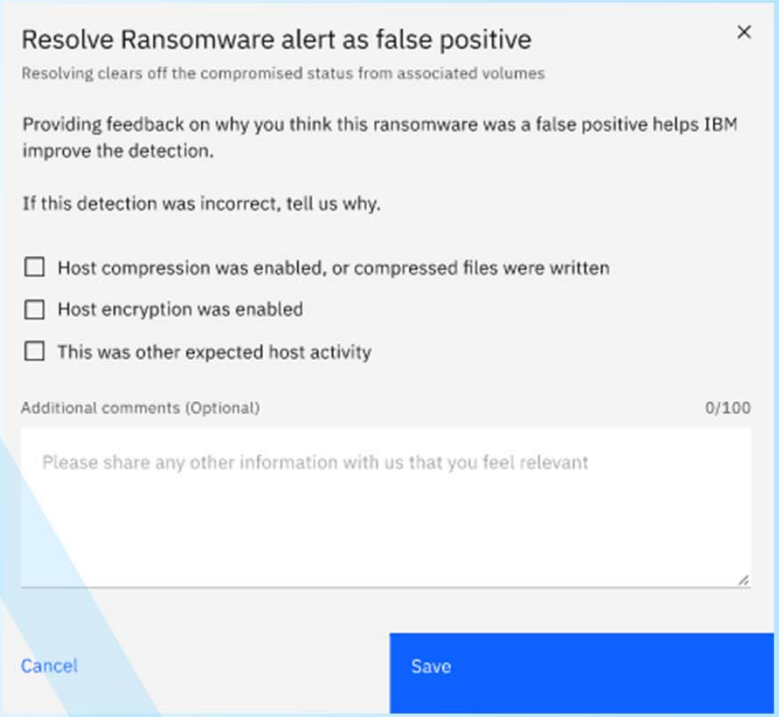

2025: Wiperware detection, AI model enhancements

Using data from the field to improve False Positives

- Original model was based on lab data and injection of ransomware and ransomware simulator
- When an alert is raised, we send back 1 hour's worth of statistics at the time of the alert
- We also sample and send back stats when no alerts – benign workload.
- We train new models with this data from real systems running in the field!
- We continue to inject ransomware to ensure the model is balanced.
- We then release patches and improved models that are more accurate and less false positives.

False Positives: The Enemy of Security Software

Field Experience

- Patch process allows upgrade to the Inference engine without affecting data path or requiring upgrade to the Storage Controller code.
- Latest model released in late November is less than 1% false positive rate.
- True positive rate remains very high.