ADVANCED TECHNOLOGY GROUP (ATG)

Accelerate with ATG Webinar: Exploring FlashCore Modules Across IBM Storage

Matt Key
Principal Storage Technical Specialist
IBM Advanced Technology Group – FlashSystem
IBM Technology
mkey@us.ibm.com

Accelerate with ATG Technical Webinar Series

Advanced Technology Group experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with ATG webinars - Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.

2025 Upcoming Webinars – Register Here!

What's New in Fusion 2.9! – April 1st, 2025

IBM Storage Ceph File System Deep Dive – April 8th, 2025

Building Affordable, Resilient Storage Solutions with IBM FlashSystem and Brocade SAN - April 24th, 2025

AI Data Assistance with IBM Content Aware Storage (CAS) - April 29th, 2025

Important Links to Bookmark:

Accelerate with ATG - Click here to access the Accelerate with ATG webinar schedule for 2025, view presentation materials, and watch past replays dating back two years. **https://ibm.biz/BdSUFN**

ATG MediaCenter Channel - This channel offers a wealth of additional videos covering a wide range of storage topics, including IBM Flash, DS8, Tape, Ceph, Fusion, Cyber Resiliency, Cloud Object Storage, and more. **https://ibm.biz/BdfEgQ**

Offerings

Client Technical Workshops

- **Cyber Resilience with IBM Storage Defender: April 9, 2025 (Virtual)**
- **IBM DS8000 G10 Advanced Functions: April 30 & May 1, 2025 (Coppell, TX)**
- **IBM Fusion & Ceph: May 14-15, 2025 (Durham, NC)**
- **IBM FlashSystem Deep Dive & Advanced Functions: May 21-22, 2025 (Chicago, IL)**
- **IBM Storage Scale & Storage Scale Functions: June 4-5, 2025 (NYC)**

TechZone Test Drive / Demo's

- IBM Storage Scale and Storage Scale System GUI
- IBM Storage Virtualize Test Drive
- IBM DS8900F Storage Management Test Drive
- Managing Copy Services on the DS8000 Using IBM Copy Services Manager Test Drive
- IBM DS8900F Safeguarded Copy (SGC) Test Drive
- IBM Cloud Object Storage Test Drive - (Appliance based)
- IBM Cloud Object Storage Test Drive - (VMware based)
- IBM Storage Protect Live Test Drive
- IBM Storage Ceph Test Drive - (VMware based)

Please reach out to your IBM Representative or Business Partner for more information.

***IMPORTANT* The ATG team serves clients and Business Partners in the Americas, concentrating on North America.**

Accelerate with ATG Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via Menti.com with code 5151 0447 or

Direct link https://www.menti.com/alhsf3bgvxu6
Or

QR Code

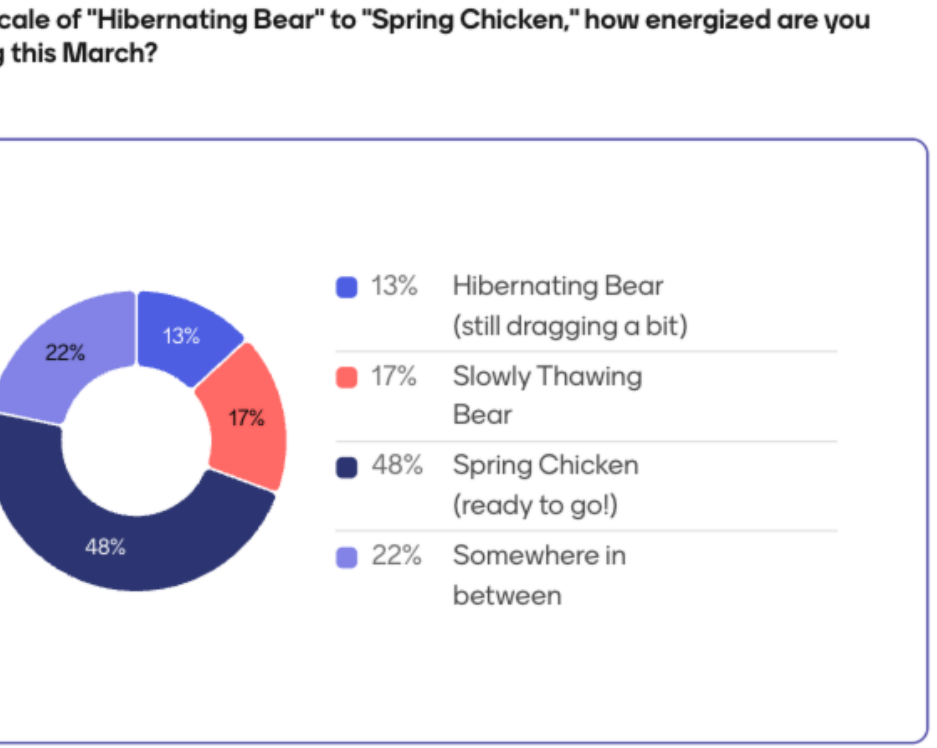

ADVANCED TECHNOLOGY GROUP (ATG)

Accelerate with ATG Webinar: Exploring FlashCore Modules Across IBM Storage

Matt Key
Principal Storage Technical Specialist
IBM Advanced Technology Group – FlashSystem
IBM Technology
mkey@us.ibm.com

Meet the Speakers – Matt Key

Matt Key is a 18-year veteran of solid state storage and came into IBM through the 2012 acquisition of Texas Memory Systems, the Houston-based engineering group behind the FlashCore Modules in IBM's distributed storage portfolio. Currently, he is the technical lead for Flash in the Advanced Technology Group (ATG), the client-facing group of subject matter experts across the portfolios of storage, servers, and software. Matt graduated Texas A&M (WHOOP!) in 2006 with an engineering degree in telecommunications.

What is a FlashCore Module (FCM)

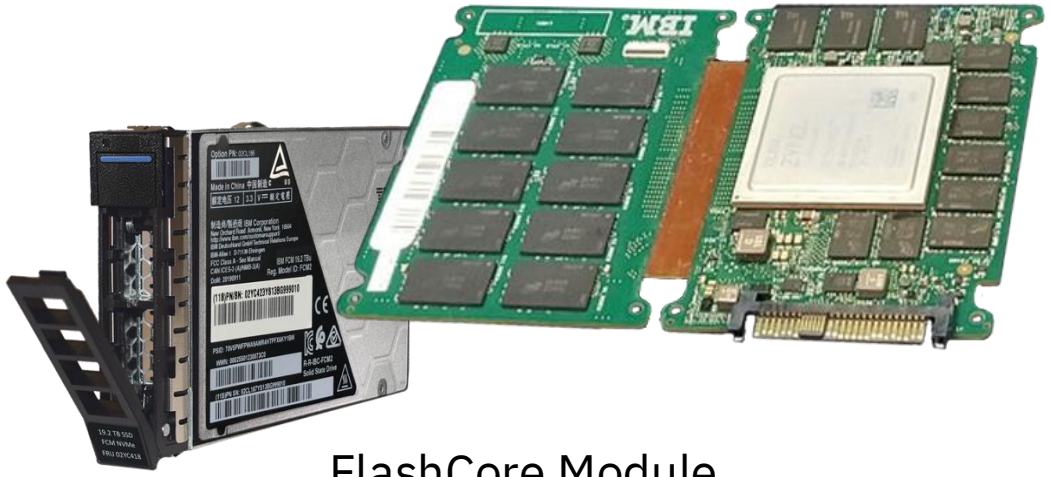

FlashCore Module

IBM FlashCore Module

- Computational Storage Devices
- Offload heavy lifting from software
- FPGA allows for ‘quick’ and subsequent enhancements
- Value-based development model
  - Performance
  - Security
  - Resiliency

Example 2.5” SSD

Traditional NVMe SSD

- Most value-add provided by storage/x86 software
- Cost-based
- ASIC controller (cost effective, but mostly static IP)
- Usually defined to a specific storage tier (Tier 0/1/2 Flash)

History of FlashCore Technology

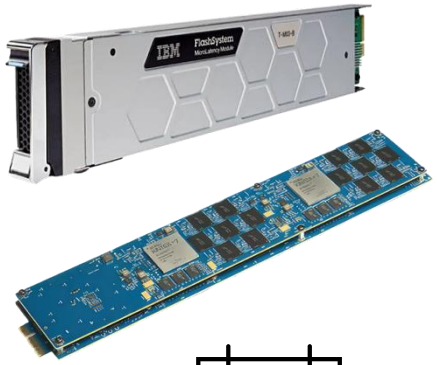

2014

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 across PCIe 4.0 interface for all FCM sizes, 176-Layer with SLC and QLC abilities, Quantum-safe encryption, and Ransomware Threat Detection

Backwards Compatible with FCM3

IBM FlashCore™

At the Heart of Quad Layer Cell (QLC) Enablement

Characterization

Read Calibration

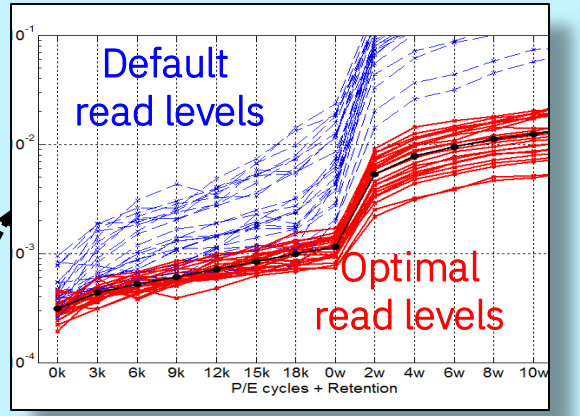

Error Correction

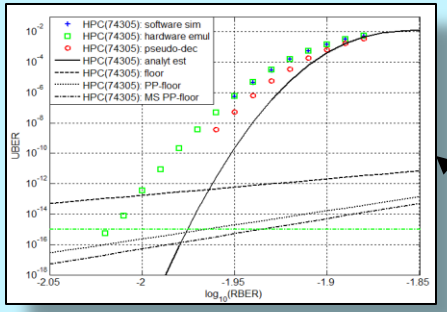

Flash Chips

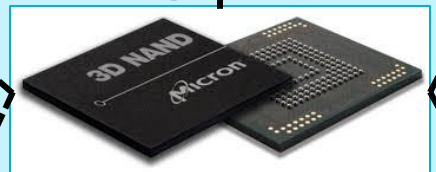

Health Binning

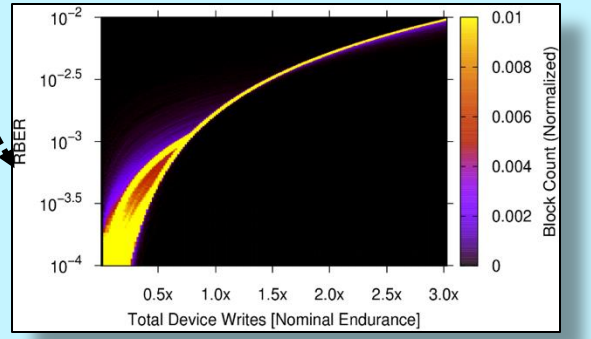

Garbage Collection

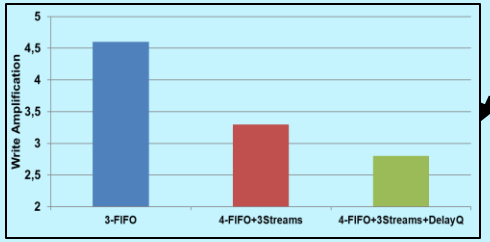

SLC/QLC Tiering

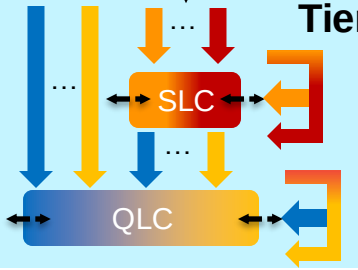

FlashCore Modules (FCM)

- FlashCore Technology drives FCM
- FCM are ‘stack competitors’ vs software-based SSD arrays

FlashCore Module 4 Specifications

FCM	Usable Capacity	Endurance (TDW)	PCIe
Small	4.8 TBu	12.2 PBu	PCIe 4.0 x4
Medium	9.6 TBu	24.5 PBu	PCIe 4.0 x4
Large**	19.2 TBu	49.0 PBu	PCIe 4.0 x4
Xlarge**	38.4 TBu	98.1 PBu	PCIe 4.0 x4

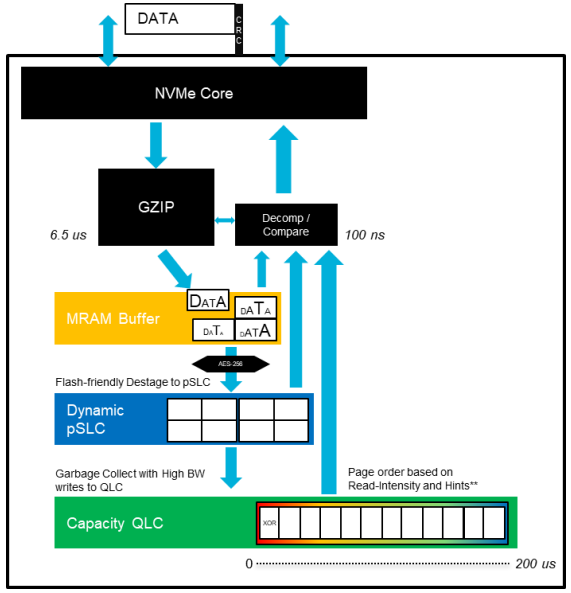

FCM Data Path

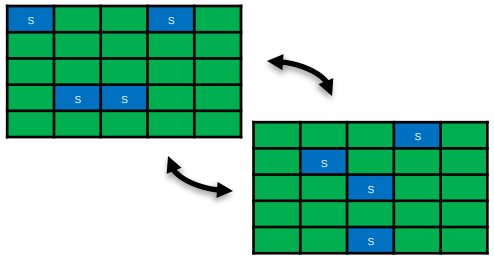

Dynamic SLC Layer

FlashCore Hints*

- XOR and Metadata Acceleration
- Turbo Power
- Proactive LBA Recovery
- Flash block Sweeper
- FCM RAID Sweeper
- Internal End-to-end CRC
- Volume I/O Tagging
- Secure Key Passing

*Not all hints are in all storage implementations
** L, XL FCMs run dual-compressors for higher performance

Where to Find FCM4 in 2025

Distributed Block

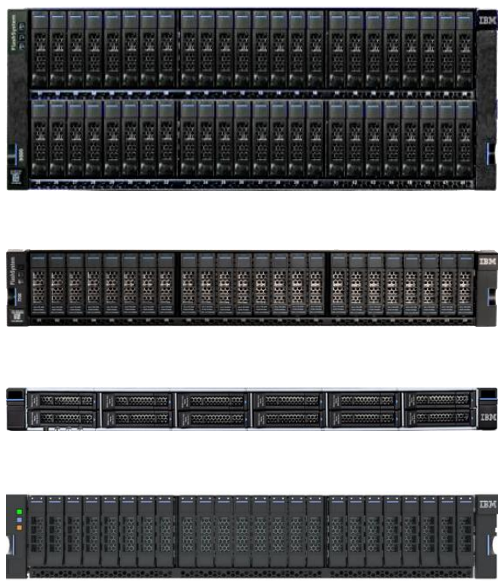

FlashSystem 9500,
7300, 5300, and C200

Enterprise Block

DS8A00 “G10”

Storage Scale

Storage Scale System 6000

IBM Storage Scale System 6000 with FCMs

NVIDIA Validated BasePOD and SuperPOD Storage Reference Architectures

FASTER, SIMPLIFIED DEPLOYMENT

NVIDIA CERTIFIED

NVIDIA-Certified Systems

EDGE | ON-PREMISES | COLOCATION | CLOUD

- **Accelerate AI Results** — Scalable Parallel Performance (Time to Value)
- **Access the Data for AI** — Single source of data – Global Data Platform (Agility)
- **Abstraction to Optimize AI Economics** — data lineage preservation (Cost Efficiency)
- **Assurance to Govern, Secure and Protect** — NIST framework compliant (Risk Reduction)

- a Global Data Platform -

6000 Flash, FCM4, Hybrid w/ HDD

Ultimate Performance

310+ GB/s read performance per node

155+ GB/s write performance per node

Ultimate Efficiency with FCMs

1.23 PiB of Usable Capacity per node

19.2TB and 38.4TB FCMs

Ultimate Economics

Global Data Platform allows singular namespace with SSDs, FCMs, HDD & Tape. In array, across arrays, and cloud

FCMs in the IBM DS8A00 G10 Family

- NVMe and Compression on THE enterprise storage system
- Now with HPFE Gen 3 (though connected with PCIe Gen 4)
  - Max 4 enclosures on DS8A10
  - Max 8 enclosures on DS8A50, DS8A80
- Supporting both Fixed Block and Count-Key Data (CKD) Formats
- Focused on SafeGuarded Copy
  - Still supports primary workloads of z/OS, AIX, IBM i, LinuxOne
  - FCMs are not targeted for TPF installations
  - Extended virtual address space for FCMs allows for MUCH high copy counts and capacity management for SGC
  - Large Extents for 32PB of capacity vs small extents with ~8PB provisioned
- 19.2TB FCMs GA'd in 2024
  - 4.8TB and 9.6TB available through RPQ

FCM4 Hints In DS8A00 – Example: Proactive Sector Repair

- NVMe is a two-way protocol
  - SCSI is more ‘reactive’ to the hosts
  - Allows for smarter devices that can request to the controller

- Corrupt LBA Recovery (CLR) in FCM4
  - Internal Sweeper to each FCM tracks uncorrectable sectors
  - Proactively sends list to RAID driver to repair sector
  - Prevents host from encountering a sector error upon I/O

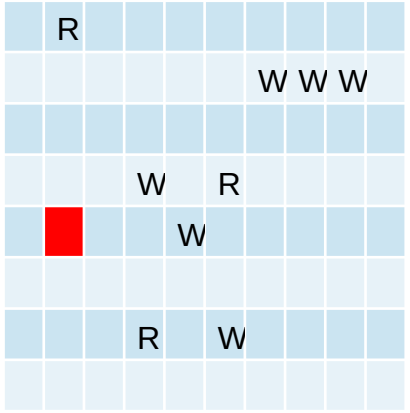

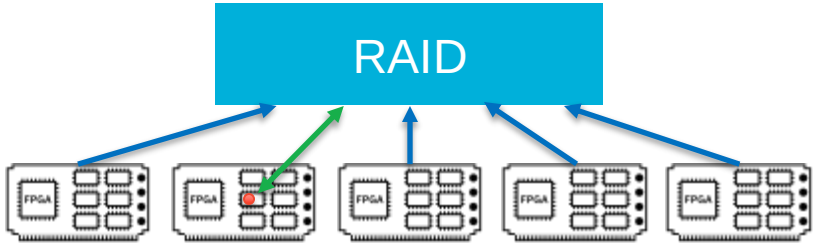

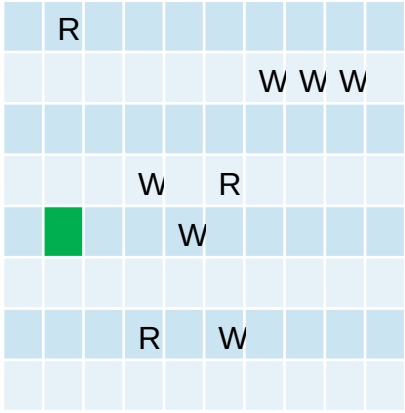

FCM4 with Always-On Hardware Compression

- IBM FlashCore Modules (FCMs)
  - Thin Provisioned and Over Provisioned NVMe Flash Drives with Hardware Compression/Decompression
    - Using LZ77 + Dynamic Huffman Encoding on 16K Blocks
    - Capacities Supported
      - 4.8 uTB Physical (4.58x over provisioning or 21.9TB Logical)
      - 9.6 uTB Physical (3x over provisioning or 28.8TB Logical)
      - 19.2 uTB Physical (3x over provisioning or 57.7TB Logical)
    - Average Compression Ratio achieved on FlashSystems is 2:1 with distributed data
      - Compression ratios will vary depending on client data. Using host compression and/or host encryption will make the data highly uncompressible
      - Many zOS clients already use zEDC/Application compression and/or encryption

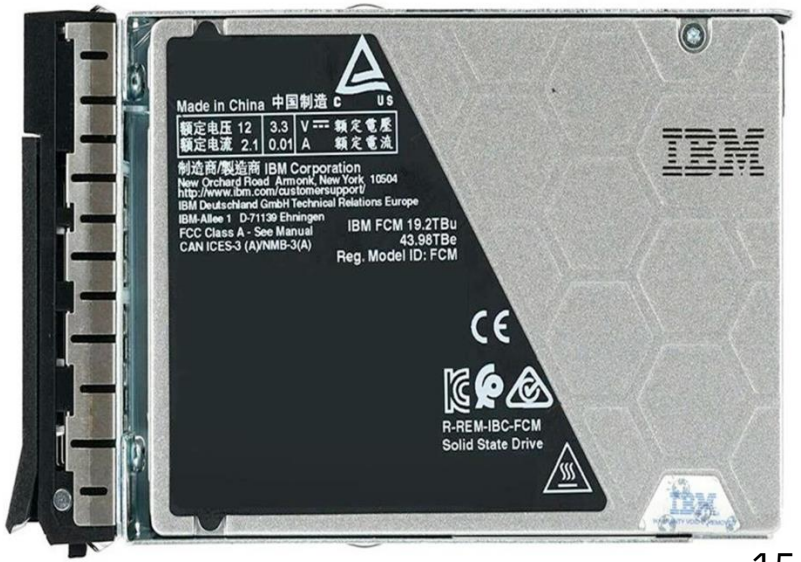

DS8000 G10 – z/OS DFSMSdss Compression Tool

APAR – OA66417

- Provide a z/OS utility that will provide zEDC compression savings for tracks on a volume
  - Initially:
    - FULL Volume image
    - 16K block increments
  - **Future:**
    - Data set Level
    - Different block sizes

- Compression Tool is invoked at Volume level
- Clients can choose a sample size of their storage group, run tool against that subset of volumes
- CSV output available
  - example requests CSV output
  - Output of CSV not shown.

- Available for zOS 2.5 and 3.1 end of September 2024, eFix possible
- https://www.ibm.com/support/pages/node/7159368

DS8000 G10 – Distributed Systems Compression Tool

Comprestimator is a host based utility for a fast estimation of a block device compression ratio

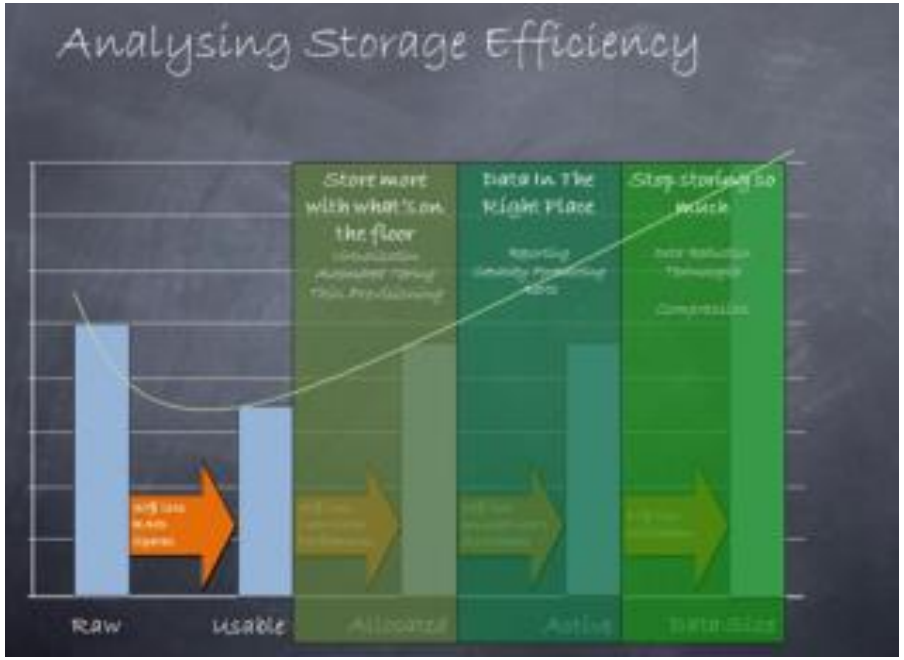

Objectives:

- Run over a block device
- Estimate:
  - Portion of non-zero blocks in the volume
  - Compression rate of non-zero blocks

Performance:

- Runs FAST! < 60 seconds, no matter what the volume size is
- Provides a guarantee on the estimation: ~5 % max error
  - Can improve guarantee with more samples (longer running time)

Method:

- Random sampling and compression throughout the volume
- Collect enough non-zero samples to gain desired confidence
  - More zero blocks → slower (takes more time to find non-zero blocks)
- Mathematical analysis gives confidence guarantees

https://www.ibm.com/support/pages/ibm-flashsystem-comprestimator

FCM Thin and Over Provisioned Flash Drives with Compression

Logical Capacity (57.6 TB)

Compression Engine

Physical Capacity (19.2 TB)

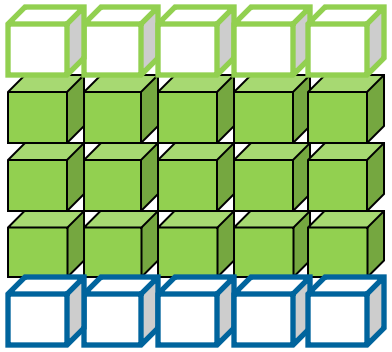

- Data of All-0's or Pre-allocated capacity consume no Usable Capacity
- Physical Capacity is the primary value to monitor
- Alerts for all capacity levels Logical Capacity is required for expanding ESE

At a 2:1 compression, 30 TB of data consumes 15 TB of Physical Capacity or 78% Used Physical

Capacity beyond Usable Capacity
- Reserved | Recovery capacity
- ECC
- Bad-block management
- Performance Overprovisioning

Non compressible scenarios and Safeguarded Copies

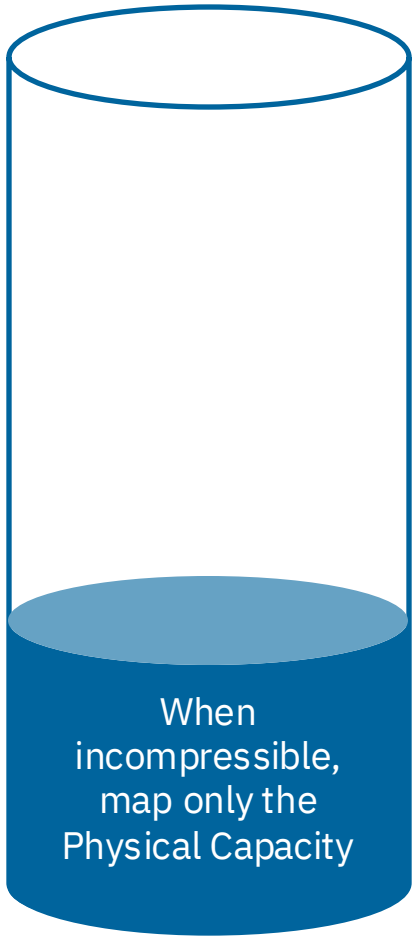

To avoid overprovisioning, map only the Physical Capacity (85%) amount worth of volumes

- ESE will still be active
- Compression will still be active
- Volumes cannot grow beyond the Physical Capacity of the pool

When leveraging SGC as the primary capacity contributor

- Embrace large extents - will consume more Logical Capacity, but allows for more copies
- Allow for variable physical capacity
- Allow for recovery space (logical & physical)

FCMs in Distributed Storage - FlashSystem

1Q2025

NEW FlashSystem C200

Capacity optimized, fixed config, for lower performance demanding

FlashSystem 5300

Dense 1U entry level NVMe

FlashSystem 7300

Balanced performance and value

FlashSystem 9500

Extreme performance and scaling for mission critical workloads

with 4th gen FlashCore Module (FCM4) computational storage device

Hardware compression (GZIP) with no performance impact

Hardware encryption (FIPS 140-3) with no performance impact

Real-time Ransomware Threat Detection

Leverages all hints for security, endurance, and performance

FlashCore Hints

- XOR and Metadata Acceleration
- Turbo Power
- Proactive LBA Recovery
- Flash block Sweeper
- FCM RAID Sweeper
- Internal End-to-end CRC
- Volume I/O Tagging
- Secure Key Passing

Post-Quantum Cryptography

- Data at Rest is continued as AES-256 Encryption
- IBM Storage Virtualize protects the security PINs of each drive through an IBM-specific protocol of Secure Key Passing (SKP)
- FCM4 Embrace CRYSTALS-Kyber algorithm and RSA for SKP
- Transparent to admin & intermixes with FCM3s cleanly

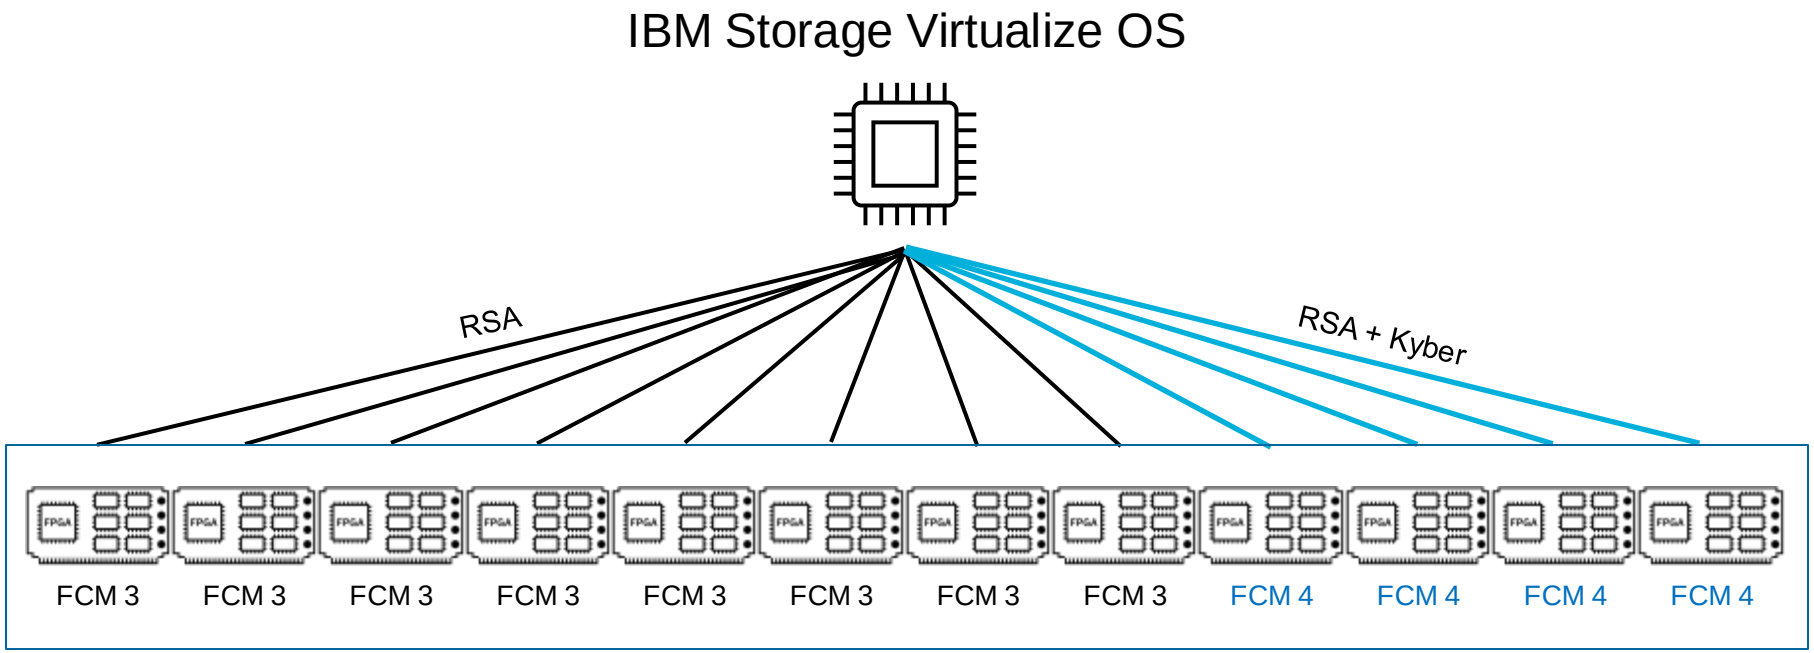

FlashCore Modules 4 are Volume Aware

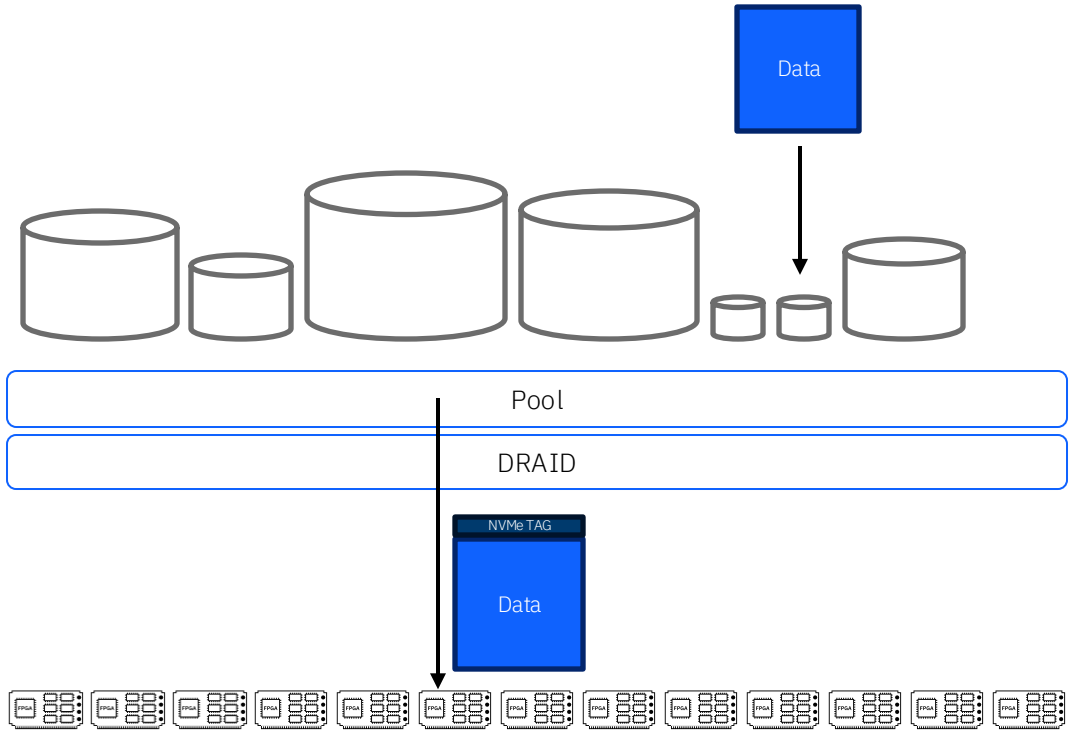

FCM-4

FLASHCORE MODULE 4

RANSOMWARE

THREAT DETECTION

Inline Data Corruption Detection (Introduced June 2023 w/ Virtualize 8.6.0)

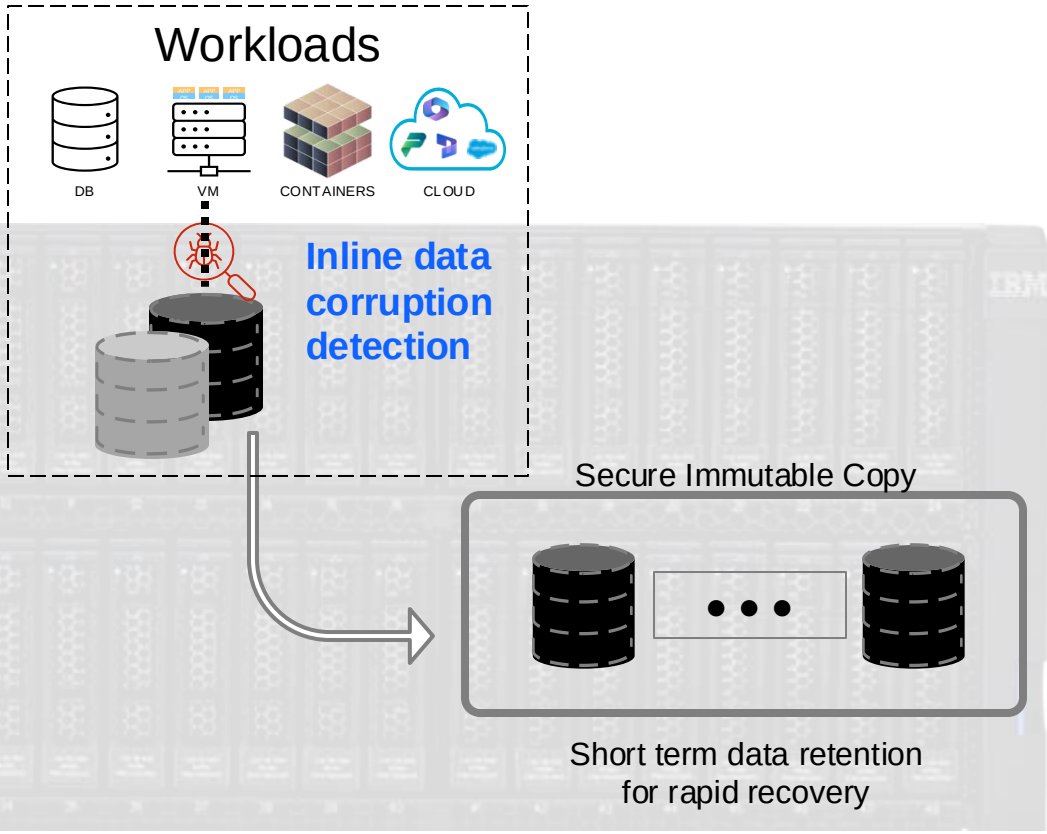

Early detection is essential to containing corruption

- Actively scanning IO before written to flash can alert of intrusion much earlier, potentially limiting impact to a single application or volume.
- Dramatically reduce the impact of a landed cyber attack, saving time to recovery.

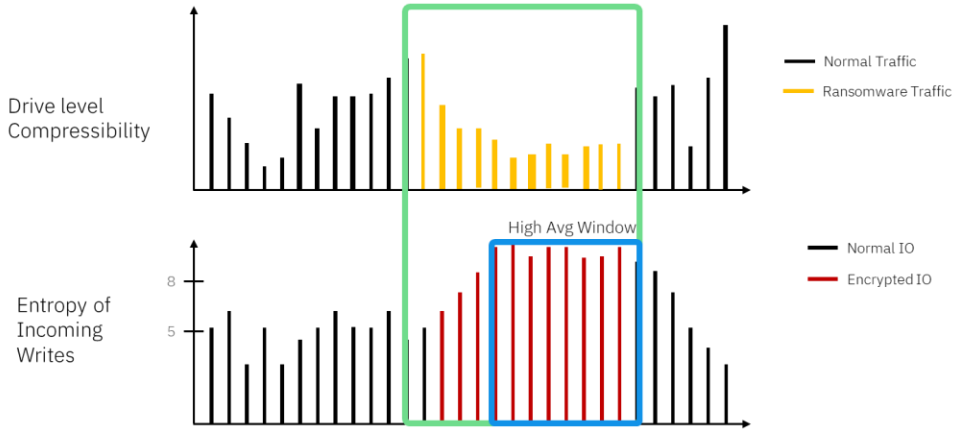

Cyber Attacks: Similar IO Access Sequences

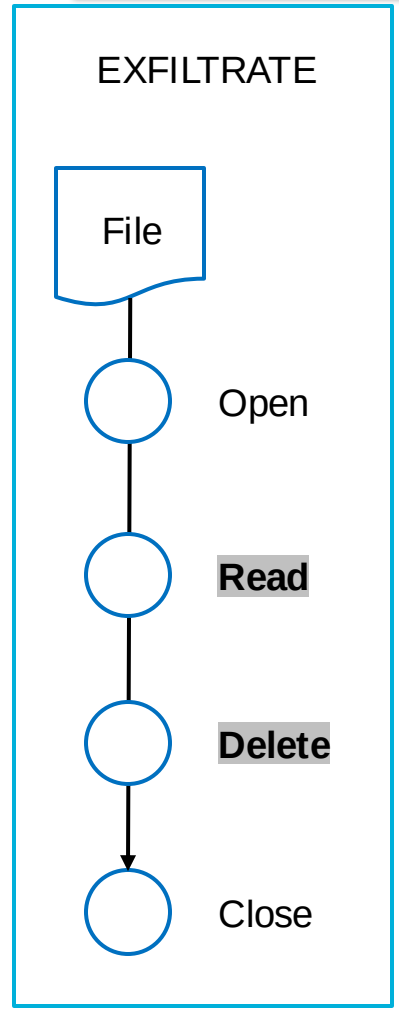

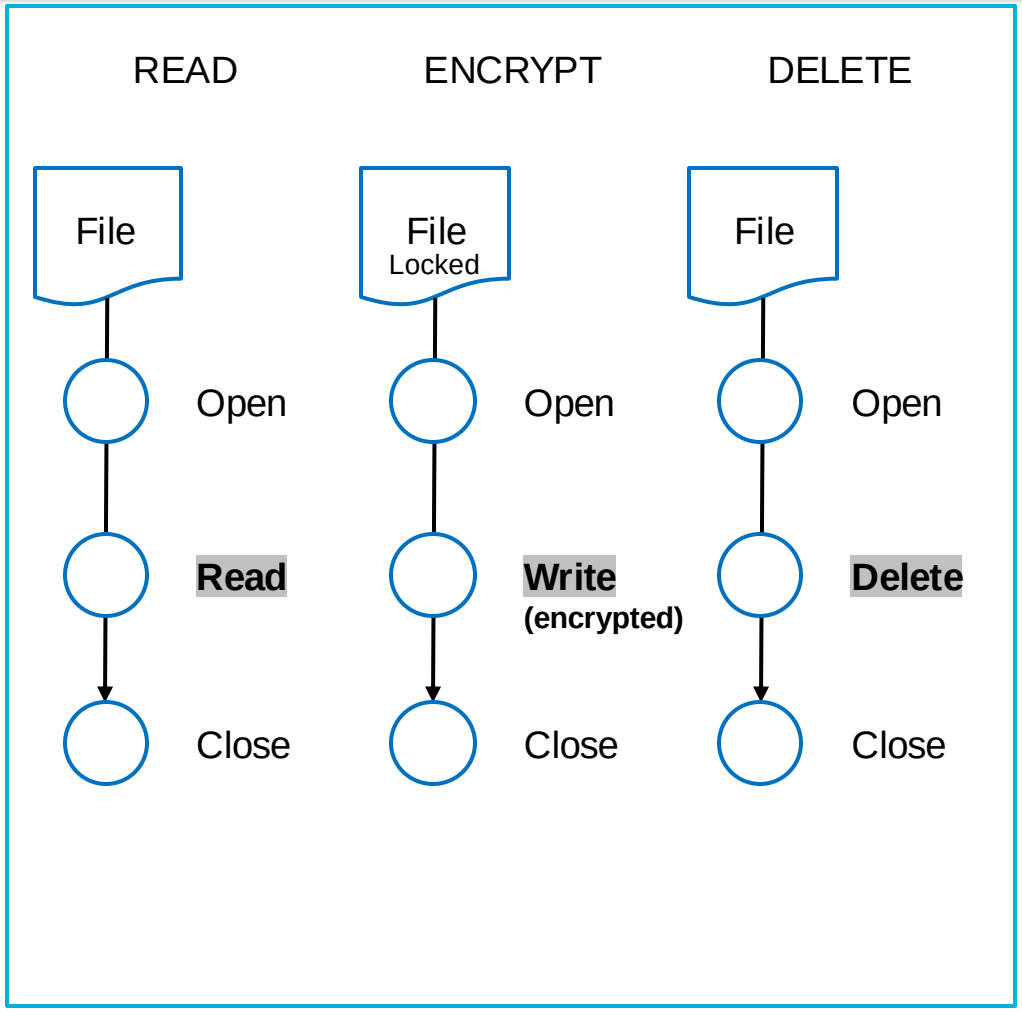

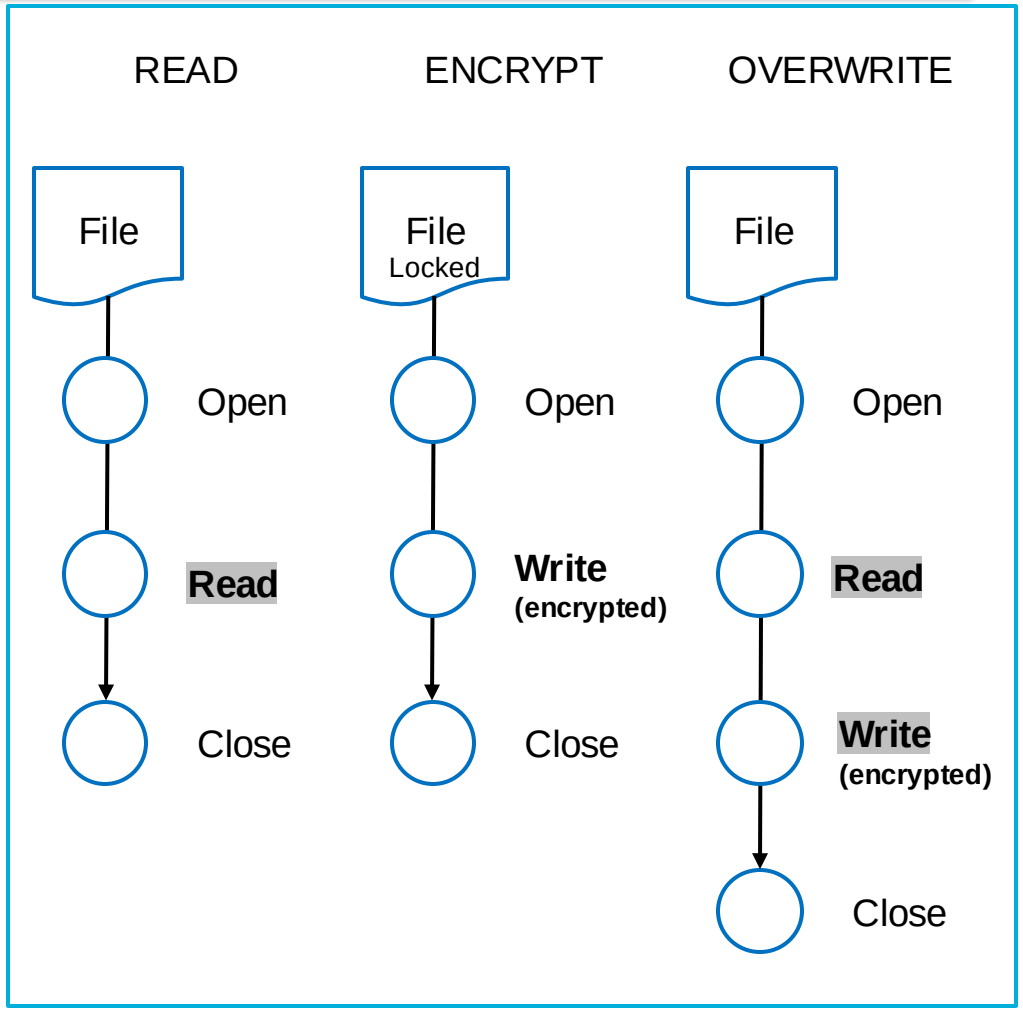

Dimensions of Analysis (What, Where, and How)

1) Cyber Attacks: Similar IO Access Sequences

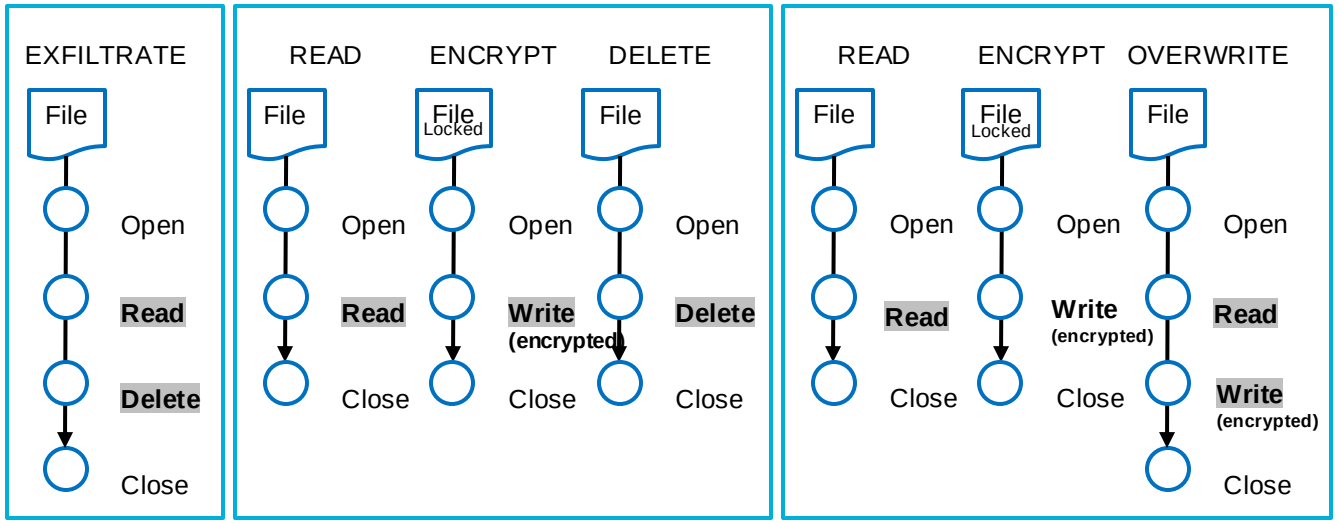

2) Logical Block Addressing (Spatial Locality)

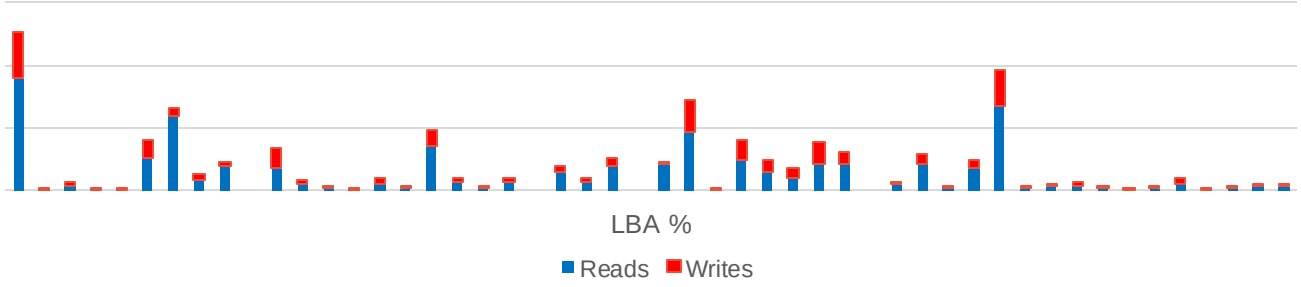

3) Data Characteristics

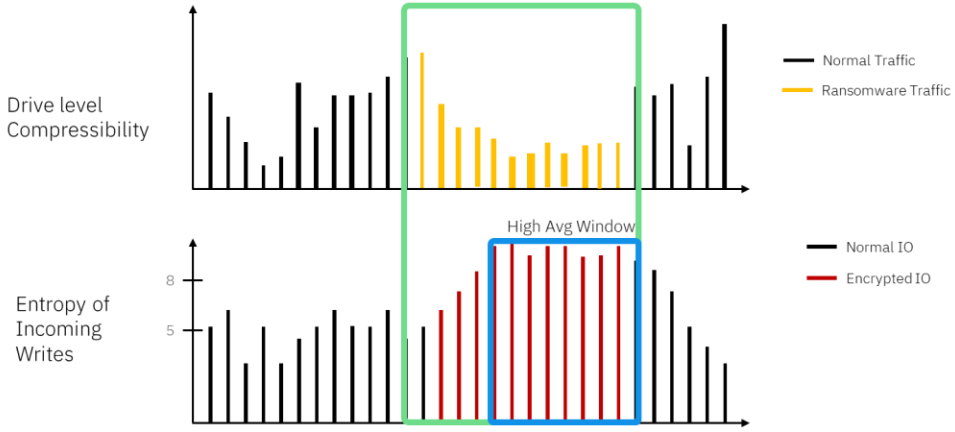

Ransomware Detection With FCMs

40+ data statistics analyzed in detection engine

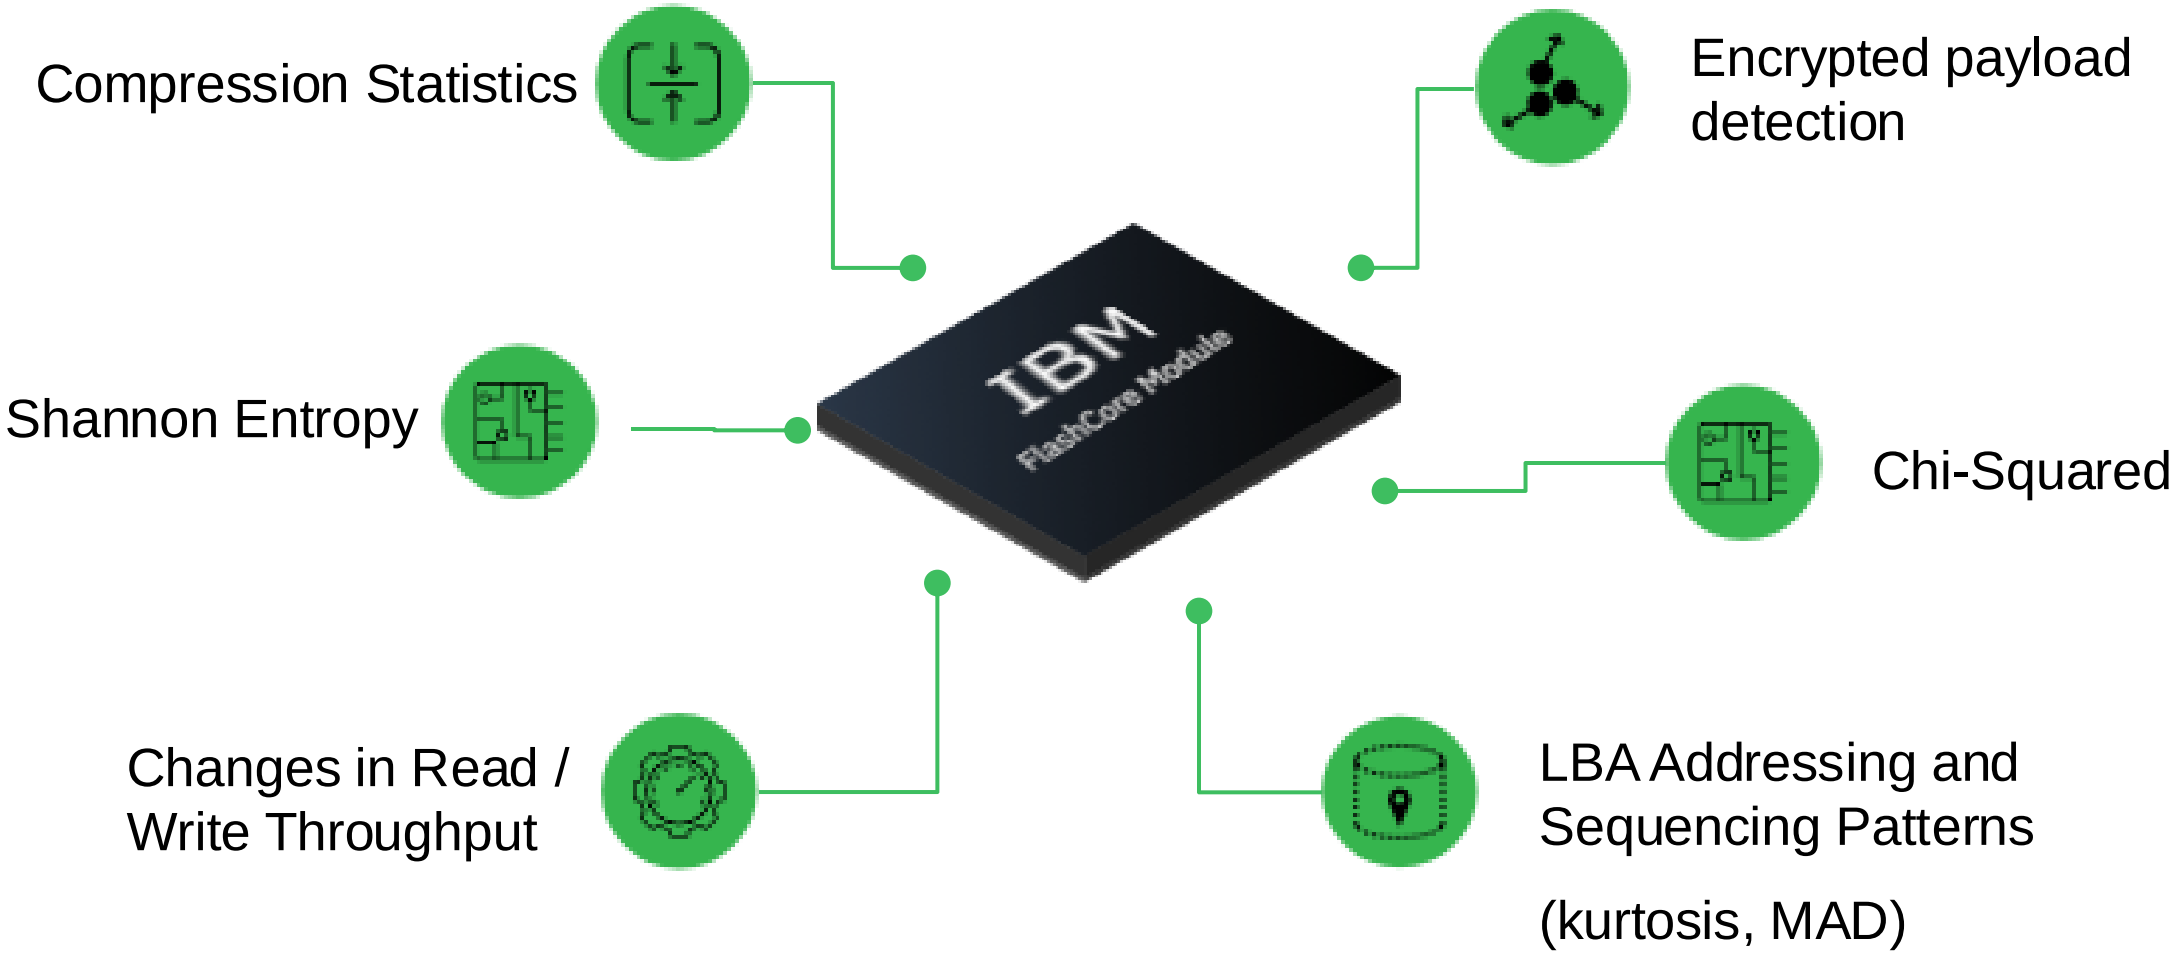

Prerequisites for Ransomware Threat Detection

- Pool must be ONLY FCM4s, updated to FW v4.1
- Pool must be created with v8.6.2+
- Only a single FCM DRAID in the pool (pre-existing req)
- DRAID6
- Storage Insights Pro alerting
- 128GB+ RAM per node
- Standard Pools only (and Fully Allocated in DRP is OK)

Ransomware Monitoring Architectural Overview

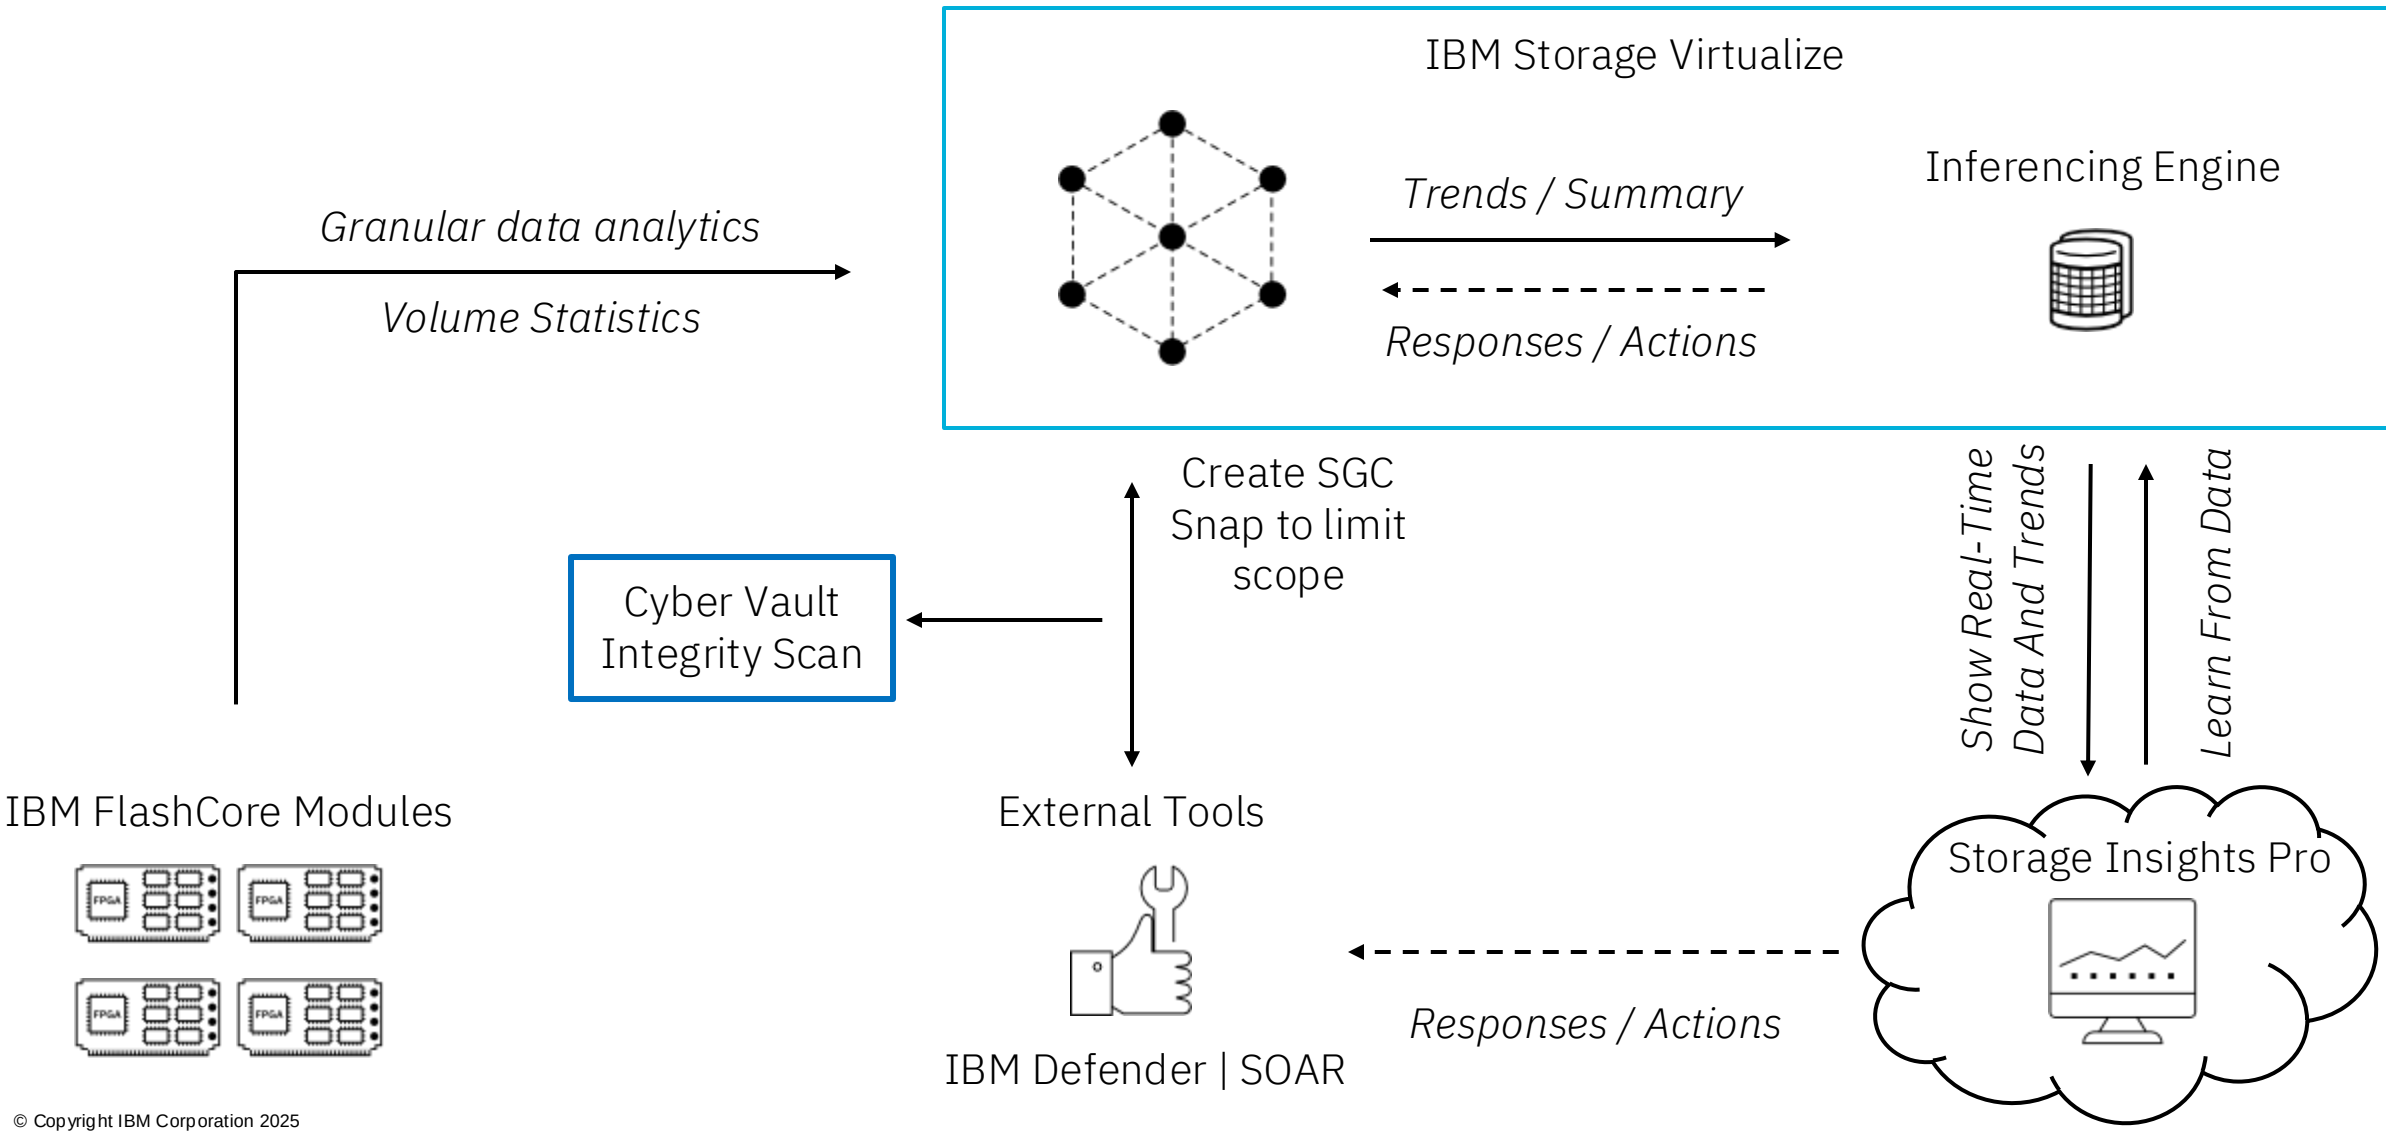

Ransomware Threat Detection in Action

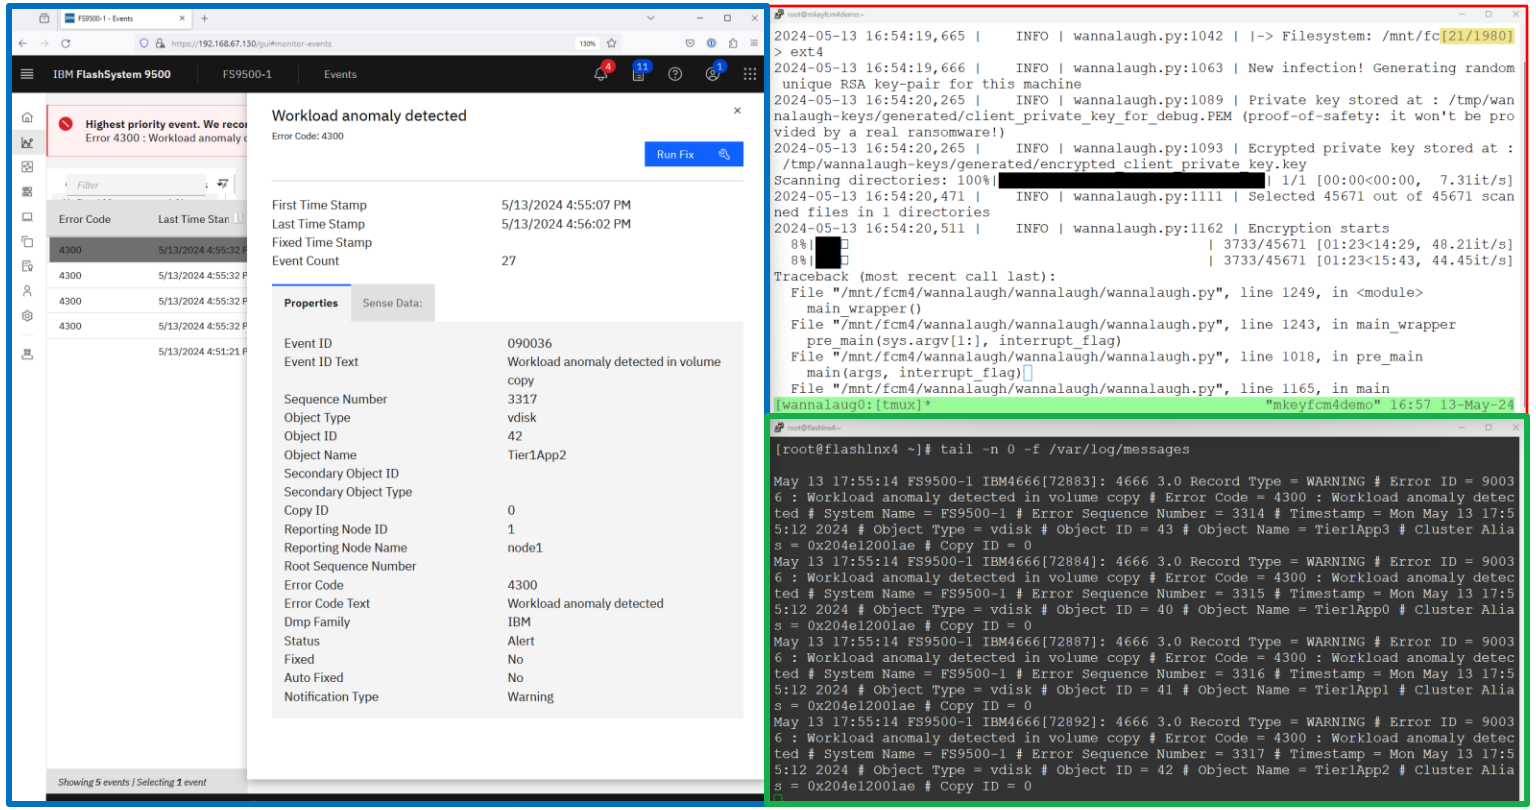

16:54:20 EST – Attack Starts

16:55:14 EST – Alert Received

54 seconds!

Ransomware Threat Detection

Brought to you by

IBM Research

IBM Storage Virtualize

IBM FlashCore

IBM Storage Insights Pro

Data Validation – Extending Ransomware Threat Detection with Cyber Vault

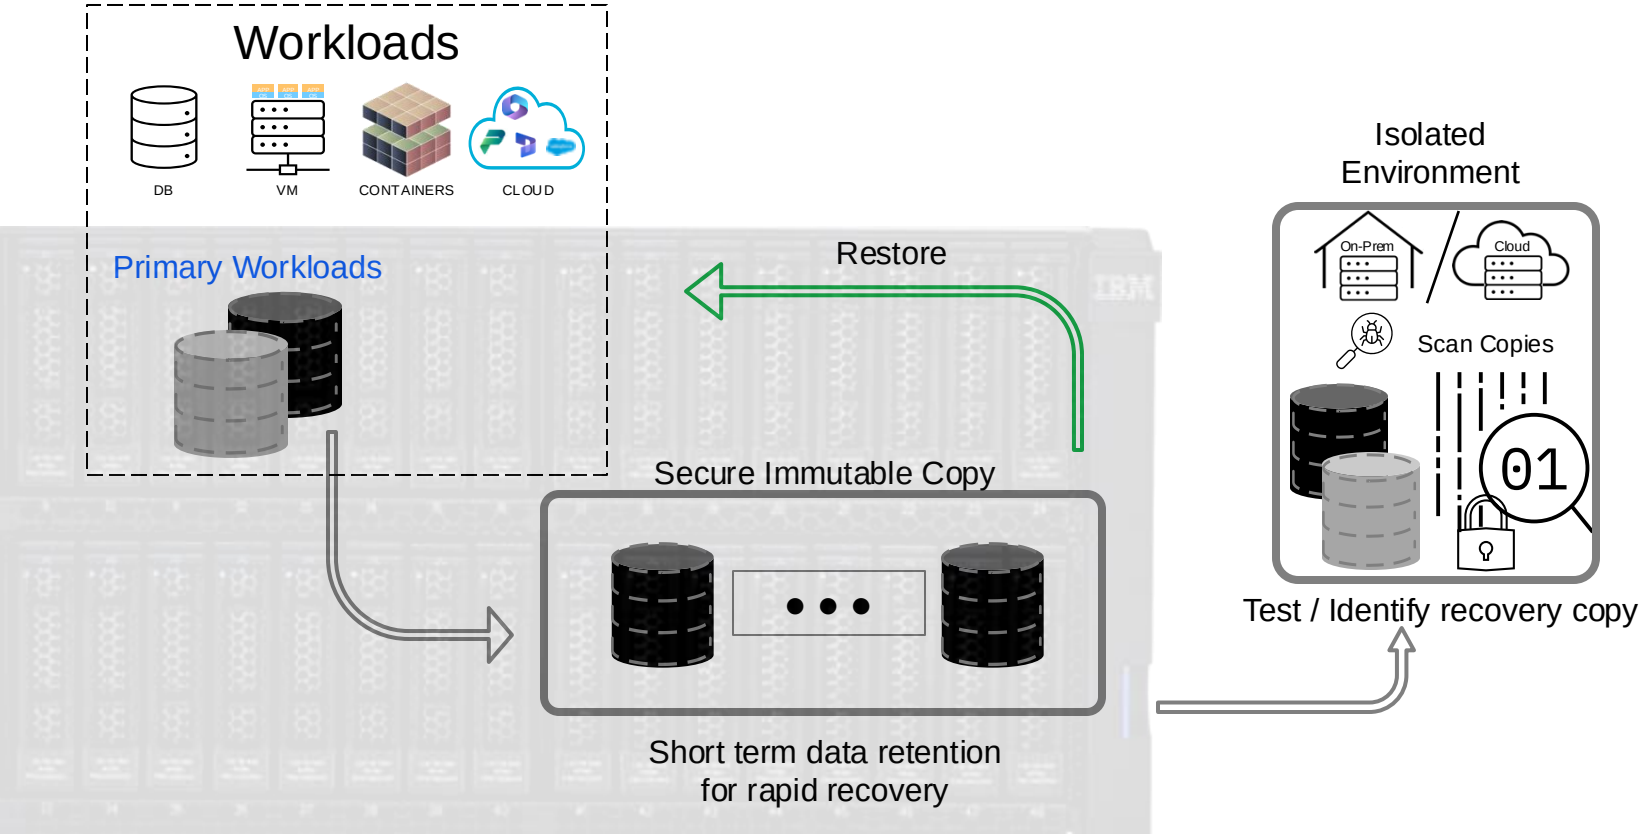

- Application Aware-ness and App Consistency
- Deep scan data with anomaly detection software
- Automate and catalog with Copy Data Manager
- Identify known-good copy
- Testing builds confidence and ensures compliance

Accelerate with ATG Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via Menti.com with code 5151 0447 or

Direct link https://www.menti.com/alhsf3bgvxu6
Or

QR Code

Thank you!