# IBM FlashSystem Announcement

## Designed to Protect, Engineered to Adapt, Built to Perform

10th February 2026

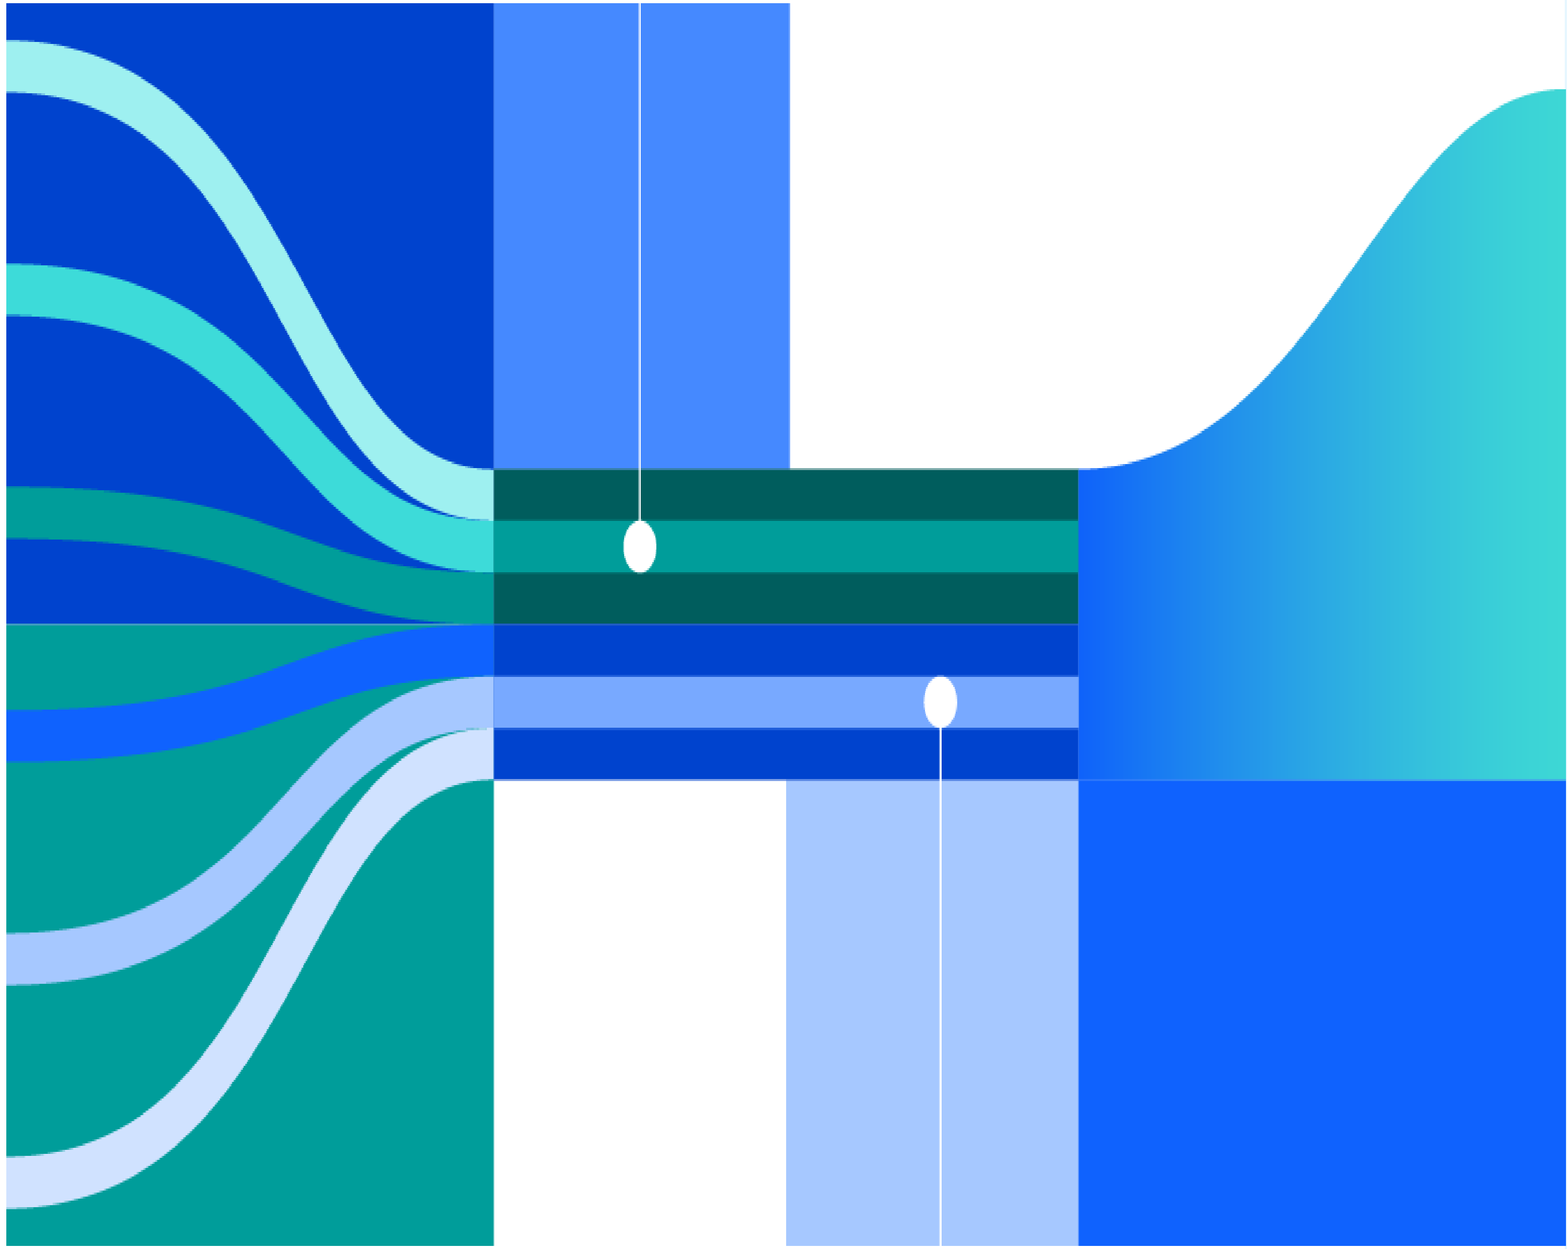

2025
IBM
IBM FlashSystem

# IBM FlashSystem

Virtualize development has made major improvements in helping to solve some of todays critical business challenges.

Trouble free performance with data reduction through IBM's unique FPGA based NVMe FlashCore technology - now including FCM4 Ransomware Threat Detection

Policy-based Management providing operational efficiency and integrated AIops business automation with Partition Mobility and Flash Grid

Security and Cyber Resiliency with box-level security hardening, Safeguarded Copy, Cyber Vault and near-real-time threat detection

# Why Clients Choose IBM FlashSystem

**Built in Cyber Resilience**
Protects data with immutable backup copies

**Trusted Innovation**
Backed by decades of enterprise storage leadership

**Operation simplicity**
Easy to deploy, manage and scale

**Seamless cloud integration for flexibility**
Integrates seamlessly across hybrid cloud environments

Simplicity

Cyber Resilience

Cloud-Ready

Innovation

# IBM FlashSystem Grid. Scalable, non-disruptive migration

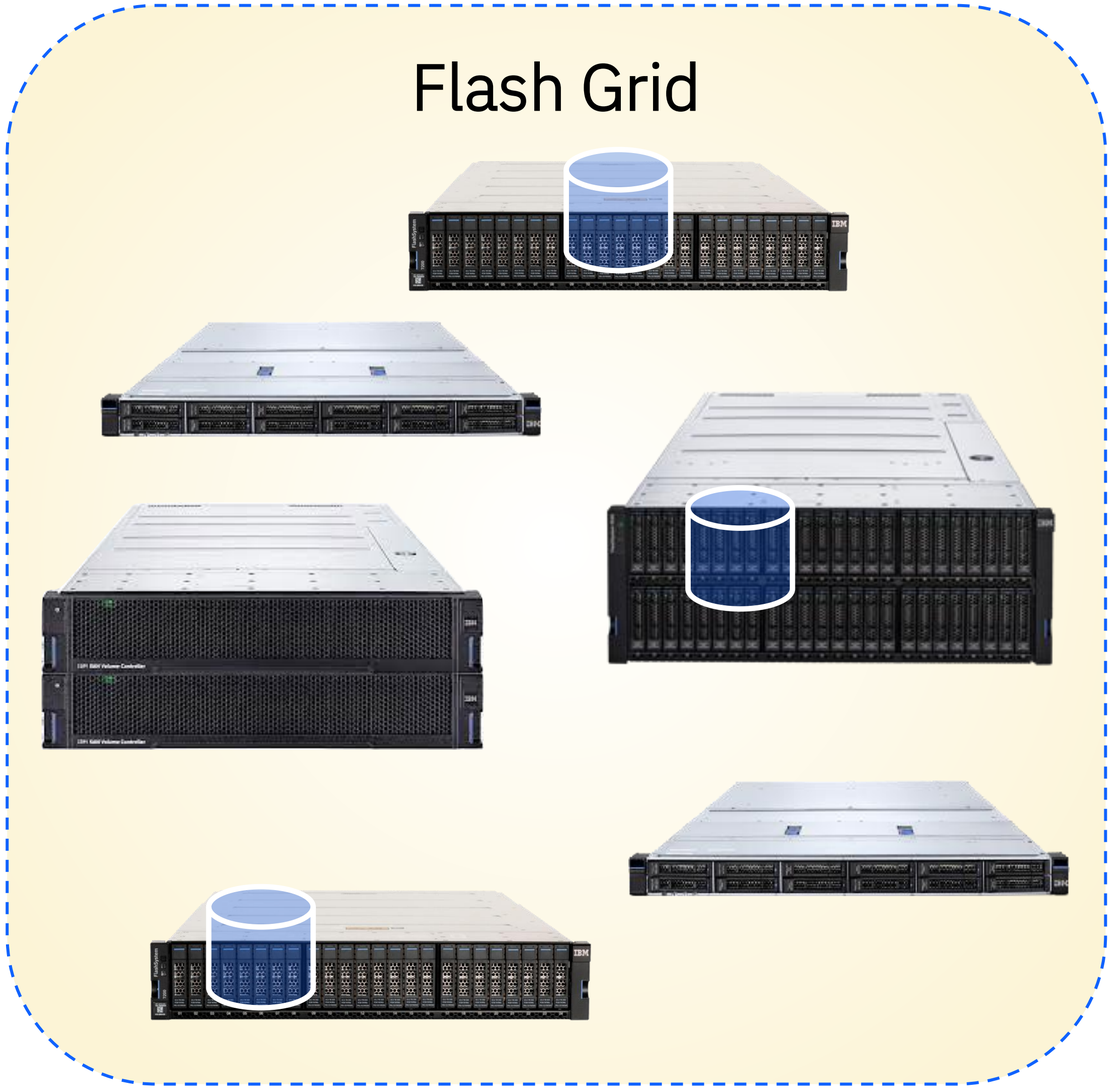


## A new approach to clustering

Single point of management

Flexible, scalable federated cluster

Independent devices and failure domains

Scale to 10s (or 100s) of devices

## Flexible asset utilisation

Provide applications with HA and/or DR resilience

Move “applications” (storage partitions) between devices

Manual or automated *non-disruptive* movement

Enables device migration and consolidation

Rebalance for capacity, performance or business needs

# **New data protection rules.** Effortless, resilient, performant
# Provisioning, replication & snapshot policies

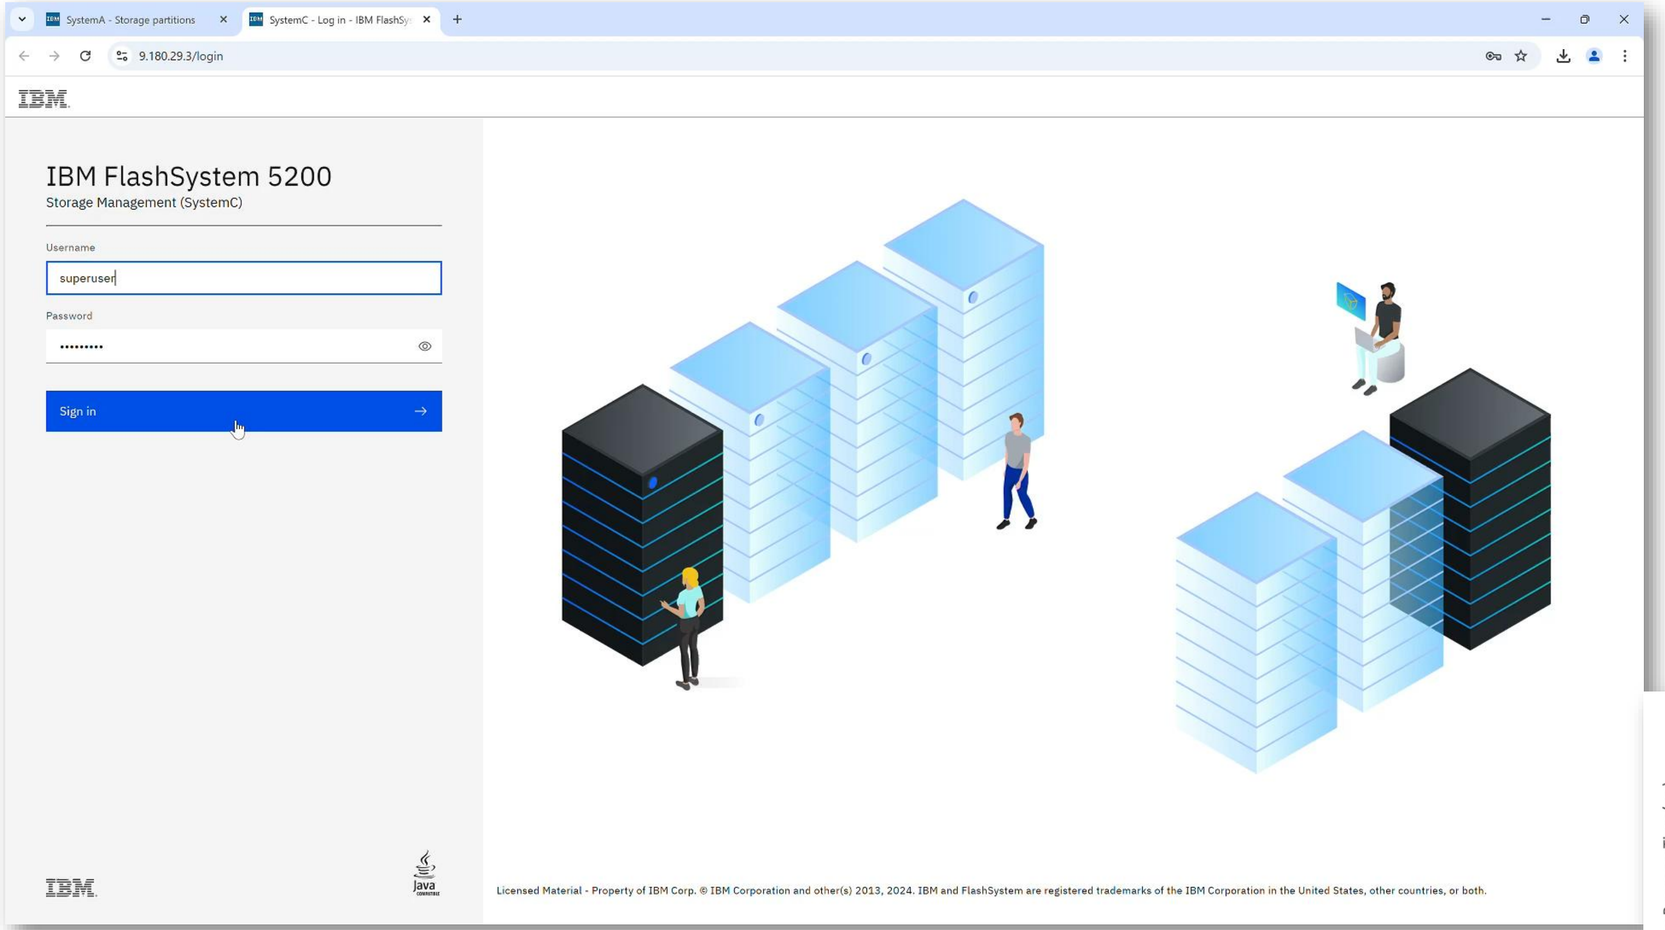


PBHA + PBR =
Disaster Recovery (3 site)

In under 5 minutes, configure a replication policy between high availability systems to establish a disaster recovery site.

**Delivering unparalleled performance with FlashSystem**

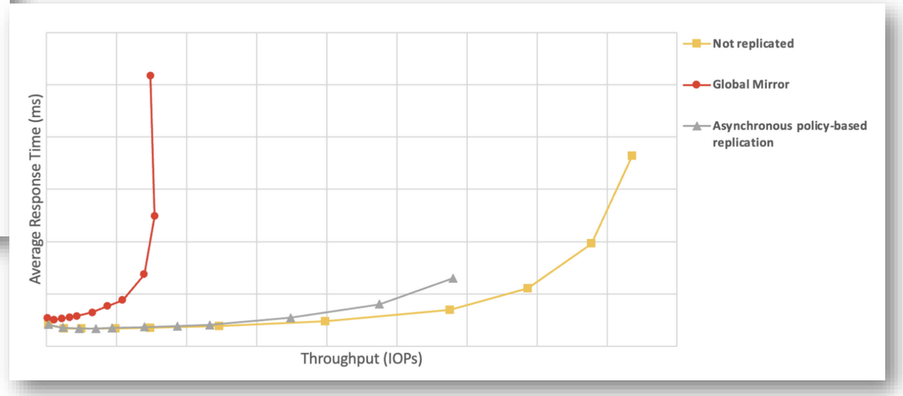

# Ransomware detection

Lower Cost Compared to DRAM, SCM and Competitive Flash Media

**100% Hardware-based Ransomware detection inline**

Based on the latest SLC and QLC 3D Memory Process Technology

Computational Storage (Operational Offloading)

16KB Page Size (More Efficient Reads/Writes)

Highest Cell Reliability (Integrated Error Correction)

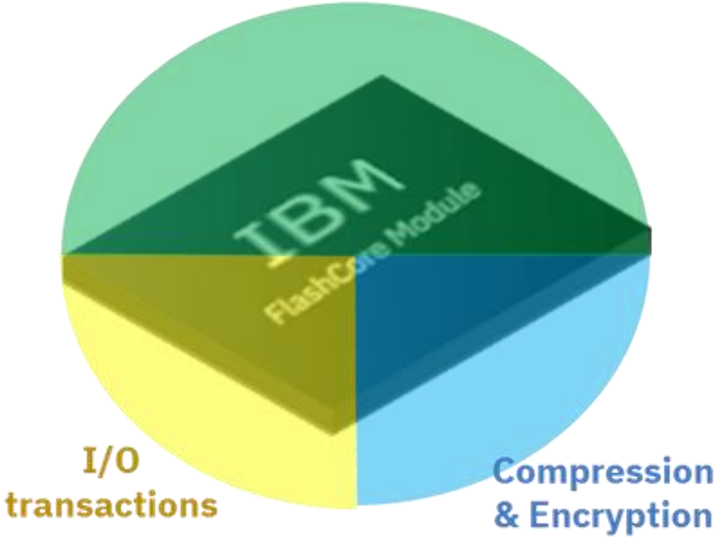


Changes in Read / Write Throughput Patterns

Compression Entropy Measurements

Information Entropy (Shannon Entropy)

Chi Squared Local Differencing Measurements

LBA Addressing and Sequencing Patterns

# Get notified within 60 seconds of a ransomware attack

so you can act fast, protect your data, and keep your business moving.

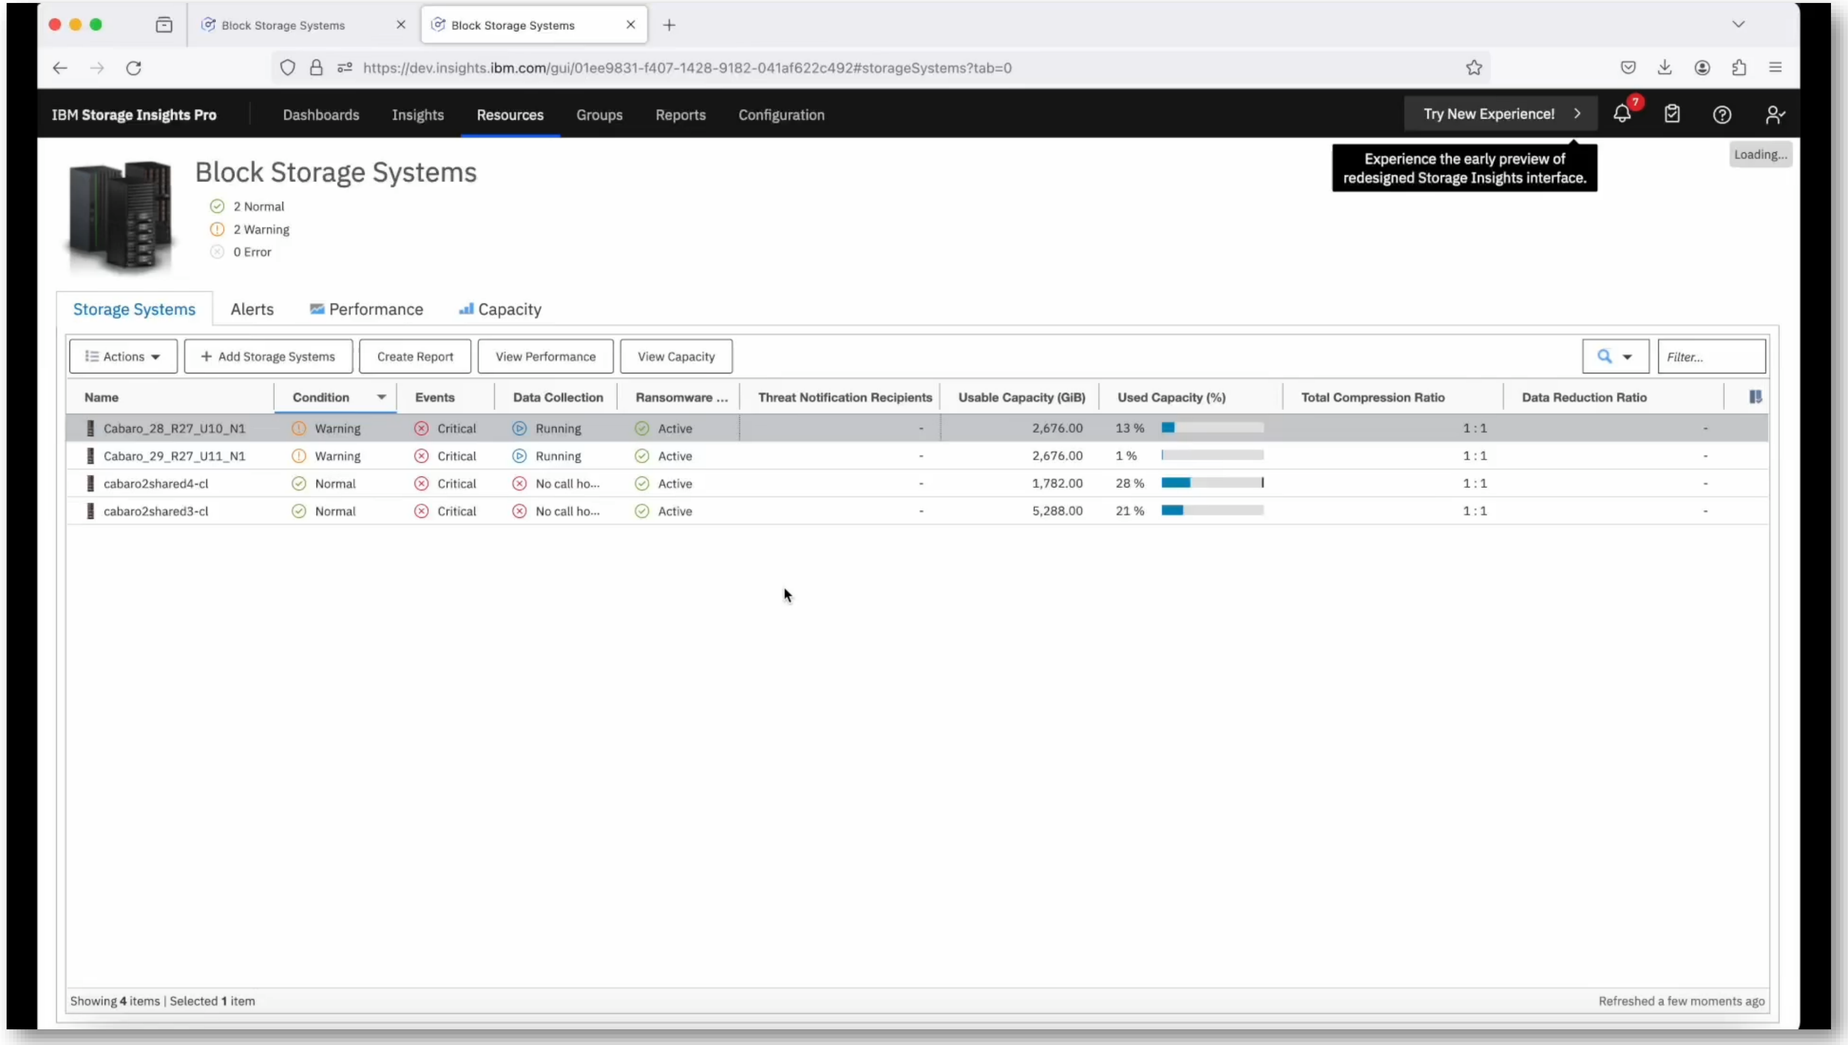


You are seeing the result of the perfect integration between:

- Drives (FCMs)
- Controller (FlashSystem)
- Cloud (Storage Insights)

1. You will receive and alert on Storage Insights Pro

2. This will guide you specifically to the affected volume(s)

3. You can take actions like:

- Validate
- Scan your immutable snapshots
- Start recovery from the selected snapshot in seconds

IBM and Business Partner Internal Use Only

CYBER RESILIENCE
2026
AI-DRIVEN STORAGE
IBM
EFFICIENCY
IBM FlashSystem

# 2026 Introducing FCM5

*Evolving the form factor from U.2 to E3.L to meet the next generation of computational storage requirements with a form factor designed for NAND rather than spinning disk…*

## Performance

*Better power delivery and signal integrity*

## Cooling

*Optimised shapes allow for more efficient airflow*

## Scalability

*Supports future needs for higher capacities, faster speeds*

# FCM5

*E3.L 1T and U.2 Comparison*

- Taller / Longer / Thinner Form Factor
- Designed for Solid State not Spinning Disks
- Increased capacity without increased appliance footprint
- 20% more raw capacity
- New card capacity FCM5-XXL

FCM5

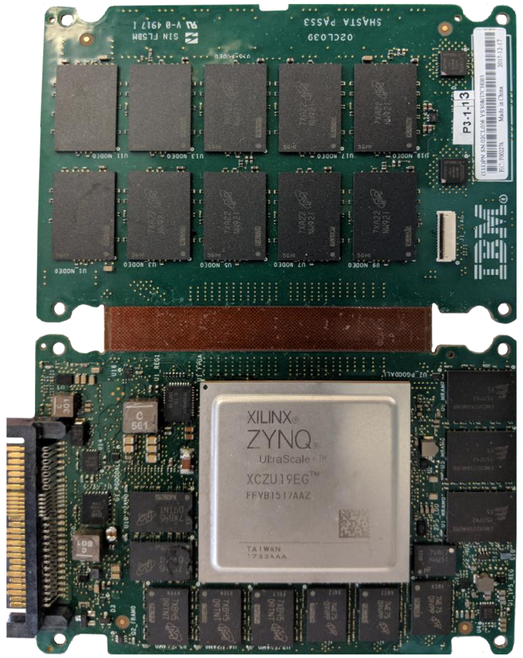

FCM4

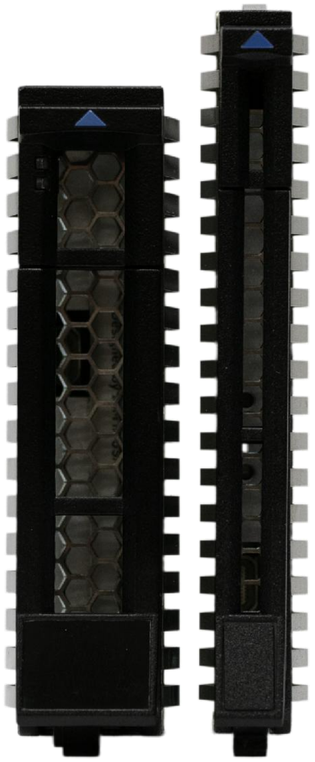

# U.2 to E3 Drive Form Factor Transition

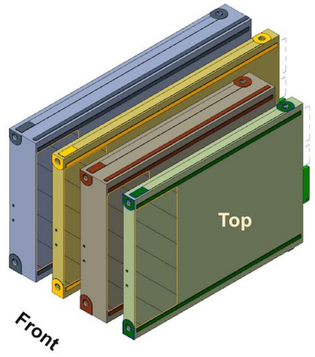


- Improved Performance: Better power delivery and signal integrity.
- Better Cooling: Optimized shapes allow for more efficient airflow.
- Scalability: Supports future needs for higher capacities, faster speeds

U.2 Form Factor

~69mm x ~15mm x ~100mm

E3.L Form Factor

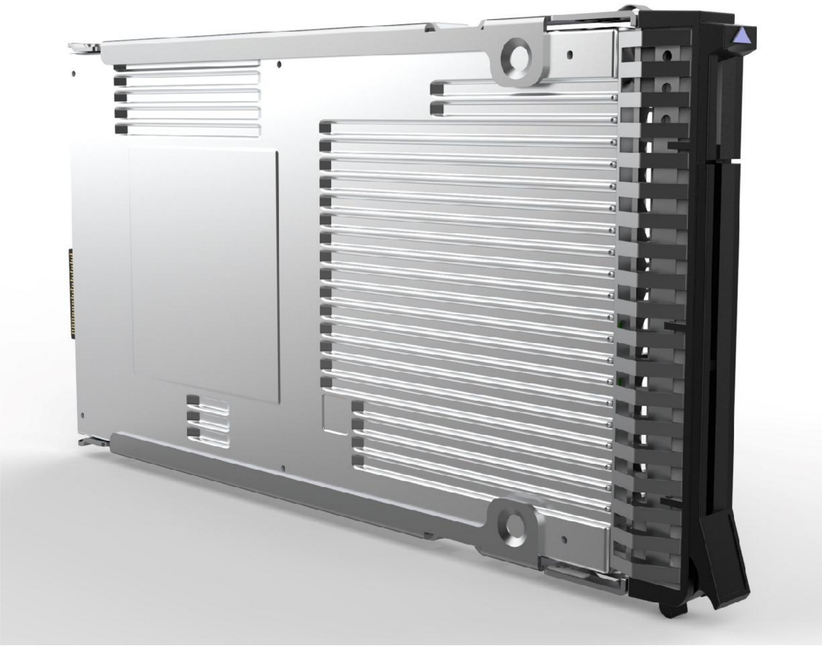

~76mm x ~7.5mm x ~112mm

EDSFF = Enterprise and Datacenter Standard Form Factor

# FCM5

- New Enterprise Datacenter Form Factor (EDSFF)
- Latest 2Tb QLC NAND Flash Technology
- Increased Raw Capacity
- Increased Effective Capacity
- Increased Performance
- Increased Computational Offload
- Increased Security
- Same FlashCore Reliability

** FCM 2XL is capable of 6:1 addressability, however FlashSystem DRAID addressability is restricted to 4.88:1*

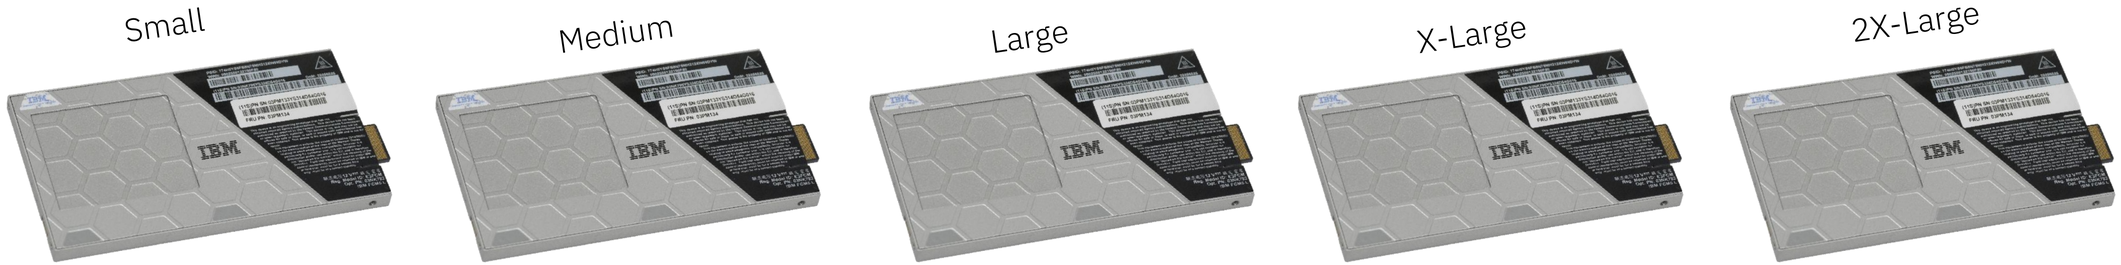


| | Small | Medium | Large | X-Large | 2X-Large |
|---|---|---|---|---|---|
| Raw Capacity | 6.60 TB | 13.20 TB | 26.40 TB | 52.80 TB | 105.60 TB |
| Physical Capacity | 5.58 TB | 11.16 TB | 22.33 TB | 44.66 TB | 89.32 TB |
| Effective Capacity | 33.60 TB | 67.20 TB | 134.40 TB | 268.80 TB | 438.00 TB |
| Data Reduction | 6:1 | 6:1 | 6:1 | 6:1 | 4.88:1* |

*+ 3XL soon… 1PBe!*

# All the existing FCM capabilities…

*"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."*

*Steve McDowell - Forbes - 2024*

*"IBM's FlashCore Module is not an ordinary SSD.*

*It doesn't just serve bytes across an NVMe interconnect. Instead, the IBM FlashCore Module is a complex computational storage device. The logic inside FlashCore makes QLC flash perform better and more reliably, with higher densities than most TLC-based solutions. This is a big win not just for IBM, but for its storage customers overall."*

*Steve McDowell - Forbes - 2022*

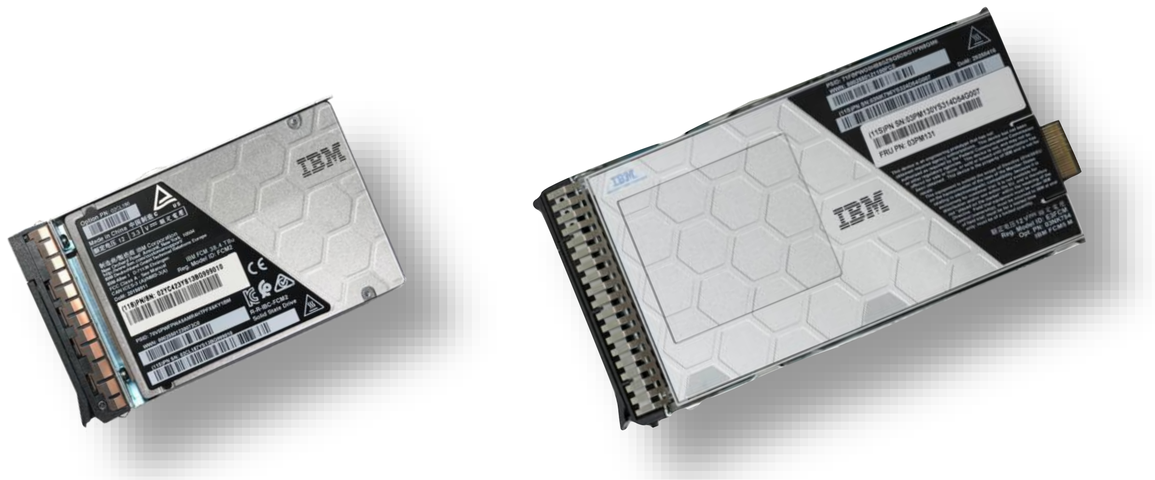


https://www.forbes.com/sites/moorinsights/2022/06/10/flashcore-v3-enables-ibms-storage-differentiation/?sh=1cc1605b225d
https://www.forbes.com/sites/stevemcdowell/2024/02/27/ibms-introduces-data-path-ransomware-detection-in-new-flashsystem/?sh=341ee60ede93

## Efficiency

- FPGA hardware paths (no firmware)
- Zero impact inline hardware compression

## Protection

- Highest in class ECC encoding
- Zero impact inline hardware quantum same encryption
- Zero impact inline anomaly detection

## Performance

- Pseudo-SLC read cache
- Read/Write heat segregation
- Smart data placement - hinting architecture
- PCIe Gen4
  - 3.5GB/s
  - 175K IOPs

## Endurance

- Health and heat binning - asymmetric wear levelling
- Lowest in class write amplification
- Variable voltage levels - lifespan
- Highest in class endurance
- *IBM has had zero FCM's wear out in the field!*

# Computational Storage

*FCM5*

- Custom IBM Controller
  - NVMe Management Processor
  - Flash Management Processor
  - x2 Ransomware Analysis Processors
  - Configurable Hardware Datapath and Controller
  - Hardware Accelerators
- Offload computation functions to the FCM
  - Hardware Compression / Decompression
  - Ransomware Threat Detection
  - AES-256 Hardware Encryption
  - FCM Deduplication Assists
- Reduce System Cost
- Increase System Performance

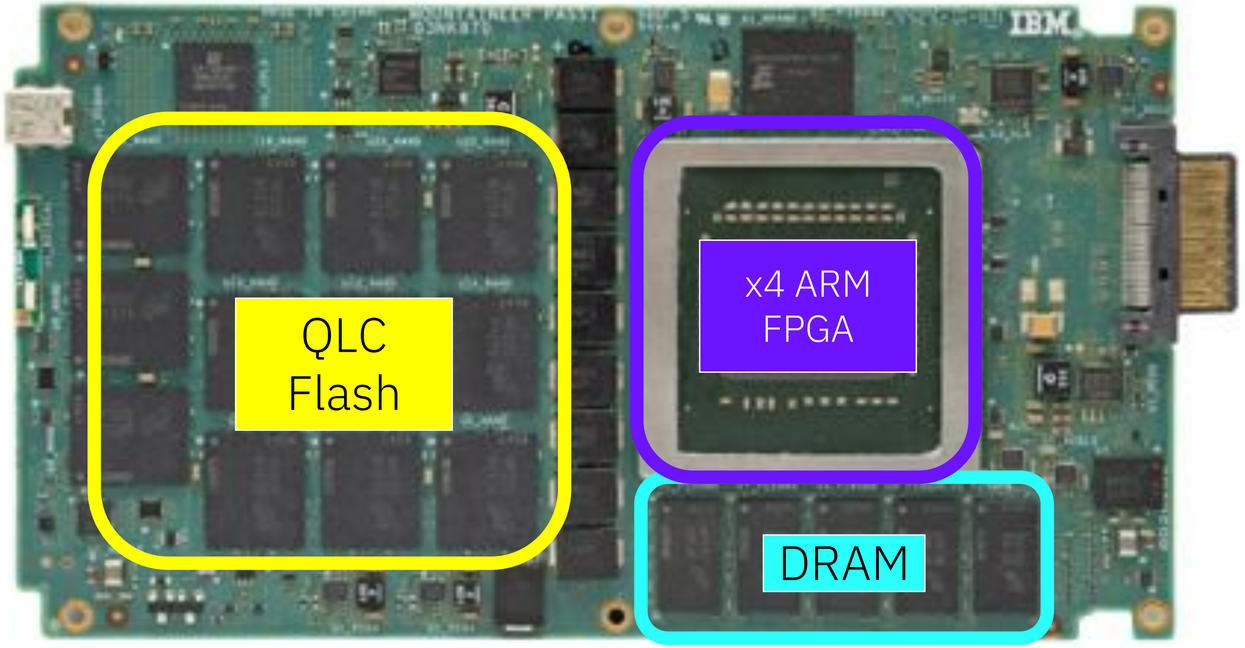

*Adding additional FCMs gives your system a “boost”*

# Compression

*IBM FlashSystem with FCM*

- FCM does All the hard(ware) work
- Decreased system hardware cost
- Lower memory requirements
- No system performance loss
- No Controller CPU cycles or add-in cards
- No performance penalty for incompressible data
- Optimized Cost, Power, and Performance

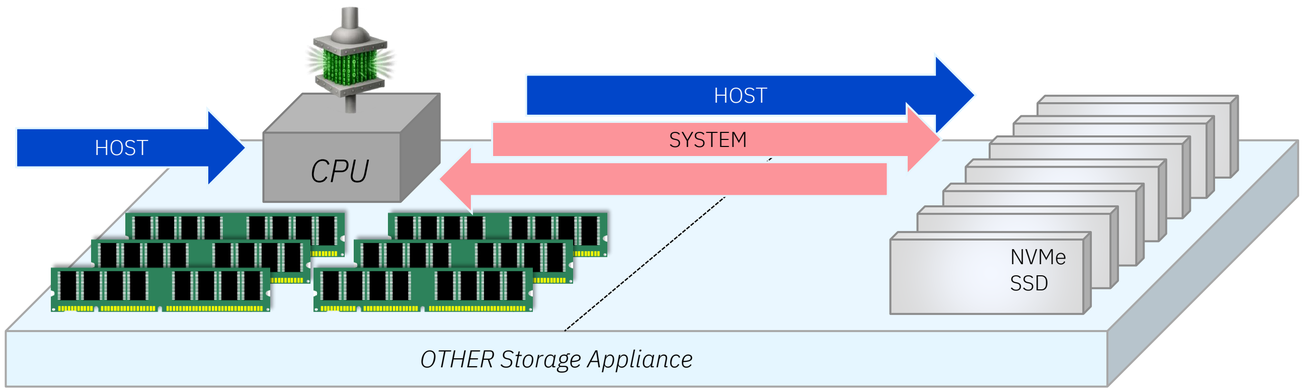


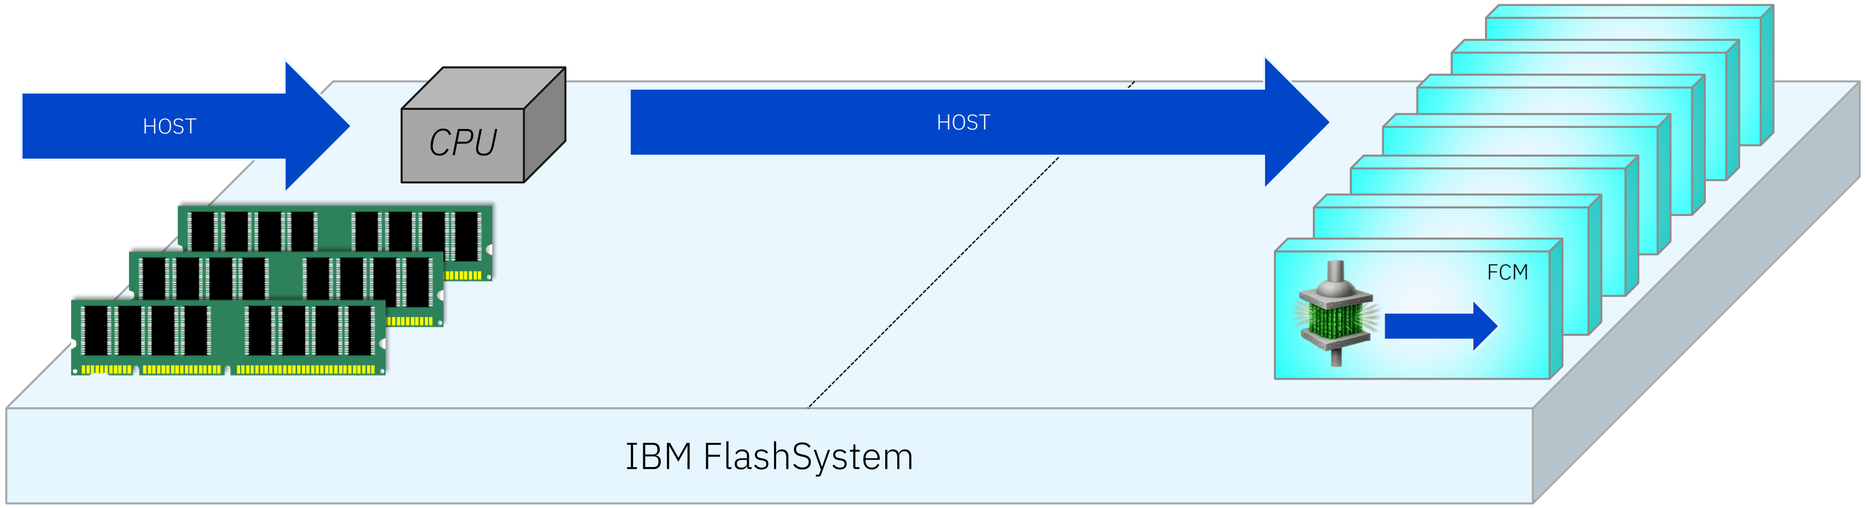

# FCM Differentiation. “Tiering” inside every FCM

- FCM tracks the most recent and frequently used LBAs
- LBA metadata stored in SLC Flash
- Hot LBA metadata cached in DRAM
- Enterprise performance, lower cost

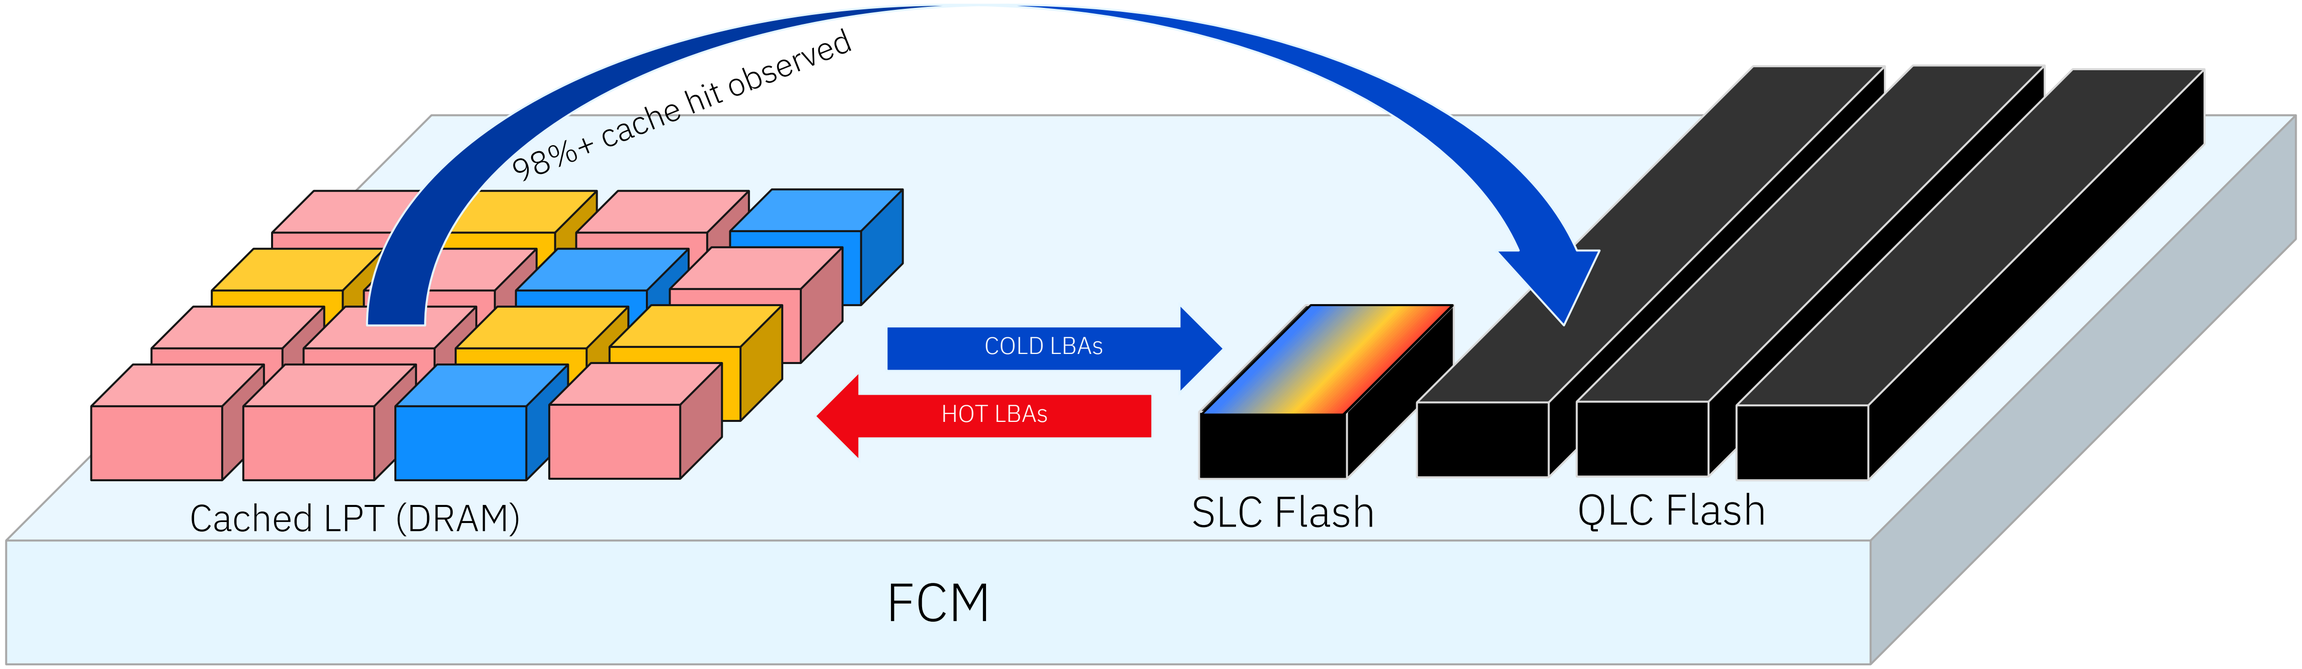

# FCM Security

*IBM Quantum Safe Technology*

FCM5 uses a hybrid implementation of Post-Quantum Cryptography algorithms for Key Transport and Image Authentication:

- FIPs 203 and 204 Compliant Algorithms

Customer data is stored in flash memory after being encrypted by a hardware XTS-AES-256 encryption engine.**

*** A quantum computer capable of breaking AES-256, at least with any algorithm known today (including Grover's), is likely several decades away at least*

# FCM Differentiation

*Ransomware Threat Detection*

- FCM is the ideal place to track and analyze data for intrusion detection
- Data is analyzed alongside the normal data flow, so there is no performance penalty, and no metadata storage needed.
- Statistics are collected and summarized based on the volume for every I/O
- Data analysis is fault resilient, stalls or issues in the analysis will not affect the storage functions of the FCM
- New Statistics added to FCM5 **NEW**
- Current models *still* do a great job of detecting encryption attacks
- Future models aim to target Exfiltration and Wiperware

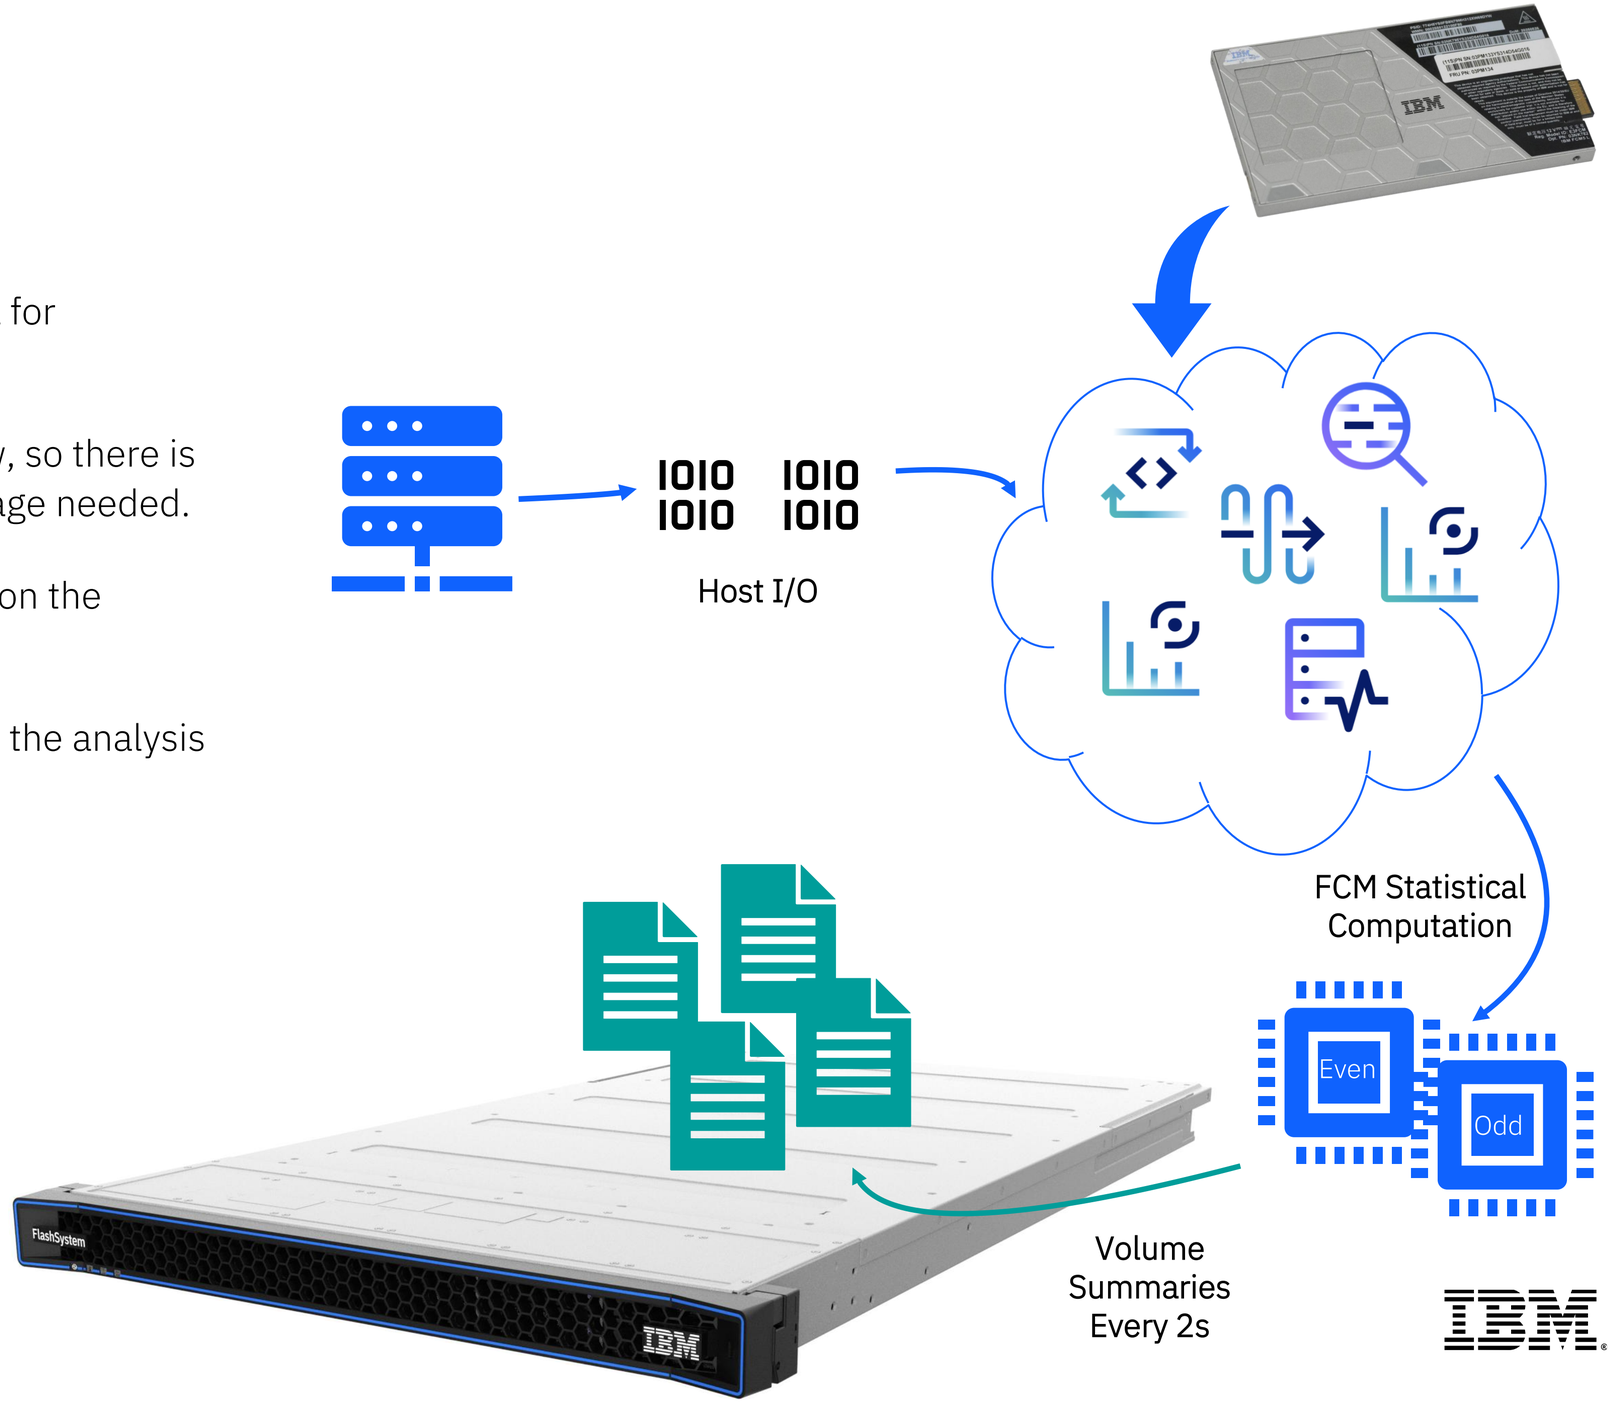

# But wait…
# FCM5 adds more …

Why don’t we extend the FCM offloads to include Deduplication…

*Another industry first innovation…*

# Deduplication

*The Concept and the Challenge*

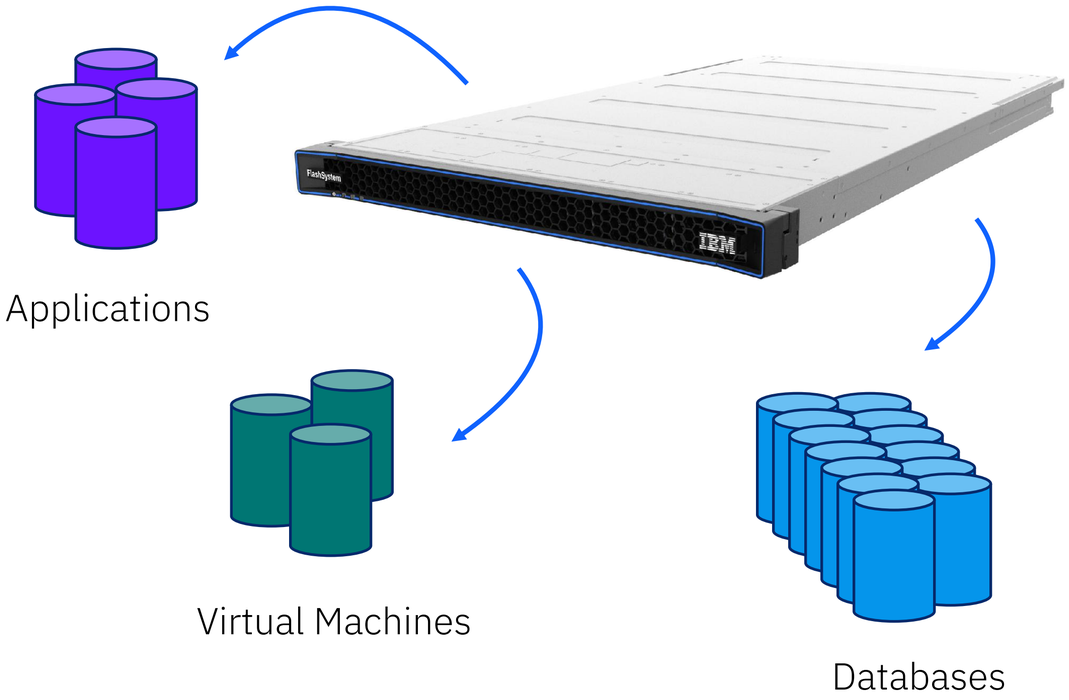


## Concept

- Deduplication conceptually is straightforward
  - *Hash chunks of data*
  - *Compare hashes to find duplicates*
  - *Save fewer copies of the data*
  - *Keep a list of duplicates that reference the saved data*
- VM / VDI / VSI typically have the best deduplication ratio

## Challenge

- As System and SSD capacity grows, there is **more and more data** to analyse
- "Brute Force" dedupe requires a LOT of processing and a LOT of memory for metadata
  - *For example: 1PiB dataset, 32B SHA-2 Hash, 4KiB Data Chunk per Hash*
  - *Requires ~8TiB of memory to store the hashes*
  - *New host writes need to search through a LOT of hashes to dedupe inline*
- Increased system cost and performance degradation quickly erode the value of deduplication

*There must be a better way to do dedupe…*

# Deduplication

*Background Scan*

- Canisters will “scan” each FCM in the background to preserve host performance
- FCM will provide Metadata for each allocated logical page
  - Light-weight 8B custom hash
  - Compressibility
  - Write heat
- Spectrum Virtualize will analyze the hashes and sort duplicates
  - Minimal processing required
  - Only ~2.5GiB of System DRAM !!!

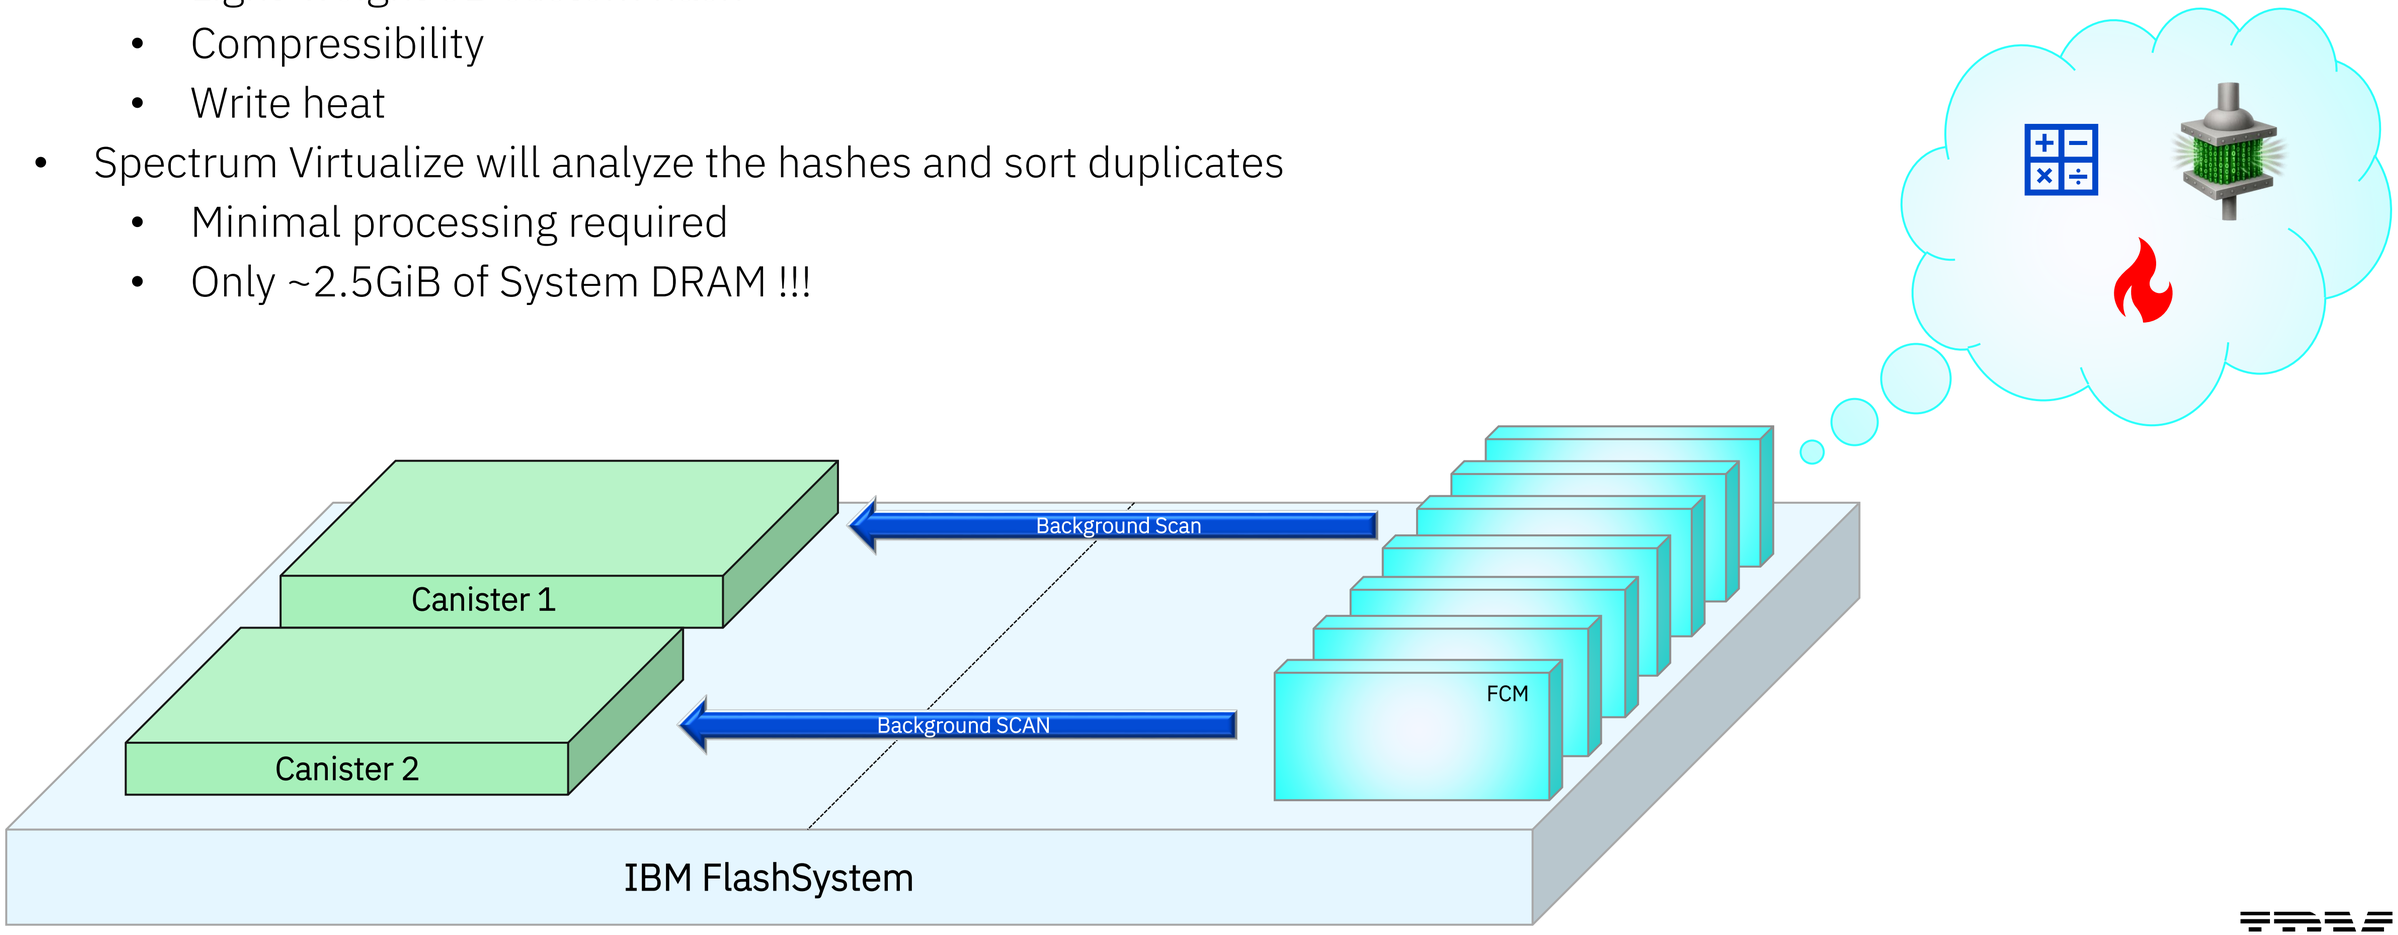

# Deduplication

*Mapping*

- Single copy of each duplicate is moved to a “Dedupe Slot” in the FCM
  - Dedupe slots cannot be overwritten by the host, but look like a normal page to FCM
- All other copies are then “mapped” to the corresponding Dedupe Slot
  - SV provides a list and FCM does the work
  - FCM will check each candidate internally to ensure it is an exact duplicate
  - FCM will modify existing Logical to Physical Table with a reference or pointer to the dedupe slot
- Mapping data is completely stored and maintained in existing FCM metadata
- Host reads to deduped data are done in-line with no performance degradation (one extra DRAM read in the FCM)

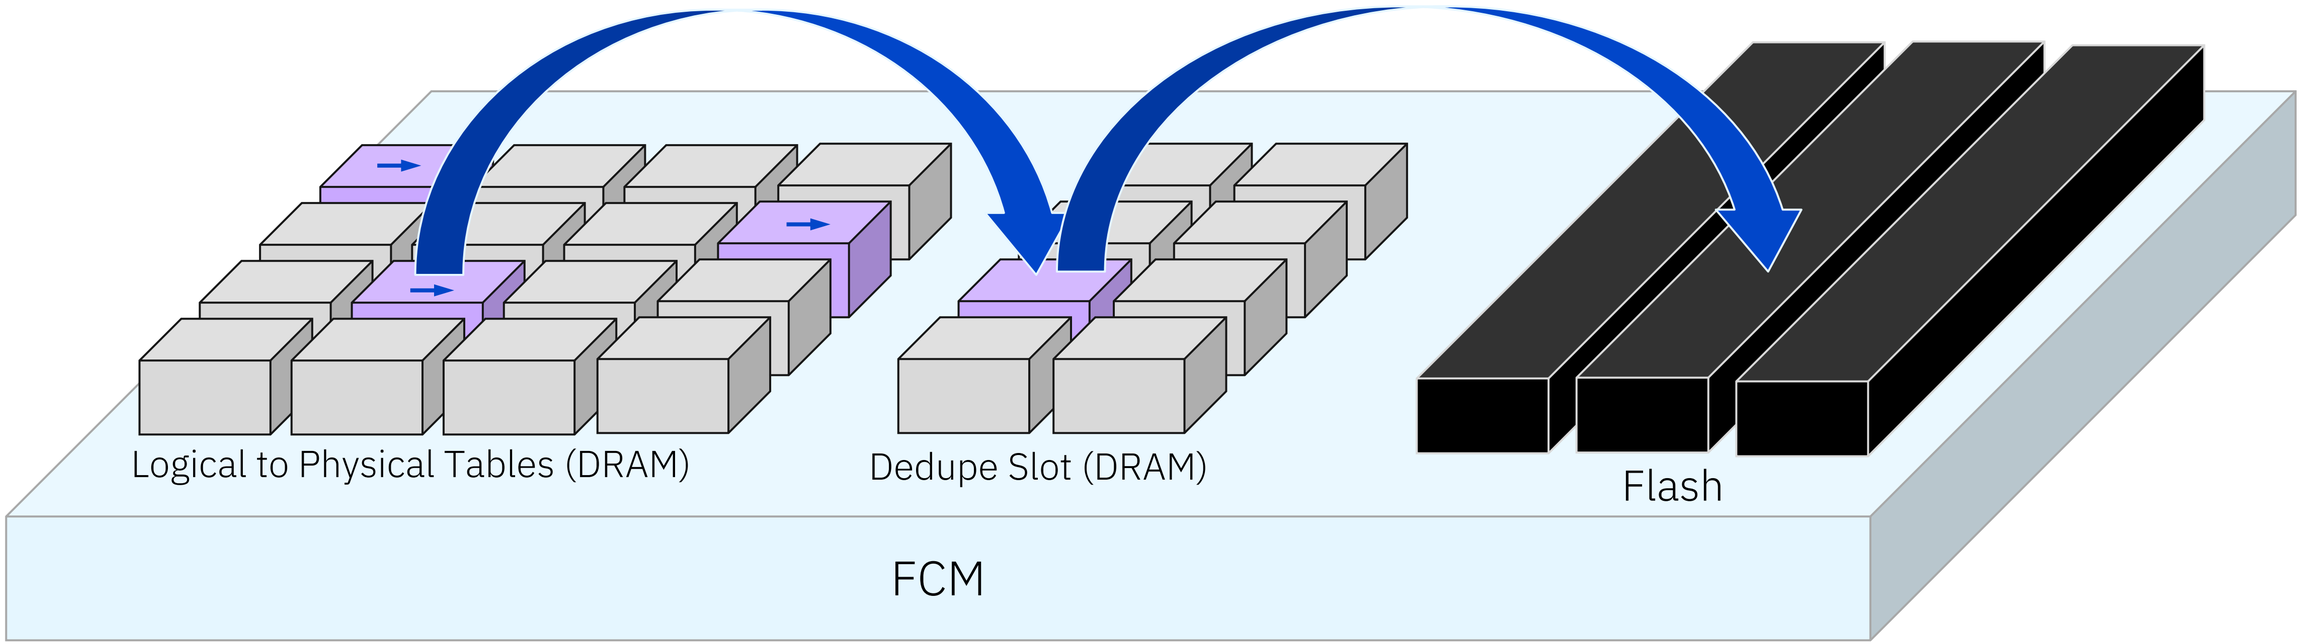

# Deduplication

*Summary*

- Deduplication scanning, analysis, and mapping all done in the background to preserve host performance
  - Incoming data will still get the benefit of inline hardware compression
- Appliance controller uses minimal processing to perform dedupe scan and analysis
- Controller only uses ~2.5GiB DRAM
- All Dedupe Mapping information is contained in existing  FCM Logical to Physical Tables
- Reads to pages that have been deduped are done in-line by the FCM, no added Read Latency
- No workload restrictions
- No capacity restrictions
- A Modern / Optimized approach to Deduplication
  - Enterprise Data Reduction without the cost and performance loss of “Brute Force” Dedupe!

# FCM Assisted Deduplication

*Case Study*

Setup:

- FlashSystem 5600
  - 6 Small FCM5, DRAID6
- 300 Fedora Core Virtual Machines
- Average card capacity before deduplication = 2.635 TB

Results:

- Average card capacity after deduplication = 1.074 TB
- Deduplication Ratio = 2.45:1
- Physical Capacity Savings = 60%

- *No Extra / Expensive Hardware*
- *No Performance Loss*

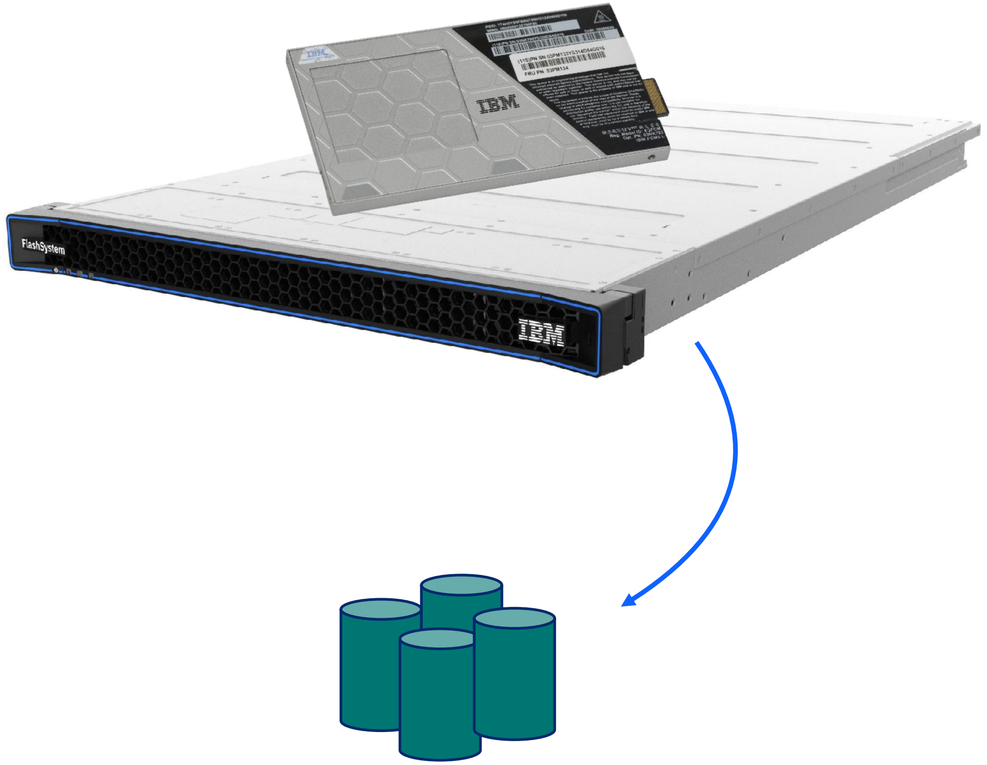

# FCM Assisted Deduplication

*Detailed Notes*

- System will scan only 1 FCM per node at a time, sequentially.
- If something is degraded or non-redundant - processing will pause or abort :

***Pause*** :

- When one node goes away (any non committed scanning work the missing node was doing is discarded)

***Resume*** :

- One minute after the node re-joins the system

***Abort :***

- If the DRAID array is unhealthy (i.e. not fully redundant)
- If the DRAID is expanding
- If software upgrade has started - re-assessed two minutes after upgrade is committed

No scanning interim data (within the nodes) is held in NV memory - any non committed work, or scanning results will be discarded and the process will begin again from scratch as appropriate.

# With the new FCM5 E3 form factor…

*Don’t we need some new enclosures?*

# IBM FlashSystem 2025

**FlashSystem C200**
Capacity optimised, fixed config (2.3PBe)
for lower performance requirements
<200K IOPS, 1-2ms

**FlashSystem 5300**
1U entry level NVMe for general
purpose workloads
>400K IOPs, <0.5ms, >40GB/s, 1PBe

**FlashSystem 7300**
2U balanced performance and value
workhorse system
>500K IOPs, <0.5ms, >50GB, 2.2PBe

**FlashSystem 9500**
Extreme performance and scaling for
mission critical workloads
>1.5M IOPs, <0.5ms, >90GB/s, 4.4PBe

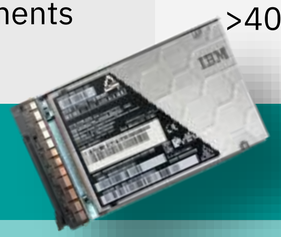

IBM FlashCore Modules

High Density, Low Footprint, Lowest Power in the industry - sustainable, less hardware to achieve performance
Best and consistent performance with FCM compression and QLC 176-layer NAND
Cyber Vault, Safeguarded Copy and inline data corruption detection - unique in the industry

Powered by IBM Storage Virtualize

Managed, monitored and proactive support by IBM Storage Insights

**Storage Virtualize**
**for Public Cloud**

Hybrid Cloud integration
TCT scheduled offload to S3

**SAN Volume Controller (SVC)**

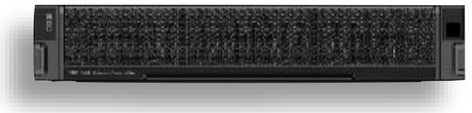

SVC Storage Engine

**FlashSystem 5015 & 5045**

# IBM FlashSystem 2026

**FlashSystem C200**
Capacity optimised, fixed config (2.3PBe)
for lower performance requirements
<200K IOPS, 1-2ms

NEW

**FlashSystem 5600**
1U entry level NVMe for general
purpose workloads
>450K IOPs, <0.5ms, >30GB/s, 2.4PBe

NEW

**FlashSystem 7600**
2U balanced performance and value
workhorse system
>800K IOPs, <0.5ms, >70GB/s, 7.2PBe

NEW

**FlashSystem 9600**
Extreme performance and scaling for
mission critical workloads
>3.0M IOPs, <0.5ms, >120GB/s, **11.8PBe**

IBM FCM Gen4

IBM FlashCore Gen5 EDSFF

NEW

High Density, Low Footprint, Lowest Power in the industry - sustainable, less hardware to achieve performance
Best and consistent performance with FCM compression, deduplication, RTD and QLC 176-layer NAND
Cyber Vault, Safeguarded Copy and inline data corruption detection - unique in the industry

Powered by IBM Storage Virtualize

Managed, monitored and proactive support by IBM Storage Insights

**Storage Virtualize for Public Cloud**

Hybrid Cloud integration
TCT scheduled offload to S3

**SAN Volume Controller (SVC)**

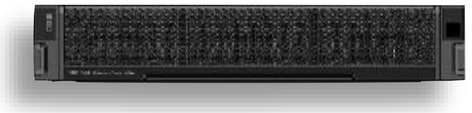

SVC Storage Engine

**FlashSystem 5015 & 5045**

# FlashSystem 5600 is a combination of new and familiar

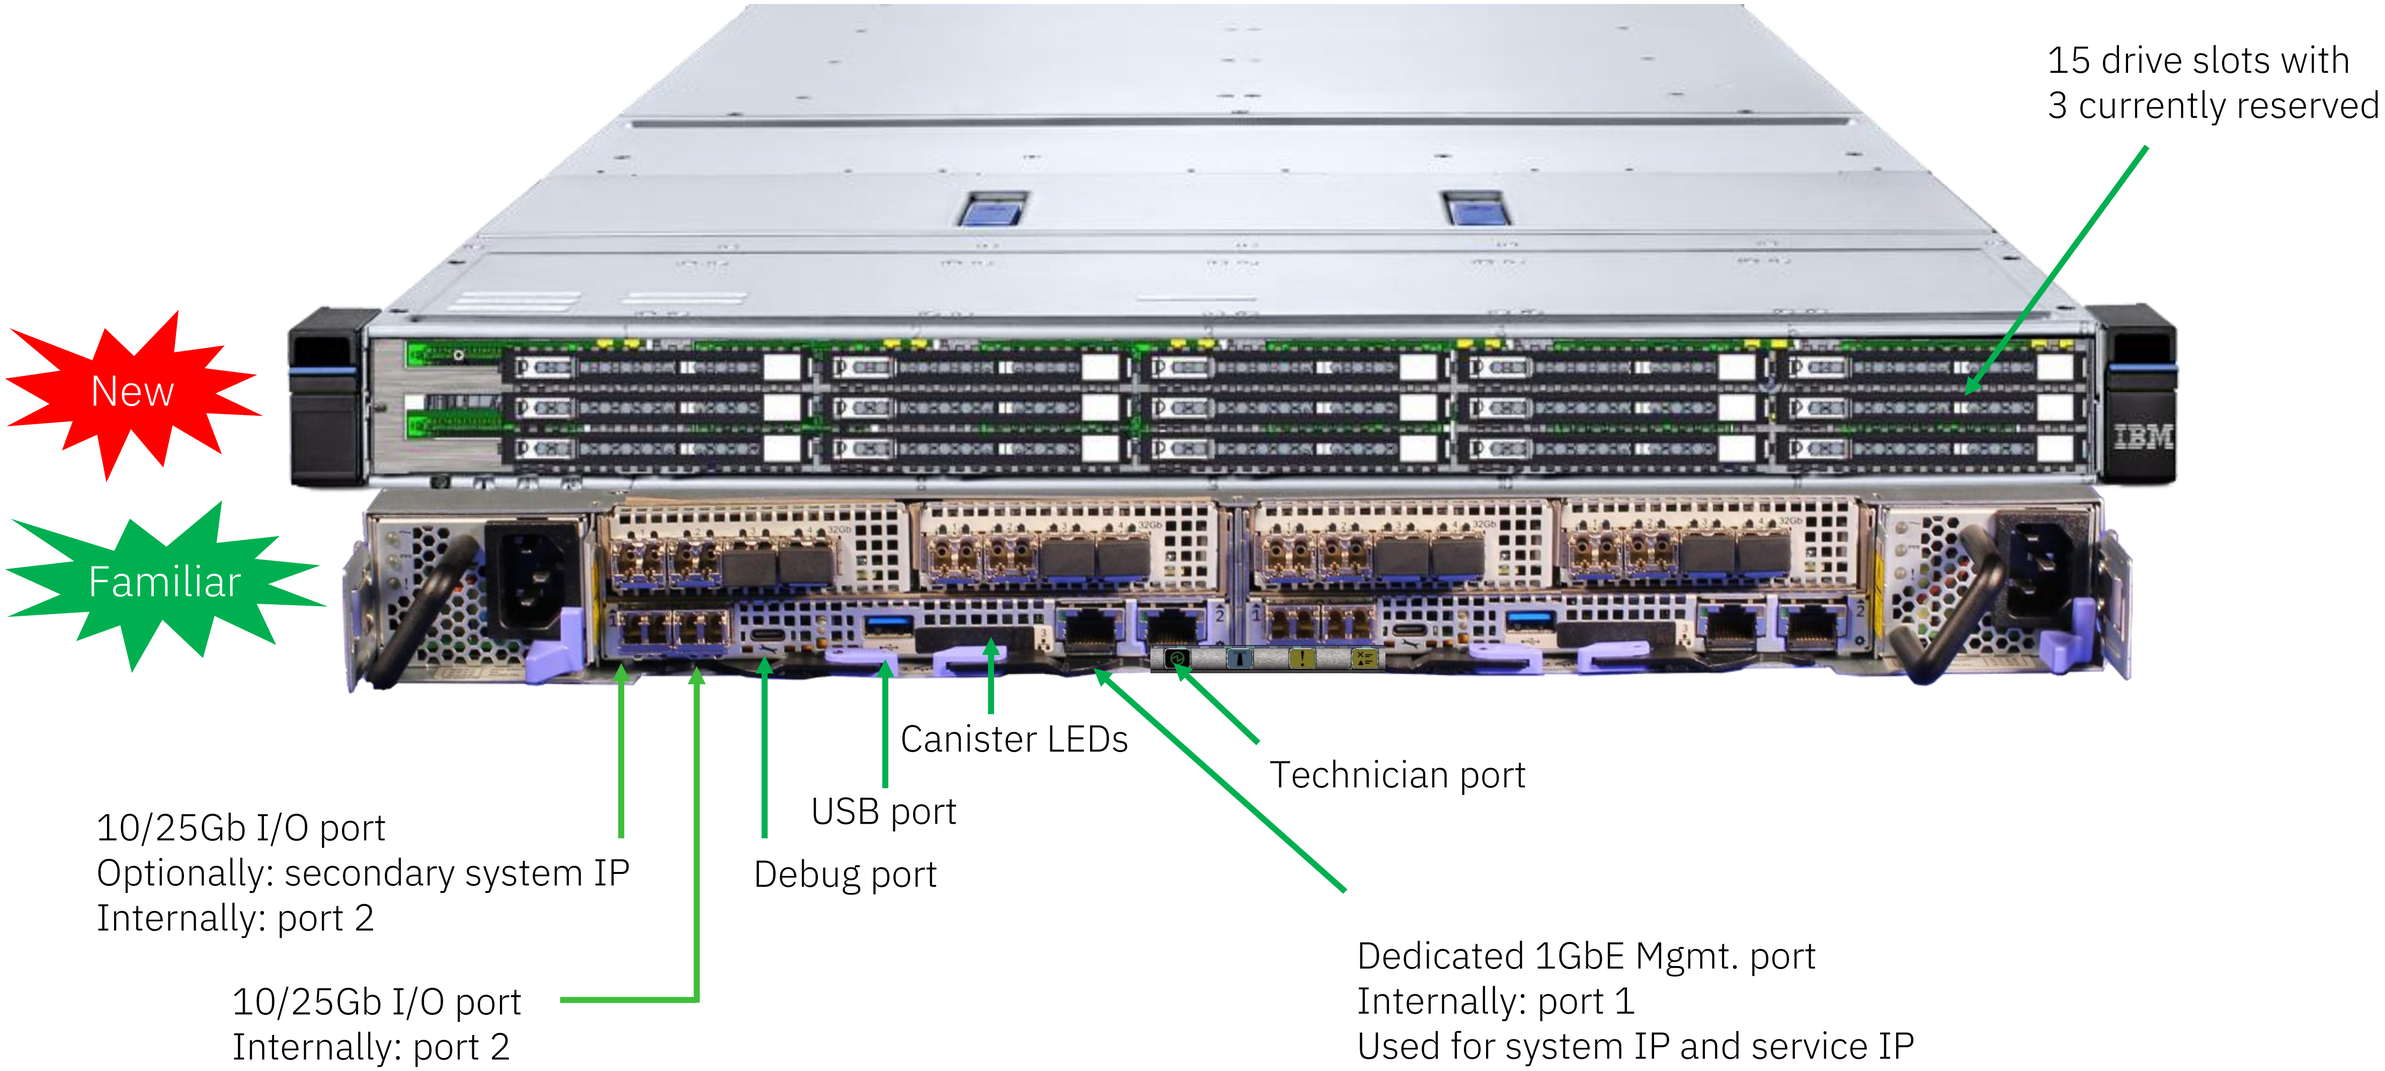

# FlashSystem 5600 - FCM5 Capacity Examples

*New EDSFF enclosure, exactly the same node canisters as existing 5300*

| Capacity Points in TB | Small | Medium | Large | XL |
|---|---|---|---|---|
| x1 FCM5 RAW | 6.6 | 13.2 | 26.4 | 52.8 |
| x1 FCM5 physical | 5.6 | 11.2 | 22.4 | 44.8 |
| x1 FCM5 effective | 33.6 | 67.2 | 134.4 | 268.8 |
| x32 FCM5 physical w/RAID | 50.4 | 100.8 | 201.6 | 403.2 |
| x32 FCM5 effective w/ RAID | 302.4 | 604.8 | 1209.6 | 2419.2 |

# A quick review of FlashSystem 5600

All inclusive Licensed Machine Code with support for all Advanced Services including

- Volume Group Snapshot and Safeguarded Copy (SGC)
- Policy Based Replication (PBR)
- Policy Based High Availability (PBHA)
- Data at Rest Encryption on all drives (using SED or the SAS)
- FCM Compression and Dedupe (requires memory upgrade)
- Ransomware Threat Detection
- Secure Boot and Measured Boot for Linux environments
- Storage Insights integration
- FlashSystem grid scale out for up to 32 FS arrays and SVCs
- Support for 8 Arrays of DRAID 1, 5 and 6

Protocol support

- NVMe/TCP and iSCSI with high-speed Ethernet
- FC-NVMe and FC-SCSI
- NVMe/RoCE is Post-GA TBD

Drive support

- NVMe Support for up to 12 drives of FCM5 Sm - XL capacities
- NVMe Support for up to 12 industry standard SSDs in 1.9-7.68TB capacity
- SAS Expansion up to 4 5U92 or 6 2U12 on 2 SAS chains

Memory support

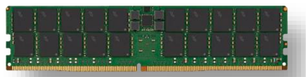

- Base - 128GiB per node for 256GiB per system with 50% bandwidth
- Upgrade 1 – 256GiB per node for 512GiB per system with 100% bandwidth

# The 7600/9600 has incredible density, resiliency and simplicity

Two 2T-E3 slots for potential future use-cases

Thirty-two 1T-E3 slots for FCM5s or, for industry-standard commodity NVMe SSDs

Two 2T-E3 slots for potential future use-cases

# FlashSystem 7600/9600 offers connectivity, and resilience

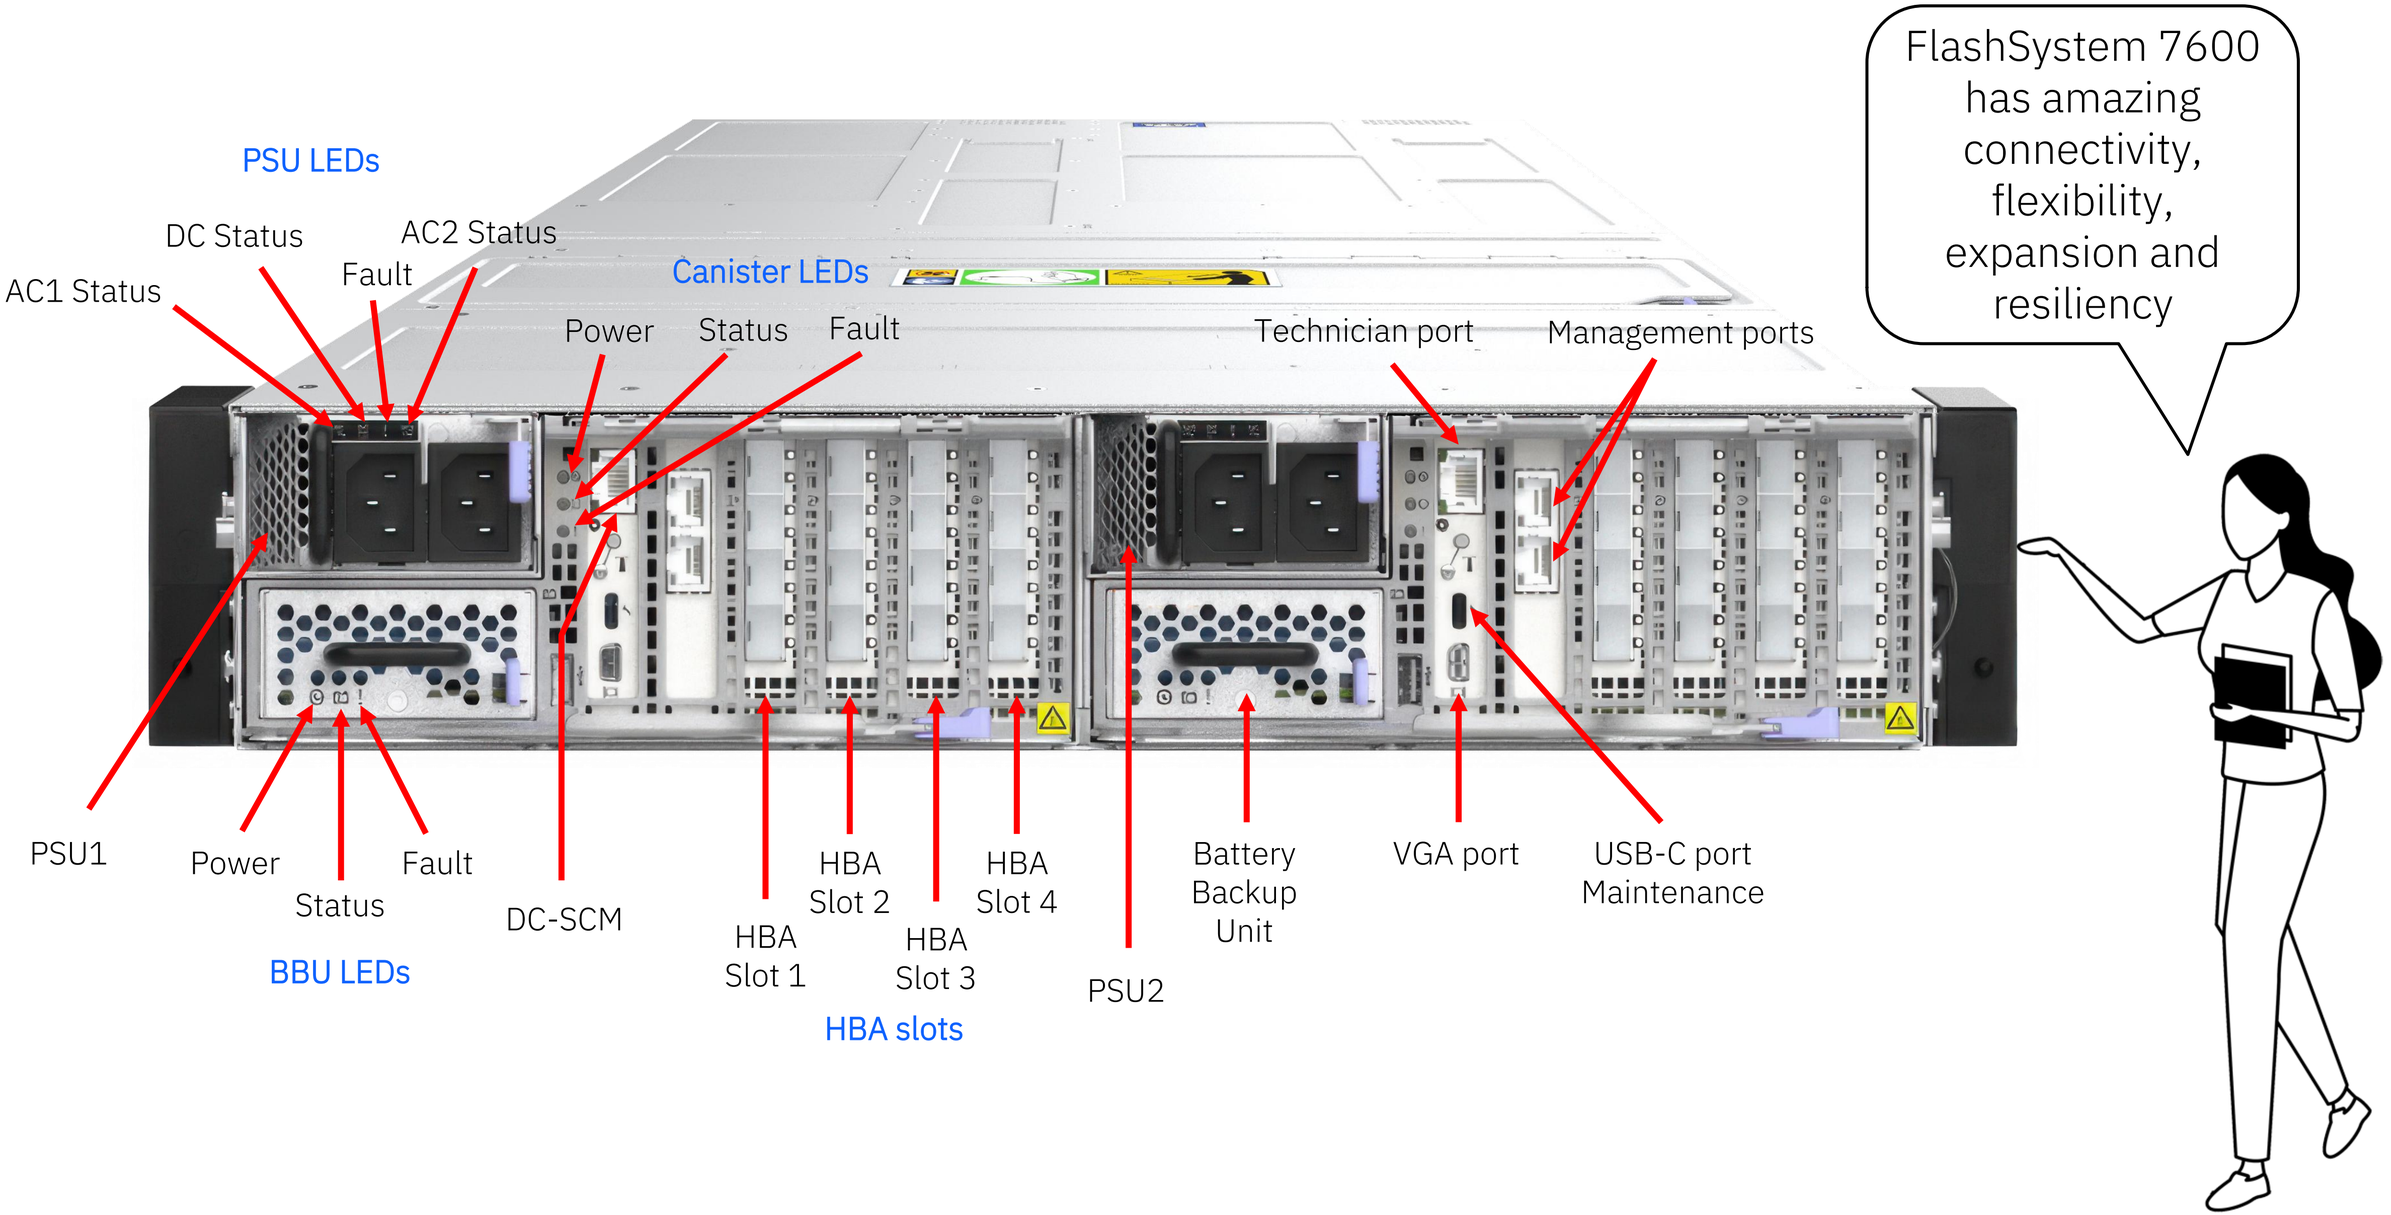

# FlashSystem 7600 - FCM5 Capacity Examples

| Capacity Points in TB | Small | Medium | Large | XL |
|---|---|---|---|---|
| x1 FCM5 RAW | 6.6 | 13.2 | 26.4 | 52.8 |
| x1 FCM5 physical | 5.6 | 11.2 | 22.4 | 44.8 |
| x1 FCM5 effective | 33.6 | 67.2 | 134.4 | 268.8 |
| x32 FCM5 physical w/RAID | 151.9 | 303.8 | 607.6 | 1215.2 |
| x32 FCM5 effective w/ RAID | 911.4 | 1822.8 | 3645.6 | 7291.2 |

# FlashSystem 9600 - FCM5 Capacity Examples

The new 2XL 105.6TB FCM5 is exclusive to to the 9600

*~4.88:1*

| Capacity Points in TB | Small | Medium | Large | XL | 2XL |
|---|---|---|---|---|---|
| x1 FCM5 RAW | 6.6 | 13.2 | 26.4 | 52.8 | 105.6 |
| x1 FCM5 physical | 5.6 | 11.2 | 22.4 | 44.8 | 89.6 |
| x1 FCM5 effective | 33.6 | 67.2 | 134.4 | 268.8 | 438.0 |
| x32 FCM5 physical w/RAID | 151.9 | 303.8 | 607.6 | 1215.2 | 2430.4 |
| x32 FCM5 effective w/ RAID | 911.4 | 1822.8 | 3645.6 | 7291.2 | 11880.8 |

# A quick review of FlashSystem 7600

All inclusive Licensed Machine Code with support for all Advanced Services including

- Volume Group Snapshot and Safeguarded Copy (SGC)
- Policy Based Replication (PBR)
- Policy Based High Availability (PBHA)
- Data at Rest Encryption on all drives (using SED or the SAS)
- FCM Compression and Dedupe (requires memory upgrade)
- Ransomware Threat Detection
- Secure Boot and Measured Boot for Linux environments
- Storage Insights integration
- FlashSystem grid scale out for up to 32 FS arrays and SVCs
- Support for 8 Arrays of DRAID 6

Protocol support

- NVMe/TCP and iSCSI with high-speed Ethernet
- FC-NVMe and FC-SCSI
- NVMe/RoCE is Post-GA TBD

Drive support

- NVMe Support for up to 32 drives of FCM5 Sm - XL capacities
- NVMe Support for up to 32 industry standard SSDs in 1.9-7.68TB capacity
- With 32 FCM5 slots in the 7600, SAS expansion is not supported

Memory support

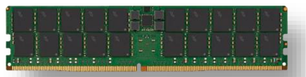

- Base - 384GiB per node for 768GiB per system with 50% bandwidth
- Upgrade 1 – 768GiB per node for 1.5TiB per system with 100% bandwidth

# A quick review of FlashSystem 9600

All inclusive Licensed Machine Code with support for all Advanced Services including

- Volume Group Snapshot and Safeguarded Copy (SGC)
- Policy Based Replication (PBR)
- Policy Based High Availability (PBHA)
- Data at Rest Encryption on all drives (using SED or the SAS)
- FCM Compression and Dedupe
- Ransomware Threat Detection
- Secure Boot and Measured Boot w/ Hardware Root of Trust
- Storage Insights integration
- FlashSystem grid scale out for up to 32 FS arrays and SVCs
- Support for 8 Arrays of DRAID 6

Protocol support

- NVMe/TCP and iSCSI with high-speed Ethernet
- FC-NVMe and FC-SCSI
- NVMe/RoCE is Post-GA TBD

Drive support

- NVMe support for up to 32 drives of FCM5 Sm - XXL capacities
- NVMe support for up to 32 industry standard SSDs in 1.9-7.68TB capacity
- With 32 FCM5 slots in the 9600, SAS expansion is not supported

Memory support

- Base - 768GiB per node for 1.5TiB per system with 50% bandwidth
- Upgrade 1 - 1.5TiB per node for 3TiB per system with 100% bandwidth

# 2U FlashSystem 9600 vs 4U FlashSystem9500

Up to 400% better density per RU

*9500 Facts :*

Rack Space : **4RU**

Max capacity :
Physical usable : **1.5PB**
Effective max : **4.7PB**

Per RU :
**394.8 TB** physical
**1.18 PB** effective

*9600 Facts :*

Rack Space : **2RU**

Max capacity :
Physical usable : **2.43PB**
Effective max : **11.8PB**

Per RU :
**1215 TB** physical
**5.9 PB** effective

*Comparison :*

Rack Space : 50% reduction

Max capacity :
Physical usable : 62% increase
Effective max: 151% increase

Per RU :
Physical usable : 207% increase
Effective : 400% increase

# FlashSystem x600 Platform Hardware Summary

## FlashSystem 5600

- Replacement for the 5300
- 1U 12 NVMe drive enclosure
  - *15 total slots, 3 are reserved for future use*
- New EDSFF enclosure using existing 5300 canisters
- 1T E3.L EDSFF form Factor
  - FCM5 S, M, L, XL
- PCIe Gen 4 canister / DDR4
- 2x12 core Intel Ice Lake D
- 256/512G memory options
- Flexible 10/25GbE on board I/O
- NVMe boot device
- 2 Gen 4 HBA slots per canister
- Internal canister battery
- One piece bezel with illuminated LED strip in steady blue

## FlashSystem 7600

- Replacement for the 7300
- 2U 32 NVMe Drive Enclosure
- 1T E3.L EDSFF form Factor
  - FCM5 S, M, L, XL
- PCIe Gen 5 canister / DDR5
- 2x 16 core AMD Epyc
- 768/1536G Memory options
- On Board Management only Ports with dedicated Technician Port
- 4 Gen 5 HBA Slots Per Canister
- Externally Replaceable Battery
- Side by Side Canister placement for improved Serviceability
- NVMe Boot Drive
- DC-SCM Hardware Root of Trust
- One piece bezel with illuminated LED strip with user selectable colour

## FlashSystem 9600

- Replacement for the 9500
- 2U 32 NVMe Drive Enclosure
- 1T E3.L EDSFF form Factor
  - FCM5 S, M, L, XL, 2XL
- PCIe Gen 5 canister / DDR5
- 2x 48 core AMD Epic
- 1.5/3.0 TiB Memory options
- On Board Management only Ports with dedicated Technician Port
- 4 Gen 5 HBA Slots Per Canister
- Externally Replaceable Battery
- Side by Side Canister placement for improved Serviceability
- Striped NVMe Boot Drives
- DC-SCM Hardware Root of Trust
- One piece bezel with illuminated LED strip with user selectable colour

# NVMe FlashSystem Family at a glance

## FlashSystem grid

| Details: | FlashSystem C200 | FlashSystem 5600 | FlashSystem 7600 | FlashSystem 9600 |
|---|---|---|---|---|
| Type of media | NVMe | NVMe | NVMe | NVMe |
| Rack Units | 2U | 1U | 2U | 2U |
| Drive slots | 24 drive slots | 12 drive slots | 32 drive slots | 32 drive slots |
| Usable / Effective | 1.1PBu / 2.3 PBe | 400TBu / 2.4 PBe | 1.1PBu / 7.2 PBe | 2.3PBu / 11.9 PBe |
| Achieved with: | FCM 4 – 46 TB | FCM 5 – 52.8 TB (XL) | FCM 5 – 52.8TB (XL) | FCM 5 – 105.6 TB (XXL) |
| CPU info per system | 4x10 core | 2x12 core | 2x16 core | 2x48 core |
| CPU cores per system | 40 cores | 24 cores | 32 cores | 96 cores |
| High-speed bus | PCIe Gen 4 | PCIe Gen 4 | PCIe Gen 5 | PCIe Gen 5 |
| Max IOPS: | 1,500,000 | 2,600,000 | 4,300,000 | 6,300,000 |
| Max Read Miss Bandwidth: | 23 GB/s | 22GB/s | 60 GB/s | 85.8 GB/s |
| Max Bandwidth per drive | 0.95 GB/s per drive | 1.83 GB/s per drive | 1.87 GB/s per drive | 2.68 GB/s per drive |
| Max HBA slot per system | 4 | 4 | 8 | 8 |
| Max FC ports/system | 16 – 32Gb FC ports | 16-32Gb FC ports | 32 – 32Gb FC ports | 32 - 32Gb FC ports |
| Mac Ethernet ports/system | 8 -10BgE on-board<br>8 – 25GbE ports | 4 - 25GbE on-board<br>16 – 25GbE ports | 32 – 25GbE ports | 32 – 25GbE ports |

# FlashSystem 2026 Platform Software

## FlashSystem 5600

All inclusive licensed machine code with support for advanced services

- Data at Rest Encryption always on for all drives (using SED or the SAS expansion card)
- Ransomware Threat Detection
- Secure Boot and Measured Boot
- Storage Insights integration
- DRAID-1, 5, 6 support
- SAS expansion up to 4 5U92 or 12 2U12, up to 12U per SAS chain, 2 chains in total

## FlashSystem 7600

In addition to FS 5600, all inclusive licensed machine code with support for all advanced services:

- DRAID-6 support only
- No SAS expansion support

## FlashSystem 9600

In addition to FS 7600, all inclusive licensed machine code with support for all advanced services:

- DRAID-6 support only
- No SAS expansion support

# FlashSystem 2026 Software Highlights

## FlashSystem 5600

Single standard pool type with support for thin provisioning

Data reduction enabled by default on FCM5:
– FCM5 Data reduction
– FCM5 compression

IS drives support fully allocated or thin provisioning in the standard pool

Single NVMe boot device

Additional replication capacity (up from 1PiB to 2PiB)

## FlashSystem 7600

In addition to FS 5600:

Additional Replication Bitmap space (up from 2PiB to 5PiB)

PCIe Gen 5 switch for inter canister link and enclosure services

Switchless NVMe drive attach

DC-SCM hardware provisioned root of trust

## FlashSystem 9600

In addition to FS 7600:

Additional Replication Bitmap space (up from 4PiB to 10PiB)

Dual NVMe boot devices

# FlashSystem 2026 Scaling & Connectivity

| Feature | 5600 A20 | | | 7600 A30 | | | 9600 A40 | | |
|---|---|---|---|---|---|---|---|---|---|
| Max Capacity per enclosure | 2.1 PiB effective | | | 6.4 PiB effective | | | 10.5 PiB effective | | |
| Max Production Capacity for snapshot | 4 PiB | | | 20 PiB | | | 40 PiB | | |
| Max Snapshot Capacity* | 200 PiB | | | 400 PiB | | | 800 PiB | | |
| Max Replication Capacity** | 2 PiB | | | 5 PiB | | | 10 PiB | | |
| Max Host mappable Volumes + Snapshots | 15,863 | | | 32,100 | | | 65,000 | | |
| Max Host mappable Volumes | 8,192 | | | 16,050 | | | 32,500 | | |
| Max 32Gb Fibre Channel Ports*** | 8 per canister | 16 per system | 512 per FlashSystem grid | 16 per canister | 32 per system | 1024 per FlashSystem grid | 16 per canister | 32 per system | 1024 per FlashSystem grid |
| Max 64Gb Fibre Channel Ports*** | 4 per canister | 8 per system | 256 per FlashSystem grid | 12 per canister | 24 per system | 768 per FlashSystem grid | 12 per canister | 24 per system | 768 per FlashSystem grid |
| Max 10/25GbE Ethernet Ports | 8 per canister | 16 per system | 512 per FlashSystem grid | 16 per canister | 32 per system | 1024 per FlashSystem grid | 16 per canister | 32 per system | 1024 per FlashSystem grid |
| Max 40/100GbE Ethernet Ports | 4 per canister | 8 per system | 256 per FlashSystem grid | 8 per canister | 16 per system | 512 per FlashSystem grid | 8 per canister | 16 per system | 512 per FlashSystem grid |
| Data Reduction | FCM5: thin-provisioning, compression and deduplication support, IS: software thin-provisioning | | | | | | | | |
| Max FlashSystem grid Members | 32 | | | | | | | | |

* This assumes Maximum Memory is installed and will scale with the memory options.
** Number is halved with HA + DR
*** Numbers assume same hardware across the FlashSystem grid. FlashSystem grid can mix hardware types, so this number can differ

IBM Confidential until Announce on 10th February 2026

# FlashSystem.AI

## Designed to Adapt, Engineered to Perform, Built to Protect

### FlashCore Modules

Intelligent density on a new form factor

Announcing the 5th generation
New form factor with enhanced data reduction.
Increased capacity for each drive size.

### Next-Gen Hardware

Industry-leading density and efficiency

Controllers delivering up to 6 times more capacity than competitors.

Systems that support up to. 32 drives in just 2U.

### FlashSystem.ai

Self-managing storage powered by AI.

Manage, optimize and secure storage using natural-language interactions, turning intent into action with Agentic AI.

For 1Q’26 we will be able to provision volumes by a request on natural language.

# IBM FlashSystem in 2026

## FlashSystem.ai

Self-managing storage powered by AI.

- Amplify human expertise with IBM trained AI agents.
- Zero-click operations accelerate tasks, increasing admin productivity and minimizing risk.
- Manage, optimize and secure storage using natural-language interactions, turning intent into action with Agentic AI.

## Next-Gen Hardware

Industry-leading density and efficiency

- Deliver consistent, high performance for mission-critical workloads with fifth-generation FlashCore Modules, providing up to 105TB per module.
- Slash data center footprint and power consumption by at least half with a next-gen hardware platform offering up to 6× denser storage.
- No compromises data reduction and threat detection with IBM-exclusive computational storage.

## Operational Resiliency

Business continuity without compromise

- Detect attacks in seconds and recover fast using embedded AI; IBM-unique drive-level threat detection; and immutable Safeguarded Copies.
- Agentic AI detects resiliency gaps and identifies remediation delivering operational resiliency and minimizing downtime, data loss, and business risk.
- Achieve nonstop operations and peace of mind through advanced data resiliency and non-disruptive data mobility.

# IBM FlashSystem.ai

For Storage professionals, a Storage Admin colleague that never sleeps

## Amplify Human Expertise

*Learns from experts, provides feedback, and delivers results.*

## Transform IT Roles

*Shift from a troubleshooting to a strategic innovation mindset.*

## Reduce Complexity

*Manage your complex storage estate & execute routine tasks.*

## Enhance Performance & Compliance

*Deliver consistent & reliable performance and compliance.*

# Talk to your storage where you have an admin that never sleeps

Three fundamentals:

Human in the loop

- Human must initiate and authorize the action via natural language

Actions are on-prem

- One way communication to pull relevant storage information from the cloud

FlashSystem.ai never sleeps

- … so you can!

FlashSystem.ai is just part of the FlashSystem… you don’t have to pay for it. It’s powered by watsonx.ai

# Human-AI Collaboration
# AI-powered user experiences

Accelerate decision-making with AI-driven storage optimization. A team of cooperative AI agents explains their reasoning, suggests performance improvements, and learns from administrator feedback—enhancing human expertise while keeping you in control.

- Transform storage into an intelligent, autonomous system.
- Deliver seamless, self-service operations with minimal manual effort.
- Automatically meet performance, capacity, and availability guarantees.
- Align storage with business goals through Data Contracts for compliance and resilience.

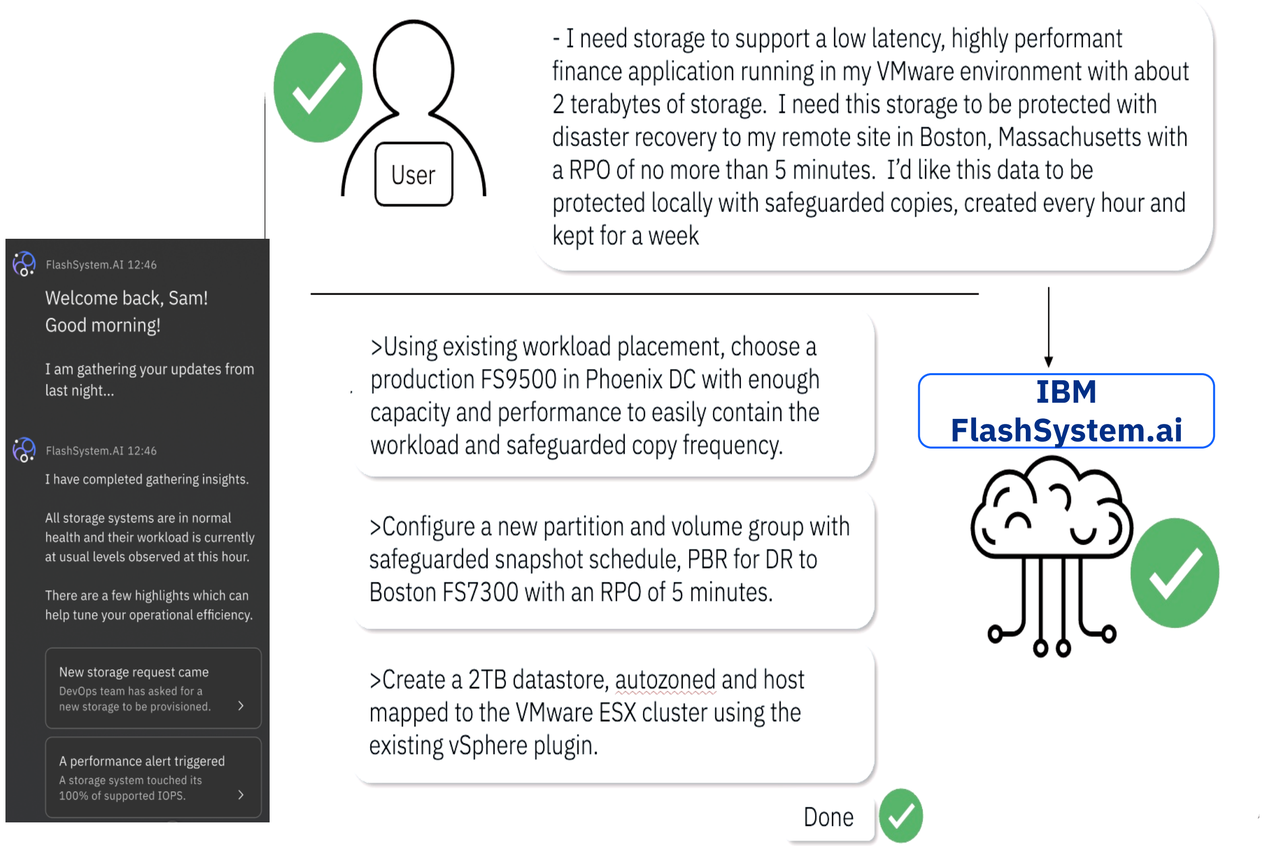


***From command- to conversational-driven, storage that listens, learns, and acts with precision.***

# Self-hosted AI container

- A container hosted in the client's compute provides the AI capabilities for a FlashSystem or grid
- No on-prem GPUs needed
- FlashSystem must be able to reach the FlashSystem.ai container and the AI container must be able to access the REST API of the FlashSystem / SVC
- The container must be able to contact Storage Insights
- Administrator user account for the AI needed on each system, to make any changes, but can have lesser rights.
- Container needs registering with each FlashSystem / SVC
- All actions taken on the FlashSystems are on prem ONLY

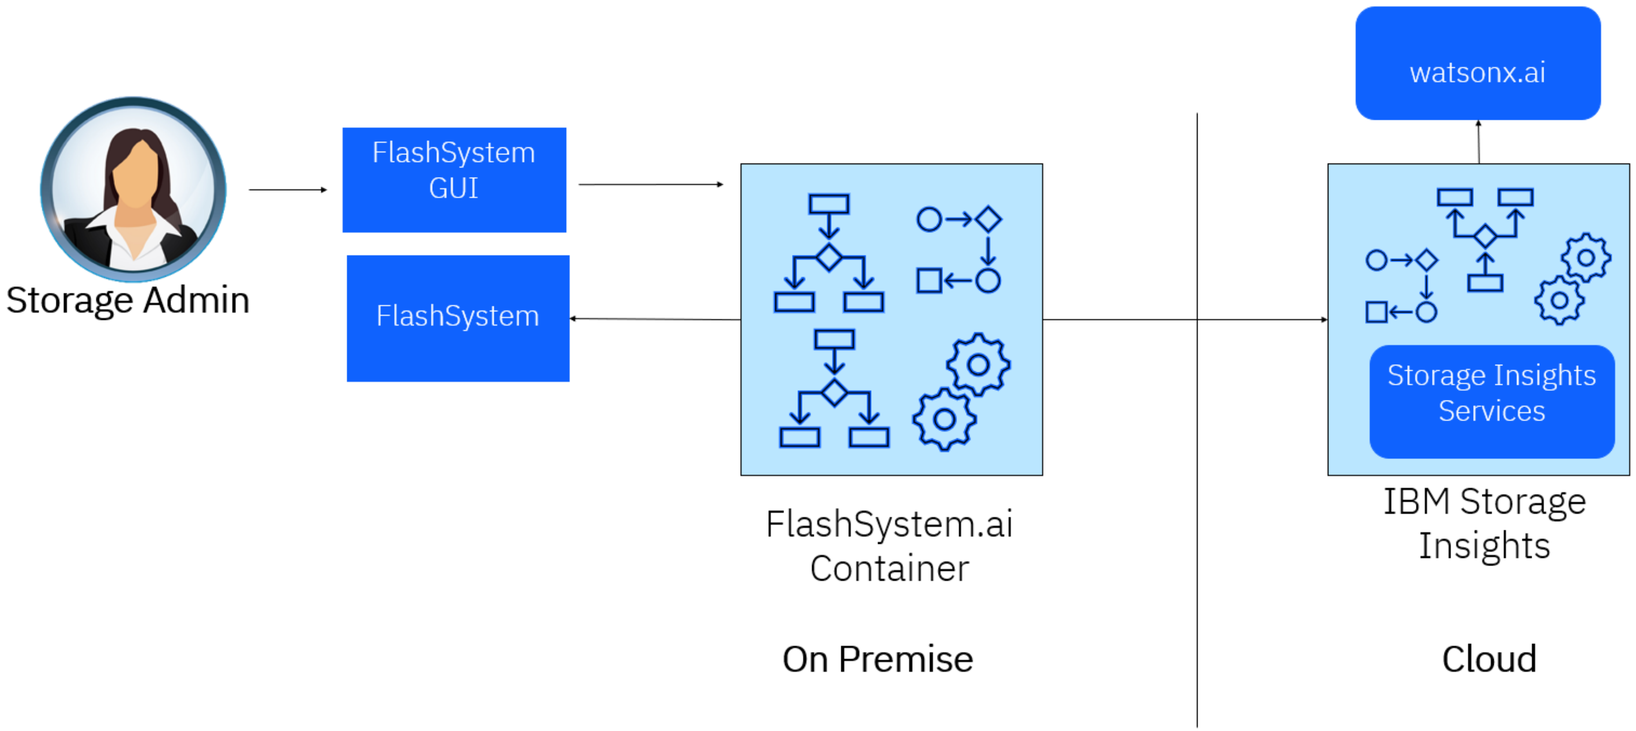

# Initial FlashSystem.ai Agentic capabilities

It is like benefiting from a storage admin who never sleeps.

Outline, inform, curate your information

- Show detailed system overview
- List volumes
- List volume groups
- List snapshot policies
- List replication policies
- List hosts
- List snapshots
- List storage pools
- List partitions and portsets
- List host-volume mappings

Zero Click Provisioning.
You can ask...

Create; act for you – respond to request with ‘human in the loop’

- Create / Add volume
- Create / Add host
- Create / Add volume and attach to host
- Create / Add volume group with snapshot policy
- Create / Add volume group with snapshot policy safeguarded
- Create a snapshot

Update; act for you – respond to request with ‘human in the loop’

- Rename/resize volumes, change volume groups (Can only increase in size can't reduce)
- Modify volume group policies
- Host cluster management (create, list, map volumes)

# IBM FlashSystem.ai

powered by **watsonx**

## What you can't do with FlashSystem.ai

- Delete any object on the system
- Resize any object on the system to be smaller than it is.
- Carry out any configuration change without human in the loop approval
- Does not work with any FlashSystem that is below 9.1.0
- Does not work with the 5015/5045 family

## Why – remember this is based on 'cause no harm' first principal

- FlashSystem.ai cannot perform any actions which can affect the availability of data, such as deleting system objects or performing any security-related tasks.

## Using FlashSystem.ai

- Administrator to take action
- At least 1 FlashSystem/SVC must be at 9.1.2
  - That FlashSystem hosts the FlashSystem.ai interface
  - Additional FlashSystem/SVCs can be at 9.1.0 and above (they can have Administrator or lower rights such as Monitor if no action is expected)
  - Participating FlashSystems can be stand-alone or part of an FlashSystem grid.
- There is a simplified experience for clients will be using FlashSystem.ai with a FlashSystem grid, as they are already a set of established trusted systems
- *Note: Storage Insights Pro is not a pre-requiste for FlashSystem.ai, but Storage Insights connection is a requirement*

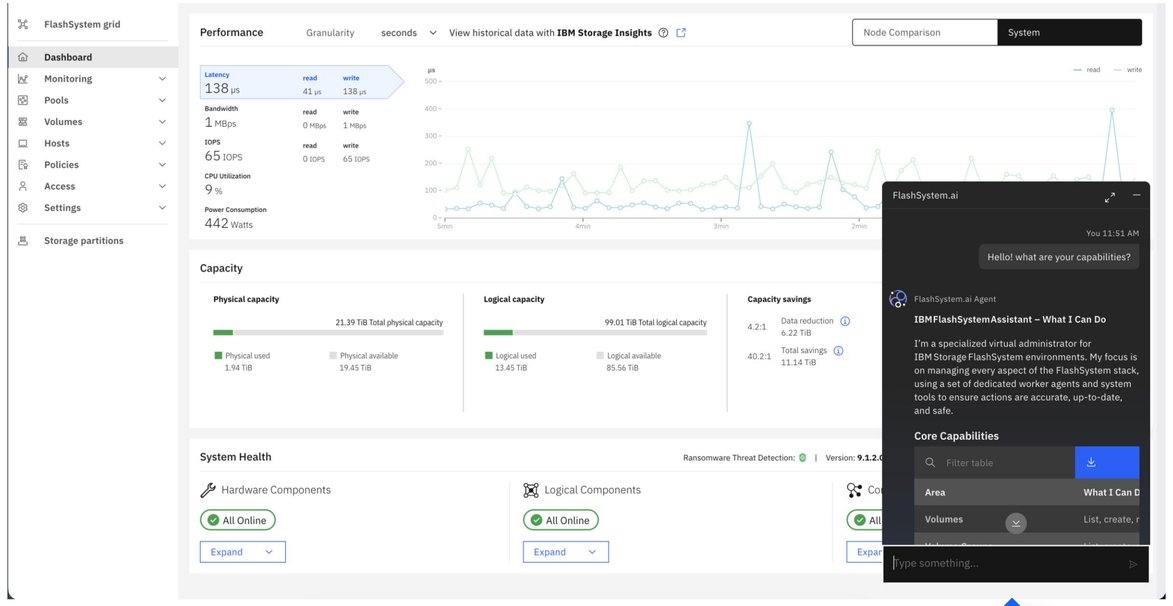


Embedded FlashSystem.ai discussion window

IBM