

FELHÚZNI-FELHÚZNI!

..avagy hogyan lett TOP GUN
az IBM FCM 5. generációja és hagyta állva
a legújabb FlashSystem termékcsalád a konkurenciát

—
IBM FlashSystem Champions – Q1 Kickoff
2026.03.25. – Arrow ECS Kft.

Seres Gábor

Storage Ecosystem Sales

gabor.seres@ibm.com



We have come a
long way in 70
years...

*From Spinning Rust (Fe_2O_3) to
numerous
Silicon Compounds (Si_3N_4) (SiO_2)*



Sept 13th - 1956 - IBM unveils the RAMAC-305
which includes the world's first ever random access storage device



The drive unit itself had a footprint of 1.5m^2 , housing a stack of $50 \times 24''$ platters with a single read/write head.



Capacity of
 $\sim 3.75\text{MB}^*$ in
today's
8bit/byte world

** Often mis-reported as 5MB, the 350 was capable of storing 5 million characters, where each character needed just 6bits + 1 parity*

Weighing over a ton, and only available to lease, for \$3200 per month...

Equivalent to ~\$38,000 per month today!



In the 70 years since,
we've never stopped
leading the way...



350 - First ever HDD - started the industry





1962 - IBM 1311 - First removable disk pack, like tape drives/media





**1965 - IBM 2314 - First disk storage “facility”
- the birth of the storage controller system**






**1971 - First ever FDD - 8" floppy disk -
enabled the birth of the software industry**






**1973 - IBM 3340 - First sealed HDD - as used
to this day - “Winchester” drive**







**1980 - IBM 3380 - First HDD to break
the 1GB barrier - 2.52GB**






1990 - IBM 9345 - First HDD to use MR heads, and separate read/write heads
1997 - First HDD range to use GMR heads, within 3 years 100% of HDD used GMR





**2003 - IBM 2145 - First Software
defined Storage Appliance - SVC**





2018 - IBM FCM - First enterprise computational storage drive



IBM's computational flash drives



Abilities	FCM	SSD
Extensive Built-in Encryption	✓	✗
Extensive Built-in Compression	✓	?
Extensive Health Binning	✓	✗
Extensive Heat Segregation	✓	✗
Variable Voltage	✓	✗
Variable Stripe RAID (Intra Module RAID)	✓	✗
~70µs latency	✓	✗
Ransomware Threat Detection	✓	✗



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



Efficiency

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

Protection

- Highest in class ECC encoding
- Zero impact inline hardware quantum safe 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!*

2026...

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*


EDSFF = Enterprise and Datacenter Standard Form Factor



~76mm x ~7.5mm x ~112mm

Comparing FCM5 and FCM1-4

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




FCM5

FCM4

FCM5



FCM4



FCM4

FCM5

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
- Hybrid use of ML-KEM for secure key transport of unlock PIN transmitted by FlashSystems storage controller to FCM
- Hybrid use of ML-DSA signatures for verification of firmware authenticity

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*



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...


FlashSystem Deduplication with FCM5

The design principles/goals

- Provide True Data Reduction value to Customers with FCM Hardware Compression and Deduplication
 - *Lessons learnt from DRP - need to think differently so allow pervasive “always on” approach*
- Spectrum Virtualize and FCM Co-design
- No additional appliance cost increase for computation
- No additional memory cost in controller for metadata
- Scalable at any capacity point
- Little to No performance impact



Deduplication *Summary*



A modern optimised approach to Deduplication lets FCMs provide full Enterprise Data Reduction without the cost and performance loss of competitors “*Brute Force*” methods!


- 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, **no added Write Latency**
- 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**



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



Raw Capacity	6.60 TB	13.20 TB	26.40 TB	52.80 TB	105.60 TB	+ 3XL soon... 1PBe!
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*	

With the new FCM5
E3 form factor...
erm...

*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)



SA2 Storage Engine



SV3 Storage Engine

FlashSystem 5015 & 5045



IBM FlashSystem 2026

IBM Confidential until Announce on 10th February 2026



FlashSystem C200

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



FlashSystem 5600

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



FlashSystem 7600

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



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



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)



SA2 Storage Engine




SV3 Storage Engine

FlashSystem 5015 & 5045



5600 has truly stunning density and sharable features


Clients benefit from the density, speed and abilities, while simultaneously having the ability to share all the features with their existing 3rd-party storage



Capacity Point	Small	Medium	Large	XL
FCM5 Raw Capacity	6.6 TB	13.2 TB	26.4 TB	52.8 TB
FCM5 Physical Capacity	5.6TB	11.2 TB	22.4 TB	44.8 TB
FCM5 Max Logical “Effective” Capacity	33.6 TB	67.2 TB	134.4 TB	268.8 TB
X12 FCM5 physical w/ RAID	50.4 TB	100.8 TB	201.6 TB	403.2 TB
X12 FCM5 effective w/ RAID	302.4 TB	604.8 TB	1209.6 TB	2419.2 TB



FlashSystem 5600 adapter configs



- 4 Port 32Gb Fibre Channel
- 2 Port 64Gb Fibre Channel
- 4 Port 25Gb Ethernet
- 2 Port 100 Gb Ethernet
- 2 Port SAS Expansion

Plug in rules – When mixed, Fibre Channel is added from left (slot 1) to right and Ethernet from right (slot 2) to left

SAS Expansion can only be applied in Slot 2

Canisters must have identical card configurations

Slot	Adapter	Speed	Max Ports Per Enclosure	Connector	Supported Protocol	Function
On-Board	Flexible IO	1Gb 10Gb 25Gb	4	RJ45 SFP+ SFP28	- iSCSI - NVME/TCP	- IP Replication long distance with TCP/IP - Host I/O - Optional Secondary Management Port
On-Board	Fixed Ports	1Gb	4	RJ45	- Management - Technician Port	Management Functions
1, 2	Fibre Channel	32Gb	16	SFP+	- SCSI - FC-NVMe	- Host I/O - Replication and HA - Migration - Virtualized Storage Attach
1,2	Fibre Channel	64Gb	8	SFP+	- SCSI - FC-NVMe	- Host I/O - Replication and HA - Migration - Virtualized Storage Attach
1,2	Ethernet	10Gb 25Gb	16	SFP+ SFP28	- iSCSI - NVMe/TCP	- Host I/O - IP Replication long distance with TCP/IP
1,2	Ethernet	40Gb 100Gb*	8	QSFP+ QSFP28	- iSCSI - NVMe/TCP - iWARP (rep RDMA)	- Host I/O - IP Replication long distance with TCP/IP - IP Replication short distance with RDMA
2	SAS	12Gb	4	Mini SAS-HD	SCSI	SAS Expansion Attach

* denotes oversubscription



Two 2T-E3 slots for potential future use-cases	32x 1T-E3 slots for FCM5s or for small capacity (1.9TB 3.84TB, 7.68TB) industry-standard commodity NVMe E3 SSDs 16x 2T-E3 compatibility from day1!	Two 2T-E3 slots for potential future use-cases
---	--	---



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



Two 2T-E3 slots for potential future use-cases	32x 1T-E3 slots for FCM5s or for small capacity (1.9TB 3.84TB, 7.68TB) industry-standard commodity NVMe E3 SSDs 16x 2T-E3 compatibility from day1!	Two 2T-E3 slots for potential future use-cases
---	--	---



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



Capacity Points in TB	Small	Medium	Large	XL	2XL	~4.88:1
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 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 S, M, L, XL, 2XL 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

All NVMe FlashSystem arrays as of the first quarter of 2026




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 ¹	32 drive slots ¹	32 drive slots ¹
Usable / Effective	1.1PBu / 2.3 PBe	403TBu / 2.42 PBe ²	1.22PBu / 7.29 PBe	2.43PBu / 11.88 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 Cache Hit IOPs:	1.5M	2.6M	4.33M	6.3M
Max 70/30/50 IOPs 16KB		0.7M	1.2M	3.00M
Max Read Miss Bandwidth:	23 GB/s	30GB/s	55 GB/s	86 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 12 – 64Gb FC ports	32 – 32Gb FC ports 24 – 64Gb FC ports	32 - 32Gb FC ports 24 – 64Gb FC ports
Mac Ethernet ports/system	8 -10BgE on-board 8 – 25GbE ports	4 - 25GbE on-board 16 – 25GbE ports	32 – 25GbE ports	32 – 25GbE ports

FlashSystem DRAID array configurations and enhancements

FlashSystem 5600	FlashSystem 7600	FlashSystem 9600
DRAID-1	DRAID-1	DRAID-1
Max drive size FCM 5-L, limit between 2 and 12 drives, SAS is 2-16 drives	NA	NA
DRAID-5	DRAID-5	DRAID-5
Max drive size FCM 5-L, limit between 4 and 8 drives	NA	NA
DRAID-6	DRAID-6	DRAID-6
Limit between 6 and 12 FCM 5 drives	Limit of between 6 and 32 drives per enclosure	Limit of between 6 and 32 drives per enclosure
	Minimum of 9 FCM XL drives	- Minimum of 9 FCM XL drives - Minimum of 13 FCM XXL drives
Max distributed arrays - 8	Max distributed arrays - 4	Max distributed arrays - 4
Homogeneous drive type/capacity per enclosure	Homogeneous drive type/capacity per enclosure	Homogeneous drive type/capacity per enclosure
Min default extent size = 2GB	Min default extent size = 8GB	Min default extent size = 8GB

Fast Drive Rebuilds
IBM's extent awareness technology gives the IO stack knowledge of where the client has written data and uses this knowledge to dramatically reduce rebuild times. This is important because FCM5 drives have the ability to roughly double effective capacity and the impact to rebuild time is proportional, but now the array will only rebuild where client data has been written. Therefore, even extents allocated to volumes can be skipped if unwritten.

No comment 😊



FlashSystem is value added with IBM Storage Insights Pro

Fleet-wide **observability**

See **what needs attention** with a **prioritized** view of the health and security of the devices



Cyber **Resilience**

Early warning of potential **ransomware** attacks and **integration** with the security tooling. It has inbuilt proactive Security Posture Scanning

AI Driven **advisories**

Get advice specific to the environment on **best practices**, known issues and upgrades



Capacity **Planning**

Save money and avoid surprises with **lifecycle alerts** and intelligent capacity planning

Intelligent **Alerting**

Proactive intelligence that detects, prioritizes, and alerts on storage **anomalies** before they **impact** performance.



Proactive **support**

AI assisted engagement with IBM support for an earlier diagnosis and quicker time to resolution


IBM Storage Insights Pro is the observability and operations platform for IBM flash storage infrastructure

Introducing the new storage framework FlashSystem 2026!

The catalyst to help move from yesterday's passive and reactive storage to tomorrow's **AI-Driven, Dynamic** Storage


The old traditional approach

A collection of *passive* data repositories



The new storage framework

An *integrated, active* storage solution



Looking back at yesterday's approach

Looking forward to tomorrow's solution



Data Contracts

Embedding business intent into your storage infrastructure


A Data Contract is a mechanism for establishing your storage requirements in IBM FlashSystem. Instead of manually adjusting storage settings, your team can **simply state business needs in plain language, and the system automatically aligns performance, protection, and behavior** to match.



Create Contracts Easily

Start from scratch or use **best-practice templates** that you can easily adjust to suit your needs.

Simplify storage management.



Apply Contracts Automatically

IBM FlashSystem.ai can **automatically apply changes and execute corrective actions** mandated by your Data Contracts.

Reduce manual intervention.



Enforce Contracts Continuously

IBM FlashSystem.ai ensures Data Contracts are enforced **consistently. Administrators are notified** whenever issues are detected.


Improve compliance & integrity.

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**. Your storage now is **aware** of regulatory requirements by **Geo-localization**




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

FlashSystem.ai is the storage admin 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 the complex storage estate and *execute* routine tasks.



Enhance Performance and Compliance

Deliver consistent and *reliable* performance and *compliance*.

FlashSystem.ai



Designed to Adapt, Engineered to Perform, Built to Protect

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.

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.

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.

IBM FlashSystem
1Q'26 Launch
10th February 2026

Köszönöm a megtisztelő figyelmet!



Seres Gábor

Storage Ecosystem Sales

gabor.seres@ibm.com

+36 20 221 6928