

IBM Storage IBM FlashSystem

Adapt, Protect, Perform

10th February 2026

Andrew Martin
Support Architect for IBM Storage Virtualize



Introducing FlashSystem.ai




Context Aware
Security



Dynamic
Compliance



Amplify Human
Expertise



Proactive
Optimisation


What if you could
have a conversation
with your Storage?

*...and your Storage could
take action for you?*

Agentic AI

Key Features of an AI agent

- Performs tasks (goal oriented) defined by incoming user query, using “skills” - tools you give your model access to, such as:
- Feedback loop to refine results



An AI Agent connects an LLM to your real-time operations data to not just to answer questions, but to perform actions


Natural Language Administration

AI agents assist users in day-to-day administration tasks, including:

- Provisioning workloads by creating volumes, volume groups, hosts and snapshots.
- Enabling disaster recovery solutions by configuring replication and immutable snapshots.
- Conversational interactions to query and update the system configuration, such as expanding volumes, renaming objects and more.




Enterprise Level AI Safeguards



- Actions executed on FlashSystem devices are done with agents that are on-premise and not in the Cloud
- Configuration changes performed by AI agents must always be confirmed by a human first (Human-in-the-Loop)
- AI agents cannot perform any actions which can affect the availability of data, such as deleting system objects or performing any security-related tasks
- AI assistance is only enabled if deployed by an administrator


Integration with Storage Insights



- Storage Insights AI agents seamlessly integrate with on-premise agents to deliver a unified AI experience.
- Storage Insights provides workload placement recommendations directly to the on-premise agents to action.
- Storage Insights is mandatory for FlashSystem.ai.

FlashSystem.ai Setup

- A container hosted in the compute infrastructure* provides the AI capabilities for a FlashSystem or grid
- *No on-prem GPUs needed*
- *No watsonx subscription needed*
- FlashSystem.ai container and FlashSystems must be able to communicate to each other
- The container must be able to contact Storage Insights for outbound traffic only
- Container needs registering with each FlashSystem
- Chat interface is available to the FlashSystem via the Graphical User Interface (GUI)



**The functions provided by the container will be integrated into the FlashSystem itself at a future date.*

An MCP server is key to FlashSystem.ai

MCP = Model Context Protocol

Key aspects of MCP servers:

Standardization:

Functionality:

Use Cases:

Architecture:


What are the benefits of MCP servers:

Context & Accuracy:

Interoperability:


Security:

What is MCP ?



FlashSystem.ai Upcoming Use Cases

.... As seen from the Launch Keynote



Human-AI Collaboration

*Storage configuration
and setup*



Proactive Optimisation & Balancing

Performance evaluation

*Partition migration within a
FlashSystem grid*



Predictive Maintenance & Knowledge

*Notifications on issues in the
environment*



Security & Regulation Compliance

*Apply Safeguarded Snapshot
policies to volume groups*

2026...

A close-up photograph of a person's hand pointing at a server rack. The server rack is black and features a honeycomb-patterned ventilation grille. A small, dark rectangular label with the text 'FlashSystem' in white is affixed to the grille. The background is slightly blurred, showing more of the server rack and a person's arm on the right side.

FlashSystem

2026 Introducing FCM5

Performance

*Better power delivery
and signal integrity*

Cooling

*Optimised shapes allow
for more efficient airflow*

Scalability

*Supports future needs for
higher capacities, faster speeds*

*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...*



FCM4 vs FCM5


Longer / Thinner / Taller / Faster Form Factor

Designed for Solid State not Spinning Disks

New additional card capacity FCM5-XXL

20% more raw capacity across the range

PCIe Gen5 interface



FCM5

Latest 2Tb QLC NAND, same FlashCore Reliability

Increased :

- Raw Capacity
- Effective Capacity
- Performance
- Computational Offload
- Security
- Data Reduction



Capacity Points in TB	Small	Medium	Large	XL	2XL
Form Factor	1T	1T	1T	1T	1T
RAW Capacity	6.6	13.2	26.4	52.8	105.6
Physical Capacity (1:1)	5.6	11.2	22.4	44.8	89.6
Max Data Reduction	6:1	6:1	6:1	6:1	4.88:1*
Effective Capacity	33.6	67.2	134.4	268.8	438.0*
Platform Support	x600	x600	x600	x600	9600

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

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.](https://www.forbes.com/sites/stevemcdowell/2024/02/27/ibms-introduces-data-path-ransomware-detection-in-new-flashsystem/?sh=341ee60ede93)”

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!](https://www.forbes.com/sites/stevemcdowell/2024/02/27/ibms-introduces-data-path-ransomware-detection-in-new-flashsystem/?sh=341ee60ede93)***

FCM Security

IBM Post Quantum Computing (PQC) 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*



FCM5 doesn't stop there...

Why don't we extend the FCM
offloads to include **Deduplication**...

Another industry first innovation...

Deduplication - approaches

Conceptually deduplication is straightforward

Hash chunks of data, compare hashes to find duplicates, keep a list of duplicates that reference one copy.

Simplistic Approach




Brute-force - software implemented

Dumb, scan everything always, over and over...

Massive system drain, duplicated efforts

Intelligent Approach



Computational storage approach


Stable data chunks only, hardware assistance

Minimise system usage, do things once!

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


** NB DRAM in the FCM is really NV-RAM (MRAM)*

- 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



A modern intelligent & **optimised** approach to **Deduplication** lets FCMs provide full Enterprise Data Reduction **without the cost and performance loss** of competitors “*Brute Force*” traditional 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 in spare cycles
 - Controller only uses ~2.5GiB DRAM, consolidated single garbage collection using existing FCM smarts
- 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**



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 system CPU and memory hardware

No capacity lost with reference tables

No performance loss




Deduplication activation will be a phased introduction.
Initial GA will scan data, but will not enable deduplication processing in the FCM.

To thin or not to thin...


The world before thin-provisioning...

- Contiguous allocation
- 1:1 mapping logical to physical




Thin-provisioning...

- Non-contiguous allocation - consolidated free space
- Software abstraction layer
- LPT - Logical to Physical Table maps fine grain chunks
- Enabled software features :
 - EasyTier, Data Reduction etc



Enter the SSD...

- Non-contiguous allocation - consolidated free space
- Always on hardware abstraction to manage NAND
- LPT - Hardware LPT maps fine grain chunks
- Enabled on device hardware features :
 - Internal heat tiering, Data Reduction



An SSD is inherently a thin-provisioned DEVICE...

Traditional software based volume level abstraction is redundant!

So why did we have both until now?

Volume Thin-provisioning...

- Allows for “**over-allocation**” of capacity...
- Worked well for things like **EasyTier** where sub-LUN allocation of grains could be allocated to different physical devices



FCM (Gen1,2,3,4) Thin-provisioning...

- Required to manage the NAND blocks
- Required to **enable** compression / **data reduction**
- **Limited** until now, by **3:1** addressability maximum
 - i.e. only over-allocate by at most 3x
- EasyTier use cases in early days of FCMs still valid

FCM5 no longer necessitates volume level thin-provisioning!



FCM5 Thin-provisioning...

- Required to manage the NAND blocks
- Required to **enable compression / deduplication**
- Now can address **6:1**
 - Caters for almost all “**sensible**” over-allocation needs
- Allows **system level simplification** by removing volume level thin-provisioning, **only one LPT for all data-reduction needs!**
- **Unmap** handled in one place by FCM5 **where it matters!**

With the new FCM5
E3 form factor...

*Don't we need some
new enclosures?*



IBM FlashSystem 2025

IOPs - OLTP 16KB @ 1ms : GB/s Read Miss Rdm 256KB @ 1ms



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
>380K IOPS, <0.5ms, >30GB/s, 1PBe



FlashSystem 7300

2U balanced performance and value
workhorse system
>500K IOPS, <0.5ms, >40GB, 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



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
>380K IOPS, <0.5ms, >30GB/s, 2.4PBe



FlashSystem 7600

2U balanced performance and value
workhorse system
>700K IOPS, <0.5ms, >55GB/s, 7.2PBe



FlashSystem 9600

Extreme performance and scaling for
mission critical workloads
>1.7M IOPS, <0.5ms, >85GB/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



FlashSystem 9600



Two 2T-E3
for future use

32x 1T-E3 slots

(Only supports one drive type/capacity per enclosure)

FCM5 - all capacities up to **2XL**

OEM NVMe E3 SSD- small capacity (1.9TB 3.84TB, 7.68TB)

Future 16x 2T-E3 ready!

Two 2T-E3
for future use

FlashSystem 9600 - FCM5 Capacity Examples



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



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	~4.88:1
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	

2U FlashSystem 9600 vs 4U FlashSystem9500

Up to 400% better density per RU

9500 Facts :

Rack Space : **4RU**

Max capacity :

Physical usable : **1.58PB**

Effective max : **4.74PB**

Per RU :

394.8 TB physical

1.18 PB effective



9600 Facts :

Rack Space : **2RU**

Max capacity :

Physical usable : **2.43PB**

Effective max : **11.88PB**

Per RU :

1215 TB physical

5.94 PB effective



Comparison :

Rack Space : **50% reduction**

Max capacity :

Physical usable : **54% increase**


Effective max: **151% increase**

Per RU :

Physical usable : **208% increase**

Effective : **403% increase**

FlashSystem 9600 - Rear View




7600/9600 Enclosure System Overview

2U Enclosure Dimensions:

- 438.0(W) x 87.9(H) x 836.8(D) mm
D = 'Enclosure bezel' to 'rear faceplate'
- 6x 6056mm dual rotor fan module

PSU, (1+1) redundant: Custom formfactor, up to 3.6kw



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 4 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

FlashSystem 7600



Two 2T-E3
for future use

32x 1T-E3 slots

(Only supports one drive type/capacity per enclosure)

FCM5 - all capacities up to **XL**

OEM NVMe E3 SSD- small capacity (1.9TB 3.84TB, 7.68TB)

Future 16x 2T-E3 ready!

Two 2T-E3
for future use

FlashSystem 5600

New EDSFF enclosure, the same node canisters as existing 5300



15 drive slots with
3 currently reserved
designated by [X]

10/25Gb I/O port
Optionally: secondary system IP
Internally: port 3

10/25Gb I/O port
Internally: port 2

USB port
Debug port

Canister LEDs

Technician port

Dedicated 1GbE Mgmt. port
Internally: port 1
Used for system IP and service IP

12x 1T-E3 slots
(Only supports one drive type/capacity per enclosure)

FCM5 - all capacities up to **XL**
OEM NVMe E3 SSD- small capacity
(1.9TB 3.84TB, 7.68TB)

FlashSystem 5600 - FCM5 Capacity Examples

The existing node canister has only 12 drive connections, hence the three unused slots...




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
x12 FCM5 physical w/RAID	50.4	100.8	201.6	403.2
x12 FCM5 effective w/ RAID	302.4	604.8	1209.6	2419.2

FlashSystem x600

Platform Hardware Summary


FlashSystem 5600

- Replacement for the 5300
- 1U 12 NVMe drive enclosure
- 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
- New EDSFF enclosure using existing 5300 canisters
 - 15 total slots, 3 are reserved
- 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



Spec-sheet Summary - FlashSystem NVMe 2026

	FlashSystem grid			
	 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
Physical / Effective	1.1PB / 2.3 PBe	403TB / 2.42 PBe	1.22PB / 7.29 PBe	2.43PBu / 11.88 PBe
Achieved with	FCM 4 – 46 TB (XL)	FCM 5 – 52.8 TB (XL)	FCM 5 – 52.8TB (XL)	FCM 5 – 105.6 TB (XXL)
CPU info per system	4x8 core	2x12 core Ice Lake D	2x16 core AMD Epyc	2x48 core AMD Epyc
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	6.37M
Max 70/30/50 IOPs 16KB	-	388K	700K	1.75M
Max Read Miss Bandwidth	23GB/s	30GB/s	60GB/s	86GB/s
Max HBA slot per system	4	4	8	8
Max FC ports/system	16 – 32Gb FC ports	16 - 32Gb FC ports 8 - 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

Performance Highlights - 2026 FlashSystem

All measurements taken at a maximum of 1ms latency

Metric	Parameters	FS5300	FS5600	Gain	FS7300	FS7600	Gain	FS9500	FS9600	Gain	Gain per RU
Interface		16x 32Gb FC	16x 32Gb FC		12x 32Gb FC	12x 32Gb FC		12x 64Gb FC	12x 64Gb FC		
Drive Count	Type	12x FCM4-S	12x FCM5-S		24x FCM3-S	32x FCM5-S		40x FCM4-S	32x FCM5-L	-8	8
Max Hit IOPS	4KB Read Cache Hit	2.5M	2.6M	4%	3.3M	4.3M	30%	4.2M	6.3M	50%	214%
Max Miss IOPS	4KB Rdm Read Miss <i>from Flash</i>	650K	700K	8%	950K	1.2M	26%	2.4M	3.0M	25%	150%

“Hero numbers” - the gains from previous generation to new ...
The 5300 and 5600 use the same node canisters - any gains are from HBA or FCM efficiency
The 9500 supported 48x drives, the 9600 32x drives - therefore the “gain per U” is a better measure.

Performance Highlights - 2026 FlashSystem

All measurements taken at a maximum of 1ms latency

Metric	Parameters	FS5300	FS5600	Gain	FS7300	FS7600	Gain	FS9500	FS9600	Gain	Gain per RU
Interface		16x 32Gb FC	16x 32Gb FC		12x 32Gb FC	12x 32Gb FC		12x 64Gb FC	12x 64Gb FC		
Drive Count	Type	12x FCM4-S	12x FCM5-S		24x FCM3-S	32x FCM5-S		40x FCM4-S	32x FCM5-L	-8	8
Max Hit IOPS	4KB Read Cache Hit	2.5M	2.6M	4%	3.3M	4.3M	30%	4.2M	6.3M	50%	214%
Max Miss IOPS	4KB Rdm Read Miss from Flash	650K	700K	8%	950K	1.2M	26%	2.4M	3.0M	25%	150%
OLTP IOPs	Measured at 1ms	388K	388K	-	530K	700K	32%	1.5M	1.75M	17%	133%

The hero numbers are one thing, but these are the most real-life from a synthetic test - OLTP IOPs. This simulates a transaction processing workload and is a consistent comparison measurement.

Performance Highlights - 2026 FlashSystem

All measurements taken at a maximum of 1ms latency

Metric	Parameters	FS5300	FS5600	Gain	FS7300	FS7600	Gain	FS9500	FS9600	Gain	Gain per RU
Interface		16x 32Gb FC	16x 32Gb FC		12x 32Gb FC	12x 32Gb FC		12x 64Gb FC	12x 64Gb FC		
Drive Count	Type	12x FCM4-S	12x FCM5-S		24x FCM3-S	32x FCM5-S		40x FCM4-S	32x FCM5-L	-8	8
Max Hit IOPS	4KB Read Cache Hit	2.5M	2.6M	4%	3.3M	4.3M	30%	4.2M	6.3M	50%	214%
Max Miss IOPS	4KB Rdm Read Miss <i>from Flash</i>	650K	700K	8%	950K	1.2M	26%	2.4M	3.0M	25%	150%
OLTP IOPs	Measured at 1ms	388K	388K	-	530K	700K	32%	1.5M	1.75M	17%	133%
Interface		16x 32Gb FC	16x 32Gb FC		12x 32Gb FC	12x 32Gb FC		24x 32Gb FC*	12x 64Gb FC		
Drive Count	Type	12x FCM4-S	12x FCM5-S		24x FCM3-S	32x FCM5-S		48x FCM4-L	32x FCM5-L	-16	8
Max Bandwidth	256KB Rdm Read Miss <i>from Flash</i>	30 GB/s	30 GB/s	-	43GB/s	55GB/s	28%	97GB/s	86GB/s	-11%	77%
	Per rack U	30 GB/s	30 GB/s	-	21.5GB/s	27.5GB/s	28%	24.25GB/s	43GB/s	77%	

When discussing the bandwidth numbers for 9600, it is essential to understand the change in enclosure size (4U-2U), and the number of FCM that can be installed has reduced, yet capacity increased! Here the “number” is lower - however the bandwidth per RU and per drive has dramatically increased!

* The 9500 test uses more 32Gb ports to ensure controller saturation for an apples to apples test.

Automatic Zoning

Old

Add Host

NPIV Enabled

Because NPIV is enabled on this system, host traffic is only allowed over the storage system's virtual ports. Ensure that SAN zoning allows connectivity between virtual ports and the host.

Name

host1

Host Connections

Select Connection

Host Port (WWPN)

Select WWPNs

☐

100000109BA1CBA0

☐

100000109BA1CBA1

Rescan

Add WWPN

Host Type

Generic

Advanced

Cancel

Save

Since 8.7.2
9.1.2 adds
automatic zoning

Add Host

All the inputs are prefilled based on the selected server.

Server name

R49B_U30_ESXi8_host

Paths

2, from 1 fabrics

Operating system

VMware_ESXi-8.0.0_(Releasebuild-20513097)

Logged in

Yes

Host name

R49B_U30_ESXi8_host

Select host ports (WWPNs)

1 × WWPNs selected

Rescan

Enter WWPN manually

Host type

Generic

Advanced

Cancel

Add

