New Gen of IBM FlashSystem x600

StorComp Technical & Market Analysis

What has changed?

This report provides an in-depth analysis of the new generation IBM FlashSystem portfolio (5600, 7600, 9600) launched in February 2026, together with the fifth generation IBM FlashCore Modules 5 (FCM5).
It goes beyond marketing claims to evaluate real technical changes, architectural trade-offs, performance characteristics, pricing trends, and practical deployment implications for enterprise customers.
Key topics include advantages and limitations of the new platforms, comparison with the previous generation and selected competitors, estimated pricing per configuration and per terabyte, as well as tips, tricks, and macroeconomic context (NAND and DRAM price spikes in 2026).

1. Purpose and Market Context

The new IBM FlashSystem generation consists of three NVMe models – 5600, 7600 and 9600 – all built on the new IBM FlashCore Modules 5 (FCM5).
FCM5 implements a computational-storage concept with embedded AI, hardware-accelerated data reduction, and security functions at the media level (VSR), moving more intelligence directly into the flash modules.
IBM's official communications emphasize higher density, improved system efficiency, and more autonomous management via FlashSystem.ai, with hardware-offloaded services such as compression, deduplication, and real-time ransomware detection running inside FCM5.

The objective of this report is to verify what truly lies behind these claims, separate solid innovation from marketing, and provide StorComp customers with a pragmatic view of how the new systems behave in real-world scenarios.

Q1 2026 saw DRAM contract prices surge 90-95% QoQ due to AI-driven server demand shortages, directly impacting all-flash array costs like IBM FlashSystem.
Enterprise SSD prices rose 53-58% as NAND production shifted to high-margin data center QLC, tightening supply.
Server DRAM faces the tightest shortages as vendors prioritize hyperscaler contracts over spot market.

IBM FlashSystem counters with new FlashCore modules (up to 105TB/disk) to mitigate NAND volatility.

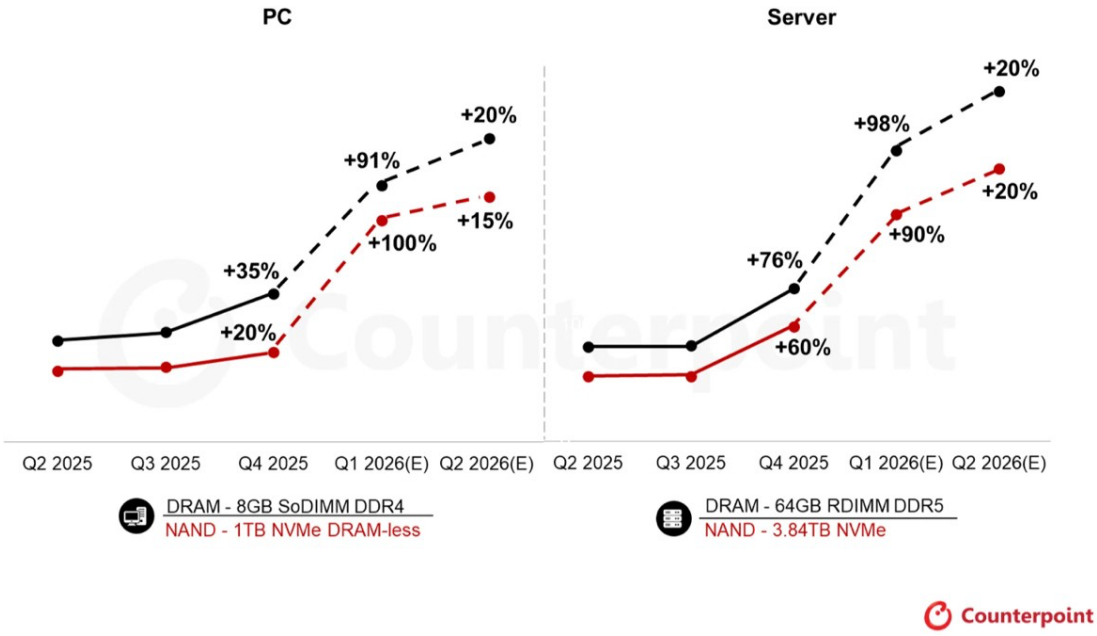

https://counterpointresearch.com/en/insights/Memory-Prices-Surge-Up-to-90-From-Q4-2025

2. Overview of the New FlashSystem Platforms

2.1 Architecture and Models

The new line replaces the previous FlashSystem 5300/7300/9500 with three models:

1. **FlashSystem 5600** – 1U chassis, 12 drive bays.
2. **FlashSystem 7600** – 2U chassis, 32 drive bays.
3. **FlashSystem 9600** – 2U chassis, 32 drive bays.

All systems use NVMe drives in the EDSFF (E3) form factor, which improves power and cooling characteristics and enables higher flash density per rack unit.

2.2 CPU Platforms and PCIe Generations

FlashSystem 5600 remains on Intel Xeon processors (2 × 12 cores) and PCIe Gen 4, positioning it in the economy+ / early-midrange segment.
FlashSystem 7600 and 9600 move to AMD EPYC CPUs and PCIe Gen 5, providing significantly more cores and higher I/O and memory bandwidth.

- **FS7600**: 2 × 16-core AMD EPYC NAPLES + PCIe Gen 5.
- **FS9600**: 2 × 48-core AMD EPYC NAPLES + PCIe Gen 5.

The transition from PCIe Gen 4 to Gen 5 effectively doubles bandwidth per lane, which is critical for fully exploiting the parallelism and throughput capabilities of FCM5-based configurations.

2.3 Key Technical Parameters – New vs Previous Generation

Official IBM spec sheets and documentation provide the following headline metrics:

Model	Max IOPS	Max Throughput	Max Raw Capacity
FS5600 (new)	2.7 M	42 GB/s	633 TB
FS5300 (previous)	2.0 M	36 GB/s	461 TB
FS7600 (new)	4.2 M	75 GB/s	1.69 PB
FS7300 (previous)	3.0 M	58 GB/s	1.23 PB
FS9600 (new)	6.3 M	86 GB/s	3.37 PB
FS9500 (previous)	6.0 M	100 GB/s	1.84 PB

Table 1: Comparison of key technical parameters for the new and previous IBM FlashSystem generations.

The high-end FS9600 nearly doubles raw capacity in 2U (from around 1.8 PB to approximately 3.37 PB) at the expense of fewer host ports (48 → 32) and a reduction in maximum read throughput (100 → 86 GB/s).
This reflects a deliberate trade-off toward higher density and efficiency, accepting a moderate drop in peak front-end bandwidth, which is often acceptable in environments where rack footprint and capacity per RU are the primary constraints.

3. IBM FlashCore Module 5.0 – Media Architecture

3.1 FCM5 Architecture and NAND Technology

FCM5 is the fifth generation of IBM's proprietary flash modules in NVMe EDSFF (E3) format, with capacities from 6.6 TB up to 105.6 TB per drive.
The largest 105.6 TB variant is currently supported only in FlashSystem 9600, enabling unprecedented raw density for a 2U system.

The key architectural change is a hybrid SLC-QLC design that aims to deliver TLC-class endurance and performance while achieving QLC-like cost and density, enabled by IBM's own algorithms for cell-voltage management and controlled write paths.
These capacities are achieved using 176-layer NAND and sophisticated wear-management mechanisms.

Historically IBM has been an early adopter of QLC in high-end arrays (maintaining DWPD = 1) even in DS8A0 mainframe storage, while many competitors are only now bringing their first QLC-based enterprise platforms to market.
The current surge in NAND prices further accelerates market adoption of high-density QLC solutions.

3.2 Hardware-Offloaded Data Services: Compression, Deduplication, Security

With FCM5, IBM moves even more functions from the controller level into the drive:

1. Hardware-accelerated compression.
2. Hardware encryption.
3. Hardware-assisted I/O pattern analysis.
4. A redesigned, hardware-offloaded deduplication engine.

All these services run **always-on** at the module level and cannot be disabled, which is a major departure from earlier designs where enabling Data Reduction Pool often had visible performance penalties (2-3 ms latency delay).
Thanks to this, the array can maintain high IOPS and low latency even for mixed workloads with high data-reduction potential (databases, VDI), addressing the performance issues previously seen with software-only deduplication in FlashSystem and Data Reduction Pools.

Capacity Points in TB	Small	Medium	Large	XL	2XL (only FS9600)
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 DRAID6	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

3.3 QLC Endurance, TLC-Class Performance and Rebuild Behaviour

IBM claims up to **5.5× more write cycles** compared with typical enterprise QLC SSDs, enabling FCM5 to support not only read-intensive but also write-intensive workloads.
Combined with the new data-layout scheme, this allows up to approximately 3.37 PB of raw capacity in 2U (FS9600) while still maintaining acceptable rebuild characteristics, though actual RTO/RPO will depend heavily on DRAID6 layout and pool sizing.

A practical configuration example is FlashSystem 5600 populated with 52.8 TB FCM5 drives.
In a 1U chassis this provides up to 633 TB of raw capacity without data reduction; with typical 5:1 data-reduction ratios, more than 2.5 PB of effective usable space can be achieved in suitable workloads.

IBM FlashSystem Grid functionality allows up to 32 arrays to be federated and managed through a single interface, enabling multi-PB scale with consistent data services.

4. Software Stack – Storage Virtualize 9.1.2

The new arrays ship with IBM Storage Virtualize 9.1.2, required to enable all new features on the platforms.
Downgrading from this code level is not supported, which is important for lifecycle and risk planning.

Version 9.1.2 is a non-LTS (non long-term support) release, meaning many customers will prefer to wait for the next LTS version 9.2.0 before they broadly deploy the new generation into mission-critical environments.

4.1 Deduplication Architecture Shift

A central change is moving deduplication from the controllers into the FCM5 modules.
Deduplication, compression, and encryption are now executed directly at the drive level, offloading CPU cycles from the controllers and making data-reduction efficiency less dependent on controller CPU utilization.
This should allow high data-reduction ratios with lower performance overhead compared with purely software-based schemes.

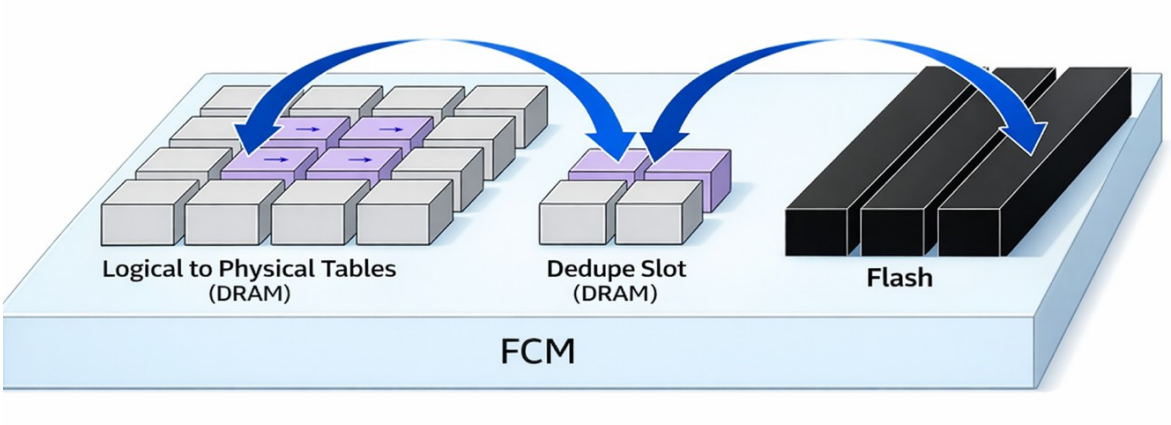

In practice, further clarification is still needed on the exact deduplication block size (for example 4 KB blocks, extents or segments), the placement of fingerprint tables, the deduplication scope (within pool, array or grid), and recommended workload profiles (VDI, VM, databases, backup, file) for optimal results.

4.2 5:1 Data-Reduction Commitment

IBM now offers a formal 5:1 Data Reduction Commitment, up from the previous 2:1 or 3:1 guarantees for suitable workloads.
The guarantee is tied to specific conditions such as data types, deduplication and compression settings, configuration best practices, and continuous monitoring through IBM Storage Insights.

4.3 Real-Time Ransomware Detection in FCM5

Another key differentiator is **Real-Time Ransomware Detection (RTD)**, which analyzes every I/O using a hardware-accelerated RISC module from AMD and machine-learning models, with minimal performance impact.

Core RTD characteristics include:

- Detection of ransomware attacks in **less than one minute**.
- Detection based on entropy patterns and abnormal I/O profiles.
- Correlation of events with Safeguarded Copy snapshots and IBM Storage Insights Pro telemetry.
- Rapid rollback to uncompromised WORM snapshots after an attack is confirmed.

Public IBM sources describe how anomaly events are propagated to Storage Insights, where snapshots are tagged with anomaly IDs, helping customers quickly identify clean recovery points.

5. Product Strengths and Limitations

5.1 Strengths

5.1.1 AMD EPYC NAPLES and PCIe Gen 5 in FS7600/FS9600

The move to AMD EPYC NAPLES and PCIe Gen 5 in FS7600 and FS9600 provides a tangible performance uplift in I/O bandwidth, CPU cores and memory throughput compared with earlier Intel/PCIe Gen 4 platforms.
Headline values of up to 6.3 million IOPS and 86 GB/s throughput can be maintained under mixed workloads, while leaving more CPU headroom for software features such as AI-driven placement and FlashSystem.ai analytics.

5.1.2 Higher FCM Capacities

FCM5 increases maximum drive capacity from 38.4 TB (FCM4) to 105.6 TB, nearly doubling maximum raw capacity in 2U from about 1.8 PB to 3.37 PB in FS9600.
This delivers higher density, lower power and cooling costs per TB, and enables consolidation of more workloads onto fewer arrays.

5.1.3 Higher (Mostly Marketing) Performance Claims

IBM advertises up to **40% better efficiency** and higher IOPS compared with the previous generation, thanks to the combination of FCM5, PCIe Gen 5 and hardware-accelerated data-reduction engines.
At this stage most data points come from vendor collateral and lab tests but given the 5:1 data-reduction commitment and current NAND pricing, these improvements are relevant for TCO modelling.

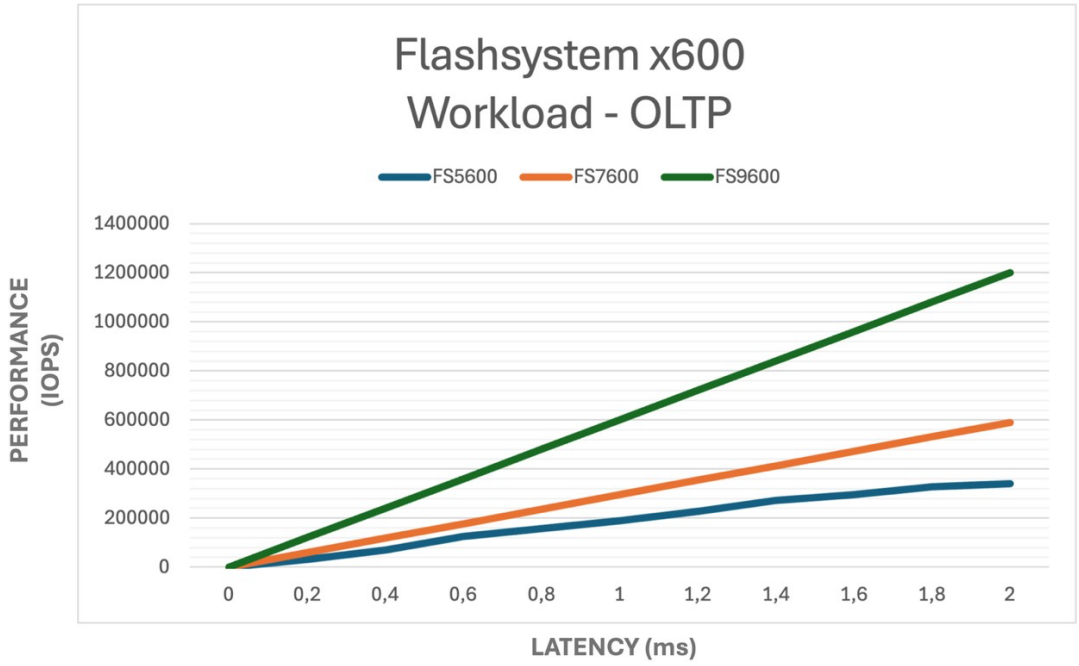

Chart 1. OLTP WORKLOAD (Block size – 16kB, workload profile 70/30 – Read/Write)

5.1.4 Evolution, Not Revolution – Architectural Continuity

The new models retain the familiar Storage Virtualize/SVC architecture and FlashSystem Grid federation concept, focusing on evolutionary upgrades: newer CPUs, FCM5 media, EDSFF form factor and FlashSystem.ai services.
For existing IBM customers this means smoother migration, re-use of existing tooling and integrations (vSphere, remote copy, backup integrations), and incremental improvements in performance, density, efficiency and cyber-resilience rather than disruptive change.

5.1.5 AI-Assisted Management – FlashSystem.ai

FlashSystem.ai bundles data-services and analytics intended to act as a co-administrator, helping with best-practice configuration, workload placement, diagnostics and recommendations, including visibility into some non-IBM arrays.
This can reduce the operational burden on storage admins and speed up anomaly detection and remediation, but also introduces an additional “black-box” layer whose recommendations must be understood and governed, especially in regulated environments.

5.2 Weaknesses and Risks

5.2.1 Initial Code Release is Non-LTS (9.1.2)

Launching the new hardware on a non-LTS code level typically implies higher risk of bugs and regressions, as well as a lack of long-term best-practice guidance.

A cautious rollout strategy is recommended, starting with test and development environments before migrating tier-1 production workloads.

5.2.2 Deduplication Initially in Scan-Only Mode

According to the 9.1.2.0 release notes, deduplication initially operates only in scan-only mode and does not yet provide full data-reduction benefits.
Data-reduction via dedupe will be enabled in a later software update (v9.1.3), which means early adopters may temporarily see lower effective data-reduction ratios than the headline 5:1 marketing figures suggest.

5.2.3 FS5600 as a Mostly Cosmetic Upgrade

In contrast to 7600/9600, FS5600 remains on PCIe Gen 4 and Intel Xeon with only 12 drives in 1U, so the change versus FS5300 is largely incremental rather than transformative.
For customers expecting a "small 7600" with PCIe Gen 5, FS5600 may be disappointing – the real architectural jump occurs at 7600 and 9600.

5.2.4 Lack of Native File and Object Protocols

FlashSystem still does not offer native file (NFS/SMB) or object (S3) protocols on the same controllers, while several AFA competitors already deliver unified block + file + object on a single platform.
This forces customers to add external gateways or SDS layers for file/object services, increasing complexity and total cost of ownership.

5.2.5 No NVMe-oF Expansion Shelves (SAS Only in FS5600)

If 7600/9600 remain limited to internal EDSFF media and only FS5600 offers SAS-attached expansion shelves, this constrains the ability to scale capacity in a modular way while keeping end-to-end NVMe.

For some designs this means moving more quickly toward federated grids of arrays rather than scaling a single system, which aligns with IBM's grid philosophy but differs from competitors that provide NVMe-oF-attached expansion shelves.

6. Practical Tips & Behavior of the New Generation

6.1 New Volumes: Thin + Data-Reduced by Default, No DRP

If new volumes are provisioned by default as thin-provisioned and data-reduced (compression + deduplication) without explicit Data Reduction Pool selection, administrators must adjust sizing and monitoring habits.
Key impacts include the relationship between logical over-provisioning and physical footprint, reserve planning, and interpretation of capacity alerts where a "full" logical volume does not necessarily mean immediate exhaustion of physical space.

6.2 Host Auto-Zoning (Brocade Gen 8 Only)

Host auto-zoning, currently limited to Brocade Gen 8 switches, can simplify provisioning in dynamic SAN environments by automatically creating zones based on detected WWPNs and LUN mappings. In tightly segmented SANs this automation must be carefully governed, with documented exceptions, rollback procedures and alignment to existing naming and labelling conventions.

6.3 Compliance Mode and Auto-Updates

Compliance mode typically enforces stricter security policies for passwords, TLS, auditing, downgrade blocking and mandatory encryption.
Combined with auto-updates, customers must clearly define the code lifecycle process, maintenance-window controls and interactions with regulatory requirements (for example ISO, NIS2, sector-specific regulations).

6.4 Lack of DRAID1 and DRAID5 on 7600/9600

If high-end models support only a subset of DRAID levels (for example DRAID6 only), configuration flexibility for performance vs capacity vs rebuild-time trade-offs is reduced.
Given individual drive sizes of 26–105 TB, the minimum pool size and stripe width are crucial for rebuild times and overall risk assessment.

7. Pricing and Market Trends in 2026

7.1 NAND and DRAM Price Shock in 2026

The year 2026 has brought an unprecedented surge in NAND flash and DRAM prices, driven by AI data-center demand, the transition to DDR5, and cautious capacity investments following years of volatility.

According to TrendForce and other industry reports, NAND prices rose by 33–38% quarter-on-quarter in Q4 2025 and a further 55–60% in Q1 2026, more than doubling overall from late Q3 2025 levels.

Samsung reported about a 100% increase in contract NAND pricing in Q1 2026, directly impacting SSD costs for end customers.

For DRAM, the situation is even more extreme: DDR4/DDR5 prices increased by 38–43% in Q4 2025 and an additional 105–110% in Q1 2026, with server DRAM up over 60% quarter-on-quarter in Q1 2026.

The main drivers include aggressive capacity reservations by hyperscale CSPs (AWS, Azure, Google Cloud) for AI workloads, the DDR4→DDR5 migration reducing supply of older technologies, vendor prioritization of high-margin server applications over consumer/enterprise, and limited new-fab investments from 2023–2025.

Analysts expect elevated pricing at least through the end of 2026, with possible extension into 2027 as Big Tech pre-books 2027 NAND allocations.

For enterprise storage vendors this translates into higher TCO for high-capacity flash configurations, making strong data reduction and QLC-based architectures economically critical rather than merely “nice to have”.

7.2 Indicative Pricing for FlashSystem 5600/7600/9600

IBM's public marketplace provides indicative pricing for FS5600 and FS7600 at different effective capacities, with price per TB decreasing at higher capacities.

FlashSystem 5600 (effective capacity, May 2026):

Raw Capacity	Effective Capacity	Price per raw TB	Segment
52 TB	~80 TB	2 100 USD	Entry, small workloads
158 TB	~300 TB	1 200 USD	Midrange, production
422 TB	~700 TB	850 USD	Midrange, production, high-capacity

Table 2: Approximate pricing for IBM FlashSystem 5600 with 5 years Expert Care Advanced.

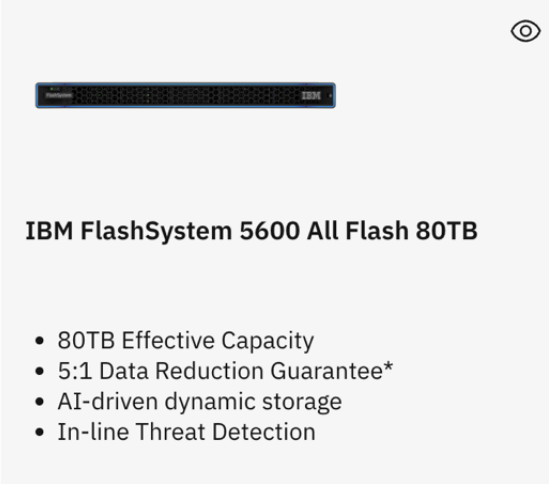

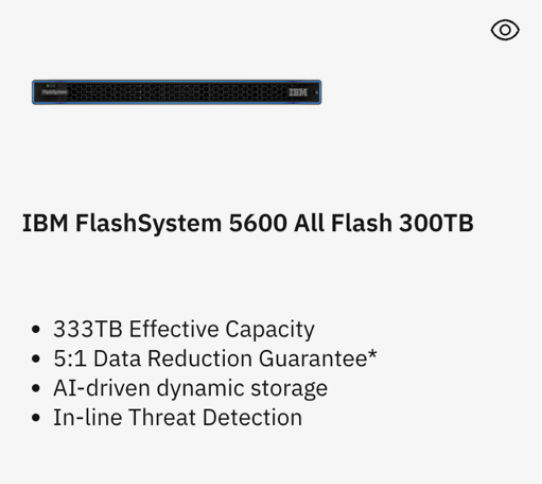

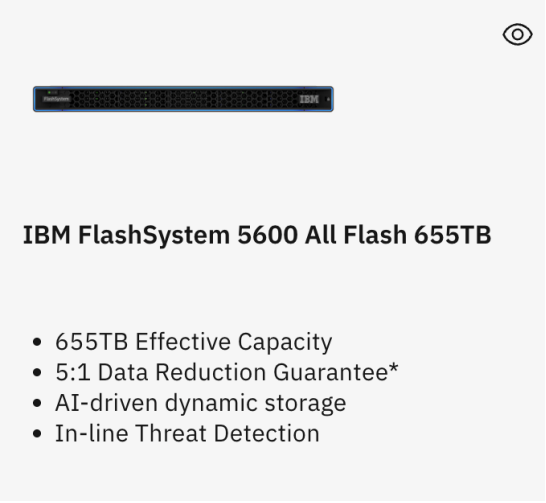

FlashSystem 7600 (effective capacity, May 2026):

Raw Capacity	Effective Capacity	Price per TB	Segment
52 TB	~80 TB	4 400 USD	Low-midrange, production
158 TB	~300 TB	2 000 USD	Standard, production
317 TB	~600 TB	1 300 USD	Standard, production, high-capacity

Table 3: Approximate pricing for IBM FlashSystem 7600 with 5 years Expert Care Advanced.

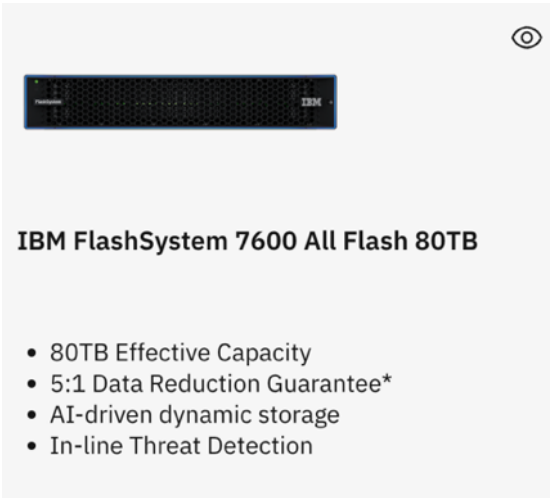

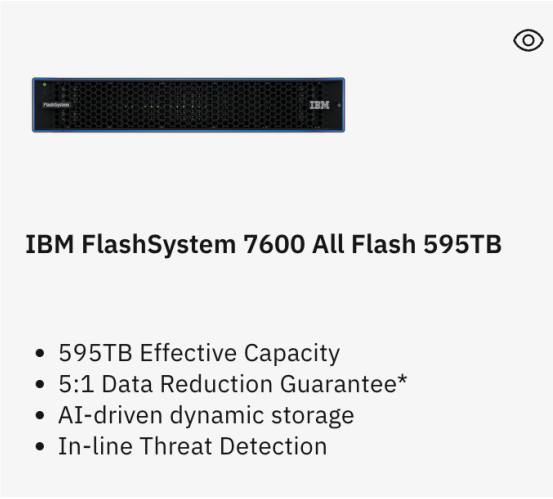

For **FlashSystem 9600** with 5 years Expert Care Advanced, IBM lists only a single reference configuration **of about 470 TB effective** and **238TB raw capacity** an average price around **3 500 USD per raw TB**.

Comparing similar TB levels across generations, **FlashSystem prices have increased by roughly 25–35%,** depending on configuration parameters such as cache size, number and capacity of drives, and port counts.

Given the high volatility of NAND pricing, customers must always re-validate current pricing at the time of project planning.
https://www.ibm.com/products/flashsystem/pricing/us-en?5300=

8. Product Ratings – How Much Does the Refresh Change?

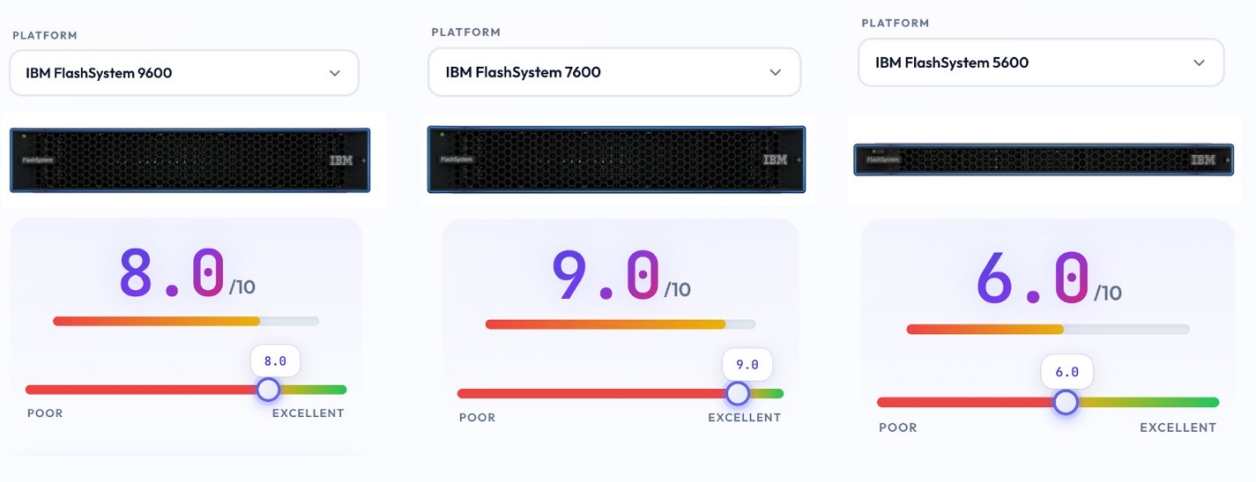

8.1 IBM FlashSystem 5600 – Rating 6/10

FS5600 receives an overall rating of 6/10, mainly because it represents an incremental rather than transformational update.

Retaining older-generation Intel CPUs, PCIe Gen 4 and a 12-drive 1U layout, even with EDSFF, leaves room for further improvement (3 slots are free).

Positive changes:

- Availability of new FCM5 media and associated performance/latency gains.
- Future benefit from hardware-accelerated deduplication once fully enabled in code.
- Higher data density from larger FCM capacities.
- SAS expansion shelves for SSD and HDD.

Limitations:

- No NVMe-oF shelves for expansion, limiting truly end-to-end NVMe scale-out.

8.2 IBM FlashSystem 7600 – Rating 9/10

The seventh-generation successor to the popular V7000 line, FS7600, earns **9/10**.
It gets substantially improved components versus its predecessor:

- AMD EPYC NAPLES CPUs (2x 16-core),
- PCIe Gen 5,
- More FC and iSCSI ports (max 32x FC 32Gbps),
- larger FCM5 capacities with compression and deduplication
- 32 instead of 24 modules – resulting in a nearly ideal evolutionary successor in the FlashSystem family.

- Max 1.5TB cache (same size as previously)

Areas for improvement:

- No SAS or NVMe-oF expansion shelves, although the absence of SAS ports may signal a future move toward NVMe-oF as the standard shelf-interconnect model.
- Four free physical slots hint that IBM may still extend maximum capacity in future refreshes.

8.3 IBM FlashSystem 9600 – Rating 8/10

FS9600, the fourth generation of the FS9000 high-end line, is rated **8/10**.

It adopts the same architectural improvements as FS7600 – dual 48-core AMD EPYC CPUs, PCIe Gen 5 and FCM5 modules – while moving from 4U to 2U without reducing maximum raw capacity.

The drop from 48 to 32 drive slots is offset by 2XL FCM5 modules of over 105 TB raw, preserving and even improving total raw capacity in 2U.

- AMD EPYC NAPLES CPUs (2x 48-core)
- PCIe Gen 5,
- Ports – FC and iSCSI ports (max 32x FC 32Gbps),
- larger FCM5 capacities with compression and deduplication (max 105TB drive size FCM5)
- 32 instead of 48 modules – resulting in a nearly ideal evolutionary successor in the FlashSystem family.

Notable drawbacks:

- No increase in maximum cache (still up to 3 TB).
- Fewer FC ports compared with FS9500 (max 32 instead of 48)
- No expansion-shelf ports, mirroring FS7600's limitations.

9. Outlook – What to Expect Next?

Several products remain untouched by this refresh and can be expected to evolve in coming cycles:

1. Entry SAS arrays FS5015 and FS5045.
2. FS C200, relatively new but still based on older-generation drives.
3. DS8A00 mainframe systems using FCM4 today, likely to transition to EDSFF and FCM5 in the future.
4. SAN Volume Controller (SVC) – future changes will likely focus on new CPUs, RAM type/quantity and NVMe port speeds rather than radical architecture shifts.
5. Looking at empty disk shelves (3 for FS5600 and 4 for FS7/9600) more drives are coming.

10. Conclusion

The new IBM FlashSystem generation delivers evolutionary yet meaningful improvements for enterprise flash storage buyers.
Density doubles at the high end (up to 3.37 PB in 2U), 7600/9600 benefit from AMD EPYC and PCIe Gen 5, and FCM5 brings hardware-accelerated data reduction and ransomware detection directly into the media layer.

Key advantages include high density, strong AI-assisted operations, 5:1 data-reduction commitments and advanced cyber-resilience capabilities, all of which become more important as NAND and DRAM prices have more than doubled between late 2025 and Q1 2026.
At the same time, there are clear limitations: FS5600 is only a cosmetic upgrade, deduplication is initially scan-only in 9.1.2, there are no native file/object services, and high-end models lack NVMe-oF expansion shelves and LTS code at launch.

For StorComp users, these systems represent strong candidates where high density, data-reduction efficiency and cyber-resilience are critical, but they should be evaluated carefully against competitors that already deliver unified file/object, mature NVMe-oF expansion architectures and fully GA LTS code streams.

Parameter	 FlashSystem 5600	 FlashSystem 7600	 FlashSystem 9600
Architecture	NVMe	NVMe	NVMe
Rack Units	1U	2U	2U
Drive slots per enclosure	12	32	32
Usable / Effective Capacity	400TBu / 2.4 PBe	1.1PBu / 7.2 PBe	2.3PBu / 11.9 PBe
Number of drives	FCM 5 – 52.8 TB (XL)	FCM 5 – 52.8TB (XL)	FCM 5 – 105.6 TB (XXL)
CPU per system	2x12 core (Intel Ice Lake)	2x16 core (AMD Epyc)	2x48 core (AMD Epyc)
High-speed bus	PCIe Gen 4	PCIe Gen 5	PCIe Gen 5
Max IOPS (4k Block size):	2,600,000	4,300,000	6,300,000
Max Read Miss Bandwidth:	22GB/s	60 GB/s	85.8 GB/s
Max HBA slot per system	4	8	8
Max FC ports/system	16-32Gb FC ports	32 – 32Gb FC ports	32 - 32Gb FC ports
Max Ethernet ports/system	4 - 25GbE on-board 16 – 25GbE ports	32 – 25GbE ports	32 – 25GbE ports

It's also worth remembering that IBM is one of the few providers that provides complete Flashsystems documentation, pricing information, and actual performance and capacity data. It also provides current Redbooks and best practices for system implementation.

Frankly, working with IBM documentation is truly enjoyable.

Explore More with StorComp

Liked this report and want to compare IBM arrays with competitors?

Sign in to **StorComp** for free and explore detailed technical parameters of the new IBM FlashSystem series, side-by-side with offerings from **HPE, Pure Storage, Hitachi Vantara, Dell, NetApp and Huawei** and more!

StorComp – vendor-neutral and reliable source of storage data

StorComp LLC
131 Continental Dr
Newark, Delaware
https://storcomp.com