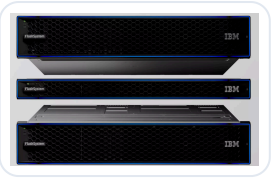
IBM August 2026

StorComp Analysis - IBM FlashSystem x600 (2026)



What changed in the 5600, 7600 and 9600 — architecture, code, pricing, and where each model fits.

TL;DR

7600 is the one to shortlist. 9600 when the job is density in 2U. 5600 only if 1U and PCIe Gen 4 are enough.

9/10

FlashSystem 7600

- The one to shortlist
- EPYC and PCIe Gen 5
- Up to 32 × 32 Gb FC
- 32 modules instead of 24
- FCM5 with reduction in hardware
- No expansion shelf
- Four slots still free

8/10

FlashSystem 9600

- Same platform as 7600, in 2U instead of 4U
- 105 TB modules cover the fewer drive slots
- Fewer FC ports than 9500
- No extra cache
- No expansion shelf

6/10

FlashSystem 5600

- New FCM5 and higher density
- SAS shelves for SSD and HDD
- Still Intel, PCIe Gen 4 and 12 drives
- Three slots left empty
- Incremental, not the new platform

What this briefing covers

- A StorComp reading of the February 2026 generation: 5600, 7600 and 9600, plus FlashCore Module 5
- Past the launch claims: architecture, the code these arrays ship with, indicative pricing, and where each model is a real step forward
- They replace FlashSystem 5300, 7300 and 9500
- All three use NVMe drives in the EDSFF (E3.S) form factor
- FCM5 is the media change: compression, encryption, deduplication and ransomware detection in the module itself

Three models, two platforms

The lineup is not one architecture in three sizes. FlashSystem 5600 stays on Intel Xeon (2 × 12 cores) and PCIe Gen 4, in a 1U chassis with 12 drive bays. FlashSystem 7600 and 9600 move to AMD EPYC and PCIe Gen 5, both in 2U with 32 bays. The report lists 7600 as 2 × 16 cores and 9600 as 2 × 48 cores.

PCIe Gen 5 doubles bandwidth per lane versus Gen 4. That is the difference between a controller that can actually feed FCM5 and one that cannot. On 5600 the new media arrives, but the platform around it does not. The architectural jump starts at 7600.

- 5600 — 1U, 12 bays, Intel Xeon, PCIe Gen 4. Economy and early midrange.
- 7600 — 2U, 32 bays, 2 × 16-core AMD EPYC, PCIe Gen 5. The practical successor to the V7000 line.
- 9600 — 2U, 32 bays, 2 × 48-core AMD EPYC, PCIe Gen 5. High-end density, down from a 4U 9500.

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

How the three models differ: chassis, CPU, bus, ports and the drive each one is built around.

Headline numbers, new versus previous

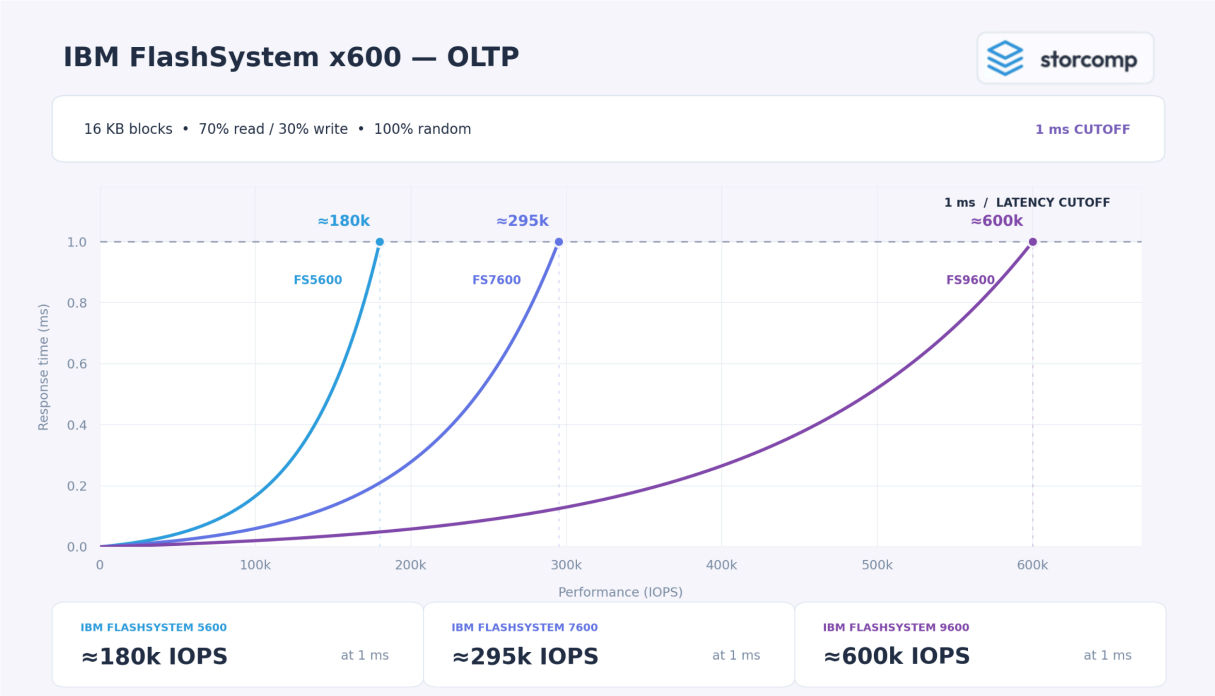
Official spec sheets put the new models next to the ones they replace. Read the 9600 row carefully: capacity almost doubles, peak read throughput does not.

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

Maximum IOPS, throughput and raw capacity. New models are highlighted.

FS9600 takes raw capacity in 2U from about 1.84 PB to about 3.37 PB. The cost of that density is fewer host ports (48 down to 32) and a lower maximum read throughput (100

GB/s down to 86 GB/s). That is a deliberate trade: more capacity per rack unit, less peak front-end bandwidth. It fits designs that are rack-bound. It does not fit a design that was already using every 9500 port.



OLTP, 16 KB, 70/30, 100% random, held at a 1 ms cutoff. About 180k IOPS on the 5600, 295k on the 7600 and 600k on the 9600. This is not a saturated maximum.

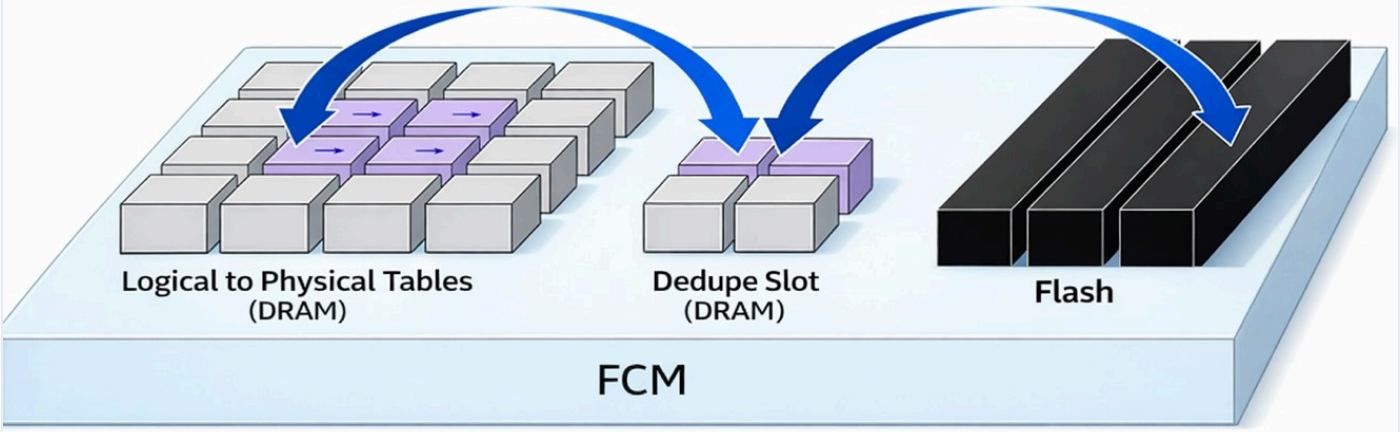
[Download chart (PNG)]

FlashCore Module 5

FCM5 is an NVMe EDSFF module from 6.6 TB up to 105.6 TB. The 105.6 TB drive is supported only in FlashSystem 9600. The design is hybrid SLC-QLC: IBM's claim is TLC-class endurance and performance at QLC-class density, using 176-layer NAND and its own cell-voltage and write-path management. IBM has used QLC in high-end arrays, including DS8000-class mainframe storage, while many competitors are only now shipping a first enterprise QLC platform. With NAND prices where they are in 2026, that density is the economic argument, not a side feature.

Compression, encryption, I/O pattern analysis and a redesigned deduplication engine run in the module, always on. They cannot be switched off. On earlier FlashSystem code, turning on a Data Reduction Pool often added about 2–3 ms. The point of FCM5 is to keep high IOPS and low latency on mixed workloads — databases, VDI — without that software penalty.

IBM also claims up to 5.5× more write cycles than typical enterprise QLC, so the modules are positioned for write-heavy work, not only read-mostly archives.



Deduplication now lives in the module: logical-to-physical tables and the dedupe slot sit in DRAM on the FCM, and only the reduced data is written to flash.

A full 5600 with 52.8 TB FCM5 drives holds up to 633 TB raw in 1U. At a 5:1 reduction ratio that is more than 2.5 PB effective, on workloads that actually reduce. FlashSystem Grid can federate up to 32 arrays behind one interface when a single chassis is not the scale point.

	Small	Medium	Large	XL	2XL
1× FCM5 raw	6.6	13.2	26.4	52.8	105.6
1× FCM5 physical	5.6	11.2	22.4	44.8	89.6
1× FCM5 effective	33.6	67.2	134.4	268.8	438.0
32× physical, DRAID6	151.9	303.8	607.6	1,215.2	2,430.4
32× effective, with RAID	911.4	1,822.8	3,645.6	7,291.2	11,880.8

FCM5 capacity points, in TB. Effective figures assume IBM's stated reduction. 2XL is 9600 only.

The code these arrays ship with

New systems ship on IBM Storage Virtualize 9.1.2. Downgrade is not supported. 9.1.2 is not a long-term support release. Many estates will wait for the 9.2.0 LTS stream before these

arrays carry tier-1 workloads. Treat the launch code as a reason to start in test and development, not as a reason to skip the hardware.

Deduplication has moved off the controllers and into FCM5, which is the right place for it if the controller was the bottleneck. Release notes for 9.1.2.0 say deduplication starts in scan-only mode. It does not yet return the full reduction. That benefit is scheduled for 9.1.3. Early systems can therefore land well short of the 5:1 figure used in the business case. Ask which code level the quote assumes, and what the array reports before dedupe is fully on.

IBM's data-reduction commitment moves from 2:1 or 3:1 to 5:1 for suitable workloads. It is conditional: data type, reduction settings, configuration practice, and ongoing monitoring in IBM Storage Insights. It is not a blanket promise on every volume.

Real-time ransomware detection in FCM5 inspects I/O in hardware, using an AMD RISC accelerator and machine-learning models. IBM's public description is detection in under a minute from entropy and abnormal I/O patterns, with events tied to Safeguarded Copy snapshots and Storage Insights Pro so a clean recovery point can be identified quickly. Worth testing. Worth not treating as a substitute for an isolated copy.

What is actually better

- 7600 and 9600 gain EPYC cores, PCIe Gen 5 and more memory bandwidth. Headline figures cited for the high end are 6.3 million IOPS and 86 GB/s, with CPU left for FlashSystem.ai and placement.
- Drive capacity goes from 38.4 TB on FCM4 to 105.6 TB on FCM5. In 2U that is the jump from about 1.8 PB raw to 3.37 PB, and fewer arrays for the same estate.
- The family is still Storage Virtualize. Existing copy services, vSphere integration and operational habits carry over. This is an evolution of the grid, not a new operating model.
- FlashSystem.ai is pitched as a co-administrator: configuration, placement, diagnostics, and some visibility beyond IBM arrays. Useful, and another layer whose recommendations need an owner in a regulated environment.

IBM's "up to 40% better efficiency" line is still mostly collateral and lab data. Combined with a 5:1 commitment and 2026 NAND prices, it belongs in a TCO model. It does not belong in a design as a measured result until the configuration, the code level and the workload are the ones you would buy.

What to treat carefully

- Launch code is non-LTS. Higher chance of bugs, thinner long-term guidance. Do not make 9.1.2 the first home of a tier-1 workload.
- Dedupe is scan-only at 9.1.2. The 5:1 story is not fully available on day one.
- 5600 is a modest step from 5300: same PCIe generation, Intel, 12 drives in 1U. It is not a small 7600.
- Still no native NFS, SMB or S3 on these controllers. File and object need a gateway or another platform. Several all-flash competitors already do block, file and object in one system.
- 7600 and 9600 do not take expansion shelves. Only 5600 has SAS expansion, and that is not end-to-end NVMe. Scale means another array in the grid, not a shelf on the same pair.
- 7600 and 9600 are described without DRAID1 and DRAID5. With 26–105 TB drives, pool width and rebuild time are the design, not a checkbox.

Notes for anyone who will run it

New volumes come up thin and data-reduced, without a separate Data Reduction Pool choice. Logical “full” is not the same as physical full. Capacity alerts, reserves and over-provisioning need a new reading, or the first incident will be a surprise.

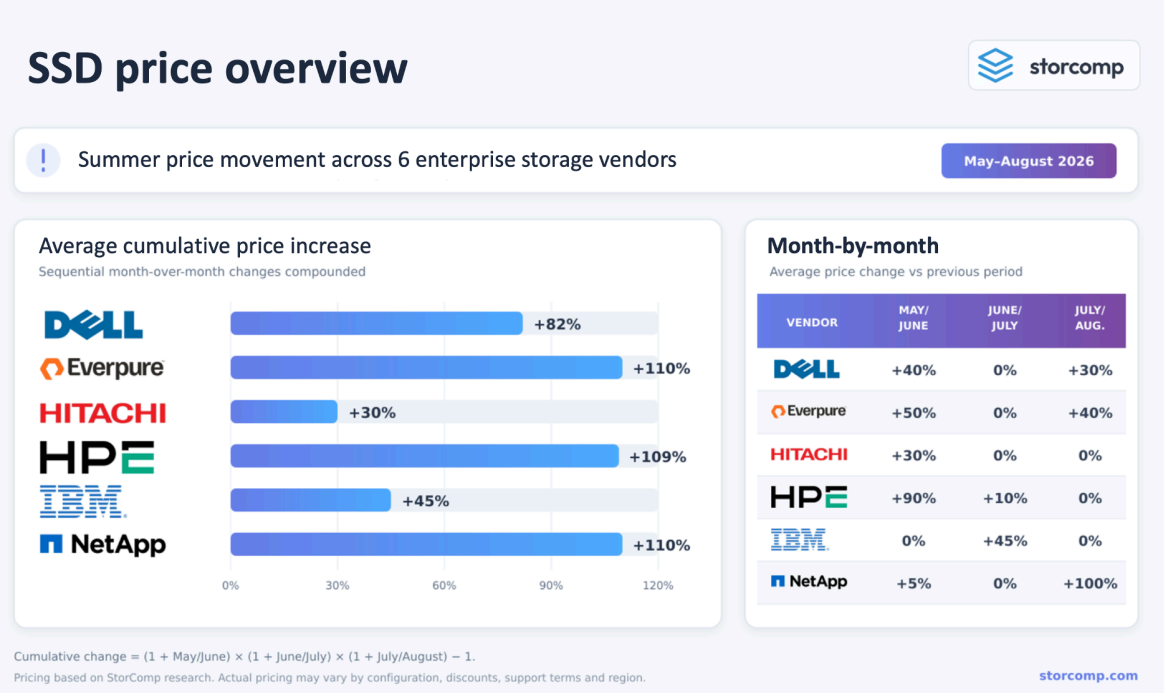
Host auto-zoning is limited to Brocade Gen 8. It can build zones from detected WWPNs and LUN mappings. In a tightly segmented SAN that automation needs an exception list, a rollback and the same naming rules the fabric already uses.

Compliance mode tightens passwords, TLS, audit, downgrade blocking and encryption. Together with automatic code updates it has to sit inside a real maintenance window, especially where ISO, NIS2 or a sector rule says who may change production code.

Pricing in a year when flash got expensive

NAND contract prices rose about 33–38% quarter on quarter in Q4 2025 and another 55–60% in Q1 2026, more than doubling from late Q3 2025. Samsung’s contract NAND pricing was reported at about a 100% increase in Q1 2026. DRAM moved further: roughly 38–43% in Q4 2025 and 105–110% in Q1 2026, with server DRAM up more than 60% in that first quarter. Hyperscalers reserved capacity for AI, DDR4 gave way to DDR5, and vendors

favoured high-margin server parts. Elevated pricing is expected at least through the end of 2026. Data reduction and dense QLC are how an array stays affordable, not a feature slide.



Summer 2026 movement across six vendors, from quotes StorComp collected. IBM's step was concentrated in June–July. This is a market average, not a tariff. See the appendix.

Indicative IBM marketplace configurations StorComp recorded around September 2026, with five years of Expert Care Advanced. They are not an IBM offer and not a price a reseller has to match. Across similar capacities, FlashSystem prices in that set are up roughly 25–35% versus the previous generation. Re-check them when the project is real. NAND is still moving. The limit on how to use these numbers is in the appendix.

Capacity	Price
5600	
about 80 TB effective	112,600 USD
about 300 TB effective	204,500 USD
7600	
about 80 TB effective	150,000 USD
about 300 TB effective	207,800 USD
9600	
about 470 TB effective	~3,500 USD / raw TB (238 TB raw reference)

Indicative IBM marketplace configurations, September 2026, with five years of Expert Care Advanced.

On a raw-TB basis from the same quote set, 5600 ran about 2,100 → 1,200 → 850 USD/raw TB as capacity stepped from ~80 to ~700 TB effective; 7600 about 4,400 → 2,000 → 1,300 USD/raw TB from ~80 to ~600 TB effective. Cache on 7600 stays at a 1.5 TB maximum; on 9600 at up to 3 TB. Neither model gained memory capacity in this refresh.

What this refresh did not touch

Entry SAS systems 5015 and 5045 are unchanged. FlashSystem C200 is still on the older drive generation. DS8000 mainframe systems are on FCM4 today and are the likely next stop for EDSFF and FCM5. SAN Volume Controller will move when CPUs, memory and NVMe port speeds move, not because the array architecture changed. Empty bays — three on 5600, four on 7600 and 9600 — are the hint that larger drive counts are still coming.

Where this leaves a shortlist

7600 and 9600 are credible when the requirement is density, reduction that does not sit on the controller, and cyber recovery tied to the media. 5600 is the same family only if the

workload fits 1U and PCIe Gen 4. Hold production cutover for an LTS code level, and do not sign a 5:1 business case against 9.1.2 scan-only dedupe.

Put them next to arrays that already do file and object on the same controllers, or that grow with an NVMe shelf instead of another node. IBM remains one of the few vendors that publishes the full FlashSystem documentation set, list-style pricing, and current Redbooks. Use that. Then compare the configuration, not the family name, in StorComp.

Appendix — how to read the prices

Every USD-per-terabyte figure in this briefing, and the SSD price chart, is an average of quotes StorComp collected. It is a picture of the market at the time those quotes were written. It is not an IBM offer, not a formal quotation, and not a number a reseller is obliged to honour.

Configuration, drive mix, support term, region, discount and the week of the quote all move the price. Use the average to see the shape of a shortlist. Put a purchase order against a quote written for the configuration you would actually buy.

> Based on StorComp research from public and market-observed sources. Actual results and pricing may vary by configuration, workload, region, discounts and support terms.

storcomp

StorComp LLC
Enterprise Storage Intelligence Platform
131 Continental Dr
Newark, Delaware
https://storcomp.com