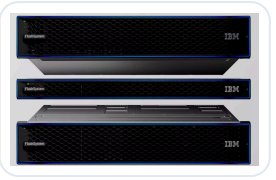
IBM Updated November 2026

StorComp Analysis - IBM FlashSystem x600 (2026)



What changed in the C300, 5600, 7600 and 9600: architecture, code 9.2.0 with full deduplication, pricing, and where each model fits.

TL;DR

7600 is the one to shortlist. 9600 when the job is density in 2U. C300 when you need a lot of flash capacity 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

8/10

FlashSystem C300

- New: capacity-optimised flash
- About $850–950/TB at 250–500 TB
- FCM5, dedupe in hardware
- Up to 3.4 PB raw in 2U
- Starts at 422 TB, grows in 211 TB steps

6/10

FlashSystem 5600

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

Update, November 2026: code 9.2.0

IBM Storage Virtualize 9.2.0 is the code the x600 generation was waiting for. The same release runs on FlashSystem C300, 5600, 7600 and 9600, so everything below applies to the whole family, not to one model.

Deduplication is finally fully available. At launch on 9.1.2 it ran in scan-only mode and did not return the full reduction. On 9.2.0 deduplicated volumes in data reduction pools deliver the

savings the 5:1 business case was built on. The usual condition still applies: the ratio depends on the data, so measure it on your own workload before it goes into a quote.

Other changes worth knowing from IBM's 9.2.0 release notes:

- Resilience: high availability plus synchronous disaster recovery (three-site), snapshot policies for HA partitions, and multiple snapshot schedules per policy, including Safeguarded schedules.
- Snapshot buffer to keep snapshot workloads with high peak I/O away from production volumes.
- Cyber resilience: wiperware detection and data exfiltration detection, on top of ransomware detection.
- FlashSystem.ai can now run onboard the array, manage FlashSystem grid systems, connect to external AI platforms and open IBM support cases.
- Operations: automatic compression for asynchronous replication, jumbo frames (MTU 9000) for NVMe/RoCE host attach, automatic fix pack download, and Fibre Channel portsets for HA and DR replication.
- Ceph object storage as a target for cloud snapshots.

What it changes for buyers: the main reason to keep these arrays in test and development was the non-LTS launch code and scan-only deduplication. With 9.2.0 both are addressed, so the 7600 and 9600 are now realistic candidates for tier-1 workloads. Ask that quotes and data-reduction assumptions are based on 9.2.0.

New in the family: FlashSystem C300

FlashSystem C300 is the capacity model of the new generation, the FCM5 counterpart to C200. It takes the same FlashCore Module 5 and the same Storage Virtualize code as 5600, 7600 and 9600, so it gets 9.2.0 and full deduplication on day one.

- 2U, 32 slots for FCM5 in EDSFF: 52.8 TB (XL) or 105.6 TB (XXL) modules, up to about 3.4 PB raw per enclosure.
- Compression, deduplication and encryption always on in the modules, plus ransomware detection.
- Entry configuration of 422 TB (8 × 52.8 TB) or 1.47 PB (14 × 105.6 TB); capacity grows in 211 TB steps, not single drives.

- Two node canisters with 128 GB each (256 GB optional); 64 Gb or 32 Gb FC, 25/10 GbE and 100 GbE with iSCSI and NVMe/TCP.

StorComp pricing at the standard discount: about $950/TB at 250 TB, $850/TB at 500 TB and $750/TB at 1 PB. 7600 sits at about $2,200/TB at 250 TB, $1,200/TB at 500 TB and $700/TB at 1 PB, so C300 costs well under half as much as a 7600 at 250 TB and about 30% less at 500 TB. At 1 PB a 7600 is slightly cheaper per raw TB, so C300 earns its place where density in 2U and data reduction matter more than the lowest $/TB.

Our rating: 8/10. It brings up to 3.4 PB of all-flash with hardware data reduction into 2U, on the same code and the same tools as the rest of the family. It is not a 10 because it is built for capacity, not peak IOPS, the entry point of 422 TB is too large for small sites, and expansion comes in 211 TB steps. For backup targets, archives, analytics and large virtual estates it is worth pricing next to 7600.

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.

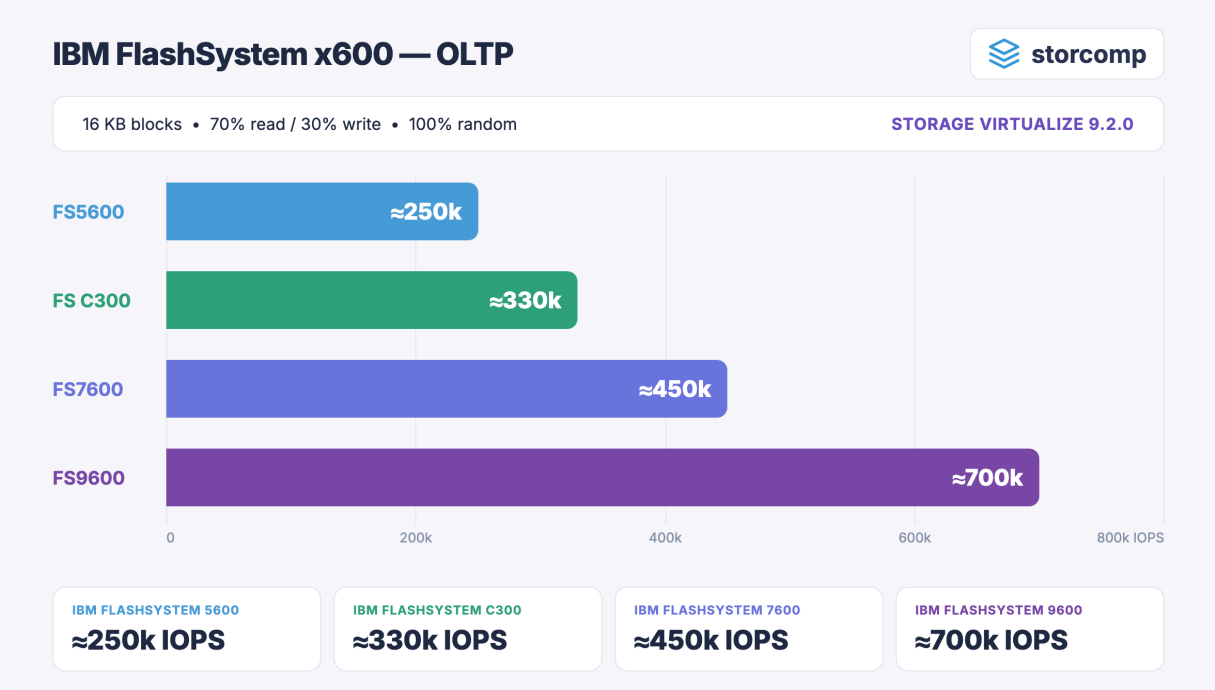
Headline numbers, new versus previous

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 flash (more with HDD shelves)**
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.



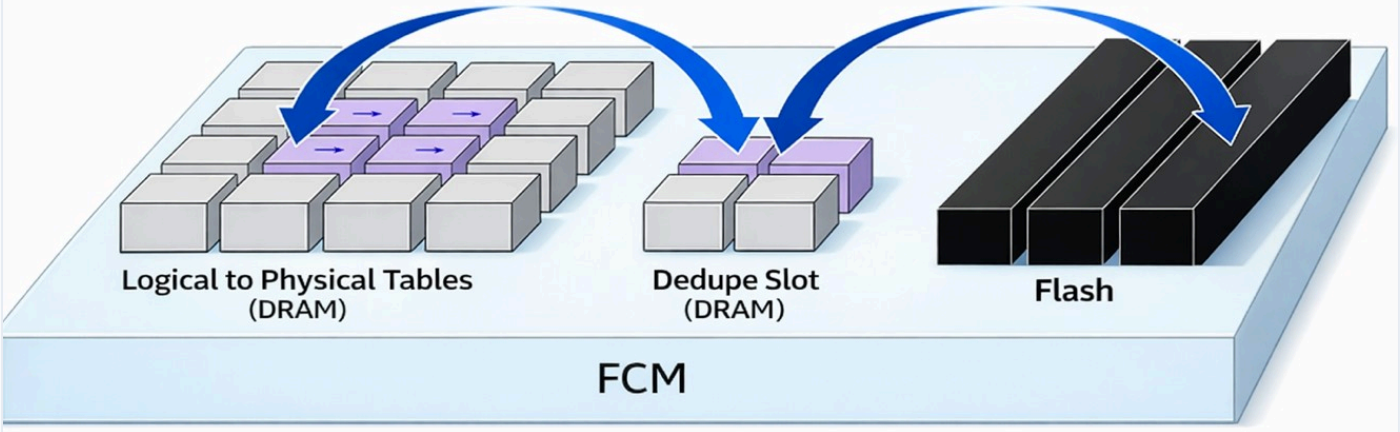
OLTP on Storage Virtualize 9.2.0: 16 KB blocks, 70/30 read/write, 100% random. About 250k IOPS on the 5600, 330k on the C300, 450k on the 7600 and 700k on the 9600.

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 and C300 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 was non-LTS. Do not make 9.1.2 the first home of a tier-1 workload; plan on 9.2.0 instead (see the November 2026 update).
- Dedupe was scan-only at 9.1.2. Fixed in 9.2.0; the 5:1 ratio still depends on the data.
- 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 (2U12 and 5U92 enclosures with NL-SAS HDD, for a capacity tier), 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

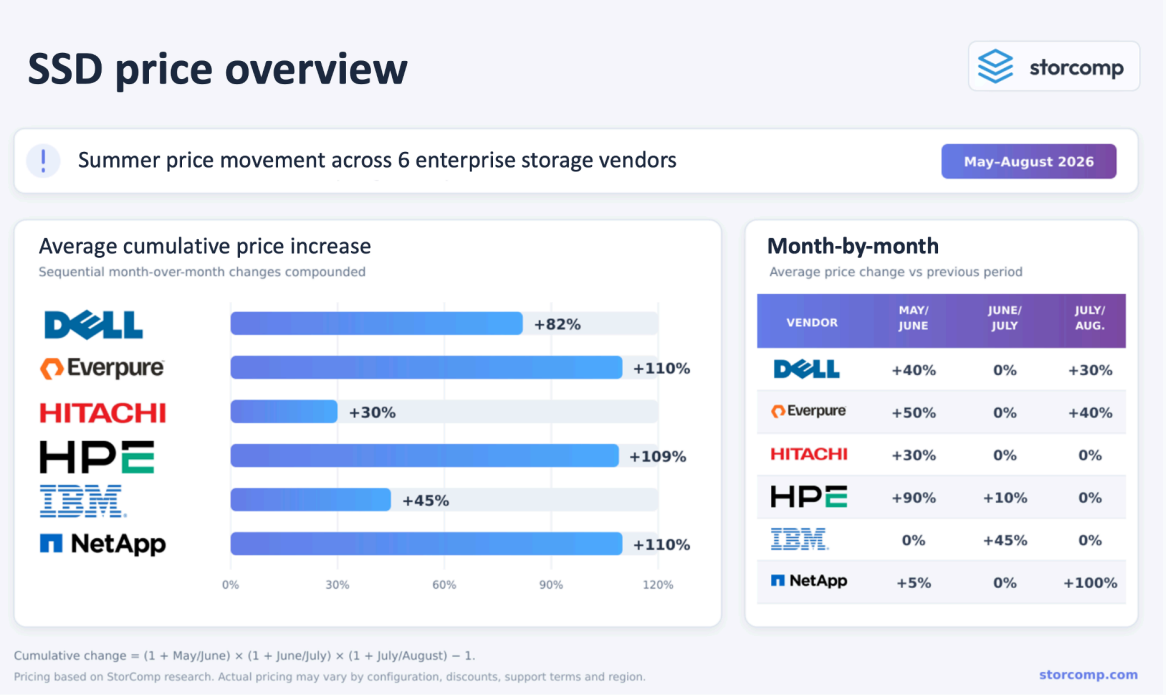
New volumes come up thin and data-reduced, without a separate Data Reduction Pool choice. Logical “full” is not the same as physical full. Capacity alerts, reserves and over-provisioning need a new reading, or the first incident will be a surprise.

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.

StorComp list-price levels for the whole family, including the new C300, at the standard discount (November 2026). These are $/TB raw by configuration size, so they compare the models with each other rather than replace a quote.

Model	~250 TB	~500 TB	~1 PB
FlashSystem 5600	$1,400	$1,000	$800
FlashSystem 7600	$2,200	$1,200	$700
FlashSystem 9600	$4,500	$2,300	$1,100
FlashSystem C300 (new)	**$950**	**$850**	**$750**

StorComp $/TB (raw) at the standard discount, November 2026.

What this refresh did not touch

IBM has announced that the entry SAS systems 5015 and 5045 will be withdrawn from sale in 2027. FlashSystem C200 was withdrawn from sale in October 2026; its FCM5 counterpart is the new C300 (see above). FS5300, FS7300 and FS9500 are withdrawn, replaced by 5600, 7600 and 9600. 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.

IBM documentation

Primary IBM sources for the new arrays, code 9.2.0 and deduplication. Check them for your exact model and code level before a design or a quote.

- IBM Documentation: What's new in Storage Virtualize 9.2.0 — Full list of 9.2.0 features (FlashSystem 7000 family docs, same code on 5600 and 9600).
- IBM Documentation: Deduplicated volumes (9.2.0) — How deduplication works in data reduction pools.
- IBM Documentation: Data reduction (9.2.0) — Compression, deduplication, thin provisioning and FCM data reduction.
- IBM Documentation: FlashSystem 5600 — Hardware overview.
- IBM Documentation: FlashSystem 7600 — Hardware overview.
- IBM Documentation: FlashSystem 9600 — Hardware overview.
- IBM Redpaper: IBM FlashCore Module Product Guide (REDP-5725) — FCM generations, including FCM5.
- IBM Newsroom: next-generation FlashSystem portfolio — Launch announcement of 5600, 7600 and 9600, 10 February 2026.

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