MARKET INSIGHT BRIEF · ADI

The All-Flash Tax

AI made flash the default. The market just made it unaffordable.

Our view

You do not have to pay flash prices for every terabyte to feed your GPUs. You have to put the right data on the right media. The all-flash data center is a pricing argument dressed as an architecture argument, and the 2026 supply shock has made that visible: every tier is rationed, so the hedge is not a cheaper medium, it is the freedom to use them all and overbuy none.

18.6x	+397%	+55-60%	Sold out	+10%
flash-to-disk cost per TB, August 2026	25 PB all-flash estate, Q2 2025 to Q2 2026	NAND contract prices, 1Q 2026 QoQ	Western Digital, calendar 2026	HDD trade prices, Q2 2026 QoQ

Memory is on the same curve: DRAM rose **+172% YoY** at the end of Q3 2025. There is no medium sitting this out.

Source: Cost multiple and all-flash estate model: VDURA Flash Volatility Index and Storage Economics Optimizer, via Blocks and Files, 2026-08-11. NAND: TrendForce, 2026-02-02. HDD supply and trade prices: Western Digital CEO via The Register and TechRadar, 2026-02, and TrendForce. DRAM: NAND Research, Memory and Flash Crisis March 2026 Update, citing TrendForce.

How far flash pulled away from disk, per terabyte

INDEX DATE	30 TB TLC SSD	30 TB HDD	MULTIPLE
Q3 2025	$3,460	$495	7.0x
Q1 2026 (peak)	not published	not published	23.2x
August 2026	$22,600	$1,216	18.6x

Source: VDURA Flash Volatility Index, published index values at three dates. Prices are list $/drive as published. Not a Scality measurement.

> **Storage is 2.4% of the AI infrastructure bill, and it decides whether the other 97.6% runs**
>
> In Q4 2025, enterprises spent $2.2B on AI storage against $87.7B on AI servers. Storage is the small line item that gates the large one. Underbuy it and the GPUs idle. Overbuy it in the wrong medium, at 2026 prices, and the budget goes to capacity that never needed to be fast.

Source: IDC Worldwide Quarterly AI Infrastructure Tracker, Q4 2025.

Four tiers, one S3 namespace

Scality ADI, AI Data Infrastructure, aligns every workload to the right media across a four-tier performance model, on a single S3 namespace spanning every media class, vendor and generation.

TIER	MEDIA	BUILT FOR
Top tier	TLC flash, S3 and S3 over RDMA, sub-50 µs latency	Training, validation, KV cache
Hot	QLC and NL-SSD	Data preparation, cleaning, normalization
Warm	NL-SSD and HDD	RAG, AI data lakes, geo-distributed access
Cold	Tape and cloud, via XDM lifecycle tiering	Archive and long-term retention

The data path does not change when the media does. High-concurrency S3 and S3 over RDMA run across the full media ladder, on NVMe and on HDD-based servers alike, so moving a workload down a tier is a placement decision rather than a re-architecture.

Where tiering wins, and where it does not

WHERE TIERING WINS	WHERE ALL-FLASH STILL WINS
Mixed workloads where most bytes are read rarely	Small estates where the management overhead of tiers outweighs the media saving
Estates growing faster than the storage budget	Workloads with no cold tail, where every byte is genuinely hot every day
Any environment where one medium being rationed must not stall the project	Environments where rack space or power, not media cost, is the binding constraint
Archive and retention obligations measured in years	

© 2026 Scality · scality.com

Performance at scale, not benchmark peak

420 GB/s	1500 TB/h	up to 20x	12 EB+
sustained reads on HDD, two-hour window	restore throughput, recovery in hours not days	better power efficiency than all-flash	in production, across 70+ countries

Source: Throughput and restore measured across a continuous two-hour window on just over 100 x86 nodes in three availability zones, with nearly 900 TB/hour backup: Scality HDD validation, solved.scality.com, 2026-05-26. Power efficiency, for typical enterprise data distributions: Scality ADI solution overview. Scale: scality.com/adi. 1000+ enterprises in production, continuously since 2009: scality.com/products/ring.

This is not a pilot architecture

10 of 20	7 of 15	6 trillion+	1B+
of the top global telcos and CSPs	of the largest banks	objects under management	users served

Source: scality.com/adi and the Scality ADI overview on solved.scality.com.

What to do before the next purchase order

1. **Measure the access distribution before you size the media.** Most estates cannot say what fraction of their bytes were read in the last 90 days. That number, not a vendor's tier chart, decides how much flash you actually need.
2. **Buy placement, not capacity.** A policy that moves an object down a tier and reclaims the fast capacity behind it is worth more per dollar in 2026 than the same spend on more flash.
3. **Keep more than one medium qualified.** Every tier is on long-term agreements into 2027 and 2028. An architecture that can only absorb one medium is exposed to whichever one is rationed next.

For the security side of the table

Tiering changes where data sits, not how well it is protected. CORE5 applies five layers of protection against ransomware across the namespace, and ADI is pre-integrated with Veeam, Commvault and Rubrik, so a colder tier is not a weaker one. Scality operates under an ISO 27001-certified information security management system.

Source: scality.com/adi, the Scality RING product page, and artesca.scality.com/en/security-cyber-resilience.

The two questions we get asked

Isn't this just hierarchical storage management with a new name?

No, because the namespace does not change when the data moves. HSM moved files between systems and left applications to follow. Here the object keeps its S3 key and its permissions on every tier, so moving a byte from flash to tape is invisible to the application reading it.

QLC wears out faster. Are we trading a cost problem for a failure problem?

Endurance is a sizing input, not a disqualifier, and it is why QLC sits on the Hot tier rather than under training. Size it against the write pattern of data preparation, which is write-once and read-many, and keep training checkpoints on TLC where the rewrite rate lives.

> This brief cannot tell you how much of your estate is cold. Any vendor who puts a percentage on your data without measuring it is showing you their average customer, not your environment. That number has to be measured.

WHAT WOULD CHANGE THIS VIEW

- → **NAND contract pricing reverses.** Signal: two consecutive quarters of QoQ decline. TrendForce already has gains moderating, so the direction is worth watching even though the level is not falling.
- → **HDD supply loosens.** Signal: long-term agreements stop extending past 2028, or a major vendor reopens uncommitted nearline capacity.
- → **QLC endurance claims hold at scale in the field.** Signal: two or more public multi-year deployments. That moves bytes down from TLC rather than changing the tiering case.

SCALITY

Built for GPU hunger, deep cold, and everything in between.

Request an ADI TCO and tiering assessment. Scality ADI is available now.