

**PROCUREMENT SPECTRUM***Insight-led. Practice-driven. Globally relevant.*

# The 2026 Memory and Storage Price Crunch

*What is happening, who is driving it, and what procurement teams should do now*

**A Procurement Spectrum Research Report • April 2026**

procurement-spectrum.com • [consult@procurement-spectrum.com](mailto:consult@procurement-spectrum.com)

## Executive Summary

Memory and storage prices are rising at a speed that most organisations have not experienced in over a decade. Between October 2025 and March 2026, conventional DRAM contract prices climbed approximately ninety-two percent quarter-on-quarter, and LPDDR5X mobile memory tracked a near-identical rise. TrendForce forecasts a further sixty percent increase in Q2 2026 for conventional DRAM and a seventy-to-seventy-five percent rise for NAND flash. High-Bandwidth Memory — the specialist chip that sits on top of every AI accelerator — is reported as effectively sold out at SK hynix and Samsung through the end of 2026.

Three forces are driving the shock. The first is the unprecedented 2026 capital expenditure commitment from five hyperscalers, which alone totals approximately six hundred and sixty billion dollars, roughly three quarters of it directed at AI infrastructure. The second is wafer capacity reallocation: Samsung, SK hynix, and Micron have each shifted fab capacity toward HBM and server-grade memory, leaving consumer DRAM and NAND with structural shortage. The third is a thin channel: DRAM inventory at distributors is down to two-to-four weeks.

This report sets out what is happening, who the supplier base looks like, how prices cascade to enterprise procurement budgets, and — most importantly — what procurement leaders should do in the next ninety days. The short answer is that most organisations should be accelerating their 2026 hardware and cloud purchases now, not waiting for relief that is unlikely to arrive before late 2027.

## Glossary — Terms used in this report

**DRAM (Dynamic Random-Access Memory):** the working memory inside servers, PCs, and mobile devices. Volatile — data is lost when power is removed.

**NAND flash:** non-volatile memory used in solid-state drives (SSDs) and mobile storage. Retains data when power is removed.

**HBM (High-Bandwidth Memory):** stacked DRAM dies mounted directly next to a GPU or AI accelerator. Required for training and serving large AI models. Produced by only three companies today.

**LPDDR5X:** low-power DDR memory used in smartphones, laptops, and increasingly in AI servers that need dense memory in low-power packages.

**UFS (Universal Flash Storage):** the high-performance NAND used in modern smartphones.

**RDIMM (Registered DIMM):** the server memory module format that sits in data-centre racks.

**QLC / TLC NAND:** Quad-Level Cell and Triple-Level Cell NAND — the density grades used in modern enterprise SSDs.

**Wafer:** a thin disc of silicon that is the raw material for chip manufacturing. A single wafer produces hundreds of chips.

**Foundry:** a contract chip manufacturer. TSMC is the dominant foundry.

**Hyperscaler:** a very large cloud services provider — in 2026, primarily AWS, Microsoft Azure, Google Cloud, Meta, and Oracle.

**Capex (capital expenditure):** long-term investment in assets such as data centres and equipment.

**QoQ / YoY:** Quarter-on-quarter and year-on-year price or volume change.

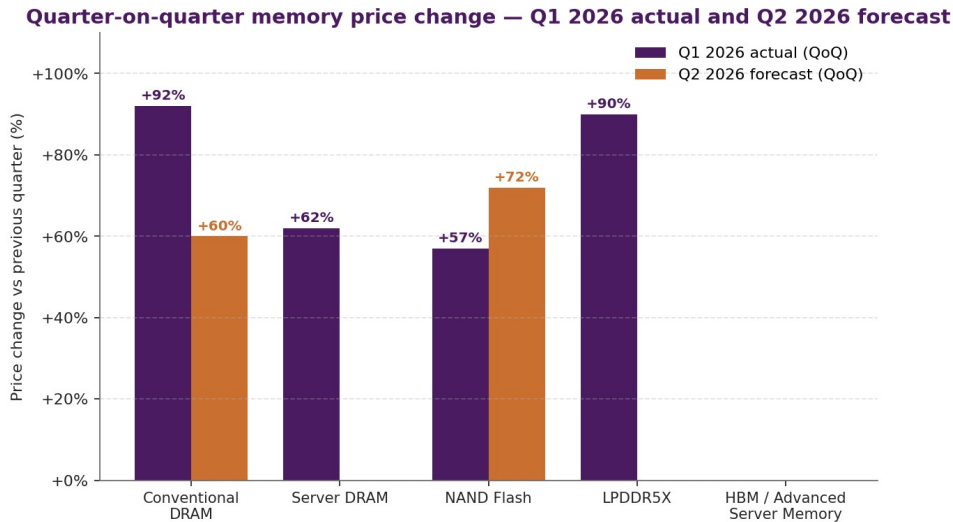
**Contract price vs. spot price:** the contract price is what large customers pay under framework agreements; the spot price is what small buyers pay on the open market. Spot is almost always more expensive.

**Channel inventory:** stock held by distributors and resellers between the chip maker and the final buyer. Low channel inventory signals tight supply.

## 1. What is happening to prices — 2026 so far

---

Memory and storage contract prices moved sharply upward in the first quarter of 2026 and are forecast to continue rising in Q2. The scale of the move is visible in the chart below.



Source: TrendForce contract-price reporting (Q1 2026) and published guidance for Q2 2026. HBM pricing reported as 'effectively sold out' by SK hynix; QoQ range not published.

Figure 1 — Quarter-on-quarter memory price change: Q1 2026 actual and Q2 2026 forecast. Source: TrendForce.

Conventional DRAM rose approximately ninety-two percent quarter-on-quarter in Q1 2026. Server DRAM — the grade used in every data-centre rack — was up sixty-two percent. NAND flash rose fifty-seven percent. LPDDR5X, the low-power memory used in smartphones and increasingly in AI servers, tracked a ninety-percent jump. TrendForce, which tracks contract prices, has published forecasts that indicate another sixty percent for conventional DRAM and seventy-to-seventy-five percent for NAND flash in Q2.

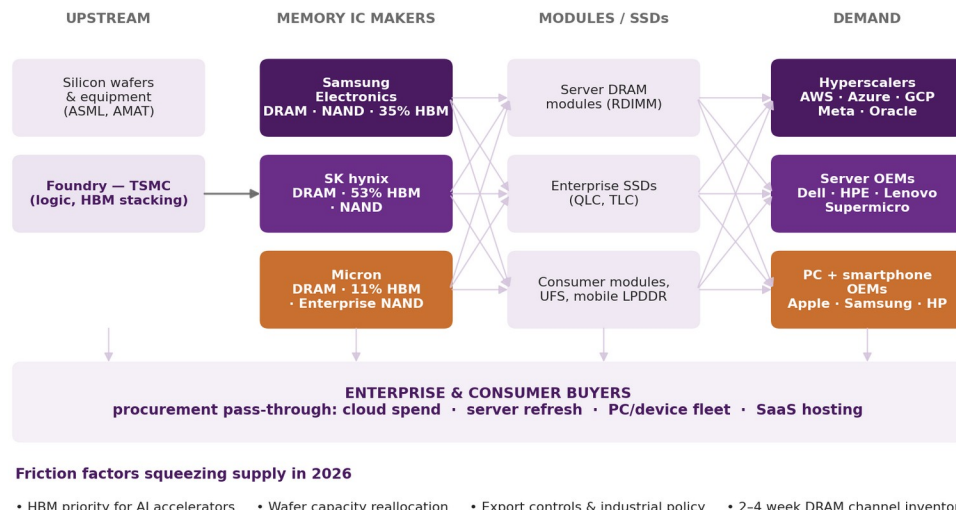
HBM is not shown on the chart because SK hynix and Samsung have described their 2026 HBM capacity as sold out. In a sold-out market, the variable is not price — it is allocation.

The practical consequence is visible downstream. Gartner's February 2026 commentary reported finished-PC price increases of seventeen percent and smartphone increases of thirteen percent; AWS announced a fifteen-percent increase on its EC2 Capacity Block ML (its on-demand GPU-server rental) in January 2026.

## 2. Who manufactures memory and storage — the 2026 supplier base

The memory industry is the most concentrated supplier base of any major electronic component. Three companies account for virtually all production-grade DRAM and HBM. The same three, plus a smaller cluster, account for the bulk of enterprise NAND.

### Memory and storage supply chain — 2026



Sources: TrendForce, Counterpoint Research, and supplier public disclosures. Supplier details illustrative; no endorsement implied.

Figure 2 — Memory and storage supply chain map: upstream materials, IC makers, modules, and demand segments.

Samsung, SK hynix, and Micron are the three memory IC makers whose decisions set prices globally. Upstream, they depend on silicon wafers (Shin-Etsu, Sumco), lithography and etch equipment (ASML, Applied Materials, Lam Research, Tokyo Electron), and on TSMC for advanced logic chips and for the interposer stacking that goes into HBM. Downstream, their output is packaged into server DRAM modules and enterprise SSDs by the same three companies plus Kioxia and SanDisk on the NAND side, and then reaches hyperscalers, server OEMs (Dell, HPE, Lenovo, Supermicro), and PC and smartphone OEMs.

### Supplier profile — the three memory IC makers

The table below captures what matters most about the three companies whose production and pricing decisions are setting the 2026 market.

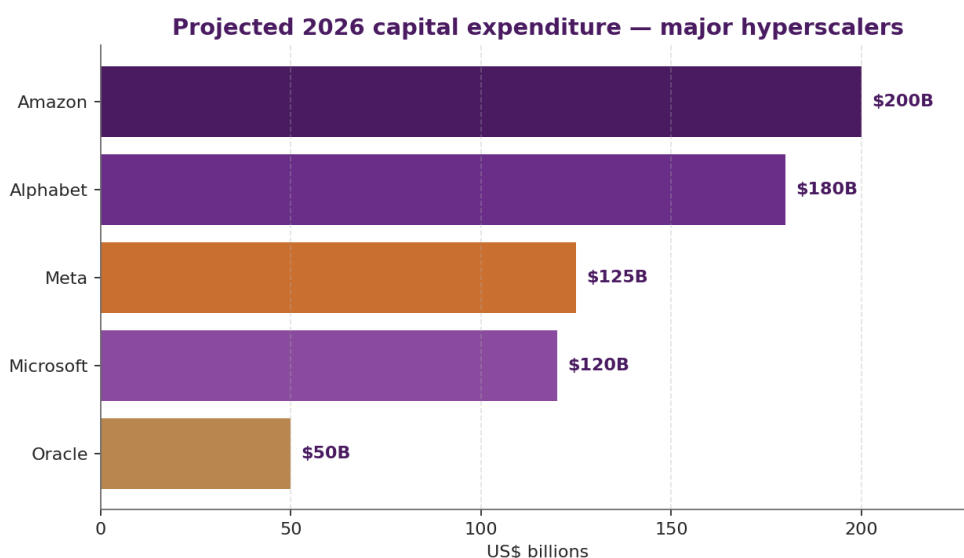
Attribute	Samsung Electronics	SK hynix	Micron Technology
Headquarters	Suwon, South Korea	Icheon, South Korea	Boise, Idaho, USA
2025 HBM market share	~35%	~53% (market leader)	~11%
DRAM position	#1 globally	#2 globally	#3 globally
NAND flash position	Major (second to Kioxia / SanDisk combined)	Top three	Top four
2026 capacity	HBM3E and HBM4	HBM3E at full run-rate	HBM3E (production),

Attribute	Samsung Electronics	SK hynix	Micron Technology
allocation bias	qualification for NVIDIA Blackwell-generation accelerators; conventional DRAM secondary	for NVIDIA; HBM capacity fully booked through 2026	HBM4 (sampling), enterprise NAND for hyperscalers
Key 2026 customers	NVIDIA (HBM qualification), AWS, Apple, Samsung Electronics mobile division	NVIDIA (primary HBM supplier), Microsoft, AWS	NVIDIA, Microsoft, Google, Oracle
2026 pricing action	90%+ contract increases on conventional DRAM; allocation prioritised to HBM customers	90-95% contract increases on conventional DRAM; HBM “effectively sold out”	~90% contract increases; increased HBM allocation; enterprise NAND prioritised over consumer

The concentration matters for procurement in a specific way. A buyer cannot negotiate around a three-supplier market the way it can around a fifteen-supplier one. The leverage sits with the sellers. The practical implication is that procurement should be negotiating on allocation and contract terms — priority, volume commitments, multi-year indexing — rather than relying on competitive tension to move unit prices.

### 3. Why prices are rising — the 2026 demand shock

The price move is not a conventional cyclical squeeze. It is a structural demand shock driven by AI infrastructure investment.



Source: Company disclosures and analyst consensus, midpoints of published 2026 capex ranges. Approximately 75 percent of combined spend is AI-related.

Figure 3 — Projected 2026 capital expenditure, major hyperscalers (USD billion). Source: Company guidance, Q4 2025 and Q1 2026.

Amazon's 2026 capital expenditure guidance is approximately two hundred billion US dollars. Alphabet's is between one hundred and seventy-five and one hundred and eighty-five billion. Meta's is one hundred and fifteen to one hundred and thirty-five billion. Microsoft's is approximately one hundred and twenty billion. Oracle's is approximately fifty billion. The combined total from these five companies is in the six-hundred-and-sixty-to-six-hundred-and-ninety-billion-dollar range, and public commentary from the companies themselves indicates that approximately seventy-five percent of this spend is directed at AI infrastructure.

The hyperscaler capex is only one side of the story. Every AI server they buy is a demand signal that propagates upstream — more HBM per GPU, more server DRAM per rack, more enterprise SSD for training data. Samsung and SK hynix have responded by shifting wafer capacity toward HBM, which mechanically reduces the supply of conventional DRAM available to everyone else. Goldman Sachs has described the resulting imbalance as the largest DRAM undersupply in more than fifteen years, at approximately 4.9 percent.

There is no quick relief lever. New DRAM fabs take two to three years to reach volume production. Samsung's Pyeongtaek expansions, SK hynix's M15X and M16, and Micron's New York campus are all on long time horizons. The supply response will come, but not in time for 2026 buying decisions.

## 4. How prices cascade to your organisation

Whichever part of the supply chain your organisation purchases from, the price move reaches you. The cascade is short.

### How memory pricing cascades through enterprise cost base



*Spot buying is the most expensive way to transact in this cycle. Long-term, indexed agreements are the most effective hedge.*

Sources: TrendForce, AWS pricing updates (Jan 2026), Gartner (Feb 2026).

Figure 4 — How memory pricing cascades through the enterprise cost base.

A memory fab raises its contract price. Within one quarter, server and SSD OEMs revise their pricing for new orders. Within two quarters, cloud service providers raise on-demand rates for GPU- and memory-intensive instance types — AWS, Azure, and GCP have all done so. Within

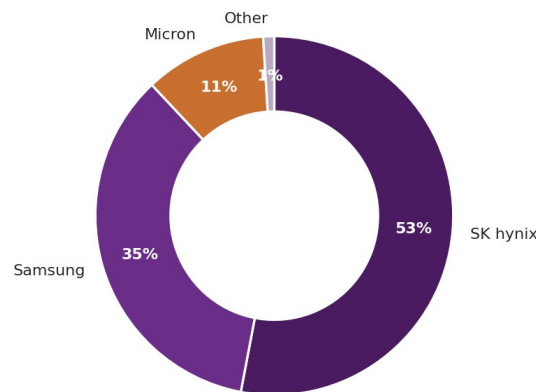
two to three quarters, PC and smartphone OEMs raise retail prices. Enterprises feel this on the infrastructure-refresh cycle, the cloud bill, the device fleet, and the SaaS hosting pass-through.

The organisations most exposed are those with three characteristics: upcoming server or storage refresh cycles, material consumption of memory-intensive cloud instances, and large PC or device estates due for replacement.

## 5. The HBM bottleneck — special case within the special case

High-Bandwidth Memory deserves separate treatment because it behaves differently from every other memory category in 2026.

High-Bandwidth Memory (HBM) market share — Q3 2025



Source: Counterpoint Research and TrendForce public disclosures. Totals may not sum to 100% due to rounding.

Figure 5 — HBM market share, Q3 2025. Source: Counterpoint Research.

SK hynix holds approximately fifty-three percent of the HBM market. Samsung holds approximately thirty-five percent. Micron holds approximately eleven percent. A sliver of other production accounts for the remainder. This three-company concentration is tighter than even conventional DRAM, and the 2026 capacity is allocated almost entirely to NVIDIA, AMD, and the hyperscalers who design custom accelerators (Google TPU, AWS Trainium, Microsoft Maia).

For procurement teams, the HBM implication is indirect but important. If your organisation is sourcing AI servers — whether NVIDIA, AMD, or a hyperscaler-managed service — you are not buying HBM directly, but your delivery timeline is governed by HBM allocation. Lead times of twenty-six to forty weeks on AI-class GPU systems are now normal. That is a procurement question, not an IT question, and it shifts the calculus on when to place orders.

## 6. What procurement teams should do now — a practical action plan

The most important question is not “what is happening” but “what should we do about it.” The analysis above supports a clear set of recommendations. They are tiered by category, because different categories have different risk profiles and different supply dynamics.

### The core recommendation: accelerate, do not wait

For most categories and most organisations, the evidence supports accelerating 2026 hardware purchases into the first half of the year rather than waiting. Three facts underpin this recommendation. First, Q2 2026 contract prices are forecast to rise further, not decline. Second, hyperscaler capex commitments are locked in through 2026 and are not reversing. Third, supply-side capacity additions will not arrive in time — new fab output is late-2027 and 2028 business.

### Tiered action plan

**Tier 1 — Buy now, accelerate by one to two quarters.** Applies to: server DRAM and server memory upgrades, enterprise SSD storage refreshes, standard corporate PCs (laptops and desktops), standard smartphone fleets for 2026 refresh, and any memory-heavy workstation. If your organisation has a 2026 refresh plan for any of these, place the order in Q2 2026 at the latest. Negotiate fixed-price contracts now, even if delivery is staggered. Consider pulling forward ten to fifteen percent of planned 2027 volumes where warranty and depreciation policy permit.

**Tier 2 — Queue now, price is secondary.** Applies to: AI-class GPU servers (NVIDIA H100, H200, B100, B200; AMD MI325X, MI350; hyperscaler custom-accelerator bare-metal offerings). These are allocation-constrained. Price negotiation matters far less than position in the delivery queue. The procurement action is to place firm orders early in 2026, pre-commit to volumes, and accept the twenty-six-to-forty-week lead times. Organisations that are still specifying or benchmarking should compress that process. For a subset of workloads, consider committed-capacity cloud reservations (AWS Capacity Reservations, Azure Reserved Instances, GCP Committed Use) as a hedge against on-premises delivery slippage.

**Tier 3 — Restructure contracts, move off spot.** Applies to: all memory-intensive cloud consumption and SaaS services billed by compute-unit. Where spot or on-demand purchasing remains the default, shift to reserved or committed-use pricing for 2026 workloads. Lock in twelve-to-twenty-four-month commitments on the workloads where consumption is predictable. The AWS EC2 Capacity Block ML fifteen-percent increase is a visible example; the less-visible increases on memory-heavy general-purpose instances follow the same logic.

**Tier 4 — Stage purchases, use indexed contracts.** Applies to: consumer-grade equipment where 2026 refresh is not mission-critical, and to product-embedded memory in hardware the organisation sells. Where deferral is possible, stage the purchase across 2026 and early 2027. Use indexed contracts that share the risk of further price moves with the supplier, rather than

fixed-price contracts that force the supplier to price in all the uncertainty. Long-term indexed agreements are the most effective hedge against further shocks.

**Tier 5 — Hold only in narrow cases.** There is a narrow case for holding off purchases into late 2027 or 2028: discretionary refresh projects where the existing estate is healthy, where business case is marginal, and where the organisation can bear a further price rise without operational impact. This is a deliberate risk decision, not a default.

### Negotiation posture for 2026

In a three-supplier, tight-allocation market, traditional competitive-tender dynamics are weaker than usual. Four levers matter more than unit-price negotiation.

The first is multi-year volume commitment in exchange for price indexing and priority allocation. Suppliers will accept visible demand in return for contract stability. The second is supplier diversification at the module and system level, not the IC level — Dell, HPE, Lenovo, and Supermicro all source from the same three memory makers, but their aggregate allocation positions and their willingness to bundle are different. The third is payment-term flexibility: extended terms and upfront deposits have become a meaningful lever when capacity is scarce. The fourth is contract clauses that protect the buyer: force-majeure language that does not void delivery commitments for “market conditions,” clear escalation paths if allocation slips, and renegotiation windows in 2027 when supply relief is expected.

### What to avoid

Three practices materially destroy value in this cycle. The first is relying on spot or unstructured purchase for any material volume — the spot-to-contract gap has widened sharply. The second is signing short-term fixed-price contracts without volume protection: the supplier's incentive is to prioritise other customers if your order becomes uneconomic for them. The third is treating the AI-server delivery queue as a pricing negotiation; it is not, and attempting to negotiate on price rather than queue position will leave the organisation without hardware.

## 7. Emerging considerations — what to watch in H2 2026 and 2027

Four indicators will signal whether the environment is shifting.

Capacity announcements from Samsung, SK hynix, and Micron will signal the pace of supply-side response. Watch for fab ramp-dates and specific volume guidance rather than headline capex announcements, which tend to overstate near-term impact.

Hyperscaler capex revisions mid-2026 are the single most powerful demand signal. A material cut would ease supply, though current commentary from all five companies implies continued increases into 2027.

Export controls and industrial policy, particularly US restrictions on memory-grade equipment to China and any reciprocal response, would shift global allocation. The 2026 baseline assumes current restrictions remain in place.

Consumer electronics demand, particularly smartphone replacement cycles, is a secondary indicator. Weak consumer demand frees a small amount of LPDDR capacity but does little for server DRAM or HBM.

A realistic base case is that contract prices stabilise at new elevated levels in early-to-mid 2027, with material relief in late 2027 as new fab capacity ramps. Organisations should plan on this basis and build renegotiation windows into any multi-year agreements signed in 2026.

## Conclusion

The 2026 memory and storage price move is structural rather than cyclical, driven by a demand shock that will persist as long as hyperscaler AI capex persists. The supply base is concentrated to a degree that limits procurement leverage. The most important procurement decisions are the ones made in the first half of 2026 — whether to accelerate refreshes, how to negotiate allocation, where to move off spot pricing, and how to structure contracts that survive a further price move.

Procurement teams that treat this as a classic cyclical negotiation will pay more and receive less. Teams that treat it as an allocation-and-contract-engineering problem will protect both budget and operational delivery.

---

*For advisory or consulting services on memory and storage procurement strategy in 2026, contact [consult@procurement-spectrum.com](mailto:consult@procurement-spectrum.com) | [procurement-spectrum.com](https://procurement-spectrum.com)*

*Sources drawn upon in this report: TrendForce quarterly contract-price reporting (Q4 2025 and Q1 2026); Gartner commentary, February 2026; Counterpoint Research HBM quarterly shares; SK hynix, Samsung Electronics, and Micron Technology earnings commentary; Amazon, Alphabet, Meta, Microsoft, and Oracle 2025 Q4 and 2026 Q1 capex guidance; AWS EC2 Capacity Block ML pricing notice, January 2026; Goldman Sachs semiconductor research, March 2026.*

*© 2026 Procurement Spectrum. This report is provided for informational purposes only and does not constitute legal, financial, technical, or commercial advice. All third-party data, platform features, pricing, and capabilities are attributed to publicly available sources as at the date of publication; readers should verify current details with each source or vendor before committing to action. Product names and trademarks are the property of their respective owners.*