

The Memory Supercycle

Researched by Hey Lefty

Automated research briefings on topics you choose — heylefty.com

TL;DR

Conventional DRAM prices skyrocketed in early 2026, handing memory giants massive leverage to demand steep price hikes in upcoming negotiations for next-generation HBM DRAM Contract Pricing Surge. Meanwhile, a high-stakes capital expenditure race is underway as SK Hynix, Samsung, and Micron build out multi-billion dollar gigafabs [AI Memory Capacity Expansion Race](/topics/019e8ec8-ffb3-71d7-981c-e48354ab25e7/notes/ai-memory-capacity-expansion-race-computex-2026), driving a structural re-rating for semiconductor equipment providers Semiconductor Equipment Valuation Re-rating.

The DDR5/HBM Profitability Arbitrage

A massive surge in conventional DRAM contract prices has inverted HBM's traditional margin advantage, handing memory makers unprecedented leverage for upcoming supply negotiations.

"HBM wafer revenue and profitability were overtaken by high-capacity DDR5... RDIMMs" — DRAM vs HBM Profitability Arbitrage

Because HBM contracts are negotiated annually, suppliers were temporarily locked out of the explosive pricing spikes seen in conventional DRAM during 2026 DRAM Contract Pricing Surge. By threatening to divert wafer capacity back to standard high-capacity memory, manufacturers are successfully forcing AI chip designers to accept steep hikes for 2027 HBM supply contracts DRAM vs HBM Profitability Arbitrage.

What to watch: Watch whether major hyperscalers capitulate to these pricing demands during negotiations or risk losing guaranteed memory allocations.

The Late-Stage Capacity Expansion Race

Leading memory manufacturers are racing to build massive new fabrication facilities to capture the AI demand boom, triggering a classic late-stage supply expansion that could eventually lead to a glut.

"...the simultaneous buildout of multi-billion dollar fabs... sets the stage for a potential supply glut once these facilities come online" — AI Memory Capacity Expansion Race

While memory executives argue that HBM's severe wafer penalty makes this market wave structurally different, the massive capital expenditure plans laid out by SK Hynix, Samsung, and Micron at Computex 2026 mimic historical peaks [AI Memory Capacity Expansion Race](/topics/019e8ec8-ffb3-71d7-981c-e48354ab25e7/notes/ai-memory-capacity-expansion-race-computex-2026). If demand cools even slightly before these multi-billion dollar gigafabs come fully online in late 2027 or 2028, the industry will face a severe capacity overhang [AI Memory Capacity Expansion Race](/topics/019e8ec8-ffb3-71d7-981c-e48354ab25e7/notes/ai-memory-capacity-expansion-race-computex-2026).

What to watch: Watch for any signs of delays or capital expenditure cuts in SK Hynix's Yongin fab or Micron's Boise facility as a tell for shifting demand.

Samsung's Vertical Integration Play

Samsung is leveraging its unique, fully integrated manufacturing model to bypass external bottlenecks and challenge SK Hynix's dominant position in the next-generation HBM roadmap.

"Samsung's core strategic weapon is its vertically integrated structure, spanning in-house memory, foundry logic, and advanced packaging." — Samsung Vertical Integration

By utilizing its own 4-nanometer and 2-nanometer foundry logic dies rather than relying on external fabs, Samsung is attempting to solve thermal and power efficiency bottlenecks faster than its rivals [Samsung Vertical Integration](/topics/019e8ec8-ffb3-71d7-981c-e48354ab25e7/notes/samsung-vertical-integration-hbm4e-hbm5-leap). If successful, this full-stack offering could rapidly erode SK Hynix's dominant high-bandwidth market share.

What to watch: Watch for Nvidia's official qualification of Samsung's multi-layer HBM4E chips, which began shipping as samples in May 2026 [Samsung Vertical Integration](/topics/019e8ec8-ffb3-71d7-981c-e48354ab25e7/notes/samsung-vertical-integration-hbm4e-hbm5-leap).

Semiconductor Equipment Valuation Re-rating

The extreme manufacturing complexity of stacking high-bandwidth memory is structurally re-rating the valuations of wafer fabrication equipment leaders.

"...AMAT is securing up to eight quarters (2 years) of forward demand visibility..." — Semiconductor Equipment Valuation Re-rating

Because high-bandwidth memory requires significantly more wafer area and manufacturing steps than conventional DRAM, equipment suppliers like Applied Materials and Lam Research are capturing far more value per wafer [Semiconductor Equipment Valuation Re-rating](/topics/019e8ec8-ffb3-71d7-981c-e48354ab25e7/notes/semiconductor-equipment-re-rating-hbm-complexity-gains). This shift has transformed these historically volatile toolmakers into stable infrastructure plays with multi-year order backlogs [Semiconductor Equipment Valuation Re-rating](/topics/019e8ec8-ffb3-71d7-981c-e48354ab25e7/notes/semiconductor-equipment-re-rating-hbm-complexity-gains).

What to watch: Watch whether Lam Research and Applied Materials can sustain their record-high valuation multiples if memory capex begins to cool.

What surprised us

- **DDR5 beat HBM in profitability.** High-bandwidth memory is typically the high-margin crown jewel, but standard DDR5 RDIMMs actually surpassed HBM in per-wafer profitability during early 2026 DRAM vs HBM Profitability Arbitrage. Because HBM contracts are negotiated annually, suppliers couldn't adjust prices to match the DRAM price spike, creating a fascinating leverage point for upcoming negotiations DRAM vs HBM Profitability Arbitrage.
- **Equipment visibility has stretched to two years.** Wafer fabrication equipment providers like Applied Materials historically operated with mere months of forward order visibility. Now, they are securing up to eight quarters of forward demand, structurally transforming them from volatile toolmakers into stable infrastructure providers Semiconductor Equipment Valuation Re-rating.
- **Samsung's aggressive foundry leap.** Despite trailing in current HBM market share, Samsung is playing its vertical integration card by targeting a cutting-edge 2-nanometer foundry process for its next-generation mockup Samsung Vertical Integration. This bypasses external foundry bottlenecks entirely, showing Samsung intends to win the next packaging war through sheer in-house scale Samsung Vertical Integration.

Open threads worth a vote

- [Samsung HBM Nvidia Qualification and 2027-2028 Supply Glut Risk](/topics/019e8ec8-ffb3-71d7-981c-e48354ab25e7#threads) — Help us track Samsung's qualification progress with Nvidia and whether the massive multi-billion dollar capex plans risk triggering a severe supply glut when new capacity comes online.

Appendix: Findings

DRAM Contract Pricing Skyrockets to Record Highs in Q1 2026

DRAM Contract Pricing Skyrockets to Record Highs in Q1 2026

The memory market has entered an explosive upcycle. In the first quarter of 2026, conventional DRAM contract prices experienced a near-doubling, rising between **93% and 98% quarter-over-quarter**. This price surge drove global DRAM industry revenue up by **81% QoQ to \$97 billion**, establishing massive windfalls for the top three memory suppliers: Samsung, SK Hynix, and Micron.

This extreme price environment is expected to persist through the second quarter of 2026. TrendForce projects conventional DRAM contract prices to rise by an additional **58% to 63% QoQ** in 2Q26. The primary driver of this "chipflation" is an acute supply shortage. Low supplier inventories and a structural reallocation of wafer capacity toward high-capacity RDIMMs and High-Bandwidth Memory (HBM) for AI servers have severely crowded out PC OEMs and smartphone clients, forcing double-digit percentage price hikes on consumer hardware.

The massive financial impact is visible in the public financial data. For example, Micron (MU) reported TTM revenue of \$58.12B, with its latest quarterly revenue ending February 28, 2026, exploding **196.3% YoY** to \$23.86 billion and generating \$13.79 billion in net income (a 756.0% YoY earnings increase).

Instance of `[[cd053c2792413]]`{why="Memory suppliers now dictate the pace of AI innovation through capacity arbitrage."}

Sources

- Rapid Contract Price Surge Drives 1Q26 DRAM Industry Up 81% QoQ, Says TrendForce
- Expect more of those DRAM price hikes as memory shortage continues to bite

The AI Memory Capacity Expansion Race and the Late-Cycle Capex Boom

The AI Memory Capacity Expansion Race and the Late-Cycle Capex Boom

At Computex 2026, the world's leading memory manufacturers announced massive capital expenditure and capacity expansion plans to address a structural supply gap. Industry sources estimate that current memory capacity can only satisfy **around 50%** of the medium- to long-term demand from global Big Tech firms.

The scale of the announced investments is staggering, presenting the classic "late-cycle tell" where all major players aggressively add capacity simultaneously to capture high-margin demand:

- **SK Hynix:** Chairman Chey Tae-won announced plans to **double total wafer capacity over the next five years**, predicting the memory shortage could persist through 2030. The company's 2026 capex is projected to surpass **KRW 30 trillion (~\$21 billion)**. It is investing **KRW 31 trillion (~\$22.5 billion)** in its first Yongin cluster fab (Y1), targeting **300,000 wafers/month** by the end of 2027. Its Cheongju **M15X** fab is ramping in 2026 to reach a total DRAM capacity target of **620,000 wafers per month by H2 2026** (nearly double 2023 levels).
- **Samsung:** Projected 2026 capex is expected to reach **~\$20 billion** (up 11% YoY), focused on advanced process nodes and its **P4L** facility.
- **Micron:** Projected 2026 capex is **\$13.5 billion** (up 23% YoY), with initial wafer output at its Boise, Idaho, fab slated for mid-2027.

While executives argue that this upcycle is structurally different due to the immense wafer penalty of HBM—which requires 3x to 4x more wafer capacity than DDR5 for the same bits—the simultaneous buildout of multi-billion dollar fabs (like Yongin Y1, Cheongju M15X, Samsung P4L, and Micron Boise) sets the stage for a potential supply glut once these facilities come online fully in 2027–2028.

Sources

- SK Hynix pledges to double wafer capacity within five years as AI memory shortage deepens toward 2030
- Samsung Applies 2nm to HBM5 First; SK to Accelerate Output via AI Factory
- AI Memory Shortage Locked Through 2030: Computex 2026 Brings Agent Economy, HBM Crisis

AI Memory Bottlenecks Drive Epic Valuation Re-rating for Semiconductor Equipment Leaders

AI Memory Bottlenecks Drive Epic Valuation Re-rating for Semiconductor Equipment Leaders

The massive capital expenditure boom in the memory sector (detailed in [\[\[ai-memory-capacity-expansion-race-computex-2026\]\]](#)) is driving an unprecedented valuation re-rating for front-end and back-end wafer fabrication equipment (WFE) providers. Applied Materials (AMAT) and Lam Research (LRCX) are key beneficiaries of this transition.

The fundamental driver is that high-bandwidth memory (HBM) is vastly more complex to manufacture than conventional DRAM. This complexity dramatically increases the equipment intensity per wafer:

- **Wafer Area & Steps:** According to Applied Materials management, each HBM unit consumes approximately **three times the wafer area** of standard DRAM and involves **19 manufacturing steps**, 15 of which require advanced semiconductor equipment. Applied Materials captures over half of the equipment value linked to those 15 steps.
- **Advanced Packaging Surge:** AMAT expects its advanced packaging revenue—driven by 3D stacking and chiplet integration of HBM and logic dies—to surge by **over 50% in calendar year 2026**.
- **Demand Visibility:** Historically, equipment makers operated with 3 to 6 months of forward order visibility, leading to highly cyclical valuations. Today, AMAT is securing up to **eight quarters (2 years) of forward demand visibility**, altering its investment profile from a cyclical toolmaker to a stable infrastructure provider.

This structural shift is reflected in public-market valuations. Applied Materials (AMAT) is trading at its 52-week high with a market cap of **\$389.08 billion** and a trailing P/E of 46.1 (with shares up 51.1% over the past three months). Similarly, Lam Research (LRCX) is at its 52-week high with a market cap of **\$418.20 billion** and a trailing P/E of 63.34 (with shares up 67.8% over the past three months).

Sources

- Does Applied Materials Stock Deserve Its Epic Rally?

HBM vs DDR5 Profitability Arbitrage Grants Suppliers Leverage for 2027 HBM4 Negotiations

HBM vs DDR5 Profitability Arbitrage Grants Suppliers Leverage for 2027 HBM4 Negotiations

A fascinating structural anomaly has emerged in the HBM market. While conventional DRAM contract prices spiked nearly 100% in Q1 2026 (as tracked in [\[\[dram-contract-pricing-skyrockets-to-record-highs-q1-2026\]\]](#)), High-Bandwidth Memory (HBM) contract prices did not keep pace. HBM contracts are negotiated annually, which prevented suppliers from adjusting their prices dynamically to reflect the quarterly market spikes.

As a result, TrendForce's analysis of per-wafer revenue and profitability—using die size, yield rates, and per-Gb pricing—revealed that **HBM wafer revenue and profitability were overtaken by high-capacity DDR5 64GB RDIMMs** in 1Q26.

This profitability inversion has completely shifted the power dynamics of the current Q2 2026 negotiations for 2027 HBM4 supply agreements. The top three memory giants (Samsung, SK Hynix, and Micron) are using the high profitability of DDR5 RDIMMs as a powerful cudgel. If AI chip designers and cloud service providers (CSPs) do not agree to massive price hikes for HBM4 in 2027, the suppliers can simply threaten to reallocate their tight wafer capacity to standard DDR5 RDIMMs. This arbitrage is expected to drive HBM4 contract prices significantly higher in 2027.

Instance of [[cd053c2792413]]{why="Memory suppliers now dictate the pace of AI innovation through capacity arbitrage."}

Sources

- Tight DRAM Supply Gives Suppliers Greater Pricing Power in HBM, with HBM Contract Prices Expected to Surge Multiples Higher in 2027, Says TrendForce
- Memory Giants Gain Pricing Power Over AI Firms as DDR5 RDIMMs Beat HBM Profitability, Setting Stage for 2027 Hikes

Samsung Leverages Vertical Integration to Leapfrog Competitors in HBM4E and HBM5

Samsung Leverages Vertical Integration to Leapfrog Competitors in HBM4E and HBM5

Samsung Electronics is executing an aggressive technology roadmap to regain momentum in the HBM market, where SK Hynix currently holds a leading **57% revenue share** (as of late 2025, with Samsung at 22% and Micron at 21%).

Samsung's core strategic weapon is its **vertically integrated structure**, spanning in-house memory, foundry logic, and advanced packaging. This allows it to offer a unified, full-stack solution that bypasses external foundry bottlenecks.

At the end of May 2026, Samsung began shipping samples of its **12-layer HBM4E** chip, pulling ahead of its rivals in the sixth-generation HBM race. The HBM4E uses Samsung's sixth-generation 10nm-class (**1c DRAM**) process technology coupled with its own **4-nanometer foundry logic base die**.

Furthermore, at Computex 2026, Samsung unveiled a physical mockup of its eighth-generation **HBM5** chip, which will apply a base/logic die built on its own cutting-edge **2-nanometer foundry process** and feature stacked configurations of up to 20 layers. By fabricating the base die on its 2nm node, Samsung aims to maximize power efficiency and heat dissipation, addressing the core thermal bottlenecks of high-layer HBM stacks.

If Samsung successfully completes qualification for HBM4E and HBM5, its massive in-house manufacturing capacity could trigger a significant shift in the competitive HBM landscape.

Sources

- Samsung Electronics ships faster HBM4E chip samples to customers; shares jump
- Samsung Applies 2nm to HBM5 First; SK to Accelerate Output via AI Factory