

MEMORY SUPER CYCLE

How in the world did I miss it?

February 2026

Kevin Cook
Zacks Investment Research

Micron Technology, Inc. · 1D · NASDAQ

O405.19 H420.88 L392.71 C411.66 -2.31 (-0.56%)

USD

411.66
SELL

0.00

411.66
BUY

Vol 33.81 M

MA 50 close 330.21

SMA 200 close 192.69

SMA 20 close 400.32

^

MU 411.66



Sep

Oct

Nov

Dec

2026

Feb

Ma

⊞



Kevin Cook

7

Dec '25 (edited) • Chip R&D

Save Post



Micron, the Megalodon of Memory

I'm a frustrated Micron (MU) investor going back at least 8 years. As Sanjay Mehrotra moved the company into all facets of the advanced computing revolution -- from self-driving cars, mobile and medical to industrial, robotics, and HBM for AI -- the stock never got any respect relative to the rest of the semiconductor industry.

I did a presentation in 2020 titled "The Megalodon of Memory" to highlight these advances to make memory "intelligent" and "embedded" -- and more than a mere commodity.

Des... [See more](#)

Evolution of DRAM

- 1992: 16Mb DRAM
- 2002: 1Gb DRAM
- 2014: 8Gb DRAM
- 2020: 1α

Memory Technologies by Segment

Segment	LPDDR, GDDR & HBM	NAND and NOR
Automotive ADAS	<ul style="list-style-type: none"> High-performance volatile memory Modules & components Low power consumption 	<ul style="list-style-type: none"> Long-term memory storage Security Protection Low Power Consumption Reprogrammable Bootling Software
Data Center		
Personal Computing		
PC Gaming		
Portable Storage		

[Jump to latest comment](#)

Kevin Cook · 20d

Wowza! I knew that David Tepper had been an unreasonably bullish fan of Micron in the early 2020s. I never checked to see what he had held on to for the recent surge, until I just read this from a Motley Fool piece:

“Micron has been a longtime holding for Tepper. At one point during the semiconductor shortage of 2020 and 2021, the stock accounted for roughly 10% of his portfolio. After Appaloosa trimmed the position, it accounted for just over 1% of the fund's equities as of the end of the third quarter.”

So I double checked their math on Whale Wisdom for the Sep quarter and found even his June position was only 1.58% of the portfolio when he held 825k shares.

As of Sep quarter he held 500k shares or 1.13% of portfolio.

He has done well and was clearly vindicated. But he might be kicking himself now and then about emptying his big bags too soon.

Sandisk Corporation · 1D · NASDAQ ≈ O610.53 H661.50 L586.37 C626.56 -3.73 (-0.59%)

USD

626.56
SELL

0.00

626.56
BUY

Vol 23.64M

MA 50 close 0 ↻ 382.53

SMA 200 close ↻ 164.16

SMA 20 close ↻ 552.07



The Micron Beat-and-Scare Report

Sep 23rd: Record Q4 Earnings and MU's NVIDIA Moment

- **Revenue:** \$11.32B (46% YoY growth).
- **Gross Margins:** Surpassed **50%** for the first time since 2018.
- **The Stat that Changed the Narrative:** 56% of Micron's revenue now came from the Data Center, fundamentally decoupling the stock from the "boring" PC/Smartphone cycle.

The 2025 Memory Pivot

The memory sector transitioned from a traditional cyclical recovery to a **structural "Super-Bull" cycle** in late 2025.

Investors who identified the "September Pivot" captured a vertical rerating of the industry, driven by an unprecedented supply-demand imbalance that has now left the market "sold out" through late 2026... and probably beyond.

The "September Pivot": Why the Market Missed It

In Sep 2025, several independent variables—AI hyperscale hoarding, wafer “cannibalization,” and the “Crucial Exit” — converged to create a perfect storm.

The surge wasn't a gradual climb; it was a **repricing event**. In early September, memory was still viewed through a “recovery” lens. By October, it was a “shortage” lens.

3 Key Drivers of the Late-2025 Surge

1. The OpenAI/Stargate "Supply Grab" (Sep-Oct 2025): Rumors—later confirmed by procurement data—showed OpenAI and other LLM developers moving to secure up to **half of the projected 2026 HBM4 and DDR5 supply** through pre-payments. In a move dubbed the "Dirty DRAM Deal," OpenAI secured an estimated **40% of global DRAM wafer capacity** (~900,000 wafers/month) through 2026. This tactical cornering of the market forced hyperscalers into a "panic-buying" frenzy to secure remaining supply.

The “Dirty Dram Deal” (DDD)

In October 2025, internal procurement leaks revealed a "Direct-to-Fab" prepayment model that bypassed traditional supply chain visibility.

- **The \$20B Prepayment:** Microsoft/OpenAI reportedly issued \$20B in "non-refundable capacity reservations" to the Big Three (Samsung, SK Hynix, MU).
- **Market Blindspot:** Most investment advisors were looking at **PC and Smartphone weakness** (which was real). They missed that the "Stargate Grab" was large enough to offset a total collapse in the consumer sector.
- **The Multiplier:** Every HBM4 module produced for Stargate consumed the equivalent silicon of **four** standard DDR5 server sticks.



Kevin Cook

7

Oct '25 • NVIDIA

Save Post



A Dinner to Remember: NVIDIA, Samsung, & Hyundai

As per usual, Jensen runs the world and makes the important deals across borders.

Last night's dinner in South Korea during the APEC Summit may go down in tech history as the most important of 2025. Dr. Singularity on X noted "When the CEOs of Nvidia, Samsung, and Hyundai meet and have a shot, you know the future is being negotiated."

Jensen Huang, CEO of NVIDIA, Jay Y. Lee, Chairman of Samsung Electronics, and Euisun Chung, Chairman of Hyundai Motor Group, got together for a "chimaek" (chicken an... [See more](#))



Jump to latest comment



SKC

+6 subsidiaries

SK discovery

+5 subsidiaries

SK gas

+6 subsidiaries

SK Inc.

SK innovation
E&S

+11 subsidiaries

SK siltron

+1 subsidiary

SK biopharmaceuticals

+2 subsidiaries

SK ecoplant

+2 subsidiaries

SK geo centric

SK Inc.
C&C

SK networks

+7 subsidiaries

SK square

+15 subsidiaries

SK broadband

+1 subsidiary

SK energy

SK hynix

+12 subsidiaries

SK Inc.
materials

+6 subsidiaries

SK on

+6 subsidiaries

SK telecom

+7 subsidiaries

SK chemicals

+3 subsidiaries

SK enmove

+6 subsidiaries

SK ie technology

SK innovation

+7 subsidiaries

SK pharmteco

+5 subsidiaries

3 Key Drivers of the Late-2025 Surge

2. The transition to HBM4 acted as a supply vacuum: Because HBM requires approximately **3x the wafer area** of standard DRAM, every bit of AI memory produced removed three bits of conventional capacity from the market.

Inventory Capitulation: Many analysts expected high inventory levels to cap price growth. Instead, **DRAM spot prices tripled** between September and November as "just-in-time" supply chains collapsed.

ISPPDR48 17.2500 +1.3100 17.2500 / 17.2500
At 5:01 Op 17.2500 Hi 17.2500 Lo 17.2500 Prev 15.9400

ISPPDR48 Index 94 Vorgeschlagene Charts 96 Aktion 97 Änd. Linien-Chart

11/03/2024 - 11/03/2025 Letzter Preis Lokalwähru MOVs Wichtige Ergebnisse

1T 3T 1M 6M YTD 1J 5J Max. Täglich Tabelle + Verwandte Daten hinzufügen Chart änd.

- Letzter Kurs 17.2500
- Hoch am 11/03/25 17.2500
- Durchschnitt 5.4017
- Tief am 02/06/25 4.3500

Fadenkreuz Anmerken News Zoom



Kevin Cook · Dec '25



From a trader I follow...

DRAM spot per BBG - this is insane, prices are rising so fast that Samsung stopped quoting

Micron , Hynix ... life is good in the memory sweet spot ... and Hynix is already sold out for FY26 in HBM, DRAM and NAND



2×16 GB DDR5 6000 CL30 — Average price (daily mean)



Kevin Cook · Dec '25

From Dec 3, translated from French...

It's a real hold-up that's currently taking place on DDR4 and DDR5 RAM memories and NAND Flash memories, whose prices have already tripled since September, with peaks of over 500% expected by next summer.



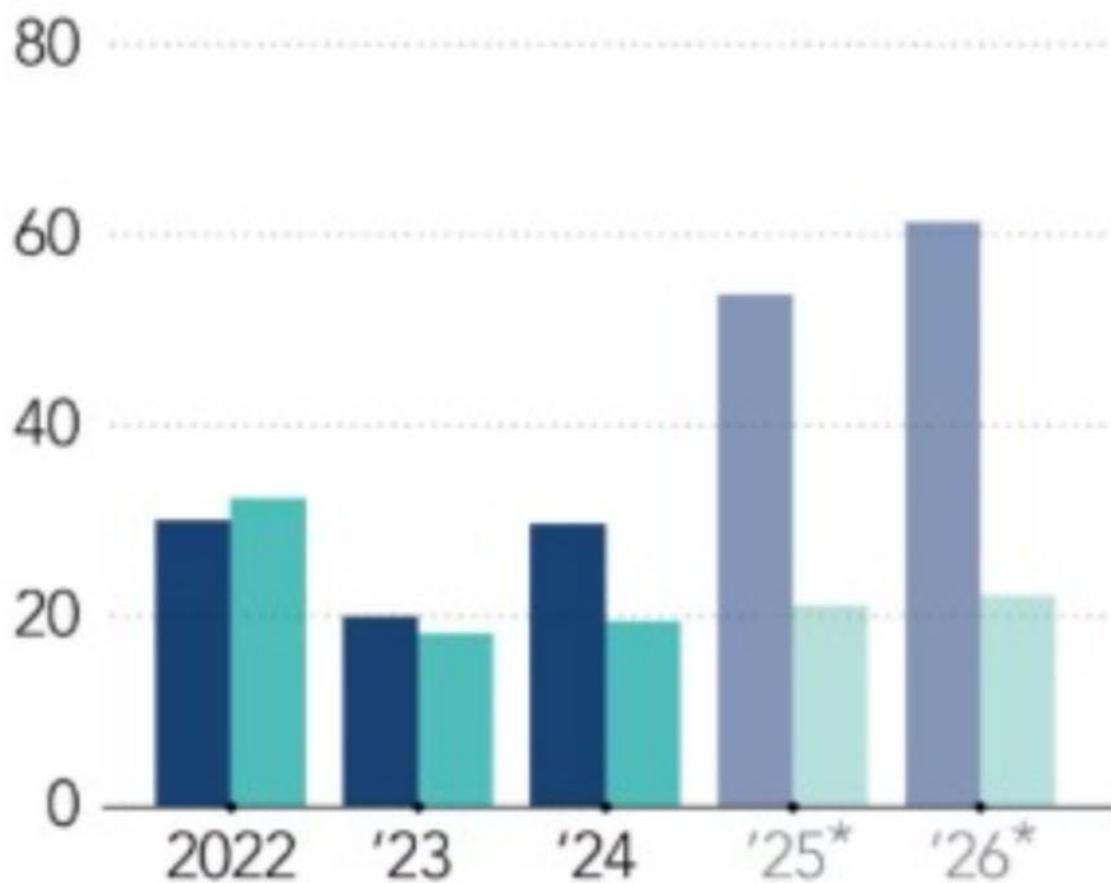
Impact on the "Big Three" Memory Ecosystem

Player	The "Stargate" Role	Market Realization (Q4 '25)
SK Hynix	The Sole Source (initially) for 12-Hi HBM3E.	Investors realized Hynix had 100% of its 2026 HBM capacity sold out by Nov 1st.
Micron (MU)	The "Swing Producer."	Their HBM4 breakthrough in late Sept allowed them to capture the overflow from Hynix, leading to the 42% Q4 stock rally.
Samsung	The Capacity King.	After failing early HBM3E yields, Samsung's Q4 pivot to "Massive-Scale NAND" for Stargate's storage layer saved their fiscal year.

Memory chip capital expenditure is growing sharply

(In billions of dollars)

■ DRAM ■ NAND



*Forecast

Source: TrendForce

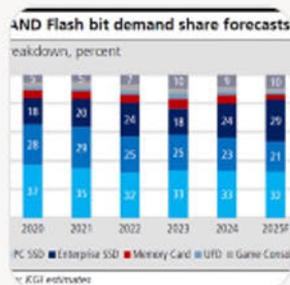
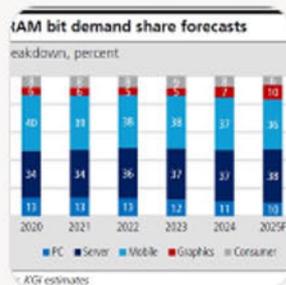
Trendforce revised up its DRAM and NAND contract price forecast again in November, now forecasting blended DRAM price to increase by 30-35%/ 20-25% QoQ and blended NAND price to increase by 25-30%/ 20-25% QoQ in 4Q25/1Q26 respectively:

<https://substack.com/profile/97916202-ayz/note/c-174552008>

Jump in DRAM/NAND demand for servers next year (2026):

38% to 47% for DRAM

29% to 40% for NAND

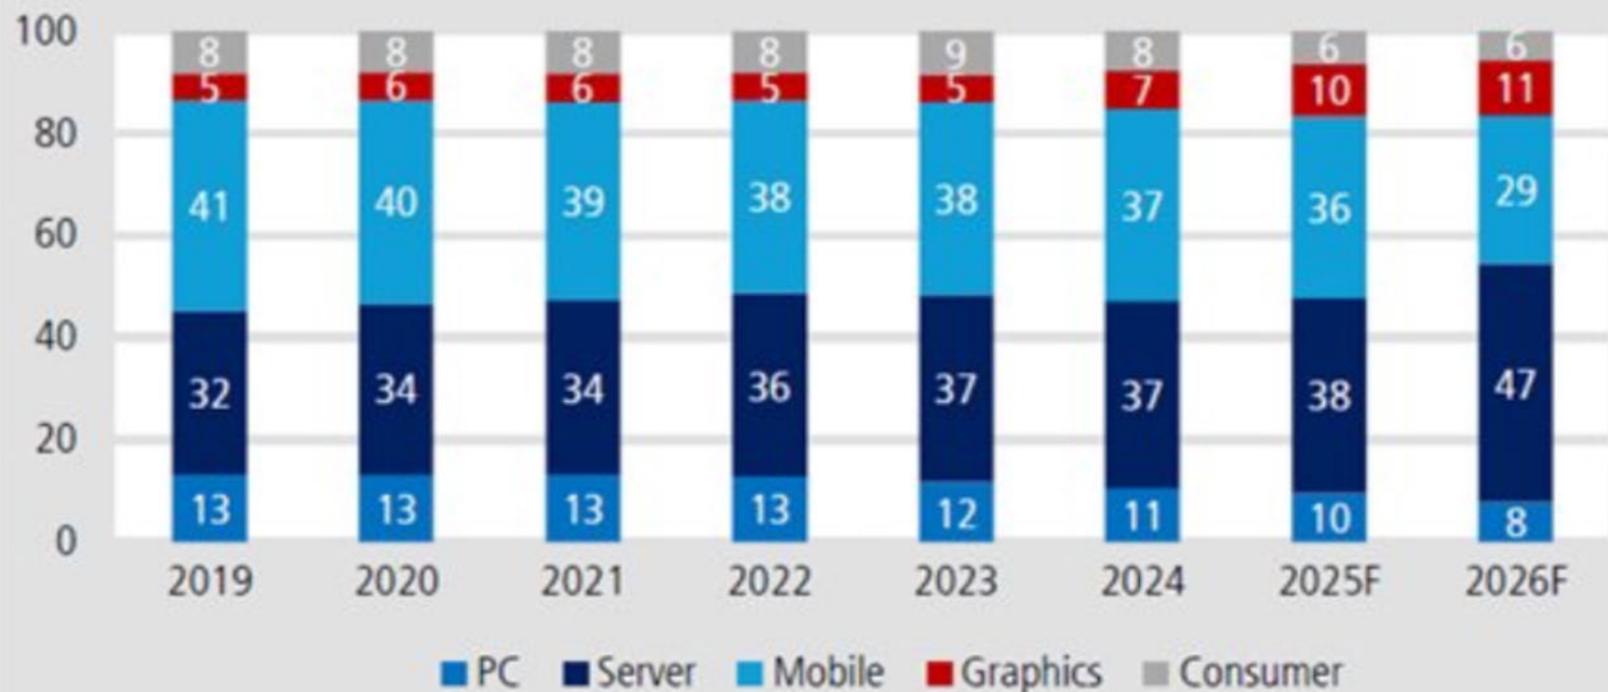


QOQ	Revised %		Original %	
	4Q25F	1Q26F	4Q25C	1Q26C
PC	15% up (17.5%)	1.5% up (18.5%)	15% up (17.5%)	15% up (17.5%)
Server	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)
Mobile	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)
Graphics	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)
Consumer	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)

QOQ	Revised %		Original %	
	4Q25F	1Q26F	4Q25C	1Q26C
PC SSD	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)
Enterprise SSD	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)
Memory Card in UFD	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)
Game Console	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)	15% up (17.5%)

Figure 3: DRAM bit demand share forecasts

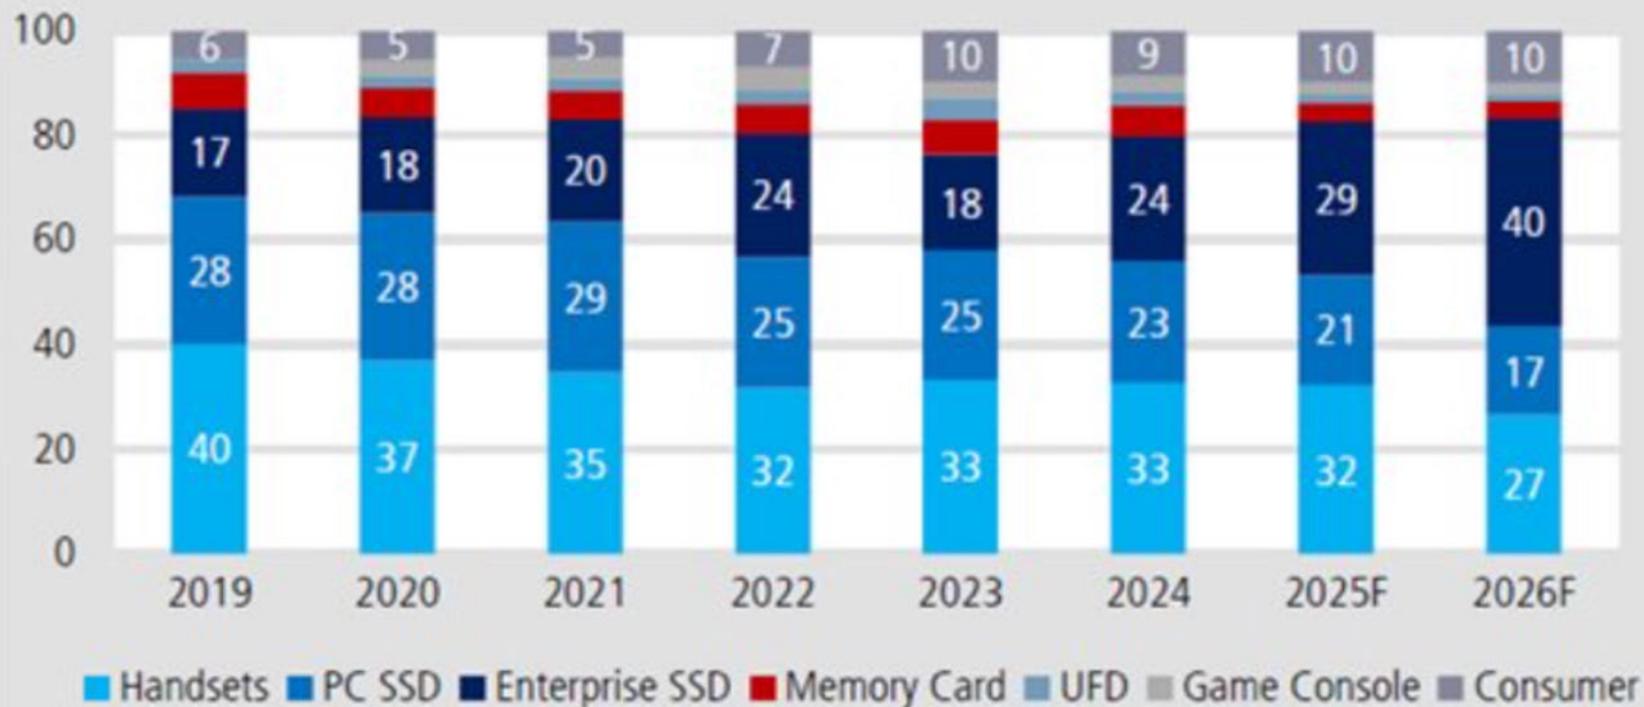
Bit demand breakdown, percent



Source: TrendForce; KGI estimates

Figure 4: NAND Flash bit demand share forecasts

Bit demand breakdown, percent



Source: TrendForce; KGI estimates

3 Key Drivers of the Late-2025 Surge

3. The “**Crucial Shock**” (Dec 2025): Micron’s decision to exit the consumer memory business (phasing out the legendary **Crucial** brand) signaled to the market that margins in the data center were so high that “commodity” retail was no longer worth the silicon. By abandoning the DIY/Retail market, Micron reallocated 100% of its leading-edge capacity to high-margin AI Data Center products, permanently tightening consumer supply.



Tom's Hardware @tomshardware · 3h



Micron outlines grim outlook for DRAM supply in first earnings call since killing Crucial memory and SSD brand — CEO says it can only meet half to two-thirds of demand



From tomshardware.com

The Spill-Over to Good Old Fashioned Storage

AI Cold Storage Demand: WDC's UltraSMR drives became the standard for "LLM Training Archives." AI needs massive amounts of training data, which lives on high-capacity HDDs (32TB+).

WDC at Nasdaq Investor Conference:

"One hyperscaler is talking to us about extending a contract to calendar year '29; Exabyte demand growth trending mid-20s% year-over-year, with Q3 hitting 30%, well above the 15% base and 23% AI-upside cases outlined at Investor Day."

Western Digital Corporation · 1D · NASDAQ ≈ O273.09 H289.66 L266.46 C281.58 -2.52 (-0.89%)

USD

281.58 SELL 0.00 281.58 BUY

Vol 9.9M
MA 50 close 215.16
SMA 200 close 122.14
SMA 20 close 261.43

WDC 281.58



Sep Oct Nov Dec 2026 Feb

320.00
240.00
200.00
160.00
120.00
80.00

⊞

iShares MSCI South Korea ETF • 1D • NYSE Arca

O131.69 H134.32 L128.33 C133.97 +3.17 (+2.42%)

USD

133.97
SELL
0.00
133.97
BUY

EWY 133.97

Vol 16.58M
MA 50 close 107.84
SMA 200 close 83.86
SMA 20 close 122.08

Sep Oct Nov Dec 2026 Feb



140.00
120.00
100.00
80.00
60.00



SK Hynix: Crowning a New Memory King

Mass production of 12-Hi HBM3E; 70% DRAM revenue growth.

HBM4 Readiness: World's first to complete HBM4 development; NVIDIA exclusive partner.

Became the "AI Memory King," beating Samsung in 2025 operating profit.

Samsung regained pricing power; forced 60% price increases across NAND/DRAM to manage demand backlog; announced HBM4 shipments for Nvidia *Rubin*.

tom's **HARDWARE** premium

[Tech Industry](#) > [Manufacturing](#) > [Semiconductors](#)

Samsung's Taylor, Texas fab could herald a breakthrough for the chipmaker, company plans 2026 risk production — new production flows, pellicles for EUV patterning as site targets 50,000 WSPM

50,000 wafer starts per month?

Kevin Cook · Dec '25



NVIDIA (NVDA) is teaming with SK Hynix and Phison on a new "AI SSD" ("Storage Next") targeting ~100M IOPS (Input-Output operations per second), roughly 10x current AI server SSDs.

The idea is a "memory-like" tier between DRAM/HBM and storage to ease inference bottlenecks. Prototype eyed for 2026.

HBM/DRAM are tight, and normal enterprise SSDs are still too slow / too high-latency for real-time AI workloads. So Nvidia wants a memory-like storage layer sitting between server memory and traditional storage.

Our man in Taiwan Dan Nystedt (@dnystedt on X) writes...

Nvidia is working with SK Hynix on a brand-new AI solid-state drive (SSD) with 10x the performance of current AI server SSDs, a project code-named "Storage Next." A prototype is expected in 2026. It's being designed as a memory layer between an AI server's memory and storage, while also preparing for the future needs of data centers in Space. The performance target of the new AI SSD is 100 million Input/Output operations per second (Wow!). Taiwan's Phison is a co-developer.

Sandisk: How the market finally valued FLASH correctly

SNDK became a pure-play bet on Enterprise SSDs, which saw price hikes of **90%** in Q4 as HDD shortages forced buyers toward NAND flash.



BARRON'S

Sign In

TECHNOLOGY

Sandisk Stock Soars 23%. The Reason Why Is Unclear.

By [Nate Wolf](#)

Jan 06, 2026 10:17 am EST

CES Jan 6: SNDK shares surge 27.5% (above \$300) on Jensen truth bombs

"With existing high-bandwidth memory (HBM), it's far from sufficient to support GPUs, and the memory bottleneck is only getting worse."

Huang unveiled the concept of a 'new memory storage platform.' In his slides explaining the new memory platform, SSDs were mentioned alongside HBM, which the industry interpreted as suggesting that SSDs could potentially be used.

Citi: We estimate that approximately 1,152TB of additional SSD NAND will be required per Vera Rubin server system to support NVIDIA's ICMS operations.

Accordingly, assuming Vera Rubin server shipments of 30,000 units in 2026 and 100,000 units in 2027, NAND demand driven by ICMS is projected to reach 34.6 million TB in 2026 and 115.2 million TB in 2027.

This represents 2.8% of expected global NAND demand in 2026 and 9.3% in 2027, a meaningful scale that is likely to create significant upside potential for demand.

As a result, with this structural demand increase factored in, the global NAND supply shortage is expected to intensify further.

MODELS

ARE

MEMORY

HOGS!



Rohan Paul  @rohanpaul_ai · Jan 10



New Google paper says LLM inference is limited by memory and networking, and it proposes 4 hardware shifts.

Inference is separated into Prefill, which processes input in parallel, and Decode, which outputs 1 token (a small text chunk) at a time while constantly reading a Key

[Show more](#)

Challenges and Research Directions for Large Language Model Inference Hardware

Xiaoyu Ma and David Patterson, Google

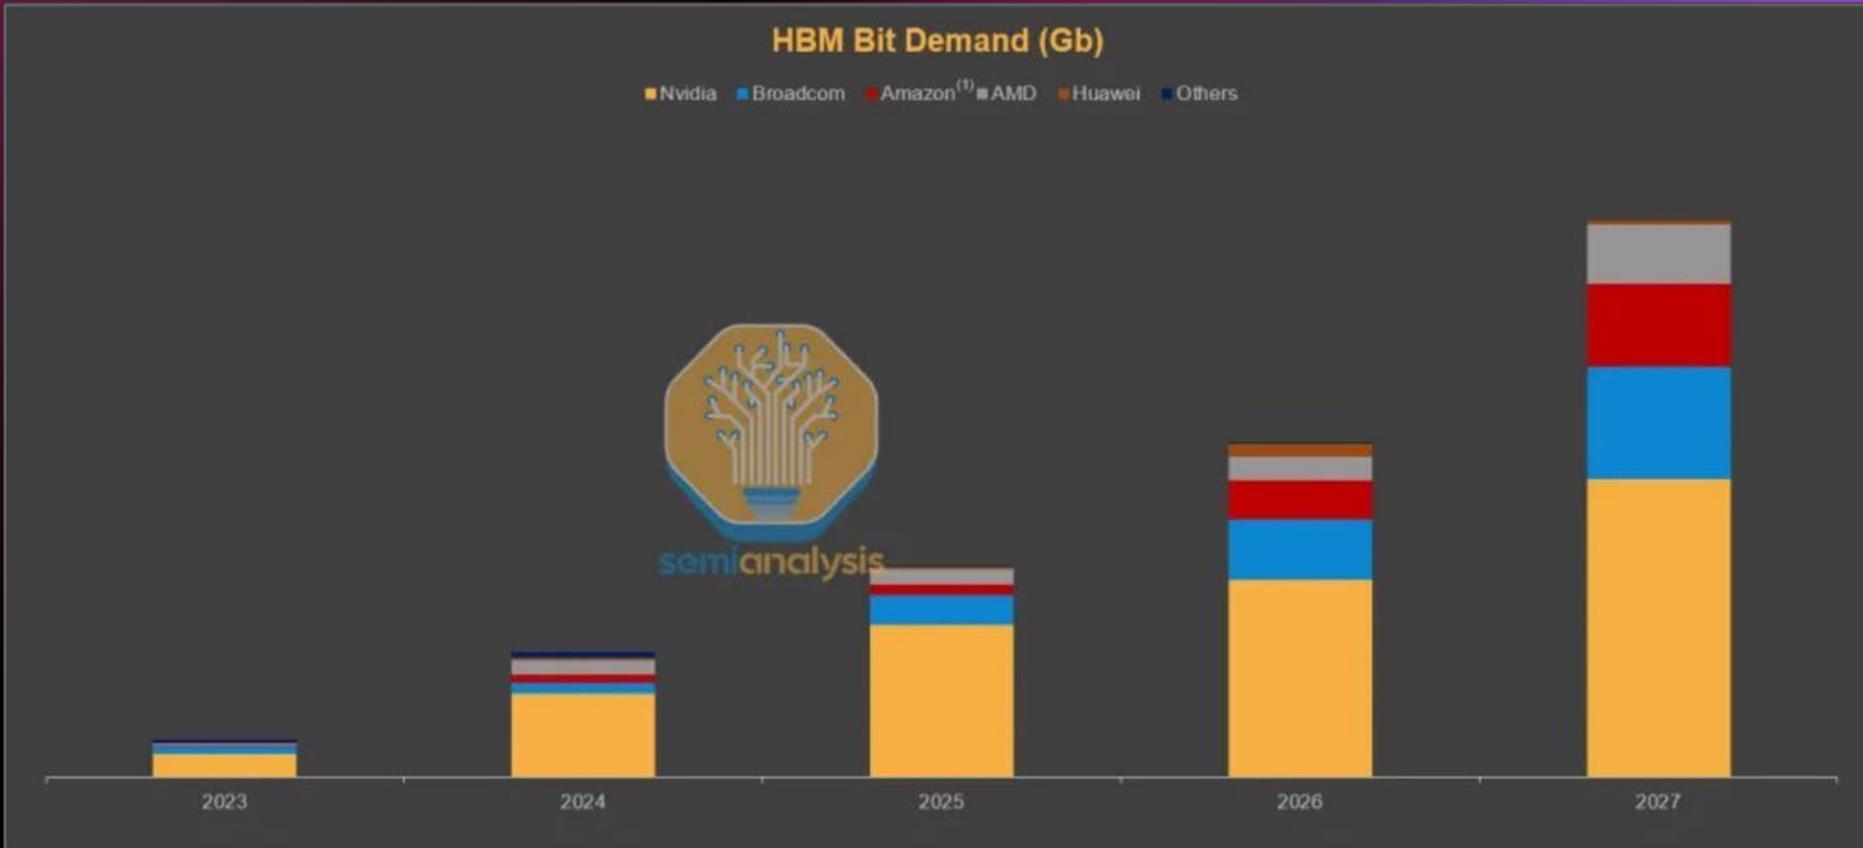
Large Language Model (LLM) inference is hard. The autoregressive Decode phase of the underlying Transformer model makes LLM inference fundamentally different from training. Exacerbated by recent AI trends, the primary challenges are memory and interconnect rather than compute. To address these challenges, we highlight four architecture research opportunities: High Bandwidth Flash for 10X memory capacity with HBM-like bandwidth; Processing-Near-Memory and 3D memory-logic stacking for high memory bandwidth; and low-latency interconnect to speedup communication. While our focus is datacenter AI, we also review their applicability for mobile devices.

Outlook: Why this supercycle has "legs" through 2027 (high capex lag)

The industry has moved beyond 'cyclicalities.' We are in a regime where memory is no longer a component, but the primary bottleneck for sovereign AI compute.

With prices up 90% since September, the barrier to entry for new AI players has shifted from 'Who has the GPUs?' to 'Who has the Wafers?'

HBM consumes 3-4x more wafers per bit than standard DRAM, forcing production capacity to scale exponentially to match surging annual demand. Thus Micron's bold move to build \$100 billion megafab in New York.



Note: (1) Amazon procures the own HBM instead of relying on their Custom ASIC partners.

Risk Factors: Power grid constraints and the "AI CapEx Bubble" debate.



Kevin Cook

1d · Chip R&D

Save Post



Cloud Capex Up Sharply, But Makes Market Even More Risky

The following article excerpt comes from LightCounting, the Semi/Optical research firm.

They have some ideas about the massive capex announcements from the hyperscalers that have particular read-through for the Semi/Optical industry wing.

This is short. Read it and consider that we may have to be more tactical with positions in the second half of 2026.

LightCounting releases a research note on the early financial results for Q4 2025...

by Vlad Kozlov, Stelyana Baleva, Carol Cao, Roy Rubenstein, Igor... [See more](#)



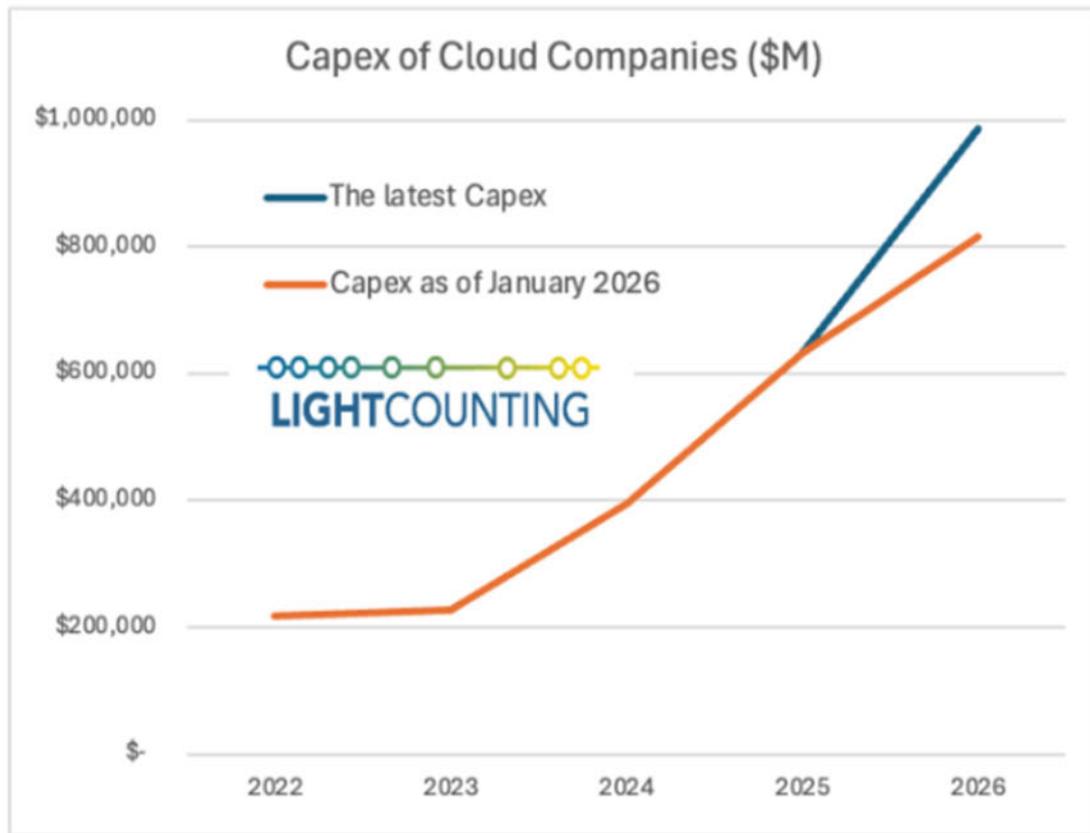
↓ Jump to latest comment

Risk Factors: Power grid constraints and the "AI CapEx Bubble" debate... plus SHORTAGE WOES

(Bloomberg, FEB 15, 2026) – A growing procession of tech industry leaders including Elon Musk and Tim Cook are warning about a global crisis in the making: A shortage of memory chips is beginning to hammer profits, derail corporate plans and inflate price tags on everything from laptops and smartphones to automobiles and data centers — and the crunch is only going to get worse.

Since the start of 2026, Tesla Inc., Apple Inc. and a dozen other major corporations have signaled that the shortage of DRAM, or dynamic random access memory — the fundamental building block of almost all technology — will constrain production. Cook warned it will compress iPhone margins. Micron Technology Inc. called the bottleneck “unprecedented.” Musk got to the intractable nature of the problem when he declared Tesla is going to have to build its own memory fabrication plant.

Figure: Changes in 2026 capex of cloud companies



Source: Companies' publicly reported financials and estimates



Kevin Cook

27d (edited) · Chip R&D

Save Post



High Bandwidth Flash (HBF): Looking where the memory puck is going

I've shared frequent Semi intel from X account "Jukan" (@jukan05 on X). Bio reads: Analyst @Citrini7 | Korean. He obviously does deep research for Citrini based on his voluminous posts.

And he will provide a good primer for us on HBF. Because I'd rather start with a human involved in the industry than an LLM.

Brief: (based on summary notes from Jukan)

💡 HBF is a new memory architecture created by vertically stacking 3D NAND flash. It is similar to HBM, which stacks DRAM, in terms of its stacked str... [See more](#)



Jump to latest comment

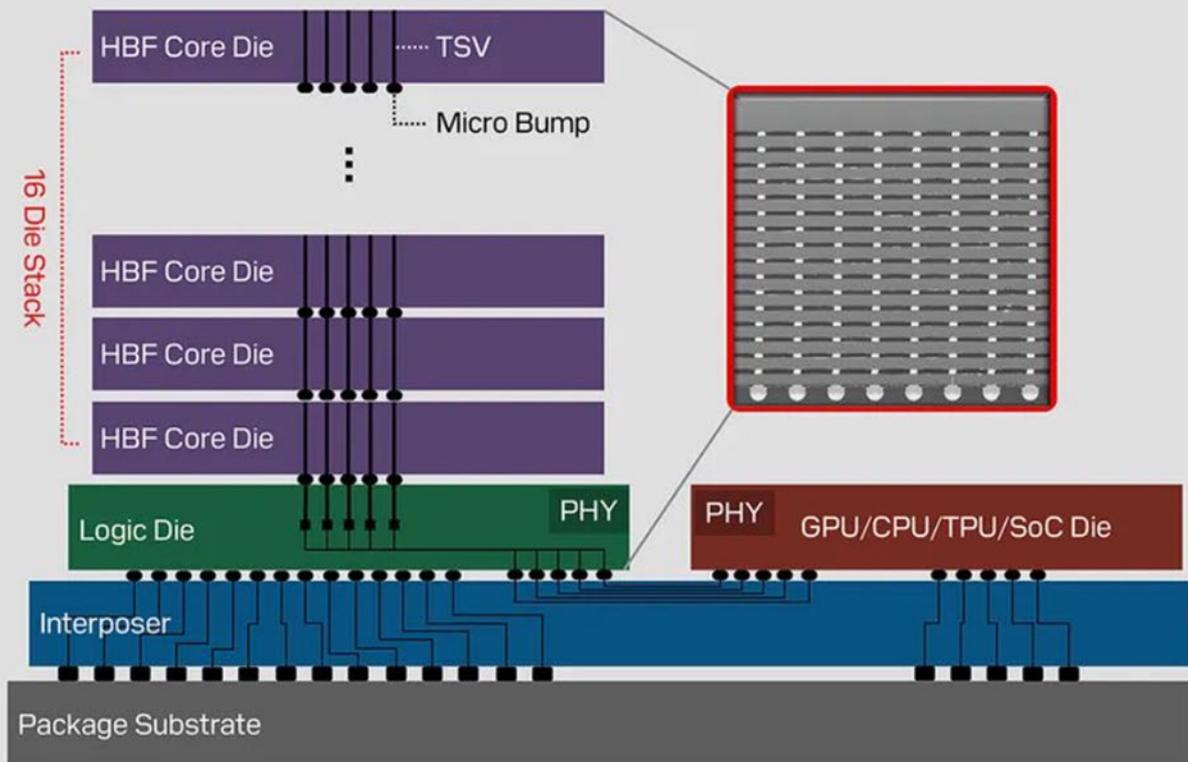
SanDisk and SK hynix are collaborating to standardize High Bandwidth Flash (HBF) for AI inference workloads in data centers, enterprises, and edge devices.

Technology: HBF is a NAND-based memory technology that aims to offer up to 8-16 times the capacity of traditional HBM at a similar cost

HBF uses stacked layers of NAND flash memory, similar to the 3D stacking used in modern SSDs, and connects them to a logic die using Through-Silicon Vias (TSVs).

This architecture allows for a high-bandwidth connection to the GPU, much faster than traditional PCIe-linked SSDs, but with much greater capacity and lower cost per gigabyte compared to HBM.

HBF STACK



Kevin Cook · 27d

Gemini 500-word brief on HBF

➔ Semiconductor Sector Analysis: The Rise of High Bandwidth Flash (HBF)

Date: January 16, 2026

💡 Subject: Bridging the AI Capacity Wall with HBF Memory

As we enter 2026, the semiconductor industry has reached a critical inflection point. While High Bandwidth Memory (HBM) remains the gold standard for peak performance in AI training, its capacity limits and soaring costs have created a "memory wall" for large-scale inference. Enter High Bandwidth Flash (HBF)—a transformative architecture that combines the massive density of 3D NAND with the parallel interface of HBM.

💡 The Architecture: NAND Meets HBM

HBF is not merely faster flash; it is a fundamental architectural shift. Unlike traditional SSDs that communicate via the relatively narrow PCIe/NVMe bottleneck, HBF utilizes Through-Silicon Vias (TSVs) and 3D die stacking (typically 16-Hi stacks) directly on a logic base die. This allows HBF to be placed on the same interposer as the GPU or AI accelerator, mimicking the parallel interface of HBM.



[Jump to latest comment](#)

Kevin Cook · 1d (edited)

SK hynix Unveils AI Chip Architecture with HBF, Reportedly Boosts Performance per Watt by Up to 2.69x

from TrendForce

SK hynix has unveiled a new semiconductor architecture concept centered on High Bandwidth Flash (HBF), a memory technology that stacks multiple NAND flash dies. According to [Hankyung](#), the company recently outlined the concept—dubbed “H3”—in a paper presented at the Institute of Electrical and Electronics Engineers (IEEE). H3 refers to a hybrid architecture that integrates both HBM and HBF into a single design.

As the report notes, the H3 architecture places both HBM and HBF alongside the GPU responsible for computation, whereas in current AI chips—including NVIDIA’s Rubin platform set for release in the second half of this year—only HBM is positioned next to the GPU.



Jump to latest comment

To assess the feasibility of the H3 architecture, SK hynix conducted simulation tests. In these

Pure Storage (PSTG)



Kevin Cook

Jan 15 • Chip R&D

Pure Storage (PSTG): Opportunistic Notes from a Small Cap Strategist

They are outsourcing the headache of NAND management to \$PSTG. 7/ \$PSTG also killed the "Forklift Scam." ... \$PSTG is coming for the nearline market. 9/ But fear the



12



13



Last comment Jan 16



Jorge Eduardo Rodriguez claret 1 like • Jan 16

If we assume PSTG has a bright future, why does ZACKS' ranking give it a 4 (Sell)???

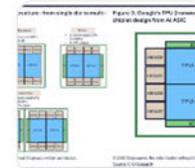


Kevin Cook

Jan 12 • Chip R&D

Memory Shortage Updates: MU, SNDK, PSTG

Micron on the DRAM Shortage: "Memory shortages won't improve before 2028, despite massive investments." AI and data center demand has raised the enterprise/data



12



141



Last comment 3h ago