

## Navigating the 2026 Memory Supercycle

- We projected the Memory shortage to last through 2027E–2028E.
- This cycle will run parallel to emerging bottlenecks in Power and Liquid Cooling.
- Buy **MICRON01** (TP THB 9.80) to capture this multi-year margin expansion.

### The Bottleneck Has Moved

The "Compute Shortage" of 2023 is effectively over. The data center industry has now slammed into a harder, more structural wall: **The Memory Supercycle**. As of 1Q26, the primary bottleneck for AI deployment has shifted from the GPU (Logic) to the HBM (Bandwidth) required to feed it. This is not a standard cyclical upturn; it is a supply shock driven by a fundamental restructuring of the global supply chain. While the market was fixated on securing NVIDIA chips, the manufacturing capacity for the memory required to run them was quietly deteriorating, creating a vacuum that is now driving prices to historic highs.

### The Mechanism: The "3-to-1" Wafer Penalty

The driver of this shortage is **Capacity Cannibalization**. The physics of AI manufacturing are brutal: industry data confirms that producing **one wafer of HBM (High Bandwidth Memory) consumes the same factory capacity as three wafers of standard DDR5**. To meet the insatiable demand for NVIDIA's Blackwell architecture, manufacturers have been forced to convert their production lines, effectively "deleting" ~30% of the world's capacity for standard server DRAM. This "Zero-Sum" environment means that every AI cluster deployed actively shrinks the supply of memory available for the rest of the economy, triggering a severe scarcity in standard DDR5 and enterprise storage.

### Chasing the Bottleneck

This Supercycle is the third phase of a sequential "Chain Reaction" that began with ChatGPT. The industry is not solving bottlenecks; it is merely pushing them downstream. We have moved from the "**Silicon Panic**" of 2023 (GPU shortages), through the "**Storage Trap**" of 2024 (HDD shortages), and have now arrived at the "**Memory Wall**" of 2026. Understanding this "Law of Constraint Migration" is critical: capital flows to the tightest point in the chain, and right now, that point is Memory.

### The Future: The Physical Wall (2027E+)

While investors chase memory today, the smart money is already positioning for the next, more immutable barrier: **The Physics of Power**. As rack densities breach 100kW, the industry faces a violent rotation toward **Liquid Cooling**, creating a new breakage point in the supply of Coolant Distribution Units (CDUs) and transformers. You can buy a GPU in 52 weeks, but you cannot build a substation in less than three years. This "Physical Wall" will define the winners of the next decade, shifting value from chip designers to infrastructure builders like Vertiv and Eaton.

### The Verdict: Buying the Scarcity Premium

We are currently in the "Sweet Spot" of the Memory Supercycle, where pricing power has consolidated entirely into the hands of the manufacturing oligopoly. The capacity is gone, the order books are full, and the margins are expanding. We believe the most actionable trade is to own the "First Mover" in this space. **We recommend MICRON01 (DR) with a TP of THB 9.80** (MU at \$500, USD/THB at THB31.3).

#### Analyst

Suwat Sinsadok, CFA, FRM, ERP  
 suwat.s@globlex.co.th  
 +662 687 7026

#### Assistant Analyst

Peerayu Sirivorawong

## The 2026 Memory Supercycle

The "Compute Shortage" of 2023 is effectively over. The data center industry has now slammed into a harder, more structural wall: **The Memory Supercycle**.

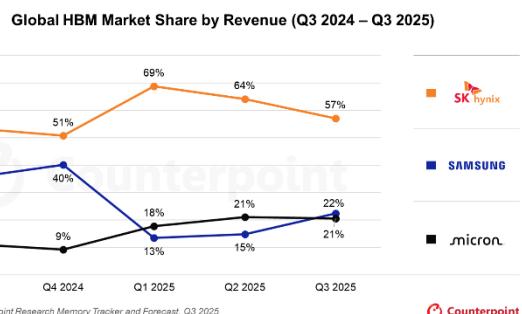
As of 1Q26, the primary bottleneck for AI deployment has shifted from the GPU (Logic) to the HBM (Bandwidth) required to feed it. This is not a standard cyclical upturn; it is a supply shock driven by **manufacturing cannibalization**.

The physics of this shortage are brutal: **Producing 1 wafer of HBM (for AI) consumes the factory capacity of 3 wafers of standard DDR5.**

To meet the insatiable demand for NVIDIA's HBM3e, manufacturers have been forced to convert their production lines, effectively "deleting" 30% of the world's capacity for standard server DRAM. This trade-off has created a "Zero-Sum" manufacturing environment. Every wafer allocated to an AI cluster is three wafers removed from the broader server market, triggering a severe scarcity in standard DDR5 and enterprise storage.

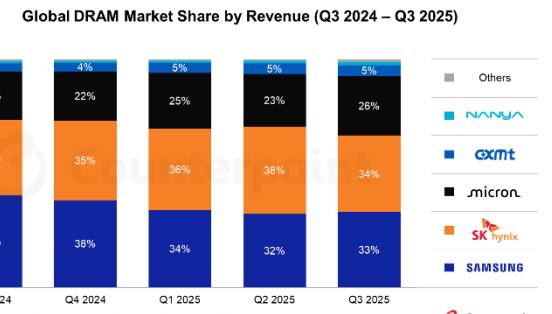
The winners of this dislocation are no longer the chip designers, but the manufacturing oligopoly: **Samsung Electronics, SK Hynix, and Micron Technology**. By exercising strict supply discipline and prioritizing high-margin AI wafers, these three players have engineered a pricing environment that benefits them exclusively. We believe this Supercycle is still in its early stages, with pricing power firmly consolidating in the hands of this "Big Three."

**Exhibit 1: Global HBM Market Share by Revenue**



Sources: Counterpoint

**Exhibit 2: Global DRAM Market Share by Revenue**



Sources: Counterpoint

This Memory Supercycle is not an isolated event. It is the third phase of a sequential supply chain failure that began with the launch of ChatGPT.

The AI infrastructure build-out follows a brutal "**Law of Constraint Migration**." In a system growing at exponential speed, capital does not solve bottlenecks; it merely pushes them downstream. As soon as the industry solved the GPU shortage (Logic), the pressure immediately transferred to the next weakest link (Memory), and is already beginning to fracture the layer beneath that (Power).

To forecast the duration of this cycle, we must understand the "Whack-a-Mole" dynamic that has defined the last three years. The industry is not fixing the supply chain; it is chasing the bottleneck.

## Timeline of Scarcity: A Rolling Crisis (2023–2026)

The market has not faced a single shortage, but rather a relay race of constraints. As capital floods into one bottleneck to solve it, the pressure immediately exposes the next structural weakness in the chain.

### Phase 1: The Silicon Panic (2023 – Now)

**The Bottleneck:** Compute (Logic)

**The Constraint:** NVIDIA H100 (Hopper) GPUs.

**What Happened:** Following the launch of ChatGPT, demand for training clusters exploded. The choke point was not the GPU die itself, but TSMC (chip manufacturer) could not package chips fast enough (CoWoS bottleneck). Lead times extended to 52+ weeks.

**The Resolution:** TSMC aggressively expanded CoWoS capacity, allowing H100 supply to catch up by late 2024. However, this "resolution" is a mirage: while you can now buy an H100 easily, the shortage simply migrated to the next generation. NVIDIA's new B200 (Blackwell) is currently sold out for 12+ months, meaning if you want state-of-the-art compute, you are back in the same queue as 2023.

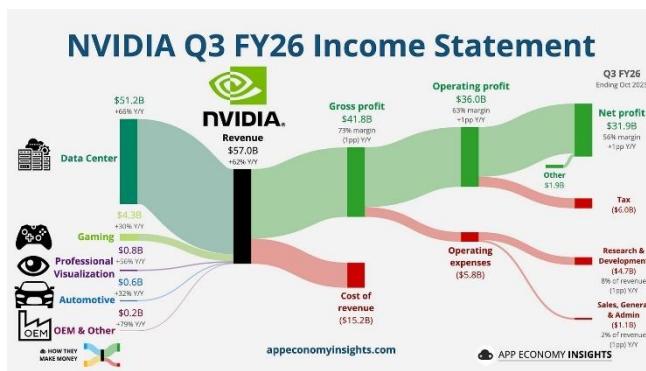
**Primary Beneficiaries:**

**NVIDIA (DR: NVDA80):** Selling the scarce asset.

**TSMC (DR: TAIWAN19):** The monopoly on packaging.

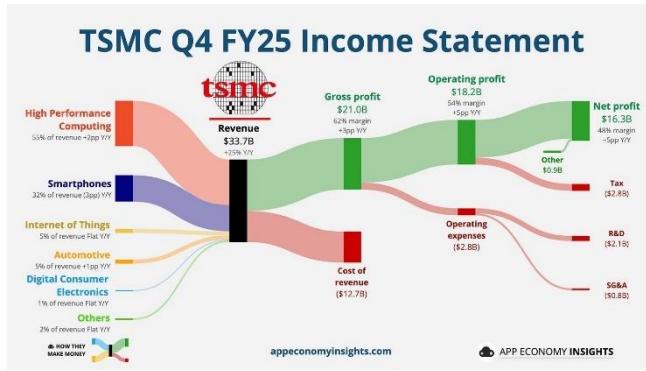
**CoreWeave (CRWV):** The "Neocloud" that stockpiled GPUs early and rented them at a premium.

### Exhibit 3: NVIDIA's 3Q26 Income Statement



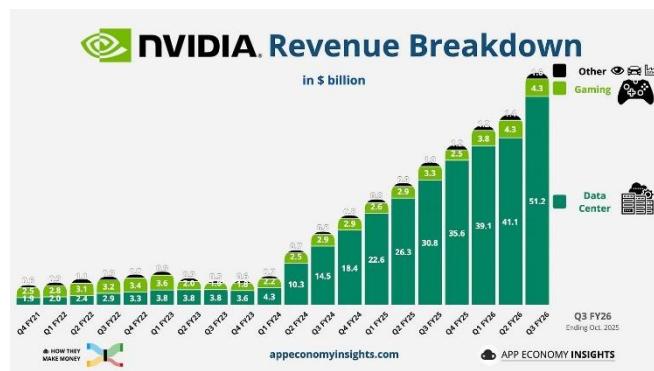
Sources: App Economy Insights

### Exhibit 5: TSMC 4Q25 Income Statement



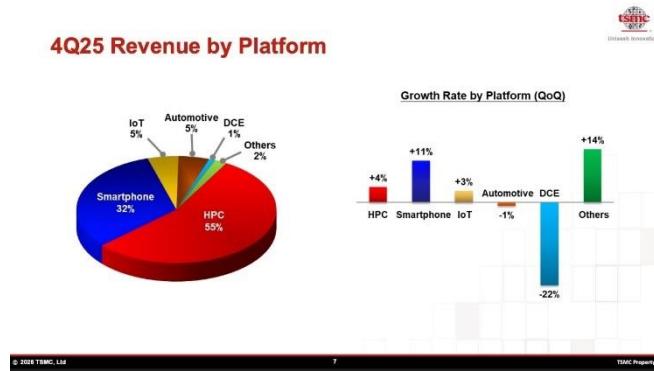
Sources: App Economy Insights

### Exhibit 4: NVIDIA's Revenue Breakdown



Sources: App Economy Insights

### Exhibit 6: TSMC 4Q25 Revenue by Platform



Sources: TSMC

## Phase 2: The Storage Trap (Late 2024 – 2025)

**The Bottleneck:** Capacity (Storage)

**The Constraint:** Nearline HDDs (22TB+).

**What Happened:** This shortage was self-inflicted by a massive capital diversion.

While the entire industry fought for GPUs in 2023, the HDD market collapsed. Cloud giants (AWS/Azure) paused all storage procurement to divert their entire CAPEX budget to NVIDIA H100s.

Facing ruin, the HDD duopoly (Seagate & Western Digital) entered survival mode. They permanently shut down factories (including Seagate's Suzhou plant) and slashed global production capacity by ~30% to stop the bleeding.

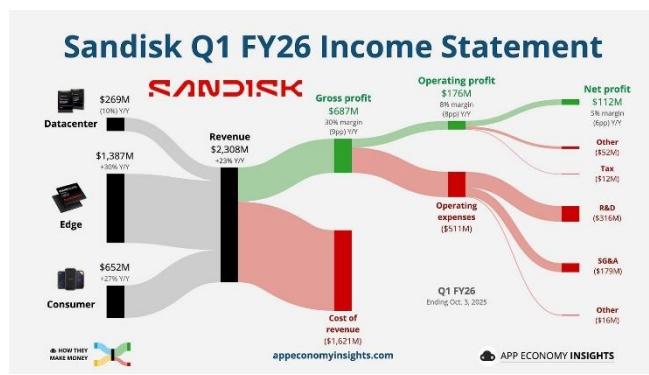
By late 2024, companies finally received their GPUs, only to realize their AI models were orders of magnitude larger than expected. Training these models generated Exabytes of "checkpoint" data and logs that had to be stored immediately.

Demand violently returned to a supply chain that had been dismantled. With no "buffer stock" left in the channel, lead times for high-capacity drives ballooned from **4 weeks to 52 weeks**

**Primary Beneficiaries:**

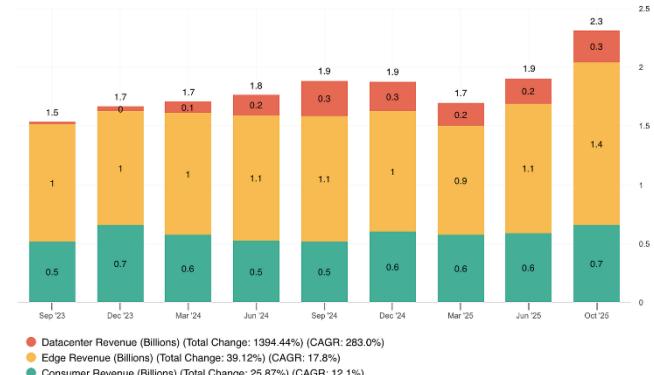
**Seagate (STX) & Western Digital (WDC) & Sandisk (SNDK):** Saw massive margin expansion as they utilized "Build-to-Order" discipline to force higher prices on desperate cloud providers.

Exhibit 7: Sandisk 1Q26 Income Statement



Sources: App Economy Insights

Exhibit 8: Sandisk Revenue Breakdown



Sources: App Economy Insights

### Phase 3: The Memory Supercycle (Current: 2025 – Now)

**The Bottleneck:** Bandwidth (Speed)

**The Constraint:** HBM3e (High Bandwidth Memory) & Standard DDR5.

**What Happened:** As outlined in the introduction, this is a "Zero-Sum" manufacturing crisis driven by physics.

Producing HBM is exponentially more complex than standard RAM. It requires vertical stacking and "Through-Silicon Vias" (drilling holes through the chip).

To chase the massive profit margins of HBM (sold to NVIDIA and Hyperscalers), manufacturers like SK Hynix and Samsung have converted their standard production lines into HBM lines.

By prioritizing AI, they have effectively "deleted" ~30% of the world's capacity to make standard server RAM.

The shortage is now structural. **Micron CEO Sanjay Mehrotra** recently admitted the company can currently meet only **50% to 66% of demand** for its top customers.

**The Consequence:** We are now in a dual-shortage. HBM is strictly allocated to Hyperscalers (Microsoft/Meta) until 2027E. Meanwhile, the "crowding out" effect has left the rest of the world fighting for a shrinking supply of standard DDR5, sending prices soaring +60% q-q in 1Q26 (now 3x-4x vs the normal price).

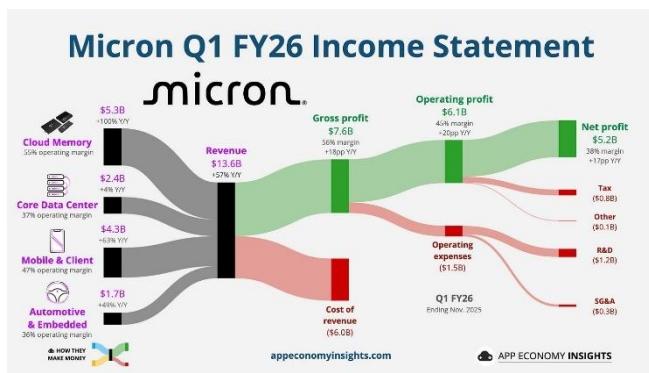
#### Primary Beneficiaries:

**Micron (DR: MICRON01) & SK Hynix (KRX: 000660):** Both companies have seized a critical "First-Mover Advantage" over Samsung by qualifying their HBM3e chips faster with NVIDIA.

Micron has sold out its supply for all of 2026. In response, they raised their FY2026 CapEx forecast to **\$20b** (up from \$18b) to aggressively build new capacity in Idaho and New York

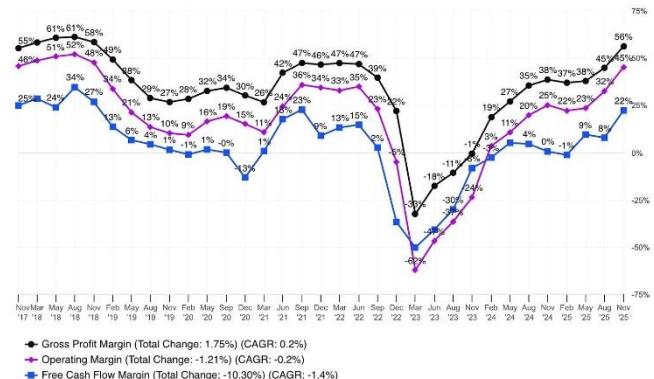
**Samsung (KRX: 005930):** As the market leader in volume, they lost the initial HBM sprint. However, they remain the only player with enough idle capacity to absorb the "overflow" demand that Micron and Hynix physically cannot fulfill.

Exhibit 9: Micron's 1Q26 Income Statement



Sources: App Economy Insights

Exhibit 10: Micron's Margin Trend



### The Future Outlook: The "Physical Wall" (2027E-2028E)

While the market is currently fixated on memory prices, the "Smart Money" is already positioning for the next, more immutable barrier: **The Physics of Power.**

#### The Bottleneck A: The Thermal Cliff (Liquid Cooling)

**The Constraint:** Coolant Distribution Units (CDUs) & Connectors.

**What we are expected:** The industry is undergoing a "Violent Rotation." Air cooling physically fails at rack densities above 50kW, but NVIDIA's Blackwell racks run at 100kW-120kW. This compels a 100% adoption rate of liquid cooling for high-end clusters.

The supply chain for CDUs (the heart of the loop) has fractured. Lead times have exploded from 12 weeks to 40+ weeks. Simultaneously, we are seeing a critical shortage of "Quick Disconnect" couplings—the aerospace-grade valves required to prevent catastrophic leaks.

#### Primary Beneficiaries:

**Vertiv (VRT):** The "Pure Play" on thermal management. They dominate the market for high-capacity CDUs.

**nVent (NVT):** The "Connector King." They specialize in the liquid manifolds and leak-proof couplings that are now mandatory for every AI rack.

#### The Bottleneck B: The Grid Cliff (Transformers)

**The Constraint:** High-Voltage Substations.

**What we are expected:** You can buy a GPU in 52 weeks, but you cannot build a substation in less than 3 to 4 years.

We are currently seeing completed data center "shells" sitting empty because the local utility cannot deliver the step-down transformers to energize them. This equipment requires massive amounts of copper and manual assembly; production cannot be easily automated or rushed.

#### Primary Beneficiaries:

**Eaton (ETN) & Schneider Electric (SU):** The global duopoly controlling the supply of medium-and-high voltage switchgear. They effectively own the queue for grid connection.

#### The Bottleneck C: The Generation Cliff (Baseload Power)

**The Constraint:** Gigawatts of Electricity.

**What we are expected:** The US electrical grid is expanding capacity at <0.5% per year, while AI power demand is projected to grow 160% by 2030. Renewables (Solar/Wind) are too intermittent for AI training runs that require 24/7 uptime.

This has forced Hyperscalers to bypass utilities entirely. Microsoft's deal to reopen Three Mile Island and Amazon's acquisition of nuclear-powered campuses prove that Nuclear Baseload is the ultimate strategic asset.

#### Primary Beneficiaries:

**Constellation Energy (CEG) & Vistra (VST) & Oklo (OKLO):** The owners of the nuclear fleet. They are the only ones with the "always-on" clean power that Hyperscalers are desperate to buy.

## Investment Strategy: Seizing the Scarcity Premium

The liquid cooling crisis is a story of future infrastructure; the memory crisis is a story of **immediate cash flow**.

We are currently standing at the foothills of the **Memory Supercycle**. This is not a hypothetical shortage; it is a mathematical certainty driven by the "3-to-1" wafer penalty. As capacity is physically deleted to serve AI demand, pricing power has shifted entirely to the manufacturers.

We believe we are still in the **early stages** of this dislocation, where the "Oligopoly of Three" will enjoy historic margin expansion that the market has **not yet fully priced in**.

**Investment Strategy: Seize the Current Constraint** The most actionable trade today is not to guess the winner of the future power wars, but to own the clear winner of the current bandwidth scarcity.

**Top Pick: MICRON01 (DR).** We recommend **MICRON01** with a **TP of THB 9.80** (MU at \$500, USD/THB at THB31.3) as our top selection for this cycle.

**The Catalyst:** Micron is the "First Mover" in HBM3e, with its entire 2026 production capacity already sold out at premium rates.

**The Moat:** As the primary US-based manufacturer, it commands a "Sovereign Premium" from key hyperscalers like Amazon and Google.

**The Valuation:** With supply strictly capped and prices rising +60%, we expect significant earnings revisions in the coming quarters.

## GENERAL DISCLAIMER

### Analyst Certification

Suwat Sinsadok, Register No. 020799, Globlex Securities Public Company Limited

The opinions and information presented in this report are those of the Globlex Securities Co. Ltd. Research Department. No representation or warranty in any form regarding the accuracy, completeness, correctness or fairness of opinions and information of this report is offered by Globlex Securities Co. Ltd. Globlex Securities Co. Ltd. Accepts no liability whatsoever for any loss arising from the use of this report or its contents. This report (in whole or in part) may not be reproduced or published without the express permission of Globlex Securities Co. Ltd.

## RECOMMENDATION STRUCTURE

### Stock Recommendations

Stock ratings are based on absolute upside or downside, which we define as  $(\text{target price}^* - \text{current price}) / \text{current price}$ .

**BUY:** Expected return of 10% or more over the next 12 months.

**HOLD:** Expected return between -10% and 10% over the next 12 months.

**REDUCE:** Expected return of -10% or worse over the next 12 months.

Unless otherwise specified, these recommendations are set with a 12-month horizon. Thus, it is possible that future price volatility may cause temporary mismatch between upside/downside for a stock based on market price and the formal recommendation.

\* In most cases, the target price will equal the analyst's assessment of the current fair value of the stock. However, if the analyst doesn't think the market will reassess the stock over the specified time horizon due to a lack of events or catalysts, then the target price may differ from fair value. In most cases, therefore, our recommendation is an assessment of the mismatch between current market price and our assessment of current fair value.

### Sector Recommendations

**Overweight:** The industry is expected to outperform the relevant primary market index over the next 12 months.

**Neutral:** The industry is expected to perform in line with the relevant primary market index over the next 12 months.

**Underweight:** The industry is expected to underperform the relevant primary market index over the next 12 months.

### Country (Strategy) Recommendations

**Overweight:** Over the next 12 months, the analyst expects the market to score positively on two or more of the criteria used to determine market recommendations: index returns relative to the regional benchmark, index sharpe ratio relative to the regional benchmark and index returns relative to the market cost of equity.

**Neutral:** Over the next 12 months, the analyst expects the market to score positively on one of the criteria used to determine market recommendations: index returns relative to the regional benchmark, index sharpe ratio relative to the regional benchmark and index returns relative to the market cost of equity.

**Underweight:** Over the next 12 months, the analyst does not expect the market to score positively on any of the criteria used to determine market recommendations: index returns relative to the regional benchmark, index sharpe ratio relative to the regional benchmark and index returns relative to the market cost of equity.