

SECOND ORDER

AI, Capital & Work · A Second Order research briefing

ISSUE 05 · JULY 7, 2026

Memory Is the Capex Constraint

Hyperscaler capital spending is guided toward about \$635B in 2026, and a rising share of that number is memory price that buys no added capacity. High-bandwidth memory is a three-firm oligopoly, two of them Korean, which is why the AI trade now reprices through Seoul.

FOR DISCUSSION · NOT INVESTMENT ADVICE

One question: how much of the capex number is memory price?

THE THROUGH-LINE

Since Issue 01 the series has treated compute as the binding constraint. In 2026 that constraint is memory. SemiAnalysis puts memory at about 30% of hyperscaler capex, up from ~8% in 2023-24; Microsoft says ~\$25B of its 2026 capex is higher memory and chip prices. Part of reported capex growth is oligopoly memory price, and only the remainder is new capacity.

WHY NOW

Memory-share estimates circulated in late June and the complex sold off through the memory makers and the Korean index while the US megacaps moved less. On June 22 SK Hynix passed Samsung as Korea's most valuable company.

WHAT YOU WILL LEARN

1. The number: memory is a third of capex (sec. 01)
2. Why the share climbs on price as much as volume (sec. 01)
3. The three-firm HBM oligopoly (sec. 02)
4. The read-through to Seoul (sec. 03)
5. Two operator moves, a watchlist, and what would change our mind

IF YOU REMEMBER ONE LINE: a rising share of AI capex growth is memory price set by three firms, and two of them are in Korea.

THE CLAIM

The capex number every client quotes is partly a memory-price number

Microsoft attributes about \$25 billion of its raised 2026 capital spending to higher prices for memory and components, a price effect rather than additional machines.

Per CFO Amy Hood on Microsoft's fiscal-Q3 2026 earnings call, April 29, 2026 (as reported by CNBC, The Register, Motley Fool)

~30%

memory as a share of hyperscaler capex in 2026, up from ~8% in 2023-24 (SemiAnalysis estimate)

~\$25B

of Microsoft's raised 2026 capex is higher memory and chip prices, buying no more machines


~79%

of HBM bits ship from Korea; SK Hynix and Samsung, two of the three HBM makers (Q2 2025)

Memory has gone from 8% of capex to a third

EXHIBIT 1**Memory has gone from 8 percent of hyperscaler capex to a third in three years**

Memory as a share of hyperscaler capital spending, percent, by estimate source



Note: Both series are analyst estimates for a category (all memory, not HBM alone). 2026-2027 are forecasts, shown open.

Source: SemiAnalysis (CY23-CY24 ~8%, CY26 ~30%), posted Apr 2026; CLSA (35% 2026, 48% 2027), via Crypto Briefing Jun 23 2026.

SO WHAT


- SemiAnalysis puts memory at ~30% of hyperscaler capex in 2026, up from ~8% in 2023-24, a near-4x rise in four years.
- CLSA carries the same category to 48% by 2027.
- These are analyst estimates for all memory, not HBM alone. 2026-27 are forecasts.

The share climbs on price as much as volume

EXHIBIT 2

The big four are guiding to about 635 billion dollars of capex in 2026

Combined capital spending, Amazon, Microsoft, Alphabet, Meta, USD billions



Note: 2022-2025 actuals are approximate (methodologies differ). 2026 is summed company guidance; third-party aggregates run to ~\$725B.

Source: Second Order sum of company 2026 guidance midpoints; 2022-2025 from reported aggregates (Visual Capitalist, Tom's Hardware).

SO WHAT


- The big four guide to about \$635B of capex in 2026 (summed company guidance).
- DRAM contract prices rose 171.8% YoY in Q3 2025 (TrendForce).
- When a large input's price more than doubles, dollar capex grows even at flat capacity.

Three firms make the HBM, two of them Korean

EXHIBIT 3

Three firms supply all the HBM, and two of them are Korean

HBM bit-shipment share by vendor, Q2 2025 vs a year earlier, percent



Note: Third-party estimate. SK Hynix and Samsung are South Korean; Micron is US. Korea supplies roughly four-fifths of the bits.

Source: Counterpoint Research, via Mark LaPedus / Semiecosystem (Aug 5 2025). Q3 2025 prints shift toward Samsung (~22%).


SO WHAT

- Counterpoint puts SK Hynix at 62% of HBM bits in Q2 2025, Micron 21%, Samsung 17%.
- Only three firms can stack the 12-16 DRAM dies 2027 accelerators need (Baker, attributed).
- All three are sold out of 2026 HBM; new supply is not expected before 2028.

The AI trade reprices through memory and Seoul

EXHIBIT 4**The AI trade now reprices through memory, and much of it through Seoul**

Equity return since the August 2025 low, percent



Note: Investor-published chart. US memory is Micron; Korean memory is SK Hynix/Samsung. Benchmark points are approximate (read from chart).

Source: IO Fund / Beth Kindig (Jun 26 2026). Memory figures are the article's stated numbers; benchmark values read off the same chart.

SO WHAT

- Since the Aug 2025 low: US memory (Micron) >1,200%, Korean memory >700%, both far above the Nasdaq-100 (IO Fund).
- Samsung + SK Hynix were a record 42.2% of the KOSPI in May 2026.
- The June drawdown moved the memory makers and the Korean index more than the US megacaps.

Reported capex growth overstates the buildout

The first-order reading of a rising capex number is more capacity.

Part of the increase is price. Memory is about a third of the bill and in a supercycle, so a chunk of reported capex growth is the same physical order at a higher unit cost. Microsoft has quantified about \$25B of that for 2026; if memory is a third of the bill and its price has more than doubled, the gap between dollars spent and capacity added widens across the cohort.

For allocation, part of AI's contribution to growth through capex is an oligopoly transferring rents from the hyperscalers to three memory makers, and only the remainder is new compute in the ground.

The operator's and allocator's playbook

1. Net capex to capacity

Read the hyperscaler headline as a gross figure and subtract memory price before treating it as capacity. The physical buildout and the dollar total have separated.

2. Size the memory beta

Long the AI-capex complex is long a memory cycle set by three firms and short their pricing power. Size it deliberately and watch Korea.

3. Watch two companies

Samsung and SK Hynix are 42% of the KOSPI. The AI trade is now unusually sensitive to two firms in one country.

4. Price memory escalators

Suppliers guide the shortage into 2028 and sign five-year contracts. Write memory-price escalators into multi-year infrastructure plans.

5. Assume no unit relief

New fabs do not reach meaningful output before 2028 (Micron). Do not plan on falling memory unit cost before then.

6. Hedge with efficiency

Where a workload runs on less memory-hungry inference, that is a direct hedge against the most inflationary input in the stack.

WATCHLIST

What would change our mind

New supply before 2028: a faster Micron or SK Hynix ramp, or a Samsung yield recovery that adds a fourth HBM source and softens the floor.

Demand step-down: a hyperscaler capex cut, or an efficiency gain such as Google's TurboQuant that cuts memory per unit of inference.

DRAM contract series: a quarter that breaks the run of increases (171.8% YoY in Q3 2025; +93-98% QoQ in Q1 2026).

2026 capex actuals: the share independent trackers attribute to memory as estimates are replaced by reported figures.

Falsifier: a 2027 in which memory's capex share falls toward 2024 levels while dollar capex keeps climbing.

NEXT ISSUE · THE POWER BILL. The series follows the buildout past memory to electricity, and to whether the grid can deliver the megawatts the 2026 capex plans assume. We separate contracted data-center capacity from energized capacity.