NATHAN RESEARCH GROUP

MARKET STUDY
AI MEMORY
SEOUL · 2026-08

MARKET RESEARCH MONOGRAPH

The HBM Market where AI memory is actually made

A study of the global high-bandwidth-memory market — what it is, how it was built, what changed in mid-2026, where the forecast can honestly reach, and how a chain running from a DRAM wafer to a rack is held at three chokepoints, two of which sit in Korea.

2026 MARKET — SPECIALIST HOUSES	THE SAME YEAR — GENERALIST PUBLISHERS	SHARE OF TOP-THREE WAFER INPUT, 2027E	KOREAN SHARE OF GLOBAL HBM
$45–60bn	**~$4bn**	**~30%**	**>80%**
Goldman Sachs, BofA, Yole	Mordor, TBRC, R and M	TrendForce; ~18% end-2025	Samsung + SK hynix; 78–92%

Evidence base: research houses and trade bodies on market sizing, both house families preserved as separate bands and never averaged · SK hynix SEC Form 424B4 (July 2026 Nasdaq listing), carrying dated IDC and Gartner data under securities-law liability · JEDEC, SEMI and WSTS publications · results releases from SK hynix, Samsung Electronics and Micron · nine DART filings from five Korean back-end and design issuers. Every event between 2026-05-01 and 2026-08-10 was verified against a dated source.

NOT INVESTMENT ADVICE
NO PRICE TARGET
PREPARED 2026-08-10, SEOUL

Contents

THE ANALYTICAL SPINE

The spine of this study is temporal: the market as it stands, the five years that set the present positions, the inflection mid-2026 delivered, and the last year a forecast can honestly reach. Competitive structure follows, then the participants, each appearing as a position in the market at three points in time; three market environments and one transferable lesson close.

Sections, and the exhibits that carry them

§	SECTION	WHAT IT COVERS	EXHIBITS
—	**Executive Summary**	A claim on manufacturing capacity	E1–E2
1	**The Market Today**	Two house families, structure, value chain	E3–E8
2	**How This Structure Was Built**	The demand engine, 2022–2026	E9–E14
3	**What Is Changing Now**	May to August 2026	E15–E17
4	**The Next Five Years**	A supply forecast, and where the axis stops	E18–E21
5	**Competitive Structure**	Concentration, share bands, the chokepoint	E22–E24
6	**The Players Within the Market**	Three distances from one market	E25–E28
7	**Market-Environment Scenarios**	Three paths for the market itself	–
8	**What This Market Teaches**	The transferable lesson	–
—	**Working With Nathan Research**	The bench, and the open questions	–
—	**Partner With Nathan Research**	Contacts, and how an engagement starts	–
A	**Appendix**	Sources · Methodology · Glossary	–

HOW TO READ THIS DOCUMENT

This is a study of a market. Companies appear as the lens that makes it legible: their disclosures evidence how the market behaves, never what they are worth. There is no valuation, no price target and no buy or sell framing here, and where no credible figure exists an exhibit stops.

Executive Summary

A CLAIM ON
MANUFACTURING CAPACITY

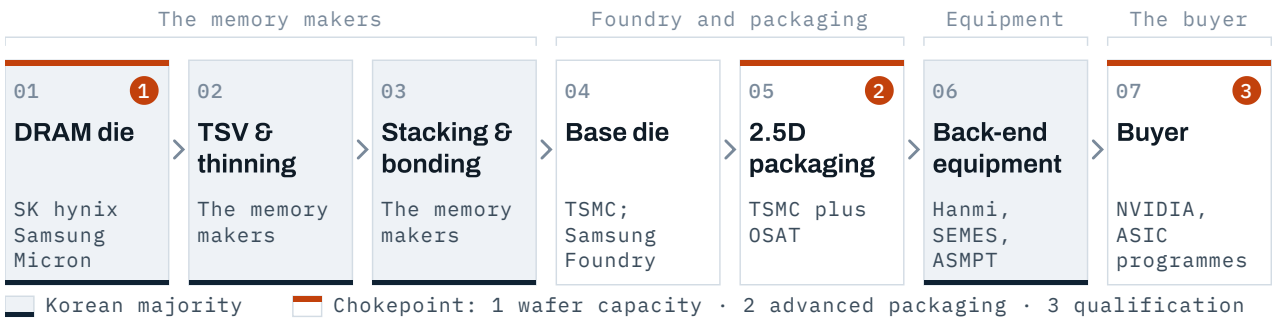
In 2026 high-bandwidth memory will take about **22%** of the DRAM wafer starts of the world's three largest memory makers and return about **9%** of the world's DRAM bits[1]. A fifth of the top three's wafer starts produces under a tenth of the world's bits.

HBM MARKET, 2026 — SPECIALIST HOUSES	THE SAME YEAR, GENERALIST PUBLISHERS	SHARE OF TOP-THREE DRAM WAFER INPUT, 2027E	KOREAN SHARE OF GLOBAL HBM
$45–60bn	**~$4bn**	**~30%**	**>80%**
Goldman, BofA, Yole	Mordor, TBRC, R and M	TrendForce; ~18% end-2025	Samsung + SK hynix

That ratio is what this market is. It sets which numbers carry information: wafer allocation, qualification dates and stack height say more than any revenue forecast, and 2027 becomes a supply question before a demand one.

Korean suppliers hold the majority position in four of the seven layers below — the die, the vias, the stack and the machines that do the bonding. The two they do not hold are the base die and the interposer, where a foundry stands between the memory and the accelerator; the seventh layer is the buyer. A market that is more than four-fifths Korean is therefore Korean where the product is made and foreign where it is finished.

EXHIBIT 1 **The market on one line — seven layers, three chokepoints, and where Korea holds the majority**



Sources: IDC, Q1 2026, via SK hynix SEC Form 424B4; TrendForce, 2026-03-20 and 2026-06-15; Asia Business Daily, 2025-12-22; industry compilations, 2026. Shading marks the layers where Korean suppliers hold the majority position, which is not the same as a layer being wholly Korean — layer 1 lists Micron alongside SK hynix and Samsung. Compilations of the Korean share of global HBM range from 78% to 92%.

WHY IT MATTERS Every quantity in this study attaches to one of these seven layers. E7 works through the chain layer by layer; E8 takes the three marked chokepoints in turn.

What the evidence establishes

FIVE FINDINGS
AND TWO OMISSIONS

Five findings survive the sourcing rules this study applies, each with the evidence that carries it.

EXHIBIT 2 **The five findings, and where each one is established**

FINDING	WHAT THE RECORD SHOWS	WORKED THROUGH IN
There is no agreed size for this market, and the disagreement is definitional	Specialist houses put 2026 between US$45bn and US$60bn; generalist publishers put the same year near US$4bn. Averaging them produces a number no house holds.	E5 Six houses
Beyond 2027 there is no market size this study can plot	Gartner's US$86bn for 2027 is the last specialist point on the record. Two later observations exist, and neither can end a plotted line.	E19 Gartner via SK hynix 424B4
Demand follows a physical ratio	Each accelerator generation must carry enough bandwidth to keep its own arithmetic busy, so HBM content per accelerator — 192GB on Blackwell B200, 288GB on Rubin, 384GB targeted for Rubin Ultra — is the forward variable that holds.	E10 NVIDIA and vendor materials
2027 is a supply forecast wearing a demand forecast's clothes	Bit shipments are projected to grow 50–60% next year and still fall short of demand growth, and NVIDIA has been evaluating shorter stacks than the ones it specified for Rubin Ultra.	E18 TrendForce, 2026-08-04
The tightest single-company hold sits outside the memory layer	The tightest hold by any single company sits in back-end equipment, where roughly seven-tenths of the world's thermocompression bonders come from one supplier.	E24 Asia Business Daily, 2025-12-22

The five do not add up to a forecast. Two of them say what the record cannot support; the three that remain are physical quantities rather than commercial ones, which is what lets the study be specific about 2027 and silent about 2031.

TWO NUMBERS THIS STUDY WILL NOT PRINT

Two numbers are not used anywhere in this study. The first is a single blended HBM market size for any year: the two house families measure different things, and the midpoint between them describes nothing. The second is any HBM market size beyond 2027 — in particular the "~US$170bn by 2031" figure circulating on aggregator sites. US$170bn is Yole Group's **2024 total memory market** (DRAM ~US$97bn plus NAND ~US$68bn) from its 2025-06-19 release, transposed onto HBM at a 2031 horizon by an aggregator, and it has been removed from the evidence base.

1.1 What is sold, and what it costs to make

SIZING
STRUCTURE
VALUE CHAIN

High-bandwidth memory is DRAM sold as a **stack**, and the stack is the unit of trade. Dies are thinned, drilled with through-silicon vias and bonded onto a logic base die; the tower stands beside an accelerator on a silicon interposer.

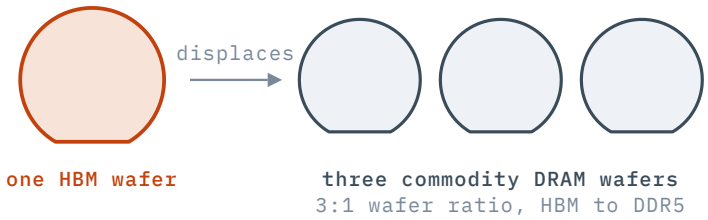
EXHIBIT 3 **Anatomy of a stack — the five parts every figure in this study refers to**



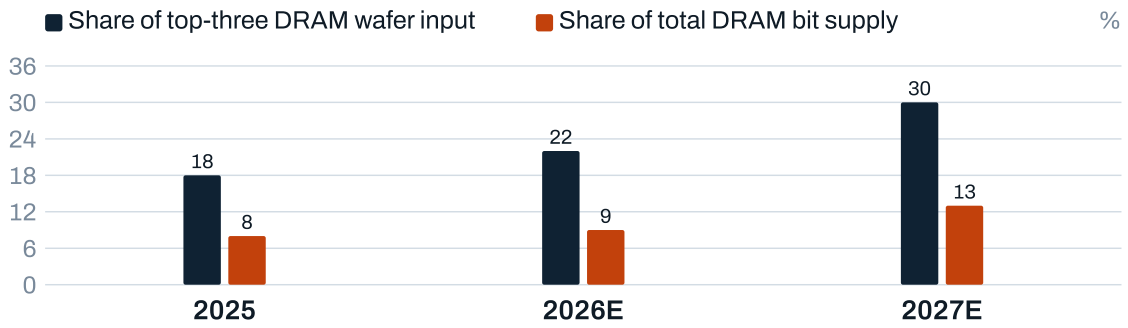
Sources: JEDEC JESD270-4, 2025-04-16; SK hynix Newsroom and industry analyses, 2024–2026. Not to scale.

Three properties separate HBM from commodity DRAM. It destroys capacity: wafer cost per gigabyte runs roughly three times DDR5, on an approximate 3:1 HBM-to-DDR5 wafer ratio at Micron[2]. It carries a per-gigabyte premium above five times conventional DRAM in 2025, after double-digit price growth through 2024 and 2025[3]. And it sells before it is built: Micron booked its whole calendar-2026 supply, HBM4 included, by December 2025[4].

EXHIBIT 4 **What HBM takes, and what it returns**



~22% of wafer input, top three, 2026 → returns → **~9%** of DRAM bits supplied, 2026



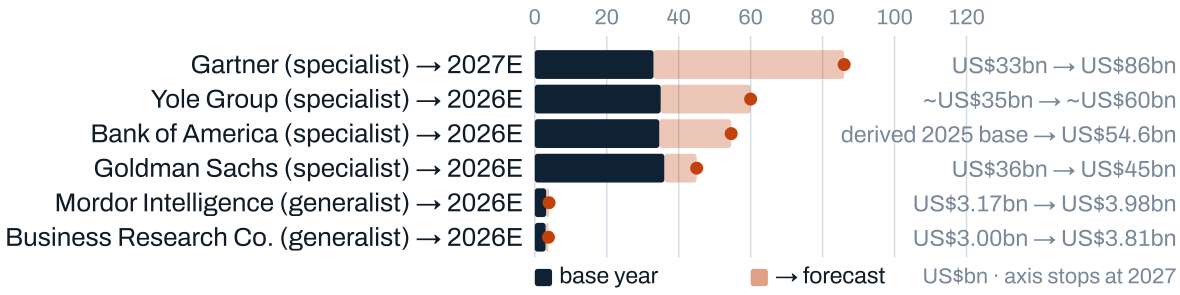
Sources: TrendForce, 2026-06-02 (2026 and 2027 are projections); Tom's Hardware, 2026-06; TechTimes, 2026-06-04. The wafer-input series covers Samsung Electronics, SK hynix and Micron; the bit-supply series covers the whole DRAM industry.

1.2 Two house families, an order of magnitude apart

SIZING
STRUCTURE
VALUE CHAIN

There is no consensus size for this market in 2026. Two families of estimate sit roughly a factor of ten apart, and the distance is definitional.

EXHIBIT 5 **Global HBM market by research house, base year to latest published point**



Sources: Gartner (2026-03-26) via SK hynix SEC 424B4; Yole Group, Chiplet Summit 2026; BofA via SK hynix Newsroom, 2026-01-05; Goldman Sachs, 2026; Mordor Intelligence and The Business Research Company, 2026. Every base is 2025 and each row carries that house's own end year, so bar lengths are not like-for-like. BofA's base is an **[NRG estimate]** from its +58%.

WHY IT MATTERS The specialist band governs: a market size in this study is stacked-DRAM revenue at the maker.

The usable planning band for 2026 among specialist houses is **US$45bn to US$60bn**, and it moves: Goldman Sachs cut 13% from US$51bn inside one revision cycle while Yole Group raised its own. The generalist family clusters near US$3.7–4.0bn, and every published CAGR across both families sits between 25% and 33% except Grand View Research's 13.1%, on a narrower data-centre-chip definition.

EXHIBIT 6 **Where HBM sits inside the memory market, and how the generations differ**

TWO HOUSES' CUTS, US$

LAYER	2025	2026E	2027E
Semiconductors	—	1.32tn	1.56tn
Memory	216bn	633bn	~748bn
Memory · WSTS	—	>800bn	+32%
DRAM	143bn	—	401bn
HBM	33bn	—	86bn

GENERATIONS, AND WHAT CHANGED AT HBM4

GEN.	INTERFACE AND STACK	STATUS, MID-2026
HBM3	1024-bit, 16 ch; ~0.8 TB/s	legacy; CXMT's target
HBM3E	1024-bit; ~1.2 TB/s; 36GB	~2/3 of 2026 shipments
HBM4	2048-bit, 32 ch; 2 TB/s	in mass production
HBM4E	2048-bit; 14–16 Gbps	samples shipped 1H26
HBM5	hybrid bonding at scale	as early as 2029

Sources: Gartner (2026-03-26) via SK hynix SEC 424B4; WSTS Spring 2026-05; JEDEC JESD270-4, 2025-04-16; Tom's Hardware, 2025-04; TrendForce, 2025-12-26 and 2026-08-04; Businesskorea, 2026-04-06. DRAM includes HBM. The WSTS row is that body's Spring cut, which also puts 2026 semiconductors at US$1.51tn and publishes 2027 as growth only; its autumn-2025 cut had put 2026 memory at US$294.8bn, the revision itself evidence. Capacity per stack: 24GB (HBM3), 24–36GB (HBM3E), 64GB (HBM4, from February 2026). HBM4 peaks at 2 TB/s; extended configurations reach 3.3 TB/s (vendor materials, 2026). HBM is ~23% of Gartner's DRAM revenue in 2025, ~21% in 2027; Yole projects above 50% by 2030.

WHY IT MATTERS The change at HBM4 is width. Doubling the interface to 2048 bits and the channel count to 32 forced a wider, logic-heavy base die — the one specification decision that pulled a foundry into a memory market (E7).

1.3 The value chain, and where it is actually held

SIZING
STRUCTURE
VALUE CHAIN

Seven layers separate a DRAM wafer from a rack. Three of them turn commodity DRAM into HBM, and those are the layers keyed below. A different three ration what the market can supply, and that set begins at the die itself.

EXHIBIT 7 The HBM value chain in seven layers — who holds each, and how tightly

	LAYER	WHAT HAPPENS	WHO HOLDS IT	CONCENTRATION
1	**DRAM die**	Leading 10nm-class nodes; yield compounds over twelve dies	SK hynix, Samsung, Micron	Top three >90% of DRAM revenue
2	**TSV & thinning**	Vias drilled, dies thinned; the scrap unit is a whole stack, not a die	The memory makers	In-house
3	**Stacking & bonding**	MR-MUF against TC-NCF; hybrid bonding at 16-high HBM4E	The memory makers	Process choice is the differentiator
4	**Base die**	From HBM4 a logic part on a foundry node, not a passive interface	TSMC for SK hynix & Micron; Samsung Foundry in-house	One foundry serves two of three
5	**2.5D packaging**	CoWoS interposer places the stack beside the accelerator	TSMC plus OSAT partners	~10% residual gap at end-2026
6	**Back-end equipment**	Thermocompression bonders assemble every stack in the world	Hanmi, SEMES, ASMPT, Yamaha	~70% held by one supplier (E24)
7	**Buyer**	Accelerator vendors and custom-silicon programmes	NVIDIA, Broadcom-designed ASICs, four hyperscalers	~US$620bn guided 2026 capex (E11)

Sources: IDC, Q1 2026, via SK hynix SEC 424B4; Semiconductor Engineering, 2026; SK hynix Newsroom and industry analyses, 2024–2026; TrendForce, 2026-01-21, 2026-03-20 and 2026-06-15; Asia Business Daily, 2025-12-22. Row 7 sums the four hyperscalers' own capex guidance at range midpoints; competing compilations of the same four run higher (E11).

Only 2.5D packaging appears in both sets; scarcity and differentiation sit apart.

EXHIBIT 8 The three chokepoints that set the economics

CHOKEPOINT 1

Wafer capacity

Wafers are not scarce. The allocation is two-sided: commodity DRAM re-prices hard enough that it is no longer a low-return fallback.

~30% of top-three input by 2027E

CHOKEPOINT 2

Advanced packaging

CoWoS runs from roughly 35,000 wafers a month in late 2024 toward 120,000–140,000 by end-2026, plus 50,000–60,000 from OSAT.

Gap narrows to ~10% at end-2026

CHOKEPOINT 3

Qualification

The least visible. On 2026-06-05 NVIDIA confirmed Samsung, SK hynix and Micron had all passed HBM4 certification for Vera Rubin.

First: three qualified sources

Sources: TrendForce, 2026-06-02 and 2026-06-15 (CoWoS gap from ~20%); NVIDIA chief-executive remarks, 2026-06-05.

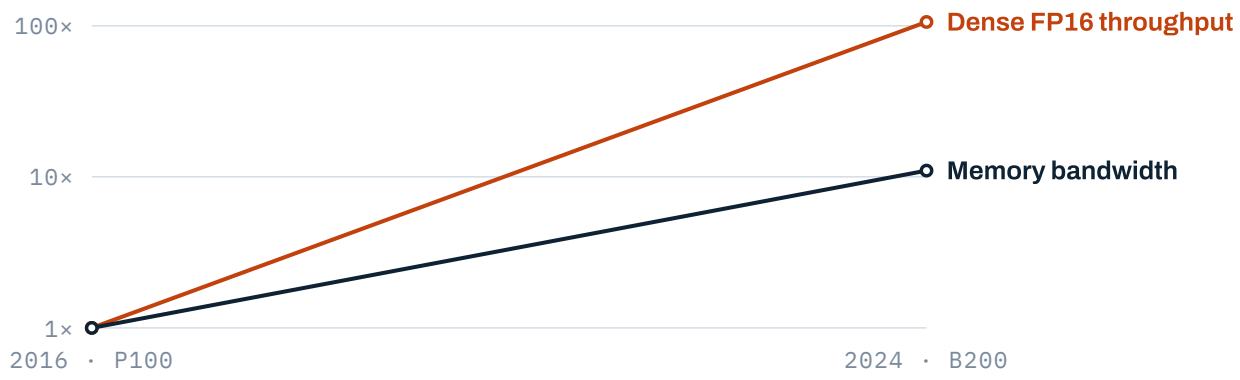
2.1 Compute outran the ability to feed it

DEMAND ENGINE
THE FIVE-YEAR RECORD

Between NVIDIA's P100 in 2016 and its B200 in 2024, dense FP16 throughput grew roughly **106 times** while the memory bandwidth feeding it grew about **11 times**[5]. HBM closes that gap.

EXHIBIT 9 **The gap that created the product — compute against memory bandwidth, P100 to B200**

Indexed · P100 (2016) = 1× · logarithmic scale

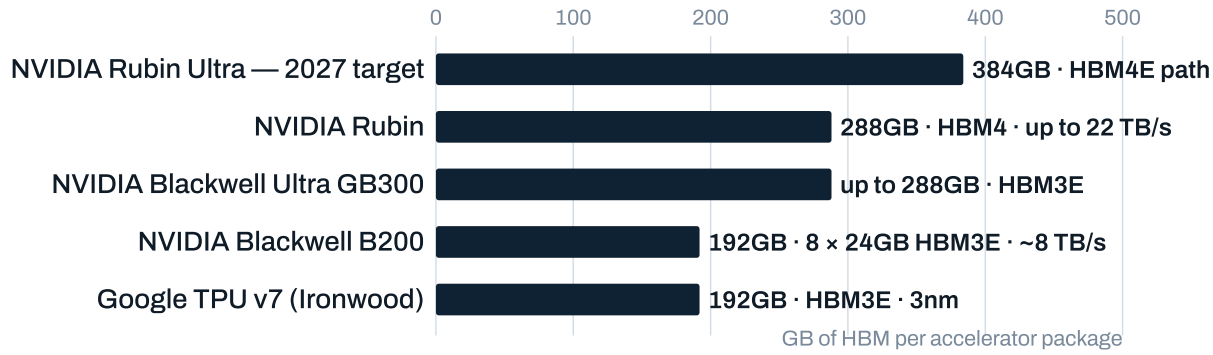


Sources: published architecture analyses, 2024–2026. Both series are indexed to the 2016 part, so the axis carries multiples and not absolute figures. Over a recent two-year window the same divergence appears at a smaller scale: AI model compute rose about 3× against 1.6× for memory bandwidth and 1.4× for interconnect.

WHY IT MATTERS Demand here follows a ratio: each accelerator generation must carry enough bandwidth to keep its own arithmetic busy, making HBM content per accelerator (E10) the forward variable that holds where unit counts cannot.

Capacity per accelerator roughly doubles across two generations while bandwidth rises nearly threefold from B200 to Rubin. A flat unit count still produces substantial bit growth while content per unit climbs, which is how TrendForce could project HBM demand growth above 70% for 2026[6] in the season accelerator unit forecasts were being cut. The lever runs downward too, and has since the third quarter of 2026 (§4).

EXHIBIT 10 **HBM content per accelerator — the multiplier that holds when unit counts do not**



Sources: NVIDIA and vendor materials, 2025–2026; Google, 2025; TrendForce, 2026-06-02 (the Rubin Ultra 2027 target). Custom AI ASICs are generally moving from 96–192GB to 216–288GB during 2026 (TrendForce, 2026-06-02) and are not plotted, because that figure is a range across an unnamed set of parts.

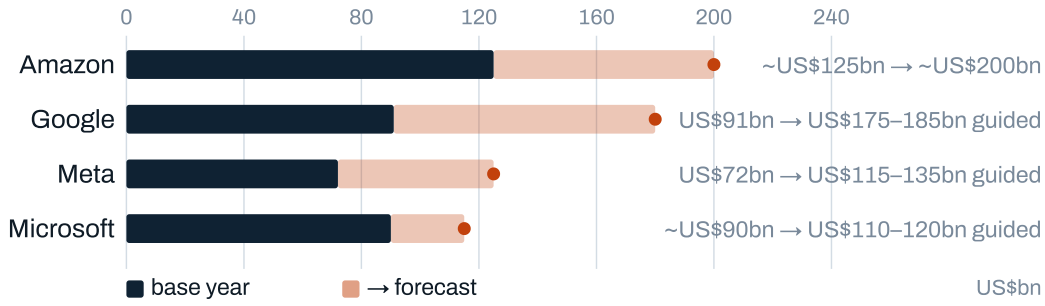
2.2

Who is paying, and through which channel

DEMAND ENGINE
THE FIVE-YEAR RECORD

Four buyers of consequence sit at the end of this chain. Their own 2026 capital expenditure guidance, summed at range midpoints, runs about two-thirds above what the same four spent in 2025.

EXHIBIT 11 **Hyperscaler capital expenditure — 2025 outturn against 2026 guidance**



Source: company guidance, 2026; bars end at range midpoints and the notes carry the ranges as guided. The bars sum to **~US$620bn against a ~US$378bn 2025 base, +64%**. Compilations of the same four circulate at US$690bn, US$700bn (CNBC, 2026-02-06), US$725bn (valueaddvc compilation of company guidance, 2026, on a ~US$410bn 2025 base) and US$760bn (Statista, 2026) — a **~US$700–760bn** band drawn from later guidance vintages than the bars above. The spread tracks vintage.

WHY IT MATTERS Capital expenditure reaches HBM demand only indirectly, through accelerator procurement; AI semiconductors at roughly 30% of 2026 semiconductor revenue[7] is the closer proxy. Guidance moves — Alphabet's second-quarter report drew a reaction that put the other three under scrutiny days later[8].

Since 2023 the buyer set broadened. ASIC-based AI servers were projected at 27.8% of 2026 shipments, custom-ASIC shipments growing 44.6% against 16.1% for merchant GPUs[9], and Goldman Sachs put HBM demand from custom-ASIC chips growing 82% to a third of the market[10]. Broadcom, at roughly 70% of custom accelerator design services[11] and Google's design partner to 2031, reported US$8.4bn of AI semiconductor revenue in its first fiscal quarter, up 106%, against a US$73bn backlog[12]; TPU v7 was projected above 3m units in 2026 at US$12,000–15,000 apiece[13]. That channel is not in the national series below.

EXHIBIT 12 **Korean semiconductor exports by month, US$bn — the cycle with no HBM line inside it**



Source: Korean trade ministry via press, 2026-02, 2026-04, 2026-06 and 2026-08-01, at the precision each release used. These are **total semiconductor** values: Korean trade statistics publish no HBM-only export line. The July release carries a level and a run-length marker, and no year-on-year figure.

2.3 The record, and a second market behind a wall

DEMAND ENGINE
THE FIVE-YEAR RECORD

The five-year record is deliberately not drawn as a line: the 2024-to-2025 step below is re-basing across houses, not growth.

EXHIBIT 13 What the record shows, 2022–2026 — separately sourced points, never one series

YEAR	MARKET READING	HOUSE	STRUCTURAL MARKER
2022	first separate HBM line	Gartner	First HBM3 ships, 6.4 Gbps per pin
2023	~8.4% of DRAM revenue	TrendForce	SK hynix ~90%; Micron enters HBM3
2024	~US$8.9bn, +127% YoY	TrendForce	HBM3E commercialised; bits +260%
2025	US$33bn / ~US$35bn / US$36bn	Gartner·Yole·Goldman	~18% of top-three wafer input
2026	US$45bn / US$54.6bn / ~60bn	Goldman·BofA·Yole	HBM4 in mass production

Sources: Gartner via SK hynix SEC Form 424B4; TrendForce, 2024-03-18, 2024-07-22 and 2026-06-02; Yole Group, Chiplet Summit 2026; Goldman Sachs, 2026; BofA via SK hynix Newsroom, 2026-01-05. Samsung's 2024 share is omitted for want of a primary source: compilations give 39% and others materially lower. HBM was roughly 20.1% of DRAM revenue by end-2024.

WHY IT MATTERS Told through capacity, these five years are one reallocation: HBM went from about 8.4% of DRAM revenue in 2023 to about 23% in 2025 on Gartner's implied numbers, never exceeding a tenth of the industry's bits (E4).

The reallocation showed up in price: consumer and enterprise memory rose over 200% from early 2025[2], and Gartner put 2026 DRAM up 125% and NAND up 234%, relief unlikely before late 2027[7]. Then 1Q26 inverted it: HBM wafer revenue fell below DDR5 64GB RDIMM pricing, and suppliers began setting the split by how HBM pricing resolved[1].

EXHIBIT 14 The export-control fork — one market, and a second behind a wall

REGION 1

The market this study measures

DRAM dies stacked on a logic base die, packaged beside the accelerator. Three qualified suppliers, one dominant packaging route, and a buyer list short enough to name.

Open to every buyer outside the boundary

EXPORT CONTROL

REGION 2

Inside the sanctions boundary

- ~~HBM~~ — unavailable in China under semiconductor sanctions, so the architecture above cannot be bought.
- **3D DRAM** — logic and DRAM dies stacked vertically with the connections distributed across the die instead of concentrated at its periphery. Products launched; use cases elsewhere still limited.

Disclosed by Semifive, Annual Report FY2025[14]

Source: Semifive, Annual Report FY2025, which also records revenue rising from ₩5.4bn in FY2021 to ₩41.9bn, ₩40.1bn, ₩76.0bn and ₩90.5bn since — the second buyer channel from another disclosure regime. Seasonality has gone from half the market: HBM ships on long-term contracts, forming a demand curve off the consumer cycle[15].

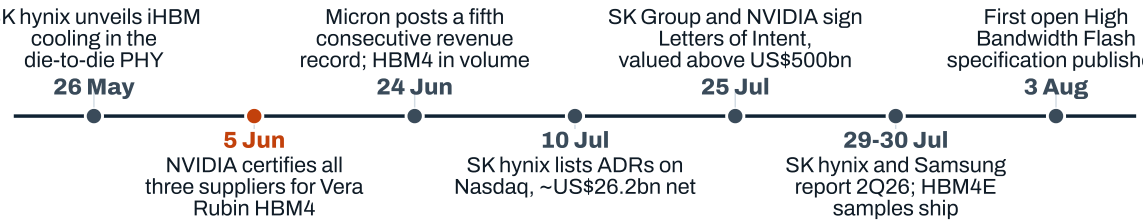
3.1 Three months that changed the question

MAY TO AUGUST 2026

Between May and August 2026 the competitive question in this market moved from who can supply high-bandwidth memory to how it is allocated.

EXHIBIT 15 **The mid-2026 sequence — certification, capital, and a new memory tier**

each item dated and web-verified to 2026-08-10



Sources: TrendForce, 2026-05-26; NVIDIA chief-executive remarks, 2026-06-05; Micron, 2026-06-24; CNBC and SK hynix SEC Form 424B4, 2026-07-10; NVIDIA newsroom and SK hynix Newsroom, 2026-07-25; SK hynix Newsroom, 2026-07-29 and Samsung Newsroom, 2026-07-30; SanDisk and SK hynix, 2026-08-03/04. The 5 June certification is marked in accent as the event the rest of this section turns on.

The inflection is the 5 June certification (E17). Reported HBM4 allocation for Vera Rubin runs SK hynix at roughly 60–70%, Samsung at 25–30% and Micron the remainder[16], with UBS putting SK hynix nearer 70% and TrendForce at about two-thirds — estimates, since no contract discloses the split. An allocation on one platform is not a share of the market, and E22 orders those two the other way round.

Three sets of results landed inside eight weeks, cited as market evidence, not as comment on any issuer. Micron's fiscal third quarter to 2026-05-28 recorded US$41.46bn of revenue against US$9.30bn a year earlier, with the fourth quarter guided to US$50.0bn ± US$1.0bn at roughly 86% gross margin[17]. SK hynix reported second-quarter revenue of ₩79.32tn and operating profit of ₩60.54tn, a 76% margin[18]; Samsung ₩171.5tn and ₩89.5tn, its device solutions division ₩127.5tn and ₩89.2tn[19]. Margins at these levels across three independent issuers in one quarter are what a supply-constrained market looks like from the inside.

Customer coupling deepened without becoming a contract: the SK Group–NVIDIA partnership of 2026-07-25, valued above US$500bn, spans AI-factory construction and a long-term memory relationship including HBM for Vera Rubin, with SK Telecom planning a 2 GW AI factory from 2027 — and it was signed as **Letters of Intent**.

3.2

Pricing power moved to the supplier

MAY TO AUGUST 2026

Samsung and SK hynix planned an HBM3E price increase of about **20%** for 2026[20]. TrendForce expects 2027 contract prices to rise several-fold, with suppliers holding pricing power through the year[21]. A buyer trading stack height down while prices climb like that is in a seller's market.

EXHIBIT 16 **Two units, kept apart — what a gigabyte costs, and what a stack costs**



PER-STACK QUOTES — A DIFFERENT UNIT, KEPT OFF THE LADDER ABOVE

~US$500 HBM4 · 48GB stack, early HBM4

~US$300 HBM3E · 36GB 12-high stack, Aug 2026

Sources: Fubon channel checks, cited 2026 (per gigabyte); Silicon Analysts, 2026-08 (per stack), citing an HBM3E stack ASP range of US$200–700. Negotiations for 2027 supply opened in 2Q26. Every figure here comes from channel checks and a data aggregator, with no research house or issuer behind it, so all are a direction of travel, not a price series.

EXHIBIT 17 **The qualification gate — who cleared it, when, and who has not**

Samsung	2025-09-19	12-layer HBM3E clears NVIDIA qualification
Samsung and SK hynix	2026-02	HBM4 mass production begins, simultaneously
All three suppliers	2026-06-05	NVIDIA certifies HBM4 for Vera Rubin
Micron	2026-06-24	HBM4 in high-volume shipment for the lead platform
CXMT	April 2026	Reported in sampling; HBM3 target end-2026

Sources: KED Global, 2025-09-19; NVIDIA chief-executive remarks via press, 2026-06-05; Micron, 2026-06-24; press compilation, 2026-04 and TechPowerUp, 2026. CXMT states an end-2026 domestic HBM3 target on roughly 20% of its capacity against a 300k wpm ambition, discloses no output, and is reported to have yield as its binding constraint. The February 2026 mass-production start is the date E6 carries; no single dated release stands behind it.

WHY IT MATTERS Share here has always been downstream of qualification: SK hynix held roughly 90% through 2023 because it was the only cleared source at the buyer that mattered. Read every share figure in E22 against these dates.

ATTRIBUTION CAVEAT

Samsung's second-half HBM4 guidance — sales tripling quarter on quarter in Q3, above 60% of HBM revenue in 2H — comes from press summaries of the 2026-07-30 call. It is **company-reported**, absent from the newsroom release, and used as an input nowhere.

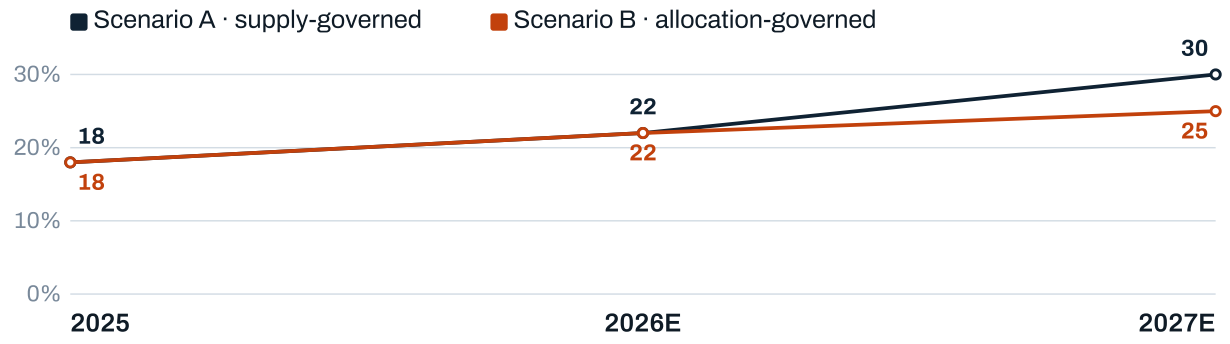
4.1

The last quantity that can honestly be plotted

A SUPPLY FORECAST AND WHERE IT STOPS

Every credible forward view for 2027 is a supply view. Bit shipments are projected to grow 50–60% and still fall short of demand; the DRAM shortage will itself cap the wafer capacity open to HBM; contract prices should rise sharply; and the largest buyer is evaluating shorter stacks than it specified[21].

EXHIBIT 18 **HBM's claim on top-three DRAM wafer input, % — two scenarios that part only at 2027**



Scenario A follows the supply-governed path (TrendForce, 2026-06-02: 18% / 22% / 30%); Scenario B ends at the upper bound of the 22–25% plateau the evidence supports (TrendForce, 2026-06-02 and 2026-08-04). Scenario C, architecture-governed, carries no 2027 quantity because its effects begin around 2029–2030. All three are set out in full in §7.

WHY IT MATTERS Wafer-input share stands here in place of market size because it is the one forward quantity on which the two scenarios carry different 2027 readings; market size for that year is a single figure from a single house, and the two observations beyond it cannot end a line (E19).

Advanced packaging, the constraint that capped 2024 and 2025, is being relieved as the CoWoS gap narrows toward 10% at end-2026 (E8) while DRAM wafer capacity becomes the new binding one — the bottleneck moving upstream from the packaging house to the fab.

EXHIBIT 19 **Why the axis stops at 2027 — one plotted point, two isolated ones**

2027 · GARTNER

The last specialist point

US$86bn for 2027 is the only semiconductor-specialist figure on the record for that year, and it is the last point E5 plots.

US$86bn · plotted

2028 · MICRON

Company-reported

A supplier's view of the market it is building into, compounding about 40% from ~US$35bn in 2025. It is useful context and a poor axis.

~US$100bn · company-reported

2030 · YOLE

Superseded vintage

From a 2025-06-19 release at 33% a year. The same house's 2026 cut raises 2026 to ~US$60bn without restating 2030.

~US$98bn · 2025-06-19 cut

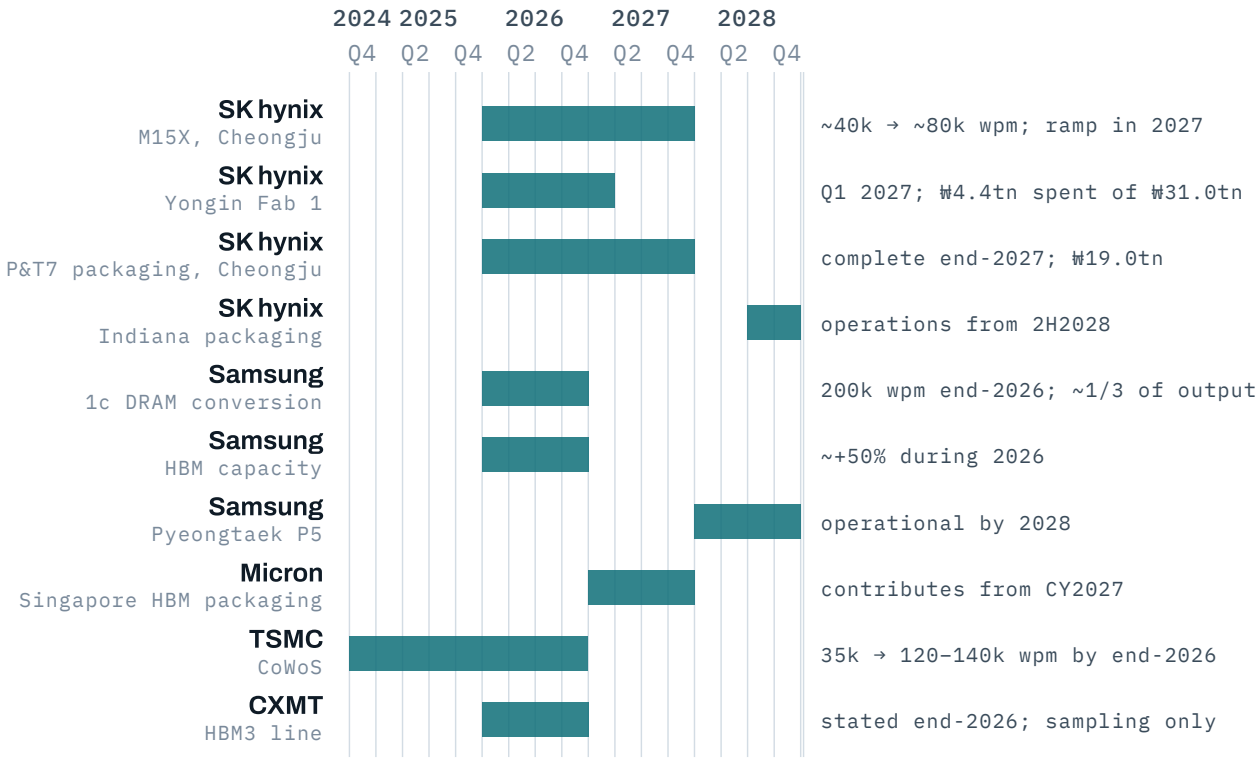
Sources: Gartner via SK hynix SEC Form 424B4 (data dated 2026-03-26); Micron FQ3 FY2026 call, 2026-06-24; Yole Group press release, 2025-06-19 and Chiplet Summit 2026; Korea JoongAng Daily, 2026-08-05. Yole's 33% rate is coherent only from a 2024 base near US$17.7bn **[NRG estimate]** and implies about 24% a year from 2025; carried forward from the 2026 cut it would imply about 13%, which no reading of this market supports. Its August 2026 briefing uses a 2031 horizon for Korean memory leadership and publishes no HBM size beside it. Near-term rates run far above the 25–33% long-run band because they start at an inflection: Gartner has HBM compounding 60.5% from 2025 to 2027 inside a DRAM market compounding 67.3%, Yole +70% in 2026, BofA +58%, and TrendForce demand growth above 70% on shipments above 30bn gigabits.

4.2 What is actually being built, and when it lands

A SUPPLY FORECAST
AND WHERE IT STOPS

Read the right-hand edges. The capacity funded to relieve the 2027 shortage arrives after it.

EXHIBIT 20 Announced HBM-relevant capacity — the assets, and the quarter each one arrives



Sources: SK hynix SEC Form 424B4, filed 2026-07 — Yongin's ₩4.4tn is spend recognised to 2026-05-31 against a ₩31.0tn project running to 2030; TechTimes, 2026-06-06; TrendForce, 2025-11-19, 2025-12-18 and 2025-12-30; industry reports, 2026 (Pyeongtaek P5); Micron, 2026-06-24; TechPowerUp, 2026. wpm = wafers per month. Each bar covers only the window its source states — a disclosed ramp, a conversion year, a filed remaining-spend window, or the quarters from which an asset operates — and no construction start is inferred where none is disclosed. Micron's fiscal 2026 capital expenditure was raised to about US$20bn from US$18bn; the CoWoS bar excludes the OSAT capacity E8 counts beside it; Micron's Boise site, which it states will bring advanced HBM packaging to the United States for the first time, carries no bar because the only timing disclosed is post-completion; CXMT's line is a stated target of roughly 20% of its mass-production capacity against a 300k wpm ambition, with no disclosed output.

Only four bars land inside 2026 — Samsung's two, the CoWoS ramp E8 already counts as the packaging chokepoint's relief, and CXMT's stated target. For scale, AI as a whole was projected to take about 20% of global DRAM wafer capacity in 2026[6], against HBM's ~22% claim on the top three.

Two entries are not capacity in the same sense as the rest: CXMT's line sits inside the sanctions boundary drawn in E14, and Micron's Singapore packaging is described only as contributing from calendar 2027. Both are drawn anyway; leaving an announced asset off a capacity chart is the more misleading choice.

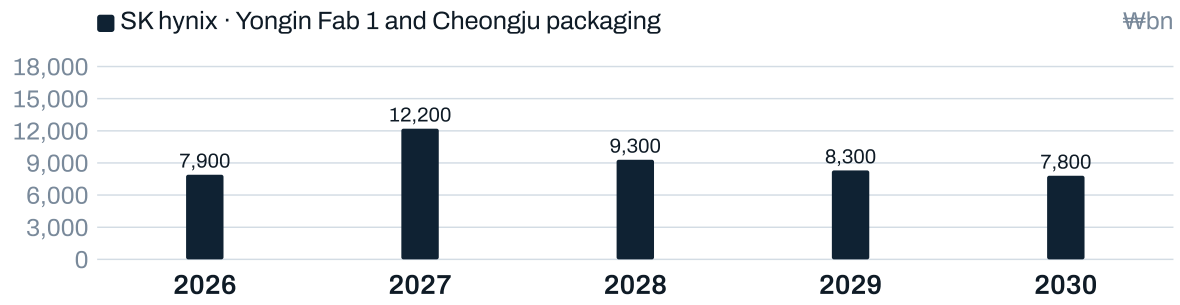
4.3

A filed budget, and the perimeter of the forecast

A SUPPLY FORECAST AND WHERE IT STOPS

One issuer publishes a year-by-year construction schedule under securities-law liability — the only exhibit in this study that plots a quantity past 2027.

EXHIBIT 21 **A filed construction budget — the heaviest years land after the shortage**



Source: SK hynix SEC Form 424B4, filed 2026-07 — one issuer's remaining-spend schedule for two named construction projects, together ₩50.0tn with ₩45.5tn unspent at the filing date. This is not a market-wide capital expenditure figure, and no other participant's construction schedule is disclosed at this resolution.

WHY IT MATTERS The axis here runs to 2030 because a filed spending schedule is a fact about one company, not a forecast about a market. The forward axes that carry the market itself still stop at 2027 (E5, E18).

The shape matters more than the level. Spending peaks in 2027 — a year after the shortage begins and a year before the assets it funds can serve it — and stays above the 2026 figure through 2029. A schedule filed under securities-law liability cannot be revised in a press cycle. The same filing applies the US$26.2bn of net proceeds from the 2026-07-10 Nasdaq listing to general corporate purposes including capital expenditure (E15).

The constraint that governs 2027 was therefore decided in 2024 and 2025, when these projects were approved (E18, E19, E20).

THE FORECAST PERIMETER

Four things this study declines to forecast. A single market size for any year — the two house families sit roughly a factor of ten apart on definition (E5). Accelerator unit volumes — estimates for Rubin in 2026 span 200,000 to 5.7m units on inconsistent unit definitions, and the production target was reportedly cut from 2m to 1.5m on HBM4 verification delays; demand is expressed here through content per accelerator (E10) and bit growth instead. Supplier-level HBM revenue — no issuer discloses it, so share is the only citable participant-level unit. CXMT capacity or share — stated targets exist, disclosed output does not.

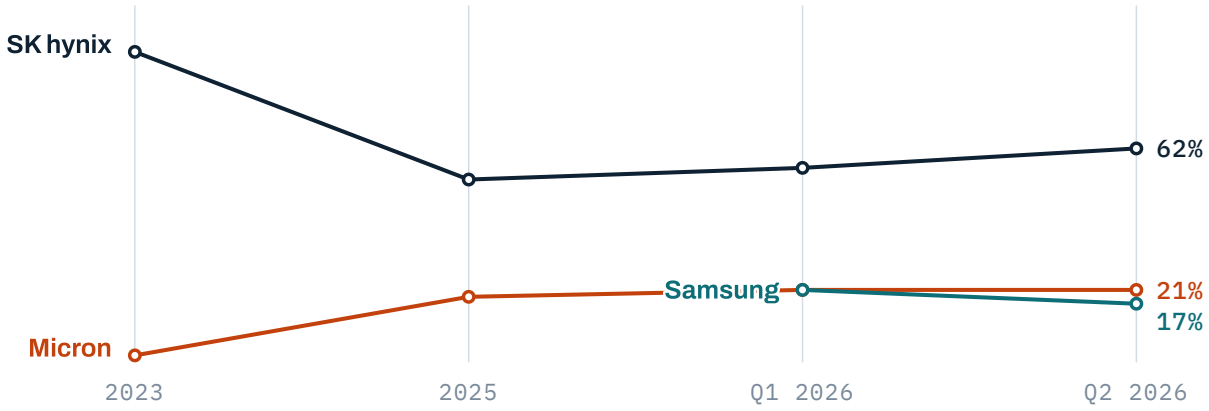
5.1

One leader, and a disputed second place

CONCENTRATION
SHARE BANDS
THE CHOKEPOINT

This is a three-supplier market inside a three-supplier market: the same three companies hold more than 90% of DRAM revenue and are the only volume HBM suppliers. Sources disagree by close to six percentage points on the leader, and on the ordering below it.

EXHIBIT 22 **HBM share by supplier, 2023 to Q2 2026 — a monopoly decaying into an oligopoly**



Sources: TrendForce, 2024-03-18 and an industry compilation (2023); industry compilation, 2025–26 (2025); IDC, 2026-05-27, via SK hynix SEC Form 424B4 (the leader at Q1 2026); Counterpoint Research, 2026 (Samsung and Micron at Q1 2026, and the leader at ~58%); Astute Group compilation, 2026 (Q2 2026). Readings are on a revenue basis where the source states one. The Q1 2026 column mixes houses because the filed IDC figure breaks out only the leader. The Q2 2026 column carries a period caveat: the identical 62% appears elsewhere attributed to Q2 2025 shipment share, and the two cannot both be right.

WHY IT MATTERS Samsung's line starts at 2026 because no primary source places its earlier level — the same omission E13 makes, for the same reason. The gap is an omission, not a zero.

The usable band is **SK hynix at 56–62%**, with Samsung and Micron each in the high teens to low twenties and the ordering between them disputed. The filed IDC reading anchors it, under securities-law liability. Both the decay from roughly 90% to roughly 53% and the partial re-consolidation after it track qualification dates rather than capacity (E17).

Against DRAM overall, Samsung leads at 38.6–39% across the first two quarters of 2026 while sitting at 17% of HBM; SK hynix is second in DRAM at 29.1% and first in HBM; Micron holds 22.4% of DRAM revenue against roughly 21% of HBM[22]. The company that leads DRAM does not lead HBM — separate contests, settled by a stacking process, a base-die route and a customer relationship; scale decides none of them.

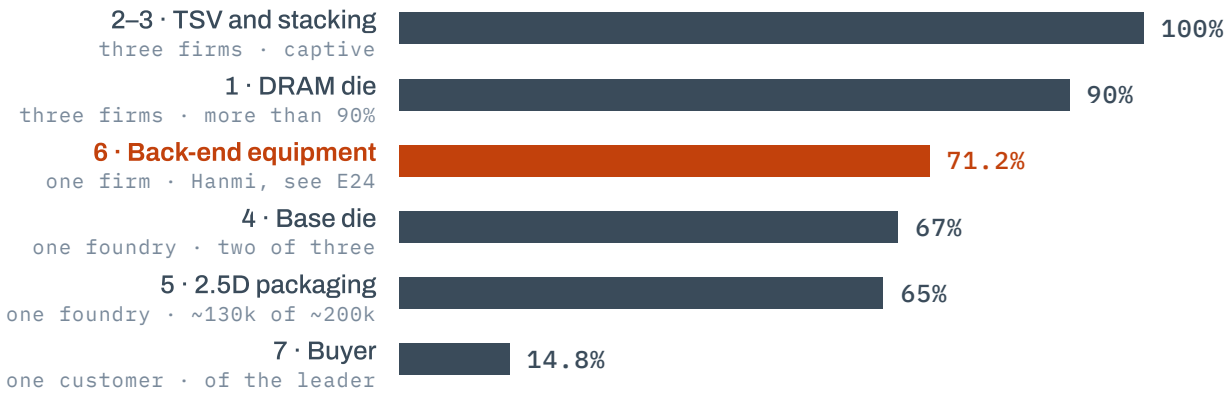
Nothing contractual keeps a fourth supplier out; complexity across design, fabrication and packaging is the barrier, and the capital to carry all three. CXMT is the observable test — a funded national champion whose stated end-2026 target and unresolved yield constraint E17 records, and whose process level Yole put at 10nm-class 1z, a node the leading makers first reached in 2019[23].

5.2 The tightest hold sits outside the memory layer

CONCENTRATION
SHARE BANDS
THE CHOKEPOINT

Two of the layers below are wholly held because nobody sells them. Of the rest, the one held most tightly by a single company makes no memory at all.

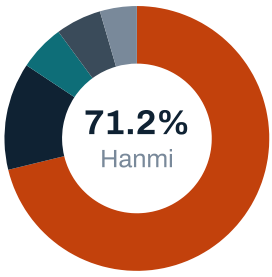
EXHIBIT 23 **Concentration, layer by layer**



Sources: IDC, Q1 2026, via SK hynix SEC Form 424B4 (layers 1–3 and 7); TrendForce, 2026-01-21 and 2026-03-20 (layer 4); TrendForce and industry reports, 2026 (layer 5 capacity); Asia Business Daily, 2025-12-22 (layer 6). Each bar is a published concentration; E7 draws the same layers unlabelled, and only its layer 5 differs. Layers 2 and 3 share a bar: both are captive at the same level. Layer 1 is plotted at the published floor of "more than 90%"; layer 5 is an [NRG estimate], TSMC's 120–140k wafers a month against an industry total approaching 200k; layer 6 is the figure E24 plots and carries the caveat stated there; layer 7 is the leader's largest customer as a share of its own Q1 2026 revenue.

WHY IT MATTERS Only three measure a company's hold on the layer it sits in, and all three sit outside the layers the memory makers keep in-house. The tightest of the three is back-end equipment (E24).

EXHIBIT 24 **Thermocompression bonder market share**



Hanmi Semiconductor — the tool every HBM stack passes through	**71.2**
SEMES — the second Korean position in the tool layer	**13.1**
ASMPT — took an SK hynix HBM4 order in December 2025	**5.6**
Yamaha Robotics — the fourth named position	**5.6**
Residual to 100% — unallocated in the source	**4.5**

Sources: Asia Business Daily, 2025-12-22 and 2026 (segment growth); TrendForce, 2025-12-12 (the ASMPT order, amid a Hanmi–Hanwha patent dispute). Shares are as published; the residual closes the ring. The 71.2% carries a period caveat — cumulative through a third quarter whose end is unconfirmed, from a single Korean press report — and reads as roughly seven-tenths. Cumulative revenue on that basis was US$247.7m; the segment should compound ~13.0% a year to 2030. Hanwha Semitech put a 0.1 micrometre hybrid bonder into an SK hynix line in April 2026 on ~₩80bn and ~₩9bn of orders.

With JEDEC thickness rules relaxed and 16-high demand soft, hybrid bonding is now seen no earlier than 16-high HBM4E[24], extending the incumbent tool's runway.

6.1 Five positions, read at three points in time

THREE DISTANCES FROM ONE MARKET

Read down a column for one participant, across a row for one moment in this market.

EXHIBIT 25 **Five participants at three points in time — the market read through the positions inside it**

	SK HYNIX	SAMSUNG	MICRON	HANMI	THE KOREAN BACK-END RING
FIVE YEARS AGO	A niche part for high-performance computing, held by the first company to stack DRAM on through-silicon vias. It shipped the first HBM3 in 4Q22 and held roughly 90% through 2023.	The largest memory company in the world and a marginal participant here. Its stacking route was the harder one to scale, and its share through 2023–2024 is this study's most contested number.	The smallest of the three DRAM producers and effectively absent, at roughly 2% in 2023. It entered HBM3 in 2Q23, six to nine months behind the buyer's only qualified source.	It supplied thermocompression bonders into a back-end market that was not yet strategic. HBM3 arrived from 4Q22 with an assembly step that ran through the machine it already built.	Outsourced assembly and test for commodity memory. Hana Micron signed a long-term back-end contract with the HBM market leader in November 2021, before the demand arrived[25].
WHERE IT STANDS NOW	56.4% of HBM revenue in Q1 2026 and second in DRAM at 29.1% — the only participant stronger in HBM than in DRAM. HBM4 mass shipments began in 2Q26.	Re-entered on the current generation: 12-layer HBM3E qualified September 2025, HBM4 in mass production from February 2026, certified for Vera Rubin in June 2026. Company-reported pin speed 11.7 Gb/s.	Roughly a fifth of the market, from about 2% three years earlier — the largest positional change made here. Its 2026 supply was booked by December 2025; HBM4 on 1-beta DRAM ships in volume.	Roughly seven-tenths of the bonder market (E24), on cumulative third-quarter revenue of US$247.7m. Recent orders: ₩44.2bn of HBM4 bonders and a first TC Bonder 4.5 contract near US$28.7m.	Five issuers, one split: Hana Micron's packaging line grew 27.0% in FY2025 while SFA Semicon, weighted to commodity memory packages, shrank 8.3% (E26, E27).
WHAT IT IS BUILDING TOWARD	Capacity and thermal engineering, no longer primacy. ₩50.0tn committed across Yongin Fab 1 and Cheongju packaging (E21), part-funded by US$26.2bn net raised on Nasdaq.	The only fully integrated position — stack, base die and packaging in one company. The base die moves from SF4 to 2nm; HBM capacity is planned about 50% higher during 2026.	Geographic optionality on a TSMC-coupled base die. Capital spending rose to about US$20bn for fiscal 2026; Singapore packaging contributes from 2027, then Boise.	Holding the position through two transitions at once: a Wide TC Bonder for HBM5 and HBM6 shown at SEMICON Korea in February 2026, and a second-generation hybrid bonder targeted by end-2026.	LB Semicon funds named 2.5D and chiplet package programmes for AI accelerators, and Hana Micron sits under a long-term contract with the market leader. Both positions were taken before 2026.

CXMT, the only announced fourth entrant, has no column: it discloses no output, no share and no revenue line, and enters this study as a stated end-2026 target with an unresolved yield constraint (E17).

Every figure here is established elsewhere (E17, E21, E22, E24, E26, E27, §3) except: Samsung's base die moving to 2nm (TrendForce, 2026-01-21) and its 11.7 Gb/s pin speed against the 10 Gb/s bar NVIDIA and AMD set, company-reported via Korean press and used as an input nowhere; Micron's 2Q23 entry (TrendForce, 2024-03-18) and Boise (Micron, 2026-06-24); Hanmi's orders and roadmap (Asia Business Daily, 2026-02-11; Seoul Economic Daily, 2026-06-08; TechTimes, 2026-06-09; Semiconductor Digest, 2026).

WHY IT MATTERS These are market positions, never valuations of the companies holding them.

6.2

The split, visible in audited disclosure

THREE DISTANCES FROM ONE MARKET

Five Korean back-end and design issuers file segment revenue that reads directly on where packaging and test demand is going. In one national industry and one fiscal year, two of them moved in opposite directions.

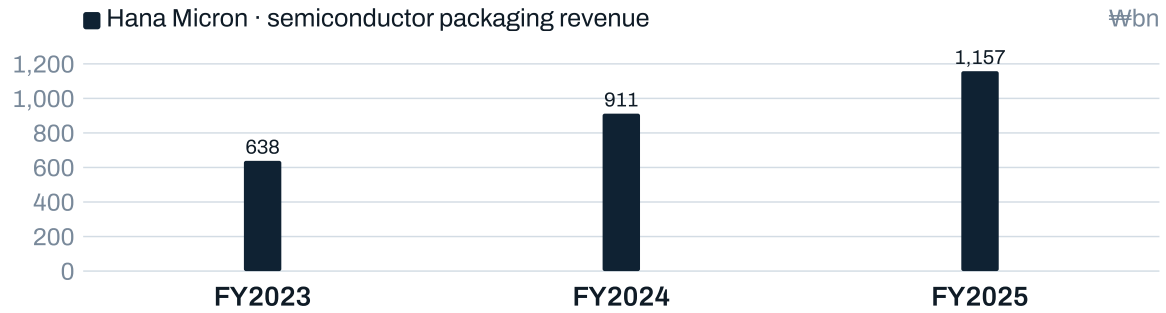
EXHIBIT 26 **Five Korean issuers, FY2025 — scale, core segment, and how often each names the product**

COMPANY	FY2025 REVENUE	CORE SEGMENT	CORE SHARE	HBM MENTIONS
Hana Micron	₩1,534.4bn	Semiconductor packaging	75.4%	0
LB Semicon	₩479.8bn	Bumping, test and back-end	98.8%	11
SFA Semicon	₩367.4bn	Memory packages	81.2%	0
Unisem	₩273.3bn	Chiller units	44.8%	3
Semifive	₩90.5bn	Design services	78.3%	2

Sources: Annual Reports FY2025 for Hana Micron[25], LB Semicon[26], SFA Semicon[27], Unisem[28] and Semifive[14]; receipt numbers are in the Appendix source register. Mentions count the literal string in the Annual Report body. Core share is the named segment's share of the revenue column beside it — Semifive's ₩90.5bn is the whole company, its design-services line ₩70.9bn of that. SFA Semicon's consolidated revenue was down 8.3% on FY2024.

WHY IT MATTERS The mention count is an observation with a caveat attached, not proof of exposure either way. What it does establish is that legibility and participation are different things in this chain: LB Semicon discloses named programmes, Hana Micron discloses growth, and both are participation in this market.

EXHIBIT 27 **One packaging line, three years — the third-tier read on where HBM demand lands**



Source: Hana Micron, Annual Report FY2025[25]. The packaging line grew 42.7% in FY2024 and 27.0% in FY2025, was 75.4% of FY2025 revenue, and was 97.0% exported. It reached ₩385.9bn in Q1 2026[29] against ₩226.7bn a year earlier[30], up 70.2%, and 76.0% of that quarter's total. This is the segment line, not the company: total FY2025 revenue was ₩1,534.4bn.

LB Semicon, the one issuer that frames its industry section around this market, puts global outsourced assembly and test at about US$99.3bn for 2026, raised its bumping price 42.8% per wafer in FY2025 against 12.3% more capacity[26], and filed a ₩49.8bn rights offering in May 2026, ₩30.0bn of it for facilities[31]. None of these issuers discloses an HBM revenue line.

6.3

Four positions, three distances, one product

THREE DISTANCES FROM ONE MARKET

The positions in E25 sit at different distances from the thing being sold, and another sits between them and files nothing in Korea at all. Distance is what set each position's exposure.

EXHIBIT 28 **Four positions at three distances from the product**

0 · THE STACK ITSELF	1 · UNDER THE DIES	1 · THE MACHINE	2 · DISPLACED WORK
SK hynix, Samsung, Micron	**TSMC**	**Hanmi Semiconductor**	**The Korean back-end ring**
The three qualified suppliers. No two are constrained by the same thing — the leading stacking yield, the only fully integrated chain, and the fastest positional gain.	From HBM4 the logic base die beneath the DRAM stack, plus the interposer that seats it beside the accelerator: the two layers Korea does not hold (E1).	Roughly seven-tenths of the bonders on which every stack in the world is assembled (E24) — a tighter single-company hold than any position inside the memory layer (E22).	Five issuers packaging what HBM pushes out of the fab. The split reaches their audited segment revenue before it reaches anyone's estimate (E26, E27).

Sources: as cited at each figure — IDC via SK hynix SEC Form 424B4 and Counterpoint Research, 2026 (the three suppliers); TrendForce, 2026-01-21, 2026-03-20 and 2026-06-15 (base die and interposer — TSMC N12 for SK hynix with 3nm under evaluation for HBM4E, Samsung Foundry SF4 in-house, and TSMC engaged by Micron for an HBM4E base die targeting 2027); Asia Business Daily, 2025-12-22 (bonders); Annual Reports FY2025 for the five issuers (E26). TSMC has no column in E25 because it discloses nothing about this market specifically; it appears here because two of the seven layers run through it.

WHY IT MATTERS Distance is measured from the product, and it does not rank the positions: the one at distance 1 that Korea does not hold is what the market cannot route around.

SK hynix's stacking yield is reported above 80% on 12-high HBM3E, and 75–80% on 1bnm DRAM against 60–65% on a competing 1cnm process, a gap that compounds across twelve dies[32]. Constraints that differ by supplier are why responses to the 2027 shortage are unlikely to arrive together. The June 2026 certification (E17) is where the current generation became an oligopoly rather than a duopoly with a third name attached.

The buyer side moved the other way. SK hynix's largest customer accounted for 23.9% of its 2025 revenue; by the first quarter of 2026 its two largest accounted for 14.8% — the bottom bar in E23 — and 12.4%. That is the custom-ASIC channel in a supplier's revenue mix. The open question sits under Samsung's position alone: whether the custom base die becomes the axis of competition, and with it how much of this market's value migrates toward foundry (E6).

THE OBSERVATION THAT CLOSES THIS SECTION

This market re-priced every position in the chain according to proximity to the buyer's qualification list, and that proximity was set years before the demand arrived. Nobody in this chain moved into position after the demand showed up, because the qualification cycles run longer than the demand cycle.

7 Three paths for the market itself

THREE PATHS
FOR THE MARKET ITSELF

The three scenarios below describe how the market environment could evolve, each anchored to published figures. What separates them is which variable governs the clearing quantity.

Three market environments, and the variable that governs each

A · SUPPLY-GOVERNED

~30% of wafer input

The consensus path, and the one every specialist house is implicitly modelling. The market clears at whatever the fabs deliver: bit shipments grow 50–60% in 2027 against faster demand growth, contract prices rise sharply, and HBM reaches about 30% of top-three wafer input (E18). Buyers manage availability by trading stack height down, which is already visible in the largest buyer's evaluation of shorter stacks. Market size tracks the specialist band and stops at 2027, because the capacity funded to relieve the shortage lands in 2027 and 2028 (E20).

Markers: HBM4E validation timelines; whether the CoWoS gap closes to ~10% at end-2026; the M15X ramp toward ~80k wafers a month.

B · ALLOCATION-GOVERNED

a 22–25% plateau

The underrated path. The binding question is whether making HBM is the best use of a wafer at all. §2.3 records the quarter in which it was not, and E8 reads that quarter as the first chokepoint. Suppliers re-weigh the split every quarter on the same arithmetic, and wafer-input share plateaus nearer 22–25% instead of climbing to 30%.

Markers: the HBM-against-RDIMM wafer-revenue comparison each quarter; whether a supplier publicly re-weights allocation; commodity prices through 2027.

C · ARCHITECTURE-GOVERNED

a ceiling, from ~2030

The slow-moving path. The memory hierarchy widens and HBM stops being the only answer to the bandwidth problem. The first open High Bandwidth Flash specification, published through the Open Compute Project in August 2026, puts 8-high and 16-high NAND stacks behind a UCIe chiplet interface and was co-developed with Google DeepMind and Tenstorrent. It gives inference systems near-compute capacity without relying entirely on HBM, with commercialisation signalled around 2030. In parallel, 3D DRAM is developing inside the boundary drawn in E14, and heat is becoming the binding physical constraint: SK hynix's iHBM puts cooling into the die-to-die interface for a 30% reduction in thermal resistance, targeted at HBM5.

Markers: HBF silicon against specification; whether a hyperscaler designs an HBF tier into a production platform; CXMT yield disclosure.

Sources: TrendForce, 2026-06-02 and 2026-08-04 (bit-shipment growth, wafer-input share, contract prices and the shorter-stack evaluation); TrendForce, 2026-06-15 (CoWoS); TechTimes, 2026-06-06 (M15X); SanDisk and SK hynix, 2026-08-03/04 (High Bandwidth Flash); SK hynix via TrendForce, 2026-05-26 (iHBM); Semifive, Annual Report FY2025[14] (3D DRAM). Scenario C carries no 2027 quantity, which is why E18 plots two lines and not three.

SCOPE OF THESE SCENARIOS

These scenarios describe the market. They carry no price target, no multiple and no buy or sell implication, and they are not mutually exclusive: A and B compete for the same two years while C runs on a different clock. Only A and B can be quantified on the present record. None of the three is settled by demand.

8

Three lessons that travel

THE TRANSFERABLE LESSON

Three findings survive being lifted out of high-bandwidth memory and applied to another manufacturing market. Each surfaced only when this market was read as a manufacturing system.

The three lessons, and where each one is established

LESSON 1

When a product destroys capacity, the supply curve bends back

Ordinary growth markets pull supply toward themselves. This one removes roughly three units of alternative output for every unit it adds, which re-prices the alternative and puts the allocation decision back in play — which is how a premium product comes to lose a quarterly allocation contest to the commodity it displaced. A forecast that models demand without modelling the allocation choice has modelled half the problem.

Worked through in E4, E8 and scenario B

LESSON 2

Qualification cycles run longer than demand cycles

Every position in this chain was fixed years before the market that rewarded it existed. The leader's advantage dates from a packaging decision taken when HBM was a graphics part; the tool maker's share from having already built the right machine; the outsourced packager's growth from a contract signed in November 2021. In a market gated by qualification, the useful leading indicator is the qualification calendar rather than the order book.

Worked through in E17, E25 and E28

LESSON 3

The concentration that matters is rarely where everyone looks

Three suppliers hold more than 90% of DRAM revenue, and that is the concentration that gets discussed. One equipment supplier holds roughly seven-tenths of the tool on which every stack in the world is assembled, and one foundry supplies the base die for two of the three memory makers. Asking of each layer how many credible suppliers exist puts the fragility somewhere else entirely.

Worked through in E23 and E24

THE MARKET IN ONE LINE

The HBM market in one line: a capacity market with a price attached, gated by qualification, concentrated at the layers nobody names, and honestly forecastable only as far as 2027. Every structural feature in this study follows from the first clause — a fifth of the top three's DRAM wafer starts spent to produce under a tenth of the world's bits, by choice, and the choice revisited every quarter.

Working With Nathan Research

THE BENCH
AND THE OPEN QUESTIONS

Research houses establish the **shape** of this market. They do not settle the questions that decide a diligence outcome — how HBM4 allocation is really split, what stacking yields are, whether hybrid bonding resets the tool layer. That detail sits with the people who design, qualify, build, tool and package these products, and reaching them compliantly is what Nathan Research does.

NRG operates **Korea's first dedicated expert-network service, established in 2013**, built for the global private-equity, hedge-fund or corporate-strategy team carrying a thesis on a Korean or Asian semiconductor asset.

Who we put in the room

BENCH	WHERE THEY SIT IN THE CHAIN
Former memory executives	DRAM and HBM product, process and planning leadership from the Korean suppliers and their US peer
Packaging and stacking engineers	Through-silicon via, molded underfill, non-conductive film, hybrid bonding
Back-end equipment specialists	Bonder qualification, tool economics, the hybrid-bonding transition
Foundry and base-die contacts	The layer HBM4 pulled into a memory market
Outsourced assembly and test operators	Where packaging demand actually lands in Korea and Southeast Asia
Accelerator and hyperscaler procurement	Qualification timetables, allocation practice, stack-height decisions

How an engagement works

STEP 1

Scope

We translate your thesis into a precise expert profile and question set, mapped to the diligence decisions you need to close.

STEP 2

Source and vet

We identify, screen and compliance-clear each expert, confirming relevance, recency and the absence of conflicts before any call takes place.

STEP 3

Convene and synthesise

We arrange interviews on your timeline and, where useful, deliver written synthesis tied back to the questions this study raises.

THE QUESTIONS THIS STUDY LEAVES OPEN — AND WE CAN HELP CLOSE

Allocation. No contract discloses the split of Vera Rubin HBM4; every reported figure is an analyst estimate. **Yield.** Every stacking-yield figure in circulation is third-party. **The tool transition.** Whether hybrid bonding resets the bonder position or postpones the contest. **The back-end ring.** Which Korean packagers sit behind HBM demand, when no filing names the product.

Partner With Nathan Research

CONTACTS
AND HOW AN ENGAGEMENT STARTS

If your team is studying the **HBM market** — the memory suppliers, the packaging and equipment layers, or the Korean back-end ring behind them — or the wider **Korean and Asian semiconductor supply chain**, we would welcome the conversation. Tell us the decision you are trying to make, and we will tell you candidly whether and how our network can help you make it.

What to expect when you reach out. A partner replies to you directly, and the first conversation is about scope: the questions you need answered, and the people who can answer them. From there, a compliant and conflict-cleared expert panel is assembled to your timeline and, where useful, written synthesis that builds on the analysis in this study.

WEB	www.nathanrg.com	
EMAIL	team@nathanrg.com General enquiries	nathan@nathanrg.com Nathan Choi · CEO
PHONE	+82 10 3059 3630 Nathan Choi · CEO	+82 2 939 3630 Office · Seoul

COMPLIANCE This study is based solely on public disclosures, regulatory filings and named third-party research. Nathan Research does not request or facilitate the exchange of material non-public information, trade secrets, confidential customer or supply contracts, or undisclosed pricing, allocation, yield or capacity data, and runs every engagement through a documented compliance protocol with expert attestations and client-defined restricted-topic lists agreed before the first call.

A Appendix — method, and the market-sizing sources

AUTHORITY ORDER
WHAT IS CITED

METHOD

Market sizing is web-research-first and cross-checked across houses to 2026-08-10. Where houses disagree the range is printed with attribution and never averaged; where the record thins to observations that cannot be joined, the exhibit stops. Company disclosure grounds the market-evidence sections only. Authority order: SK hynix's SEC registration statement, which carries dated IDC and Gartner data under securities-law liability, then DART filings and company results releases, then JEDEC, SEMI, WSTS and Yole publications, then research-house notes, then single-outlet press and aggregators — and where a filing and a press compilation disagree, the filing governs.

Where an attribution names no house — an industry compilation or report, published architecture analyses, supply-chain analysts, an analyst or press compilation, vendor materials, channel checks or a third-party estimate — it marks a figure that circulates across several trade outlets with no single publisher behind it, and it is printed as a direction of travel rather than as a series.

KRX market data is near-unused in this study. Nothing here turns on what a listed participant is worth, so no price, volume or short-interest series appears anywhere and none was used to form a view. The one place a market-side cross-check would have applied — a Korean issuer's May 2026 rights offering against a listed-share-count change — could not be completed: final pricing fell after the panel data on hand, so the filing is cited alone and no share-count change is asserted.

NAMED SOURCES

MARKET SIZING AND FORECAST — SPECIALIST

Gartner — HBM, DRAM, memory and total semiconductor lines, data dated 2026-03-26, via SK hynix's 424B4; press cut 2026-04-08. **Yole Group** — Chiplet Summit 2026 near-term cut; press release 2025-06-19; FMS 2026 briefing via Korea JoongAng Daily, 2026-08-05. **Bank of America** (BofA), **Goldman Sachs** and **UBS** via SK hynix Newsroom, 2026-01-05. **TrendForce** — 2024-03-18, 2024-07-22, 2025-11-19, 2025-12-12, 2025-12-18, 2025-12-24, 2025-12-26, 2025-12-30, 2026-01-21, 2026-01-28, 2026-03-20, 2026-05-26, 2026-06-02, 2026-06-15, 2026-07-07 and 2026-08-04. **IDC** *Worldwide DRAM Demand and Supply*, 2026-05-27, and the Q1 2026 share reading, both via that filing. **Counterpoint Research**, 2026, and **Omdia** via press, 2026-06. **Astute Group** compilation, 2026. **Wedbush**, 2026-02 (design-services share). **Silicon Analysts**, 2026-08, and **Fubon** channel checks, cited 2026 (stack and per-GB prices). **Statista** and a **valueaddvc** compilation of company guidance, 2026 (hyperscaler capital expenditure).

MARKET SIZING — GENERALIST PUBLISHERS

Mordor Intelligence; **The Business Research Company** (TBRC); **Research and Markets** (R and M); **Market Research Future** (MRFR), whose US$3.73bn is the low end of the 2026 cluster; **Grand View Research**; **Precedence Research**. All 2026 cuts except Grand View's 2024–2030 data-centre-chip outlook. Carried as a separate band, never merged with the specialist family (E5).

STANDARDS, TRADE BODIES AND EQUIPMENT

JEDEC JESD270-4, published 2025-04-16. **WSTS** spring 2026 forecast, published 2026-05, and the autumn 2025 cut it revised. **SEMI**, 2026 — the US$145bn fab-equipment figure Unisem cites as its planning baseline, recorded here and nowhere else. The first open **High Bandwidth Flash** specification, from SanDisk and SK hynix, 2026-08-03/04.

Appendix — issuer and press sources, and the register

WHAT IS CITED
THE FOOTNOTE MAP

NAMED SOURCES, CONTINUED

ISSUER DISCLOSURE AND RESULTS

SK hynix SEC Form 424B4, filed 2026-07 — the registration statement for the Nasdaq ADR listing, carrying dated IDC and Gartner market data, capacity and remaining-spend schedules, customer concentration and listing proceeds under securities-law liability. **SK hynix Newsroom**, 2026-01-05, 2026-07-25 and 2026-07-29. **Samsung Newsroom**, 2026-07-30, and press summaries of the same day's call, labelled company-reported wherever used. **Micron** press release and FQ3 FY2026 call, 2026-06-24, and its 2025-12-18 statement. **NVIDIA** newsroom and chief-executive remarks, 2026-06-05 (via Reuters) and 2026-07-25. **Google**, 2025. **Broadcom**, 2026. Capital-expenditure guidance from **Amazon**, **Google**, **Meta** and **Microsoft**, 2026. Nine **DART** filings from five Korean issuers, listed by footnote in the register below.

PRESS AND TRADE

Asia Business Daily, 2025-12-22, 2026-02-11 and 2026. **Seoul Economic Daily**, 2026-06-08. **TechTimes**, 2026-06-04, 2026-06-06 and 2026-06-09. **Tom's Hardware**, 2025-04 and 2026-06. **TechPowerUp**, 2026. **Businesskorea**, 2026-04-06. **KED Global**, 2025-09-19. **CNBC**, 2026-02-06, 2026-07-10 and 2026-07-28. **Korea JoongAng Daily**, 2026-08-05. Korean trade ministry export releases via press, 2026-02, 2026-04, 2026-06 and 2026-08-01. **Semiconductor Engineering**, 2026. **Semiconductor Digest**, 2026. **The Elec** and **Korea Herald**, 2026 — the Hanwha Semitech hybrid bonder's 0.1 micrometre alignment specification and its April 2026 entry into an SK hynix line, both printed in E24 with no outlet named inline. The two order values beside them, ~₩80bn and ~₩9bn, come from Korean press reporting with no single outlet behind it.

SOURCE REGISTER — EVERY FOOTNOTE MARKER IN THE BODY

Each marker in the body resolves here, numbered in order of first appearance. Two kinds of entry share the series: the house, issuer or outlet standing behind a figure, set out at length in the named-source sections on this page and the last; and nine **DART** filings, whose receipt numbers are printed here and nowhere else. Exhibits carry their own sources under the plate, so a chart can be read without turning to this page. All five DART issuers are Korean back-end or design companies whose segment disclosure evidences where packaging, test and design demand is going.

[1] TrendForce, 2026-06-02 · **[2]** TechTimes, 2026-06-04 · **[3]** Gartner, via SK hynix SEC Form 424B4 · **[4]** Micron, 2025-12-18 · **[5]** published architecture analyses, 2024–2026 · **[6]** TrendForce, 2025-12-26 · **[7]** Gartner, 2026-04-08 · **[8]** CNBC, 2026-07-28 · **[9]** industry compilation, 2026 · **[10]** Goldman Sachs via SK hynix Newsroom, 2026-01-05 · **[11]** Wedbush, 2026-02 · **[12]** Broadcom, 2026 · **[13]** analyst compilation, 2026 · **[14]** Semifive, Annual Report FY2025 (receipt 20260323001690, 2026-03-23) · **[15]** LB Semicon, Quarterly Report Q1 2026 (receipt 20260515001798, 2026-05-15) · **[16]** supply-chain analysts, 2026 · **[17]** Micron, 2026-06-24 · **[18]** SK hynix Newsroom, 2026-07-29 · **[19]** Samsung Newsroom, 2026-07-30 · **[20]** TrendForce, 2025-12-24 · **[21]** TrendForce, 2026-06-02 and 2026-08-04 · **[22]** IDC and Counterpoint compilations, 2026 · **[23]** Yole Group via Korea JoongAng Daily, 2026-08-05 · **[24]** TrendForce, 2026-07-07 · **[25]** Hana Micron, Annual Report FY2025 (receipt 20260319000032, 2026-03-19) · **[26]** LB Semicon, Annual Report FY2025 (receipt 20260320000806, 2026-03-20) · **[27]** SFA Semicon, Annual Report FY2025 (receipt 20260323000826, 2026-03-23) · **[28]** Unisem, Annual Report FY2025 (receipt 20260313001168, 2026-03-13) · **[29]** Hana Micron, Quarterly Report Q1 2026 (receipt 20260514001268, 2026-05-14) · **[30]** Hana Micron, Quarterly Report Q1 2025 (receipt 20250515002716, 2025-05-15) · **[31]** LB Semicon, Report on Material Facts — rights offering (receipt 20260515002719, 2026-05-15) · **[32]** industry analyses, 2026; third-party estimates

These readings sit here because no market argument rests on them. Unisem supplies the WSTS and SEMI planning baselines, and its own revenue is a front-end capital-expenditure proxy. LB Semicon reports wafer output up 12.8% in FY2025. The May 2026 rights offering is ₩49.8bn over 12,000,000 shares, ₩30.0bn of it for facilities and ₩19.8bn for working capital.

Appendix — glossary, data gaps and scope

TERMS
WHAT IS MISSING
DISCLAIMER

GLOSSARY

HBM — high-bandwidth memory; DRAM dies stacked vertically on a logic base die and sold as a stack. **HBM3 / HBM3E / HBM4 / HBM4E / HBM5** — successive JEDEC generations. HBM4 doubles the interface to 2,048 bits and the channel count to 32 per stack. **TSV** — through-silicon via; the vertical connection between stacked dies. **Base die** — the logic die at the bottom of a stack, a foundry part from HBM4 onward. **MR-MUF** — mass reflow with molded underfill; bonds and fills a whole stack in one low-temperature thermal cycle. **TC-NCF** — thermocompression with non-conductive film; Samsung's advanced variant applies a film per layer at around 300°C under high force, with warpage risk rising as stacks grow taller. **Hybrid bonding** — direct copper-to-copper bonding without solder, the expected successor process.

TC bonder — the thermocompression tool on which stacks are assembled. **CoWoS** — chip-on-wafer-on-substrate; the 2.5D packaging route that places a stack beside an accelerator on an interposer. **OSAT** — outsourced semiconductor assembly and test. **Wafer input** — wafer starts committed to a product; the capacity measure this study uses instead of a revenue split. **Bit supply** — industry output measured in memory bits rather than in units or revenue. **ASIC** — application-specific integrated circuit; here, a custom AI accelerator designed for one buyer. **HBF** — High Bandwidth Flash; a NAND-based tier specified in August 2026 to sit between HBM and SSD, reaching 512GB per stack at bandwidth tiers of about 0.4 to 3.0 TB/s. **RDIMM** — registered dual in-line memory module; the mainstream server DRAM format.

DATA GAPS CARRIED FORWARD

(1) No single 2026 market size — two house families roughly ten times apart on definition. (2) No market size beyond 2027; the "~US$170bn by 2031" figure in circulation is excluded, for the reason the executive summary gives. (3) No continuous 2021–2025 revenue series — the 2024-to-2025 step is re-basing. (4) Samsung's 2024 share omitted for want of a primary source. (5) Two share readings no exhibit carries: Omdia's ~57% at mid-2026, and a ~70% attributed to TrendForce and flagged as an outlier. (6) No supplier-level HBM revenue — no issuer discloses it. (7) No accelerator unit volumes — estimates span 200k to 5.7m on inconsistent definitions. (8) No CXMT capacity or share. (9) Per-stack prices are channel checks, printed as ranges. (10) The 71.2% bonder share is one report on a cumulative-through-Q3 basis. (11) Korean trade statistics publish no HBM-only export line.

SCOPE AND DISCLAIMER

This is a **market-research study** prepared to support diligence on a market. It is **not investment advice**, not a solicitation, and contains **no price target, no valuation, no multiple and no buy or sell recommendation**. Company disclosures and results are cited only as evidence of market dynamics, never as a value argument. Forward bands are house-disagreement ranges rather than NRG forecasts of record, and figures labelled **[NRG estimate]** are our own arithmetic on published inputs, identified at the point of use. Prepared by Nathan Research Group, Seoul, 2026-08-10.