The Silicon Supercycle: How AI Infrastructure Is Rewriting Asia's Industrial Map

Semiconductors, Data Centers, Power, and the $1 Trillion Bet on Asian Hardware

Darmine Capital · Macro Research Division · Singapore

January 2026

For research and educational purposes only. Not financial advice (NFA). Not investment advice.

Abstract

The global semiconductor market hit $791.7 billion in 2025 — a 25.6% increase — and WSTS projects it will approach $975.5 billion in 2026, effectively reaching $1 trillion within the year. Asia manufactures over 75% of the world's chips and commands the critical nodes of the AI hardware supply chain: TSMC's foundry monopoly (77% of wafer revenue from sub-7nm), SK Hynix's 62% share of High Bandwidth Memory, and the APAC data center pipeline that now exceeds 19 GW. Meanwhile, the four US hyperscalers — Amazon, Alphabet, Microsoft, and Meta — have committed approximately $650 billion in capital expenditure for 2026, a 60%+ increase from the already-record 2025. This capex tsunami flows overwhelmingly through Asian semiconductor supply chains before it builds a single server rack. The resulting supercycle is reshaping equity markets (Korea's KOSPI +76–79% in 2025), creating new infrastructure bottlenecks (power constraints in Singapore, Japan, and Malaysia), and forcing a geopolitical reckoning over who controls the physical layer of artificial intelligence. This paper maps the AI infrastructure value chain across APAC and identifies the trade ideas we believe best capture this multi-year capex cycle.

1. The $1 Trillion Semiconductor Market

The Semiconductor Industry Association confirmed global sales of **$791.7 billion in 2025**, up 25.6% from $630.5 billion in 2024. Q4 2025 alone generated $236.6 billion, up 37.1% year-over-year, reflecting accelerating momentum as the year progressed. The WSTS Autumn 2025 forecast projects **$975.5 billion for 2026** — growth of over 26% — with memory and logic each expected to increase by more than 30%. The industry is on track to breach $1 trillion before year-end.

Segment	2025 Revenue	YoY Growth	2026 Forecast
Logic	$301.9B	+39.9%	>30% growth
Memory (DRAM + NAND)	$223.1B	+34.8%	>30% growth
Analog	Recovery	+7.5%	Moderate growth
Sensors	Recovery	+10.4%	Moderate growth
Total	$791.7B	+25.6%	~$975.5B

The critical insight: this is **not a broad-based recovery**. The semiconductor boom is overwhelmingly concentrated in AI-related logic and memory. Deloitte estimates the AI chip market alone will reach approximately $500 billion in 2026 — over half the total semiconductor market. Non-AI segments (automotive, consumer electronics, industrial) remain in sluggish recovery. This concentration creates both opportunity and risk: the entire supercycle depends on sustained hyperscaler capex.

Regionally, the Americas and Asia-Pacific remain the strongest contributors. Asia-Pacific sales grew 24.8% year-over-year in October 2025, while the Americas surged 59.6%. Japan was the sole regional decliner at -10%. The geographic mismatch — Asia manufactures the chips but the Americas generate the most revenue growth — reflects the hyperscaler-driven demand pattern where US companies design and buy the most advanced chips, but Asian fabs and memory makers produce them.

The unit economics tell a revealing story. SIA reported approximately 1.05 trillion chip units sold in 2025, implying an average selling price of $0.74. But this average obscures a radical bifurcation: a single NVIDIA H100 GPU sells for $25,000–40,000, while billions of commodity chips sell for pennies. The ASP for AI-relevant chips (advanced logic + HBM) is orders of magnitude higher than the market average and rising, while legacy chip ASPs are flat or declining. This pricing dynamic concentrates revenue and margins in companies serving the AI supply chain, particularly in Northeast Asia.

Memory led the structural shift. The DRAM market grew 72% in 2024, followed by an estimated 16.5% in 2025, with HBM as the sole high-growth driver. WSTS forecasts memory growing over 39% in 2026 to exceed $310 billion. Logic IC growth of 32% will be driven almost entirely by advanced foundry (TSMC 3nm/2nm) and custom ASIC design for cloud providers. The semiconductor industry is approaching $1 trillion, but it would be more accurate to describe this as two markets: a $500 billion AI semiconductor market growing 50%+ annually, and a $500 billion

legacy market growing low single digits.

2. TSMC: The Irreplaceable Chokepoint

TSMC's Q4 2025 results crystallize the AI hardware supercycle. Full-year revenue hit **$122.4 billion**, up 35.9% year-over-year in USD terms. Net income for Q4 reached NT$505.7 billion ($16.3 billion), up 35% YoY, with a gross margin of 62.3% and operating margin of 54%. Full-year EPS surged 46.4% to NT$66.25. Advanced technologies (7nm and below) accounted for 77% of Q4 wafer revenue, with 3nm alone at 28%.

Metric	Q4 2025	Full Year 2025	2026 Guidance
Revenue	$33.7B	$122.4B	~30% growth (~$159B)
Gross margin	62.3%	59.9%	63–65% (Q1)
Operating margin	54.0%	50.8%	54–56% (Q1)
Capex	$10.5B (Q4)	$40.9B	$52–56B
3nm revenue share	28%	N/A	Growing (N2 ramp)

The 2026 capex guidance of **$52–56 billion** — up from $40.9 billion in 2025 — represents the single largest annual capital expenditure in semiconductor history. Approximately 70–80% is allocated to advanced node technologies. TSMC is simultaneously ramping its Arizona fabs (second fab tool move-in planned for 2026, third fab construction started), expanding Kumamoto (Japan), and preparing 2nm (N2) for high-volume production. AI accelerator revenue reached high-teens percent of total revenue in 2025, with a forecast CAGR approaching mid-to-high 30s percent through 2029.

CEO C.C. Wei stated definitively: "We believe AI is real. Not only real, it is starting to grow into our daily life." TSMC's revenue grew 35.9% in 2025 versus foundry industry growth of 16% — meaning TSMC is taking share even as the market expands. Advanced packaging (CoWoS, SOIC) contributed approximately 8% of 2025 revenue and is expected to exceed 10% in 2026. The scarcity of CoWoS packaging capacity remains the binding constraint for AI chip production, not wafer fabrication.

The N2 (2nm) ramp is the next critical milestone. TSMC expects initial 2nm margin dilution of 2–3% for full-year 2026, a manageable cost given the node's premium pricing. The company is already accepting permits for a fourth Arizona fab and its first advanced packaging facility in the US. By platform, HPC represented 55% of Q4 revenue, smartphones 32%, IoT 5%, and automotive 5%. The HPC share has expanded from approximately 40% two years ago, reflecting the structural shift toward AI and data center workloads. Overseas fab ramp is expected to dilute gross margin by 2–3% initially and 3–4% in later stages, but TSMC's pricing power more than compensates.

Our view: TSMC remains the single most important equity in the AI infrastructure trade. Its Q1 2026 guidance of $34.6–35.8 billion with 63–65% gross margins implies further sequential acceleration, consistent with ~30% full-year revenue growth (~$159 billion). However, at current valuations, the stock prices in much of the good news. We prefer to express the TSMC thesis through less crowded supply chain names — particularly memory makers and packaging equipment suppliers.

3. The HBM Revolution: Memory Becomes the New GPU

High Bandwidth Memory has transformed from a niche product into the **most critical and highest-margin component** in AI infrastructure. BofA estimates the 2026 HBM market will reach **$54.6 billion**, a 58% increase from 2025. Some forecasts project the HBM market will surpass the entire 2024 DRAM market by 2028. Memory prices are expected to rise another 40% through Q2 2026 according to Counterpoint Research, and Samsung and SK Hynix have already raised HBM3E supply prices by approximately 20% for 2026.

Supplier	HBM Share (Q2 2025)	2025 Op. Profit	HBM4 Status
SK Hynix	62% (shipments)	■47.2T (record)	Dev complete, mass prod 2026
Micron	21%	Strong ($13.6B Q1 rev)	Samples shipped, prod 2026
Samsung	17%	■24.9T (memory)	Catching up, P5 fab by 2028

SK Hynix has emerged as Asia's most important AI stock. The company surpassed Samsung in operating profit for the first time in 2025, posting a record ■47.2 trillion versus Samsung's ■43.6 trillion across its entire business. SK Hynix holds over 60% of HBM shipments and 57% of HBM revenue. NVIDIA accounts for approximately 90% of its HBM supply. Goldman Sachs projects SK Hynix will maintain over 50% total HBM share through at least 2026, while UBS predicts approximately 70% share in HBM4 for NVIDIA's next-generation Rubin platform.

The memory industry's profitability has been transformed. Samsung's memory division and SK Hynix delivered gross margins of 63–67% in Q4 2025, surpassing TSMC's 60% guidance for the first time in seven years. This inversion — memory margins exceeding foundry margins — is historically unprecedented and reflects the structural shift in value toward HBM. SK Hynix boosted its 2025 capex to ■29 trillion (from an initial ■22 trillion target), a 30% increase driven by clear 2026 HBM demand visibility. Samsung plans to expand HBM production capacity by 50% in 2026. The entire HBM industry is sold out through 2026.

The transition from HBM3E to HBM4 is the next inflection point. Revenue mix in 2026 is forecast at approximately 55% HBM4 and 45% HBM3E. HBM4 features a 2,048-bit interface with data rates exceeding 2 TB/s per stack — double HBM3E throughput. The competitive implications are profound: SK Hynix's lead in HBM4 development gives it a structural advantage in securing NVIDIA Rubin and custom ASIC orders, while Samsung's catch-up effort will determine whether the HBM duopoly holds or fragments.

Our view: SK Hynix is the highest-conviction single name in Asian AI infrastructure. Its dominance in HBM, sold-out capacity through 2026, and positioning for HBM4/Rubin create a multi-year earnings visibility that few cyclical companies have ever enjoyed. Samsung's memory recovery is the higher-beta trade for those willing to bet on its HBM4 catch-up.

4. The $650 Billion Capex Tsunami

The four US hyperscalers — Amazon, Alphabet, Microsoft, and Meta — collectively plan to spend approximately **$635–665 billion in capital expenditure in 2026**, up roughly 67% from $381 billion in 2025, which itself was nearly 50% above 2024's $240 billion. The vast majority flows to AI chips, servers, and data center infrastructure. This spending is now large enough to be visible in US GDP.

Company	2025 Capex	2026 Guidance	Primary Use
Amazon	$131.8B	~$200B	AWS AI + cloud infrastructure
Alphabet	$91.4B	$175–185B	Servers (60%), DCs (40%)
Meta	$72B	$115–135B	DCs, servers, networking
Microsoft	~$83B	~$116–145B	Azure AI + Copilot infra
Total	~$381B	~$635–665B	+67% YoY

The scale is staggering. Amazon alone plans $200 billion — more than the GDP of 140 countries. JP Morgan noted that AI-related capital expenditures contributed 1.1% to US GDP growth in H1 2025, outpacing the contribution from consumer spending. This is not discretionary — Amazon's CEO described demand backlog at all-time highs, and Microsoft's unfulfilled order backlog doubled to $625 billion. The capex is chasing existing demand, not speculative build.

The free cash flow concern is real. Morgan Stanley projects Amazon's FCF to turn negative (–$17 billion) in 2026. Pivotal Research sees Alphabet's FCF plummeting nearly 90% to $8.2 billion from $73.3 billion. Barclays estimates Meta's FCF will drop up to 90%. Aggregate capex for the Big Five (including Oracle) now exceeds projected operating cash flows after buybacks and dividends, necessitating external funding — Alphabet held a $25 billion bond sale in November 2025, and its long-term debt quadrupled to $46.5 billion. Amazon disclosed it may seek to raise equity and debt.

For Asian supply chains, the implications are unambiguous: every dollar of hyperscaler capex passes through TSMC, SK Hynix, Samsung, ASE Technology, and dozens of Taiwanese, Korean, and Japanese component suppliers before it materializes as a functioning server. The capex flow is the revenue flow for Asian hardware. Even if the AI investment thesis is ultimately questioned, the 2026–2027 hardware orders are already committed.

The structural nature of this spend is key: aggregate capex for the Big Five, after buybacks and dividends, now exceeds projected operating cash flows — reaching 94% of FCF in 2025–2026 vs 76% in 2024. The hyperscalers collectively held $420 billion in cash at end-2025, but are increasingly tapping debt markets. Debt-funded capex is stickier than cash-funded capex because the infrastructure commitments are already contractual. Alphabet's stated strategy of building a "meaningful moat" through infrastructure creates a competitive dynamic where cutting capex means losing market share.

Our view: The hyperscaler capex wave is the most reliable forward indicator for Asian hardware earnings through 2027. We track quarterly capex guidance as the single most important macro variable for our APAC semiconductor positioning. Any deceleration in capex growth (not absolute level) would be the signal to reduce exposure. For now, the trajectory remains parabolic.

5. APAC Data Centers: The Physical Buildout

Asia-Pacific's data center market has reached **approximately 14 GW of operational capacity**, with a development pipeline of 19.4 GW (including 3.7 GW under construction and 15.7 GW in planning). Moody's projects the region's capacity will more than double by 2030, accounting for 40% of the global total. Mordor Intelligence estimates the APAC data center market will grow from $29.5 billion in 2025 to $79 billion by 2030, a CAGR of 21.8%.

Market	Status	Key Dynamic
Japan	Rapid growth	Grid power constraint, Osaka expanding
Singapore	1.2 GW new tranches	Moratorium lifted, sustainability mandates
Malaysia	Fastest growing	Johor corridor, Chinese hyperscaler demand
India	Highest growth CAGR	$67.5B Big Tech pledges, Hyderabad focus
Australia	Hyperscale boom	OpenAI/NEXTDC AUD 7B campus in Sydney
Korea	Constrained	Community energy concerns pausing projects

Singapore's decision to lift its data center moratorium with sustainability-linked permitting — releasing 1.2 GW of new development capacity — positions the city-state as a premium hub for specialized AI workloads. Singapore offers 10-year tax holidays for facilities achieving PUE below 1.3. But power and land scarcity push larger requirements to Malaysia's Johor corridor, which is attracting massive Chinese hyperscaler investment. Malaysia emerged as a global hot spot, receiving over $14 billion in data center pledges from US and Chinese firms in 2025.

India's data center market is projected to be the fastest-growing in APAC with a 20.5% CAGR to 2030. Amazon, Microsoft, and Google pledged a combined $67.5 billion in Indian investments since October 2024, with 80% announced in December 2025 alone. Japan continues to grow rapidly, with Microsoft committing $2.9 billion and SoftBank converting manufacturing plants to AI data centers in Osaka. However, grid power remains the biggest constraint across every market.

Malaysia and India together accounted for 58% of new operational data center capacity in H2 2025 and recorded the most significant increases during the period. The Edge data center segment is projected to grow at a 25% CAGR from 2025 to 2030 (Savills), driven by IoT, streaming, and 5G use cases in dense urban markets where land for large-scale facilities is unavailable. This creates investment opportunities beyond the hyperscale segment. Hong Kong's market is absorbing existing capacity with activity remaining flat, while Korea faces community pushback over data center energy consumption, pausing several major developments.

OpenAI's partnership with NEXTDC to build an AUD 7 billion hyperscale AI campus in Sydney sets the tone for 2026. Sovereign wealth funds are allocating unprecedented capital to digital real estate, and CBRE expects capital recycling to accelerate as owners dispose of stabilized assets to fund new AI-ready facilities providing 100 MW+ capacity.

Our view: APAC data center REITs and developers (Digital Core REIT, Keppel DC REIT, NTT DATA) are structural beneficiaries. However, the real opportunity is in the power infrastructure

required to feed these facilities. Electricity consumption from APAC data centers is projected to climb from 320 TWh in 2024 to 780 TWh by 2030 — a 165% increase. Utilities with grid-proximate generation capacity and cooling technology providers will capture disproportionate value.

6. Power: The Binding Constraint on AI Growth

Every major APAC data center market faces the same fundamental problem: **power grids built for distributed residential and industrial loads cannot support concentrated multi-hundred-megawatt data center campuses**. A single NVIDIA GB200 deployment consumes 30 MW continuously — more than entire business districts in most Asian cities. Singapore's 230kV transmission network cannot handle the 400kV connections that 100MW+ facilities require; upgrading requires $2 billion and five years of construction. Samsung's Pyeongtaek semiconductor complex already consumes 1 GW continuously, with AI chip production adding another 500 MW by 2026.

Only 32% of projected APAC data center demand will be met by renewable energy. Thailand's data center power demand grew 400% between 2020–2024, while generation capacity increased only 8%. Vietnam attracts hyperscale investment but suffers weekly power cuts during summer peaks. Microsoft opted to build its own power plants in Indonesia rather than wait for grid upgrades. Japan is strategically relocating data pools near offshore wind and nuclear sites. South Korea leverages nuclear for 28% of generation, but the government's renewable pivot creates uncertainty about future nuclear expansion.

Bank of America predicts APAC data center capacity will double within five years, adding 2 GW annually — double the 2018–2023 growth rate. The APAC data center power market is expected to grow from $8.6 billion in 2025 to $13.6 billion by 2030. The bottleneck has shifted from chips to physical shells to energy — and energy is the hardest to solve due to decade-long infrastructure lead times.

7. DeepSeek and the Efficiency Paradox

DeepSeek's January 2025 release of R1 — a reasoning model matching OpenAI's o1 at a fraction of the training cost — temporarily wiped over $600 billion from NVIDIA's market cap and sent AI supply chain stocks plummeting. The market feared efficient models would reduce hardware demand. Eleven months later, those stocks have not just recovered but continued to grow. The hyperscalers have collectively increased capex guidance by over 60%.

The resolution of the DeepSeek paradox follows Jevons' Paradox: efficiency gains lower the cost per unit of inference, which expands total demand for inference. DeepSeek's techniques (MLA, MoE) are now standard across all major frontier labs, but the result is more AI deployment, not less hardware. Alibaba responded by forecasting AI capex over the next three years exceeding its total spend in the previous decade. Chinese hyperscalers are building out infrastructure at scale using DeepSeek-derived models. The net effect: more chips consumed globally, not fewer.

DeepSeek's longer-term impact is the pluralization of AI development. The open-source R1 model was adopted heavily across Southeast Asia, Africa, and developing markets. DeepSeek reportedly holds 89% market share among Chinese AI users. The upcoming V4 model (expected February 2026) is reportedly a coding-optimized model that could outperform Anthropic's Claude 3.5 and OpenAI's GPT-4o on coding benchmarks. Wedbush's Dan Ives expects more market shocks from DeepSeek in 2026.

The sovereign AI movement — where governments invest in domestic AI infrastructure and models — is directly catalyzed by DeepSeek's demonstration that frontier capability does not require frontier spending. Japan's emphasis on physical AI and robotics, South Korea's focus on AI-safe semiconductors, and India's massive infrastructure pledges all reflect this shift. TSMC noted sovereign AI as a growing demand segment alongside consumer, enterprise, and cloud AI. The APAC AI market is projected to surge from $102.6 billion in 2025 to $816 billion by 2032 (34.5% CAGR), positioning it as the most consequential economic force in Asia since 1990s export-led manufacturing.

8. Export Controls and the Parallel AI Ecosystem

US semiconductor export controls have created a bifurcated AI ecosystem. The latest policy shift — replacing long-standing waivers with annual tool licences for Samsung, SK Hynix, and TSMC's China fabs — signals a tighter, more conditional framework. NVIDIA's China-specific H20 chip received conditional approval, while 80 new Chinese firms were added to the entity list in February 2026. The practical effect: China can access some AI hardware but not frontier chips, forcing domestic alternatives.

China's self-sufficiency push is accelerating. Ant Group developed AI training techniques using Chinese-made chips from Alibaba and Huawei that cut costs by 20%, achieving results comparable to NVIDIA's H800. SiCarrier, a semiconductor equipment maker with ties to Huawei, is developing lithography tools benchmarking against ASML. CXMT has reached 300,000 wafers per month in DRAM capacity and plans to increase to 500,000. However, competing with SK Hynix and Samsung in HBM remains years away — Chinese suppliers struggle with thermal management and operating speed at the cutting edge.

The investment implication is twofold. First, export controls create a structural premium for US-allied supply chains — TSMC, SK Hynix, Samsung, and ASML benefit from restricted competition. Second, China's parallel ecosystem creates a separate addressable market for domestic semiconductor equipment and mature-node foundries. South Korea's semiconductor blueprint includes $518 billion in investment through 2030, including SK Hynix's Yongin Cluster (potentially ■600 trillion / ~$407 billion). India Semiconductor Mission approved four additional fabs in 2025, pursuing domestic manufacturing.

Japan is also investing aggressively in semiconductor sovereignty. Rapidus prototyped 2nm GAA technology at its new foundry, Micron is building an advanced memory fab with government support, and the Ministry of Economy, Trade and Industry identified AI infrastructure as a core priority. ASML invested $164 million in a South Korea office complex, deepening its commitment to the Korean semiconductor ecosystem. Semiconductor equipment suppliers reported significantly increased collaborative research resources in 2025 — 3D heterogeneous integration, digital twins for manufacturing, and advanced packaging are the primary areas of focus.

Our view: The bifurcation is structural and investable. Long the "US-allied" supply chain (TSMC, SK Hynix, ASML, Tokyo Electron) as export controls create competition-free demand. Selectively long Chinese domestic semiconductor equipment makers (NAURA, AMEC) for the parallel ecosystem thesis, but size smaller given technology gaps and regulatory uncertainty.

9. Valuations and the Crowding Problem

Company / Index	Current P/E	Status	Risk
TSMC	~25–28x fwd	Premium, justified by moat	Overseas fab margin dilution
SK Hynix	~8–10x fwd	Cheap vs earnings power	HBM cycle peak concerns
Samsung Elec.	~12–14x fwd	Discount for conglomerate	HBM catch-up execution
KOSPI	10.4x fwd	Below historical avg	Political risk, crowding
TAIEX (Taiwan)	~18–20x fwd	TSMC-driven premium	Geopolitical tail risk
Tokyo Electron	~25–30x fwd	Priced for WFE cycle	Decel risk if capex slows

Korea's KOSPI at 10.4x forward earnings is the cheapest major semiconductor market in the world despite hosting the top two memory producers and delivering 76–79% returns in 2025. The Korea discount reflects political instability, corporate governance concerns, and the market's historical cyclicality. But this cycle is different: HBM creates a structural, multi-year demand floor that previous memory cycles lacked. SK Hynix's sold-out capacity through 2026 means earnings visibility is unprecedentedly high for a "cyclical" company.

Taiwan carries geopolitical tail risk that prevents full valuation convergence with US tech peers. However, TSMC's global fab diversification (Arizona, Japan, Germany) gradually reduces the concentration risk premium. The Japanese semiconductor equipment complex (Tokyo Electron, Screen Holdings, Advantest) trades at cycle-peak multiples but benefits from the TSMC and memory capex supercycle. Advantest, which dominates HBM testing equipment, is particularly well-positioned as HBM4 testing requirements are significantly more complex than HBM3E.

The ASML factor deserves attention. Bernstein raised its ASML price target to €1,300, noting ASML "stands to benefit enormously from capacity expansion planned for 2026–2027." While Dutch-listed, its order book is a direct proxy for Asian fab capex. SEMI reports that global silicon wafer shipments will rebound 5.4% in 2025 with new records expected by 2028, and sub-2nm wafer value alone is projected to reach $106.6 billion by 2030.

The crowding risk is real. BofA's Asia FMS shows semiconductors as the most overweight sector in Asia. Any sign of capex deceleration — whether from AI revenue disappointment, energy constraints, or simply CFO caution — could trigger violent de-grossing. The playbook should emphasize names with the highest earnings visibility and lowest valuation multiples relative to their structural positioning.

10. Portfolio Positioning Summary

Our conviction trade ideas for the AI infrastructure supercycle:

Highest conviction: Long SK Hynix — 62% HBM market share, record profits, sold-out capacity through 2026, HBM4 leadership for NVIDIA Rubin. Trading at ~8–10x forward earnings versus US memory peers at 15–20x. The most asymmetric risk/reward in Asian tech.

Core position: Long KOSPI as the cheapest semiconductor market — 10.4x forward P/E with the highest structural earnings growth driven by the AI memory supercycle. Korea NPS increasing equity allocation to 55% provides a domestic flow catalyst.

Supply chain beta: Long Taiwan semiconductor ecosystem (ASE Technology for packaging, Advantest for HBM testing) as TSMC's $52–56B capex flows through the supply chain. These names trade at lower multiples than TSMC itself while capturing proportional revenue growth.

Infrastructure play: Long APAC data center REITs and Singapore utilities as structural beneficiaries of the physical buildout. Keppel DC REIT and Digital Core REIT capture the capacity expansion, while Singapore's 1.2 GW data center release creates a decade-long growth runway.

Tactical: Long Samsung Electronics for the HBM4 catch-up trade. If Samsung can close the gap with SK Hynix in HBM4 qualification (expected mid-2026), the conglomerate discount could narrow significantly. Higher risk, higher beta.

Tail hedge: Short positions in consumer electronics semis exposed to memory price inflation. As AI workloads absorb HBM and DDR5 supply, consumer memory prices have surged (246% in 2025). NVIDIA cut gaming GPU production 30–40% in H1 2026 due to GDDR7 constraints. Consumer-facing companies without AI offsets face margin compression.

Key catalysts to monitor: Quarterly hyperscaler capex guidance (the single most important variable). TSMC monthly revenue data as a real-time demand indicator. HBM pricing trajectory — any sign of price stabilization would signal supply catching up to demand. Samsung's HBM4 qualification timeline with NVIDIA. NVIDIA's Rubin platform launch timeline and its HBM4 content specifications. APAC power grid expansion approvals, particularly in Singapore, Malaysia, and Japan.

11. Conclusion

The Silicon Supercycle is real, it is concentrated, and it flows through Asia. Three insights define our framework. First, the $650 billion in hyperscaler capex is not speculative — it is chasing documented demand backlogs at all-time highs. The question is not whether the money will be spent but whether Asian supply chains can absorb it fast enough. Second, HBM has permanently altered the memory industry's economics: SK Hynix posting higher margins than TSMC would have been unthinkable two years ago, yet it reflects a structural shift in where value accrues in the AI hardware stack. Third, power — not chips, not capital, not land — is the ultimate binding constraint. The markets that solve their energy equation (Japan with nuclear, Singapore with efficiency mandates, Australia with renewables) will capture disproportionate data center investment.

The supercycle will not end in 2026. TSMC's AI accelerator revenue CAGR of mid-to-high 30s percent through 2029, SK Hynix's projection of 30% annual HBM growth through 2030, and the physical reality that APAC data center capacity must double by 2030 all point to a

multi-year structural investment theme. The correction risk is real — crowded positioning, FCF concerns, and potential AI revenue disappointment could trigger sharp pullbacks. But pullbacks in a structural bull market are buying opportunities, not exit signals. Position accordingly.

Sources and Further Reading

SIA: "Global Annual Semiconductor Sales Increase 25.6% to $791.7 Billion in 2025" (February 6, 2026).

WSTS: Autumn 2025 Global Semiconductor Market Forecast. $772.2B (2025), $975.5B (2026).

Deloitte: "2026 Global Semiconductor Industry Outlook" (February 2026). AI chip market estimates.

TSMC: Q4 2025 Earnings Release and Earnings Call Transcript (January 15, 2026).

SK Hynix: Q4 2025 Earnings Release (January 29, 2026). Record operating profit data.

CNBC: "SK Hynix overtakes Samsung in annual profit" (January 29, 2026). Memory market dynamics.

BofA: 2026 Memory Supercycle Outlook. HBM market size ($54.6B), DRAM/NAND growth forecasts.

Goldman Sachs: HBM Market Analysis. ASIC demand growth, SK Hynix share projections.

Counterpoint Research: HBM market share data (Q2–Q3 2025). Memory price forecasts.

TrendForce: HBM3E pricing analysis, Samsung/SK Hynix capacity expansion reports (December 2025).

CNBC: "Tech AI spending approaches $700 billion in 2026" (February 6, 2026).

Sherwood News: "How Big Tech companies are spending their huge capex budgets" (February 10, 2026).

Fortune: "Big Tech’s $630 billion AI spree" (February 6, 2026). FCF impact analysis.

CBRE: "Asia Pacific Data Centre Boom to Continue in 2026" (2026). Regional market analysis.

Cushman & Wakefield: "APAC Data Centre Update H2 2025." 19.4 GW pipeline, 14 GW operational.

Mordor Intelligence: APAC Data Center Market ($29.5B to $79B by 2030, 21.8% CAGR).

Introl: "APAC Data Center Power Crisis" (February 2026). Power constraint analysis.

CSIS: "DeepSeek, Huawei, Export Controls, and the Future of the U.S.-China AI Race" (March 2025).

Eastspring: "Deep Seeking Opportunities in Asia’s AI Revolution" (March 2025).

BBF Digital: "Can AI Be Asia’s Next Growth Engine?" (February 2026). $102.6B to $816B by 2032.

Futurum Group: TSMC Q4 2025 Analysis. Advanced node mix, AI demand assessment.

SemiEngineering: Annual Global IC Fabs and Facilities Report 2025.