

The AI Supercycle and Bottlenecks: A Multi-year Tailwind for EM

(Abridged and updated)

As demand for artificial intelligence (AI) compute scales rapidly, concerns that hyperscalers may be ‘overbuilding’ have receded. The question increasingly in focus is whether supply bottlenecks around energy and chips will soon force a slowdown in the investment cycle. In this month’s Emerging View, we chart these bottlenecks and explore how they might affect the trajectory of capex, and the implications for emerging market (EM) assets.

The Semiconductor Supercycle: A key tailwind for EM


Global data centre capacity today stands at around 100 gigawatts (GW). McKinsey estimates AI will drive 125-205GW of incremental data centre demand by 2030. The AI capex cycle is following this demand. In the process, the largest transfer of capital from developed markets (DM) to EM in history is taking place. With US hyperscalers the key driver – Goldman Sachs estimates global AI capex spend between 2026 and 2030 to top USD 6trn – most of this spend is on semiconductors, for which the supply chain is overwhelmingly situated in EM countries.

Picks-and-shovels: From DM to EM

The first-order beneficiaries from the capex boom the prime contractors that manufacture AI chips, Taiwan Semiconductor Manufacturing Corporation (TSMC), SK Hynix and Samsung. TSMC delivered FY25 revenue of USD 122bn (+35% yoy) and net income of USD 55bn (+40% yoy), while Nomura sees Samsung/SK Hynix 2027 operating profit at USD 258bn and USD 214bn.

The result is that, based on 2027 forecasts, four of the world’s 10 most profitable companies are EM-listed: Samsung #1, Hynix #2, Aramco #4, TSMC #9.

Fig 1: World's most profitable companies in 2027(E)



Source: Nomura (Samsung & SK Hynix 2027E via BigGo February 2026); company filings, consensus 2027E for other names; Fortune Global 500 for Aramco. Four of the world's top ten profit pools sit in EM.

AI Investment has vastly exceeded expectations. Can this continue?

AI leaders speak as though no amount of capital is too much to pave the road that will fast-track society into the future. But ultimately investment cycles must generate an acceptable return, or they end. For now, adoption and revenue are still moving fast enough to justify further investment. Large language models (LLMs) have been the fastest major technology adoption cycle on record by far, even adjusting for today's larger population base.

Across all models, cumulative AI revenue stands at roughly USD 80bn-100bn today, from close to zero just four years

ago. This revenue trajectory has no precedent in the history of capitalism. However, the scale of investment necessitates that revenue continues expanding at a rapid clip. If not, capex is likely to be rationalised.

Perhaps the key reason, in our view, that AI is unlikely to resemble prior capex-driven bubbles is the tightness of the bottleneck around leading-edge chip production. This all comes back to there currently being only one company capable of manufacturing the most advanced AI chips at scale: TSMC.

Physical Constraints: Chips (Global Constraint)

Chip supply has historically scaled to cyclical consumer electronics demand, producing notorious booms and busts – a history that instils caution in producers. TSMC, SK Hynix and Samsung are run by leadership teams bruised by past oversupply and focused on not making the same mistakes again. But the AI era is shifting the nature of semiconductor cyclicality, with agentic AI likely to provide a multi-year demand boom for compute. Supply caution against vertical demand growth leads many analysts to conclude we may be entering a multi-year chip deficit.

a. TSMC: Leading-edge logic

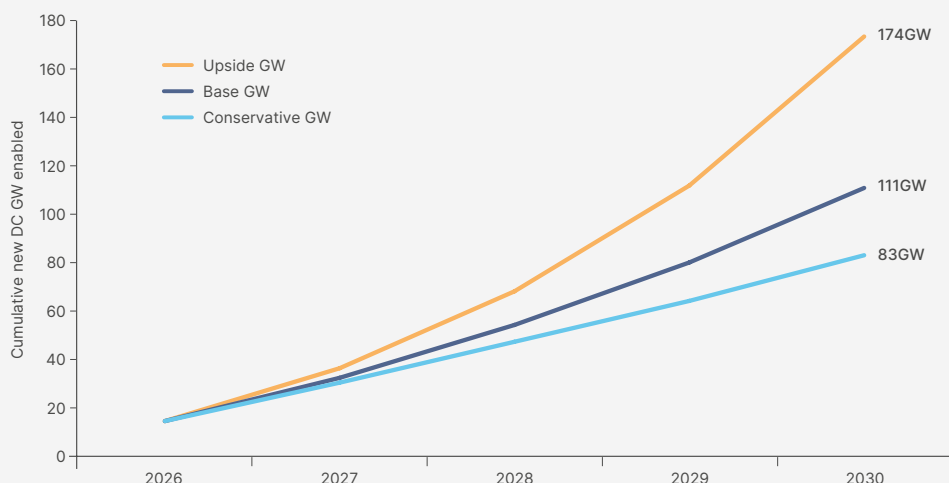
Dylan Patel of SemiAnalysis estimates that 1GW of Nvidia frontier AI capacity requires 55k 3nm wafers for graphics processor unit (GPU) logic.¹ Using this benchmark, three key variables determine how much new capacity leading-edge logic can enable: TSMC's 2nm/3nm supply growth, AI's allocated share (vs consumer/industrial electronics) and the replacement rate of the existing AI compute base. Using these three inputs, a crude benchmark for annual data centre enablement can be expressed as:

$$\text{Gross GW enabled} = (\text{TSMC 3nm/2nm wafer supply} \times \text{AI allocation share}) / 55k$$

$$\text{Net new GW enabled} = \text{Gross GW enabled} \times (1 - \text{replacement drag})$$

Scenarios with conservative, base and upside cases produce the results in Fig 2. Each scenario combines three assumptions: compound annual growth rate (CAGR) (10%, 15%, 25%), AI allocation (45%, 55%, 70%) and the replacement drag from refreshing the existing compute base (10%, 20%, 30%).

Fig 2: Illustrative net new data centre capacity enabled by leading-edge logic supply



Source: Ashmore estimates. Uses 55k 3nm/2nm wafers per GW of Rubin-class capacity. AI allocation and replacement drag scale linearly from anchors of 36%/0%.

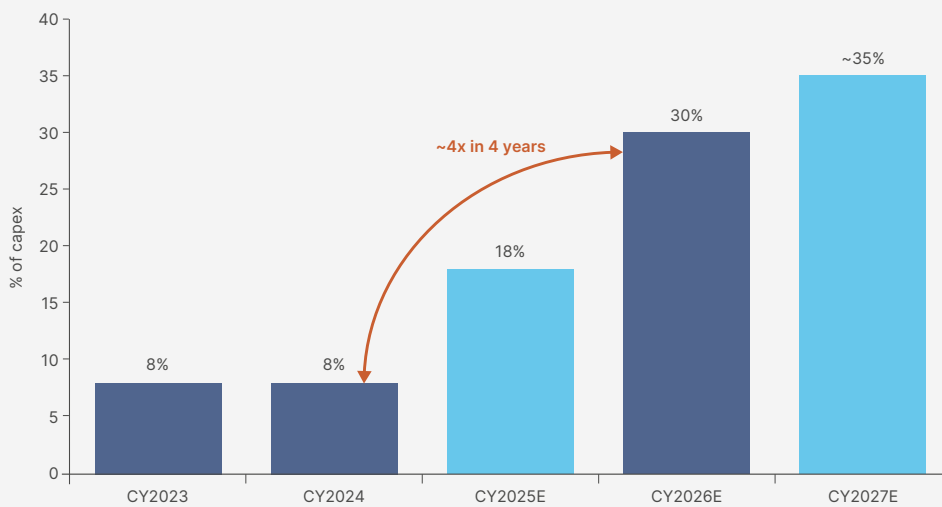
¹ See – <https://www.dwarkesh.com/p/dylan-patel>

b. Memory

As AI workloads shift further towards inference and agentic use cases, demand is rising across the full memory stack: HBM (high bandwidth memory), most specialised product and important for AI, DRAM (dynamic random-access memory) and NAND (not and). Supply is tight across all three and prices are soaring. Memory chips now represent an estimated 30% of AI capex, from 8% in 2024. Global HBM capacity is sold out through 2026 and most of 2027. SK Hynix forecasts ~30% annual demand growth through 2030.

Supply is expanding, but not at that pace. This points to higher-for-longer memory pricing, structural rent capture by the HBM oligopoly. This makes Samsung and SK Hynix still appear very cheap on a cash flow basis. Consensus free cash flow over the next three years represents a large share of current market capitalisation before assigning any terminal value to structurally higher HBM demand.

Fig 3: **Memory as a share of capex**



Source: SemiAnalysis via Tom's Hardware.

<https://www.tomshardware.com/tech-industry/memory-will-consume-30-percent-of-hyperscaler-spending-this-year>

As more value accrues to semiconductor bottleneck companies in EM, supply chains across Asia benefit disproportionately.

Physical (US) constraints: Energy

The second constraint is energy, most acute in the US, where power availability is already delaying data centre projects. The US Department of Energy's July 2025 Resource Adequacy Report sees only around 22GW of new firm capacity nationally by 2030; RAND's 2025 study puts net front-of-the-meter additions at 33GW on a fuller basis, including more solar and wind additions which are less useful for data centres. For our illustrative analysis, we take the midpoint, 28GW and assume data centres capture 70% of extra capacity, around 20GW.

The grid constraint has pushed most projects toward generating power 'behind the meter' (BTM) to bypass it. Longer term, BTM could meaningfully ease the bottleneck; for now, it shifts rather than solves it, since the power equipment itself, mostly gas turbines is also supply-constrained. We model that BTM could facilitate around 35GW in the US by 2030, though adjacent bottlenecks such as transformers could prevent even this.

Combined, that points to roughly 55GW of incremental US supply by 2030 – well short of demand. 451 Research forecasts 100GW of US demand growth between 2025 and 2030, while Synmax shows nearly 80GW of projects scheduled for 2025–2027 alone. Given the bottleneck, few will arrive on schedule.

Closing the gap requires innovation, and history cautions against betting against it. Better peak management is one route: the grid has ample spare capacity off-peak, with the binding constraint summer evenings, when air conditioning and lighting coincide. But reorganising US grid management around data centres will not happen overnight.

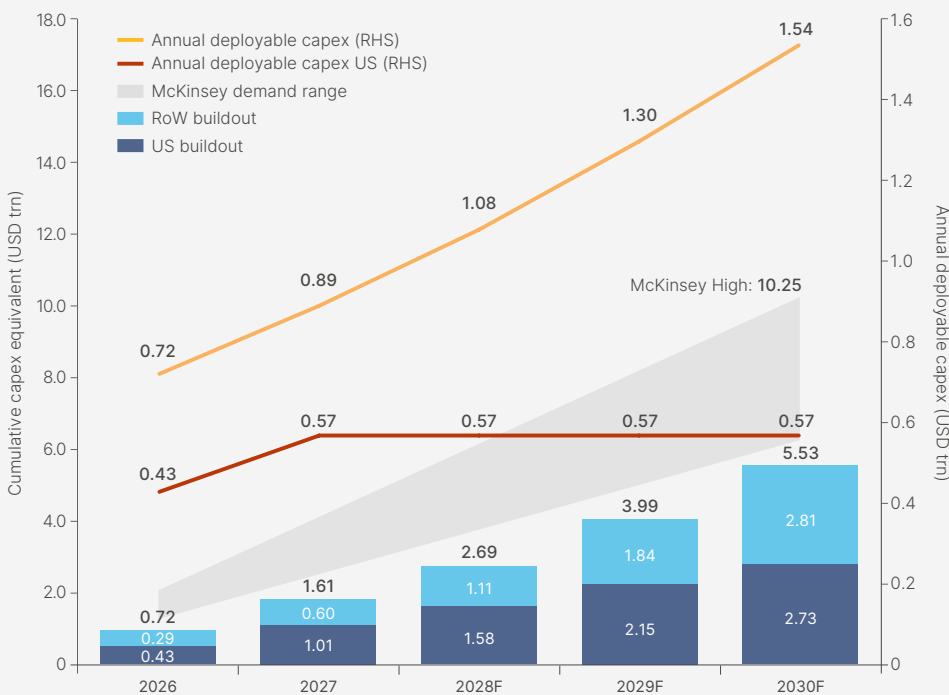
Constraints vs Capex

The simplest way to frame the limit is to convert buildable GW into deployable capex. At our USD 50bn per GW anchor, 55GW of effective US data centre power by 2030 implies around USD 2.75trn of deployable US AI capex. Still enormous – but it means US commitments are unlikely to keep growing at the recent pace, with USD 1trn already expected to be announced over 2026-2027.

If US power is the binding constraint, power-rich regions – Latin America, the Gulf Cooperation Council,

Scandinavia, and parts of Southeast Asia – can absorb some slack, and we expect more hyperscaler capex to flow there. But that does not solve the semiconductor constraint. Our base case of 111GW of AI-enabling production through 2030 undershoots even the low end of McKinsey's demand estimate (125GW). As TSMC's CEO put it last November, advanced-node capacity falls "about three times short" of current AI demand.

Fig 4: Global and US commitment of data centres against power and semiconductor ceilings



Source: Ashmore, 2026.

US 'annual deployable' capex in 2028/2029/2030 is a yearly average of the difference between what we model power enabling by 2030 and forecasts of US AI capex announcements across 2026 and 2027.

Macro implications: Three paths for the AI capex cycle

Scenario 1: Boom continues

Constraints ease faster than expected: BTM power scales, grid optimisation works, Chip on Wafer Substrate (CoWoS) and HBM supply expand, and non-US regions absorb the US power bottleneck. Investment stays above consensus and the cycle accelerates. Demand stays strong across the bottleneck complex – chips, memory, copper, power equipment, construction labour – and EM benefits broadly. Bottleneck-supplier margins may compress on greater supply, but profits stay high. USD impact mixed: strong US nominal growth supports it, a wider external deficit and stronger EM terms of trade lean against.

Scenario 2: Bottleneck-constrained cycle – base case

Demand stays strong but physical constraints bite. US power and global chip supply prevent the buildout scaling at the pace announcements imply; some projects move offshore, some defer, and spend tilts toward chips, powered shells and refits over speculative greenfield. Capex growth slows but does not collapse – fewer GW added than demand justifies, while bottleneck inflation keeps nominal spend elevated. This is the best relative setup for EM: the US bears more of the slowdown, EM captures the rents, and the USD weakens at the margin as EM balances improve. The trade is selective – own the constraints, not generic AI beta.

Scenario 3:

Demand-led rationalisation – bear case

The risk is not that supply constraints bite, but that end demand disappoints — revenues fail to scale, LLM commoditisation compresses token pricing, regulation raises compute costs, or confidence in returns fades. Capex slows as expected returns fall, delays become cancellations, order books normalise, bottleneck pricing power fades and the complex de-rates. DM may outperform defensively if lower inflation enables rate cuts, and the USD likely strengthens initially on risk aversion.

Bottom line

The AI buildout is here to stay, and is likely the transformational technology of our lifetime.

Unlike the dotcom-era fibre glut, GPU and data centre utilisation runs near 100% with demand higher still. Physical constraints will remain a headwind until innovation circumvents them, but this is not collapse: projects shift right, redirect to power-rich regions, and scarce suppliers keep pricing power. If demand keeps exceeding what chips, memory, power and grid can deliver, the implication is to own the constraints – now consensus. Less discussed is the multi-year EM tailwind this provides even if headline US capex flattens: a strong case for North Asian semiconductors, selected Asian hardware chains, power equipment, copper and power-rich EM regions.

Read the full version of this paper on our website

Emerging View

The AI Supercycle and Bottlenecks: A Multi-year Tailwind for EM

31 May 2026

By Gustavo Medeiros and Ben Underhill



As AI compute demand scales, the worry has shifted from overbuilding to scarcity. Can energy and chip supply keep pace, or will bottlenecks throttle the capex cycle? We map the constraints and their implications for emerging markets.

[Find out more →](#)

Head office

Ashmore Investment Management Limited, 16 Palace Street, London, SW1E 5JD T: +44 (0)20 3077 6000

Local offices

Bogota T: +57 601 589 9000	Lima T: +511 391 0396	New York T: +1 212 661 0061	Tokyo T: +81 03 4571 1601	Fund prices www.ashmoregroup.com Bloomberg FT.com Reuters S&P Lipper
Dublin T: +353 1588 1300	Mexico T: +52 55 4065 8898	Riyadh T: +966 11 483 9100	Qatar	
Jakarta T: +6221 2953 9000	Mumbai T: +9122 6269 0000	Singapore T: +65 6580 8288		

www.ashmoregroup.com  @AshmoreEM

No part of this article may be reproduced in any form, or referred to in any other publication, without the written permission of Ashmore Investment Management Limited © 2026.

Important information: This document is issued by Ashmore Investment Management Limited ('Ashmore') which is authorised and regulated by the UK Financial Conduct Authority and which is also, registered under the U.S. Investment Advisors Act. The information and any opinions contained in this document have been compiled in good faith, but no representation or warranty, express or implied, is made as to their accuracy, completeness or correctness. Save to the extent (if any) that exclusion of liability is prohibited by any applicable law or regulation, Ashmore and its respective officers, employees, representatives and agents expressly advise that they shall not be liable in any respect whatsoever for any loss or damage, whether direct, indirect, consequential or otherwise however arising (whether in negligence or otherwise) out of or in connection with the contents of or any omissions from this document. This document does not constitute an offer to sell, purchase, subscribe for or otherwise invest in units or shares of any Fund referred to in this document. The value of any investment in any such Fund may fall as well as rise and investors may not get back the amount originally invested. Past performance is not a reliable indicator of future results. All prospective investors must obtain a copy of the final Scheme Particulars or (if applicable) other offering document relating to the relevant Fund prior to making any decision to invest in any such Fund. This document does not constitute and may not be relied upon as constituting any form of investment advice and prospective investors are advised to ensure that they obtain appropriate independent professional advice before making any investment in any such Fund. Funds are distributed in the United States by Ashmore Investment Management (US) Corporation, a registered broker-dealer and member of FINRA and SIPC.