Not for redistribution without written consent of Morgan Stanley

Morgan Stanley | RESEARCH

INSIGHT

July 27, 2026 06:30 AM GMT

Global Thematics

Playing the AI Infrastructure Dip: Where to Invest and Where We See Risk

Following large moves in AI infrastructure stocks, we revisit key debates and concerns (Chinese models, tokenmaxxing, political pushback, "model lockdowns," physical bottlenecks) and ways to play the dip. Bottom line: we're bullish on the "Intelligence Superhighway" but see key speed bumps ahead.

Key Takeaways

- Our highest conviction view: the demand for compute is likely to significantly exceed supply for many years to come.
- We disagree with the "tokenmaxxing" argument that Enterprise limits on token spending will suppress AI revenue – we see the opposite playing out.
- Chinese LLM advances represent a real competitive threat to frontier model developers, but only bolster our "Jevon's Paradox" bullishness on compute demand.
- We are fundamentally bullish on the rate of improvement in AI capabilities, the benefits of AI adoption, and associated capex.
- We list stocks to play each part of the AI infrastructure stack and explore important potential Multipolar World policy dynamics with respect to AI competition.



Tech Diffusion

A Morgan Stanley Research Key Theme of 2026

Given the significant selloff in AI Infrastructure stocks over the past few weeks, we assess the biggest investor concerns driving at least a portion of the weakness. To be clear, we believe a meaningful driver of weakness has been technical rather than fundamental.

Concern #1: Reduction in "Tokenmaxxing" may limit revenue for LLM developers. The issue here is the belief that companies will place compute budget limits on their employees, as we have seen in a few high-profile examples. The data suggest that this is not a valid concern.

(i) Ramp estimates find the current median token spend by Enterprise users is extremely low (<$11/month),

(ii) the economics of AI usage are deeply in the money (see our $55 labor cost saving/ $2-3 token cost math here), and

(iii) Enterprise failure to adopt the most advantageous AI capabilities will lead to large competitive disadvantages, a dynamic we expect to become increasingly clear

MORGAN STANLEY & CO. LLC
Stephen C Byrd
Equity Analyst
Stephen.Byrd@morganstanley.com

Michelle M. Weaver, CFA
Equity Strategist
Michelle.M.Weaver@morganstanley.com

Ariana Salvatore
Equity Strategist
Ariana.Salvatore@morganstanley.com

MORGAN STANLEY ASIA (SINGAPORE) PTE +
Daniel K Blake
Equity Strategist
Daniel.Blake@morganstanley.com

MORGAN STANLEY & CO. LLC
Robert S Kad
Equity Analyst
Robert.Kad@morganstanley.com

Joseph Moore
Equity Analyst
Joseph.Moore@morganstanley.com

Brian Nowak, CFA
Equity Analyst
Brian.Nowak@morganstanley.com +1 212 761-3365

Adam Wood
Equity Analyst
Adam.Wood@morganstanley.com

Adam Jonas, CFA
Equity Analyst
Adam.Jonas@morganstanley.com +1 212 761-1726

MORGAN STANLEY MEXICO, CASA DE BOLSA, S.A. DE C.V.+
Nikolaj Lippmann
Equity Strategist
Nikolaj.Lippmann@morganstanley.com

MORGAN STANLEY & CO. LLC
David Arcaro, CFA
Equity Analyst
David.Arcaro@morganstanley.com

Joe Laetsch, CFA
Equity Analyst
Joe.Laetsch@morganstanley.com

Michael J. Cyprys, CFA, CPA
Equity Analyst
Michael.Cyprys@morganstanley.com +1 212 761-7619

Cameron McVeigh, CFA
Equity Analyst
Cameron.McVeigh@morganstanley.com

Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

+= Analysts employed by non-U.S. affiliates are not registered with FINRA, may not be associated persons of the member and may not be subject to FINRA restrictions on communications with a subject company, public appearances and trading securities held by a research analyst account.

over time.

Concern #2: Chinese open-weight models may become a greater competitive threat to American frontier LLMs, which in turn could drive lower spend on compute to train LLMs. While reserving judgment on the capabilities of recently released Chinese LLMs, we would cite the following caveats: (i) these models may actually be fairly compute intensive to achieve a given targeted result, implying that the price of a token is not equal to the value achieved by token usage (see analysis here), and (ii) we expect to see a range of LLM advances that drive more efficiencies, which will sustain the drive for better AI results at lower cost (e.g., increasing the usage of "orchestration layers" that allocate work to the most cost-effective LLM). Importantly, this single-minded pursuit of efficiency by both Chinese and American LLM developers reinforces our fundamental view that demand for compute is likely to vastly exceed supply.

Concern #3: The data center industry may face challenges in building out the needed volumes of compute, driven by the combination of People, Power and Politics. While we view this broadly as a valid concern, we see these issues more as speed bumps than barriers.
(i) People: Skilled data center construction workers, especially electricians, welders and pipefitters, are in short supply.
(ii) Power: Grid access is increasingly difficult to secure.
(iii) Politics: Pushback against data center growth is gaining traction in multiple US states and jurisdictions. These limitations may result in a US shortfall of data center capacity that could exceed 10% by 2028 relative to the amount needed based on our global Semiconductor analysts' AI chip sales forecasts.

We reiterate the core findings of our work:

1. The value of Intelligence, and compute, will continue to rise rapidly as AI capabilities increase at a non-linear rate.

2. Enterprise AI use cases are deeply in the money – the token cost is a small fraction of the benefit.

3. The ROI on AI capex is attractive. Our token economics model suggests both large LLMs and more efficient models generate strong ROIs on the underlying AI infrastructure.

4. We will continue to see Jevon's Paradox play out. This creates a massive quantum of economic value, the division of which is unclear and depends on which businesses own the most powerful "Intelligence Bottlenecks." To briefly explain: Jevons paradox describes a counterintuitive outcome: greater efficiency in using a resource lowers its effective cost, which can increase total demand enough to raise overall resource consumption. The mechanism operates through (1) greater consumption by existing users and (2) broader adoption across new users and applications. The key takeaway is that efficiency improvements alone do not guarantee lower aggregate consumption, especially when demand responds strongly to lower costs. This extends to AI: more efficient chips and models can reduce the cost of generating each unit of intelligence. Lower costs may expand AI adoption, increase the frequency and complexity of AI workloads, and ultimately

Josh Baer, CFA
Equity Analyst
Josh.Baer@morganstanley.com +1 212 761-4233

MORGAN STANLEY ASIA (SINGAPORE) PTE +
Jonathan F Garner
Equity Strategist
Jonathan.Garner@morganstanley.com +65 6834 8717

MORGAN STANLEY TAIWAN LIMITED+
Charlie Chan
Equity Analyst
Charlie.Chan@morganstanley.com +886 2 7305 7125

MORGAN STANLEY & CO. INTERNATIONAL PLC+
Shawn Kim
Equity Analyst
Shawn.Kim@morganstanley.com +44 20 7677 7018

MORGAN STANLEY TAIWAN LIMITED+
Howard Kao
Equity Analyst
Howard.Kao@morganstanley.com +886 2 7305 7389

MORGAN STANLEY & CO. INTERNATIONAL PLC+
Marina Zavolock
Equity Strategist
Marina.Zavolock@morganstanley.com +44 20 7425 5318

MORGAN STANLEY ASIA LIMITED+
Laura Wang
Equity Strategist
Laura.Wang@morganstanley.com +852 2848 6453

MORGAN STANLEY & CO. LLC
Devin McDermott
Equity Analyst and Commodities Strategist
Devin.McDermott@morganstanley.com +1 212 761 1725

MORGAN STANLEY ASIA (SINGAPORE) PTE +
Mayank Maheshwari
Equity Analyst
Mayank.Maheshwari@morganstanley.com +65 6834 6719

MORGAN STANLEY & CO. INTERNATIONAL PLC+
Lee Simpson
Equity Analyst
lee.simpson@morganstanley.com +44 20 7425 3378

MORGAN STANLEY ASIA LIMITED+
Robin Xing
Chief China Economist
Robin.Xing@morganstanley.com +852 2848 6411

Gary Yu
Equity Analyst
Gary.Yu@morganstanley.com +852 2848 6418

Yang Liu
Equity Analyst
Yang.Liu.CFA@morganstanley.com +852 2239 7911

MORGAN STANLEY & CO. INTERNATIONAL PLC+
Max R Yates
Equity Analyst
Max.Yates@morganstanley.com +44 20 7425 7417

Cedar Ekblom, CFA
Equity Analyst
Cedar.Ekblom@morganstanley.com +44 20 7425 3453

drive total demand for compute and electricity higher despite the efficiency gains.

5. We see an increasing probability of both US and Chinese AI policy intervention, which could lead to a greater bifurcation of the global market.

Our core investment recommendations:

(i) Own the AI Infrastructure bottlenecks: labor, "time to power" (fuel cells, Bitcoin players turned Powered Shell Providers, turbines, energy storage, power developers including energy companies that have become involved, and data center REITs with strong growth).

(ii) Own the compute manufacturing ecosystem, given the rapidly rising value of Intelligence and the shortage of compute relative to demand.

(iii) Own leading Chinese AI solution providers, given how capable and cost-competitive these companies' products are and how little related upside is priced into many of these stocks.

(iv) Own Energy Security assets that are important to AI Infrastructure development. We would flag Energy Storage as a standout asset class here (see Jack Lu's excellent report).

(v) Own hyperscalers with the capabilities to achieve scale benefits and drive attractive ROI on their AI capex (we are Overweight META, GOOGL, MSFT and AMZN).

Terence Tsui
Equity Analyst
Terence.Tsui@morganstanley.com +44 20 7425 3095

MORGAN STANLEY & CO. LLC
Cesar A Medina
Equity Analyst
Cesar.Medina@morganstanley.com +1 212 761 4411

MORGAN STANLEY TAIWAN LIMITED+
Daniel Yen, CFA
Equity Analyst
Daniel.Yen@morganstanley.com +886 2 7305 8803

MORGAN STANLEY C.T.V.M. S.A.+
Fernando P Amaral
Equity Analyst
Fernando.Amaral@morganstanley.com +55 11 3048 6011

Julia M Leao Nogueira
Equity Strategist
Julia.Nogueira@morganstanley.com +55 11 3048 6036

MORGAN STANLEY & CO. INTERNATIONAL PLC+
Ronald Ho, CFA
Quantitative Analyst
Ronald.Ho@morganstanley.com +44 20 7425 2722

MORGAN STANLEY & CO. LLC
Gregory Gao
Research Associate
Greg.Gao@morganstanley.com +1 212 296 3125

MORGAN STANLEY & CO. INTERNATIONAL PLC+
Stephan M Kessler
Quantitative Analyst
Stephan.Kessler@morganstanley.com +44 20 7425 2854

MORGAN STANLEY & CO. INTERNATIONAL PLC, SEOUL BRANCH+
Joon Seok
Equity Analyst
Joon.Seok@morganstanley.com +82 2 399 4934

MORGAN STANLEY ASIA LIMITED+
Daisy Dai, CFA
Equity Analyst
Daisy.Dai@morganstanley.com +852 2848 7310

MORGAN STANLEY ASIA (SINGAPORE) PTE.+
Kristal Ji
Equity Strategist
Kristal.Ji@morganstanley.com +65 6834 6919

MORGAN STANLEY ASIA LIMITED+
Vicky Wu
Equity Strategist
Vicky.Wu@morganstanley.com +852 3963 3028

Henry Zhao
Research Associate
Henry.Zhao@morganstanley.com +852 2239 7731

MORGAN STANLEY & CO. LLC
William Tackett, CFA
Research Associate
William.Tackett@morganstanley.com +1 212 761 0108

MORGAN STANLEY ASIA LIMITED+
Ehsernta Fu
Equity Strategist
Ehsernta.Fu@morganstanley.com +852 3963 3311

MORGAN STANLEY & CO. LLC

Anna Feldman

Equity Strategist

Anna.Feldman@morganstanley.com +1 212 761 3284

MORGAN STANLEY & CO. INTERNATIONAL PLC+

Cindy Huang

Equity Analyst

Cindy.Huang@morganstanley.com +44 20 7425 2415

MORGAN STANLEY & CO. INTERNATIONAL PLC, SEOUL BRANCH+

Heewon Choi

Equity Analyst

Heewon.Choi@morganstanley.com +82 2 399 4836

Minseo Kang

Research Associate

Minseo.Kang@morganstanley.com +82 2 399 4816

MORGAN STANLEY & CO. LLC

Joy Golub

Research Associate

Joy.Golub@morganstanley.com +1 212 761 4924

Garo K Amerkanian

Research Associate

Garo.K.Amerkanian@morganstanley.com +1 212 761 1207

Executive Summary

The single most important driver of our bullishness on the value of AI Infrastructure is the non-linear rate of improvement in AI capabilities. AI Infrastructure will, over time, become an "Intelligence Superhighway" that provides significant net benefits to economies around the world. It is difficult for all of us to think on a non-linear basis, and we notice that there is not enough focus on extrapolating the implications of AI adoption over time. While it may sound daunting for LLM developers to accumulate 10x the compute to achieve a 2x increase in capabilities, an important mitigant is the rapid improvement in GPU and TPU computational power per chip. For example, for NVIDIA, we project that each new chip generation will have 250% more exaFLOPs, with only 60% more power usage. For a tracker of AI training facilities, see the EpochAI tracker here.

The other element of the 10x:2x math that is not fully appreciated is the magnitude of importance relating to the doubling of model capabilities. Given that frontier LLMs can currently perform >50% of a wide range of human knowledge tasks at or better than humans (see analysis here), what might the economic implications be of a LLM that is twice as capable as the current best LLM? While difficult to determine precisely, we believe it is likely that, over time, the vast majority of knowledge tasks could be performed by frontier LLMs at or better than humans. A recent example, in the field of math: a team at DeepMind developed a LLM framework, AlphaProof Nexus, that "autonomously solved nine Erdos problems out of 353 attempted, including two problems that had been open for 56 years, at an inference cost of a few hundred dollars per problem." Another frontier LLM solved a long-standing Erdos problem last month.

A dynamic that might further accelerate the rate of increase in AI capabilities is the development work at the major LLM labs to incorporate recursive self-improvement (RSI) into LLM architectures. To ground yourself in the basic mathematical principles of how neural networks actually function, we recommend watching a video developed by 3Blue1Brown — this video is quite a beautiful way to understand what neural networks actually do.

What is so powerful about recursive LLM self-improvement is that this technology can result in:

- Much more rapid improvement in LLM capabilities, because the improvement can occur at "machine speed" rather than limited by the speed of human work to improve LLM capabilities, and
- Improvement without the necessity for the same magnitude of additional compute — this video describes the dynamic very well.

While we do not know when the models will incorporate full recursive self-improvement capabilities, we would note that a recent LLM was the first that partially wrote its own code, and various estimates suggest we may see this at some point between spring 2027 and 2030, though this development is subject to significant risk and uncertainty. When and if RSI is fully realized, such a development may bring with it both staggeringly large benefits of AI adoption, as well as far more intense concerns about AI disruption and risk of misuse and weaponization. It is also possible that only a few LLM developers achieve recursive self-improvement at a similar point in time, which would put other LLM

developers at risk of falling behind rapidly in terms of LLM capabilities.

Moonshot K3 has driven significant market weakness across a wide range of AI Infrastructure stocks. In response, we cite Morgan Stanley's Asia Internet and Telecoms analyst Gary Yu:

- *"K3, the largest open-weight Chinese LLM so far, may suggest the scaling law still holds. Its high pricing is also positive to LLM market dynamics. We do not view K3 as an overnight miracle but rather as the result of cumulative progress across China's AI model industry. GLM-5.2, released in June, was already outstanding. We expect more Chinese LLMs with larger sizes, higher pricing, and better performance, with strong global competitiveness."*

Further, our Intelligence Factory model projects that ROIC for a 200MW data center operating with Kimi K3 (as opposed to the extant frontier model) sees a modest ~3% increase in margin. Kimi is clearly a strong model, evidenced by its performance on the benchmarks, but it appears to continue the trend of increasingly strong frontier models (reflected by its parameter count and price) rather than a smaller open-weight model like Qwen or MiniMax.

Urgency to spend on AI Infrastructure spend remains high among key players, and we expect the upcoming earnings season to reflect this, along with bullishness on returns on AI capex. As noted by US Internet analyst Brian Nowak in a recent report, "we expect capex spend to continue to head higher (now expected to reach to $1.2tr/$1.4tr across the five primary companies we track) and are laser focused on companies' ability to bring on capacity to drive and benefit from incremental AI-enabled revenue (from bare metal compute to fuller stack first/third party software and services)...we see this spend leading to more compute capacity...with total hyperscaler available compute expected to approach 120GW by '28, up from ~30 in '25."

- While GOOGL is expected to add the most compute capacity (9 GW/11 GW in '27/'28) we see AWS having the most available compute in '28 (at 35 GW vs GOOGL at 31 GW).
- We see META reaching 14/21 total GW of capacity by '27/'28, up from ~3.5 GW at YE25.
- In this report, Brian highlights key upcoming AI monetization data points for AMZN, META and GOOGL.

Given the above takeaways, and given the recent market pullback affecting a range of AI Infrastructure stocks, we believe this point in time represents an unusually attractive buying opportunity.

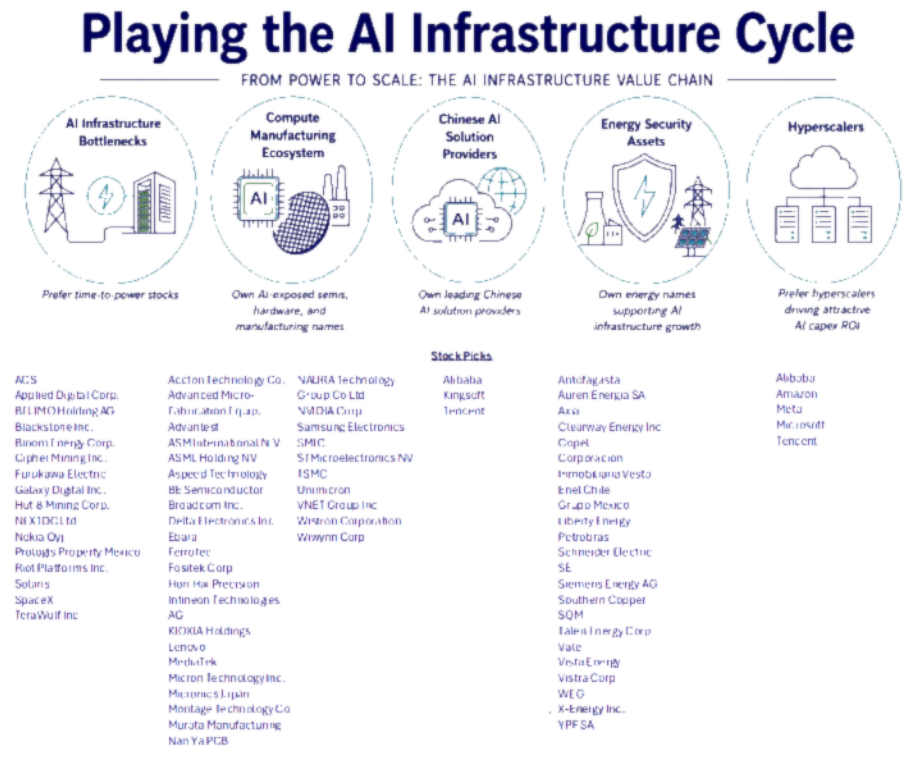
Our core investment recommendations:

- **Own the AI Infrastructure bottlenecks:** labor, "time to power" (fuel cells, Bitcoin players turned Powered Shell Providers, turbines, energy storage, power developers including energy companies that have become involved, and data center REITs with strong growth).
- **Own the compute manufacturing ecosystem,** given the rapidly rising value of Intelligence and the shortage of compute relative to demand.

- **Own leading Chinese AI solution providers,** given how capable and cost-competitive these companies' products are, and given how little related upside is priced into many of these stocks.
- **Own Energy Security assets that are important to AI Infrastructure development** — and we would flag Energy Storage as a standout asset class here (see Jack Lu's excellent report here).
- **Own hyperscalers with the capabilities to achieve scale benefits and drive attractive ROI on their AI capex**
  - **We are Overweight META, GOOGL, MSFT, and AMZN.**

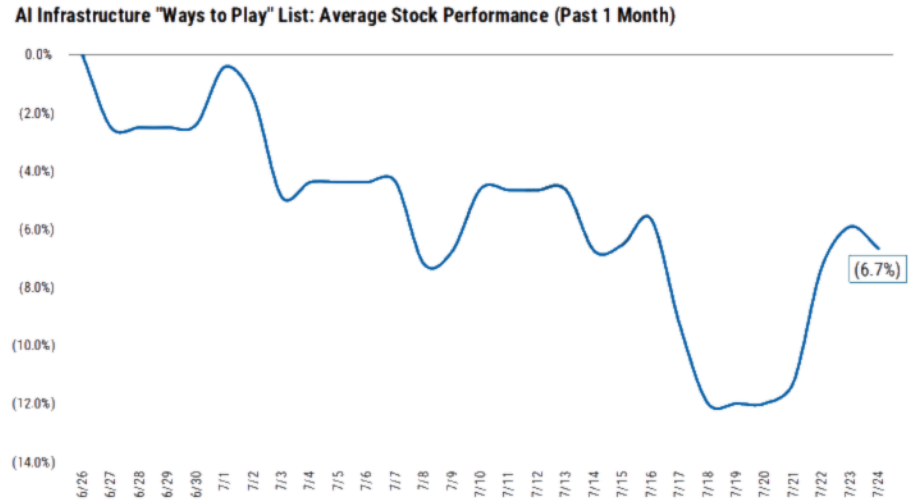
The following table highlights a selection of Overweight-rated stocks in each category (see more in Ways to Play). In aggregate, these stock have been underperformers over the past month, underperforming indices by ~7 points, despite strong secular positioning.

Exhibit 1: We prefer AI infrastructure bottleneck players, computing manufacturers, Chinese AI solution providers, energy security assets, and hyperscalers to play this theme.



Source: Morgan Stanley Research

Exhibit 2: On average, these names have been underperformers over the past month, despite strong fundamentals and AI infrastructure exposure.



Source: Morgan Stanley Research, Factset

What is included in this report

A Closer Look includes an assessment of the key AI debates we are hearing from investors, an explanation of why "time to power" mathematics is such a powerful alpha driver, and a quantitative approach to capitalizing on periods of market weakness, such as the one we are currently experiencing.

Key Morgan Stanley Analytical Tools: many investors may not appreciate the breadth and depth of Morgan Stanley analytical tools that can be used to assess the key AI debates. We would flag:

- Our Token Economics model, showing the economics for both AI Adopters and owners of the compute infrastructure,
- Our analysis of the US power shortfall facing data center developers,
- An analysis of US pushback against data center developers,
- Our survey of Enterprise AI Adopters, which includes data on employment dislocation driven by AI. We also include consumer survey work.
- Our AI Mapping database, which provides an assessment for all stocks covered by Morgan Stanley globally and likely AI-driven impacts to these companies' businesses.
- Our AI and the Future of Work analysis, which maps task-specific AI usage data to S&P 500 companies to estimate the magnitude of potential benefit to these companies from AI adoption.

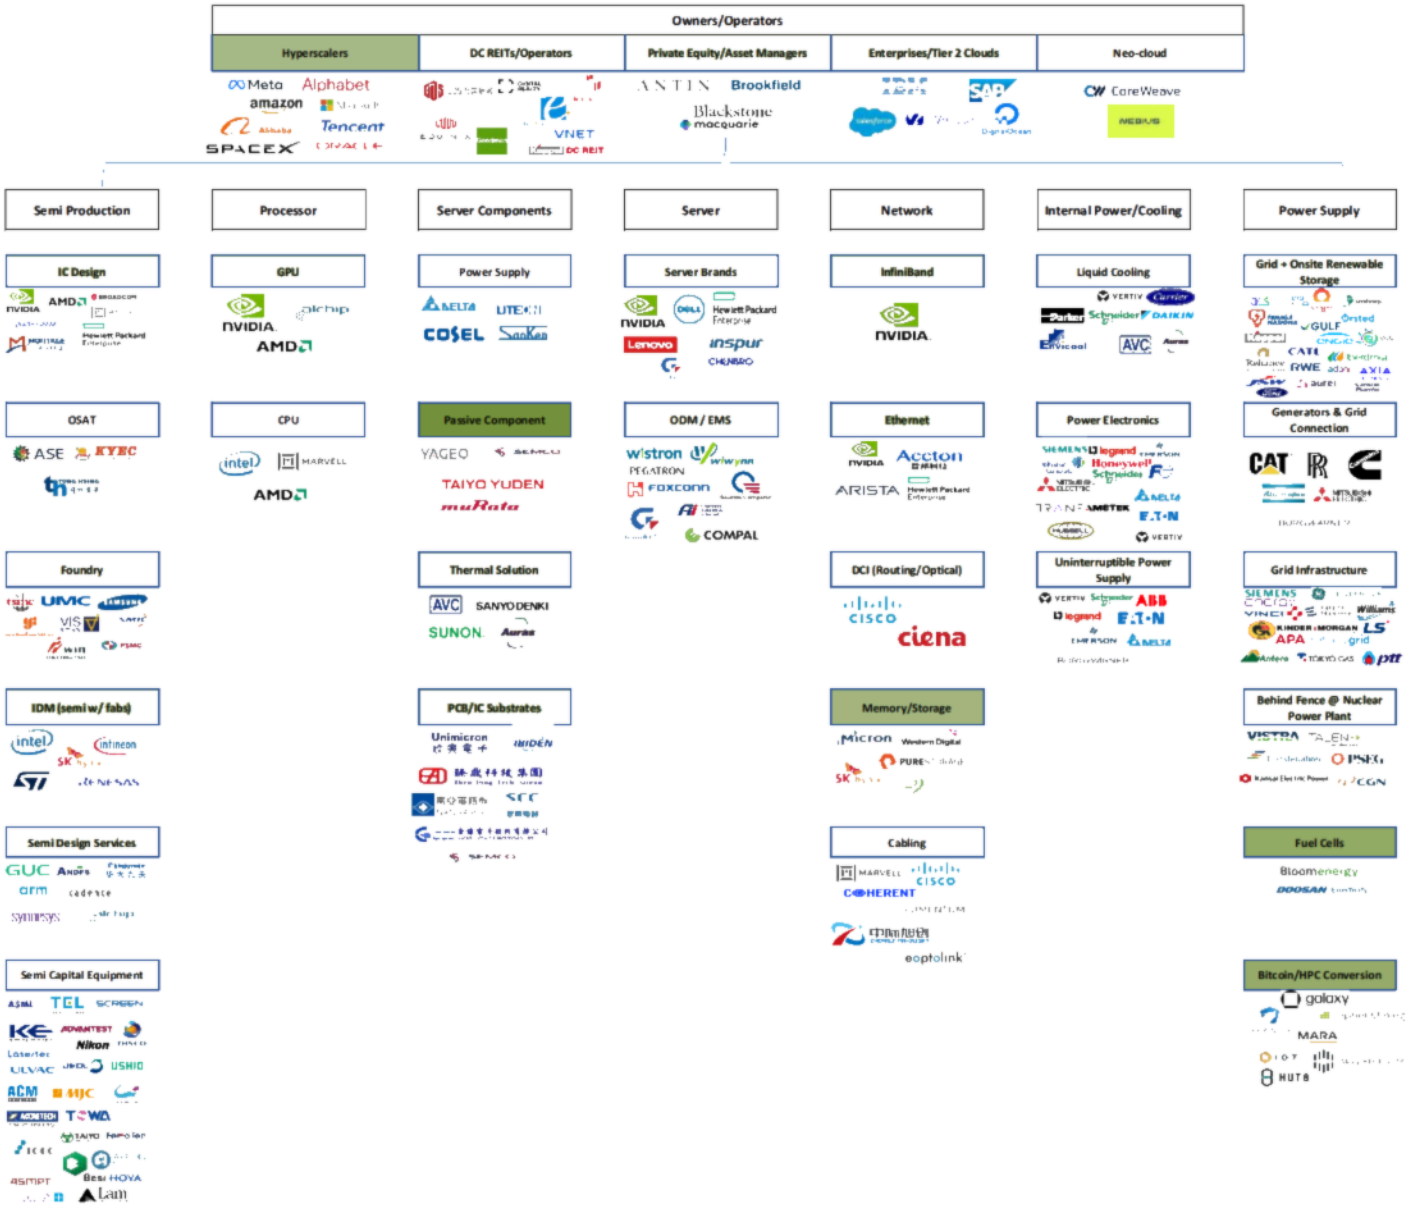
Overview of Key Morgan Stanley Reports: Morgan Stanley research analysts have been prolific in assessing the key AI debates; in this section of the report, we provide links to some of the most important of these reports.

The Potential for "Two Worlds" of AI Policy and Adoption: we explore the potential for US and Chinese AI policy that results in a sharper divide in approaches to AI access and

adoption.

AI Infrastructure stocks, especially those that can de-bottleneck data center growth, rise in value. If the portion of global GDP that AI can address at lower cost and better performance increases, the infrastructure that enables such value creation would rise in value. We have developed a complete map of AI Infrastructure, along with a list of stocks across that value chain and a "heat map" showing the relative attractiveness of stocks across the value chain at any given point in time (corresponding to the degree of upside to Morgan Stanley analyst price target — the deeper the green shading, the more attractive the median upside to Morgan Stanley price targets is) — **please click here to ask us for this analysis.**

Exhibit 3: AI Infrastructure Value Chain Heatmap



Source: Morgan Stanley Research, Data as of 7/15/2026

A Closer Look

A deeper dive into the big AI capex debates — please reach out to us to discuss in more depth

Concern #1`

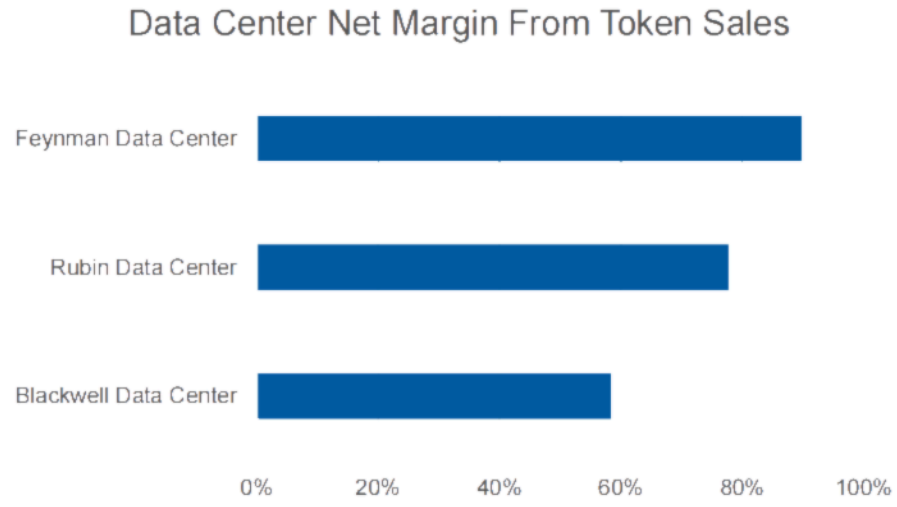
Reduction in "Tokenmaxxing" may limit revenue for LLM developers

The concern here would be that companies limit token spending among employees, as we have seen in a few high-profile examples. The data suggests that average Enterprise employee spend on tokens is actually quite low, with the potential to increase significantly in the coming months and years. Click here to request a copy of our Intelligence Factory model.

From the perspective of an Enterprise adopter of AI, the economics are compelling. For example, in a study of a wide range of Enterprise AI use cases, the average single use would save an Enterprise customer $55 in costs. While it is not possible to precisely know how many tokens will be used to achieve this outcome (and token usage will likely vary significantly across Enterprise use cases), we can ground ourselves in some helpful math. One analysis suggests that a single AI Agent may consume ~15x the tokens relative to a query-based usage of a LLM. A typical query-based usage of a LLM would consume ~5,000 tokens. If an Enterprise customer used 5-7 AI Agents to execute an economic task, the token usage would be 375,000-525,000 tokens.

Token pricing can be found here — you will see that pricing can range widely among models and between input and output tokens. At an assumed average token price of $3-10 per million tokens (in general, input tokens will outweigh output tokens, but not always), the cost to execute the economic task would be $2-5, a fraction of the $55 benefit. From the perspective of the LLM developer, based on token pricing that is publicly available, the Blackwell data center would generate margins of ~60%, factoring in all data center costs (but not the cost to develop the LLM). Interestingly, when we run the same math for Rubin and Feynman data centers, the margins would rise to ~80% and 90%, respectively. This suggests that, in the future, token pricing could fall as the share of global compute is comprised of increasingly capable chips.

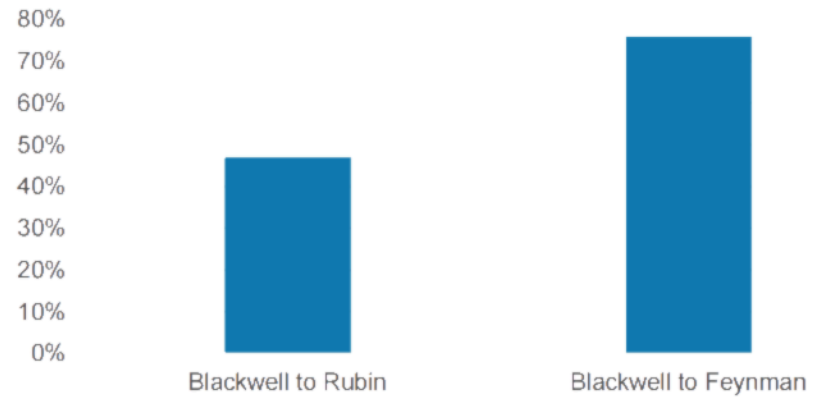
Exhibit 4: Our Intelligence Factory Model Projects 58-90% Margins for Token Sales, Depending on GPU



Source: Morgan Stanley Research

Framing our token economics math differently, our model shows that, as we move from the Blackwell to the Feynman, hyperscalers could price tokens 75% lower and achieve the same margins on a leased data center:

Exhibit 5: Subsequent generations of NVIDIA GPUs could lower token pricing by ~75%



Source: Morgan Stanley Research

Investment implications: Bullish for AI adopters with strong moats; bullish for ROI on AI capex.

Concern #2

Chinese open weight models may become a greater competitive threat to American frontier LLMs, which in turn could drive lower spend on compute to train LLMs.

While we reserve judgment regarding the capabilities of recently released Chinese LLMs,

we would point to a few initial conclusions:

- These models appear to be highly capable,
- These models may actually be fairly compute intensive to achieve a given targeted result, implying that the price of a token is not equal to the value achieved by token usage (see a helpful analysis here), and
- We expect to continue to see a range of LLM advances that result in better efficiencies (an example being the rise in the usage of "orchestration layers" that allocate work to the most cost-effective LLM).

These efficiencies will continue the relentless move to better AI results at lower cost.

Importantly, this relentless pursuit of efficiency (by both Chinese and American LLM developers) reinforces our fundamental view that demand for compute is likely to vastly exceed supply. For example, a Google exec recently stated that the company likely needs to double compute every 6 months, with a result of "1,000x in 4-5 years" (a doubling every 6 months would in fact result in >1,000x compute within 5 years). To put this growth rate in perspective, in 2025-28, our projected CAGR in compute sold by NVIDIA is ~140% per year; extrapolated over a 5-year period to line up with the Google statement, in 5 years the compute accumulated would be <10% of Google's projected growth in compute demand.

Investment implications: Bullish for AI inference-focused infrastructure broadly (chip manufacturers, neoclouds, data center developers), equipment that is critical to data center growth (memory, optical, backup power, energy storage, transformers, cooling), power access "de-bottlenecking solutions" providers (fuel cells, turbines, energy storage, Bitcoin site repurposing, and select repurposing/usage of existing power plants by data centers). Bullish for key Chinese AI solutions providers, given signs of increasing competitiveness and a strong value proposition; in many markets around the world, we believe adoption of Chinese LLMs is likely to increase rapidly given the performance of models like Kimi K3.

Concern #3

The data center industry may face challenges in building out the needed volumes of compute, driven by the combined challenges with respect to People, Power and Politics. We view this broadly as a valid concern, though we see this as more of a "speed bump" rather than a brick wall.

Let's start our discussion with an analysis of political pushback to data center development. Throughout the year, issues such as utility rate freezes, heightened scrutiny of customer bills, shielding residential ratepayers from costs associated with data center infrastructure, lowering electricity bills, and slowing data center development have increasingly become part of gubernatorial campaign platforms, and we would expect this to be a an important voter issue in November. Further, over the past decade, states competed with one another to offer generous packages to attract data centers, including sales and use tax exemptions, property tax abatements, long-duration windows, and direct subsidy packages. We are now beginning to see a reversal, with states now moving to

pause, condition, or repeal data center tax breaks. While largely not permanent in nature, we still see these actions as a risk to future data center build-out, potentially contributing to pauses or delays while emphasizing parts of the country which still have incentive mechanisms in place.

Exhibit 6: Recent State Rollbacks of Data Center Tax Incentives

State	Status	Action
Alabama	Enacted	Legislation to modify (limit) current data center exemptions
Arizona	Enacted	3-year moratorium on new sales tax exemptions for data centers
Delaware	Enacted	Legislation to eliminate state tax credits and license fee reductions for most large data centers
Illinois	Enacted	Pause of state tax incentives for new data center developments
Iowa	Enacted	Reformed limits on new sales tax exemptions
Massachusetts	Enacted	Pause of the acceptance of applications for the state's sales and use tax exemption, pending stronger ratepayer and environmental protections
Nebraska	Enacted	Eliminated a dedicated sales and use tax exemption; tax breaks still exist
New Jersey	Enacted	Pause on applications for the Next New Jersey Program - AI
North Carolina	Enacted	Provision to remove sales tax exemption for electricity used on site; complete repeal proposed
Ohio	Enacted	Pause of new data center sales tax exemptions
Oklahoma	Enacted	Removal of tax incentives for any data centers not operational by Jan 1, 2027
Oregon	Enacted	1-year moratorium to block new data centers in urban and suburban areas from accessing property tax incentives
Washington	Enacted	Removal of tax exemptions for equipment replacements and refurbishments
Michigan	Introduced/Pending	Legislation to repeal sales and use tax exemptions
Montana	Introduced/Pending	Package of bills to repeal tax break structure
Pennsylvania	Introduced/Pending	Legislation to eliminate or restructure sales tax exemption
Wisconsin	Introduced/Pending	Bills to add strict requirements for tax break qualification

Source: Morgan Stanley Research

On the national level, the Ratepayer Protection Act is being considered in the House as one of the first federal attempts to allocate infrastructure build-out costs — the bill would require state utilities to consider creating a "large load standard" requiring data centers to pay for grid upgrades. Recall, in March, Amazon, Google, Meta, Microsoft, Oracle, xAI and a frontier lab signed the Ratepayer Protection Pledge, a voluntary agreement from the White House committing to shielding existing consumers from the data center infrastructure build-out. The Ratepayer Protection Act would codify elements of the pledge and would help essentially create a national data center tariff, which is currently being implemented in a piecemeal manner across utility jurisdictions and at the state level. As grid interconnection lead times extend beyond five years in certain regions and pressure increases for data centers to directly secure or fund the power they consume, onsite generation is emerging as an attractive solution to address these concerns. These "de-bottlenecking" solutions include small scale gas turbines, reciprocating engines, fuel cells, and battery storage and allow data centers to get access to electricity ahead of grid interconnection. With an estimated ~38 GW potential shortfall in power for US data centers through 2028 (growing to ~122GW through 2030), we see an expanding TAM for such solutions that can accelerate access to power for large loads. :

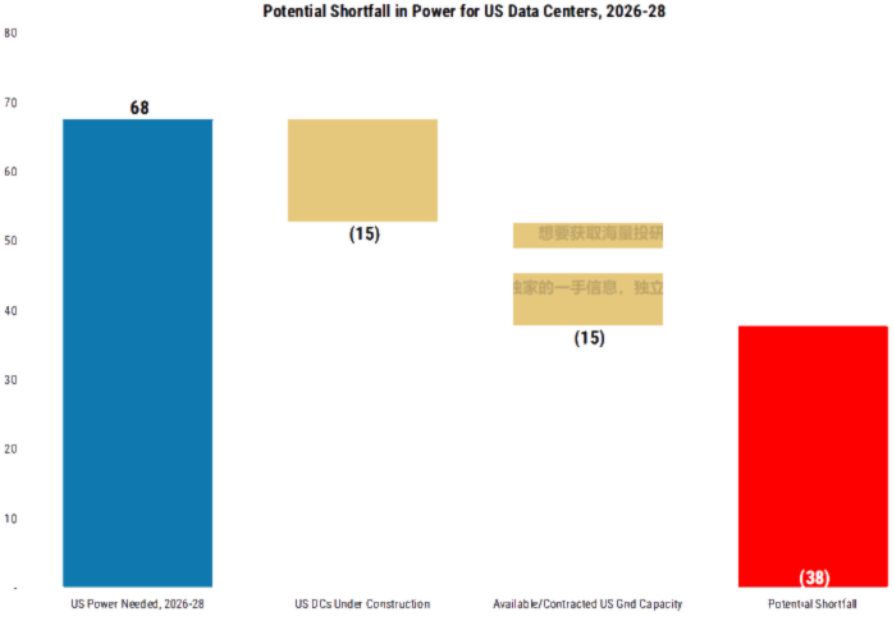
Time to Power Solutions

Let's shift our focus to power grid bottlenecks facing data center developers. We estimate a 38 GW potential shortfall in available US power grid access relative to the magnitude of data center demand in the next several years. With a substantial ramp in data center power demand we expect lengthy grid interconnection delays with some regions already experiencing 5- to 7-year wait times. With increasing grid constraints, we see an expanding TAM for "de-bottlenecking" solutions that can accelerate the access to power for large loads.

The two fastest "time to power" strategies are:

1. Repurpose the existing grid access possessed by ex-Bitcoin companies and
2. Quick-to-deploy, scalable power generation, which can offer 1-3 years of time-to-power advantage vs. the grid in many cases.

Exhibit 7: We Project a 38GW Power Shortfall in 2026-2028.



Source: Morgan Stanley Research

Exhibit 8: Even When Including Time-to-Power Solutions, We Still Project a 1-11GW Shortfall

Total US Data Center Power Demand (GW), 2026-28	**68**
Less Data Centers Under Construction	(15)
Less Utility Grid Access	(15)
Power Shortfall Before New "Time to Power" Solutions	**38**

	Low End	Midpoint	High End	Probability of Success
Solution #1: Nat Gas Turbines	15	18	20	90%
Solution #2: Bloom Energy Fuel Cells	5	7	8	90%
Solution #3: Site DC at Op'l Nuclear Plant	-	3	5	75%
Solution #4: Convert Bitcoin Sites	10	14	19	90%
Probability-Weighted "Time to Power" Solutions thru 2028	27	37	46	
Net Shortfall Through 2028	**(11)**	**(1)**	**None**	
Shortfall as a % of 2026-28 US DC Deployments	*16%*	*2%*	*0%*	

We believe the most likely outcome

Source: Morgan Stanley Research

The Value of Time to AI Players

The math relating to "time to power" contracts entered into between power developers and data center developers/hyperscalers is not fully understood by the market, in terms of ultimate impacts to shareholders. For example, the powerful magnifying effects of long-term contracts with investment grade hyperscalers is not fully appreciated. The analogy here would be to renewables contracts with investment grade tech sector customers. It would not be unusual to see a 15- to 25-year Power Purchase Agreement (PPA) signed for solar power to have an unlevered Free Cash Flow (FCF) yield of 9%. Such a project would typically cost ~$2/watt, and these projects can often be levered at 80-90% at the project level, resulting in levered equity returns in the mid- to high-teens. This, in our view, is an attractive ROE, given the low-risk, long-term nature of the contracted cash flows.

What is so striking to us is how attractive the "powered shell" data center contracts are, and how misunderstood this value creation is. For example, several Bitcoin companies have recently signed agreements in which they will build data center "powered shells" at $10-12/watt, with 15- to 25-year, fixed price contracts (with escalators), at unlevered FCF yields of 15-19%. These contracts will generate extremely attractive returns, in our view, and the Bitcoin stocks do not reflect this value creation. The following table walks through the math:

Exhibit 9: Value Creation Comparison: Renewable PPAs and Data Center Powered Shell Leases

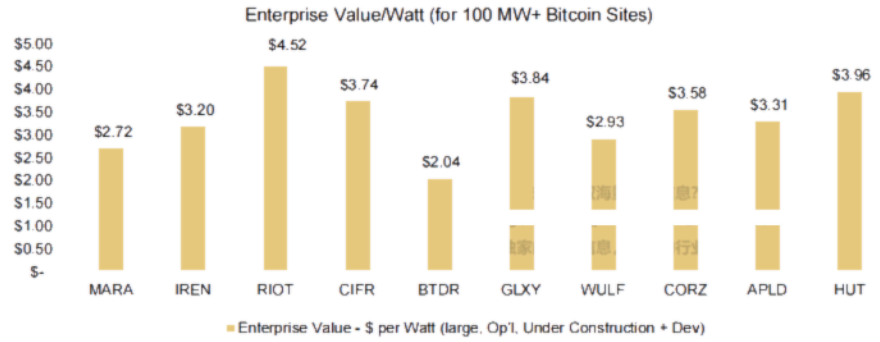
	Renewable PPAs	Data Center "Powered Shell" Leases
Cost/watt to build	$1.50-2/watt	$10-12/watt, with a few brownfield projects at a lower cost
Unlevered FCF yield on capex	Varies, but 10% would be a typical level	Typically 15-19%
Leverage potential (typically project level)	80%+	80%+
At 15x EV/EBITDA, value creation potential	$1.00/watt on a project costing $2/watt (math is: ($2 * .1 * 15x) - $2 cost) - this is a 50% net value creation on aggregate investment, or a >200% value creation on the equity portion of the investment, assuming 80% leverage	For a $10/watt project cost and a 17% unlevered ield, net value creation would be $15.50/watt (math is: ($10 * .17 * 15x) - $10 cost) - this is a 155% net value creation on aggregate investment, or a >700% value creation on the equity portion of the investment, assuming 80% leverage
Net value creation per gigawatt ($m)	$1,000	$15,500
Net value creation as a % of equity invested	250%	775%

Source: Morgan Stanley Research

What is striking about this comparison: The value creation per gigawatt is ~15x higher for the powered shell relative to the renewable project (with much greater capex required for the powered shell), and the value creation as a percentage of the equity investment is also much higher for the power shell.

Investment implications: We believe the Bitcoin stocks whose management teams are building out data center capacity for hyperscalers are not properly pricing in the value creation benefit — the following chart shows the EV/watt multiple for the Bitcoin stocks (keeping in mind the example transaction in the prior exhibit generated $15.50/watt in net value):

Exhibit 10: EV/Watt for the Powered Shell Provider Cohort from $2.04-$4.52 per watt. We Remain Bullish on the Conviction of These Names Trading at $15 EV/Watt.



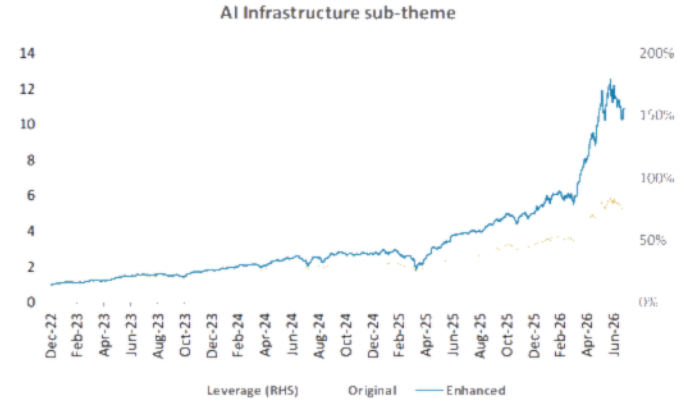
Source: Morgan Stanley Research

Taking Advantage of the Mathematics of Volatility

As we laid out in our 10 Thematic Predictions report, one practical way thematic conviction can translate into alpha is through disciplined buying during periods of theme-specific volatility. Over the past few years, we have seen multiple episodes in which certain thematic categories underperform — driven either by broad market dynamics or by idiosyncratic concerns tied to a given theme. A particularly sharp example was the sell-off in Powering AI stocks following the "DeepSeek moment" in 2025, though there have also been many smaller and shorter-lived drawdowns. From a thematic investing perspective, such dislocations can create attractive entry points for investors with conviction in the underlying structural trend. As compute demand continues to grow and will likely exceed supply, AI infrastructure and power supply are among the most critical constraining factors for AI adoption.

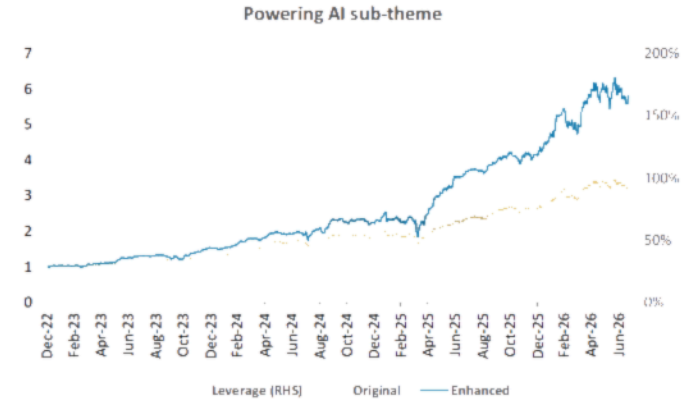
That said, we have seen significant periods of volatility around our Thematic stock calls in 2026, and this can cause corrections that are not driven by underlying fundamentals. To illustrate, we developed a "buy-the-dip" overlay on the AI Infrastructure (Bloomberg: MSRSTAII) and Powering AI (Bloomberg: MSRSTPAI) subthemes from 2023 onward. The back-test assumed an initial 100% capital allocation at end-2022. We increase exposure to 150% when the thematic portfolio declines by 5%, and to 200% when it declines by 15%. Each increase in leverage is held for three months (63 trading days), after which exposure reverts to 100%, unless a new signal is triggered earlier. The cumulative performance of the strategy versus the underlying thematic portfolio can be found here:

Exhibit 11: Buy-the-dip strategy cumulative performance since 2023 - AI Infrastructure sub-theme.



Source Bloomberg (ticker: VSRSTAII), Morgan Stanley Research. Data as of July 23rd, 2026

Exhibit 12: Buy-the-dip strategy cumulative performance since 2023 - Powering AI sub-theme.

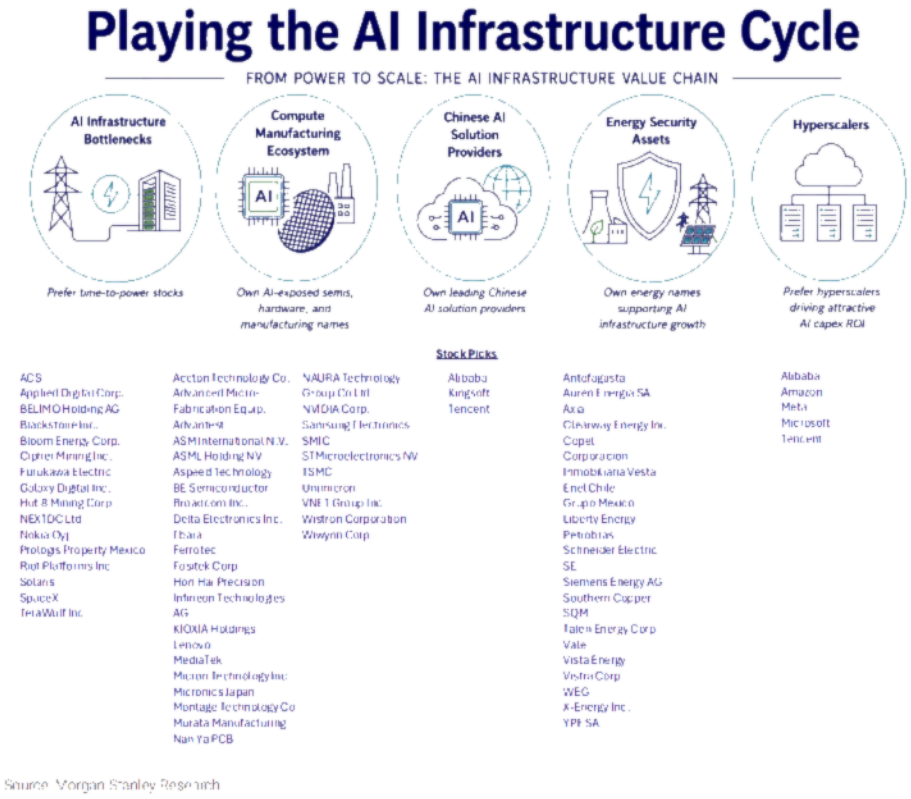


Source Bloomberg (ticker: VSRSTPAI), Morgan Stanley Research. Data as of July 23rd, 2026

Ways to Play

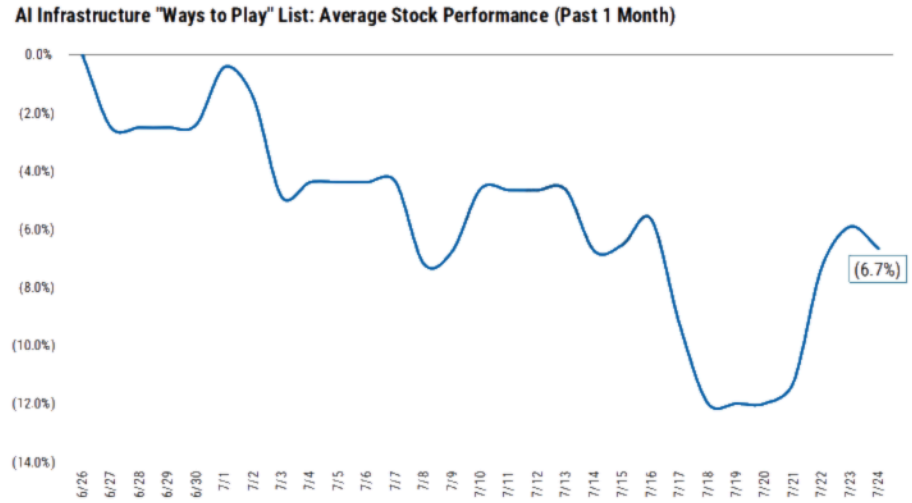
Here we present our ways to play the AI Infrastructure dip. Our global teams ran stocks screens to identify ways to play the AI Infrastructure theme, particularly focusing on names with positive risk-reward skews after the recent pullback. These names range across AI / Infrastructure Bottlenecks (including financers), Compute Manufacturing Ecosystem, Chinese AI Solution Providers, Energy Security Assets, and Hyperscalers. In aggregate, these names have underperformed world indices by 7 pts over the past month despite strong secular tailwinds, presenting a buying opportunity.

Exhibit 13: We prefer AI infrastructure bottleneck players, computing manufacturers, Chinese AI solution providers, energy security assets, and hyperscalers to play this theme.



Source: Morgan Stanley Research

Exhibit 14: On average, these names have been underperformers over the past month, despite strong fundamentals and AI infrastructure exposure.



Source: Morgan Stanley Research, Factset

North America

Exhibit 15: North America Beneficiaries

Ticker	Company Name	GICS Industry	Market Cap USD ($M)	Upside to PT	Bull/Bear Skew	Price (USD)
NVDA.O	NVIDIA Corp.	Semiconductors & Semiconductor Equipment	5005528	38%	2.5	208.76
MSFT.O	Microsoft	Software	2835434	57%	3.1	381.58
AMZN.O	Amazon	Broadline Retail	2496830	42%	7.6	249.99
AVGO.O	Broadcom Inc.	Semiconductors & Semiconductor Equipment	1817015	28%	2.9	392.47
META.O	Meta	Interactive Media & Services	1307064	30%	2.8	645.85
MU.O	Micron Technology Inc.	Semiconductors & Semiconductor Equipment	1040115	21%	2.1	990.21
SPCX.O	SpaceX	Diversified Telecommunication Services	871241	154%	11.1	118.24
BX.N	Blackstone Inc.	Capital Markets	97838	48%	3.1	124.50
VST.N	Vistra Corp.	Independent Power and Renewable Electricity Producers	55089	23%	5.6	168.98
BE.N	Bloom Energy Corp.	Electrical Equipment	52591	43%	3.0	217.30
TLN.O	Talen Energy Corp.	Independent Power and Renewable Electricity Producers	16338	34%	3.1	379.27
HUT.O	Hut 8 Mining Corp.	Software	12384	124%	16.0	117.67
CIFR.O	Cipher Mining Inc.	Software	9469	82%	2.9	25.82
WULF.O	TeraWulf Inc.	Software	9148	256%	1719.0	20.05
RIOT.O	Riot Platforms Inc.	Software	8520	51%	2.3	23.86
APLD.O	Applied Digital Corp.	IT Services	7770	22%	2.3	29.89
GLXY.O	Galaxy Digital Inc.	Capital Markets	4343	52%	5.1	24.32
XFR.O	X-Energy Inc.	Electrical Equipment	4186	158%	12.8	15.92
CWEN.N	Clearway Energy Inc.	Independent Power and Renewable Electricity Producers	3888	79%	3.3	33.57
SEI.N	Solaris Energy Infrastructure	Energy Equipment & Services	3156	49%	2.3	60.38
LBRT.N	Liberty Energy	Energy Equipment & Services	2833	84%	2.7	19.62

Source: Morgan Stanley Research, Factset. Data as of 7/24/26.

Latin America

Exhibit 16: Latin America Beneficiaries

Ticker	Company Name	Country	Sector	Rating	Mkt Cap. (US$ bn)	PT (US$, unless noted)	Last price	Upside	2027 Fwd P/E Estimate
SCCO.N	Southern Copper Corp.	Peru	Copper mining	Underweight	149.8	146.0	176.3	-17.6%	27.5
GMEXICOB.MX	Grupo Mexico S.A.B. de C.V.	Mexico	Copper mining	Equal-Weight	52.9	11.4	208.3	+94.5%	25.4
ANTO.L	Antofagasta	United Kingdom	Copper mining	Underweight	47.7	47.5	3,831.0	+98.7%	17.4
SQM.N	Sociedad Quimica y Minera de Chile S.A.	Chile	Lithium	++	9.8	++	68.7	na	na
VALE.N	Vale	Brazil	Mining	Equal-Weight	63.1	16.5	14.8	31.8%	6.5
AKAST.SA	AKA Energia	Brazil	Electricity generation	Overweight	22.7	10.4	50.3	-79.3%	7.9
ENIC.N	Enel Chile	Chile	Electricity generation	Overweight	6.0	4.3	4.3	-0.2%	11.1
CPLE3.SA	Companhia Paranaense de Energia	Brazil	Electricity generation	Overweight	8.5	3.0	14.5	+79.7%	10.5
PBR.N	Petrobras	Brazil	Energy	Overweight	89.8	26.5	18.6	51.8%	5.5
VIST.N	Vista Energy SAB de CV	Mexico	Energy	Overweight	6.5	94.0	69.5	35.3%	5.7
YPF.N	YPF SA	Argentina	Energy	Overweight	20.2	59.0	51.4	+4.8%	7.7
FIBRAPL14.MX	Prologis Property Mexico SA de CV	Mexico	Industrial Real Estate/Hard Assets	Overweight	7.2	4.9	75.8	-93.5%	12.1
VESTA.MX	Corporacion Inmobiliaria Vesta SAB de C	Mexico	Industrial Real Estate/Hard Assets	Overweight	3.0	4.5	58.5	+52.4%	30.5
WIEGF3.SA	WEG	Brazil	Industrial goods	Equal-Weight	38.1	10.5	46.0	+75.9%	28.2
AUAS3.SA	Auren Energia SA	Brazil	Electricity generation	Equal-Weight	2.3	2.8	11.2	+75.4%	n/m

Morgan Stanley Research, Factset. Data as of 7/24/2026

++ Rating and estimates for this company have been removed from consideration in this report because, under applicable law and/or Morgan Stanley policy, Morgan Stanley may be precluded from issuing such information with respect to this company at this time.

Europe

Exhibit 17: European AI List

#	Ticker	RIC	Company	MS Rating	Industry Classification	Top AI Enabler Materiality	Up/down side to PT	Price (USD)
1	ENR GR	ENR1n.DE	Siemens Energy AG	OW	Real Assets / AI Capex	Significant	33%	171.09
2	ASML NA	ASML.AS	ASML Holding NV	OW	Compute	Core to Thesis	22%	1795.70
3	ASM NA	ASMI.AS	ASM International N.V.	OW	Compute	Core to Thesis	26%	993.37
4	IFX GR	IFXGn.DE	Infineon Technologies AG	OW	Compute	Core to Thesis	39%	74.47
5	BESI NA	BESI.AS	BE Semiconductor Industries N.V.	OW	Compute	Core to Thesis	13%	261.42
6	STMPA FP	STMPA.PA	STMicroelectronics NV	OW	Compute	Core to Thesis	36%	54.52
7	NOKIA FH	NOKIA.HE	Nokia Oyj	OW	Compute/Goods	Moderate	61%	9.89
8	ACS SM	ACS.MC	ACS, Actividades de Construccion y Servicios SA	OW	Real Assets / AI Capex	Core to Thesis	36%	130.42
9	SU FP	SCHN.PA	Schneider Electric SE	OW	Real Assets / AI Capex	Significant	19%	301.79
10	BEAN SW	BEAN.S	BELIMO Holding AG	OW	Real Assets / AI Capex	Core to Thesis	37%	1010.92

Source: LSEG Data Analytics, Morgan Stanley Research. Data as of 7/23/2026

Asia

Exhibit 18: Preferred Asia Pacific plays on a recovery in AI Infrastructure sentiment

Ticker	Company Name	Analyst	Top Pick	Market	Mkt Cap	Share Price	Price Target	% Upside	MS Rating	Price to Earnings		Return On Equity	
					US$bn	Close, LCY	LCY			FY1	FY2	FY1	FY2
2330.TW	TSMC	Chan, Charlie		Taiwan	1,923.6	2405	2988	24%	OW	21.7	16.6	53.0%	48.4%
005930.KS	Samsung Electronics	Kim, Shawn	Top	Korea	1,129.2	270000	381000	41%	OW	NA	NA	85.5%	65.9%
0700.HK	Tencent Holdings Ltd.	Yu, Gary		China	511.0	445	650	46%	OW	14.9	13.0	20.8%	19.9%
BABA.N	Alibaba Group Holding	Yu, Gary	Top	China	277.9	114	180	58%	OW	25.5	17.0	7.1%	9.9%
285A.T	KIOXIA Holdings	Yoshikawa, Kazuo	Top	Japan	215.9	61880	110000	78%	OW	7.0	6.3	342.8%	103.2%
2454.TW	MediaTek	Chan, Charlie	Top	Taiwan	190.9	3875	5588	44%	OW	55.3	26.5	27.3%	48.5%
2308.TW	Delta Electronics Inc.	Yang, Derrick	Top	Taiwan	150.9	1880	2700	44%	OW	45.3	28.9	33.1%	41.9%
6857.T	Advantest	Tamura, Suzune	Top	Japan	132.8	30800	36000	17%	OW	41.0	27.5	68.5%	63.5%
0981.HK	SMIC	Chan, Charlie		China	114.4	70.75	85.00	20%	OW	43.5	30.0	7.7%	10.4%
2317.TW	Hon Hai Precision	Kao, Howard		Taiwan	109.0	258	310	20%	OW	15.7	12.9	11.7%	13.3%
6981.T	Murata Manufacturing	Sato, Shoji	Top	Japan	100.1	8538	12500	46%	OW	46.0	31.4	12.3%	17.3%
002371.SZ	NAURA Technology Group Co Ltd	Chan, Charlie		China	81.8	740	818	11%	OW	58.2	39.8	21.0%	26.3%
3037.TW	Unimicron	Kao, Howard		Taiwan	44.5	896	1465	64%	OW	60.3	26.8	21.4%	40.7%
6809.HK	Montage Technology Co Ltd	Yen, Daniel		China	40.8	312	432	39%	OW	NA	NA	27.9%	32.6%
2345.TW	Accton Technology Corporation	Yang, Derrick	Top	Taiwan	39.6	2325	3800	63%	OW	27.1	20.1	83.2%	68.6%
0992.HK	Lenovo	Kao, Howard		China	38.1	24.62	30.00	22%	OW	13.8	11.0	35.7%	35.6%
688012.SS	Advanced Micro-Fabrication Equip	Chan, Charlie		China	36.3	375	550	47%	OW	61.3	44.9	16.9%	19.6%
6669.TW	Wiwynn Corp	Kao, Howard	Top	Taiwan	31.6	5470	7500	37%	OW	15.3	11.7	53.3%	52.8%
8046.TW	Nan Ya PCB	Kao, Howard		Taiwan	26.1	1230	1550	26%	OW	78.6	37.2	21.9%	38.9%
5274.TWO	Aspeed Technology	Yen, Daniel		Taiwan	19.9	15540	20888	34%	OW	82.8	57.0	108.5%	99.6%
3231.TW	Wistron Corporation	Kao, Howard		Taiwan	16.1	174	210	21%	OW	12.4	9.6	16.7%	19.5%
6361.T	Ebara	Tamura, Suzune		Japan	15.7	5752	7100	23%	OW	26.3	20.4	19.7%	22.9%
5801.T	Furukawa Electric	Shirakawa, Yu		Japan	15.3	3513	6600	88%	OW	29.5	18.4	20.0%	27.6%
NXT.AX	NEXTDC Ltd	McLeod, Andrew		Australia	7.4	13.92	18.00	29%	OW	NM	NM	-3.6%	-8.7%
6871.T	Micronics Japan	Hasegawa, Yoshihito		Japan	3.8	16370	19000	16%	OW	29.7	22.0	31.0%	34.5%
KC.O	Kingsoft Cloud Holdings	Liu, Yang		China	3.0	9.69	15.00	55%	OW	NM	521.1	-9.4%	-3.4%
6805.TW	Fositek Corp	Yang, Derrick		Taiwan	3.0	1455	2426	67%	OW	21.5	15.8	64.4%	56.8%
6890.T	Ferrotec	Hasegawa, Yoshihito		Japan	2.5	8940	13000	45%	OW	17.2	13.5	9.4%	11.3%
VNET.O	VNET Group Inc	Tang, Tom	Top	China	2.2	7.77	16.00	106%	OW	53.9	23.4	4.3%	7.2%

Source: FactSet, Morgan Stanley Research. Data as of 23 July 2026

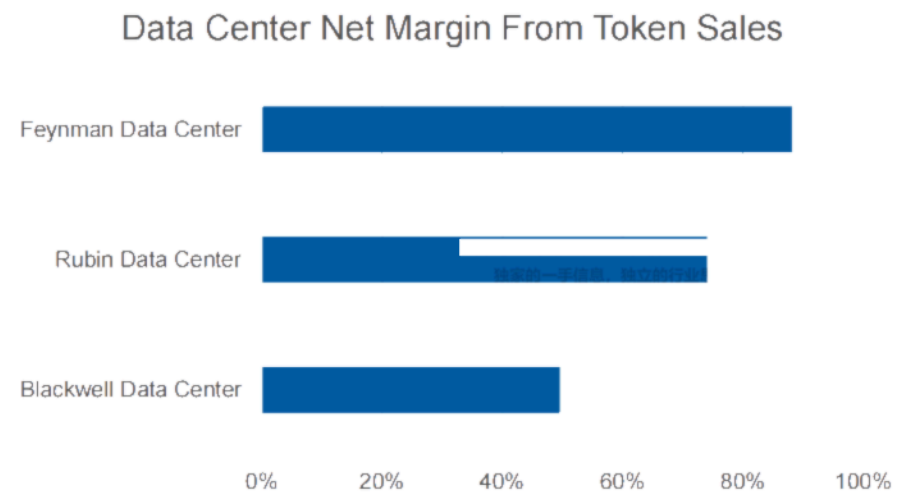
Our Analytical Tools for Assessing AI Debates

Our Tokenomics Model

Our Intelligence Factory reveals favorable token economics for both Enterprise customers and for AI players selling token usage via their LLMs. Click here to request a copy of our Intelligence Factory model. For enterprise adopters of AI, the economics of operating a data center are compelling. In a study of a wide range of Enterprise AI use cases, utilizing AI to complete a task in the workplace would save an Enterprise customer $55 in costs. While it is not possible to precisely know how many tokens will be used to achieve this outcome (and token usage varies significantly across enterprise use cases), we can ground ourselves in some helpful math.

One analysis suggests that a single AI Agent may consume ~15x the tokens relative to a query-based usage of a LLM. Further, a typical query based usage of a LLM consumes ~5,000 tokens. If an Enterprise customer used 5-7 AI Agents to execute an economic task, the total token usage would be 375,000-525,000 tokens. Token pricing can be found here —it varies widely among models and between input and output tokens. At an assumed average token price of $3-10 (in most cases, input tokens outnumber output tokens), the cost to execute the economic task would be $2-5, only a fraction of the $55 benefit. From the perspective of an LLM developer, based on token pricing that is publicly available, the Blackwell data center would generate margins of ~50%, factoring in all data center costs (but not the cost to develop or host the LLM). Interestingly, when we run the same math for Rubin and Feynman data centers, the margins would rise to ~75% and 90%, respectively. This suggests that, in the future, token pricing could fall as the share of global compute is comprised of increasingly capable chips.

Exhibit 19: Our Intelligence Factory Model Projects 58-90% Margins for Token Sales, Depending on GPU

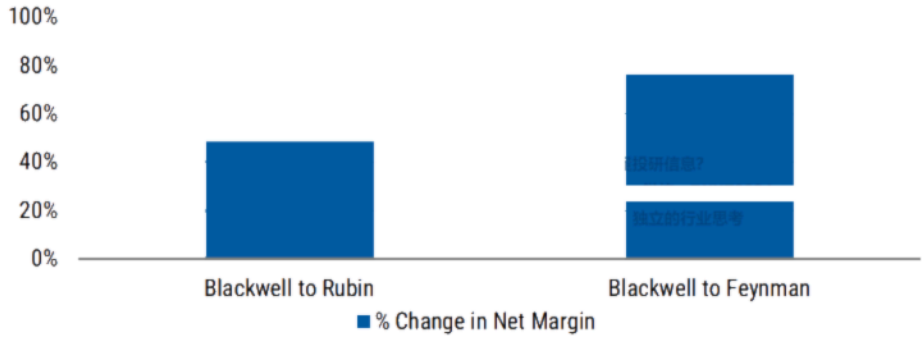


Source: Morgan Stanley Research

Framing our token economics math differently, our model shows that, as we move from the Blackwell to the Feynman, hyperscalers could price tokens ~75% lower and achieve

the same margins on a leased data center:

Exhibit 20: Subsequent generations of NVIDIA GPUs could lower token pricing by ~75%



Source: Morgan Stanley Research

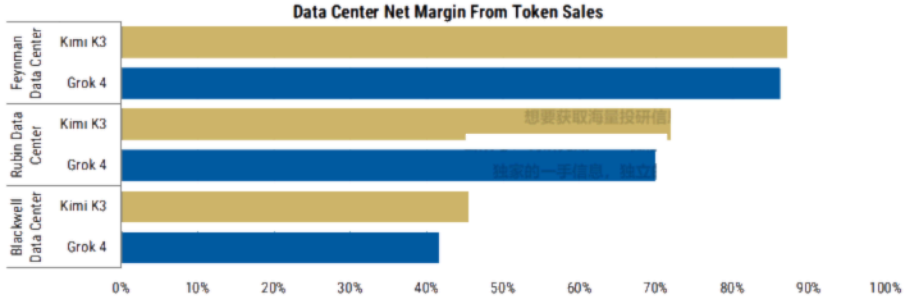
Further, we update our model to include Kimi K3; we find that, given the estimated $3/million input tokens and $15/million output tokens, margins in the same data center utilizing Blackwells are 45%. While K3 performs well on benchmarks, we iterate that it is another strong frontier model and is priced accordingly. Kimi K3 may be an important inflection point for AI. While increased competition on the production of high-performing frontier models may be negative for the frontier labs of the US, advancements in AI could be a net positive for companies in general.

A notable milestone for development of AI would be an open-source model that:

1. Operates at or above benchmarks,
2. Shows token efficiency gains over existing frontier models, and
3. Is radically more effective on a dollar-per-token basis.

While not the revolutionary moment described, we see improved margins for K3 over our prior frontier model, Grok 4. The model delivered stronger benchmarking but maintained the same pricing structure.

Exhibit 21: Our Projections Show Improved Margins (+1-3%) For A Data Center Utilizing K3



Source: Kimi, OpenRouter, Morgan Stanley Research

US Power Shortfall Analysis

As discussed earlier, our Intelligence Factory model evaluates the current US power shortfall. We utilize sales numbers and projections from our semis team to build chip-by-chip server level MW demands and work our ways towards a year-over-year evaluation of

the total demand of power. This is summarized in our Power Summary sheet, which includes our current graph on the US shortfall on energy.

The US shortfall analysis in Exhibit 9 provides the information on our time-to-power solutions; we provide our expectations on 1) nat gas turbines, 2) Bloom Energy fuel cells, 3) nuclear power, and 4) BTC-to-HPC site conversions in a concise chart.

Data Center Pushback Analysis

Our recent report, Power Struggle: AI, Data Centers & Local Backlash, discusses public sentiment towards data centers and the potential for moratoriums. This is timely: a growing number of localities have passed or proposed guardrails and restrictions on data center projects. These include infrastructure cost-sharing requirements, ratepayer protection initiatives, environmental disclosure rules and zoning changes. Restrictions have mostly been in the form of temporary moratoriums, impacting future projects or those lacking vested status. Further, although the explicit public opposition has tended to concentrate more in blue states and among Democratic-leaning voters, we're seeing state-wide data center moratorium proposals emerge roughly evenly in Democratic trifectas, Republican trifectas, and divided governments.

Importantly, most efforts are aimed at slowing — not reversing — the build-out, as many of these policy initiatives are temporary.

Exhibit 22: Proposed State-Level AI Data Center Moratorium Bills & Political Breakdown

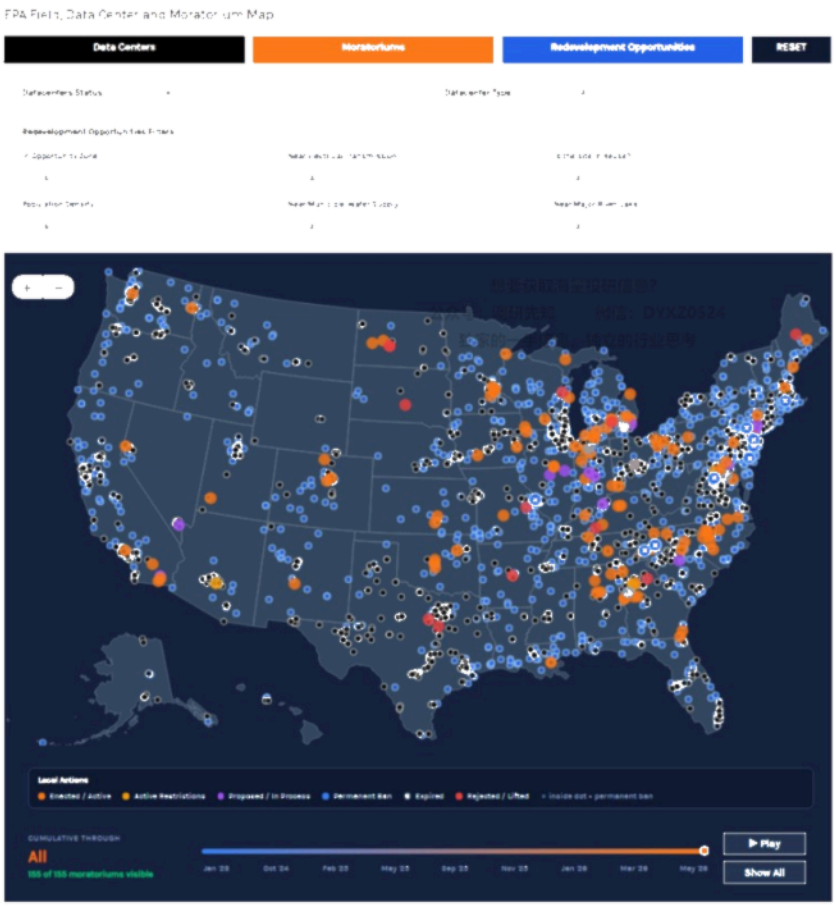
State	Bill(s)	Status	Description	Governor	Senate	House	Alignment
Delaware	SB 353	Introduced	Moratorium through Jan. 31, 2027 on applications and permits for data centers able to use 100 MW or more.	Matt Meyer (D)	D	D	Aligned
Georgia	HB 1059	Introduced	Forbids local governments from permitting data centers until Dec. 31, 2028.	Brian Kemp (R)	R	R	Aligned
Maine	LD 307 / HP 207	Vetoed	Creates Maine Data Center Coordination Council and temporarily limits certain data centers over 20 MW until Nov. 2027.	Janet Mills (D)	D	D	Aligned
Maryland	HB 120	Failed	Prohibits construction or permitting of new data centers under stated co-location/generation conditions.	Wes Moore (D)	D	D	Aligned
Michigan	HB 5594 / HB 5595	Introduced	Local governments and state agencies may not permit/approve certain new data centers until Apr. 1, 2027.	Gretchen Whitmer (D)	D	R	Divided
Minnesota	HF 4888 / SF 4298	Failed	Bans new data-center permits until a Public Utilities Commission report, no earlier than July 2027.	Tim Walz (D)	D	NP	Divided
New Hampshire	HB 1265	Failed	Establishes a one-year moratorium on data-center construction.	Kelly Ayotte (R)	R	R	Aligned
New York	**A11560 / S10642**	**Enacted**	**Places a one-year moratorium on large data-center permits and adds service-classification, energy-efficiency and host-community provisions.**	**Kathy Hochul (D)**	**D**	**D**	**Aligned**
Oklahoma	SB 1488	Failed	Establishes a moratorium on data-center construction until Nov. 2029 and requires a study.	Kevin Stitt (R)	R	R	Aligned
Pennsylvania	SB 1359 / HB 2533	Introduced	Creates a statewide three-year moratorium on hyperscale data centers and impact studies, or authorizes municipal moratoriums.	Josh Shapiro (D)	R	D	Divided
South Carolina	H 5526	Introduced	Blocks state/local acceptance or action on data-center applications until a comprehensive oversight framework is enacted.	Henry McMaster (R)	R	R	Aligned
South Dakota	SB 232	Failed	Imposes a one-year moratorium on construction or expansion of hyperscale data centers.	Larry Rhoden (R)	R	R	Aligned
Vermont	S 205	Introduced	Creates a moratorium on AI data centers until 2030 and requires a report on impacts.	Phil Scott (R)	D	D	Divided
Virginia	HB 1515	Continued	Prohibits final local approval for a new data center until interconnection requests are fulfilled or July 1, 2028.	Abigail Spanberger (D)	D	D	Aligned
Wisconsin	SB 1061 / AB 1099	Failed	Prohibits data-center operation until 14 listed conditions are satisfied.	Tony Evers (D)	R	R	Divided

Source NCSL, VulliState, Morgan Stanley Research

Tracking Data Center Moratoriums

We track opposition efforts at the local, state, and federal levels in the text below and encourage readers to click through our interactive map below to get more detail on the opposition efforts.

Exhibit 23 Data Center Moratorium Interactive Map



Source: AlphaWise Web Intelligence, Morgan Stanley Research, 451 Research, Interconnected Capital, EPA.

Enterprise Adoption Surveys

2Q26 CIO Survey

We make a point to ground our expectations of AI in real-world applications; Our 2026 CIO survey shows that AI/ML remained CIOs' top spending priority in 2Q26 and reached another survey high in net prioritization. AI/ML ranked as the top priority for 18.0% of respondents, up modestly from 17.7% in 1Q26. That said, net defensibility declined to +4% from +10% QoQ, suggesting that CIOs are applying greater scrutiny to which AI initiatives they would protect in a downturn.

CIO objectives remain centered on productivity. In 2Q26, 30% of CIOs cited improving broader internal employee productivity through GenAI and LLMs as their primary objective, up from 29% in 2Q25. Customer-facing applications designed to drive incremental revenue ranked second at 21%, followed by worker labor savings and productivity improvement at 17% and applications designed to improve customer satisfaction at 12%.

AI deployment timelines continue to accelerate. In 1Q26, 43% of CIOs expected their first AI/LLM projects to reach production by YE26, up from 17% in 4Q25, while 11% expected deployment in 2027 and 20% expected deployment after 2027. By 2Q26, ~76% of CIOs

had already placed their first AI/LLM project into production or expected to do so within the next nine months.

Future of Work Surveys

Earlier this year, we conducted two surveys to assess the implications of AI on employment across various industry groups we had determined to be most impacted by AI adoption. Our surveys indicate that AI adoption is reshaping workforces across industries, with broadly consistent dynamics observed in both waves. Survey respondents in the five new sectors ('Wave 2') said that in the past year, AI adoption had led to 12% of jobs to be eliminated and a further 15% to not be backfilled, partially offset by 22% new hires. This yields a 5% net job loss for companies included in Wave 2, broadly consistent with the 4% net job loss reported by companies in Wave 1 sectors. However, Wave 2 companies experienced a higher combined share of roles eliminated or not backfilled at 27% versus 22% in Wave 1, balanced by a higher level of new hires, indicating that whilst net losses may be broadly consistent, labor disruption is higher in Wave 2 sectors.

Consumer Survey

AI use has also become broad-based among consumers, reflected in our recent June consumer survey. Approximately half of respondents reported using AI at least weekly for both personal and work-related purposes. Looking ahead over the next 12 months, 42% expect to increase their personal use of AI tools, while nearly half of employed adults expect to increase their use of AI at work.

Global AI Stock Mapping

Click here to request the latest AI Mapping Database, which maps Morgan Stanley's global coverage across three dimensions:

1. AI exposure, including AI Adopter, AI Enabler, and both AI Enabler and Adopter;
2. The materiality of AI to the investment thesis, including Core to Thesis, Significant, and Moderate; and
3. Company pricing power.

Our analysis shows that stocks with greater AI materiality and stronger pricing power have outperformed. We have completed five waves of the mapping, with each survey conducted approximately six months apart. The database therefore shows how analysts' views of AI's impact on their coverage have evolved across the five waves.

Future of Work Analysis

Our Future of Work model from last summer estimated that agentic AI can lead to economic value creation for the S&P 500 totaling ~$490 billion annually (net of indicative costs to implement agentic AI solutions, which we estimated to be 5% of the customer benefit). We acknowledged tasks can be augmented or automated with AI, and considered solely the benefits associated with the latter for estimated results. The long-term AI "automation-driven economic value creation" as a percentage of 2026e pretax income was highest for Consumer Staples Distribution/Retail, Real Estate Management & Development, Transportation, and Health Care Equipment & Services.

The Potential for "Two Worlds" of AI Policy and Adoption

Ariana Salvatore

Multipolarity in motion: we see an increasing probability of both US and Chinese AI policy intervention, which could lead to a greater bifurcation of the global market. We have been of the view for years that the US and China are gradually de-risking, rather than de-coupling, vis-à-vis policy tools (like export controls, tariffs, and domestic incentives for manufacturing) employed over the past ~8-10 years. That de-risking is most stark in the higher-tech segments of the economy, where hawkishness on China appears to be indicative of the overall direction of travel. But it's not just high tech: across a number of policy vectors, we expect further divergence between the two economies through the implementation of policy tools like export controls/restrictions on commodities like memory and a higher effective tariff rate on China vs our other trading partners, even for non-strategic goods. Similarly, we think this backdrop – and the geopolitical realities that accompany the progression of AI into the physical layer – is a key reason that the US is unlikely to enact a nation-wide data center ban or moratorium in the near term.

This is the "Two Worlds" prediction we included in January as part of our 10 Thematic predictions for 2026. Given the competition between the US and China, we see a larger global TAM for Chinese AI solution providers than the market appreciates, a dynamic explored in depth in our China 2.0 report. As the team highlighted in that report, "China's AI is entering a new phase – one defined less by catching up in capability and more by capturing value. The narrative has shifted decisively from training to inference, from technology to application, and from potential to real earnings. Rather than competing purely at the frontier, China is optimizing for speed, cost efficiency, and system-level integration, enabling AI to diffuse rapidly across the real economy."

This competitive backdrop underscores our view for post-midterms AI policy: we think AI policy initiatives seem likely to take shape regardless of the election outcome, reflecting bipartisan alignment around AI as a strategic priority, but with emerging regulatory and societal constraints. This shift toward AI as a strategic, capital-intensive 'arms race' with policy (across permitting, energy, and industrial support) as an enabler should continue to support the AI investment cycle and associated infrastructure demand, albeit with a slightly different approach depending on congressional composition. We see election outcomes as more likely to shape the pace and geographic distribution of deployment across friction points like energy availability, permitting, and labor constraints. At the same time, public pushback is emerging as a binding constraint, particularly around data center buildout (see here). Regulation, meanwhile, is likely to remain incremental and fragmented: core areas like AI safety, data governance, and infrastructure siting may see targeted action, but broader efforts appear unlikely near-term.

US Restrictions: Form, Feasibility, and Timing

In response to this reality, we see a rising probability of restrictions on Chinese model usage, and we could see similar restrictions from the Chinese government.

What form could these or similar restrictions take and how could they work in

practice? We think the executive branch is likely to lead on this front – given Congressional gridlock and competing agenda items (the NDAA, CR, and potential third reconciliation bill among others before the August recess). Without any timing friction at all, BIS under the Commerce Department could expand its existing Entity List, under which US firms must require a licence to transact with a designated Chinese AI developer. In practice, this could force cloud providers to remove officially hosted models, enterprise-software companies to terminate licensing agreements, and model marketplaces to de-list products associated with the designated laboratory. Commerce could also use the Information and Communications Technology and Services supply-chain regime (the "connected software applications" review) to review and/or prohibit any transaction involving "information and communications technology and services" (ICTS) and block it if it's tied to a certain country or poses an undue security risk.

Other options include Treasury sanctions under OFAC (which would likely require a catalyst or explicit justification), Justice Department restrictions on transactions exposing bulk personal data, federal procurement exclusions and Federal Trade Commission (FTC) enforcement requiring the deletion of models trained on unlawfully obtained data. Precedent for these options can be found here, here, here, here, and here.

Executive-led restrictions are one option, but they are inherently temporary.

Although we see continuation of these policies from one administration to the next, Congress could also make the restrictions more durable – by incorporating AI and semiconductor provisions into the annual National Defense Authorization Act (NDAA), which is the must-pass annual defense bill. We've highlighted in recent research the potential to include the AI OVERWATCH Act and the Multilateral Alignment of Technology Controls on Hardware, or MATCH Act, in this bill (see here and here). The AI OVERWATCH Act focuses on the export of advanced AI chips to foreign adversaries: versions of the proposal would codify restrictions on the highest-performance chips and subject certain future export approvals to congressional review. The MATCH Act operates further upstream by tightening controls on semiconductor-manufacturing equipment and components and seeking greater alignment between US restrictions and those maintained by allied equipment-producing countries (like the Netherlands). Neither measure directly bans Americans from downloading or deploying a Chinese model, but they would restrict China's ability to acquire the chips and manufacturing machinery required to train future frontier systems, complementing rather than replacing Entity List designations, application restrictions or procurement bans aimed at models already available in the United States.

We see potential for at least some extent of these provisions to be included in the FY27 NDAA, which should pass by the end of this year, but note the uncertainty surrounding the legislative process. House-Senate negotiations on this front will be key.

Recent statements from administration officials and US AI companies suggest that policy is moving toward tighter restrictions, even if an outright ban remains uncertain. Frontier labs have accused Chinese developers of using distillation, fraudulent accounts and proxy infrastructure to extract capabilities from leading US models, framing the issue as both intellectual-property theft and a national-security risk. White House science adviser Michael Kratsios stated in April 2026 that Chinese entities were conducting deliberate, industrial-scale extraction campaigns and said the government would develop

mechanisms to punish offenders, while Treasury Secretary Scott Bessent stated in July that the United States had the ability to sanction developers using stolen US technology. Reporting on internal deliberations also indicated that officials had considered Entity List additions, security advisories, supply-chain rules and an executive order imposing security and liability conditions on US companies hosting Chinese models.

China Restrictions: Focused on Integration & Market Access

China could similarly restrict US AI models primarily through administrative regulation, cybersecurity controls and market access requirements. Under the 2023 Interim Measures for Generative Artificial Intelligence Services, any company offering a model, chatbot or API to the Chinese public must comply with requirements concerning training-data legality, personal-information protection, content controls, transparency, algorithm registration and, for services possessing "public-opinion attributes or social-mobilisation capacity," a security assessment. The rules expressly apply to API-based services and authorize the Cyberspace Administration of China to instruct relevant institutions to take "technical measures and other necessary measures" against noncompliant services supplied from outside China.

In practice, this could mean blocking a foreign model or API, preventing its distribution through Chinese app stores and cloud platforms, requiring a locally controlled deployment, or permitting access only after the service has been adapted to Chinese content, data-storage and security requirements. Implementation would principally be executive and regulator-led through the CAC, the Ministry of Industry and Information Technology, public-security bodies and other sectoral regulators within China's centralized party-state system.

China also has several broader national-security and industrial-policy mechanisms that could be applied to US AI companies or the infrastructure supporting their models. A foreign provider, cloud service or AI-enabled product could be subjected to a cybersecurity review on the ground that it creates supply-chain, data-security or critical-infrastructure risks, as evidenced by the 2023 review of Micron. The government could also employ government-procurement standards to require "safe and reliable" domestic technology, extending the localization approach already used to phase foreign processors, operating systems and database products out of government computers and servers. That syncs with the stated ambitions toward inidigenization that featured at the US-China summit in May of this year. Lastly, the Ministry of Commerce maintains a mirroring Unreliable Entity List, which has previously been used to prohibit listed American companies from China-related import and export activity and from making new investments. Hence, the Chinese approach likely reflects an ambition to restrict US models from government systems, state-owned enterprises, regulated sectors and approved distribution channels while promoting domestic substitutes.

Public evidence indicates that the current direction is toward regulatory localization and effective exclusion of unapproved foreign models, rather than an expressly articulated plan to ban all American AI technology. The clearest example is Apple Intelligence: Apple was unable to launch its global AI offering in mainland China on the same terms used elsewhere and instead pursued regulatory approval through partnerships integrating models and technology from Alibaba and Baidu. Reporting

characterized the local partnership as essential because foreign generative-AI services must pass Chinese reviews and satisfy domestic content- and data-governance requirements; China ultimately approved a version relying on Alibaba's Qwen rather than a wholly US-provided model stack. At the same time, the Interim Measures encourage indigenous innovation in models, chips, software, computing infrastructure and data resources, and specifically encourage the use of "safe and trustworthy" technologies. This combination of mandatory approval, local model integration, domestic data handling, procurement localization and support for technological self-reliance suggests that US models may be confined to tightly controlled or locally mediated deployments even without a formal blanket prohibition.

US Ways to Play Perspectives

US Internet

Brian Nowak

We view the hyperscalers as the key AI infrastructure beneficiaries this cycle. While there has been lower confidence in GenAI ROIC and/or concerns that **META**'s $380bn+ of capex ('27+'28) and ~$25bn of annual hyperscaler investment will not generate returns, we believe this misses the company's improving moats – including its 3.5bn daily active users and litany of GPU enabled improvements in the pipeline. We think META's engagement and monetization moats are deeper than ever and remain bullish on the neocloud optionality, which would allow the company to monetize excess capacity in a compute constrained world. Our price target of $775 (30% upside) reflects implied ~23X base case blended '27/'28 P/E.

AMZN continues to be an AI infrastructure beneficiary through the scaling of Trainium, accelerating demand for AWS compute, and the expansion of its multi-model offering strategy, which positions the company to capture enterprise AI spending across both model training and inference. Within AWS, the company has services built on top of the infrastructure (Bedrock) where the long tail of enterprises who have existing relationships with AWS can access frontier/open source LLMs. Our price target of $330 (42% upside) reflects ~25X an average of $13 of base '27/'28 EPS.

US Semis

Joeseph Moore

Nvidia (NVDA) remains our Top Pick in semis. While the stock has underperformed a sharp runup in the peer group, our conviction remains high. We understand that there is more leverage in other parts of the supply chain, but we now feel that this is the best value in the group. While the multiple is likely constrained by the large market cap and index weighting, our view is that the value disparity will revert over time. Our price target of $288 (39% upside) reflects ~22x our MW CY27 EPS of $13.08.

Broadcom (AVGO) has the second largest AI exposure (in absolute dollars) in our coverage and is poised to grow as hyperscale capex increases. While there has been an investor preference for growthier "bottleneck" stories across the AI semiconductor ecosystem and the debate around MediaTek versus Broadcom share on Google TPU is an overhang, we believe in the continued strength of the company's AI growth trajectory and that AVGO should remain the majority TPU supplier over time, with ~80% share. Our price target of $502 (31% upside) reflects 28x 2027e ModelWare Base EPS of $17.92.

We continue to like the memory names, as the longer-term concerns that the memory shortage will intensify in 2027 and again in 2028 are still as strong as ever. There isn't enough memory vs. AI requirements, and we just don't see that changing. Conversations with our purchasing contacts in the data center space show no signs of shortages abating, despite some concerns of a second derivative deceleration, higher capex, and de-speccing. We see prices up at least 25% on a like-for-like basis from 2Q to 3Q, above our estimates

and above 3rd-party estimates. As a result, we like **Micron (MU)** with a price target of $1200 (30% upside), reflecting 30x through-cycle earnings of $40.

US Software

Adam Wood & Josh Baer

We believe **Microsoft (MSFT)** (covered by Adam Wood) combines one of the strongest competitive moats in software with a highly attractive AI opportunity set, supported by its unparalleled enterprise distribution, broad product portfolio, and presence across both the application and infrastructure layers of the AI stack. Microsoft's position at the center of enterprise workflows, coupled with its ability to embed AI capabilities across products including Microsoft 365, Azure, and Security, leaves the company well positioned to benefit as AI adoption continues to scale. Importantly, despite its leadership position, we continue to view Microsoft's valuation as compelling relative to its growth durability and strategic AI positioning, keeping us firmly Overweight. Our price target of $600, representing 57% upside, reflects approximately 25x our base-case FY28e EPS estimate of $23.86.

US Humanoids/Robotics/Space

Adam Jonas

We believe **SpaceX** remains uniquely positioned across launch, connectivity, and AI, with a dominant competitive moat in space launch and satellite connectivity (>3/4 of global mass to orbit and maneuverable satellites in orbit) and meaningful long-term optionality from Grok, Cursor, and neocloud. At less than $120/share, we believe investors are paying a reasonable multiple for the core Space + Connectivity business while receiving limited-to-no value for the company's AI opportunity, despite multiple monetization paths to our estimated ~$100bn in 2028 AI revenues on 6 GW average compute. We see the current valuation as an attractive entry point, with our $300 price target implying substantial upside and supported by a 15-year DCF that values Space + Connectivity alone at approximately $136/share.

US Asset Managers

Michael Cyprys

Blackstone (BX) is one of the largest global investors in AI-related infrastructure, with a large footprint in data center and supporting energy assets and a best-in-class private markets franchise with significant brand power and scale. We see AI as a core driver of BX's fundraising, deployment and realizations, with **its data center platform reaching $185bn of total value in 2Q26, including facilities under construction, up more than 40% from $130bn at the start of 2026**. We view BX as best-placed amongst the alts to capitalize on the AI theme supported by mgmt's expectations to lease more than 3x as much capacity this year as in any prior year with scope for the platform to double over the next few years. **BX's newly launched data center REIT, BXDC, adds another leg to the opportunity here.** This $2b stabilized data center REIT provides a scalable vehicle to acquire operating assets as developers and hyperscalers recycle capital, giving BX attractive exposure to a TAM that mgmt believes could exceed $1tr. We see BXDC and a

number of direct investments and strategic partnership enhancing BX's access to capital while creating incremental fee streams and providing another monetization avenue for its development pipeline **supporting our view that BX offers differentiated exposure to AI infrastructure across fundraising, deployment, investment performance and realizations**. We see a long runway for growth that's not currently priced into shares, with shares trading at just 15x '27 PE vs the broader market at 19.5x. Our $184 price target reflects 22.5x Base Case 2027e EPS of $8.17 and offers ~48% upside.

US Data Center REITs

Cameron McVeigh

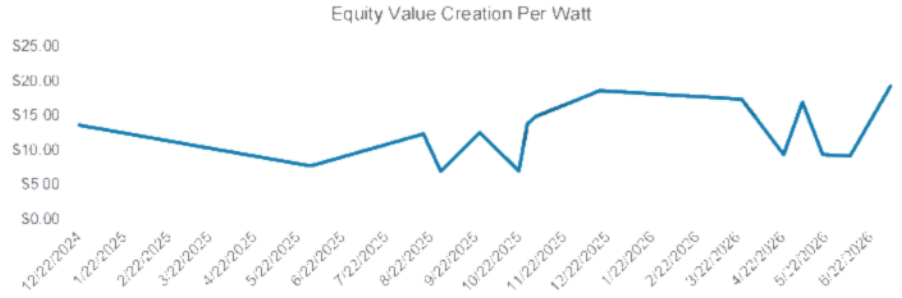
Data Center REITs (**EQIX, DLR and BXDC**) rent out space and power to Fortune 500 enterprises, cloud service providers, neoclouds and AI labs. We see this key infrastructure enabling the proliferation of AI workloads across the broader economy. OW EQIX, EW DLR and EW BXDC will continue to benefit from the inflection of AI inference related demand as these operators are located in Tier 1 markets and facilitate lower latency workloads. We expect accelerating demand to show in bookings and lease renewal rates, recently apparent in DLR 2Q26 results with 0-1MW record bookings and a $1.4bn backlog (see here). We continue to prefer EQIX with a scaled interconnection network, attractive pricing power and AI inference related upside not reflected in consensus estimates, in our view (see here for more). Our mid'27 price target of $1,250, representing 15% upside, implies ~25x fwd. AFFO/share.

Powered Shell Providers

Stephen Byrd

Powered Shell Provider (PSP) stocks—WULF, CIFR, HUT, RIOT, APLD, and GLXY— develop some of the most lucrative real estate projects in the US, but their shares declined following recent news from META and Kimi. Across 17 data center lease transactions that these companies signed with AI customers, value creation ranges from ~ $7 to $19 per watt and continues to trend higher.

Exhibit 24: Equity Value Creation on PSP Deals Trends Higher Over Time



Source: Morgan Stanley Research

We believe PSP stocks will eventually trade at a 15x EV/Watt multiple, a discount to the 20-25x multiple omature data center providers such as EQIX and DLR, which Cameron McVeigh covers. Given the recent downturn, we reiterate that the market undervalues the PSP names on an EV/Watt basis as shown in Exhibit 10 .

US Power

David Arcaro, Robert Kad, and Joe Laetsch

We see onsite power generation equipment manufacturers such as **BE** gaining traction as grid interconnection lead times exceed five years in certain regions. As data centers face increasing pressure to secure or fund the power they consume directly, we see onsite generation as another key solution. This approach improves "time to power" and supports a "bring your own generation" model, two critical requirements for the next phase of data center construction. We therefore expect onsite power to remain a major strategic priority for the industry. The opportunity for growth expands to spans small modular turbines, like those produced by **XE**. While not onsite power, XE is another name we see as a winner despite recent underperformance because some data centers may sign SMR deals to secure power in the long-term.

The downturn also pressured onsite power project developers such as **SEI, WMB, and LBRT**, but we continue to see attractive opportunities across this group. For SEI, we see a disconnect between the stock price and the value of future growth. We estimate that SEI's existing procured fleet of 3,200 MW through 2029 is worth ~$66/share on a standalone basis; within that total, we see a high likelihood that SEI contracts its available 800-900 MW in the near term. The stock trades modestly below this value, implying 10% upside to the base business without assuming any growth beyond 2028. Grid bottlenecks will likely persist and worsen through 2030, creating several additional growth avenues: (1) contracting additional turbine supply for 2029 and beyond; (2) adding engines, fuel cells, or batteries to its offerings; and (3) re-contracting capacity at higher prices over time. We assign $24/share to post-2028 growth, assuming SEI adds 500 MW annually through 2035. We view that assumption as well within available turbine supply and conservative relative to both the 140 GW data center TAM by 2030 and the 38 GW power shortfall by 2028.

WMB, covered by Robert Kad, and **LBRT**, which Joe Laetsch covers, represent a differentiated group of onsite and behind-the-meter power project developers that can serve data centers and other power-intensive customers. WMB offers the broadest integrated model, combining natural gas pipeline laterals and supply—including, in some cases, upstream production, gathering, and processing volumes—with gas-fired generation, fiber-optic cable, batteries, and potentially carbon capture and sequestration. LBRT is building a complementary Power-as-a-Service platform through Liberty Power Innovations and targets ~3 GW of onsite capacity by 2029 across data centers, mining, energy, and industrial markets. The company has also agreed to deliver 1 GW of power solutions to Vantage Data Centers, including a 400 MW firm reservation by 2027. Further, LBRT's strategic alliance with SLB provides commercial support and access to modular data center infrastructure, cooling systems, and related hardware.

Finally, Grid-Connected Power Producers (Utilities)—**CEG, VST, TLN, BEP, CWEN**—all generate and sell electricity through the traditional grid via wholesale markets or long-term PPAs. The argument on these names is clear: with a projected 38GW power shortfall, the companies reliably generating and selling electricity stand to benefit.

Asia Equity and Thematic Strategy

Daniel Blake, Jonathan Garner, Ehsernta Fu, Kristal Ji

Our Take on Recent Events and Debates – No Break in the Capex Race

The current paradigm is characterised by ever-intensifying competition between the US and China AI ecosystems and underlying tech stacks. The backdrop of the Multipolar World means even more to play for in the global TAM. Outside the US & China, Sovereign AI and economic security initiatives are also driving incremental spending as countries look to maintain flexibility and build a greater degree of national resilience.

US-China and North Asian supply chains are still far too deeply interconnected to allow for any sudden decoupling arising from economic and technology security policies, as seen in the truce struck with the Busan agreement. This should continue act as a hurdle for any further step changes in market access.

However, the policy initiatives in recent years to protect leadership positions ('narrow yard, high fence') or leverage chokepoints (including rare earths licensing) are catalysing huge investment programs. Key examples include the acceleration of China's semi localisation efforts and the various critical minerals and supply chain initiatives, including Pax Silica. These capex programs will run for years and potentially decades (supply chain diversification has already been underway since tariffs picked up in 2017).

So amid the debate on potential further competitive actions, as discussed in the "Two Worlds" section, we would see an extension of these same secular investment trends.

In our view, the further growth of China's largely open-weight AI ecosystem will continue to target a leading global position, consistent with its development strategy and agenda laid out in the 2026 WAIC (see discussion by our China Equity Strategist, Laura Wang below: China Equity Strategy – State-Backed Infra Roadmap Driving Growth for Economy and Earnings). Meanwhile the more closed-source and capital intensive system in the US is faced with a stronger competitive landscape and a race to embed the technology in enterprise workflows and pursue agentic opportunities, such as E-Commerce.

Overarching all of this is the pursuit of Recursive Self Improvement/Artificial General Intelligence, which many in the industry appear to frame in a winner-takes-all mindset.

The bottom line

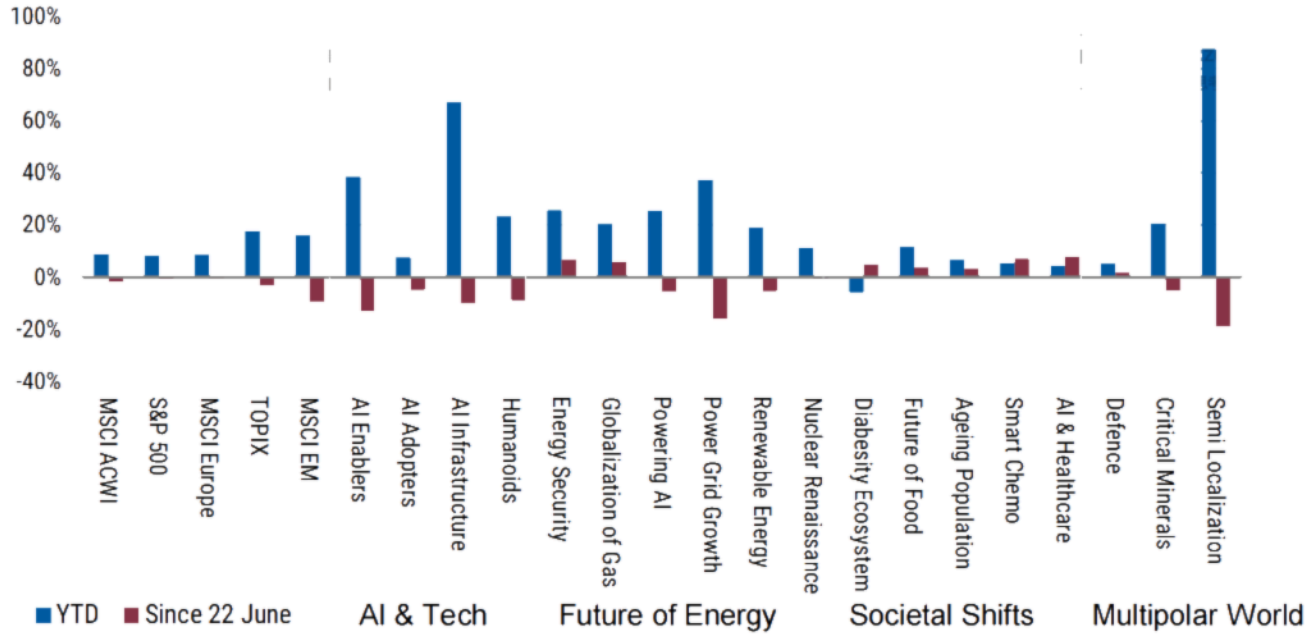
We see recent developments reinforcing the Morgan Stanley house view of an Investment Super Cycle, with a bullish view on AI capex in the coming years.

AI Supply Chain – Key Picks After the Correction

We see the recent market correction as primarily a technical event, driven by profit-taking and an unwinding of crowding and retail leverage, rather than a shift in the fundamentals for the AI capex outlook. In particular, the 21% fall in the MSCI APAC Semis/Equipment

index between 22 June and 17 July created a better setup for the AI Compute Infrastructure complex, although we acknowledge there may be some lingering impact from the unwinding of margin financing and levered investment products.

Exhibit 25: Performance across themes – some momentum reversal since 22 June, but YTD returns (equal-weighted) still positive for almost all themes, with majority >10%



Source: FactSet, Morgan Stanley Research. Data as of 23 Jul 2026. Performance available based on CIS Indices. See AI Enablers (MSRSTAIE), AI Adopters (MSRSTAIA), AI Infrastructure (MSRSTAII), Humanoids (MSRSTHUM), Energy Security (MSRSTENS), Globalization of Gas (MSRSTGNS), Powering AI (MSRSTPAI), Power Grid Growth (MSRSTPGE), Renewable Energy (MSRSTRES), Nuclear Renaissance (MSRSTNUC), Diabesity Ecosystem (MSRSTDSE), Future of Food (MSRSTFFB), Ageing Population (MSRSTPAP), Smart Chemo (MSRSTSMC), AI & Healthcare (MSRSTAIH), Defence (MSRSTDFN), Critical Minerals (MSRSTCRM), Semi Localization (MSRSTASR).

In this report, our Asia Tech Research team has discussed the latest developments and opportunities across the supply chain in the Asia Technology Research section below. We have also partnered with the team to identify a list of 30 key picks after the correction. This is far from an exhaustive list, but factors in a combination of analyst conviction, coverage 'Top Pick' designation, upside to price target, capitalisation, and the extent of the correction since June 2026.

Exhibit 26: Preferred Asia-Pacific plays on a recovery in AI Infrastructure sentiment

Ticker	Company Name	Analyst	Top Pick	Market	Mkt Cap	Share Price	Price Target	% Upside	MS Rating	Price to Earnings		Return On Equity	
					US$bn	Close, LCY	LCY			FY1	FY2	FY1	FY2
2330.TW	TSMC	Chan, Charlie		Taiwan	1,923.6	2405	2988	24%	OW	21.7	16.6	53.0%	48.4%
005930.KS	Samsung Electronics	Kim, Shawn	Top	Korea	1,129.2	270000	381000	41%	OW	NA	NA	85.5%	65.9%
0700.HK	Tencent Holdings Ltd.	Yu, Gary		China	511.0	445	650	46%	OW	14.9	13.0	20.8%	19.9%
BABA.N	Alibaba Group Holding	Yu, Gary	Top	China	277.9	114	180	58%	OW	25.5	17.0	7.1%	9.9%
285A.T	KIOXIA Holdings	Yoshikawa, Kazuo	Top	Japan	215.9	61880	110000	78%	OW	7.0	6.3	342.8%	103.2%
2454.TW	MediaTek	Chan, Charlie	Top	Taiwan	190.9	3875	5588	44%	OW	55.3	26.5	27.3%	48.5%
2308.TW	Delta Electronics Inc.	Yang, Derrick	Top	Taiwan	150.9	1880	2700	44%	OW	45.3	28.9	33.1%	41.9%
6857.T	Advantest	Tamura, Suzune	Top	Japan	132.8	30800	36000	17%	OW	41.0	27.5	68.5%	63.5%
0981.HK	SMIC	Chan, Charlie		China	114.4	70.75	85.00	20%	OW	43.5	30.0	7.7%	10.4%
2317.TW	Hon Hai Precision	Kao, Howard		Taiwan	109.0	258	310	20%	OW	15.7	12.9	11.7%	13.3%
6981.T	Murata Manufacturing	Sato, Shoji	Top	Japan	100.1	8538	12500	46%	OW	46.0	31.4	12.3%	17.3%
002371.SZ	NAURA Technology Group Co Ltd	Chan, Charlie		China	81.8	740	818	11%	OW	58.2	39.8	21.0%	26.3%
3037.TW	Unimicron	Kao, Howard		Taiwan	44.5	896	1465	64%	OW	60.3	26.8	21.4%	40.7%
6809.HK	Montage Technology Co Ltd	Yen, Daniel		China	40.8	312	432	39%	OW	NA	NA	27.9%	32.6%
2345.TW	Accton Technology Corporation	Yang, Derrick	Top	Taiwan	39.6	2325	3800	63%	OW	27.1	20.1	83.2%	68.6%
0992.HK	Lenovo	Kao, Howard		China	38.1	24.62	30.00	22%	OW	13.8	11.0	35.7%	35.6%
688012.SS	Advanced Micro-Fabrication Equip.	Chan, Charlie		China	36.3	375	550	47%	OW	61.3	44.9	16.9%	19.6%
6669.TW	Wiwynn Corp	Kao, Howard	Top	Taiwan	31.6	5470	7500	37%	OW	15.3	11.7	53.3%	52.8%
8046.TW	Nan Ya PCB	Kao, Howard		Taiwan	26.1	1230	1550	26%	OW	78.6	37.2	21.9%	38.9%
5274.TWO	Aspeed Technology	Yen, Daniel		Taiwan	19.9	15540	20888	34%	OW	82.8	57.0	108.5%	99.6%
3231.TW	Wistron Corporation	Kao, Howard		Taiwan	16.1	174	210	21%	OW	12.4	9.6	16.7%	19.5%
6361.T	Ebara	Tamura, Suzune		Japan	15.7	5752	7100	23%	OW	26.3	20.4	19.7%	22.9%
5801.T	Furukawa Electric	Shirakawa, Yu		Japan	15.3	3513	6600	88%	OW	29.5	18.4	20.0%	27.6%
NXT.AX	NEXTDC Ltd	McLeod, Andrew		Australia	7.4	13.92	18.00	29%	OW	NM	NM	-3.6%	-8.7%
6871.T	Micronics Japan	Hasegawa, Yoshihito		Japan	3.8	16370	19000	16%	OW	29.7	22.0	31.0%	34.5%
KC.O	Kingsoft Cloud Holdings	Liu, Yang		China	3.0	9.69	15.00	55%	OW	NM	521.1	-9.4%	-3.4%
6805.TW	Fositek Corp	Yang, Derrick		Taiwan	3.0	1455	2426	67%	OW	21.5	15.8	64.4%	56.8%
6890.T	Ferrotec	Hasegawa, Yoshihito		Japan	2.5	8940	13000	45%	OW	17.2	13.5	9.4%	11.3%
VNET.O	VNET Group Inc	Tang, Tom	Top	China	2.2	7.77	16.00	106%	OW	53.9	23.4	4.3%	7.2%

Source: FactSet, Morgan Stanley Research; data as of 23 July 2026

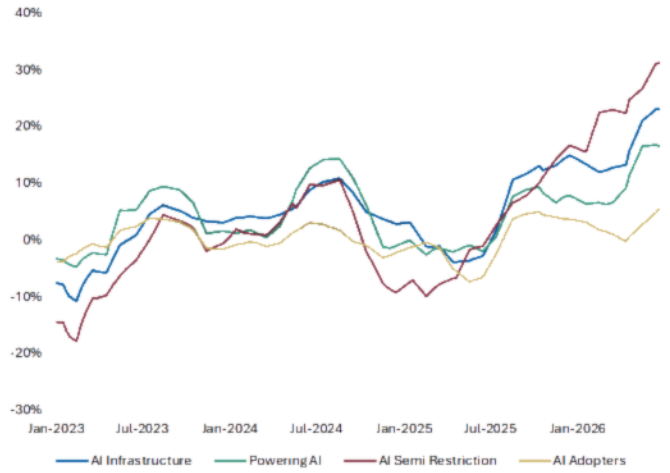
Technical unwinding provides better entry to other key themes

More broadly, we have also seen a correlated pullback in many related-themes, including across the Powering AI, Energy Security and Defence groups (which we acknowledge were well-owned by early 2026). From here, we see improved risk-reward from a share price perspective, while the fundamental imperative for Future of Energy and Multipolar World related themes has strengthened further as Middle East conflict has re-escalated.

For a roadmap to our bullish outlook for capex-related themes beyond AI, see:

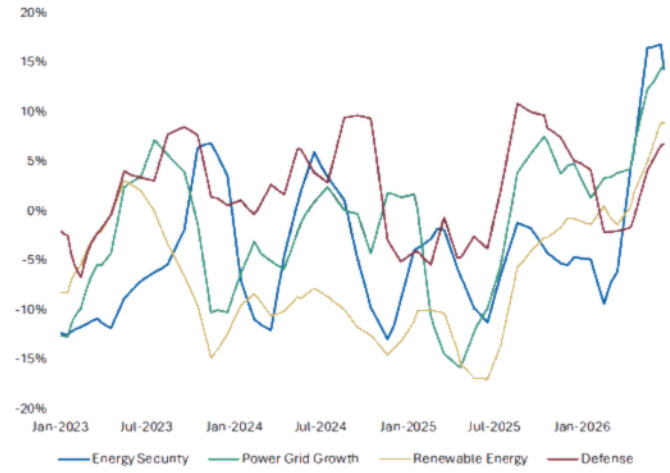
- Asia Energy Security and AI: Energy Meets Compute: Supercycle Recharges (21 May 2026)
- Asia Economics Mid-Year Outlook: The industrial super-cycle is outweighing the energy shock (12 May 2026)
- Asia Thematic Alpha: Asian Plays on the Capex Super-Cycle (13 May 2026).

Exhibit 27: AI Related Themes Earnings Estimate Revision Breadth (3mma) – AI-related themes showing ongoing strength



Source FactSet, Morgan Stanley Research. Weekly data as of 17 July 2026

Exhibit 28: Future of Energy and Defence Earnings Estimate Revisions Breadth (3mma) – Strong recovery since March



Source FactSet, Morgan Stanley Research; weekly data as of 17 July 2026

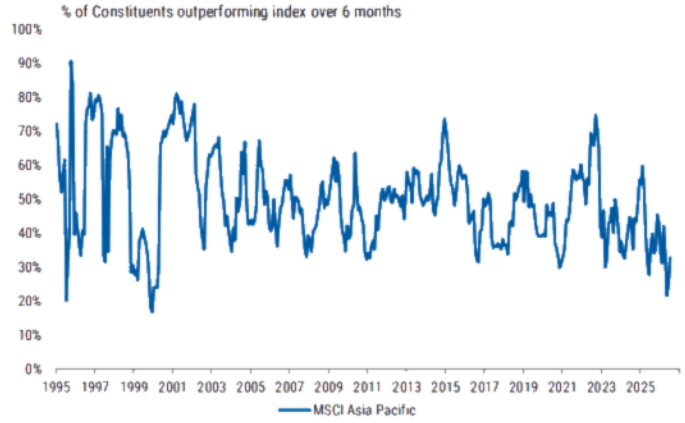
Exhibit 29: Growth vs. Valuation by Theme – China's AI Path offers strong prospective growth



Source FactSet, Morgan Stanley Research. Data as of 27 July 2026

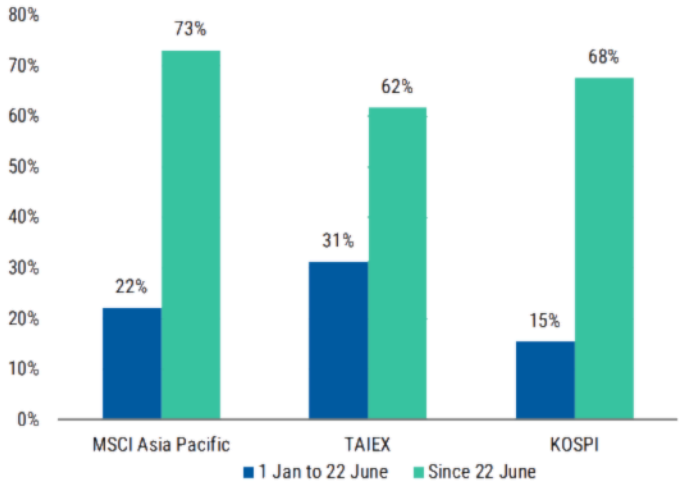
In assessing how far we are through the correction, we note that sentiment had become exuberant, with broad-based investment in AI bottlenecks and memory in particular. Back on 22 June, we highlighted a record-low share of index outperformers, outside of 2H 1999, but since then we've seen a solid improvement in market breadth.

Exhibit 30: Asia Pacific market breadth is picking up from 2H 1999 lows



Source: Bloomberg, Morgan Stanley Research; Data as of 24 July 2026

Exhibit 31: Share of stocks outperforming EM in 2026, before and after 22 June



Source: Bloomberg, Morgan Stanley Research; Data as of 24 July 2026

Crowding was particularly acute in the HF community, as reported by our PB Content team, but we had also seen extensive overweights built across LO investors in Semis/ Equipment and closely-related Capital Goods names (including CATL and suppliers of bottlenecked power equipment, such as Mitsubishi Heavy Industries and HD Hyundai Electric). The prior systematic analysis of flows and positioning signals led by Kristal Ji had been flagging a cautious near-term outlook, particularly in Korea and Taiwan (see Asia EM Equity Strategy: Flows & Positioning Guide, 2 Jul 2026).

Exhibit 32: Flows & Positioning Signals Monitor – as published on 2 July 2026

Indicator	Source	N3M Returns	N3M Vol
KOREA			
GEM Positioning	EPFR		-
Foreign Flows	Korea Exchange (Bloomberg)	-	
Retail Flows	Korea Exchange (Bloomberg)		
CHINA			
Southbound Flows vs HSI	HKEX Stock Connect (CEIC)		-
National teams flow vs. CSI 300 Financials	EPFR		-
Active long-only foreign Flows vs MSCI China	EPFR		-
Passive long-only foreign flows vs MSCI China	EPFR		-
GEM Positioning vs MSCI China	EPFR		-
TAIWAN			
Foreign Flows	Taiwan Stock Exchange	-	
Retail Margin Balance	Taiwan Stock Exchange	-	
GEM Positioning	EPFR		
INDIA			
Foreign Flows	Central Depository Services (Bloomberg)		-
Domestic Institutional Investor Flows	SEBI, CDSL		
GEM Positioning	EPFR		-
JAPAN			
Foreign Flows	Tokyo Stock Exchange		
Individual Flows	Tokyo Stock Exchange		
Long-Only Positioning (Global)	EPFR	-	-

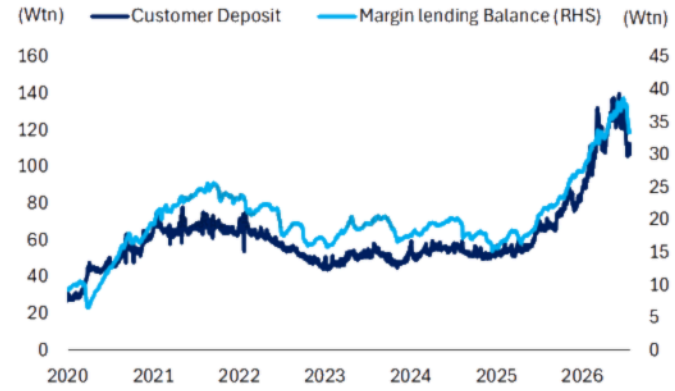
Signals from historical backtesting:	
Better Forward Return / Lower Volatility	
Neutral/Mixed	
Weaker Forward Return / Higher Volatility	
No Signal	-

Source: Morgan Stanley Research. Note: For additional details on the framework and market-specific dynamics, see Asia ex-J Equity Strategy: Exploring Signals from Flows & Positioning - A Practical Guide (21 Apr 2026). The signals are derived from historical backtesting and may not be indicative of future performance. They are intended as a supplementary input and do not override our fundamental views within the regional market allocation framework.

Another area of concentration is Korean retail leverage and a buildup of flows into levered ETFs, which contributed into a short-gamma concentration that amplified volatility. This concern around crowding and leverage had been one of the factors holding us back from a more bullish stance on Korean and Taiwanese equities, despite strong earnings momentum. Instead, we have been arguing for a broadening out in regional equities, with diversification into other upstream beneficiaries of the investment super-cycle and Multipolar trends (see Asia Thematic Alpha: Multipolarising Opportunities into 3Q, 22 June 2026).

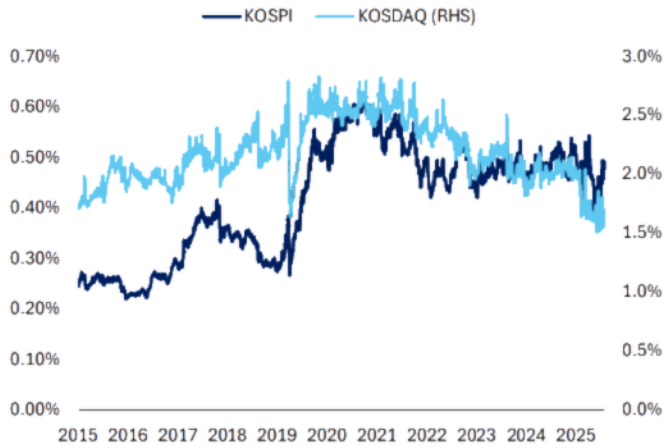
More recently, we have seen margin lending balances correct somewhat. These remain well above 2025 levels in absolute terms, but have fallen back towards historic lows as a share of market cap. Our South Korea strategy team discusses their market outlook and technical signals in more details below in Korea Equity Strategy .

Exhibit 33: Korean security account deposits and margin lending balances



Source: WiseFN, KOFIA, Morgan Stanley Research. Data as of 22 July 2026

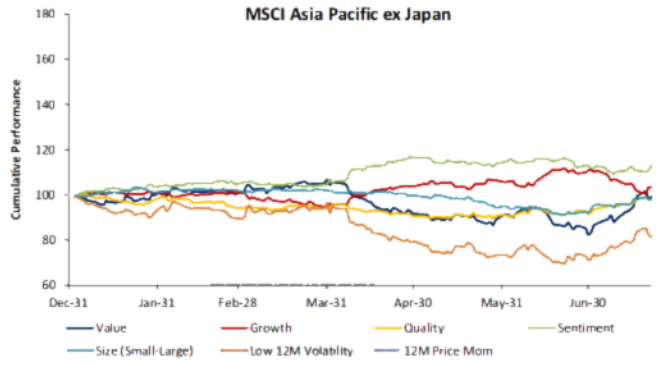
Exhibit 34: Korean margin balances as a percentage of market cap



Source: WiseFN, KOFIA, Morgan Stanley Research. Data as of 22 July 2026

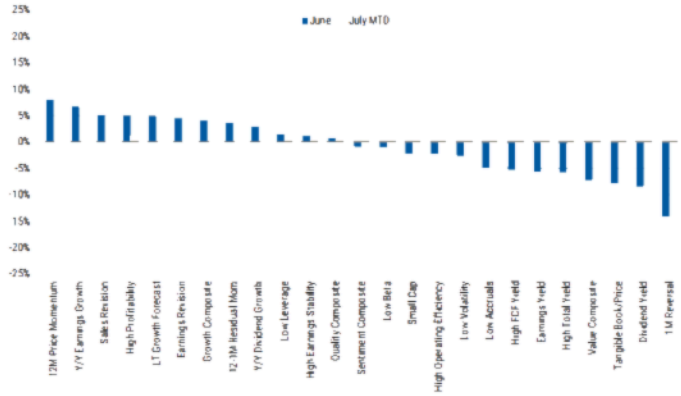
From a factor perspective, our APxJ Quant Strategist, Simon Jin, calculates a 28% reversal of the very strong momentum performance seen YTD. Growth has also come under pressure, underperforming by 9ppt during this pullback. In contrast, Value and Low Volatility have been the key beneficiaries, outperforming by 22.8ppt and 19.6ppt, respectively, with Quality also registering a solid 8ppt outperformance. Although Momentum has rebounded sharply since July 21, it remains too early to conclude that the correction has fully run its course. For investors looking to lock in gains while maintaining market exposure, Simon recommends a balanced multi-factor approach combining Value, Growth, and Earnings Estimate Revisions.

Exhibit 35: YTD factor performance – MSCI APxJ



Source: FactSet, Morgan Stanley Research. Data as of 22 Jul 2026

Exhibit 36: Factor performance in June vs July MTD – MSCI APxJ



Source: FactSet, Morgan Stanley Research. Data as of 22 Jul 2026

Asia/EM Market and Sector Recommendations

From here, we recommend investors reinforce a preference for upst
the investment super-cycle. This means a barbell of positioning acro
Infrastructure with Energy, Materials and Industrials, as reflected in
Renewable Energy/Storage and Defence cohorts.

人在草木间

We recommend an underweight position across Consumer Discretionary, Staples, and Autos. For Financials, we recommend a long position in trade/investment-driven sectors and beneficiaries of reflation (Japan, Singapore) versus Australian Financials, which face near-term stagflationary conditions and an adjustment in the housing market after major tax changes in the May Budget.

Exhibit 37: MSCI EM Sector Dashboard – We prefer a combination of Tech and Energy/Materials/Industrials and remain cautious on Consumer/Services sectors broadly

	MSCI EM Sectors & Industry Groups	Automobiles & Components	Cons. Discret. Retailing	Durables & Apparel	Consumer Staples	Consumer Services	Telecoms	Media & Entertainment	Software & Services	Semiconductors & Equipment	Tech Hardware & Equipment	Financials	Real Estate	Health Care	Utilities	Energy	Materials	Industrials
MS Rating		UW	UW	UW	UW	UW	EW	EW	UW	OW	OW	EW	EW	EW	EW	OW	OW	OW
Index Weight	% of MSCI EM	2.6%	3.7%	0.7%	2.8%	1.2%	2.1%	4.2%	1.1%	25.5%	15.4%	19.4%	1.0%	2.5%	2.0%	3.4%	5.8%	6.6%
Earnings Revisions & Growth	CY2 EPS Revision Breadth (3ma)	-8%	-10%	4%	-8%	-14%	-7%	-15%	-15%	34%	43%	2%	-1%	-7%	-8%	6%	7%	1%
	3m Δ of 12mf EPS	-1%	-3%	3%	-1%	8%	0%	1%	-2%	31%	28%	1%	-2%	0%	-7%	2%	5%	13%
	12mf EPS growth	25%	31%	15%	11%	57%	17%	9%	15%	65%	81%	9%	8%	16%	11%	6%	110%	35%
ROE	Level	8.8%	10.1%	14.8%	15.8%	1.5%	15.3%	15.5%	15.6%	32.4%	15.8%	12.5%	8.8%	10.8%	11.2%	12.6%	11.5%	13.7%
	10yr %ile	55%	43%	60%	93%	15%	98%	91%	4%	98%	88%	92%	33%	60%	90%	65%	69%	88%
12mf PE	Level	13.2x	12.3x	14.4x	17.7x	18.7x	16.2x	13.1x	19.4x	11.3x	6.7x	9.3x	10.3x	26.3x	12.1x	9.0x	10.2x	11.9x
	10yr %ile	63%	24%	27%	2%	21%	74%	4%	8%	12%	1%	84%	73%	45%	79%	69%	26%	43%
P/Book	Level	1.6x	1.4x	2.4x	3.2x	1.8x	2.9x	2.3x	3.3x	11.8x	5.5x	1.2x	1.0x	3.7x	1.4x	1.2x	1.9x	2.9x
	10yr %ile	59%	1%	18%	2%	9%	82%	2%	1%	99%	99%	77%	34%	54%	70%	64%	83%	99%
P/Sales	Level	0.9x	1.1x	1.1x	1.0x	1.8x	2.3x	4.0x	2.7x	11.3x	1.9x	1.8x	1.7x	3.4x	1.2x	0.9x	1.5x	1.6x
	10yr %ile	63%	11%	54%	2%	2%	95%	20%	2%	98%	98%	97%	76%	82%	78%	88%	87%	95%
Div Yield	Level	1.7%	1.2%	2.6%	3.3%	1.2%	3.3%	1.2%	3.3%	0.8%	0.9%	3.7%	4.0%	1.0%	3.4%	4.2%	2.8%	1.7%
	10yr %ile	40%	76%	87%	99%	79%	35%	98%	99%	2%	2%	44%	57%	69%	67%	42%	26%	24%
Technicals	30 day RSI	34	48	34	55	55	56	51	41	44	38	62	55	60	44	46	39	40
	50dma / 200 dma	1%	-15%	2%	-3%	-13%	2%	-16%	-16%	43%	47%	1%	-3%	-4%	0%	1%	-1%	10%
Perf.	1m	-7%	8%	0%	2%	8%	-1%	6%	-2%	-15%	-26%	-1%	-1%	7%	0%	3%	-8%	-14%
	3m	-14%	-12%	-2%	-5%	-8%	0%	-8%	-16%	22%	10%	-1%	-2%	-3%	-10%	-6%	-16%	-9%
	YTD	-7%	-16%	0%	-5%	-19%	3%	-21%	-31%	68%	65%	3%	-3%	0%	3%	5%	-6%	7%

Source: FactSet, Morgan Stanley Research. Data as of 23 July 2026

Japan remains our core overweight recommendation at an index level, given ongoing strong earnings momentum from the reflation and reform backdrop. The recently released economic and fiscal policy is consistent with our bullish outlook for investment-related sectors, as laid out by our strategist, Sho Nakazawa, in Japan Equity Grand Strategy under Takaichi Administration (23 Feb 2026). Alongside a tech recovery, we could also see broader North Asia moving higher, noting that our equal-weight stance for Taiwan, Korea, and China is relative to our EM outlook, which has 11% upside to our base case in June 2027.

Exhibit 38: Asia/EM Market Dashboard – Japan and Singapore remain core overweight recommendations, but North Asia should perform well in the near term via recovery in sentiment towards AI and broader capex

	MSCI Indices	Japan	China	Hong Kong	Taiwan	Korea	India	Australia	Singapore	Malaysia	Thailand	Indonesia	Philippines	Brazil	Mexico	South Africa	Saudi Arabia	UAE	MSCI EM	MSCI APxJ	USA	Europe
MS Rating	vs Asia/EM	OW	EW	EW	EW	EW	EW	EW	OW	EW	EW	UW	UW	OW	EW	EW	EW	EW	EW	-	-	-
Index Weight	% of APAC-EM	26%	13%	2%	17%	13%	7%	7%	2%	1%	1%	0%	0%	2%	1%	2%	2%	1%	11.9% (ACWI)	12.2% (ACWI)	63.6% (ACWI)	13.9% (ACWI)
Earnings Revisions & Growth	CY2 EPS Rev Breadth (3ma)	17%	-2%	-3%	59%	26%	-8%	-20%	7%	-2%	11%	-29%	-26%	-7%	-16%	-23%	0%	-10%	2%	2%	26%	10%
	3m Δ of 12mf EPS	8%	0%	0%	13%	25%	1%	1%	2%	3%	2%	0%	1%	1%	2%	-6%	0%	-1%	10%	10%	6%	0%
	12mf EPS growth	14%	14%	10%	35%	78%	14%	7%	10%	9%	8%	5%	6%	15%	11%	21%	11%	6%	36%	37%	21%	13%
ROE	Level					15.1%	13.9%	12.8%	11.5%	10.6%	10.9%	14.8%	14.9%	18.5%	16.3%		14.1%	19.1%	13.9%	12.7%	20.9%	13.9%
	10yr %'ile	72%	23%	59%	90%	98%	49%	67%	79%	92%	60%	50%	98%	68%	98%	91%	93%	98%	98%	98%	98%	85%
12mf PE	Level	16.0x	10.8x	14.1x	19.1x	5.2x	20.6x	18.0x	17.4x	14.1x	18.3x	9.2x	11.1x	8.2x	12.3x	8.4x	13.4x	8.6x	10.3x	11.2x	20.3x	14.8x
	10yr %'ile	82%	31%	39%	94%	1%	51%	79%	98%	34%	88%	2%	18%	39%	26%	7%	2%	6%	3%	3%	58%	64%
P/Book	Level	2.0x	1.3x	1.1x	6.2x	4.2x	3.2x	2.6x	2.0x	1.5x	2.1x	1.3x	1.5x	1.7x	2.4x	2.3x	2.1x	1.8x	2.7x	2.7x	5.8x	2.5x
	10yr %'ile	98%	12%	26%	99%	99%	49%	95%	89%	31%	79%	1%	2%	50%	77%	77%	41%	80%	99%	99%	98%	99%
P/Sales	Level	1.6x	1.4x	2.6x	3.6x	2.4x	2.5x	2.5x	3.3x	2.3x	1.4x	1.8x	2.4x	1.1x	1.2x	1.6x	3.0x	2.3x	2.0x	2.3x	3.6x	1.7x
	10yr %'ile	97%	53%	46%	98%	98%	66%	83%	99%	65%	85%	4%	58%	45%	65%	49%	18%	17%	98%	98%	98%	98%
Div Yield	Level	1.7%	2.3%	3.5%	1.4%	0.8%	1.2%	3.3%	3.0%	3.9%	3.0%	7.7%	3.2%	5.6%	3.9%	3.7%	3.8%	4.9%	2.0%	1.9%	1.1%	2.8%
	10yr %'ile	1%	77%	68%	2%	3%	45%	12%	7%	67%	61%	98%	92%	54%	96%	80%	76%	63%	9%	8%	4%	21%
Technicals	30 day RSI	51	50	57	48	45	63	61	78	58	56	67	67	61	53	45	43	58	47	48	51	59
	50/200d MA	9%	-8%	-1%	29%	48%	-2%	0%	5%	1%	12%	-24%	2%	-2%	2%	-4%	-1%	-4%	13%	13%	7%	4%
Rel. Pos.	vs MSCI EM	NA	0.3%	0.0%	-1.1%	0.8%	-1.0%	-0.9%	-0.1%	-0.4%	0.0%	0.1%	0.0%	1.7%	0.5%	-0.4%	-0.7%	NA	NA	NA	NA	NA
Perf.	1m	-6%	2%	4%	-10%	-27%	-1%	1%	8%	2%	0%	2%	6%	4%	-3%	-4%	-3%	-6%	-9%	-9%	-1%	0%
	3m	1%	-9%	-5%	14%	12%	-3%	-3%	11%	-2%	2%	-22%	10%	-14%	-5%	-14%	-6%	-2%	1%	2%	5%	0%
	YTD	11%	-10%	3%	48%	70%	-11%	9%	15%	2%	23%	-36%	9%	11%	8%	-9%	3%	-5%	15%	16%	9%	5%

Source: FactSet, Morgan Stanley Research. Data as of 23 July 2026

All that said, we continue to see a wide bull-bear range for markets, as we called in our 2026 Asia EM Equity Outlook: A Robust Approach for an Uncertain World (16 Nov 2025). We still see a long list of potential risks, including:

- The impact of Middle East developments on energy and product markets,
- The rise in long-dated real yields and broader credit market conditions (particularly for tech issuance).

Meanwhile, bottlenecks in the data center buildout may mean a slower pull-through for the supply chain, eroding pricing power and margins.

Finally, AI-related policy remains a source of uncertainty in the coming months, both in terms of regulatory approvals as well as in the US-China competitive dimension.

Asia Technology Research

The combination of a return of confidence in global AI capex and the ongoing milestones across the Chinese AI ecosystem is likely to put the focus back on a number of upstream beneficiaries. Here we highlight several key segments and opportunities across our coverage.

The Rise of China's AI Chip Ecosystem

Charlie Chan, Daisy Dai, Henry Zhao

China's AI compute industry could be competitive globally, given strong system design and infrastructure

More broadly, we believe China's AI GPU race is no longer just a chip specification contest. While domestic silicon still trails the US by roughly two generations at the chip level, the effective gap is narrowing through multi-die design, advanced packaging, rack-scale system architecture, optical networking, and software-hardware co-optimization.

This is why we think system-level competitiveness matters more than ever. In a market increasingly dominated by inference and utilization, vendors that provide the best real-world token economics at acceptable software migration cost are likely to win customer budgets, even without access to leading-edge process technology.

From an investment perspective, this leads to a simple conclusion: we believe the sector should not be valued as a monolithic policy theme.

Instead, we believe investors should differentiate between vendors with a realistic path to shipment scale, ecosystem credibility, and pricing discipline, and those that may face challenges translating technical potential into sustainable revenues and margins.

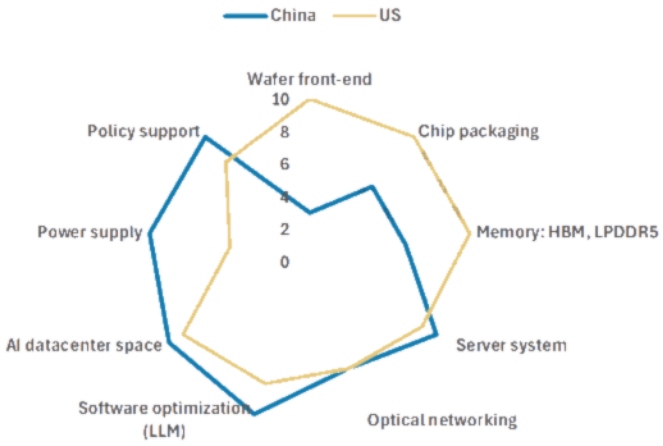
We therefore evaluate the group through a two-dimensional framework of economics × execution, combining TCO, token cost, TPS, and performance-per-dollar with qualitative factors such as foundry access, software ecosystem maturity, CSP relationships, and roadmap credibility. In our view, this framework remains the most effective approach to distinguishing companies that are well positioned to benefit from industry growth from those that could face greater challenges as the industry consolidates.

What has become clearer through our recent channel checks is that economics, not ideology, is driving adoption. At our China Summit, several major LLM developers indicated a willingness to deploy domestic GPUs as long as token cost is competitive. This aligns with our core finding that domestic accelerators already offer materially lower TCO than Nvidia products currently available in China, and that leading Chinese chips can achieve broad cost-per-token parity in inference workloads.

In other words, purchasing decisions appear to be increasingly made on deployable economics rather than absolute peak silicon performance. This matters because China's AI demand is becoming more inference-heavy, more recurring, and more utilization-driven, which structurally favors solutions optimized around cost efficiency, software adaptation,

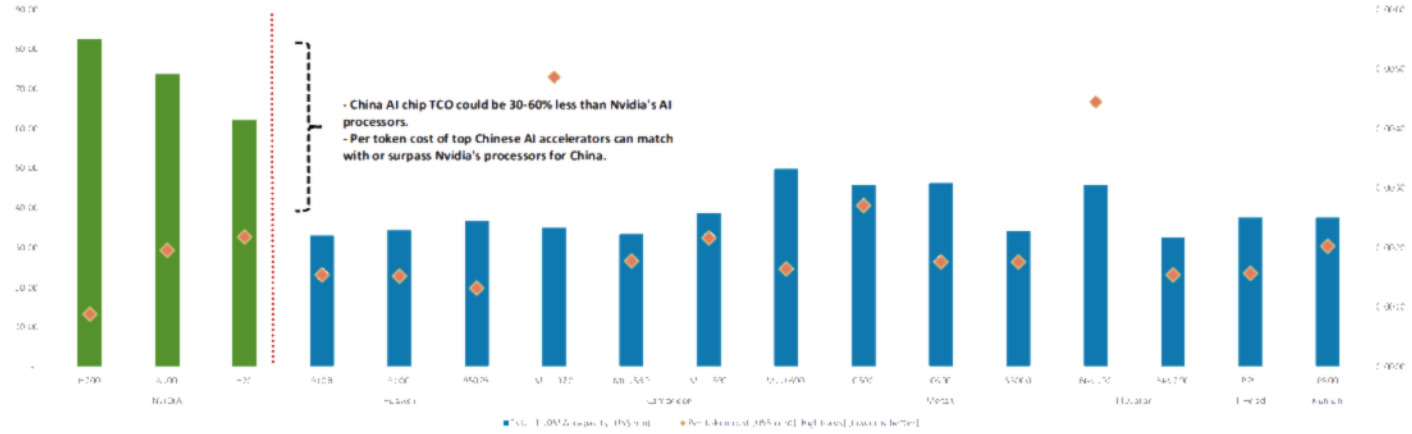
and availability rather than headline benchmark leadership.

Exhibit 39: Relative strengths of AI industries in the US and China



Source: Morgan Stanley Research

Exhibit 40: Domestic chips have lower TCO and comparable per token cost (AI LLM inference) vs. NVIDIA's processors for China



Source: Company data, Morgan Stanley Research estimates

China – AI GPU TAM

China's AI GPU demand is concentrated across a small number of large buyer groups, whose capex decisions ultimately define the size of the addressable market.

- The first group comprises China's CSPs – including ByteDance, Alibaba, and Tencent – which procure AI chips both to train and run inference on proprietary models and to deploy AI infrastructure for external cloud customers.
- The second group includes China's telecom operators, SOEs and municipal governments – the so-called sovereign AI buyers – where demand is driven by national AI infrastructure build-out, data sovereignty, and public-sector applications.
- AI start-ups (e.g., DeepSeek, MiniMax) and auto OEMs (e.g., XPeng, Xiaomi) also purchase AI chips, although their volumes currently remain smaller than those of

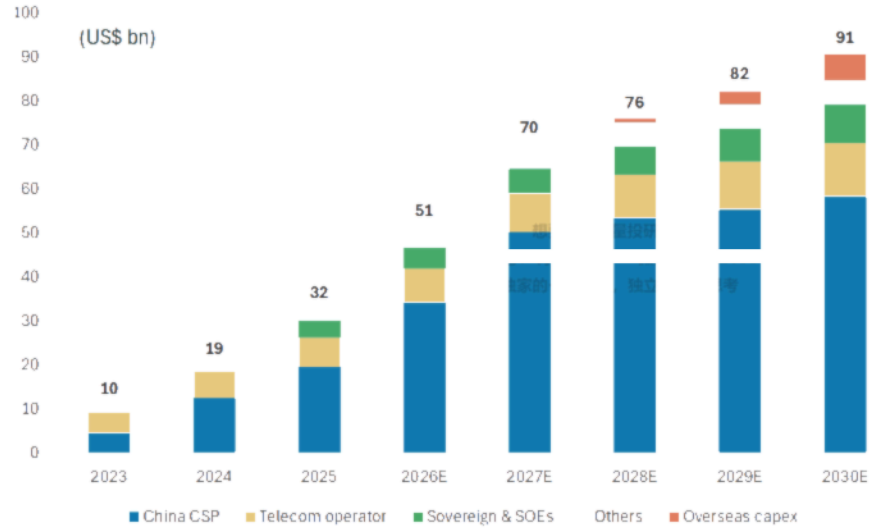
the first two groups.

We forecast China's AI chip TAM to reach US$91bn by 2030 up 36% from US$67bn previously, implying a 23% CAGR over 2025-30. The increase is mainly because:

- We added a new category (Chinese CSPs' overseas AI data centers using local GPUs). In the short term, China CSPs may turn to more GPU rental to fulfill the strong computing demand, while in the medium to longer term, it is likely that Chinese AI GPUs may see overseas adoption. Therefore, we assume that part to be zero until 2027, but from 2028 onwards, we expect 3%, 10%, and 20% of overseas capex to be addressed by local CPUs in 2028, 2029, and 2030.
- ByteDance is planning to significantly increase its capital spending in 2026 and 2027 as it seeks to strengthen its position in China's AI market and expand its competitiveness internationally. Media reports have suggested that 2027 capex could be US$100bn if economic and business conditions are favourable; for our forecasts, we assume a more conservative US$80bn (~Rmb542bn).
- We added Kingsoft Cloud to our database; its capex surged to Rmb3bn, with full-year 2026 capital investments projected to exceed Rmb15-20bn to meet explosive AI and cloud demand.
- We raised sovereign and SOE-related TAM to US$9bn up from US$7bn. Recent news (*Bloomberg*, June 9) reported that China is preparing Rmb2tn over the next five years for building data centres across the country. So far we have not seen specific government guidance, but directionally SOE and local government will spend more on AI infrastructure.

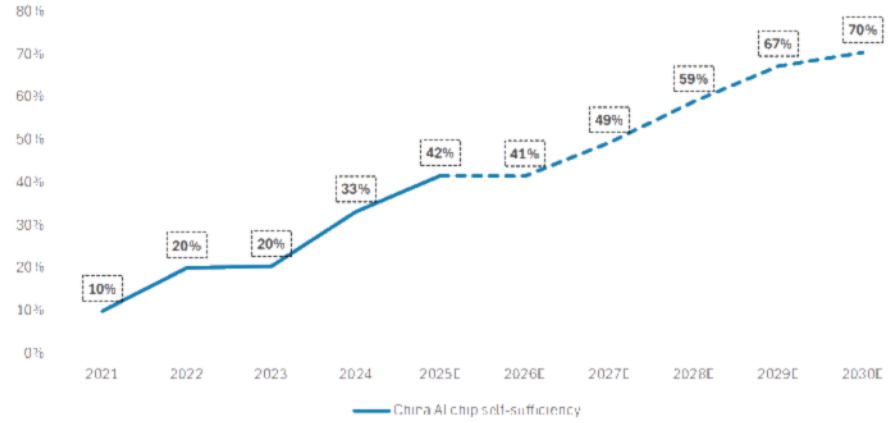
We expect China's AI-chip self-sufficiency to rise from 42% in 2025 to 70% in 2030e. We expect leading node capacity expansion and continued chip performance improvement to drive local AI chip revenue growth.

Exhibit 41: We expect China's AI chip TAM to grow to US$91bn by 2030F



Source: Company data, Morgan Stanley Research (E) estimates

Exhibit 42: We expect China's AI chip self-sufficiency to reach 70% in 2030E



Source: Company data, Morgan Stanley Research (e) estimates

Cloud Semis – Larger TAM Supports Outlook

Daniel Yen

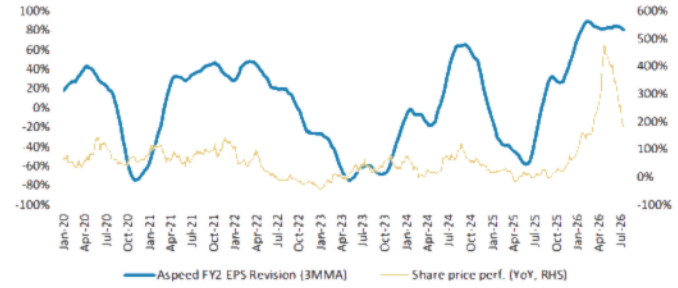
We expect a larger TAM for both CPU and GPU servers through 2027. Reiterate OW on both Aspeed and Montage into earnings. Build positions into cloud capex prints.

CPU servers – showing more upside, especially in China: During our recent WAIC trip in China, we observed increasing adoption of AMD, Nvidia and Hygon. AMD's multicore architecture, Nvidia's Vera bundle sales strategy, and Hygon's government support and performance catch-up with legacy 6 gen Intel CPU are now contributing to rising demand. High-end X86 CPU pricing continued to trend up in China, nearly 2x Y/Y.

GPU servers – strong demand supported by SuperPods solutions: Nearly every domestic GPU vendor displayed a 64- or 128-card scale-up system. Chinese AI GPUs continue to face manufacturing process constraints at the chip level, but are trying to connect each GPU (link to WAIC report). Nvidia may also broaden its rack-scale offerings over time, as Kyber may not be the optimal solution. We believe SuperPods could be one such variance, connecting multiple racks with optics (multi-racks scale-up). Kimi K3 also demonstrates the importance of Nvidia GPUs for training workloads.

Margins likely to remain at a high level: We think cloud semis vendors will benefit from the spec migration, driven by strong demand for high-performance servers. Gross margins on newer solutions (e.g. AST 2700 BMC, legacy BMC with software offering and DDR5 gen 3 interface) are usually higher than those of prior-generation products, which could ffset the higher cost burden, especially from OSAT.

Exhibit 43: Aspeed: Estimate revision breadth



Source: Company data, Morgan Stanley Research estimates

Exhibit 44: Aspeed: Forward P/E ratio



Source: Company data, Morgan Stanley Research estimates

Memory

Shawn Kim, Cindy Huang

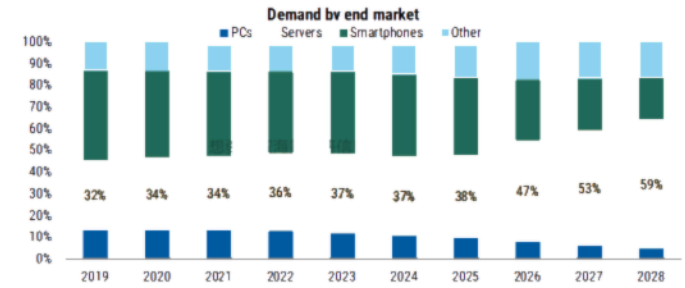
For more details, please see Global Technology: Chipflation – Navigating A Memory Crisis (2 Jun 2026)

A Multi-Year Bottleneck

The extremely robust demand for AI infrastructure has caused supply constraints in all types of memory (DRAM, HBM, NAND, HDD). As AI increasingly moves beyond data centers into edge devices and the physical world, it could trigger a further surge in demand for memory chips. In that scenario, these components would no longer be seen as commodity chips but as critical pillars of the global digital infrastructure.

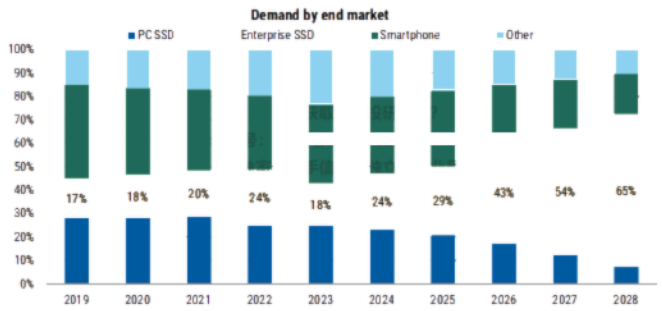
AI and data centers are becoming the dominant buyers of memory – we expect the server share of DRAM demand to rise from 37% in 2023 to 59% in 2028e, while enterprise SSDs jump from 18% to 65% of NAND demand. The result is a structurally tighter market where non-AI buyers may face weaker access to supply for an extended period.

Exhibit 45: DRAM – AI server demand becomes dominant...



Source: Trendforce estimates, Morgan Stanley Research

Exhibit 46: ...as well as NAND enterprise SSDs as key growth drivers

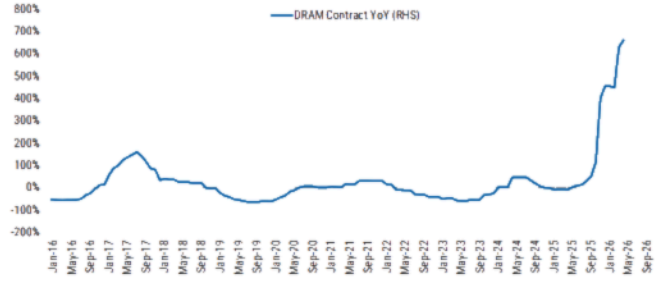


Source: Trendforce estimates, Morgan Stanley Research

Prices for memory have risen more than six-fold over the last year (Exhibit 47) – a sharp discontinuity from the multi-decade price declines. The price of a gigabyte of DRAM fell by around a factor of 10 every five years from 1957 to 2020. Since 2010, the price has fallen more slowly as Moore's Law drove sustained price declines. However, this trend no

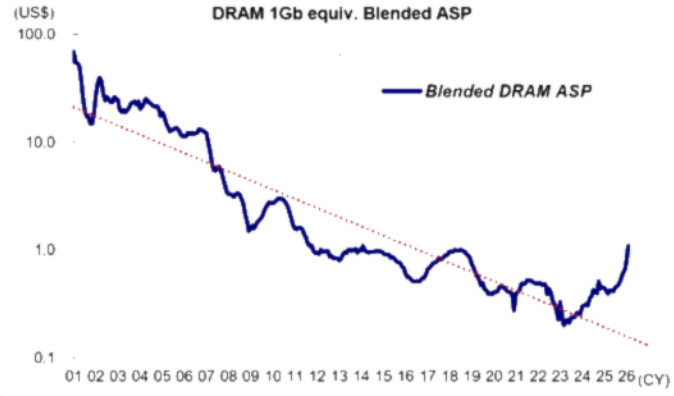
longer applies in the AI economy, which has driven DRAM ASPs sharply higher (Exhibit 48). Building memory fabrication plants (fabs) takes years, so there is no quick response to address demand. This means sustained upside risks to consumer goods prices in the coming years.

Exhibit 47: DRAM prices YoY



Source: Trendforce, Morgan Stanley Research

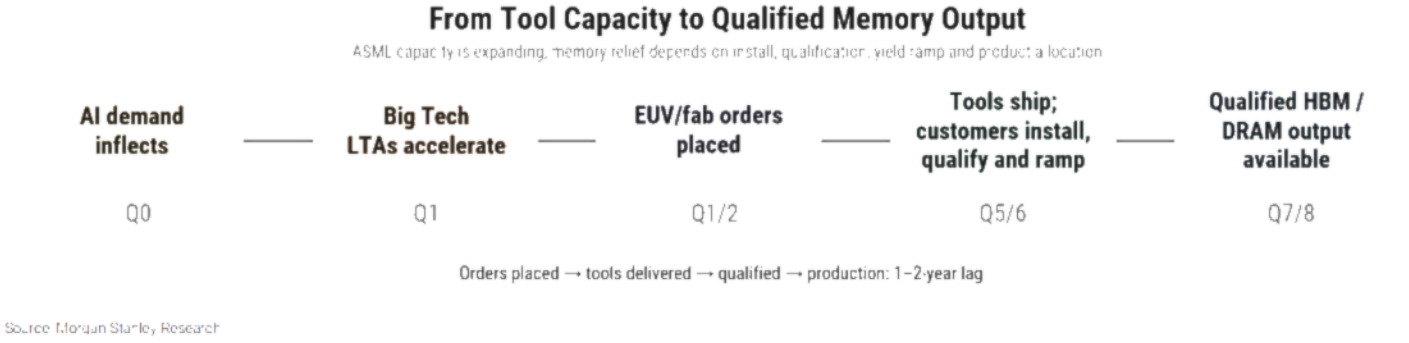
Exhibit 48: Long-term DRAM price/Gb



Source: WSTS

Exhibit 49:

Memory supply relief depends on install, qualification, yield ramp and product allocation – a ~2 year process

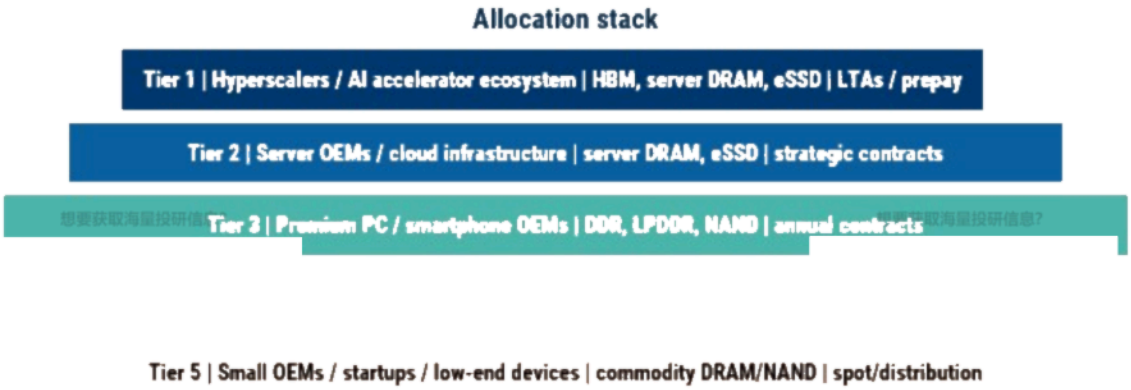


Source: Morgan Stanley Research

What's different this time? Previous chip shortages during COVID demonstrated how even minor supply disruptions can halt entire industries. However, this memory chip shortage looks more structural, and AI-related demand could amplify such effects on a much larger scale.

- **Mobile AI agents and physical AI agents – another very large additional surge in demand for memory chips is likely as these are rolled out:** As AI increasingly moves beyond data centers into edge devices and the physical world, the computational requirement for AI is set to grow exponentially. These components would no longer be seen as commodity memory chips but as critical pillars of the global digital infrastructure.
- **Long-Term Agreements (LTAs) are now absolute necessities to secure capacity** – given long lead times and the persistence of the supply-demand mismatch. Memory producers are establishing strategic supply chain partnerships with all hyperscale customers via LTAs, through which memory producers are able to lock in long-term orders and prices, establish a stable high-margin model, and mitigate the industry's cyclical volatility.

Exhibit 50: Buyer Hierarchy: Who gets memory supply first?



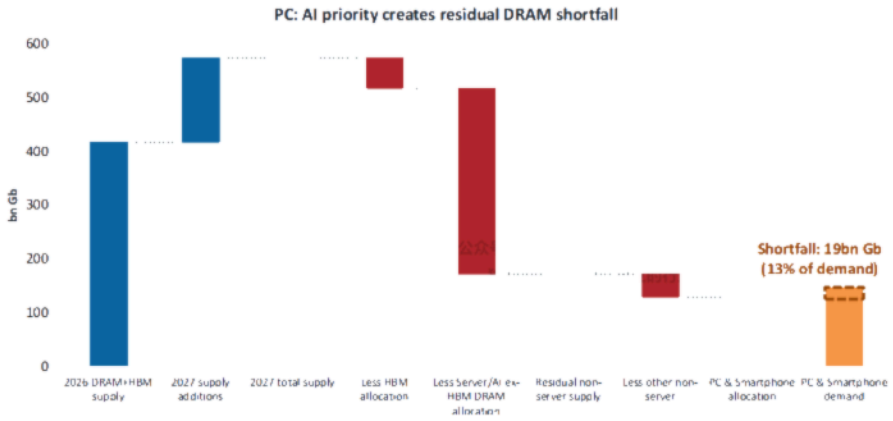
Source: Morgan Stanley Research

Memory inflation is turning into a broad hardware-sector margin risk, extending beyond PCs and smartphones into gaming consoles, servers, cloud, networking, autos, medical devices, broadband equipment, industrial systems and low-cost education/ embedded devices. The impact is not only higher DRAM/NAND pricing, but also allocation priority:

- Suppliers are prioritizing HBM, server DRAM and enterprise SSD demand tied to AI infrastructure, while traditional consumer and industrial memory users are increasingly coping with significantly higher prices, longer lead times and weaker bargaining power.
- The more affected group consists of companies that are not directly monetizing AI but still have to pay AI-driven memory prices. There is accounting asymmetry, where hyperscalers can capitalize memory-heavy AI servers and depreciate them over time, while consumer hardware and industrial OEMs see memory inflation flow more directly through inventory and COGS. This will require choices on whether to pass on price increases to consumers, reduce device specifications, or accept reduced profits.
- Memory inflation is not just a financial transfer from non-AI sectors to memory producers. It is increasingly a constraint on what gets built and when for many non-tech industries.

Applying our memory sufficiency framework separately to PC and smartphones, we find both segments face supply shortfalls in 2027. On our assumptions (PCs absorbing 17% of non-server DRAM and smartphones 58%) total addressable memory falls ~15% and -12% short of demand, respectively. This implies unit downside of - 58mn PCs (-15% vs MS shipment forecast) and ~134mn smartphones (-12% vs. Morgan Stanley's forecast). These numbers are sensitive to assumptions on memory content per device, but our analysis indicates demand rationing risk if the supply build does not accelerate.

Exhibit 51: AI prioritization turns 2027 supply growth into a consumer memory shortfall



Source: Morgan Stanley Research estimates

Greater China Tech Hardware

Howard Kao

ABF Substrates

In February this year, we turned bullish on Taiwan's ABF substrate plays in our coverage (report), on the strength of our bottom-up analysis that suggested ABF substrates will be in an up-cycle until 2030, with industry undersupply starting from 2027 onwards, all driven by AI. Since then, we have raised our estimates for companies' earnings and the entire ABF substrate TAM multiple times – *see here, here, and here.*

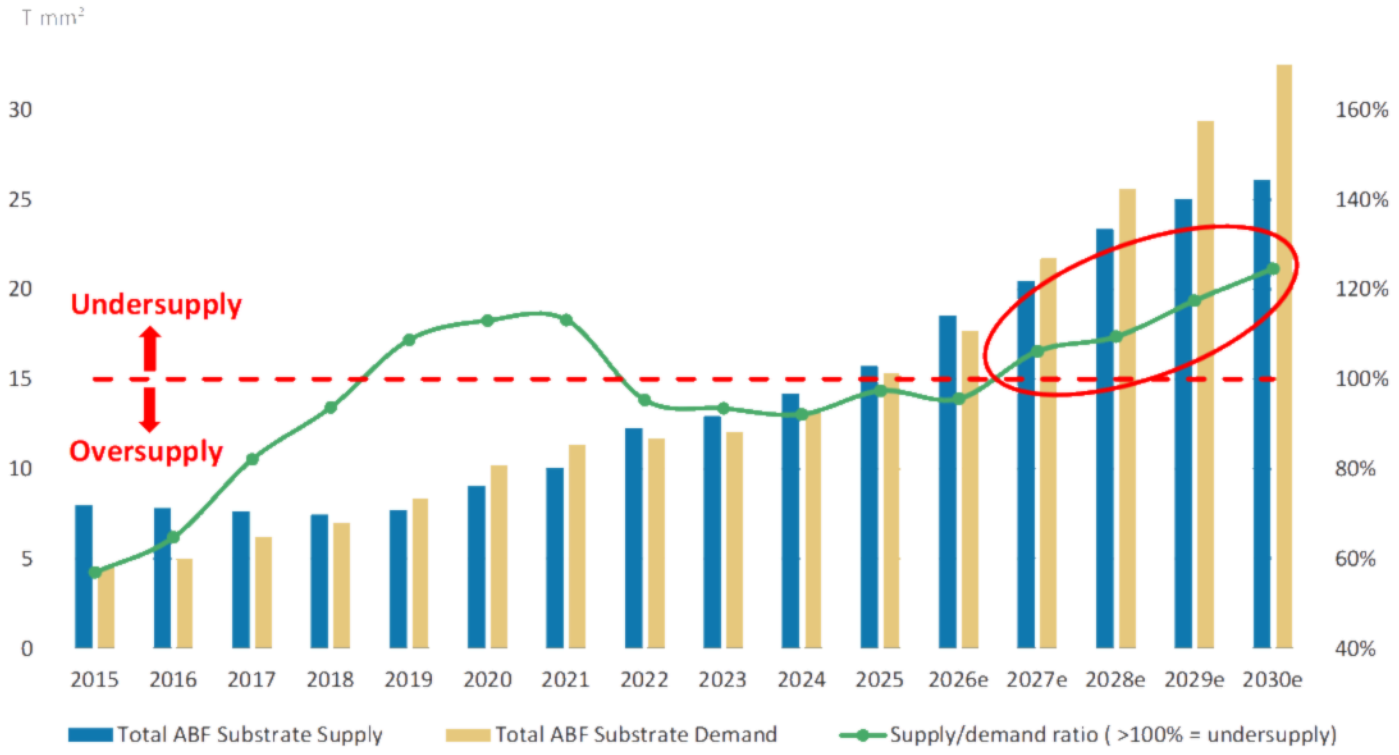
Why is AI demand the key driver of tightness?

- End market mix is flipping structurally. We expect PC-related demand (once ~70% of ABF value in 2015) will fall to <15% by 2030, while server/AI GPU/AI ASIC/ networking demand will rise from ~10% to >75% over the same period. AI chips consume roughly 10x more ABF than PC-related packages because of larger physical form factors and higher manufacturing complexity.
- We expect structural undersupply of ABF substrates to return from 2027, persisting through 2030 (~25% gap by 2030, more severe at the high end), driven by larger chip sizes, higher layer counts, and new packaging complexity (embedded passives, EMIB-T) that lowers manufacturing yields and effective output.
- Upstream T-glass bottleneck also reinforces the squeeze: We estimate a ~20% supply/demand gap in low-CTE glass cloth (needed for thicker, warpage-resistant CCL cores in AI packages). Coupled with increased precious metal pricing (primarily gold and copper), this is already leading to inflationary material pricing for substrate makers. Hence, their cost pass-through will support superior pricing power, in our view. This has also been evidenced by improving profitability of the substrate suppliers in the past few months, which we highlighted here, here, and here, despite an inflationary cost environment.
- Unlike the 2020-22 WFH-driven cycle, this up-cycle has much clearer demand visibility because suppliers can engage directly with AI end customers (Nvidia, AMD, Google, Amazon, Meta, Broadcom), giving them confidence to invest in capacity — and end-customers are reportedly offering subsidies to secure supply, a

sign of prolonged supply tightness.

Stock implications: We like Unimicron (3037.TW), NYPCB (8046.TW), and ZDT (4958.TW) in view of this story of growing AI demand globally. For a specific play on ABF in the Chinese market, we would focus on ZDT (4958.TW) because all of its ABF revenue today comes from supplying to China (report here).

Exhibit 52: Expecting ABF substrate undersupply gap to widen by 2030 vs. our prior estimates

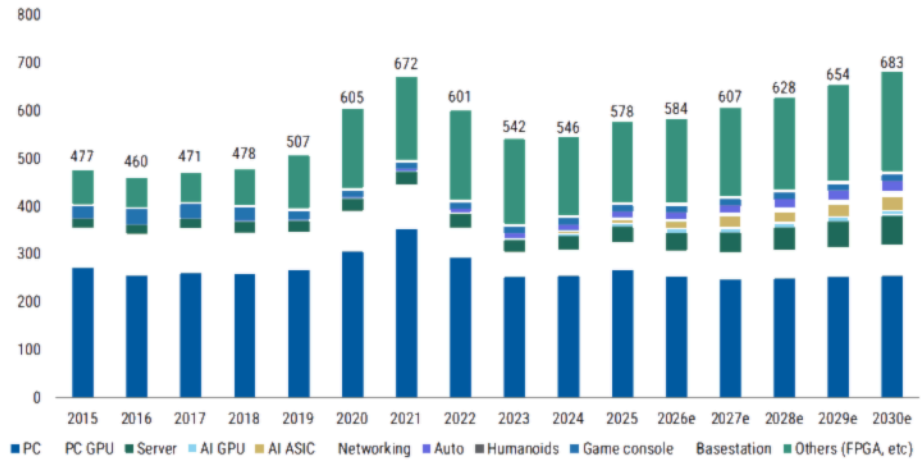


Source: Morgan Stanley Research (e) estimates.

Periods	Description
Period I (2010-2012): PC-centric equilibrium	ABF substrate demand is primarily driven and dominated by CPU and GPU demand for desktops, notebooks, and discrete GPUs, with a meaningfully smaller exposure towards server CPUs.
Period II (2012-16): PC decline and demand dilution	Global PC demand entered a structural down-cycle as consumer behavior changed, with smartphones and tablets absorbing more consumer time spent. This resulted in lower volumes and increased pricing competition as suppliers fought to fill idle capacity.
Period III (2016-18): The quiet inflection - servers and early HPC	During these years, the PC market matured, while cloud computing started to take off. At the same time, server CPU substrates quietly grew larger with more layers. We also started to see early accelerators being adopted in high performance computing.
Period IV (2018-2022): WFH demand driven structural shortage	The Covid-19 pandemic drove massive work from home demand for consumer electronics and cloud computing. This sudden surge in PC demand and increased time spent on gaming and online entertainment, bundled with a continuous increase in substrate package size and layer count, resulted in a structural shortage, which also resulted in longer lead times and raised prices. It was during this period (2021-22) that every ABF supplier globally started to invest in new capacity.
Period V (2022-2024): Post-pandemic correction and new capacity arrival	As the world came out of the Covid-19 pandemic, WFH demand faded and PC and cloud computing demand were negatively affected. But at the same time, the new investments and capacity additions were just starting to slowly come online, resulting in a drop in utilization rates and pricing.

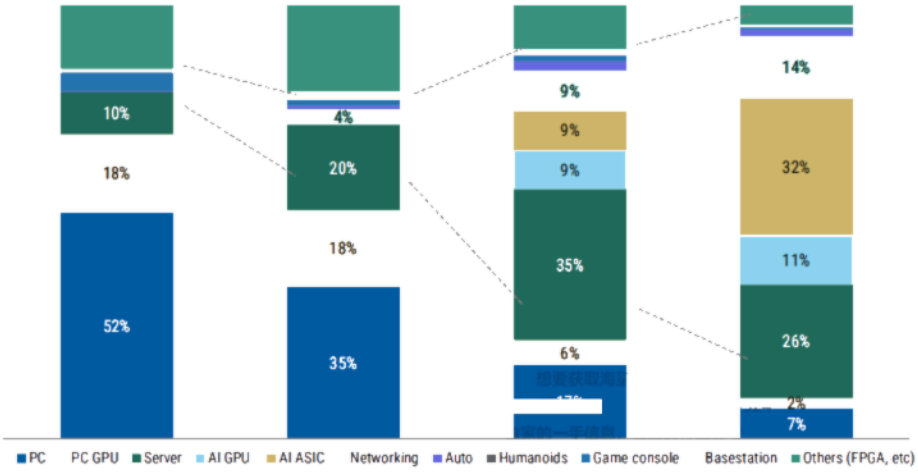
Period VI (2024-2026): Bifurcated recovery	PC demand normalized at the pre-Covid levels, with some replacement demand driven by Windows upgrade, while the recovery and growth of AI and general server demand for cloud picks up at a meaningfully higher growth trajectory.
Period VII (2026-30): AI diffusion	Demand for AI proliferates from training to inference, driving ABF substrate market growth to re-accelerate over the next five years after a two-year down-cycle following the Covid-19 pandemic. This drives increased demand for AI computing, networking, and cloud computing, while demand for PCs remains soft.

Exhibit 53: Global ABF Substrate Market by End Application



Source: Morgan Stanley Research estimates Note: unit: sales (M).

Exhibit 54: Server/AI/Networking/HPC Applications expected to rise to 75%+ by 2030, up from ~10% in 2015



Source: Morgan Stanley Research estimates

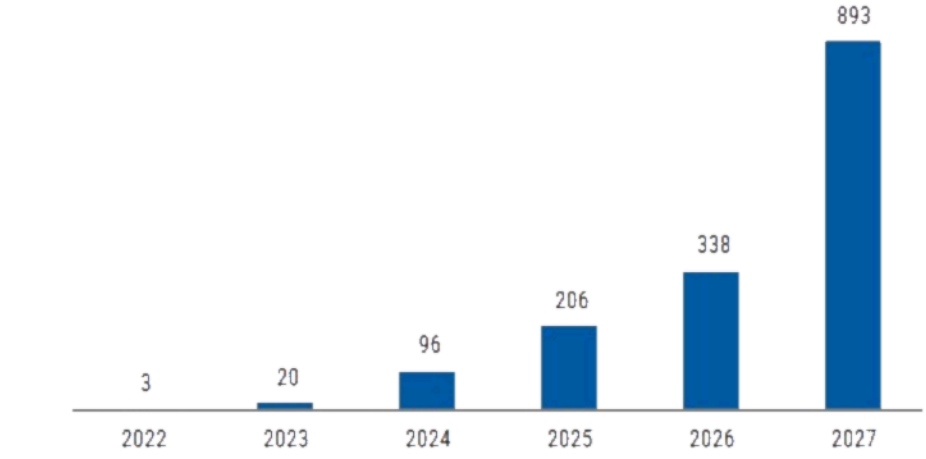
MLCC

MLCC is another "AI squeeze" story – there has been fear that midrange/low-end MLCCs will face undersupply because leading Japanese and Korean MLCC suppliers are shifting more capacity to cater to the high-end, high-capacitance MLCCs used in AI infrastructure hardware. See our research on AI MLCC opportunities here, here, and here.

At the stock level, we like Yageo (2327.TW), covered by Howard Kao, SEMCO (009150.KS), covered by Shawn Kim and Murata (6981.T), covered by Shoji Sato. In terms

of exposure to midrange/low-end, Yageo has the highest exposure to those segments amongst this trio.

Exhibit 55: By 2027, AI server MLCC demand will likely already reach close to US$1bn, tracking meaningfully above our previous forecast of ~US$550mn by 2030



Source: Company data, Morgan Stanley Research estimates

AI MLCC demand is growing. We estimate MLCC dollar content will rise ~182% for VR200 NVL72 vs. GB300 NVL72. VR200 MLCC content is ~US$4.3K vs. ~US$1.5K for GB300 — a jump that likely explains the current strength in high-end AI server MLCC demand and ODMs' aggressive inventory build ahead of the Rubin rack ramp from 2H26. Content per computing board and switch board is rising meaningfully, with computing boards seeing the larger increase. New BlueField and ConnectX modules add further MLCC content per rack.

Total MLCC unit demand is rising ~80% for a VR200 NVL72 rack vs. GB300. We estimate ~570K MLCC units per Rubin NVL72 rack vs. ~320K for GB300, driven by higher content per computing/switch board plus added PCB content (e.g., the CX9 Orchid Module). This unit growth is a main driver of the 182% dollar-content increase.

The remaining ~100% growth comes from a richer mix of high-cap MLCCs. High-cap MLCCs (47uF+) are rising to 30%+ of the mix in Rubin systems, vs. <20% for GB300 (Exhibit 3). By our estimate, 47uF+ MLCC unit demand rises ~210% for VR200 vs. GB300, while demand for 47uF-and-below content rises only ~50%.

Exhibit 56: VR200 MLCC content is 180%+ higher compared to GB300

MLCC content per board (US$)	GB300	VR200
Compute PCB	$25	$90
Switch PCB	$20	$45
BlueField DPU Module	$5	$5
ConnectX Orchid Module	$5	$5
Other peripheral PCB	$5	$5
Units per rack	**GB300**	**VR200**
Compute PCB	36x	36x
Switch PCB	9x	9x
BlueField DPU Module	0x	18x
ConnectX Orchid Module	0x	72x
Other peripheral PCB	90x	45x
Total MLCC Content per rack	**GB300**	**VR200**
Compute PCB	$900	$3,240
Switch PCB	$180	$405
BlueField DPU Module	$0	$90
ConnectX Orchid Module	$0	$360
Other peripheral PCB	$450	$225
Total MLCC Content per rack	**$1,530**	**$4,320**

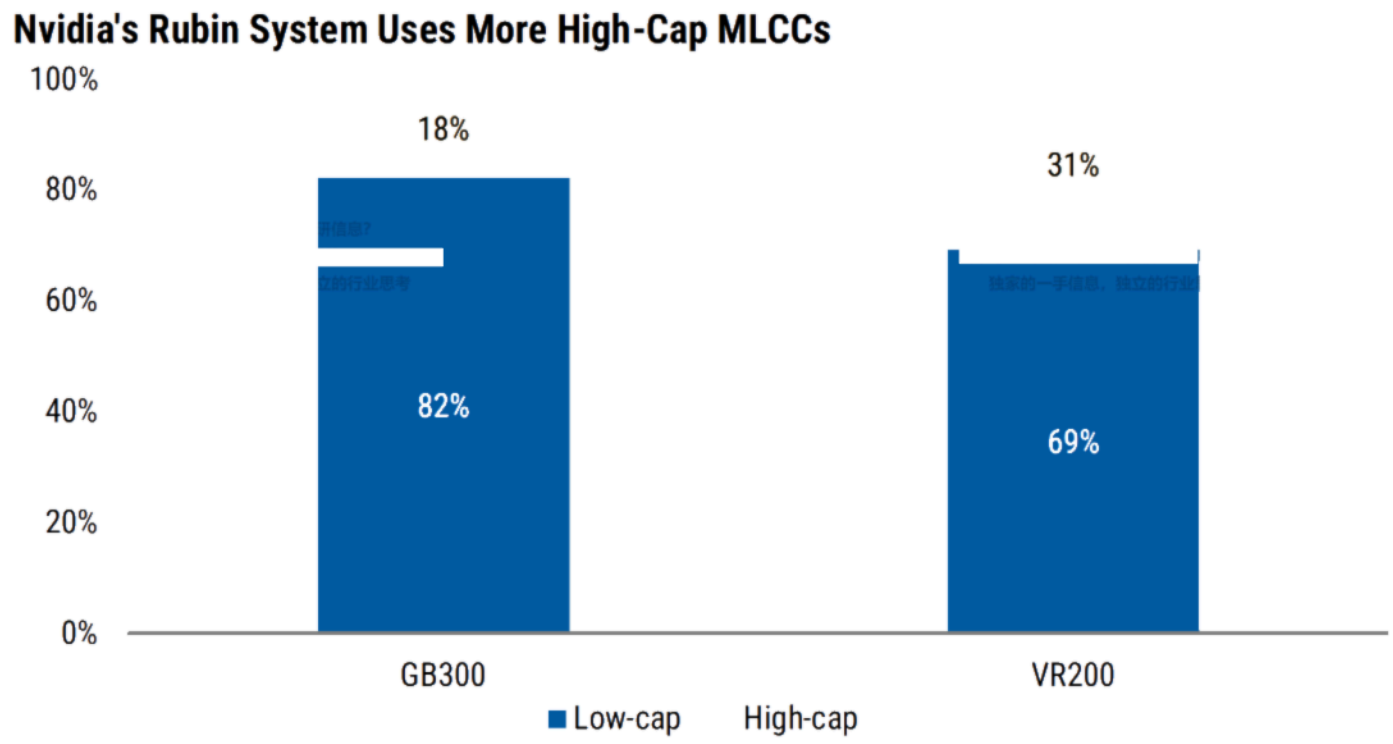
Source: Morgan Stanley Asia Pacific Research estimates based on supply chain checks

Exhibit 57: VR200 MLCC unit demand is rising to 570K+, up close to 80% vs. a GB300 rack

Server Assumptions		# of MLCCs		US$	
General server			**2,000**		**30**
HGX Hopper AI server (8x GPU)			**22,000**		**190**
GPU module (OAM)	8x	1,300	10,400	10	80
GPU baseboard (UBB)	1x	4,900	4,900	45	45
CPU motherboard	1x	1,500	1,500	25	25
Peripheral boards	8x	650	5,200	5	40
HGX B200 AI server (8x GPU)			**36,000**		**317**
GPU module (OAM)	8x	3,000	24,000	24	192
GPU baseboard (UBB)	1x	5,300	5,300	60	60
CPU motherboard	1x	1,500	1,500	25	25
Peripheral boards	8x	650	5,200	5	40
HGX B300 AI server (8x GPU)			**48,000**		**435**
GPU module (OAM)	8x	4,400	35,200	40	320
GPU baseboard (UBB)	1x	6,100	6,100	50	50
CPU motherboard	1x	1,500	1,500	25	25
Peripheral boards	8x	650	5,200	5	40
GB200 (NVL72)			**337,500**		**1,710**
Bianca	36x	6,800	244,800	30	1,080
Switch board	9x	3,800	34,200	20	180
Peripheral boards	90x	650	58,500	5	450
GB300 (NVL72)			**319,500**		**1,530**
Bianca	36x	6,300	226,800	25	900
Switch board	9x	3,800	34,200	20	180
Peripheral boards	90x	650	58,500	5	450
VR200 (NVL72)			**571,050**		**4,320**
Strata	36x	11,000	396,000	90	3,240
Switch board	9x	7,000	63,000	45	405
BlueField Module	18x	2,000	36,000	5	90
CX9 Orchid Module	72x	650	46,800	5	360
Peripheral boards	45x	650	29,250	5	225
Vera Rack			**494,048**		**3,060**
Vera MGX	128x	3,316	424,448	20	2,560
Switch board	4x	7,000	28,000	45	180
Peripheral boards	64x	650	41,600	5	320
STX Rack			**326,256**		**2,415**
Mainboard	16x	4,791	76,656	31	495
Peripheral boards	384x	650	249,600	5	1,920
LPX Rack			**581,024**		**4,326**
LPX UBB	16x	5,374	85,984	35	559
LPX Module	256x	1,705	436,480	14	3,527
BF4	16x	2,000	32,000	5	80
x86 CPU Module	16x	1,660	26,560	10	160

Source: Morgan Stanley Asia Pacific Research estimates based on supply chain checks

Exhibit 58: High-cap (47uF+) MLCC usage to increase to 30%+ in VR200 (vs <20% for GB300)



Source: Morgan Stanley Asia Pacific Research estimates based on supply chain checks. Note: High-cap refers to 47uF + MLCCs

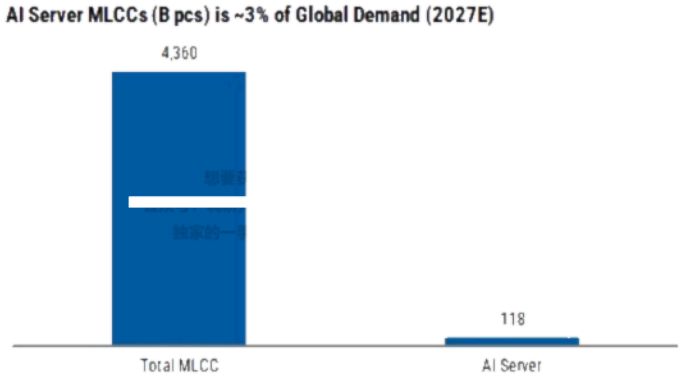
According to our colleague Shoji Sato, the global MLCC market is expected to reach approximately 4.36 trillion units and US$18.3 billion in value by CY27. To put the AI MLCC opportunity into context, AI servers would account for roughly 3% of total MLCC unit demand and approximately 5% of industry revenue by that time, based on our estimates.

Although these percentages may appear modest in the context of the overall MLCC market, we believe they understate the significance of the AI opportunity. AI server demand has emerged from a near-zero base just five years ago and is growing at a pace well above the broader industry. More importantly, the impact extends beyond its direct contribution to market size and increasingly influences industry capacity allocation.

In our view, the more important consideration is the pressure AI demand places on manufacturing capacity and the potential crowding-out effect on other MLCC categories. High-capacitance MLCCs used in AI infrastructure are substantially more manufacturing-intensive than mainstream MLCC products, often requiring several hundred to well over one thousand dielectric layers depending on capacitance, voltage rating, and package size. By comparison, many commodity MLCCs used in consumer electronics utilize materially fewer layers. As a result, each high-capacitance MLCC consumes a disproportionate amount of manufacturing capacity relative to its unit volume.

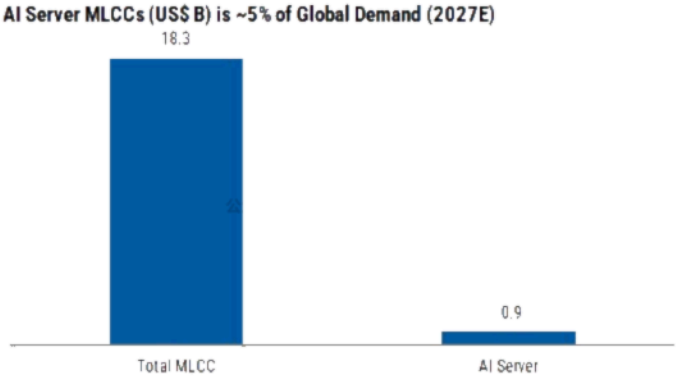
Consequently, even if AI servers represent only a mid-single-digit percentage of industry revenue, the associated capacity consumption is likely substantially higher. As AI-related demand continues to accelerate, suppliers may increasingly reallocate production capacity toward higher-value products, potentially tightening supply conditions and supporting pricing across broader segments of the MLCC market.

Exhibit 59: We estimate that AI server MLCC demand accounts for ~3% of total MLCC demand, in units



Source: Morgan Stanley Research (E) estimates

Exhibit 60: We estimate that AI server MLCC demand accounts for ~5% of total MLCC demand, in value



Source: Morgan Stanley Research (E) estimates

From a bottom-up perspective, since late 2Q, we have also observed:

- **Increasing supplier utilization rates:** Yageo attended our Asia AI Summit in late May, and mentioned that 2Q utilization for all standard and premium products is running higher than its guidance during the 1Q earnings call (standard products at 75%+ and premium at 85%+). Yageo Chairman Pierre Chen even mentioned at Yageo's AGM that utilization rates will reach 80%+ for standard products and 90% + for premium, though he did not mention a specific time for when this will happen.
- **Price hikes for the midrange/low-end products:** Our latest supply chain checks indicate stronger pricing momentum in 2H26. We now expect direct-customer pricing to increase 30-40% on average in 2H26, vs. our previous estimate of 10–20%.
- **Distributors and ODMs aggressively building up inventories:** Our checks indicate that both distributors and ODMs are asking for more MLCC and buying as much as they can to build inventory. The distributors are scared of supply shortage, while the ODMs are trying to make sure that MLCC does not become the next bottleneck in the supply chain.

China Equity Strategy – State-Backed Infra Roadmap Driving Growth for Economy and Earnings

Laura Wang, Vicky Wu

Policy Framework Evolution: from Blueprint to Operationalization and Globalization, Aiming for Infra and Public Utility access

China's AI policy architecture has evolved through three distinct stages, each building on the prior, with infrastructure consistently being at the core.

The foundational blueprint was the State Council's "New Generation AI Development Plan" in 2017: State Council Document [2017] No. 35, issued on 20 July 2017, set three cascading milestones:

1. Parity with global peers by 2020,
2. World-leading position in select AI fields by 2025, and
3. Becoming the primary global AI innovation center by 2030, with a core AI industry TAM target set for Rmb1tn and related industries of Rmn10tn.

Between 2021 and 2024, the "East Data, West Computing" strategy operationalized the compute-infrastructure side, routing hyperscale build-out to western provinces with cheaper power while preserving Tier 1 city latency hubs — by end-2025, China's installed data center capacity reached ~37 GW, ~27% of global capacity.

Since day one, the Chinese authorities had tried to install a regulatory framework. The regulatory scaffold was set in 2023: the Interim Measures for the Management of Generative AI Services were jointly issued by the CAC, NDRC, MOE, MOST, MIIT, MPS, and NRTA on 10 July 2023, effective 15 August 2023, adopting a "development-and-security balanced," classified and graded supervision approach (softer than the April 2023 draft), covering training data legality, algorithm filing, content labeling. and security assessments for public-facing services.

The second stage started in mid-2025 and reframed AI as an operating system for the whole economy: On August 26 2025, the State Council issued Guofa [2025] No. 11 — "Opinions on Deeply Implementing the AI Plus Initiative" — setting concrete penetration targets:

- Intelligent terminals and AI agents >70% by 2027
- >90% by 2030, and
- *Full transition into an "intelligent economy and society" by 2035,* prioritizing six sectors (science, industry, consumption, livelihood, governance, global cooperation).

This was reinforced by draft ethics rules (Administrative Measures for Ethical Management of AI, trial) and the AI-Generated Content Labeling Measures released in the same window. Funding was material:

- A ~ national AI fund of Rmb60bn (~US$ 8.2bn).
- ~Rmb138bn in local venture guidance funds, and
- Hangzhou alone earmarking up to Rmb300bn/year for its "Six Little Dragons" cluster (which incubated DeepSeek, Moonshot, Zhipu/Knowledge Atlas, MiniMax, and peers).

The third stage — infrastructure elevation — is the most consequential for the next 5-10 years: On 28 April 2026, the Politburo elevated the *"Six Networks" — computing, water, power, communications, urban underground, logistics* — to national strategic-infrastructure tier, *effectively classifying AI computing as a public utility on par with water and electricity.*

The State Council made this operational on 9 May 2026, and the NDRC confirmed that >Rmb35tn of Six Networks spending would be concentrated in the 15th Five-Year Plan (2026-2030), with >Rmb1tn marked down for the computing network in 2026 alone. Bloomberg reported on 9 June 2026 the drafting of a nationwide interconnected AI data center blueprint worth ~Rmb2tn (~US$295 bn) over five years, operated primarily by China Mobile and China Telecom, with a hard ≥80% domestic-hardware mandate (Huawei, Cambricon, etc.), sourced from sovereign debt, state-backed vehicles, banks, and hyperscalers like Alibaba.

On the AI agent governance front, the CAC/NDRC/MIIT "Implementation Opinions on Standardized Application and Innovative Development of Intelligent Agents" (8 May 2026) became the world's first dedicated national framework treating agents as a distinct regulated class, defining three decision-authority tiers (user-only, user-authorized, autonomous) with 19 priority scenarios. Most recently, the Interim Measures for Anthropomorphic AI Interaction Services took effect 15 July 2026, requiring distress detection, crisis intervention and use-time curbs on companion/chatbot services, enforced by the CAC with penalties up to Rmb200,000.

Looking back at the multi-decade timeline set out by the Chinese policymakers, China has actually largely been hitting most major milestones as planned: Despite a temporary setback on parity with global peers (the US) by 2020, select leadership by 2025 (the DeepSeek moment) was achieved. The country looks to be well on track to become a primary innovation center by 2030 (continuously spearheading technology self-sufficiency and striking a balance between model performance and cost, providing affordable solutions for mass market usage).

Exhibit 61: China's AI policy framework evolution – from blueprint to operationalization and globalization

Date	Policy / Event	Issuing Body	Key Initiatives
Stage 1 (2017 - 2024): Foundational blueprint & regulatory guidance			
Jul-17	**New Generation AI Development Plan** (State Council [2017] No. 35)	State Council	3-stage roadmap: parity by 2020 + world-leading in select fields by 2025 + global AI innovation hub by 2030; RMB 1 trn core AI industry + RMB 10 trn related industries target
2021–2024	**"East Data, West Compute"** rollout	NDRC + 3 ministries	8 national compute hubs; routes hyperscale build-out to western provinces (cheap power) while preserving Tier-1 latency clusters
10 Jul 2023 (eff. 15 Aug 2023)	**Interim Measures for Management of Generative AI Services**	CAC + NDRC, MOE, MOST, MIIT, MPS, NRTA	Classified/graded supervision; training-data legality; algorithm filing; content labelling; security assessments for public-facing services
Stage 2 (2024-2025): Whole economy operating system "Intelligent economy and society"			
Jul-25	**Global AI Governance Action Plan** (World AI Conference, Shanghai)	MFA / State Council	6 principles. innovation, infrastructure, open ecosystems, high-quality data, green AI, multi-stakeholder governance
Aug-25	Draft **Ethics Measures (Trial) + AI-Generated Content Labelling Measures**	CAC	Ethics review regime; mandatory watermarking/labelling of synthetic content
Aug-25	**"AI Plus" Initiative Opinions** (Guofa [2025] No. 11)	State Council	Penetration targets: intelligent terminals & AI agents >70% by 2027, >90% by 2030; 6 priority sectors (science. industry. consumption. livelihood. governance. global cooperation); full "intelligent society" by 2035
2025	**National AI Fund + Local Venture Guidance Funds**	Central + local govts	RMB 60 bn (~USD 8.2 bn) national fund, RMB 138 bn local, Hangzhou up to RMB 300 bn/yr for "Six Little Dragons" (DeepSeek, Zhipu, MiniMax, Moonshot, etc.)
Stage 3 (2026 onwards): AI compute as public utility			
Apr-26	**Politburo "Six Networks" (六张网) Directive**	Politburo	Elevates compute network to national strategic infrastructure — same tier as water. power, logistics. communications, urban underground
May-26	**Implementation Opinions on Intelligent Agents**	CAC + NDRC + MIIT	World's first dedicated AI-agent framework; 3 decision-authority tiers (user-only / user-authorized / autonomous); 19 priority scenarios; open-source & domestic-chip compatibility
May-26	**State Council operationalizes Six Networks; NDRC briefing**	State Council / NDRC	>RMB 35 trn total Six Networks capex in 15th FYP (2026–2030); >RMB 1 trn compute network in 2026 alone
Jun-26	**Nationwide AI Data-Center Network Blueprint** (Bloomberg-reported)	NDRC + MIIT (drafting)	~RMB 2 trn (~USD 295 bn) over 5 years; operated by China Mobile / China Telecom; **≥80% domestic hardware mandate** (Huawei et al.)
Jul-26	**Interim Measures for Anthropomorphic AI Interaction Services**	CAC	Mandatory emotional distress detection, crisis intervention, addictive-use curbs; bans training on companion sensitive data; fines up to RMB 200,000
In parallel, China's initiatives to expand its AI footprint globally			
Oct-23	**Global AI Governance Initiative**	State Council / MFA	People-centered dev; AI-for-good; opposes tech monopolies & unilateral barriers, elevates developing-country voice
26-Jul-25	**Global AI Governance Action Plan** (WAIC 2025)	MFA / State Council (WAIC Shanghai)	6 themes: innovation, infra. open ecosystems, data, green AI, multi-stakeholder governance; Premier Li Qiang proposes idea for a China-led AI intl org
16-Jul-26	**WAICO Founding Agreement Signed** (World Artificial Intelligence Cooperation Organization)	29 founding countries. Shanghai	HQ in Shanghai; membership: Africa (10), ASEAN & Central Asia (11), Europe (Belarus, Russia, Serbia), LatAm (Brazil, Cuba, Nicaragua, Venezuela)
17-Jul-26	**President Xi's Global South Pledges**	WAIC 2026 keynote	**5,000 AI training/seminar spots** for developing countries over 5 yrs; AI cooperation platforms with **ASEAN, Arab League, African Union, BRICS**; AI weather-warning system extended to **30 countries**

Source: Morgan Stanley Research

China's AI diplomacy track – China-standard AI ecosystem runs in parallel with the US framework

Alongside the domestic three-wave build-out, China has run a parallel, policy-guided global footprint expansion that has quietly become the fourth pillar of the framework: It began with the Global AI Governance Initiative (October 2023) — Beijing's first coherent international AI pitch — and matured through the Global AI Governance Action Plan at WAIC 2025, which positioned *China as norm-setter around six themes: innovation, infrastructure, open ecosystems, high-quality data, green AI, and multi-stakeholder governance.*

Open source and economical access for the mass market vs. the US: China's strategy is deliberately asymmetric to the US approach:

- Washington leans on export controls and closed-source frontier models along with the "Pax Silica" coalition (~35 countries).
- Beijing leans on open-source models, subsidized computing access, and capacity building for the Global South, framing AI as an "international public good" rather than a strategic asset to be fenced.

The economic backdrop is real: according to the speech delivered by President Xi at the latest 2026 WAIC, China's core AI industry gross revenue reached Rmb1.2tn (~US $176.6bn) in 2025 with >6,200 AI enterprises, and its open-source models have been downloaded >10 billion times cumulatively.

Moreover, the latest OpenRouter data show that Chinese AI models now contribute roughly 18 trillion tokens as of June 2026, versus around 5.5 trillion tokens by the US models. Chinese models' high performance-to-price ratio is a key driver (*note: OpenRouter tracks ~2% of global token consumption*).

The latest WAIC marks the moment that AI moved from industry-ministry portfolio to head of state diplomatic tool: WAIC 2026 is the ninth edition and the first time President Xi Jinping attended in person and delivered the keynote since the event's 2018 inauguration. The conference produced the World AI Cooperation Organization (WAICO) — the world's first intergovernmental body dedicated exclusively to AI, headquartered in Shanghai, signed by 29 founding members on 16 July (largely Global South: Africa, ASEAN, Central Asia, plus Russia/Belarus/Serbia and Brazil/Cuba/Venezuela).

China also made concrete pledges — 5,000 AI training spots for developing countries over five years, AI cooperation platforms with ASEAN / Arab League / African Union / BRICS, and an AI-powered weather warning system extended to 30 countries — signalling that Beijing can build AI infrastructure alone as well as govern it internationally.

WAICO's Shanghai headquarters and Global South membership base are designed to institutionalize a China-centric AI standards ecosystem that runs in parallel with the US/G7 framework — effectively creating an "AI Belt and Road" that pairs domestic computing overcapacity (from the Six Networks build-out) with subsidized cloud access, open-source model exports, and capacity-building diplomacy across Africa, Southeast Asia, Central Asia, the Middle East, and Latin America.

Domestic and Global AI Blueprint Warrants Continuous AI Infrastructure Spending on Computing, Semi Self-sufficiency, Deployment and More

As China aims to achieve a China-centric Global South ecosystem with subsidized AI accessibility, a combination of policy, regulation, and capital will be necessary to ensure continuous capex expansion and infrastructure-related building, not only domestically but also on a global basis.

Supply-side constraints being addressed with policy and funding support

1. Computing abundance to become a state-guaranteed utility. With Rmb2tn planned over five years and 2026 and 2027 hyperscaler capex projected at US$85bn (+40% YoY) and US$100bn (+20% YoY), respectively (by Morgan Stanley's China Internet & Tech team). Revision risk is skewed towards the upside as Chinese hyperscaler capex is only a fraction of US peers' — China's cloud capex is ~11–14% of US hyperscalers' level as of 2026E, and Chinese hyperscalers' 2027 capex of ~Rmb706bn (~US$101bn) would account for only ~12% of US hyperscalers' capex. Alibaba and Tencent's capex intensity (~15%) still lags vs. Google and Meta (40–50%) and Amazon and Microsoft (~20%).
2. Model and application deployment to lead the US system by volume, not frontier

capability. China's daily AI token consumption reached ~140tn in March 2026 (from ~100bn in early 2024), and four of OpenRouter's top-five most-used models globally in February 2026 were Chinese (DeepSeek, Qwen, Kimi, GLM), reflecting the "AI Plus" penetration mandate at work. This will continue to drive global demand for Chinese models and data centers in addition to related tech hardware and power related equipment. Chinese memory companies are also expected to become a powerful global supplier/competitors on a three- to five-year basis, especially with the committed capacity expansion by companies like CXMT and YMTC, etc.

3. Embedded self-reliance bottlenecks warrant continuous R&D and capex investment — HBM memory, EUV lithography, CUDA-equivalent software ecosystems, and advanced packaging remain the binding constraints. We expect USD70bn+ in additional subsidies for domestic chip fabrication and an accelerating open-source push (DeepSeek, Qwen, GLM) as Beijing's asymmetric response to the US CUDA/hyperscaler stack.

To provide long-term support for the aforementioned choke points, the government has taken multiple initiatives to support funding:

- *National Big Fund:*
  - Big Fund Phase III plus six major state-owned banks: US$7.8bn capital injection into SMIC South — the single largest government-backed semiconductor funding event documented in 2026.
  - Big Fund Phase I and the Beijing Integrated Circuit Fund previously held 49% of SMIC North, which is now being consolidated into SMIC's wholly owned structure — a governance simplification that improves capital allocation efficiency.
- *STAR Market IPOs as capital formation mechanism:* The STAR board on Shanghai Stock Exchange is functioning as the primary capital formation vehicle for China's semiconductor self-reliance drive. CXMT's Rmb29.5bn IPO and YMTC's pre-IPO coaching represent the two most significant capital market events in China's semiconductor sector in 2026.
- *Local government co-investment:* CXMT's IPO prospectus explicitly notes that "the company could also cooperate with local government to expand capacity" — indicating that local government co-investment remains a key supplementary funding mechanism alongside the STAR Market IPO proceeds.

Sovereign and SOE-driven demand: the government as anchor customer

Beyond supply-side support, the Chinese government is acting as a major demand-side anchor for domestic AI chips through sovereign and SOE procurement mandates:

- *State-funded data center projects are required to use only domestic AI chips:* Reuters reported in early November 2025 that China required new state-funded data center projects to use only domestic AI chips, and that projects less than 30% complete were ordered to remove installed foreign chips. Post-APEC, China's chip self-sufficiency drive has become more coercive on the demand side and more structured on the planning side, including reported domestic chip requirements for state-backed AI infrastructure, reported localization thresholds for new fab

equipment, and explicit 2026-2030 national planning around core-technology self-reliance etc.

- *Sovereign and SOE TAM raised to US$9bn by 2030 (from US$7bn),* reflecting government advocacy to increase spending on AI infrastructure. Morgan Stanley raised its projection of sovereign and SOE-related TAM in view of government advocacy to increase spending on AI infrastructure, pushing total Chinese AI chip TAM to US$91bn by 2030.
- *Rmb2tn over five years for data center construction:* As discussed above, Bloomberg reported that China is preparing Rmb2tn over the next five years for building data centers across the country.

Five-Year Plan and US export control to keep capex, R&D, funding, and policy support in the long run as self sufficiency ratio goes up

The 15th Five-Year Plan reaffirms self-reliance and economic security, with persistent geopolitical tension and supply chain vulnerability underpinning these priorities, alongside a growing emphasis on quality, sustainability, and resilience of growth. Meanwhile, export controls from the US side will continue to serve as an accelerant. These constraints have materially sped up China's efforts to build a fully domestic semiconductor supply chain, with leading-node foundry capacity emerging as the key bottleneck. A symbiotic relationship is forming between domestic AI chip design houses and local leading-node foundries.

As a result, Morgan Stanley expects China's AI chip self-sufficiency to rise from 42% in 2025 to 70% in 2030, underpinned by capacity build-out at advanced nodes, steady yield improvements, and continued government support. A structural 30% gap will remain, reflecting the remaining fundamental technology constraints in HBM, EUV, and advanced software ecosystems.

State-backed AI Spending as Crucial GDP Driver and Investment Opportunities – How to Position

China's AI infrastructure spending trajectory is one of the most powerful structural forces reshaping the country's growth trajectory as well as investment landscape.

AI to become China's key productivity engine over the medium term... For 2026-27, the primary growth impulse from AI stems from the capex expansion cycle in computing infrastructure, data centers, and energy, estimated to add 0.2-0.3ppt to annual real GDP growth. Over the longer term, AI could lift China's total factor productivity growth approximately 3ppt cumulatively over the next decade, potentially raising the level of potential GDP ~3.5ppt above a no-AI trajectory by 2035.

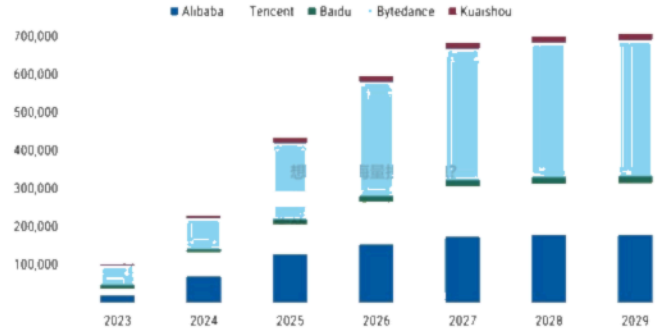
The 15th Five-Year Plan has institutionalized "AI+" as a national productivity strategy, explicitly linking digital upgrading to the target of maintaining faster labor productivity growth than GDP expansion — providing a durable policy anchor for AI infrastructure investment.

Exhibit 62: AI's impact on China's GDP growth: near-term neutral, long-term positive



Source: CEIC, Morgan Stanley Research estimates

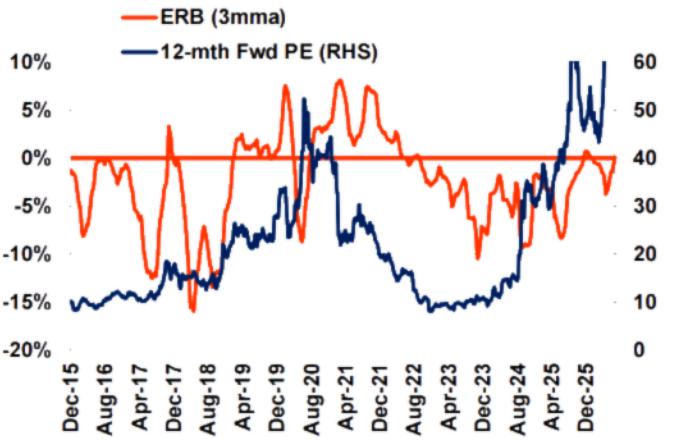
Exhibit 63: Hyperscalers to step up AI-related capex significantly in the next two years



Source: Company data, Morgan Stanley Research estimates

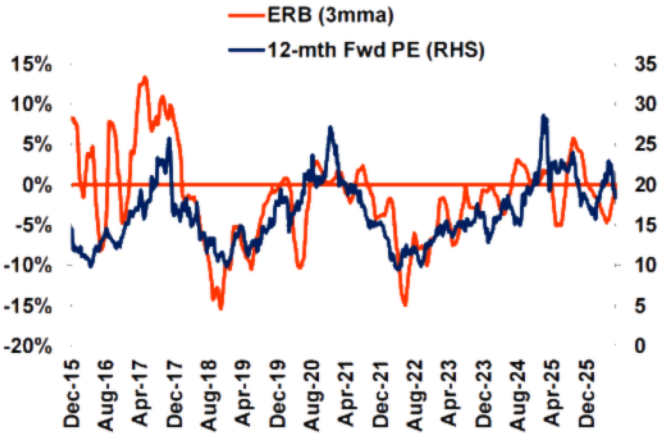
...while it has already been driving upside surprises for the Chinese equity market's earnings growth. As discussed in our 2Q earnings preview report, AI-related sectors (Semis, Tech Hardware) are among the few sectors that have released more net positive alerts,. They also exhibit positive earnings estimate revision breadth trend while the overall MSCI China ERB remains under water.

Exhibit 64: Semi – Earnings estimate revision breadth (ERB) vs. 12-month P/E trend



Source: IBES Rimes, MSCI, Morgan Stanley Research. Weekly data as of July 22, 2025. Earnings Estimate Revision Breadth (ERB) = (number of upward revisions - number of downward revisions) / total number of estimates

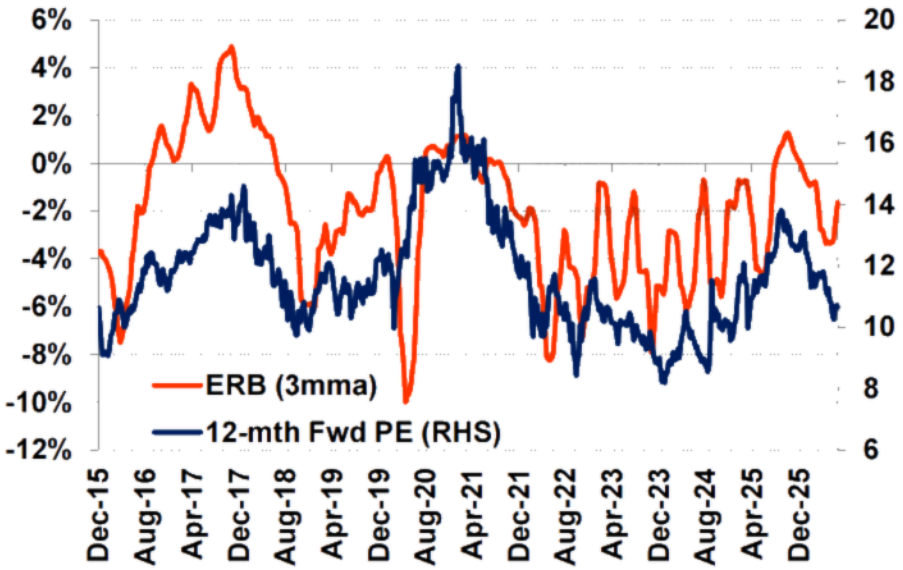
Exhibit 65: Tech Hardware – Earnings estimate revision breadth (ERB) vs. 12-month P/E trend



Source: IBES Rimes, MSCI, Morgan Stanley Research. Weekly data as of July 22, 2025. Earnings Estimate Revision Breadth (ERB) = (number of upward revisions - number of downward revisions) / total number of estimates

Exhibit 66:

MSCI China earnings estimate revision breadth (ERB) vs. 12-month P/E trend – ERB saw some improvement since July

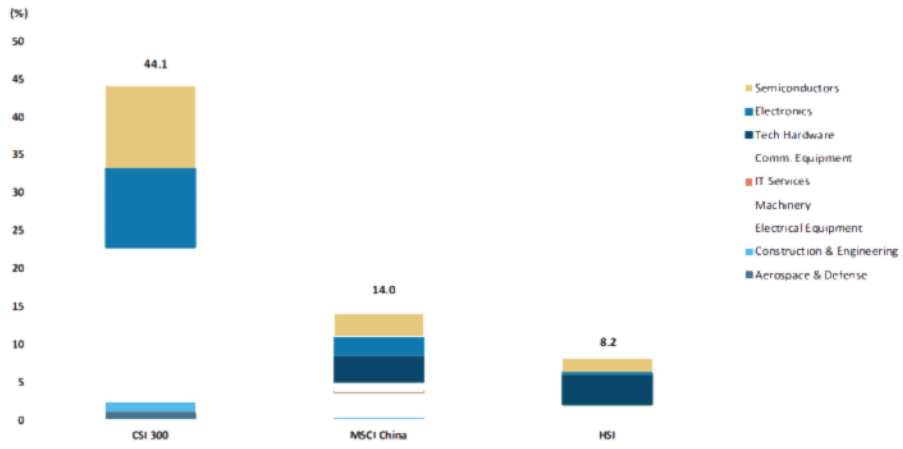


Source: IBES, Refinitiv, MSCI, Morgan Stanley Research. Weekly data as of July 22, 2026. Earnings Estimate Revision Breadth (ERB) = (number of upward revisions - number of downward revisions) / total number of estimates

We strongly prefer the onshore A-share market over the long run: As the Chinese economy continues to exhibit a K-shaped trajectory in the next couple of years, we believe allocating to the sources of earnings growth certainty is the right, sustainable strategy. The onshore A-share market is a much more real-time reflection of China's economic growth model transition (from property/infrastructure and consumption more towards tech and innovation as well as high-end manufacturing). These so-called "New Quality Productivity" sectors represent 44% of A-shares' CSI 300 index, compared to a mere 14% in MSCI China, providing far greater concentration in high-end upstream manufacturing and hard-core tech companies that are the primary beneficiaries of the AI capex cycle.

Select exposure to the hyperscalers and LLMs via Hong Kong is also crucial: We discuss this below.

Exhibit 67: Snapshot of high-end upstream manufacturing and hard-core tech companies – much bigger representation in the A-share market than in MSCI China



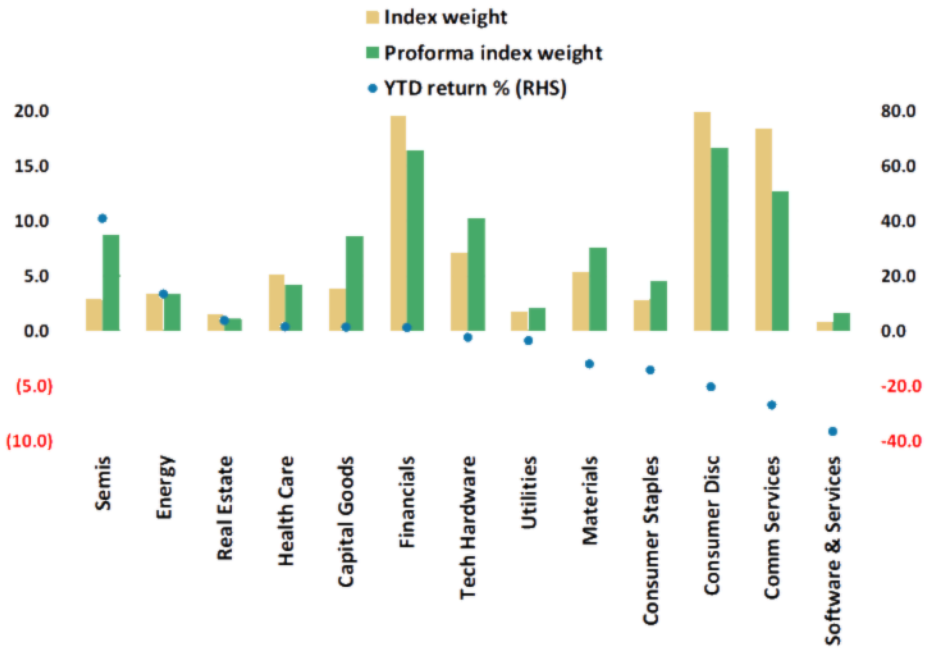
Source: FactSet, Morgan Stanley Research

Thematic stock-picking over passive index investing in view of index compilation rule limitation: The sectors best positioned to benefit from AI are underrepresented in MSCI China and other global indices because of composition rule limits:

1. The US government's NS-CMIC list of Chinese companies particularly affects Information Technology, Industrials/Capital Goods, Energy as well as Materials, which are some of China's best performers YTD and also directly integrated as part of the AI/Energy capex super-cycle.
2. Chinese A-shares have an inclusion factor for MSCI index weight calculations for qualified market cap of merely 20%.

Our analysis shows that these sectors collectively could gain 19.5ppt index weight within MSCI China if these two hurdles were removed. Therefore, picking the right themes with the higher certainty of fast earnings growth would be much more effective, vs. passively investing only at index level, especially the offshore indices. Themes along the infrastructure and power super-cycle value chain recommended in this report by the sector teams fit such criteria perfectly.

Exhibit 68: MSCI China's pro forma index weight by sector – Semis, Capital Goods, Tech Hardware, Materials and Energy would schow the biggest rise in weighting



Source: US government websites, Factset, Morgan Stanley Research. Note: Data as of July 23, 2026

Hyperscalers and LLMs – exclusive opportunity sets only available in Hong Kong now: Despite the overarching concentration of AI and technology value chain in the A-share market, there are two groups of very relevant components that are entirely missing in the A-share market right now – the hyperscalers and the LLM companies.

The hyperscalers are Tencent and Alibaba: dominant Internet players investing heavily in the AI value chain, with diversified business lines progressing in LLMs, chips, cloud services, data centers and more. We have been advising investors to add back such exposure since summer this year, seeing better earnings growth trajectory as well as more

AI progress from here. (For details, please see: China Market-Wise: China Outperforming in a Volatile Market – What's Next?).

On the other hand, the two listed LLM plays – Z.AI (Knowledge Atlas) and MiniMax – are both listed in the Hong Kong market, having just made it into the Hang Seng and Hang Seng Tech Indices in June this year. The earliest they could be added to the MSCI China Index would be August this year. As the market gradually digests their IPO share unlocking pressure (unlocking happened in early July 2026), we believe more inflows would be attracted to the unique Chinese LLM exposure they represent.

LatAm Equity Strategy – Minerals, Energy and Hardware

Nikolaj Lippmann, Julia Nogueira

We have never believed that LatAm was going to lead AI development, yet many investors have adopted a glass-half-empty view, effectively giving up on the AI theme in the region. If the objective is to identify the next frontier AI lab, we agree that Latin America is not the place to look. However, this misses where the investment opportunity lies.

We are already seeing a shift in equity alpha away from domestic and regional investors toward global AI-focused investors across a range of sectors, from copper mining to utilities. While US investors have developed a compelling narrative around data centers, power demand, and enabling infrastructure – and are now beginning to grapple with the practical constraints on that buildout – LatAm investors have often dismissed the opportunity altogether, arguing that the region is unlikely to meet even its own data center needs because the obstacles are too great.

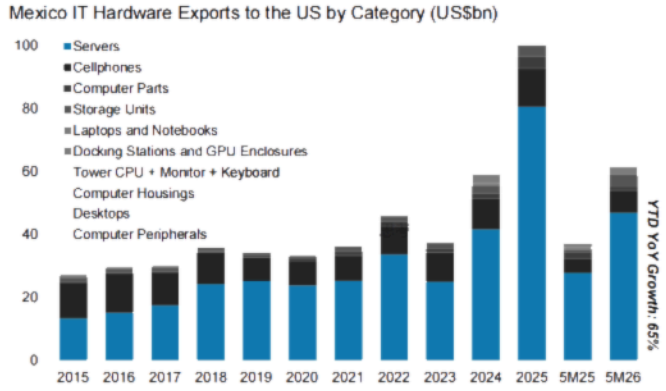
We challenge this pessimistic view, while acknowledging that the investment case will take time to play out. Latin America's role in the AI value chain is already clear as a supplier of critical minerals, manufacturing capacity (particularly in Mexico), and increasingly energy. Looking further ahead, we believe electricity – and eventually data infrastructure itself – could become attractive sources of value creation, as the region leverages its abundant hydro and natural gas resources to convert low-cost energy into higher-value digital infrastructure.

The role of LatAm in the AI supply chain. AI is highly relevant to Latin America, even if the region lacks a globally significant technology sector. Latin America produces around 10% of the world's oil, holds roughly 20% of global energy reserves, accounts for 30% of global copper production, and possesses around 60% of the world's lithium reserves. As a result, the global AI capital expenditure cycle has become an important catalyst for the region, influencing commodity prices, terms of trade, exchange rates, and equity markets.

Beyond its role as a supplier of critical inputs, Latin America also offers significant long-term optionality as AI reshapes global energy demand. Brazil's abundant hydro resources and Argentina's vast natural gas reserves could support the development of electricity-intensive data centers, creating new avenues for value creation as digital infrastructure expands.

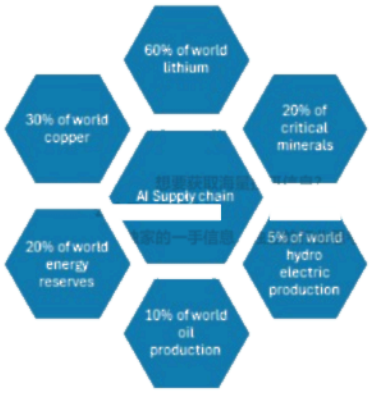
Mexico provides a complementary exposure. Beyond being home to Grupo México, the parent company of Southern Copper (one of the world's largest publicly listed copper producers) its strategic importance to AI increasingly stems from its role as a manufacturing and hardware hub. Mexico has become a key exporter of servers and other computing equipment to the United States, a position that is likely to strengthen further as supply chains continue to realign under US strategic decoupling. - listen to How AI Will Power LatAm; Opportunities & Risks.

Exhibit 69: Mexico is increasingly a supplier of IT Hardware to the US...



Source: UN Comtrade and Morgan Stanley Research

Exhibit 70: ...while LatAm is a supplier of minerals and energy to the AI supply chain



Source: Morgan Stanley Research

How to Play AI in LatAm

Key themes

- Electricity price convergence
- Commodity prices
- Terms of trade
- Strategic Decoupling and reorganization of supply chains

The global AI capex cycle is one of the defining investment themes underpinning our LatAm, Brazil, and Mexico Model Portfolios. AI-driven demand for copper, lithium, and other critical commodities reinforces our Overweight position in materials. While debate continues over the pace of data center development in Latin America, particularly the timing of the base versus bull case, the key question is increasingly when, not if. This enhances the terminal value and risk-reward profile of regional electricity companies, particularly Axia (our preferred exposure), as well as Copel and Enel. Over time, we expect electricity prices to converge upward from the structurally low levels characteristic of today's relatively low-tech regional economy, creating meaningful upside for power producers.

Mexico is emerging as a critical manufacturing hub in the global AI supply chain. The country's rapid expansion in AI server production has become a macroeconomic game changer. At current growth rates, exports of IT hardware to the United States are on track to surpass auto exports by the end of 2026. We believe this marks the early stages of a much broader manufacturing ecosystem, supported not only by the AI investment cycle but also by the geopolitical realities of a more multipolar world. In this sense, AI is providing Latin America with an opportunity to reshape its investment cycle around higher-value manufacturing and strategic industrial capabilities.

The AI capex cycle is also reshaping the macroeconomic backdrop for Latin America. Strong global investment has improved the region's terms of trade, supported currency appreciation, and helped anchor inflation expectations, creating scope for lower interest

rates. At the same time, the renewed global capex cycle has contributed to a reversal in international investor positioning toward Emerging Markets, with a decade-long Overweight in developed markets narrowing substantially during 2025 and the first half of 2026. For Latin America, this has created an exceptionally supportive environment for equity markets and materially improved the region's overall risk-reward profile. None of this eliminates long standing fiscal, judicial, or regulatory challenges, which remain significant across many countries. However, a stronger external backdrop provides policy makers, companies, and investors with a considerably better starting point and, as in investing, having a better hand meaningfully improves the odds of success.

Exhibit 71: Morgan Stanley list of AI plays in Latin America

Ticker	Company Name	Country	Sector	Rating	Mkt Cap. (US$ bn)	PT (US$, unless noted)	Last price	Upside	2027 Fwd P/E Estimate
SCCO.N	Southern Copper Corp	Peru	Copper mining	Underweight	149.6	146.0	179.3	-18.6%	27.5
GMEXICOB.MX	Grupo Mexico S.A.B. de C.V.	Mexico	Copper mining	Equal-Weight	92.9	11.4	208.3	-94.5%	251.4
ANTO.L	Antofagasta	United Kingdom	Copper mining	Underweight	47.7	47.5	3,631.0	-98.7%	17.4
SQM.N	Sociedad Quimica y Minera de Chile S.A.	Chile	Lithium	++	9.8	++	68.7	na	na
VALE.N	Vale	Brazil	Mining	Equal-Weight	63.1	19.5	14.8	31.8%	6.5
AXIA3.SA	AXIA Energia	Brazil	Electricity generation	Overweight	22.7	10.4	50.3	-79.3%	7.9
ENIC.N	Enel Chile	Chile	Electricity generation	Overweight	6.0	4.3	4.3	-0.2%	11.1
CPLE3.SA	Companhia Paranaense de Energia	Brazil	Electricity generation	Overweight	8.5	3.0	14.5	-79.7%	10.5
PBR.N	Petrobras	Brazil	Energy	Overweight	69.8	28.5	18.8	51.8%	5.5
VIST.N	Vista Energy SAB de CV	Mexico	Energy	Overweight	6.6	94.0	69.5	35.3%	5.7
YPF.N	YPF SA	Argentina	Energy	Overweight	20.2	59.0	51.4	14.8%	7.7
FIBRAPL14.MX	Prologis Property Mexico SA de CV	Mexico	istrial Real Estate/Hardw	Overweight	7.2	4.9	75.8	-93.5%	12.1
VESTA.MX	Corporación Inmobiliaria Vesta SAB de C	Mexico	istrial Real Estate/Hardw	Overweight	3.0	4.5	58.9	-92.4%	303.5
WEGE3.SA	WEG	Brazil	Industrial goods	Equal-Weight	38.1	10.6	46.0	-76.9%	25.2
AURE3.SA	Auren Energia SA	Brazil	Electricity generation	Equal-Weight	2.3	2.8	11.2	-75.4%	nm

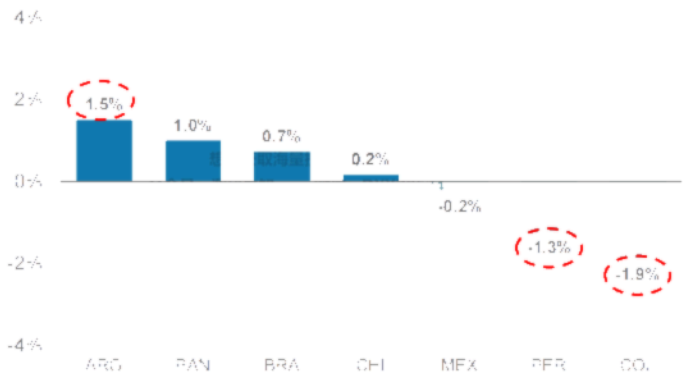
Source: Morgan Stanley Research, Factset; data as of 7/24/2026

++ Rating and estimates for this company have been removed from consideration in this report because, under applicable law and/or Morgan Stanley policy, Morgan Stanley may be precluded from issuing such information with respect to this company at this time.

Most of our Overweight positions in our LatAm Model Portfolio are linked to the global AI capex cycle, as the region supplies minerals, energy and hardware for the process. We favor electricity generation through AXIA (on the back of convergence of electricity prices globally), need-to-have minerals to the AI cycle through GMEXICO and SQM, and industrial real estate for IT hardware export through FIBRAPL.

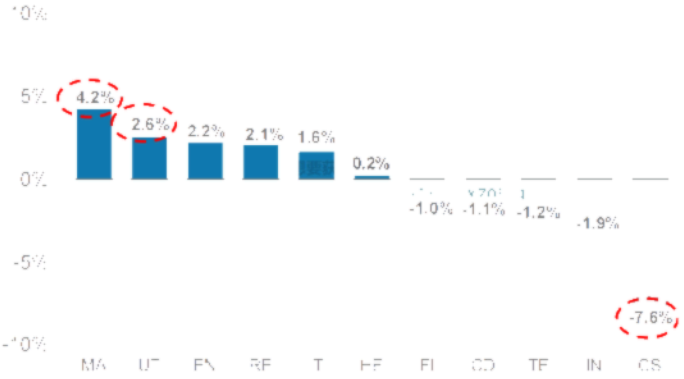
Axia Energia is our preferred way to play LatAm's data center theme. As highlighted in Latam Data Centers: Aiming for a Bull Case, power generators are both beneficiaries and enablers of data-center growth, making them the most compelling opportunity within utilities. AXIA3 (OW) operates LatAm's largest generation platform, with ~45GW of installed capacity and more than 50% of its portfolio available for commercialization beyond 2027, providing unmatched leverage to a higher power-price environment. Importantly, we believe the outlook for power prices remains highly skewed to the upside even without incremental data-center demand. The stock combines an attractive risk-reward profile with solid fundamentals, including a robust balance sheet, ongoing operational improvements, and attractive returns from transmission retrofitting investments.

Exhibit 72: We are OW Brazil, Argentina and Panama at the country level...



Source Bloomberg and Morgan Stanley Research

Exhibit 73: ...and Materials, Utilities and Energy at the sector level



Source Bloomberg and Morgan Stanley Research

Exhibit 74:

Most of our Overweight positions in our LatAm Model Portfolio are linked to the global AI capex cycle

07/23/26 Ticker	Company	Country	Sector	Rating	Weight	MSCI Latam	Over/ Under	Ccy	Price	12-mth PT/FV	Implied Return	Mkt Cap (USDm)
FIBRAPL14.MX	Prologis Property	Mexico	RE	OW	3.6%	0.6%	**3.0%**	MXN	76.6	90.0	24%	7,301
XP.O	XP Inc	Brazil	FI	OW	3.9%	0.9%	**3.0%**	USD	16.8	26.0	58%	8,842
AXIA3.SA	Axia Energia	Brazil	UT	OW	5.6%	2.6%	**3.0%**	BRL	51.4	72.0	48%	29,689
BVN.N	Buenaventura	Peru	MA	OW	3.6%	0.6%	**3.0%**	USD	31.3	34.5	15%	7,955
NU.N	Nubank	Brazil	FI	OW	8.1%	5.6%	**2.5%**	USD	14.2	21.0	48%	71,181
EMBJ.N	Embraer	Brazil	IN	OW	3.7%	1.2%	**2.5%**	USD	65.5	90.0	38%	12,120
MELI.O	Mercadolibre	Brazil	IT	OW	2.0%	0.0%	**2.0%**	USD	1,798.7	2,450.0	36%	91,189
PBR.N	Petrobras	Brazil	EN	OW	11.8%	9.8%	**2.0%**	USD	19.0	26.0	47%	122,443
CPLE3.SA	Copel	Brazil	UT	OW	2.4%	0.7%	**1.7%**	BRL	14.9	18.5	31%	8,730
OMAB.O	OMA	Mexico	IN	OW	2.0%	0.4%	**1.6%**	USD	104.1	125.0	26%	5,078
B3SA3.SA	B3	Brazil	FI	OW	3.7%	2.2%	**1.5%**	BRL	15.7	22.0	48%	15,534
SQM.N	Soquimich	Chile	MA	NA	3.0%	1.5%	**1.5%**	USD	69.3	na	NA	18,614
BSAC.N	Santander Chile	Chile	FI	EW	2.1%	0.6%	**1.5%**	USD	33.7	37.0	15%	15,895
VIST.N	Vista Oil & Gas	Argentina	EN	OW	1.5%	0.0%	**1.5%**	USD	69.1	97.0	40%	7,511
GCC.MX	Cem. de Chihuahua	Mexico	MA	OW	1.5%	0.0%	**1.5%**	MXN	208.1	250.0	22%	3,933
BPAC11.SA	BTG Pactual	Brazil	FI	OW	3.1%	1.6%	**1.5%**	BRL	56.3	65.0	18%	48,405
JBS.N	JBS	Brazil	CS	OW	2.2%	0.7%	**1.5%**	USD	12.1	19.0	65%	12,955
ENELCHILE.SN	Enel Chile	Chile	UT	OW	1.6%	0.3%	**1.3%**	CLP	80.7	85.0	10%	5,896
GMEXICOB.MX	Grupo Mexico	Mexico	MA	EW	4.9%	3.9%	**1.0%**	MXN	209.3	222.0	10%	92,975
TBBB.N	Tiendas 3B	Mexico	CS	OW	1.0%	0.0%	**1.0%**	USD	39.9	44.0	10%	4,804
CPA.N	Copa Holdings	Panama	IN	OW	1.0%	0.0%	**1.0%**	USD	135.7	180.0	38%	5,558
LAUR.O	Laureate	Mexico	CD	EW	1.0%	0.0%	**1.0%**	USD	36.6	36.5	7%	5,018
ORBIA.MX	Orbia	Mexico	MA	OW	1.0%	0.0%	**1.0%**	MXN	24.1	40.0	66%	2,885
CX.N	Cemex	Mexico	MA	OW	3.1%	2.1%	**1.0%**	USD	12.3	15.0	22%	17,888
SAUD3.SA	BradSaude	Brazil	HE	EW	0.9%	0.0%	**0.9%**	BRL	14.9	16.5	17%	1,600
GFNORTEO.MX	Banorte	Mexico	FI	EW	3.7%	3.1%	**0.6%**	MXN	189.3	230.0	29%	30,387
BBAS3.SA	Banco do Brasil	Brazil	FI	OW	1.4%	0.9%	**0.5%**	BRL	20.9	39.0	90%	23,592
SBSP3.SA	Sabesp	Brazil	UT	OW	1.9%	1.9%	**0.0%**	BRL	29.0	32.3	15%	20,136
AMX.N	America Movil	Mexico	TE	OW	2.7%	2.7%	**0.0%**	USD	26.0	30.0	18%	78,110
VALE.N	Vale	Brazil	MA	EW	6.7%	6.7%	**0.0%**	USD	14.8	16.5	19%	60,505
ITUB.N	Itau Unib.	Brazil	FI	OW	5.3%	5.3%	**0.0%**	USD	8.4	10.0	27%	92,043
Model portfolio					**100.0%**						**32%**	
MS universe											20%	
Model portfolio vs MS universe											**11%**	

Source: Bloomberg and Morgan Stanley Research
SQM = ++
++ = Rating and estimates for this company have been removed from consideration in this report because, under applicable law and/or Morgan Stanley policy, Morgan Stanley may be precluded from issuing such information or with respect to this company at this time.

LatAm Spring?

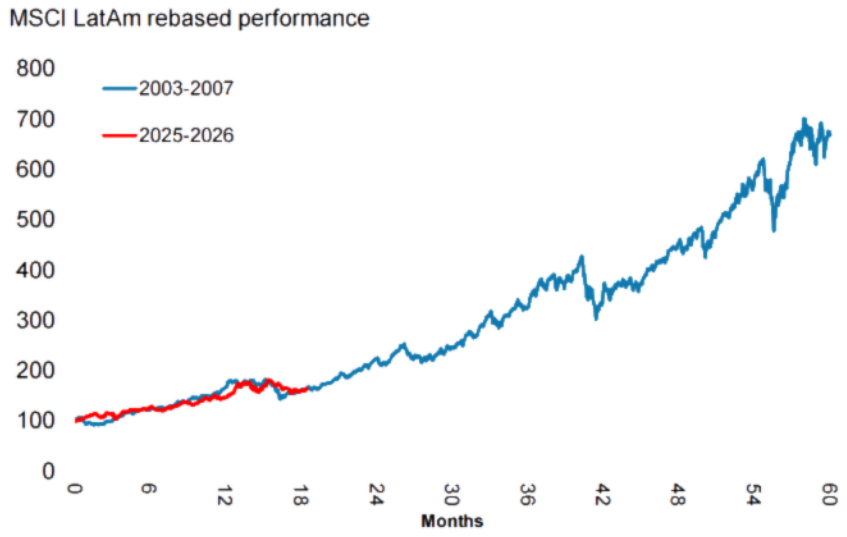
Our LatAm Spring Thesis represents the bull case for both the region's economies and its equity markets. We believe the global AI capex cycle provides Latin America with a rare opportunity to rebalance its growth model toward investment-led expansion. Over the past decade, rising populism and policy uncertainty have crowded out private investment, weakening both the real economy and capital markets. The scale of the AI investment cycle has the potential to reverse part of that trend. As a reference point, a single US$50 billion data center investment is equivalent to roughly 1% of the region's GDP. When

combined with the upstream investment required in copper, lithium, power infrastructure, and IT hardware manufacturing, the cumulative economic impact becomes considerably larger.

Infrastructure bottlenecks should be viewed as catalysts for investment rather than reasons for pessimism. Electricity remains a major constraint in countries such as Argentina and Mexico, but from a bullish perspective these shortages are precisely what create the need—and the incentive—for a new wave of capital expenditure. As global demand for power-intensive AI infrastructure grows, investment in generation, transmission, and supporting infrastructure could become a powerful driver of economic activity and productivity across the region.

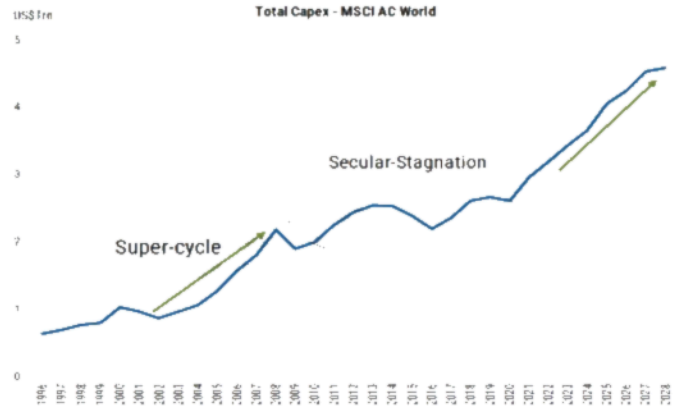
The parallels with the China-led commodity supercycle are increasingly compelling. Just as China's investment boom in the early 2000s transformed Latin America's terms of trade and underpinned a prolonged bull market in regional equities, we believe the AI capex cycle has the potential to become the defining external growth engine of the coming decade. While the beneficiaries will differ—extending beyond commodities to include electricity, infrastructure, and advanced manufacturing—the underlying mechanism is similar: a sustained global investment cycle that materially improves Latin America's growth prospects and investment opportunity set.

Exhibit 75: Could LatAm have another multiyear bull market driven by the AI Capex cycle – globally and eventually locally?



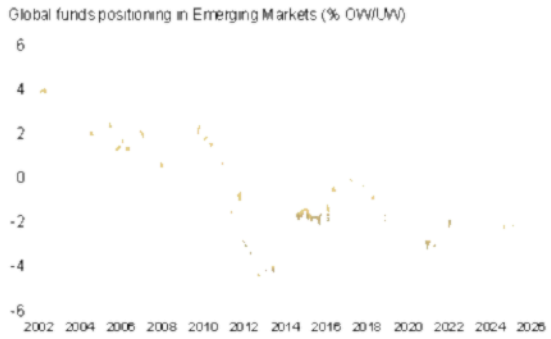
Source: Bloomberg, Morgan Stanley Research

Exhibit 76: Capex is back...



Source: FactSet, Morgan Stanley Research

Exhibit 77: ...and so are flows to emerging markets



Source: EPFR Global and Morgan Stanley Research

Key advantages under ground, key risks above ground

Latin America's AI opportunity is ultimately constrained less by its natural resources than by its institutional environment. As our Latin America Energy analyst Bruno Montanari often remarks, the region has been dealt an excellent hand below ground, while many of its challenges lie above ground. That observation is particularly relevant for the AI investment cycle. Latin America possesses many of the physical ingredients required to participate in global AI infrastructure, including abundant hydroelectric resources in Brazil, vast natural gas reserves in Argentina, and critical minerals across the region, but converting these advantages into data centers, digital infrastructure, and long-term investment will depend on permitting, electricity expansion, financing, judicial certainty, and the evolution of AI, privacy, and data regulation. There remains a considerable distance between recognizing that Brazil generates around 45% of its electricity from hydroelectric power (approximately 110 GW) or that Argentina possesses world-class gas reserves, and successfully converting these resources into globally competitive AI infrastructure.

Many of Latin America's AI risks stem from above-ground constraints rather than resource availability. As discussed throughout this report, the region's comparative advantages are largely physical, while its principal risks are institutional. These include political uncertainty, judicial intervention, regulatory complexity, infrastructure bottlenecks, and shortages of skilled labor and capital. Electricity constraints are particularly acute in Argentina and Mexico. Although these bottlenecks are solvable, doing so requires sustained investment, policy consistency, financing, equipment, and time. Likewise, legal and regulatory uncertainty – particularly in Brazil – could either accelerate or impede AI investment depending on the direction policy makers ultimately choose.

Legal risks – Legal and regulatory risk remains the most immediate uncertainty for AI investment in the region. Legal frameworks across Latin America, particularly in Brazil, are generally more interventionist and less predictable than those in the United States. The Brazilian Supreme Court has demonstrated a willingness to mandate content removal, impose liability on digital platforms, and require disclosure of user data in ways that differ materially from US legal principles. Whereas Section 230 of the US Communications Decency Act largely shields platforms from liability for user-generated content, Brazilian

jurisprudence has increasingly shifted responsibility toward the platforms themselves. Moreover, the Brazilian Supreme Court has assumed a de facto regulatory role in several areas where legislative guidance has been limited, including digital speech and platform governance. The temporary suspension of X in Brazil illustrates how judicial intervention can directly affect technology companies and increase regulatory uncertainty for AI developers and data center operators.

Data security risk – Data governance is likely to become an increasingly important geopolitical consideration for Latin America's AI ecosystem. Companies investing in data centers across the region must navigate uneven cybersecurity standards, evolving privacy regulation, and growing geopolitical competition between US- and Chinese-led technology ecosystems. Over time, we expect increasing restrictions on where sensitive data can be stored, processed, and transferred, as well as greater scrutiny over ownership of digital infrastructure. US policy is likely to remain the dominant influence for countries closely integrated into the North American technology ecosystem. At the same time, Brazil has joined the China-led World Artificial Intelligence Cooperation Organization (WAICO), alongside several other countries including Venezuela and Cuba, highlighting that competing regulatory frameworks for AI are beginning to emerge. As the global AI ecosystem fragments, Latin American countries may increasingly face pressure to align with one regulatory architecture or another.

Equipment risk – Technology supply chains are becoming a national security issue, creating additional risks around equipment and infrastructure. Data centers in Latin America will rely on globally sourced hardware spanning servers, networking equipment, semiconductors, and power systems. As technology competition between the United States and China intensifies, the origin of this equipment is likely to receive greater regulatory scrutiny. Concerns regarding cybersecurity, embedded software vulnerabilities, supply chain resilience, and potential foreign influence—whether ultimately substantiated or not—may influence procurement decisions, financing, and regulatory approvals. We believe Mexico will sit at the center of these debates given its deep integration into the US manufacturing ecosystem and its growing role in AI server production. While some equipment standards may eventually be incorporated into the USMCA framework, we expect much of the practical guidance to emerge through bilateral agreements, executive actions, and national security policy rather than formal trade agreements.

Korea Equity Strategy

Joon Seok, Heewon Choi, Minseo Kang

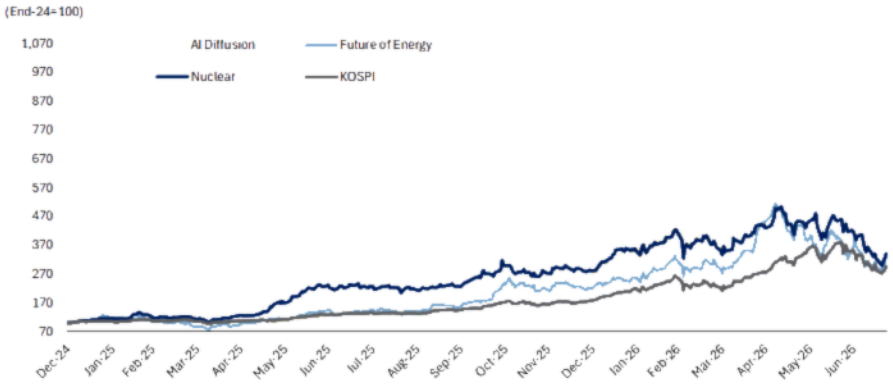
A key global partner for AI infra build-out

The focus is on memory and peripheral areas of tech that are needed to satiate the growing demand for computing, as outlined by our Asia/Europe Tech analyst, Shawn Kim. However, to support the build-out of AI data centers there are other pressing needs such as power generation. We believe Korea has many of the needed industries to build out AI infrastructure and will be a key exporter in this area.

Korea's equity market has been seen as cyclical and asset-heavy, which has been a disadvantage for the KOSPI. However, we see a multi-year cycle for Industrials, with key sub-themes being:

- Energy security (power generation, nuclear, LNG)
- Defense/shipyards (active deterrence),
- Reconstruction/infra build-out
- Autos/robotics

Exhibit 78: Share Price Performance for AI-Linked Thematic Groups



Source: Wisenut, Morgan Stanley Research. Past performance is no guarantee of future results.

Exhibit 79: AI-Linked Thematic Stocks

	Company Name	Ticker	Rating	TP	Close	Mkt. Cap (Wtn)	Share Price Performance (%) YTD	6M	3M
	KOSPI				7,097	5,632.0	68.4	42.2	9.6
AI Tech Diffusion	Samsung Electronics	005930.KS	Overweight	381,000	270,000	1,578.5	125.2	77.5	20.3
	SK Hynix	000660.KS	Overweight	2,600,000	1,919,000	1,367.7	194.8	150.2	56.7
	Samsung Electro-Mechanics	009150.KS	Overweight	2,560,000	1,448,000	108.2	467.8	432.4	87.1
	Leeno Industrial Inc.	058470.KQ	Overweight	77,000	73,900	5.6	22.6	7.6	-40.1
	Hanmi Semiconductor	042700.KS	Equal-Weight	230,000	216,000	20.6	69.5	26.2	-26.4
	FADU	440110.KQ	Overweight	150,000	73,700	3.7	246.8	246.8	7.8
	Wonik IPS	240810.KQ	Overweight	160,000	113,300	5.6	66.9	47.9	0.1
	Isu petasys	007660.KS	Overweight	114,000	93,400	6.9	-21.6	-22.1	-37.8
Future of Energy	Sung Kwang Bend	014620.KQ	Overweight	50,000	27,500	0.7	6.8	-5.8	-31.2
	T.K.	023160.KQ	Equal-Weight	48,000	23,000	0.6	6.0	-9.3	-47.8
	LS ELECTRIC	010120.KS	Overweight	270,000	223,500	33.5	142.9	119.6	3.0
	HD HYUNDAI ELECTRIC	267260.KS	Overweight	1,500,000	869,000	31.3	12.3	-1.7	-23.0
Nuclear	Doosan Enerbility	034020.KS	Overweight	120,000	73,200	46.9	-2.8	-21.5	-40.3
	Samsung C&T Corporation	028260.KS	Overweight	340,000	360,000	58.4	50.3	17.8	12.5
	Hyundai Eng. & Constr.	000720.KS	Not Covered		108,900	12.1	55.4	-0.3	-37.9

Source: Wisenut, Morgan Stanley Research. Past performance is no guarantee of future results (as of July 24, 2026)

Power generation: Powering AI now intertwined with Energy Security

The rapid expansion of AI-driven data center demand is bringing power availability to the forefront as a critical bottleneck in the AI value chain, elevating "Powering AI" into a strategic investment theme. This is driving investor focus across the power generation and electrical infrastructure ecosystem, particularly on Korean electrical equipment, which benefits from visible backlog and margin expansion tied to robust US transformer and grid equipment orders. Furthermore, the US-Iran conflict has brought energy security to the forefront, supporting investments across reconstruction, gas, petrochemicals, LNG, and clean energy.

Korea's power generation and grid equipment sector has entered a structurally distinct phase — transitioning from a domestically focused industrial base to a globally competitive export engine. Korean power equipment makers are now leveraging cost competitiveness and proven technology to win large-scale tenders in the US, Middle East, and Europe, where grid modernization and renewable integration are accelerating.

What makes this cycle distinct from prior up-cycles is its global reach and the convergence of multiple structural demand drivers: AI-driven data center power demand, aging grid replacement, renewable integration, and on-site power generation — all arriving simultaneously and against a backdrop of severe supply constraints:

- South Korea has been adding wallet share in the US transformer import market significantly since 2023, rising to 14% (from 8% in 2022) and further to 19% in 2024, fueled by aging infrastructure replacement demand and timely capacity expansion.
- As of 2024, the US is the largest destination for Korean transformer exports, accounting for 45% of the total mix. Europe and the Middle East account for 12% and 10%, respectively, with Middle East exports growing faster.

Korean power equipment companies benefit from several structural advantages in the current environment:

- *Product orientation over solutions:* Companies like HD Hyundai Electric are typically more product-oriented businesses (ultra-high voltage transformers) rather than solutions-focused, which is favorable for margin mix.
- *Reliability premium:* US regulatory frameworks create an asymmetric incentive for utilities — they face more long-term risk from underinvestment than from incremental post-event spending. Post-outage responses typically involve earlier replacement cycles, higher spare inventories, and substation hardening, all of which benefit Korean suppliers with >40% of their order backlog exposed to the US.
- *UHV technology leadership:* Only a limited number of global manufacturers (3-4 global players) can produce 765kV ultra-high-voltage transformers, and Korean players are among the most preferred suppliers among US utilities.

Refer to the below reports:

- LS Electric: 2Q26 Earnings – Order Momentum Continues to Accelerate; Q&A Summary (23 Jul 2026)

- HD Hyundai Electric Co Ltd: CEO Meeting Takeaways + Q&A Summary (9 Jul 2026)
- Investor Presentation: Korea Tech and EV Materials – Key Issues and Debates (14 Apr 2026)

Exhibit 80: Korean Power Generation Companies

	Core product categories	Major manufacturing & production sites	Revenue exposure by region	Revenue exposure by product	Key segment focus
HD Hyundai Electric	Power & distribution transformers, high voltage GIC & switchgear, low/medium-voltage breakers	South Korea (Ulsan), US (Montgomery, Alabama); Sales and service offices globally	North America 30%, Domestic and others 45%, Middle East 25%	Power: 48%, Distribution 26%, Rotating 16%, Subsidiaries 10%	Primarily utility grid equipment (UHV transformers, substations) and industrial power systems
LS Electric	Comprehensive T&D and automation range: HV/MV/LV switchgear and circuit breakers, power transformers (inc. HV transformers).	South Korea (Cheongju smart factory for LV components, Busan for HV transformers), China (Dalian, Wuxi for manufacturing subsidiaries), Vietnam and regional offices in Middle East and the US.	Domestic: 50%, North America 17%, Asia (ex-China): 14%, China: 9%, Europe 8%	Electric Infra: 32%, Electric equipment: 21%, Renewables: 8%, Automation: 7%, Subsidiaries: 32%	Dual focus on utilities (power transmission & distribution equipment) and industrials (factory automation and smart energy solutions).
Hyosung Heavy Industries	HV/MV power equipment, power system solutions, rotating machinery & industrial equipment	Primary production complex in Changwon, South Korea, while having local factories in Nantong, China, and Memphis, USA.	Korea: 65%, Asia (ex-Korea): 15%, North America: 10%, Europe: 5%, Others 5%	Heavy industries: 63%, construction 37%	Primarily serves electric utilities and transmission grid operators, industrial and commercial sectors, and renewable energy projects

Source: Company reports, Morgan Stanley research

Nuclear power: Korea to be a backbone for AI power demand

Korea's nuclear power industry is well-positioned to capitalize on the AI-driven electricity demand surge. The country's fully localized nuclear value chain, government-backed export strategy ("Team Korea"), early-mover SMR partnerships with US developers, and explicit policy linkage between nuclear energy and AI sovereignty create a structurally differentiated competitive position. Doosan Enerbility — as Korea's sole NSSS/STG supplier and a multi-role EPC provider — is the primary vehicle for this exposure. The key near-term catalysts are large-scale reactor order wins (Middle East, Eastern Europe) and SMR commercialization, and Korea's new government has explicitly linked nuclear energy to its AI sovereignty initiative, supporting nuclear as a stable baseload source to ensure the country's AI competitiveness.

Korea nuclear value chain players are seeking opportunities in US-based nuclear projects as well, particularly in equipment supply and EPC services by leveraging their partnerships with US companies in both the large-scale NPPs and SMRs. Looking further ahead, as demand for powering data centers continues to grow and SMRs emerge as a commercially viable source of reliable power, Korean companies are actively collaborating with major SMR developers to establish themselves as key players in the global SMR market; to secure a first-mover advantage, they are supporting first-of-a-kind SMR deployments across multiple countries including the US.

Korea's nuclear power and power equipment supply chain remains well positioned given limited global supplier capacity with proven track record across both large-scale reactors and SMR projects, while supply shortages in gas turbines are further reinforcing the strategic value of diversified dispatchable power solution providers. Korea holds a distinctive and strategically superior position in the global nuclear landscape, underpinned by three structural pillars:

- *Decades-long domestic track record:* Korea has developed one of the most competitive domestic nuclear value chains through several decades of domestic nuclear power plant (NPP) projects, including an overseas track record (e.g., UAE

Barakah).

- *Full vertical integration:* From forgings and major equipment manufacturing to EPC and decommissioning and back-end solutions, Doosan Enerbility has become one of the few global players with vertical integration across the entire nuclear value chain.
- *Geopolitical strategic value:* Korea's strategic position as an alliance partner in an increasingly multipolar world, where energy sovereignty and security are becoming more important, gives it a distinctive edge.

Refer to the below reports:

- Doosan Enerbility: Re-entering the Catalyst Window (16 Jun 2026)
- Asia Energy Security and AI: Energy Meets Compute: Supercycle Recharges (21 May 2026)

Exhibit 81: Nuclear Power: Key Things to Look for in 2H26

Segment	Potential Catalyst	Current Status	What would be new/meaningful
Large-scale Reactor	Korea-US Strategic Investment Corp. project selection	The Korea-US Strategic Investment Corporation launched on June 18, 2026, to support the US investment package. - Separately, Japan is reportedly mulling investments into US SMR developments, potentially including GE-Hitachi's BWRX-300 and Nuscale	- Inclusion of nuclear / SMR / power infrastructure among first selected projects could be positive for Korea nuclear supply chain players - Positive surprise if Korean supply-chain role expands beyond legacy scope (RPV & SG for Doosan Enerbility), or if Korea offers procurement packages or JV/cooperation structure for the US nuclear buildout
	Korean participation in US AP1000 / Westinghouse buildout	- US govt announced a strategic partnership to construct at least US$80bn of new reactors in the US using AP1000 technology - Meanwhile, Fermi America signed MOUs with Doosan Enerbility and Hyundai E&C related to four planned AP1000 reactors for Project Matador	- EPC contracts within 2H26~1H27 for the Poland and Bulgaria AP1000 units will be the key milestone for the projects, likely followed by tangible order for nuclear equipment
SMR	Nuscale project visibility	-ENTRA1 continues to work with TVA on planning for up to 6GW of NuScale SMR capacity; Market expectation points to 2H26 as an approval window. - Nuscale-to-Doosan Enerbility order visibility could improve once the PPA is made	TVA board approval, PPA, site / plant sequencing, or procurement timeline Binding module / component order, delivery schedule, or production ramp
	Additional equipment order for TerraPower Natrium Reactor	- TerraPower and Doosan Enerbility agreed to supply core barrel, guard vessel, and internal supports for the Natrium Reactor.	Additional equipment orders / More visibility on manufacturing milestones, schedule confirmation, or follow-on Natrium deployment
	Rolls-Royce partnership	- Major projects for Rolls-Royce SMR deployment include (1) UK Wylfa, (2) Czech Temelin, and (3) Sweden Ringhals.	Conversion from pre-production / manufacturability review into tangible equipment order
Gas/Hydro, EPC	overseas gas turbine orders	- Doosan Enerbility has won orders for 10 units of 380mw-class gas turbines during 1Q26, including 7 units to a US datacenter developer	Additional export orders / More management guidance on delivery slot visibility and pricing momentum
Others	Korea 12th Basic Plan for Electricity Supply and Demand	- Korean government is set to announce the 12th Basic Plan for Electricity Supply and Demand in 2H26 (likely in 3Q26) - The updated plan will cover the mid-to-long term powering scheme during 2026 through 2040 Korea plans to proceed with two new large-scale reactors targeting 2037-2038 under the current 11th plan, while targeting domestic i-SMR development	Higher power-demand forecast, confirmation of new nuclear capacity, additional nuclear / SMR assumptions, or favorable treatment of dispatchable power

Source: Morgan Stanley research

Overseas E&C: Energy transition conjoined with reconstruction/infra buildup

Korea's overseas Engineering & Construction (E&C) sector is entering a structurally favorable multi-year up-cycle, driven by the convergence of energy security imperatives, the global energy transition, and a nuclear renaissance. There continue to be meaningful opportunities for Korean contractors overseas, with traditional hydrocarbon and infrastructure projects remaining a mainstay. In our view, however, there are increasing opportunities in energy transition and nuclear power EPC.

Korean E&C firms benefit from a set of structural competitive advantages in overseas markets, including:

1. A long history of working in key regions and healthy bilateral government relationships;

2. Technical prowess and sourcing/networking capacity;
3. A strong track record of finishing projects on or ahead of time; and
4. Strong backing from Korean export credit agencies and policy banks such as Korea Eximbank and K-SURE.

On the nuclear EPC business, we see incremental optionality for Korean overseas EPC.

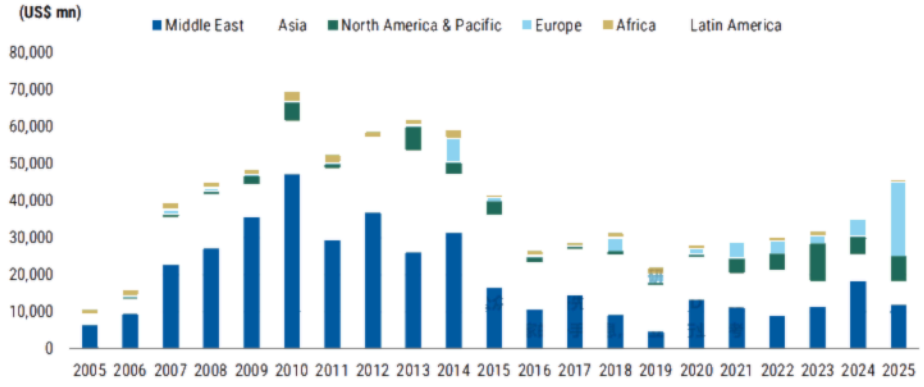
- Hyundai E&C partners with Westinghouse on large-scale nuclear projects in Romania, Slovenia, Finland, Sweden, and Ukraine, and with Holtec on deploying SMRs at the Palisades nuclear plant site in Michigan. Samsung C&T is partnering with Fluor on Romania NPP #3, 4 and with NuScale on Romania Doicesti SMR; it has teamed up with Fermi Energia to deploy SMRs in Estonia.
- For Samsung C&T, the nuclear pipeline includes Cernavoda EPC (Romania) and Ringhals SMR FEED (Sweden), potentially materializing by end-2026 or 2027. SCT expects initial progress from its SMR project pipeline, which could gradually translate into new order inflows. In addition to nuclear-related opportunities, the company is also targeting EPC contracts across broader energy infrastructure, including solar power and data centers.

Refer to following reports:

- Samsung E&A Co Ltd: 2Q26 Review: Earnings Beat; Order Momentum Remains Strong (24 Jul 2026)

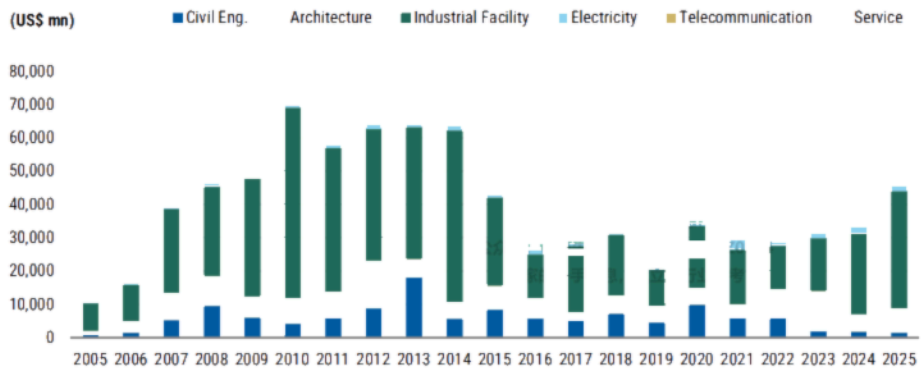
- Samsung E&A Co Ltd: Back to the Future (29 Apr 2026)
- Samsung C&T: 1Q26 Review: In Line Excluding One-off; Order Momentum to Build via Captive Demand (29 Apr 2026)

Exhibit 82: Korean E&C Sector – Overseas New Orders – by Geography



Source: International Construction Information Service, Morgan Stanley research

Exhibit 83: Korean E&C Sector – Overseas New Orders – by Project Type



Source: International Construction Information Service, Morgan Stanley research

Korea is also undertaking initiatives to build up its domestic AI ecosystem

AI infrastructure buildup already on the way

From a domestic perspective, we see a significant and structural market opportunity as Korea's first wave of AI data centers begins construction, supported by favorable government policies. Korea is emerging quickly as an early adopter of AI-enabled services, making the upgrade of supporting infrastructure a priority for both service providers and the government.

- On the enterprise side, AI is being applied to coding, video editing, demand prediction, ad spending optimization, route optimization, and DNA analysis. On the consumer side, a key US foundation model has already reached 18 million monthly active users in Korea, with further adoption expected following its integration into KakaoTalk. The expected launch of AI agents in 2026 should further lift usage beyond chatbots toward task execution
- Global AI companies such as Amazon are requesting large-scale, AI-optimized data center build-outs in Korea. Korean conglomerates, recognizing the need to secure AI computing power, have also begun constructing internal data centers, led by Samsung SDS.

Supporting AI infrastructure is a key initiative for the Lee administration, which took office in June: Core policy pillars include:

1. Developing a "sovereign AI" model;
2. Major investment and infrastructure build-out;
3. Regulatory reform; and
4. Talent development.

The government targets W100tn of investment — a mix of public and private capital — including spending on GPUs and AI semiconductors. Data centers designated as national strategic technology facilities will qualify for special tax incentives and broader policy support.

A connected initiative under the AI Superhighway program involves upgrading the

national grid via High-Voltage Direct Current (HVDC) technology: MOTIE projects data center power demand to grow from ~8 TWh in 2025E to 30 TWh by 2038E, implying a ~10% CAGR vs. ~2% CAGR for total power demand. The Energy Highway is a multi-phase, multi-decade project to build a "U-shaped" HVDC backbone spanning the nation's coasts.

- *Phase 1 (by 2030):* West coast HVDC corridor from the renewable-rich southwest to Seoul, capable of delivering up to 20 GW of clean electricity. A 2 GW HVDC line from Saemangeum to Hwaseong (~220 km) is already in the demonstration stage.
- *Phase 2 (2030s):* Southern coast expansion linking power hubs of the south, targeting a 30%+ increase in total transmission capacity.
- *Phase 3 (2040–50):* East coast completion forming a nationwide HVDC loop. The Donghaean–Shin Gapyeong HVDC line (~230 km) is already under construction.

Refer to the below report:

- S. Korea IT Services and Software: AI and robotics – is the hype for real? (22 Jun 2026)
- Tech Diffusion and GenAI: S. Korea: The Emerging AI Infrastructure Opportunity (21 Jan 2026

Exhibit 84: List of "Neocloud" data centers announced or under consideration in Korea

Project / Partnership	Location	Planned Capacity	Estimated Investment	Timeline / Status	Details
SK Telecom / AWS	Ulsan	**100 MW** (Phase 1: 40 MW)	~W7 trillion	Phase 1 targeted by 2027	- Planned AI data center with approx. 60,000 GPUs - Site has sufficient land for future GW-scale expansion.
Samsung SDS	Gumi	**120 MW**	Up to W8 trillion	Stage 1 operations targeted for early 2029	- Large-scale AI data center serving both Samsung affiliates and external enterprise customers.
Samsung SDS / National AI Computing Center	Solaseado (Jeollanam-do)	**40 MW**	~W2.5 trillion	Full opening and completion targeted by Oct 2028	- AI computing center with approx. 15,000 GPUs. - Structured as a special purpose company (SPC) with 30% government funding.
SK Telecom / US Foundation Model	Solaseado (Jeollanam-do)	TBD	TBD	Feasibility study underway	- Proposed AI data center being evaluated as part of collaboration between SK Telecom and US foundation model.
AWS	Incheon	**~100 MW**	Not disclosed	Announced at the 2026 AWS APAC Summit	- Planned hyperscale facility supporting both AWS cloud infrastructure and AI computing workloads.

Source: Company data, media sources, Morgan Stanley Research

A new industrial policy super-cycle anchored on AI, semiconductors and regional development

South Korea has unveiled one of the largest industrial investment programmes globally, centred on approximately KRW2,000tn (~US1.5tn) of regional investment and around KRW1,500tn of cumulative long-term domestic investment commitments by Samsung Electronics, SK Group/SK hynix, and other companies.

The programme aims to position South Korea as the global leader across the AI infrastructure value chain by accelerating investment in:

- Advanced semiconductor manufacturing
- AI data centres
- Physical AI (robotics and industrial automation)

We believe the initiative represents Korea's most significant industrial policy since the

development of its semiconductor industry in the 1980s–1990s. It could materially reinforce the country's competitive position in AI and high-tech over the next decade.

Our economist for Korea, Kathleen Oh, sees meaningful upside risk to the growth outlook: a minimum of 50-70bp over the next three years; 100-150bp beyond 2031. However, she believes a more concrete investment allocation and specific execution plan are needed at this point.

- Refer to South Korea: A Landmark Industrial Strategy Unveiled – Execution is Key (1 Jul 2026)

The key strategic objectives are as follows:

Maintain global semiconductor leadership: Korea currently dominates global memory semiconductor production through Samsung Electronics and SK hynix. As AI computing demand accelerates, the government intends to expand manufacturing capacity before global supply constraints become binding.

- Four new memory semiconductor fabrication plants
- Expanded High Bandwidth Memory (HBM) manufacturing
- Advanced packaging facilities
- Semiconductor materials ecosystem expansion

This strategy seeks to reinforce Korea's competitive advantage in AI memory, where Samsung and SK Hynix together account for the majority of global HBM supply.

Build national AI infrastructure: Unlike previous semiconductor-focused policies, this programme extends across the entire AI value chain. SK Group alone plans approximately KRW1,000tn of AI data center investment over the coming decade as part of its broader AI strategy.

- Hyperscale AI data centres
- Cloud infrastructure
- National AI computing capacity
- Power transmission
- Renewable energy integration
- Telecommunications infrastructure

Develop physical AI: A notable feature of the initiative is its focus on "Physical AI", encompassing a wide range of areas. Rather than competing directly with US foundation-model developers, Korea aims to leverage its manufacturing strengths to become a leader in industrial AI applications.

- Industrial robotics
- Smart manufacturing + Autonomous production
- Logistics automation
- AI-enabled mobility

This report references export controls and/or entities that may be subject to export control restrictions. Readers are solely responsible for ensuring that their investment or trade activities are carried out in compliance with applicable laws.

This report references U.S. Executive Order 14032 and/or entities or securities that are designated thereunder. U.S. persons may be prohibited from buying certain securities of entities named in this report. Readers are solely responsible for ensuring that their investment activities are carried out in compliance with applicable laws.