This material is neither intended to be distributed to Mainland China investors nor to provide securities investment consultancy services within the territory of Mainland China. This material or any portion hereof may not be reprinted, sold or redistributed without the written consent of J.P. Morgan.

China AI Infra Ecosystem

Unprecedented growth; lingering constraints

China's AI ecosystem is seeing a fast take-off, as local open-sourced models catch up with global frontier models, large CSPs increase investments in AI infrastructure and local AI Infra ecosystem sees rapid development, in response to restrictions on overseas chips. Domestic AI compute demand is likely to show an ~80% CAGR, to support strong inference token consumption. Looking into 2028, foundry and HBM capacity are likely the key bottlenecks, despite significant capacity growth. We expect select AI Chip vendors and SPEs to benefit from the local buildout. We believe that China can support 80% of its AI infra demand with local chips by 2028, up from 40% in 2025. To overcome single chip performance gap, Chinese AI chip vendors are likely to aggressively embrace advanced packaging, interconnects and system-level optimization – benefiting back-end service providers and server ODM/OEMs. Beyond tech hardware, this aggressive localization buildout should also benefit Power, Cooling, and ESS vendors, and rekindle growth for local Datacenter providers. Our top picks are Iluvatar CoreX, NAURA, AMEC, JCET, V-Test, Huaqin, VNET, Sungrow, CATL, TGOOD.

- **Local AI compute growing rapidly, benefiting AI Chip vendors and SPE.** We expect China local compute demand to show an 80% CAGR, driven by rising token consumption. This should help support a 40% CAGR in AI Chip units into 2030E. With overseas GPU availability remaining constrained, the localization ratio should rise materially, reaching 80%/90% of total supply by 2028E/30E, respectively. We prefer domestic AI chip vendors with secure foundry capacity and inventory (Iluvatar Core X), and leading domestic WFE vendors (AMEC, Naura).
- **Foundry and HBM primary bottlenecks into 2028, despite strong capacity growth.** We believe local foundries are likely to add 150-200kwpm of advanced logic capacity over the next 2-3 years, although yield challenges and technology migration remain bottlenecks, which could keep chip supply below domestic demand. HBM is another chokepoint, in our view: the transition to local HBM3/3e production is currently constrained by qualification timelines and equipment access. The demand supply gap is likely to bridge only gradually, given slower yield progress in both advanced logic and HBM manufacturing.
- **Advanced packaging, interconnect and system optimization to bridge chip supply and performance gap.** With chip supply remaining constrained, we expect Chinese vendors to pursue chiplet-based, multi-die designs. Advanced packaging is likely to grow quite fast, as domestic OSATs can handle CoWoS-S-like solutions today. Supernode solutions such as CloudMatrix384 that can boost scale-up density can also help level the single-chip performance gap, benefiting server ODM/OEMs.
- **Local ecosystem building up in Datacenters, Power and Cooling.** Wholesale IDC is robust: strong AI-related orders at leading vendors and tighter power-quota screening sustain a benign pricing environment. We forecast a ~50% CAGR (2026-30E) in incremental AI power demand, driving a structural shift toward liquid cooling, 800V DC/SST architecture and prefabricated power modules, while ESS and backup-power suppliers gain from "grid-friendly" configurations.

Technology and Telecoms

Gokul Hariharan AC
(852) 2800-8564
gokul.hariharan@jpmorgan.com
J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

Billy Feng
(86-21) 6106 6359
billy.feng@jpmorgan.com
SAC Registration Number: S1730520030005
J.P. Morgan Securities (China) Company Limited

Ri Xu
(86-21) 6106 6318
ri.xu@jpmchase.com
SAC Registration Number: S1730522100001
J.P. Morgan Securities (China) Company Limited

Jocelyn Gao
(852) 2800-8529
jocelyn.gao@jpmorgan.com
J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

Cherry Liu
(86-21) 6106-6200
ye.liu@jpmorgan.com
SAC Registration Number: S1730521090001
J.P. Morgan Securities (China) Company Limited

Albert Hung
(886-2) 2725-9875
albert.hung@jpmchase.com
J.P. Morgan Securities (Taiwan) Limited/ J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

See page 30 for analyst certification and important disclosures, including non-US analyst disclosures.

Contributing Authors

Technology and Telecoms

Gokul Hariharan [AC]
(852) 2800-8564
gokul.hariharan@jpmorgan.com
J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

Billy Feng
(86-21) 6106 6359
billy.feng@jpmorgan.com
SAC Registration Number: S1730520030005
J.P. Morgan Securities (China) Company Limited

Ri Xu
(86-21) 6106 6318
ri.xu@jpmchase.com
SAC Registration Number: S1730522100001
J.P. Morgan Securities (China) Company Limited

Jocelyn Gao
(852) 2800-8529
jocelyn.gao@jpmorgan.com
J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

Cherry Liu
(86-21) 6106-6200
ye.liu@jpmorgan.com
SAC Registration Number: S1730521090001
J.P. Morgan Securities (China) Company Limited

Albert Hung
(886-2) 2725-9875
albert.hung@jpmchase.com
J.P. Morgan Securities (Taiwan) Limited/ J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

Power Equipment and Utilities

Alan Hon
(852) 2800-8573
alan.hon@jpmorgan.com
J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

Stephen Tsui, CFA
(852) 2800-8592
stephen.tsui@jpmorgan.com
J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

Infrastructure, Industrials & Transport

Karen Li, CFA
(852) 2800-8589
karen.yy.li@jpmorgan.com
J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

Asia Autos & EV Battery

Rebecca Wen
(852) 2800-8505
rebecca.y.wen@jpmorgan.com
J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

J.P.Morgan

Table of Contents

Executive summary 4

Exponential token growth driving domestic AI chip demand 8

Rapidly expanding AI inference token usage in China 8
Token demand supports 80% computing growth CAGR 10
This translates into AI chip unit growth of 40% CAGR 11

Supply of computing chips the key bottleneck 13

Landscape of China's AI chip market 13
Foundry supply assumptions suggest tight supply into 2028 15

System-level efforts to deliver higher AI capability 20

Computing performance is not the only metric to measure domestic AI competitiveness 20
Structural shifts in cooling, power and energy stack from the AI infrastructure cycle 23

Key beneficiaries of China's indigenous AI supply chain 25

Local semiconductor supply chain set to benefit immediately 26
Sourcing accessory chips and peripherals is not difficult, but localization is likely to increase 26
Ecosystem-wide efforts to facilitate China's data center infrastructure build-out 27

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

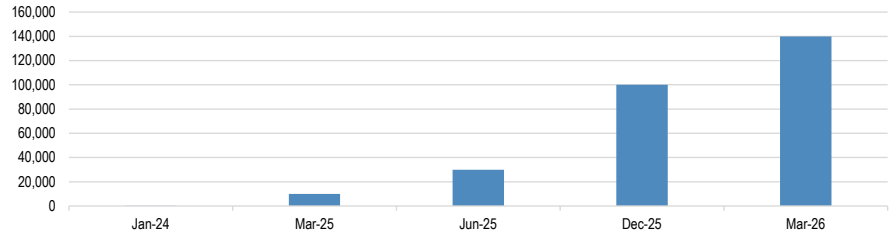
Executive summary

Exponential token consumption growth driving computing power demand

The emergence of DeepSeek and the proliferation of low-cost, open-source models has led to a surge in daily token consumption from approximately 100 billion in early 2024 to over 140 trillion by March 2026, a 1,000-fold increase in under two years.

Figure 1: Daily token consumption in China

Billion

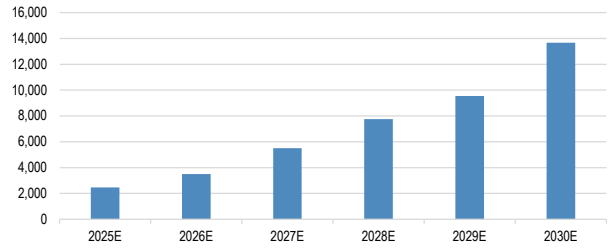


Source: National Data Administration (26 Mar 2026).

We forecast that inference-driven computing power demand will show an ~80% CAGR through 2030, accounting for 85% of total requirements. We therefore estimate chip demand required to support the strong token expansion will show a ~40% CAGR during the same period, reaching ~14mn units by 2030E. Based on our supply evaluation, we expect 2mn/3mn/5mn domestic AI chips to be supplied in 2026/27/28E. The increasing adoption of multi-die packages with smaller die sizes could improve output levels or use wafer capacity more efficiently but does not remove the near-term supply constraint.

Figure 2: Incremental chip demand (annual basis)

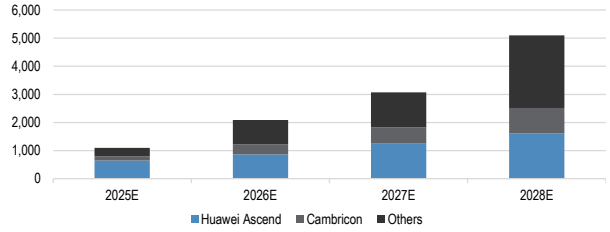
k units (700mm^2 equivalent)



Source: J.P. Morgan estimates.

Figure 3: Shipment forecast of AI chips supplied by domestic companies

k units (700mm^2 equivalent)



Source: J.P. Morgan estimates.

Foundry and HBM the primary supply bottlenecks

At the foundry level, the availability of advanced wafer fabrication capacity represents the most critical chokepoint. Despite ambition to add 150-200k wafers per

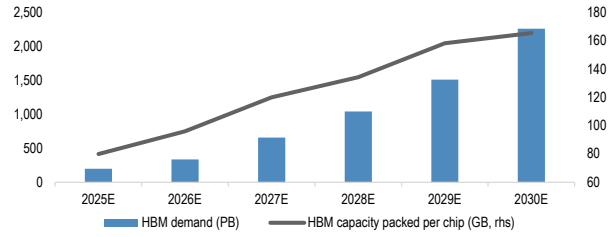
Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

month of advanced capacity over the next three years, effective yield improvement and technology migration present formidable challenges.

Beyond compute dies, High Bandwidth Memory (HBM) availability is another major bottleneck, in our view. Existing domestic AI chips rely mainly on HBM2e, while the transition to HBM3/3e remains constrained by qualification timelines and equipment access limitations. Our HBM demand forecast suggests a ~60% CAGR during 2026-30E, implying ~200kwpm HBM capacity required to completely meet the domestic demand by 2030. Although CXMT's listing could be a positive catalyst, we are not expecting the self-reliance gap to be rapidly closed, as product competitiveness is as critical as capacity expansion.

Figure 4: HBM demand forecast vs capacity per compute die

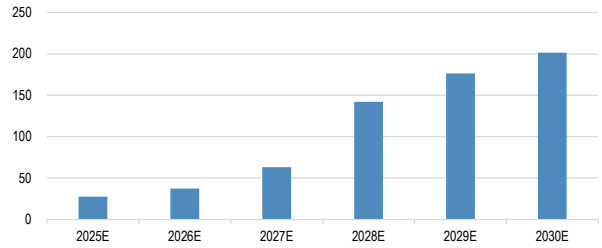
PB, GB



Source: J.P. Morgan estimates.

Figure 5: Effective wafer capacity required for HBM base wafer

kwpm



Source: J.P. Morgan estimates.

Further, owing to front-end constraints, we believe back-end processes are likely to be an important driver for chip performance improvement. Advanced packaging capabilities at key domestic vendors are still transitioning toward CoWoS-L-like technology, with mass production timelines from late 2026 or beyond.

Indigenous packaging and system solutions to mitigate constraints

Chinese players across the supply chain are pursuing a multi-pronged strategy to mitigate hardware constraints. On the architecture front, chiplet-based designs and multi-die packaging are being adopted to circumvent yield limitations and improve wafer output efficiency. Huawei's Tau Scaling Law and LogicFolding approach demonstrate that meaningful performance gains remain achievable even when geometric scaling and advanced lithography access are constrained.

- For instance, rather than pursuing a significantly powerful compute die, Huawei launched its next generation of AI chips, including Ascend 950PR and Ascend 950DT, which use the same 950 compute die but different memory configurations to match various workloads - 950PR is dedicated for prefill and recommendation purposes, and 950DT with higher memory capacity/ bandwidth will be used for the decode phase of inference and training.

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Table 1: Comparison of Huawei Ascend chips

	910C	950PR	950DT
Availability	1Q25	1Q26	4Q26
e	SIMD	SIMD/SIMT	
Data formats	FP32/HF32/FP16/BF16/INT8	FP32/HF32/FP16/BF16/FP8/MXFP8/HiF8/MXFP4	
Interconnect bandwidth	784 GB/s	2 TB/s	
Computing power	800 TFLOPS FP16	1 PFLOPS FP8, 2 PFLOPS FP4	
Memory	128 GB, 3.2TB/s	128 GB, 1.6TB/s	144 GB, 4TB/s
Proprietary HBMs		HiBL 1.0	HiZQ 2.0

Source: Huawei (Company website), J.P. Morgan.

At the system level, supernode architectures, exemplified by Huawei CloudMatrix384, Alibaba Panjiu AL128, and Sugon scaleX, are increasing scale-up density to offset individual chip performance gaps through superior y.

On the software front, proprietary toolkits including Huawei CANN, Cambricon NeuWare, MetaX MXMACA and other proprietary software stacks are being developed to optimize hardware utilization. Looking ahead, we expect players with robust inventory positions, secured capacity, and demonstrated software optimization capability are more likely to outperform in the near term, while product performance and cost efficiency will be the key differentiators in the medium-to-long term.

Mapping the indigenous AI supply chain

We map out key beneficiaries of AI demand proliferation below (Figure 6) with focus on infrastructure in China, covering upstream suppliers of tools/services and manufacturing capacity, midstream component/assembly vendors of server/switch equipment, and downstream datacenter hosting partners and suppliers for power generation/storage.

Under JPM coverage, we would like to highlight below names across the ecosystem, and key theses are discussed in the Key Beneficiaries section:

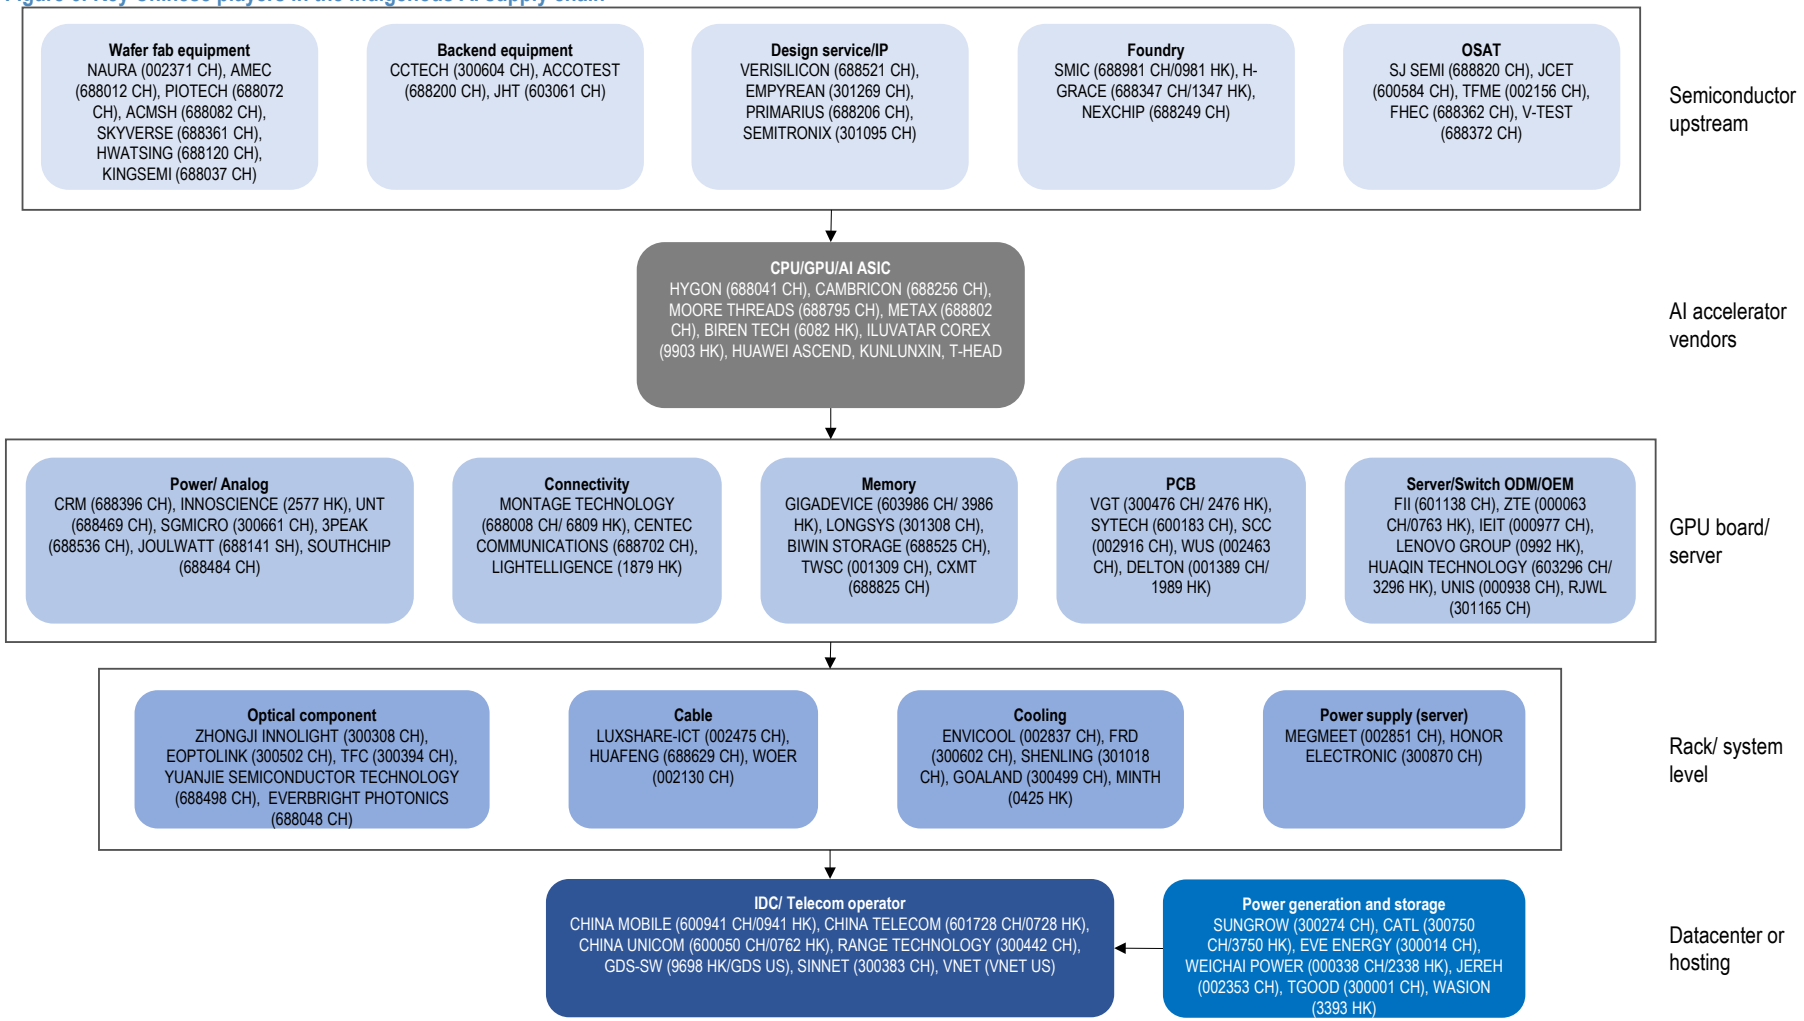
Table 2: Top picks - Valuation

Company	Ticker	Price (LC)	Price target (LC)	Mkt. cap ($, mn)	Rating	Return 2026	PE		PB		ROE (%)		Price upside (%)
							FY26E	FY27E	FY26E	FY27E	FY26E	FY27E	
Iluvatar CoreX	9903 HK	380	920	13,045	OW	163%	NM	92.5	16.2	13.8	-10%	16%	142%
NAURA	002371 CH	668	830	72,078	OW	46%	62.0	34.7	10.6	8.2	19%	27%	24%
AMEC	688012 CH	351	420	49,943	OW	92%	52.6	47.0	9.8	8.3	22%	19%	20%
JCET	600584 CH	73	115	19,509	OW	99%	53.1	33.3	4.3	3.9	8%	12%	57%
V-Test	688372 CH	128	260	3,223	OW	18%	52.2	25.3	4.8	4.2	11%	18%	103%
Huaqin	603296 CH	79	118	17,672	OW	24%	21.0	16.4	3.5	3.0	19%	20%	49%
VNET	VNET US	6	15	1,771	OW	-26%	NM	33.4	2.8	1.7	-44%	7%	141%
Sungrow	300274 CH	89	140	27,549	OW	-48%	12.9	12.8	3.2	2.7	28%	23%	57%
CATL-A	300750 CH	358	520	251,310	OW	-1%	17.1	14.1	4.1	3.5	26%	27%	45%
TGOOD	300001 CH	34	42	5,361	OW	34%	23.0	18.2	3.7	3.1	17%	19%	23%

Source: J.P. Morgan estimates, Bloomberg Finance L.P. Note: share prices as of Sep 1, 2026.

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Figure 6: Key Chinese players in the indigenous AI supply chain



Source: J.P. Morgan. Note: an interactive dashboard (link) to assist with the screening of the indigenous supply chain in China

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

J.P.Morgan

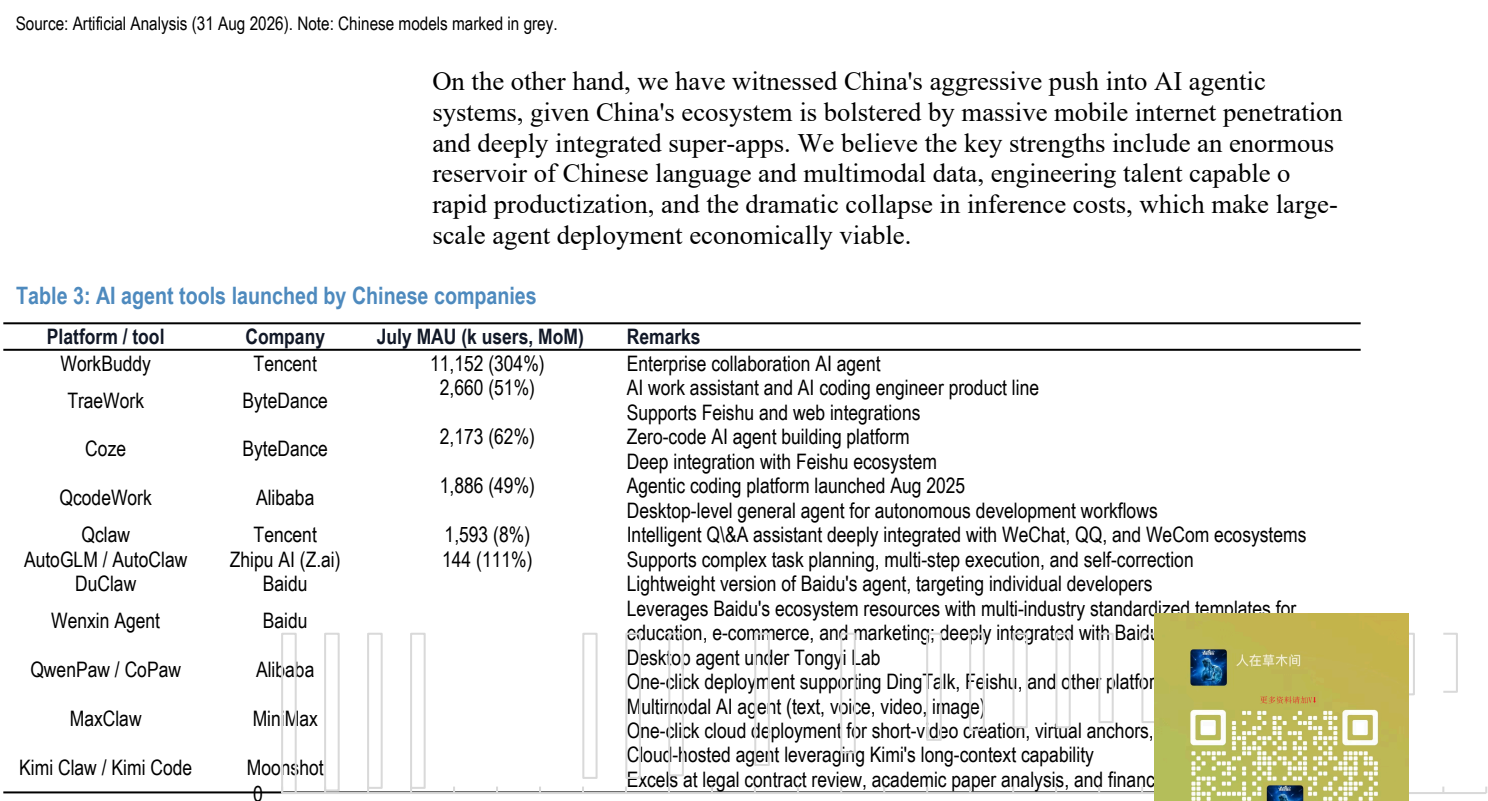
Exponential token growth driving domestic AI chip demand

Rapidly expanding AI inference token usage in China

Chinese AI models are increasingly capable

In the GenAI era, we have soar not observed a major disparity between Chinese players and global leaders in model capability, despite escalating hardware constraints. Daily token consumption has grown exponentially, even as the long-term proliferation of AI demand has yet to fully play out.

Figure 7: Artiicial Analysis Intelligence Index of select models



Source: Artificial Analysis (31 Aug 2026). Note: Chinese models marked in grey.

On the other hand, we have witnessed China's aggressive push into AI agentic systems, given China's ecosystem is bolstered by massive mobile internet penetration and deeply integrated super-apps. We believe the key strengths include an enormous reservoir of Chinese language and multimodal data, engineering talent capable o rapid productization, and the dramatic collapse in inference costs, which make large-scale agent deployment economically viable.

Table 3: AI agent tools launched by Chinese companies

Platform / tool	Company	July MAU (k users, MoM)	Remarks
WorkBuddy	Tencent	11,152 (304%)	Enterprise collaboration AI agent
TraeWork	ByteDance	2,660 (51%)	AI work assistant and AI coding engineer product line Supports Feishu and web integrations
Coze	ByteDance	2,173 (62%)	Zero-code AI agent building platform Deep integration with Feishu ecosystem
QcodeWork	Alibaba	1,886 (49%)	Agentic coding platform launched Aug 2025 Desktop-level general agent for autonomous development workflows
Qclaw	Tencent	1,593 (8%)	Intelligent Q\&A assistant deeply integrated with WeChat, QQ, and WeCom ecosystems
AutoGLM / AutoClaw	Zhipu AI (Z.ai)	144 (111%)	Supports complex task planning, multi-step execution, and self-correction
DuClaw	Baidu		Lightweight version of Baidu's agent, targeting individual developers
Wenxin Agent	Baidu		Leverages Baidu's ecosystem resources with multi-industry standardized templates for education, e-commerce, and marketing; deeply integrated with Baidu
QwenPaw / CoPaw	Alibaba		Desktop agent under Tongyi Lab One-click deployment supporting DingTalk, Feishu, and other platfor
MaxClaw	Min Max		Multimodal AI agent (text, voice, video, image) One-click cloud deployment for short-video creation, virtual anchors,
Kimi Claw / Kimi Code	Moonshot		Cloud-hosted agent leveraging Kimi's long-context capability Excels at legal contract review, academic paper analysis, and financ

Source: J.P. Morgan, Company data, AI CPB (17 Aug 2026).

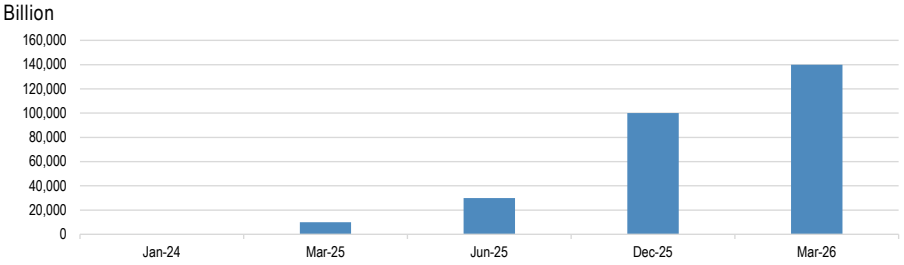
Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

We believe the agentic surge will fundamentally reshape token consumption dynamics, as agents engage in persistent, multi-turn sessions involving long-context reasoning and repeated tool-calling loops, causing per-server token volume to rise by orders of magnitude.

Token consumption surged on unprecedented inference demand

In early 2024, national daily token consumption was approximately 100 billion, before DeepSeek had launched. DeepSeek's late-2024 release then fundamentally transformed the industry ecosystem: with its low cost, lightweight architecture, and open-source characteristics, it lowered the AI application threshold from big-tech exclusive to SME-accessible, triggering a nationwide surge in inference compute demand. By end-2025, the figure had surged to 100 trillion and further breached 140 trillion by March 2026. In two years, daily token consumption grew by over 1,000x, not merely high growth, but a near-vertical expansion curve.

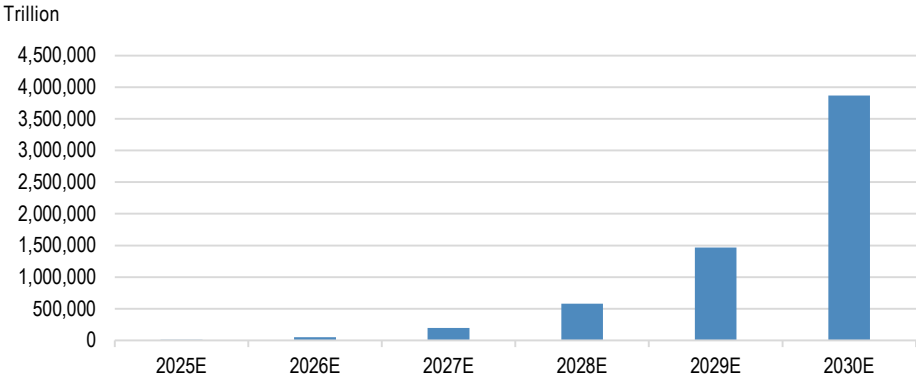
Figure 8: Daily token consumption in China



Source: National Data Administration (26 Mar 2026).

JPM China Internet team is forecasting (link) China token consumption to show a roughly 330% CAGR over 2025-30, driven by 1) rising penetration across both consumer users and enterprise workloads, and 2) use cases expanding from conversational AI toward agents and multimodal output (increasing tokens per session).

Figure 9: China inference token demand (annual)



Source: National Data Administration, J.P. Morgan estimates.

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

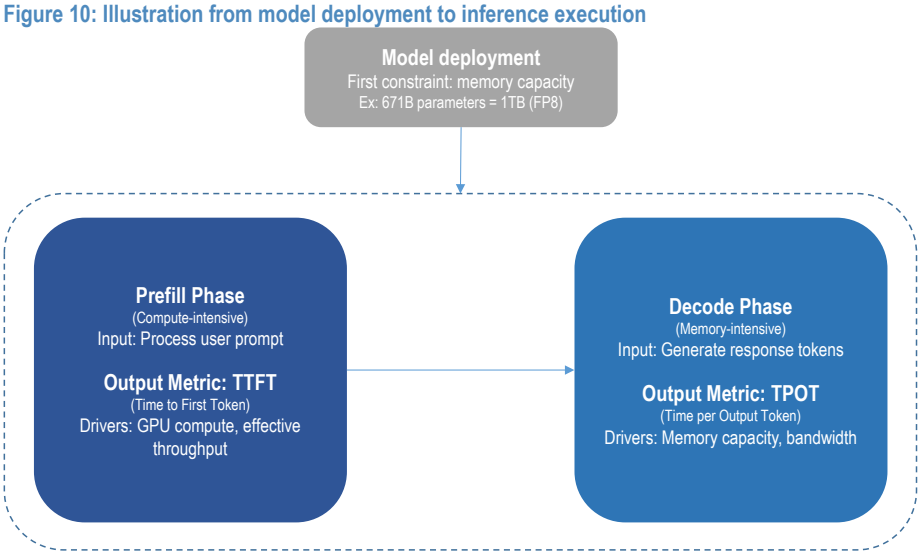
Consequently, the LLM development trajectory is pivoting from training-centric scaling laws toward inference-optimized architectures, such as longer context windows, efficient reasoning, and edge-cloud hybrid deployment, while shifting the bulk of AI infrastructure investment from training clusters to distributed and low-latency inference infrastructure.

Token demand supports 80% computing CAGR

Although token consumption (or processing) could consist of both training and inference demand, the proliferation of AI applications is driven by inference or external usage. We believe the first constraint on deploying large LLMs is memory capacity, which determines deployability. For example, the theoretical memory requirement of a DeepSeek R1 model with 671 billion parameters is approximately 1TB (for FP8). Inference speed or service level is primarily limited by GPU computing power and hardware utilization, which is in turn determined by interconnection and memory bandwidth.

The basics are: 1) the prefill phase of inference is compute-intensive, as it processes the user's input prompt and determines Time to First Token (TTFT), 2) the decode phase is memory-intensive, as it generates response tokens and determines service level or Time per Output Token (TPOT), and 3) in practice, a higher ratio of compute resources is allocated to the decode queue because decode has more intensive memory-capacity and memory bandwidth requirements.

Figure 10: Illustration from model deployment to inference execution



Source: J.P. Morgan.

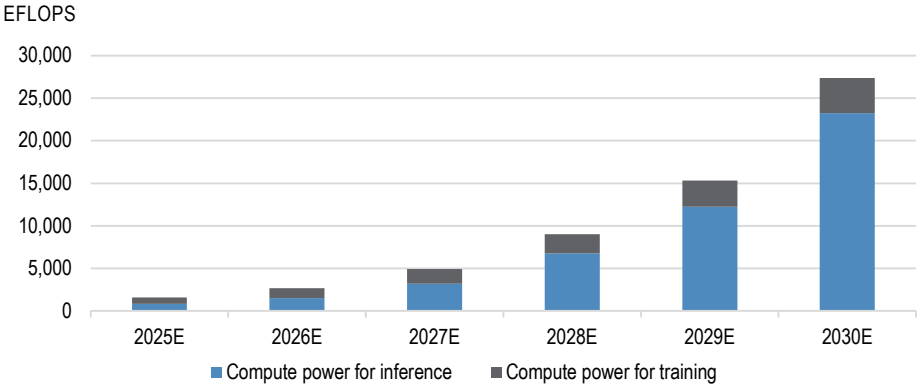
The simplified framework we use to quantify the computing capacity required for inference demand focuses on throughput (tokens/s) at a given batch size or concurrency level, and on throughput (tokens/s) per TFLOPS, which we view as a relatively practical approach. The key prerequisites include deploying a 671B model and meeting certain service levels, such as TPOT below 50ms, with a batch size of 128 and KV cache length of 4,096.

As rising token consumption is largely driven by inference, based on the token demand forecast, we expect the incremental compute power required for inference to

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

show an ~80% CAGR until 2030E (accounting for 85% of total compute power by then, remainder being allocated for Training). There is some leeway related to data format, model optimization, hardware utilization, and average model size deployed. On the other hand, the incremental demand from training should be marginal (increasing parameters and multimodal as key catalysts), as the preliminary requirements have been largely resolved by installed computing capacity, which could be released or reused.

Figure 11: Compute power required for inference and training



Source: J.P. Morgan estimates. Note: installation base

This translates into AI chip unit CAGR of 40%

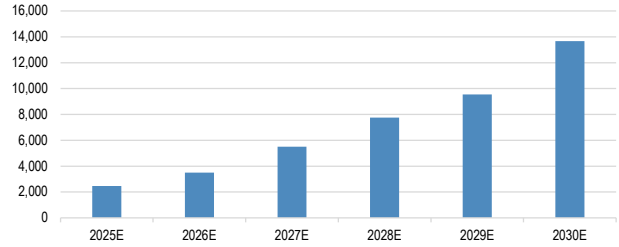
Our forecast of compute power that is needed to support the token demand growth, translates into a ~40% CAGR of AI chip demand during 2026-30E, assuming a single-die package and steady performance improvement. To gauge hardware utilization and algorithmic efficiency, we currently assume 1 Throughput (tokens/s) per TFLOPS and ~30% CAGR of efficiency gain to 3.3 by 2028E, driven largely by model architecture innovation and custom hardware optimization. For instance, despite larger parameters, DeepSeek V4 claims to reduce single-token inference compute (FLOPs) to 27% and memory (KV cache) usage down to 10% of V3 levels at long-context lengths through the new Hybrid Attention architecture.

With consideration of diminishing overseas GPU supply, we expect domestic leading-edge foundry capacity allocated to AI GPU/ASIC chips to reach 50kwpm by 2030. Given physical limits and manufacturing challenges, we believe foundry technology should continue to migrate to more advance process nodes, which is likely to keep yields at sub-optimal levels.

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Figure 12: Incremental chip demand (annual basis)

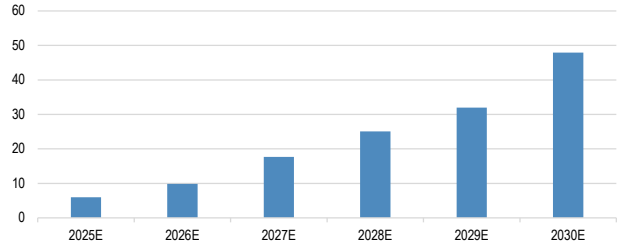
k units (700mm² equivalent)



Source: J.P. Morgan estimates.

Figure 13: Effective wafer capacity required for leading edge

kwpm



Source: J.P. Morgan estimates.

From a top-down view, we believe these forecasts are achievable, as Chinese players are targeting to add 150-200k of advanced capacity (in total) in the next 3 years. However, effective capacity, satisfactory yield achievement, and technology migration will likely remain challenging. Separately, we observe the increasing adoption of multi-die packages with smaller (single) die size to overcome yield limitations and improve output per wafer, which we believe would be a more efficient way to consume wafer capacity in the near term.

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Supply of computing chips the key bottleneck

Landscape of China's AI chip market

After Enflame filed for IPO on the STAR board in late January, all so-called 'Four Little Dragons' of Chinese GPU players are in increasing focus for public market investors. Together with increasing demand and investment for AI infrastructures, we believe we are now entering the phase of n or n of domestic GPUs or AI ASICs, the system-level hardware, as well as the ecosystem.

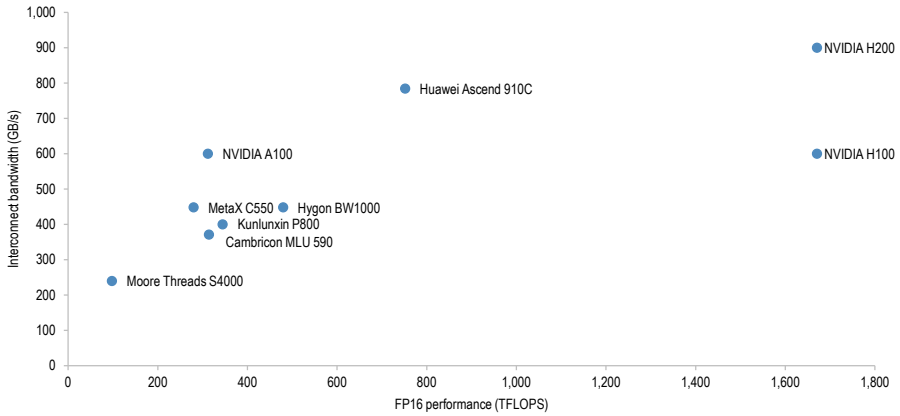
Figure 14: Key Chinese suppliers of AI chips separated by architecture

GPGPU	Non-GPGPU
Biren (6082 HK)	Cambricon (688256 CH)
Hygon (688041 CH)	Enflame
Iluvatar CoreX (9903 HK)	Huawei Ascend
MetaX (688802 CH)	Kunlunxin (Baidu)
Moore Threads (688795 CH)	T-head (Alibaba)

Source: Moore, MetaX, Biren, Enflame, Iluvatar IPO prospectus, J.P. Morgan

While comparing mainstream products provided by key Chinese companies with legacy NVIDIA chips, we believe the computing performance (FP16) of the majority is similar to A100 but y is still behind, which indicates a relatively low bandwidth across GPUs (limiting the utilization of computing power). Also, most existing AI chips on the market have not yet supported FP8 data format directly, even though we believe it would likely be introduced with newly designed chips.

Figure 15: Comparison of computing performance and interconnect bandwidth



Source: Moore, MetaX, Biren, Enflame, Iluvatar IPO prospectus, J.P. Morgan

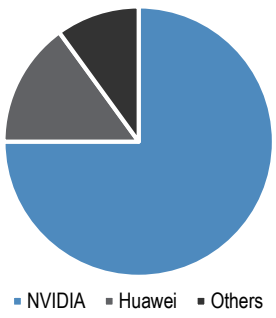
Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Thanks to the upcoming proliferation of AI applications, we expect China CSPs, telecom operators, and LLM enterprises to secure computing power proactively, largely for inference demand. Due to restraints on access to the most advanced AI chips and fabrication capacity, Chinese companies and institutions will increasingly need to rely on the indigenous AI supply chain.

However, given the difficulty of building leading-edge capacity and inferior production yield, we believe the availability or allocation of wafer fabrication would be the key bottleneck at the current stage. Therefore, AI computing chips including GPU, ASICs supplied by local companies should remain in a seller's market in the next 6 to 12 months, and chip players with secured capacity or safety inventory could be more preferred by customers in the near term, other than performance, compatibility and cost efficiency.

The lower-tier suppliers are small in business size, compared to the leading AI ASIC player Cambricon (688256 CH, Not Covered) and the hardware giant Huawei (private). But we expect substantial investment in AI hardware in the early stage of building the AI infrastructure to be unlocked for tier-2 players gradually, due to the limited availability of NVIDIA's advanced GPU products, which contributed a significant portion of China's chip consumption in the past.

Figure 16: Market share of AI computing chips in China by vendor (2024)



Source: Moore, MetaX, Biren, Enflame, Iluvatar IPO prospectus, J.P. Morgan

Figure 17: Market cap of key AI chip vendors in China

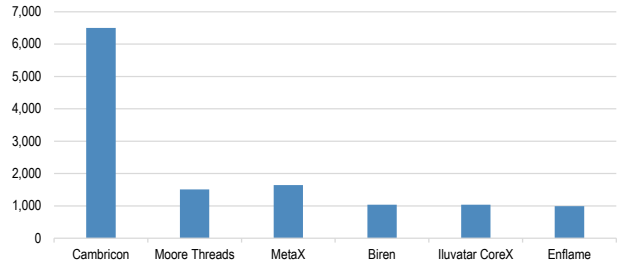
Rmb mn



Source: Wind (28 Aug 2026).

Figure 18: Revenue size of key AI chip vendors in China (2025)

Rmb mn



Source: 2025 annual reports (Cambricon, MetaX, Moore Threads, Biren, Iluvatar CoreX), IPO prospectus (Enflame).

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Currently, we believe telecom operators mainly collaborate with Huawei and CUDA-like solutions such as MetaX and Moore Threads, and Chinese CSP customers are also seeking to develop in-house AI chip solutions.

Both Alibaba and Baidu have discussed listing their own chip departments (T-head and Kunlunxin), Tencent invested in Enflame and also has a chip for cloud AI called Zixiao, and ByteDance is also working with some partners to develop its own AI ASIC.

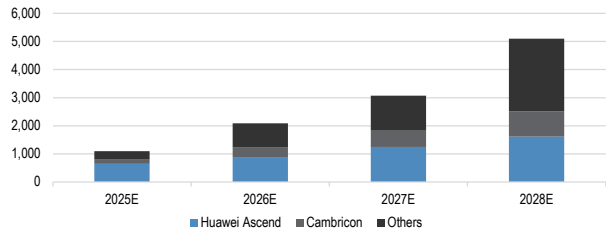
Foundry supply assumptions suggest tight supply into 2028

We provide the shipment forecast for AI chips supplied by domestic companies, and our numbers are based on the package of single-die (700mm^2 equivalent), which also include chips provided by local companies but fabricated at overseas foundries. We now expect 2mn/3mn/5mn AI chips including GPU and ASIC to be supplied in 2026/27/28E, even though the numbers are quite dependent on the size of compute die and package (ie dual-die). So we expect more effective (domestic) capacity allocation to other Chinese chip vendors from this year, although wafer-out should be expected from 2H26 onwards.

Our shipment forecast suggests approximately Rmb230bn/326bn/532bn AI server/hardware revenue in 2026/27/28E, on the assumption of a stabilized memory price and an increasing mix of supernode delivery. However, the approval to procure H200 or other overseas chips could introduce more uncertainties.

Figure 19: Shipment forecast of AI chips supplied by domestic companies

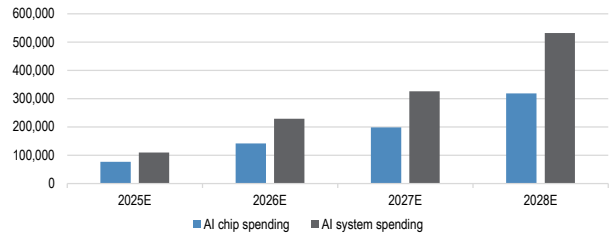
k units (700mm^2 equivalent)



Source: J.P. Morgan estimates.

Figure 20: Spending allocated to domestic AI chips/hardware

Rmb mn



Source: J.P. Morgan estimates.

We believe the supply/demand of domestic AI chips in China will remain tight in the next 1-2 years, although the gap (especially inference demand) could likely be filled through improved utilization/efficiency of hardware installed and software/model optimization. For instance, DeepSeek-V4-Pro claimed 27% of single-token inference FLOPS and 10% KV cache, compared to DeepSeek-V3.2 at a one-million-token context window.

Although demand from model training is expected to grow marginally, improving the efficiency of training and optimization remains a key challenge. In particular, most domestic AI chips currently available in China do not support lower-precision data format such as FP8/FP4. The limitation is likely to undermine the competitiveness of foundation models and slow the proliferation of AI applications.

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

At the same time, we also need to monitor customers' attitudes toward adding domestic AI hardware, including allocated capital spending, to assess the progress of optimization and fine-tuning and their satisfaction level with localized AI chips, given the increasing investment density and price increases of components such as memory and CPU.

Domestic foundry leading edge capacity ramp critical in the long-run

Despite some chips or wafers being fabricated at overseas fabs (mainly TSMC) with lower density of transistors or smaller die sizes, the majority will depend on domestic foundries' leading-edge capacity. While the flexibility of capacity addition is somewhat stringent, given the inferior production yield caused by a lack of crucial equipment (and materials) as well as the longer learning curve with varied customers and chip designs.

In the past several quarters, Cambricon has delivered significant revenue growth, and market consensus is forecasting over 140% revenue growth into 2026, following a five-fold expansion in 2025, largely attributed to previously piled-up inventory and the ramp-up of capacity allocation by domestic foundry partners. Though the release of advanced capacity by domestic foundries is likely at a very moderate pace, we believe a certain level of capacity reservation will provide high visibility on revenue generation (beyond inventory depletion) as well as smooth technology migration in the future.

Figure 21: Cambricon's quarterly revenue trend and inventory

Rmb mn

Source: Cambricon (Quarterly financial reports, 1Q22 – 2Q26).

Advanced packaging a key variable in performance optimization

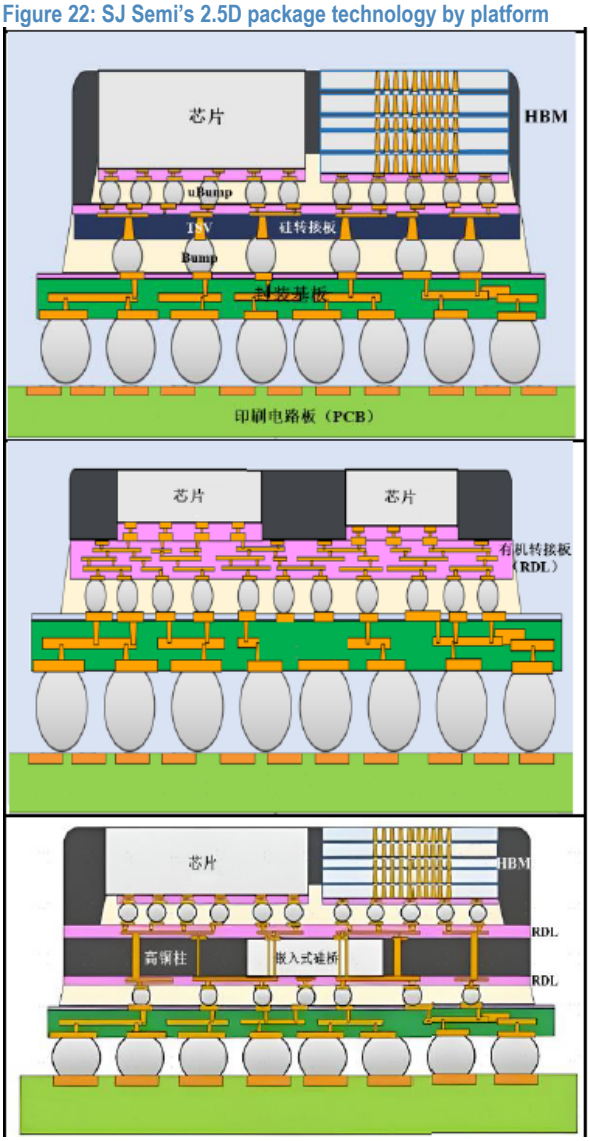
We provide our shipment forecast for domestic AI chips largely based on accessible wafer fabrication capacity; however, we expect CoWoS capability to materially impact output numbers and chip performance. Our supply chain study suggests domestic vendors (mainly OSATs) can handle packaging of compute dies, HBMs, or I/Os with CoWoS-S like solutions (or a mix of CoWoS-R), while the transition to CoWoS-L technology should be expected next year or later, to couple with customers' product roadmaps and clear engineering issues for mass production.

SJ Semiconductor (688820 CH, Not Covered), a leading advanced-packaging player in China, disclosed in its IPO prospectus published in October 2025 that its 2.5D

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

technology platforms include SmartPoser-Si, SmartPoser-RDL, and SmartPoser-BD. Our understanding is that SmartPoser-BD as a CoWoS-L like solution is still under verification. Given traditional OSAT players such as Tongfu (002156 CH, Not Covered), JCET (600584 CH), and Forehope (688362 CH, Not Covered) have claimed roadmaps and production readiness for advanced packaging (with heavy investment planned), we believe capacity may not be an issue to match wafer-out from upstream foundry, but the progress of technology transition should be closely monitored.

Figure 22: SJ Semi's 2.5D package technology by platform



SmartPoser-Si

SmartPoser-RDL

SmartPoser-BD

Source: SJ Semi (IPO prospectus).

Availability of competitive HBM chips another constraint

Memory is another key differentiator of AI chip performance, especially High Bandwidth Memory (HBM), which offers higher bandwidth for faster data access and retrieval and makes processing more efficient. While we believe HBM or HBM3/3e availability would be a potential bottleneck in near term, given:

1) The existing AI chips in China are mainly packaged with HBM2e, as the chips were designed 2-3 years ago, but the new chips are designed to pair with HBM3/3e.

2) Our supply chain check suggests CXMT's HBM3 product is still under qualification, so the availability is likely to depend on inventory depletion or sourcing from overseas vendors.

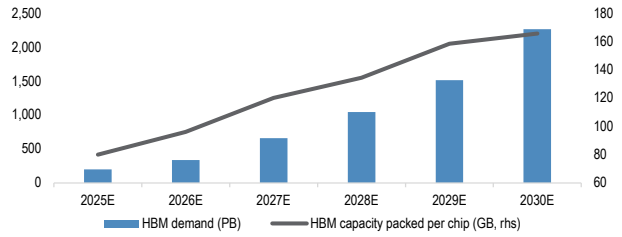
Table 4: Comparison of Samsung's HBM products

	Bandwidth (per stack)	Stack	Capacity option	Speed
HBM2e	410 GB/s	8H	16GB	up to 3.2 Gbps
HBM3	819 GB/s	12H, 8H	24GB, 16GB	up to 6.4 Gbps
HBM3e	1,180 GB/s	12H, 8H	36GB, 24GB	up to 9.2 Gbps

Source: Samsung (Company website), JP. Morgan.

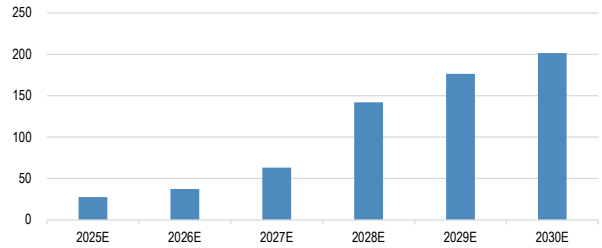
Our forecast suggests a 60% CAGR in HBM demand from 2026-2030E. We forecast a 60% CAGR in HBM demand from 2026-30E, based on chip demand and HBM capacity per compute die. Our underlying assumptions are: 1) domestic AI chips mainly adopt HBM2e now, and 2) the full transition to HBM3e should be expected from 2028 onwards. Our trial calculation indicates we will need approximately 200kwpm capacity for front-end wafer (HBM3e memory die) installed by 2030E to fully fulfill localized demand. Even though we have quite low visibility on domestic capacity status, the key challenge is product competitiveness rather than capacity expansion alone.

Figure 23: HBM demand forecast vs capacity per compute die



Source: J.P. Morgan estimates.

Figure 24: Effective wafer capacity required for HBM base wafer



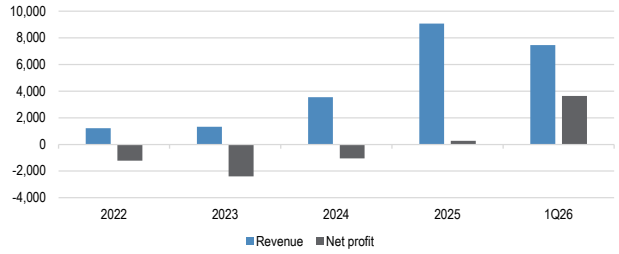
Source: J.P. Morgan estimates.

CXMT (688825 CH, Not Covered), the largest DRAM manufacturer in China and a single-digit percentage share player in the global market by value, was listed this July. Due to its high exposure to consumer electronics applications such as mobile devices and the increasing difficulty of procuring advanced equipment overseas, we do not expect the self-reliance gap in HBM products to close rapidly.

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Figure 25: CXMT's revenue and net profit trend

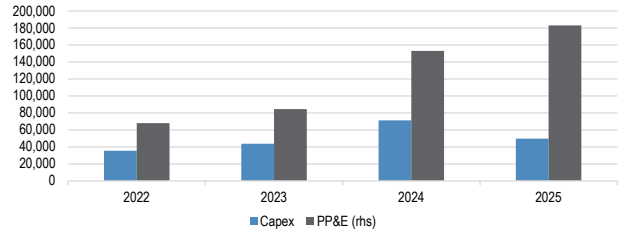
Rmb mn



Source: CXMT (IPO prospectus).

Figure 26: CXMT's Capex vs PP&E

Rmb mn



Source: CXMT (IPO prospectus).

In addition, Huawei announced two proprietary HBMs for its Ascend 950 series in September 2025: HiBL 1.0 (1.6TB/s bandwidth and 128GB total capacity) and HiZQ 2.0 (4TB/s bandwidth and 144GB total capacity) for prefill chip and decode chip respectively. Previously, Huawei also introduced a software solution or algorithm called Unified Cache Manager (UCM), which allocates data across memory types such as HBM, DRAM, and SSDs based on latency requirements. We believe this hierarchical strategy aims to improve the overall system's efficiency while reducing HBM consumption under constrained resources. The UCM solution has been open-source since November 2025, even though it might not be easy for external partners to deploy.

Adopting multi-die package for AI chip design offers upside to AI chip supply

Due to technology and equipment restrictions, domestic AI chips are being forced to accelerate adoption of chiplet architectures or multi-die packaging by splitting a large chip into multiple smaller dies, each manufactured at its optimal process node and then integrated into one package. We have also observed increasing adoption of multi-chip packages for newly launched or upcoming AI chips, with smaller die sizes rather than a single 700mm^2 die, to overcome yield limitations and improve output per wafer. Therefore, the importance of advanced packaging (2.5D/3D) has increased significantly, as communication efficiency among chips or dies materially relies on back-end processes.

Furthermore, Huawei recently presented Tau (τ) Scaling Law in May 2026, and LogicFolding is considered a proof point for τ Scaling, demonstrating that meaningful performance gains remain achievable even when geometric scaling and advanced lithography access are constrained. For Kirin 2026 (2-tier LogicFolding), transistor density is expected to improve to 238MTr/mm2 (+55% in one generation) with +41% SoC performance-core power efficiency. Based on transistor density alone, Kirin 2026 is demonstrated to be at a similar, or slightly lower, level than TSMC N3 node. The long-term targets are 400MTr/mm2 and beyond, on top of moving towards full-scale multi-layer folding. Beyond foundry availability (closely coupled manufacturing) and highly customized EDA tools, we believe incremental processes (such as hybrid bonding, TSV) and complexity would be introduced at back-end.

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

System-level efforts to deliver higher AI capability

Computing performance is not the only metric to measure domestic AI competitiveness

Looking beyond short-term capacity constraints, Chinese players are pursuing open-platform and open-source approaches to improve utilization and optimize limited resources, covering chip design, y, supernode architecture, software toolkits, and more. Key highlights include: 1) separating chip offerings for the prefill and decode phases of AI inference demand, with different memory products packaged; 2) increasing scale-up density, such as through supernode solutions, to avoid efficiency loss in scale-out networks; and 3) cultivating proprietary protocols and toolkits compatible with mainstream solutions to optimize hardware performance and improve user-friendliness.

Huawei Ascend 950 series

Instead of launching the 910D as the successor of 910C (as speculated), Huawei announced its novel product roadmap in September 2025, and its next generation of AI chips includes Ascend 950PR and Ascend 950DT, which use the same 950 compute die but different memory products. Although some may question the decision not to design the chip with higher transistor density or package multiple compute dies (i.e., four 910 dies vs. two for 910C), we view this as a desirable approach to optimize hardware performance under existing circumstances.

Similar to the concept of NVIDIA Rubin CPX, 950PR is dedicated to prefill and recommendation purposes, and 950DT with higher memory bandwidth (up to 4TB/s) will be used for the decode phase of inference and training. Given the prefill phase is compute-intensive with higher demand for parallel computing, Huawei introduced the proprietary HBM called HiBL 1.0 (claimed as cost-effective compared to HBM3e), while HiZQ 2.0 with higher memory capacity/bandwidth would be provided for 950DT, aiming at balancing performance and hardware investment required. In addition, FP8 data format will be included, targeting more efficient training of LLMs; similar practices are being observed at other AI Chip companies.

Table 5: Comparison of Huawei Ascend chips

	910C	950PR	950DT
Availability	1Q25	1Q26	4Q26
e	SIMD	SIMD/SIMT	
Data formats	FP32/HF32/FP16/BF16/INT8	FP32/HF32/FP16/BF16/FP8/MXFP8/HiF8/MXFP4	
Interconnect bandwidth	784 GB/s	2 TB/s	
Computing power	800 TFLOPS FP16	1 PFLOPS FP8, 2 PFLOPS FP4	
Memory	128 GB, 3.2TB/s	128 GB, 1.6TB/s	144 GB, 4TB/s
Proprietary HBMs		HiBL 1.0	HiZQ 2.0

Source: Huawei (Company website), J.P. Morgan.

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Huawei CloudMatrix384, Alibaba Cloud Panjiu AL128, and Sugon scaleX Supercluster

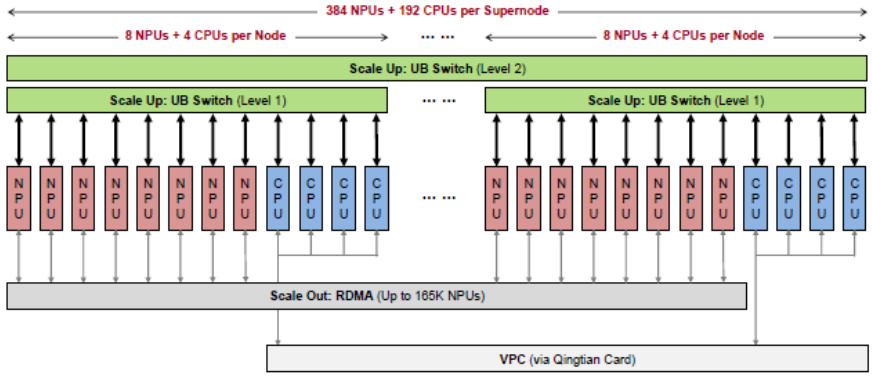
We would also like to briefly present how Huawei (chip/hardware vendor), Alibaba (CSP), and Sugon (supercomputer veteran) combine scale-up and scale-out with proprietary and open-source protocols/solutions.

- **CloudMatrix384:**

Launched in March 2025, a single supernode with over 300 PFLOPS computing power comprises 384 Ascend 910 chips and 192 Kunpeng CPUs, interconnected via Huawei's own network solution unified bus (UB). The CloudMatrix384 supernode spans 16 racks: 12 compute racks, which collectively host 48 Ascend 910 nodes, and 4 communication racks housing the second-tier (L2) UB switches that interconnect all nodes within the supernode.

The UB network enables direct all-to-all data exchange across all compute and memory components. This peer-to-peer serving architecture facilitates dynamic resource pooling, streamlined memory management, and better inter-node communication, effectively addressing the scalability and efficiency challenges common in traditional datacenter architecture.

Figure 27: Peer-to-peer hardware architecture of a CloudMatrix384 supernode



Source: Huawei (Serving Large Language Models on Huawei CloudMatrix384).

Huawei has developed the software ecosystem named compute architecture for neural networks (CANN) for its chips like NVIDIA's CUDA, which functions as an intermediary software layer, enabling efficient integration between high-level AI frameworks (like PyTorch and TensorFlow) and the low-level hardware interfaces of Ascend chips.

At MWC Barcelona 2026, Huawei then unveiled the latest SuperPoD product Atlas 950 SuperPoD, which integrates 64 NPUs per cabinet and can scale up to 8,192 NPUs.

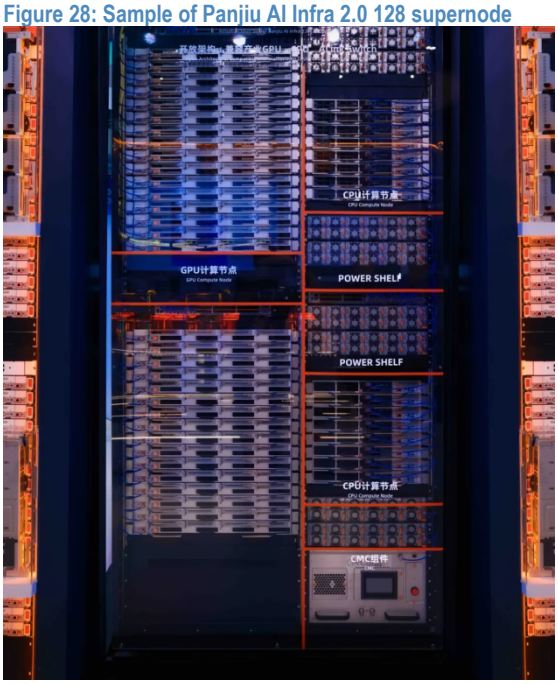
- **Panjiu AL128:**

Alibaba Cloud unveiled its new-generation Panjiu AI Infra 2.0 AL128 supernode AI server at the Apsara Conference 2025 last November. The customized double-wide

rack can support 128 to 144 GPUs (each GPU node contains four GPUs), and adopts a flexible, modular, multi-dimensional decoupling system architecture.

The supernode can adapt heterogeneous GPUs for different scenarios and nodes, which requires higher bandwidth and lower latency for inter-GPU communication. The system adopts the orthogonal connectors to connect ALink Switch chips, and it supports the international open standard UALink, and also supports the mainstream interconnect protocols such as NVLink, xLink, UB, and xCN.

Figure 28: Sample of Panjiu AI Infra 2.0 128 supernode



Source: Alibaba Cloud (Company website).

- **scaleX640 and Supercluster:**

Sugon showcased its 10,000+ GPU supercluster at HAIC2025 in December, deploying 10,240 AI accelerator cards with 16 ultra-nodes interconnected (640 GPUs per ultra-node), which supports multiple accelerator brands and mainstream computing ecosystems, and claims to deliver over 5 EFLOPS of total computing power and achieves a PUE as low as 1.04. The inter-communication across the ultra-nodes is based on scaleFabric, the self-developed 400G native RDMA high-speed network.

Pushing optimization through interconnected hardware and software

From a hardware perspective, Chinese companies are temporarily not pursuing significantly more powerful compute dies because of restrictions and physical limits, while they focus on improving performance efficiency to tackle real-world tasks (e.g., outperforming NVIDIA processors under specific conditions). Meanwhile, due to inferior performance of the single GPU or compute die, they are exploring ways of boosting utilization of computing power via increasing scale-up density with private or common network protocols or firmware, to offset efficiency loss from scale-out

connections. However, deployment of supernode or cluster would not be easier for customers, given the intensive capital spending on hardware and higher power capacity requiring redesign of traditional data centers (i.e., adding liquid cooling systems), and the incremental efficiency of cost/performance should also be further demonstrated.

From a software and ecosystem perspective, NVIDIA CUDA has been an insurmountable obstacle for Chinese companies given its dominance, as CSPs or data centers have to cohost various GPU clusters, and CUDA compatibility was the key enabler to make this cohosting happen. Nowadays, chip companies prefer to launch their proprietary software stacks, such as Huawei CANN (compute architecture for neural networks), Cambricon NeuWare, MetaX MXMACA, and Moore Threads MUSA, to optimize hardware performance for AI applications and enhance customer loyalty. Huawei also announced plans to open-source the CANN toolkit, although adoption still depends on how well CANN supports existing AI frameworks in the long run.

Another key variable beyond technology development in China is the preference on hardware, software, and ecosystem by domestic CSP customers. We believe Chinese CSPs are not just customers, they are also shaping the AI industry with hardware and software capability, and their acceptance of existing chips/hardware might trigger market reshuffle.

Structural shifts in cooling, power and energy stack from the AI infrastructure cycle

Compared to traditional data centers, hosting AI server drives a structural step-up in rack-level power density and total power capacity. The mainstream air- or fan-cooling paradigm is increasingly challenged, creating a durable shift toward liquid-cooling systems for server/switch equipment and building-level thermal management, motivated not only by thermal limits but by the structural need to lower operating costs and energy expense as power-per-rack rises.

Beyond cooling, higher requirements on more efficient power generation and large-scale energy storage are emerging as integral features of AI data center infrastructure rather than optional add-ons. Stability requirements, carbon-free mandates and operational ROI are converging to demand a re-architected power stack: prefabricated, factory-tested power modules that compress deployment timelines and improve reliability; a migration toward 800V DC and solid-state transformer (SST) architectures to serve higher-density, faster-to-deploy footprints; and a multi-layered backup-power ecosystem that spans conventional engines, fuel cells and battery-based ESS. In addition, the growing share of green power in data center energy mixes is structurally driving investment in distribution networks and smart-grid infrastructure to support dispatch and grid integration.

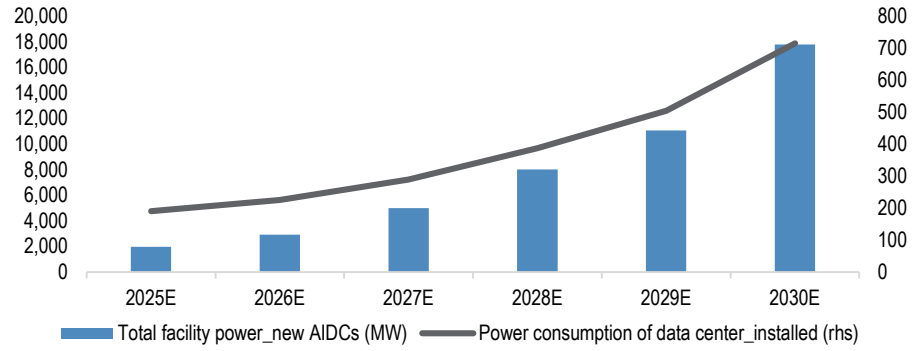
Though we believe aggregate power supply in China will not be a binding constraint, we are now forecasting a ~50% CAGR (2026–30E) for incremental power demand from AI servers/ data centers, with all data center usage representing ~5% of total power consumption in China by 2030E. This magnitude underscores a structural re-coupling of compute and energy infrastructure: AI buildout is no longer merely a demand driver for semiconductors and servers, but a catalyst for a full-stack

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

infrastructure super-cycle spanning advanced thermal solutions, next-generation power architectures, captive energy storage and intelligent grid integration.

Figure 29: Incremental power demand of AIDC vs power usage of data centers in China

MW, TWh



Source: CAICT, NEA, J.P. Morgan estimates.

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Key beneficiaries of China's indigenous AI supply chain

Our previous mapping of key beneficiaries of AI demand proliferation with focus on infrastructure in China covers upstream tools/services and manufacturing capacity (WFE, foundry, advanced packaging, testing), midstream component/assembly vendors (server/switch OEM/ODM), and downstream datacenter hosting partners alongside power generation, storage, cooling, and smart-grid infrastructure.

We prefer domestic AI chip suppliers with sufficient inventory and capacity to produce leading edge AI chips in the medium term. Key enablers such as advanced logic/memory WFEs, foundries/back-end partners, and server/switch OEM/ODMs should outperform. Concurrently, wholesale IDC demand remains robust, supported by high-power AI cabinet deployments and tighter power-quota screening that sustains a benign pricing environment.

In the longer run, the increasing complexity of AI hardware systems, growing localization of peripheral chips/components, and the architectural shift toward low-latency clustering will sustain supply-chain momentum. This extends beyond semiconductors to 800V DC/SST-based power architectures and prefabricated power solutions, AI server liquid cooling, and ESS/backup power, while green-power adoption drives smart-grid investment.

Table 6: Potential beneficiaries under JPM coverage universe

Sector	Company name	Ticker	JPM rating	JPM analyst	Market cap (US$ mn)	ADV (US$mn)	Last close (trading currency)	P/E (2027E)	Share price return	
									5D	YTD
Semiconductor	Iluvatar CoreX	9903 HK	OW	Billy Feng	13,045	155	380	92.5	4%	163%
Semiconductor	MetaX	688802 CH	N	Billy Feng	40,039	240	673	499.1	2%	16%
Semiconductor	NAURA	002371 CH	OW	Billy Feng	72,078	1310	668	34.7	-5%	46%
Semiconductor	AMEC	688012 CH	OW	Billy Feng	49,943	1554	351	47.0	-5%	92%
Semiconductor	SMIC	981 HK	N	Gokul Hariharan	100,586	1127	69	45.4	-1%	-4%
Semiconductor	JCET	600584 CH	OW	Billy Feng	19,509	2353	73	33.3	0%	99%
Semiconductor	V-Test	688372 CH	OW	Billy Feng	3,223	223	128	25.3	6%	18%
Server / Datacenter	Huaqin	603296 CH	OW	Cherry Liu	17,672	359	79	16.4	3%	24%
Server / Datacenter	VNET	VNET US	OW	Albert Hung	1,771	40	6	33.4	-10%	-26%
Cooling	Envicool	002837 CH	UW	Alan Hon	12,520	574	66	80.1	6%	-20%
Cooling	Minth	425 HK	OW	Rebecca Wen	3,807	19	25	7.2	-14%	-18%
Power & ESS	Sungrow	300274 CH	OW	Alan Hon	27,549	1245	89	12.8	-21%	-48%
Power & ESS	CATL-A	300750 CH	OW	Rebecca Wen	251,310	1850	358	14.1	-6%	-1%
Power & ESS	Weichai Power-A	000338 CH	OW	Karen Li	35,078	465	27	13.3	1%	61%
Electrical equipment	Wasion	3393 HK	OW	Stephen Tsui	2,046	13	15	9.8	-11%	-9%
Electrical equipment	TGOOD	300001 CH	OW	Stephen Tsui	5,361	156	34	18.2	-5%	34%
Smart Grid / Green Power	Nari Tech	600406 CH	OW	Stephen Tsui	27,384	220	23	17.9	-2%	4%
Smart Grid / Green Power	Xuji Electric	000400 CH	OW	Stephen Tsui	3,256	68	21	14.5	-4%	-16%
Average									**-3%**	**23%**

Source: J.P. Morgan estimates, Bloomberg Finance L.P. (market close as of Sep 1, 2026).

Gokul Hariharan
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Local semiconductor supply chain set to benefit immediately

We expect **AI chip fabless** companies to be key beneficiaries of domestic AI infrastructure buildout in China (link to report), thanks to 1) opportunities to commercialize products in a market historically dominated by overseas companies, 2) soaring demand driven by accelerating AI proliferation across LLMs and agentic AI, and 3) likely continued controlled access to advanced chips from overseas. Owing to US restrictions, we expect companies with high inventory buffers or sufficient capacity support and customer breakthroughs at CSPs to outperform. We are positive on Iluvatar CoreX (9903 HK, OW, covered by Billy Feng) and MetaX (688802 CH, N, covered by Billy Feng), thanks to their superior product and supply chain strategies.

Given manufacturing capability or capacity as the key bottleneck, we believe **WFE** names in China could sustain the strong sales momentum supported by both continuous addition of advanced logic and memory capacity (also higher density of capital spending) and increasing localization rate. We believe NAURA (002371 CH, OW, covered by Billy Feng) and AMEC (688012 CH, OW, covered by Billy Feng) could benefit from the upbeat demand.

Leading edge foundry players in China such as SMIC (981 HK, N, covered by Gokul Hariharan) are key enablers for the buildout of advanced logic manufacturing and should continue to add leading edge foundry production. A key challenge for vendors like SMIC remains low yield rates (which keeps blended ASPs at quite low levels despite rising leading edge Foundry production) and rapidly rising depreciation costs, which keep Gross Margins under check (link to report). As a result, we remain Neutral on SMIC, and would recommend investors to gain more exposure through capex plays and OSAT vendors.

Chinese **OSAT** players tend to provide full-stack solutions with CoWoS-like technology. Even though the transition to CoWoS-L-like technology remains a key uncertainty, we expect top OSAT players to benefit from the buildout of the local AI supply chain and higher value-added services, and we believe JCET (600584 CH, OW, covered by Billy Feng and Ri Xu) is gaining share (link to report). In addition, we believe **third-party testing** service providers such as V-Test (688372 CH, OW, covered by Billy Feng and Ri Xu) are also favored (link to report), thanks to the rapidly increasing volume and testing time of high-end AI chips manufactured by local foundries. Inferior manufacturing yield and accelerating adoption of multi-chip packages will likely expand the multiplier effect for the testing market.

We also provide an interactive dashboard of China Semiconductor Ecosystem (link) to assist with the screening of the indigenous supply chain in China.

Sourcing accessory chips and peripherals is not difficult, but localization is likely to increase

Because of insufficient local capacity and restricted access to overseas supply, we believe AI computing chips and HBM could be the key bottlenecks for governments, institutions, SOEs, and enterprises to build AI infrastructure and power AI innovation in the visible future. With that said, the availability of other accessory chips or components to construct AI server/clusters would not be an issue, if there is no further escalation on import/export regulations, in our view.

For example, power and analog chips are largely sourced from global IDMs, and connectivity chips are mainly supplied by international industry leaders. Moreover, Chinese vendors of other key components like PCB, optical transceivers, and so on are already on the list of suppliers to top-tier US (CSP/hardware) customers. Particularly, AI hardware in China is not necessarily adopting the most advanced products due to less competitive computing chips or smaller cluster scale. Especially, we observe rising interest by Chinese customers in NPO as the next-generation solution (3.2T and beyond) for scale-up demand, rather than closely decoupled CPO solutions.

On the other hand, we believe increasing self-sufficiency in accessory chips or components should be well expected, given Chinese power/analog IDM/fabless companies are actively seeking business expansion into higher-margin AI-related products, as they have swiftly gained shares on consumer and automotive applications (from overseas peers), which has made the competitive landscape in the China market quite intensified. However, the progress of localizing connectivity or networking chips might not be as smooth, due to fewer players with mature or commercialized products.

Also, we should bear in mind that Huawei is a comprehensive hardware company, which also owns one of the leading domestic IC design houses HiSilicon offering a broad product portfolio and its investment subsidiary Hubble with extensive coverage across the supply chain.

Ecosystem-wide efforts to facilitate China's data center infrastructure build-out

Servers

Driven by robust domestic CSPs' capex and improving localized GPU supply, we project structural growth for China **server ODM/OEM** industry. Concurrently, we anticipate a pronounced architectural shift toward "Superpod" deployments, which is designed to minimize inter-node latency and effectively mitigate individual GPU performance gaps through low-latency clustering. This architecture evolution demands deeper cross-domain hardware integration between high-density compute platforms and ultra-low latency networking. Underpinned by a rare combination of full-stack R&D and manufacturing across AI servers, general servers, and high-speed switches, Huaqin (603296 CH, OW covered by Cherry Liu) commands a superior competitive moat to emerge as a primary beneficiary of this infrastructure super-cycle.

Data centers

We have seen strong wholesale IDC momentum across China **IDC** vendors over the past 12–18 months, driven by robust domestic demand, evidenced by strong order bookings, improving visibility, and faster customer move-ins. Management at VNET indicated ~70%+ of current new orders are AI-related. We believe high-power AI cabinet deployments should support better pricing and margins, while tighter power quota screening by regulators (e.g., NDRC) is helping sustain a more benign pricing environment, which is supportive for leading players. Overall, we believe secular AI-driven server demand can sustain earnings momentum for China IDC vendors over the next 2–3 years. Within the space, VNET (VNET US, OW, covered by Albert Hung, link to report) is our top pick. Key downside risks include domestic AI chip supply constraints and lower-than-expected China CSP capex.

Cooling solutions and systems

We believe cooling solution providers may benefit from AI data center deployment in China, as they are able to provide room cooling solutions to AI datacenters and electronics cooling solutions to AI chipsets/servers, both of which are incremental to their earnings.

Envicool (002837 CH) is a major cooling solutions provider in China, focusing on data center and energy storage cooling solutions, while it has also expanded into electronics cooling solutions. Into 1Q26 results, Envicool's earnings declined by 82% yoy, driven by a decline in GPM (-2.2pcts yoy) and elevated operating expenses. This partly confirmed our view that it is difficult for Envicool to obtain market share in global server thermal market. Meanwhile, the margin decline suggests that the competitive landscape of domestic data center and energy storage cooling business may have deteriorated. Thus, we cut FY26-27 earnings by 9-15%, factoring in the weak 1Q26 operating results. Our current assumption implied Envicool having 3%/6% market share in the global server cooling market in 2027/28, while the current share price still implies a 25-30% market share in 2027/28 in our view. We maintain our UW rating on Envicool (Covered by Alan Hon, link to report).

We also expect Minth (425 HK) to benefit from AI server liquid cooling because of 1) its first batches of AI server liquid cooling product orders from Taiwan-based AI server manufacturers, and 2) strong synergies between its automotive battery housing and liquid cooling capabilities. It recently raised its AI liquid cooling revenue guidance to Rmb300mn in 2026 and Rmb1.5bn in 2027. We believe Minth is well positioned because its products can achieve high performance while being clean, leak-free, safe, and reliable, and the company is actively developing opportunities with Taiwan AI server ODMs. Minth is our top pick in China auto parts (OW, covered by Rebecca Wen, link to report) and AI liquid cooling support a potential valuation re-rating.

Power and ESS

We expect data center deployment may potentially lift green power demand, which may be positive for energy storage system (ESS) manufacturers. Although China currently has sufficient installed power capacity relative to demand, the benefit of installing ESS in AIDCs seems incremental in China. Per our checks with industry experts, most new AIDCs in the US will require ESS for reliability and optimization (link to report), marking the natural role of ESS in AIDCs' power solutions. Per our discussion with Sungrow (300274 CH, OW, covered by Alan Hon), management framed AIDCs as a new opportunity set, where tight power supply creates a role for ESS in the DM, including configurations to enable "grid-friendly" data centers, support rapid-start of AIDCs, etc. We have an OW rating on Sungrow (link to report).

We expect CATL to be another key beneficiary, thanks to its superior ESS battery cell and system technologies. Its sodium-ion ESS is structurally well suited for AIDC use cases where power quality, response time, and redundancy are prioritized. Its strategic investment in VNET secures a captive, nationwide ESS deployment channel. We believe CATL can transform from a battery supplier into a strategic infrastructure partner for AI compute buildout through three moves: a 49% stake in Zhongheng Electric's parent for DC-side HVDC and IDC power architecture

expertise, the world's largest sodium-ion battery order (60GWh) with Hyperstrong targeting AIDC deployment from 4Q26, and a ~38.1% strategic stake in VNET Group to secure a captive, nationwide ESS deployment channel across 889MW of operational IDC capacity. CATL-H / A is our top pick in China's battery supply chain (3750 HK/300750 CH, OW, covered by Rebecca Wen, link to report).

We are also positive on Weichai Power (2338 HK/000338 CH, OW, covered by Karen Li), given its positioning across AIDC backup engines, gas engines, SOFC, and integrated onsite power solutions. We understand Weichai supplies diesel engines into domestic backup gensets serving customers including Alibaba, ByteDance, Huawei and telco companies, supporting orders as AI capex accelerates and customers value lead time plus local service. Weichai's AIDC engine shipments are expected to reach 3,500–4,000 units in 2026. SOFC capacity is targeted to ramp from 30MW in 2026 to 1GW by 2030, adding medium term upside if
n stays on track.

Electrical / power equipment

China's AI infrastructure buildout should also support demand for low-voltage electrical equipment, including distribution transformers, circuit breakers, etc. The shift toward 800V DC and SST-based power architectures in AI data centers should also accelerate the buildout of a localized AI power-equipment supply chain, as customers seek higher-density, faster-to-deploy and more cost-efficient power infrastructure. This could provide opportunities to domestic companies with exposure to the new architecture. TGOOD (300001 CH, OW, covered by Stephen Tsui) has exposure to this theme.

Pre-fabricated power solutions should be a key beneficiary of China's AIDC buildout, as AI data centers require faster deployment, higher power reliability and lower total system cost than traditional substation construction. Wasion (3393 HK, OW, covered by Stephen Tsui) and TGOOD are our preferred AIDC power plays given their direct leverage to these needs (link to report). Wasion supplies critical digital infrastructure including power distribution containers/modules, IT containers/skids, power distribution equipment and HVDC systems, while its prefabricated, factory-tested solutions help customers scale faster and improve reliability. It is also expanding into HVDC, data center BESS and liquid cooling, broadening its wallet-share potential. TGOOD is differentiated in high-voltage prefabricated infrastructure, with AI Powerhouse integrating key substation equipment into factory-built modules to reduce delivery time, capex and power costs. Together, they offer strong exposure to the ramping local AI supply chain and AIDC power infrastructure buildout.

In addition to data center buildouts, rising adoption of green power by China's data centers should drive further investment in distribution networks and smart-grid infrastructure to support power dispatch and integration. Potential beneficiaries include Nari Tech (600406 CH, OW, covered by Stephen Tsui) and Xuji Electric (000400 CH, OW, covered by Stephen Tsui).

Analyst Certification: The Research Analyst(s) denoted by an "AC" on the cover of this report certifies (or, where multiple Research Analysts are primarily responsible for this report, the Research Analyst denoted by an "AC" on the cover or within the document individually certifies, with respect to each security or issuer that the Research Analyst covers in this research) that: (1) all of the views expressed in this report accurately reflect the Research Analyst's personal views about any and all of the subject securities or issuers; and (2) no part of any of the Research Analyst's compensation was, is, or will be directly or indirectly related to the specific recommendations or views expressed by the Research Analyst(s) in this report. For all Korea-based Research Analysts listed on the front cover, if applicable, they also certify, as per KOFIA requirements, that the Research Analyst's analysis was made in good faith and that the views reflect the Research Analyst's own opinion, without undue influence or intervention.

All authors named within this report are Research Analysts who produce independent research unless otherwise specified. In Europe, Sector Specialists (Sales and Trading) may be shown on this report as contacts but are not authors of the report or part of the Research Department.

Important Disclosures

Company-Specific Disclosures: J.P. Morgan does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. Important disclosures, including price charts and credit opinion history tables (if applicable), are available for compendium reports and all J.P. Morgan–covered companies, and certain non-covered companies, by visiting https://www.jpmm.com/research/disclosures, calling 1-800-477-0406, or e-mailing research.disclosure.inquiries@jpmorgan.com with your request.

Explanation of Equity Research Ratings, Designations and Analyst(s) Coverage Universe:
J.P. Morgan uses the following rating system: Overweight (over the duration of the price target indicated in this report, we expect this stock will outperform the average total return of the stocks in the Research Analyst's, or the Research Analyst's team's, coverage universe); Neutral (over the duration of the price target indicated in this report, we expect this stock will perform in line with the average total return of the stocks in the Research Analyst's, or the Research Analyst's team's, coverage universe); and Underweight (over the duration of the price target indicated in this report, we expect this stock will underperform the average total return of the stocks in the Research Analyst's, or the Research Analyst's team's, coverage universe. NR is Not Rated. In this case, J.P. Morgan has removed the rating and, if applicable, the price target, for this stock because of either a lack of a sufficient fundamental basis or for legal, regulatory or policy reasons. The previous rating and, if applicable, the price target, no longer should be relied upon. An NR designation is not a recommendation or a rating. Some stocks under coverage have a rating but no price target; in these cases, we expect the stock will outperform/perform in line/underperform the average total return of the stocks in the Research Analyst's, or the Research Analyst's team's, coverage universe of the relevant duration of the region. In our Asia (ex-Australia and ex-India) and U.K. small- and mid-cap Equity Research, each stock's expected total return is compared to the expected total return of a benchmark country market index, not to those Research Analysts' coverage universe. If it does not appear in the Important Disclosures section of this report, the certifying Research Analyst's coverage universe can be found on J.P. Morgan's Research website, https://www.jpmorganmarkets.com.

Coverage Universe: Hariharan, Gokul : ASE Technology Holding Co Ltd (3711.TW), ASMPT Ltd (0522) (0522.HK), AirTAC (1590.TW), Alchip Technologies (3661.TW), Chipbond Technology (6147.TWO), GDS Holdings (GDS), GUC (3443.TW), Hiwin (2049.TW), Hon Hai Precision (2317.TW), MediaTek Inc. (2454.TW), Nien Made (8464.TW), Novatek Corp. (3034.TW), Powerchip Semiconductor Manufacturing Corp. (6770.TWO), SMIC (0981) (0981.HK), Silicon Motion (SIMO), Sinbon (3023.TW), TSMC (2330.TW), Teco Electric & Machinery (1504.TW), UMC (2303.TW), Vanguard International Semiconductor Corp. (5347.TWO), Voltronic Power Technology (6409.TW), Xiaomi (1810) (1810.HK)

J.P. Morgan Equity Research Ratings Distribution, as of July 04, 2026

	Overweight (buy)	Neutral (hold)	Underweight (sell)
J.P. Morgan Global Equity Research Coverage*	53%	36%	12%
IB clients**	83%	80%	73%
JPMS Equity Research Coverage*	51%	37%	12%
IB clients**	95%	92%	87%

*Please note that the percentages may not add to 100% because of rounding.
**Percentage of subject companies within each of the "buy," "hold" and "sell" categories for which J.P. Morgan has provided investment banking services within the previous 12 months.
For purposes of FINRA ratings distribution rules only, our Overweight rating falls into a buy rating category; our Neutral rating falls into a hold rating category; and our Underweight rating falls into a sell rating category. Please note that stocks with an NR designation are not included in the table above. This information is current as of the end of the most recent calendar quarter.

Equity Valuation and Risks: For valuation methodology and risks associated with covered companies or price targets for covered companies, please see the most recent company-specific research report at http://www.jpmorganmarkets.com, contact the primary analyst or your J.P. Morgan representative, or email research.disclosure.inquiries@jpmorgan.com. For material information about the proprietary models used, please see the Summary of Financials in company-specific research reports and the Company Tearsheets, which are available to download on the company pages of our client website, http://www.jpmorganmarkets.com. This report also sets out within it the material underlying assumptions used.

History of Investment Recommendations:
A history of J.P. Morgan investment recommendations disseminated during the preceding 12 months can be accessed on the Research & Commentary page of http://www.jpmorganmarkets.com where you can also search by analyst name, sector or financial instrument.

Analysts' Compensation: The research analysts responsible for the preparation of this report receive compensation based upon various factors, including the quality and accuracy of research, client feedback, competitive factors, and overall firm revenues.

Registration of non-US Analysts: Unless otherwise noted, the non-US analysts listed on the front of this report are employees of non-US affiliates of J.P. Morgan Securities LLC, may not be registered as research analysts under FINRA rules, may not be associated persons of J.P. Morgan Securities LLC, and may not be subject to FINRA Rule 2241 or 2242 restrictions on communications with covered companies, public appearances, and trading securities held by a research analyst account.

Other Disclosures

J.P. Morgan is a marketing name for investment banking businesses of JPMorgan Chase & Co. and its subsidiaries and affiliates worldwide.

UK MIFID FICC research unbundling exemption: UK clients should refer to UK MIFID Research Unbundling exemption for details of J.P. Morgan's implementation of the FICC research exemption and guidance on relevant FICC research categorisation.

All research material made available to clients are simultaneously available on our client website, J.P. Morgan Markets, unless specifically permitted by relevant laws. Not all research content is redistributed, e-mailed or made available to third-party aggregators. For all research material available on a particular stock, please contact your sales representative.

Any long form nomenclature for references to China; Hong Kong; Taiwan; and Macau within this research material are Mainland China; Hong Kong SAR (China); Taiwan (China); and Macau SAR (China).

J.P. Morgan Research may, from time to time, write on issuers or securities targeted by economic or financial sanctions imposed or administered by the governmental authorities of the U.S., EU, UK or other relevant jurisdictions (Sanctioned Securities). Nothing in this report is intended to be read or construed as encouraging, facilitating, promoting or otherwise approving investment or dealing in such Sanctioned Securities. Clients should be aware of their own legal and compliance obligations when making investment decisions.

Digital assets, which may also be referred to as virtual or crypto assets or crypto currencies (e.g. bitcoin), discussed in this research report are subject to a rapidly changing regulatory landscape. The application of relevant laws in respect of such assets may be unsettled, and the legality of any digital asset investment may be unclear. Prior to making any investment decision regarding any digital asset, investors should seek their own legal advice and be aware of any legal considerations regarding the purchase, holding or selling of such an investment. For relevant regulatory advisories on digital assets, including bitcoin and ethereum, please see https://www.jpmorgan.com/disclosures/cryptoasset-disclosure.

The author(s) of this research report may not be licensed to carry on regulated activities in your jurisdiction and, if not licensed, do not hold themselves out as being able to do so.

Exchange-Traded Funds (ETFs): J.P. Morgan Securities LLC ("JPMS") acts as authorized participant for substantially all U.S.-listed ETFs. To the extent that any ETFs are mentioned in this report, JPMS may earn commissions and transaction-based compensation in connection with the distribution of those ETF shares and may earn fees for performing other trade-related services, such as securities lending to short sellers of the ETF shares. JPMS may also perform services for the ETFs themselves, including acting as a broker or dealer to the ETFs. In addition, affiliates of JPMS may perform services for the ETFs, including trust, custodial, administration, lending, index calculation and/or maintenance and other services.

Options and Futures related research: If the information contained herein regards options- or futures-related research, such information is available only to persons who have received the proper options or futures risk disclosure documents. Please contact your J.P. Morgan Representative or visit https://www.theocc.com/components/docs/riskstoc.pdf for a copy of the Option Clearing Corporation's

Characteristics and Risks of Standardized Options or https://www.finra.org/sites/default/files/2020-08/Security_Futures_Risk_Disclosure_Statement_2020.pdf for a copy of the Security Futures Risk Disclosure Statement.

Changes to Interbank Offered Rates (IBORs) and other benchmark rates: Certain interest rate benchmarks are, or may in the future become, subject to ongoing international, national and other regulatory guidance, reform and proposals for reform. For more information, please consult: https://www.jpmorgan.com/global/disclosures/interbank_offered_rates

Private Bank Clients: Where you are receiving research as a client of the private banking businesses offered by JPMorgan Chase & Co. and its subsidiaries ("J.P. Morgan Private Bank"), research is provided to you by J.P. Morgan Private Bank and not by any other division of J.P. Morgan, including, but not limited to, the J.P. Morgan Corporate and Investment Bank and its Global Research division.

Legal entity responsible for the production and distribution of research: The legal entity identified below the name of the Reg AC Research Analyst who authored this material is the legal entity responsible for the production of this research. Where multiple Reg AC Research Analysts authored this material with different legal entities identified below their names, these legal entities are jointly responsible for the production of this research. Where more than one legal entity is listed under an analyst's name, the first legal entity is responsible for the production unless stated otherwise. Research Analysts from various J.P. Morgan affiliates may have contributed to the production of this material but may not be licensed to carry out regulated activities in your jurisdiction (and do not hold themselves out as being able to do so). Unless otherwise stated below in the legal entity disclosures, this material has been distributed by the legal entity responsible for production, or where more than one legal entity is listed under the analyst's name, the first legal entity will be responsible for distribution. If you have any queries, please contact the relevant Research Analyst in your jurisdiction or the entity in your jurisdiction that has distributed this research material.

Legal Entities Disclosures and Country-/Region-Specific Disclosures:
Argentina: JPMorgan Chase Bank N.A Sucursal Buenos Aires is regulated by Banco Central de la República Argentina ("BCRA"- Central Bank of Argentina) and Comisión Nacional de Valores ("CNV"- Argentinian Securities Commission - ALYC y AN Integral N°51).
Australia: J.P. Morgan Securities Australia Limited ("JPMSAL") (ABN 61 003 245 234/AFS Licence No: 238066) is regulated by the Australian Securities and Investments Commission and is a Market Participant of ASX Limited, a Clearing and Settlement Participant of ASX Clear Pty Limited and a Clearing Participant of ASX Clear (Futures) Pty Limited. This material is issued and distributed in Australia by or on behalf of JPMSAL only to "wholesale clients" (as defined in section 761G of the Corporations Act 2001). A list of all financial products covered can be found by visiting https://www.jpmm.com/research/disclosures. J.P. Morgan seeks to cover companies of relevance to the domestic and international investor base across all Global Industry Classification Standard (GICS) sectors, as well as across a range of market capitalisation sizes. If applicable, in the course of conducting public side due diligence on the subject company(ies), the Research Analyst team may at times perform such diligence through corporate engagements such as site visits, discussions with company representatives, management presentations, etc. Research issued by JPMSAL has been prepared in accordance with J.P. Morgan Australia's Research Independence Policy which can be found at the following link: J.P. Morgan Australia - Research Independence Policy.
Brazil: Banco J.P. Morgan S.A. is regulated by the Comissao de Valores Mobiliarios (CVM) and by the Central Bank of Brazil. Ombudsman J.P. Morgan: 0800-7700847 / 0800-7700810 (For Hearing Impaired) / ouvidoria.jp.morgan@jpmchase.com.
Canada: J.P. Morgan Securities Canada Inc. is a registered investment dealer, regulated by the Canadian Investment Regulatory Organization and the Ontario Securities Commission and is the participating member on Canadian exchanges. This material is distributed in Canada by or on behalf of J.P.Morgan Securities Canada Inc.
Chile: Inversiones J.P. Morgan Limitada is an unregulated entity incorporated in Chile.
China: J.P. Morgan Securities (China) Company Limited has been approved by CSRC to conduct the securities investment consultancy business.
Colombia: Banco J.P. Morgan Colombia S.A. is supervised by the Financiera de Colombia (SFC). Any reference in this material to products or services offered abroad by entities other than the Bank in Colombia is included exclusively for descriptive purposes. Such references do not constitute, and should not be construed as, promotional activity or the provision of financial products or services within Colombian territory, as defined under applicable Colombian regulation.
Dubai International Financial Centre (DIFC): JPMorgan Chase Bank, N.A., Dubai Branch is regulated by the Dubai Financial Services Authority (DFSA) and its registered address is Dubai International Financial Centre - The Gate, West Wing, Level 3 and 9 PO Box 506551, Dubai, UAE. This material has been distributed by JP Morgan Chase Bank, N.A., Dubai Branch to persons regarded as professional clients or market counterparties as defined under the DFSA rules.
European Economic Area (EEA): Unless specified to the contrary, research is distributed in the EEA by J.P. Morgan SE ("JPM SE"), which is authorised as a credit institution by the Federal Financial Supervisory Authority (Bundesanstalt für tungsaufsicht, BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB). JPM SE is a company headquartered in Frankfurt with registered address at TaunusTurm, Taunustor 1, Frankfurt am Main, 60310, Germany. The material has been distributed in the EEA to persons regarded as professional investors (or equivalent) pursuant to Art. 4 para. 1 no. 10 and Annex II of MiFID II and its respective implementation in their home jurisdictions ("EEA professional investors"). This material must not be acted on or relied on by persons who are not EEA professional investors. Any

investment or investment activity to which this material relates is only available to EEA relevant persons and will be engaged in only with EEA relevant persons.

Hong Kong: J.P. Morgan Securities (Asia Pacific) Limited (CE number AAJ321) is regulated by the Hong Kong Monetary Authority and the Securities and Futures Commission in Hong Kong, and J.P. Morgan Broking (Hong Kong) Limited (CE number AAB027) is regulated by the Securities and Futures Commission in Hong Kong. JP Morgan Chase Bank, N.A., Hong Kong Branch (CE Number AAL996) is regulated by the Hong Kong Monetary Authority and the Securities and Futures Commission, is organized under the laws of the United States with limited liability. Where the distribution of this material is a regulated activity in Hong Kong, the material is distributed in Hong Kong by or through J.P. Morgan Securities (Asia Pacific) Limited and/or J.P. Morgan Broking (Hong Kong) Limited.

India: J.P. Morgan India Private Limited (Corporate Identity Number - 68724), having its registered office at J.P. Morgan Tower, Off. C.S.T. Road, Kalina, Santacruz - East, Mumbai – 400098, is registered with the Securities and Exchange Board of India (SEBI) as a 'Research Analyst' having registration number INH000001873. J.P. Morgan India Private Limited is also registered with SEBI as a member of the National Stock Exchange of India Limited and the Bombay Stock Exchange Limited (SEBI Registration Number – INZ000239730) and as a Merchant Banker (SEBI Registration Number - MB/INM000002970). Telephone: 91-22-6157 3000, Facsimile: 91-22-6157 3990 and Website: http://www.jpmipl.com. JPMorgan Chase Bank, N.A. - Mumbai Branch is licensed by the Reserve Bank of India (RBI) (Licence No. 53/ Licence No. BY.4/94; SEBI - IN/CUS/014/ CDSL : IN-DP-CDSL-444-2008/ IN-DP-NSDL-285-2008/ INBI00000984/ INE231311239) as a Scheduled Commercial Bank in India, which is its primary license allowing it to carry on Banking business in India and other activities, which a Bank branch in India are permitted to undertake. For non-local research material, this material is not distributed in India by J.P. Morgan India Private Limited. Compliance Officer: Ashutosh Sharma; ashutosh.j.sharma@jpmchase.com; +912261575002. Grievance Officer: Ramprasadh K, jpmipl.research.feedback@jpmorgan.com; +912261573000. Registration granted by SEBI and certification from NISM in no way guarantee performance of the intermediary or provide any assurance of returns to investors. Please visit Terms and Conditions and Most Important Terms and Conditions (MITC). The annual Compliance audit report is available at http://www.jpmipl.com/#research.

Indonesia: PT J.P. Morgan Sekuritas Indonesia is a member of the Indonesia Stock Exchange and is registered and supervised by the Otoritas Jasa Keuangan (OJK).

Korea: J.P. Morgan Securities (Far East) Limited, Seoul Branch, is a member of the Korea Exchange (KRX). JPMorgan Chase Bank, N.A., Seoul Branch, is licensed as a branch office of foreign bank (JPMorgan Chase Bank, N.A.) in Korea. Both entities are regulated by the Financial Services Commission (FSC) and the Financial Supervisory Service (FSS). For non-macro research material, the material is distributed in Korea by or through J.P. Morgan Securities (Far East) Limited, Seoul Branch.

Japan: JPMorgan Securities Japan Co., Ltd. and JPMorgan Chase Bank, N.A., Tokyo Branch are regulated by the Financial Services Agency in Japan.

Malaysia: This material is issued and distributed in Malaysia by JPMorgan Securities (Malaysia) Sdn Bhd (18146-X), which is a Participating Organization of Bursa Malaysia Berhad and holds a Capital Markets Services License issued by the Securities Commission in Malaysia.

Mexico: J.P. Morgan Casa de Bolsa, S.A. de C.V., J.P. Morgan Grupo Financiero is member of the Mexican Stock Exchange ("Bolsa Mexicana de Valores") and the Institutional Stock Exchange ("Bolsa Institucional de Valores"), and it is authorized to act as a broker dealer by the National Banking and Securities Exchange Commission ("Comisión Nacional Bancaria y de Valores").

New Zealand: This material is issued and distributed by JPMSAL in New Zealand only to "wholesale clients" (as defined in the Financial Markets Conduct Act 2013). JPMSAL is registered as a Financial Service Provider under the Financial Service providers (Registration and Dispute Resolution) Act of 2008.

Philippines: J.P. Morgan Securities Philippines Inc. is a Trading Participant of the Philippine Stock Exchange and a member of the Securities Clearing Corporation of the Philippines and the Securities Investor Protection Fund. It is regulated by the Securities and Exchange Commission.

Singapore: This material is issued and distributed in Singapore by or through J.P. Morgan Securities Singapore Private Limited (JPMSS) [MDDI (P) 070/08/2026 and Co. Reg. No.: 199405335R], which is a member of the Singapore Exchange Securities Trading Limited, and/or JPMorgan Chase Bank, N.A., Singapore branch (JPMCB Singapore), both of which are regulated by the Monetary Authority of Singapore. This material is issued and distributed in Singapore only to accredited investors, expert investors and institutional investors, as defined in Section 4A of the Securities and Futures Act, Cap. 289 (SFA). This material is not intended to be issued or distributed to any retail investors or any other investors that do not fall into the classes of "accredited investors," "expert investors" or "institutional investors," as defined under Section 4A of the SFA. Recipients of this material in Singapore are to contact JPMSS or JPMCB Singapore in respect of any matters arising from, or in connection with, the material.

South Africa: J.P. Morgan Equities South Africa Proprietary Limited and JPMorgan Chase Bank, N.A., Johannesburg Branch are members of the Johannesburg Securities Exchange and are regulated by the Financial Services Conduct Authority (FSCA).

Taiwan: J.P. Morgan Securities (Taiwan) Limited is a participant of the Taiwan Stock Exchange (company-type) and regulated by the Taiwan Securities and Futures Bureau. Material relating to equity securities is issued and distributed in Taiwan by J.P. Morgan Securities (Taiwan) Limited, subject to the license scope and the applicable laws and the regulations in Taiwan. **To the extent that J.P. Morgan Securities (Taiwan) Limited produces research materials on securities not listed on the Taiwan Stock Exchange or Taipei Exchange ("Non-Taiwan Listed Securities"), these materials shall not constitute securities recommendations for the purpose of applicable Taiwan regulations, and, for the avoidance of doubt, J.P. Morgan Securities (Taiwan) Limited does not act as broker for Non-Taiwan Listed Securities.** According to Paragraph 2, Article 7-1 of Operational Regulations Governing Securities Firms Recommending Trades in Securities to Customers (as amended or supplemented) and/or other applicable laws or regulations, please note

that the recipient of this material is not permitted to engage in any activities in connection with the material that may give rise to conflicts of interests, unless otherwise disclosed in the "Important Disclosures" in this material.
Thailand: This material is issued and distributed in Thailand by JPMorgan Securities (Thailand) Ltd., which is a member of the Stock Exchange of Thailand and is regulated by the Ministry of Finance and the Securities and Exchange Commission. The registered address is 548 One City Center Building, 50th Floor, Ploenchit Road, Lumphini, Pathum Wan, Bangkok 10330.
UK: Research is produced in the UK by J.P. Morgan Securities plc ("JPMS plc") which is a member of the London Stock Exchange and is authorised by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority or J.P. Morgan Markets Limited ("JPMML Ltd") which is authorised and regulated by the Financial Conduct Authority. Unless specified to the contrary, this material is distributed in the UK by JPMS plc and is directed in the UK only to: (a) persons having professional experience in matters relating to investments falling within article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) (Order) 2005 ("the FPO"); (b) persons outlined in article 49 of the FPO (high net worth companies, unincorporated associations or partnerships, the trustees of high value trusts, etc.); or (c) any persons to whom this communication may otherwise lawfully be made; all such persons being referred to as "UK relevant persons". This material must not be acted on or relied on by persons who are not UK relevant persons. Any investment or investment activity to which this material relates is only available to UK relevant persons and will be engaged in only with UK relevant persons. A description of J.P. Morgan EMEA's policy for prevention and avoidance of conflicts of interest related to the production of Research can be found at the following link: J.P. Morgan EMEA - Research Independence Policy.
U.S.: J.P. Morgan Securities LLC ("JPMS") is a member of the NYSE, FINRA, SIPC, and the NFA. JPMorgan Chase Bank, N.A. is a member of the FDIC. Material published by non-U.S. affiliates is distributed in the U.S. by JPMS who accepts responsibility for its content.

General: Additional information is available upon request. The information in this material has been obtained from sources believed to be reliable. While all reasonable care has been taken to ensure that the facts stated in this material are accurate and that the forecasts, opinions and expectations contained herein are fair and reasonable, JPMorgan Chase & Co. or its affiliates and/or subsidiaries (collectively J.P. Morgan) make no representations or warranties whatsoever to the completeness or accuracy of the material provided, except with respect to any disclosures relative to J.P. Morgan and the Research Analyst's involvement with the issuer that is the subject of the material. Accordingly, no reliance should be placed on the accuracy, fairness or completeness of the information contained in this material. There may be certain discrepancies with data and/or limited content in this material as a result of calculations, adjustments, translations to different languages, and/or local regulatory restrictions, as applicable. These discrepancies should not impact the overall investment analysis, views and/or recommendations of the subject company(ies) that may be discussed in the material. Artificial intelligence tools may have been used in the preparation of this material, including assisting in data analysis, pattern recognition, and content drafting for research material. J.P. Morgan accepts no liability whatsoever for any loss arising from any use of this material or its contents, and neither J.P. Morgan nor any of its respective directors, officers or employees, shall be in any way responsible for the contents hereof, apart from the liabilities and that may be imposed on them by the relevant regulatory authority in the jurisdiction in question, or the regulatory regime thereunder. Opinions, forecasts or projections contained in this material represent J.P. Morgan's current opinions or judgment as of the date of the material only and are therefore subject to change without notice. Periodic updates may be provided on companies/industries based on company-specific developments or announcements, market conditions or any other publicly available information. There can be no assurance that future results or events will be consistent with any such opinions, forecasts or projections, which represent only one possible outcome. Furthermore, such opinions, forecasts or projections are subject to certain risks, uncertainties and assumptions that have not been verified, and future actual results or events could differ materially. The value of, or income from, any investments referred to in this material may fluctuate and/or be affected by changes in exchange rates. All pricing is indicative as of the close of market for the securities discussed, unless otherwise stated. Past performance is not indicative of future results. Accordingly, investors may receive back less than originally invested. This material is not intended as an offer or solicitation for the purchase or sale of any financial instrument. The opinions and recommendations herein do not take into account individual client circumstances, objectives, or needs and are not intended as recommendations of particular securities, financial instruments or strategies to particular clients. This material may include views on structured securities, options, futures and other derivatives. These are complex instruments, may involve a high degree of risk and may be appropriate investments only for sophisticated investors who are capable of understanding and assuming the risks involved. The recipients of this material must make their own independent decisions regarding any securities or financial instruments mentioned herein and should seek advice from such independent financial, legal, tax or other adviser as they deem necessary. J.P. Morgan may trade as a principal on the basis of the Research Analysts' views and research, and it may also engage in transactions for its own account or for its clients' accounts in a manner inconsistent with the views taken in this material, and J.P. Morgan is under no obligation to ensure that such other communication is brought to the attention of any recipient of this material. Others within J.P. Morgan, including Strategists, Sales staff and other Research Analysts, may take views that are inconsistent with those taken in this material. Employees of J.P. Morgan not involved in the preparation of this material may have investments in the securities (or derivatives of such securities) mentioned in this material and may trade them in ways different from those discussed in this material. This material is not an advertisement for or marketing of any issuer, its products or services, or its securities in any jurisdiction.

Confidentiality and Security Notice: This transmission may contain information that is privileged, confidential, legally privileged, and/or exempt from disclosure under applicable law. If you are not the intended recipient, you are hereby notified that any disclosure,

copying, distribution, or use of the information contained herein (including any reliance thereon) is STRICTLY PROHIBITED. Although this transmission and any attachments are believed to be free of any virus or other defect that might affect any computer system into which it is received and opened, it is the responsibility of the recipient to ensure that it is virus free and no responsibility is accepted by JPMorgan Chase & Co., its subsidiaries and affiliates, as applicable, for any loss or damage arising in any way from its use. If you received this transmission in error, please immediately contact the sender and destroy the material in its entirety, whether in electronic or hard copy format. This message is subject to electronic monitoring: https://www.jpmorgan.com/disclosures/email

MSCI: Certain information herein ("Information") is reproduced by permission of MSCI Inc., its affiliates and information providers ("MSCI") ©2026. No reproduction or dissemination of the Information is permitted without an appropriate license. MSCI MAKES NO EXPRESS OR IMPLIED WARRANTIES (INCLUDING MERCHANTABILITY OR FITNESS) AS TO THE INFORMATION AND DISCLAIMS ALL LIABILITY TO THE EXTENT PERMITTED BY LAW. No Information constitutes investment advice, except for any applicable Information from MSCI ESG Research. Subject also to msci.com/disclaimer

Any mentions of third-party trademarks, brand names, products, and services are for referential purposes only and any mention thereof is not meant to imply any sponsorship, endorsement, or affiliation.

Sustainalytics: Certain information, data, analyses and opinions contained herein are reproduced by permission of Sustainalytics and: (1) includes the proprietary information of Sustainalytics; (2) may not be copied or redistributed except as specifically authorized; (3) do not constitute investment advice nor an endorsement of any product or project; (4) are provided solely for informational purposes; and (5) are not warranted to be complete, accurate or timely. Sustainalytics is not responsible for any trading decisions, damages or other losses related to it or its use. The use of the data is subject to conditions available at https://www.sustainalytics.com/legal-disclaimers. ©2026 Sustainalytics. All Rights Reserved.

"Other Disclosures" last revised August 29, 2026.

Copyright 2026 JPMorgan Chase & Co. All rights reserved. This material or any portion hereof may not be reprinted, sold or redistributed without the written consent of J.P. Morgan. It is strictly prohibited to use or share without prior written consent from J.P. Morgan any research material received from J.P. Morgan or an authorized third-party ("J.P. Morgan Data") in any third-party artificial intelligence ("AI") systems or models when such J.P. Morgan Data is accessible by a third-party.