

April 26, 2026 09:50 PM GMT

Greater China Semiconductors | Asia Pacific

China's AI Accelerators – Who's Poised to Win?

We initiate coverage on Cambricon and Iluvatar at Overweight, and MetaX at Equal-weight.

Bull vs. bear cases for China AI accelerators: Who will capture China's inference-driven substitution cycle? In a [recent report](#), we argued that China is narrowing the US lead in AI compute not simply at the chip level but also through system-level innovation, supply-chain localization, and increasingly attractive inference economics. We continue to believe this path will lift China's domestic AI accelerator self-sufficiency rate to 86% by 2030, reshaping the global competitive landscape for AI semiconductors over the next decade.

Over the past month, our channel checks have turned incrementally more constructive on China's AI accelerator industry. 1) At our China Summit, major AI LLM developers such as MiniMax and Zhipu signaled their willingness to adopt domestic AI chips as long as token economics are competitive. 2) Our recent field trip indicated that the Nvidia GPU supply in China tightened after the [SMCI-related disruption](#), redirecting incremental demand toward domestic alternatives. 3) Strong spot demand for Nvidia RTX 5090 in China suggests AI inference demand remains robust. 4) Rising token prices and GPU rental prices also point to a still-tight compute market. The main negative datapoint: price competition appears to be arriving earlier than we had expected, as some vendors have started cutting prices to gain share. 5) WAIC will be held in Shanghai in July 2026, where we expect to see next-generation Chinese AI accelerator products, especially from Iluvatar.

Stock calls: We initiate coverage on Cambricon and Iluvatar at Overweight, and MetaX at Equal-weight. We believe all three are well positioned to benefit from China's accelerating AI chip localization trend, though each offers a differentiated investment case.

- **Cambricon (Overweight; PT Rmb1,588):** We view Cambricon as the leading domestic AI inference chip play, supported by strong CSP customer anchoring, proven hardware-software co-optimization, and solid positioning in large-scale cloud inference deployments.
- **Iluvatar (Overweight; PT HK\$600):** We view Iluvatar favorably for its diversified supply chain strategy, stronger supply visibility, and growing exposure to cloud customers.
- **MetaX (Equal-weight; PT Rmb758):** We see MetaX as a differentiated domestic GPGPU vendor, with relatively strong CUDA-like software compatibility, and a more scalable near-term manufacturing path. However, its valuation is less attractive than peers.

Key risks to our view include: 1) slower AI demand growth, 2) earlier-than-expected pricing pressure, 3) policy/export headwinds.

Charlie Chan Equity Analyst Charlie.Chan@morganstanley.com	+886 2 2730-1725
Gary Yu Equity Analyst Gary.Yu@morganstanley.com	+852 2848-6918
Daniel Yen, CFA Equity Analyst Daniel.Yen@morganstanley.com	+886 2 2730-2863
Daisy Dai, CFA Equity Analyst Daisy.Dai@morganstanley.com	+852 2848-7310
Henry Zhao Research Associate Henry.Zhao@morganstanley.com	+852 2239-7731
Joanne Lau Research Associate Joanne.CY.Lau@morganstanley.com	+852 3963-1592
Tiffany Yeh Equity Analyst Tiffany.Yeh@morganstanley.com	+886 2 7712-3032
Lucas Wang Research Associate Lucas.Wang@morganstanley.com	+886 2 2730-2875
Ethan Jia Research Associate Ethan.Jia@morganstanley.com	+852 3963-2287

GREATER CHINA TECHNOLOGY SEMICONDUCTORS	
Asia Pacific	
Industry View	Attractive

investors should be aware that the firm may have a conflict of

Research as only a single factor in making their investment decision.

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

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

Executive Summary – Don't Underestimate China's AI Compute Ecosystem

Booming China AI GPU market

China's AI GPU market is entering a more commercially grounded phase, with the debate shifting from whether domestic chips can participate to which vendors will win meaningful share as inference demand scales. In our view, two structural forces shape the market: (1) a rapid rise in AI inference demand, driven by commercialization across consumer and enterprise applications, and (2) persistent export controls, making localization a long-duration feature of China's AI compute market rather than a temporary policy response. Together, these forces expand the addressable market for domestic AI accelerators and improve the probability of sustained substitution. This aligns with our framework that China's AI chip TAM could reach US\$67bn by 2030, with domestic self-sufficiency rising to 86%.

Our core thesis remains that China's localization strategy is gaining traction: scaling domestic chips, foundries, packaging, and equipment capabilities to partially offset process-node disadvantages. In the bull case, domestic AI semis broaden from inference into selected training workloads, software ecosystems improve faster than expected, and some vendors achieve overseas adoption or indirect export opportunities. In the bear case, product differentiation fades, pricing pressure intensifies earlier than expected, and the sector moves toward commoditization and consolidation.

Exhibit 1: China AI accelerators' market cap trend




Exhibit 2: China AI accelerators' revenue trend




Exhibit 3:

"10 Dragons" of Chinese AI GPGPU vendors. We focus on Cambricon, MetaX, Iluvatar company research in this report

Company	Ticker	GPU Product	Node	GPGPU/ASIC	Foundry sourcing viability	Sovereign background	Affiliated design house	AI inference performance
 HUAWEI	Private	950, 910C, 910B	7nm	ASIC	✓	✓	✗	➡
 Cambricon 寒武纪	688256-S5	MLU series	7nm	ASIC	✓	✗	✗	➡
HYGON	688041-S5	DCU	7nm	GPGPU	✓	✓	✗	➡
 META X 沐曦	688802-SH	MXC series; MXG series	7nm/12nm	GPGPU	✓	✗	✗	➡
 摩尔线程 MOORE THREADS	688795-SH	MUSA MTT S5000, S4000	7nm/12nm	GPGPU	▲	✗	✗	➡
 壁仞科技 BIREN TECHNOLOGY	6082-HK	BR100	7nm/12nm	GPGPU	▲	✗	✗	➡
 天数智芯 Iluvatar CoreX	9903-HK	TianGai-100 series, Zhikai-100 Series	7nm	GPGPU	✓	✗	✗	➡
 T-HEAD	Currently under H-share listing counseling process	Hangguang 800, PPU	12nm	GPGPU	✓	✗	✓	➡
 昆仑芯 KUNLUNXIN	Currently under H-share listing counseling process	R and P series	6nm	ASIC	▲	✗	✓	➡
 Enflame 燧原科技	Currently under A-share listing counseling process	S60, I20, L600	12nm	ASIC	▲	✗	✓	➡

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

More broadly, we believe China's AI GPU race is no longer just a chip-specification contest. While domestic silicon still trails the US by roughly two generations at the chip level, the effective gap is narrowing through multi-die design, advanced packaging, rack-scale system architecture, optical networking, and software-hardware co-optimization. This is why we think system-level competitiveness matters more than ever. In a market increasingly dominated by inference and utilization, the vendor that delivers the best real-world token economics at acceptable software migration cost is likely to win customer budgets, even without leading-edge process technology.

From an investment perspective, this leads to a simple conclusion: the sector should not be valued as a monolithic policy theme. Instead, investors need to distinguish between vendors with a realistic path to shipment scale, ecosystem credibility, and pricing discipline, and those that may struggle to convert technical potential into durable revenues and margins. We therefore evaluate the group through a two-dimensional framework of economics × execution, combining TCO, token cost, TPS, and performance-per-dollar with qualitative factors such as foundry access, software ecosystem maturity, CSP relationships, and roadmap credibility. In our view, that framework remains the most effective way to separate likely winners from those at risk of being marginalized as the industry consolidates.

What has become clearer through our recent channel checks is that economics, not ideology, is driving adoption. At our China Summit, major LLM developers indicated they are willing to deploy local GPUs as long as token cost is competitive. This aligns with our

core finding that domestic accelerators already offer materially lower TCO than Nvidia products available in China and that leading Chinese chips can achieve broad cost-per-token parity in inference workloads. In other words, purchasing decisions are increasingly made on deployable economics rather than absolute peak silicon performance. This matters because China's AI demand is becoming more inference-heavy, more recurring, and more utilization-driven, which structurally favors solutions optimized around cost efficiency, software adaptation, and availability rather than headline benchmark leadership.

Exhibit 4: Relative strengths of US and China AI industries





Exhibit 5: Domestic chips have lower TCO and comparable per token cost (AI LLM inference) vs. NVIDIA's processors for China



Who's poised to win in China AI GPUs?


China's AI accelerator ecosystem spans merchant chip vendors, sovereign-backed players, and captive chip design houses linked to major cloud service providers (CSPs). We assess this ecosystem's competitors in a global GPU/ASIC context, compare relative positioning across performance, cost, and execution, and apply a consistent valuation framework to identify stocks with the most attractive risk/reward. Our field work with CSPs suggests that while per-token cost is the single most important KPI, software optimization and

strategic customer partnerships matter even more than we had assumed.

Based on recent shipment trends, customer allocation, market share evolution, and the earlier-than-expected onset of price erosion, we believe the next phase will be defined less by theoretical peak performance and more by commercial execution, software readiness, and customer capture.

From a competitive perspective, we believe the market should be segmented by customer type. For top-tier CSPs and leading LLM developers, the primary decision metric is increasingly per-token cost, but that KPI alone is not sufficient. In practice, software maturity, framework compatibility, cluster-level optimization, and the depth of strategic partnerships play a decisive role in order allocation. For sovereign AI, telecom, SOE, and government-linked demand, supply security, domestic controllability, and policy alignment carry greater weight. This creates room for different winners across end-markets. In our view, vendors with strong CSP co-development relationships and credible software stacks are better positioned to win high-volume cloud inference deployments, while vendors with stronger domestic supply chain visibility or government relationships may be better positioned in sovereign and public sector projects.

Exhibit 6: Order placement and potential orders for domestic AI accelerators developers, according to our industry checks



Within this framework, we see meaningful differentiation among Cambricon, MetaX, Iluvatar. **Cambricon** stands out on the ASIC/DSA path, where its inference performance, customer anchoring, and hardware-software co-optimization support strong deployment economics, particularly in large-scale cloud use cases. **Iluvatar** is differentiated by its diversified foundry strategy, better supply visibility, and a pragmatic path to customer migration through software compatibility. **MetaX** is one of the more credible domestic GPGPU players, in our view, due to relatively stronger CUDA-like software compatibility and a manufacturing path that may prove more scalable near term.

In short, Cambricon looks strongest on current cloud inference traction (second only to

Huawei Ascend), Iluvatar on supply-chain resilience plus commercial optionality, and MetaX on scalable GPGPU positioning.

Near-term market tracker for China AI GPU demand

At the same time, near-term industry conditions have turned more favorable for domestic vendors. Our recent field trip suggests Nvidia GPU availability in China has tightened, creating more room for local substitution. Strong spot-market demand for Nvidia 5090 cards, alongside higher token prices and GPU rental prices, also points to resilient downstream inference demand. These datapoints reinforce our view that the demand environment remains robust, especially for customers that need immediate deployment rather than waiting for supply normalization. The caveat: competition is intensifying faster than expected. We are already seeing price cuts in parts of the market, implying the sector may move into a market-share phase earlier than we had assumed. As a result, we expect execution quality to matter more than ever – particularly in software optimization, customer support, and strategic account penetration.

The World Artificial Intelligence Conference (WAIC) will be held in Shanghai in July 2026, where we expect to see next-generation Chinese AI accelerator products, especially from Iluvatar.

Exhibit 7: Nvidia's 5090 price keeps rising in China



Exhibit 8: Average token price for China's mainstream AI LLMs






Exhibit 9: Surge in ByteDance (Volcano Engine/Doubao) tokens indicates high AI demand



on daily numbers.

Exhibit 10: China CSP's capex will be a key demand driver for China AI GPU



Identifying Winners With Our Framework: Cambricon and Iluvatar Are Our Preferred Plays

Building on our performance and cost analysis, we apply a structured framework to assess domestic AI chip vendors' relative positioning, focusing on quantitative economics and qualitative execution.

Our framework: economics × execution

We evaluate vendors across two key dimensions:

- **Inference economics (quantitative)** – including TCO, cost per token, TPS performance, and performance per watt/dollar
- **Execution capability (qualitative)** – including access to leading-node capacity, software ecosystem maturity, depth of CSP relationships, and product roadmap soundness

In our view, sustained leadership requires strength in both. Vendors that excel in only one – e.g., strong in silicon but with a weak ecosystem – are unlikely to achieve durable share.

Exhibit 11: Comparison among Cambricon, MetaX and Iluvatar

			
Ticker	688256-SH	688802-SH	9903-HK
Product	MLU 220/270/370/580/590/690 (AI Training + inference)	C Series (AI training + inference) N Series (AI Inference) G Series (Graphic rendering)	Tiangai 100/150/200/300 (AI training + inference) Zhikai 100 (AI inference)
GPGPU/ASIC	ASIC	GPGPU	GPGPU
Chip suppliers			
Process node for latest products	7nm/N+2	12nm/N+1	7nm
Secured orders from major CSPs	✓	X	✓
Sovereign fund as major shareholder	X	✓	X
Per token cost performance	↑	→	→
2025 Revenue (Rmb mn)	CNY 6,497	CNY 1,644	CNY 1,034
Profitability	✓	X	X

Cambricon: Leading in inference performance and customer anchoring

Within this framework, we see Cambricon as one of the strongest positioned players on the ASIC (DSA) pathway.

Quantitatively, Cambricon's latest generation (e.g., MLU590) delivers competitive inference performance, with our TPS analysis suggesting meaningful outperformance vs. NVIDIA H20 under certain DeepSeek R1 scenarios. Combined with competitive pricing, this supports strong cost-per-token economics, which we view as the key CSP decision metric.

Qualitatively, Cambricon benefits from deep customer integration. Based on our industry checks, its multi-year collaboration with ByteDance enabled continuous hardware–software co-optimization and real-world deployment validation, providing an advantage in application-level tuning and commercialization readiness.

Taken together, we view Cambricon as a near-term leader in inference-driven deployments, particularly where efficiency and customer-specific optimization are critical.

Iluvatar: Leveraging supply chain resilience with strong order

visibility

We believe Iluvatar is well positioned to benefit from accelerating domestic AI chip substitution in China, supported by supply chain resilience, software compatibility, and improving commercial traction.

Based on our industry checks, leading Chinese CSPs have placed sizable pre-orders for Iluvatar's TianGai-150 AI chips, with shipments expected to commence in 2H26. Importantly, Iluvatar's diversified foundry strategy – including export-compliant production at TSMC – offers greater capacity visibility vs. peers relying solely on domestic fabs or non-compliant manufacture overseas, reducing supply disruption risk.

On software, Iluvatar's GPGPU architecture offers high CUDA compatibility, lowering migration friction. The company has helped clients to migrate LLM stacks from NVIDIA platforms to TianGai-150. In our view, this positions Iluvatar favorably as enterprises seek pragmatic NVIDIA alternatives.

MetaX: Positioning for scalability through software and supply

Within the GPGPU pathway, we view MetaX as a credible domestic participant, supported by its focus on improving CUDA ecosystem compatibility. While CUDA remains NVIDIA's key moat – given deep integration across compilers, libraries (e.g., cuDNN, NCCL), and a large developer base – it also creates structural switching costs that are not easily replicated.

Against this backdrop, MetaX's strategy of building a CUDA-like software stack and compatibility layer provides a reasonable adoption pathway for domestic customers. Based on our industry checks, the company has made steady progress in compiler adaptation, framework compatibility (e.g., PyTorch), and runtime optimization, although overall ecosystem maturity and stability still lag global leaders.

In addition, MetaX adopts a pragmatic manufacturing strategy, leveraging relatively mature nodes (e.g., N+1/12nm) to support yield stability and supply availability. While this may limit peak performance vs. leading-edge products, it offers a more balanced trade-off among performance, cost, and manufacturability. Overall, we view MetaX as a company with improving execution and scalability potential, though further validation in large-scale commercial deployments remains a key factor to monitor.

Performance and Cost: Which Domestic AI Chips Stand Out?

Inference economics more important than foundation model training for China AI GPU market


In a prior [China AI Insights note](#), we conducted a comprehensive comparison of domestic AI accelerators across key performance and economic metrics, including total cost of ownership (TCO), total processing performance (TPP), token output per second (TPS) under DeepSeek R1 inference, and performance per watt.

1. TPS - the revenue

We believe TPS (tokens per second) is another relevant metric for China’s inference driven market. Unlike peak FLOPS, TPS captures end to end system performance, reflecting hardware capability (compute throughput, memory and interconnect bandwidth) and software efficiency under real world workloads. Using DeepSeek R1 as our benchmark and calibrating assumptions against NVIDIA’s disclosed H200 result (5,899 TPS in Feb 2025), we find that leading domestic accelerators – such as Huawei’s Ascend 950PR/DT and Cambricon’s MLU690 – can outperform NVIDIA H20 by 50-150% in our scenarios. This reflects improvements in compute capability and system level optimization and compute to network balance.


Domestic vendors have made meaningful progress through memory, interconnect, and system architecture improvements, enabling competitive inference performance despite process node disadvantages. This reinforces our view that performance leadership is increasingly workload dependent, with domestic chips already competitive in inference scenarios, even as NVIDIA maintains an edge at the technology frontier.

Exhibit 12: TPS (tokens per second) analysis for China AI accelerators



layer: 61; Batch size: 1; Computing power: FP8, and if FP8 is not available, we will use FP16).


Exhibit 13: Our TPS result matches with NVIDIA's published result for H200 (Feb 2025), with same token length and data format assumptions



Source: NVIDIA

Exhibit 14: GB300 could deliver 50x of token per watt performance vs. H200 with DeepSeek R1

GB300 NVL72 Delivers 50x Tokens per Watt Leap



Source: NVIDIA

TPS Methodology

Our TPS framework uses core hardware and workload variables that determine token throughput in inference. Hardware inputs include: (1) effective compute throughput (FP8, or FP16 where FP8 is unavailable), (2) memory bandwidth, (3) interconnect bandwidth, and (4) chip utilization rate (UTR). Workload inputs include model size (671GB), layers (61), active experts (9/257 under MoE), input/output token length (1,024/1,024), and batch size (1).

We select DeepSeek R1 as the benchmark model given its representativeness within China's current LLM ecosystem and its MoE architecture. To calibrate our framework, we benchmarked NVIDIA H200 output figures disclosed in 1Q25 with DeepSeek R1. Specifically, in February 2025, NVIDIA's H200 achieved 5,899 TPS when running DeepSeek R1 inference.

Exhibit 15: Our key TPS formulas

$$\text{processing time} = \text{decode time} + \text{prefill time}$$

$$\text{prefill time} = \text{prefill computing time} + \text{prefill communication time}$$

$$\text{decode time} = \frac{\text{num of active experts}}{\text{num of total experts}} \times (\text{decode computing time} + \text{decode communication time}) \times KV$$

$$\text{prefill computing time} = \frac{\text{num of input tokens} \times \text{num of active parameters}}{\text{computing power} \times \text{UTR}} \times \frac{1}{1000}$$

$$\text{prefill communication time} = \frac{2 \times \text{num of parameters}}{\text{memory bandwidth}}$$

$$\text{decode computing time} = \frac{\text{num of output tokens} \times \text{num of active parameters}}{\text{computing power} \times \text{UTR}} \times \frac{1}{1000}$$

$$\text{decode communication time} = \frac{2 \times \text{num of parameters}}{\text{network bandwidth}}$$

Limitations

Our TPS estimates rely primarily on hardware specifications (compute throughput, memory bandwidth, networking bandwidth).

A critical variable in our model is chip utilization rate (UTR), which reflects how effectively theoretical hardware throughput translates into sustained real-world performance. We fine-tuned our UTR assumptions such that our modeled TPS approximates NVIDIA's published results. This calibration step helps anchor our model to observable market data rather than relying purely on theoretical peak specifications.

It is also important to note that NVIDIA's published results likely incorporate multiple layers of software optimization, including kernel-level tuning, TensorRT graph optimization, improved memory scheduling, communication library optimization, and refined expert routing for MoE models. As such, the NVIDIA benchmark (Feb 2025 result) we adopted could be viewed as a point-in-time reference that may improve further with subsequent driver, framework or firmware updates.

Accordingly, while our framework provides a structured, hardware-grounded basis for comparing inference performance across vendors, realized TPS may vary depending on software maturity and optimization, workload mix, and cluster configuration. Our analysis assumes a fixed inference workload (input/output tokens of 1,024/1,024 and batch size of 1), which may not reflect all deployment scenarios.

Of note, when assessing NVIDIA's next-generation platforms, such as the GB300, the performance gap widens materially. In NVIDIA's latest publication ([link](#)), the company indicates that GB300 delivers up to 50x improvement in token-per-watt performance versus H200 with DeepSeek R1, driven by significantly higher compute throughput, enhanced networking bandwidth, next-generation HBM performance, and support for lower-precision formats, such as FP4.

Importantly, our comparative analysis in this report is confined to products that China can currently procure or may reasonably access in the near term. As such, our benchmarking framework does not incorporate NVIDIA's latest frontier platforms. In absolute terms, NVIDIA's most advanced systems remain meaningfully ahead of the products included in our modeling. Accordingly, our conclusions should be interpreted within the context of China-accessible supply rather than global technology leadership at the cutting edge.

2. TCO - the cost

TCO remains one of the most compelling advantages for domestic AI chips. On an all-in basis – including chip acquisition, power, and supporting infrastructure – we estimate domestic accelerators can deliver 30-60% lower TCO than NVIDIA solutions currently accessible in China. Lower upfront chip pricing and structurally lower electricity and infrastructure costs in China drive this edge. The advantage becomes more pronounced at scale, particularly in inference-heavy deployments where high utilization pushes operating expenses to dominate lifecycle costs.

When we translate system performance into cost per token, the gap narrows. While NVIDIA retains an absolute performance lead – especially at the high end – leading domestic accelerators have reached broad cost-per-token parity with NVIDIA's A100/H20-class products and in certain configurations may outperform. In our view, this marks a critical inflection as CSPs optimize for monetization and utilization rather than peak

silicon capability.

Exhibit 16: AI GPU pricing: China vs. the US – China generally sells at a lower price given lower margin requirements




Exhibit 17: Electricity costs: China enjoys much cheaper power than other economies





Exhibit 18: Domestic chips have lower TCO and comparable per token cost (AI LLM inference) vs. NVIDIA's processors for China



3. Energy efficiency - the power consumption

On energy efficiency, domestic chips have largely closed the gap with NVIDIA's A100 and H20, although they still lag newer platforms such as H100 and H200. Meanwhile, on performance per dollar, domestic accelerators already show clear advantages, driven by lower pricing and improving real world performance.

Incorporating acquisition cost, domestic chips deliver stronger performance per dollar due to materially lower pricing. On this basis, leading domestic accelerators already outperform the A100 and narrow the gap with the H200, reinforcing their appeal for inference heavy deployments.

Exhibit 19: Performance/watt comparison




Exhibit 20: Performance/cost comparison




Exhibit 21: Cost efficiency and power efficiency comparison

Brand Type Product	Nvidia				Huawei Ascend			Cambricon				MetaX		Illuvatar		Moore Thread	T-Head	Kunlun
	GP GPU A100	GP GPU H20	GP GPU H100	GP GPU H200	ASIC (DSA) 910B	ASIC (DSA) 910C	GP GPU 950PR*	ASIC (DSA) MLU 370	ASIC (DSA) MLU 580*	ASIC (DSA) MLU 590*	ASIC (DSA) MLU 690*	GP GPU C500	GP GPU C600*	GP GPU BI-V100*	GP GPU BI-V200*	GP GPU S5000*	GP GPU PPU*	ASIC (DSA) P800*
Foundry																		
Die size (mm²)	826	814	814	814	666	666	480	450	650	750	650	800	700	400	500	700	750	750
Dies per package	1	1	1	1	1	2	2	1	1	1	2	1	2	1	2	1	1	1
Gross die per wafer	62	63	63	63	80	80	116	125	82	69	82	64	75	142	142	75	69	69
Yield (%)	70%	50%	50%	50%	70%	70%	30%	70%	45%	70%	25%	70%	45%	75%	80%	70%	45%	55%
Wafer price (US\$)	\$11,000	\$18,000	\$18,000	\$18,000	\$11,000	\$11,000	\$12,000	\$11,000	\$9,000	\$11,000	\$12,000	\$12,000	\$9,000	\$11,000	\$11,000	\$11,000	\$9,000	\$12,000
Die cost per chip (US\$)	\$253	\$571	\$571	\$571	\$196	\$393	\$690	\$126	\$244	\$228	\$1,171	\$268	\$533	\$103	\$194	\$210	\$290	\$316
Memory																		
Memory size (GB)	80	96	80	144	64	96	128	48	96	96	144	96	144	64	96	80	96	96
Memory type	HBM2e	HBM3	HBM3	HBM3	HBM2e	HBM3	LPDDR5	LPDDR5	HBM3	HBM3	HBM3e	HBM3	HBM3e	HBM2e	HBM3	HBM3	HBM3	HBM3
US\$/16GB	\$160	\$200	\$200	\$200	\$160	\$200	\$150	\$150	\$200	\$200	\$220	\$200	\$220	\$160	\$200	\$200	\$200	\$200
Memory cost (US\$)	\$800	\$1,200	\$1,000	\$1,800	\$640	\$1,200	\$1,200	\$450	\$1,200	\$1,200	\$1,980	\$1,200	\$1,980	\$640	\$1,200	\$1,000	\$1,200	\$1,200
Packaging																		
Chips per CoWoS/2.5D packaging	30	30	30	30	30	15	15	30	30	30	15	30	15	30	15	30	30	30
CoWoS wafer price (US\$)	\$10,000	\$10,000	\$10,000	\$10,000	\$10,000	\$10,000	\$7,000	\$10,000	\$7,000	\$10,000	\$7,000	\$10,000	\$7,000	\$10,000	\$10,000	\$10,000	\$7,000	\$7,000
Packaging and testing cost (US\$)	\$567	\$567	\$567	\$567	\$567	\$1,133	\$793	\$567	\$397	\$567	\$793	\$567	\$793	\$567	\$1,133	\$567	\$397	\$397
Chip costs																		
Chip manufacturing costs (US\$)	\$1,620	\$2,338	\$2,138	\$2,938	\$1,403	\$2,726	\$2,683	\$1,142	\$1,841	\$1,994	\$3,944	\$2,035	\$3,307	\$1,310	\$2,527	\$1,776	\$1,887	\$1,913
Other overheads (US\$)	\$688	\$558	\$662	\$262	\$197	\$474	\$488	\$143	\$659	\$606	\$856	\$251	\$408	\$747	\$559	\$601	\$685	\$659
Gross margin assumption (%)	80%	69%	86%	84%	60%	60%	63%	55%	65%	65%	72%	60%	60%	64%	55%	68%	76%	70%
Estimated chip cost (US\$)	\$2,308	\$2,896	\$2,800	\$3,200	\$1,600	\$3,200	\$3,171	\$1,286	\$2,500	\$2,600	\$4,800	\$2,286	\$3,714	\$2,057	\$3,086	\$2,377	\$2,571	\$2,571
Chip price to China CSP																		
Market price (Rmb)	NA	CNY 85,000	CNY 210,000	CNY 210,000	CNY 35,000	CNY 70,000	CNY 60,000	CNY 25,000	CNY 50,000	CNY 65,000	CNY 120,000	CNY 50,000	CNY 65,000	CNY 50,000	CNY 60,000	CNY 65,000	CNY 75,000	CNY 75,000
Market price (US\$)	\$15,000	\$12,143	\$30,000	\$30,000	\$5,000	\$10,000	\$8,571	\$3,571	\$7,143	\$9,286	\$17,143	\$7,143	\$9,286	\$7,143	\$8,571	\$9,286	\$10,714	\$10,714
Import cost markup (%)	30%	30%	50%	50%	20%	20%	0%	20%	0%	20%	0%	20%	0%	20%	20%	20%	0%	20%
Original chip price in USD	\$11,538	\$9,341	\$20,000	\$20,000	\$4,000	\$8,000	\$8,571	\$2,857	\$7,143	\$7,429	\$17,143	\$5,714	\$9,286	\$5,714	\$6,857	\$7,429	\$10,714	\$8,571
Performance																		
Performance (TFLOPS at FP16)	624	148	1,979	1,979	400	800	500	96	280	315	700	240	300	128	380	500	350	316
Performance (TFLOPS at FP8)	X	296	3,958	3,958	X	X	1,000	X	560	630	1,400	X	600	X	760	1,000	700	632
Performance/Cost	0.04	0.02	0.13	0.13	0.08	0.08	0.12	0.03	0.08	0.07	0.08	0.03	0.06	0.02	0.09	0.11	0.07	0.06
Power consumption (W)	400	400	700	700	400	750	580	250	560	580	750	350	450	350	700	700	700	750
Performance/watt	1.56	0.74	5.65	5.65	1.00	1.07	1.72	0.38	1.00	1.09	1.87	0.69	1.33	0.37	1.09	1.43	1.00	0.84

4. GPGPU vs. ASIC product definition, does that matter?

In assessing domestic AI chip competitiveness, we view architectural choice – GPGPU vs. ASIC (DSA) – as a fundamental strategic trade-off in China's AI ecosystem. Unlike global peers with access to leading edge nodes and mature software, domestic vendors must optimize across manufacturing limitations, software maturity, and evolving demand structures.

Architecturally, GPGPUs offer superior programmability and flexibility, suiting a rapidly evolving model landscape and compatibility with mainstream frameworks such as PyTorch. This flexibility is particularly valuable in China, where software ecosystems are still converging and model architectures continue to iterate. However, GPGPUs are less power- and area-efficient and typically require more advanced nodes to achieve competitive performance.

In contrast, ASIC-based designs enable tighter hardware-software co-optimization, allowing vendors to partially offset process disadvantages and deliver higher efficiency in targeted workloads. This makes ASICs attractive for inference-heavy or relatively stable use cases. The trade-offs include reduced generality, higher software porting costs, and greater sensitivity to changes in model architecture.

Exhibit 22: GPGPU vs. ASIC: What's the difference?

	Nvidia H200	Huawei Ascend 950PR
Type	GPGPU	ASIC(DSA)
General computing capability	Strong (AI, graphic, scientific computing, etc)	Relatively weaker (Focus on AI acceleration)
AI Matrix computation	Implemented via Tensor Cores	Natively implemented via 3D Cube
Ecosystem	CUDA	CANN (Compatible with PyTorch, etc.)
Main use case	General-purpose AI acceleration	AI LLM training and inference

Cambricon (688256.SS): Leading in inference performance and customer anchoring; Initiate at OW

Cambricon is a leading domestic AI chip designer, focusing on cloud and edge AI processors with proprietary architecture and full-stack capabilities. We view it as one of China's core players in AI accelerators, positioned to benefit from rising domestic AI compute demand and ongoing localization trends under geopolitical constraints. Its MLU series chips have achieved meaningful deployment in major CSP workloads (e.g., search, advertising, recommendation), demonstrating improving product-market fit and supporting future order visibility. The company's ongoing transition toward a chip-centric revenue model, combined with progress in domestic supply chain substitution, enhances revenue quality and scalability. In addition, continued product iteration (e.g., next-generation MLU roadmap) and software ecosystem development (NeuWare) provide potential upside as adoption deepens across large-scale AI clusters.

Earnings: Our EPS estimates for 2026 are above consensus. We expect revenue growth to remain solid, supported by continued CSP order ramp and product iteration, although profitability may be influenced by supply chain dynamics and R&D intensity.

Valuation: We derive our price target using a residual income model. Our price target implies a 110x 2026e P/E multiple and 32x 2026e P/S, which we believe is justified given the company's 90% revenue CAGR over 2025-28 and its positioning in China's AI semiconductor market.

Key downside risks include: 1) slower-than-expected AI demand, 2) customer concentration and order volatility, and 3) supply chain and yield constraints.

Exhibit 23: Cambricon: Summary of key metrics

688256.SS Greater China Technology Semiconductors				
Price target	Rmb1,588			
Up/downside to price target (%)	19%			
Shr price, close(Apr 26,2026)	1352.5			
52-Week Range	Rmb520.67-1595.88			
Sh out, dil, curr (mn)	425			
Mkt cap, curr (mn)	570,280			
EV, curr (mn)	566,051			
Avg daily trading value (bn)	12			
Fiscal Year Ending	12/25	12/26e	12/27e	12/28e
EPS for consensus(Rmb)	4.87	14.49 e	22.50 e	31.88 e
Consensus EPS(Rmb)\$		11.88 e	25.04 e	46.96 e
Revenue, net (Rmb mn)	6,497	20,944	33,186	44,256
EBITDA (Rmb mn)	2,211	6,240	11,097	15,772
ModelWare net inc (Rmb mn)	6,497	20,944	33,186	44,256
P/E	278.3	91.8	59.1	41.7
P/BV	48.4	32.3	22.1	15.4
RNOA (%)	57.8	78.5	82.5	124.5
ROE (%)	37.9	52.0	54.6	53.0
EV/EBITDA	256.5	89.4	49.5	34.3
Div yld (%)	0.0	0.1	0.3	0.4
FCF yld ratio (%)	0.5	0.3	1.8	1.9
Leverage (EOP) (%)	(49.6)	(42.4)	(62.8)	(66.3)
Unless otherwise noted, all metrics are based on Morgan Stanley ModelWare framework				
\$ = Consensus data is provided by Thomson Reuters Estimates				
e = Morgan Stanley Research estimates				
Company description				
Cambricon Technologies is a leading artificial intelligence (AI) chip designer in China, specializing in intelligent processors for AI computing. Founded in 2016 and listed on the STAR Market in 2020, Cambricon originated from research at the Chinese Academy of Sciences, where its founding team developed early neural network processors.				

Source:

Cambricon: Financial Summary

Exhibit 24: Cambricon: Quarterly financials

(Rmb mn)	1Q25	2Q25	3Q25	4Q25	1Q26E	2Q26E	3Q26E	4Q26E	2025	2026E	2027E	2028E
Total revenues	1,111.4	1,769.2	1,726.8	1,889.8	1,992.2	2,799.9	4,913.2	11,238.7	6,497.2	20,944.1	33,185.7	44,255.5
<i>Q/Q Change</i>	12.4%	59.2%	-2.4%	9.4%	5.4%	40.5%	75.5%	128.7%	0.0%	0.0%	0.0%	0.0%
<i>Y/Y Change</i>	4229.6%	4424.9%	1332.5%	91.0%	79.2%	58.3%	184.5%	494.7%	453.2%	222.4%	58.4%	33.4%
Cost of Sales	489.1	780.5	790.2	854.0	936.3	1,343.9	2,456.5	5,619.3	2,913.9	10,356.1	17,177.0	23,455.0
<i>Percent of Revenues</i>	44%	44%	46%	45%	47%	48%	50%	50%	45%	49%	52%	53%
Gross Profit	622.3	988.7	936.6	1,035.7	1,055.9	1,456.0	2,456.7	5,619.4	3,583.3	10,588.0	16,008.8	20,800.5
<i>Gross Margin</i>	56.0%	55.9%	54.2%	54.8%	53.0%	52.0%	50.0%	50.0%	55.2%	50.6%	48.2%	47.0%
<i>Incremental Margin</i>	47.9%	55.7%	NM	60.8%	19.7%	49.5%	47.4%	50.0%	54.8%	48.5%	44.3%	43.3%
Total Opex	330.0	338.0	355.2	593.7	583.7	820.4	1,292.2	1,652.1	1,616.9	4,348.4	4,911.5	5,028.8
<i>Percent of Revenues</i>	29.7%	19.1%	20.6%	31.4%	29.3%	29.3%	26.3%	14.7%	24.9%	20.8%	14.8%	11.4%
R&D	272.6	269.3	300.9	508.0	518.0	728.0	1,130.0	1,461.0	1,350.8	3,837.0	4,314.1	4,232.2
<i>Percent of Revenues</i>	24.5%	15.2%	17.4%	26.9%	26.0%	26.0%	23.0%	13.0%	20.8%	18.3%	13.0%	9.6%
General & Adm Exp.	43.1	54.9	39.8	60.4	49.8	70.0	122.8	134.9	198.1	377.5	431.4	575.3
<i>Percent of Revenues</i>	3.9%	3.1%	2.3%	3.2%	2.5%	2.5%	2.5%	1.2%	3.0%	1.8%	1.3%	1.3%
Selling Expenses	14.4	13.7	14.6	25.4	15.9	22.4	39.3	56.2	68.0	133.8	165.9	221.3
<i>Percent of Revenues</i>	1.3%	0.8%	0.8%	1.3%	0.8%	0.8%	0.8%	0.5%	1.0%	0.6%	0.5%	0.5%
Operating Income	292.2	650.8	581.3	442.1	472.2	635.6	1,164.5	3,967.3	1,966.4	6,239.7	11,097.3	15,771.7
<i>Operating Margin</i>	26.3%	36.8%	33.7%	23.4%	23.7%	22.7%	23.7%	35.3%	30.3%	29.8%	33.4%	35.6%
Total Non-operating Income (loss)	63.1	31.6	(15.2)	13.5	28.5	30.5	20.7	36.4	93.0	116.2	147.0	162.2
Profit Before Taxes	355.4	682.4	566.1	455.5	500.7	666.2	1,185.2	4,003.7	2,059.4	6,355.9	11,244.3	15,933.9
<i>Percent of Revenues</i>	32.0%	38.6%	32.8%	24.1%	25.1%	23.8%	24.1%	35.6%	31.7%	30.3%	33.9%	36.0%
Taxes	0.1	(0.0)	(0.3)	1.1	0.0	0.0	0.0	200.2	0.9	200.2	1,686.6	2,390.1
<i>Tax Rate</i>	0.0%	0.0%	0.0%	0.2%	0.0%	0.0%	0.0%	5.0%	0.0%	3.1%	15.0%	15.0%
Total Net Income to Parent	355.5	682.6	566.6	454.6	500.8	666.3	1,185.4	3,803.7	2,059.2	6,156.2	9,558.2	13,544.4
<i>Percent of Revenues</i>	32.0%	38.6%	32.8%	24.1%	25.1%	23.8%	24.1%	33.8%	31.7%	29.4%	28.8%	30.6%
EPS for consensus (Rmb)	0.8	1.6	1.3	1.1	1.2	1.6	2.8	9.0	4.9	14.5	22.5	31.9

Exhibit 25: Cambricon: Financial summary

Income Statement					Cash Flow Statement				
Rmb mn (Years End Dec)	2025	2026E	2027E	2028E	Rmb mn (Years End Dec)	2025	2026E	2027E	2028E
Net sales	6,497.2	20,944.1	33,185.7	44,255.5	Cashflow from Operations	(498.5)	2,117.2	10,223.2	10,783.2
COGS	2,913.9	10,356.1	17,177.0	23,455.0	Net profits	2,059.2	6,156.2	9,558.2	13,544.4
Gross profit	3,583.3	10,588.0	16,008.8	20,800.5	Depreciation	116.7	0.0	0.0	0.0
Operating expenses	1,616.9	4,348.4	4,911.5	5,028.8	Working Capital Change	(3,126.6)	(4,039.0)	665.0	(2,761.1)
Operating income	1,966.4	6,239.7	11,097.3	15,771.7	Other adjustments	452.2	0.0	0.0	0.0
Non-operating income	93.0	116.2	147.0	162.2	Cashflow from Investing	(4,142.0)	(80.0)	(80.0)	(80.0)
Pre-tax income	2,059.4	6,355.9	11,244.3	15,933.9	Capex	(142.8)	0.0	0.0	0.0
Income tax	0.9	200.2	1,686.6	2,390.1	Change of LT Investment	(53.1)	0.0	0.0	0.0
Minority Interest	(0.7)	(0.5)	(0.5)	(0.5)	Change of ST Investment	(3,501.0)	0.0	0.0	0.0
Reported net Income	2,059.2	6,156.2	9,558.2	13,544.4	Other adjustments	(445.1)	(80.0)	(80.0)	(80.0)
Adj. wtd. avg. shrs(m)	419.6	421.7	421.7	421.7	Cashflow from financing	3,986.0	(500.0)	(1,500.0)	(2,500.0)
Reported EPS (Rmb)	4.9	14.6	22.7	32.1	Increase in L/T debt	0.0	0.0	0.0	0.0
EPS for consensus (Rmb)	4.9	14.5	22.5	31.9	Increase in S/T debt	(100.1)	0.0	0.0	0.0
					Cash Dividend Paid	0.0	(500.0)	(1,500.0)	(2,500.0)
					Issuance of stock	1,465.7	0.0	0.0	0.0
					Other adjustments	2,620.4	0.0	0.0	0.0
					Exchange rate adjustment	(0.0)	0.0	0.0	0.0
					Net change in cash	(654.6)	1,537.2	8,643.2	8,203.2
Balance Sheet					Financial Ratios				
Rmb mn (Years End Dec)	2025	2026E	2027E	2028E		2025	2026E	2027E	2028E
Cash	1,317.4	2,854.6	11,497.8	19,701.0	Growth(%)				
Mkt Securities	4,261.2	4,261.2	4,261.2	4,261.2	Revenue	453.2	222.4	58.4	33.4
AR/NR	670.6	2,295.2	3,636.8	4,849.9	Operating profits	NA	217.3	77.9	42.1
Inventory	4,943.5	8,511.8	7,059.0	9,639.0	Pretax profits	NA	208.6	76.9	41.7
Other	887.8	887.8	887.8	887.8	Net profits	NA	199.0	55.3	41.7
Current Assets	12,080.4	18,810.5	27,342.5	39,338.9	EPS	NA	197.5	55.3	41.7
Long-term investments	299.7	299.7	299.7	299.7	Margins (%)				
Fixed assets	372.7	372.7	372.7	372.7	Gross Margin	55.2	50.6	48.2	47.0
Intangible Assets	189.1	169.1	149.1	129.1	Operating Margin	30.3	29.8	33.4	35.6
Other assets	495.8	595.8	695.8	795.8	Pretax Margin	31.7	30.3	33.9	36.0
Total Assets	13,437.7	20,247.8	28,859.8	40,936.1	Net Profit	31.7	29.4	28.8	30.6
S/T borrowings	0.0	0.0	0.0	0.0	Return (%)				
AP/NP	1,115.9	2,269.8	2,823.6	3,855.6	ROAE	23.8	42.0	44.4	43.6
Other ST liabilities	217.0	217.0	217.0	217.0	ROAA	20.4	36.6	38.9	38.8
LT debt	0.0	0.0	0.0	0.0	Gearing (%)				
Other LT liabilities	261.5	261.5	261.5	261.5	Net Debt/Equity	(11.1)	(16.3)	(45.0)	(53.8)
Total Liabilities	1,594.5	2,748.4	3,302.1	4,334.1	Liabilities/Equity	13.5	15.7	12.9	11.8
Common shares	421.7	421.7	421.7	421.7	Ratios (X)				
Additional capital	11,270.3	11,270.3	11,270.3	11,270.3	Current ratio	9.1	7.6	9.0	9.7
Retained earning	(11.8)	5,644.4	13,702.6	24,747.0	Quick ratio	1.5	2.1	5.0	6.0
Other shareholders' equity	163.0	163.0	163.0	163.0	Others				
Total Equity	11,843.3	17,499.5	25,557.6	36,602.0	AR/NR Turnover (days)	40.0	40.0	40.0	40.0
Total Liab. & Shrhldr's Equity	13,437.7	20,247.8	28,859.8	40,936.1	Inventory Turnover (days)	500.0	300.0	150.0	150.0
					AP Turnover (days)	90.0	80.0	60.0	60.0
					Cash Conversion (days)	450.0	260.0	130.0	130.0

E = Morgan Stanley Research Estimates

Source: Morgan Stanley Research, Company Data


Risk Reward – Cambricon Technology Corporation (688256.SS)

Leading in inference performance and customer anchoring

PRICE TARGET Rmb1,588.00

Key valuation assumptions underpinning our model include: an 8.4% cost of equity (derived from a beta of 1.06, risk-free rate of 2.0%, and equity risk premium of 6.0%), a long-term payout ratio of 26%, a medium-term growth rate of 16%, and a perpetual terminal growth rate of 6.0%.

RISK REWARD CHART




Key: — Historical Stock Performance ● Current Stock Price ◆ Price Target

OVERWEIGHT THESIS

- Deep engagement with major CSPs with proven product-market fit, with MLU590 widely deployed in SAD workloads, supporting strong order visibility as AI infrastructure scales.
- Domestic supply chain transition is underway, with production shifting to SMIC; while yields remain a challenge, gradual improvement is expected through 2026.
- Aggressive product roadmap, with next-generation MLU690 expected in 4Q26, potentially delivering a ~2.2x performance uplift and sustaining technology leadership.
- Our price target implies 32x of 2026e P/S, which is slightly lower than its historical average P/S multiple.

Consensus Rating Distribution



Risk Reward Themes

New Data Era: *Positive*
 Secular Growth: *Positive*
 Technology Diffusion: *Positive*

View descriptions of Risk Rewards Themes [here](#)

BULL CASE

Rmb3,000.00

61x 2026e P/S

We assume (1) >120% revenue CAGR in 2025-28 driven by stronger-than-expected domestic generative AI infrastructure build-out and large-scale procurement of cloud training & inference chips; (2) sustained share gains in China's domestic high-performance AI chip market; (3) gross margin improves to over 55% in 2026 and 2027 thanks to production scale effects, mature technology iteration and optimized high-value product structure.

BASE CASE

Rmb1,588.00

32x 2026e P/S

We assume (1) 90% revenue CAGR in 2025-28 driven by domestic generative AI infrastructure spending from CSP clients; (2) gross margin reaches 51% and 48% in 2026 and 2027 given lower selling prices due to products with smaller die size.

BEAR CASE

Rmb800.00

16x 2026e P/S

We assume (1) <60% revenue CAGR in 2025-28 given slower-than-expected domestic AI capital expenditure and LLM commercialization progress; (2) market share loss in China's AI chip market, and failure to secure continuous mass orders from top-tier customers; and (3) gross margin falls below 40% in 2026 and 2027 due to fierce price competition, rising wafer fabrication costs and underutilized production capacity.

Risk Reward – Cambricon Technology Corporation (688256.SS)

KEY EARNINGS INPUTS

Drivers	2025	2026e	2027e	2028e
Gross Profit (YoY) (%)	438.0	195.5	51.2	29.9


INVESTMENT DRIVERS

- CSP capex expansion
- Large model deployment growth
- Domestic substitution tailwind
- Product iteration (MLU roadmap)
- Software ecosystem improvement
- Policy support for AI chips
- New customer penetration
- Scaling of AI clusters


RISKS TO PT/RATING

- RISKS TO UPSIDE**
- Stronger-than-expected AI demand
 - CSP order ramp-up
 - Accelerating localization
- RISKS TO DOWNSIDE**
- Capacity and yield constraints
 - Customer concentration risk
 - Slower technology iteration

MS ESTIMATES VS. CONSENSUS



GLOBAL REVENUE EXPOSURE



View explanation of regional hierarchies [here](#)



Cambricon: Investment Positives & Concerns

Positives

Deeply entrenched with major CSP customers with proven product-market fit.

According to our industry checks, Cambricon's MLU590 chips are widely deployed in search, advertising, and recommendation (SAD) systems across major cloud service providers. These deployments demonstrate strong performance and reliability, providing visibility into future orders as CSPs scale AI infrastructure.

Domestic supply chain transition is progressing. While historically reliant on overseas foundries, our industry checks confirm Cambricon began shifting wafer production to SMIC in 2025. Although yields remain challenged on advanced nodes, we expect steady improvement through 2026 as processes mature.

First profitable domestic AI chip company validates business model. Cambricon achieved full-year profitability in 2025, generating Rmb20.6bn net income on Rmb65bn revenue. This milestone shows domestic AI chip companies can achieve sustainable profitability at scale, a critical validation for the sector.

Aggressive product roadmap maintains technology leadership. According to our industry checks, the company may launch and ship its next-generation MLU690 chip in 4Q26, delivering ~2.2x performance improvement (as gauged by our TPS result).

Concerns

Sanctions and capacity risk: The company was added to the U.S. Entity List in 2022, cutting off access to TSMC's 7nm manufacturing. It has since shifted to SMIC's N+2 process, but yield rates for large-die chips remain low, leading to supply constraints.

SMIC yield constraints limit capacity and margin upside. Our industry checks indicate MLU580 adopted local wafer manufacturing and is still experiencing limitations from SMIC's capacity and yields. This highlights the risk that slower-than-expected yield improvement and lower-than-expected capacity allocation could delay order fulfillment and pressure gross margins.

Rising competition in AI semi market. Competitors like Huawei Ascend and Iluvatar are gaining traction with CSP and government customers, which may compress Cambricon's market share, with potential price competition further squeezing margins.

High customer concentration creates revenue volatility. Given Bytedance is a key account, its capex and allocation decisions could create significant uncertainty for order flow and revenue. In addition the top five customers accounted for 94.63% of revenue in 2024, with Bytedance contributing 79.15%. The loss of any key customer could have a material impact on revenue.

EDA tool restriction risk: Chip design relies on EDA tools from Cadence and Synopsys. Further restrictions could impact product development cycles and iteration speed.

Cambricon: Where We Are Vs Consensus

We are more bullish than consensus on the company's revenue trajectory in 2026-27.

We believe Cambricon's 2026 revenue will still be partially constrained by domestic advanced logic foundry capacity, which will result in limited MLU580 shipments. That said, we still project 222% revenue growth for 2026. As domestic wafer production capacity constraints and advanced packaging bottlenecks are gradually alleviated, we forecast Cambricon's revenue will rise to Rmb33bn in 2027 and Rmb44bn in 2028, implying a 90% CAGR in 2025-28.

Turning to profitability, Cambricon is among the few domestic AI chip companies that have successfully navigated the early-stage high-investment and order-scarce phase, and has now achieved sustainable profitability traction. We project net income attributable to shareholders will reach Rmb6.2bn in 2026, Rmb9.6bn in 2027 and Rmb13.5bn in 2028, implying net margins of 29.4%, 28.8% and 30.6%, respectively.

Exhibit 26: Cambricon earnings estimates

US\$ mn	2026E			2027E			2028E		
	Mse	Consensus	Diff.	Mse	Consensus	Diff.	Mse	Consensus	Diff.
Net sales	20,944	14,057	49%	33,186	24,196	37%	44,256	47,904	-8%
Gross profit	10,588	7,703	37%	16,009	13,284	21%	20,801	25,606	-19%
Operating profit	6,240	5,221	20%	11,097	9,386	18%	15,772	17,994	-12%
Pretax income	6,356	5,111	24%	11,244	9,336	20%	15,934	17,994	-11%
Net income	6,156	5,042	22%	9,558	8,643	11%	13,544	14,361	-6%
EPS for consensus	14.49	12.19	19%	22.50	20.84	8%	31.88	34.37	-7%
Margins									
Gross margin	50.6%	54.8%		48.2%	54.9%		47.0%	53.5%	
Operating margin	29.8%	37.1%		33.4%	38.8%		35.6%	37.6%	
Pretax margin	30.3%	36.4%		33.9%	38.6%		36.0%	37.6%	
Net margin	29.4%	35.9%		28.8%	35.7%		30.6%	30.0%	
Opex %	20.8%	17.7%		14.8%	16.1%		11.4%	15.9%	

Cambricon: Valuation

We set our 12-month price target at Rmb1,588. With Cambricon now consistently profitable, we value the company using a residual income model, which we believe appropriately captures intrinsic value by fully incorporating our earnings forecasts.

Key assumptions: an 8.4% cost of equity (beta 1.06, risk-free rate 2.0%, equity risk premium 6.0%), a 26% long-term payout ratio, 16% medium-term growth, and a 6.0% perpetual terminal growth rate. We note that due to economic burdens, China is still in a low-interest-rate environment. Therefore, we use a relatively low 2% as the risk-free rate. Regarding the perpetual growth rate, the three Chinese AI accelerator companies covered in this report are all relatively young, with their establishment less than a decade ago. Amid the rapid development of AI, we believe these firms will maintain a phase of high growth and volume expansion over the next ten years, justifying a perpetual growth rate of 6%. Furthermore, when compared with the AI semiconductor leader NVIDIA, the market capitalization of companies such as Cambricon may only equate to 1.5% of NVIDIA's level. This also appears low, given China's strong push to develop large language models and the broader AI market.

Our price target implies a 110x 2026e P/E multiple and 32x 2026e P/S.

Exhibit 27: Cambricon: Residual income valuation

Rmb million	2026E	2027E	2028E	2029E	2030E	2031E	2032E	2033E	2034E	2035E	2036E	2037E
Total Equity	17,499	25,558	36,602	48,228	61,715	77,360	95,507	116,559	140,978	169,305	202,164	240,280
Net Profit	6,156	9,558	13,544	18,711	24,225	31,141	39,524	48,448	58,999	70,279	83,404	99,509
ROAE	42.0%	44.4%	43.6%	37.0%	33.2%	30.4%	28.4%	26.8%	25.6%	24.7%	23.9%	23.3%
Residual Income	3,979	6,306	9,001	10,498	11,958	13,604	15,482	17,639	20,126	22,999	26,322	30,170
Spread	33.6%	36.0%	35.2%	28.7%	24.8%	22.0%	20.0%	18.5%	17.3%	16.3%	15.5%	14.9%
Ending Equity Capital	17,499											
PV of Forecast Period	96,661											
PV of Continuing Value	560,292											
Equity Value	674,452											
No. of Shares	425											
Price Target	1,588											

Bull & bear cases

Our bull case value is Rmb3,000, and assumes a stronger revenue ramp with accelerated gross margin expansion and domestic AI chip substitution.

We assume (1) >120% revenue CAGR in 2025-28 driven by stronger-than-expected domestic generative AI infrastructure build-out and large-scale procurement of cloud training & inference chips; (2) sustained market share gains in China's high-performance AI chip market, with breakthroughs in long-term mass procurement contracts from key internet giants and state-owned enterprise customers; (3) gross margin above 55% in 2026 and 2027 thanks to production scale effects, mature technology iteration, and an optimized high-value product mix.

Our bear case value is Rmb800, and assumes slower revenue growth with gross margin erosion amid intensifying competition.

We assume (1) <60% revenue CAGR in 2025-28 given slower-than-expected domestic AI capex and LLM commercialization; (2) market share loss in China's AI chip market and failure to secure continuous mass orders from top-tier customers; (3) gross margin below

4.0% in 2026 and 2027 due to fierce price competition, rising wafer fabrication costs, and underutilized capacity.

Peer comparison

We expect Cambricon's EPS to increase at an 87% CAGR from 2025 to 2028. We also project revenue to deliver a 90% CAGR over the same period. We attribute the higher earnings and revenue growth to surging capex in domestic AI computing infrastructure and generative AI capacity build-out by major clients – including CSPs, telecom operators, and state-owned enterprises – as well as accelerated domestic substitution of overseas high-end AI chips amid supply chain restrictions. In our view, the higher earnings growth justifies Cambricon's high multiples of 110x 2026e P/E and 32x 2026e P/S.

Exhibit 28: China AI semi peer comparison table

Ticker	Name	Market Cap (US\$ bn)	Fab capacity implied revenue (2026, US\$ bn)	Consensus revenue (2026, US\$ bn)	2026 P/S	2026 P/E
Global						
NVDA-US	Nvidia	4,851	269	302	18	25
AMD-US	AMD	498	13	16	38	47
China						
688256-CN	Cambricon*	82	3.0	2.1	27	112
688041-CN	Hygon*	89	1.1	n.a.	79	126
688802-CN	MetaX*	39	0.6	n.a.	68	n.a.
688795-CN	Moore Threads*	42	0.3	n.a.	146	n.a.
6082-HK	Biren*	14	0.3	n.a.	48	n.a.
9903-HK	Iluvatar CoreX*	14	0.4	n.a.	31	n.a.

Exhibit 29: Cambricon P/S multiple





Exhibit 30: Cambricon P/E multiple



Iluvatar (9903.HK): Leveraging supply chain resilience with strong order visibility; Initiate at OW

Iluvatar is a domestic GPU designer focused on general-purpose AI computing, with products spanning AI training and inference. We initiate coverage with an Overweight rating, as we believe the company is well positioned to benefit from accelerating domestic substitution, supported by improving supply chain resilience, secured CSP orders, and a visible path to profitability. In particular, upcoming shipment ramp of the TianGai-150 chips from 2Q26, with a stronger ramp in 2H26, provides clear near-term revenue visibility. The company's diversified foundry strategy, including access to TSMC, offers greater supply stability relative to peers, while its high compatibility with the CUDA ecosystem lowers customer migration barriers. In addition, a clear trajectory toward profitability, supported by operating leverage, and a product portfolio covering both training and inference workloads position Iluvatar to capture broader demand across China's AI computing market.

Earnings: Our EPS estimates for 2026–27 are Rmb-1.03 and Rmb5.7. We expect strong revenue growth driven by TianGai shipment ramp and CSP demand, with margin expansion supported by operating leverage as scale improves. (Note that Iluvatar's IPO was in Jan 2026, so there are, as yet, no meaningful consensus estimates.)

Valuation: We derive our price target using a residual income model. Our price target implies 44x 2026e P/S, which we believe is justified given the 122% revenue CAGR that we forecast for 2025-28 and an improving profitability profile.

Key downside risks include: 1) delays in TianGai shipment ramp, 2) tightening export controls impacting supply chain access, and 3) intensifying competition from both domestic peers and global incumbents.

Exhibit 31: Iluvatar: Summary of key metrics

9903.HK Greater China Technology Semiconductors				
Price target				HK\$800
Up/downside to price target (%)				42%
Shr price, close(Apr 26, 2026)				457
52-Week Range				HK\$148.9-468.4
Sh out, dil, curr (mn)				254
Mkt cap, curr (mn)				116,223
EV, curr (mn)				85,840
Avg daily trading value (bn)				2
Fiscal Year Ending	12/25	12/26e	12/27e	12/28e
EPS for consensus(Rmb)	(5.31)	(1.03) e	5.70 e	10.10 e
Consensus EPS(Rmb)\$		NM	NM	NM
Revenue, net (Rmb mn)	1,034	3,460	7,828	11,192
EBITDA (Rmb mn)	(876)	52	1,802	2,891
ModelWare net inc. (Rmb mn)	(1,004)	(262)	1,449	2,729
P/E	NM	NM	56.7	36.2
P/BV	-	16.7	12.9	9.5
RNOA (%)	(88.8)	(6.2)	30.8	29.6
ROE (%)	(145.7)	(4.3)	29.5	35.7
EV/EBITDA	NM	1775.3	51.7	31.8
Div yld (%)	-	0.0	0.0	0.0
FCF yld ratio (%)	-	1.1	(1.4)	1.2
Leverage (EOP) (%)	(24.3)	(30.1)	(5.9)	(16.2)

Unless otherwise noted, all metrics are based on Morgan Stanley ModelWare framework
 § = Consensus data is provided by Thomson Reuters Estimates
 e = Morgan Stanley Research estimates

Company description
 Iluvatar CoreX Semiconductor Co., Ltd. is a leading Chinese general-purpose GPU (GPGPU) design company headquartered in Shanghai. Founded in 2018 by a team with an AMD background, the company has established itself as one of the most advanced domestic AI chip developers, with end-to-end capability spanning architecture design, chip development, and software stack optimization.

Source:

Iluvatar: Financial Summary

Exhibit 32: Iluvatar: Financial summary

Income Statement				
Rmb mn (Years End Dec)	2025	2026E	2027E	2028E
Net sales	1,033.6	3,060.0	7,500.8	11,378.9
COGS	(475.6)	(1,488.4)	(3,650.5)	(5,608.1)
Gross profit	558.0	1,571.6	3,850.2	5,770.8
Operating expenses	(1,607.5)	(1,880.0)	(2,450.0)	(3,090.0)
Operating income	(1,049.5)	(308.4)	1,400.2	2,680.8
Non-operating income	45.9	46.5	49.3	48.0
Pre-tax income	(1,003.7)	(261.9)	1,449.5	2,728.8
Income tax	0.0	0.0	0.0	159.3
Minority Interest	0.0	0.0	0.0	0.0
Reported net income	(1,003.7)	(261.9)	1,449.5	2,569.5
Adj.wtd.avg.shrs(m)	189.0	254.3	254.3	254.3
Reported EPS (Rmb)	(5.3)	(1.0)	5.7	10.1
EPS for consensus (Rmb)	(5.3)	(1.0)	5.7	10.1

Balance Sheet				
Rmb mn (Years End Dec)	2025	2026E	2027E	2028E
Cash	1,504.7	2,655.4	1,051.5	1,891.7
Mkt Securities	0.0	0.0	0.0	0.0
AR/NR	630.1	830.1	1,030.1	1,230.1
Inventory	709.8	1,631.1	3,200.5	3,841.1
Other	587.5	1,352.3	2,887.9	3,751.9
Current Assets	3,432.0	6,468.9	8,170.0	10,714.9
Long-term investments	76.9	76.2	76.2	76.2
Fixed assets	187.7	287.7	387.7	487.7
Intangible Assets	189.8	289.8	389.8	489.8
Other assets	25.5	25.5	25.5	25.5
Total Assets	3,912.0	7,148.1	9,049.2	11,794.1
S/T borrowings	643.6	643.6	643.6	643.6
AP/NP	291.0	448.5	900.1	1,075.5
Other ST liabilities	177.8	177.8	177.8	177.8
LT debt	365.4	365.4	365.4	365.4
Other LT liabilities	81.1	81.1	81.1	81.1
Total Liabilities	1,558.9	1,716.5	2,168.1	2,343.5
Common shares	228.9	254.3	254.3	254.3
Additional capital	0.0	3,315.0	3,315.0	3,315.0
Retained earning	2,162.1	1,900.2	3,349.7	5,919.2
Other shareholders' equity	(37.9)	(37.9)	(37.9)	(37.9)
Total Equity	2,353.1	5,431.6	6,881.1	9,450.6
Total Liab. & Shrhldr's Equity	3,912.0	7,148.1	9,049.2	11,794.1

Cash Flow Statement				
Rmb mn (Years End Dec)	2025	2026E	2027E	2028E
Cashflow from Operations	(1,161.6)	(1,790.4)	(1,204.0)	1,240.2
Net profits	(1,003.7)	(261.9)	1,449.5	2,569.5
Depreciation	97.0	100.0	100.0	100.0
Working Capital Change	(823.6)	(1,728.5)	(2,853.4)	(1,529.3)
Other adjustments	568.8	100.0	100.0	100.0
Cashflow from Investing	(132.3)	(399.3)	(400.0)	(400.0)
Capex	(59.7)	(200.0)	(200.0)	(200.0)
Change of LT Investment	(49.1)	0.7	0.0	0.0
Change of ST Investment	0.1	0.0	0.0	0.0
Other adjustments	(23.6)	(200.0)	(200.0)	(200.0)
Cashflow from financing	2,490.9	3,340.4	0.0	0.0
Increase in L/T debt	323.4	0.0	0.0	0.0
Increase in S/T debt	77.6	0.0	0.0	0.0
Cash Dividend Paid	0.0	0.0	0.0	0.0
Issuance of stock	38.0	3,340.4	0.0	0.0
Other adjustments	2,051.9	0.0	0.0	0.0
Exchange rate adjustment	(5.9)	0.0	0.0	0.0
Net change in cash	1,197.1	1,150.7	(1,604.0)	840.2

Financial Ratios				
	2025	2026E	2027E	2028E
Growth(%)				
Revenue	91.6	196.1	145.1	51.7
Operating profits	18.3	(70.6)	NA	91.5
Pretax profits	12.5	(73.9)	NA	88.3
Net profits	12.5	(73.9)	NA	77.3
EPS	(2.6)	(80.6)	NA	77.3
Margins (%)				
Gross Margin	54.0	51.4	51.3	50.7
Operating Margin	(101.5)	(10.1)	18.7	23.6
Pretax Margin	(97.1)	(8.6)	19.3	24.0
Net Profit	(97.1)	(8.6)	19.3	22.6
Return (%)				
ROAE	(66.0)	(6.7)	23.5	31.5
ROAA	(35.9)	(4.7)	17.9	24.7
Gearing (%)				
Net Debt/Equity	(21.1)	(30.3)	(0.6)	(9.3)
Liabilities/Equity	66.2	31.6	31.5	24.8
Ratios (X)				
Current ratio	3.1	5.1	4.7	5.6
Quick ratio	1.9	2.7	1.2	1.6
Others				
AR/NR Turnover (days)	168.4	160.0	140.0	120.0
Inventory Turnover (days)	403.8	400.0	320.0	250.0
AP Turnover (days)	29.5	10.0	10.0	10.0
Cash Conversion (days)	542.8	550.0	450.0	360.0

Risk Reward – Iluvatar CoreX Semiconductor Co., Ltd. (9903.HK)


Leveraging supply chain resilience with strong order visibility

PRICE TARGET HK\$600.00

Key valuation assumptions include:

- An 8.3% cost of equity (derived from a beta of 1.05, risk-free rate of 2.0%, and equity risk premium of 6.0%)
- A long-term payout ratio of 33%
- A medium-term growth rate of 16%
- A perpetual terminal growth rate of 6%

RISK REWARD CHART



OVERWEIGHT THESIS

- Strong order visibility from domestic CSPs, with TianGai-150 shipments ramping from 2Q26 and contributing >Rmb4bn revenue over 2026–27.
- Diversified foundry strategy ensures supply security, with access to TSMC capacity mitigating risks from export controls and domestic yield constraints.
- High CUDA compatibility enables low-friction migration, with successful LLM deployment validating product-market fit.
- Clear path to profitability, with breakeven expected in 2026 and full-year profitability in 2027 driven by operating leverage and scaling.
- Our price target implies 44x 2026e P/S, which is lower than the peer average of 75x 2026e P/S.

Risk Reward Themes

New Data Era: *Positive*

Technology Diffusion: *Positive*

View descriptions of Risk Rewards Themes [here](#)

BULL CASE

HK\$1,100.00

80x 2026e P/S

We assume: 1) >150% revenue CAGR in 2025-28e driven by a more aggressive domestic AI infrastructure build-out and large-scale procurement of both training and inference chips from Chinese cloud service providers and state-owned enterprises. 2) Gross margin improves to over 60% in 2026e and 2027e thanks to production scale effects, mature technology iteration and optimized high-value product structure. 3) The company achieves full-year profitability in 2026, ahead of our base case.

BASE CASE

HK\$600.00

44x 2026e P/S

We assume: 122% revenue CAGR in 2025-28e given strong orders from CSP clients and solid supply chain capacity. 2) Gross margin at 51.4% in 2026e and 51.3% 2027e due to slight price competition. 3) The company achieve full-year profitability in 2027.

BEAR CASE

HK\$300.00

22x 2026e P/S

We assume: <80% revenue CAGR in 2025-28e given slower-than-expected domestic AI capital expenditure and large model commercialization progress. 2) Nvidia relaxes export restrictions on mid-range AI chips to China, leading to intensified price competition. 3) Gross margin falls below 45% in 2026e and 2027e due to fierce price competition. 4) The company does not achieve full-year profitability until 2028.

Risk Reward – Iluvatar CoreX Semiconductor Co., Ltd. (9903.HK)


KEY EARNINGS INPUTS

Drivers	2025	2026e	2027e	2028e
Operating profit (YoY) (%)	18.3	(70.6)	(554.1)	91.5

INVESTMENT DRIVERS

- CSP AI capex growth
- TianGai shipment ramp
- CUDA migration trend
- Supply chain advantage (TSMC)
- AI inference expansion
- Policy support
- Customer base expansion
- Margin improvement trajectory

GLOBAL REVENUE EXPOSURE



View explanation of regional hierarchies [here](#)

RISKS TO PT/RATING

RISKS TO UPSIDE

- Stronger-than-expected CSP orders
- Faster CUDA replacement with Iluvatar's software
- Expansion of overseas and domestic capacity

RISKS TO DOWNSIDE

- Order ramp below expectations
- Escalation of sanctions
- Intensifying competition

MS ESTIMATES VS. CONSENSUS

FY 2026e

Sales / Revenue (Rmb, mn) ◆ 3,060
Note: There are not sufficient brokers supplying consensus data for this metric

Net income (Rmb, mn) ◆ (262)
Note: There are not sufficient brokers supplying consensus data for this metric

EPS (Rmb) ◆ (1.03)
Note: There are not sufficient brokers supplying consensus data for this metric

◆ Mean ◆

Iluvatar: Investment Positives & Concerns

Positives

Upcoming order inflection from domestic cloud providers. According to our industry checks, major Chinese cloud service providers placed significant preorders for Iluvatar's TianGai-150 AI chips, with shipments set to begin in Q2 2026 and ramp sharply in 2H26. We estimate these orders will contribute over Rmb4bn to Iluvatar's 2026 and 2027 revenue.

Industry-leading supply chain security. Unlike domestic peers that rely exclusively on SMIC for advanced nodes, Iluvatar maintains a diversified foundry strategy and produces fully compliant AI chips at TSMC. This provides greater capacity certainty and eliminates production-disruption risk from further export restrictions.

Superior software compatibility and migration capabilities. Iluvatar's GPGPU chips offer near-seamless compatibility with NVIDIA's CUDA ecosystem, significantly reducing customer migration costs. The company has helped clients to migrate LLM stacks from NVIDIA platforms to TianGai-150.

Clear and imminent profitability inflection. We expect Iluvatar is on track to achieve positive adjusted net income in Q3 2026 and full-year profitability in 2027. Strong revenue growth with operating leverage should drive this outcome, i.e., R&D and SG&A as a percentage of revenue decline further in 2026-2027.

Comprehensive product portfolio covering training and inference. Iluvatar also offers edge AI inference products, the Zhikai series, enabling it to capture broader customer demand and provide end-to-end AI computing solutions.

Concerns

TianGai-150 order execution and Tiangai-300 validation risk. Approximately 65% of our 2026 revenue forecast depends on successful TianGai-150 deployments. Also, future growth depends on the Tiangai-300 product. Rigorous 3–6 month customer validation and TSMC CoWoS packaging constraints could delay order placement and revenue recognition.

TSMC supply chain remains vulnerable to US export controls. While Iluvatar's TSMC partnership provides a competitive advantage, further US restrictions on advanced node manufacturing for Chinese companies would severely impact operations.

Profitability inflection could be delayed. Aggressive R&D spending on TianXuan and Tianji architectures and intensifying price competition could push full-year profitability beyond 2027.

Domestic competition intensifies across segments. Major competitors – including Huawei Ascend 950/960, Cambricon MLU580/690, and MetaX's upcoming C600 – will directly compete for cloud and government orders. We estimate Iluvatar's top line to reach Rmb7.5bn by 2027, but aggressive peer pricing could limit upside.

Iluvatar: Earnings Estimates

We believe Iluvatar's 2026 revenue will be driven by strong demand for its TianGai-150 chip, with supply constraints far less severe than for domestic peers due to diversified foundry partnerships. We forecast revenue of Rmb3.06bn in 2026, up 196% YoY. As the TianGai-300 chip enters full-scale production in 2H26 and mass shipment in 2027, we project revenue of Rmb7.5bn in 2027, and Rmb11.4bn in 2028, implying a 122% CAGR over 2025-28.

Turning to profitability, we expect positive net income in 2H26 and full-year profitability in 2027. We project net income attributable to shareholders of -Rmb262mn in 2026, Rmb1.45bn in 2027, and Rmb2.73bn in 2028, implying net margins of -8.6%, 19.3%, and 22.6%, respectively.

Exhibit 33: Iluvatar: Earnings estimates

(Rmb mn)	1H25	2H25	1H26E	2H26E	1H27E	2H27E	1H28E	2H28E	2025	2026E	2027E	2028E
Total revenues	324.3	709.3	976.0	2,084.0	3,004.6	4,496.2	5,097.1	6,281.8	1,033.6	3,060.0	7,500.8	11,378.9
HH Change	-5.2%	118.8%	37.6%	113.5%	44.2%	49.6%	13.4%	23.2%				
YY Change	64.2%	107.4%	201.0%	193.8%	207.5%	115.7%	69.6%	39.7%	91.6%	196.1%	145.1%	51.7%
Cost of Sales	(161.8)	(313.8)	(490.1)	(998.3)	(1,469.2)	(2,181.3)	(2,515.2)	(3,092.9)	(475.6)	(1,488.4)	(3,650.5)	(5,608.1)
Percent of Revenues	50%	44%	50%	48%	49%	49%	49%	49%	46%	49%	49%	49%
Gross Profit	162.4	395.6	485.9	1,085.7	1,535.3	2,314.9	2,581.9	3,188.9	558.0	1,571.6	3,850.2	5,770.8
Gross Margin	50.1%	55.8%	49.8%	52.1%	51.1%	51.5%	50.7%	50.8%	54.0%	51.4%	51.3%	50.7%
Incremental Margin	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Total Opex	(793.7)	(813.8)	(920.0)	(960.0)	(1,160.0)	(1,290.0)	(1,470.0)	(1,620.0)	(1,607.5)	(1,880.0)	(2,450.0)	(3,090.0)
Percent of Revenues	244.8%	114.7%	94.3%	46.1%	38.6%	28.7%	28.8%	25.8%	155.5%	61.4%	32.7%	27.2%
R&D	(451.5)	(522.7)	(550.0)	(580.0)	(610.0)	(640.0)	(700.0)	(750.0)	(974.2)	(1,130.0)	(1,250.0)	(1,450.0)
Percent of Revenues	139.2%	73.7%	56.4%	27.8%	20.3%	14.2%	13.7%	11.9%	94.2%	36.9%	16.7%	12.7%
General & Adm Exp.	(274.6)	(207.2)	(220.0)	(230.0)	(300.0)	(350.0)	(420.0)	(470.0)	(481.8)	(450.0)	(650.0)	(890.0)
Percent of Revenues	84.7%	29.2%	22.5%	11.0%	10.0%	7.8%	8.2%	7.5%	46.6%	14.7%	8.7%	7.8%
Selling Expenses	(67.6)	(83.9)	(150.0)	(150.0)	(250.0)	(300.0)	(350.0)	(400.0)	(151.5)	(300.0)	(550.0)	(750.0)
Percent of Revenues	20.9%	11.8%	15.4%	7.2%	8.3%	6.7%	6.9%	6.4%	14.7%	9.8%	7.3%	6.6%
Operating Income	(631.3)	(418.3)	(434.1)	125.7	375.3	1,024.9	1,111.9	1,568.9	(1,049.5)	(308.4)	1,400.2	2,680.8
Operating Margin	-194.7%	-59.0%	-44.5%	6.0%	12.5%	22.8%	21.8%	25.0%	-101.5%	-10.1%	18.7%	23.6%
Total Non-operating Income (loss)	21.9	23.9	23.5	23.0	25.6	23.7	23.9	24.0	45.9	46.5	49.3	48.0
Profit Before Taxes	(609.3)	(394.3)	(410.6)	148.7	400.9	1,048.5	1,135.9	1,592.9	(1,003.7)	(261.9)	1,449.5	2,728.8
Percent of Revenues	-187.9%	-55.6%	-42.1%	7.1%	13.3%	23.3%	22.3%	25.4%	-97.1%	-8.6%	19.3%	24.0%
Taxes	0.0	0.0	0.0	0.0	0.0	0.0	0.0	159.3	0.0	0.0	0.0	159.3
Tax Rate	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	10.0%	0.0%	0.0%	0.0%	5.8%
Total Net Income to Parent	(609.3)	(394.3)	(410.6)	148.7	400.9	1,048.5	1,135.9	1,433.6	(1,003.7)	(261.9)	1,449.5	2,569.5
Percent of Revenues	-187.9%	-55.6%	-42.1%	7.1%	13.3%	23.3%	22.3%	22.8%	-97.1%	-8.6%	19.3%	22.6%
EPS for consensus (Rmb)	(3.5)	(1.8)	(1.6)	0.6	1.6	4.1	4.5	5.6	(5.3)	(1.0)	5.7	10.1

Iluvatar: Valuation

We set our price target at HK\$600. Given Iluvatar's near profitability and high-growth phase, we value the company with a residual income model, which we believe appropriately balances near-term growth momentum with long-term intrinsic value.

Key valuation assumptions include:

- An 8.3% cost of equity (derived from a beta of 1.05, risk-free rate of 2.0%, and equity risk premium of 6.0%)
- A long-term payout ratio of 33%
- A medium-term growth rate of 16%
- A perpetual terminal growth rate of 6%

Our price target implies a 44x 2026e P/S multiple.

Exhibit 34: Iluvatar: Residual income model

Rmb million	2026E	2027E	2028E	2029E	2030E	2031E	2032E	2033E	2034E	2035E	2036E	2037E
Total Equity	5,432	6,881	9,451	11,463	13,796	16,504	19,644	23,287	27,513	32,414	38,101	44,697
Net Profit	(262)	1,449	2,570	2,981	3,458	4,011	4,652	5,397	6,260	7,262	8,424	9,772
ROAE	NM	23.5%	31.5%	28.5%	27.4%	26.5%	25.7%	25.1%	24.6%	24.2%	23.9%	23.6%
Residual Income		828	1,594	1,909	2,187	2,507	2,878	3,308	3,807	4,384	5,054	5,831
Spread		15.2%	23.2%	20.2%	19.1%	18.2%	17.4%	16.8%	16.3%	15.9%	15.6%	15.3%
Ending Equity Capital	5,432											
PV of Forecast Period	17,004											
PV of Continuing Value	111,790											
Equity Value	134,226											
No. of Shares	254											
HKD/RMB	0.88											
Price Target (HK\$)	600											

Peer comparison


We expect Iluvatar's revenue to grow at a 122% CAGR from 2025 to 2028, quite high among domestic peers. We attribute this superior earnings and revenue growth to:

- Iluvatar's diversified supply chain strategy, allowing use of both domestic and international foundries, resulting in significantly fewer capacity constraints than competitors
- Leading software stack compatibility and Prefill and Decode separation technology, significantly reducing customer migration costs and improving inference efficiency.

In our view, superior top-line momentum, a leading position in China's domestic general-purpose GPU segment, and improving revenue visibility from long-term framework adaptation and customer lock-in support a premium multiple of 44x 2026e P/S.

Exhibit 35: China AI semi peer comparison table

Ticker	Name	Market Cap (US\$ bn)	Fab capacity implied revenue (2026, US\$ bn)	Consensus revenue (2026, US\$ bn)	2026 P/S	2026 P/E
Global						
NVDA-US	Nvidia	4,851	269	302	18	25
AMD-US	AMD	498	13	16	38	47
China						
688256-CN	Cambricon*	82	3.0	2.1	27	112
688041-CN	Hygon*	89	1.1	n.a.	79	126
688802-CN	MetaX*	39	0.6	n.a.	68	n.a.
688795-CN	Moore Threads*	42	0.3	n.a.	146	n.a.
6082-HK	Biren*	14	0.3	n.a.	48	n.a.
9903-HK	Iluvatar CoreX*	14	0.4	n.a.	31	n.a.

Exhibit 36: Iluvatar P/S multiple

Bull & bear cases

Our bull case value is HK\$1,100, and assumes accelerated domestic substitution with faster-than-expected product performance and market share gains

We assume:

- A >150% revenue CAGR in 2025-28, driven by more aggressive domestic AI infrastructure build-out and large-scale procurement of training and inference chips from Chinese cloud service providers and state-owned enterprises.
- Gross margin rises to over 60% in 2026 and 2027, driven by scale, mature technology iteration, and an optimized high-value product mix.
- Full-year profitability in 2026, ahead of 2027 in our base case.

Our bear case value is HK\$300, and assumes slower revenue growth with gross margin compression amid intensifying competition and delayed product launches.

We assume:

- A <80% revenue CAGR in 2025-28 given slower-than-expected domestic AI capital expenditure and large model commercialization.
- Nvidia relaxes export restrictions on mid-range AI chips to China, intensifying price competition.
- Gross margin falls below 45% in 2026 and 2027 due to fierce price competition, rising wafer fabrication costs, and underutilized capacity.
- Full-year profitability not achieved until 2028, one year later than our base case.

MetaX (688802.SS): Positioning for scalability through software and supply; Initiate at EW

MetaX Integrated Circuits is a domestic GPU designer focused on high-performance training chips, with a growing presence in China's AI computing market. We initiate coverage with an Equal-weight rating, as we see balanced risk-reward driven by solid execution in software compatibility and supply chain, alongside uncertainties in scale-up and customer diversification. The company's proprietary MXMACA platform offers high compatibility with NVIDIA's CUDA ecosystem, lowering switching costs and supporting gradual customer adoption. In addition, MetaX has secured relatively stable demand from government and enterprise customers, which tends to carry higher margins than CSP orders. Its production on SMIC's N+1 process provides supply chain visibility and yield stability. That said, further validation in large-scale deployments and broader customer penetration remain key to sustaining growth.

Earnings: Our EPS estimates for 2026–27 are Rmb0.05 and Rmb1.33. We expect revenue growth to be supported by continued C500 shipments and initial C600 ramp, with margin trends dependent on product mix and scale effects. (Note that MetaX's IPO was in Dec 2025, so there are, as yet, no meaningful consensus estimates.)

Valuation: We derive our price target using a residual income model. Our price target implies a 75x 2026e P/S multiple, which we believe reflects the 66% revenue CAGR that we project for 2025–28, balanced against execution risks and competitive pressures.

Key downside risks include: 1) slower-than-expected C600 commercialization, 2) customer concentration and demand variability, and 3) intensifying competition in the domestic GPU market.

Exhibit 37: MetaX: Summary of key metrics

688802.SS Greater China Technology Semiconductors				
Price target				Rmb758
Up/downside to price target (%)				9%
Shr price, close(Apr 26, 2026)				718
52-Week Range				Rmb479.01-895
Sh out, dil, curr (mn)				400
Mkt cap, curr (mn)				287,272
EV, curr (mn)				277,972
Avg daily trading value (bn)				5
Fiscal Year Ending	12/25	12/26e	12/27e	12/28e
EPS for consensus(Rmb)	(2.11)	0.05 e	1.33 e	2.50 e
Consensus EPS(Rmb)§		NM	NM	NM
Revenue, net (Rmb mn)	1,644	4,030	6,209	7,483
EBITDA (Rmb mn)	(826)	(78)	413	861
ModelWare net inc. (Rmb mn)	(781)	22	533	1,001
P/E	NM	12948.7	523.6	278.9
P/BV	16.3	21.2	20.4	19.0
RNOA (%)	(104.5)	(1.6)	5.8	9.1
ROE (%)	(66.3)	0.2	4.0	7.3
EV/EBITDA	NM	NM	659.7	316.0
Div yld (%)	0.0	0.0	0.0	0.0
FCF yld ratio (%)	(1.9)	(0.6)	(0.5)	0.1
Leverage (EOP) (%)	(71.4)	(59.2)	(48.5)	(48.6)
Unless otherwise noted, all metrics are based on Morgan Stanley ModelWare framework				
§ = Consensus data is provided by Thomson Reuters Estimates				
e = Morgan Stanley Research estimates				
Company description				
MetaX Integrated Circuits (Shanghai) Co., Ltd. is a leading domestic GPU design company in China, specializing in high-performance general-purpose GPUs (GPGPU) for AI training, inference, and graphics rendering.				

Source:

MetaX: Financial Summary

Exhibit 38: MetaX: Financial summary

Income Statement					Cash Flow Statement				
Rmb mn (Years End Dec)	2025	2026E	2027E	2028E	Rmb mn (Years End Dec)	2025	2026E	2027E	2028E
Net sales	1,644.1	4,029.9	6,208.6	7,482.9	Cashflow from Operations	(3,927.5)	(1,571.4)	(1,139.2)	523.9
COGS	741.6	1,861.3	2,905.6	3,540.3	Net profits	(781.5)	21.6	533.0	1,000.7
Gross profit	902.5	2,168.6	3,303.0	3,942.7	Depreciation	0.0	0.0	0.0	0.0
Operating expenses	1,728.1	2,247.0	2,890.0	3,082.0	Working Capital Change	(3,146.1)	(1,593.0)	(1,672.2)	(476.8)
Operating income	(825.6)	(78.4)	413.0	860.7	Other adjustments	0.0	0.0	0.0	0.0
Non-operating income	(54.2)	100.0	120.0	140.0	Cashflow from Investing	(3,969.8)	(20.0)	1,480.0	480.0
Pre-tax income	(771.4)	21.6	533.0	1,000.7	Capex	0.0	0.0	0.0	0.0
Income tax	10.0	0.0	0.0	0.0	Change of LT Investment	(20.0)	(20.0)	(20.0)	(20.0)
Minority Interest	3.0	4.0	5.0	6.0	Change of ST Investment	(3,949.8)	0.0	1,500.0	500.0
Reported net income	(781.5)	21.6	533.0	1,000.7	Other adjustments	0.0	0.0	0.0	0.0
Adj. wtd. avg. shrs (m)	370.0	400.0	400.0	400.0	Cashflow from financing	12,439.8	0.0	0.0	0.0
Reported EPS (Rmb)	(2.1)	0.1	1.3	2.5	Increase in L/T debt	0.0	0.0	0.0	0.0
EPS for consensus (Rmb)	(2.1)	0.1	1.3	2.5	Increase in S/T debt	(314.6)	0.0	0.0	0.0
					Cash Dividend Paid	0.0	0.0	0.0	0.0
					Issuance of stock	12,753.9	0.0	0.0	0.0
					Other adjustments	0.5	0.0	0.0	0.0
					Exchange rate adjustment	0.0	0.0	0.0	0.0
					Net change in cash	4,542.5	(1,591.4)	340.8	1,003.9
Balance Sheet					Financial Ratios				
Rmb mn (Years End Dec)	2025	2026E	2027E	2028E		2025	2026E	2027E	2028E
Cash	5,571.7	3,980.2	4,321.0	5,324.9	Growth(%)				
Mkt Securities	4,000.0	4,000.0	2,500.0	2,000.0	Revenue	121.3	145.1	54.1	20.5
AR/NR	900.9	1,656.1	2,551.5	2,050.1	Operating profits	(31.6)	(90.5)	(626.5)	108.4
Inventory	1,828.6	2,549.8	3,184.2	3,879.7	Pretax profits	(45.1)	(102.8)	2,372.9	87.7
Other	1,199.6	1,699.6	2,199.6	2,699.6	Net profits	(44.5)	(102.8)	2,372.9	87.7
Current Assets	13,500.7	13,885.8	14,756.4	15,954.4	EPS	(69.9)	(102.6)	2,372.9	87.7
Long-term investments	1.2	1.2	1.2	1.2	Margins (%)				
Fixed assets	183.8	183.8	183.8	183.8	Gross Margin	54.9	53.8	53.2	52.7
Intangible Assets	195.8	215.8	235.8	255.8	Operating Margin	(50.2)	(1.9)	6.7	11.5
Other assets	103.7	103.7	103.7	103.7	Pretax Margin	(46.9)	0.5	8.6	13.4
Total Assets	13,985.2	14,390.2	15,280.9	16,498.9	Net Profit	(47.5)	0.5	8.6	13.4
S/T borrowings	180.0	180.0	180.0	180.0	Return (%)				
AP/NP	254.0	637.4	995.1	1,212.4	ROAE	(10.9)	0.2	4.0	7.0
Other ST liabilities	300.0	300.0	300.0	300.0	ROAA	(8.7)	0.2	3.6	6.3
LT debt	0.0	0.0	0.0	0.0	Gearing (%)				
Other LT liabilities	100.0	100.0	100.0	100.0	Net Debt/Equity	(41.0)	(28.8)	(30.2)	(35.0)
Total Liabilities	834.0	1,217.4	1,575.1	1,792.4	Liabilities/Equity	6.3	9.2	11.5	12.2
Common shares	400.0	400.0	400.0	400.0	Ratios (X)				
Additional capital	14,347.7	14,347.7	14,347.7	14,347.7	Current ratio	18.4	12.4	10.0	9.4
Retained earning	(1,596.5)	(1,574.9)	(1,041.9)	(41.2)	Quick ratio	8.8	5.0	4.7	4.4
Other shareholders' equity	0.0	0.0	0.0	0.0	Others				
Total Equity	13,151.2	13,172.8	13,705.8	14,706.5	AR/NR Turnover (days)	200.0	150.0	150.0	100.0
Total Liab. & Shrhldr's Equity	13,985.2	14,390.2	15,280.9	16,498.9	Inventory Turnover (days)	900.0	500.0	400.0	400.0
					AP Turnover (days)	125.0	125.0	125.0	125.0
					Cash Conversion (days)	975.0	525.0	425.0	375.0
E = Morgan Stanley Research Estimates									
Source: Morgan Stanley Research, Company Data									


Risk Reward – MetaX Integrated Circuits (688802.SS)

Positioning for scalability through software and supply

PRICE TARGET Rmb758.00

Key valuation assumptions include: an 6.5% cost of equity (derived from a beta of 0.75, risk-free rate of 2.0%, and equity risk premium of 6.5%), a long-term target payout ratio of 50%, a medium-term CAGR of 18% through 2038e, and a perpetual terminal growth rate of 6%.

RISK REWARD CHART



EQUAL-WEIGHT THESIS

- Strong CUDA compatibility via MXMACA, lowering migration barriers and supporting faster adoption among existing NVIDIA users.
- Stable government and enterprise orders provide revenue visibility, with higher-margin projects supporting near-term gross margin profile.
- Mature domestic supply chain with SMIC N+1 production, ensuring capacity availability and relatively stable yields versus peers.
- Our target price implies 75x 2026e P/S, which is higher than the peer average.

Risk Reward Themes

New Data Era: Positive
Technology Diffusion: Positive

View descriptions of Risk Rewards Themes [here](#)

BULL CASE

Rmb1,500.00

149x 2026e P/S

We assume (1) >100% revenue CAGR in 2025-28e on the back of faster-than-expected mass production and delivery of flagship GPGPUs, as well as booming demand for domestic general computing and AI training chips; (2) significant market share expansion in China's domestic AI semi market, with full adaptation to mainstream AI frameworks and sovereign customers; (3) gross margin improves to over 60% in 2026e.

BASE CASE

Rmb758.00

75x 2026e P/S

We assume (1) 66% revenue CAGR in 2025-28e on the back of mass production and delivery of C600 and C700 AI GPU; (2) Market share expansion in China's domestic AI semi market, with rising orders from sovereign customers and entry into core CSPs; (3) gross margin reach 54% in 2026e.

BEAR CASE

Rmb380.00

38x 2026e P/S

We assume (1) <40% revenue CAGR in 2025-28e given delayed mass production schedule of flagship GPU products, slower-than-expected customer validation and weaker domestic AI computing capital expenditure. (2) Share loss in China's AI semi market amid intensifying domestic and overseas competition. (3) Gross margin falls below 45% in 2026e due to high R&D, rising wafer and packaging costs, and fierce price competition in the market.

Risk Reward – MetaX Integrated Circuits (688802.SS)


KEY EARNINGS INPUTS

Drivers	2025	2026e	2027e	2028e
Operating profit (YoY) (%)	(31.6)	(90.5)	(626.5)	108.4

INVESTMENT DRIVERS

- CUDA migration trend
- Government AI capex
- C500 shipment ramp and C600 commercialization
- Domestic substitution tailwind
- Customer expansion to CSP clients
- Yield and cost improvement

GLOBAL REVENUE EXPOSURE



View explanation of regional hierarchies [here](#)

RISKS TO PT/RATING

- RISKS TO UPSIDE**
- Faster CUDA migration
 - Stronger government and CSP orders
 - Better-than-expected C600 ramp
- RISKS TO DOWNSIDE**
- Product concentration risk
 - Weaker-than-expected demand
 - Intensifying competition

MS ESTIMATES VS. CONSENSUS

FY 2026e	
Sales / Revenue (Rmb, mn)	◆ 4,029.9 Note: There are not sufficient brokers supplying consensus data for this metric
Net income (Rmb, mn)	◆ 22 Note: There are not sufficient brokers supplying consensus data for this metric
EPS (Rmb)	◆ 0.1 Note: There are not sufficient brokers supplying consensus data for this metric
◆ Mean ◆	

MetaX: Investment Positives & Concerns

Positives

Industry-leading software compatibility with NVIDIA CUDA. MetaX's proprietary MXMACA software stack offers near-seamless CUDA compatibility, supporting over 6,000 CUDA APIs and enabling most applications to migrate to MetaX hardware. This could significantly reduce switching costs and accelerates adoption.

Stable government and enterprise orders. With a strong government investment background, MetaX has secured contracts with multiple national and provincial intelligent computing centers, delivering higher gross margins (over 60%) vs. cloud service provider orders.

Fully operational domestic supply chain with reasonable yields. Unlike peers still transitioning to domestic manufacturing, MetaX has ramped production on SMIC's N+1 process with sufficient yields, providing significant capacity advantages and supply chain security.

Focused strategy on training chips creates differentiation. While most domestic players compete across training and inference, MetaX has concentrated resources on high-performance training chips, enabling faster technology iteration and better product-market fit.

Concerns

Slow CSP penetration limits revenue upside. MetaX has secured major orders from top Chinese CSPs, such as Bytedance, yet it lags peers. Failure to gain meaningful share in this largest, fastest-growing segment would cap long-term growth.

Process node trade-off undermines core performance. MetaX's flagship GPUs use SMIC's N+1 process, a choice that delivers stable mass production and high yield, but carries structural limitations in transistor density and power efficiency vs. more advanced industry-standard nodes.

Over-reliance on SMIC creates supply chain concentration. The company now depends almost entirely on SMIC's N+1 process. Any disruption to SMIC's operations or capacity constraints could materially impact order fulfillment.

Competitive pressure weighs on market share. China's domestic GPGPU market has entered a hyper-competitive phase, with over 10 mass-production vendors triggering significant price-based competition at the mid-to-low-end, while rising expectations of a marginal relaxation of US GPU export controls to China pose additional downside risk.

MetaX: Earnings Estimates

We think revenue momentum will stay strong in 2026-28, driven by C600 and future product mass production, growing orders for the company's high-performance GPGPU products, full-stack adaptation to mainstream AI frameworks, and faster order conversion from core clients, including cloud service providers and sovereign platforms amid the ongoing domestic expansion cycle.

Our 2026/27/28 revenue growth forecasts are 145%/54%/21%, and we also expect the company to turn profitable for the first time in 2026, with EPS of Rmb0.05, followed by Rmb1.33 EPS in 2027.

Exhibit 39: MetaX: Earnings estimates

(Rmb m)	1Q25	2Q25	3Q25	4Q25	1Q26E	2Q26E	3Q26E	4Q26E	2025	2026E	2027E	2028E
Total revenues	320.4	594.5	321.2	408.0	506.3	775.3	1,195.4	1,552.9	1,644.1	4,029.9	6,208.6	7,482.9
Q/Q Change	-38.4%	85.5%	-46.0%	27.0%	24.1%	53.1%	54.2%	29.9%				
Y/Y Change			665.4%	-21.5%	58.0%	30.4%	272.2%	280.6%	121.3%	145.1%	54.1%	20.5%
Cost of Sales	143.4	267.5	145.1	185.6	231.9	356.6	552.3	720.5	741.6	1,861.3	2,905.6	3,540.3
Percent of Revenues	45%	45%	45%	46%	46%	46%	46%	46%	45%	46%	47%	47%
Gross Profit	177.1	327.0	176.1	222.4	274.4	418.6	643.1	832.3	902.5	2,168.6	3,303.0	3,942.7
Gross Margin	55.3%	55.0%	54.8%	54.5%	54.2%	54.0%	53.8%	53.6%	54.9%	53.8%	53.2%	52.7%
Incremental Margin	NM		NM	53.3%	53.0%	53.6%	53.4%	52.9%	56.1%	53.1%	52.1%	50.2%
Total Opex	401.5	370.0	336.6	620.0	440.0	497.0	554.0	756.0	1,728.1	2,247.0	2,890.0	3,082.0
Percent of Revenues	125.3%	62.2%	104.8%	152.0%	86.9%	64.1%	46.3%	48.7%	105.1%	55.8%	46.5%	41.2%
R&D	217.9	230.0	240.1	420.0	300.0	340.0	380.0	530.0	1,108.0	1,550.0	1,985.0	2,065.0
Percent of Revenues	68.0%	38.7%	74.8%	102.9%	59.2%	43.9%	31.8%	34.1%	67.4%	38.5%	32.0%	27.6%
General & Adm Exp.	152.7	100.0	51.9	150.0	85.0	100.0	115.0	165.0	454.6	465.0	645.0	725.0
Percent of Revenues	47.6%	16.8%	16.2%	36.8%	16.8%	12.9%	9.6%	10.6%	27.6%	11.5%	10.4%	9.7%
Selling Expenses	30.9	40.0	44.6	50.0	55.0	57.0	59.0	61.0	165.5	232.0	260.0	292.0
Percent of Revenues	9.7%	6.7%	13.9%	12.3%	10.9%	7.4%	4.9%	3.9%	10.1%	5.8%	4.2%	3.9%
Operating Income	(224.4)	(43.0)	(160.5)	(397.6)	(165.6)	(78.4)	89.1	76.3	(825.6)	(78.4)	413.0	860.7
Operating Margin	-70.0%	-7.2%	-50.0%	-97.5%	-32.7%	-10.1%	7.5%	4.9%	-50.2%	-1.9%	6.7%	11.5%
Total Non-operating Income (loss)	7.9	(92.3)	(2.3)	32.5	25.0	25.0	25.0	25.0	(54.2)	100.0	120.0	140.0
Profit Before Taxes	(232.3)	49.3	(158.2)	(430.2)	(140.6)	(53.4)	114.1	101.3	(771.4)	21.6	533.0	1,000.7
Percent of Revenues	-72.5%	8.3%	-49.2%	-105.4%	-27.8%	-6.9%	9.5%	6.5%	-46.9%	0.5%	8.6%	13.4%
Taxes	0.2	0.0	1.5	8.4	0.0	0.0	0.0	0.0	10.0	0.0	0.0	0.0
Tax Rate	-0.1%	0.0%	-0.9%	-2.0%	0.0%	0.0%	0.0%	0.0%	-1.3%	0.0%	0.0%	0.0%
Total Net Income to Parent	(232.5)	49.3	(159.6)	(438.6)	(140.6)	(53.4)	114.1	101.3	(781.5)	21.6	533.0	1,000.7
Percent of Revenues	-72.6%	8.3%	-49.7%	-107.5%	-27.8%	-6.9%	9.5%	6.5%	-47.5%	0.5%	8.6%	13.4%
EPS for consensus (Rmb)	(0.6)	0.1	(0.4)	(1.1)	(0.4)	(0.1)	0.3	0.3	(2.1)	0.1	1.3	2.5

MetaX: Valuation

Our price target is Rmb758, with an Equal-weight rating. As MetaX is moving from a pre-profit development stage to sustained, scalable profitability, we also value the company using a Residual Income Model. Unlike relative multiples with limited applicability at profitability inflection, we believe this framework best captures intrinsic long-term value, fully incorporating our earnings forecasts through breakeven and the subsequent profit ramp-up amid the booming domestic substitution cycle for high-performance GPGPU chips.

Key assumptions: a 6.5% cost of equity (beta 0.75, risk-free rate 2.0%, equity risk premium 6.5%), a long-term target payout ratio of 50%, a medium-term CAGR of 18% through 2038, and a perpetual terminal growth rate of 6%.

Our price target implies a 75x 2026e P/S multiple, supported by accelerating mass commercialization, full-stack software ecosystem competitiveness, and improving long-term revenue visibility as it scales through the profitability inflection. As the stock trades over 70x 2026e P/S now, we think MetaX is less attractive than Cambricon and Iluvatar.

Exhibit 40: MetaX: Residual income valuation

Rmb million	2026E	2027E	2028E	2029E	2030E	2031E	2032E	2033E	2034E	2035E	2036E	2037E
Total Equity	13,173	13,706	14,706	15,297	15,994	16,816	17,786	18,930	20,281	21,875	23,755	25,975
Net Profit	22	533	1,001	1,181	1,393	1,644	1,940	2,289	2,701	3,188	3,761	4,438
ROAE	0.2%	4.0%	7.0%	7.9%	8.9%	10.0%	11.2%	12.5%	13.8%	15.1%	16.5%	17.9%
Residual Income	(833)	(334)	75	202	368	563	793	1,062	1,378	1,749	2,184	2,696
Spread	-6.3%	-2.5%	0.5%	1.4%	2.4%	3.5%	4.7%	6.0%	7.3%	8.6%	10.0%	11.4%
Ending Equity Capital	13,173											
PV of Forecast Period	3,998											
PV of Continuing Value	285,924											
Equity Value	303,095											
No. of Shares	400											
Price Target	758											

Bull & bear cases

Our bull case value is Rmb1,500, and assumes robust mass production delivery with rapid customer penetration and gross margin expansion.

We assume (1) >100% revenue CAGR in 2025-28 on faster-than-expected mass production and delivery of flagship GPGPUs, plus booming demand for domestic general computing and AI training chips; (2) significant market share expansion in China's domestic GPGPU market, with full adaptation to mainstream AI frameworks and successful entry into core CSPs and sovereign customers; (3) gross margin above 60% in 2026, driven by mass production scale effects, optimized supply chain management, and an upgraded high-value product mix.

Our bear case value is Rmb380, and assumes lower product commercialization with revenue miss and gross margin contraction.

We assume (1) <40% revenue CAGR in 2025-28 given delayed mass production of flagship GPU products, slower-than-expected customer validation, and weaker domestic AI computing capex; (2) share loss in China's general-purpose GPU market amid intensifying domestic and overseas competition, failure to complete full adaptation to mainstream AI frameworks, and inability to secure mass procurement orders from top-tier customers;

and (3) gross margin below 40% in 2026 due to high R&D, rising wafer and packaging costs, and fierce price competition.


Peer comparison

We expect MetaX's revenue to increase at a 66% CAGR from 2025 to 2028. We attribute the superior top-line growth to booming AI computing capex from core clients, including national supercomputing centers and potentially CSPs, as well as accelerated mass production, full-stack software ecosystem adaptation, and customer validation of flagship GPU products amid ongoing domestic penetration. In our view, although MetaX could experience faster revenue expansion given core customer penetration, MetaX's premium P/S multiple at 75x 2026e looks too high for us vs. peers.

Exhibit 41: China AI semi peer comparison table

Ticker	Name	Market Cap (US\$ bn)	Fab capacity implied revenue (2026, US\$ bn)	Consensus revenue (2026, US\$ bn)	2026 P/S	2026 P/E
Global						
NVDA-US	Nvidia	4,851	269	302	18	25
AMD-US	AMD	498	13	16	38	47
China						
688256-CN	Cambricon*	82	3.0	2.1	27	112
688041-CN	Hygon*	89	1.1	n.a.	79	126
688802-CN	MetaX*	39	0.6	n.a.	68	n.a.
688795-CN	Moore Threads*	42	0.3	n.a.	146	n.a.
6082-HK	Biren*	14	0.3	n.a.	48	n.a.
9903-HK	Iluvatar CoreX*	14	0.4	n.a.	31	n.a.

Exhibit 42: MetaX P/S multiple



Baidu & Alibaba

Covered by Gary Yu & Joanne Lau

For detailed analysis, please refer to our deep-dive report [China's AI Path: Owning the Full AI Stack via In-house Chips](#).

T-Head (Alibaba)

T-Head Semiconductor (Pingtougou) was founded in September 2018 as a wholly owned chip business under Alibaba. It originated as a spin-off from Alibaba's DAMO Academy, with a strategic mandate to build custom ASIC products and support China's semiconductor self-sufficiency. Alibaba is reportedly preparing T-Head for a [potential IPO](#), and it plans to restructure T-Head into a standalone business with partial employee ownership ahead of commencing any listing process.

T-Head produces both GPUs and CPUs, with CPU products remaining a core part of the portfolio.

Exhibit 43: T-head key products

Product	海龙 810E (Zhenwu 810E PPU)	含光 800 (Hanguang 800)	倚天 710 (Yitian 710)	倚岳 510 (Zhenyue 510)	玄铁 C910 (Xuantie C910)	玄铁 C920 (Xuantie C920)	羽阵 611 (Yuzhen 611 / TH6101)
Category / Use	Data-center AI accelerator (training + inference)	Data-center AI inference chip	Cloud server CPU	Enterprise SSD controller	RISC-V CPU IP core (high-perf)	RISC-V CPU IP core w/ vector	UHF RFID tag chip
Compute	N/A	Peak compute disclosed as 820 TOPS	N/A	N/A	Supports FP16 / BF16 / FP32 / FP64 + INT8/16/32/64 N/A (IP core)	Vector unit supports FP16/FP32/FP64 + INT8/16/32/64 N/A (IP core)	N/A
Memory / Storage	96GB HBM2e	Not publicly disclosed on the product page	Uses system DRAM; supports 8-channel DDR5	Supports TLC/QLC NAND (1x1L2x4L), NVMe features	N/A (IP core)	N/A (IP core)	Memory disclosed: 128-bit EPC, 96-bit TID (incl. 48-bit SN), 32-bit access/kill shared
I/O / Interconnect	Inter-chip bandwidth 700GB/s	Not publicly disclosed	96-lane PCIe 5.0	PCIe Gen5 x4, NVMe 1.4b, ZNS	N/A (IP core)	N/A (IP core)	RFID air interface; protocol compliant with EPC Global G2 V2 / ISO/IEC 18000-6C
Other key disclosed spec	Reported to be positioned as comparable to NVIDIA H20	12nm, 17B transistors	Arm v9 compatible, 60B transistors, SPECint2017 score disclosed	Supports enterprise features like multi-stream / atomic write / IO virtualization, E2E data protection; hot-plug	Vector units designed per RVV 0.7.1	RV64GCV; high-perf core with vector computing ability	Sensitivity: read -24 dBm / write -20 dBm; operating temp -40°C to +85°C

Valuation

In its [latest results call](#), management indicated >60% of T-head revenue is from external cloud customers, and has delivered >470k units so far. Management did not rule out a potential spin-off opportunity, which could further unlock SOTP value, but no timeline was noted. On both the demand and supply side, support remains strong: AliCloud continues to drive enormous training and inference needs, and increasing production capacity supply.

We value T-head at a valuation range of US\$28-83bn, based on 12-24x F28e P/S, with reference to key peers Cambricon and Iluvatar, applied an estimated revenue range of Rmb16-24bn (CPU + GPU). This translates to US\$8-24/share in BABA's SOTP assuming a 30% holdco discount.

We revised midrange SOTP to US\$250 (from US\$245), T-head at US\$15/share (from US\$22), implying T-head market cap of US\$52bn, based on 18x F28e EV/sales:

- This is based on Rmb20bn F28e revenue (the midpoint of our estimated F28e revenue range of Rmb\$16-24bn).
- Multiple is at a 10% discount to key A-share peer Cambricon's 20x 2027e EV/sales given A/H share discount and similar to smaller H-share peer Iluvatar's 18x 2027e.

We also applied a 30% holdco discount which is consistent with broad Internet peers.

Exhibit 44: T-head market cap (US\$ mn)

		T-head F2028 revenue (Rmbm)					
		16,000	18,000	20,000	22,000	24,000	
P/S (x)	24x	55,252	62,158	69,065	75,971	82,878	
	21x	48,345	54,388	60,432	66,475	72,518	
	18x	41,439	46,619	51,799	56,978	62,158	
	15x	34,532	38,849	43,165	47,482	51,799	
	12x	27,626	31,079	34,532	37,986	41,439	


Exhibit 45: T-head market cap to BABA (US\$ mn)

		T-head F2028 revenue (Rmbm)					
		16,000	18,000	20,000	22,000	24,000	
P/S (x)	24x	38,676	43,511	48,345	53,180	58,014	
	21x	33,842	38,072	42,302	46,532	50,763	
	18x	29,007	32,633	36,259	39,885	43,511	
	15x	24,173	27,194	30,216	33,237	36,259	
	12x	19,338	21,755	24,173	26,590	29,007	

Exhibit 46: T-head market cap to BABA/share (US\$)

		T-head F2028 revenue (Rmbm)					
		16,000	18,000	20,000	22,000	24,000	
P/S (x)	24x	16.3	18.3	20.4	22.4	24.4	
	21x	14.2	16.0	17.8	19.6	21.4	
	18x	12.2	13.7	15.3	16.8	18.3	
	15x	10.2	11.4	12.7	14.0	15.3	
	12x	8.1	9.2	10.2	11.2	12.2	

Exhibit 47: BABA mid point SOTP



Kunlunxin (Baidu)

Kunlunxin is Baidu's semiconductor subsidiary. Baidu began chip development in 2011, and in April 2021 spun the project out as a separate entity led by its chief chip architect, Ouyang Jian. Kunlunxin is CUDA compatible, enabling developers to migrate from NVIDIA with minimal changes. Baidu has a strategic partnership with Samsung, which manufactured the first Kunlun chip in 2018 using 14nm. We expect both Samsung and SMIC to remain key suppliers as domestic demand grows.

Kunlunxin focuses on producing ASICs and it is CUDA-compatible.

Exhibit 48: Kunlunxin key products

Product	Generation	Process Node	Key Specifications
Kunlun AI Chip	1st Gen (2018)	14nm	30x faster than original FPGA-based processor
Kunlun II	2nd Gen (2021)	NA	Comparable to Nvidia AI00
P800	3rd Gen (2025)	NA	Used in 30,000-chip cluster for training large models

Valuation

Baidu proposed the spin-off of KLX in December 2025 while we expect on track IPO progress with an expected listing timeline by late 2Q to early 3Q. The proposed listing is part of management's plan to unlock shareholder value.

We value Kunlunxin at a range of US\$20-62bn, assuming 2027 revenue range of Rmb12-18bn, using P/S multiples range of 12-24x with reference to key peers Cambricon and Iluvatar.

We revised midrange SOTP to US\$225 (from US\$215), KLX at US\$47/share (from US\$45), implying market cap of US\$38.8bn, based on 18x 2027e EV/sales:

- This is based on Rmb15bn 2027e revenue (the midpoint of our estimated 2027 revenue range of Rmb\$12-18bn).
- Multiple is at a 10% discount to key A-share peer Cambricon's 20x 2027e EV/sales given A/H share discount and similar to smaller H-share peer Iluvatar's 18x 2027e. We also applied a 30% holdco discount which is consistent with broad Internet peers.

Exhibit 49: Kunlunxin market cap (US\$bn)

		Kunlunxin 2027 revenue (Rmbm)				
		12,000	13,500	15,000	16,500	18,000
P/S (x)	24x	41,439	46,619	51,799	56,978	62,158
	21x	36,259	40,791	45,324	49,856	54,388
	18x	31,079	34,964	38,849	42,734	46,619
	15x	25,899	29,137	32,374	35,612	38,849
	12x	20,719	23,309	25,899	28,489	31,079


Exhibit 50: Kunlunxin market cap to Baidu (US\$bn)

		Kunlunxin 2027 revenue (Rmbm)				
		12,000	13,500	15,000	16,500	18,000
P/S (x)	24x	17,245	19,401	21,556	23,712	25,867
	21x	15,089	16,975	18,862	20,748	22,634
	18x	12,934	14,550	16,167	17,784	19,401
	15x	10,778	12,125	13,473	14,820	16,167
	12x	8,622	9,700	10,778	11,856	12,934

Exhibit 51: Kunlunxin market cap to Baidu/share (US)

		Kunlunxin 2027 revenue (Rmbm)				
		12,000	13,500	15,000	16,500	18,000
P/S (x)	24x	50.3	56.6	62.9	69.2	75.5
	21x	44.1	49.6	55.1	60.6	66.1
	18x	37.8	42.5	47.2	51.9	56.6
	15x	31.5	35.4	39.3	43.3	47.2
	12x	25.2	28.3	31.5	34.6	37.8


Exhibit 52: Baidu mid point SOTP



Cambricon: Company background

Cambricon Technologies is a leading artificial intelligence (AI) chip designer in China, specializing in intelligent processors for AI computing. Founded in 2016 and listed on the STAR Market in 2020, Cambricon originated from research at the Chinese Academy of Sciences, where its founding team developed early neural network processors. In 2016, the company launched its first commercial AI chip, Cambricon-1A, for terminal devices. In 2018, it introduced the cloud MLU100 and edge MLU200 series, expanding into data centers. In 2020, Cambricon went public, marking a major commercialization milestone. In subsequent years, it iterated its AI chips and software ecosystem, strengthening its position in China's AI semiconductor industry. Cambricon is now advancing next-generation AI processors, including the MLU580/590/690 series, aimed at higher performance and efficiency for large-scale AI computing.

Exhibit 53: Cambricon's development timeline



Key products

Cambricon has built a comprehensive AI portfolio spanning cloud, edge, and terminal applications, supported by proprietary hardware and software. Offerings broadly include cloud AI chips, edge AI chips, and intelligent processing modules.

Hardware platform:

Cloud AI chips: Cambricon's cloud products center on the MLU series, including MLU100/200/300/500/600. These chips target large-scale data center workloads such as AI LLM training and inference. According to our industry check, the latest MLU590/690 significantly improve computing power, memory bandwidth (HBM3/3e), and inter-chip connectivity, making them suitable for high-performance AI clusters.

Edge AI chips: Cambricon offers edge processors such as MLU220 for intelligent surveillance, industrial automation, and smart transportation. These chips emphasize low power and real-time processing, enabling efficient deployment in resource-constrained edge scenarios.

Software platform:

NeuWare: Cambricon has developed a full-stack proprietary software ecosystem centered on its NeuWare platform, which serves as both a key competitive moat and a potential limitation due to ecosystem maturity. The platform includes the BANG programming language, a self-developed parallel computing language compatible with C99/C++11

syntax, which claims to achieve up to 85% of peak performance with one-tenth of development time.

MagicMind, an inference acceleration engine based on MLIR, supports model conversion from PyTorch, TensorFlow, and ONNX. In addition, CNToolkit functions as a comprehensive development toolkit for the MLU architecture, similar to NVIDIA’s CUDA Toolkit, while CNCL provides a multi-card and multi-node communication library to support distributed training across AI clusters.

Exhibit 54: Cambricon's products


	Product	Type	Manufacture	Node	Computing power (TFLOPS/TOPS)			Memory type	Memory size (GB)	Memory bandwidth (GB/s)	Chip-to-chip networking bandwidth (GB/s)	TDP (W)
					FP16 (TFLOPS)	FP8 (TFLOPS)	INT8 (TOPS)					
AI Training	MLU 370	DSA	TSMC	7nm	96		X	X	LPDDR5	48	614	200
	MLU 580*	DSA	SMIC	12nm	280		560	X	HBM3	30	1,600	425
	MLU 590*	DSA	TSMC	7nm	315		630	X	HBM3	30	1,600	425
	MLU 690*	DSA	SMIC	7nm	702		1,404	X	HBM3	30	2,600	750
Edge AI	MLU 220	DSA	TSMC	16nm		X	X	16	LPDDR4	8	400	15

Exhibit 55: Built on a 7nm process, MLU370 is Cambricon’s first AI chip to adopt chiplet, integrating 39bn transistors



Source: Cambricon

Exhibit 56: MLU220 is an AI accelerator specifically designed for edge computing scenarios.



Source: Cambricon

Exhibit 57:

Cambricon's NeuWare platform







Exhibit 58: Cambricon's MagicMind platform



Cambricon has transformed its revenue structure and business model in recent years, shifting toward higher-quality, scalable growth. In 2023, intelligent computing cluster systems contributed 86% of revenue, positioning the company as a system integrator. By 2025, cloud AI chips and accelerator cards contributed over 99%, underscoring Cambricon's evolution into a pure-play AI chip company with improved revenue quality and stronger technological differentiation.

The company also delivered rapid revenue growth, reaching Rmb6.4bn in 2025 with a three-year CAGR of 107%, primarily driven by large-scale commercialization of MLU590 and ramping orders from major customers such as ByteDance.

Cambricon's gross margin improved alongside its pivot to high-value AI chips. Historically, margins were constrained by system integration, which carries lower profitability due to higher hardware and assembly costs. As the mix shifted to cloud AI chips and accelerator cards – over 99% of revenue by 2025 – gross margin expanded meaningfully. Large-scale deployment of higher-end products such as MLU590, combined with economies of scale and a more standardized product offering, further supported margin uplift. Gross margin increased to 45% in 2025 from 31% in 2023.

Exhibit 59: Cambricon's revenue shifted to cloud AI products**Exhibit 60:** Cambricon's gross margin declined in 2025 due to product mix change

Key competitors

Cambricon currently holds a small share of the global AI chip market, which remains dominated by international players like Nvidia. However, we expect the company to gradually gain share in China's domestic AI computing market, supported by localization trends and rising demand from cloud service providers and state-owned enterprises amid geopolitical tensions and export restrictions.

In AI chips, Cambricon primarily competes with NVIDIA in China, which dominates globally with its GPU ecosystem. Among Chinese peers, key competitors include Huawei (Ascend), MetaX, Illuvatar, and Kunlun Xin, which are expanding in AI training and inference chips. We note that while NVIDIA continues to lead in performance and ecosystem maturity, Cambricon is narrowing the gap domestically through product iteration and policy-driven demand, with revenue growth accelerating alongside increased adoption by major Chinese customers.

Exhibit 61: Among "10 China AI Semi Dragons", we view Cambricon as a technology-driven AI chip leader with strong architectural capabilities

Company	Ticker	GPU Product	Node	GGPU/ASIC	Foundry sourcing viability	Sovereign background	Affiliated design house	AI inference performance
 HUAWEI	Private	950, 910C, 910B	7nm	ASIC	✓	✓	X	➡
 寒武纪	688256-SS	MLU series	7nm	ASIC	✓	X	X	➡
 HYGON	688041-SS	DCU	7nm	GGPU	✓	✓	X	➡
 沐曦	688802-SH	MXC series; MXG series	7nm/12nm	GGPU	✓	X	X	➡
 摩尔线程	688795-SH	MUSA MTT 55000, 54000	7nm/12nm	GGPU	▲	X	X	➡
 壁仞科技	6082-HK	BR100	7nm/12nm	GGPU	▲	X	X	➡
 天数智芯	9903-HK	TianGai-100 series, Zhikai-100 Series	7nm	GGPU	✓	X	X	➡
 T-HEAD	Currently under H-share listing counseling process	Hangguang 800, ppu	12nm	GGPU	✓	X	✓	➡
 昆仑芯	Currently under H-share listing counseling process	R and P series	6nm	ASIC	▲	X	✓	➡
 燧原科技	Currently under A-share listing counseling process	560, 120, L600	12nm	ASIC	▲	X	✓	➡

Key suppliers and customers

Cambricon's supply chain and customer base reflect its positioning within China's domestic AI semiconductor ecosystem amid ongoing geopolitical and technology restrictions.

On the supplier side, Cambricon previously relied on TSMC for advanced-node manufacturing. However, following tightening export controls and limited access to leading-edge foundry capacity, we believe the company has transitioned production to domestic foundries, primarily SMIC. Based on our industry checks, Cambricon's MLU580 is expected to be a fully localized AI chip, marking a key milestone in supply chain independence. That said, domestic manufacturing still faces yield and process-maturity challenges, which could impact cost structure and product consistency in the near term.

In addition to wafer fabrication, high-bandwidth memory (HBM) is critical for AI accelerators, and Samsung remains a key supplier. Reliance on advanced HBM underscores memory bandwidth's importance in AI workloads, while also introducing potential supply constraints and cost considerations given the concentrated global HBM market.


On the customer side, Cambricon primarily serves cloud service providers (CSPs) and large enterprises. ByteDance stands out as the most strategic partner, with deep collaboration in procurement, joint R&D, and software-hardware co-optimization. ByteDance has actively adapted Cambricon's chips within its AI infrastructure. Other key customers include major Chinese CSPs and telecom operators such as China Mobile, as well as select government and enterprise clients, reflecting growing adoption of domestic AI chips across multiple scenarios.

Exhibit 62: Cambricon's suppliers and customers

Shareholding structure

Cambricon maintains a founder-led structure with control concentrated in core management. Chen Tianshi, founder, chairman, and CEO, directly holds ~28.35% and remains the ultimate controlling shareholder. In addition, he controls Aixi Technology and Aijaxi Technology, both employee shareholding platforms with 7.27% and 1.15% stakes, bringing his total voting rights to ~36.8% on a combined basis.

The Chinese Academy of Sciences–affiliated Institute of Computing Technology (through CAS-related entities such as Zhongke Suanyuan) holds ~15.6%, reflecting the company's academic origins. Other notable shareholders include Zhang Jianping (~1.62%), a well-known individual investor. Overall, the structure highlights a balance between founder control, institutional backing, and strategic investors.

Exhibit 63: Cambricon: Shareholder structure

Management team

Cambricon's management team combines strong academic credentials and deep technical expertise in AI and processor architecture. Chen Tianshi, founder, chairman, and CEO, entered the University of Science and Technology of China's Special Class for the Gifted Young at 16 and later earned a PhD from the Chinese Academy of Sciences; he is widely recognized as a leading figure in China's AI chip industry. Chen Yunji, Chief Scientist and Chen Tianshi's elder brother, is a researcher at the Institute of Computing Technology under the Chinese Academy of Sciences and a co-chief architect of the Loongson processor, bringing extensive CPU and AI architecture experience.

Liu Shaoli, Deputy General Manager and co-founder, is responsible for terminal IP and cloud chip R&D, playing a key role in product development.

Exhibit 64: Cambricon: Management team


Name	Age	Position	Prior Experience
Tianshi Chen	41	Chairman & CEO	Former Researcher and PhD Supervisor at Institute of Computing Technology (CAS); founder of Cambricon in 2016
Zai Wang	42	Employee Director & Vice President	Former engineer at Zhengzhou Commodity Exchange; IT role at Zhongyuan Bank; later researcher at ICT (CAS)
Shaoli Liu	39	Director & Vice President	Former Associate Researcher at ICT (CAS); founding team member of Cambricon
Haoyin Ye	38	Director, VP, CFO & Board Secretary	Former investment manager at China High-Tech Investment Group and SDIC Venture Capital
Shuai Chen	40	Vice President	Former engineer at ICT (CAS); Postdoctoral researcher at University of Toronto
Yi Liu	41	Vice President	Former engineer at Loongson and Senior Engineer at NVIDIA Shanghai
Yao Zhang	40	Vice President	Former roles at ICT (CAS), Loongson, and Xiaomi Pinecone

Iluvatar: Company Background

Iluvatar CoreX Semiconductor Co., Ltd. is a leading Chinese general-purpose GPU (GPGPU) design company headquartered in Shanghai. Founded in 2018 by a team with an AMD background, the company has established itself as one of the most advanced domestic AI chip developers, with end-to-end capability spanning architecture design, chip development, and software stack optimization.

The company listed on the Hong Kong Stock Exchange on January 8, 2026, raising approximately HK\$3.7bn (US\$476mn) in its IPO.

Exhibit 65: Iluvatar's development timeline



Product portfolio & technology roadmap

Iluvatar's product strategy focuses on high-performance, cost-effective AI computing solutions that address mainly AI inference workloads.

Iluvatar has outlined an aggressive multi-year roadmap aimed at closing and eventually surpassing NVIDIA's performance lead:

- 1H26: Launch TianXuan architecture products, benchmarked against NVIDIA's Blackwell platform
- 2H26: Launch TianJi architecture products, designed to outperform NVIDIA's Blackwell platform
- 2027: Launch TianQuan architecture products, projected to surpass NVIDIA's next-generation Rubin platform
- Post-2027: Focus on developing "breakthrough computing architectures" that will redefine AI computing capabilities

Exhibit 66: Iluvatar: Product roadmap



Historical financial performance

Iluvatar CoreX has delivered rapid top-line growth in recent years, reflecting accelerating commercialization of its GPU products, albeit still in a loss-making phase typical of early-stage AI chip companies. Revenue increased from Rmb189mn in 2022 to Rmb1.03bn in 2025, a 76% CAGR, driven primarily by rising AI GPU shipments.

Despite strong revenue momentum, the company remains loss-making due to heavy R&D investments and early-stage scaling dynamics. Net losses expanded from Rmb554mn in 2022 to Rmb1.0bn in 2025, broadly in line with revenue scale.

That said, early signs point to improving financial quality. Gross profit grew faster than revenue, reaching Rmb558mn in 2025. Gross margin improved to 54% (up 4.9ppts Y/Y), while adjusted loss margin has begun to narrow, suggesting operating leverage is gradually emerging as volumes scale.

Overall, Iluvatar exhibits a typical “high-growth, high-investment” profile: rapid revenue expansion from product commercialization, coupled with sustained losses from intensive R&D and ecosystem build-out. The key inflection point to watch is whether continued revenue scaling can translate into margin improvement and an eventual path to profitability.

Exhibit 67: Iluvatar's historical revenue





Exhibit 68: Iluvatar's gross profit and gross margin



Key suppliers and customers

As a fabless GPU company, Iluvatar relies on external foundries and advanced semiconductor supply chains for manufacturing and packaging. The company is adopting TSMC's 7nm node and producing fully compliant products. Unlike other Chinese AI semi players, access to advanced nodes and high-end components (e.g., HBM) could be easier for Iluvatar, potentially affecting yield, cost structure, and product performance.

On the customer side, Iluvatar has built a relatively diversified base vs. early-stage peers. As of 2025, it had served over 300 customers across internet, finance, healthcare, education, and transportation.

Key competitors

Iluvatar operates in the highly competitive global and domestic AI accelerator market, where NVIDIA remains the dominant player globally.

Within China, competition is more intense, with key domestic players including Huawei (Ascend series), Cambricon, Biren Technology, Moore Threads, and MetaX. These companies all target the fast-growing AI computing market under localization trends but differ in technical routes and ecosystem strategies. For example, Huawei started pursuing hybrid GPU and ASIC design and Cambricon remains ASIC, while MetaX and Iluvatar adopt GPU-based architectures with varying degrees of CUDA compatibility.

Exhibit 69: We view Iluvatar as one of the best plays among Chinese AI GPGPU vendors


Company	Ticker	GPU Product	Node	GPGPU/ASIC	Foundry sourcing viability	Sovereign background	Affiliated design house	AI inference performance
 HUAWEI	Private	950, 910C, 910B	7nm	ASIC	✓	✓	✗	➡
 寒武纪	688256-SS	MLU series	7nm	ASIC	✓	✗	✗	➡
 HYGON	688041-SS	DCU	7nm	GPGPU	✓	✓	✗	➡
 META X 沐曦	688802-SH	MXC series; MXG series	7nm/12nm	GPGPU	✓	✗	✗	➡
 摩尔线程 MOORE THREADS	688795-SH	MUSA MTT S5000, S4000	7nm/12nm	GPGPU	▲	✗	✗	➡
 壁仞科技 BIREN TECHNOLOGY	6082-HK	BR100	7nm/12nm	GPGPU	▲	✗	✗	➡
 天数智芯 Iluvatar CoreX	9903-HK	TianGai-100 series, Zhikai-100 Series	7nm	GPGPU	✓	✗	✗	➡
 T-HEAD	Currently under H-share listing counseling process	Hanguang 800, PPU	12nm	GPGPU	✓	✗	✓	➡
 昆仑芯 KUNLUNXIN	Currently under H-share listing counseling process	R and P series	6nm	ASIC	▲	✗	✓	➡
 Enflame 燧原科技	Currently under A-share listing counseling process	S60, I20, L600	12nm	ASIC	▲	✗	✓	➡

Shareholding structure

Post its January 2026 Hong Kong IPO, Iluvatar maintains a diversified, institutionally

dominated shareholding structure with 254mn shares outstanding. Major pre-IPO shareholders include Dazheng Capital (18.28%), with founder and employee platforms collectively controlling 23.6%. The company has implemented a comprehensive share-based incentive plan covering over 300 key technical and management personnel, aligning their interests with long-term shareholder value creation. No single shareholder holds a controlling stake, ensuring balanced governance and strategic decision-making.

Exhibit 70: Iluvatar: Shareholding structure



Management team

Iluvatar CoreX's management combines global GPU industry experience with domestic telecom and commercialization expertise. Chairman and CEO Lujiang Gai has a finance and investment background, signaling a stronger focus on capital operations and business execution rather than pure R&D. On the technical side, executives such as Yile Sun (ex-AMD) bring deep experience in GPU architecture and high-performance computing, supported by a broader team covering chip implementation, software, and supply chain.

Compared with Cambricon's research-driven team and MetaX's AMD-centric engineering team, Iluvatar's management is more balanced between technology and commercialization, with stronger go-to-market and ecosystem integration capabilities. However, its relatively less founder-led technical identity may imply weaker differentiation in core architecture, making execution a key determinant of long-term competitiveness.

Exhibit 71: Iluvatar: Management team


Name	Age	Position	Prior Experience
Lujiang Gai	44	Chairman of the Board & CEO	Former accountant/senior accountant at Baker Tilly and PwC Zhongtian; senior manager at Deloitte Beijing; assistant to chairman at Beijing Titan Source Natural Gas; partner and head of risk control at Beijing Zhongjin Taian Asset Management.
Yile Sun	45	Executive Director & Vice President	Former engineer at MacroSynergy Technology and ATI Visual Technology; senior manager at AMD Shanghai from 2006 to 2017; currently responsible for chip R&D.
Zheng Liu	51	Executive Director & COO	Former roles at Shanghai telecom operations, China Unicom/China Netcom system and management functions; former VP of China-ASEAN Information Harbor; former GM of China Unicom Innovation and Entrepreneurship Investment (Shanghai).
Lei Yang	41	Executive Director, CFO, Board Secretary	Former assistant manager at Deloitte; internal audit manager at Home Depot Investment Management (Shanghai); senior audit executive at China Fortune Land Development; executive director in investment banking at Tianfeng Securities.
Yuan Liu	44	Vice President	Former PMTS ASIC/layout design engineer at AMD Shanghai; currently responsible for chip mass production and implementation.
Jiasheng Shi	41	Vice President	Former software engineer at Huiguo Software; senior software engineer at AMD Shanghai; later worked at Jaunt Shanghai; currently responsible for software R&D.
Bin Liang	58	Vice President	Former principal solutions architect at Oracle; former director at Inspyrus; former VP at Biren Technology; now responsible for co-technology development.

MetaX: Company Background

MetaX Integrated Circuits (Shanghai) Co., Ltd. is a leading domestic GPU design company in China, specializing in high-performance general-purpose GPUs (GPGPU) for AI training, inference, and graphics rendering. Founded in September 2020 and listed on the STAR Market in December 2025, MetaX is one of the few Chinese companies with fully proprietary GPU IP, instruction sets, and architecture.

A core team with deep AMD experience founded the company, bringing end-to-end expertise in GPU architecture and mass production. Since inception, MetaX rapidly scaled R&D; by end-2025, 73% of ~925 employees worked in R&D. Within five years, it moved from startup to public listing, demonstrating strong execution.

Exhibit 72: MetaX's development history



Product portfolio

MetaX has a comprehensive GPU portfolio covering AI training, inference, and graphics rendering across three core lines: Xiyun (C), Xisi (N), and Xicai (G). Its general-purpose GPU architecture enables broad applicability across computing scenarios.

AI training / general-purpose GPUs (C-series): The Xiyun C500 is the core revenue driver and has entered large-scale production and customer deployment. Built on a 7nm process, the C500 delivers ~240TFLOPS FP16 and supports a CUDA-compatible ecosystem. The next-generation C600 represents a major upgrade, adopting SMIC's 12nm process and a chiplet architecture. The C600 integrates HBM3e with bandwidth >3.35TB/s and is positioned as a domestically manufactured high-end GPU comparable to NVIDIA H100.

AI inference and edge GPUs (N-series): The N-series targets inference workloads and edge deployments, focusing on power efficiency and cost-performance optimization for cloud inference, edge computing, and enterprise AI.

Graphics GPUs (G-series): The G-series focuses on graphics rendering and visualization for digital content creation, industrial simulation, and virtual reality, differentiating MetaX from domestic peers focused primarily on AI acceleration.

Exhibit 73: MetaX's C600 will adopt SMIC's N+1 process



Source: MetaX

Exhibit 74: MetaX's N-series targets inference workloads and edge deployment




Source: MetaX

Software to help MetaX to stand out

Software platform (MXMACA): MetaX developed its proprietary MXMACA (MetaX Advanced Computing Architecture), a unified heterogeneous computing platform highly compatible with the CUDA ecosystem. As of mid-2025, MXMACA supports over 6,000 mainstream open-source applications and more than 2,200 high-performance operators, with native adaptation to PyTorch, TensorFlow, and large models like DeepSeek and Qwen. The CUDA compatibility strategy is a key differentiator vs. competitors such as Huawei Ascend, which relies on its proprietary CANN ecosystem. By lowering migration costs for existing NVIDIA users, MXMACA enables faster customer adoption and ecosystem integration. However, the depth and stability of CUDA compatibility still require further optimization, particularly in complex computing scenarios.

Exhibit 75: MetaX MXMACA



Revenue growth and business model transition

MetaX has delivered explosive revenue growth, driven by commercialization of the Xiyun C500. Revenue rose from Rmb53mn in 2023 to Rmb408mn in 2025, indicating strong product ramp and customer demand.

Revenue remains highly concentrated in C500-based GPU accelerator cards, a typical early-stage profile with reliance on a flagship product. As the C600 moves toward commercialization, MetaX is expected to shift to a more diversified, higher-performance mix, critical for sustaining growth and improving revenue quality.

Gross margin

MetaX's gross margin is in the early stage of improvement, reflecting scaling production and heavy R&D investment. As a fabless GPU company, gross margin benefits structurally from high-value chip design; however, near-term profitability is constrained by low initial volumes, high tape-out costs, and supply chain inefficiencies.

Looking ahead, gross margin is expected to improve as shipment volumes scale, yield rates stabilize (particularly for domestically manufactured chips such as C600), and mix shifts toward higher-end GPUs. That said, margin expansion will depend on stable mass production and cost control across the evolving supply chain.

Exhibit 76: MetaX's revenue surged given rising GPU shipment





Exhibit 77: MetaX's gross margin remained over 50%



Key suppliers and customers

On suppliers, MetaX initially relied on advanced foundries such as TSMC for 7nm C500 manufacturing according to our industry check. As it advances localization, the next-generation C600 adopts SMIC's 12nm (N+1) process, a key step toward domestic production. The chip also integrates HBM3e memory, with major suppliers including Samsung, highlighting continued reliance on global leaders for high-bandwidth memory. While localization enhances supply security, challenges remain in yield rates and process maturity, particularly for advanced packaging and chiplet integration.

On customers, MetaX primarily serves enterprise clients and cloud-related infrastructure providers. Its base includes companies such as Super Telecom, with orders largely concentrated in C500-based accelerator cards. As the company scales, expansion into major cloud service providers (CSPs) could be a key growth driver, although customer concentration remains relatively high.

Key competitors

MetaX currently holds a very small share of the global AI accelerator market, which remains dominated by NVIDIA. However, we expect the company to gradually gain share in China's domestic GPU market, supported by localization demand and supply chain restructuring.

Domestically, key competitors include Huawei (Ascend series), Cambricon, Moore Threads, etc. Compared to peers, MetaX differentiates through its GPGPU focus, CUDA-compatible

software ecosystem, and progress toward supply chain localization (e.g., C600).

We note that Huawei maintains a dominant position in China's AI accelerator market due to its ecosystem and customer base. MetaX's strengths lie in its AMD-origin team and technology stack, while challenges include limited scale, a narrow portfolio, and an early-stage customer base.

Exhibit 78: We view MetaX as one of the best plays among Chinese AI GPGPU vendors


Company	Ticker	GPU Product	Node	GPGPU/ASIC	Foundry sourcing viability	Sovereign background	Affiliated design house	AI inference performance
 HUAWEI	Private	950, 910C, 910B	7nm	ASIC	✓	✓	X	➡
 寒武纪	688256-SS	MLU series	7nm	ASIC	✓	X	X	➡
HYGON	688041-SS	DCU	7nm	GPGPU	✓	✓	X	➡
 META-X 沐曦	688802-SH	MXC series; MXG series	7nm/12nm	GPGPU	✓	X	X	➡
 摩尔线程 MOORE THREADS	688795-SH	MUSA MTT S5000, S4000	7nm/12nm	GPGPU	▲	X	X	➡
 壁仞科技 BIREN TECHNOLOGY	6082-HK	BR100	7nm/12nm	GPGPU	▲	X	X	➡
 天数智芯 iluvatar CoreX	9903-HK	TianGai-100 series, Zhikai-100 Series	7nm	GPGPU	✓	X	X	➡
 T-HEAD	Currently under H-share listing counseling process	Hangguang 800, PPU	12nm	GPGPU	✓	X	✓	➡
 昆仑芯 KUNLUNXIN	Currently under H-share listing counseling process	R and P series	6nm	ASIC	▲	X	✓	➡
 Enflame 燧原科技	Currently under A-share listing counseling process	S60, I20, L600	12nm	ASIC	▲	X	✓	➡

Shareholding structure

MetaX has a founder-led ownership structure with diversified institutional backing. As of Dec 2025, CEO Weiliang Chen directly and indirectly holds ~22.94%, serving as the controlling shareholder. The company has completed eight financing rounds since inception, attracting top-tier investors.

Notably, Weidong Ge, founder of Chaos Investment, holds ~6.73% and is the largest external shareholder. Other institutional investors include Sequoia China, etc. The shareholder base reflects strong support from industrial and financial investors, providing capital backing for long-term R&D.

Exhibit 79: MetaX: Shareholder structure



Management team

MetaX's management team is a key competitive advantage. Compared with many domestic GPU startups, leadership stands out for technical depth and an unusually strong "AMD pedigree," with multiple senior executives having worked on global GPU products from architecture through tape-out and volume production. This gives the company stronger credibility in high-performance GPU design, software stack construction, and commercialization.

Founder and Chairman/General Manager Chen Weiliang is the company's core figure. Before founding MetaX in 2020, he spent over a decade at AMD Shanghai and earlier worked at Trident, Techland, and other graphics-chip companies. His career is highly concentrated in GPU design and product engineering, supporting the market view that MetaX is among the more "pure-play GPU DNA" teams domestically.

The technical leadership around Chen is also unusually strong. Peng Li, co-founder and director/deputy general manager/CTO, previously worked at AMD Shanghai and is often highlighted as a key figure in architecture and productization.

Yang Jian, another co-founder and senior technical executive, brings a broader resume spanning Trident, S3-related background, VIA, AMD, Huawei Shanghai, and HiSilicon.

Exhibit 80: MetaX: Management Team

Name	Age	Position	Prior Experience
Weiliang Chen	49	Chairman & CEO	Former Senior Director at AMD Shanghai; prior roles at Trident, Techland, and other GPU-related companies
Li Peng	48	Director, VP & CTO	Former engineer at Shanghai AP Computer Systems; long-term technical leader at AMD Shanghai
Jian Yang	52	Director, VP & CTO	Former roles at Trident, S3/VIA, AMD, Huawei Shanghai, and HiSilicon; extensive GPU and SoC architecture experience
Zhongwei Wei	45	CFO & Board Secretary	Former roles at BOC International, Haitong Securities, CITIC Securities, Shenwan Hongyuan, and China Merchants Securities
Shuang Wang	40	Director & Senior R&D Executive	Former engineer at ZTE Microelectronics and R&D manager at HiSilicon