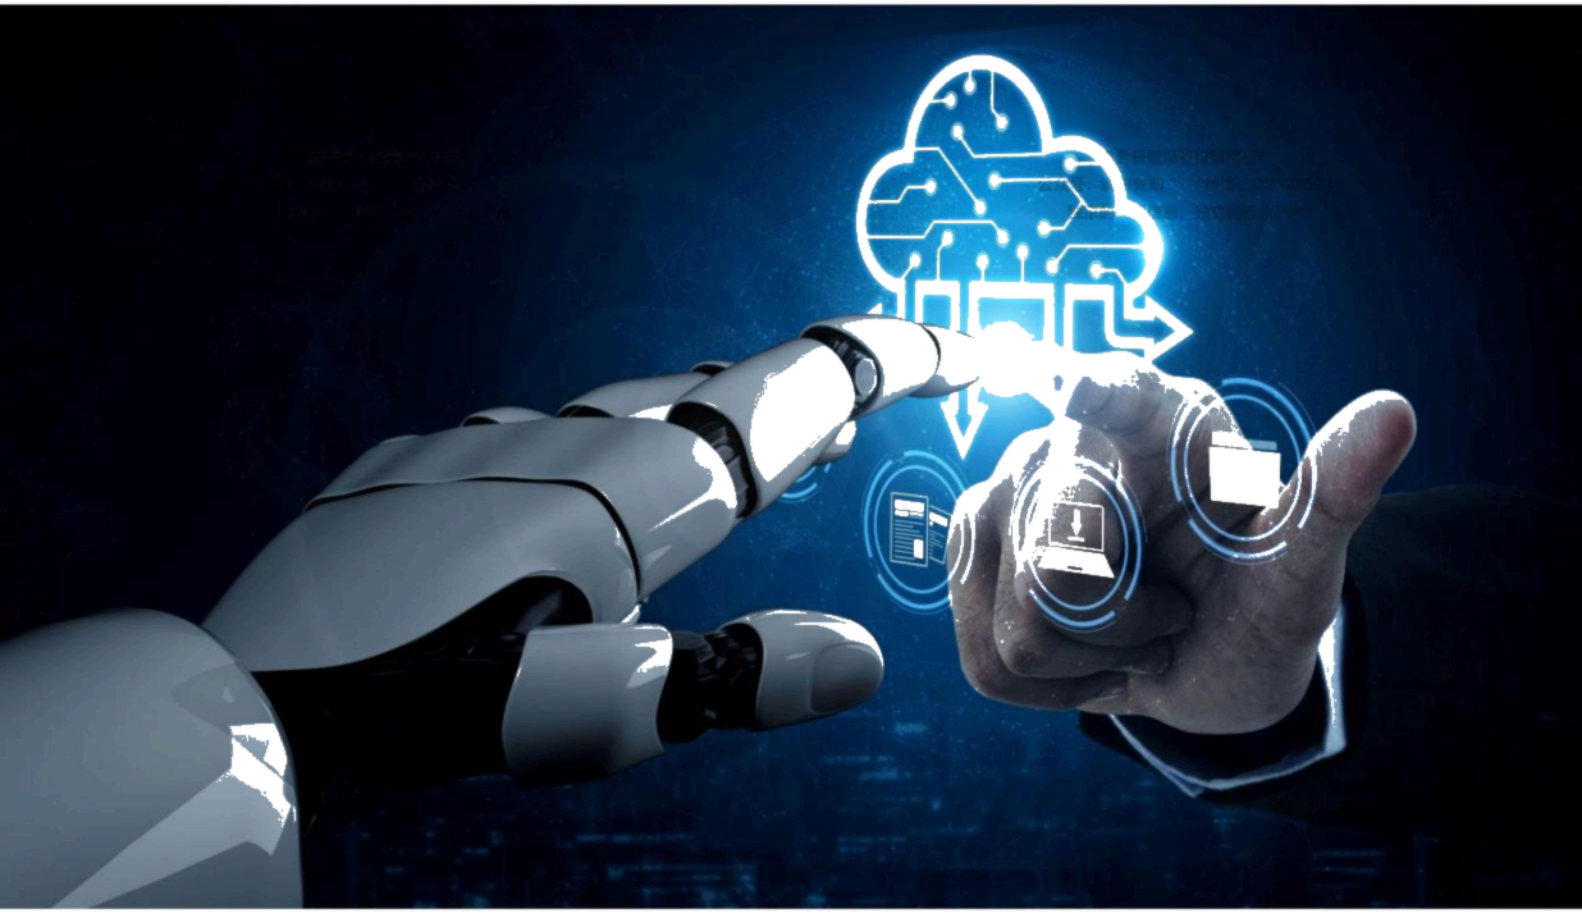
NOMURA

Connecting Markets East & West

ANCHOR REPORT

Global Markets Research

11 September 2026



China AI value chain – sustainable AI investment drives better ROI

Accumulate AI leaders Alibaba, Kingdee, VNET

We believe the global generative AI market is at a turning point, with wider adoption in our society and the economy, and all major AI players are striving to build more sustainable monetization models to drive better ROI. We believe global LLM players will continue to train closed-source LLMs to target AGI (Artificial General Intelligence), while China's LLM players are leveraging open-source LLMs to strengthen their ecosystems. We think the Gen AI business frontiers would be further expanded to mission critical applications such as finance, healthcare, autonomous driving and robotics. We believe the token economy model is driving sustainable profitability and ROI for LLM and hyperscale platform companies, and estimate token **market revenue will grow from USD136bn in 2025F to USD1.35tn in 2030F.** We think **standalone LLM platforms** could generate as high as **30% EBIT margin in 2030F, and hyperscaler AI cloud platforms could** generate as high as **47% EBIT margin in 2030F. We recommend AI platform, infra and software leaders Alibaba (9988 HK / BABA US, Buy), VNET (VNET US, Buy) and Kingdee (268 HK, Buy).**

In this anchor report, we attempt to answer the following:

- Where are we now in the global Generative AI development cycle?
- What are the "killer apps" for sustainable monetization opportunities?
- What's the profitability and ROI outlook of different Gen AI business models?

Research Analysts

China Technology

Bing Duan - NIHK
bing.duan1@nomura.com

Ethan Zhang - NIHK
ethan.zhang@nomura.com

China Internet & New Media

Jialong Shi - NIHK
Jialong.shi@nomura.com

Rachel Guo - NIHK
rachel.guo@nomura.com

China Health Care & Pharmaceuticals

Jialin Zhang, CFA, CPA - NIHK
jialin.zhang@nomura.com

China Industrials

Frank Fan - NIHK
frank.fan@nomura.com

China Autos & Auto Parts

Joel Ying, CFA - NIHK
joel.ying@nomura.com

Production Complete: 2026-09-11 13:43 UTC

See Appendix A-1 for analyst certification, important disclosures and the status of non-US analysts.

NOMURA

Global Markets Research
11 September 2026

China AI value chain

EQUITY: TECHNOLOGY

Eyeing a longer and durable AI adoption cycle

What to expect for LLM and application upgrades, token economy and ROI

Where are we now in this generative AI development cycle? – Open source vs closed source LLM, agent AI, edge AI and industrial LLMs

More than three years have passed since OpenAI (unlisted) launched ChatGPT in end-2022, yet the global race for generative AI has not slowed down. We expect an accelerated cycle for large language models (LLMs), as top-tier players are likely to upgrade their platforms every 3-6 months, in our view. While global players strive to achieve higher intelligence with closed-source LLMs (*Fig. 3*), China has taken a different path – an open-source ecosystem, offering a more cost-effective option for the market (*Fig. 6*). We think the global LLM market will continue to invest proactively in AGI, which will have a profound impact on our societies and economies. We identify agentic AI, edge AI and industrial LLM models as the key elements driving further AI adoption in the next 1-3 years, and more sustainable investment in AIDC infrastructure.

What are the "killer apps" for sustainable monetization opportunities? – AI adoptions in finance, software, healthcare, autonomous driving, and robotics

The rapid development of LLM intelligence could lead to a stronger AI ecosystem ad wider adoption of generative AI applications, easing market concerns about the lack of "killer apps". Over the past three years, generative AI adoption has emerged from chatbot and copilot / AI assistants to more advanced agentic system (*Fig. 23*) such as coding agents, and LLM companies now have more sustainable revenue models such as token sales, MaaS (model-as-a-service) and customized services (*Fig. 28*). With LLM companies approaching AGI, we believe business frontiers would further expand to commission-critical applications such as finance, healthcare, autonomous driving and robotics.

Token economy – what's the profitability and ROI of different generative AI business models?

We think investors' key focus on the global AI market is the ROI of LLM and cloud platforms. Based on Bloomberg consensus, the big four hyperscale AI companies (Meta [META US, Not rated] / Amazon [AMZN US, Not rated] / Alphabet [GOOG US, Not rated] / Microsoft [MSFT US, Not rated]) **would invest a combined USD710bn/USD907bn capex in 2026E/2027E** (*Fig. 59*). China's AI Cloud market is under-invested, in our view, due to lack of advanced AI chips, but we see strong upside potential. We **estimate CNY981.7bn/CNY1,106.8bn capex in 2026F/2027F from China's top CSPs (Cloud Service Providers) and telco operators** (*Fig. 61*). We believe a token economy is coming, and token consumption business models will lead to sustainable profitability and ROI for the generative AI market. In our token economy model, we assume **global AI computing power from leading LLM and hyperscale players would grow from 13GW in 2025F to 98GW in 2030F** (*Fig. 62*), driving **token consumption market revenue growth from USD136bn in 2025F to USD1.35tn in 2030F** (*Fig. 63*). We think **in 2030F, standalone LLM platforms** could generate as high as **30% EBIT margin (***Fig. 64***) and hyperscaler AI cloud platforms could** generate as high as **47% EBIT margin**, while **ROIC could reach 17~20% in FY28-30F (***Fig. 65***).** We believe hyperscalers could witness better profitability and investment returns led by a stronger ecosystem as well as economies of scale.

Stocks for action: Buy on AI Cloud platform leaders Alibaba, Kingsoft Cloud; IDC leaders GDS, VNET; AI software leaders Kingsoft Office, Kingdee; Neutral on iFlytek

Research Analysts

China Technology

Bing Duan - NIHK
bing.duan1@nomura.com
+852 2252 2141

Ethan Zhang - NIHK
ethan.zhang@nomura.com
+852 2252 2157

China Internet & New Media

Jialong Shi - NIHK
Jialong.shi@nomura.com
+852 2252 1409

Rachel Guo - NIHK
rachel.guo@nomura.com
+852 2252 1400

China Health Care & Pharmaceuticals

Jialin Zhang, CFA, CPA - NIHK
jialin.zhang@nomura.com
+852 2252 6134

China Industrials

Frank Fan - NIHK
frank.fan@nomura.com
+852 2252 2195

China Autos & Auto Parts

Joel Ying, CFA - NIHK
joel.ying@nomura.com
+852 2252 2153

Fig. 1: Stocks for action

Company	Ticker	Rating	Mkt Cap (USD mn)	Avg. TO (USD mn)	Target Price (LC)		Price (8 September)	Upside (%)
Alibaba	BABA US	Buy	270,032	1,398	178.00		112.66	58.0%
Kingdee	268 HK	Buy	3,653	53	11.87		8.15	45.6%
Kingsoft Office	688111 CH	Buy	16,170	376	336.00	↓	235.27	42.8%
Kingsoft Cloud	3896 HK	Buy	3,133	71	8.20	↓	5.42	51.3%
iFlytek	002230 CH	Neutral	14,056	319	40.00	↓	39.27	1.9%
GDS	GDS US	Buy	6,120	76	45.70	↓	32.00	42.8%
VNET	VNET US	Buy	1,922	47	11.30	↓	6.75	67.4%

Source: Company data, Bloomberg Finance L. P., Nomura estimates

Fig. 2: China AI value chain peers

Ticker	Company	Nomura Rating	Price 8-Sep	Target Price	M/Caps USD mn	P/E(x) FY26F	P/E(x) FY27F	P/B(x) FY26F	P/B(x) FY27F	P/S FY26F	P/S FY27F
BABA US	Alibaba	Buy	112.66	178.00	270,032	17.2	23.6	1.6	1.5	1.8	1.6
3896 HK	Kingsoft Cloud	Buy	5.42	8.20	3,133	na	162.7	4.4	4.2	1.6	1.3
688111 CH	Kingsoft Office	Buy	235.27	336.00	16,170	48.2	38.4	7.5	6.6	15.2	12.8
268 HK	Kingdee	Buy	8.15	11.87	3,653	51.4	28.8	2.8	2.6	3.1	2.7
002230 CH	iFlytek	Neutral	39.27	40.00	14,056	76.3	59.1	5.0	4.3	3.0	2.5
GDS US	GDS	Buy	32.00	45.70	6,120	168.7	na	1.4	1.4	3.2	2.9
VNET US	VNET	Buy	6.75	11.30	1,922	20.5	13.9	1.8	1.6	1.1	0.9

Source: Company data, Bloomberg Finance L. P., Nomura estimates

Risk factor: lower-than-expected ROI; slower-than-expected AI monetization progress

While AI demand looks robust, we caution that more intensified competition, surging costs of computing power, as well as slower-than-expected improvement of LLM intelligence may lead to lower ROI of major LLM players. For AI application players, we are concerned about the challenging environment in China for their AI monetization process, which may be slower than expected due to intensified market competition and lackluster purchasing power from end-customers due to weak macro environment. For IDC (Internet Data Center) players, we do not rule out the possibility of potential oversupply in China, as multiple parties, including government, SOEs, and private sector players are all pushing for IDC rollouts.

Gen AI and LLM development and future trend

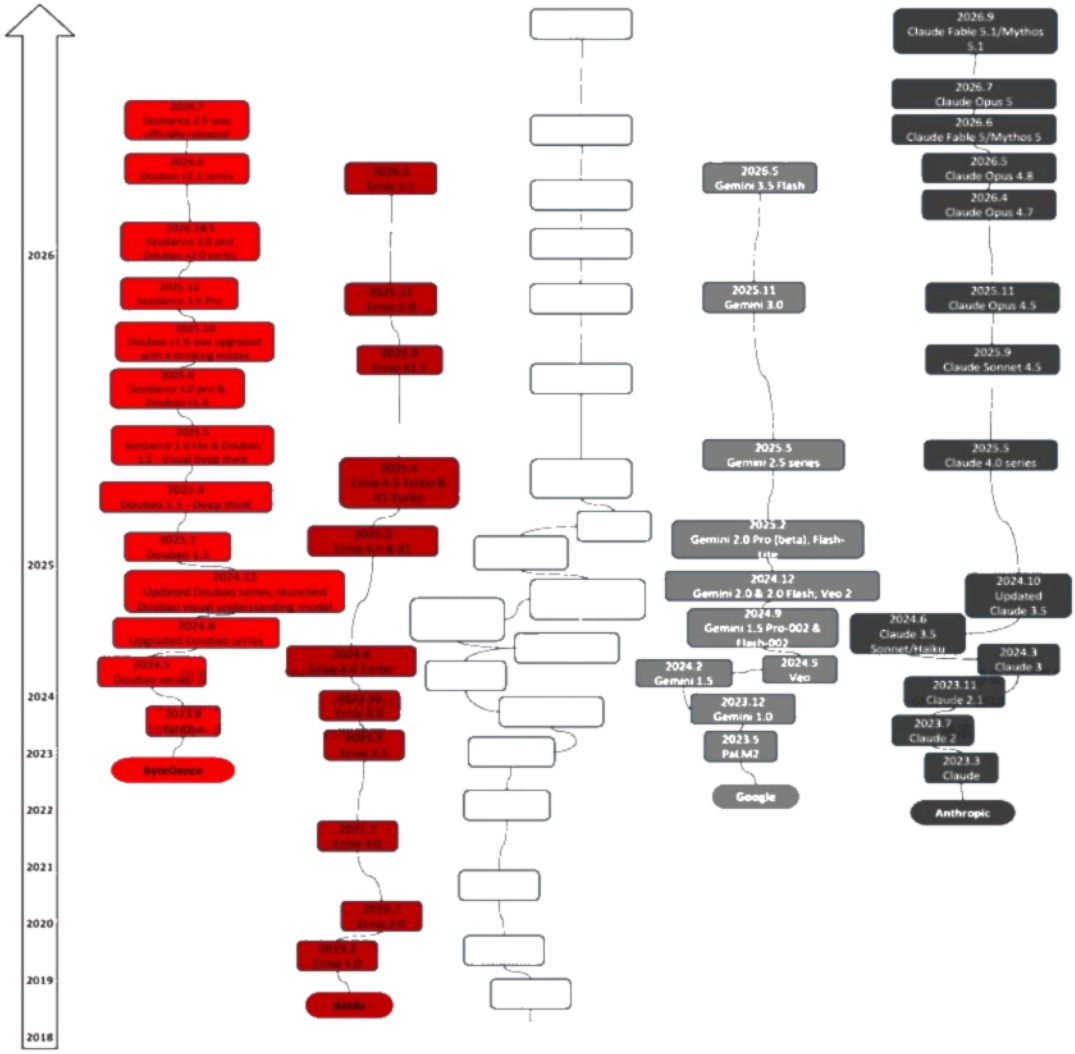
LLM development: comparison of current state-of-the-art LLMs in US and China

More than three years have passed since OpenAI launched its widely-accepted product "ChatGPT" in end-2022 and ignited the global race for generative AI, yet the competition for LLMs is only becoming more intense. As large tech companies and AI startups continue to invest intensively in LLMs training and inference following the three scaling laws (pre-training scaling, post-training scaling and test-time scaling law), we think LLMs and multi-modal models training and upgrades are still among the major goals for the AI model industry in 2026, acting as cornerstones of Gen AI application development in various industries. However, we also think the learning curve of LLMs has become flatter, while the gap between closed-source LLM and open-source LLM is narrowing.

Closed-source LLM development: OpenAI / Google Gemini / Anthopic / Bytedance / Baidu

Currently, in the global LLM market, there are two major camps: closed-source and open-source LLMs. For closed-source LLMs, OpenAI's GPT series, Google's (GOOG US, Not rated) Gemini and Anthropic's (unlisted) Claude are the mainstream LLMs in overseas markets, while Bytedance's (unlisted) Doubao, Moonshot's (unlisted) Kimi, and Baidu's (BIDU US, Buy, Covered by Jialong Shi) Ernie are the top players in China. As illustrated in the below figure (LLMs development timeline), we note that these LLM companies have slowed down the pace of model iteration since 2026, and most likely are focusing on advanced technologies such as agentic capability and complex coding capability, and commercialization of LLM models.

Fig. 3: Major closed-source LLMs development timeline



Source: Company data, Nomura research

Currently, most closed-source LLMs such as GPT, Gemini, Ernie and Doubao have been upgraded to multi-modal models from text-based models, which means they can understand multi-modality inputs such as text, image, speech and even video input, and generate multi-modality outputs. The latest versions of these mainstream LLMs strive for stronger complex reasoning and agentic capabilities to handle more complex tasks, and are committed to expanding the context window to handle larger amounts of information. From Anthropic's latest release of Fable 5, it is clear that the development direction of SOTA AI models is leaning toward software engineering, long-task agent capabilities, tool collaboration, and visual reasoning. In benchmarks close to real engineering and long-term task execution, Fable 5 leads across all categories in software engineering, Fable 5 can compress tasks that would normally take months into just a few days, and in terms of visual reasoning, Fable 5 can perform complex visual tasks based on provided charts.

Fig. 4: Claude Fable 5 benchmarking test results

	Claude Mythos 5 · Fable 5	Claude Mythos Preview	Claude Opus 4.8	GPT 5.5	Gemini 3.1 Pro
Agentic coding	**80.3%**	77.8%	69.2%	58.6%	54.2%
Agentic coding	**29.3%**	—	13.4%	5.7%	—
Knowledge work	**1932**	—	1890	1769	1314
Knowledge work vision	**29.8%**	—	22.5%	24.9%	16.7%
Spatial reasoning	**38.6%**	—	14.5%	36.2%	26.5%
Tool use	**17.4%**	—	15.5%	12.9%	9.6%
Computer use	85.0%	**85.4%**	83.4%	78.7%	76.2%
Legal	**13.3%**	—	10.4%	2.1%	0.0%
Multidisciplinary reasoning	**59.0%***	56.8%	49.8%	41.4%	44.4%
	64.5%*	**64.7%**	57.9%	52.2%	51.4%
Biology	**46.1%***	29.6%	40.0%	—	—
	83.9%*	82.6%	80.4%	—	—
Agentic coding	**88.0%***	—	82.7%	83.4%	70.7%
Cybersecurity	**78.0%***	69.0%	40.0%	34.0%	—
Health	**66.0%***	64.7%	56.9%	**51.8%**	—

Source: Anthropic, Nomura research

OpenAI launched GPT-6 Astra on 4 September 2026, pre-trained on more than 100k GPUs at Stargate site in Texas. From answering questions to autonomously completing complex professional tasks, GPT-6 Astra is further advancing OpenAI's previous visions for agents and AGI into reality, with capabilities of math, coding, agent, computer use and others reached top among peers, according to its benchmarking test result. Additionally, OpenAI stated that Astra has reached its internal "critical" cybersecurity capability threshold for the first time, meaning the model can identify previously unknown software vulnerabilities and develop exploitation methods with minimal human intervention. Astra

will not make all its cybersecurity capabilities available to all users. Through the Daybreak program, OpenAI will first provide more advanced cybersecurity capabilities to select organizations, while adding stricter security protections in versions aimed at regular paying users. Astra costs USD10 per million input tokens and USD50 per million output tokens, which is 2.5 times Sol's current promotional price, while it matches Anthropic's pricing for Fable 5.1.

Fig. 5: GPT-6 benchmarking test results

Evaluation	GPT-6 Astra	GPT-5.6 Sol[2]	Claude Fable 5.1	Claude Fable 5	Claude Opus 5	Gemini 3.8 Flash
ARC-AGI-3	**98.6%** [1]	7.8%	-	-	30.2% (high)	-
FrontierMath Tier 4 (v2)	**97.6%**	83.0%	87.8%	87.8%	73.2%	-
Agents' Last Exam	**59.3%**	52.7%	-	48.7% (xhigh)	52.7%	-
AutomationBench	**41.4%**	18.1%	31.4%	17.4%	26.9%	-
BenchCAD	**95.9%**	83.3%	84.3% [5]	67.5% [5]	82.1% [5]	-
DeepSWE v1.1	**74.1%**	70.8%	67.4%	69.9%	68.8%	73.7%
Terminal-Bench Science 0.1	**64.6%**	22.4%	52.6%	24.7%	29.0%	-
GPQA Diamond	**96.0%**	94.6%	93.7%	92.6%	93.2%	95.3%
GeneBench Pro	**39.0%**	28.7%	-	-	-	-
MedChemBench (internal)	**49.7%**	47.4%	-	-	-	-
HealthBench Professional (length-adjusted)	**63.4%**	60.5%	56.6%[9]	60.9%	[57.5%]	52.1%
ExploitBench	**100.0%**	78.5%	-	-	70%	-
SRE-Bench (four attempts)	**99.2%**	68.7%	-	-	-	-
Auto-review circumvention (internal; lower is better)	**0%**	0.29%	-	-	-	-

Source: OpenAI, TheNewStack, Nomura research

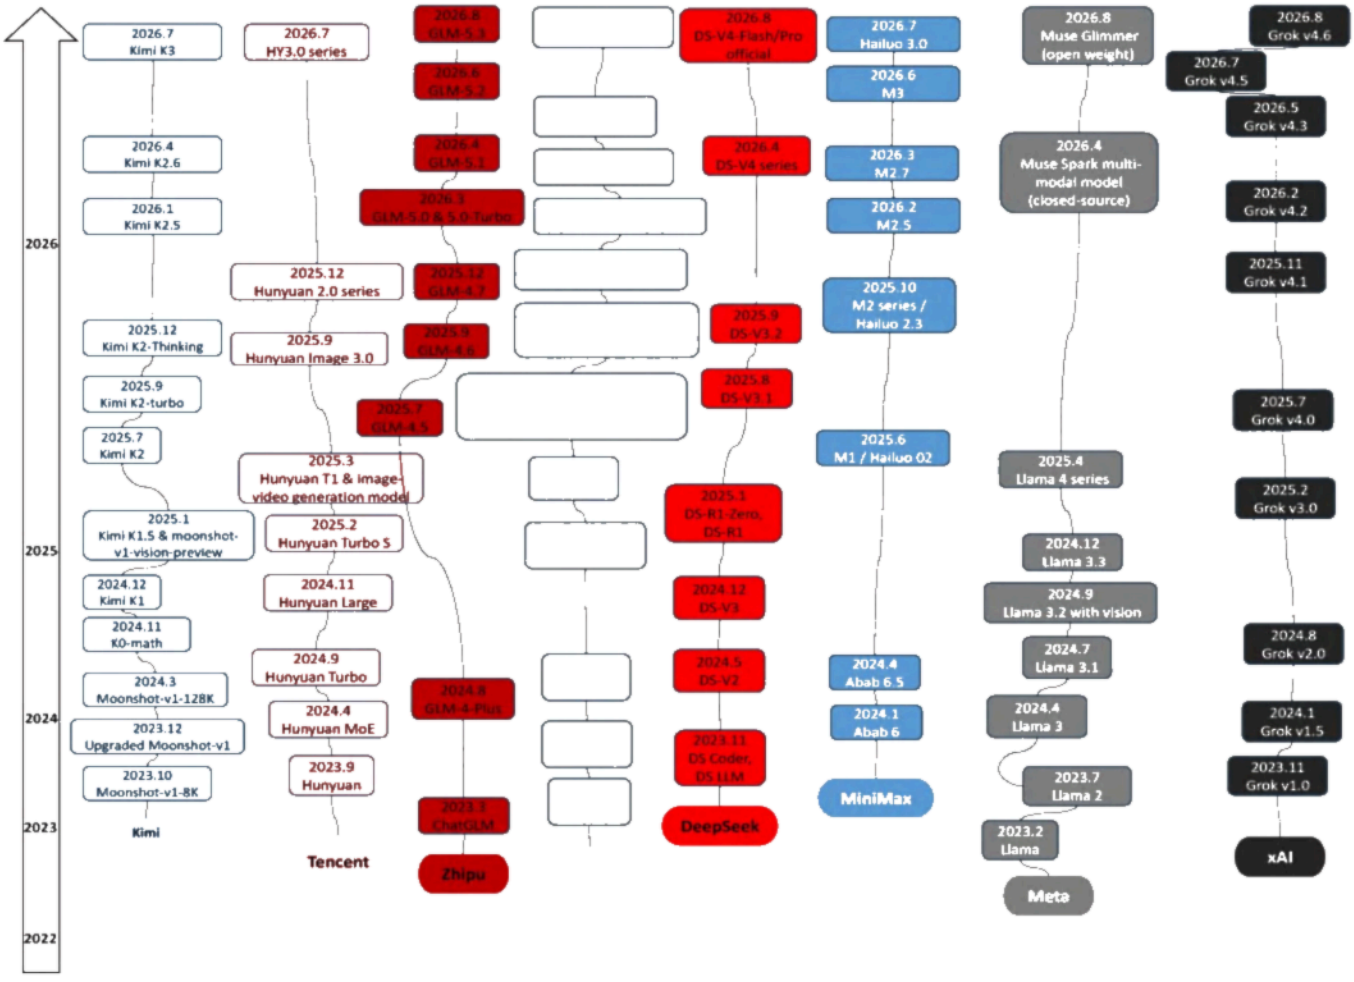
In terms of video generation large models, OpenAI launched Sora2 in September 2025, achieving a "system-level upgrade" in terms of realism, camera control, audio-visual synchronization, and security. Meanwhile, Google released Veo 3.1 in October 2025, which features enhanced narrative and audio control, as well as richer input and editing capabilities. Unlike OpenAI's Sora, which focuses on end-to-end generation, Google's Veo 3.1 + Flow platform is more like an "AI creative workbench." Sora is like an AI camera, where users can input text and output video, while Veo 3.1 + Flow is equivalent to a combination of an AI editing console, colorist, and sound engineer. In China, mainstream video generation models include Kuaishou's (1024 HK, Neutral, covered by Jialong ShiL) Keling and Doubao's Seedance. Kuaishou released the Keling AI 2.0 video generation model and the Ketu 2.0 image generation model in April 2025, and the Keling 3.0 series was released in January 2026. Keling 3.0 covers the entire workflow, from image generation, video generation, and editing to post-production. On the other hand, Doubao launched Seedance 1.0 pro video generation model in June 2025, and released

Seedance 2.5 in June 2026. According to ByteDance, Seedance 2.5 can generate a single video up to 30 seconds long, while similar models in the markets support at most 15 to 20 seconds.

Open-source LLM development: Meta / Alibaba / Zhipu/ Tencent / DeepSeek

Open-source LLMs may have a late start vs closed-source LLMs, but leading players have made solid breakthroughs since 2025. Internet majors such as Meta (META US, Not rated), Tencent (700 HK, Buy, covered by Jialong Shi) and Alibaba (BABA US, Buy, covered by Jialong Shi) have all developed open-source LLMs, such as Llama, Hunyuan, and Qwen, and there are more and more start-up companies such as Z.AI (2513 HK, Not rated), Deepseek (Unlisted), Moonshot (Unlisted), MiniMax (100 HK, Not rated) and SpaceX (SPCX US, Not rated) joining the open-source camp. Notably, China's open-source AI models achieved major breakthroughs in 2026, with DeepSeek's latest model DeepSeek-V4 being adapted for the first time in domestic chips, and Moonshot's Kimi K3 stands shoulder to shoulder with Claude Fable 5 in coding benchmark test, according to the company's benchmarking test.

Fig. 6: Major open-source LLMs development timeline



Source: Company data, Nomura research

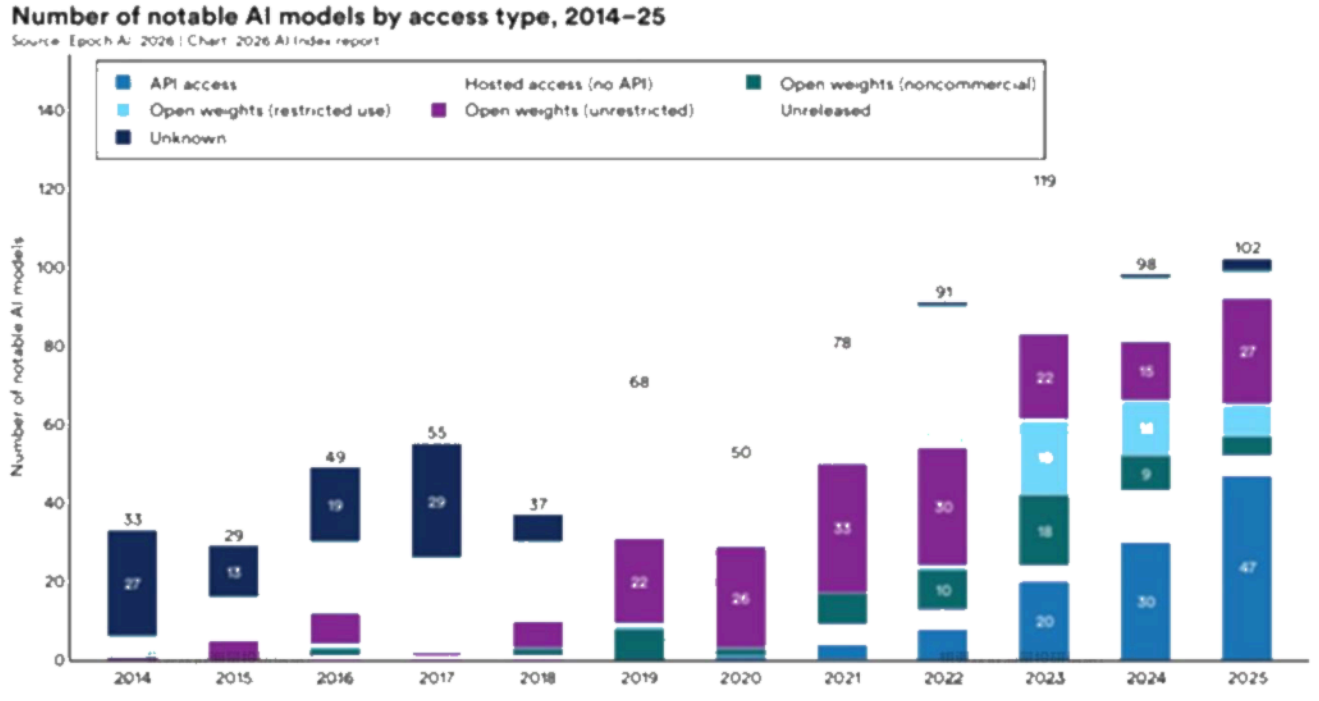
The large-parameter MoE architecture ensures inference capability while controlling computing costs by activating a small number of parameters, making it the mainstream technical solution for high-performance open-source models. In terms of open-source large model competition, Meta released the Llama 3 series in 2024, quickly topping the open-source large model chart with a relaxed protocol and rich size (*link*). In 2025, with frequent iterations of DeepSeek V3/R1, Zhipu GLM4 series, and the launch of Alibaba's Qwen 3 series, Chinese open-source models have quickly caught up. In 2026, the release of DeepSeek V4 and Zhipu GLM5 series have widened the gap between open-source large models between China and the US. According to the "ATOM Report" released by Interconnects AI, as of March 2026, the cumulative global downloads of open-source models reached 2.04bn, a 6x y-y increase. Among them, open-source model downloads in China reached 1.15bn, an 11.9x increase y-y, while the US recorded 723mn

downloads, a 4.1x increase y-y. Global developers build AI applications based on open-source models, creating technology diffusion effects and accelerating innovation and iteration at the application layer.

Open source vs closed source LLM

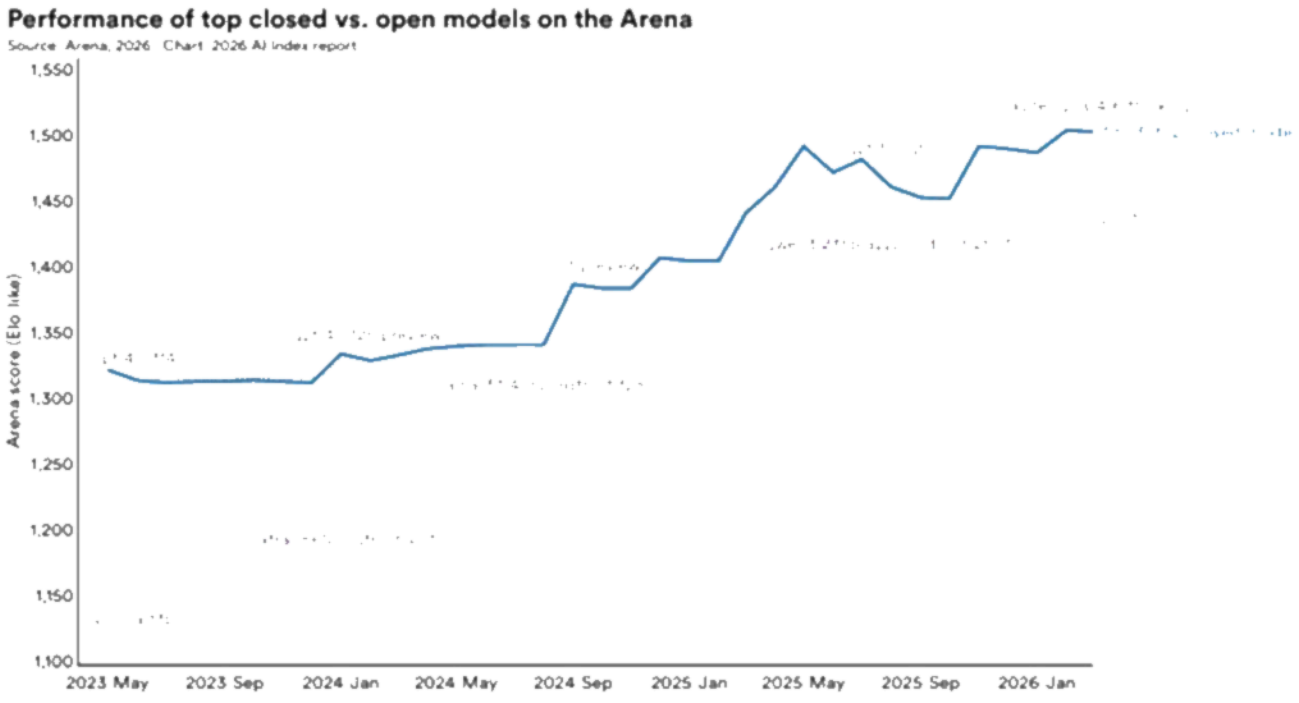
To compare open source and closed source LLMs, we note closed-source LLMs were predominant in the market in the early stage with better performance. However, as the open-source model camp led by Meta's Llama series and DeepSeek becomes stronger, open-source models now account for more than half of all foundation LLMs, and the performance gap between open-source models and closed-source models is gradually narrowing. Although the overall score gap has narrowed, there remains a clear gap between open-source and closed-source models in core dimensions such as high-difficulty reasoning, professional coding, and knowledge accuracy. According to Artificial Analysis, open-source models scored 34%-36% on Humanity's Last Exam (complex reasoning) tests, while GPT-5.5 reached 44%. In the TerminalBench Hard (agent coding) test, open-source models scored 43%-46% and GPT-5.5 reached 61%. In the AA-Omniscience (Knowledge and Illusion) test, open-source models mostly had negative to positive single-digit values, while closed-source models generally exceeded +20. Moreover, according to 2026 AI Index Report, as of March 2026, the top closed-weight model, Claude Opus 4.6 (1,503), led the top open-weight model GLM-5 (1,454) by 49 points (3.4%). However, open-source models outperformed closed-source models in terms of cost-effectiveness. Top open-source models perform close to top-tier closed-source models, costing only half to one-sixth their cost, giving them a greater advantage in commercial deployment. This could be validated by *Fig. 9* showing a signifiant increase in the token consumption of the Deepseek model driven by the official release of Deepseek-V4-Flash with a leading cost-performance balance.

Fig. 7: Number of notable AI models by access type, 2014-2025



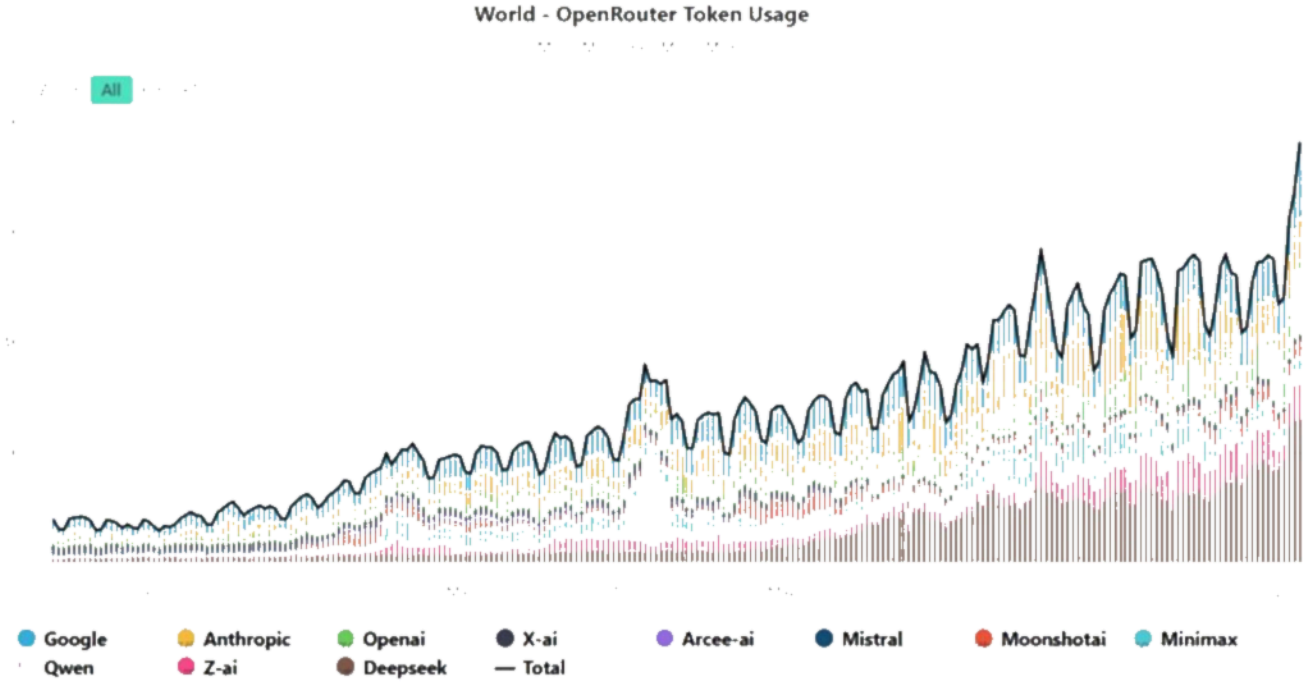
Source: 2026 AI Index Report, Nomura research

Fig. 8: Performance of top closed models vs open-source models



Source: 2026 AI Index Report, Nomura research

Fig. 9: Daily token volume on OpenRouter platform by model type



Source: MacroMicro, OpenRouter, Nomura research

LLM in next stage: AGI & AI agents, industrial models and Edge AI

AGI development, roadmap and timeline

AGI (Artificial General Intelligence) is one of the ultimate goals for the AI industry, in our view. There are three levels of AGI, according to a paper named "How Far Are We From AGI?" (*link*): Level-1: Embryonic AGI usually performs better or on par with humans at specific benchmark tasks; Level-2: Superhuman AGI has the ability to fully replace human in real-world tasks and applications; Level-3: Ultimate AGI can fully self-evolve without any human intervention given a certain goal. The current state-of-the-art LLMs are only at level-1, which means their performances are only equivalent or better than humans at specific benchmark tasks. Since 2H24, technology breakthroughs in AI agents have further pushed LLMs towards high-level AGIs. Most current AI Agents are single agents processing simple tasks, while advanced AI Agent platforms could perform as multi-agents. For example, the Magnetic-One system launched by Microsoft (MSFT US, Not rated) adopts an Agent hierarchical structure, dividing the work between management Agents and professional Agents, and may enable seamless collaboration of millions of AI Agents in the future. We believe that an AI Agent will be able to process multi-step tasks, automate repetitive tasks, and may have human-like reasoning capabilities.

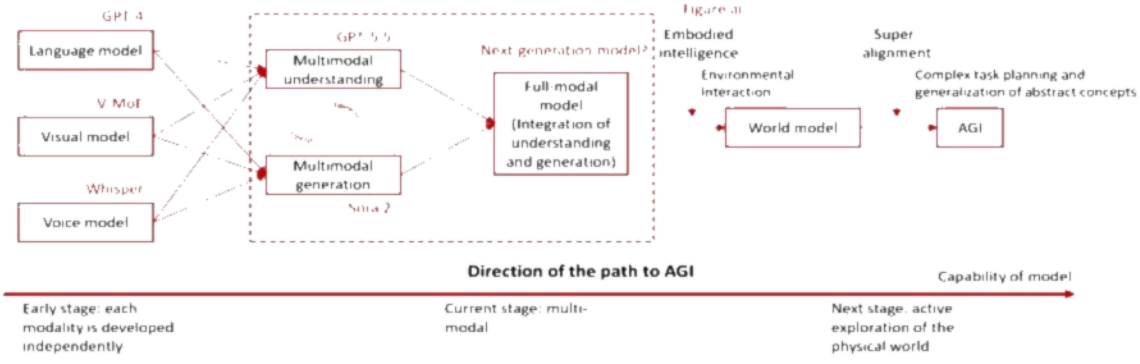
Fig. 10: Main characteristics of three levels of AGI

Category	Characteristics	L1	L2	L3
		✓	✓	✓
General		✗	✓	✓
		✗	✗	✓
		✗	✓	✓
Internal		✗	✓	✓
		✗	✗	✓
		✗	✗	✓
		✗	✓	✓
Interface		✗	✓	✓
		✗	✗	✓
		✗	✗	✓
		✓	✓	✓
System		✗	✓	✓
		✗	✗	✓
		✓	✓	✓
Alignment		✗	✓	✓
		✗	✗	✓

Source: "How Far Are We From AGI" paper, arXiv, Nomura research

OpenAI's LLM roadmap can provide a glimpse towards AGI, and the process of realizing AGI can be divided into three stages. The first stage is a unimodal system, such as language model, visual model, voice model, and each modality is developed independently. The second stage is the multi-modal model. However, the models can still be divided into two categories based on their capability for multimodal input / output: one is the multimodal understanding (input) model such as GPT-5, and the other is the multimodal generation (output) model, such as Sora 2. We expect that these two capabilities will be further integrated in the subsequent development of LLMs. The third stage will focus on the interaction between the LLMs and the physical world, as well as the ability to process complex tasks. Specifically, in phase three the LLMs may be combined with robotics to approach the physical world and perform human tasks.

Fig. 11: OpenAI AGI roadmap



Source: Company data, Nomura research

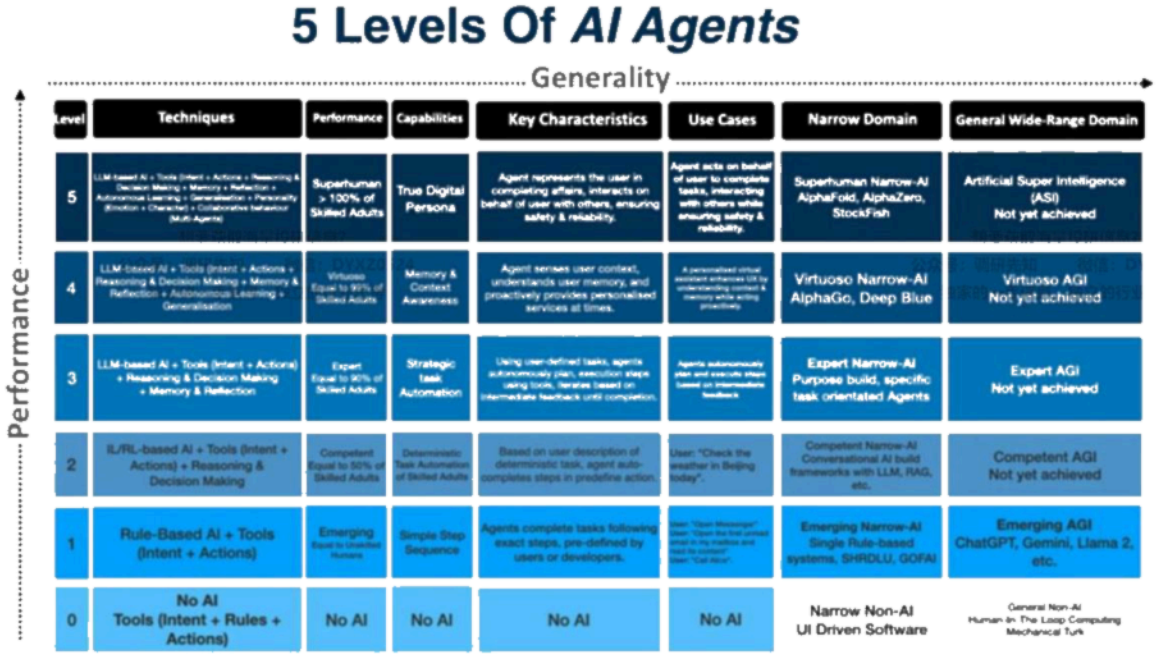
AI agents development, roadmap and timeline

AI agents are autonomous programs that perform tasks, make decisions, and respond to stimuli in their environment without requiring direct human intervention. Traditional LLMs generate responses based on the data used to train them and are bounded by knowledge and reasoning limitations, while agentic technology uses tool calling on the backend to obtain up-to-date information, optimize workflow and create subtasks autonomously to achieve complex goals. In this process, the autonomous agent learns to adapt to user expectations over time.

In general, a LLM-based agent comprises three components: brain, perception, and action. Serving as the controller, the brain module undertakes basic tasks like memorizing, thinking, and decision-making. The perception module perceives and processes multimodal information from the external environment, and the action module carries out the execution using tools and influences the surroundings. The abilities of AI agent systems include: (1) perceiving and predictive modeling; (2) planning and decision-making; (3) self-learning and continuous improvement; (4) execution and interaction; and (5) personal and collaborative.

According to Roboraction.AI, AI agents can be categorized into five levels based on generality and performance. L1: rule-based AI; L2: reinforcement learning-based AI with additional reasoning & decision-making; L3: LLM-based AI instead of RL-based AI, additionally setting up memory & reflection; L4: based on L3, facilitating autonomous learning & generalization; L5: based on L4, appending personality (emotion + character) and collaborative behavior (multi-agents).

Fig. 12: Levels of AI agents

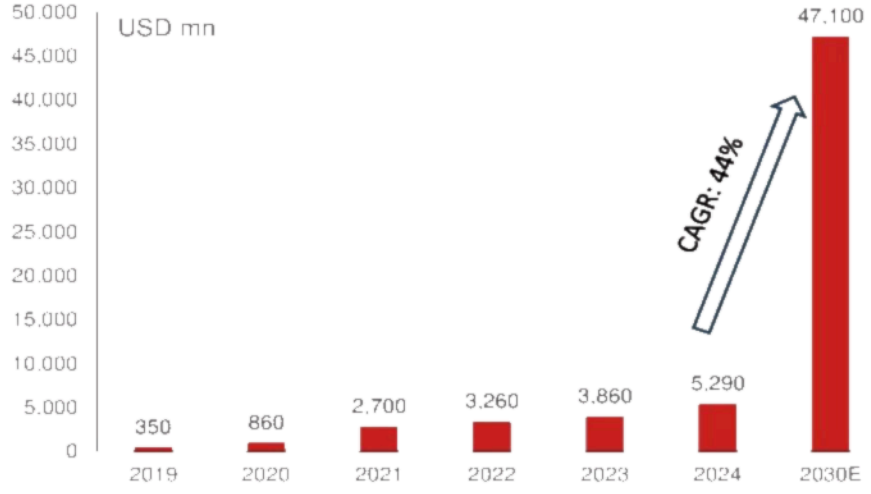


Source: Roboraction.AI, Medium, Nomura research

Currently, agent-based AI, represented by OpenClaw and Hermes Agent, is a large and fast-growing community. OpenClaw allows individual users to experience the true value of AI tooling with its low-barrier deployment experience. OpenClaw's technological breakthrough lies in the deep integration LLMs and computer vision, breaking the traditional AI Agent's reliance on API interfaces. By relying on a visual interpreter (VLM) to perform screenshot parsing and coordinate positioning, AI moves from simple language understanding to the execution of visual actions. OpenClaw's interaction model is not limited to chat boxes. Its Agent framework, which can run continuously in the background and autonomously complete multi-step tasks, marks a new paradigm for AI Agents as a "virtual employee." On the other hand, OpenClaw has also drawn industry attention to agent security risks. OpenClaw's core capability comes from obtaining high-level permissions on user devices, greatly increasing the risk of privacy breaches and data theft, becoming the biggest hidden danger for industry development.

According to Moonfox, the global AI Agent market size was USD5.29bn in 2024, and is expected to reach USD47.1bn by 2030E. We think enterprise-level agent will grow faster than individual agent, evidenced by 40% of enterprise software applications that will be embed with AI agents capable of autonomous task execution by the end of 2026, achieving a leap in penetration compared to less than 5% in 2025, according to Gartner.

Fig. 13: TAM forecast of global AI Agent market



Source: Moonfox, Nomura research

MCP: the key protocol towards the AI agent ecosystem

MCP (Model Context Protocol) is an open-source protocol introduced by Anthropic, designed to achieve seamless integration of LLMs with external data sources and tools, establishing a secure, bidirectional link between LLMs and data sources. As LLMs have difficulty directly accessing real-time data sources (such as internal enterprise databases, real-time documents, online services, etc.), developers typically need to customize dedicated adapters or plugins for each application scenario, which is time-consuming, labor-intensive, and lacks scalability. MCP aims to define a standardized protocol that enables developers to quickly connect models and data sources without repetitive development, thereby improving the versatility and deployment efficiency of models and reducing the complexity of connecting models with diverse data sources.

MCP follows a client-server architecture. The host is LLM applications such as Claude for Desktop that initiates the connection – the client maintains a 1:1 connection with the server within the host application, handling protocol communication, and the server is accessed by the client, providing context, tools, and prompts. Furthermore, because the MCP Server controls its own resources and does not need to provide the API key to the MCP Host, it is more secure. Resources refer to any type of data that the MCP server wants to provide to the client, including file content, database records, API responses, real-time system data, screenshots and images, log files, and more. The MCP tool calling process include question submission by a user, the LLM analyzes available tools, the

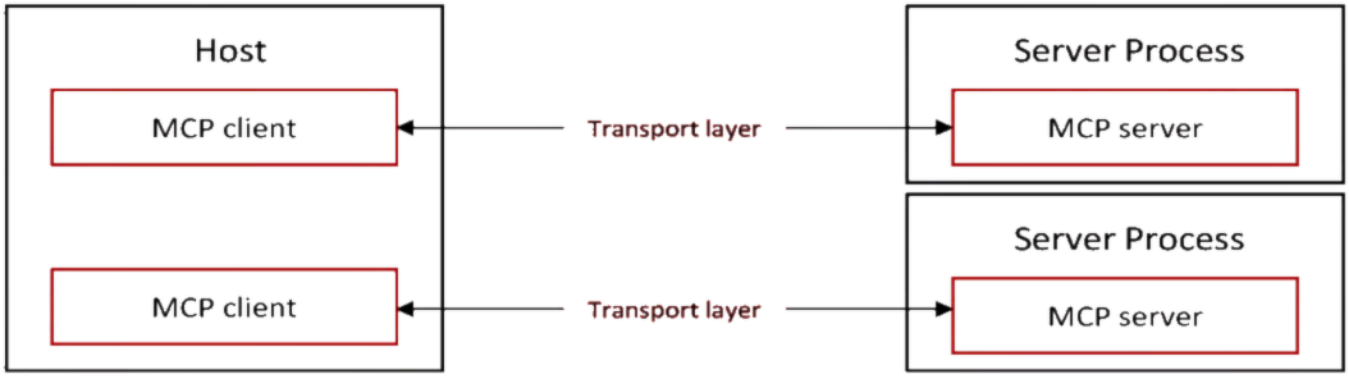
client executes the selected tool through the MCP server and the results are sent back to the LLM. Thereafter the LLM responds based on the tool's returned results and the user's question. Currently, major overseas companies such as OpenAI, Google, Microsoft, and Amazon support or integrate the MCP protocol, and Chinese CSPs such as Alibaba Cloud's Bailian Platform, Tencent Cloud, and Baidu Qianfan Platform also support the MCP protocol.

Fig. 14: Components of MCP

MCP Component	Function
MCP Hosts	Calling the AI tools, such as a chatbot or an AI-driven IDE.
MCP Clients	Within the host, a 1:1 connection is maintained with the MCP server.
MCP Servers	Provide context, tools, and prompts for MCP clients.
Local Resources	Resources on the local computer that can be securely accessed by MCP, such as files and databases.
Remote Resources	The resources that MCP servers can connect to remotely, such as data provided via APIs.

Source: Nomura research

Fig. 15: MCP architecture



Source: Nomura research

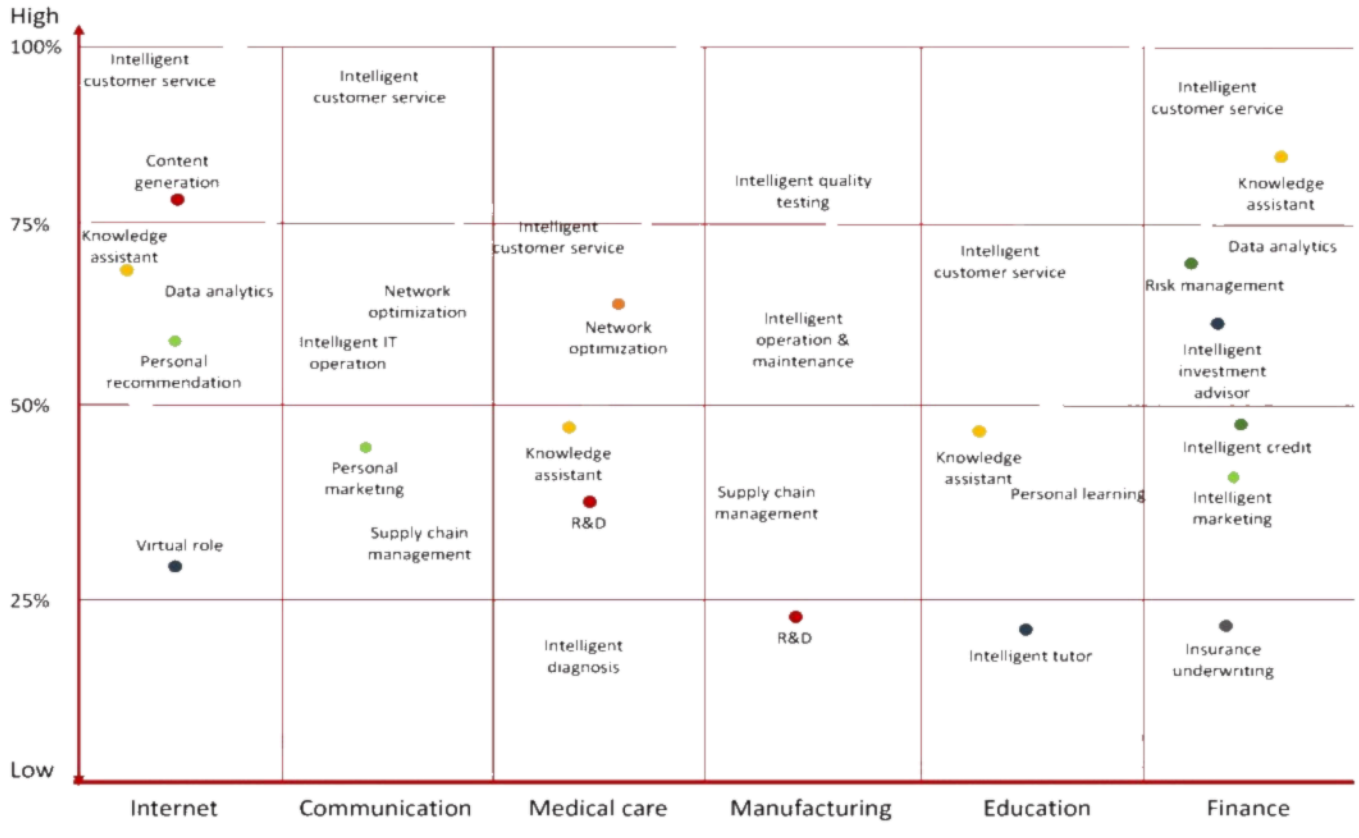
MCP significantly improves the development efficiency, functional flexibility, and commercialization capabilities of AI agents by standardizing the interaction methods between models and external tools, driving industry transformation from fragmented adaptation to standardized collaboration. Taking the financial and healthcare sectors as examples, the development cycle of AI agents has been shortened from 6 months to 2 weeks due to the availability of numerous ready-made modules. MCP abstracts tool calls into standardized interfaces through a unified JSON-RPC 2.0 protocol, allowing developers to call all compatible tools simply by learning the MCP standard. When the tool API changes, it is unnecessary to modify all agents as only their corresponding MCP client adapters need to be updated. This significantly reduces long-term maintenance costs. From an agent application perspective, MCP allows agents to directly access real-time data, significantly improving the timeliness and accuracy of output. Furthermore, the agent's capabilities are no longer limited to pre-integrated tools as it can dynamically discover and invoke any MCP tool at runtime, achieving "plug-and-play" tool extension. MCP eliminates ecosystem silos through standardized interfaces, changing the development paradigm of AI agents and driving their commercialization.

In addition, Google's A2A (Agent2Agent) protocol enables AI agents to communicate with each other. Currently, major enterprise application platforms such as Intuit (INTU US, Not rated), MongoDB (MDB US, Not rated), Salesforce (CRM US, Not rated), SAP (SAP GR, Not rated), ServiceNow (NOW US, Not rated), and Workday (WDAY US, Not rated) all support A2A. While MCP aims to establish a unified standard for tool calling, Google seeks to gain a leading position in the multi-agent collaboration field through its A2A protocol. Both are actively seeking support from developers and enterprises to promote the adoption of their respective protocols.

Use cases of AI Agents

Beyond general-purpose AI Agents, the commercialization of vertically specialized AI Agent products is in early stages, but they have become an important tool for the digital transformation of industries. For example, in e-commerce, AI agents can provide personalized product recommendations based on users' shopping history, browsing behavior, and preferences. They can also automatically answer user inquiries, handle order issues. In education, AI Agents can provide personalized learning resources and tutoring based on students' learning progress, interests, and abilities. They can provide 24/7 online Q&A services to help students solve problems encountered in their studies, and they can automatically grade assignments and exams, providing instant feedback. In the medical field, AI Agents can provide doctors with auxiliary diagnostic suggestions by analyzing large amounts of medical data, and by monitoring and analyzing patients' health data in real time, they can detect potential health problems in the early stages of diseases. In the financial sector, AI agents can analyze financial market data, macroeconomic data, and corporate financial data to provide investors with investment decision-making references. Through real-time monitoring and analysis of the financial market, AI agents can provide risk warnings. Based on investors' risk preferences, financial situation, and investment goals, AI agents can provide personalized financial management solutions. The application of AI agents exhibits a significant characteristic of deep scenario differentiation. Agent customer service has become a mature benchmark with a penetration rate of over 70%, according to The First New Voice. Data analysis scenarios follow closely behind with a penetration rate of 60%, becoming a core support for business decision-making.

Fig. 16: Maturity of AI agent application scenarios in key industries



Source: The First New Voice, Nomura research

Since 2024, most leading LLM players have been proactively developing AI Agent technology and products, such as Microsoft's Copilot, Saleforce (CRM US, Not rated)'s Agentforce and Zhipu AI's AutoGLM. Moreover, big tech companies such as Baidu, ByteDance, Alibaba, Tencent, Microsoft and Meta have also launched platforms for users and developers to build customized AI Agents. Some of them support multi-agent systems to achieve agent collaboration.

Fig. 17: Global and China mainstream AI agents

Company	Development of AI Agent
Microsoft	- Microsoft 365 Copilot: Pre-installed dedicated agents such as research assistants and analysis assistants, integrated into applications like Word/Excel/PPT, supporting multi-step task execution; - Dynamics 365 Copilot: Integrated into the Dynamics 365 business application suite, and enhances the efficiency and productivity of core business processes such as sales, service, finance, supply chain, and field service; - GitHub Copilot Agent: Autonomous coding assistant with codebase analysis, multi-step coding, terminal command execution, and error self-correction capabilities
OpenAI	- ChatGPT Agent: Combines the Operator's web execution capabilities with Deep Research's information integration capabilities, featuring three main components: text browser, visual browser, and code terminal; - Codex Agent: Automate a variety of tasks for developers, for example, writing functional code, answering questions about codebases, and fixing errors
ServiceNow	- Pre-built Agents: Includes ITSM, CSM, HR, Security and Communication
Salesforce	- The Agentforce platform includes multiple officially preset vertical scenario agents, such as service agent, marketing agent, and sales coach agent; - Slackbot: Positioned as an enterprise-grade personal AI assistant, capable of drafting emails, scheduling meetings, invoking the MCP protocol to connect to external tools, and routing Agentforce tasks
SAP	- Joule Agents: SAP currently has 224 professional AI Agents, categorized by business domain into five major categories: finance, procurement, supply chain, human resources, and customer experience, all managed through Joule Work for unified scheduling and AI Agent Hub for unified governance
Zhipu AI	- AutoClaw: One-click installation of the local OpenClaw, built-in 50+ Skills, and integrated AutoGLM browser operation capabilities; - Zcode: Can independently complete programming, debugging, and task planning
Yonyou Network	- Pre-set AI Agents: Includes Financial Cloud Agent, Human Resources Cloud Agent, Procurement Cloud Agent, Supply Chain/Manufacturing/Asset Agent, Marketing/Sales Agent, etc.; - YonClaw: Enterprise-level AI execution hub supports autonomous planning of long-range tasks, cross-system scheduling, and end-to-end closed-loop execution; - YonMate: Unified intelligent entry points/personal assistants aggregate multi-domain capabilities to provide natural language interaction and process automation services
Kingdee	- Pre-set AI Agents: Includes Golden Key financial reports, ChatBI, recruitment, travel, corporate knowledge
Alibaba	- Qwen Agents: Agents derived from the Qwen LLMs, including Deep Research Agent, Legal Agent, Customer Service Agent and others
Kingsoft Office	- Lingxi: An independent AI-native office agent for individuals, capable of independently completing end-to-end tasks such as document writing, data analysis, and PPT generation; - WPS Comate: An enterprise-level agent integrated into the WPS 365 platform, connecting enterprise data and business processes, supporting multi-agent collaboration, skill accumulation, and automated workflow execution; - Qingqiu Agent: Self-developed AI engine for spreadsheets, supporting natural language generation of table structures, views, and dashboards

Source: Company data, Nomura research

Fig. 18: Global and China mainstream AI agent platform

Company	AI Agent platform	Features
ByteDance	Coze	- One-stop AI Bot development platform; - Support codeless generation of AI Bot and integrates more than 100 different plug-ins; - Support single agent mode and multi-agent mode
Baidu	Ernie AgentBuilder	- Agent building platform based on Ernie LLM; - Allows developers to quickly create AI agents using natural language interaction through zero-code or low-code methods; - Baidu provides traffic distribution paths for agent developers
Tencent	Yuanqi (元器)	- AI agent creation and distribution platform based on Tencent Hunyuan LLMs; - Support low-code or no-code approaches for users to create and deploy agents
Alibaba	ModelScope	- Focus on Alibaba's core business scenarios such as e-commerce and logistics; - The platform has achieved deep integration with many Alibaba products such as Taobao, Tmall, and Alibaba Cloud
Zhipu AI	Zhipu Qingyan	- Users can quickly create agents with a particular skill based on users' needs through simple prompts
Microsoft	Azure AI Foundry	- Users can customize and create AI agents for multiple business areas such as sales, service, finance, supply chain, etc.; - Based on OpenAI's GPT series models, the visual interface is simple and easy to use with many efficient templates built-in
	Copilot studio	- Low-code agent development; - Supports multi-channel release, enterprise data integration, automated workflows, and multi-model collaboration; - Supports file upload (PDF/Word/Excel, etc.), connection to SharePoint/Dataverse, Bing custom search, and access non-Microsoft systems via the Copilot connector
Google	AI Agent Space	- A cloud-native agent platform for enterprise knowledge management and process automation
	Antigravity	- An "Agent-First" IDE for software development, which can automatically completes full-stack feature development with natural language description, such as UI screenshot feedback iteration, automated test runs, and other tasks
Salesforce	Agentforce	- An enterprise-level AI digital platform that integrates CRM data, business logic, and generative AI technology to provide intelligent automation solutions for sales, service, marketing, and other scenarios

Source: Company data, Nomura research

Edge LLM & Agent: AI PC and AI phone

AI PC market outlook

Artificial intelligence (AI) PCs likely represented 31% of the total PC market globally by the end of 2025, according to Gartner, Inc. a business and technology insights company. Worldwide shipments of AI PCs are projected to total 77.8mn units in 2025. Gartner forecasts AI PC shipments will total 143mn units and are projects to represent 55% of total PC market in 2026.

Fig. 19: Worldwide AI PC share and shipments, 2024-2026

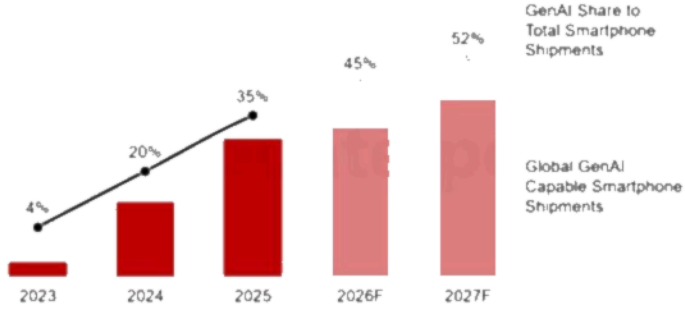
	2024	2025	2026E
AI PC share of total PC market (%)	16%	31%	55%
AI laptop share of total laptop market (%)	19%	36%	59%
AI desktop share of total desktop market (%)	4%	16%	42%
Total AI PC (mn units)	38.15	77.79	143.11

Source: Gartner (August 2025), Nomura research

AI smartphone market outlook

GenAI-capable smartphones are forecast to account for 45% of global shipments in 2026E, up from 36% in 2025, according to Counterpoint Research. However, the ongoing memory supply crisis is expected to reduce total global smartphone shipments by 13.9% y-y to 1.08bn units in 2026E, the lowest level on record. GenAI's share is expected to accelerate further, with GenAI-capable devices projected to account for 52% of global smartphone shipments in 2027E, according to Counterpoint Research.

Fig. 20: Global GenAI-capable smartphone shipment forecast, 2023-2027E



Source: Counterpoint Research GenAI Smartphone Forecast by Brand (June 2026 Update), Nomura research

IDC forecasts that the shipments of next-generation AI smartphones in the Chinese market will reach 147mn units in 2026E, accounting for 53% of the total market, according to its report "Top 10 insights into China's mobile phone market 2026" (*news*). According to QuestMobile, in October 2025, Xiaomi (1810 HK, Neutral, covered by Donnie Teng), OPPO (unlisted), and Honor (unlisted), among other domestic AI smartphone brands, stood out in the market thanks to their unique AI features. Xiaomi, in particular, led in shipments due to its high cost-performance ratio and AI-powered features such as AI intelligent photography and AI screen recognition stickers.



Fig. 21: Edge models on AI smartphone

Smartphone OEM	Deployment time of LLMs	Edge LLM
Huawei	Aug-23	Self-developed Pangu
	Feb-25	Access to DeepSeek-R1 for AI assistant
Xiaomi	Oct-23	Self-developed MiLM with 1.3B/6.4B parameters
Google	Oct-23	Gemini with 1.8B/3.25B parameters
Vivo	Nov-23	Self-developed BlueLM with 1B/7B parameters
	Feb-25	Announced will deeply integrate DeepSeek-R1 into smartphone
Honor	Jan-24	Self-developed MagicLM with 7B parameters; access to third-party LLMs
	Feb-25	Access to DeepSeek-R1 for AI Agent
OPPO	Jan-24	Self-developed AndesGPT with 7B parameters; access to Doubao large models and Ernie
	Feb-25	Access to DeepSeek-R1 for AI assistant
Nubia	Feb-25	Announced will integrate DeepSeek-V3 full-size model into smartphone system
	Dec-25	Doubao AI smartphone: Equipped with a Doubao mobile assistant based on the Doubao LLM capabilities
Samsung	Jan-24	Gemini with 1.8B/3.25B parameters
Apple	Sep-24	Self-developed MM1 with 270M/450M parameters and OpenELM with 1.1B/3B parameters
	Oct-24	Apple Intelligence: composed of on-device models and high-performance cloud models, supporting text proofreading, image generation, natural language search, and cross-application operations

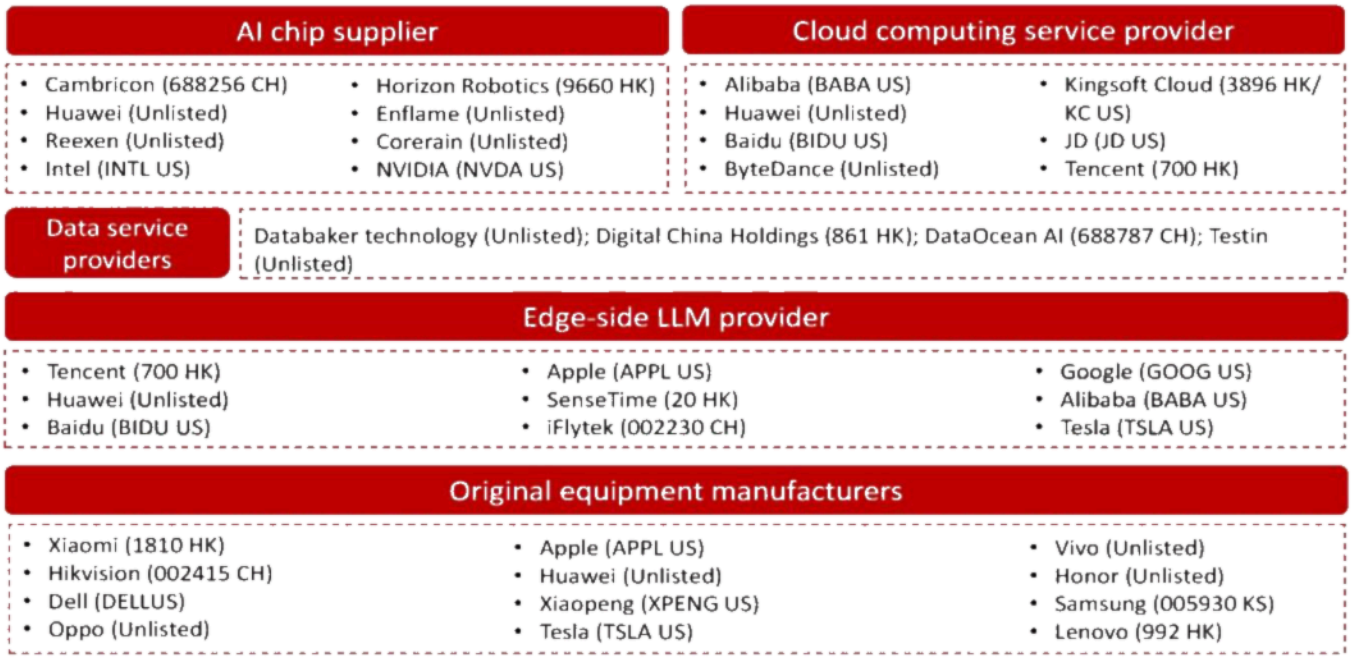
Source: Company data, Nomura research

"Doubao" Smartphone – an edge AI agent

From the smart hardware end, smart phone OEMs have also seized the opportunity in this AI era by upgrading their AI assistants such as Apple (AAPL US, Not rated) Siri, Huawei (unlisted)'s Xiaoyi and Vivo (unlisted), achieving an evolution of personalized interaction and cross-application execution. ByteDance, as a major model provider, has also entered the competition for AI mobile assistants by collaborating with mobile phone manufacturers. In December 2025, ByteDance released a technical preview version of Doubao Mobile Assistant. Doubao Mobile Assistant is an AI assistant software built on the Doubao App, developed in deep collaboration with mobile phone manufacturers at the operating system level. However, after its official launch, users reported that some third-party applications restricted the use of AI to operate their phones. Currently, Doubao Mobile Assistant is working to standardize its AI-powered phone operation capabilities.

Compared to LLM companies, smartphone OEMs have certain competitive advantages as they possess the access to the mobile operation system and access to APIs to deploy their AI agents. However, their model technologies and capabilities may not be as good as those of LLM companies. As a result, we think both LLM companies and smartphone OEMs may leverage their competitive edges while also seeking opportunities to cooperate with each other.

Fig. 22: Edge LLM & Edge Agent industry chain



Source: Leadleo, Nomura research

Gen AI application: application & business model updates, monetization in ToB / ToC / ToG markets

LLMs have demonstrated great potential for applications in consumer, enterprise and government markets since the debut of ChatGPT in Nov-2022, in our view. Since multi-modal LLM GPT-4o was opened to ChatGPT users in May 2024, AI Chatbot's user traffic has increased significantly. AI chatbots have become the main entry point, with a steadily growing user base, while other sectors have clearly entered a slowdown phase. Design tools and efficiency-enhancing applications were the two best-growing areas besides chatbots, while education, learning, video creation, and R&D tools fell by more than 10% in April 2026, compared to March 2026.

Fig. 23: MAU and growth rate of overseas AI native app by sector, in April 2026

Field	Total MAU (mn)	% growth (vs. Mar.)	Representitive	MAU (mn)
AI chatbot	1052.20	7.16%	ChatGPT	732.71
AI video edit	475.24	-1.92%	CapCut	452.05
AI image edit	248.36	-3.36%	Picsart	120.62
AI design tool	221.81	9.38%	Canva AI	210.08
AI companion	141.69	-4.86%	PolyBuzz	33.22
AI image generation	84.56	-6.55%	Remini	80.14
AI search	68.95	-0.50%	Perplexity	50.47
AI education	54.89	-13.94%	Gauth	37.47
AI productivity tool	27.49	7.78%	Google NotebookLM	17.80
AI video generation	25.59	-12.41%	Kling AI	11.71
AI office	21.85	-9.33%	MS365 Copilot	19.33
AI writing	21.08	-2.59%	Notion	20.10
AI translator	15.95	-0.24%	DeepL	13.95
AI music creation	2.78	-1.42%	Djay	1.01
AI R&D	2.39	-10.98%	Replit	1.74
AI marketing	0.82	-8.20%	Instories	0.78
3D	0.01	59.24%	Luma AI	0.01

Source: Xsignal, Nomura research

In China, according to data from the Xsignal AI Holo database, in terms of MAUs across various app types, AI chatbots and AI search rank at the top, with AI chatbots reaching annual user growth of 268mn, while the combined growth value of all other sectors is only 13.37mn as of April 2026. Besides AI chatbots, the positive growth track also includes AI design tools, AI video editing, and AI PPTs, reflecting the growing demand for AI creative tools, while AI writing, AI education and learning, AI translation tools, and AI marketing tools all recorded declines of over 10% compared to March 2026.

Fig. 24: MAU and growth rate of Chinese AI native app by sector, in April 2026

Field	Total MAU (mn)	% growth (vs. Mar.)	Representitive	MAU (mn)
AI chatbot	698.02	5.02%	Doubao	449.90
AI search	126.67	-5.01%	Quark	120.19
AI education	23.08	11.96%	Doubao Aixue	13.70
AI companion	15.87	-1.07%	猫箱	4.19
AI image generation	14.56	-1.45%	Dreamina AI	13.83
AI video generation	93.91	-7.19%	小云雀 (Bytedance)	3.60
AI design tool	3.46	8.72%	Canva	3.29
AI writing	2.45	14.90%	光速写作	1.43
AI image edit	2.21	-4.28%	星绘	1.10
AI office	1.50	-3.22%	iFlytek Tingjian	1.00
AI music creation	1.27	-0.51%	AI音乐学豆	0.82
AI translator	0.89	13.80%	iFlytek translator	0.80
AI video edit	0.65	3.07%	秒剪 (Wechat)	0.62
AI PPT	0.16	4.63%	AiPPT.cn	0.09

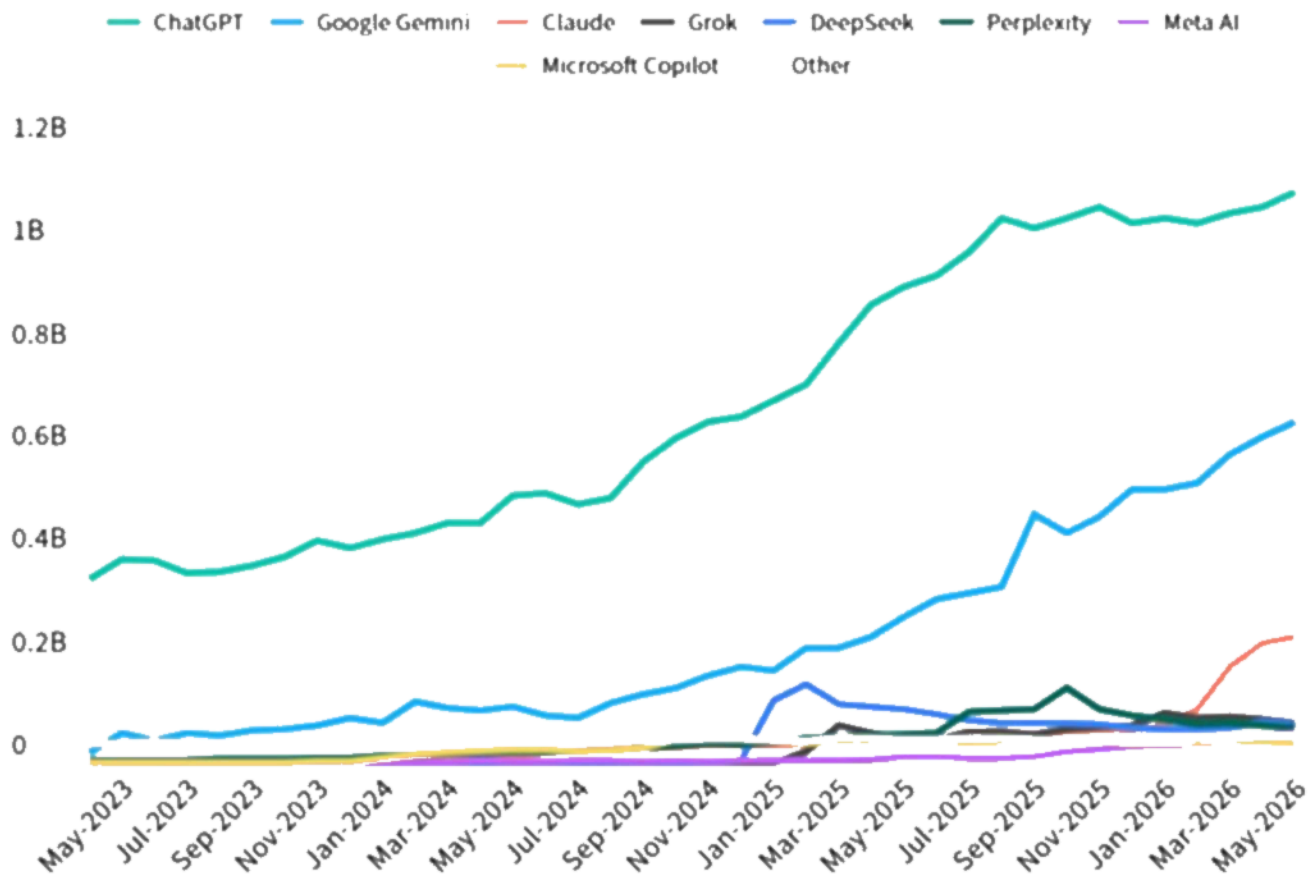
Source: Xsignal, Nomura research

Development of major large language model applications in China and global markets

ChatGPT continues to hold a commanding lead in the global market, though several competitors have narrowed the gap in recent months. Google Gemini has steadily expanded its users, supported in part by a strong Android user base and deep integrations across Google's ecosystem. Meanwhile, Claude has experienced rapid growth since the start of 2026. ChatGPT's share of the AI assistant market, based on MAU, has fallen below 50% since March 2026. By May 2026, ChatGPT held a 46% share, ahead of Google Gemini (28%) and Claude (10%). Claude had seen average monthly revenue per user (ARPU) rising from less than USD0.50 in September 2025 to USD2.76 in May 2026, with its focus on research and professional workflows.

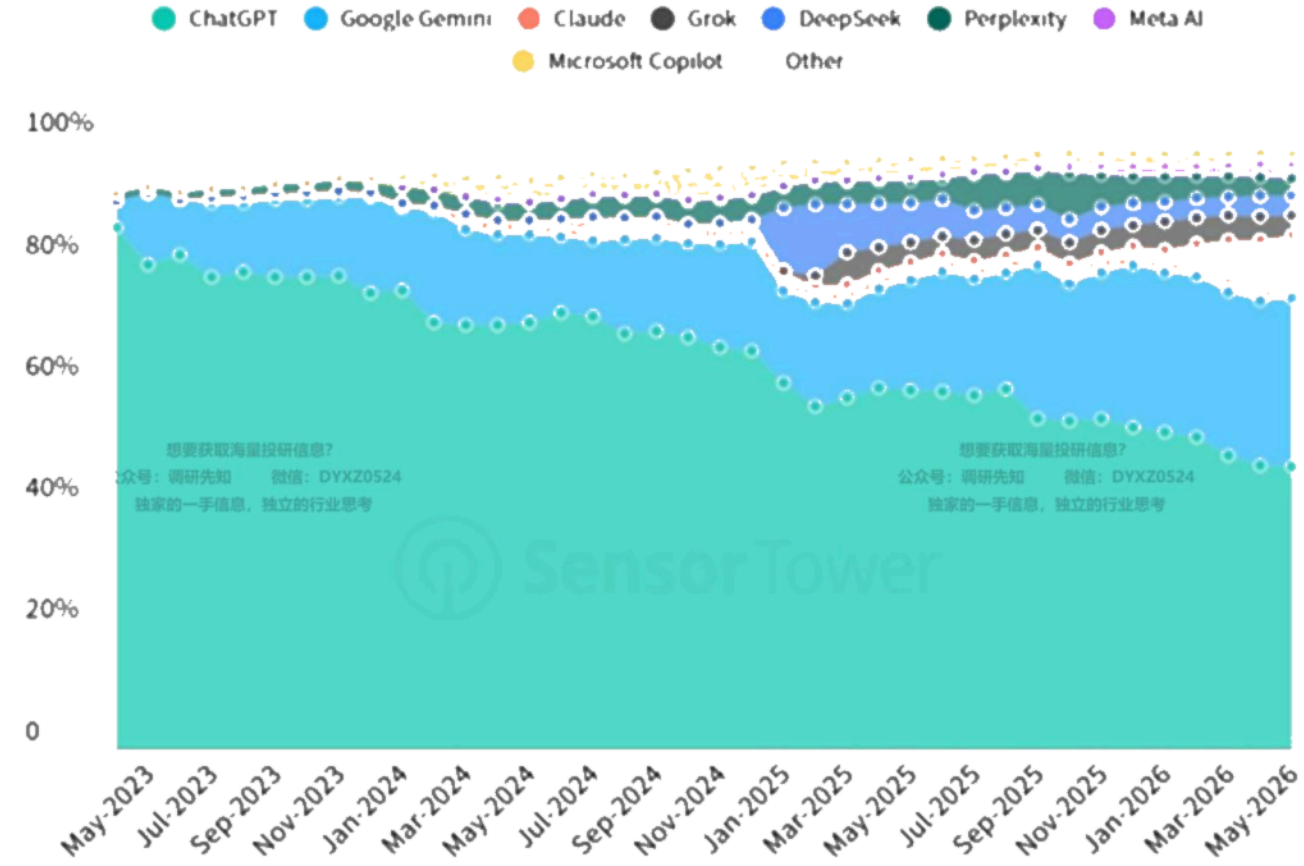
Although the monetization progress of AI models in China is slower than in overseas markets, it has already accelerated compared to 2025. Doubao has launched a paid subscription plan, with monthly pricing ranging from CNY68 to CNY500 in three tiers, catering to different user groups. Moreover, in 1Q 2026, Zhipu GLM API prices were raised by about 83%, while overseas subscription prices approached Claude Code and API call still grew by about 400%, reflecting accelerated commercialization progress, in our view.

Fig. 25: MAU across web and mobile app for top global AI assistants



Source: Sensor Tower, Nomura research

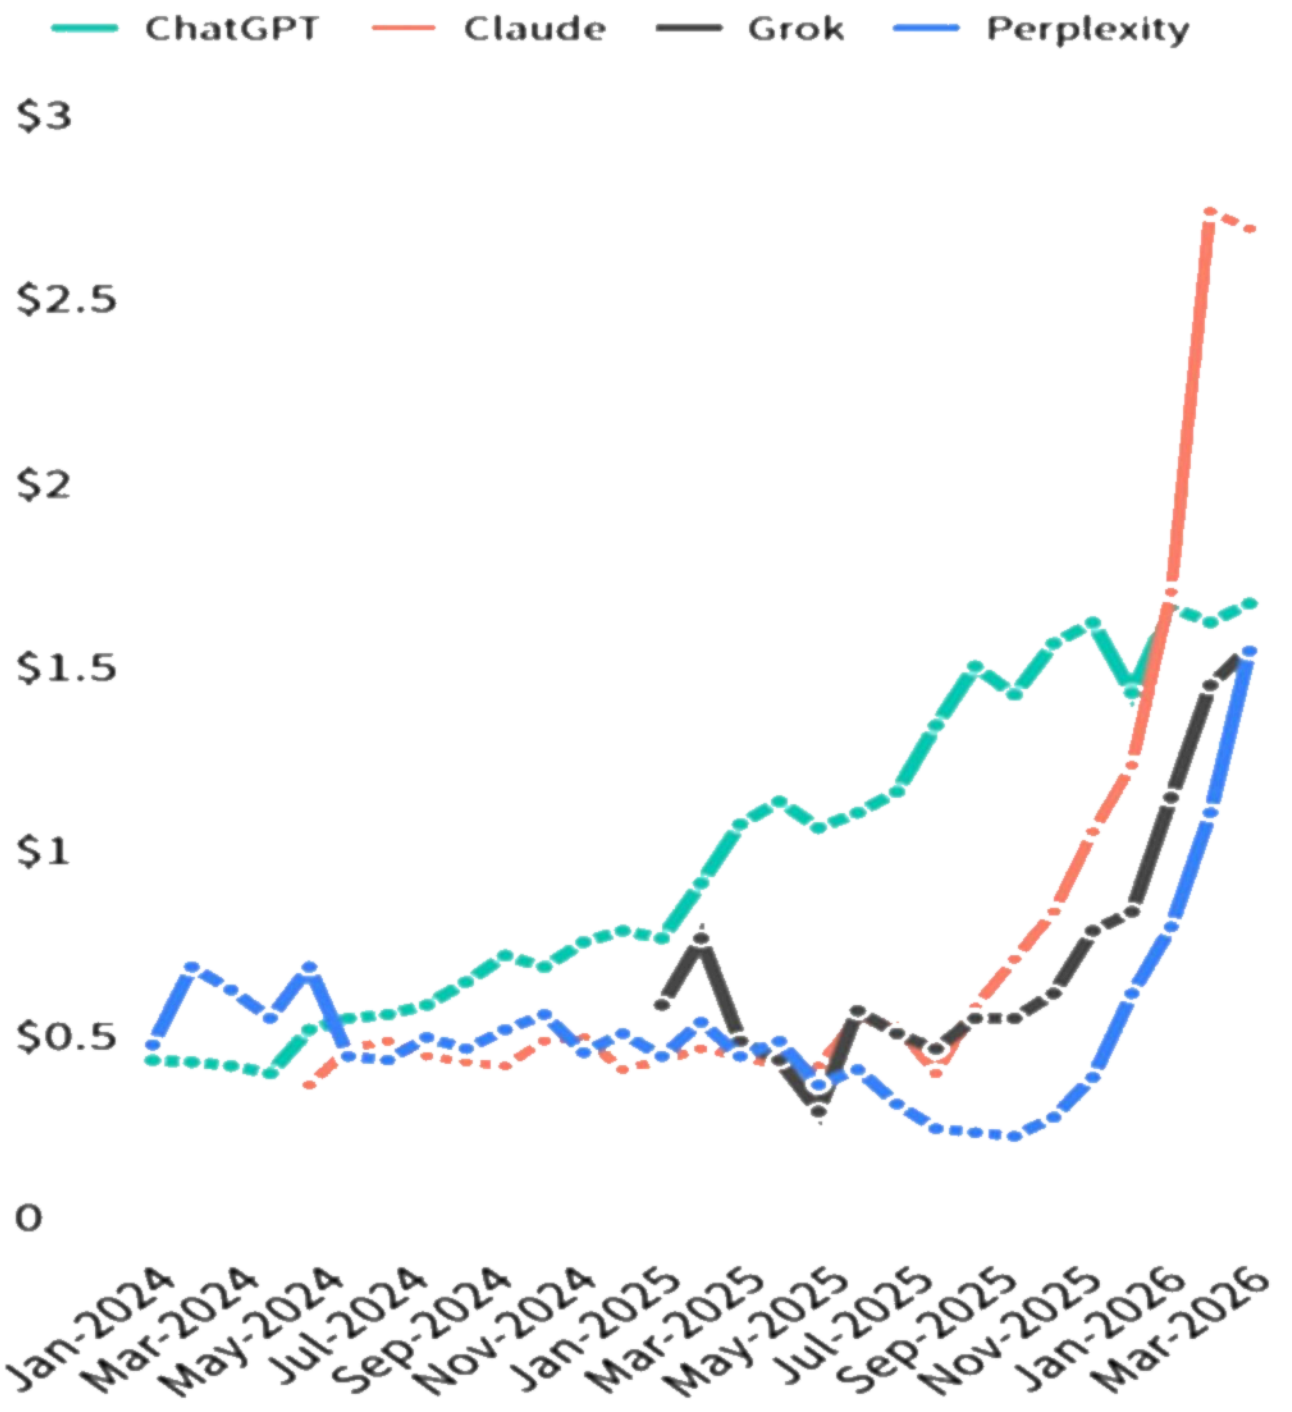
Fig. 26: Market share of top AI assistants by MAU across web and mobile app



想要获取海量投研信息?
公众号：调研先知 微信：DYXZ0524
独家的一手信息，独立的行业思考
SensorTower

Source: Sensor Tower, Nomura research

Fig. 27: Average Monthly Revenue per User (ARPU) across iOS and Google Play for select global AI assistant apps



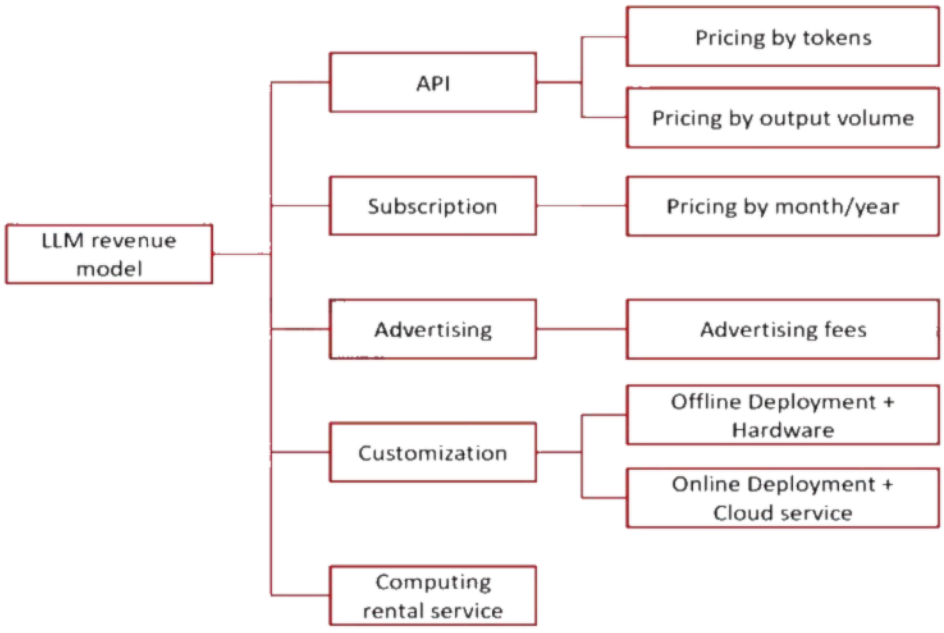
Source: Sensor Tower, Nomura research

Comparison of business models among LLM players

We think LLM companies have several types of strategies for monetization, including computing power rental services, API fees, subscription / membership fees, free usage + advertising fees, and project customization. First, LLM companies could rent out their computing power (GPU / ASICs) to smaller AI / LLM startups or enterprise customers who want to train their own models. Secondly, LLM companies can provide model-as-a-service (MaaS) solutions, similar to the SaaS (Software as a Service) business model of the software industry, and they usually charge for the API calling fees, based on the volume of input and output tokens (pay-as-you-go), or based on a fixed subscription plan. Consumers and enterprise customers can also use LLMs for free (which may attract advertisers), or subscribe to a membership plan, for accessing the functions of the commercial LLMs. In addition, large enterprise and government customers can request

LLM companies to build customized models or provide customized LLM solutions for them.

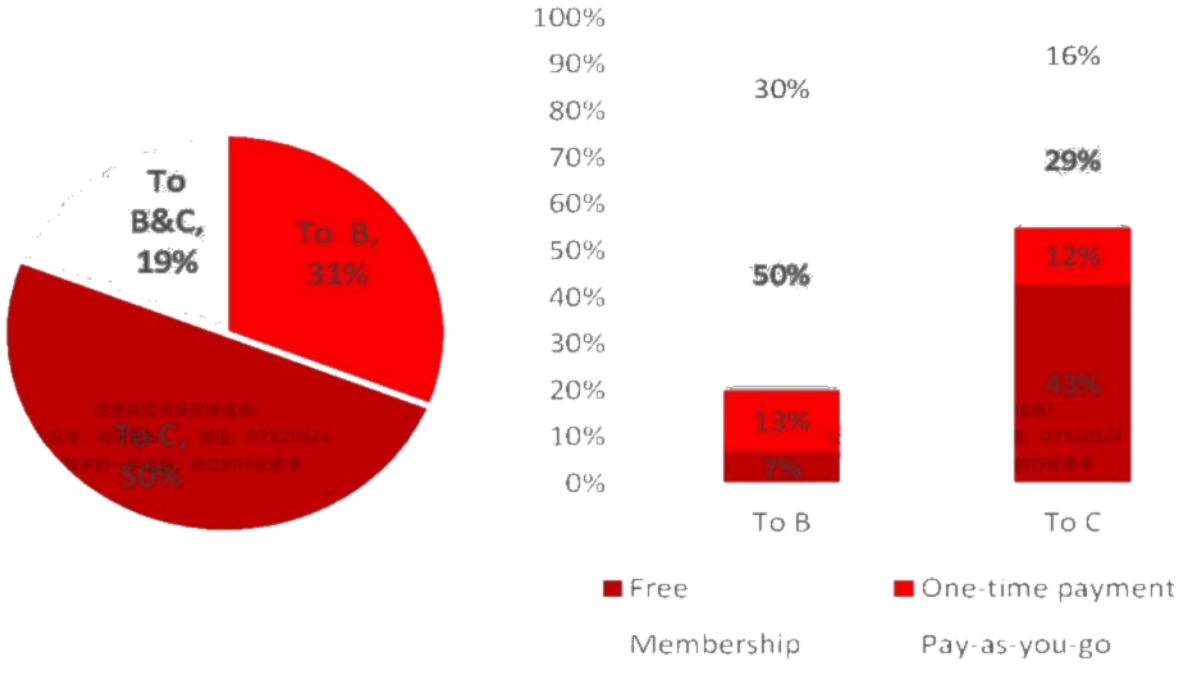
Fig. 28: LLM company revenue model



Source: Nomura research

In terms of commercialization progress of AI products, the monetization model in the enterprise market looks clearer than in the consumer market, as a larger portion for consumer AI products are still offered for free. For enterprise AI products, the most common revenue models are the membership fee (50%) and pay-as-you-go model (30%), according to Quantumbit. In the consumer AI market, over 40% of the products currently still do not have a clear revenue model, which offer free usage.

Fig. 29: User structure and proportion of revenue type of AI applications

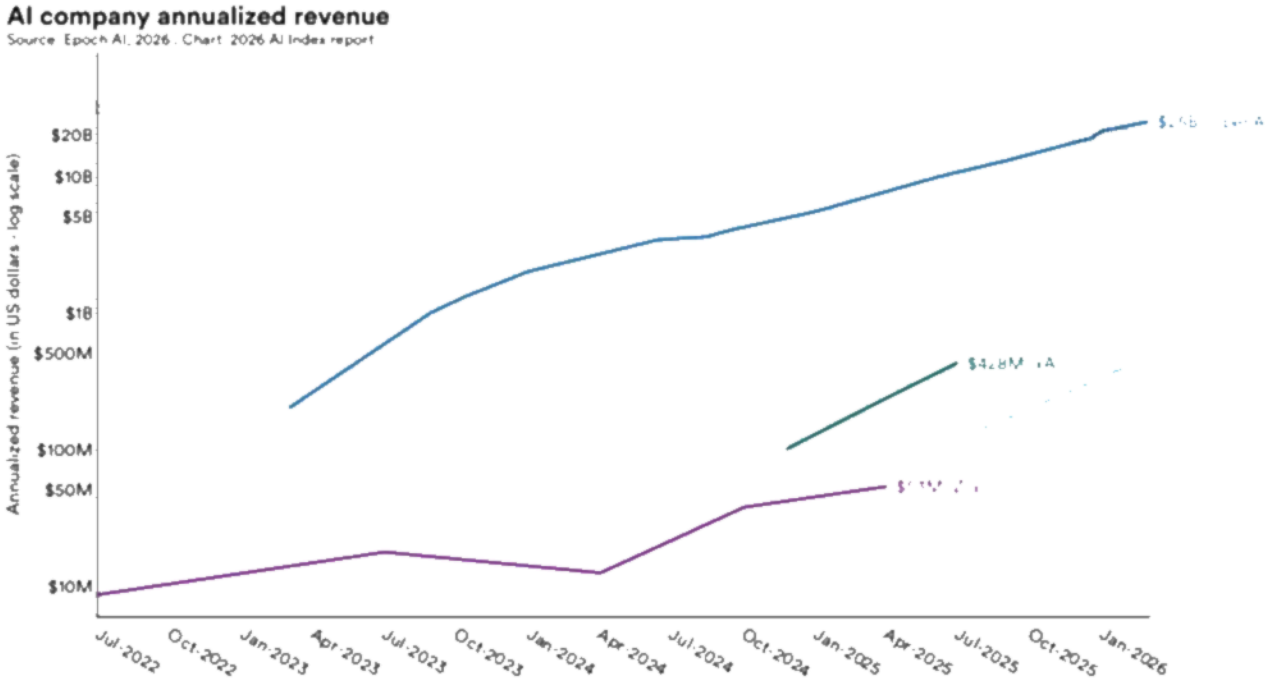


Source: "China AIGC Application Overview Report 2024" by QuantumBit AI, Nomura research

Progress of commercialization of Global and China LLM players

In terms of business model for LLMs, OpenAI has already established several revenue models, which mainly include ChatGPT subscription, API call fees, and customization fees (strategic cooperation). According to The Information., OpenAI topped USD25bn in annualized revenue as of the end of February 2026. ChatGPT's revenue growth rate significantly outpaced user growth during the same period, indicating that the growth logic has shifted from purely "traffic scale-driven" to "paid penetration-driven".

Fig. 30: Global LLM company annualized revenue



Source: AI Index Report 2026, Nomura research

For another key LLM player Anthropic, the company announced that its annualized revenue (ARR) exceeded USD30bn in April 2026, a significant increase from USD9bn at the end of 2025, and Anthropic also guides 2Q2026 revenue to more than double to USD10.9bn. OpenAI leverages ChatGPT's massive user base to hold an advantage in the consumer market, but individual users have less purchasing power than enterprise users, while Anthropic has deeply cultivated the enterprise market, building a high-margin moat with most of revenue coming from enterprise customers. We believe that leading overseas LLMs players have all achieved closed-loop commercialization, with the progress faster in the enterprise market than in the consumer market, making it a strong contributor to profitability for LLM players, in our view.

Fig. 31: Commercialization performance of overseas AI companies

Company	AI product	AI business model	Commercialization performance
OpenAI	ChatGPT	Subscription	50mn+ paid subscribers; OpenAI generated USD5.7bn in 1Q26, according to The Information
	GPT models	API; Enterprise projects	
	Codex	Subscription	5mn+ weekly active users, up more than 6x since the launch of the desktop app in February 2026, according to OpenAI
Anthropic	Claude app	Subscription	Anthropic topped USD30bn in ARR in as of April 2026, according to the company
	Claude models	API; Enterprise solutions	
	Claude Code	Subscription	
Microsoft	Copilot	Subscription	30mn+ Microsoft 365 Copilot paid seats; 50mn+ Github Copilot users and Copilot revenue accelerated over 60% q-q, according to 4QFY26 earnings call
Google	AI search	Free	-
	Gemini	Subscription; API	Google's model APIs are now processing ~22 billion tokens per minute, up from 16 billion a quarter ago. Over the last 12 months, more than 2,000 enterprises consumed over 100 billion tokens, according to 2Q26; Gemini app has 950 million MAUs as of 2Q26.
Amazon	Amazon Q	Subscription	AI revenue run-rate is over USD25bn, growing triple-digit percentage y-y, according to 2Q26 earnings call
	Nova	API; Enterprise projects	
Meta	Meta AI	Free	Meta is developing new personal agents that will be priced with several tiers
	Llama/Muse Spark	API; Enterprise projects	
	GenAI tools	-	9 million small businesses on platforms are now using at least one of Meta's AI ad creative tools, as of 2Q26; Advantage+ end-to-end solutions reached over USD75bn ARR as of 2Q26
	AI glass	Sale	Reality Labs segment revenue grew 16% y-y due to strong growth in AI Glasses revenue

Source: Company data, Nomura research

In China, with the launch of many state-of-the-art LLMs in 2026, the commercialization progress is set to accelerate, in our view. However, we think China's LLM companies' commercialization progress is still behind overseas' peers as most of China's AI applications are free-of-charge, or charge very low fees for consumer and enterprise customers. We think B-end AI applications also grow faster than C-end in China, with better commercialization progress. For example, Zhipu AI, primarily focused on enterprise business, achieved a revenue of CNY724mn in 2025, a y-y increase of 131.9%. In 1Q26, Zhipu API call pricing increased by 83%, with demand still outstripping supply, and call volume grew by 400%.

Fig. 32: Commercialization performance of China AI companies

Company	AI product	AI business model	Commercialization performance
Baidu	Ernie bot	Free	-
	Ernie LLM	API; Enterprise projects	
	AI marketing agents and digital human		Revenue from AI-native marketing services reached CNY2.3bn in 1Q26, up 36% y-y.
	AI applications (Baidu Wenku, Baidu Drive, Digital Employee, etc.)	Subscription	Revenue from AI Applications was CNY2.5bn in 1Q26, approximately flat y-y.
Bytedance	Doubao	Free + Subscription	Standard: CNY68/month; Enhanced: CNY200/month; Pro: CNY500/month
	Doubao LLM	API; Enterprise projects	As of June 2026, the Doubao models' daily average token calls reached 180 trillion, a more than tenfold increase over the past year
	Dreamina AI	Free + Subscription	Basic membership: CNY79/month; Standard: CNY239/month; Premium: CNY649/month
Alibaba	Tongyi Qianwen	Free	Alibaba Cloud's AI business ARR has exceeded CNY35.8bn, continuing to maintain triple-digit growth. AI-related product revenue accounted for 30% of Alibaba Cloud's external commercialization revenue, according to 4QFY26 earnings call.
	Qwen LLM	API; Enterprise projects	
Tencent	Hunyuan LLM	API; Enterprise projects	The number of daily API calls exceeds 300 million in Jun-24. Benefiting from the application of AI technology, the revenue of the marketing services segment in 3Q24 was CNY30 bn with +17% y-y
	Yuanbao	Free	-
	Workbuddy	Free+Subscription	Standard:CNY99/month, Premium: CNY199/month, Flagship: CNY999/month
SenseTime	Sense Chat	Free	-
	Sense Nova	API; Enterprise projects	-
	Raccoon	Free + Subscription	Free; CNY49 for Premium; CNY899 for Enterprise
Zhipu AI	Zhipu Qingyan	Free	-
	Zcode	Free + coding plan	Lite:CNY49/month, Pro: CNY149/month, Max: CNY469/month
	Autoclaw	Free	-
iFlytek	Spark	Free	-
	Spark LLM	API; Enterprise projects	In the medical field, the imaging cloud based on the Spark Medical large model generated revenue of CNY100mn+ in 1H26, +158% y-y.
	Smart hardware	Sale	Spark reading machine sales increased by 1471% y-y in 1H26
Moonshot	Kimi app and Kimi Work	Free + Subscription	As of June 2026, the company's ARR surpassed USD300mn, according to Financial News; The K3's output price increased to CNY100/million tokens, about 3.7x higher than the previous model.
	Kimi LLMs	API; Enterprise projects	

Source: Company data, Nomura research

From the perspective of LLM's pricing, we note that Chinese LLMs shift towards the high end with a higher pricing, driven by better performance, while global LLM players with closed source LLMs have started to reduce API prices to compete with low-cost open-source models. For example, Zhipu simultaneously announced a 30% overall price increase for Coding Plan when GLM-5 was released in February 2026 and as of 1Q26, Zhipu API price increased 83%. Similarly, the Kimi K3 API is priced at CNY20/mn tokens for input and CNY100/mn tokens for output, which is 3.1x and 3.7x higher than Kimi K2.6, respectively. Deepseek also changed its pricing strategy, with the use of peak-valley pricing. The output price of V4-Pro has risen from CNY6 per mn tokens to CNY27 and V4-Flash output price increased from CNY2 to CNY9 during peak periods, while it would be at a 50% discount during non-peak time. On the other hand, OpenAI announced in July 2026 that the API price of GPT-5.6 Luna dropped up to 80% and GPT-5.6 Terra dropped by 20%. We think Chinese LLM players are accelerating LLM monetization, thanks to the breakthrough in LLM capabilities. Moreover, given that coding plans are usually sold out in a short span of time, we think consumers are willing to accept price hikes with better LLM performance, which forms a positive feedback loop and could benefit the Chinese LLM market with a healthier competitive environment, in our view. In the global market, we think Chinese open-source LLMs with a lower pricing bring price pressure to overseas top closed source LLMs, which force them to adjust the pricing strategy to maintain global market share.

Fig. 33: API price of overseas major LLMs, as of September 2026

Company	Model	Input price (USD/mn token)	Output price (USD/mn token)
OpenAI	GPT-6 Astra	$10.00	$50.00
	GPT-5.6 Solar	$5.00	$30.00
	GPT-5.6 Terra	$2.00	$12.00
	GPT-5.6 Luna	$0.20	$1.20
	GPT-5.5	$5.00	$30.00
Anthropic	Claude Fable 5.1/Mythos 5.1	$10.00	$50.00
	Claude Opus 4.8	$5.00	$25.00
	Claude Sonnet 4.6	$3.00	$15.00
	Claude Haiku 4.5	$1.00	$5.00
Google	Gemini 3.5 flash	$1.50	$9.00
	Gemini 3.1 pro	$2.00	$12.00
	Gemini 3.1 flash	$0.50	$3.00
Meta	Muse Spark 1.3	$1.25	$4.25
	Llama 4-Maverick	$0.24	$0.77
	Llama 4-Scout	$0.15	$0.40

Source: Company data, Nomura research

Fig. 34: API price of China major LLMs, as of September 2026

Company	Model	Input price (CNY/mn token)	Output price (CNY/mn token)
Alibaba	Qwen3.6-Max-Preview	CNY 9.00	CNY 54.00
	Qwen3.6-Plus	CNY 2.00	CNY 12.00
	Qwen3.6-Flash	CNY 1.20	CNY 7.20
Baidu	Ernie 5.0 pro	CNY 1.50	CNY 6.00
	Ernie 5.0	CNY 0.80	CNY 3.20
Tencent	Hunyuan 2.0-thinking	CNY 1.00	CNY 4.00
	Hunyuan 2.0-instruct	CNY 0.80	CNY 2.00
ByteDance	Doubao-Seed-2.1-Pro	CNY 6.00	CNY 30.00
	Doubao-Seed-2.1-Turbo	CNY 3.00	CNY 15.00
Zhipu	GLM 5.2	CNY 8.00	CNY 28.00
	GLM 5.1	CNY 6.00	CNY 24.00
	GLM 5.0	CNY 1.00	CNY 3.20
Deepseek	DS-v4-pro	CNY3.00	CNY6 (peak period: CNY27)
	DS-v4-flash	CNY 1.00	CNY2 (peak period: CNY9)
Moonshot	K3	CNY 20.00	CNY 100.00
	K2.7 Code	CNY 6.50	CNY 27.00
	K2.6	CNY 6.50	CNY 27.00
MiniMax	M3	CNY 2.10	CNY 8.40
	M2.7	CNY 2.10	CNY 8.40

Source: Company data, Nomura research

AI + Finance

Previously, AI use cases in finance were mostly in process automation, as a way to enable efficiency and improve productivity at the back-office (operations) and middle-office (compliance and risk management) of financial service providers. We expect Gen AI technologies to support production and deployment of additional front-end use cases and applications and to further accelerate the existing use of AI tools at the operational level such as risk management, fraud detection, personalized financial product recommendations, robo-advisory, and automated customer service.

AI + Finance market TAM analysis

According to Statista, the global AI spending in financial services industry is expected to grow from USD75bn in 2026E to USD125bn by 2028E, representing a 29% CAGR. The banking sector, in particular, has allocated USD53bn in 2026 for AI and Gen-AI investments, and will grow to USD90bn+ by 2028E, with a two-year CAGR of 30%, according to Statista.

In the China market, led by DeepSeek's powerful semantic understanding and logical reasoning, an increasing number of domestic banks, securities firms, and insurance institutions have connected to or deployed AI into their operations and businesses. Financial industry users are showing a relatively positive approach towards AI investment. According to IDC, the Gen-AI investment in China's financial industry was about CNY9bn in 2025, and the market size is expected to climb to CNY45bn by 2029E, with Agent being the key AI product form.

From an application scenario perspective, AI is currently penetrating various aspects of the financial industry, including intelligent customer service, intelligent marketing, internal operations, intelligent investment advisory/wealth management, intelligent investment research, intelligent risk control, and internal R&D. However, in terms of application level, according to IDC's survey on the application of AI in different scenarios by financial institutions, apart from the relatively in-depth application and mature function implementation of AI in customer service scenarios, the application of AI in other scenarios, especially core business scenarios, is still mainly concentrated at a relatively shallow function level.

AI + Finance value chain

Currently, there are many AI startups working on gen AI products including financial information search, financial analysis, risk management, for banking, insurance and investment management clients.

Fig. 35: Select financial LLM applications

Company	Financial LLM	Features
Bloomberg	BloombergGPT	Relying on Bloomberg's extensive financial data sources, it has built a dataset of 363 billion tags, supporting various tasks within the financial industry, with outstanding performance in financial Q&A, sentiment analysis, and text classification.
SenseTime	Finance LLM	It is leading in multimodal interaction, supporting "document + image + voice" full-modal financial scenarios.
iFlytek	Spark Finance	Customer service models supporting dialect recognition; Multi-turn dialogue technology in the education sector is reapplied to financial advisory scenarios.
Industrial and Commercial Bank of China	ICBC AI model	Self-developed models covering core scenarios such as transaction automation and credit risk control.

Source: Company data, Nomura research

Fig. 36: Value chain of Gen-AI applications in finance

Intelligent finance				Business empowerment	
Smart investment research	Robo-advisory	Smart claims	Smart risk control	Sales & Marketing	Customer services
Datayes (Unlisted) Ant Group (Unlisted) Duxiaoman (Unlisted) iFinD (300033 CH) HundSun (600570 CH) CICC (601995 CH) Rabyte Technologies (Unlisted) Value Simplex (Unlisted) Bloomberg (Unlisted) Portrait Analytics (Unlisted)	iFinD (300033 CH) CreditEase (YRD US) JF Smart Invest (9636 HK) China Galaxy Securities (6881 HK) CITIC Securities (6030 HK) Nutmeg (Unlisted) Wealthfront (Unlisted) Betterment (Unlisted) Schwab Intelligent Portfolios (Unlisted)	CPIC (601601 CH/2601 HK) Taikang (Unlisted) PingAn (2318 HK) China Life (2628 HK) Huize (HUIZ US) Zhong An insurance (6060 HK) Mayibao (Unlisted) CLARA Analytics (Unlisted) Shift Technology (Unlisted)	4Paradigm (6682 HK) Huawei (Unlisted) GoUpSec (Unlisted) MEGVII (Unlisted) Inspur (000977 CH) DataCanvas (Unlisted) AHI (Unlisted) DataVisor (Unlisted) OpenGamma (Unlisted) Credit Benchmark (Unlisted) Sino vancco (Unlisted)	OceanEngine (Unlisted) Mmoser (Unlisted) Convertlab (Unlisted) TalkingData (Unlisted) Marketingforce (2556 HK)	Tencent (700 HK) iFlytek (002230 CH) T&I Net Communication (835653 CH) Zhichi Tech (Unlisted) Active.AI (Unlisted) Glia (Unlisted) Boost.ai (Unlisted)

Financial large model	Agricultural Bank of China (601288 CH), Industrial and Commercial Bank of China (601398 CH), Huawei (Unlisted), 360 (601360 CH), HundSun (600570 CH), Duxiaoman (Unlisted), Sino vancco (Unlisted), Ping An Bank (000001 CH), PICC (601319 CH), Baidu (9888 HK/BIDU US), iFinD (300033 CH), Bloomberg (Unlisted), Ant Group (Unlisted)
Solution provider for finance	Alibaba (BABA US), Tencent (700 HK), Baidu (9888 HK/BIDU US), AWS (AMZN US), CloudWalk (688327 CH), ICBC Tech (Unlisted, subsidiary of ICBC, 601398 CH), GienTech (Unlisted), Yusys Technologies (300674 CH), iFlytek (002230 CH), Yonyou Networks (600588 CH)

Source: Nomura research

AI + Software

AI + Software market TAM

AI software is defined as a software that enables intelligent behavior. It is widely used in industries such as finance, manufacturing, retail, and the public sector, and is a key technological support for enterprises' digital transformation. AI software spending will grow to USD297.9 bn by 2027E, according to Gartner. Market growth is expected to accelerate from 17.8% in 2022 to 20.4% in 2027E, with a 19.1% CAGR. Government sector has the largest spend of over USD70bn by 2027, while oil and gas is growing fastest with a 25.2% CAGR.

Fig. 37: AI software forecasts across vertical industries (USD mn), worldwide, 2022-2027E

Vertical market (USD mn)	2022	2023	2024	2025	2026	2027	CAGR 2022-2027 (%)
Banking and investment services	$22,269	$26,349	$31,311	$37,505	$45,379	$55,161	20%
		18.3%	18.8%	19.8%	21.0%	21.6%	
Communication, media and services	$16,873	$19,893	$23,514	$27,961	$33,515	$40,310	19%
		17.9%	18.2%	18.9%	19.9%	20.3%	
Education	$5,678	$6,622	$7,766	$9,197	$11,020	$13,292	19%
		16.6%	17.3%	18.4%	19.8%	20.6%	
Government	$29,972	$35,225	$41,578	$49,375	$59,016	$70,666	19%
		17.5%	18.0%	18.8%	19.5%	19.7%	
Healthcare and life science	$8,768	$10,187	$11,857	$13,837	$16,182	$18,894	17%
		16.2%	16.4%	16.7%	16.9%	16.8%	
Insurance	$6,929	$8,109	$9,523	$11,251	$13,390	$15,983	18%
		17.0%	17.4%	18.1%	19.0%	19.4%	
Manufacturing	$13,857	$16,452	$19,633	$23,571	$28,488	$34,495	20%
		18.7%	19.3%	20.1%	20.9%	21.1%	
Oil and gas	$953	$1,177	$1,464	$1,834	$2,315	$2,929	25%
		23.5%	24.4%	25.3%	26.2%	26.5%	
Power and utilities	$6,836	$8,160	$9,788	$11,834	$14,458	$17,763	21%
		19.4%	20.0%	20.9%	22.2%	22.9%	
Retail	$5,818	$6,711	$7,769	$9,046	$10,609	$12,484	17%
		15.3%	15.8%	16.4%	17.3%	17.7%	
Transportation	$4,068	$4,864	$5,836	$7,051	$8,592	$10,509	21%
		19.6%	20.0%	20.8%	21.9%	22.3%	
Wholesale trade	$2,296	$2,688	$3,161	$3,747	$4,494	$5,429	19%
		17.1%	17.6%	18.5%	19.9%	20.8%	
All industries	$124,316	$146,437	$173,199	$206,207	$247,459	$297,914	19%
		17.8%	18.3%	19.1%	20.0%	20.4%	

Source: Gartner, Nomura research

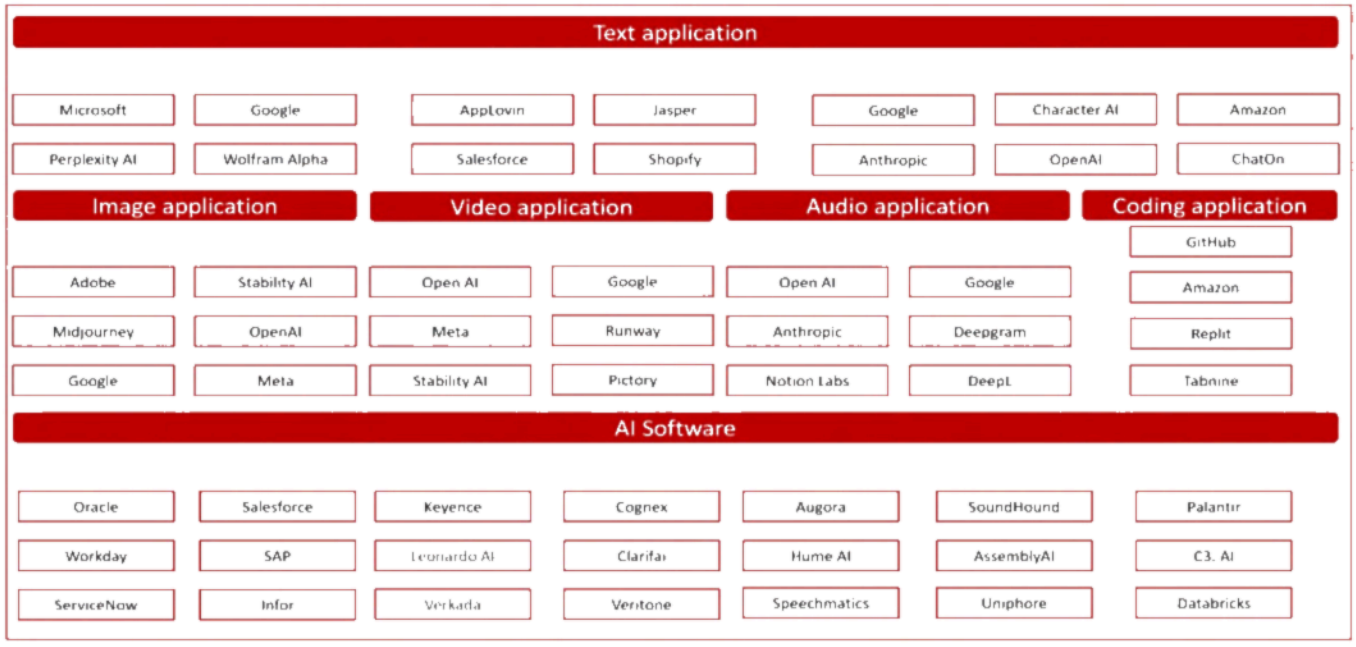
AI + Software use cases

- AI-Driven Code Review: Gen AI-powered code review tools are streamlining the traditionally cumbersome and time-consuming process. These tools analyze code comprehensively, identifying bugs and security vulnerabilities while providing developers with valuable insights and improvement suggestions.
- AI-Powered Chatbots and Virtual Assistants: AI-powered chatbots and virtual assistants are becoming increasingly popular in CRM systems. These AI tools can handle customer inquiries, provide instant responses, and assist with common tasks,

improving customer satisfaction and reducing response times.

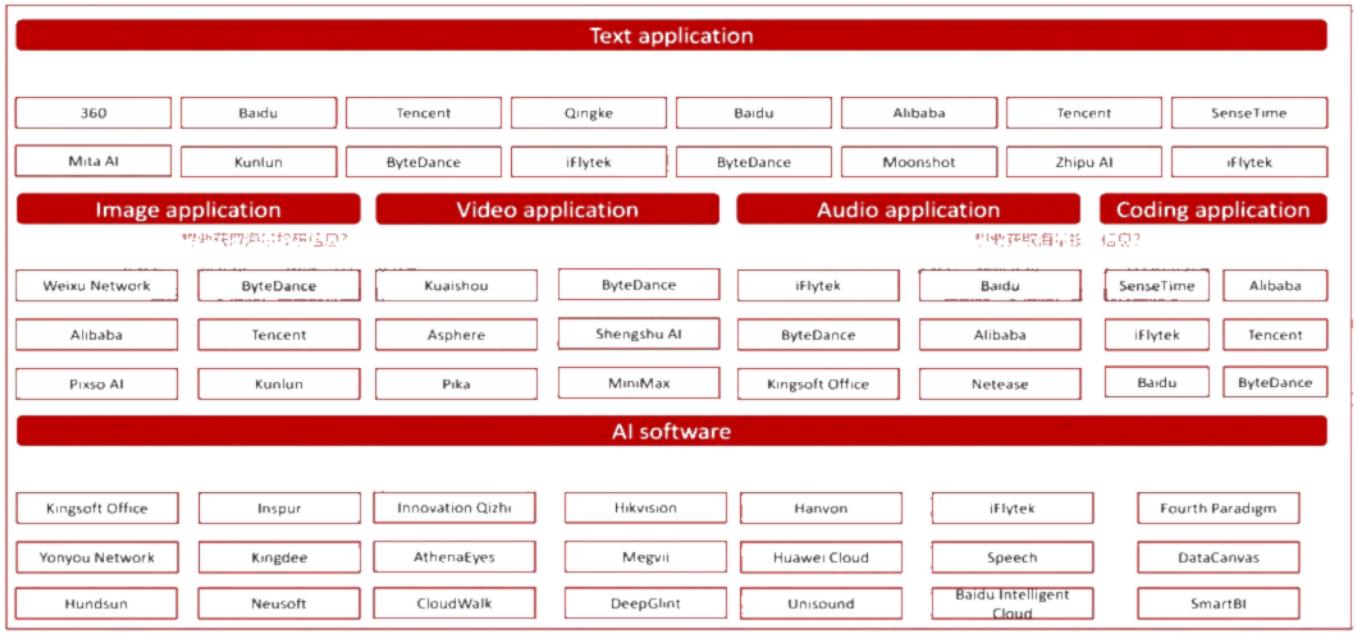
- Enhanced Diagnostics for Healthcare with AI: AI diagnosis tools are trained on massive datasets of medical images, enabling them to identify subtle patterns and anomalies that might be missed by a human radiologist.
- Data and Learning Analytics for Education: AI helps analyze data from online learning portals, classroom attendance and grades. This data provides insights into student performance, helping educators identify trends and tailor instruction to address gaps in understanding and performance.
- Bolstering Research and Development Activities: Generative AI fuels predictive modeling and helps create potential outcomes and trends, enabling R&D teams to make informed decisions about product development; it also accelerates the overall pace of research by automating routine and time-consuming tasks.

Fig. 38: Global software & application company landscape



Source: Nomura research

Fig. 39: China software & application company landscape



Source: Nomura research

AI + Autonomous driving

Software solution, remains the ultimate puzzle for the industry

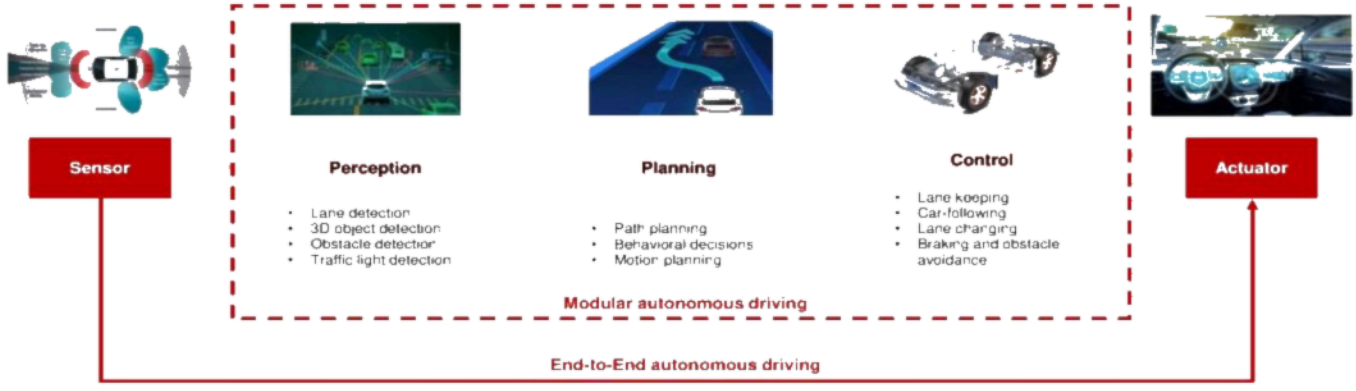
Transitioning from rule-based to End-to-End solution

Traditional autonomous driving solutions follow a sequential modular pipeline, including sensors, perception, prediction and rule-based planning. While this approach offers modularity, interpretability and structured debugging, it also introduces significant drawbacks, which are information loss, error accumulation and scalability shortage.

End-to-end (E2E) autonomous driving is an approach that uses a single, integrated deep learning system to translate raw sensor data directly into vehicle control signals, bypassing traditional modular pipelines that separate tasks like perception, prediction, and planning.

Compared to rule-based solution, end-to-end system eliminates the need for multiple interconnected modules, potentially reducing complexity and latency. More importantly, end-to-end systems become highly adaptable by learning from vast datasets, potentially handling complex scenarios that are challenging for rule-based designs. As OEMs focus on the scalability of their smart driving solutions, end-to-end become the mainstream tech-route to deal with autonomous driving in complex urban scenarios nowadays in the passenger vehicle market.

Fig. 40: Modularized vs end-to-end autonomous driving solutions



Source: Nomura research

Over the past few years, we see a gradual improving process for OEMs moving from rule-based to end-to-end systems. Generally, we believe there are four stages in this evolution trend.

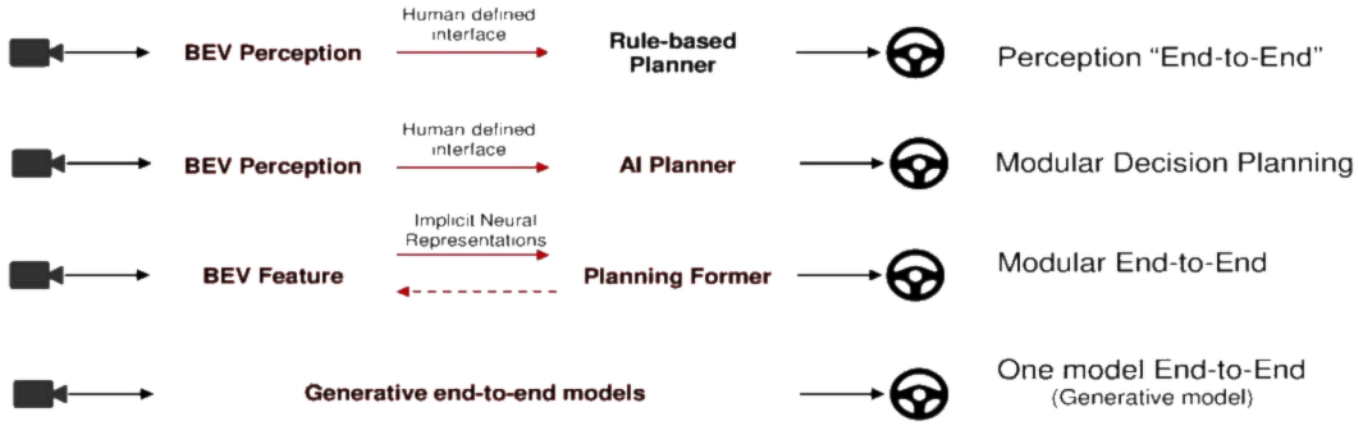
Stage 1 can be taken as end-to-end perception. With the perception module achieving sensor fusion in the BEV (Bird's Eye view) system by adopting transformer and cross attention mechanism, the accuracy and stability of detection results have been greatly improved. Meanwhile, at this stage, the prediction, decision and planning modules are still rule-based.

Stage 2 is called modular decision and planning. While the perception module remains, the prediction, decision and planning functions are integrated into one neural network. At the stage 2, both perception and planning modules are based on deep learning but needed to be trained separately. And the interface between the two is still defined by human interpretation.

Stage 3 is modular end-to-end. The perception and planning system can be jointly trained in an end-to-end manner towards the ultimate task ensuring minimum information loss, while they are still two modules but communicating together. Instead of human-defined representations of the world, the perception system block outputs feature representations and support back-propagation.

Stage 4 is the true one model end-to-end system. In the final stage, perception and planning are eventually combined into one deep learning network. The model takes raw sensor input to generate ego-motion plans.

Fig. 41: Technology evolution of end-to-end solution



Source: Cherish Capital, Nomura research

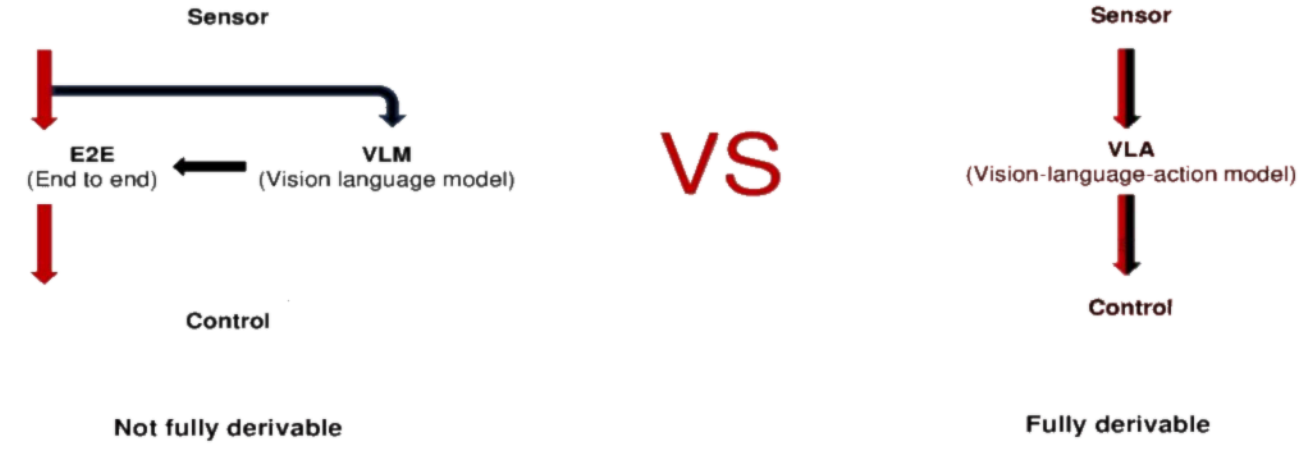
Other development trend within E2E era

At current stage, we believe active smart driving players in the market are already at the stage 3 or 4 with the evolution, while some players may still use rule-based solutions as a safety redundancy. Considering an aggressive development approach within the market, we also see new buzzwords appearing in the industry starting from 2025, including VLM, VLA and World model.

VLM refers to the use of Vision Language Model to enhance self-driving systems by integrating visual perception with language-based reasoning, allowing a vehicle to understand complex scenes and make decisions in a more human-like way. Combining VLM with E2E system, that is what companies called a "fast-slow" system. Currently we see a clear trend of players previously using VLM now further upgrading into VLA.

VLA, which is called Vision Language Action, is a fusion of E2E and VLM methodologies. While VLMs focus primarily on environmental modeling, VLAs extend this focus to include planning and control, marking the key difference between the two. Theoretically, a VLA-equipped autonomous vehicle can read road signs through text and other inputs, determine from multiple sources whether the reversible lane is usable, and interact with other vehicles via turn signals. As a further enhancement from VLM system, VLA has usually been taken as "end-to-end 2.0". Based on the definition of VLA system, we take it as the Stage 4 of E2E development.

Fig. 42: E2E evolution— VLM vs VLA



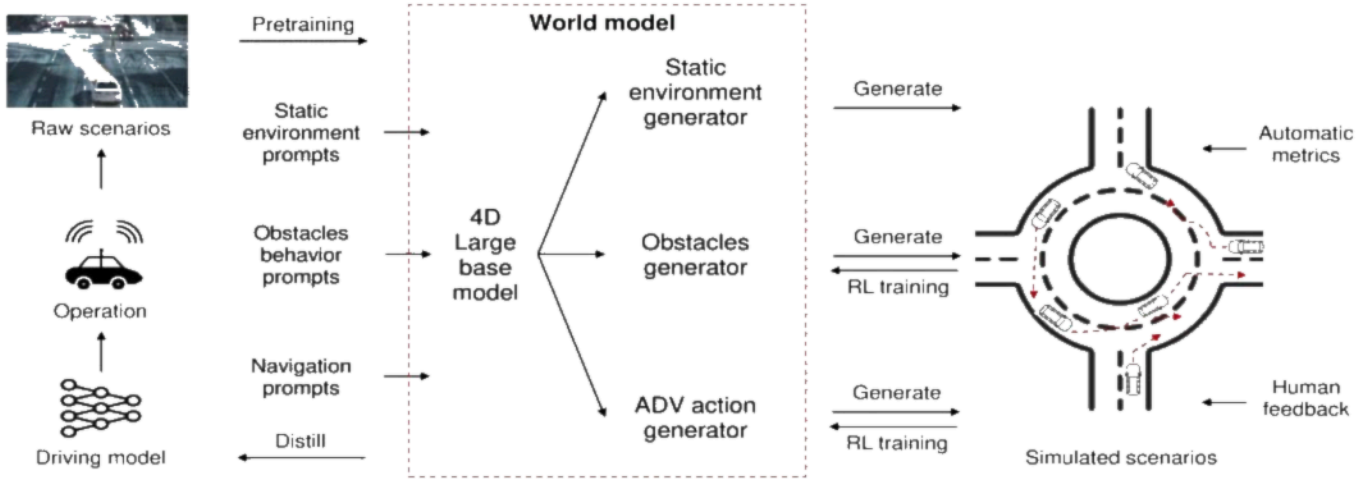
Source: Deeproute.AI, Nomura research

With the fast development in the past few years, we understand that players can use two distinct approaches to train the models, which are imitation learning and reinforcement learning. While all the above-mentioned E2E solutions are mainly based on imitation learnings, solutions focusing on reinforcement learnings are also been brought up by robotaxi companies and leading smart driving OEMs in the market, one of which is World

model learning. And World model learning now shows a clear direction that can be further explored, in our view, as the majority of market players quickly embrace such methodology.

World model learning for autonomous driving involves building internal predictive models that allow a vehicle to imagine and anticipate future scenarios based on sensor data, enhancing safety and decision-making by simulating outcomes without real-world interaction. These models learn how the environment evolves, enabling "what if" simulations for planning and training in virtual environments that can replicate real-world complexity, which is fundamentally important for autonomous driving. We recognize world model learning is also within the Stage 4 of E2E development.

Fig. 43: World model for autonomous driving

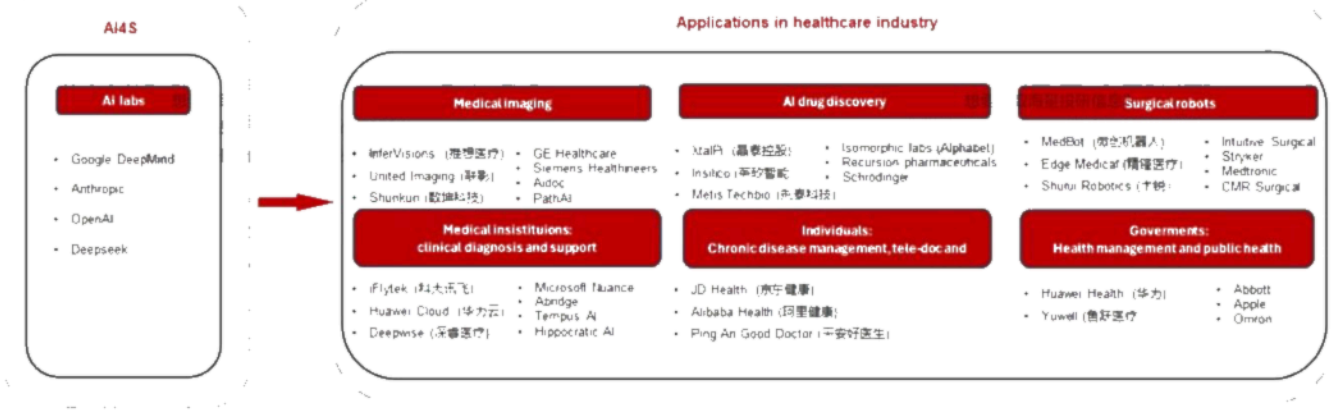


Source: Pony.AI, Nomura research

AI + Healthcare

Artificial intelligence (AI) is not something new for the healthcare sector. Companies and researchers have applied AI for years in medical imaging analysis, online medical consultation, molecule discovery and design, among others. Yet, generative AI in recent years has 'turbocharged', and is increasingly involved in medical processes. For example, more medical institutions are applying AI for clinical diagnosis and support to reduce the workload of medical professionals and streamline the process. The recently-emerged and hotly-discussed AI for Science (AI4S), thanks to more and more frontier AI labs including Anthropic (unlisted), OpenAI (unlisted), and Google (GOOGL US, not rated) DeepMind are racing to build models capable of autonomous scientific discovery, attracting unprecedented interest regarding AI's potential disruptive role in the fundamental healthcare space.

Fig. 44: AI in healthcare industry



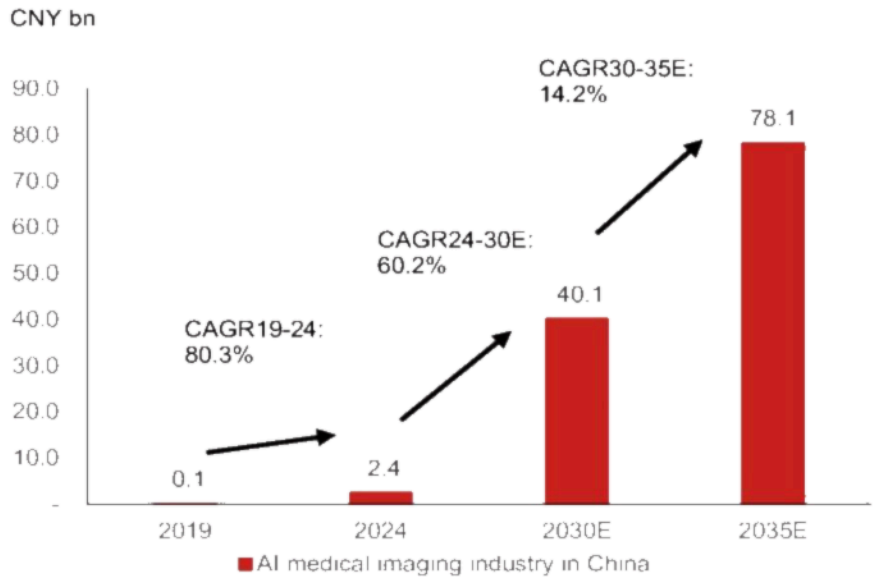
Source: Nomura research

Medical imaging

Medical imaging is one of the most well-applied areas for AI in the healthcare space. AI-empowered medical imaging spans from pulmonary nodule and breast lesion detection to coronary CTA (Computed Tomography Angiography), stroke, and cancer screening. Beyond diagnostics, AI enhances image quality via intelligent reconstruction and de-noising, automates lesion quantification, and streamlines the entire imaging workflow – from scan preparation to structured reporting and multi-center research.

According to GrandView research, the global AI medical imaging market was estimated at USD1.36bn in 2024 and is projected to reach USD19.8bn in 2033E, with a CAGR of 34%. Frost and Sullivan(F&S) estimates China's AI medical imaging market recorded CNY2.4bn in 2024E and will reach CNY78.1bn in 2035E.

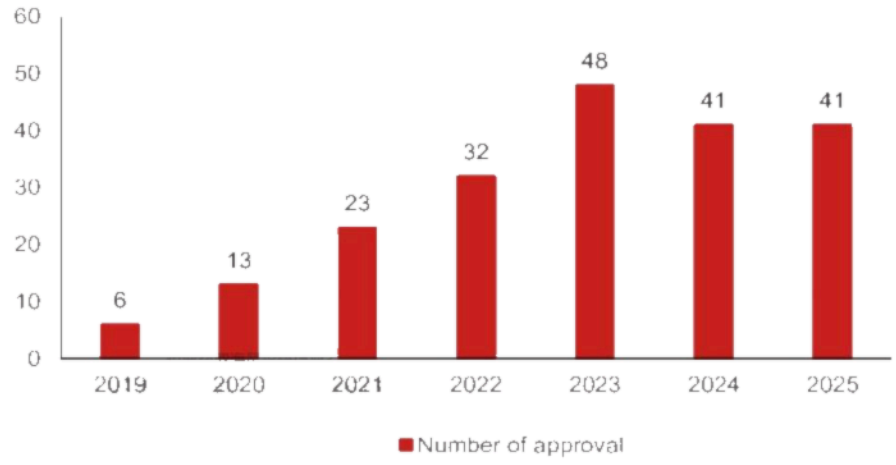
Fig. 45: China AI medical imaging market size, 2019-2035E



Source: Frost and Sullivan, Nomura research

Overseas companies prioritize foundational model development (Paige [unlisted]'s Virchow, Tempus [TEM US, not rated]'s multi-modal platforms), global enterprise-grade deployment (Aidoc [unlisted], Viz.ai [unlisted] across 1,000+ hospitals), and FDA-approved workflow-integrated solutions. China has competitive edge in hardware-AI integration and economies of scale, in our view, as it has thousands of hospitals that generate vast real-world data. According to NMPA, companies including medical imaging manufactures (United Imaging [688271 HK, Buy]), AI firms (InferVision [unlisted], Deepwise [unlisted], ShuKun [unlisted]) and internet companies (Tencent [700 HK, Buy]) obtained 41 (vs 6 in 2019) NMPA Class III certificates for AI medical imaging device in 2025. We believe that the higher accessibility of open-source models could help Chinese companies narrow the algorithmic gap.

Fig. 46: No. of Class III AI medical imaging device approved in China, 2019-25



Source: NMPA, Nomura research

AI drug discovery (AIDD)

Drug R&D is another rapid evolving area in which AI plays an increasingly active role. Firms leverage machine learning, natural language processing and generative models to identify targets, design new molecules, select optimal compounds, predict pharmacokinetic properties and optimize clinical trials via patient recruitment analytics and real-world evidence. Compared to conventional drug R&D (10-15 years required and USD2.6bn cost per approved molecule and 10% success rate, according to Tufts CSDD and Biomedtracker), AI can potentially shorten development timelines with lower costs and achieved higher success rate, in our view.

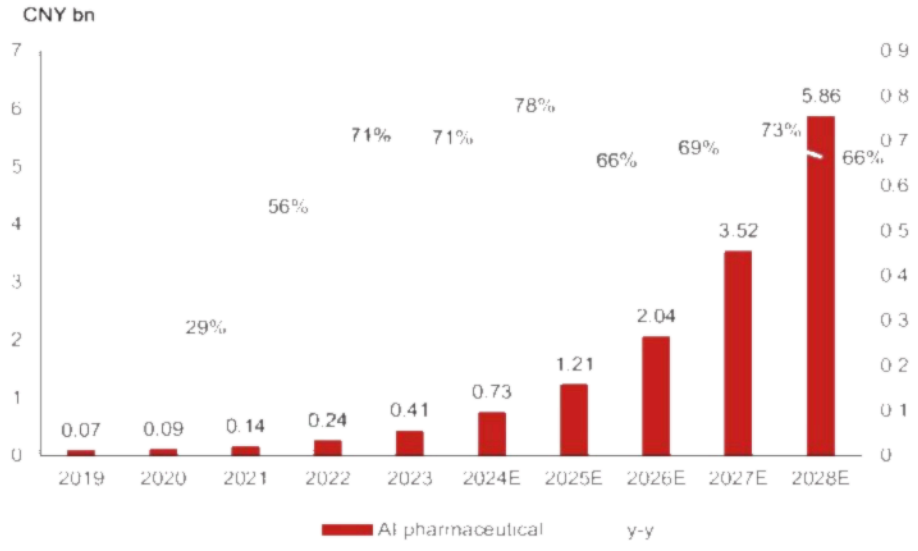
Fig. 47: AI drug discovery applications

Application Area	AI Technology	Impact
Target identification	ML-based analysis of genomic/proteomic data	Novel drug target discovery from biological datasets
Molecular design	Generative AI (de novo drug design)	Novel chemical entities with optimized pharmacokinetics
Drug optimization & repurposing	Deep learning on drug-target interactions	Identify new uses for existing compounds
Preclinical testing	AI+robotics automated experiments	Reported improvement in synthesis success rates
Clinical trial design	Predictive analytics	Optimize patient recruitment, predict trial outcomes
Protein structure prediction	AlphaFold, BioEmu	Protein structure prediction in near real-time
Virtual screening	NIM microservices	Accelerated small molecule screening

Source: Nomura research

The global AI in drug discovery market was valued at approximately USD2.5bn in 2024E by Statista and is expected to reach USD20bn in 2032E, with a CAGR of 30%. China's market is forecast to record a significantly faster CAGR of 68% from CNY0.73bn in 2024E to CNY5.86bn in 2028E, according to F&S.

Fig. 48: China AI drug discovery market size, 2019-2028E



Source: Frost and Sullivan, Nomura research

We notice companies from the US and China lead in the global AI drug development race: US firms maintain leadership in first-in-class discovery, rare disease pipelines, holding roughly 50% of global AI pharmaceuticals patents since 2020 (source:GlobalData patent analytics). We have observed Chinese firms' R&D capacities is recognized by MNCs, evidenced by 157 out-licensing deals that were worth USD135.7bn in 2025. Among these deals, AI-related transactions include AstraZeneca (AZN LN, not rated)'s USD5bn deal with CSPC (1093 HK, Buy), Sanofi (SAN FP, not rated)'s USD2.56bn licensing agreement with Earendil Labs (unlisted), and XtalPi's USD3.45bn collaboration with Eli Lilly (LLY US, not rated) and USD5.99bn deal with DoveTree (unlisted). This reflects the collaboration pattern across the Pacific: US companies leverage China's clinical data for model training, while Chinese peers rely on US deep-pocket markets.

Fig. 49: Major China AI pharma out-licensing deals

Date	Company	Partner	Deal Value	Collaboration Focus
8/4/2026	Alphamab Oncology	Pathos AI	USD 2.218bn (USD 125mn upfront)	Global exclusive license for JSKN016, a first-in-class TROP2/HER3 bispecific ADC
7/10/2026	CSPC	AstraZeneca	USD 1.77bn	siRNA discovery platform and extrahepatic targeting delivery for novel siRNA candidates
7/1/2026	Insilico Medicine	Bora Pharmaceuticals	>USD 3bn (combined)	Multi-target strategic collaboration leveraging the Pharma.AI platform
7/1/2026	Insilico Medicine	Takeda	Up to USD 600mn	AI-driven drug discovery collaboration
6/30/2026	J&T Pharma (via Boulevard Bio)	Deerfield	>USD 1.6bn (USD 20mn upfront)	Global development, manufacturing and commercialization of trispecific TCE MTS-128
6/1/2026	Insilico Medicine	SK Biopharmaceuticals	USD 2.5bn	AI pharmaceutical collaboration
3/29/2026	Insilico Medicine	Eli Lilly	USD 2.75bn (USD 115mn upfront)	Global exclusive license for preclinical pipeline and R&D collaboration using Pharma.AI
1/30/2026	CSPC	AstraZeneca	USD 18.5bn (USD 1.2bn upfront)	Strategic R&D and license for long-acting peptide drugs, leveraging CSPC's sustained-release technology and peptide AI discovery platform
1/5/2026	Insilico Medicine	Servier	USD 888mn	Pharma.AI platform focusing on challenging oncology targets to identify and develop novel therapeutics
1/1/2026	Earendil Labs	Sanofi	USD 2.56bn (USD 160mn upfront + near-term)	Platform-level strategic collaboration covering multiple autoimmune and inflammatory programs, focusing on bispecific antibodies
11/5/2025	XtalPi (Ailux)	Eli Lilly	USD 345mn (multi-million upfront)	Multi-target strategic collaboration and platform license to accelerate bispecific antibody discovery using Ailux's proprietary platform
11/1/2025	Insilico Medicine	Eli Lilly	>USD 100mn	Generation, design and optimization of novel candidate compounds for innovative targets using Pharma.AI (third collaboration)
8/5/2025	XtalPi	DoveTree Medicines	USD 5.99bn (USD 51mn upfront)	End-to-end "AI+robotics" platform to develop small molecule and antibody candidates in oncology, immunology and neuroscience
6/13/2025	CSPC	AstraZeneca	USD 5.33bn (USD 110mn upfront)	AI engine-driven discovery platform for oral small molecule candidates
4/17/2025	Earendil Labs	Sanofi	USD 1.845bn (USD 125mn upfront)	Global rights for two bispecific antibody drugs HXN-1002 and HXN-1003
1/10/2025	Insilico Medicine	Menarini	USD 550mn (USD 20mn upfront)	Out-licensing of preclinical oncology pipeline candidate MEN2501 discovered using AI

Source: Company data, Nomura research

Fig. 50: Listed AI pharmaceutical companies

Company	Ticker	Up-to-date Mkt Cap (USD bn)	Mkt at IPO / Listing (USD bn)	Mkt Cap Change Since IPO	IPO Proceeds (USD bn)	2025 Revenue & y-y Growth	Major AI Deals
XtalPi（晶泰控股）	2228 HK	4.55	2.53	80%	0.13	RMB802.6mn (+201.2%)	Eli Lilly (USD345mn); DoveTree (up to USD5.99bn); Gan & Lee
Insilico Medicine（英矽智能）	3696 HK	3.46	2.18	59%	0.29	USD56.2mn (-34.5%)	Eli Lilly (USD2.75bn); Menarini (>USD550mn); Servier (up to USD888mn); SK Biopharm (>USD2.5bn); Bora (>USD2.5bn, proposed alliance); Takeda
Metis TechBio（剂泰科技）	7666 HK	1.83	3.52	-48%	0.27	RMB105mn (+6,985%)	Deerfield / Boulevard Bio (USD20mn upfront; up to USD1.6bn milestones)
Schrödinger	SDGR US	1.33	0.97	37%	0.23	USD256mn (+23.3%)	Broad platform licensing; computational drug discovery
Recursion Pharmaceuticals	RXRX US	1.71	3.29	-48%	0.5	USD74.7mn (+26.9%)	Sanofi (up to 15 programs; USD134mn received); Roche/Genentech; Bayer; Merck KGaA; Takeda; Exscientia acquisition

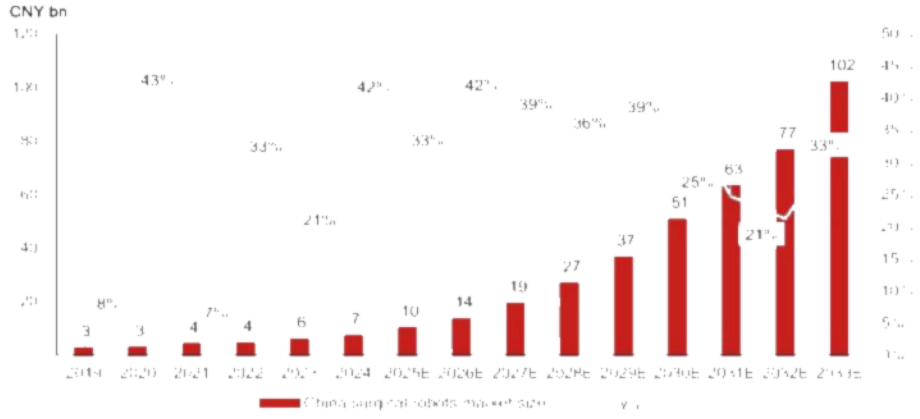
Source: Bloomberg Finance L.P., Nomura research

Surgical robots

More and more surgical robot companies have integrated AI into their products, with the aim of providing advanced features such as more accuracy, real-time surgical image navigation, sub-millimeter servo control, and force-feedback sensing.

In China, more and more companies (currently more than 50) enter to compete the high-growing surgical robots market that F&S forecasts to jump tenfold from CNY10bn in 2025 E to CNY102bn in 2033E.

Fig. 51: China surgical robots market size, 2019-2033E



Source: Frost and Sullivan, Nomura research

In the ex-China markets, the market is still largely dominated by Intuitive Surgical (ISRG US, not rated), which integrates AI into its da Vinci 5 ecosystem that includes AI-driven case insights via My Intuitive platform and real-time assistance for anatomical identification during surgery; and for Ion system, AI helps segment CT scans and plans biopsy trajectories, enabling patient-specific navigation. Management has also highlighted a layered AI strategy that it confirmed to sustain investments to support long-term digital roadmap.

Clinical diagnosis and support for medical institutions

AI-powered medical diagnosis, applied across cancer early screening, pathology diagnosis, clinical decision support, and automated report generation, delivers comparable and increasing quantifiable clinical diagnosis. For example, Ambient systems use in-room microphones to listen to doctor-patient conversations in real time, automatically converting spoken dialogue into structured electronic health records (EHR), draft diagnostic notes, and referral letters — eliminating manual typing.

In China, the focus is clinical deployment (primary-care integration, reimbursement, medical imaging), while that in American is foundational AI models (protein folding, multi-

modal databases) and vertical SaaS. The DeepSeek-enabled cost revolution and together with the implementation of favorable policy is positioning China's diagnostic AI ecosystem for accelerated global competitiveness, in our view.

Physician use seek AI tools to update knowledge and learn complex medical cases in a rapid pace.

Chronic disease management, remote medical consultation, online drug sales and personalized medicine for individuals

Medical resources are unevenly distributed almost in every country, but people in remote areas can get access to top physicians in urban regions, and patients with non-urgent diseases can do a self-assessment before heading to hospitals through the help of AI tools. From an individual user perspective, we observe more and more people are turning to AI chatbot for preliminary diagnosis, medical knowledge, updated guidelines on certain drugs usage. In the long run, AI can also be used for chronic disease management and even for precision medicine, if AI can further breakdown the cause of disease at the gene level, in our view.

The current mature business model is for drugs sales, which AI can serve as a doctor or pharmacist. China's online pharmaceutical retail market exceeded CNY80bn in 2025, as per the National Medical Products Administration (NMPA).

In May 2026, China's NMPA issued the "Compliance Guidelines for Online Retail of Prescription Drugs." Key provisions in the table below:

Fig. 52: Regulation requirement and impact

Regulatory Requirement	Impact
Prescription reviews must be conducted by licensed pharmacists only	Eliminates "AI-second-prescription" practices; raises personnel costs
AI prohibited from replacing human prescription review	Constrains automated scaling; increases compliance costs
Daily prescription review volume must be within "reasonable range"	Creates throughput bottleneck for high-volume platforms
Marketing promotions for prescription drugs prohibited	Ends discount-driven, buy-one-get-one strategies
Short videos/live streaming to induce drug purchases banned	Restricts customer acquisition channels
Real-name verification for prescription drug purchases	Enhances traceability but adds friction
Third-party platforms bear greater compliance responsibility	Increases platform governance costs

Source: Nomura research

Health management and public health monitor for governments

Medical device makers also use AI in their wearable health monitoring devices that can capture ECG, continuous glucose (CGM), blood oxygen, and sleep indicators. AI models issue proactive warnings before acute events — atrial fibrillation onset, nocturnal hypoglycemia, heart-failure deterioration risk — and can connect directly to family doctors or care teams.

From a public health perspective, WHO or CDC can apply AI to combine fever-clinic visits, OTC cold-medicine sales, anomalous social-media signals, weather, and population-mobility data into multi-point trigger models to predict flu/infectious-disease outbreaks weeks in advance, supporting precision prevention and ICU/vaccine resource allocation.

As developed countries are facing rapid population aging, we believe it would be in the interest of the governments to turn to AI and automative to provide eldercare solutions.

AI4S: the next big thing

AI4S refers to the application of frontier AI models to accelerate scientific discovery across biology, chemistry, materials science, physics, math, etc. Unlike earlier "AI in research" approaches that used machine learning as a narrow analytical tool, AI4S envisions AI as a “human-like” agent in the scientific research setting: generating hypotheses, designing experiments, interpreting results, and iterating.

Currently US-leading AI labs which are frontrunners in the AI4S area are:

In 2024, Google DeepMind's Demis Hassabis and John Jumper received the Nobel Prize

in Chemistry for AlphaFold, which has predicted over 200mn protein structures and is used by more than 2mn researchers across 190+ countries. This validated AI's capacity for foundational scientific contribution, not merely incremental optimization.

Anthropic pursues aggressive and comprehensive AI4S strategies, embedding biological capabilities directly into its general-purpose Claude models while building in-house drug discovery expertise. It has co-worked with multiple companies in pharmaceutical and clinical areas, besides building its own wet-lba. It also attracted John Jumper to join in June 2026.

OpenAI's AI4S strategy emphasizes specialized vertical models, plugin ecosystems, and trusted-access partnerships rather than building in-house wet lab. GPT-Rosalind is not publicly available. Access is restricted to qualified enterprise customers (e.g., Amgen [AMGN US, not rated, Moderna [MRNA US, not rated], Thermo Fisher [TMO US, not rated]) through HIPAA-aligned regulated workspaces with a commitment not to train on customer data.

Google DeepMind, the pioneer of AI4S via AlphaFold, has taken a third path: spinning out an independent AI drug discovery company while pursuing broader AI4S research internally.

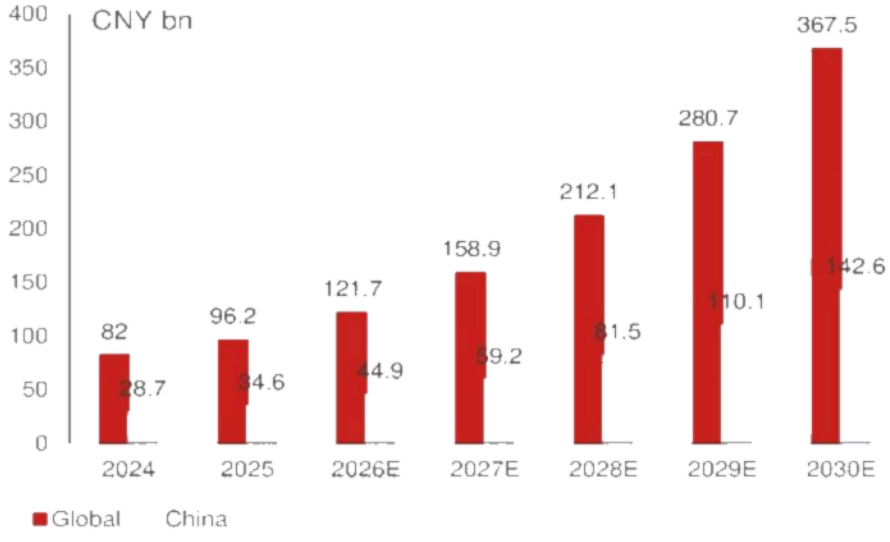
Chinese peers, despite strong policy backing, national computing infrastructure have yet to emerge as a "Deepseek moment" for AI4S. However, some Chinese CRDMO companies have the potential to involve AI4S development by providing services for these AI labs.

AI + Robotics

AI + Robotics market TAM

The global market size for embodied AI robotic solutions reached CNY82.0bn in 2024 and is projected to grow further to CNY367.5bn by 2030E, representing a CAGR of 28.4% from 2024 to 2030E, according to Frost & Sullivan. The market size for embodied AI robotic solutions in China reached CNY28.7bn in 2024 and is expected to further grow to CNY142.6bn by 2030, with a CAGR of 30.6% from 2024 to 2030E.

Fig. 53: Market size of embodied AI robotics solutions (CNY bn)



Source: Frost & Sullivan, Nomura research

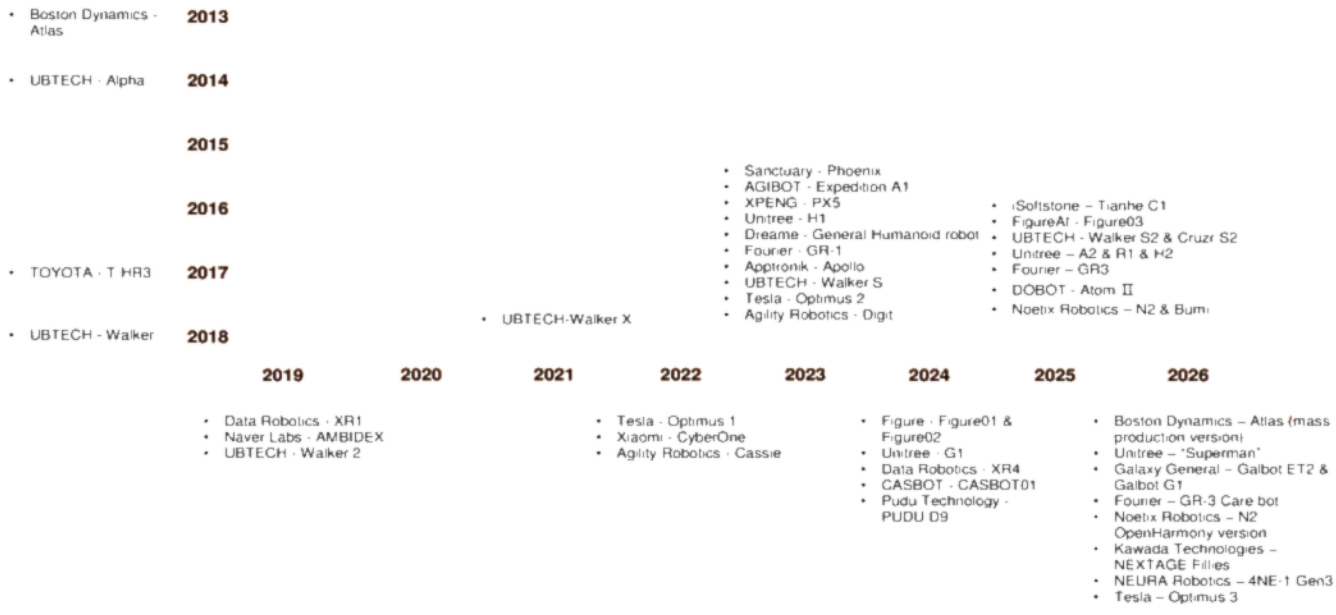
Select embodied robotic products

The pace of practical training of humanoid robots in automobile factories has accelerated significantly. The automobile production process covers a large number of subdivided scenarios such as assembly, handling, and quality inspection, providing an ideal "test field" for robots. Tesla (TSLA US, Not rated) released its second-generation humanoid robot Optimus in mid-December 2023, with an increase in walking speed. In early 2024, Figure 01, a humanoid robot released by Figure AI (unlisted), entered the BMW factory for

standardized task training and execution; UBTECH's (9880 HK, Buy) Walker S entered NIO's (9866 HK, Buy) factory in February and completed its first on-site training and performed assembly auxiliary tasks.

At the application level, the latest iteration of Optimus demonstrated its ability to perform basic factory tasks, such as battery cell sorting, power tool use, and delicate operations such as picking up eggs. These application examples highlight Optimus's technological breakthroughs in precision operation and autonomous working capabilities, in our view, which could give the market more confidence. The robot is expected to be trial-produced in 2024 and mass-produced in 2025. Figure 02, launched by Figure AI in 2024, integrates OpenAI's GPT-4o multi-modal large model in the brain. Compared with Figure 01 which uses GPT4, Figure 02 further improves the robot's common-sense reasoning capabilities, allowing it to better understand and respond to complex tasks with on-board computing and AI reasoning capabilities improved by three times, according to the company. Figure 03 uses the Helix, an end-to-end VLA (Vision-Language-Motion) general model, to target home service scenarios and can complete complex household chores such as folding clothes, washing dishes, and tidying up rooms. Figure AI aims to produce 100,000 units of it within four years.

Fig. 54: Launch of mainstream humanoid robotics timeline (selected)



Source: Company data, Nomura research

Key trends in the AI robotics market

NeuralAxis by NXP: The reflex-first blueprint for physical AI

NXP (NXPI US, Not rated) set out the NeuralAxis (Neural Axis Architecture) framework at the COMPUTEX 2026 keynote, presented by President and CEO Rafael Sotomayor. NXP's premise is that what limits Physical AI is not scaling up language-model reasoning but engineering a low-latency reflex layer, the machine equivalent of unconscious human response, which is consistent with Moravec's paradox. NeuralAxis splits the stack the way a nervous system does, into three tiers that run decoupled but coordinated: reasoning at roughly 300ms (cortex), coordination for motion control and balance (cerebellum), and reflex as fast as 40ms (spinal cord), sited at the edge next to the actuators. That split is what buys low latency, distributed control and energy efficiency together instead of trading one against another. Humanoids are where it bites the hardest. Instead of a central brain, reflex processors sit in the joints, hands and feet and decide locally on grip force and ankle balance, chaining the recovery of balance, grasp, posture and gait inside 40ms. Decoupling reasoning from motion control also means a new skill can be added without destabilizing the walk. NXP applies the same blueprint beyond humanoids, bringing drone glass-to-glass latency inside 20ms and dividing software-defined-vehicle control into reasoning, coordination and safety-critical zonal execution. On commercialization, our industry checks point to a real productivity gain over conventional automation and a likely sharp pickup in diagnostic-robot sales.

From blueprint to product: WVLA2.0's world-model–VLA fusion and co-design

If NeuralAxis is the blueprint and the safety doctrine, Unitree's (688836 CH, Buy) WVLA2.0 is what the blueprint looks like once it ships. WVLA2.0 is the first commercially deployable version after two years of development, and it fuses the predictive capability of the WMA (World-Model Action) model with VLA's end-to-end action generation, which sets Unitree apart from peers that have bet on VLA alone. The fusion buys better high-level task comprehension, 2D/3D spatial-semantic reasoning, action generation that respects dynamics constraints, and more tolerance of disturbance. On perception, four visual streams run in parallel — a RealSense (Unlisted) depth camera, a Livox (Unlisted) MID360 LiDAR (Light Detection and Ranging) and two side cameras — and resolve into a single 360-degree representation, with position updating inside 10ms even under interference. The co-design shows on the output side: action parameters go out over the CAN (controller area network) bus to the G1's 23 degrees of freedom, handled by Unitree's own "cerebellum" motion control, and single-arm grasping of objects under 2kg holds positioning error inside 5mm. Lightweighting holds on-board compute under 100 TOPS, which is what lets the model run entirely on the G1 EDU's NVIDIA (NVDA US, Not rated) Jetson Orin NX with no cloud dependence; our industry survey indicates this is what avoids task interruption from latency or disconnection.

Embodiment-free data capture becomes mainstream

Keeping the model on the robot also shapes how the data gets collected, and collection is moving to "embodiment-free" capture as the default. In a single-take demo, a WVLA2.0-equipped G1 worked through six tasks in a deliberately disturbed meeting room with no teleoperation, which shows long-horizon execution and a perception-prediction-decision-action loop closing at about 90ms per inference, or roughly ten iterations a second. Our industry checks are less flattering on the detail: blind spots and gaps in rear-area perception, noticeable noise, slow execution, and imprecise fine manipulation, with no quantified success-rate benchmark over continuous operation to check any of it against. Management guides industrial manufacturing as the first landing zone, specifically joint-motor assembly, loading and fixture handling, given Unitree's own factories. Logistics sorting and flexible 3C assembly follow, with home and healthcare later, where open and unstructured environments make the problem much harder. Management sees the differentiation as full-stack in-house integration plus the scale of real-machine data coming off the global fleet, an asset it says model vendors working from the cloud alone cannot replicate.

Rolling local buffers make Unitree's fleet-data advantage physical, not commercial

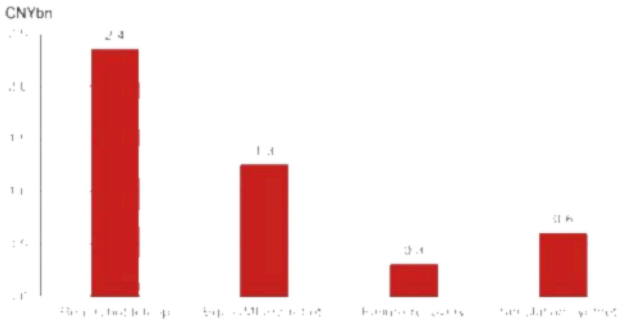
The barrier to replicating that data asset is a physical one. Based on our industry checks, the standard industry design is inference on the robot, training in the cloud, and logs kept in tiers, so the operating loop closes on the machine and the training loop closes remotely. Raw capture sits in a rolling buffer on the robot for anywhere between 6–8 hours and three days, then gets overwritten; what actually goes up is a compressed selection of event clips, exception tickets and success/failure samples. Uploading a full day of raw video is not economic over cellular or Wi-Fi backhaul. Same-day training data splits about 1:3 between cloud and robot, and our survey shows that the ratio as stable because training volumes are planned, not opportunistic. File formats have settled across vendors; however, what data gets uploaded, what gets overwritten and on what schedule have not been settled. This also puts a boundary on the no-cloud-dependence claim above: it holds for the operating loop, while the training loop stays in the cloud at that 1:3 split. Leading robotics firms decide what to keep before the rest is detected, and that is where we read as the fleet advantage.

Fig. 55: Unitree G1 bill-of-materials

Parts	Line item	Cost (CNY)	% of BOM
Joint modules	9 large joints	13,500	32.5%
Joint modules	14 small joints	14,000	33.7%
Sensing components	3D LiDAR - DJI Livox MID360	3,840	9.2%
Sensing components	Depth camera - Intel RealSense D435i	1,869	4.5%
Structure parts	Structural parts	2,000	4.8%
Structure parts	Outer shell	500	1.2%
Structure parts	Other structure/misc	1,000	2.4%
Power	Battery pack (13S, 9,000mAh, 0.42kWh, 2.5kg)	2,000	4.8%
Power	PDU	500	1.2%
Power	DC-DC	480	1.2%
Power	Other power control	200	0.5%
Computing	Main SoC (8nm, 6 TOPS NPU)	380	0.9%
Computing	DRAM 8GB	75	0.2%
Computing	Storage 64GB	60	0.1%
Computing	Other main-board content	900	2.2%
Wire harness	XT30 harness	90	0.2%
Wire harness	Harness terminals	150	0.4%
Wire harness	End-of-limb capacitive buffer	30	0.1%
	Total BOM	**41,574**	**100.0%**

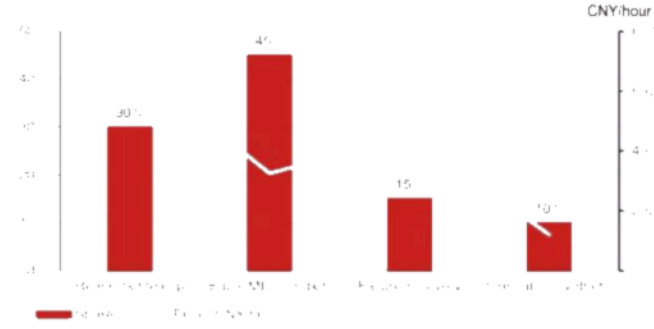
Source: Company data, Nomura research

Fig. 56: Data collection market size (CNY bn)



Source: Company data, Nomura estimates

Fig. 57: Share and ASP in different data collection types



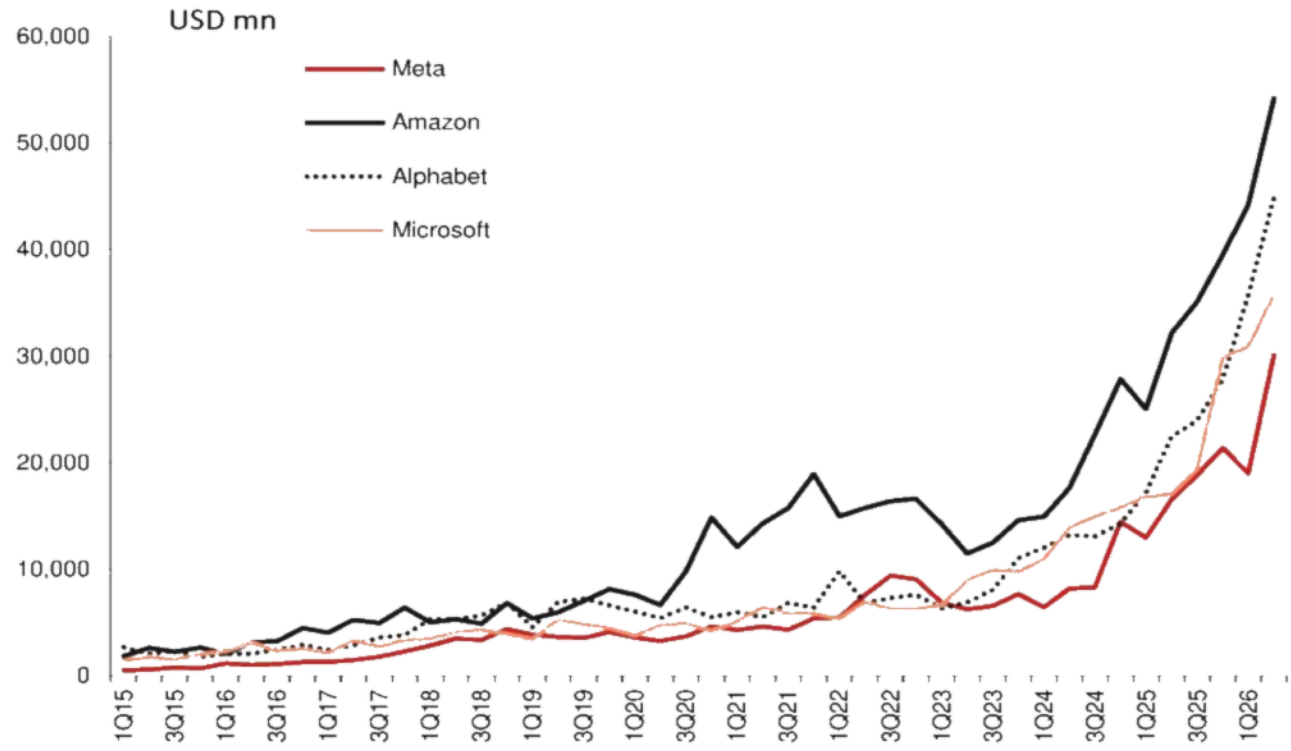
Source: Company data, Nomura estimates

Token economics – When will AI generate positive ROI and in which areas?

Compute supply model: Capex growth may slow down and but computing power supply will likely sustain longer

Since 2024, global hyperscale tech platforms have been rushing into AIDC investments for LLM training, fearing of missing out the opportunities for the generative AI market. Based on Bloomberg consensus, the big four hyperscale AI companies (Meta / Amazon / Alphabet / Microsoft) would invest a combined USD710bn / USD907bn capex in 2026E / 2027E. Most of these companies' managements reiterated the strong demand outlook for their cloud / AI platforms during their June-26 quarterly earnings calls, noting that demand cannot be fully fulfilled due to the supply gap in the infrastructure supply chain. Despite market concerns about the ROI on the intensive infrastructure spending, Amazon's management noted that it took less than three years to reach break-even on server & networking equipment investment in the June 2026 quarter earnings call, and AWS Cloud could achieve USD1tn annual revenue in the long run. We think these indicate strong commitment from the hyperscale AI companies to invest in AI computing power in the future.

Fig. 58: Global hyperscaler capex growth trend (Actual)



Source: Company data, Nomura research

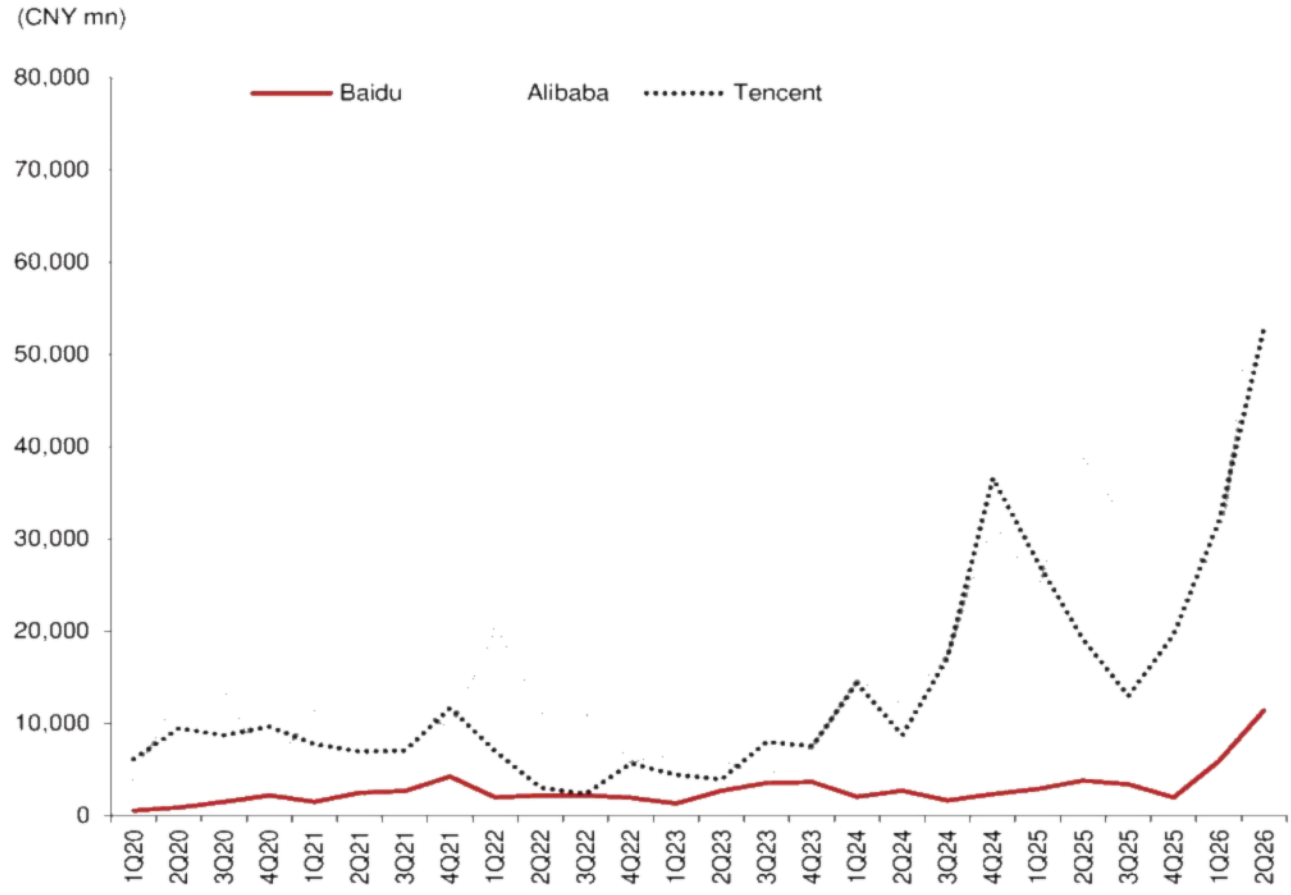
Fig. 59: Global hyperscaler capex forecast vs OCF trend

	Actual				Estimates for 2026-2027				
					Consensus (pre-2Q26 earnings)		Consensus (post-2Q26 earnings)		Guidance
(USD mn)	2022	2023	2024	2025	2026E	2027E	2026E	2027E	2026
CAPEX	**170,180**	**165,910**	**251,579**	**376,051**	**707,032**	**899,290**	**710,426**	**907,128**	**755,650**
y-y	*18%*	*-3%*	*52%*	*49%*	*88%*	*27%*	*89%*	*28%*	*101%*
Microsoft	24,769	35,202	55,552	83,094	160,629	198,127	161,936	198,885	175,000
Google	31,485	32,251	52,535	91,447	186,400	258,443	186,732	259,760	200,000
Meta	31,431	27,266	37,256	69,691	135,671	175,811	136,025	177,312	137,500
Amazon	63,645	52,729	82,999	131,819	200,521	239,969	202,497	244,864	220,000
OCF	**397,022**	**490,141**	**581,304**	**580,533**	**875,659**	**910,443**	**873,013**	**1,085,869**	
Microsoft	84,386	102,647	125,583	160,506	198,146	237,109	204,421	237,241	
Google	91,495	101,746	125,299	164,713	211,499	264,469	211,127	263,997	
Meta	50,475	71,113	91,328	115,800	137,346	168,933	137,037	169,012	
Amazon	46,752	84,946	115,877	139,514	184,128	239,933	184,120	240,261	
CAPEX/OCF (%)	43%	34%	43%		81%	99%	81%	84%	

Source: Company data, Bloomberg Finance L.P., Nomura estimates

On the other hand, we think China's AI market is under invested due to the lack of advanced AI chips, as well as concerns about the solid monetization models. However, we think China has the potential to accelerate AI spending, thanks to the proactive development of LLMs (especially open-source LLMs), as well as technology upgrades from domestically developed AI chips. Government sector, SOEs and private sectors would all participate actively in China's AIDC investment, in our view.

Fig. 60: China BAT capex trend



Source: Company data, Nomura research

Fig. 61: China hyperscalers & telcos' AI capex forecasts

China CSPs' total capex (CNY mn)	2021	2022	2023	2024	2025	2026F	2027F	2028F
Baidu	10,896	8,286	11,190	8,134	12,073	32,690	36,007	40,034
y-y		-24%	35%	-27%	48%	171%	10%	11%
Alibaba	41,450	53,309	34,330	32,087	85,972	182,188	142,763	153,392
y-y		29%	-36%	-7%	168%	112%	-22%	7%
Tencent	29,302	22,679	21,008	62,927	87,482	207,190	181,389	148,655
y-y		-23%	-7%	200%	39%	137%	-12%	-18%
BAT Total	81,648	84,274	66,528	103,148	185,527	422,068	360,159	342,080
y-y		3%	-21%	55%	80%	127%	-15%	-5%
ByteDance				80,000	150,000	300,000	500,000	650,000
y-y					88%	100%	67%	30%
BBAT Total				183,148	335,527	722,068	860,159	992,080
y-y					83%	115%	19%	15%
China Telcos' total capex (CNY mn)	**2021**	**2022**	**2023**	**2024**	**2025**	**2026F**	**2027F**	**2028F**
China Mobile	183,600	185,200	180,300	164,000	150,900	136,600	129,770	128,472
y-y		1%	-3%	-9%	-8%	-9%	-5%	-1%
China Telecom	86,723	92,528	98,838	93,513	80,400	73,000	69,350	68,657
y-y		7%	7%	-5%	-14%	-9%	-5%	-1%
China Unicom	69,000	74,200	73,900	61,400	55,000	50,000	47,500	47,025
y-y		8%	0%	-17%	-10%	-9%	-5%	-1%
China Telcos Total	339,323	351,928	353,038	318,913	286,300	259,600	246,620	244,154
y-y		4%	0%	-10%	-10%	-9%	-5%	-1%
BBAT + China Telcos Total Capex (CNY mn	**339,323**	**351,928**	**353,038**	**502,061**	**621,827**	**981,668**	**1,106,779**	**1,236,234**

China Telcos' computing power investment (CNY mn)	2021	2022	2023	2024	2025	2026F	2027F	2028F
China Mobile		33,500	39,100	37,100	23,300	37,800	49,140	61,425
y-y			17%	-5%	-37%	62%	30%	25%
China Telecom		17,886	23,496	21,582	20,100	25,550	30,660	35,259
y-y			31%	-8%	-7%	27%	20%	15%
China Unicom		6,158	7,390	8,794	11,256	17,500	19,250	20,790
y-y			20%	19%	28%	55%	10%	8%
Total		57,544	69,986	67,476	54,656	80,850	99,050	117,474

Note: Computing investment of China Telecom before 2025 is adjusted based on the proportion of FY25 computing infrastructure to industrial digitalization
Source: Company data, Nomura estimates

Based on our assumptions, **we expect the total AI computing power from leading LLM and hyperscale players to grow from 13GW in 2025F to 98GW in 2030F**. Although the net additions of capex / computing power may slow down from 2028F onwards, the enlarged AI clusters could underpin LLM training and inference demand for a longer term, in our view. As generative AI applications see wider adoption from the consumer and enterprise markets, we think inference workloads will take more share from training workloads, which could lead to better monetization opportunities, as training does not generate tokens / revenue directly.

Fig. 62: Global AI computing power size

	2025F	2026F	2027F	2028F	2029F	2030F
Compute deployment						
Cumulative AI compute (GW)	13 0	33 0	58 0	76 0	91 0	98 0
Net additions in year (GW)	6 0	20 0	25 0	18 0	15 0	7 0
Inference compute (GW)	8.5	21.5	37.7	49.4	59.2	63.7
Training compute (GW)	4 6	11 6	20 3	26 6	31 9	34 3

Source: Nomura estimates

Token economic model: pricing elasticity vs demand, profitability and ROI, LLM companies vs hyperscalers

In our token economy model, we assume the tokens generated by the total inference computing power every year would be the maximum supply for the generative AI market. We assume **token output per GW would increase 35% each year**, led by efficiency improvement in the computing hardware. Meanwhile, **in our base case scenario, we expect overall token price to decline 20% each year**, which may stimulate higher adoption rate and lead to more token demand. Based on the above-mentioned assumptions, we estimate **token revenue could grow from USD136bn in 2025F to USD1,3tn in 2030F**.

Fig. 63: Global token revenue forecasts (USD bn)

Token supply & demand	2025F	2026F	2027F	2028F	2029F	2030F
Throughput per GW (mn tok/s)	300	405	547	738	996	1,345
Annual capacity per GW (trn tok)	9,461	12,772	17,242	23,277	31,424	42,423
Max sellable tokens (trn, @util ceiling)	67,952	232,867	552,530	977,406	1,579,926	2,296,969
Token demand (trn)	90,000	251,558	703,125	1,054,688	1,582,031	2,056,641
Actual tokens sold (trn)	**67,952**	**232,867**	**552,530**	**977,406**	**1,579,926**	**2,056,641**
Implied utilization	85.0%	85.0%	85.0%	85.0%	85.0%	76.1%
Demand / max supply	**132.4%**	**108.0%**	**127.3%**	**107.9%**	**100.1%**	**89.5%**
Pricing & revenue						
Blended token price ($/mn tok)	$2.00	$1.60	$1.28	$1.02	$0.82	$0.66
Annual compute rent fee per GW ($bn/GW)	$9.7	$9.4	$9.1	$8.9	$8.6	$7.5
Platform token revenue ($bn)	**$136**	**$373**	**$707**	**$1,001**	**$1,294**	**$1,348**

Source: Nomura estimates

We think there are two types of AI players which build their monetization models based on the token economy model. The first type is **standalone LLM companies, which rent the AI computing power for LLM training and inference**; these companies capture the token revenue by offering APIs to end-customers, and pay for the computing power usage to infrastructure platforms. These players run an asset-light model, as they do not need to buy large AI clusters which are cash burning businesses; however, they need to pay the premium for computing power rental fees. The second type of player is **hyperscale AI platform companies, which have a vertically integrated business model by building the AI infrastructure** and train their models in-house; these are asset-heavy business models, but they could enjoy economies of scale as their AI Cloud ecosystem grows. In our forecast below, the asset-light LLM business model's return is highly correlated with token prices and compute rental fees, and EBIT margin could gradually improve with a declining portion of rental fees (improved GPM). Based on our assumptions, the **asset-light LLM business model could generate as high as 29.6% EBIT margin in 2030F (assuming ~70% GPM)**. Meanwhile, the hyperscaler (with LLM incorporated) model could generate **as high as 47% EBIT margin in 2030F**, as they can effectively absorb the computing cost, while **ROIC could reach 17~20% in FY28-30F, in our assumption**. We believe standalone LLM players have to take great efforts to achieve better intelligence for their models, in order to charge premium token prices and improve their margins, while hyperscalers could enjoy better ROI thanks to the stronger ecosystem as well as economies of scale.

Fig. 64: LLM players' P&L analysis

LLM players P&L ($bn)	2025F	2026F	2027F	2028F	2029F	2030F
Token revenue	$136	$373	$707	$1,001	$1,294	$1,348
Inference compute rent (to cloud)	$68	$168	$286	$365	$425	$398
Gross profit	$68	$205	$421	$636	$870	$950
Gross margin	**50.0%**	**55.0%**	**59.5%**	**63.6%**	**67.2%**	**70.5%**
Training compute rent	$37	$90	$154	$196	$229	$214
Opex + R&D	$34	$93	$177	$250	$324	$337
Platform operating profit	**($3)**	**$21**	**$90**	**$189**	**$317**	**$399**
EBIT margin %	**(1.9%)**	**5.8%**	**12.7%**	**18.9%**	**24.5%**	**29.6%**

微信: DYXZ0524 微信: DYXZ0524

Source: Nomura estimates

Fig. 65: Hyperscalers' AI business P&L analysis

Hyperscaler + LLM P&L ($bn)	2025F	2026F	2027F	2028F	2029F	2030F
Token revenue	$136	$373	$707	$1,001	$1,294	$1,348
IT Depreciation	$65	$165	$290	$380	$455	$490
Non-IT Depreciation	$9	$22	$39	$51	$61	$65
Power cost	$8	$20	$36	$47	$56	$60
Other opex	$13	$33	$58	$76	$91	$98
EBIT	$41	$132	$285	$448	$632	$634
EBIT margin %	**30.4%**	**35.5%**	**40.3%**	**44.7%**	**48.8%**	**47.1%**
Capex in year	$210	$700	$875	$630	$525	$245
Total Cumulative capex (incl. installed base)	$455	$1,155	$2,030	$2,660	$3,185	$3,430
Pre-tax ROIC (EBIT/cumulative capex)	**9.1%**	**11%**	**14%**	**17%**	**20%**	**18%**

Source: Nomura estimates

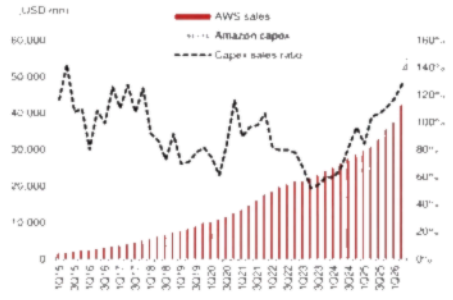
According to Bloomberg consensus, **cloud revenue from AWS//Google/Microsoft is expected to grow from USD451bn in 2026Eto USD800bn+ by 2028E, a CAGR of 36%** . In FY26-28F, cloud revenue of AWS/Google/Microsoft is expected to grow 31-35%/46-80%/23-38%, according to Bloomberg consensus. Based on AWS' management comment, **AWS Cloud could achieve USD1tn annual revenue business in the long run**, which reflects significant potential growth driven by AI demand from enterprises, in our view.

Fig. 66: AWS vs Google Cloud vs Azure Cloud future growth trend

Year	Amazon	Growth	Alphabet	Growth	Microsoft	Growth	Total cloud revenue (USD mn)	Total Growth
2019	34,986		8,918		16,778		60,682	
2020	45,370	30%	13,059	46%	25,267	51%	83,697	38%
2021	62,202	37%	19,203	47%	37,676	49%	119,082	42%
2022	80,096	29%	26,280	37%	52,425	39%	158,801	33%
2023	90,757	13%	33,088	26%	67,758	29%	191,603	21%
2024	107,634	19%	43,229	31%	90,238	33%	241,101	26%
2025	129,173	20%	58,705	36%	123,943	37%	311,821	29%
2026F	174,464	35%	105,935	80%	170,765	38%	451,165	45%
2027F	233,523	34%	173,914	64%	209,915	23%	617,352	37%
2028F	305,275	31%	253,671	46%	277,226	32%	836,172	35%

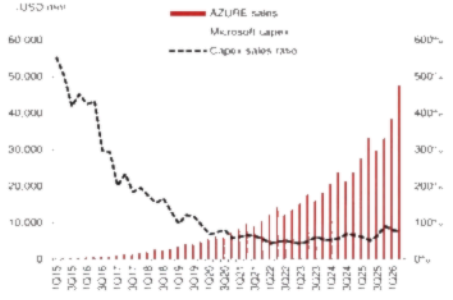
Source: Bloomberg consensus estimates, Nomura research

Fig. 67: AWS Cloud business growth trend



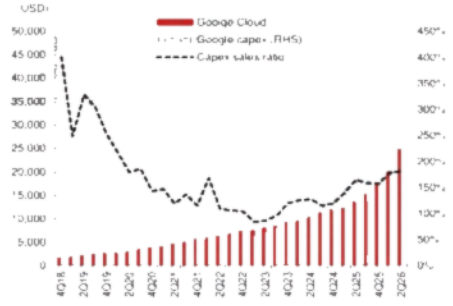
Source: Company data, Nomura research

Fig. 68: Microsoft-Azure Cloud business growth trend



Source: Company data, Nomura research

Fig. 69: Google Cloud business growth trend



Source: Company data, Nomura research

The token economy is increasingly evolving toward a "Token Factory" model, enabling the commercialization and large-scale monetization of compute capacity and AI services. Meanwhile, due to different industrial foundations and strategic priorities, China and

overseas players were adopting distinct development paths.

In overseas markets, leading AI developers such as OpenAI and Anthropic primarily pursue a closed-source strategy to preserve technological leadership and maximize commercial returns. Pricing is typically differentiated based on model capabilities (*Fig. 33*), allowing vendors to optimize monetization across customer segments.

China, by contrast, has adopted an "open source + ecosystem" strategy. Supported by structurally lower electricity costs and system-level efficiency optimizations pioneered by companies such as DeepSeek, Chinese foundation models could achieve lower training & inference costs. This cost advantage enables Chinese LLM providers to compete aggressively on API pricing and accelerate market penetration. Meanwhile, the open-source ecosystem lowers barriers to AI adoption, expands the developer community, and accelerates the commercialization of domestic AI applications. As application usage scales, token consumption increases accordingly, creating a self-reinforcing flywheel between lower costs, broader adoption, and growing token demand, in our view.

In terms of LLM's pricing, we note that Chinese leading LLM labs are covering the full pricing range (*Fig. 34*), from low-end to high-end market. The SOTA models from these China labs start to compete with top tiered models in the global market. For example, Zhipu simultaneously announced a 30% overall price increase for Coding Plan when GLM-5 was released in February 2026; and as of 1Q26, Zhipu API price increased 83% in total, according to the company. The recently-launched Kimi K3 (on 16 July) API is priced at CNY20/mn tokens for input and CNY100/mn tokens for output, which is 3.1x and 3.7x higher than Kimi K2.6, respectively. Compared to some low pricing models, the Kimi K3's API output and input prices are 12.5x and 10x that of the MiniMax-M3, respectively.

We think China's LLM players (including both close-source and open-source) will continue to gain traction in the global market, providing highly cost-efficient models which could fulfil the demand for a wide range of IT workloads. Meanwhile, top tier global LLMs could also accelerate their innovation, targeting the most complex workloads (i.e. science) and maintaining their technology & pricing premium, in our view. Moreover, due to geopolitical risks and the decoupling trend, we think the sovereign AI trend will become more popular, as most countries/enterprises would find it too risky to rely on only a small elite club for all Gen AI workloads in the future. Therefore, we think leading LLM players from both China and US would thrive, as long as they can stay ahead of the technology curve.

Policy impact on LLM market under US-China competition

We think generative AI is the cornerstone technology of our future society and economy, therefore, LLMs are becoming the key area for competition between large economies, especially the US and China, and both the countries have introduced a series of policies to support their Gen AI technology and LLM development. The AI policy guidelines during Trump's second term ("Trump 2.0") may have several targets, in our view: 1) strengthening government participation in industrial policy and innovation systems; 2) continuing the decoupling trend with China in the advanced AI technology sector; 3) reconstruct the global AI alliance network centered on the US. We think the policies of Trump 2.0 are targeting systematically deregulation and stimulate demand through three layers to ensure the leading development of the AI industry. First, regulatory relaxation. For example, Executive Order No. 14179 revoked multiple pre-assessment provisions from the Biden era, lowering compliance thresholds. Second, tax and credit incentives. The Treasury Department provides a 25% tax credit for investments in AI accelerators and high-bandwidth memory chips, while the "One Big Beautiful Bill" allocates USD500 mn to establish a "Business AI Modernization Fund", directly subsidizing federal deployment of SaaS-based AI solutions. Third, federal demand leverage. Policies engage converting government orders into first-order effects to accelerate AI commercialization.

In the meantime, China is also vigorously developing an independent AI ecosystem to meet the challenges of escalated decoupling trend. The Chinese government is committed to improve the localization ratio of computing power, strengthen LLM technology through algorithmic innovation incentives and open-source ecosystems, provide high-quality data through standardized and compliant data governance, and encourage commercial usage of LLMs in the industrial world. In addition, to further promote global cooperation, the Chinese government has proposed the concept of establishing a "World AI Cooperation Organization" and released the "Global AI Governance Action Plan," calling on the United Nations, international organizations, governments, and enterprises to jointly build an open, inclusive, fair, verifiable AI governance path that benefits developing countries, standing alongside the closed system led by the US.

Fig. 70: Major AI-related policies issued by the governments of the US and China

U.S.			China		
Date	Policy	Highlight	Date	Policy	Highlight
Jan-25	"AI Diffusion Framework"	The weights for closed-source AI models with training volumes of $\geq 10^{26}$ times are included in export controls; A global tiered licensing system for 3A090.a chips is implemented, dividing countries into three tiers and setting TPP quota caps; Introducing the Foreign Direct Product Rule (FDP), extending weights to models trained abroad	Apr-25	Ministry of Commerce Announcement No. 18 (Rare Earth Export Controls)	Export controls are imposed on seven categories of medium and heavy rare earth items including samarium, gadolinium, terbium, dysprosium, lutetium, scandium, and yttrium, requiring an export license, which is used to directly respond to U.S. chip controls on China.
Jan-25	"Removing Barriers to American Leadership in Artificial Intelligence"	This marks a shift in the U.S. government's AI policy toward a development strategy oriented toward "dual expansion of technology and market."	Apr-25	"Action Plan for the Development of the General AI Industry"	The file proposed that 5-10 internationally competitive AI companies will be formed by 2028, with a focus on supporting 12 major directions including LLMs, AI chips, and computing power networks
Jan-25	BIS expands AI chip export controls and updates the Entity List	BIS issued interim final rules expanding export controls on high-performance computing chips, added export restrictions on advanced closed-weight, dual-use AI models, and established ECCN 4E091	May-25	"Action Plan for Building Digital China 2025"	The goal is for the added value of core digital economy industries to exceed 10% of GDP by the end of 2025, computing power scale exceeding 300 EFLOPS, and a significant improvement in government digitalization and intelligence
Apr-25	Export controls on NVIDIA H20/L20/L2 chips upgraded	Actual ban on sale of compliant alternative chips such as Nvidia's H20, L20, and L2	Aug-25	"Opinions on Deepening the Implementation of the 'AI+' Action"	A "three-step" strategic goal is proposed: by 2027, achieve broad and deep integration of artificial intelligence with six key areas, with the penetration rate of smart terminals and agents exceeding 70%; By 2030, AI application penetration will exceed 90%; By 2035, China will fully enter a new stage of intelligent economy and intelligent society development
May-25	Trump gov. revoke the chip export control in "AI Diffusion Framework" and released three guidances	Three guidances were issued: 1) Global ban on Huawei Ascend chips; 2) Industry guidelines to prevent the reuse of advanced computing chips; 3) Clearly state export control policies for chips used to train AI models, warning against providing computing power support to Chinese AI models	Jan-26	"Implementation Opinions on the 'AI + Manufacturing' Special Action"	The document focuses on manufacturing, proposing that by 2027, 3-5 general-purpose AI models will be deeply applied in manufacturing, 1,000 high-level industrial intelligent agents will be launched, and 100 high-quality datasets will be created across industrial fields
Jul-25	"Winning the AI Race: America's AI Action Plan"	Three pillars have been established: 1) Accelerating AI innovation; 2) Building AI infrastructure (simplifying data center approvals, StarGate's USD500bn investment); 3) Leading international AI diplomacy (exporting technology to allies, strengthening export controls)	Jul-26	"Action Plan for AI Cooperative Development"	It proposed an international cooperation framework covering eight areas: data, computing power, ecosystem, empowerment, talent, rules, governance, and ethics

Source: Nomura research

The governments of both the countries have introduced relevant policies to support the widespread application of LLM and the development of smart terminals and robots, as well as large-scale fiscal investment to support AI-related industries. On the data front, both the countries regard data as a strategic resource and have established data control mechanisms. In terms of security, both the countries have introduced regulatory policies to prevent risks such as misinformation, algorithmic discrimination, and national security arising from AI technology.

The core differences in the policies of the two countries include different technology roadmaps and different strategy for competition. In terms of competition, the US was pursuing technology leadership by protecting its advanced chips and algorithms, while China empowers the Gen AI applications in real economy and industrial use cases. Regarding the technology roadmaps for LLMs, the US has tightening controls on the top closed-source LLMs and restrict the exports of core technologies (i.e. model weights) to China. On the other hand, China reduces its dependence on US technology through open source LLM development.

Fig. 71: Comparison of AI-related policies between the US and China

Dimension	U.S.	China
LLM development strategy	Limited open model weights and resources to protect core technologies	Build an innovation ecosystem centered on open source and breakthrough in core technologies to achieve independent substitution
Data governance	Restricted entry + extraterritorial jurisdiction	Restricted exit
Government support	The government mainly provides tax incentives, R&D subsidies, and infrastructure, while enterprises independently decide their technical routes	The government provides policy guidelines, tax incentives, subsidies and infrastructure, and enterprises are laid out technical direction around policy orientation

Source: Nomura research

The US's AI restriction policy against China spans from chips to LLM technology and to Physical AI. In terms of hardware, the US has imposed economic sanctions on a basket of Chinese AI chips and optical module manufacturers through the Entity List, while also imposing GPU export controls to restrict advanced chips from entering China. In January 2026, the BIS (Bureau of Industry and Security) made limited adjustments to its export license review policy for specific chips such as H200, changing from "presumed refusal" to "case-by-case review", but must simultaneously meet conditions such as sufficient supply, no encroachment on the US production capacity, no total TPP (Total Processing Performance) exports to China exceeding 50% of the US market, and security controls. In June 2026, US lawmakers officially submitted the "Guarding the US Against Adversarial Robotics Dominance Act", which plans to impose import restrictions or even a complete ban on Chinese-made humanoid robots and quadruped robots. The bill requires US national security agencies to complete safety reviews within one year of humanoid robots, quadruped robots, and their control software produced by China, Russia, Iran, North Korea, and other countries. Products identified as posing national security risks will be added to the FCC's "Control List" and banned from entry and domestic use. In August 2026, *Reuters* reported that the US Federal Communications Commission (FCC) is drafting new import restrictions on Chinese data center components, focusing on optical transceiver modules, with plans to announce and implement them within the year (*news*).

In terms of LLM and applications, in January 2025, the US BIS added ECCN 4E091 to the "Artificial Intelligence Diffusion Framework," which for the first time includes the weights of specific advanced closed-source AI models in export controls, mainly targeting models with training weights of 10^{26} or more operations. In March 2025, the US Department of Commerce included cloud computing power, model weights, and high-end algorithms in its licensing controls, aiming to cut off China's external access to training resources. It even includes AI research papers, academic exchanges, and talent mobility within its control scope, aiming to build a comprehensive "AI firewall". In April 2026, US lawmakers proposed the "Deterring American AI Model Theft Act" to restrict other countries' companies from distilling US LLMs to train their own LLMs. Core provisions include authorizing the US government to precisely identify overseas entities that use "query-copy" technology to extract the US's AI model capabilities, with a focus on institutions related to China and Russia, authorizing the Ministry of Commerce to include violating entities in the export control "entity list", and may

invoke the International Emergency Economic Powers Law to impose financial sanctions and freeze assets.

Fig. 72: Major timeline of updates on BIS entity list for China AI model and hardware companies, 2025-2026

Date	Event Description	Key Targets
Jan-25	BIS issued two Final Rules adding 25 Chinese companies to the Entity List	Includes Zhipu AI, Beijing Keyi Hongyuan Optoelectronics, etc.
Mar-25	BIS added 50+ Chinese tech firms; 12 specifically for AI model development, servers and AI computing	Beijing Zhiyuan AI Research Institute, Nettrix, Inspur, etc.
Sep-25	BIS added 23 Chinese companies in semiconductor, biotech, quantum tech and aerospace	Shanghai Fudan Microelectronics, Aerospace Information Research Institute (CAS), National Time Service Center (CAS), etc.
May-26	BIS issued guidance blocking Chinese firms from acquiring advanced AI chips via overseas subsidiaries	Targets Chinese parent companies acquiring advanced AI chips such as NVIDIA Rubin/Blackwell series and AMD MI350x

Source: Bureau of Industry and Security, Nomura research

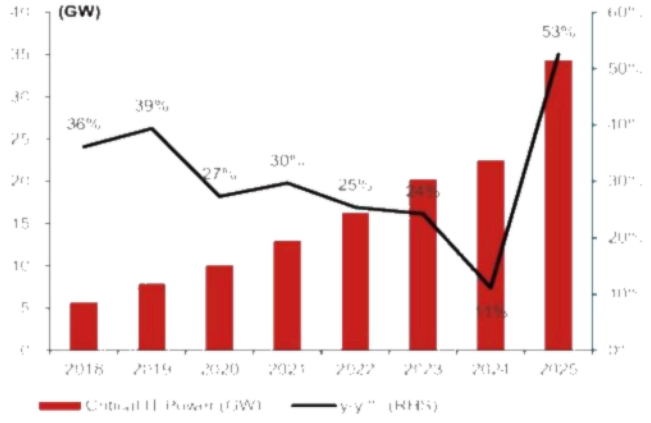
The performance of the latest LLMs from Chinese companies such as DeepSeek, Moonshot, and Zhipu reflects that despite limited access to advanced chips, China has achieved competitive LLM performance with limited computing power through multiple innovations (such as MLA attention mechanisms and FP8 mixed-precision training). Meanwhile, US chip export controls have accelerated the commercialization progress of China's domestic chips. The insufficient performance of domestic chips per card led to innovative solutions such as SuperPod, which scales up AI chip clusters effectively, making domestic chips more usable for training and inference.

We believe that the China-US AI industry competition has entered a new stage of full-chain, multi-dimensional development. The US builds a blockade system from chips to models, from hardware to software to maintain its leading position, while China focuses on open-source ecosystems and domestic substitution as the main development strategies. From a risk perspective, the supply shortage may intensify and cost may increase due to various restrictions on different parts of the value chain (i.e. chips, optical components, LLMs). But in the medium to long term, regulations are accelerating the autonomy and the resilience of China's AI industry chain.

China IDC: robust demand growth driven by surging token consumption

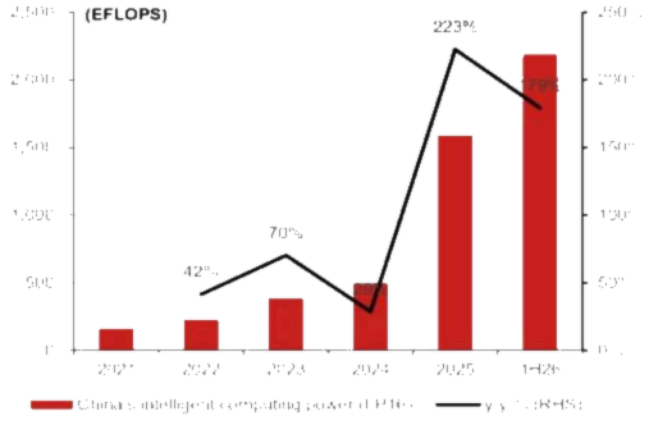
According to NDA and CAICT, China's cumulative number of racks for IDCs grew 53% y-y to 13.7mn by end-2025, equivalent to critical IT power of 34GW assuming standard rack density of 2.5kW per rack. On the other hand, China's cumulative intelligent computing power (FP16) recorded 223% y-y growth to 1590 EFLOPS by end-2025, followed by another 179% y-y to 2185 EFLOPS by the end of June 2026 (*news link in Chinese*). We see robust demand for computing power in China, driven by surging token consumption in China, which is fueling growing demand for AIDC infrastructure rollout, in our view.

Fig. 73: China's cumulative IDC critical IT power



Note: Based on equivalent density of 2.5kW per rack
Source: MIIT, CAICT, Nomura research

Fig. 74: China's intelligent computing power (FP16) trend



Source: MIIT, CAICT, Nomura research

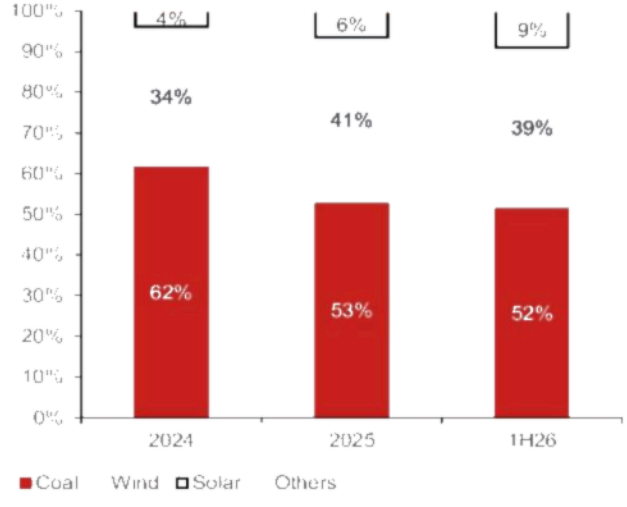
On the other hand, we note the Chinese government has enforced stricter controls over the domestic AIDC construction since early 2025, by introducing new centralized project approval procedures by NDRC, which requires large- and mega-sized AIDC projects to locate at designated hub nodes under the "Eastern Data Western Computing" (EDWC) initiative. According to Government Work Report 2026, the task of "implementing new infrastructure projects includes ultra-large computing clusters and coordination of power and computing resources (算电协同)" , and these have been included in the 15th Five-Year-Plan. We expect domestic AIDC planning to factor in more considerations of: 1) penetration of green power, with more policy support on "direct renewable connection" (绿电直连) and power purchase agreements (PPAs); and 2) penetration of domestic AI hardware for training and inferencing purposes, with lower dependence on imported chips.

Case study: accelerating AIDC rollout in Ulanqab

Ulanqab is located in Inner Mongolia, one of the eight computing power hub nodes under the EDWC initiative. Ulanqab has become one of the pioneer cities in China in terms of AIDC development, likely attributable to: 1) adequate natural cooling resource with annual average temperature of 4-5°C, leading to better PUE (power usage effectiveness) below 1.2; 2) convenient location, just ~240km from Beijing, allowing lower network latency and synergy with data center clusters in Beijing and Hebei; and 3) concentrated renewable energy installation with attractive power tariff.

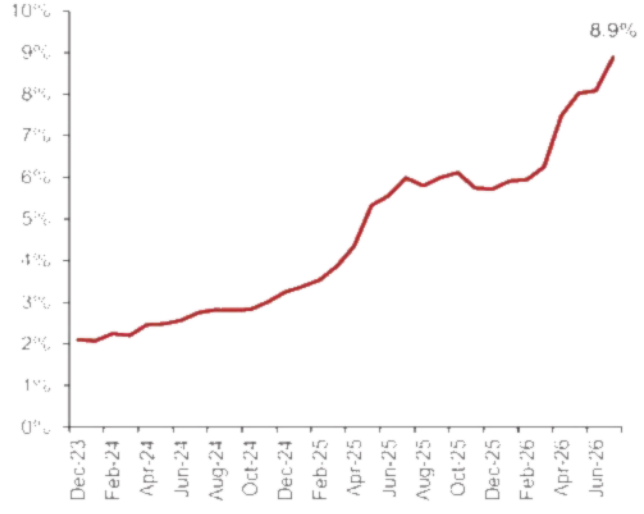
According to Ulanqab's Municipal Development and Reform Commission (MDRC), the city achieved over cumulative renewable power installation of over 20GW (wind: 15GW, solar: 5.3GW), representing 67% of total power generation capacity by end-1H26. Renewable power generation totalled 19.6TWh in 1H26, accounting for 48% of the total. AIDC has become the key economic growth driver for Ulanqab, with total computing power reaching ~165EFLOPS by mid-2026, implying a national market share of ~8%, based on our estimates. By July 2026, the monthly electricity consumption by data centers reached 732mn kWh, accounting for 8.9% of total power consumption (vs national level of ~2%). We estimate utilized IT power of 0.8GW by July 2026 in Ulanqab, while the city's government indicated that contracted capacity reached 5mn standard racks by mid-2026, equivalent to critical IT power of 12.5GW. Incumbent players in Ulanqab include VNET (VNET US, Buy), Centrin (unlisted), ZData (unlisted), UCloud (688158 CH, not rated), and we see more emerging players announcing AIDC plans in the city, including GDS (GDS US, Buy, *news link*), GLP (unlisted, *news link*) as well as potential projects by DeepSeek (unlisted, *news link*) and RedNote (unlisted, *news link*).

Fig. 75: Ulanqab's power generation by energy resource



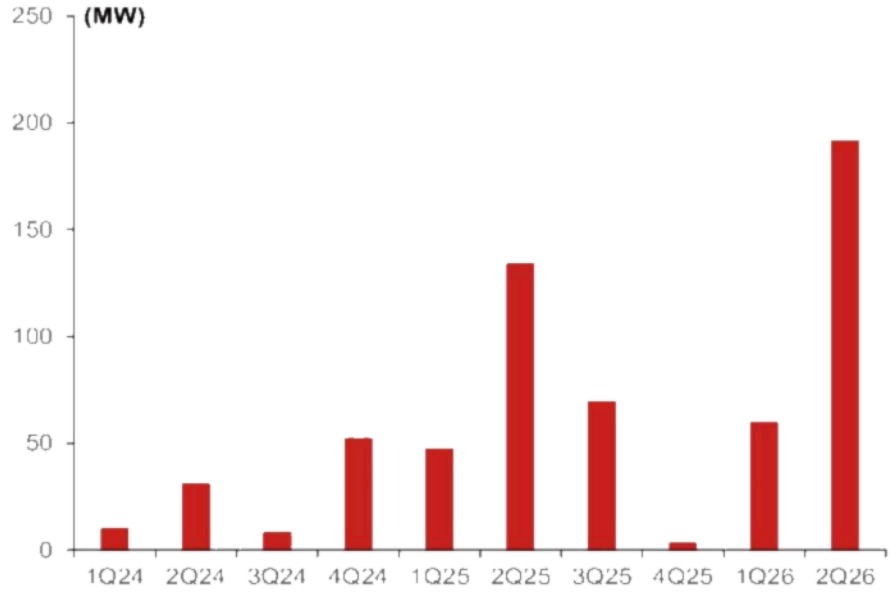
Source: Ulanqab MDRC, Nomura research

Fig. 76: Ulanqab's data center power consumption % total



Source: Ulanqab MDRC, Nomura research

Fig. 77: Net addition of utilized IT power for Ulanqab's data center industry



Note: Implied by data center power consumption, assuming an average PUE of 1.2
Source: Ulanqab MDRC, Nomura research

NOMURA

Global Markets Research
11 September 2026

Kingdee 0268.HK 268 HK

EQUITY: TECHNOLOGY

AI business enjoys solid growth momentum...

...Core ERP business continues to improve margins

Action: Maintain Buy and TP of HKD11.87; implying 45.6% upside

Kingdee reported 1H26 results on 11 Aug; top line growth accelerated to 13.6% y-y to CNY3.6bn (from +12% y-y in FY25) and net profit was at CNY54.5mn (vs. a net loss of CNY97.7mn in 1H25). Cloud services revenue grew 16.6% y-y to CNY3.1bn, and subscription revenue rose 20.4% y-y to CNY2.0bn. **Subscription annual recurring revenue (ARR) was at ~CNY4.4bn (+18.3% y-y)**, mainly driven by the growth of Kingdee AI Suite and the robust demand from customers in certain industries. GPM maintained its upward trajectory, with a **2.1pp y-y improvement to 67.7%**, due to the continued structural shift to subscription and high-value AI products, in our view. We think Kingdee is on the right track to transform into an AI native company, and AI agents are becoming its new growth engines. Kingdee reported 1H26 **AI+SaaS products** (AI Constellation, AI Galaxy, AI Stellar, etc.) **revenue of CNY2.8bn, +10.3% y-y, while AI native products** (Lingee and Kingdee AI Suite) **recorded a revenue of CNY296mn, +189.0% y-y.** We maintain our Buy rating with DCF-based TP of HKD11.87, implying 42.9x FY27F EPS of CNY0.24. The stock currently trades at 29.5x FY27F P/E.

Launch of AI OS marking Kingdee's further AI transformation; long-term AI target on track

In May-26, Kingdee launched enterprise AI operating system Lingee that has adopted a "seat + token points" sales model since July, with seat prices of CNY58-1,168/ month/seat and points charging based on token usage and product capability. Management maintains its "2030 target" of 50%/50% revenue contribution by AI+SaaS/AI native solutions in 2030E. We expect AI-native products including AI agents and AI OS to contribute higher revenue in 2H26E than in 1H26.

FY26E outlook: double-digit revenue growth; AI-native product revenue exceeding CNY1bn

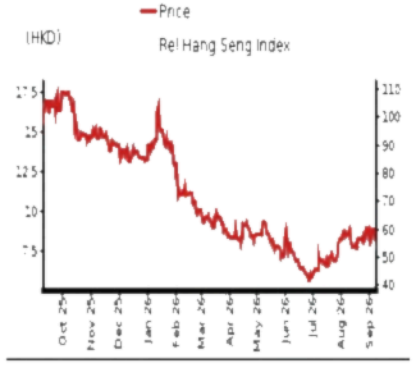
Management reiterated its full-year guidance in 1H26 earnings call: 1) total revenue to maintain double-digit growth; 2) AI-native product revenue to exceed CNY1bn in FY26E, and Lingee aims to sign 10,000 contracts by FY26E; 3) subscription revenue to maintain ~20% y-y growth; 4) overseas revenue to maintain rapid growth; 5) adjusted NPM target of ~7%; and 6) OCF to increase by more than 20% y-y.

Year-end 31 Dec	FY25		FY26F		FY27F		FY28F
Currency (CNY)	Actual	Old	New	Old	New	Old	New
Revenue (mn)	7,006	7,885	7,885	9,102	9,102	10,831	10,831
Reported net profit (mn)	93	450	450	846	846	1,426	1,426
Normalised net profit (mn)	90	100	100	97	97	89	89
FD normalised EPS	2.55c	2.81c	2.81c	2.74c	2.74c	2.52c	2.52c
FD norm. EPS growth (%)	–	10.1	10.1	-2.5	-2.5	-8.1	-8.1
FD normalised P/E (x)	274.8	–	249.6	–	256.1	–	278.8
EV/EBITDA (x)	32.5	–	31.3	–	25.7	–	18.2
Price/book (x)	3.1	–	3.0	–	2.7	–	2.3
Dividend yield (%)	–	–	–	–	–	–	–
ROE (%)	1.1	5.3	5.3	9.3	9.3	14.0	14.0
Net debt/equity (%)	net cash	net cash	net cash	net cash	net cash	net cash	net cash

Source: Company data, Nomura estimates

Rating Remains	**Buy**
Target price Remains	**HKD 11.87**
Closing price 8 September 2026	**HKD 8.15**
Implied upside	**+45.6%**
Market Cap (USD mn)	3,653.0
ADT (USD mn)	51.7

Relative performance chart



Source: LSEG, Nomura

Research Analysts

China Technology

Bing Duan - NIHK
bing.duan1@nomura.com
+852 2252 2141

Ethan Zhang - NIHK
ethan.zhang@nomura.com
+852 2252 2157

Key data on Kingdee

Performance

(%)	1M	3M	12M		
Absolute (HKD)	-4.6	8.5	-46.8	M cap (USDmn)	3,653.0
Absolute (USD)	-4.6	8.4	-47.1	Free float (%)	80.0
Rel to Hang Seng Index	-3.3	5.8	-45.5	3-mth ADT (USDmn)	51.7

Income statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Revenue	6,256	7,006	7,885	9,102	10,831
Cost of goods sold	-2,185	-2,302	-2,433	-2,717	-3,233
Gross profit	4,071	4,704	5,452	6,385	7,597
SG&A	-4,406	-4,667	-5,316	-6,058	-6,991
Employee share expense					
Operating profit	-335	37	136	327	606
EBITDA	436	708	664	757	963
Depreciation	-771	-670	-528	-430	-357
Amortisation					
EBIT	-335	37	136	327	606
Net interest expense	106	2	81	81	81
Associates & JCEs					
Other income					
Earnings before tax	-230	39	217	408	687
Income tax	-45	-35	-192	-361	-609
Net profit after tax	-275	4	25	47	79
Minority interests	-42	-19	-41	-77	-130
Other items	140	105	116	127	140
Preferred dividends					
Normalised NPAT	-177	90	100	97	89
Extraordinary items	35	2	350	749	1,337
Reported NPAT	-142	93	450	846	1,426
Dividends	0	0	0	0	0
Transfer to reserves	-142	93	450	846	1,426
Valuations and ratios					
Reported P/E (x)	–	267.6	55.3	29.4	17.4
Normalised P/E (x)	-136.6	274.8	249.6	256.1	278.8
FD normalised P/E (x)	–	274.8	249.6	256.1	278.8
Dividend yield (%)	–	–	–	–	–
Price/cashflow (x)	25.9	22.7	18.6	15.2	26.8
Price/book (x)	3.2	3.1	3.0	2.7	2.3
EV/EBITDA (x)	53.8	32.5	31.3	25.7	18.2
EV/EBIT (x)	–	614.1	153.1	59.6	29.0
Gross margin (%)	65.1	67.1	69.1	70.1	70.1
EBITDA margin (%)	7.0	10.1	8.4	8.3	8.9
EBIT margin (%)	-5.4	0.5	1.7	3.6	5.6
Net margin (%)	-2.3	1.3	5.7	9.3	13.2
Effective tax rate (%)	–	88.5	88.5	88.5	88.5
Dividend payout (%)	–	0.0	0.0	0.0	0.0
ROE (%)	-1.9	1.1	5.3	9.3	14.0
ROA (pretax %)	-3.9	0.3	1.1	2.9	5.3
Growth (%)					
Revenue		12.0	12.5	15.4	19.0
EBITDA		62.3	-6.2	14.0	27.3
Normalised EPS		–	10.1	-2.5	-8.1
Normalised FDEPS		–	10.1	-2.5	-8.1

Source: Company data, Nomura estimates

Cashflow statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
EBITDA	436	708	664	757	963
Change in working capital	1,769	337	392	428	463
Other operating cashflow	-1,272	53	279	449	-497
Cashflow from operations	934	1,097	1,335	1,634	929
Capital expenditure	-681	-587	-411	-349	-297
Free cashflow	253	511	924	1,285	633
Reduction in investments	0				
Net acquisitions					
Dec in other LT assets	0				
Inc in other LT liabilities	0				
Adjustments	-393	-21	1,355	71	1,316
CF after investing acts	-140	490	2,280	1,355	1,949
Cash dividends	0	0	0	0	0
Equity issue	0	0	0	0	0
Debt issue	0	0	0	0	0
Convertible debt issue					
Others	1,670	-325	0	0	0
CF from financial acts	1,670	-325	0	0	0
Net cashflow	1,530	165	2,280	1,355	1,949
Beginning cash	1,530	1,695	3,975	5,330	
Ending cash	1,530	1,695	3,975	5,330	7,279
Ending net debt	-2,944	-1,695	-3,975	-5,330	-7,279

Balance sheet (CNYmn)

As at 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Cash & equivalents	1,530	1,695	3,975	5,330	7,279
Marketable securities	0	0	0	0	0
Accounts receivable	416	412	392	372	353
Inventories	46	61	80	106	140
Other current assets	3,640	3,655	3,670	3,685	3,700
Total current assets	5,632	5,823	8,116	9,493	11,472
LT investments	0	0	0	0	0
Fixed assets	1,744	1,874	436	355	295
Goodwill					
Other intangible assets	1,232	1,433	1,433	1,433	1,433
Other LT assets	5,173	5,457	5,457	5,457	5,457
Total assets	13,781	14,586	15,442	16,737	18,657
Short-term debt	205	0	0	0	0
Accounts payable	188	144	110	84	64
Other current liabilities	5,086	5,494	5,934	6,409	6,923
Total current liabilities	5,480	5,638	6,044	6,493	6,987
Long-term debt	0	0	0	0	0
Convertible debt	0	0	0	0	0
Other LT liabilities	172	740	740	740	740
Total liabilities	5,652	6,377	6,783	7,233	7,726
Minority interest	79	1	1	1	1
Preferred stock	0	0	0	0	0
Common stock	86	85	85	85	85
Retained earnings	7,964	8,122	8,572	9,418	10,844
Proposed dividends					
Other equity and reserves	0	0	0	0	0
Total shareholders' equity	8,050	8,207	8,657	9,504	10,929
Total equity & liabilities	13,781	14,586	15,442	16,737	18,657
Liquidity (x)					
Current ratio	1.03	1.03	1.34	1.46	1.64
Interest cover	–	–	–	–	–
Leverage					
Net debt/EBITDA (x)	net cash	net cash	net cash	net cash	net cash
Net debt/equity (%)	net cash	net cash	net cash	net cash	net cash
Per share					
Reported EPS (CNY)	-4.12c	2.62c	12.69c	23.86c	40.21c
Norm EPS (CNY)	-5.14c	2.55c	2.81c	2.74c	2.52c
FD norm EPS (CNY)	-5.14c	2.55c	2.81c	2.74c	2.52c
BVPS (CNY)	2.27	2.31	2.44	2.68	3.08
DPS (CNY)	0.00	0.00	0.00	0.00	0.00
Activity (days)					
Days receivable	21.6	18.6	15.3	12.3	
Days inventory	8.5	10.6	12.5	14.0	
Days payable	26.3	19.0	13.0	8.3	
Cash cycle	0.0	3.7	10.2	14.9	17.9

Source: Company data, Nomura estimates

Company profile

Kingdee International is a China-based software company. The company mainly engages in the enterprise resource planning (ERP) business. Kingdee has launched its cloud-based ERP platforms, namely Kingdee Cloud Galaxy and Kingdee Cloud Cosmic to facilitate digital transformation of its enterprise customers.

Valuation Methodology

We use DCF valuation to arrive at our TP of HKD11.87. Our DCF valuation assumes a WACC of 10.1%, market risk premium of 7% and a terminal growth rate of 2.0%, with cash flows discounted back to end-2026F. The benchmark for this stock is the Hang Seng Index.

Risks that may impede the achievement of the target price

Downside risks include: 1) slower-than-expected cloud revenue growth; 2) the company being unable to turnaround its cloud business due to competition or additional R&D investments; and 3) intense competition in the next generation of cloud platforms from incumbent ERP players.

ESG

Kingdee is a leading cloud software solution provider, which is environmentally friendly and takes social responsibility by helping enterprise and industrial clients for digital transformation.

NOMURA

Global Markets Research
11 September 2026

Kingsoft Office 688111.SS 688111 CH

EQUITY: TECHNOLOGY

Aspiring to become an AI office platform...

...Leveraging a strong enterprise ecosystem to compete with large LLM players

Action: Maintain Buy; cut TP to CNY336, implying 42.8% upside

Kingsoft Office (KO) reported 1H26 results on 19 August, with top line recording 24.7% y-y growth (2Q26: +25.4% y-y), driven by the healthy growth of enterprise and government businesses, in our view. 1H26 earnings recorded 237% y-y growth due to the returns from investment fund projects. KO launched Lingxi Agent (for individuals) and WPS Comate (for enterprises) in July, and we think the AI agentic products should help the company solidify its leadership position in China's AI office software market, and underpin better AI monetization opportunities. We slightly lower our FY26-27F revenue and earnings forecasts by 1.0-1.8% and 1.0-1.4%, respectively, to reflect slower-than-expected ARPU growth of WPS AI. We also introduce FY28F revenue/earnings forecasts that suggest 26%/28% y-y growth. We maintain our Buy rating and cut DCF-based TP to CNY336 (from CNY434), assuming a WACC of 8.4% (previously 9.0%), market risk premium of 7% (unchanged) and terminal growth rate of 2.5% (previously 3.5%). Our TP implies 55.8x FY27F EPS of CNY6.03, or 17x FY27F P/S. The stock is trading at 39x FY27F EPS.

Upgrading from collaborative office to an AI office platform; token consumption likely to increase further

KO released two AI office agents in July 2026: Lingxi Professional for individuals and WPS Comate for organizations. As the AI entry point for WPS 365, WPS Comate enables enterprise-level AI capabilities and AI governance to collaborate on the same platform. The agentic platform may stimulate higher volume of token consumption, potentially accelerating ARPU growth in the enterprise segment, in our view. We think WPS 365 will maintain solid growth over FY26-28F and estimate an FY26-28F revenue CAGR of 52%; however, we estimate FY26-28F revenue CAGRs of 18%/24% for WPS Individual/WPS Software.

Risk factor: AI office sector becoming increasingly competitive

As potential competitors to KO, large language model (LLM) players are also developing AI office and collaboration products. Alibaba (BABA US, Buy, covered by Jialong Shi) established "Qwen Office Business Unit" and launched Qwen Work in July, aiming to capture the enterprise-level AI office market, by integrating agent products, including QoderWork, Wukong, and MuleRun. Tencent (700 HK, Buy, covered by Jialong Shi) launched an AI collaboration product WorkBuddy in March, which gained impressive traction with white-collar professionals, reaching 20.97mn monthly PC visits in June (*news*). We think KO faces a more challenging environment in China's office software market.

Year-end 31 Dec	FY25	FY26F		FY27F			FY28F
Currency (CNY)	Actual	Old	New	Old	New	Old	New
Revenue (mn)	5,929	7,324	7,254	9,253	9,088	0	11,413
Reported net profit (mn)	1,836	2,257	2,234	2,833	2,792	0	3,586
Normalised net profit (mn)	1,836	2,257	2,234	2,833	2,792	0	3,586
FD normalised EPS	3.37	4.88	4.82	6.13	6.03		7.74
FD norm. EPS growth (%)	15.3	22.3	43.1	25.5	25.0		28.4
FD normalised P/E (x)	69.8	–	48.8	–	39.0	–	30.4
EV/EBITDA (x)	78.7	–	62.8	–	47.6	–	35.7
Price/book (x)	8.5	–	7.6	–	6.7	–	5.8
Dividend yield (%)	0.5	–	0.6	–	0.8	–	1.0
ROE (%)	15.2	16.6	16.4	18.3	18.2		20.5
Net debt/equity (%)	net cash	net cash	net cash	net cash	net cash		net cash

Source: Company data, Nomura estimates

Rating Remains	**Buy**
Target price Reduced from CNY 434.00	**CNY 336.00**
Closing price 8 September 2026	**CNY 235.27**
Implied upside	**+42.8%**
Market Cap (USD mn)	16,169.9
ADT (USD mn)	349.6

Relative performance chart



Source: LSEG, Nomura

Research Analysts

China Technology

Bing Duan - NIHK
bing.duan1@nomura.com
+852 2252 2141

Joel Ying, CFA - NIHK
joel.ying@nomura.com
+852 2252 2153

Key data on Kingsoft Office

Performance

(%)	1M	3M	12M		
Absolute (CNY)	-11.4	5.1	-21.9	M cap (USDmn)	16,169.9
Absolute (USD)	-11.0	6.3	-17.0	Free float (%)	18.0
Rel to CSI 300	-8.5	8.4	-24.0	3-mth ADT (USDmn)	349.6

Income statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Revenue	5,121	5,929	7,254	9,088	11,413
Cost of goods sold	-761	-833	-1,005	-1,227	-1,466
Gross profit	4,360	5,096	6,249	7,861	9,946
SG&A	-3,139	-3,773	-4,631	-5,757	-7,206
Employee share expense					
Operating profit	1,221	1,322	1,618	2,104	2,740
EBITDA	1,266	1,370	1,685	2,171	2,807
Depreciation	-44	-48	-66	-67	-67
Amortisation					
EBIT	1,221	1,322	1,618	2,104	2,740
Net interest expense	16	0	-1	-1	-1
Associates & JCEs					
Other income	513	631	757	931	1,157
Earnings before tax	1,750	1,953	2,375	3,034	3,896
Income tax	-95	-131	-159	-264	-339
Net profit after tax	1,655	1,822	2,216	2,770	3,557
Minority interests	-10	15	18	22	29
Other items					
Preferred dividends					
Normalised NPAT	1,645	1,836	2,234	2,792	3,586
Extraordinary items					
Reported NPAT	1,645	1,836	2,234	2,792	3,586
Dividends	-416	-580	-705	-882	-1,132
Transfer to reserves	1,229	1,256	1,528	1,910	2,453
Valuations and ratios					
Reported P/E (x)	80.5	69.8	48.8	39.0	30.4
Normalised P/E (x)	80.5	69.8	48.8	39.0	30.4
FD normalised P/E (x)	80.5	69.8	48.8	39.0	30.4
Dividend yield (%)	0.4	0.5	0.6	0.8	1.0
Price/cashflow (x)	60.3	51.3	32.7	28.1	22.5
Price/book (x)	9.6	8.5	7.6	6.7	5.8
EV/EBITDA (x)	85.2	78.7	62.8	47.6	35.7
EV/EBIT (x)	88.3	81.6	65.4	49.1	36.6
Gross margin (%)	85.1	85.9	86.2	86.5	87.2
EBITDA margin (%)	24.7	23.1	23.2	23.9	24.6
EBIT margin (%)	23.9	22.3	22.3	23.2	24.0
Net margin (%)	32.1	31.0	30.8	30.7	31.4
Effective tax rate (%)	5.4	6.7	6.7	8.7	8.7
Dividend payout (%)	25.3	31.6	31.6	31.6	31.6
ROE (%)	15.4	15.2	16.4	18.2	20.5
ROA (pretax %)	9.6	8.1	9.1	11.4	14.3
Growth (%)					
Revenue	12.4	15.8	22.4	25.3	25.6
EBITDA	25.1	8.2	23.0	28.9	29.3
Normalised EPS	2.3	15.3	43.1	25.0	28.4
Normalised FDEPS	2.3	15.3	43.1	25.0	28.4

Source: Company data, Nomura estimates

Cashflow statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
EBITDA	1,266	1,370	1,685	2,171	2,807
Change in working capital	1,711	-4,752	184	144	270
Other operating cashflow	-782	5,883	1,461	1,558	1,759
Cashflow from operations	2,196	2,502	3,330	3,873	4,836
Capital expenditure	-175	-98	-98	-98	-98
Free cashflow	2,021	2,404	3,232	3,775	4,738
Reduction in investments	0	0	0	0	
Net acquisitions					
Dec in other LT assets	0	0	0	0	
Inc in other LT liabilities	213	347	263	317	353
Adjustments	-4,312	-2,216	-763	-817	-853
CF after investing acts	-2,079	535	2,732	3,275	4,238
Cash dividends	-406	-482	-705	-882	-1,132
Equity issue	62	66	0	0	0
Debt issue	0	0	-22	0	0
Convertible debt issue					
Others	-157	-322	-18	-22	-29
CF from financial acts	-502	-738	-745	-904	-1,161
Net cashflow	-2,580	-204	1,987	2,371	3,077
Beginning cash	3,478	898	694	2,681	5,053
Ending cash	898	694	2,681	5,053	8,130
Ending net debt	-802	-672	-2,681	-5,053	-8,130

Balance sheet (CNYmn)

As at 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Cash & equivalents	898	694	2,681	5,053	8,130
Marketable securities					
Accounts receivable	594	743	930	1,164	1,458
Inventories	1	1	1	1	1
Other current assets	4,612	9,769	10,269	10,769	11,269
Total current assets	6,105	11,206	13,881	16,987	20,858
LT investments					
Fixed assets	432	453	459	459	450
Goodwill	79	73	52	31	11
Other intangible assets	9,297	6,423	6,386	6,386	6,386
Other LT assets					
Total assets	15,913	18,156	20,779	23,863	27,704
Short-term debt	96	22	0	0	0
Accounts payable	1,124	1,249	1,612	1,884	2,282
Other current liabilities	2,243	2,671	3,180	3,787	4,452
Total current liabilities	3,462	3,942	4,792	5,670	6,734
Long-term debt	0	0	0	0	0
Convertible debt					
Other LT liabilities	1,014	1,361	1,624	1,941	2,294
Total liabilities	4,476	5,302	6,416	7,611	9,028
Minority interest	81	0	-18	-40	-68
Preferred stock					
Common stock	463	463	463	463	463
Retained earnings	10,893	12,390	13,918	15,829	18,282
Proposed dividends					
Other equity and reserves					
Total shareholders' equity	11,356	12,853	14,382	16,292	18,745
Total equity & liabilities	15,913	18,156	20,779	23,863	27,704
Liquidity (x)					
Current ratio	1.76	2.84	2.90	3.00	3.10
Interest cover	–	2,704.8	2,705.2	2,807.6	2,911.2
Leverage					
Net debt/EBITDA (x)	net cash	net cash	net cash	net cash	net cash
Net debt/equity (%)	net cash	net cash	net cash	net cash	net cash
Per share					
Reported EPS (CNY)	2.92	3.37	4.82	6.03	7.74
Norm EPS (CNY)	2.92	3.37	4.82	6.03	7.74
FD norm EPS (CNY)	2.92	3.37	4.82	6.03	7.74
BVPS (CNY)	24.55	27.75	31.05	35.17	40.47
DPS (CNY)	0.90	1.25	1.52	1.90	2.44
Activity (days)					
Days receivable	44.0	41.1	42.1	42.1	42.0
Days inventory	0.5	0.3	0.3	0.3	0.3
Days payable	531.3	519.8	519.8	519.8	519.8
Cash cycle	-486.8	-478.4	-477.4	-477.5	-477.5

Source: Company data, Nomura estimates

Company profile

Kingsoft Office is a leading office software and service provider in China. The company focuses on development of WPS office and relevant applications, while also providing one-stop, multi-platform application solutions helping customers moving onto cloud. The company's products and services apply to the government and various industries, as well as consumers on both PC and mobile platforms.

Valuation Methodology

Our TP of CNY336.00 is based on DCF valuation method. Our DCF valuation assumes a WACC of 8.4%, market risk premium of 7% and terminal growth rate of 2.5% while we discounted back cash flow to 2026 as a 12-month forward valuation. Our TP implies a 22x/17x FY26/27F P/S or 69.8x/55.8x FY26/27F P/E. The benchmark index for this stock is CSI 300.

Risks that may impede the achievement of the target price

Downside risks include: 1) weaker-than-expected demand from independent and controllable technology initiative in the China market; 2) slower-than-expected development of the WPS AI and subscription business; and 3) increasing competition for the collaboration business in the SaaS market.

ESG

As a software and service company, Kingsoft Office assists customers to improve office efficiency and efficiently produce professional and high quality office documents. Its continuous cloud business strategy will help end-customers fulfill their responsibility in energy conservation and emission reduction. The company has also established a good platform for employees and advocates a healthy work and lifestyle. KO also communicates with the communities where it operates and practices public welfare undertakings.

1H26 and 2Q26 results review

Kingsoft Office (KO) reported 1H26 results on 19 August, and recorded 1H26 revenue/ earnings growth of 25%/237% y-y, translating into changes of +25%/-6% y-y for 2Q26. We note KO's top-line maintained strong growth in 2Q26, driven by solid **growth in WPS 365 and WPS software, +60% and +34% y-y in 2Q26, respectively**. The significant increase in earnings for 1H26 was mainly driven by a sizable increase in investment income in 1Q26, and we believe both GPM and OPM were still under pressure, likely due to increased hardware cost and S&M expenses, as well as continued investment in AI capabilities, in our view. 1H26 GPM came in at 84.9%, -0.1pp y-y (2Q26: 83.3%, -0.9pp y-y), while OPM increased 1.1pp y-y (0.1pp q-q) to 22.9% in 1H26 (2Q26: +1.4pp y-y and -4.4pp q-q to 25.1%).

Fig. 78: Kingsoft Office – 1H26 and 2Q26 earnings review

(CNY mn)	1Q24	2Q24	3Q24	4Q24	1Q25	2Q25	3Q25	4Q25	1Q26	2Q26	y-y	q-q	1H24	2H24	1H25	2H25	1H26	y-y	h-h
	A	A	A	A	A	A	A	A	A	A			A	A	A	A	A		
Turnover	1,225	1,188	1,214	1,494	1,301	1,356	1,521	1,750	1,613	1,700	25%	5%	2,413	2,708	2,657	3,272	3,313	25%	1%
% chg yoy	17%	6%	11%	16%	6%	14%	25%	17%	24%	25%			11%	14%	10%	21%	25%		
COGS	(186)	(178)	(176)	(221)	(183)	(215)	(206)	(229)	(216)	(284)			(364)	(397)	(398)	(435)	(500)		
Gross Profit	1,039	1,010	1,037	1,274	1,119	1,141	1,315	1,521	1,397	1,416			2,049	2,311	2,259	2,836	2,813		
OPEX	(743)	(743)	(805)	(848)	(803)	(879)	(964)	(1,127)	(992)	(1,064)			(1,486)	(1,653)	(1,682)	(2,091)	(2,056)	22%	-2%
Operating Income	**296**	**267**	**233**	**425**	**315**	**262**	**351**	**394**	**405**	**352**	**34%**	**-13%**	**563**	**658**	**577**	**745**	**757**	**31%**	**2%**
% chg yoy	51%	3%	20%	35%	6%	-2%	51%	-7%	29%	34%			24%	29%	-2%	13%	31%		
Finance Expense	14	3	1	-	-	3	(2)	(2)	1	(1)			17		4	-	(0)		
Other gains expense	79	109	99	225	104	95	114	319	1,999	8			189	324	198	428	2,007	913%	369%
Pretax income	**389**	**379**	**332**	**650**	**419**	**361**	**463**	**710**	**2,406**	**359**	**-1%**	**-85%**	**769**	**982**	**780**	**1,173**	**2,765**	**254%**	**136%**
% chg yoy	38%	9%	8%	44%	8%	-5%	39%	9%	474%	-1%			22%	29%	1%	19%	254%		
Tax	(24)	(23)	(14)	(35)	(25)	(22)	(31)	(52)	(211)	(36)			(47)	(49)	(47)	(84)	(247)		
Minority expenses	2	(2)	0	(10)	8	6	0	0	0	0			(1)	(10)	15	0	0		
Profit to shareholders	**367**	**354**	**318**	**605**	**403**	**344**	**431**	**658**	**2,195**	**322**	**-6%**	**-85%**	**721**	**924**	**747**	**1,089**	**2,518**	**237%**	**131%**
% chg yoy	37%	7%	8%	43%	10%	-3%	35%	9%	445%	-6%			20%	29%	4%	18%	237%		
EPS (Diluted)	0.79	0.77	0.69	1.31	0.87	0.74	0.93	1.42	4.74	0.69			1.56	2.00	1.61	2.35	5.44		
Ratio Analysis																			
GPM	84.8%	85.0%	85.5%	85.2%	86.0%	84.1%	86.5%	86.9%	86.6%	83.3%	-0.9pp	-3.3pp	84.9%	85.3%	85.0%	86.7%	84.9%	-0.1pp	-1.8pp
OPEX	60.6%	62.5%	66.3%	56.8%	61.7%	64.8%	63.4%	64.4%	61.5%	62.6%	-2.2pp	1.1pp	61.6%	61.0%	63.3%	63.9%	62.0%	-1.3pp	-1.9pp
OPM	24.2%	22.5%	19.2%	28.5%	24.2%	19.3%	23.1%	22.5%	25.1%	20.7%	1.4pp	-4.4pp	23.4%	24.3%	21.7%	22.8%	22.9%	1.1pp	0.1pp
NPM	30.0%	29.8%	26.2%	40.5%	31.0%	25.4%	28.3%	37.6%	136.1%	19.0%	6.4pp	117.1pp	29.9%	34.1%	28.1%	33.3%	76.0%	47.9pp	42.7pp

Source: Company data, Nomura research

By segment, KO recorded revenue growth of 15%/61%/33% y-y for WPS individual/WPS 365/WPS software in 1H26 (2Q26: 17%/60%/34%), accounting for 61%/15%/22% of total revenue. Total number of paying users increased 15% y-y to CNY48.3mn. WPS Office MAD (monthly active device) recorded an increase of 3.9% y-y to CNY676mn, of which PC MAD grew 10.1% y-y to CNY336mn and Mobile MAD declined 1.7% y-y to CNY340mn.

Fig. 79: Kingsoft Office – 1H26 and 2Q26 results by segment

(CNY mn)	1Q24	2Q24	3Q24	4Q24	1Q25	2Q25	3Q25	4Q25	1Q26	2Q26	y-y	q-q	1H24	2H24	1H25	2H25	1H26	y-y	h-h
Revenue																			
By Product																			
WPS Individual	773	840	809	861	857	891	899	978	975	1,039	17%	7%	1,613	1,670	1,748	1,877	2,014	15%	7%
% chg yoy	25%	21%	18%	10%	11%	6%	11%	14%	14%	17%					8%	12%	15%		
WPS 365	93	97	117	129	151	158	201	210	244	253	60%	4%	190	247	309	411	497	61%	21%
% chg yoy	417%	411%	22%	200%	62%	63%	72%	63%	62%	60%					63%	67%	61%		
WPS Software	332	221	260	455	262	280	391	529	347	375	34%	8%	553	715	542	920	722	33%	-21%
% chg yoy	10%	38%	4%	16%	21%	27%	50%	16%	32%	34%					2%	29%	33%		
Others	27	28	27	49	31	27	30	32	47	33	23%	-30%	55	77	58	62	80	38%	30%
% chg yoy	-46%	-47%	-39%	-22%	17%	-4%	12%	-35%	50%	23%					6%	-19%	38%		

Source: Company data, Nomura research

Earnings forecast revisions

We slightly revise down our FY26-27F revenue/earnings forecasts by 1.0%-1.8% and 1.0%-1.4% to reflect slower-than-expected user growth of WPS AI. We raise our FY26-27F GPM by 1.4-1.6pp, to reflect higher contributions from WPS 365 and WPS software with higher gross margins. Segment-wise, we now lower FY26-27F revenues forecasts by 3.3-6.4% for the WPS Individual segment, while raising the forecasts by 6.4-9.8% for the WPS 365 segment, to reflect the slower-than-expected growth of WPS AI penetration rate and accelerated growth of WPS 365.

Fig. 80: Kingsoft Office – earnings forecast revisions

(CNY mn)	2022	2023	2024	2025	2026F	2027F	2028F	2026F	2027F	2026F	2027F
	A	A	A	A	(NOM)	(NOM)	(NOM)	(OLD)	(OLD)	% diff	% diff
Turnover	3,885	4,556	5,121	5,929	7,254	9,088	11,413	7,324	9,253	-1.0%	-1.8%
% chg yoy	18%	17%	12%	16%	22%	25%	26%	24%	26%		
COGS	(583)	(670)	(761)	(833)	(1,005)	(1,227)	(1,466)	(1,117)	(1,393)	-10.1%	-11.9%
Gross Profit	3,302	3,886	4,360	5,096	6,249	7,861	9,946	6,207	7,860	0.7%	0.0%
OPEX	(2,576)	(2,921)	(3,139)	(3,773)	(4,631)	(5,757)	(7,206)	(4,541)	(5,690)	2.0%	1.2%
Operating Income	726	965	1,221	1,322	1,618	2,104	2,740	1,667	2,170	-2.9%	-3.0%
% chg yoy	2%	33%	27%	8%	22%	30%	30%	26%	30%		
Finance Expense	13	95	16	(0)	(1)	(1)	(1)	23	29	-102.6%	-102.6%
Other gains/expense	459	330	513	631	757	931	1,157	712	881	6.4%	5.6%
Pretax income	1,198	1,390	1,750	1,953	2,375	3,034	3,896	2,401	3,080	-1.1%	-1.5%
% chg yoy	7%	16%	26%	12%	22%	28%	28%	23%	28%		
Tax	(67)	(75)	(95)	(131)	(159)	(264)	(339)	(131)	(229)	22.0%	15.3%
Minority expenses	(13)	3	(10)	15	18	22	29	(14)	(18)		
Profit to shareholders	1,118	1,318	1,645	1,836	2,234	2,792	3,586	2,257	2,833	-1.0%	-1.4%
% chg yoy	7%	18%	25%	12%	22%	25%	28%	23%	26%		
EPS (Diluted)	1.21	1.43	1.78	2.65	4.82	6.03	7.74	4.88	6.13	-1.2%	-1.6%
			25%	49%	82%	25%	28%				
Ratio Analysis											
GPM	85.0%	85.3%	85.1%	85.9%	86.2%	86.5%	87.2%	84.7%	84.9%	1.4pp	1.6pp
OPEX	66.3%	64.1%	61.3%	63.6%	63.8%	63.3%	63.1%	62.0%	61.5%	1.9pp	1.9pp
OPM	18.7%	21.2%	23.9%	22.3%	22.3%	23.2%	24.0%	22.8%	23.5%	-0.4pp	-0.3pp
NPM	28.8%	28.9%	32.1%	31.0%	30.8%	30.7%	31.4%	30.8%	30.6%	0.0pp	0.1pp

Source: Company data, Nomura estimates

Fig. 81: Kingsoft Office – revenue forecast revisions

(CNY mn)	2022	2023	2024	2025	2026F	2027F	2028F	2026F	2026F	2026F	2026F
Revenue	A	A	A	A	(NOM)	(NOM)	(NOM)	(OLD)	(OLD)	% diff	% diff
By Product											
WPS Individual		2,782	3,283	3,626	4,077	4,797	5,675	4,215	5,128	-3.3%	-6.4%
% chg yoy			18%	10%	12%	18%	18%	16%	22%		
WPS 365		175	437	720	1,151	1,784	2,674	1,083	1,624	6.4%	9.8%
% chg yoy			149%	65%	60%	55%	50%	50%	50%		
WPS Software		1,387	1,268	1,461	1,903	2,383	2,937	1,896	2,369	0.4%	0.6%
% chg yoy			-9%	15%	30%	25%	23%	30%	25%		
Others		210	132	120	120	120	120	132	132	-8.8%	-8.8%
% chg yoy			-37%	-9%	0%	0%	0%	10%	0%		
By Geography											
China	3,761	4,396	4,902	5,606	6,802	8,522	10,700	6,896	8,711	-1.4%	-2.2%
y-y growth	*18%*	*17%*	*11%*	*14%*	*21%*	*25%*	*26%*	23%	26%		
% in total	*97%*	*97%*	*96%*	*95%*	*94%*	*94%*	*94%*	*94%*	*94%*		
Overseas	123	157	218	321	448	562	705	428	541	4.7%	3.8%
y-y growth	*20%*	*28%*	*38%*	*47%*	*40%*	*25%*	*26%*	34%	26%		
% in total	*3%*	*3%*	*4%*	*5%*	*6%*	*6%*	*6%*	*6%*	*6%*		

Source: Company data, Nomura estimates

Our FY26-28F revenue/earnings forecasts are 3-17%/1-15% higher than WIND consensus estimates, as we are more positive on an acceleration of IT demand from government projects and WPS 365.

Fig. 82: Kingsoft Office – Nomura estimates vs consensus

CNY mn	FY2025	FY2026F			FY2027F			FY2028F		
		(NOM)	(CON)	diff	(NOM)	(CON)	diff	(NOM)	(CON)	diff
Revenue	5,929	7,254	7,032	3%	9,088	8,289	10%	11,413	9,718	17%
Pretax Profit	1,953	2,375	2,352	1%	3,034	2,787	9%	3,896	3,313	18%
Net Profit	1,836	2,234	2,212	1%	2,792	2,627	6%	3,586	3,126	15%
EPS	2.65	4.82	4.77	1%	6.03	5.66	6%	7.74	6.74	15%

Source: Company data, WIND, Nomura estimates

Valuation and risks

Our new TP of CNY336 is based on an unchanged DCF valuation methodology. Our DCF valuation assumes a WACC of 8.4% (previously 9.0%), market risk premium of 7.0% (unchanged), and a terminal growth rate of 2.5% (previously 3.5%), while we discount back cash flows to 2026F as a 12-month forward valuation. Our TP implies 21x/17x FY26F/27F P/S or 69.8x/55.8x FY26F/27F P/E.

Downside risks include: (1) weaker-than-expected demand from independent and controllable technology initiatives in the China market; (2) slower-than-expected development of the WPS AI and subscription businesses; and (3) intensifying competition in the AI office sector.

Fig. 83: Kingsoft Office – DCF valuation

CNY mn	2019	2020	2021	2022	2023	2024	2025	2026F	2027F	2028F	2029F	2030F	2031F	2032F	2033F	Terminal
Period							0	0	1	2	3	4	5	6	7	
Revenue	1,580	2,261	3,280	3,885	4,556	5,121	5,929	7,254	9,088	11,413	14,312	17,854	22,075	26,956	32,418	
Profit after tax	401	878	1,041	1,118	1,318	1,645	1,836	2,234	2,792	3,586	4,642	5,930	7,411	9,181	11,122	
Operating cash flow	586	1,514	1,864	1,603	2,058	2,196	2,502	3,330	3,873	4,836	5,909	7,232	8,629	10,266	12,015	
CAPEX	(63)	(55)	(149)	(176)	(166)	(175)	(98)	(98)	(98)	(98)	(98)	(98)	(98)	(98)	(98)	
FCF	522	1,459	1,715	1,427	1,892	2,021	2,404	3,232	3,775	4,738	5,812	7,134	8,531	10,168	11,917	206,329
Discount factor									0.92	0.85	0.78	0.72	0.67	0.62	0.57	0.57
PV									3,482	4,031	4,560	5,163	5,694	6,260	6,767	117,164

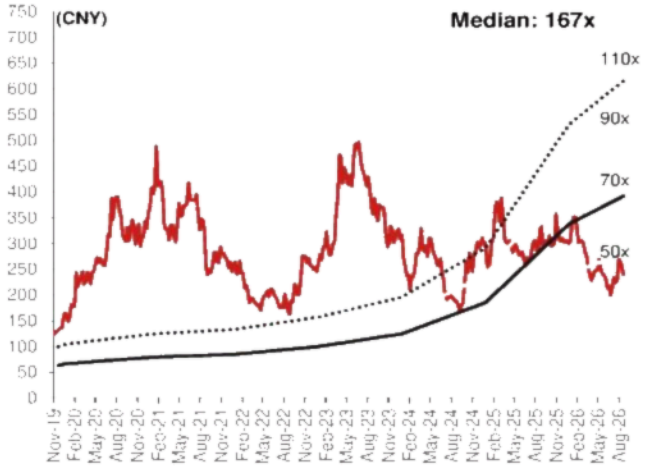
Assumptions	
Risk free rate	1.7%
Market risk premium	7.0%
Beta	0.96
Terminal growth rate	2.5%
WACC	8.4%

PV of corporate value	153,121
+ cash	2,681
- borrowing	0
- minority interest preferred shares	18
PV of equity value	155,821
# of outstanding (mn)	463
	CNY
Target price	336

CNY 336

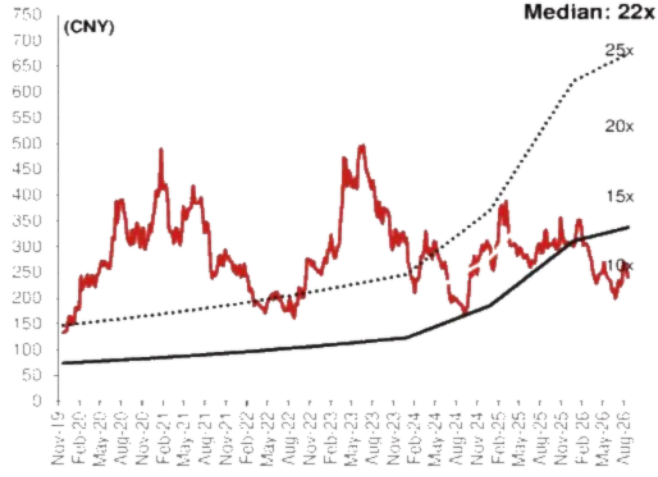
Source: Company data, Nomura estimates

Fig. 84: Kingsoft Office – 12-month forward P/E band



Source: Company data, Nomura research

Fig. 85: Kingsoft Office – 12-month forward P/B band



Source: Company data, Nomura research

NOMURA

Global Markets Research
11 September 2026

Kingsoft Cloud 3896.HK 3896 HK

EQUITY: TECHNOLOGY

A key supplier of AI computing power in China...

...Beneficiary of accelerated capex growth from LLM players

Action: Maintain Buy but cut TP to HKD8.2, implying 51.3% upside

We believe the China AI investment cycle will likely accelerate, driven by the better performance of domestic AI chips and continued strong demand from large language models (LLM) training, as well as fast growing inference demand from more token-consuming applications (i.e., AI agents). IDC predicts China's enterprise-level token consumption will increase about 20x y-y in 2026E, to reach 40,000tn, with an average daily consumption of 110tn (*link*). We think Kingsoft Cloud (KC) should benefit from the robust demand, bolstered by ecosystem customers including Xiaomi (1810 HK, Neutral, covered by Donnie Teng) and Kingsoft (3888 HK, Buy), and other AI customers (LLM, robotics, etc). We raise FY26-27F revenue by 6.0-8.1%, to reflect accelerated training & inference demand. However, we cut FY26-27F GPM by 0.4-1.5pp to reflect cost pressure (due to component price hikes, as well as rental fees from third-parties). Our FY26-27F earnings are down from -CNY234mn/CNY118mn to -CNY847mn/-CNY549mn. We also introduce FY28F revenue forecasts with 26% y-y growth, and estimate FY28F earnings to turn positive to CNY43mn. We reaffirm Buy, with a lower DCF-based TP of HKD8.2, assuming a WACC of 9.1% (from 10.0%, as we assume a lower market risk premium), implying 2.06x FY27F P/S. The stock trades at 1.4x FY27F P/S.

Continued robust demand for AI Cloud, while computing power supply still tight

We think KC's public cloud segment will continue to enjoy strong demand for its cloud infra services as well as Model-as-a-Service (MaaS), driven by an accelerated investment cycle from large tech firms such as Xiaomi, and domestic leading LLM labs, such as Moonshot (unlisted) and DeepSeek (unlisted). We estimate AI revenue would contribute 40%/45%/50% of total revenue and clock a 41% CAGR through FY26-28F. In addition, we think current supply of AI chips, particularly for inference demand, is still tight for China cloud service providers, which has hindered revenue growth. We anticipate the tight supply may gradually ease with the capacity ramp-up of domestic chips.

KC has raised prices for some cloud services; gross margin improving

KC has raised prices for some of its products since 12 July 2026, due to rising AI hardware costs. AI computing power prices have risen 15-50%, and file storage prices have increased by 30-50%, according to the company. We believe KC's price increases as well as rising AI exposure may help support GPM recovery in 2H26F.

Year-end 31 Dec	FY25		FY26F		FY27F		FY28F
Currency (CNY)	Actual	Old	New	Old	New	Old	New
Revenue (mn)	9,559	11,524	12,212	14,300	15,462	0	19,416
Reported net profit (mn)	-936	-234	-847	118	-549	0	43
Normalised net profit (mn)	-936	-234	-847	118	-549	0	43
FD normalised EPS	-22.80c	-5.63c	-18.70c	2.85c	-12.12c		0.94c
FD norm. EPS growth (%)	–	–	–	–	–		–
FD normalised P/E (x)	–	–	–	–	–	–	539.5
EV/EBITDA (x)	0.9	–	5.3	–	6.8	–	6.4
Price/book (x)	2.5	–	2.7	–	2.9	–	2.9
Dividend yield (%)	–	–	–	–	–	–	–
ROE (%)	-12.9	-5.2	-9.5	2.6	-6.7		0.5
Net debt/equity (%)	3.8	84.2	76.1	112.0	150.1		198.2

Source: Company data, Nomura estimates

Rating Remains	**Buy**
Target price Reduced from HKD 8.50	**HKD 8.20**
Closing price 8 September 2026	**HKD 5.42**
Implied upside	**+51.3%**
Market Cap (USD mn)	3,132.5
ADT (USD mn)	64.1

Relative performance chart



Source: LSEG, Nomura

Research Analysts

China Technology

Bing Duan - NIHK
bing.duan1@nomura.com
+852 2252 2141

Ethan Zhang - NIHK
ethan.zhang@nomura.com
+852 2252 2157

Key data on Kingsoft Cloud

Performance

(%)	1M	3M	12M		
Absolute (HKD)	-14.6	-8.4	-27.8	M cap (USDmn)	3,132.5
Absolute (USD)	-14.6	-8.5	-28.3	Free float (%)	43.1
Rel to Hang Seng Index	-13.7	-11.5	-27.0	3-mth ADT (USDmn)	64.1

Income statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Revenue	7,785	9,559	12,212	15,462	19,416
Cost of goods sold	-6,444	-8,055	-10,322	-12,652	-15,382
Gross profit	1,341	1,503	1,890	2,811	4,034
SG&A	-3,080	-2,276	-2,701	-3,157	-3,634
Employee share expense					
Operating profit	-1,739	-773	-810	-346	401
EBITDA	-476	1,707	1,439	1,904	2,650
Depreciation	-1,090	-2,307	-2,076	-2,076	-2,076
Amortisation	-173	-174	-174	-174	-174
EBIT	-1,739	-773	-810	-346	401
Net interest expense	-203	-417	-296	-464	-624
Associates & JCEs					
Other income	-39	242	242	242	242
Earnings before tax	-1,981	-948	-864	-568	19
Income tax	2	4	9	11	16
Net profit after tax	-1,979	-944	-855	-557	35
Minority interests	11	7	7	7	7
Other items					
Preferred dividends	0	0	0	0	0
Normalised NPAT	-1,968	-936	-847	-549	43
Extraordinary items					
Reported NPAT	-1,968	-936	-847	-549	43
Dividends					
Transfer to reserves	-1,968	-936	-847	-549	43
Valuations and ratios					
Reported P/E (x)	–	–	–	–	539.5
Normalised P/E (x)	-9.4	-22.3	-26.7	-41.3	539.5
FD normalised P/E (x)	–	–	–	–	539.5
Dividend yield (%)	–	–	–	–	–
Price/cashflow (x)	29.5	5.5	24.9	14.5	10.9
Price/book (x)	3.6	2.5	2.7	2.9	2.9
EV/EBITDA (x)	–	0.9	5.3	6.8	6.4
EV/EBIT (x)	–	–	–	–	42.4
Gross margin (%)	17.2	15.7	15.5	18.2	20.8
EBITDA margin (%)	-6.1	17.9	11.8	12.3	13.7
EBIT margin (%)	-22.3	-8.1	-6.6	-2.2	2.1
Net margin (%)	-25.3	-9.8	-6.9	-3.6	0.2
Effective tax rate (%)	–	–	–	–	-85.3
Dividend payout (%)	–	–	–	–	0.0
ROE (%)	-18.7	-12.9	-9.5	-6.7	0.5
ROA (pretax %)	-11.3	-4.3	-3.6	-1.3	1.3
Growth (%)					
Revenue	-9.7	22.8	27.8	26.6	25.6
EBITDA		–	-15.7	32.3	39.2
Normalised EPS		–	–	–	–
Normalised FDEPS		–	–	–	–

Source: Company data, Nomura estimates

Cashflow statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
EBITDA	-476	1,707	1,439	1,904	2,650
Change in working capital	1,709	-1,674	-1,746	-1,391	-1,437
Other operating cashflow	-604	3,768	1,215	1,049	894
Cashflow from operations	628	3,801	909	1,562	2,107
Capital expenditure	-3,672	-4,742	-5,742	-5,742	-4,742
Free cashflow	-3,044	-941	-4,833	-4,180	-2,635
Reduction in investments	0	0	0	0	
Net acquisitions					
Dec in other LT assets	4	349	0	0	0
Inc in other LT liabilities	-996	-367	0	0	0
Adjustments	52	213	0	0	0
CF after investing acts	-3,983	-747	-4,833	-4,180	-2,635
Cash dividends					
Equity issue	0	4,533	0	0	0
Debt issue	1,535	4,948	3,500	3,500	3,500
Convertible debt issue	0	0	0	0	0
Others	879	-5,365	-1,260	-1,260	-1,260
CF from financial acts	2,414	4,116	2,240	2,240	2,240
Net cashflow	-1,569	3,369	-2,593	-1,940	-395
Beginning cash	4,218	2,649	6,018	3,425	1,485
Ending cash	2,649	6,018	3,425	1,485	1,090
Ending net debt	-3,419	354	6,447	11,887	15,782

Balance sheet (CNYmn)

As at 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Cash & equivalents	2,649	6,018	3,425	1,485	1,090
Marketable securities					
Accounts receivable	1,469	1,740	2,050	2,596	3,259
Inventories					
Other current assets	2,723	3,265	3,265	3,265	3,265
Total current assets	6,841	11,023	8,740	7,345	7,614
LT investments					
Fixed assets	4,630	10,095	13,761	17,428	20,094
Goodwill					
Other intangible assets	695	533	359	185	11
Other LT assets	5,427	5,078	5,078	5,078	5,078
Total assets	17,593	26,729	27,938	30,036	32,798
Short-term debt	2,226	3,348	6,848	10,348	13,848
Accounts payable	1,877	2,014	1,838	2,253	2,739
Other current liabilities	5,057	4,059	2,799	1,539	279
Total current liabilities	9,159	9,421	11,485	14,140	16,866
Long-term debt	1,661	3,024	3,024	3,024	3,024
Convertible debt					
Other LT liabilities	1,267	4,971	4,971	4,971	4,971
Total liabilities	12,087	17,416	19,480	22,135	24,861
Minority interest	337	-4	-11	-19	-26
Preferred stock					
Common stock	18,967	24,104	24,104	24,104	24,104
Retained earnings	-14,292	-15,248	-16,095	-16,645	-16,602
Proposed dividends					
Other equity and reserves	493	461	461	461	461
Total shareholders' equity	5,168	9,317	8,470	7,920	7,963
Total equity & liabilities	17,593	26,729	27,938	30,036	32,798
Liquidity (x)					
Current ratio	0.75	1.17	0.76	0.52	0.45
Interest cover	-8.6	-1.9	-2.7	-0.7	0.6
Leverage					
Net debt/EBITDA (x)	–	0.21	4.48	6.24	5.95
Net debt/equity (%)	net cash	3.8	76.1	150.1	198.2
Per share					
Reported EPS (CNY)	-53.04c	-22.80c	-18.70c	-12.12c	0.94c
Norm EPS (CNY)	-53.04c	-22.80c	-18.70c	-12.12c	0.94c
FD norm EPS (CNY)	-53.04c	-22.80c	-18.70c	-12.12c	0.94c
BVPS (CNY)	1.40	2.06	1.87	1.75	1.76
DPS (CNY)	0.00	0.00	0.00	0.00	0.00
Activity (days)					
Days receivable	133.3	61.3	56.6	54.8	55.2
Days inventory	0.0	0.0	0.0	0.0	
Days payable	123.4	88.2	68.1	59.0	59.4
Cash cycle	9.9	-26.9	-11.5	-4.2	-4.2

Source: Company data, Nomura estimates

Company profile

Kingsoft Cloud is the largest independent cloud service provider in China market and they mainly focus on IaaS market in cloud business. The company offer public cloud services to its premium customers in various sectors such as Video, gaming, e-commerce, AI etc and enterprise cloud services to selected verticals including financial service, public service and healthcare sector etc.

Valuation Methodology

Our TP of HKD8.20 is based on DCF valuation method. Our DCF valuation assumes a WACC of 9.1%, market risk premium of 7.0% and terminal growth rate of 2.5% while we discounted back cash flow to 2026 as a 12 month forward valuation. Our TP implies a 2.61x/2.06x P/S 2026F/2027F. The benchmark index for this stock is the Hang Seng Composite Index.

Risks that may impede the achievement of the target price

Downside risks include: 1) slower-than-expected public cloud business (including AI cloud) expansion rate due to a slower-than-expected internet business growth rate; 2) lower-than-expected demand growth from enterprise customers, which would affect the growth rate of enterprise cloud business; and 3) potential intense competition, which may result in a faster-than-expected decline in the net dollar retention rate of public cloud service premium customers.

ESG

Kingsoft Cloud focuses on technological innovation, conducting responsible operations, talent development, building a green workplace, and strengthening supply chain management and community engagement. KC implements strict environmental management and has adopted energy-saving, low carbon technologies and products at its data centers. The company has also established and maintained effective risk management and internal control systems to ensure stable and efficient business operations. KC respects the human rights of its employees and partners, and has formulated a Kingsoft Human Rights Policy based on internationally-recognized human rights principles.

2Q26 results review

Kingsoft Cloud (KC) reported 2Q26 results on 19 August, with revenue up 31% y-y to CNY3.1bn (+14% q-q), led by revenue growth from AI cloud and enterprise cloud projects, according to the company. Segment-wise, **public cloud revenue** was up 45% y-y (+18% q-q) in 2Q26, driven by AI demand growth; **enterprise cloud revenue** declined 1% y-y (+1% q-q). Gross margin improved 0.8pp y-y and 2.4pp q-q to 15.2% in 2Q26, on an expanding revenue scale, especially intelligent computing services, according to management. KC recorded a non-GAAP net loss of CNY60mn (vs a loss of CNY237mn in 1Q26), and a net loss of CNY93mn. Capex and leased assets obtained totally amounted to CNY3.2bn in 2Q26, mainly used to support intelligent computing services and enterprise cloud projects; and, management expects to continue to invest for further business expansion in 2H26E.

Fig. 86: Kingsoft Cloud - 2Q26 results review

(CNY mn)		1Q24 A	2Q24 A	3Q24 A	4Q24 A	1Q25 A	2Q25 A	3Q25 A	4Q25 A	1Q26 A	2Q26 A	% chg y-y	% chg q-q
Revenue		1,776	1,892	1,886	2,232	1,970	2,349	2,478	2,761	2,704	3,072	31%	14%
	% y-y	-5%	3%	16%	30%	11%	24%	31%	24%	37%	31%		
Public cloud		1,187	1,235	1,176	1,410	1,353	1,625	1,752	1,902	1,996	2,358	45%	18%
	% y-y	3%	6%	16%	34%	14%	32%	49%	35%	47%	45%		
Enterprise cloud		588	657	710	822	616	724	726	859	707	714	-1%	1%
	% y-y	-17%	-3%	17%	23%	5%	10%	2%	4%	15%	-1%		
COGS		(1,482)	(1,573)	(1,582)	(1,806)	(1,652)	(2,010)	(2,097)	(2,296)	(2,358)	(2,606)		
Gross Profit		293	319	303	426	318	339	381	465	346	466	38%	35%
OPEX		(567)	(596)	(1,447)	(469)	(553)	(666)	(526)	(532)	(512)	(443)		
Operating profit		(274)	(277)	(1,144)	(44)	(234)	(327)	(145)	(66)	(166)	23		
	% y-y	-54%	-24%	42%	-87%	-15%	18%	-87%	53%	-29%	-107%		
Net finance expense		(43)	(49)	(53)	(58)	(78)	(113)	(112)	(114)	(122)	(141)		
Pretax income		(379)	(346)	(1,050)	(204)	(307)	(455)	(10)	(176)	(339)	(76)		
	% y-y	-38%	-30%	35%	-29%	-19%	31%	-99%	-14%	10%	-83%		
Net income		(359)	(354)	(1,057)	(197)	(314)	(456)	(5)	(160)	(344)	(93)		
	% y-y	-41%	-29%	35%	-31%	-13%	29%	-100%	-19%	9%	-80%		
Ratio Analysis													
GPM		16.5%	16.9%	16.1%	19.1%	16.2%	14.4%	15.4%	16.9%	12.8%	15.2%	0.8pp	2.4pp
OPEX		-32.0%	-31.5%	-76.7%	-21.0%	-28.0%	-28.3%	-21.2%	-19.3%	-18.9%	-14.4%	13.9pp	4.5pp
OPM		-15.4%	-14.6%	-60.7%	-1.9%	-11.9%	-13.9%	-5.9%	-2.4%	-6.1%	0.7%	14.7pp	6.9pp
Adjusted EBITDA		1.9%	3.2%	9.8%	16.1%	16.2%	17.3%	33.4%	28.4%	27.6%	35.8%	18.5pp	8.2pp
NPM		-20.2%	-18.7%	-56.1%	-8.8%	-15.9%	-19.4%	-0.2%	-5.8%	-12.7%	-3.0%	16.4pp	9.7pp

Source: Company data, Nomura research

Earnings forecast revisions

We raise FY26-27F revenue by 6.0-8.1% to reflect accelerated inference demand from more token-consuming applications (i.e., AI agents). We cut FY26-27F GPM by 0.4-1.5pp, to reflect the price hikes for AI hardware such as memory. As a result, we cut FY26-27F earnings to a net loss of CNY847mn / CNY549mn, respectively, from -CNY234mn / +CNY118mn. We introduce FY28F revenue forecasts with 26% y-y growth and estimate FY28F earnings to turn positive to CNY43mn.

Fig. 87: Kingsoft Cloud — earnings forecast revisions

(CNY mn)	2022 A	2023 A	2024 A	2025 A	2026F (NOM)	2027F (NOM)	2028F (NOM)	2026F (OLD)	2027F (OLD)	2026F Diff.	2027F Diff.
Revenue	8,180	7,047	7,785	9,559	12,212	15,462	19,416	11,524	14,300	6.0%	8.1%
% y-y	-10%	-14%	10%	23%	28%	27%	26%	21%	24%		
Public cloud	5,360	4,382	5,007	6,633	9,287	12,537	16,549	8,462	11,085	9.8%	13.1%
% y-y	-13%	-18%	14%	32%	40%	35%	32%	28%	31%		
Enterprise cloud	2,817	2,664	2,778	2,925	2,925	2,925	2,867	3,062	3,216	-4.5%	-9.0%
% y-y	-3%	-5%	4%	5%	0%	0%	-2%	5%	5%		
COGS	(7,751)	(6,197)	(6,444)	(8,055)	(10,322)	(12,652)	(15,382)	(9,563)	(11,638)	7.9%	8.7%
Gross Profit	430	850	1,341	1,503	1,890	2,811	4,034	1,961	2,662	-3.6%	5.6%
OPEX	(2,681)	(2,305)	(2,160)	(2,276)	(2,701)	(3,157)	(3,634)	(2,310)	(2,624)		
Operating profit	(2,251)	(1,455)	(819)	(773)	(810)	(346)	401	(349)	39		
% y-y	24%	-35%	-44%	-6%	5%	-57%	-216%	-55%	-111%		
Net finance expense	(57)	(68)	(203)	(417)	(296)	(464)	(624)	(196)	(234)		
Pretax income	(2,664)	(2,166)	(1,981)	(948)	(864)	(568)	19	(245)	105	-252.0%	-641.4%
% y-y	69%	-19%	-9%	-52%	-9%	-34%	-103%	-74%	-143%		
Net income	(2,658)	(2,176)	(1,967)	(936)	(847)	(549)	43	(234)	118	-262.2%	-565.4%
% y-y	67%	-18%	-10%	-52%	9%	35%	108%	-75%	-150%		
Ratio Analysis											
GPM	5.3%	12.1%	17.2%	15.7%	15.5%	18.2%	20.8%	17.0%	18.6%	-1.5pp	-0.4pp
OPEX	-32.8%	-32.7%	-27.7%	-23.8%	-22.1%	-20.4%	-18.7%	-20.0%	-18.3%	-2.1pp	-2.1pp
OPM	-27.5%	-20.6%	-10.5%	-8.1%	-6.6%	-2.2%	2.1%	-3.0%	0.3%	-3.6pp	-2.5pp
Adjusted EBITDA	-9.2%	-3.8%	8.2%	24.4%	15.6%	15.5%	16.3%	9.3%	7.5%	6.3pp	8.0pp
NPM	-32.5%	-30.9%	-25.3%	-9.8%	-6.9%	-3.6%	0.2%	-2.0%	0.8%	-4.9pp	-4.4pp

Source: Company data, Nomura estimates

Our FY26-28F revenue forecasts are 2-4% below Bloomberg consensus estimates, as we are cautious about topline growth due to tight supply of AI chips. Our FY26-27F earnings are 56-176% below, as we are more cautious on gross margin, due to hikes in upstream component prices as well as rental fees from third-parties.

Fig. 88: Kingsoft Cloud - Nomura vs Bloomberg consensus

CNY mn	FY2025	FY2026F			FY2027F			FY2028F		
		(NOM)	(CON)	diff	(NOM)	(CON)	diff	(NOM)	(CON)	diff
Revenue	9,559	12,212	12,744	-4%	15,462	15,931	-3%	19,416	19,875	-2%
Operating Profit	(773)	(810)	(326)	149%	(346)	70	-593%	401	364	10%
Pretax Profit	(948)	(864)	(927)	-7%	(568)	(557)	-2%	19	(276)	-107%
Net Profit	(936)	(847)	(544)	-56%	(549)	(199)	-176%	43	10	320%

Source: Company data, Bloomberg Finance L.P. consensus, Nomura estimates

Valuation and risks

Our TP of HKD8.2 continues to be based on a DCF valuation method, assuming a WACC of 9.1% (previously: 10.0%), a market risk premium of 7.0% (previously: 8.85%), and a terminal growth rate of 2.5% (previously: 3.5%), while we discount back cash flows to 2026 at a 12-month-forward valuation. Our TP implies 2.61x/2.06x 2026F/2027F P/S.

Fig. 89: Kingsoft Cloud – DCF valuation

Year	2019	2020	2021	2022	2023	2024	2025	2026F	2027F	2028F	2029F	2030F	2031F	2032F	2033F	Terminal
Period								0	1	2	3	4	5	6	7	
Revenue	3,956	6,577	9,061	8,180	7,047	7,785	9,559	12,212	15,462	19,416	24,712	31,265	39,217	48,668	59,660	
Profit after tax	(1,161)	(982)	(1,589)	(2,658)	(2,176)	(1,967)	(936)	(847)	(549)	43	657	1,665	2,942	4,218	5,724	
Operating cash flow	(439)	(290)	(709)	189	(169)	628	3,801	909	1,562	2,107	2,683	3,628	4,821	6,084	7,517	
CAPEX	(1,000)	(1,560)	(723)	(1,419)	(1,965)	(3,672)	(4,742)	(5,742)	(5,742)	(4,742)	(4,242)	(3,742)	(3,242)	(3,242)	(2,742)	
FCF	(1,439)	(1,850)	(1,432)	(1,230)	(2,134)	(3,044)	(941)	(4,833)	(4,180)	(2,635)	(1,560)	(114)	1,579	2,842	4,774	74,253
Discount factor									0.92	0.84	0.77	0.71	0.65	0.59	0.54	0.54
PV									(3,832)	(2,214)	(1,201)	(81)	1,022	1,686	2,597	40,383

Assumptions	
Risk-free rate	1.7%
Market risk premium	7.0%
Beta	2.1
Terminal growth rate	2.5%
WACC	9.1%

PV of corporate value	38,360
+ cash	3,425
- borrowing	9,872
- minority interest preferred shares	11
PV of equity value	31,924
# of shares outstanding	4,532
Target price CNY	7

HKD 8.2

Source: Company data, Nomura estimates

NOMURA

Global Markets Research
11 September 2026

iFlytek 002230.SZ 002230 CH

EQUITY: TECHNOLOGY

"National team" on China's LLM market...

...Leveraging domestic AI chips for LLM training & inference

Action: Maintain Neutral; TP cut to CNY40, implying 1.9% upside

We believe iFlytek is well-positioned as a full-stack domestic large language model (LLM) platform company, which could benefit from the acceleration of domestic AI chips to enhance LLM capabilities, and fast growing token demand from AI applications in the China market. According to iFlytek, its Spark LLM was the only mainstream LLM in China trained on domestic AI chips. We think iFlytek stands to gain from AI adoption in key enterprise verticals, i.e., education, finance, healthcare, while consumer market monetization faces competitive headwinds. We cut FY26-27F revenue/ earnings by 5-8% / 7-8%, to reflect intensified competition in the education and consumer segments. We also introduce FY28F forecasts, with revenue/earnings growth of 11%/26% y-y. We reaffirm Neutral and cut TP to CNY40, based on 65x (previously: 98x) FY27F EPS of CNY0.61, in line with the A-shares software median P/E. The stock trades at 64x FY27F P/E.

LLM commercialization accelerating in enterprise market; token consumption surging on open platforms

On 29 April, iFlytek released its first 30B medium-sized model – the Spark X2-flash – based on domestic AI chips. We expect the launch of Huawei's (unlisted) Ascend 950 DT (designed for LLM training) this year to underpin faster performance upgrades in the next-generation Spark LLM. LLM API and MaaS platform service revenue increased 70% y-y in 1H26, while token consumption in June-2026 was 2x the total volume in 2025. We expect LLM and Agent product upgrades will continue to drive token demand, and form the foundation for the company's AI commercialization growth.

Industry exposure improves; AI hardware is emerging but faces strong competition

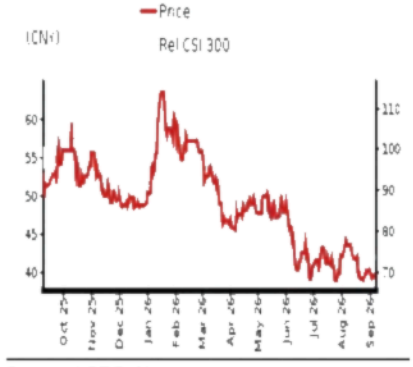
As part of its long-term strategy, the company has been actively diversifying its revenue mix away from the government sector, which declined 2.65% y-y in 1H26, and now forms <30% of total sales. In the enterprise market, the healthcare/ smart auto segments recorded solid growth (+59%/ 20% y-y), while the education sector underperformed (down 6% y-y). In the consumer market, AI hardware such as AI glasses/ Guide01 robot have been fully integrated with iFlytek's LLMs to dive into the physical world. However, we believe the consumer market remains highly competitive, as major AI tech companies are all vying for market share in this segment. We prefer Kingdee (268 HK, Buy) in the China software sector.

Year-end 31 Dec	FY25		FY26F		FY27F		FY28F
Currency (CNY)	Actual	Old	New	Old	New	Old	New
Revenue (mn)	27,105	32,663	31,018	37,908	34,963	0	38,827
Reported net profit (mn)	839	1,215	1,131	1,603	1,476	0	1,864
Normalised net profit (mn)	839	1,215	1,131	1,603	1,476	0	1,864
FD normalised EPS	36.31c	51.47c	47.88c	66.48c	61.19c		84.80c
FD norm. EPS growth (%)	50.0	34.1	31.9	29.2	27.8		38.6
FD normalised P/E (x)	108.2	–	82.0	–	64.2	–	46.3
EV/EBITDA (x)	124.9	–	88.1	–	67.8	–	55.2
Price/book (x)	4.9	–	4.8	–	4.1	–	3.8
Dividend yield (%)	0.3	–	–	–	–	–	–
ROE (%)	4.6	6.5	5.9	8.2	7.2		8.5
Net debt/equity (%)	28.1	34.1	29.9	33.6	25.7		15.3

Source: Company data, Nomura estimates

Rating Remains	**Neutral**
Target price Reduced from CNY 50.00	**CNY 40.00**
Closing price 8 September 2026	**CNY 39.27**
Implied upside	**+1.9%**
Market Cap (USD mn)	14,056.6
ADT (USD mn)	263.2

Relative performance chart



Source: LSEG, Nomura

Research Analysts

China Technology

Bing Duan - NIHK
bing.duan1@nomura.com
+852 2252 2141

Joel Ying, CFA - NIHK
joel.ying@nomura.com
+852 2252 2153

Key data on iFlytek

Performance

(%)	1M	3M	12M		
Absolute (CNY)	-9.8	-11.0	-25.5	M cap (USDmn)	14,056.6
Absolute (USD)	-9.3	-10.0	-20.8	Free float (%)	79.7
Rel to CSI 300	-7.2	-8.1	-27.9	3-mth ADT (USDmn)	263.2

Income statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Revenue	23,343	27,105	31,018	34,963	38,827
Cost of goods sold	-13,391	-15,624	-17,879	-20,154	-22,381
Gross profit	9,952	11,481	13,138	14,809	16,446
SG&A	-9,574	-11,210	-12,642	-14,040	-15,359
Employee share expense					
Operating profit	378	271	496	769	1,087
EBITDA	1,062	840	1,198	1,540	1,853
Depreciation	-685	-569	-701	-771	-766
Amortisation					
EBIT	378	271	496	769	1,087
Net interest expense	-135	-173	-229	-229	-229
Associates & JCEs					
Other income	-14	495	567	639	710
Earnings before tax	229	593	833	1,178	1,567
Income tax	278	256	307	307	307
Net profit after tax	507	849	1,141	1,486	1,874
Minorities	53	-10	-10	-10	-10
Other items					
Preferred dividends					
Normalised NPAT	560	839	1,131	1,476	1,864
Extraordinary items					
Reported NPAT	560	839	1,131	1,476	1,864
Dividends	0	-240	0	0	
Transfer to reserves	560	599	1,131	1,476	1,864
Valuations and ratios					
Reported P/E (x)	162.2	108.2	82.0	64.2	46.3
Normalised P/E (x)	162.2	108.2	82.0	64.2	46.3
FD normalised P/E (x)	162.2	108.2	82.0	64.2	46.3
Dividend yield (%)	–	0.3	–	–	–
Price/cashflow (x)	36.4	79.1	57.0	44.7	33.3
Price/book (x)	5.1	4.9	4.8	4.1	3.8
EV/EBITDA (x)	97.0	124.9	88.1	67.8	55.2
EV/EBIT (x)	273.1	387.5	212.8	135.8	94.1
Gross margin (%)	42.6	42.4	42.4	42.4	42.4
EBITDA margin (%)	4.6	3.1	3.9	4.4	4.8
EBIT margin (%)	1.6	1.0	1.6	2.2	2.8
Net margin (%)	2.4	3.1	3.6	4.2	4.8
Effective tax rate (%)	-121.3	-43.2	-36.9	-26.1	-19.6
Dividend payout (%)	0.0	28.6	0.0	0.0	0.0
ROE (%)	3.2	4.6	5.9	7.2	8.5
ROA (pretax %)	1.0	0.7	1.1	1.6	2.1
Growth (%)					
Revenue	18.8	16.1	14.4	12.7	11.1
EBITDA	396.7	-20.9	42.5	28.6	20.3
Normalised EPS	28.2	50.0	31.9	27.8	38.6
Normalised FDEPS	28.2	50.0	31.9	27.8	38.6

Source: Company data, Nomura estimates

Cashflow statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
EBITDA	1,062	840	1,198	1,540	1,853
Change in working capital	-1,266	-261	-205	-130	-34
Other operating cashflow	2,699	569	635	707	778
CF from operations	2,495	1,148	1,628	2,117	2,596
Capital expenditure	-3,220	-2,926	-2,418	-1,676	-697
Free cashflow	-725	-1,778	-790	441	1,899
Reduction in investments	0	0	0	0	
Net acquisitions					
Dec in other LT assets	0	0	0	0	
Inc in other LT liabilities	95	0	0	0	0
Adjustments	-180	0			
CF after investing acts	-810	-1,778	-790	441	1,899
Cash dividends	0	-240	0	0	0
Equity issue	544	0	4,000	0	0
Debt issue	718	2,134	0	0	0
Convertible debt issue					
Others	-627	0	-3,900	0	0
CF from financial acts	635	1,894	100	0	0
Net cashflow	-175	116	-690	441	1,899
Beginning cash	3,563	3,387	3,504	2,814	3,254
Ending cash	3,387	3,504	2,814	3,254	5,154
Ending net debt	3,151	5,169	5,859	5,419	3,519

Balance sheet (CNYmn)

As at 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Cash & equivalents	3,387	3,504	2,814	3,254	5,154
Marketable securities					
Accounts receivable	16,234	17,857	19,643	21,607	23,768
Inventories	2,847	3,131	3,444	3,789	4,168
Other current assets	1,595	1,595	1,595	1,595	1,595
Total current assets	24,062	26,086	27,495	30,245	34,683
LT investments					
Fixed assets	6,070	8,427	10,143	11,048	10,979
Goodwill	2,917	2,917	2,917	2,917	2,917
Other intangible assets	8,430	8,430	8,430	8,430	8,430
Other LT assets					
Total assets	41,479	45,860	48,985	52,639	57,009
Short-term debt	1,866	4,000	4,000	4,000	4,000
Accounts payable	10,981	12,629	14,523	16,701	19,207
Other current liabilities	2,512	2,512	2,512	2,512	2,512
Total current liabilities	15,359	19,141	21,035	23,213	25,719
Long-term debt	4,673	4,673	4,673	4,673	4,673
Convertible debt					
Other LT liabilities	2,732	2,732	2,732	2,732	2,732
Total liabilities	22,764	26,546	28,440	30,618	33,124
Minority interest	921	921	921	921	921
Preferred stock					
Common stock	17,793	18,393	19,624	21,099	22,964
Retained earnings					
Proposed dividends					
Other equity and reserves	0	0			
Total shareholders' equity	17,793	18,393	19,624	21,099	22,964
Total equity & liabilities	41,479	45,860	48,985	52,639	57,009

Liquidity (x)

Current ratio	1.57	1.36	1.31	1.30	1.35
Interest cover	2.8	1.6	2.2	3.4	4.7

Leverage

Net debt/EBITDA (x)	2.97	6.15	4.89	3.52	1.90
Net debt/equity (%)	17.71	28.11	29.86	25.68	15.33

Per share

Reported EPS (CNY)	24.21c	36.31c	47.88c	61.19c	84.80c
Norm EPS (CNY)	24.21c	36.31c	47.88c	61.19c	84.80c
FD norm EPS (CNY)	24.21c	36.31c	47.88c	61.19c	84.80c
BVPS (CNY)	7.70	7.96	8.14	9.60	10.44
DPS (CNY)	0.00	0.10	0.00	0.00	0.00

Activity (days)

Days receivable	231.2	229.5	220.6	215.3	213.9
Days inventory	72.3	69.8	67.1	65.5	65.1
Days payable	285.2	275.8	277.1	282.7	293.6
Cash cycle	18.4	23.6	10.6	-1.9	-14.7

Source: Company data, Nomura estimates

Company profile

iFlytek is a China-based software solution provider that specializes in speech intelligence and artificial intelligence technology. The company provides AI-enabled solutions in key vertical industries, including education, healthcare, juridical and consumer electronics.

Valuation Methodology

Our TP of CNY40.00 is based on 65x 2027F EPS of CNY0.61. The benchmark index is CSI 300.

Risks that may impede the achievement of the target price

Downside risks include: 1) rising competition in the AI and voice recognition market, 2) slower-than-expected AI technology ramp-up and 3) inability to monetize the AI technology in the longer term. Upside risks include: 1) faster-than-expected AI technology ramp-up; and 2) better-than-expected in monetization of the AI technology.

ESG

As a leading AI solution provider, iFlytek takes social responsibility by helping enterprise and industrial clients in digital transformation.

1H26 & 2Q26 earnings review

iFlytek recorded 7% y-y revenue growth in 1H26, and 2% y-y revenue growth in 2Q26 (+20% q-q). Its gross margin declined slightly by 0.4pp y-y in 1H26 (+0.1pp y-y in 2Q26), and net loss narrowed down by 15% in 1H26, while net loss margin improved 0.2pp y-y in 2Q26.

Fig. 90: iFlytek - 1H26 & 2Q26 earnings review

(CNY mn)	1Q25	2Q25	3Q25	4Q25	1Q26	2Q26	y-y	q-q	1H24	2H24	1H25	2H25	1H26
	A	A	A	A	A	A			A	A	A	A	A
Turnover	4,658	6,253	6,078	10,116	5,274	6,349	2%	20%	9,325	14,019	10,911	16,194	11,623
% chg yoy	28%	10%	10%	19%	13%				19%	19%	17%	16%	7%
COGS	(2,786)	(3,735)	(3,624)	(5,479)	(3,219)	(3,784)			(5,577)	(7,814)	(6,522)	(9,103)	(7,002)
Gross Profit	1,872	2,518	2,454	4,637	2,055	2,565			3,748	6,204	4,389	7,092	4,621
OPEX	(2,260)	(2,588)	(2,587)	(3,775)	(2,622)	(2,933)			(4,224)	(5,350)	(4,848)	(6,362)	(5,555)
Sales and marketing	(931)	(1,155)	(1,141)	(1,964)	(1,048)	(1,237)			(1,626)	0	(2,086)	0	(2,285)
Operating Income	(388)	(70)	(132)	862	(567)	(368)	-423%	35%	(477)	854	(459)	730	(935)
% chg yoy	-38%	-148%	63%	-8%	46%				11%	241%	-4%	-15%	104%
Finance Expense	(36)	(41)	(44)	(51)	(51)	(72)			(64)	-	(77)	-	(123)
Other gains/expense	229	(95)	337	25	444	210			(51)	(33)	134	266	654
Pretax income	(195)	(207)	160	835	(173)	(230)	-11%	-33%	(592)	821	(402)	995	(404)
% chg yoy	-40%	-22%	212%	8%	-11%				1270%	77%	-32%	21%	0%
Tax	(21)	195	(15)	98	(15)	225			122	156	174	82	211
Minority expenses	23	(34)	27	(27)	18	(29)			69	(16)	(11)	1	(11)
Profit to shareholders	(193)	(46)	172	906	(170)	(34)	25%	80%	(401)	961	(239)	1,078	(204)
% chg yoy	-36%	-54%	202%	0%	-12%				-645%	65%	-40%	12%	15%
Ratio Analysis													
GPM	40.2%	40.3%	40.4%	45.8%	39.0%	40.4%	0.1pp	1.4pp	40.2%	44.3%	40.2%	43.8%	39.8%
OPEX	48.5%	41.4%	42.6%	37.3%	49.7%	46.2%	4.8pp	-3.5pp	45.3%	38.2%	44.4%	39.3%	47.8%
OPM	-8.3%	-1.1%	-2.2%	8.5%	-10.7%	-5.8%	-4.7pp	4.9pp	-5.1%	6.1%	-4.2%	4.5%	-8.0%
NPM	-4.1%	-0.7%	2.8%	9.0%	-3.2%	-0.5%	0.2pp	2.7pp	-4.3%	6.9%	-2.2%	6.7%	-1.8%

Source: Company data, Nomura research

Segment-wise, the core education business' revenue declined 6% y-y in 1H26, owing to a tighter control of government budgets, in our view. Meanwhile, the company made good progress in AI penetration within the healthcare and automotive sectors, which recorded 59% / 20% y-y revenue growth in 1H26, respectively. The consumer sector recorded 23% y-y revenue growth, mainly driven by token consumption and developer growth in its open AI platform, in our view, while the consumer hardware sector faces more headwinds due to intensified competition.

Fig. 91: iFlytek - 1H26 earnings summary by segment

(CNY mn)	1H23	2H23	1H24	2H24	1H25	2H25	1H26
Revenue by segment	A	A	A	A	A	A	A
Education	2,421	3,430	3,012	4,534	3,696	5,618	3,490
% chg yoy	4%	-16%	24%	32%	23%	24%	-6%
% as total revenue	31%	29%	32%	32%	34%	35%	30%
Smart City / Judiciary	1,506	2,428	1,186	2,431	1,270	2,301	1,101
% chg yoy	-31%	7%	-21%	0%	7%	-5%	-13%
% as total revenue	19%	21%	13%	17%	12%	14%	9%
Healthcare	192	348	228	463	276	582	438
% chg yoy	37%	6%	19%	33%	21%	26%	59%
% as total revenue	2%	3%	2%	3%	3%	4%	4%
Consumer	2,423	3,762	3,512	4,374	3,900	5,349	4,809
% chg yoy	21%	42%	45%	16%	11%	22%	23%
% as total revenue	31%	32%	38%	31%	36%	33%	41%
Automotive	213	483	352	636	399	841	481
% chg yoy	26%	63%	65%	32%	13%	32%	20%
% as total revenue	3%	4%	4%	5%	4%	5%	4%
Finance	91	197	98	545	439	419	226
% chg yoy	22%	23%	7%	176%	350%	-23%	-49%
% as total revenue	1%	2%	1%	4%	4%	3%	2%
Telecom	969	1,127	902	999	886	1,049	859
% chg yoy	-11%	12%	-7%	-11%	-2%	5%	-3%
% as total revenue	12%	10%	10%	7%	8%	6%	7%

Source: Company data, Nomura research

Earnings forecast revisions

We cut FY26-27F revenue by 5-8% to reflect intensified competition in the education and consumer segments, and cut earnings by 7-8% to reflect the intense AI investments. We also introduce FY28F revenue/earnings growth forecasts of 11%/26% y-y.

Fig. 92: iFlytek - earnings forecast revisions

(CNY mn)	2022	2023	2024	2025	2026F	2027F	2028F	2026F	2027F	2026F	2027F
	A	A	A	(NOM)	(NOM)	(NOM)	(NOM)	(OLD)	(OLD)	diff.	diff.
Turnover	18,820	19,650	23,343	27,105	31,018	34,963	38,827	32,663	37,908	-5%	-8%
% chg yoy	3%	4%	19%	16%	14%	13%	11%	18%	16%		
COGS	(11,136)	(11,267)	(13,391)	(15,624)	(17,879)	(20,154)	(22,381)	(19,065)	(22,126)		
Gross Profit	7,684	8,384	9,952	11,481	13,138	14,809	16,446	13,599	15,782	-3%	-6%
OPEX	(7,615)	(8,562)	(9,574)	(11,210)	(12,642)	(14,040)	(15,359)	(12,581)	(14,373)	0%	-2%
Sales and marketing	(3,164)	(3,584)	(4,083)	(5,191)	0	0	0	0	0		
Operating Income	69	(178)	378	271	496	769	1,087	1,018	1,409		
% chg yoy	-91%	-357%	-312%	-28%	83%	55%	41%	46%	38%		
Finance Expense	79	(15)	(135)	(173)	(229)	(229)	(229)	(204)	(204)		
Other gains/expense	100	613	(14)	495	567	639	710	(19)	(22)		
Pretax income	248	420	229	593	833	1,178	1,567	795	1,183	5%	0%
% chg yoy	-83%	69%	-45%	159%	40%	41%	33%	50%	49%		
Tax	251	193	278	256	307	307	307	367	367		
Minority expenses	63	44	53	(10)	(10)	(10)	(10)	53	53	-119%	-119%
Profit to shareholders	561	657	560	839	1,131	1,476	1,864	1,215	1,603	-7%	-8%
% chg yoy	-64%	17%	-15%	50%	35%	31%	26%	37%	32%		
Ratio Analysis											
GPM	40.8%	42.7%	42.6%	42.4%	42.4%	42.4%	42.4%	41.6%	41.6%	0.7pp	0.7pp
OPEX	40.5%	43.6%	41.0%	41.4%	40.8%	40.2%	39.6%	38.5%	37.9%	2.2pp	2.2pp
OPM	0.4%	-0.9%	1.6%	1.0%	1.6%	2.2%	2.8%	3.1%	3.7%	-1.5pp	-1.5pp
NPM	3.0%	3.3%	2.4%	3.1%	3.6%	4.2%	4.8%	3.7%	4.2%	-0.1pp	0.0pp

Source: Company data, Nomura estimates

Fig. 93: iFlytek - revenue breakdown by segment

(CNY mn)	2022	2023	2024	2025	2026F	2027F	2028F	2026F	2027F	2026F	2027F
Revenue by segment	A	A	A	(NOM)	(NOM)	(NOM)	(NOM)	(OLD)	(OLD)	diff.	diff.
Education	6,420	5,851	7,545	9,314	11,312	13,512	15,869	11,405	13,680	-1%	-1%
% chg yoy	3%	-9%	29%	23%	21%	19%	17%	22%	20%		
% as total revenue	34%	30%	32%	34%	36%	39%	41%				
Smart City / Judiciary	4,446	3,933	3,617	3,572	3,706	4,030	4,584	3,575	3,823	4%	5%
% chg yoy	-11%	-12%	-8%	-1%	4%	9%	14%	2%	7%		
% as total revenue	24%	20%	15%	13%	12%	12%	12%				
Healthcare	467	540	692	858	979	1,019	958	884	868	11%	17%
% chg yoy	38%	15%	28%	24%	14%	4%	-6%	8%	-2%		
% as total revenue	2%	3%	3%	3%	3%	3%	2%				
Consumer	4,640	6,185	7,886	9,249	10,571	11,764	12,739	11,545	13,451	-8%	-13%
% chg yoy	-1%	33%	28%	17%	14%	11%	8%	20%	17%		
% as total revenue	25%	31%	34%	34%	34%	34%	33%				
Automotive	465	696	989	1,240	1,493	1,723	1,902	1,662	2,030	-10%	-15%
% chg yoy	3%	50%	42%	25%	20%	15%	10%	27%	22%		
% as total revenue	2%	4%	4%	5%	5%	5%	5%				
Finance	235	289	643	859	890	789	581	1,581	2,017	-44%	-61%
% chg yoy	19%	23%	123%	34%	4%	-11%	-26%	43%	28%		
% as total revenue	1%	1%	3%	3%	3%	2%	1%				
Telecom	2,097	2,096	1,901	1,936	1,971	2,006	2,043	1,901	1,901	4%	6%
% chg yoy	51%	0%	-9%	2%	2%	2%	2%	0%	0%		
% as total revenue	11%	11%	8%	7%	6%	6%	5%				

Source: Company data, Nomura estimates

Our FY26-28F revenue forecasts are 3-12% below WIND consensus, as we are more cautious about revenue growth from the consumer market. Our FY26-28F earnings forecasts are 7-9% lower than WIND consensus, as we are concerned about increasing R&D expenses and marketing expenses for LLM development and sales of AI hardware, respectively.

Fig. 94: iFlytek — Nomura vs consensus

CNYmn	FY2025	FY2026F			FY2027F			FY2028F		
		(NOM)	(CON)	diff	(NOM)	(CON)	diff	(NOM)	(CON)	diff
Revenue	27,105	31,018	31,927	-3%	34,963	37,492	-7%	38,827	43,929	-12%
Pretax Profit	593	833	923	-10%	1,178	1,113	6%	1,567	1,424	10%
Net Profit	839	1,131	1,236	-9%	1,476	1,616	-9%	1,864	2,043	-9%
EPS	0.36	0.48	0.51	-7%	0.61	0.67	-9%	0.77	0.85	-9%

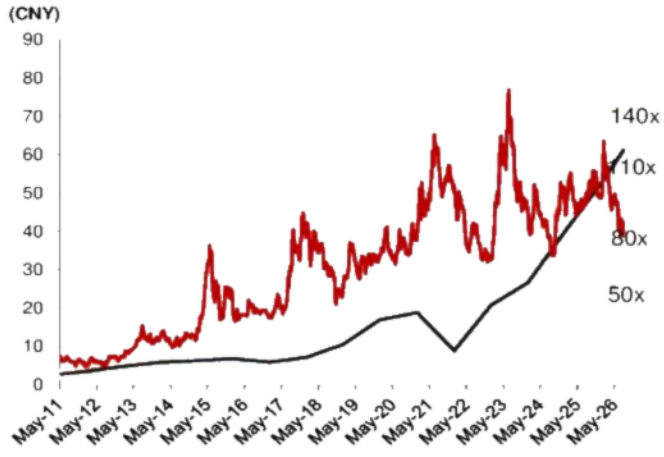
Source: Company data, Nomura estimates

Valuation and risks

We maintain our Neutral rating and raise our TP to CNY40, based on 65x (previously: 98x) rolled-forward FY27F EPS of CNY0.61, in line with the A-shares software median P/E.

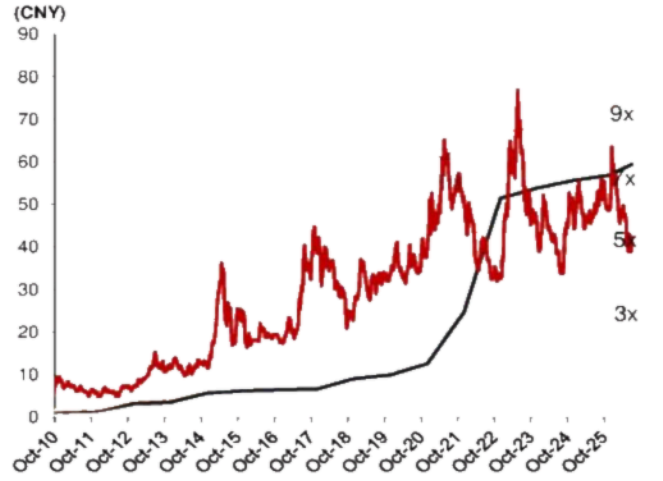
Downside risks: 1) higher-than-expected competition in the LLM market; and 2) slower-than-expected demand recovery from downstream customers.

Fig. 95: iFlytek — 12-month-forward P/E bands



Source: Company data, Nomura estimates

Fig. 96: iFlytek — 12-month-forward P/B bands



Source: Company data, Nomura estimates

NOMURA

Global Markets Research
11 September 2026

GDS Holdings GDS.OQ GDS US

EQUITY: TECHNOLOGY

Strong order backlog with long-term visibility

Reiterate Buy with TP trimmed to USD45.7

2Q26 results in line with the company's raised order intake target

GDS reported 2Q26 results with service revenue and adjusted EBITDA growing 6.5% y-y and 2.5% y-y, respectively – largely in line with market expectations (see *our note*). The company registered 263MW of new contracts in 2Q26, bringing total order intake for 1H26 to 471MW. Management has further raised its new bookings target to 1GW (previously 500MW) for FY26E, in addition to another 1GW of power reservations by customers by end-FY26E. Revenue guidance for FY26E was revised up to CNY12.7-13bn and adjusted EBITDA guidance was also raised to CNY5.9-6.1bn, to reflect the one-off items in 1Q26.

Strong order backlog, though move-in still skewed toward 2H27F and beyond

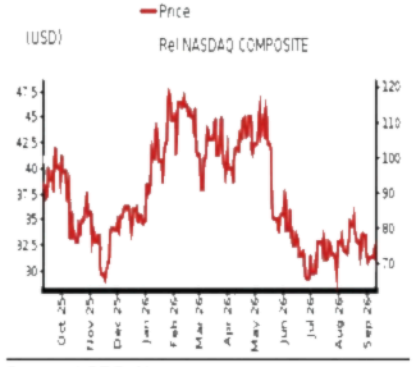
With 1GW of new bookings guided for this year, we expect GDS to achieve order backlog of ~1.2GW by end-FY26F, more than doubling from 450MW as of end-FY25. By end-1H26, the company secured total resources of 5.8GW, with ~3.1GW available for sale, which we expect to support a sustainable order-taking trend for the company over 2H26-FY27F. Meanwhile, we expect accelerating order intake will start to contribute to EBITDA growth in FY28F, with moderate growth for FY26-27F as a result of concentrated contract renewal for its existing portfolio.

Action: Reiterate Buy with a lower TP of USD45.7

We slightly lift our FY26-28F revenue forecasts by 1.4%-4.7% to reflect stronger-than-expected order intake, as we now estimate annual new commitment of 1GW for FY26-28F, and additional customer move-in of 710MW and 950MW for FY27/28F. Meanwhile, given changing market mix with more contribution from projects in new markets (e.g., Inner Mongolia), we cut our FY26-28F adjusted EBITDA margin forecasts by 1.8-3.7pp, and lower adjusted EBITDA by 2.6-6.8%. We reiterate our Buy rating, but reduce our TP to USD45.7 (from USD48.4). Our new TP is based 18.8x FY27F EV/EBITDA, which is in line with the company historical average EV/EBITDA. The stock is trading at 15.5x FY27F EV/EBITDA.

Rating Remains	**Buy**
Target price Reduced from USD 48.40	**USD 45.70**
Closing price 8 September 2026	**USD 32.00**
Implied upside	**+42.8%**
Market Cap (USD mn)	6,120.0
ADT (USD mn)	62.0

Relative performance chart



Source: LSEG, Nomura

Research Analysts

China Technology

Ethan Zhang - NIHK
ethan.zhang@nomura.com
+852 2252 2157

Bing Duan - NIHK
bing.duan1@nomura.com
+852 2252 2141

Year-end 31 Dec	FY25		FY26F		FY27F		FY28F
Currency (CNY)	Actual	Old	New	Old	New	Old	New
Revenue (mn)	11,432	12,577	12,772	13,721	13,916	16,993	17,783
Reported net profit (mn)	311	261	-94	-411	-945	230	-557
Normalised net profit (mn)	311	261	-94	-411	-945	230	-557
FD normalised EPS	1.56	1.27	-45.54c	-2.01	-4.60	1.12	-2.71
FD norm. EPS growth (%)	-91.3	-18.4	-129.2	-257.7	–	–	–
FD normalised P/E (x)	149.3	–	–	–	–	–	–
EV/EBITDA (x)	20.8	–	13.2	–	15.5	–	13.5
Price/book (x)	1.7	–	1.5	–	1.6	–	1.6
Dividend yield (%)	–	–	–	–	–	–	–
ROE (%)	1.3	1.0	-0.4	-1.5	-3.5	0.8	-2.1
Net debt/equity (%)	90.0	85.4	83.7	107.4	118.1	122.1	149.2

Source: Company data, Nomura estimates

See Appendix A-1 for analyst certification, important disclosures and the status of non-US analysts.

Key data on GDS Holdings

Performance

(%)	1M	3M	12M		
Absolute (USD)	0.1	-7.8	-4.5	M cap (USDmn)	6,120.0
Absolute (USD)	0.1	-7.8	-4.5	Free float (%)	51.0
Rel to NASDAQ COMPOSITE	1.1	-9.7	-25.7	3-mth ADT (USDmn)	62.0

Income statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Revenue	10,322	11,432	12,772	13,916	17,783
Cost of goods sold	-8,099	-8,847	-10,180	-12,467	-15,530
Gross profit	2,223	2,585	2,592	1,449	2,253
SG&A	-1,071	-2,641	-900	-1,151	-1,266
Employee share expense					
Operating profit	1,152	-56	1,692	298	987
EBITDA	4,395	3,404	5,658	5,477	7,035
Depreciation	-3,243	-3,459	-3,966	-5,180	-6,048
Amortisation					
EBIT	1,152	-56	1,692	298	987
Net interest expense	-1,835	-1,635	-1,487	-1,313	-1,509
Associates & JCEs					
Other income	68	2,404	0	0	0
Earnings before tax	-615	713	205	-1,016	-522
Income tax	-156	-470	-235	203	104
Net profit after tax	-771	243	-30	-813	-417
Minority interests	67	67	-64	-132	-139
Other items	4,075	0	0	0	0
Preferred dividends					
Normalised NPAT	3,371	311	-94	-945	-557
Extraordinary items					
Reported NPAT	3,371	311	-94	-945	-557
Dividends					
Transfer to reserves	3,371	311	-94	-945	-557
Valuations and ratios					
Reported P/E (x)	12.9	149.3	–	–	–
Normalised P/E (x)	12.9	149.3	-511.4	-50.7	-86.0
FD normalised P/E (x)	12.9	149.3	–	–	–
Dividend yield (%)	–	–	–	–	–
Price/cashflow (x)	22.5	13.8	11.9	9.5	8.6
Price/book (x)	1.9	1.7	1.5	1.6	1.6
EV/EBITDA (x)	16.3	20.8	13.2	15.5	13.5
EV/EBIT (x)	62.4	–	44.2	285.4	95.9
Gross margin (%)	21.5	22.6	20.3	10.4	12.7
EBITDA margin (%)	42.6	29.8	44.3	39.4	39.6
EBIT margin (%)	11.2	-0.5	13.2	2.1	5.5
Net margin (%)	32.7	2.7	-0.7	-6.8	-3.1
Effective tax rate (%)	–	65.9	114.4	–	–
Dividend payout (%)	0.0	0.0	–	–	–
ROE (%)	16.3	1.3	-0.4	-3.5	-2.1
ROA (pretax %)	1.7	-0.1	2.5	0.4	1.1
Growth (%)					
Revenue	5.5	10.8	11.7	9.0	27.8
EBITDA	278.5	-22.6	66.3	-3.2	28.4
Normalised EPS		-91.3	-129.2	–	–
Normalised FDEPS		-91.3	-129.2	–	–

Source: Company data, Nomura estimates

Cashflow statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
EBITDA	4,395	3,404	5,658	5,477	7,035
Change in working capital	-1,103	-1,055	-1,791	439	-749
Other operating cashflow	-1,353	1,017	168	-877	-741
Cashflow from operations	1,938	3,365	4,036	5,039	5,544
Capital expenditure	-2,965	-4,611	-10,000	-15,000	-15,000
Free cashflow	-1,027	-1,245	-5,964	-9,961	-9,456
Reduction in investments	0	0	0	0	
Net acquisitions					
Dec in other LT assets					
Inc in other LT liabilities					
Adjustments	-5,795	1,571	1,983	0	0
CF after investing acts	-6,822	325	-3,982	-9,961	-9,456
Cash dividends					
Equity issue	0	1,016	0	0	0
Debt issue	629	1,235	2,000	5,000	5,000
Convertible debt issue	0	3,853	0	0	0
Others	6,706	10	2,235	0	-139
CF from financial acts	7,335	6,113	4,235	5,000	4,861
Net cashflow	513	6,438	253	-4,961	-4,595
Beginning cash	7,355	7,868	14,306	14,560	9,599
Ending cash	7,868	14,306	14,559	9,599	5,004
Ending net debt	26,957	24,153	25,900	35,860	45,455

Balance sheet (CNYmn)

As at 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Cash & equivalents	7,868	14,306	14,560	9,599	5,004
Marketable securities					
Accounts receivable	3,022	2,467	2,991	2,956	4,644
Inventories					
Other current assets	723	1,774	3,639	3,639	3,639
Total current assets	11,613	18,547	21,190	16,194	13,287
LT investments					
Fixed assets	40,204	38,054	43,587	53,408	62,360
Goodwill					
Other intangible assets	6,367	5,461	5,461	5,461	5,461
Other LT assets	15,464	17,936	16,454	16,454	16,454
Total assets	73,649	79,998	86,692	91,517	97,562
Short-term debt	4,342	2,952	2,952	2,952	2,952
Accounts payable	2,593	1,932	2,530	2,935	3,873
Other current liabilities	2,143	2,244	2,244	2,244	2,244
Total current liabilities	9,078	7,128	7,726	8,131	9,069
Long-term debt	21,906	23,363	25,363	30,363	35,363
Convertible debt	8,577	12,144	12,144	12,144	12,144
Other LT liabilities	10,419	9,625	9,625	9,625	9,625
Total liabilities	49,980	52,261	54,860	60,264	66,202
Minority interest	130	897	897	897	897
Preferred stock	1,081	1,057	3,289	3,428	3,428
Common stock	29,597	31,707	31,707	31,707	31,707
Retained earnings	-6,044	-5,095	-3,232	-3,950	-3,843
Proposed dividends					
Other equity and reserves	-1,094	-829	-829	-829	-829
Total shareholders' equity	23,539	26,840	30,935	30,356	30,462
Total equity & liabilities	73,649	79,998	86,692	91,517	97,562
Liquidity (x)					
Current ratio	1.28	2.60	2.74	1.99	1.47
Interest cover	0.6	0.0	1.1	0.2	0.7
Leverage					
Net debt/EBITDA (x)	6.13	7.10	4.58	6.55	6.46
Net debt/equity (%)	114.5	90.0	83.7	118.1	149.2
Per share					
Reported EPS (CNY)	17.85	1.56	-45.54c	-4.60	-2.71
Norm EPS (CNY)	17.85	1.56	-45.54c	-4.60	-2.71
FD norm EPS (CNY)	17.85	1.56	-45.54c	-4.60	-2.71
BVPS (CNY)	124.62	134.77	150.49	147.67	148.19
DPS (CNY)	0.00	0.00	0.00	0.00	0.00
Activity (days)					
Days receivable	97.5	87.6	78.0	78.0	78.2
Days inventory	0.0	0.0	0.0	0.0	
Days payable	120.4	93.4	80.0	80.0	80.2
Cash cycle	-22.9	-5.7	-2.0	-2.0	-2.0

Source: Company data, Nomura estimates

Company profile

GDS develops and operates data centers in China. The company provides colocation, managed hosting and consulting services to cloud services providers, internet companies, financial institutions and other large domestic private sector or multinational corporations. GDS is the leading carrier-neutral IDC player in China in terms of market share.

Valuation Methodology

Our TP of USD45.70 is derived from 18.8x FY27F EV/EBITDA, in line with the historical average EV/EBITDA. The benchmark index for this stock is the NASDAQ Composite Index.

Risks that may impede the achievement of the target price

Downside risks include: 1) lower-than-expected data center demand related to AI business in China; 2) slower-than-expected data center business expansion in overseas markets; 3) intensified market competition in the data center sector; and 4) supply chain disruption as a result of rising geopolitical risks.

ESG

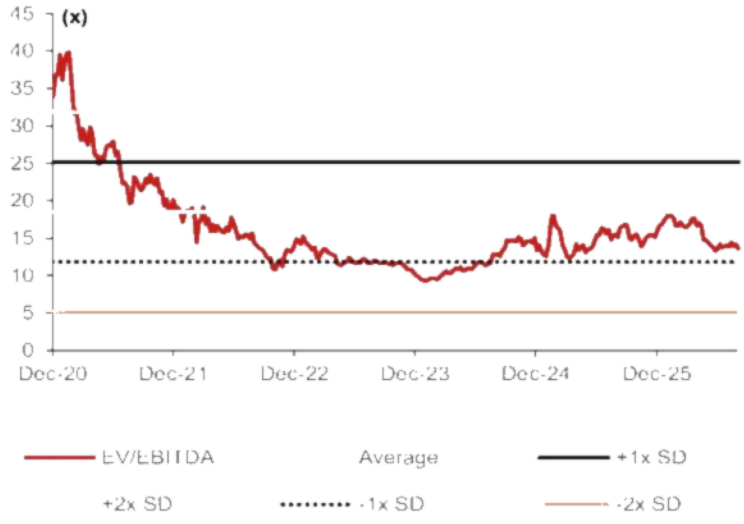
As a leading IDC player in China market, GDS uses energy conservation technology in its data centers. They onsite renewable energy systems and AI-based energy efficiency tool to precisely manage the PUE of its data centers. The company is also making progress in reducing its greenhouse gas emissions by utilizing renewable energy and over 20% of its electricity consumption in 2020 are from renewable energy sources. Meanwhile, GDS also has compliance policies in place that clearly define its compliance requirement. And they are committed to providing all the employees with equal employment opportunities and a good workplace.

Fig. 97: GDS – earnings forecast revisions

	2024 A	2025 A	2026F (NOM)	2027F (NOM)	2028F (NOM)	2026F (OLD)	2027F (OLD)	2028F (OLD)	2026F Diff.	2027F Diff.	2028F Diff.
GDS (in CNYmn)											
Net revenue	10,322	11,432	12,772	13,916	17,783	12,577	13,721	16,993	1.6%	1.4%	4.7%
% y-y	6%	11%	12%	9%	28%	10%	9%	24%			
Service revenue	10,322	11,428	12,341	13,916	17,783	12,147	13,721	16,993	1.6%	1.4%	4.7%
% y-y	6%	11%	8%	13%	28%	6%	13%	24%			
COGS	(8,099)	(8,847)	(10,180)	(12,467)	(15,530)	(9,797)	(11,359)	(13,441)	3.9%	9.7%	15.5%
Gross profit	2,223	2,585	2,592	1,449	2,253	2,780	2,362	3,552			
OPEX	(1,071)	(1,080)	(900)	(1,151)	(1,266)	(1,007)	(1,176)	(1,456)			
Sales & marketing expense	(116)	(149)	(135)	(145)	(186)	(151)	(171)	(211)			
General & admin expense	(918)	(898)	(734)	(974)	(1,040)	(816)	(960)	(1,189)			
R&D expense	(36)	(33)	(31)	(32)	(41)	(40)	(45)	(55)			
Operating profit	1,152	1,505	1,692	298	987	1,773	1,186	2,095			
% y-y	43%	31%	12%	-82%	231%	18%	-33%	77%			
Net finance expense	(1,835)	(1,635)	(1,487)	(1,313)	(1,509)	(1,487)	(1,560)	(1,668)			
Pretax income	(615)	713	205	(1,016)	(522)	386	(374)	427			
% y-y	-84%	-216%	-71%	-595%	-49%	-46%	-197%	-214%			
Net income	3,371	895	1,863	(719)	107	1,938	(678)	204	-3.9%	6.0%	-47.5%
% y-y	-178%	-73%	108%	-139%	-115%	116%	-135%	-130%			
Adjusted gross profit	5,314	5,915	6,067	6,563	8,235	6,190	6,975	8,638	-2.0%	-5.9%	-4.7%
% y-y	4%	11%	3%	8%	25%	5%	13%	24%			
Adjusted EBITDA	4,876	5,403	5,531	5,852	7,410	5,676	6,281	7,664	-2.6%	-6.8%	-3.3%
% y-y	3%	11%	2%	6%	27%	5%	11%	22%			
Ratio Analysis											
GPM	21.5%	22.6%	20.3%	10.4%	12.7%	22.1%	17.2%	20.9%	-1.8pp	-6.8pp	-8.2pp
OPEX	-10.4%	-9.4%	-7.0%	-8.3%	-7.1%	-8.0%	-8.6%	-8.6%	1.0pp	0.3pp	1.4pp
OPM	11.2%	13.2%	13.2%	2.1%	5.5%	14.1%	8.6%	12.3%	-0.8pp	-6.5pp	-6.8pp
Adjusted GPM	51.5%	51.7%	47.5%	47.2%	46.3%	49.2%	50.8%	50.8%	-1.7pp	-3.7pp	-4.5pp
Adjusted EBITDA margin	47.2%	47.3%	43.3%	42.1%	41.7%	45.1%	45.8%	45.1%	-1.8pp	-3.7pp	-3.4pp

Source: Company data, Nomura estimates

Fig. 98: GDS - 12M forward adj. EV/EBITDA bands



Source: Bloomberg Finance L.P., Company data, Nomura estimates

NOMURA

Global Markets Research
11 September 2026

VNET Group VNET.OQ VNET US

EQUITY: TECHNOLOGY

Key beneficiary of AIDC rollout in northern China

Reiterate Buy with a revised TP of USD11.3

Strong order intake announced in 2Q26

VNET delivered revenue and adjusted EBITDA growth of 14.2% and 25.4% y-y, respectively, in 2Q26, largely in line with our estimates (see *our note*). Notably, VNET announced additional order wins of 347MW in 2Q26, including one key order of 345MW from a cloud service provider, bringing YTD total new wholesale bookings to 862MW. New capacity delivery was at 99GW during the quarter, with cumulative wholesale capacity in service reaching over 1GW by the end of the quarter. VNET maintained an unchanged delivery target of 450MW for FY26E (2H26E: 333MW) followed by 252MW delivery for 1H27E. Net additions of wholesale capacity utilized was 57MW in 2Q26, with period-end utilization rate of 73.9%. Management also indicated that the company has secured new land bank of ~1.5GW in 2Q26, including ~500MW in overseas regions. We expect the company to achieve ~1GW order wins within FY26F, supporting its fundamental growth in FY27-28F.

Well-positioned in Inner Mongolia to tap fast-growing AIDC demand

As one of the early-movers in AIDC development in Inner Mongolia, we expect VNET to achieve in-service IT power capacity of ~1GW by mid-2027F in Ulanqab, , which has been mostly contracted with Chinese CSPs. We expect continued AIDC demand in the northern regions of China, owing to attractive power tariffs and low temperatures. According to SCMP (*news link*), ByteDance (unlisted) has been in preliminary talks to expand 5-6GW computing power in Ulanqab over the next two years. We believe VNET is well-positioned to tap the fast-growing AIDC demand given its secured land bank and long-term relationship with key customers.

Reiterate Buy but cut TP to USD11.3

We lower FY26-28F adjusted GPM by 0.2pp for each year, offset by a lower opex-to-sales ratio (we cut by 0.2-0.4pp), leading to a largely unchanged adjusted EBITDA at CNY3.67/4.70/5.99bn for FY26/27/28F. We reiterate our Buy rating but cut our TP to USD11.3, based on a rolled-forward DCF valuation model with cash flows discounted back to end-FY27F. We assume a WACC of 7.0% (unchanged) and terminal growth rate of 2.0% (unchanged). Our TP implies 11.6x 2027F EV/EBITDA and the stock is trading at 10x 2027F EV/EBITDA.

Year-end 31 Dec	FY25		FY26F		FY27F		FY28F
Currency (CNY)	Actual	Old	New	Old	New	Old	New
Revenue (mn)	9,949	11,556	11,552	14,385	14,383	18,028	18,062
Reported net profit (mn)	-257	-1,994	-2,260	433	432	707	657
Normalised net profit (mn)	-257	-1,994	-2,260	433	432	707	657
FD normalised EPS	-96.49c	-7.41	-8.40	1.61	1.60	2.63	2.44
FD norm. EPS growth (%)	-237.1	–	–	–	–	63.6	52.1
FD normalised P/E (x)	–	–	–	–	28.3	–	18.6
EV/EBITDA (x)	11.4	–	10.9	–	10.2	–	9.0
Price/book (x)	2.0	–	2.5	–	2.3	–	2.0
Dividend yield (%)	–	–	–	–	–	–	–
ROE (%)	-4.1	-38.1	-40.5	9.7	8.4	14.0	11.5
Net debt/equity (%)	232.0	556.1	452.5	631.0	559.5	624.5	591.4

Source: Company data, Nomura estimates

Rating Remains	**Buy**
Target price Reduced from USD 13.00	**USD 11.30**
Closing price 8 September 2026	**USD 6.75**
Implied upside	**+67.4%**
Market Cap (USD mn)	1,921.7
ADT (USD mn)	35.9

Relative performance chart



Source: LSEG, Nomura

Research Analysts

China Technology

Ethan Zhang - NIHK
ethan.zhang@nomura.com
+852 2252 2157

Bing Duan - NIHK
bing.duan1@nomura.com
+852 2252 2141

Key data on VNET Group

Performance

(%)	1M	3M	12M		
Absolute (USD)	-7.9	-23.9	-12.1	M cap (USDmn)	1,921.7
Absolute (USD)	-7.9	-23.9	-12.1	Free float (%)	26.1
Rel to NASDAQ COMPOSITE	-6.9	-25.8	-33.3	3-mth ADT (USDmn)	35.9

Income statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Revenue	8,259	9,949	11,552	14,383	18,062
Cost of goods sold	-6,427	-7,757	-9,237	-11,429	-14,171
Gross profit	1,832	2,192	2,315	2,954	3,891
SG&A	-1,163	-1,412	-1,373	-1,695	-2,116
Employee share expense					
Operating profit	669	780	941	1,259	1,775
EBITDA	2,276	2,706	3,653	4,677	5,972
Depreciation	-1,503	-1,822	-2,608	-3,360	-4,140
Amortisation	-104	-104	-104	-58	-58
EBIT	669	780	941	1,259	1,775
Net interest expense	-373	-561	-750	-982	-1,245
Associates & JCEs	8	7	0	0	0
Other income	178	198	300	400	500
Earnings before tax	483	424	491	677	1,030
Income tax	-234	-558	-652	-102	-154
Net profit after tax	248	-133	-161	575	875
Minority interests	-65	-123	-2,099	-144	-219
Other items					
Preferred dividends	0	0	0	0	0
Normalised NPAT	183	-257	-2,260	432	657
Extraordinary items					
Reported NPAT	183	-257	-2,260	432	657
Dividends	0	0	0	0	0
Transfer to reserves	183	-257	-2,260	432	657
Valuations and ratios					
Reported P/E (x)	69.1	–	–	28.3	18.6
Normalised P/E (x)	69.1	-50.3	-5.4	28.3	18.6
FD normalised P/E (x)	69.1	–	–	28.3	18.6
Dividend yield (%)	–	–	–	–	–
Price/cashflow (x)	6.3	6.7	4.6	3.4	2.7
Price/book (x)	2.0	2.0	2.5	2.3	2.0
EV/EBITDA (x)	10.7	11.4	10.9	10.2	9.0
EV/EBIT (x)	36.1	39.4	42.1	37.8	30.2
Gross margin (%)	22.2	22.0	20.0	20.5	21.5
EBITDA margin (%)	27.6	27.2	31.6	32.5	33.1
EBIT margin (%)	8.1	7.8	8.1	8.8	9.8
Net margin (%)	2.2	-2.6	-19.6	3.0	3.6
Effective tax rate (%)	48.5	131.5	132.7	15.0	15.0
Dividend payout (%)	0.0	–	–	0.0	0.0
ROE (%)	3.0	-4.1	-40.5	8.4	11.5
ROA (pretax %)	2.3	2.3	2.2	2.4	3.0
Growth (%)					
Revenue	11.4	20.5	16.1	24.5	25.6
EBITDA		18.9	35.0	28.0	27.7
Normalised EPS		-237.1	–	–	52.1
Normalised FDEPS		-237.1	–	–	52.1

Source: Company data, Nomura estimates

Cashflow statement (CNYmn)

Year-end 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
EBITDA	2,276	2,706	3,653	4,677	5,972
Change in working capital	3,111	1,281	61	-374	-500
Other operating cashflow	-3,381	-2,068	-1,082	-664	-879
Cashflow from operations	2,005	1,919	2,632	3,639	4,593
Capital expenditure	-4,924	-7,657	-11,275	-11,250	-10,125
Free cashflow	-2,918	-5,738	-8,643	-7,611	-5,532
Reduction in investments					
Net acquisitions	205	707	-187	-187	-187
Dec in other LT assets					
Inc in other LT liabilities					
Adjustments	328	-1,062	0	0	0
CF after investing acts	-2,385	-6,093	-8,830	-7,798	-5,719
Cash dividends					
Equity issue	0	0	951	0	0
Debt issue	2,751	4,834	8,000	8,000	7,000
Convertible debt issue	-4,262	3,085	0	-1,962	0
Others	3,145	2,206	0	0	0
CF from financial acts	1,634	10,124	8,951	6,038	7,000
Net cashflow	-751	4,031	121	-1,760	1,281
Beginning cash	2,244	1,492	5,524	5,644	3,884
Ending cash	1,492	5,524	5,644	3,884	5,165
Ending net debt	10,182	14,426	22,306	30,104	35,823

Balance sheet (CNYmn)

As at 31 Dec	FY24	FY25	FY26F	FY27F	FY28F
Cash & equivalents	1,492	5,524	5,644	3,884	5,165
Marketable securities	0	0	0	0	0
Accounts receivable	1,656	2,222	2,283	2,862	3,619
Inventories	0	0	0	0	0
Other current assets	3,672	3,706	3,706	3,706	3,706
Total current assets	6,820	11,452	11,634	10,453	12,490
LT investments	0	0	0	0	0
Fixed assets	17,217	22,776	31,443	39,333	45,318
Goodwill	0	0	0	0	0
Other intangible assets	1,404	2,005	2,096	2,233	2,370
Other LT assets	6,916	8,362	8,354	8,346	8,338
Total assets	32,357	44,594	53,527	60,365	68,517
Short-term debt	2,009	3,232	3,232	3,232	3,232
Accounts payable	709	742	864	1,069	1,326
Other current liabilities	6,625	8,474	8,474	8,474	8,474
Total current liabilities	9,343	12,447	12,569	12,774	13,031
Long-term debt	7,767	11,580	19,580	27,580	34,580
Convertible debt	1,898	5,139	5,139	3,177	3,177
Other LT liabilities	6,428	6,864	6,864	6,864	6,864
Total liabilities	25,436	36,030	44,152	50,395	57,651
Minority interest	555	2,346	4,445	4,589	4,808
Preferred stock	0	0	0	0	0
Common stock	17,299	17,360	18,331	18,351	18,371
Retained earnings	-10,860	-11,126	-13,386	-12,954	-12,298
Proposed dividends					
Other equity and reserves	-73	-16	-16	-16	-16
Total shareholders' equity	6,366	6,218	4,929	5,381	6,057
Total equity & liabilities	32,357	44,594	53,526	60,365	68,517
Liquidity (x)					
Current ratio	0.73	0.92	0.93	0.82	0.96
Interest cover	1.8	1.4	1.3	1.3	1.4
Leverage					
Net debt/EBITDA (x)	4.47	5.33	6.11	6.44	6.00
Net debt/equity (%)	159.9	232.0	452.5	559.5	591.4
Per share					
Reported EPS (CNY)	70.38c	-96.49c	-8.40	1.60	2.44
Norm EPS (CNY)	70.38c	-96.49c	-8.40	1.60	2.44
FD norm EPS (CNY)	70.38c	-96.49c	-8.40	1.60	2.44
BVPS (CNY)	24.46	23.37	18.32	20.00	22.52
DPS (CNY)	0.00	0.00	0.00	0.00	0.00
Activity (days)					
Days receivable	74.5	71.1	71.2	65.3	65.7
Days inventory	0.0	0.0	0.0	0.0	
Days payable	39.9	34.1	31.7	30.9	30.9
Cash cycle	34.6	37.0	39.4	34.4	34.7

Source: Company data, Nomura estimates

Company profile

VNET develops and operates internet data centers (IDCs) in China. The company provides colocation, managed hosting and consulting services to cloud services providers, internet companies, financial institutions and other large domestic private sector or multinational corporations.

Valuation Methodology

Our target price of USD11.30 is based on a DCF valuation model. We assume a WACC of 7.0% and terminal growth rate of 2.0%. Cash flows are discounted back to end-2027F. Our TP implies 11.6x 2027F EV/EBITDA. The benchmark index is NASDAQ.

Risks that may impede the achievement of the target price

Downside risks include: (1) slower-than-expected ramp-up of utilization rates for newly-built data centers; and (2) faster-than-expected decline in MRR on intensified market competition.

ESG

VNET has been focusing on building IDCs with industry-leading standards to maximize the energy utilization. The company also targets to increase the usage of renewable energy to power its data centers.

Fig. 99: VNET - earnings forecasts revisions

(CNY mn)	2024 Actual	2025 Actual	2026F Forecast	2027F Forecast	2028F Forecast	2026F Old	2027F Old	2028F Old	2026F Diff.	2027F Diff.	2028F Diff.
Net revenue	**8,259**	**9,949**	**11,552**	**14,383**	**18,062**	**11,556**	**14,385**	**18,028**	0.0%	0.0%	0.2%
y-y	*11.4%*	*20.5%*	*16.1%*	*24.5%*	*25.6%*	*16.1%*	*24.5%*	*25.3%*			
COGS	(6,427)	(7,757)	(9,237)	(11,429)	(14,171)	(8,928)	(11,042)	(13,658)			
Gross Profit	**1,832**	**2,192**	**2,315**	**2,954**	**3,891**	**2,628**	**3,343**	**4,370**			
y-y	*41.8%*	*19.7%*	*5.6%*	*27.6%*	*31.7%*	*19.8%*	*27.2%*	*30.7%*			
Adjusted cash gross profit	**3,337**	**4,218**	**4,955**	**6,242**	**7,928**	**4,980**	**6,271**	**7,950**	-0.5%	-0.5%	-0.3%
y-y	*12.1%*	*26.4%*	*17.5%*	*26.0%*	*27.0%*	*18.1%*	*25.9%*	*26.8%*			
Sales and marketing expenses	(264)	(279)	(255)	(332)	(417)	(266)	(346)	(434)			
General and administrative expenses	(659)	(797)	(717)	(907)	(1,140)	(752)	(951)	(1,192)			
R&D expenses	(247)	(261)	(326)	(406)	(510)	(321)	(385)	(464)			
Operating profit / (loss)	**669**	**780**	**941**	**1,259**	**1,775**	**1,168**	**1,662**	**2,281**			
Depreciation and amortization	1,598	1,926	2,711	3,418	4,198	2,435	2,936	3,590			
Share-based compensation	149	26	20	20	20	20	20	20			
Others	14	246	-	-	-	50	100	100			
Adjusted EBITDA	**2,430**	**2,978**	**3,673**	**4,697**	**5,992**	**3,674**	**4,718**	**5,991**	0.0%	-0.4%	0.0%
y-y	*19.1%*	*22.6%*	*23.3%*	*27.9%*	*27.6%*	*23.4%*	*28.4%*	*27.0%*			
Finance costs (net)	(373)	(561)	(750)	(982)	(1,245)	(734)	(983)	(1,171)			
Foreign exchange gain (loss)	(19)	6	-	-	-	-	-	-			
Profit before taxation	**483**	**424**	**491**	**677**	**1,030**	**434**	**678**	**1,110**			
Taxation	(234)	(558)	(652)	(102)	(154)	(565)	(102)	(166)			
Net profit	**248**	**(133)**	**(161)**	**575**	**875**	**(130)**	**577**	**943**			
Non-controlling interest	(65)	(123)	(2,099)	(144)	(219)	(1,863)	(144)	(236)			
Profit for owners of the company	**183**	**(257)**	**(2,260)**	**432**	**657**	**(1,994)**	**433**	**707**	13.4%	-0.2%	-7.2%
Key ratios											
Adjusted GPM	40.4%	42.4%	42.9%	43.4%	43.9%	43.1%	43.6%	44.1%	-0.2pp	-0.2pp	-0.2pp
OPEX-to-Sales	14.2%	13.4%	11.2%	11.4%	11.4%	11.6%	11.7%	11.6%	-0.4pp	-0.3pp	-0.2pp
Adjusted EBITDA margin	29.4%	29.9%	31.8%	32.7%	33.2%	31.8%	32.8%	33.2%	0.0pp	-0.1pp	-0.1pp

Source: Company data, Nomura estimates

Fig. 100: VNET - DCF valuation

CNYmn	2023	2024	2025F	2026F	2027F	2028F	2029F	2030F	2031F	2032F	2033F	2034F	2035F	Terminal
Period						1	2	3	4	5	6	7	8	
Revenue	7,413	8,259	9,949	11,552	14,383	18,062	21,919	25,950	29,682	33,254	36,887	40,578	44,362	
Profit after tax	(2,644)	183	(257)	(2,260)	432	657	675	976	1,269	1,583	1,919	2,257	2,676	
Operating cash flow	2,063	2,005	1,919	2,632	3,639	4,593	5,279	6,293	7,280	8,221	9,133	10,037	11,029	
CAPEX	(2,967)	(4,924)	(7,657)	(11,275)	(11,250)	(10,125)	(9,113)	(8,201)	(7,381)	(6,643)	(6,510)	(6,380)	(6,252)	
FCF	(904)	(2,918)	(5,738)	(8,643)	(7,611)	(5,532)	(3,834)	(1,908)	(101)	1,578	2,623	3,657	4,777	97,499
	130%	223%	97%	51%	-12%	-27%	-31%	-50%	-95%	-1666%	66%	39%	31%	
Discount factor						0.93	0.87	0.82	0.76	0.71	0.67	0.62	0.58	0.58
PV						(5,171)	(3,349)	(1,558)	(77)	1,125	1,748	2,278	2,781	56,758

Assumptions	
Risk-free rate	1.7%
Market risk premium	12.4%
Beta	1.1
Terminal growth rate	2.0%
WACC	7.0%

PV of corporate value		54,536
+ cash		3,884
- borrowing		(33,988)
- minority interest		(2,878)
PV of equity value		21,554
# of ADS outstanding		285
USDCNY		6.72
Valuation	CNY	76

USD
11.3

Source: Company data, Nomura estimates