VIG
ASSET MANAGEMENT

MEGATREND MONITOR

The future won't wait, and we're here to help you make sense of it.

August 2026

Welcome to the world of Megatrend Monitor!

Our newly launched monthly newsletter serves as an exclusive guide from the portfolio managers and analysts of the VIG Megatrend Fund Family in the ever-changing global market. Our goal is to filter out the noise and analyze the past month's key innovations, corporate breakthroughs, and market movements within the megatrend universe. Each month, our experts scrutinize the most important events in the innovation ecosystem, from corporate fundamentals to global trend shifts. With this publication, our investors – particularly partners in the VIG Megatrend and InnovationTrend funds – gain direct insight into the logic behind our investment decisions, uncovering the technological and social drivers propelling tomorrow's winners.

Péter Richter
Senior Portfolio Manager

Balázs Csúr
Junior Analyst

Contents

4 MegaTrends: another wave of rotation – software, copper, China advances?

7 Nvidia has fallen behind

9 The central bank of AI

13 US AI labs may be in trouble

16 A red sun in the sky

18 MegaTrend Equity Fund newsletter – What we wrote to our investors in July

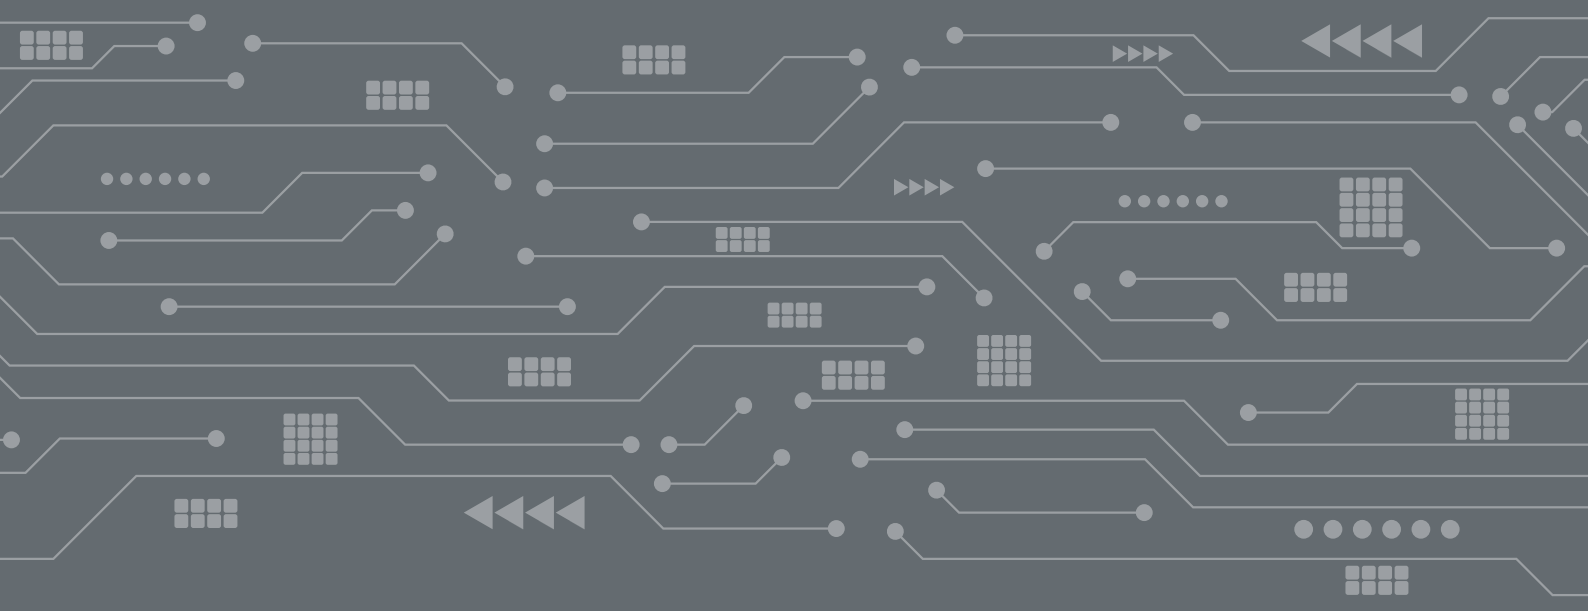

MegaTrends: another wave of rotation – software, copper, China advances?

The megatrend market has broadened further: capital is not leaving growth stories, it is flowing out of chips and into cloud services, software and the areas needed to build out networks. The global structure is stable, but for now investors are looking for the winners of the physical and software implementation phase.

Within the megatrend universe, the primary hunting ground of our VIG MegaTrend Equity Investment Fund, all three main strategies were profitable over the past 30 days, yet a substantial reshuffling took place behind the scenes. The race was won this time by demographics-focused megatrends (+4.0%), closely followed by the innovation-focused (+3.8%) and the environment-focused categories (+2.7%).

Performance of the megatrend strategies over the past 30 days

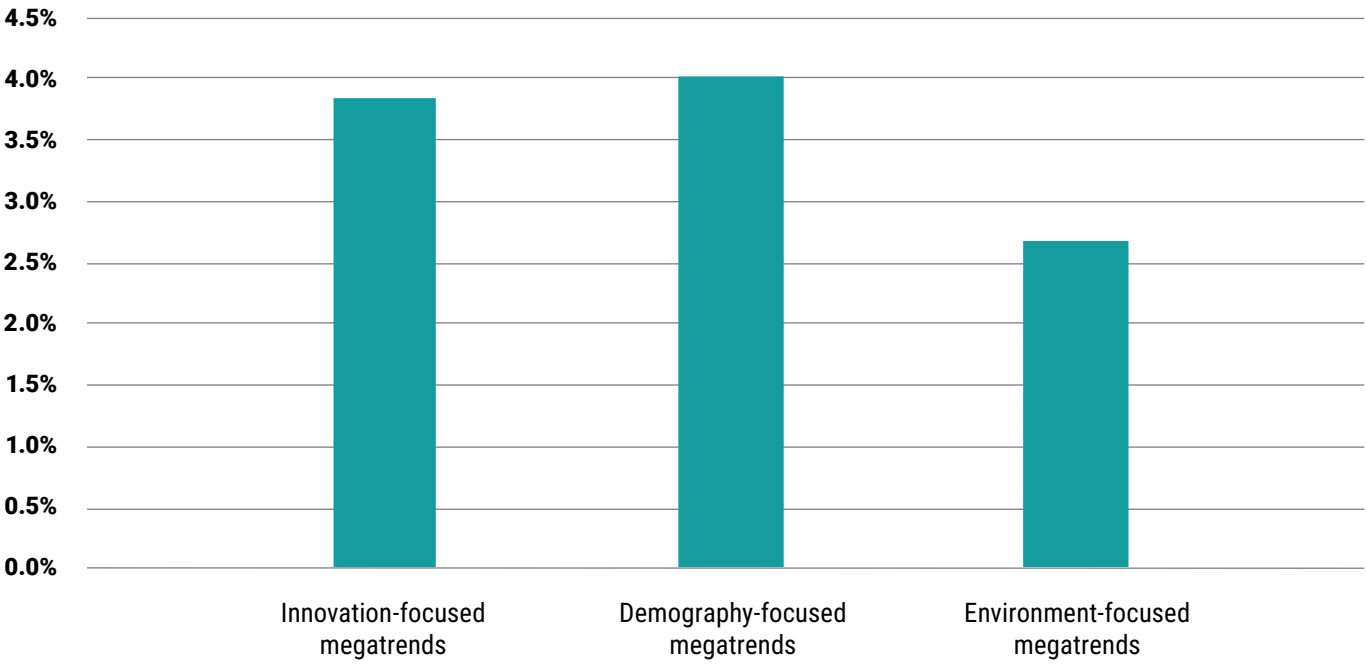

Source: Bloomberg, VIG Asset Management

Sub-trend performance and the drivers of the rotation

Behind the reshuffling relative to the previous month lay a classic, mature-market rotation wave and several structural macro drivers. Investors did not withdraw capital from megatrend investments; they began instead to reallocate existing gains from the "overheated" segments into areas that are still lagging but hold high growth potential.

- From hardware to software. The stars of the year so far, the semiconductor manufacturers (Semiconductors: -8.9%), went through a significant correction, while cloud providers (Cloud Computing, +15.0%) and software companies (Software, +13.6%) soared. Chipmaker shares had become extremely overbought in the first half of the year (even after the correction, semiconductor stocks were still up 76% year to date in mid-August), but when the second-quarter earnings reports came in, although they were excellent, they only beat analyst expectations marginally therefore, investors began to take profits. Capital did not leave the AI universe, however; it took a new direction. The market started to price in that, after the accumulation of hardware (GPUs), the software deployment and cloud infrastructure phase can begin. Capital flowed towards those enterprise software vendors and cloud providers that have started to monetise AI integrations in practice.

- Copper and the energy transition. The physical constraint on AI (Energy Transition / Copper (+16.9%)) was the strongest performing sub-trend of the month. Building AI data centres, upgrading networks and constructing electrical infrastructure requires an enormous amount of copper. And the market has come face to face with the physical reality that global copper mining is struggling to keep pace with this demand.

- The Chinese turnaround. After months of decline, oversold Chinese technology stocks took flight: the China Internet sub-trend produced a spectacular rebound (+10.4%). The Chinese technology sector had become oversold in recent times because of geopolitical tensions and the slowdown in Chinese economic growth. Investors have now realised that the valuations of Chinese tech giants (such as Alibaba and Baidu) have fallen to historic lows, even as AI development and revenue growth have picked up.

- Moderate pressure on renewable energy. The classic green sectors continued to struggle (Clean & Renewable Energy: -6.1%, Solar: -5.6%, Hydrogen: -5.0%), yet thanks to the surge in copper and raw materials, environment-focused megatrend investments still managed to close in positive territory overall. Although the energy transition is inevitable in the long run, renewable energy projects continue to suffer from the high cost of capital caused by the rising interest rate environment, and from other problems like construction delays, for example, as well.

Monthly and year-to-date (YTD) performance leaderboard

In the year-to-date (YTD) ranking the semiconductor sector (Semiconductors, +75.8%) still holds first place, but its lead has narrowed. Biotechnology and healthcare remain in the leading group after last month's rally (Genomic Revolution: +53.0%, Alternative Medicine: +43.9%). Thanks to China's "resurrection", the YTD loss of China Internet has eased from -25.1% to -14.5%.

		1 month		YTD
Top 1	Energy Transition/ Copper	**16.9%**	Semiconductors	**75.8%**
Top 2	Cloud Computing	**15.0%**	Genomic Revolution	**53.0%**
Top 3	Software	**13.6%**	Alternative Medicine	**43.9%**
Top 4	China Internet	**10.4%**	Cybersecurity	**40.8%**
Top 5	Resource Scarcity/ Natural Res.	**9.9%**	Quantum Computing	**38.3%**
Bottom 5	Electric Vehicles	**-4.7%**	Medical Devices	**-10.8%**
Bottom 4	Clean Energy/ Hydrogen	**-5.0%**	FinTech	**-11.6%**
Bottom 3	Clean Energy/ Solar	**-5.6%**	EM Internet	**-13.3%**
Bottom 2	Clean & Renewable Energy	**-6.1%**	China Internet	**-14.5%**
Bottom 1	Semiconductors	**-8.9%**	Social Media	**-14.8%**

Source: Bloomberg, VIG Asset Management

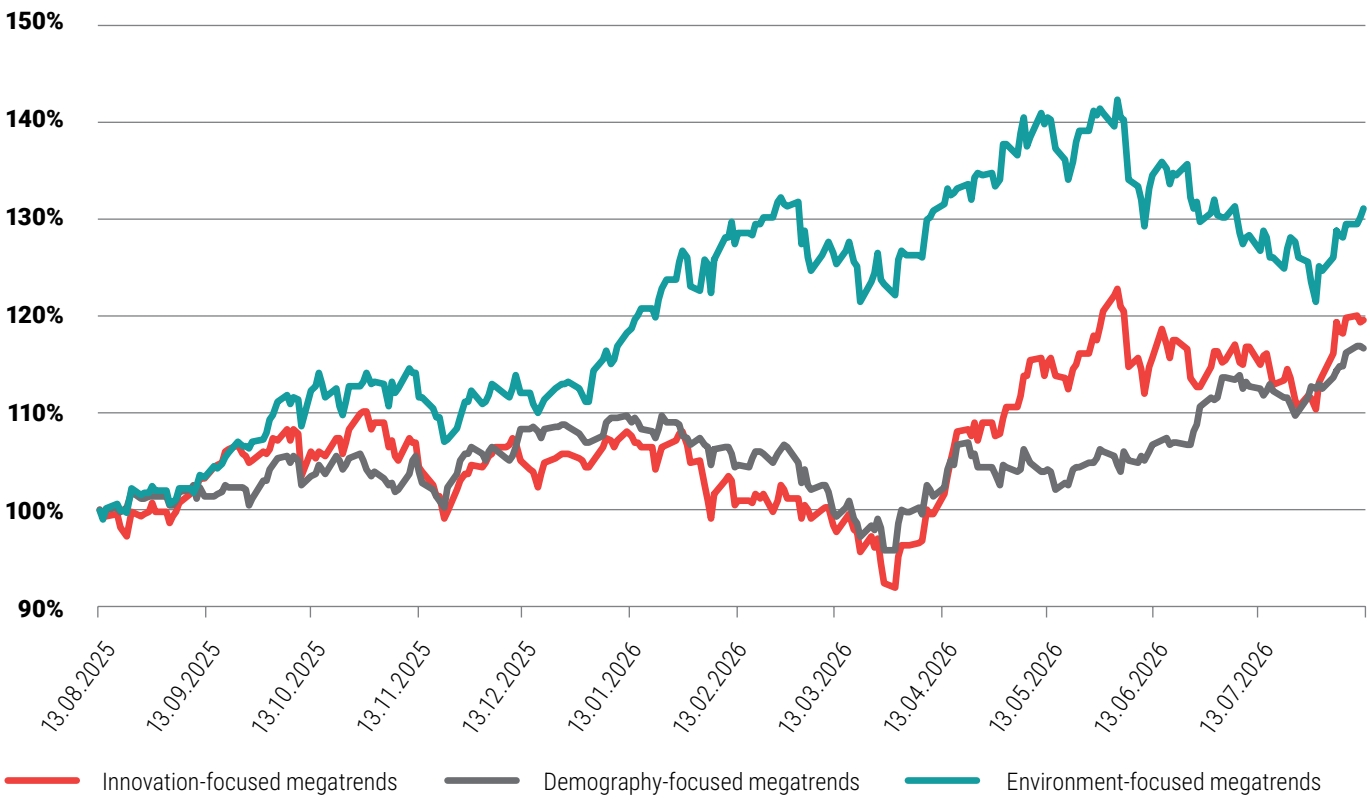

Source: Bloomberg, VIG Asset Management

Nvidia has fallen behind

For a long time the clear winner of the AI trade that took flight after the 2022 bear market was Nvidia, the American maker of high-performance processors. The company's share price rose almost seventeenfold from its low of $11.2 in the autumn of 2022 to the end of 2025. Anyone looking only at the returns of the past year, however, sees a laggard: the stock gained just 18% against a rise of more than 100% in the broader chip sector. There are several reasons for this. One is simple size mathematics: a market capitalisation of nearly $5 trillion in itself limits how much growth can realistically be expected from the world's most valuable company. There are, however, a few more interesting factors that may raise genuine doubts about Nvidia's market position and long-term return potential.

Custom hyperscaler chips

The technology giants competing in the AI race are all working on chips of their own. Nvidia's chips are very good general-purpose tools for AI computation, but there are specialised computing tasks where performance can be improved further. In order to avoid the extremely high profit margin on Nvidia's products, Google, Amazon, Microsoft, Meta and OpenAI have all created their own chip. These are typically more powerful or more economical than Nvidia's chips in some narrow field of application, and they reduce companies' dependence on the chip giant's monopoly. The value chain of ASIC (Application Specific Integrated Circuit) chips bypasses Nvidia. Tech companies design these chips jointly with the American firms Marvell and Broadcom, leaders in the manufacturing of custom artificial intelligence chips and high-speed networking equipment, while 95% of production is handled by Taiwan's TSMC. It is important to note, however, that their use is for now almost entirely confined to internal purposes, and they only represent a threat to the medium- and long-term growth of Nvidia's addressable market rather than to its current market position.

One of the leading AI labs, Anthropic, has access to more than one million Google TPUs, and there are models it trains and runs primarily on TPUs. (The TPU is Google's own application-specific chip, designed expressly to accelerate machine learning and matrix mathematics workloads.) This supports the view that Nvidia is not indispensable in the AI chip market.

Company	Custom chip	Purpose	Design partner	Manufacturer (foundry)
Google	TPU v7 (Ironwood) + v6e (Trillium)	Training + inference	Broadcom + MediaTek	TSMC
Meta	MTIA 400	Internal inference	Broadcom	TSMC
Amazon	Trainium 3	Training + inference	Marvell + in-house design	TSMC
Microsoft	Maia 100	Azure inference	Broadcom	TSMC
OpenAI	Titan (in development)	Frontier model training	Broadcom	TSMC

Source: Bloomberg, VIG Asset Management

Training vs. inference

When we talk about AI data centres, they generally do two main things:

1. Training is the process in which artificial intelligence models learn patterns from vast data sets (text, images or video) and build up the linguistic or visual knowledge they need in order to function. This involves moving, reading and writing enormous quantities of data, and it is a continuous task lasting weeks or even months. Besides raw computing power, the decisive bottleneck is how well the chips can communicate with one another. A frontier model does not fit on a single chip; it is typically distributed across more than ten thousand accelerators*, which have to synchronise with each other at every training step. Chip-to-chip communication is the area where Nvidia is strongest, which is why it is expected to remain dominant in model training over the long term.

2. Inference, by contrast, is the running of the finished model. A question comes in and an answer goes out. Inference has a compute-intensive phase, when the model reads and processes the question, and a memory-intensive phase, when it generates the answer token by token. The latter depends not only on computing power but also on memory bandwidth. To produce every new token, the chip has to read the model's activated weights from memory. So however much raw compute capacity sits on the chip, if it cannot load the weights fast enough the arithmetic units will be "bored". The inference phase requires far less that ten thousand chips to operate, as it only needs to operate a single coherent machine. Memory bandwidth, energy, and cost efficiency per unit of performance under a given response-time constraint matter far more.

The centre of gravity of AI compute has been shifting from training towards inference for years, especially thanks to agentic models and models that reason through logical steps. Agentic artificial intelligence is a program that is able to pursue goals, use software or other tools, and carry out operations with a certain degree of autonomy. Today inference accounts for roughly two-thirds of compute, and this is the workload where the hyperscalers' own chips (these are the giant tech companies that build and operate digital infrastructure on an enormous scale: millions of servers in hundreds of data centres) can compete more easily with Nvidia's ecosystem. The further we move towards the world of inference, the less Nvidia will be the only possible answer.

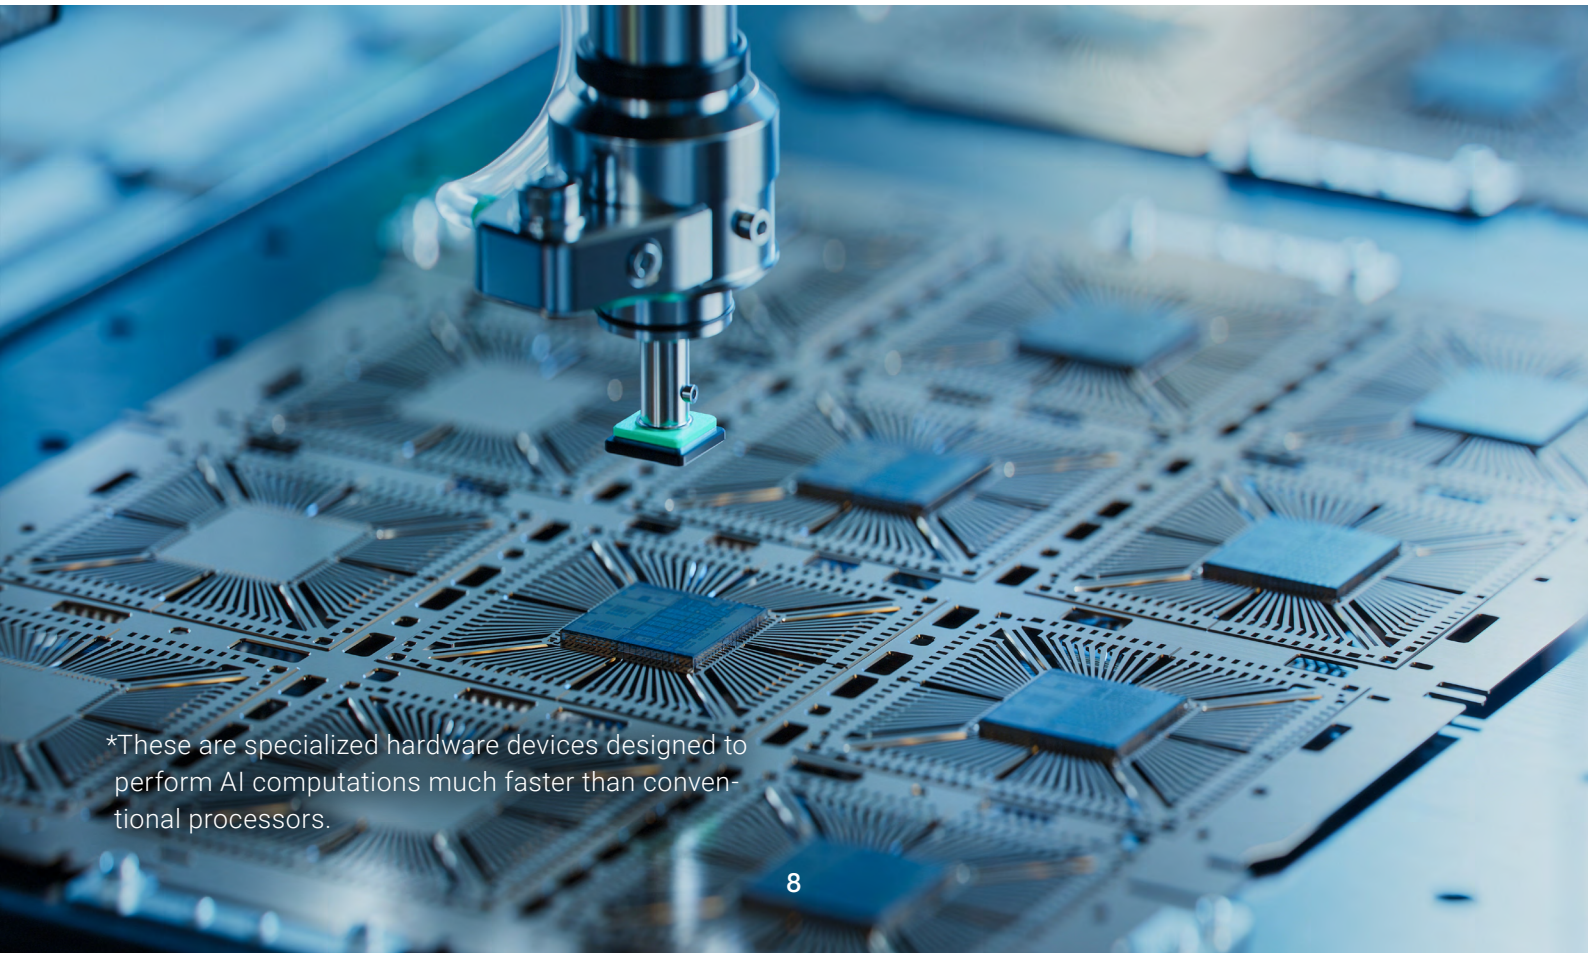

*These are specialized hardware devices designed to perform AI computations much faster than conventional processors.

The central bank of AI

On the evening of 10 August, Nvidia announced that it had signed letters of intent with six large financial institutions on the creation of so-called "compute financing" platforms. The partners are the asset managers Apollo, BlackRock, Blackstone, Brookfield and Kohlberg Kravis Roberts & Co. (KKR), along with the investment bank Goldman Sachs. The goal is to mobilise more than $500 billion of third-party capital for AI infrastructure. Jensen Huang, Nvidia's chief executive, wants to turn the AI accelerators sold by his company (these are specialised hardware devices that perform artificial intelligence, machine learning and neural network computations far faster than conventional processors) and chip systems – and, in a broader sense, computing capacity itself – into an investable asset class.

Guarantees and credit support

It is not unprecedented for Nvidia or another large AI infrastructure player to reach for creative financing methods. These typically have two purposes. One is to improve, through the provision of guarantees, a third party's (often its own customer's) ability to raise capital and the terms on which it can do so. The other is to boost short-term demand for chips through direct vendor financing and product-for-equity swaps. On the supplier side, Nvidia and its main American rival AMD are the biggest users of these techniques. At the other end of the chain stand the NeoCloud companies (specialised cloud providers optimised expressly for training and running large AI models) such as CoreWeave, Nebius and IREN, along with frontier AI labs such as OpenAI and Anthropic, they are the beneficiaries of these structures. Meta, the owner of Facebook, and Oracle, a leader in database management, form a separate category. These companies use SPVs (special purpose vehicles) to finance expensive hardware and the associated loans off their own balance sheets. The purpose of the SPVs is to keep these off the balance sheets of tech giants that operate with relatively few physical assets of their own, such as real estate, machinery or data centres.

The overwhelming majority of the AI build-out has so far been funded from the operating cash flow of the hyperscalers mentioned above, but over the past year this appears to be reversing. First at Oracle, then at Meta, and by now at Google too, free cash flow, which could be the engine of further AI investment , is approaching zero. The new financing platforms offer a solution to this.

The end of hyperscaler self-funding

Aggregate LTM operating cash flow and capex, MSFT, GOOGL, AMZN, META, ORCL

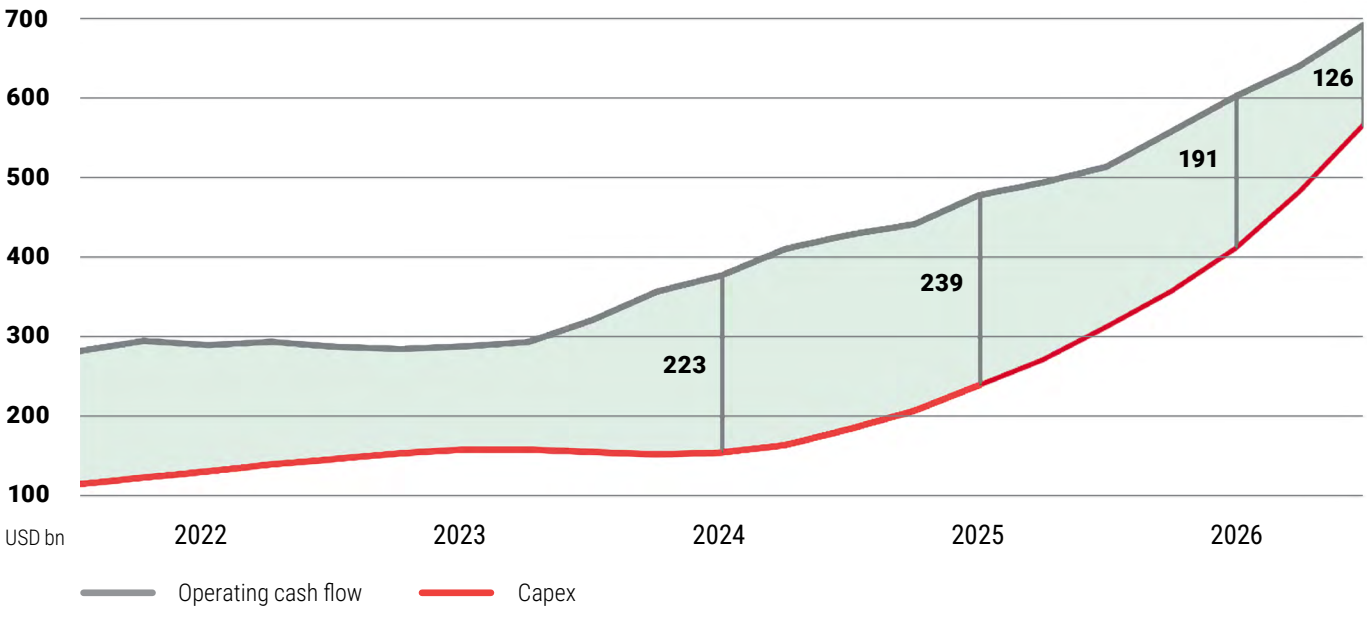

Source: Bloomberg, VIG Asset Management

This time the money comes not from Nvidia's own cash flow but primarily from insurers, pension funds and retail wealth reached through bank distribution. Guarantees are crucial here as well. An insurer can hardly buy speculative assets to cover its annuity liabilities, so securitised computing capacity has to be of IG (investment grade) quality. A NeoCloud company on its own carries a weaker ("sub-IG") rating, so through its guarantees Nvidia lends its own AA/Aa1 rating in order to make these securities representing computing capacity sellable. Even though the money comes from external players, Nvidia's guarantee burden grows further as a result.

Customer concentration is also an important consideration for Nvidia. From the latest quarterly reports we know that four direct customers (not named individually) accounted for more than 60% of revenues. Direct customers are typically manufacturers and system integrators, but behind them stand the same hyperscaler companies as indirect buyers, the ones talking in their own reports about capital expenditure running into the hundreds of billions. The declared aim of the backstop guarantee programmes and the new financing platforms is to open up the compute market beyond the large hyperscalers and AI labs. The current customer base is increasingly focused on developing and using its own chips so that it does not have to pay Nvidia's enormous margin. For Nvidia, the new platform is also a form of defence against an increasingly less loyal customer base developing TPU, Trainium and MTIA, proprietary hardware accelerators specialised for artificial intelligence and developed by the world's largest technology giants.

Pricing chaos in the market for computing capacity

Estimates suggest that by 2029 there could be more than $7 trillion of AI debt outstanding, which would make it the world's second largest asset-backed credit market after the US mortgage market ($13 trillion). According to the latest figures, annual AI capex will be well above $2 trillion in 2028, while cumulative capex between 2024 and 2029 will reach $11 trillion. The primary source of this may be the credit market.

One would therefore assume that the market in question has a mature pricing methodology, standardised futures products and a transparent market picture, like other mature asset classes of a similar trillion-dollar scale. That is not the case. There is no standard GPU rental price index, deals are bilateral, and no derivatives market exists.

The rise of AI debt

Quarterly total debt; CoreWeave includes operating lease liabilities

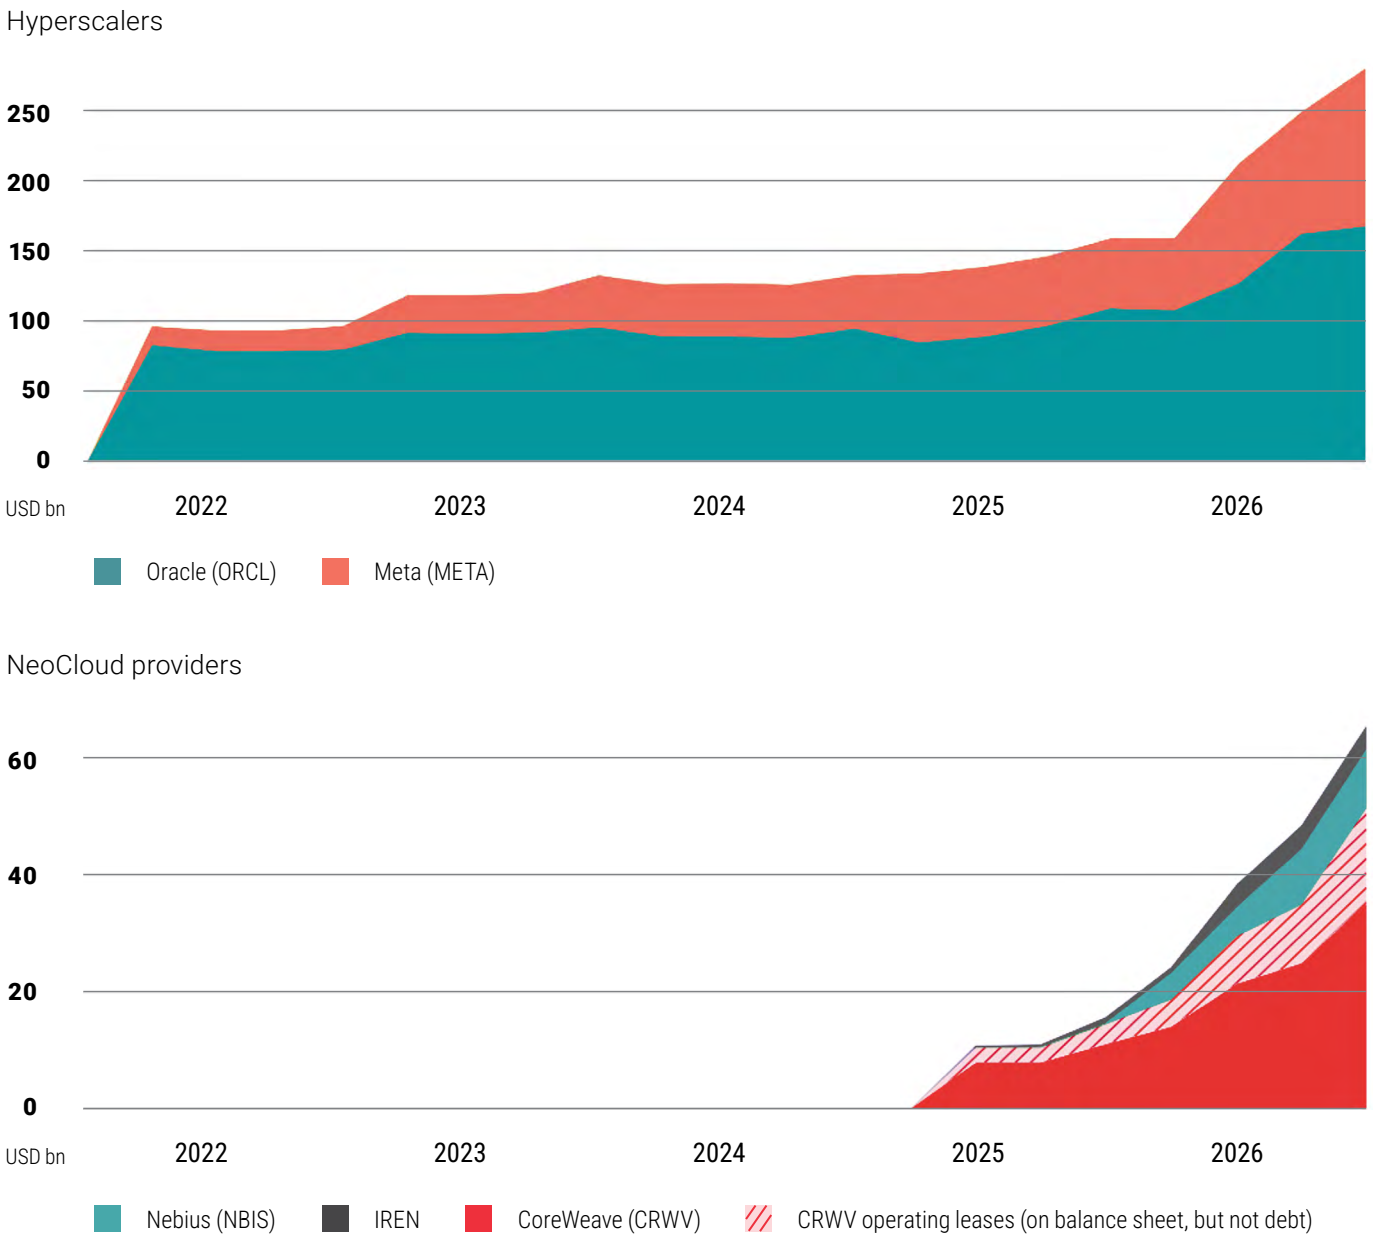

Source: Bloomberg; CoreWeave 10-Q/earnings releases (operating leases); VIG Asset Management

The example of xAI illustrates the situation well. In May, Anthropic contracted for the entire capacity of Colossus 1 for $1.25 billion a month, while in June Google contracted for roughly 110,000 GPUs (graphics chips) for $920 million a month. For Google this works out at about $11.5 per GPU per hour. In Anthropic's case the figure depends on the source, because the GPU count is variously put at between 220,000 and 325,000, so the implicit price ranges from $5.3 to $7.8. The top of the pricing range is double the bottom even for the same provider. Pricing depends primarily on the mix of GPU generations and the CPU-GPU-memory configuration, but the details are often not public, or can only be reverse-engineered with difficulty. The market for computing capacity is not transparent. The various cloud providers all lease out their capacity on markedly different terms and at different prices.

Wrong-way risk

And this is where things start to become worrying. Nvidia's guarantee becomes callable at precisely the moment when the collateral behind it is worth the least. If AI demand slows, the secondary market value of GPUs falls, the cash flow of the NeoClouds deteriorates and the guarantees all fall due at once. Credit market terminology calls this wrong-way risk. Nvidia is essentially writing a put option on the residual value of its own product, and this position turns short exactly when its own fundamentals are weakening too. As long as the guarantees are not called, they sit as off-balance-sheet contingent liabilities, so they do not show up in the balance sheet total. With expected revenue of roughly $216 billion and free cash flow of $97 billion for 2026, today's scale is still comfortably bearable, but this buffer does not scale together with the size of the platform. If the exposure grows materially, the first signal will come not from the rating agencies but from Nvidia's CDS spread (credit default swap premium) and the spread of its bonds over Treasuries. These are worth watching, because the credit market reprices structural risk faster than the equity market.

Another important question is what revenue actually measures. If Nvidia uses a guarantee to make a buyer financeable and that buyer then purchases chips from it, then part of the sales volume reflects not independent demand but Nvidia's own willingness to extend credit. The problem is that part of what the market sees as genuine demand is in fact driven by easily available financing. The two are almost impossible to separate from the outside, because Nvidia does not disclose what share of its revenues is backed by financial guarantees or other collateral.

Finally, a systemic observation

The structure creates strong mutual dependence between Nvidia and the AI ecosystem. As long as the market is growing, this benefits everyone. If, however, overcapacity develops, the large cloud providers will rein in their investments, the revenues of the NeoCloud companies will fall, and demand for Nvidia's chips will decline as well. That is precisely when the guarantees undertaken by Nvidia may be triggered: the company would have to rescue partners whose failure would further reduce its own revenues. In other words, the entire system can start moving downwards at the same time. Nvidia may be the "central bank" of the AI market, but unlike a real central bank it cannot create money.

The sharing of risk is also only partly real. Within a few years, banks, SPVs and several hundred billion dollars of financing may all rest on the same fundamental risk: the durability of demand for AI infrastructure and, ultimately, Nvidia's financial strength. In the short term this does not represent a direct danger, but the vulnerable structure is already taking shape and the guarantee mechanism is up and running. This is why in earlier analyses we called the NeoCloud providers the most unstable dominoes of the AI boom, and the latest developments reinforce that view.

US AI labs may be in trouble

The statement published on 24 July marked the culmination of the debate on open- and closed-weight models that had intensified with the arrival of the Chinese AI model Kimi K3. With a few exceptions, the leading players of the American tech and AI sector came out in support of open models.

With closed-weight models the developer does not release the trained parameters, and typically discloses neither the architecture, nor the parameter count, nor the training data. The model is accessible only through an API or an application and cannot be run on your own infrastructure. This gives the developer a strong bargaining position: it can withdraw a model at any time, update it quietly, reprice it or tighten the terms. This is how every leading model of the Western frontier labs works, and companies often pay millions of dollars for that exclusivity.

With open-weight models the parameters can be downloaded, and usually the architecture is public as well. This reverses the bargaining position, since what has once been downloaded cannot be taken away, the version can be frozen, and the provider cannot unilaterally reprice or withdraw it. This is how Llama, Mistral, DeepSeek and Qwen work. The best open-weight models set the price ceiling, because a premium can be charged for a closed model only for as long as the capability gap justifies it.

Open Weights and American AI Leadership memorandum signatories

Source: Memorandum text

On 20 July the American news site Axios reported that the administration had revived plans to ban leading Chinese models on cybersecurity grounds. Since then Moonshot AI, the Beijing-based developer of Kimi K3, has been accused of large-scale, covert industrial distillation, by copying the best models of the leading American lab Anthropic using smuggled Nvidia accelerators. This would not be unprecedented, since earlier this year Anthropic reported that DeepSeek, Moonshot and Shanghai-based MiniMax had collected more than 16 million Claude responses in distillation campaigns. According to Moonshot's own measurements, K3 sits confidently in the leading group and on some benchmarks is the most capable model in the world, ahead of Claude Mythos and GPT 5.6 Sol class models. This has largely been confirmed by independent testers such as Artificial Analysis. Demand for the model soon exceeded Moonshot's capacity to such an extent that the lab suspended registration for new subscribers.

The question is this: if a Chinese AI lab with orders of magnitude less capital and a constrained chip supply can deliver a model with frontier-level capabilities and then give away its parameter weights for free, what justifies the trillion-dollar valuations of the leading, closed Western labs?

K3's API pricing is roughly 70% lower than Anthropic's best. The model itself is huge, at around 2.8 trillion parameters, but thanks to the MoE architecture, which works with many smaller expert sub-models, it uses on average only 50 billion active parameters while making predictions based on learned patterns and logic. Its per-token compute requirement matches that of a 50-billion-parameter dense model, so at pricing of $15 per million output tokens the gross margin on the API is presumably healthy.

Ranking of the smartest AI models (28.07.2026.)

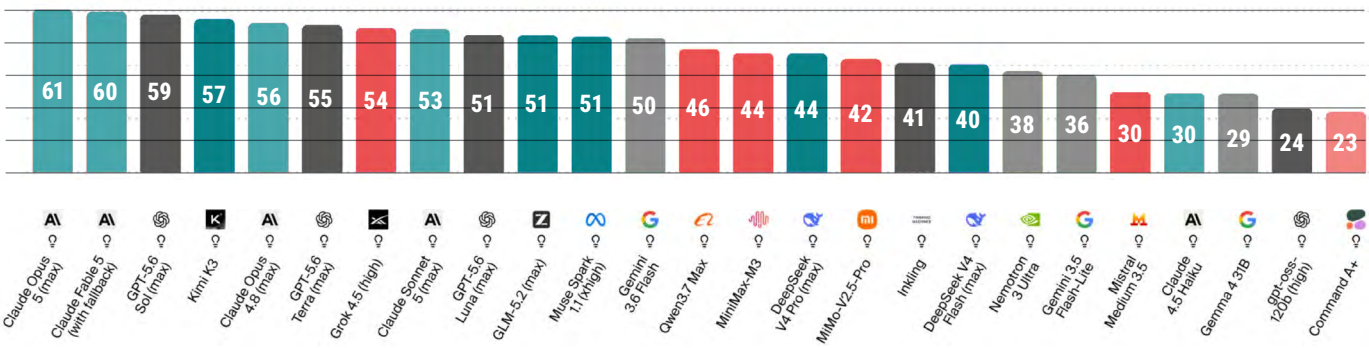

Source: Artificial Analysis • https://artificialanalysis.ai/

The signatories of the memorandum supporting open AI models include the cream of the American AI world: Nvidia, Meta, Microsoft, IBM, Dell, Palantir, several venture capital investors, Elon Musk and David Sacks (the White House AI adviser). With a slight delay, OpenAI and Google also joined the illustrious list. They all argue that open-weight models play an important role both in terms of cybersecurity and in keeping the AI race alive, and that by banning Chinese models the two leading American AI labs would in fact be stifling their competitors.

Those who regard Chinese open models as a threat are led by Anthropic, but this position is also supported by Amazon, a major player in cloud-based AI, by several leading figures at OpenAI such as president and co-founder Greg Brockman, and by a large share of AI safety researchers. In their view, releasing frontier open models freely carries serious dangers, because in the wrong hands they can be used for cyberattacks and fraud. When we use a model from a Western frontier lab, numerous filters typically restrict harmful behaviour. Anyone with access to the full weights and code of the model, however, can remove those restrictions, and the model will no longer refuse malicious requests.

In addition, the accusation still stands against Moonshot that its models were created at least in part "with the help of" Western frontier models. The opposing camp regards this practice as theft of intellectual property. Model distillation is far cheaper than conventional training. Its essence is that a smaller, and therefore cheaper, model is trained on the responses of a model with a larger parameter count, so that it achieves similar capability. This can have a severe deterrent effect on the frontier labs, since they spend tens of billions of dollars on training, which other players then copy, offering incrementally weaker but an order of magnitude cheaper alternatives to their models. A parallel can be drawn between pharmaceutical drug development and the development of frontier AI models. To encourage the development of new active substances, the established practice is multi-year patent protection and the manufacturing exclusivity that follows from it for the developer. If a compound developed at a cost of several billion dollars had to compete with generic alternatives within weeks or months, nobody would consider it worthwhile, to embark on the development. It is no different with AI models.

The gap between open-weight models and the leaders has narrowed to a few points on the main benchmarks, and in time to just a few months. If we add that the pace of model releases at the closed labs is of the same order of magnitude, the difference is practically imperceptible.

Progress in open weights vs. proprietary intelligence

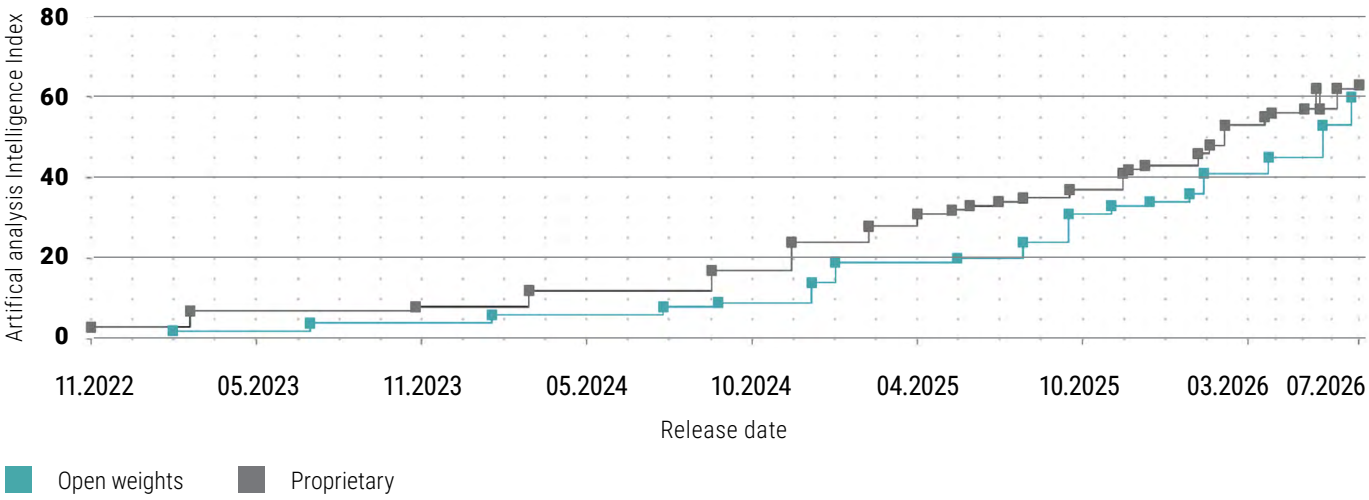

Comparison of leading open weight vs. proprietary AI models (19.08.2026)

Source: Artificial Analysis • https://artificialanalysis.ai/

A red sun in the sky

On 27 July 2026 the most valuable segments of the Western tech sector came under attack on three fronts.

In the morning in Shanghai, shares of the Chinese semiconductor manufacturer ChangXin Memory (CXMT) began trading at 49.50 yuan after an issue price of 8.66 yuan, and by the end of the day stood roughly 470% higher. Its capitalisation of 3.3 trillion yuan corresponds to nearly $490 billion. In the course of a single morning the memory manufacturer overtook ICBC and became the most valuable company on the Chinese exchanges. This market had until now been in the hands of America's Micron, South Korea's SK Hynix and Samsung, all three of which reached a $1 trillion valuation on the back of the AI trend. A Chinese entrant carries numerous risks for them. CXMT's pricing is highly competitive, and it cannot be expected to show the kind of supply discipline the other chipmakers exercise in order to avoid oversupply. Dumping by a Chinese memory manufacturer could on its own break the rally of the three chip giants. Realistically this scenario is several years away, but the long-term market position of today's leading memory manufacturers looks fragile.

The same day brought news that a state-backed company in Shanghai had begun small-series production of Chinese-developed immersion DUV lithography machines (which can draw extremely fine circuit patterns onto silicon wafers). Shares in the current market leader, the Dutch company ASML (which derives 30% of its revenue from Chinese customers) fell 6% on the news. The Dutch lithography equipment maker is one of the strongest monopolies in technology markets and Europe's most valuable company. The long-term erosion of its market position, however, is becoming increasingly likely. Chinese production volumes would allow the delivery of 5 DUV machines in 2026 and a further 20 in 2027, which will most likely land first at Chinese chipmakers and tech giants. For now this is dwarfed by the more than 100 DUV machines ASML delivers each year, and there is no evidence that this equipment is capable of the same performance (yield, throughput) as ASML's machines. In EUV lithography ASML may retain its monopoly position for the next few years. The tightness of chip supply gives cyclical tailwinds to the "wafer fabrication equipment" (WFE) sector that builds integrated circuits, so weak demand for ASML's machines is unlikely in the short term. Even so, the threat is real, and it is another data point suggesting that the Western technological lead is not invulnerable.

Thirteen months from private mark to peer

CXMT's Shanghai debut vaulted it into the same valuation bracket as the incumbent DRAM makers

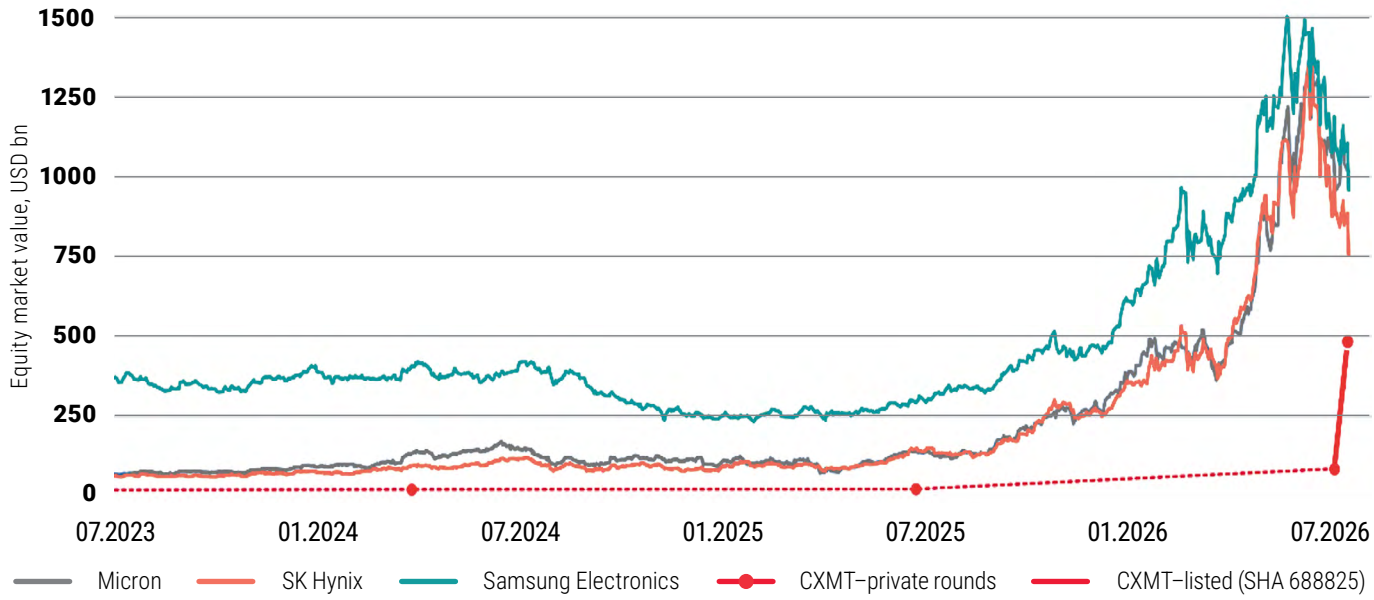

Source: yfinance, company fillings and press reports, VIG Asset Management

Also on that day, Beijing-based Moonshot AI uploaded the full weights of Kimi K3 to Hugging Face. With 2.8 trillion parameters it is the largest open-weight model ever released. The developer lab has faced accusations of distillation and chip smuggling, but K3's near-frontier capabilities are undeniable. This undermines the pricing power of the closed American AI labs and calls into question the durability of their competitive advantage. It is a kind of indicator of AI dumping: Chinese open labs, working from orders of magnitude fewer resources, offer somewhat weaker but far cheaper models, thereby squeezing out the Western labs.

All three stories reinforce the view that the Chinese tech sector's lag behind the West is minimal, and that with its often state-supported strategy, which is not primarily profit-driven, the sector is capable of disrupting the Western tech giants.

OpenRouter traffic is up 145-fold, and Chinese models now carry most of it

Anthropic, OpenAI and Google took 74% of routed tokens in January 2025 and 27% over the past month; Chinese labs went from 6% to 65%

Daily tokens routed, platform total

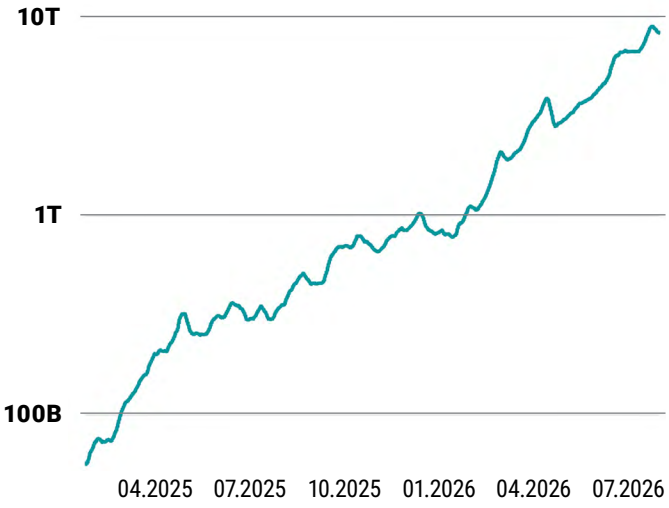

Share of tokens by model author

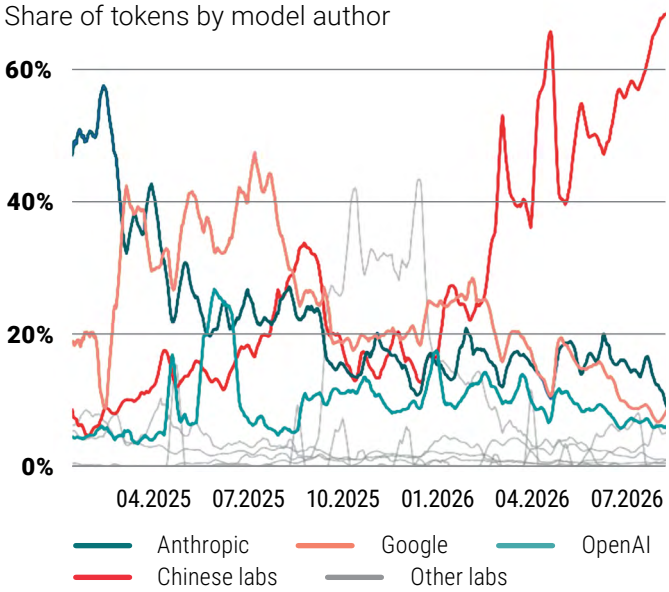

Source: OpenRouter (openrouter.ai/rankings), as of 28.07.2026. T11:45:35.598Z. 7-day moving average; author shares are of top-50 model trafffic. Chinese labs shown as one aggregate of 9 authors

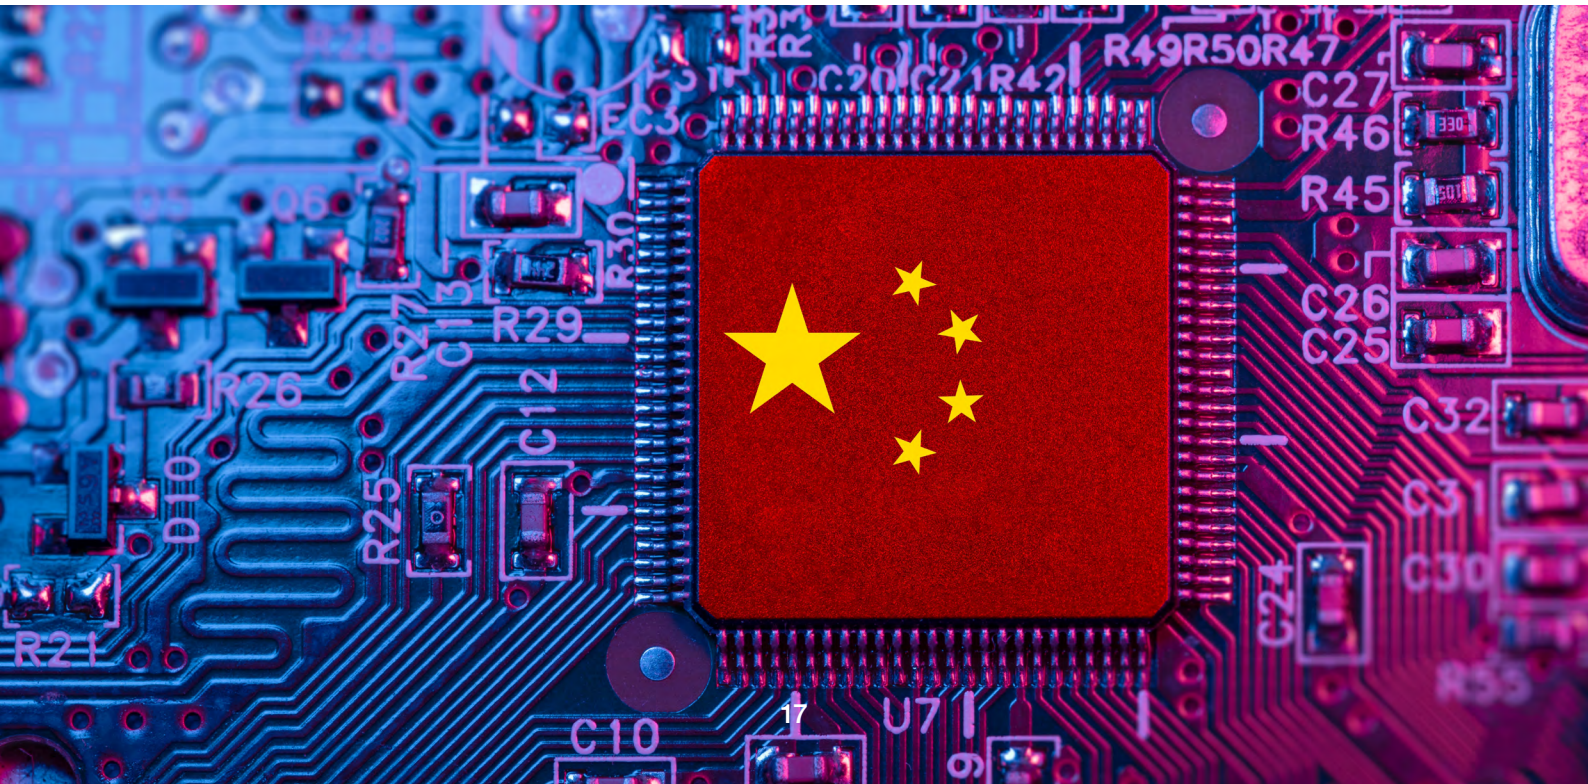

Megatrend Equity Fund newsletter – What we wrote to our investors in July

July was perhaps the greatest test yet for AI investments. The data centre infrastructure and chip sectors, which had performed outstandingly since the start of the year, went through a substantial correction: semiconductors fell 25% from their June peak, and the broader technology sector also delivered deeply negative returns. Value stocks outperformed growth stocks by a record margin, just as equal-weighted indices outperformed their market-capitalisation-weighted equivalents.

The Korean index, which holds an extremely high concentration of memory chip manufacturers, fell almost 40% from its June peak, and more than 1.2 million investor accounts hit the margin call threshold. After the historic bull market, Korean retail investors had hastily taken on large risks through leveraged ETFs created on single stocks, so in the days following the crash brokerages force-liquidated several hundred thousand accounts. On the American market, meanwhile, segments previously considered AI losers, such as the software sector, significantly outperformed during the sharp correction.

The Chinese technology sector struck on three fronts against Western companies that are finding it ever harder to defend their lead. On 27 July Reuters reported that a Chinese state-owned company had begun manufacturing immersion DUV lithography equipment. Previously only the Dutch company ASML was capable of this, and it had counted as the strongest monopoly in chip manufacturing. On the same day the Chinese DRAM manufacturer CXMT was listed on the Shanghai stock exchange and briefly became China's most valuable company after the listing. This is a good indication of how significant the appearance of a new entrant is in the memory chip market, which until now has functioned as a narrow oligopoly. Finally, also in July, Moonshot AI, one of China's leading AI labs, published the full parameter weights of its new frontier model. This undermines the market position of the Western AI labs and suggests that the gap between open-weight models and closed developments is by now negligible.

The negative sentiment and the external competitive threat were partly offset by a series of excellent quarterly reports. Of the companies reporting by the end of July, 86% beat profit expectations, and their results exceeded consensus by an average of 31.4%, compared with the usual figure of around 7%. The Magnificent 7 reports, however, did not trigger a uniform reaction: Microsoft and Amazon shares were lifted by strong cloud segment growth, while Meta and Alphabet suffered from raised investment forecasts.

In the oil market, Brent, the European benchmark, rose 24% and American WTI 21% during the month, with Brent closing at $90.12 on 31 July. The status of the war with Iran is changing ever more rapidly. Brent had fallen back to around $73 by the end of June, then on 8 July, after Trump declared the ceasefire over at the NATO summit, it jumped 5.2% to $78.02, and by the end of the month it had risen above $90. Official American communication on the war often changes several times within a week, and sometimes within a single day, but there is still no sign of a meaningful solution for reopening the Strait of Hormuz.

The financial sector was one of the best performing segments of the month. The rise in long-term yields supported interest margins, and the large banks reporting in mid-July described strong capital markets activity. The quarterly reports showed that advisory and issuance fee income at the leading investment banks grew rapidly. The return of stock market listings feeds directly into the sector's results.

In the Fund, at the beginning of the month we reduced the weight of the semiconductor sector, largely ahead of its decline, and increased the global allocation to the defensive healthcare and financial sectors. The momentum of renewable energy and solar stocks broke in June, so we took profits in that area. During the month we also increased the weight of Nvidia, which has underperformed its sector this year but may remain a key player in the long-term AI trend.

Disclaimer

This is a distribution announcement. Detailed information is needed to make a well-founded investment decision. Please inform yourself thoroughly regarding the Fund's investment policy, potential investment risks and distribution in the Fund's key investment information, official prospectus and management regulations available at the Fund's distribution outlets and on the website of VIG Asset Management Hungary Closed Company Limited by Shares (hereinafter: Asset Management) (www.vigam.hu). The costs related to the distribution of the fund (buying, holding, selling) can be found in the fund's management regulations and at the distribution outlets. Past returns do not predict future performance. Please note that in comparison with other investment funds, the return achieved may be affected by differences in the reference index and therefore the investment policy.
The future performance that can be achieved by investing may be subject to tax, and the tax and duty information relating to specific financial instruments and transactions can only be accurately assessed on the basis of the individual circumstances of each investor and may change in the future. It is the responsibility of the investor to inform himself about the tax liability and to make the decision within the limits of the law.
The information contained in this leaflet is for informational purposes only and does not constitute an investment recommendation, an offer or investment advice. Asset Management accepts no liability for any investment decision made on the basis of this information and its consequences.
This leaflet may contain trade names or brand names owned by third parties, the use of which serves illustration purposes only. All product and company names are trademarks™ or registered trademarks® of their respective owners. The use of such trademarks does not imply any business or other relationship between Asset Management and the owners of the trademarks, nor shall it be interpreted as a recommendation by Asset Management.
The Asset Management's license number for managing alternative investment funds (AIFM) is: H-EN-III-6/2015.
The Asset Management's license number for UCITS fund management (collective portfolio management) is: H-EN-III-101/2016.

VIG Asset Management Hungary

H-1091 Budapest, Üllői út 1.

vigam.hu/en

linkedin.com/company/vig-asset-management-hungary

VIG ASSET MANAGEMENT