VIG
ASSET MANAGEMENT

MEGATREND MONITOR

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

May 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

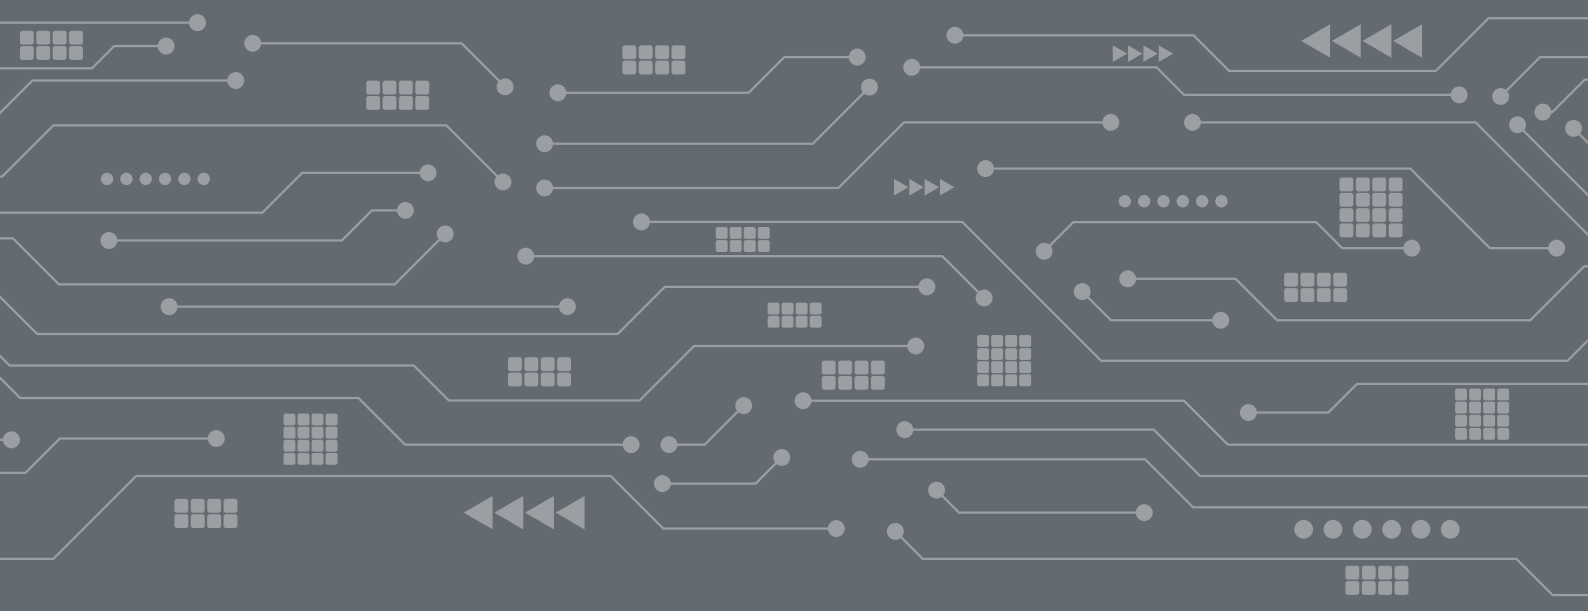

Performance of megatrends

Let's start by looking at how the three main strategies – innovation-focused, demographic-focused, and environment-focused megatrends – have performed over the past year within the megatrend universe.

Although headlines and market narratives might lead us to believe that IT themes and technology-focused innovations have dominated the recent period, a closer look at the numbers reveals surprising correlations. In fact, over the past year, it was not technology-focused assets but those related to environmental megatrends that delivered the most outstanding performance. This apparent contradiction is the result of an extremely complex global process: the geopolitical uncertainty caused by the conflict in Iran has heightened the importance of energy security, while the rise of artificial intelligence has generated unprecedented energy demand for data centers. The modernization of energy networks, the compelling demand for renewable resources, and the development of sustainable infrastructure have thus become the true, fundamental drivers of the innovation boom, proving that the green transition is no longer merely a matter of principle but can be a profitable market reality.

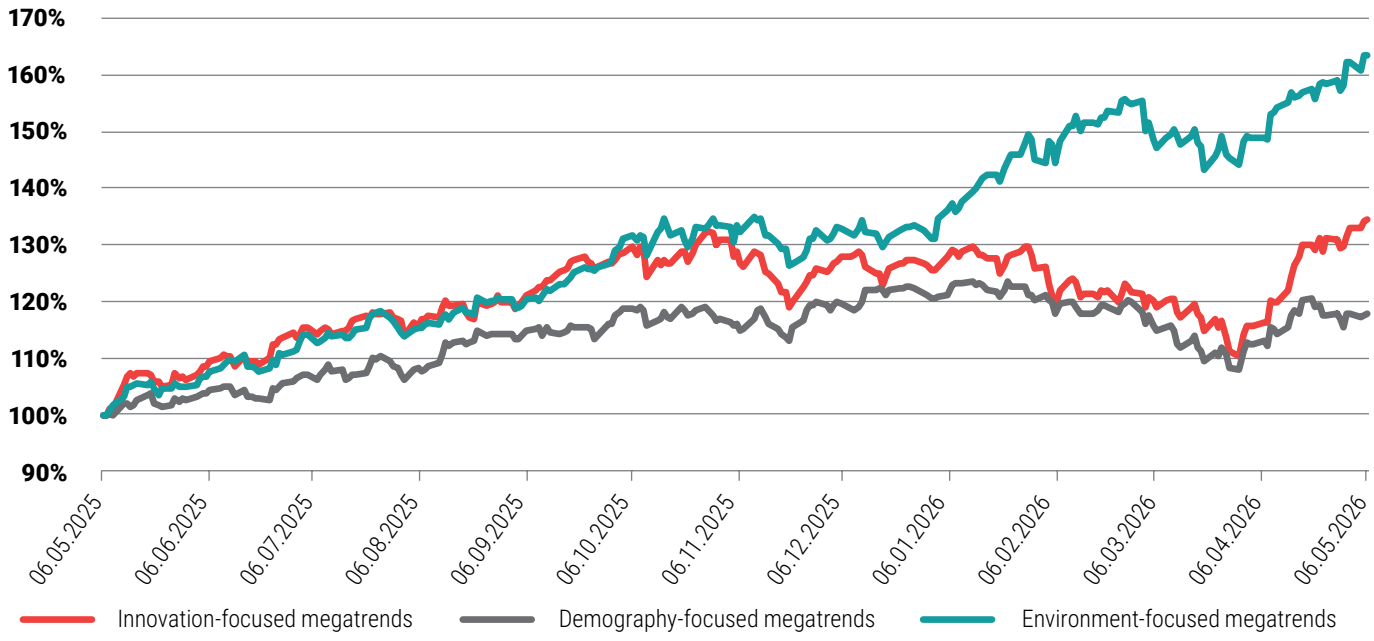

Source: based on data of Bloomberg terminal

Under the hood of the AI revolution: a new golden age for hardware manufacturing

The stock market surge so far in 2026 has been driven mainly by the technology giants (Alphabet, Amazon, Microsoft, Meta), whose total investment budgets now exceed $750 billion—$80 billion more than expected at the start of earnings season and 83% higher than spending in 2025. The surge in AI capital expenditures (capex) shows no signs of slowing down.

The surge in investment expectations is driving a similar rise in profit expectations for AI infrastructure companies, which is improving the broader market's earnings outlook and leading to upward revisions in their estimates for the S&P 500's earnings per share (EPS).

While the stock market celebrates, a serious contradiction looms in the background: the pace of software demand and capital investment has far outstripped the physical capacity of hardware manufacturers. This gap has forced companies like OpenAI and Anthropic to limit their services due to "computational hunger."

Trends in earnings forecasts for chip sector companies by year-end

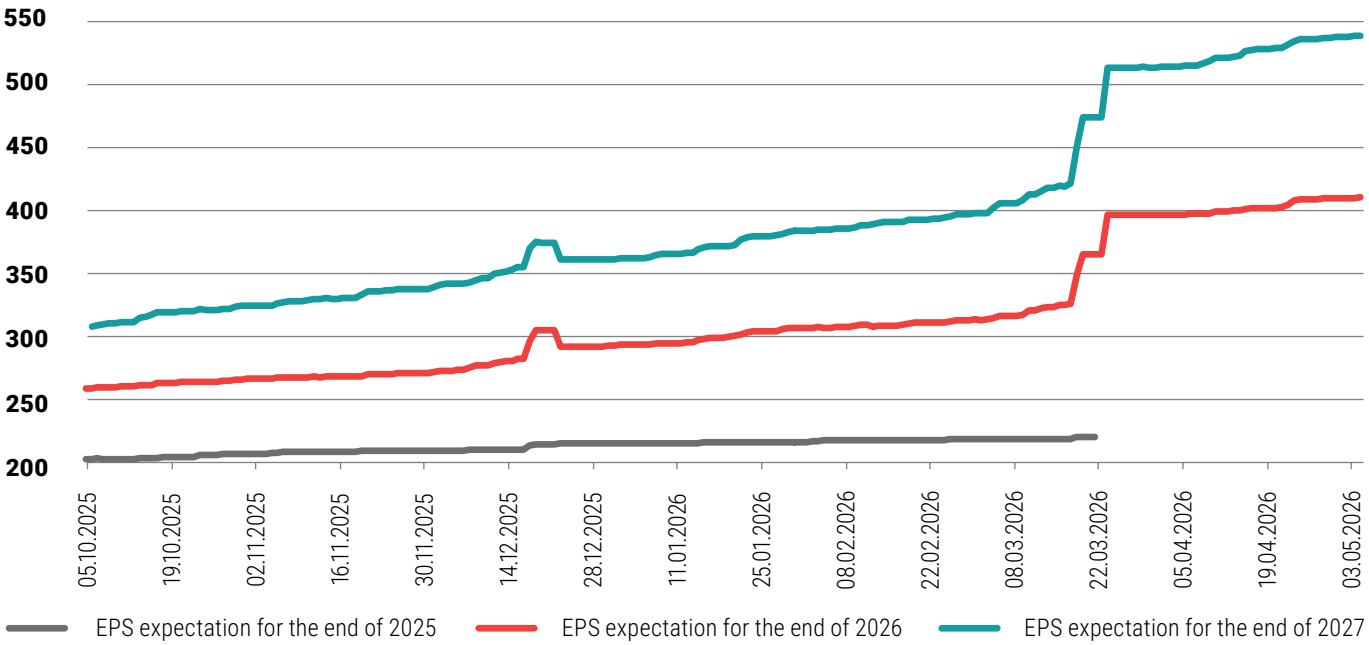

Source: Bloomberg

The hierarchy of critical bottlenecks

The industry is no longer struggling with a shortage of a single chip, but with a complex, interdependent system of constraints:

- HBM (High Bandwidth Memory) – the main obstacle: modern AI models require faster data access than ever before. Memory manufacturers (SK Hynix, Micron, Samsung) have sold out all their capacity through the end of 2026. Although software algorithms (e.g., Google TurboQuant) are attempting to reduce memory requirements, the physical shortage is expected to persist for another three years.
- GPUs and manufacturing constraints: although Nvidia dominates, the bottleneck here is no longer just design, but TSMC's "CoWoS" packaging technology. The most advanced plants are running at full capacity, and building a new factory takes 2–3 years, so supply cannot immediately keep pace with market surges.
- CPU – the unexpected comeback: while CPUs were previously thought to have been pushed into the

background in the AI era, the new type of "agent-based" AI systems (which plan, reason, and execute tasks) require much more processor coordination. While a chatbot required one CPU for every 12 GPUs, this ratio has shifted to 1:1 for agents, giving Intel a massive boost.

- Physical and political barriers: beyond chips, the construction of data centers is also hitting roadblocks. Political resistance is growing in the U.S. and Europe due to the massive power consumption, which is delaying projects. Power supply and cooling technologies (liquid cooling) have become the new critical elements of infrastructure.

Currently, hardware manufacturers are calling the shots. Since they are unwilling to risk "overbuilding," the shortage may persist. This creates a peculiar situation: the world's richest companies (Microsoft, Google) are lining up for chips, while limited supply has caused rental fees for older models (e.g., Nvidia H100) to rise by 30%. While cloud service providers have tripled their spending, hardware manufacturers (such as TSMC or memory manufacturers) are more cautious: they have increased their investments by only 50%, fearing that overcapacity will lead to losses later on. This contrast between "hardware caution" and "software euphoria" could mean that the pace of AI development in 2026 will be determined not by the genius of engineers, but by the speed at which factories are built and the number of silicon wafers produced.

Capital offence AI supply chain, capital expenditure

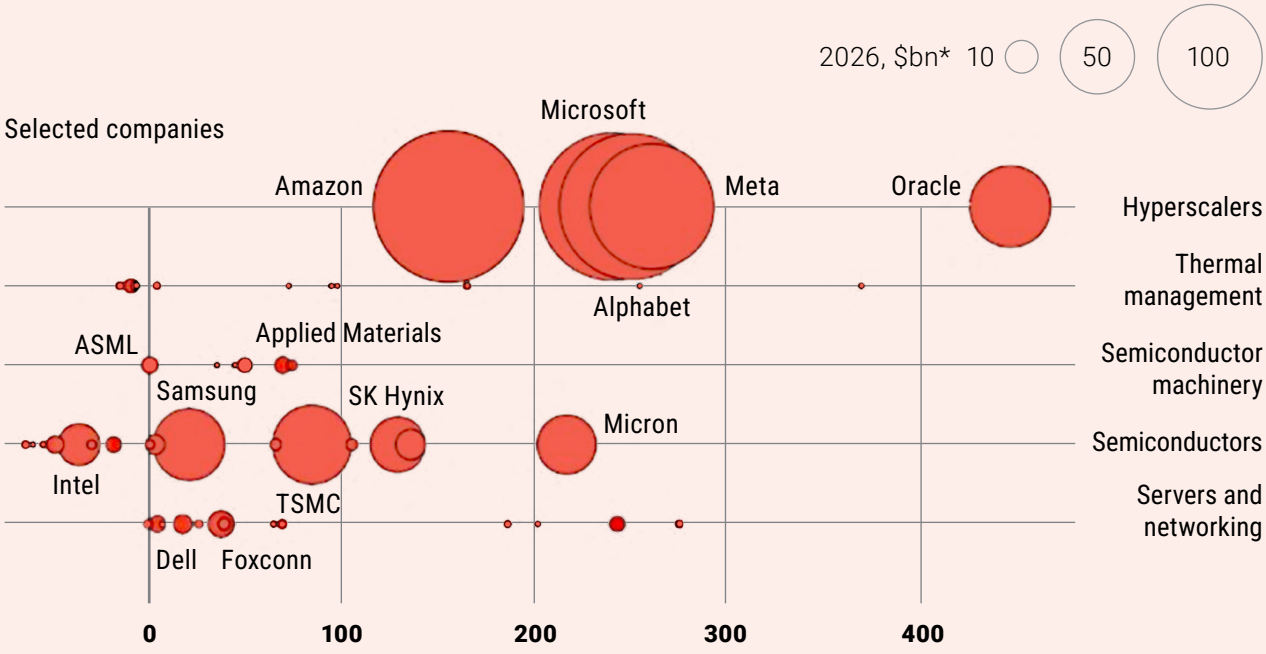

Sources: S&P Global, The Economist

Software: dark clouds on the horizon

In April, after more than six months of decline, the software sector regained its footing. The IGV ETF, which comprises leading U.S. software companies, rose nearly 10%. However, the relief is relative; with semiconductors posting a historic monthly performance of nearly 40%, the gap between the two pillars of the AI story has never been wider.

For software providers, any AI-related news is bad news these days. A new model, a new feature, or an X post from one of the leading AI labs – all serve as negative catalysts that immediately cast doubt on the sector's future revenues.
A pessimistic investor's vision might look something like this:

Source: based on free data service of Yahoo Finance
https://finance.yahoo.com/quote/IGV/ • https://finance.yahoo.com/quote/%5ESOX/

1. In the short term

AI coding assistants are upending the labor equation in software development. Where large, coordinated teams were once required, a single developer and a suite of AI assistants may now suffice. As a result, barriers to entry will collapse, and the resulting competition will erode software providers' pricing power.

2. In the medium term

the seat-based pricing model will become obsolete. Software functions will shrink to chat interfaces, and workflows will be executed by agents – even within existing software frameworks – that operate at speeds orders of magnitude faster than human pace. The role of humans will also change. They will no longer oper-

ate the software but will supervise agents that navi gate the software more skillfully than they do. This will negatively impact any service provider that has built its competitive advantage on the quality of its user interface and will necessitate a forced transition from seat-based to usage-based pricing.

3. In the long term
the concept of software will transform. It is not humans who use applications; armies of AI agents will autonomously develop and operate systems on a scale that would have been unimaginable in the era of artisanal software development. The cost of producing software will ultimately boil down to the price of tokens[1], and the primary user will no longer be a human but another agent. Probably not a single software provider today is prepared for this world, and it is difficult to say who will be able to remain relevant in such a future.

However, we are still very much at the beginning of this timeline. Most software companies view the impacts of AI as an existential issue and are among the first to adopt the latest models and features. The minority of software-optimistic investors base their view on the fact that, despite steep price declines, fundamentals remain intact, and the market tends to overreact when projecting the effects of new technologies. The future of the software sector is difficult to predict, but it is clear that the AI trend may remain the most intense battleground in the coming years.

[1]Token: The basic unit of text processing in large language models; typically represents words or word fragments

The age of physical AI

In April, records were broken in two long-distance running events, but only one of these can be attributed to the performance of a human being. On the 26th, at the London Marathon, two runners broke the two-hour barrier, which had previously been considered impossible under competitive conditions. Kenyan runner Sebastian Sawe set a new world record with a time of 1:59:30.

A week earlier, roughly 12,000 people and 300 humanoid robots lined up at the start of the Beijing E-Town Half Marathon. The runner who won the human field with a time of 1:07:47 finished more than 17 minutes behind the fastest humanoid, named Lightning. The level of the robot field has improved dramatically compared to last year's debut race, which the fastest robot completed in 2:40:42 – a time that was still well behind the human field at the time.

This story may be a sign of a larger trend.

More than 750,000 robots are now working alongside human workers in Amazon's warehouses. Most of these are not the humanoid robots familiar from sci-fi movies, but rather short, disc-shaped units that move shelves and sort packages. The presence of robots is also becoming an increasingly common sight in airport terminals. At an airport in Seoul, autonomous baggage handlers escort passengers to the gate; in Tokyo, cleaning robots patrol the terminals at night; and at Munich Airport, information robots answer passengers' questions in four languages. The message here is that robots are increasingly becoming part of our everyday lives, and the parallel development of artificial intelligence and robotics – which operate seamlessly in the real world – will only reinforce this trend.

Source: LA Times

Men's half Marathon world record From human peaks to the first humanoid robot victory (1998-2026)

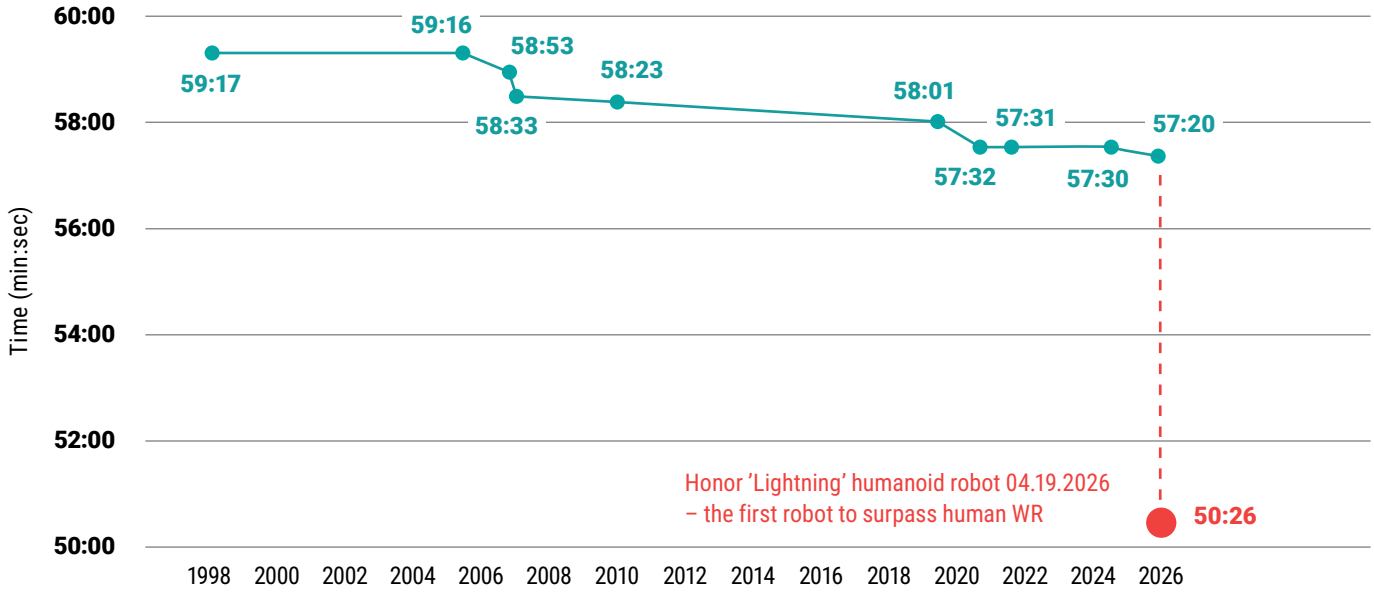

Sources: record progression certified by World Athletics (Wikipedia: Half_marathon_world_record_progression), Honor's "Lightning" robot – autonomous humanoid winner of the Beijing E-Town half marathon on April 19, 2026 (IEEE, TechCrunch, Xinhua).

https://www.telegraph.co.uk/world-news/2026/04/19/robots-beat-humans-in-half-marathon-for-first-time/
https://techcrunch.com/2026/04/19/robots-beat-human-records-at-beijing-half-marathon/
https://en.wikipedia.org/wiki/Half_marathon_world_record_progression/

A shake-up in the world of top model builders

Ever since artificial intelligence burst into the public consciousness, ChatGPT – much like Google did with the internet – has become synonymous with AI. Two years ago, OpenAI, the company behind it, was undeniably leading the AI race by a comfortable margin. In recent months, however, a new AI king has been crowned: rival AI firm Anthropic's annualized revenue exceeded $40 billion, surpassing OpenAI's latest figures of $25 billion. In secondary market trading, Anthropic's valuation has reached $1 trillion, while OpenAI hovers around $880 billion. Anthropic is also reportedly working on a $50 billion funding round, which, with a $900 billion valuation, would officially put the company ahead of its rival. It's worth examining how OpenAI might lose its dominant position and what this reveals about the future of AI.

The answer lies not in a single brilliant product or technological breakthrough, but in the fact that the two companies are following radically different paths. With ChatGPT, OpenAI has created the fastest-growing consumer application in history: by March 2026, it had reached 900 million weekly active users, with more than 50 million paying subscribers. These are impressive numbers, but consumer subscriptions are in a lower price bracket and are more prone to churn. Currently, about 60% of OpenAI's revenue comes from consumers (ChatGPT Plus, Pro, and Go subscriptions), and only 40% from enterprise customers. Anthropic, by contrast, has essentially skipped the consumer phase. The Claude chatbot exists, but it has never been positioned as a competitor to ChatGPT; rather, it is marketed as a developer tool and an enterprise solution. According to Dario Amodei, CEO of Anthropic, roughly 80% of revenue comes from enterprise clients, with over 1,000 companies spending at least $1 million annually on Claude's services. That number has doubled in less than two months. This is a significant departure from OpenAI's model, as enterprise contracts are higher-value, harder to cancel, and represent a much more predictable source of revenue.

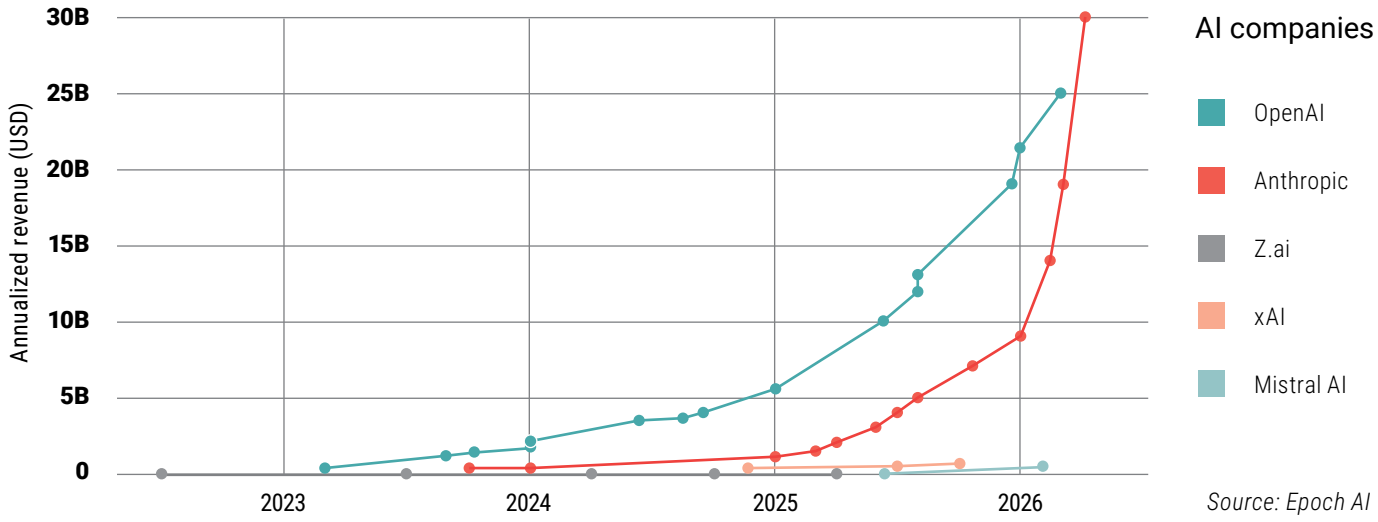

Source: Epoch AI

The economic reality of AI services is that revenue scales linearly in proportion to the number of tokens used. A company that initially uses Anthropic Claude only for code review will eventually expand its use to internal documentation, customer service, and legal text analysis, and then, with the spread of agentic models, to its entire workflow, exponentially increasing the number of tokens used over time. The classic "land and expand" dynamic (entering with a narrow use case and then naturally expanding within the organization) is one of the strongest in the AI services market and can provide a lasting competitive advantage to the player that penetrates enterprises more quickly.

However, the competition between consumer dominance and enterprise focus has not yet been decided. OpenAI's 900 million users represent the world's largest AI distribution channel. In contrast, Anthropic's 80% enterprise share and the Claude Code-driven growth spiral represent a higher-value, stickier business model. The market currently prices in the view that Anthropic's approach is more effective, but the AI market may be large enough for both models to succeed. The coming months will be decisive, as both companies are preparing for an IPO, so the IPO filings will ultimately reveal what lies behind the numbers.

The World's Most Dangerous AI

This could be the dramatic opening of a sci-fi story, for example:
Treasury Secretary Scott Bessent and Federal Reserve Chair Jerome Powell summoned Wall Street executives to an emergency meeting, fearing that the world's newest and most advanced AI model could trigger a cybersecurity collapse in the world's most critical financial IT systems.

Or perhaps:
The new model, developed by Anthropic, one of the world's leading AI labs, demonstrated capabilities during internal testing that led the company to withhold it from the public for security reasons. The decision was justified by the fact that, in the wrong hands, the new technology poses a threat to the world's digital infrastructure. Following its predecessors Sonnet (short song) and Opus (masterpiece), the model was named Mythos.

Interestingly, both news items appeared in April on official channels, including Bloomberg's platforms. At first glance, this paints a decidedly grim picture of the present, but it is worth taking a closer look at what other factors are actually shaping events.

1. According to industry reports, Mythos is the largest and, accordingly, most expensive Anthropic model to date. Based on official pricing, it is approximately five times more expensive than the previous top-tier version, Opus 4.6. Meanwhile, it is well known that most AI labs, including Anthropic, operate at a loss across nearly every tier of their subscription packages; even by conservative estimates, they spend several times the monthly fee per user.

2. In this context, making Mythos widely available would not only further increase losses, but, due to its significant computational demands, would effectively crowd out the running of other models and slow down further development as well. In a world of finite computational capacity, AI labs must constantly balance inference and training. Capital allocation between short-term inference revenue and long-term model development is one of the most critical strategic decisions, especially when running a model requires five times the capacity of the previous generation. Widespread, subscription-based availability is economically difficult to justify on current infrastructure.

3. Mythos Preview truly represents a generational leap in the fields of coding and cybersecurity, yet the picture is more nuanced in the broader benchmark landscape: the three leading frontier models – Mythos, GPT-5.5, and Gemini 3.1 Pro all dominate in different areas, and none clearly stands out above the other two in terms of general intelligence.

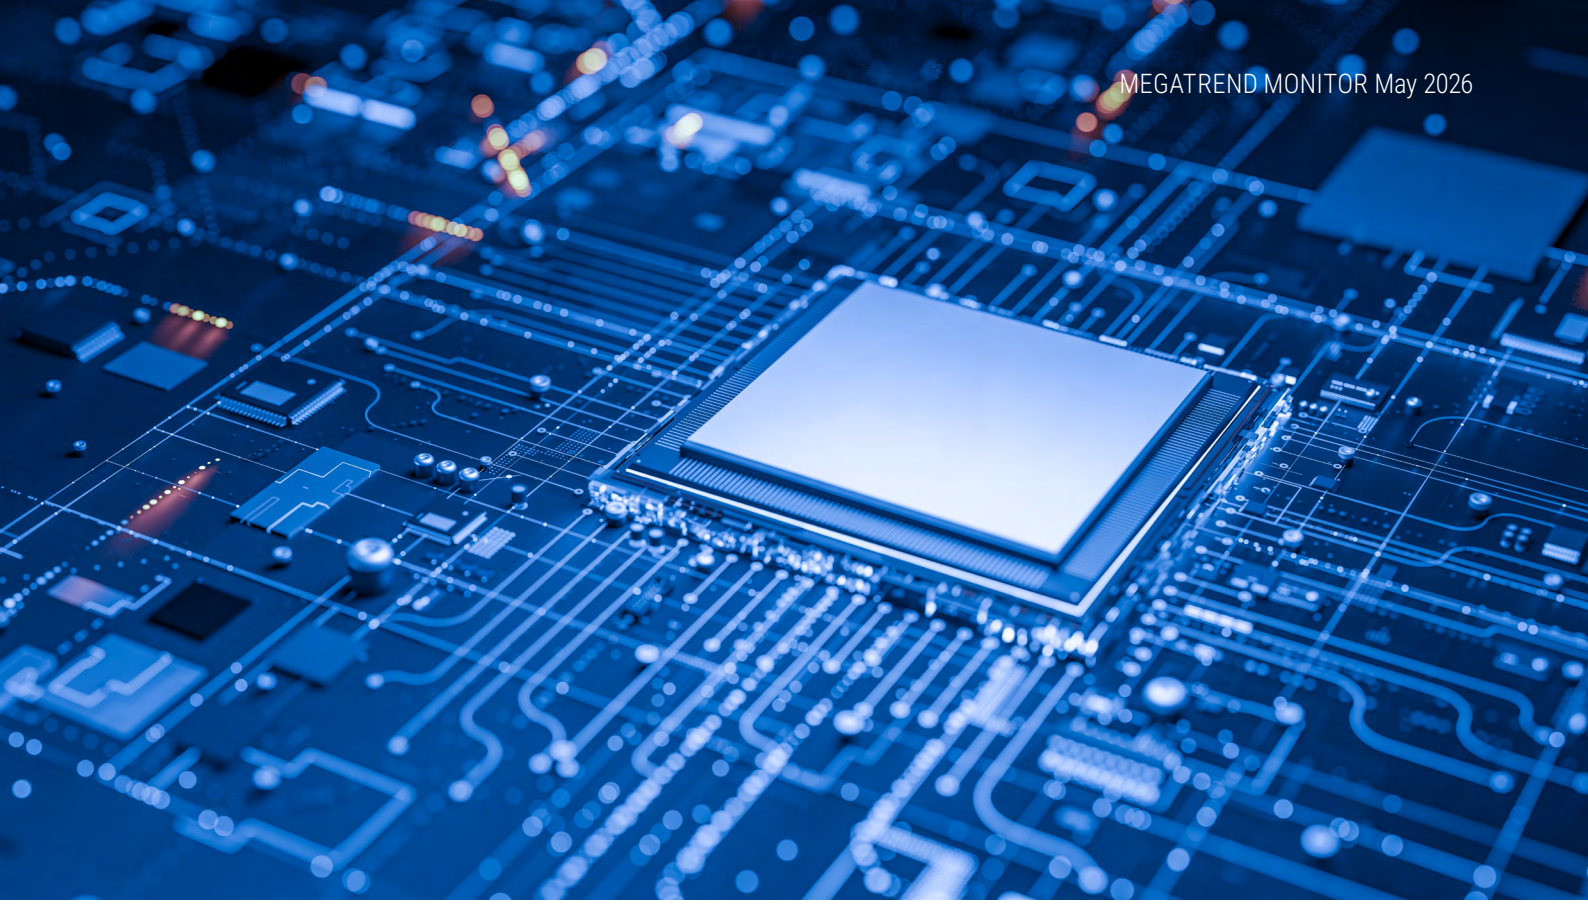

Overall, it can be said that AI model developers are highly capital-intensive companies that constantly require new and increasingly larger rounds of investment. For them, therefore, convincing the next generation of investors is of existential importance, and there are few tools more effective for this than grandiose claims about the capabilities of new models—claims that are difficult to verify.

This does not mean that the development of AI models is not real or significant, but the technological breakthrough and the narrative built around it are two different things, and a significant divergence between the two can distort the markets. When an AI lab withholds a model citing security concerns, it may be a genuine concern, but it could also be a deliberate positioning strategy ahead of the next funding round.

Ranking of the smartest AI models

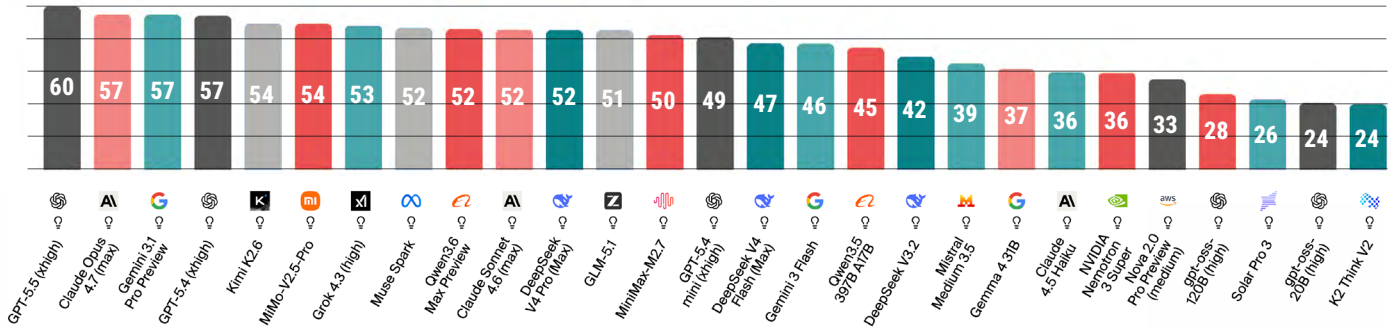

Claude Mythos Preview did not undergo standardized tests due to its closed nature.
Source: Artificial Analysis • https://artificialanalysis.ai/

Where might Apple be headed under its new leadership?

Apple's product strategy has been based on the same insight for decades: you don't have to be first; you just have to be the best. A common feature of the product lineup ranging from the 2001 iPod to the 2024 Vision Pro is that Apple rarely took on the role of first mover; instead, it created the benchmark product for each category based on the experiences of early entrants and market validation. This strategy now faces its greatest test yet in the light of the AI race. Under Tim Cook's leadership, the company has exercised a level of investment discipline unique among hyperscalers, thereby giving competitors an initial advantage that is difficult to overcome.
However, Cook will no longer bear the consequences of this decision, having announced on April 20 that, after 15 years as CEO, John Ternus will take over the CEO position starting in September. The timing does not seem coincidental. Apple's products now hit the market each year with barely noticeable improvements. Gone are the days when a newly released phone or laptop introduced significant innovations.
The choice of CEO profile is a strategic message in itself: while Cook was a man of supply chains and operational discipline, Ternus comes from the world of product innovation and hardware engineering. Cook leaves behind a nearly flawless, $4 trillion machine, and it will be Ternus's task to decide what to do with it.

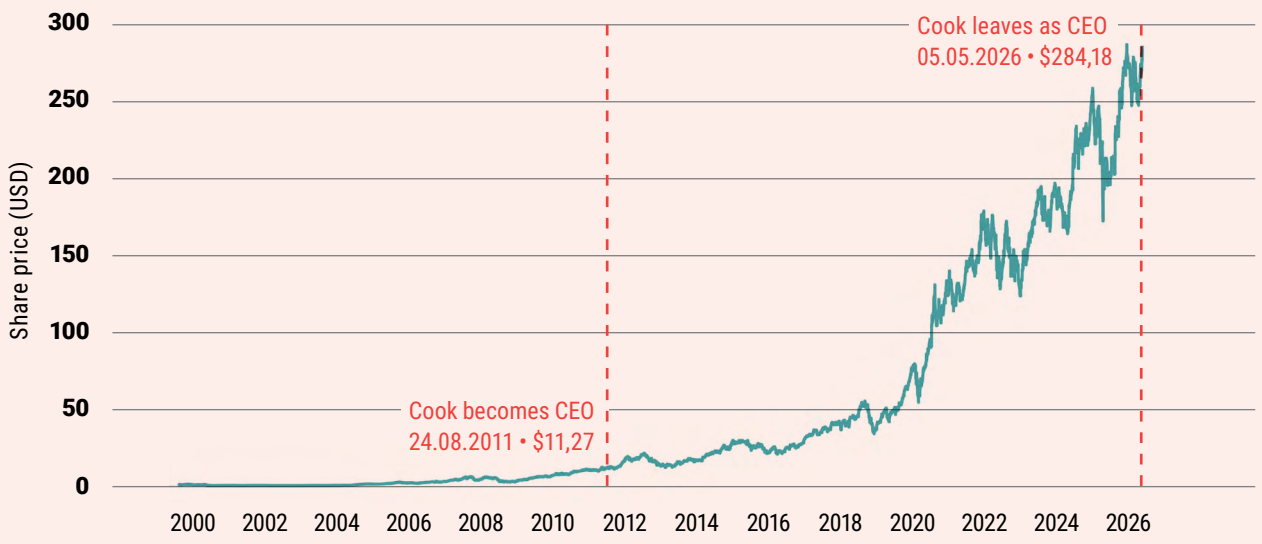

Source: based on data service of Yahoo Finance • https://finance.yahoo.com/quote/AAPL/

What do the quarterly reports from the world's largest companies reveal?

The earnings season that began in late April was one of the most important periods of the quarter. The world's five largest technology companies by market capitalization – Alphabet, Meta, Amazon, Microsoft, and Apple – all reported their quarterly results within a matter of days. Overall, the results exceeded expectations. This time, the focus was not only on revenue and earnings figures, but also on how each company is positioning itself in the AI investment race: who is spending how much, on what, and what return on investment they promise. Below, we examine what the numbers revealed for each company.

Alphabet stands out with its comprehensive AI model (chip-data center-cloud-application-data), which provides a significant cost advantage. Growth in cloud services exceeds that of competitors, while the capital expenditure ratio is more controlled compared to peers thanks to its own chips. The market reacted positively to the report.

Ticker	Period	Rev Est ($B)	Rev Actual ($B)	Rev Surprise %	EPS Est ($)	EPS Actual ($)	EPS Surprise %
GOOGLE	Q1'26	106.79	109.93	2.94%	2.63	2.62	-0.38%

Despite strong fundamentals, Meta faced a negative market reaction because management significantly increased its projected capital expenditures. The $135 billion annual capex is raising concerns due to uncertainties regarding its return on investment, especially since AI monetization appears less clear-cut for Meta than for its competitors.

In Amazon's report, AWS continued to prove itself as a robust growth engine, with operating income significantly exceeding expectations. However, high capital expenditures, weak free cash flow, and a somewhat disappointing operating outlook could cause uncertainty in the short term. Demand for its own chip series, called Trainium, could be a positive catalyst in the longer term.

Ticker	Period	Rev Est ($B)	Rev Actual ($B)	Rev Surprise %	EPS Est ($)	EPS Actual ($)	EPS Surprise %
META	Q1'26	55.45	56.31	1.55%	6.79	7.31	7.66%

Ticker	Period	Rev Est ($B)	Rev Actual ($B)	Rev Surprise %	EPS Est ($)	EPS Actual ($)	EPS Surprise %
AMAZON	Q1'26	177.30	181.50	2.37%	1.64	2.78	69.51%

In Microsoft's case, too, growth in its cloud platform, Azure, helped the company exceed expectations. The company was more cautious than its peers in forecasting the level of new investments. Despite the strong results and the controlled CAPEX plan, the market reaction was muted.

Apple's report focused less on AI infrastructure and more on the stability of its hardware and ecosystem. Growth was more moderate, and investors are looking for a compelling AI strategy. As a result, the company's valuation is based more on stable cash flow and shareholder returns than on accelerating growth.

Ticker	Period	Rev Est ($B)	Rev Actual ($B)	Rev Surprise %	EPS Est ($)	EPS Actual ($)	EPS Surprise %
MICROSOFT	FQ3'26	81.46	82.89	1.76%	4.03	4.27	5.96%

Ticker	Period	Rev Est ($B)	Rev Actual ($B)	Rev Surprise %	EPS Est ($)	EPS Actual ($)	EPS Surprise %
APPLE	FQ2'26	109.66	111.18	1.39%	1.95	2.01	3.08%

Quarterly report data according to Bloomberg terminal

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 Asset Management's website (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. VIG Asset Management Hungary Closed Company Limited by Shares accepts no liability for any investment decision made on the basis of this information and its consequences.

The Asset Management's license number for managing alternative investment funds (AIFM) is: H-EN-III-6/2015.

The Fund Manager'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