

Too big to f-AI-l



Source: AI generated

The AI story has been one of relentless momentum. ChatGPT marked the first chapter, transforming OpenAI into the central force behind the industry's expansion. The company now faces the challenge of supporting the scale of investment behind it.

Charles-Henry Monchau, CFA, CAIA, CMT
Chief Investment Officer
charles-henry.monchau@syzgroup.com

Assia Driss
Syz Research Lab Team Coordinator
assia.driss@syzgroup.com

Hugo Morel
Syz Research Lab Team
hugo.morel@syzgroup.com

Introduction

OpenAI has long stood at the epicentre of the artificial intelligence revolution. ChatGPT propelled the company into global prominence, driving one of the fastest revenue ramps in tech history. In early 2026, the company secured a massive \$122bn funding round at \$852bn valuation, placing it among the most valuable private companies globally.

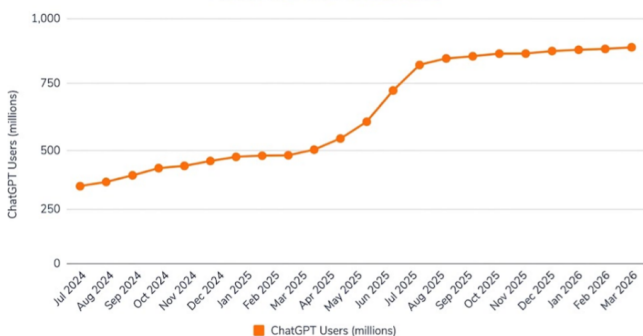
Yet a Wall Street Journal report published on 28 April introduced a more cautious tone. OpenAI missed several key internal targets. ChatGPT did not reach 1 billion weekly active users by the end of 2025, landing closer to 900 million by February 2026. Annual revenue goals were missed, along with multiple monthly targets in early 2026. The board has also begun scrutinising Sam Altman’s aggressive push for additional compute capacity. These misses come as OpenAI races toward a potential IPO later in 2026 or 2027, despite projecting heavy losses, potentially \$14bn in 2026 alone, and cumulative cash burn that could reach tens or even hundreds of billions before any hoped-for profitability in 2029 or beyond. Is this a manageable competitive hiccup in a still-growing market, or the first credible crack in the AI growth narrative?

What actually happened

OpenAI entered 2025 with extraordinary momentum. Revenue had climbed from roughly \$2bn in 2023 to \$6bn in 2024, while ChatGPT’s weekly active users surged from about 100 million to 400 million by early 2025. Internally, the assumption was simple: this pace will continue. It didn’t.

By late 2025, ChatGPT reached around 900 million weekly active users on an enormous scale, but still short of the one-billion target. More telling was the slowdown beneath the surface. Monthly user growth dropped sharply from 42% in early 2025 to just 13% by September. Around the same period, OpenAI faced a wave of online backlash tied to military-related partnerships, triggering a boycott among parts of its user base. While difficult to quantify precisely, the timing aligns with weakening growth and rising churn, suggesting that reputational pressure may have amplified an already natural deceleration.

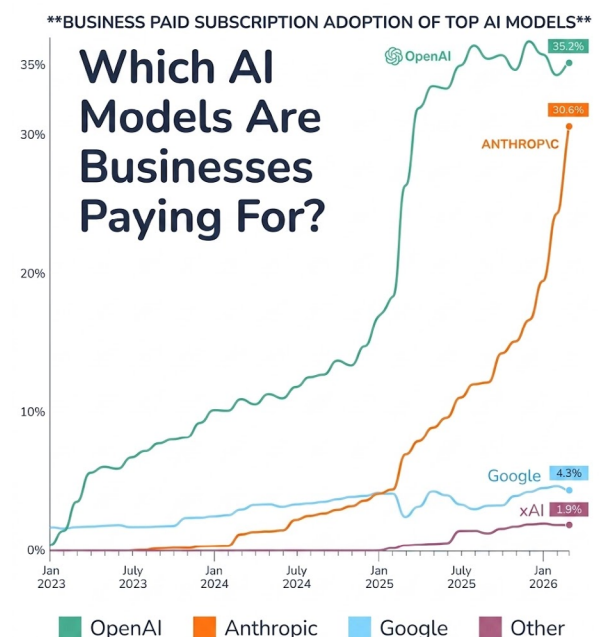
CHATGPT'S TWELVE-MONTH USER TREND
MONTHLY ACTIVE USERS



Source: Firstpagesage

From a monetisation standpoint, fragility is noticeable. Only about 5% of OpenAI’s 800-900 million weekly users are paying subscribers, leaving the company with massive scale but relatively thin revenue per user. Even with a \$20bn annualised revenue in 2025, performance fell short of internal expectations and by early 2026, OpenAI was missing multiple monthly revenue targets, a sign that demand was no longer keeping pace with projections.

On the other hand, Anthropic pulled ahead on efficiency. By March 2026, it reached roughly \$19bn in ARR, with Claude Code alone generating over \$2.5bn after 5.5× growth following Claude 4. The monetisation gap is striking. Anthropic earns about \$211 per monthly user, compared to OpenAI’s roughly \$25 per monthly user. Retention adds further pressure. AI apps retain just 21.1% of users annually, versus 30.7% for traditional software, making long-term revenue harder to sustain.



Source: Visual Capitalist

Financially, the model is under strain. Deutsche Bank projects \$143bn in negative cumulative free cash flow from 2024 to 2029, while OpenAI could post a \$74bn operating loss in 2028 alone. Still, OpenAI’s IPO ambitions remain alive. Polymarket assigns a 51.5% probability to a 2026 listing.

Market reaction

The reaction was immediate and telling. Within hours of the Wall Street Journal report raising concerns about OpenAI’s future profitability and its ability to fund ever-expanding compute needs, the broader “OpenAI ecosystem” traded sharply lower.

Oracle, deeply tied to OpenAI through the multi-year Star-gate data-centre initiative, reportedly a roughly \$300bn computing partnership, fell around 7%. CoreWeave, one of the most exposed pure-play AI infrastructure providers with contracts ranging from billions to tens of billions of

dollars, dropped 7%. In Tokyo, SoftBank, one of OpenAI's most visible financial backers, lost close to 10%.

The shockwaves extended across the semiconductor complex. Nvidia, Broadcom and AMD saw pressure despite their diversified revenue base.

At first glance, the magnitude of the move appeared disproportionate. OpenAI remains a private company, and its direct revenues are still a fraction of the global tech ecosystem. Yet, the sell-off was about what OpenAI represents and the structure underpinning the entire AI investment thesis.

Over the past two years, a powerful narrative has taken hold. AI demand is effectively limitless, and the only real constraint is compute supply. That narrative has underpinned hundreds of billions of dollars in capital allocation decisions across the technology stack.

At the centre of this ecosystem sits OpenAI, functioning as both a symbol and a demand engine. The model is circular and this is precisely the issue.

OpenAI raises capital at increasingly elevated valuations. That capital is then committed to long-term compute contracts with hyperscalers and infrastructure specialists such as Oracle and CoreWeave. These companies, in turn, invest heavily in data centres, GPUs, and networking equipment, often sourced from Nvidia and its ecosystem. Strong infrastructure demand translates into robust revenues and earnings growth for these suppliers, reinforcing high equity valuations and enabling further capital deployment into AI.

The system works as long as end-demand keeps accelerating. If OpenAI's monetisation—the most visible “end demand” signal—slows or misses targets, investors suddenly question whether the entire wheel can keep spinning at the required speed. Chief Financial Officer Sarah Friar reportedly warned leadership that the company might struggle to honour future data-centre contracts if revenue growth does not accelerate.

This is a classic case of narrative violation. OpenAI has been the poster child of the AI revolution; any sign of strain at that level forces investors to reassess the entire ecosystem.

Competition is also becoming more relevant. Models from Google (Gemini) and Anthropic are increasingly credible, particularly in enterprise and coding use cases. Pricing pressure and higher churn become more plausible outcomes, complicating the path to profitability.

The episode also amplified concerns that had already started to build. High-profile investors had begun trimming exposure to AI leaders. Figures such as Peter Thiel reduced Nvidia positions. SoftBank sold shares. Bearish voices, including Michael Burry, pointed to OpenAI as a potential “linchpin” vulnerability.

The WSJ report crystallised these concerns, resulting in a rapid repricing of expectations. The question now centres on whether the economics of AI demand can sustain the scale of investment currently underway.

Is this a bubble, an OpenAI-specific issue, or something bigger?

The market reaction reflects a broader reassessment of the AI investment narrative.

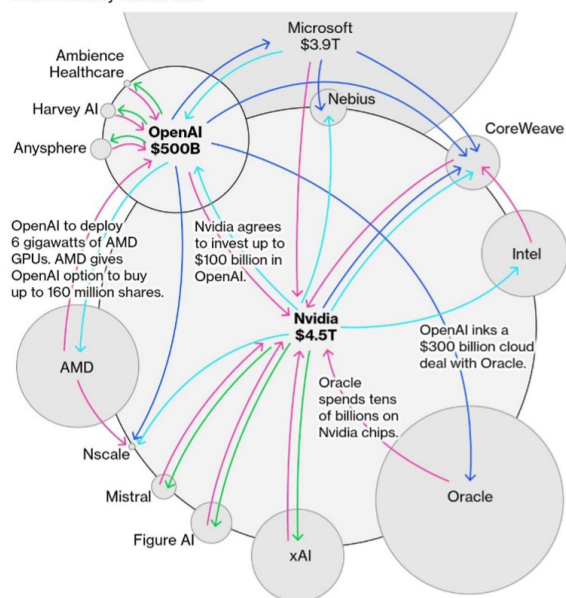
One interpretation points to a company-specific dynamic. This looks, first and foremost, like a competitive miss rather than a collapse in overall AI demand. Total usage of generative AI continues to expand at a rapid pace; the pie is simply being shared more evenly across players. Google's Gemini is gaining traction on the consumer side, and Anthropic continues to build strong positioning in enterprise and coding applications. The emergence of credible alternatives reflects a maturing market.

OpenAI itself remains a dominant force. It still commands the largest consumer base, continues to grow revenue at a pace that would be exceptional by any historical standard, and recently raised \$122bn at a valuation exceeding \$850bn, placing it among the most valuable private companies globally.

This view is reinforced at the infrastructure level. Nvidia's data-centre business remains broadly diversified, supported by demand from hyperscalers, sovereign buyers, and enterprise workloads. At the same time, Microsoft, Alphabet, Amazon, and Meta are all guiding for substantial

How Nvidia and OpenAI Fuel the AI Money Machine

Hardware or Software Investment Services Venture Capital
Circles sized by market value



Source: Bloomberg News reporting

Bloomberg

Source: Bloomberg

AI-related capital expenditures in 2026, with combined figures in the \$600-700 billion range. These investments are driven by their own ecosystems (cloud, search, advertising, and internal productivity) not solely by OpenAI's trajectory.

Private markets also remain receptive. The ability of AI companies to continue raising capital at scale suggests that investor appetite for the theme has not disappeared. Capital is still available, and the long-term narrative around AI-driven productivity gains remains intact.

On the other hand, a different interpretation focuses on the structure of the system itself. The circular financing dynamic that has powered the AI boom is also its most visible vulnerability. Capital flows from investors to AI labs, from labs to infrastructure providers, and from there into GPUs and data centres, often with limited visibility on end-user returns at scale. In that context, the distinction between a virtuous investment cycle and a fragile feedback loop becomes less clear. Cash-burn levels remain substantial even under optimistic assumptions. If inference costs fail to decline quickly enough, or if monetisation through enterprise adoption, agents, or new product layers progresses more slowly than expected the gap between investment and return becomes harder to justify. The scale of committed spending leaves little room for execution missteps.

Additionally, a significant share of major equity indices is concentrated in AI-linked names, including Nvidia, Microsoft, Alphabet, Amazon, and Meta. A repricing of growth expectations across this group would not remain contained, it would translate into broader index-level volatility.

Sentiment had already started to shift before this episode. Questions around the sustainability of capex intensity, the role of debt in financing infrastructure expansion, and the need for measurable productivity gains had been building for months.

Conclusion

This is not the moment to call the top on AI. It is, however, the moment to be more selective and more disciplined. Exposure now needs to be differentiated across the value chain. Diversified hyperscalers remain more insulated. By contrast, companies with concentrated exposure to OpenAI are more sensitive to any shift in its spending trajectory. OpenAI itself should be treated less as a sector proxy and more as a single-name risk. Revenues are scaling quickly, around a \$20bn run rate, but the issue is whether that growth can support \$1,400bn in long-term commitments.

Welcome to Syzerland®

For further information

Banque Syz SA

Quai des Bergues 1
CH-1201 Geneva
T. +41 58 799 10 00
syzgroup.com

Charles-Henry Monchau, CFA, CAIA, CMT

Chief Investment Officer
charles-henry.monchau@syzgroup.com

Assia Driss

Syz Research Lab Team Coordinator
assia.driss@syzgroup.com

Hugo Morel

Syz Research Lab Team
hugo.morel@syzgroup.com

This marketing document has been issued by Bank Syz Ltd. It is not intended for distribution to, publication, provision or use by individuals or legal entities that are citizens of or reside in a state, country or jurisdiction in which applicable laws and regulations prohibit its distribution, publication, provision or use. It is not directed to any person or entity to whom it would be illegal to send such marketing material.

This document is intended for informational purposes only and should not be construed as an offer, solicitation or recommendation for the subscription, purchase, sale or safekeeping of any security or financial instrument or for the engagement in any other transaction, as the provision of any investment advice or service, or as a contractual document. Nothing in this document constitutes an investment, legal, tax or accounting advice or a representation that any investment or strategy is suitable or appropriate for an investor's particular and individual circumstances, nor does it constitute a personalized investment advice for any investor.

This document reflects the information, opinions and comments of Bank Syz Ltd. as of the date of its publication, which are subject to change without notice. The opinions and comments of the authors in this document reflect their current views and may not coincide with those of other Syz Group entities or third parties, which may have reached different conclusions. The market valuations, terms and calculations contained herein are estimates only. The information provided comes from sources deemed reliable, but Bank Syz Ltd. does not guarantee its completeness, accuracy, reliability and actuality. Past performance gives no indication of nor guarantees current or future results. Bank Syz Ltd. accepts no liability for any loss arising from the use of this document.