

The fastest revenue ramp in tech history, and the infrastructure buildout behind it

Published 10 June 2026

Blake Heimann

Senior Associate, Quantitative Research

Key Takeaways

- AI revenue growth is rewriting the record books: OpenAI and Anthropic have reached revenue milestones in under four years that took previous technology leaders decades to achieve.
- The real AI opportunity extends far beyond model providers: Every layer of the infrastructure stack, from semiconductors and networking to power and data centres, is being rebuilt to support AI.
- A multi-year capital expenditure cycle is underway: Constrained compute, power and manufacturing capacity are driving sustained investment across the companies enabling AI at scale.
- A differentiated way to invest in AI: The WisdomTree AI Infrastructure UCITS ETF targets the companies building the semiconductors, networking, power and data centre infrastructure underpinning the next phase of AI growth.
- Related Products WisdomTree AI Infrastructure UCITS ETF - USD Acc Find out more

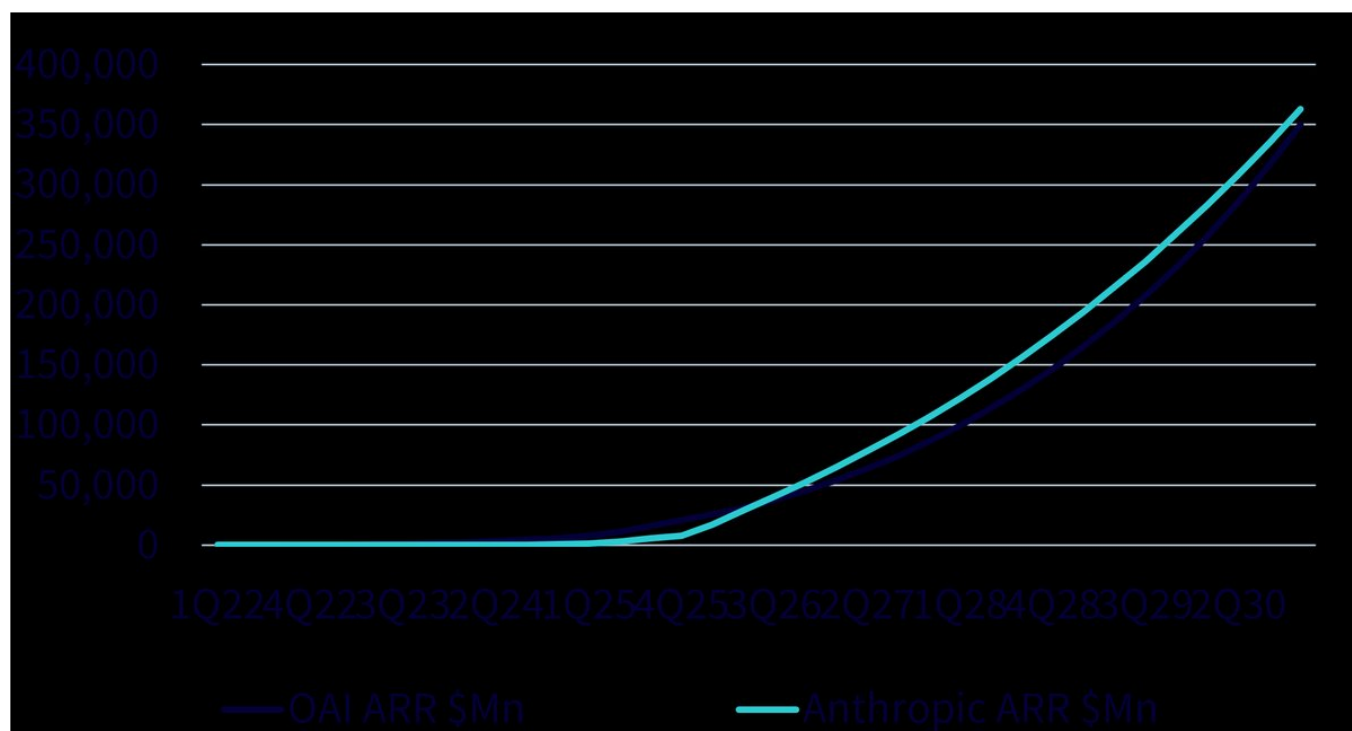
A new chapter in tech history

The leading artificial intelligence (AI) labs are scaling revenue faster than any cohort in technology history. In 2026, the frontier players are not just building smarter models; they are growing revenue at a pace that makes prior technology cycles look slow.

Consider OpenAI. The company behind ChatGPT has scaled from effectively zero in 2022 to roughly \$2 billion in monthly revenue in early 2026¹. Annualised revenue is now reported to be approaching the \$25 billion range, with enterprise deals accounting for around 40% of the business and growing².

Anthropic, the company behind Claude, is competing aggressively and, based on recent reporting, has surpassed OpenAI. It has reached approximately \$30 billion in annualised revenue around April 2026. This represents a jump from roughly \$9 billion at the end of 2025, more than tripling its run rate in just four months³.

Figure 1: Annualised revenue of leading AI model providers: OpenAI vs Anthropic, 2022-2030E



Source: SemiAnalysis, April 2026. **Forecasts are not an indicator of future performance and any investments are subject to risks and uncertainties.**

For context, it took Salesforce roughly 20 years to reach \$30 billion in annual revenue; a milestone that Anthropic and OpenAI have each roughly achieved in under four years. No prior software cycle has matched this pace, outpacing traditional software companies, early cloud giants, and every consumer technology cycle that came before it. Some industry forecasts suggest continued growth, although actual outcomes may differ materially.

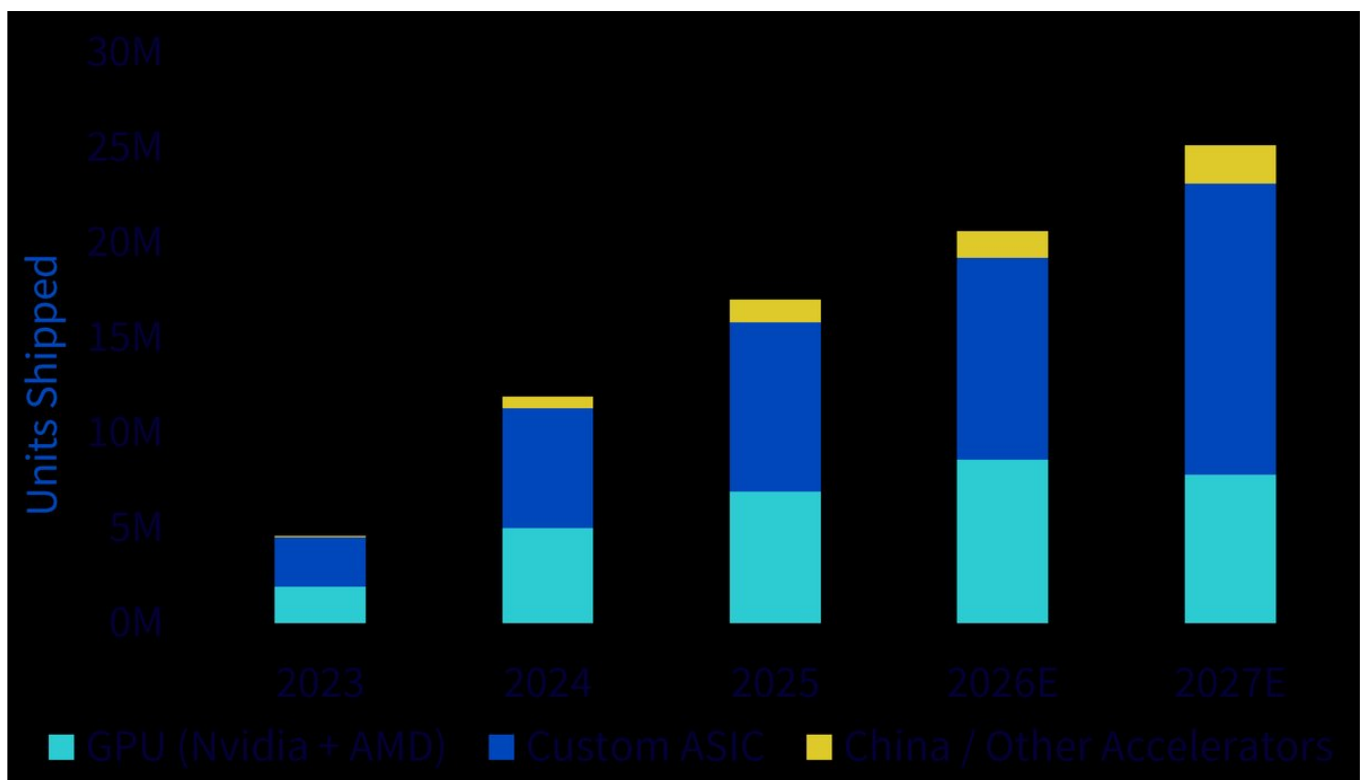
With initial public offerings (IPOs) on the horizon, 2026 could mark the moment AI fully steps into the public markets spotlight. The question for investors is: who supplies the picks and shovels?

The infrastructure beneath the revenue

All of this AI activity, from consumer chat to enterprise workflows, sits on top of a vast and increasingly constrained physical infrastructure.

Demand for AI compute continues to scale rapidly, driven by both model development and the accelerating shift toward inference. Behind this growth is a layer of hyperscalers and neocloud providers (Microsoft, Google, Amazon, Meta, Oracle and a new generation of specialised operators) on track to invest over \$700 billion in combined capital expenditures in 2026, roughly 60–70% above 2025 levels, with the vast majority directed toward AI infrastructure⁴. They are also designing custom silicon alongside leading platforms such as NVIDIA, deploying these compute accelerators, including graphics processing units (GPUs), application-specific integrated circuits (ASIC) and custom silicon chips, at scale to train and run increasingly complex artificial intelligence models.

Figure 2: XPU unit shipments - GPUs (NVIDIA + AMD) vs custom ASIC vs other accelerators, 2023-2027E

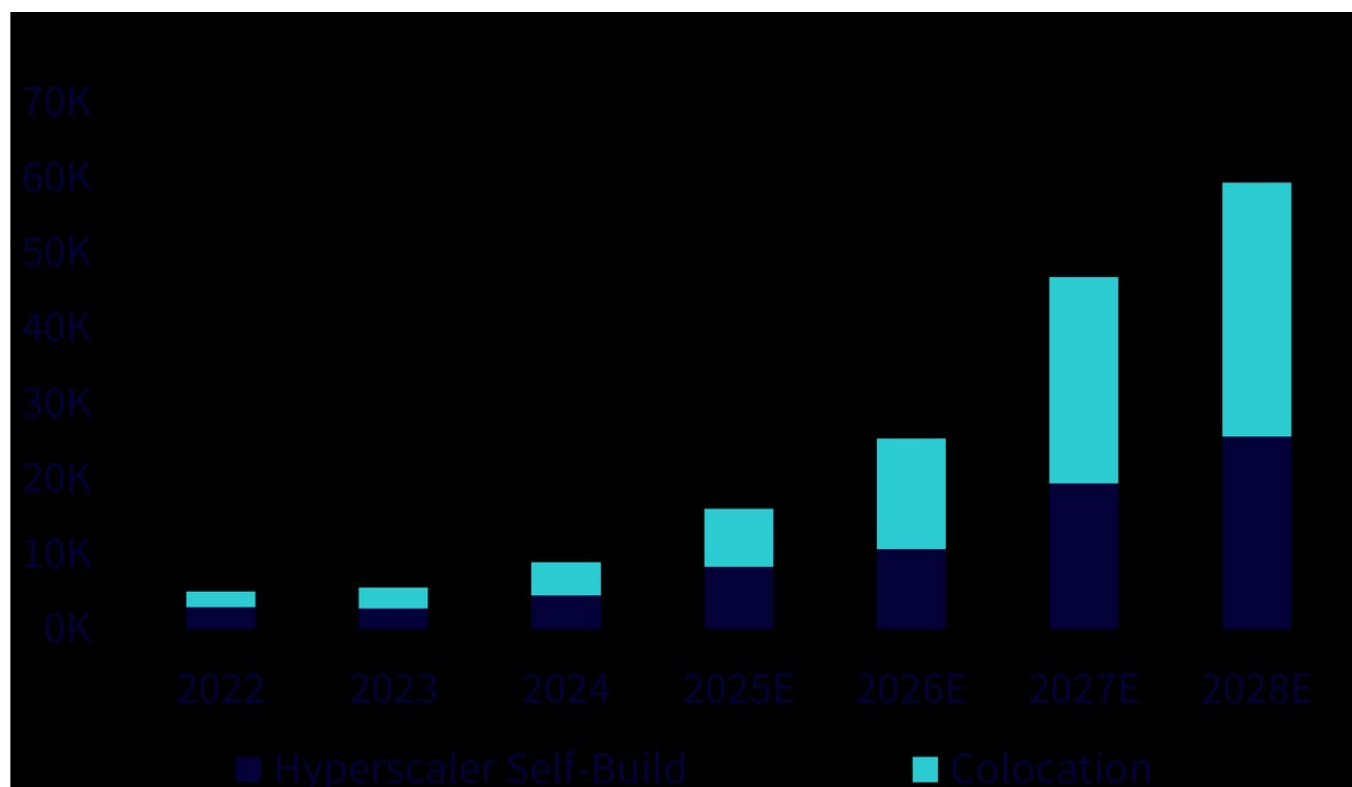


Source: SemiAnalysis Accelerator Model, April 2026. **Forecasts are not an indicator of future performance and any investments are subject to risks and uncertainties.**

This translates directly into the physical expansion of the data centre infrastructure needed to house the AI servers and their chips. AI workloads are significantly more power-intensive than traditional compute, requiring specialised 'AI factories' built for high-density systems.

Power is emerging as a critical constraint. Global data centre capacity is accelerating as hyperscalers and colocation providers expand their footprint, but grid access, cooling systems and electrical infrastructure are becoming binding bottlenecks. This may support a multi-year capital expenditure cycle across power, industrial and electrical equipment suppliers.

Figure 3: Global data centre MW additions, 2022–2028e: hyperscaler self-build vs colocation (ex-China)



Source: SemiAnalysis AI Datacentre Industry Model, April 2026. Chart excludes China. **Forecasts are not an indicator of future performance and any investments are subject to risks and uncertainties.**

Within the data centre, system architecture is becoming more complex. AI servers are built around high-performance XPUs, alongside high-bandwidth memory and specialised networking infrastructure. As AI clusters scale to tens of thousands of processors, high-speed optical and electrical connectivity becomes essential to maintain performance.

This complexity extends upstream into semiconductor manufacturing, where spending on wafer fabrication equipment is accelerating to meet demand for advanced logic and memory chips. With both fabrication and advanced packaging capacity constrained, the broader equipment ecosystem is well-positioned to benefit from the sustained investment required to support this multi-year buildout.

The scale of this buildout reaches well beyond the chipmakers most investors associate with AI. From the silicon at the core, to the servers and networking that connect it, to the power and industrial systems that keep it running, every layer of the stack is being rebuilt for AI workloads, and every layer is a place where capital is being deployed. The harder question is which companies, across an ecosystem this broad, are best positioned to capture it.

For the full case behind the AI infrastructure investment theme, read [The Case for AI Infrastructure](#).

Meet our research partner: SemiAnalysis

Identifying the right companies across the complex AI value chain requires specialist expertise. That is why WisdomTree has partnered with SemiAnalysis, a leading independent research firm, and the only firm tracking every layer of the semiconductor and AI stack.

SemiAnalysis at a glance:

- Headquartered in San Francisco, with over 25 analysts covering the full semiconductor and AI value chain
- 100+ articles published annually and 80+ industry conferences attended each year
- Coverage spans capital equipment, fabrication, chip design, networking, server architecture, data centres and model training and inference economics

What sets SemiAnalysis apart is the depth and originality of its research. They combine proprietary data sources, including satellite imagery, die-shot analysis and direct supply-chain benchmarking, with primary research that other research desks cannot replicate. Their proprietary Accelerator, WFE (Wafer Fab Equipment), AI TCO (Total Cost of Ownership), Datacenter, Foundry and GPU pricing models inform our index construction and ongoing maintenance of the strategy behind the exchange-traded fund (ETF).

Founder Dylan Patel engages directly with senior executives across the ecosystem, including NVIDIA's Jensen Huang, Lisa Su of AMD and Satya Nadella of Microsoft. SemiAnalysis's research is read across the institutional investor community and inside the companies they cover: from hyperscalers to semiconductor manufacturers.

Learn more about our research partner in [Meet the Partner: SemiAnalysis](#)

Introducing the WisdomTree AI Infrastructure UCITS ETF

The fund provides targeted exposure to the companies behind the AI infrastructure buildout. The underlying Index is designed in partnership with SemiAnalysis, spanning seven key AI infrastructure categories:

Companies in these categories are evaluated using a Composite Score combining a forward-looking qualitative assessment (80%) of strategic positioning with an AI Revenue Score (20%) measuring direct revenue exposure. Each category is assigned an industry score multiplier, reflecting differences in capital intensity, infrastructure constraints and strategic importance. Weights are determined by the product of these scores, directing higher allocations to companies with the strongest exposure to the most strategically important parts of the value chain. The Index is rebalanced quarterly to remain aligned with the evolution of the AI infrastructure stack.

Why WisdomTree?

- **Expert-driven selection** built in collaboration with SemiAnalysis.

- **Full value chain exposure** spanning compute, power, networking, semiconductors, data centres and beyond.
- **Differentiated by design**, capturing critical and often overlooked segments and bottlenecks.
- **Targeting structural growth** tied to the multi-year AI spending cycle.

For a deeper look at the index methodology and portfolio construction, read the [Investment Case](#) and the [Index Methodology](#).

In conclusion

The revenue story for AI is unprecedented. But it sits on a physical foundation that has to be built. The [WisdomTree AI Infrastructure UCITS ETF](#) provides exposure to the companies enabling the scaling of artificial intelligence at every layer of the stack.

While AI adoption continues to expand rapidly, investors should be aware that future growth may be slower than expected and that many companies across the AI ecosystem are exposed to technological, competitive, regulatory and execution risks. Valuations within AI-related sectors can be sensitive to changes in investor sentiment, earnings expectations, infrastructure bottlenecks and broader market conditions, which may result in periods of significant volatility.

1 OpenAI, as of 31 March, 2026; additional reporting from CoinDesk and The Stack.

2 Sacra analysis of OpenAI revenue; OpenAI statements, February–March 2026.

3 Bloomberg and VentureBeat reporting, April 2026; Anthropic-related announcements.

4 Combined 2026 capex guidance from Q4 2025 / FY2026 earnings: Amazon ~\$200B; Alphabet \$175–185B; Meta \$115–145B; Microsoft \$120–145B+ (with an \$80B+ order backlog reflecting power constraints); Oracle ~\$50B (+136% YoY). Sources: company earnings releases and transcripts; Reuters, Bloomberg, CNBC, Futurum Group, Yahoo Finance.

Important Risks Related to this Article

Important Information

Marketing communications issued in the European Economic Area (“EEA”): This document has been issued and approved by WisdomTree Ireland Limited, which is authorised and regulated by the Central Bank of Ireland.

Marketing communications issued in jurisdictions outside of the EEA: This document has been issued and approved by WisdomTree UK Limited, which is authorised and regulated by the United Kingdom Financial Conduct Authority.

WisdomTree Ireland Limited and WisdomTree UK Limited are each referred to as “WisdomTree” (as applicable). Our Conflicts of Interest Policy and Inventory are available on request.

For professional clients only. The information contained in this document is for your general information only and is neither an offer for sale nor a solicitation of an offer to buy securities or shares. This document should not be used as the basis for any investment decision. Investments may go up or down in value and you may lose some or all of the amount invested. Past performance is not necessarily a guide to future performance. Any decision to invest should be based on the information contained in the appropriate prospectus and after seeking independent investment, tax and legal advice.

The application of regulations and tax laws can often lead to a number of different interpretations. Any views or opinions expressed in this communication represent the views of WisdomTree and should not be construed as regulatory, tax or legal advice. WisdomTree makes no warranty or representation as to the accuracy of any of the views or opinions expressed in this communication. Any decision to invest should be based on the information contained in the appropriate prospectus and after seeking independent investment, tax and legal advice.

This document is not, and under no circumstances is to be construed as, an advertisement or any other step in furtherance of a public offering of shares or securities in the United States or any province or territory thereof. Neither this document nor any copy hereof should be taken, transmitted or distributed (directly or indirectly) into the United States.

Although WisdomTree endeavours to ensure the accuracy of the content in this document, WisdomTree does not warrant or guarantee its accuracy or correctness. Where WisdomTree has expressed its own opinions related to product or market activity, these views may change. Neither WisdomTree, nor any affiliate, nor any of their respective officers, directors, partners, or employees accepts any liability whatsoever for any direct or consequential loss arising from any use of this document or its contents.