ametra

AMPLIFY | DEEP DIVE

Beyond the Headlines

6 JULY 2026

AI ROI: Real Money or Dot-Com Redux?

Ametra Investment Managers Private Limited | SEBI Registered Portfolio Manager | Reg No. INP000008905

AMETRA INVESTMENT RESEARCH · DEEP DIVE · JULY 2026

AI ROI: Real Money or Dot-Com Redux?

The Infrastructure Bet That Success Has Already Priced In

The scoreboard says the AI trade is working. Nvidia posted Q1 FY2027 revenue of $81.6 billion, up 85% year-on-year. Microsoft, Amazon, Google and Meta collectively guide $725 billion of capital expenditure for 2026 — the largest coordinated private investment cycle in recorded history.

SpaceX listed on Nasdaq on June 12, raising $85.7 billion at a $1.75 trillion valuation — the largest IPO in history, with Anthropic and OpenAI in the pipeline behind it. The market has priced all of this as settled fact.

But it is not. Beneath the boom, three structural problems are compounding at once:

- the debt cycle now funding what free cash flow cannot;
- an enterprise adoption story where 95% of pilots fail to reach the P&L;
- a revenue-recognition architecture in the GPU supply chain that is circular by design.

The market has priced the bull case. It has not priced the plausible alternatives.

Macro Snapshot (as of 3 July 2026)

Indicator	Reading	Read-through
Big-5 AI capex guidance (2026)	~$725 billion	Up 77% from $410B in 2025; largest in corporate history
Hyperscaler debt issuance (2025)	$121 billion	4× the prior decade's annual average
AI direct revenue vs. capex	~$50–60B vs. $725B	$8–10 of investment per $1 of current revenue
Sequoia revenue gap	~$600 billion / year	Capex requires end-user revenue that does not yet exist
Capex/revenue divergence vs. 2001	46% vs. 32%	Already exceeds the excess that preceded the telecom bust
Enterprise GenAI ROI (MIT, 2025)	95% see zero P&L impact	Adoption is broad; value realisation is not
Amazon Q1 FY2026 free cash flow	Down 95% YoY to $1.2B	Capex consuming all operating surplus
AI share of US GDP growth (H1 2025)	~92%	Macro story now depends on AI capex continuing
Data center capacity: announced/built	12GW / 5GW (2026)	5-year grid transformer backlog is the constraint
CAPE ratio (S&P 500, Jul 2026)	~41x	Second-highest in 125 years; dot-com peak ~44–45x

1. The Debt That Wasn't There Before

For a decade, the hyperscalers were the envy of bond investors precisely because they did not need the bond market. Microsoft, Amazon, Alphabet and Meta were cash machines, generating $200 billion of free cash flow as recently as 2024. AI has changed the arithmetic.

In 2025 the five largest issued **$121 billion** of US corporate bonds — four times their average annual issuance over the prior decade — and Bank of America calculates that capex now consumes **94% of hyperscaler operating cash flow** after dividends and buybacks. The result is on the balance sheet below.

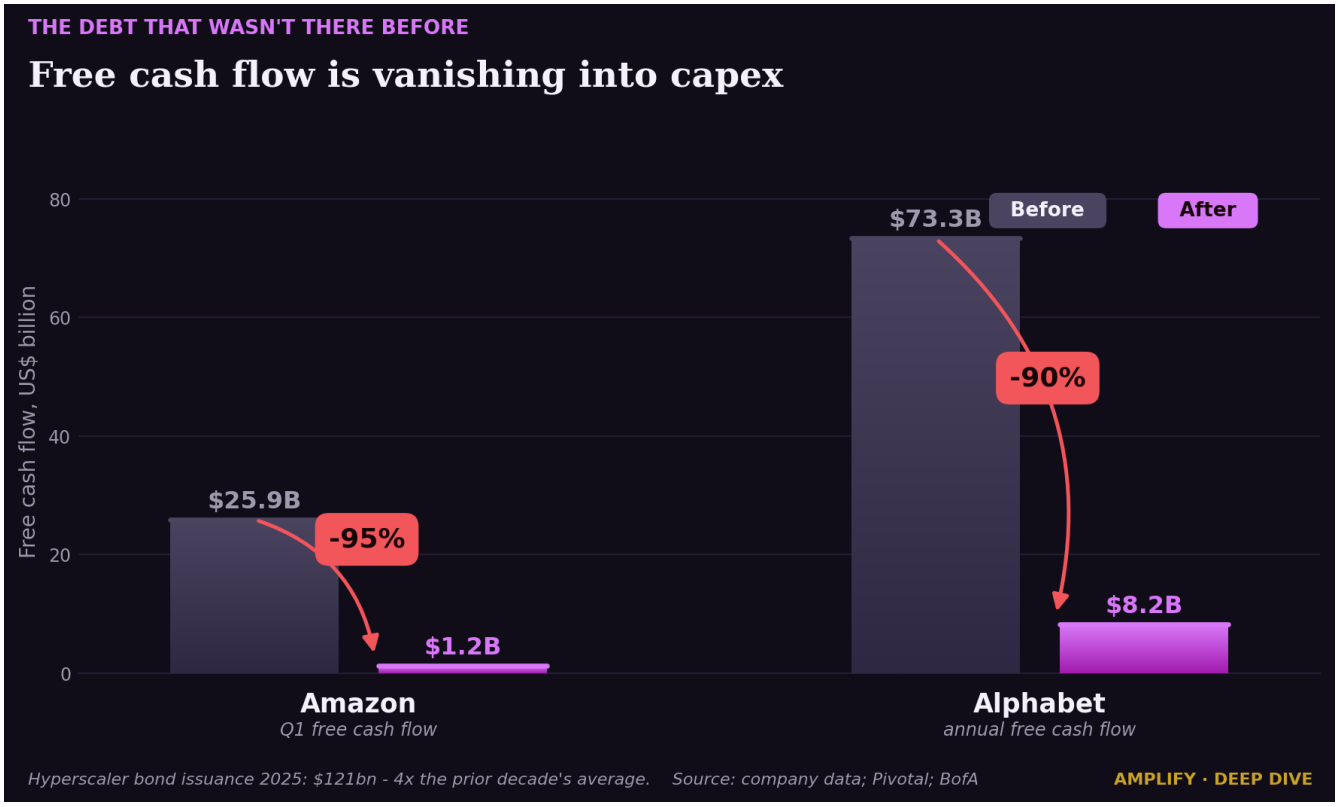

The most extreme case is Oracle. After a $300 billion data-centre deal with OpenAI, it raised 2026 capex guidance 43% to $50 billion, opening a funding gap above $27 billion a year. Its five-year credit-default-swap spread widened from ~40 to ~200 basis points — a 16% implied default probability for an investment-grade issuer — and its equity fell 60% from its September 2025 peak. The sequence matters: credit markets repriced first; equity followed.

And the risk is not confined to tech. A Harvard analysis found that information-processing equipment and software drove roughly 92% of US GDP growth in the first half of 2025. AI capex is now the primary engine of US growth, so any slowdown threatens headline GDP, not just tech stocks — which is why the correction will be politically deferred as long as possible, and deeper for it.

2. The $600 Billion Question No One Is Answering

Sequoia's David Cahn framed the arithmetic bluntly. Take Nvidia's GPU revenue run-rate, double it for total data-centre cost of ownership (GPUs are about half; the rest is energy, buildings, cooling and networking), and double again for the margin the end-buyer of compute needs. As of mid-2026 that number is roughly $600 billion a year. Current AI direct revenue across the ecosystem is $50–60 billion — a gap of $8–10 of investment for every $1 of revenue.

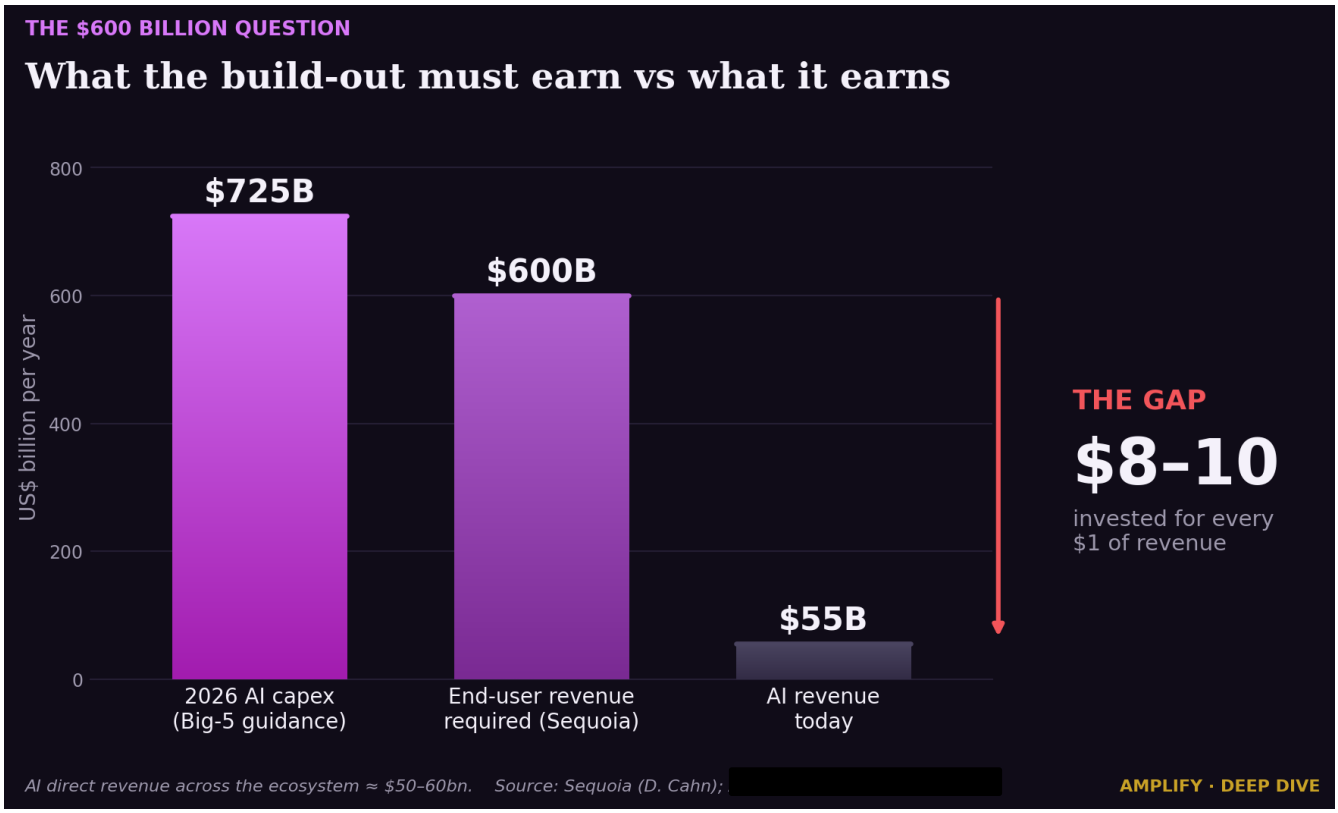

The adoption data makes the gap harder to close, not easier. MIT's Project NANDA found 95% of organisations deploying generative AI saw zero measurable P&L impact — not low returns, zero. BCG put 60% at no material value; McKinsey found 88% use AI somewhere but only 39% see any EBIT impact; Gartner expects 60% of AI projects without AI-ready data to be abandoned. Allianz pegs the capex-vs-revenue divergence at 46%, already beyond the 32% of the 2001 telecom bust.

The unit economics compound it. Token costs are rising even as enterprises are pushed to use more: GPT-5.5 is priced at $5 per million input and $30 per million output tokens, and one company reportedly spent $500 million in a single month on Claude Code alone. The buyer overpays for marginal productivity; the seller underprices a service running on borrowed margin. Neither model is sustainable.

3. The Circular Machine: How GPU Revenue Is Counted

The most uncomfortable part of the AI trade is one the earnings headlines do not surface: a significant share of the revenue booked at the supply layer is funded, directly or indirectly, by the same ecosystem that benefits from reporting it.

Nvidia's Q2 FY2026 filing shows two unnamed customers at 39% of revenue — and its largest customers (Microsoft, Google, Meta, Amazon) are also its largest AI investors and cloud partners. **When Microsoft buys GPUs, sells compute to OpenAI, and part-funds OpenAI's build-out, the revenue chain is not fully arm's-length.**

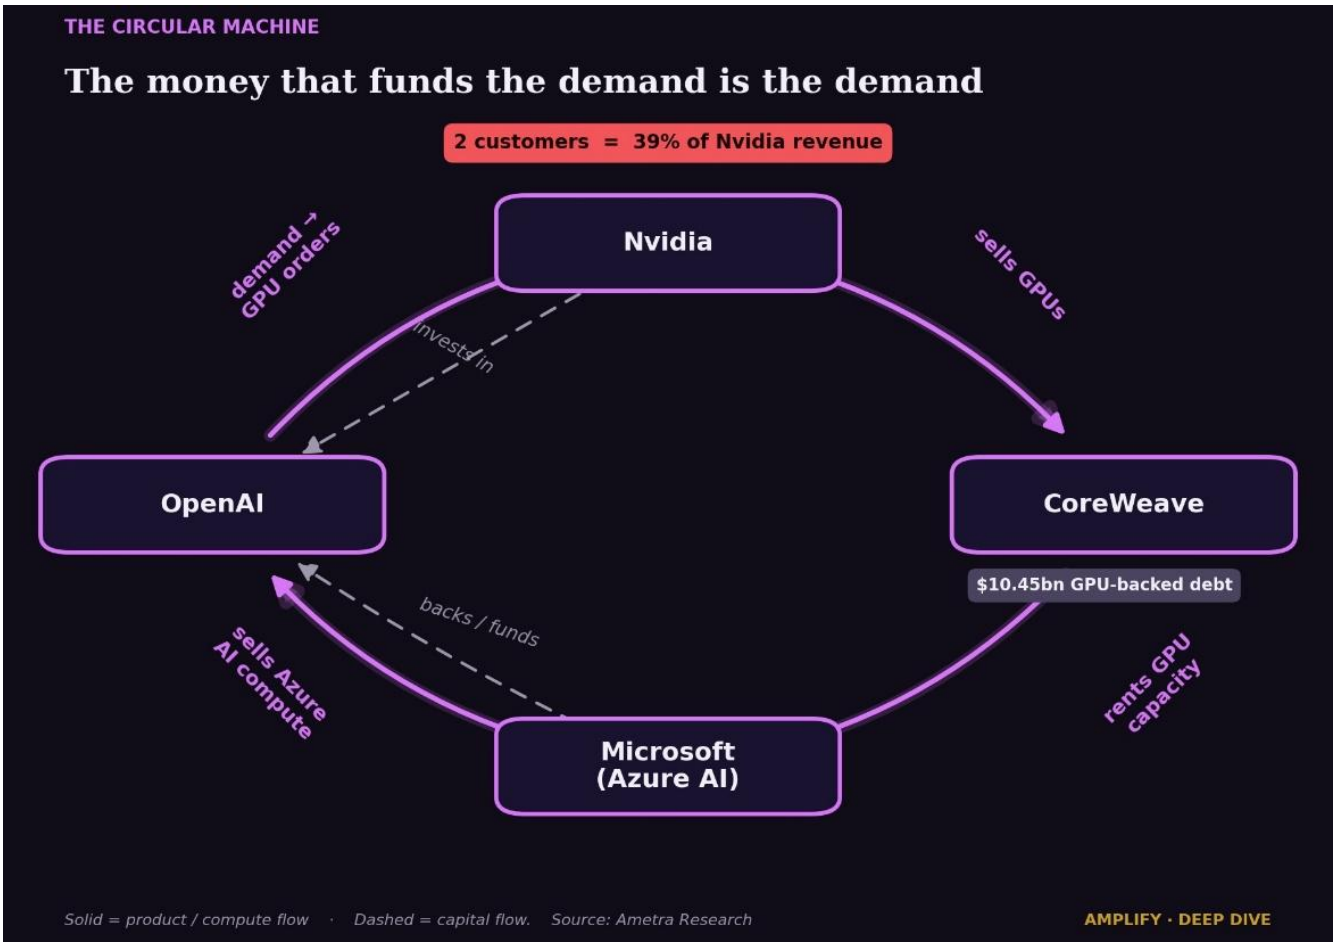

The neocloud makes it explicit. CoreWeave carries $10.45 billion of largely GPU-collateralised debt; Nvidia holds a 7% stake, is its main supplier, and has agreed to absorb unused capacity through 2032 — while CoreWeave's revenue depends on Microsoft and OpenAI, two of Nvidia's own largest customers.

The parallel to Lucent in 1999 is instructive. At its peak Lucent had $15 billion of vendor financing — lending customers the money to buy its equipment — against $300 million of operating cash flow. Jim Chanos flags the modern echo in the accounting: GPUs are depreciated over six years while they obsolesce in three to four.

4. Ametra's Read

We do not read the AI trade as a binary — technology worthless or technology transformative. The technology is real. What is not yet real is the revenue architecture that justifies current valuations and capex commitments at the same time.

The infrastructure layer has been priced as if the application layer is already delivering at scale. It is not. The 95% enterprise failure rate, the $600 billion revenue gap, the token-cost paradox and the free-cash-flow compression at the very companies funding the build-out are not secondary concerns — they are the central variable on which the entire valuation edifice rests.

There is a physical limit the financial model has not absorbed. The IEA projects data-centre electricity demand doubling to ~945 TWh by 2030 — roughly Japan's entire consumption. Yet of 12GW of US capacity announced for 2026, only 5GW is under construction; the bottleneck is high-voltage transformers, now on five-year lead times, up from 24–30 months pre-2020, with the IEA flagging 20% of projects at risk of grid delay. Hyperscalers are committing $725 billion a year to build capacity they cannot physically deploy on schedule.

The dot-com parallel crystallises here with unusual precision. In 1999, telecoms laid 80 million miles of fibre for an internet that arrived years later than their debt could tolerate. The cables are used today; they carry today's internet. The companies that built them went bankrupt in the interim. The AI equivalent is now visible: infrastructure committed faster than transformers can be built, faster than grids can be permitted, and faster than end-users are generating the revenue to justify it.

Position for the Gap, Not for the Hype

The IPO wildcard. SpaceX completed the largest IPO in history on June 12; OpenAI ($852bn) and Anthropic ($965bn) target listings in September and October. Three companies that have never turned a sustained profit are asking public markets to absorb roughly $3.6–3.8 trillion of market cap in a matter of months — capital that will come, in part, from the same portfolios that hold Nvidia, Microsoft and the rest of the AI trade. The mega-IPO wave is not a catalyst for the rally. It is a liquidity test of it.

Stay cautious on pure-play AI infrastructure. where valuation assumes revenue that has not materialised.

Favour genuine end-user revenue today. over companies whose valuation depends on the $600 billion being found.

Monitor the IPO pipeline. how SpaceX trades relative to its private valuation after listing is one of the cleaner real-time signals on whether public markets still price AI narrative ahead of AI fundamentals.

Sector Positioning Map

Stance	Where	Why
Caution — priced for perfection	Nvidia, GPU-pure plays	Revenue real but circular; 2 customers = 39% of revenue; depreciation risk unpriced
Caution — FCF collapse	Amazon, Alphabet, Oracle	Capex consuming 90–95% of FCF; credit repricing underway at Oracle
Caution — IPO liquidity risk	Existing Mag-7 holdings	SpaceX/OpenAI/Anthropic listings absorb $200B+ from portfolios already long AI
Watch — application layer	Enterprise SaaS with real AI revenue	Genuine end-user demand, not an infrastructure bet
Relative caution — neocloud	CoreWeave, Nebius	GPU-backed debt; revenue lags capex 2:1; first maturities 2026–27
Physical bottleneck	Grid, transformers	5-year transformer backlog; only 5GW of 12GW building; 20% at delay risk
Macro overhang	US equity market broadly	AI capex = 92% of US GDP growth; any slowdown is macro, not sectoral

Disclaimer

Investments in securities markets are subject to market risks. Please read all related documents carefully before investing. Past performance may not be indicative of future results.

Performance-related information provided herein has not been verified by SEBI. Any forward-looking statements, projections, estimates, or forecasts are based on current assumptions and expectations and should not be construed as guarantees of future performance.

This document has been prepared by Ametra Investment Managers Private Limited ("Ametra") based on internal analysis, proprietary models, publicly available information, and sources believed to be reliable. While reasonable care has been taken in its preparation, Ametra makes no representation or warranty, express or implied, as to the accuracy, completeness, or adequacy of the information contained herein. The information provided is for informational purposes only and does not constitute investment, legal, tax, or financial advice, nor should it be relied upon as the sole basis for any investment decision. Investors are advised to consult their financial, legal, and tax advisors and exercise independent judgment before making any investment decisions.

Ametra, its directors, employees, affiliates, and associates shall not be liable for any direct or indirect loss arising from the use of, or reliance upon, this document. Any third-party trademarks, logos, brand names, indices, or data referenced herein remain the property of their respective owners and are used solely for identification and informational purposes.

Ametra Investment Managers Private Limited | SEBI Registered Portfolio Manager | Reg No. INP000008905

ametra

AMPLIFY · DEEP DIVE

AMETRA

SEBI REGISTERED PORTFOLIO MANAGER

INP000008905

support@ametra.in

+91 90194 69258

www.ametra.in