



PERSPECTIVE · SEPTEMBER 2026

THE BILL COMES LATER

AI doesn't have a technology problem. It has a return-on-capital problem. The technology may change the world. The harder question is who earns a return financing it.

BRYAN MUSTO · CHIEF REVENUE OFFICER · 9 MIN READ · ·

From 2005 to 2010 I was an equity derivatives analyst at a global investment bank, which meant watching the financial crisis get priced in real time. In August 2007, the bank I worked for said it could no longer value three of its funds because liquidity in parts of the market had "evaporated." Many historians date the crisis from that morning. "Evaporated" may be the most honest word ever put in a bank press release.

What I remember most is how reasonable everything looked beforehand. The houses were real. The demand was real. The models rested on one assumption almost nobody questioned: that national home prices don't fall.

The AI buildout rests on assumptions too. The technology is real, and I expect it to change the global economy. But the recent sell-off isn't the market deciding AI is a fad. It's the market starting to ask a narrower question, one that can actually be answered: when does the bill arrive, and who pays it?

THE BILL

Capex shows up as someone else's revenue today, and as your own cost for years afterward.

The five largest U.S. hyperscalers (Microsoft, Alphabet, Amazon, Meta and Oracle) spent roughly \$150 billion on capital expenditures in 2023. Goldman Sachs estimates about \$800 billion this year and \$1.4 trillion by 2028, and attributes almost half of 2026 S&P 500 earnings growth to AI investment. By Goldman's estimate, a \$250 billion surprise in 2027 capex would move S&P 500 earnings growth by about six points. A handful of capital budgets now steer the earnings of the whole index.

Here is what the bull case leaves out. **Capex is the only expense that shows up as someone else's revenue before it shows up as your own cost.** When a hyperscaler buys a rack of GPUs, the chipmaker books the sale immediately. The buyer books an asset and depreciates it over the years that follow. The earnings boost arrives first. *The bill comes later.*

Goldman's own work shows the turn coming. It estimates depreciation will subtract about five points from S&P 500 earnings growth in 2027, roughly half the lift from AI spending, and could offset all of it by 2028.


BUBBLES, IN THE POPULAR IMAGINATION, POP. I SUSPECT THIS ONE WILL AMORTIZE.

MEASURING THE BILL

To see how large the bill gets, I built a simple model. It takes the five companies' capex since 2023, from their filings through 2024 and Goldman's estimates after that, and runs it through straight-line depreciation. Sixty percent goes to servers and network equipment, depreciated over the server life. Forty percent goes to buildings and power, depreciated over 25 years once in service. Then it asks how much revenue that capacity must earn to cover the depreciation plus a 10% pre-tax return on the capital still tied up.

The amortization gap

Annual depreciation from the five largest hyperscalers' capex since 2023, and the revenue needed to cover it plus a 10% pre-tax return. \$ billions. Lines show a five-year server life; shaded bands span six-year (lower) to four-year (upper) lives.



Author's model. Capex: company filings (2023–24), Goldman Sachs estimates (2025–28), held flat after 2028. Assumes 60% of capex is servers and network equipment depreciated over the stated life and 40% is buildings and power depreciated over 25 years once in service. Excludes power, staff and other operating costs, so it understates the revenue required. See methodology below.

💡 BRYAN MUSTO

The model lines up with reality. On the companies' own six-year lives, it produces about \$130 billion of 2026 depreciation from post-2023 assets, consistent with the roughly \$160 billion a year the five now report at their mid-2026 run rate, older assets included.

Three things stand out. First, by 2028 the buildout generates roughly **\$390 billion to \$530 billion a year of depreciation**, depending on server life. That is more than the five companies' entire cloud revenue today, which runs at about \$400 billion a year. Second, covering that bill plus a 10% return takes roughly **\$675 billion to \$785 billion of annual revenue**, before a dollar of electricity or payroll. Third, the bill keeps rising after the spending stops growing. Hold capex flat at \$1.4 trillion after 2028, and depreciation still climbs every year through 2030, because each year's spending stacks on top of the last.

SERVER LIFE	2027 DEPRECIATION	2028 DEPRECIATION	2028 REVENUE NEEDED (WITH 10% RETURN)
6 YEARS (COMPANY STANDARD)	\$244B	\$393B	\$675B
5 YEARS	\$288B	\$454B	\$724B
4 YEARS	\$342B	\$528B	\$783B
3 YEARS	\$410B	\$626B	\$860B


Author's model; same assumptions as the chart above. Depreciation counts only assets bought since 2023.

THE USEFUL-LIFE QUESTION

How fast the bill arrives depends on how long the chips last. A product cycle, a better chip every year or two, is not the same as economic life: older GPUs keep earning on inference and lighter workloads. The companies' own assumptions already diverge, and Michael Burry estimates that stretched lives will understate industry depreciation by about \$176 billion from 2026 to 2028. Nvidia counters that four to six years is realistic.

Useful life is a judgment call

Depreciation lives for servers and GPUs, in years: what companies book versus what the debate argues.



Sources: company filings (Amazon 2025 10-K; Microsoft, Alphabet and Meta disclosures); Nvidia; Michael Burry via CNBC.

BRYAN MUSTO

A six-year depreciation schedule against a roughly two-year product cycle isn't necessarily wrong. But it embeds an enormous assumption about residual productivity, and the model shows its price: moving from six years to four adds about \$135 billion to 2028 depreciation alone.

THE BULL CASE

The demand is real. The question is whether it compounds fast enough, for long enough.

The strongest argument against my view is in the companies' own filings, and it deserves a fair hearing. In the most recent quarter, Azure grew 43%, AWS 37%, Google Cloud 82% and Oracle's cloud infrastructure 121%. Nvidia's data-center revenue rose 117%. Microsoft, Oracle, Google Cloud and Amazon now report about \$2.35 trillion of contracted backlog between them. OpenAI's annualized revenue is reportedly nearing \$70 billion, and Anthropic told investors its run rate reached \$65 billion in July.

On that evidence, the bill is payable. Cloud revenue of about \$400 billion growing 40% a year reaches roughly \$780 billion by mid-2028, the top of the range above.

But look at what that requires. Every incremental dollar of cloud revenue would have to go to paying for the new assets, with nothing left for power, people or the older infrastructure. And 40% growth would have to hold for three more years on a base that is already enormous. Backlog is a promise, not cash. Oracle burned about \$5 billion of free cash flow last quarter while reporting \$664 billion of contracted revenue, much of it reportedly tied to a single customer, OpenAI. Anthropic's leaked draft prospectus reportedly lists \$518 billion of compute obligations.

The market isn't really pricing whether AI works. It's pricing whether 40% growth lasts until 2029.

Meanwhile the end customer is still early. McKinsey's 2025 survey found only 39% of companies reporting any enterprise-level profit impact from AI, and Gartner puts measurable ROI at about one AI initiative in five. From my seat in enterprise technology, the distance between a successful pilot and a margin line is where most AI budgets wait. Satya Nadella has proposed the right test: for this not to be a bubble, the benefits need to be "much more evenly spread."

None of this requires anyone to be foolish. Mark Zuckerberg has said that if Meta ends up "misspending a couple of hundred billion dollars," the greater risk would still have been not spending. For any one company he's probably right. But rational firms can still overbuild an industry. Through a derivatives lens, capex is an option premium paid to stay in the game, and options decay. In this buildout, theta has a more familiar name: depreciation.

WHO PAYS

If stress comes, it won't start where the spending is largest. It will start where the balance sheets are thinnest.

Equity markets debate stories. Credit markets eventually demand cash flow, and right now credit is asking harder questions than equity. Goldman counts about \$88 billion of AI-related borrowing by lower-rated companies this year. Neuberger Berman put AI-related leveraged-finance issuance at about \$20 billion in the first eleven months of 2025. SoftBank's \$11.1 billion bond sale in September, the largest high-yield deal on record, priced its dollar notes between 8.625% and 9.75%. CoreWeave's 2026 notes priced at 9.625% and 9.75%. Demand was strong. The price was the message.

A lender's upside is capped at the coupon. It gets paid from utilization, contracted cash flow and what the collateral is worth in a bad year, so problems show up in credit before they show up in a valuation model with a thirty-year tail. The broad high-yield spread widened from 266 to 312 basis points over the final seven trading days of September. That isn't a crisis. It is the market starting to charge for the question.

Anyone who worked through 2008 learned that losses land hardest on people who didn't know what they owned; most of Madoff's victims held him through feeder funds. AI is not a fraud, but its risk is layered the same way.

Where AI risk sits

Four layers of the same buildout. Further down the stack, balance sheets get weaker or harder to see.

1	The fortresses Hyperscalers funding most of their spending from operating cash flow.	LEVERAGE		TRANSPARENCY	
2	The levered builders Neoclouds, data-center developers and infrastructure vehicles borrowing at 9% to 10%.	LEVERAGE		TRANSPARENCY	
3	The opaque layer Special-purpose vehicles, private credit, data-center asset-backed securities, and vendor financing in which suppliers invest in the customers who buy their chips. Much of vendor financing's history is written in the footnotes of bankruptcy filings. Lucent, Nortel and Insull's stacked holding companies are all in there.	LEVERAGE		TRANSPARENCY	
4	Everyone else Index funds, pensions and retirement accounts carrying a concentrated position in all of the above.	LEVERAGE		TRANSPARENCY	

Leverage and transparency ratings are the author's qualitative assessment.

Leverage behaves like water. It flows downhill to the balance sheet least able to hold it. If there is a reckoning, I don't think it starts with the hyperscalers, whose balance sheets can absorb mistakes that would bankrupt almost anyone else. It starts further down: neoclouds, speculative developers, levered infrastructure vehicles and the private credit behind them, where the cost of capital collides first with utilization, construction delays and obsolescence.

THE ELECTRON

The technology can be right while the capital structure is wrong.


The binding constraint is shifting from GPUs to what surrounds them: power, interconnections, substations, permitted land and cooling. In September, Oracle sent a force-majeure notice over the 2.45-gigawatt Jupiter campus that Blue Owl's Stack Infrastructure is building for OpenAI, seeking the right to defer payments if power delays push back the planned 2028 opening. Oracle says the project remains on schedule, but the roughly \$18 billion of bank debt behind it was reported trading below 90 cents on the dollar. Nothing about AI demand changed that week. What changed was the market's view of who bears the cost when the electricity arrives late.

History is consistent on this point. Britain's railway investment peaked in 1847 at about 7% of GDP, by Andrew Odlyzko's estimate; shareholders were crushed and the network survived. After the telecom bust, a widely cited Merrill Lynch estimate put fiber capacity in use below 3%; the investors were wiped out, and the next owners built the modern internet on the same glass. And as Paul David showed, electrification took roughly four decades to lift factory productivity. Technology arrives on the vendor's timeline. Returns arrive on the customer's.

Here AI differs from fiber, in both directions. A frontier GPU bought today will still work in five years, but its residual value is far less certain than glass in the ground. Power infrastructure is the opposite. Interconnection rights, substations and powered, permitted buildings take years to secure and can serve several generations of chips.

Chips expire. Power lasts.

Approximate economic lives of the assets behind the AI buildout, in years.



Approximate, typical ranges for illustration. GPU lives from company filings; infrastructure lives are typical industry ranges and vary by asset.

THE FIBER OF THIS CYCLE ISN'T THE CHIP. IT'S THE ELECTRON.

A restructuring can destroy equity while improving the economics of the asset for its next owner. That is the difference between an investment failing and a technology failing.

WHAT I'M WATCHING, AND WHAT WOULD CHANGE MY MIND

- 1. Cloud growth against the 40% bar.** If combined cloud revenue growth at the five falls below 30% for two straight quarters before mid-2027, the gap in the model is widening, not closing.
- 2. The AI credit premium.** High-yield data-center debt traded about 220 basis points wider than the high-yield index at its December 2025 peak, by Penn Mutual's reading of Bloomberg data, and the gap has been rebuilding since July. A return above that level, held for a month, would mean stress is spreading beyond a few issuers.
- 3. Useful lives.** If Microsoft, Alphabet or Meta shortens server lives below five years, the industry is conceding the model's middle case.
- 4. Project finance.** A major data-center loan restructured, or trading below 80 cents on the dollar, would confirm that the levered middle breaks first.
- 5. Returns outside tech.** If McKinsey's next survey shows a majority of companies, up from 39%, reporting enterprise-level profit impact from AI, the demand side is catching up and I would expect the gap to close.

I'll publish a scorecard against these five in April 2027.

CONCLUSION

The recent volatility isn't the end of the AI trade. It may be the beginning of price discovery.

Two questions have been treated as one. Will AI transform the global economy?

Probably. Will everyone financing that transformation earn an acceptable return?

History says that is a much harder question, and the gap between those two answers is where fortunes are made and lost.

For three years, the question on earnings calls was *how much are you investing in AI?* The next one is *what return are you earning on it?*

THAT'S NOT THE END OF AI. THAT'S AI GROWING UP.

METHODOLOGY

Scope: Microsoft, Alphabet, Amazon, Meta and Oracle. Capex for 2023–24 is the sum of reported purchases of property and equipment (Microsoft and Oracle on fiscal years). For 2025–28 the model uses Goldman Sachs estimates (\$413B, \$800B, \$1.2T, \$1.4T), held flat at \$1.4T in 2029–30.

Asset mix: 60% servers and network equipment, 40% buildings and power. Microsoft says about two-thirds of recent capex went to short-lived assets; Alphabet says about 60%.

Depreciation: straight-line, no salvage value. Servers use a half-year convention over a 3-, 4-, 5- or 6-year life. Buildings and power use 25 years, starting the year after spending to reflect construction time.

Required revenue: annual depreciation plus 10% of average net invested capital (cumulative capex less accumulated depreciation). Power, staff, software and other operating costs are excluded, so the figure is a floor.

Not included: depreciation on assets bought before 2023, operating leases (large and growing), and revenue from conventional cloud or Meta's advertising that this capacity also serves.

Check: on six-year lives, the model gives about \$131B of 2026 depreciation from post-2023 assets, against about \$160B annualized reported by the five in mid-2026 including older assets.

SOURCES

[Goldman Sachs Research: Can AI Investment Drive S&P 500 Earnings Even Higher? \(Sept 2026\)](#)

[Microsoft FY2026 Q4 results \(July 29, 2026\)](#)

[Microsoft FY2023 Form 10-K \(server useful lives\)](#)

[Alphabet Q2 2026 results \(SEC exhibit 99.1\)](#)

[Alphabet Q2 2026 Form 10-Q \(cloud backlog\)](#)

[Amazon Q2 2026 results](#)

[Amazon 2025 Form 10-K \(server useful lives\)](#)

[Amazon Q2 2026 Form 10-Q \(backlog\)](#)

[Meta 2025 Form 10-K \(server useful lives\)](#)

[Oracle Q1 FY2027 results \(SEC exhibit 99.1\)](#)

[Motley Fool: Oracle's backlog and OpenAI concentration \(Mar 2026\)](#)

[Nvidia Q2 FY2027 results](#)

[Bloomberg: OpenAI's annualized revenue nears \\$70 billion, Axios says \(Sept 29, 2026\)](#)

[CNBC: Anthropic says annualized revenue climbed to \\$65 billion in July \(Aug 17, 2026\)](#)

[Fortune: Anthropic draft S-1 income statement \(Sept 29, 2026\)](#)

[CNBC: Michael Burry on hyperscaler depreciation \(Nov 11, 2025\)](#)

[CNBC: Nvidia pushes back on AI bubble charges \(Nov 25, 2025\)](#)

[Reuters: AI borrowers face tough sell in risky corners of U.S. credit market \(Sept 30, 2026\)](#)

[Reuters: SoftBank issues \\$11.1 billion in bonds in OpenAI financing push \(Sept 23, 2026\)](#)

[CoreWeave: pricing of senior notes \(2026\)](#)

[FRED: ICE BofA US High Yield Index option-adjusted spread](#)

[Penn Mutual Asset Management: AI infrastructure bonds, a full picture of spreads \(Aug 6, 2026\)](#)

[Bloomberg: Oracle cites force majeure to shield itself on big data center \(Sept 24, 2026\)](#)

[CNBC: Oracle data center force majeure \(Sept 24, 2026\)](#)

[McKinsey: The State of AI \(2025\)](#)

[Gartner newsroom \(Oct 27, 2025\)](#)

[Irish Times: Satya Nadella at Davos \(Jan 20, 2026\)](#)

[Fortune: Mark Zuckerberg on AI investment \(Sept 23, 2025\)](#)

[Andrew Odlyzko: Collective hallucinations and inefficient markets, the British Railway Mania of the 1840s](#)

[Light Reading: Fiber utilization figures challenged \(2002\)](#)

[Paul A. David, "The Dynamo and the Computer," American Economic Review, 1990](#)



ABOUT THE AUTHOR

Bryan Musto is a Chief Revenue Officer and go-to-market leader working at the intersection of financial services and enterprise technology. He currently serves as CRO of Northlight Solutions Group, and his background includes senior roles at Capgemini Invent, Vanguard, and BNP Paribas. He holds an MBA from the Wharton School of the University of Pennsylvania.