The AI Infrastructure Contraction

A Theory of How the Current Buildout Ends

William Foster Wood

Thesis

Artificial intelligence infrastructure — the data centers, chips, and financing structures built to power it — is following the same historical arc as every major general-purpose technology buildout before it: canals, railroads, electrification, telecommunications, and fiber-optic networks. In every prior case, the underlying technology proved to be exactly as transformative as its promoters claimed. And in every prior case, the capital structure built to fund its early expansion collapsed anyway, because capital was deployed faster than revenue could justify it, financed substantially by debt, and organized through interconnected balance sheets that turned a valuation correction into a systemic one.

AI is not exempt from this pattern. It is arguably the clearest instance of it in a century, because the financing structure underneath the current buildout has re-created, almost mechanically, the specific features that made the 1920s-30s collapse systemic rather than contained: opaque interconnection, reciprocal exposure between the same small set of players, and a growth narrative funded in part by capital that never left the circle of companies reporting it as revenue.

Part I — The Historical Pattern

The pattern repeats across every major infrastructure technology in modern economic history:

- Canals (1820s–1840s): Britain and the U.S. overbuilt canal networks on the assumption that demand would keep compounding. Railroads arrived mid-buildout and made much of the network obsolete before it was paid for. Several U.S. states defaulted on canal bonds.
- Railroads (1840s, 1870s, 1890s): Britain's 1840s "Railway Mania" wiped out investors when a large share of financed lines were never built. The U.S. repeated the pattern twice — the Panic of 1873 was triggered by railroad overbuilding and the collapse of Jay Cooke's financing house, and it happened again in the 1890s. The railroads themselves were not the failure. The debt-financed pace of construction, running ahead of freight and passenger revenue, was.
- Electrification and utility holding companies (1920s): Utility holding companies — Samuel Insull's empire being the textbook case — stacked debt on debt through pyramided corporate structures to build generation and distribution capacity years ahead of household demand. The pyramids collapsed in 1929–32 when the leverage unwound. The electric grid was the right investment. The capital structure funding it was not.
- Telegraph and telephone: overbuilt, competing wire networks in the same cities before consolidation absorbed the excess capacity.

- Automobiles (1900s–1920s): hundreds of manufacturers were funded on the assumption that each could scale; nearly all were wiped out or absorbed before the industry matured into a handful of survivors.
- Fiber-optic and telecom (1999–2001): networks were built on the assumption of relentless bandwidth growth, funded heavily by vendor financing — equipment makers lending customers the money to buy their equipment. When growth didn't arrive on schedule, the debt collapsed the sector even though the fiber itself became the backbone of the internet a decade later.

Two features are constant across all of these episodes. First, the technology's long-run value is eventually vindicated — canals, rail, the grid, and fiber all turned out to be at least as transformative as promised. Second, the capital deployed to build them out early is destroyed anyway, because it was committed on a timeline the underlying revenue could not match, financed by debt issued during the period when questioning the growth curve was professionally and socially costly. The companies doing the overbuilding are rarely the ones who profit most once the dust settles; whoever has the cleanest balance sheet when the capital gets repriced ends up buying the physical assets for cents on the dollar.

Part II — What Makes This Cycle Structurally Different, and Why It Isn't Safer

The hyperscalers financing the current AI buildout — Microsoft, Alphabet, Amazon, Meta, and Oracle — differ from the standalone railroads, utility pyramids, and telecoms of prior cycles in one respect: they are already enormously cash-generative on existing, profitable businesses unrelated to AI. That difference is real, but it changes the shape of the correction rather than preventing it. It means the capital destruction is less likely to produce outright bankruptcies of the kind seen in 1847, 1873, or 2001, and more likely to surface as a multi-year earnings and margin compression at a handful of companies, absorbed by balance sheets large enough to survive it.

It does not mean the capital isn't being destroyed, or that the free cash flow cushion is what it appears to be. That cushion is being consumed in real time by the buildout itself: hyperscaler capital expenditure is on trend to exceed hyperscaler cash inflows by the third quarter of 2026. The cash that is supposed to absorb the coming cost is being redirected into the very spending that creates that cost. And a meaningful share of current reported profitability is itself an artifact of the accounting gap described below — which means today's "awesome free cash flow" is not a defense against the correction. It is a preview of what gets restated once the accounting catches up to the economics.

Part III — The Accounting Mechanism: A Depreciation Bill Coming Due

The specific fuse is depreciation accounting. The five largest hyperscalers spent roughly $412 billion on AI-related capital expenditure in 2025, with 2026 spending estimated near $760 billion. Against that spend, they expect to recognize only about $211 billion in depreciation and amortization in 2026. The remaining roughly $549 billion sits capitalized on the balance sheet, not yet charged against earnings — deferred to future years under depreciation schedules of five to six years for servers and twenty-five to forty years for data center buildings.

Those schedules likely overstate the useful economic life of the hardware. Chips are obsoleted by the next generation, not worn out physically; critics, including investor Michael Burry, estimate the real economic life of AI hardware is closer to two to three years, which would mean current depreciation is understated by roughly $176 billion across 2026–2028 — flattering reported earnings in the near term at the cost of a larger, later correction. Amazon has already been forced to move in this direction: it revised its own useful-life assumptions and absorbed a $700 million hit to 2025 operating income as a result. Companies do not voluntarily take that kind of charge unless the operating data leaves them no real choice.

The market has not yet priced this. Full-year 2026 S&P 500 earnings growth estimates climbed from roughly 14 percent in February to over 23 percent by late summer — the opposite of the normal seasonal pattern, in which analysts trim estimates as the year progresses. That divergence indicates the depreciation wave has not yet been reflected in consensus numbers. When accounting schedules are forced toward economic reality — as Amazon's already have been, once — the earnings hit lands at the same companies that carry the largest weight in every major stock index.

Part IV — The Interconnection: Why This Time Resembles 1930, Not Just 2001

What made the 1929–32 collapse systemic rather than contained was not simply that a lot of capital had been deployed. It was that the capital structure was opaque and reciprocally interlinked. Insull's utility empire was a pyramid of holding companies, each layered on debt raised against the equity of the one beneath it, with cross-guarantees that made it impossible to isolate where one entity's risk ended and another's began. Banks had cross-lent and cross-invested with these structures and with each other, so a failure in one corner propagated through the entire credit system rather than staying contained to the company that failed.

The AI financing web being built in 2025–26 reproduces this structure with different names. Nvidia has invested tens of billions of dollars directly in OpenAI and other frontier labs. OpenAI has committed roughly $1.15 trillion across seven infrastructure vendors through 2035 — including $350 billion to Broadcom, $300 billion to Oracle, $250 billion to Microsoft, and $100 billion to Nvidia. Those vendors, in turn, use OpenAI's committed spending to justify further capital investment in the chips and data centers that Nvidia sells. Each leg of the loop books the same underlying dollars as new revenue or backlog. Analysts now estimate the scale of these round-tripping arrangements at more than $800 billion.

The mechanism that makes this dangerous is not the spending itself but the double-counting of demand it produces. When Nvidia invests in OpenAI, and OpenAI uses that capital to buy compute from Oracle, and Oracle uses that contracted revenue to justify buying more chips from Nvidia, the same pool of dollars is being counted as independent demand signals at three different points in the chain. Nvidia's own

CEO has effectively confirmed the structure while disputing the label, stating that his only regret is not having invested "more and sooner" in the companies buying his chips — an admission that Nvidia is functioning as financier to its own customer base, precisely the role banks and vendor-finance arms played in the utility and telecom buildouts that collapsed when revenue failed to arrive on the pyramid's schedule.

The exposure is not contained to equity markets. A significant share of the buildout is debt-financed rather than cash-financed — technology-sector debt issuance to fund AI infrastructure is estimated at roughly $1.5 trillion over 2026–2028 — which means bondholders and, by extension, the broader credit system now carry direct exposure to whether OpenAI's revenue materializes anywhere near its committed scale. OpenAI is currently projected to lose approximately $14 billion in 2026, nearly triple its 2025 loss, against a bet on $100 billion in revenue by 2029. Bain & Company estimates the AI ecosystem needs roughly $2 trillion in annual revenue by 2030 to justify current infrastructure spending, against a trajectory running about $800 billion short of that mark.

Part V — How the Correction Propagates

If the revenue gap does not close, the failure will not stay contained to OpenAI, or to any single company. Because the same dollars have been counted multiple times across the chain, a shortfall at the demand end runs back up through Oracle's contract backlog, into Nvidia's reported growth, and into the earnings of every hyperscaler whose capital expenditure was justified by that backlog. Because a growing share of that spending is debt-financed, the shortfall also runs into the credit markets that financed it — at a moment when long-end Treasury yields are already elevated and the Treasury's own efforts to hold down borrowing costs are struggling against the bond market. A stumble in AI infrastructure credit would not be an isolated equity-market event; it would land on top of an already fragile rate environment, compounding rather than offsetting it.

This is the structural echo of 1930 rather than simply of 2001. The dot-com telecom collapse destroyed capital broadly but the failures were largely idiosyncratic and sequential — one overleveraged telecom at a time. A collapse built on an $800 billion web of circular financing among a handful of interlocked companies, a quarter of whose combined value Peter Wildeford has estimated is effectively a leveraged bet on AGI arriving on schedule, has the shape of a 1930s pyramid: opaque, reciprocal, and capable of propagating through the system all at once rather than company by company.

Part VI — What Doesn't Break, and Who Profits from the Correction

Consistent with every prior cycle, the technology itself is not the thing at risk. AI's long-run utility is not in serious question any more than the electric grid's was in 1932 or fiber's was in 2002. What is at risk is the capital structure built to pull years of future revenue forward into the present — and the companies most exposed to that structure are not necessarily the ones that will dominate the technology once the correction has run its course.

In every previous cycle, the entities that emerged strongest were not the original overbuilders but whoever held the cleanest balance sheet when the capital got repriced, able to acquire distressed physical assets — track, generation capacity, fiber routes — for a fraction of their original cost. The open question for this cycle is which of the current players is positioned to be that buyer: the hyperscaler with the least circular financing exposure and the most durable cash flow from businesses outside AI, rather than the one with the largest headline commitments today.

Summary

The AI buildout is following the same arc as every general-purpose infrastructure technology since canals: real long-run value, funded by capital deployed faster than revenue can justify, financed substantially by debt, and organized this time through a circular financing web among a small number of interlocked companies that mirrors the opacity and reciprocal exposure of the 1920s utility pyramids more closely than it resembles the more contained, sequential telecom failures of 2001. The depreciation accounting gap — roughly $549 billion in 2026 capital spending not yet charged against earnings — is the mechanical fuse. The circular financing web — upward of $800 billion in round-tripped commitments among Nvidia, OpenAI, Oracle, Microsoft, and a handful of others — is the transmission mechanism that would carry a shortfall in AI revenue across the system rather than containing it to one company. The debt financing underneath both is what connects the outcome to the broader credit markets, at a moment when those markets are already under separate strain. None of this requires AI to fail as a technology. It only requires the revenue to arrive later, or smaller, than the financing assumed — which is exactly what happened to canals, railroads, utilities, and fiber before it.