

# The AI Debt Illusion: Inside Big Tech’s Trillion-Dollar Shadow Build-out

The tech giants have quietly pivoted from cash-rich self funding to a complex web of shadow debt. Here is how the the off-balance-sheet vehicles are funding it.

Valentina Dardano

June 27, 2026

---

## Contents

1	The End of the Self-Funding Era	2
2	The SPV Loophole: Control Without Consolidation	4
3	The Duration Mismatch	4
4	The Web of Counterparty Risk	5
5	When Shadow Debt Reprices	5
6	A Race Between Three Clocks	6

---

## Introduction

For the better part of a decade, the defining feature of Big Tech was an invincible balance sheet. The self-funding model allowed the world’s largest tech firms to build sprawling empires entirely out of their own operating cash flow, leaving plenty of room for buybacks and a famously fattening cash pile.

To fund the artificial intelligence supercycle, Amazon, Alphabet, Microsoft, Meta, and Oracle are projected to deploy an estimated \$725 billion in capital expenditures by 2026 [2, 8]. If the current trajectory holds, that number will cross \$1 trillion by 2027 [10]. For context, this means a handful of tech firms are executing a capital expenditure cycle roughly equivalent to the entire GDP of a G20 nation.

But the most interesting part of this buildout isn’t the size of the bill; it’s how it is being

paid for. In just a few years, the group has transitioned from self-funding, to borrowing, to borrowing in ways the public markets cannot easily see.




Figure 1: Combined Capex of the five biggest AI spenders (US\$ bn). Source: Company guidance; analyst estimates [2, 8, 10].

## 1. The End of the Self-Funding Era

Follow the cash and the shift makes sense. For most of the past decade, capex ran near 40% of these companies' operating cash flow. That left plenty of room for buybacks, dividends and a fattening cash pile.

In 2026, capex net of dividends and buybacks is running at about 94% of operating cash flow across the group, and on some measures north of 100% [2]. The self-funding model that defined Big Tech's balance sheets has run out of room. Historically, Big Tech offered a reliable free cash flow (FCF) yield that acted as a bond-like proxy for investors. By pushing CapEx past 100% of operating cash flow, that yield is evaporating, forcing the market to re-evaluate these stocks from cash cows back to highly speculative growth engines.




Figure 2: Aggregate capex as a share of operating cash flow. Source: Bank of America; UBS [2].

However, there is a wide range of variation within this overall metric. The same dollar

amount invested in technology spend produces significantly varying degrees of strain based upon how much cash each individual business already generates as well as the degree of leverage each company is willing to absorb.

The table below provides a view of the data when viewed across the rows and examines a single dimension of a shift from the one company that can still easily self-finance to the one that is pushing the model the farthest.

Table 1: Big Tech Financing Shift [3, 8]

Company	2026 capex	Free cash flow	Leverage signal	Primary financing route
<b>Microsoft</b>	\$~190bn	Positive; +33% YoY last quarter	Lowest in the group	Internal cash, finance leases, fund-level vehicle
<b>Amazon</b>	\$~200bn	Turning negative	Rising	Operating cash plus record bond issuance
<b>Alphabet</b>	\$~185bn	Falling sharply	\$40.7bn of off-balance-sheet vehicle commitments	Cash, bonds, and new Blackstone joint venture
<b>Meta</b>	\$125–145bn	Approaching zero or negative	Spreads widening on rising debt	SPV structures plus corporate bonds
<b>Oracle</b>	\$~35bn	Negative	Total debt over \$130bn; order book \$~455bn (~7x sales)	On-balance-sheet debt and SPVs

Two companies sit at the edges of the range. Microsoft still turns its buildout into rising free cash flow. That figure was up roughly a third YoY last quarter even as Microsoft carried the biggest capex line of the five [8], and it funds itself with internal cash, finance leases and an equity stake in a fund-level partnership rather than issuing debt directly. Oracle sits at the other end. Its free cash flow has gone negative, its total debt has climbed past \$130 billion (the heaviest load of any of these firms), and its order book of contracted future revenue has ballooned to about \$455 billion, roughly seven times annual sales [3].

Amazon, Alphabet and Meta land in between, free cash flow squeezing down toward zero or below it. Alphabet used to be the poster child of the cash-funded model. Its latest 10-Q discloses \$40.7 billion of future funding commitments to off-balance-sheet vehicles [8], including a roughly \$30 billion equity derivative that works as AI-infrastructure financing, plus a new joint venture with Blackstone that shifts part of its capex onto private capital. A year ago you could draw a clean line between the firms that paid cash and the firms that moved spending off-book. Nowadays that line has mostly dissolved.

## 2. The SPV Loophole: Control Without Consolidation

The real change appears in the structures beneath the bond issuance, and they're worth slowing down on, because they're built to be hard to see. Meta's Hyperion facility is the template transaction. Instead of raising the money itself, Meta and a consortium capitalised a dedicated entity that took on about \$27 billion of senior debt, with Pimco as anchor lender and roughly \$2.5 billion of equity from Blue Owl [10].

Meta holds only a minority stake. It's enough to control and run the site, but deliberately below the threshold that would drag the entity's debt onto Meta's own books, and the vehicle is ring-fenced so its liabilities stay separate from the parent [1]. By carefully structuring these as Variable Interest Entities (VIEs) under US GAAP, the parent company maintains operational control but stays just below the accounting threshold that requires consolidating the special purpose vehicle's liabilities onto the master balance sheet. As a result, the debt shows up on no balance sheet a Meta shareholder ever reads.

In its place sits a long-dated lease. Meta is on the hook for years of payments on the facility no matter what the AI economics do. That's debt service in all but name, booked as an operating commitment. Count these leases for what they are and the group's real leverage runs well above the headline numbers.

Others have copied the Hyperion blueprint. More than \$120 billion of data-centre spending has moved off balance sheets in under two years [9], and Moody's puts hyperscalers' lease commitments sitting outside reported debt at around \$662 billion [7]. Reported leverage now captures only a shrinking slice of what these companies have actually promised to pay.

## 3. The Duration Mismatch

Forget the accounting for a second and look at what the money actually buys. Most of it goes on hardware. McKinsey puts hardware at roughly 60% of data-centre investment [6], and that hardware ages fast. The financing runs the opposite way. Bonds at five years and up, leases and separate-vehicle structures longer than that, and real-estate-style debt under some sites written over fifteen or twenty years.

This is a classic duration mismatch: short-lived collateral paid for with long-dated obligations. The same fault line ran under the late-1990s telecom buildout and the 2008 funding markets. When the asset wears out faster than the liability amortises, the borrower hits a refinancing wall against collateral worth a fraction of what it cost. The lender, meanwhile, is left with recovery risk and no deep secondary market for used accelerators.

Depreciation makes it worse. You spread an asset's cost over an assumed useful life, and the operators book servers at roughly five to six years. If the real working life is shorter, depreciation is understated and today's operating income looks better than it is. While depreciation is a non-cash accounting charge, a shortened hardware lifespan turns what was supposed to be one-off growth CapEx into recurring maintenance CapEx. If a server becomes obsolete before the bond financing it matures, the operator must bleed real cash to replace it just to keep existing revenues flat.

Moody's puts the gap between booked and actual depletion at about \$176 billion across 2026 to 2028 [7]. Epoch AI runs the math on a single one-gigawatt facility: cut the equipment life from five years to three and annualised cost of ownership jumps from about \$8.5 billion to \$12 billion, a 41% increase [5]. The operators can't even agree with each other. Amazon has shortened parts of its server fleet, citing how fast the gear goes obsolete, while Meta has lengthened most of its servers. Which tells you the reported figures rest on estimates a careful reader should poke at.

Put both halves together and the return hurdle is steep. Hardware that has to be written off in three or four years also has to earn back its purchase price plus the cost of the debt and

leases behind it, all inside that same short window, or it destroys value. Set against that bar, the revenue is still thin. Industry-wide AI revenues came in near \$60 billion in 2025 while capital spending approached \$400 billion [9], and Sequoia's estimate that the buildout eventually has to clear roughly \$600 billion a year shows how far there is to go [4]. The depreciation argument and the revenue gap are really the same question asked twice: can these assets throw off cash faster than they wear out?

## 4. The Web of Counterparty Risk

If the debt is leaving the operators' books and the collateral might not last, the risk has to land somewhere. Increasingly it lands away from the names on the buildings. The lenders are mostly non-banks. Blackstone, Blue Owl, Apollo, Pimco and BlackRock originate most of this paper, and Morgan Stanley expects another \$800 billion of private-credit data-centre financing over the next two years [9].

Behind those funds sit the pension and insurance balance sheets whose money they're putting to work. And the same unit of compute often gets counted more than once: as chip demand, as operator revenue, and as loan collateral, with chip vendors sometimes financing the buyers. That double counting makes aggregate demand look sturdier than the cash flows underneath it.

This creates a fragile structure of intertwined counterparty risk, where the same unit of compute is booked as revenue by the operator, collateral by the private credit fund, and a yield-generating asset by the pension fund.

The exposure is concentrated, too. Oracle's standing rests heavily on a single customer: a \$300 billion, five-year contract to supply OpenAI underpins much of its \$455 billion order book, so an investment-grade balance sheet now leans on the solvency of one venture-funded counterparty. That concentration has already drawn legal fire. In January 2026, bondholders led by the Ohio Carpenters' Pension Plan filed a proposed class action alleging Oracle concealed how much it would need to borrow; after raising \$18 billion in September 2025, the company returned weeks later for roughly \$38 billion more to fund OpenAI data centres, and the earlier bonds fell to trade like lower-rated debt [3]. The largest operators can absorb a shock the smaller ones cannot, but the structures, and the institutions holding the paper, are where fragility is currently concentrating.

## 5. When Shadow Debt Reprices

For all the financial engineering, the obligations don't disappear. They just move to investors who haven't yet marked them at a stress price. Morgan Stanley's own behaviour is the tell. The same bank that built Meta's Hyperion vehicle has reportedly looked at shedding parts of its data-centre exposure through significant risk transfers, the arranger hunting for an exit even while it's still selling the deals [10].

The arithmetic explains the instinct. Of the roughly \$3 trillion these companies plan to invest through 2028, about half (some \$1.5 trillion) has to come from debt, and the investment-grade market has never had to swallow a single theme on that scale [10].

The early evidence says the market reprices fast once the full obligation comes into view. Oracle's first bonds were treated as solidly investment-grade right up until the surprise \$38 billion follow-on. Within weeks, their credit spreads blew out, pricing the debt closer to high-yield territory despite the investment-grade label [3]. Nothing was missed, no payment skipped. What moved them was the disclosure of how much more Oracle owed. If that pattern repeats across the group as off-balance sheet commitments come to light, the cost of the next trillion in financing climbs at exactly the moment the sector can least afford it.

The risk has also been parked with the holders least able to move quickly. Private-credit

funds, whose AI lending went from near zero to more than \$200 billion in a few years, sit between the buildout and the pension and insurance money behind them. The Federal Reserve Bank of Chicago found banks' direct exposure modest, around 0.8% of assets, but flagged a bigger indirect exposure through their lending to those funds [9]. What looks contained at any one institution is, in the aggregate, piled into a handful of lenders financing the same assets, for the same few tenants, against the same fast depreciating chips.

## 6. A Race Between Three Clocks

Bubble might not be the appropriate word for this. These are cash-generative businesses putting up infrastructure that paying customers already use. AWS is running near a \$150 billion annual revenue rate, Google Cloud around \$80 billion and growing over 60%, Azure's AI line at a \$37 billion run rate [8]. We know we have actual demand and growth; what is being questioned is how fast and how much, with revenue still chasing commitments that are now mostly locked in.

What is actually changing is where the centre of gravity for analysis is going. A year ago, Big Tech split cleanly: firms that paid cash, firms that borrowed. That split is gone. Outside Microsoft, the group is converging on financing built to stay clear of headline leverage. If you're underwriting these companies, the numbers that matter aren't on the face of the balance sheet anymore. They're in the lease footnotes, the commitments to unconsolidated Variable Interest Entities, and the depreciation assumptions. The reported debt figure tells you less than it ever has.

It appears that the biggest risk associated with these investments is related to timing. To understand the exact failure mode of this buildout, consider these firms as operating three clocks simultaneously.

The first clock is the hardware clock; the amount of time during which the economic value from a chip deteriorates, approximately three to four years. The second is the revenue clock; how quickly the new AI workload becomes a stable, recurring source of cash flow. The third is the refinancing clock; when the bonds, loans, and leases expire and must be rolled over at prevailing market rates.

This is a race between Return on Invested Capital (ROIC) and the Weighted Average Cost of Capital (WACC). By pushing debt into off-balance sheet vehicles, these firms are extending the refinancing clock and artificially suppressing their WACC. They are taking advantage of the fact that maximizing leverage in shadow vehicles keeps the corporate-level cost of capital optically low.

However, financial engineering only solves for the cost of capital. The hardware and revenue clocks ultimately dictate the return. If the hardware clock expires before the debt amortizes, the invested capital must be aggressively written down and physically replaced, which crushes net operating profit. No amount of off-balance sheet maneuvering can save a negative ROIC if the AI revenue fails to clear the accelerated depreciation hurdle. As long as the revenue clock runs faster than the other two clocks, ROIC stays above WACC and the buildout succeeds.

If this system breaks, the collapse will not start with the tech giants. Their balance sheets are specifically engineered to hide the risk. Instead, the first cracks will appear at the periphery: a major AI tenant missing a lease payment, a private credit fund writing down a data-center loan, or a rating agency downgrading a highly concentrated player like Oracle. From there, financial contagion will spread through shared lenders and insurers much faster than Big Tech's reported leverage would ever suggest.

Investors must therefore look past the optical comfort of massive cash balances and focus entirely on how the market prices this shadow debt. Oracle has already shown how fast that pricing can snap. The day AI credit spreads widen across the entire group, the footnotes will become the story. At that point, financing engineered for an AI supercycle becomes a

conventional distressed-debt crisis, and only the cash flow matters.

## References

- [1] Bank for International Settlements. Mechanics of off-balance-sheet data-centre financing structures. *BIS Quarterly Review*, March 2026. URL <https://www.bis.org/publ/quarterly.htm>.
- [2] Bank of America and UBS. Aggregate hyperscaler capital expenditure estimates: 2024–2027. Institutional Research Reports, 2026. Company guidance and analyst estimates as of Q1 2026.
- [3] J. Bursztynsky. Oracle sued by bondholders over losses tied to AI buildout. CNBC, January 14 2026.
- [4] D. Cahn. AI’s \$600b question. Sequoia Capital Insights, 2025.
- [5] Epoch AI. Hardware depreciation and annualized cost of ownership in 1-gigawatt AI facilities. Institutional Research, 2026.
- [6] McKinsey & Company. Data-centre investment hardware allocation estimates. Institutional Research, 2026.
- [7] Moody’s Ratings. Hyperscaler lease commitments and asset depletion across 2026–2028. Institutional Research, 2026.
- [8] J. Novet and A. Capoot. Microsoft (MSFT) q3 earnings report 2026. CNBC, 2026. URL <https://www.cnbc.com/2026/04/29/microsoft-msft-q3-earnings-report-2026.html>. Includes Alphabet Form 10-Q disclosures and aggregate hyperscaler capital expenditure guidance.
- [9] Quinn Emanuel Urquhart & Sullivan. Emerging litigation risks in financing the AI data centers boom. Client Alert, March 16 2026. URL <https://www.quinnemanuel.com/the-firm/publications/client-alert-emerging-litigation-risks-in-financing-ai-data-centers-boom/>.
- [10] H. Son. Morgan Stanley significant risk transfer loans data center AI infrastructure exposure. *Fortune*, December 4 2025.