

Hidden Leverage in the AI Infrastructure Boom: A Forensic Investigation and Historical Comparison

Kilambi Research

arvin@kilambi.com

October 2026

Abstract. This paper conducts a forensic investigation of how the global AI infrastructure buildout is financed, and whether its financing structures resemble six historical episodes of hidden leverage: Enron, WorldCom, Tyco, the 2007–2009 global financial crisis, Lehman Brothers’ Repo 105, and Greensill Capital. We reconstruct obligation-level exposure maps for Microsoft, Amazon, Alphabet, Meta, Oracle, OpenAI, Anthropic, CoreWeave, Nvidia, Broadcom, and xAI from a five-vintage dataset of 116 obligations across 84 entities with primary committed exposure of \$4,315.4 billion; catalog seven financing mechanisms (SPVs, sale-leasebacks, residual-value guarantees, operating and finance leases, purchase commitments, supplier-arranged customer financing) with their treatment under US GAAP and IFRS; and run nine-scenario coverage stress tests over 2027–2030 with a collateral-depreciation model for GPU-secured debt. We find no evidence of fraud: the structures are disclosed in filings, and the post-2002 accounting regime (ASC 810, ASC 842) blocks the specific tricks used in the historical cases. The genuine parallels are structural, not criminal. The funding architecture of the buildout resembles pre-crisis shadow banking (private-credit intermediation outside the bank perimeter, SPV layering, maturity transformation, originate-to-distribute incentives), and Greensill is the closest single-firm analogue (private-credit intermediation of future economic claims, guarantee dependence, arranger opacity). The solvency core holds: Microsoft, Amazon, and Alphabet never breach 1.0x fixed-obligation coverage in any scenario, including the severe case, and Meta breaches only in 2027 of the severe case (0.71x) against a \$90.3 billion cash buffer. The periphery is fragile: CoreWeave breaches from 2027 in every case, Oracle compresses to 1.33x base coverage by FY2029 and breaches in the severe case, Meta faces a one-year 2027 squeeze, and GPU-secured facilities go underwater on hardware value within two years under every depreciation assumption. The verdict is a solvent core funding a fragile periphery through increasingly circular structures: a credit cycle with genuine systemic-fragility elements in its transmission architecture, not a fraud story and not a clean bill of health.

Keywords: artificial intelligence; infrastructure finance; off-balance-sheet leverage; special-purpose vehicles; private credit; financial crises; accounting standards; stress testing; Enron; shadow banking

JEL Classification: G21, G32, G33, M41, M48, E44

Non-technical summary

Companies are spending unprecedented sums on the data centers and chips behind artificial intelligence, and much of the financing sits in structures that do not look like ordinary debt. This paper asks a plain question: how much fixed financial obligation have the builders taken on once every lease, purchase commitment, guarantee, and special-purpose vehicle is counted, and does any of it resemble the hidden leverage behind past financial scandals?

We rebuilt the obligation stacks of eleven builders and funders from their filings and a dataset of 116 individual obligations. The headline finding is reassuring on the core question of fraud: we found no hidden losses, no fabricated earnings, and no executive self-dealing. Every large structure we examined is disclosed somewhere in a filing, and the accounting rules written after Enron and the 2008 crisis specifically block the old tricks.

What is not reassuring is the shape of the funding. A growing share of the buildout is financed outside the regulated banking system by private-credit funds and by the chipmakers themselves, who are arranging loans so their customers can buy their chips. The weakest borrowers depend entirely on capital markets staying open: one large operator cannot cover its 2027 fixed payments from its own cash flow in any scenario we ran, and Oracle’s cushion thins to the edge under stress. Loans secured by the chips themselves are underwritten on contracted payments, not on the hardware, because the hardware loses most of its value within two years under every assumption we tested.

The historical comparison lands in the middle. This is not Enron and not 2008: the assets are real infrastructure with contracted customers, not fraudulent loans. But the plumbing, private credit, layered vehicles, vendor-arranged financing, has genuine echoes of both. The paper ends with eight specific, observable warning signs to monitor.

Contents

1	Introduction	6
2	Institutional background	8
2.1	The buildout: who spends, who operates, who funds	8
2.2	The seven financing mechanisms, in plain language	9
2.3	What “off-balance-sheet” means after ASC 842	11
3	Related literature	12
4	Data	14
4.1	The obligation-level dataset	14
4.2	SEC filings and company disclosures	14
4.3	Banking and market data	16
4.4	Historical primary sources	16
4.5	Evidence grading	17

5	Methodology	17
5.1	The five-tier leverage framework	17
5.2	Counting rules	18
5.3	The coverage model	18
5.4	The collateral-depreciation model	19
5.5	The historical-comparison rubric	19
6	Results	21
6.1	Company-level exposure maps	21
6.2	Financing mechanisms and accounting treatment	30
6.3	Historical episodes: forensic summaries	33
6.4	The resemblance matrix	36
6.5	Stress tests	42
6.6	Propagation and catalysts	46
7	Limitations and robustness	52
7.1	Falsifiability	52
7.2	Null findings	54
7.3	Data gaps	54
7.4	Robustness	54
8	Conclusion	55
A	Counting rules and tier definitions	57
B	Data appendix: series definitions and source table	58
C	Additional company tables	59
D	AI assistance disclosure	59

List of Tables

1	The October 2026 chip-financing wave: chronology.	6
2	2026 capex guidance versus operating cash flow (\$B).	9
3	The honest taxonomy: four buckets of fixed exposure, with cohort examples.	12
4	Summary statistics: the obligation dataset (vintage 2026-10-11).	15
5	Filing vintages and operating-cash-flow bases by company.	15
6	Market stress marks, October 2026.	16
7	Primary sources for the six historical cases.	17
8	Counting rules (summary).	18
9	Coverage-model assumptions (A1–A10): justification and sensitivity.	20
10	Collateral-model assumptions (C1–C5): justification and sensitivity.	21
11	Company tier ladders: from GAAP debt to economic fixed obligations (\$B).	22
12	Tier walk, Microsoft: from GAAP debt to economic fixed obligations (\$M).	23
13	Flagship financing structures of the AI buildout.	27
14	Counterparty concentrations by company.	28
15	Liquidity buffers against the 2027 fixed-obligation burden (\$B).	29
16	Financing mechanisms and their treatment under US GAAP and IFRS.	31
17	Disclosure quality by mechanism: what outsiders can and cannot see.	32
18	The resemblance matrix: six historical cases against the AI financing evidence (main results table).	37
19	Warning signs then and now: historical signals mapped to today’s observables.	41
20	Fixed-obligation coverage ratios (OCF / annual burden), 2027–2030.	42
21	2027 base-case burden composition (\$M).	42
22	The nine stress scenarios: definitions.	44
23	Material exposures not in the coverage model.	44
24	Scenario-breach summary: all nine scenarios, six borrowers.	45
25	Collateral LTV paths: loan balance / hardware value by GPU economic life.	46
26	Propagation severity matrix: five shock origins by eight affected nodes.	47
27	Highest-severity propagation cells: shock, channel, and basis.	48
28	Eight observable catalysts for forced deleveraging or contagion.	49
29	Falsifiers: the evidence that would overturn each main conclusion.	53
30	Data gaps: where each binds and the remedy.	55
31	Sensitivity of the stress verdicts to key assumptions.	55
32	Maturity schedule inputs used in the coverage model (\$M).	59

List of Figures

1	Tier ladders: GAAP recognized debt (Tier A) versus economic fixed obligations (Tier E), by company. The gap between the bars is the disclosed-but-unrecognized commitment stack. Labels show the Tier E / Tier A multiple.	22
2	All-in fixed-obligation burden: Tier E / operating cash flow, by company. Oracle at 13.9x (14.8x with the \$288B uncommenced update) is the cohort's highest by a wide margin; CoreWeave at 11.9x is second.	28
3	Debt-maturity wall: scheduled debt principal by year, where the monitor carries a schedule. Microsoft, Meta, Oracle, Broadcom, and CoreWeave shown; Amazon, Alphabet, and Nvidia have no debt-maturity schedule in the monitor (gap).	29
4	Historical timeline, 1999–2026: the six comparison episodes against AI buildout financing milestones. Enron (2001), WorldCom/Tyco (2002), GFC (2007–2009), Lehman Repo 105 (2008), Greensill (2021), and the AI chip-financing wave (2025–2026).	33
5	Resemblance heatmap: the six historical cases graded Strong / Partial / Weak / Misleading against the AI financing evidence. No case grades Strong; Tyco is a null finding.	37
6	Coverage waterfall, 2027–2030: Oracle, CoreWeave, and Meta fixed-obligation coverage under base, adverse, and severe cases. The 1.0x line is the breach threshold. Microsoft, Amazon, and Alphabet (never below 1.27x even in the severe case) are omitted for scale.	43
7	Collateral LTV erosion curves: F1 (CoreWeave-type DDTL) and F2 (Oracle construction-debt template) under 2–5 year GPU economic lives. The 1.0x line marks underwater on hardware value.	47
8	Propagation network: shock origins (left) to affected nodes (right); line weight encodes severity. The heaviest edges run neocloud default → private credit, private-credit stress → operators, and bank underwriting loss → bond markets.	48

1 Introduction

The question this paper answers is narrow: how much economic leverage sits outside reported debt in the AI infrastructure buildout, and does it resemble the hidden leverage behind six historical episodes of concealed risk? The answer is a measured one. The builders’ recognized debt is a small fraction of their fixed obligations (Tier E economic exposure runs 2.5x to 17x reported debt across the rated names), the structures financing the gap are disclosed but sit in footnotes and off-balance-sheet vehicles that credit metrics built on debt miss, and the closest historical parallels are structural rather than criminal: the funding architecture of pre-crisis shadow banking and the private-credit mechanics of Greensill, not the fraud of Enron, WorldCom, or Tyco.

Three facts make the question urgent now. First, scale. Our obligation dataset carries primary committed exposure of \$4,315.4 billion (OBSERVED, A) across 116 obligations and 84 entities as of the October 11, 2026 vintage, and the single largest disclosed-but-unrecognized numbers in the cohort are extraordinary by any historical standard: \$329.1 billion of uncommenced leases at Microsoft (OBSERVED, A), a \$707 billion fixed-or-guaranteed commitment portfolio at Alphabet (OBSERVED, A), and \$260–288 billion of uncommenced data-center leases at Oracle (OBSERVED, A)/(OBSERVED, B). Second, market stress is already visible. Oracle’s five-year CDS printed a record 261bp in October 2026 (OBSERVED, B), Meta’s \$27.3 billion Beignet/Hyperion SPV notes traded down to about 91 cents (OBSERVED, B), Oracle’s 2055 bonds traded near 77 cents (OBSERVED, B), and JPMorgan and MUFG spent more than six months unable to distribute \$38 billion of Oracle-leased construction debt (OBSERVED, B). Third, the financing frontier is moving fast: in October 2026 alone, the vendor-arranged “chip-financing wave” produced Broadcom’s \$60 billion Anthropic package (a \$42 billion Class A bank syndication plus an \$18 billion junior tranche), Broadcom arranging more than \$50 billion of private financing for OpenAI’s co-developed chips, Oracle negotiating a chip SPV with Apollo and Goldman Sachs, and Nvidia backstopping a \$105 billion facility for SB Energy (OBSERVED, B)/(OBSERVED, C).

Table 1. The October 2026 chip-financing wave: chronology.

Date	Event	Grade	
Oct 2	Bloomberg reports Broadcom’s \$60B Anthropic debt package (\$42B Class A + \$18B junior)	(OBSERVED, B)/(OBSERVED, C)	
Oct 5	\$42B Class A bank syndication begins (BofA/Citi/Morgan Stanley); package status update	(OBSERVED, B)/(OBSERVED, C)	
Oct 5	TeraWulf 8-K doubles Muskie power to 1 GW on expected AI demand	(OBSERVED, B)	
Oct 7	WSJ reports Broadcom arranging >\$50B “Nexus” financing for OpenAI; Oracle chip SPV talks with Apollo/Goldman	(OBSERVED, B)/(OBSERVED, C)	Source: AI
Oct 8	FT reports OpenAI ~\$50B annualized revenue run rate (investor documents); the \$68–70B “gross-up” figure debunked	(OBSERVED, B)	
Oct 9	Goldman prints Oracle 5Y CDS at record 261bp, Broadcom at 136bp; Anthropic S-1 still not filed (direct EDGAR check)	(OBSERVED, B)	
Oct 10	FT reports banks exploring SRTs on AI debt	(OBSERVED, B)	
Oct 14	Anthropic investor day (scheduled)	n.a.	

leverage monitor, vintage 2026-10-11. The wave is the empirical basis for the supplier-arranged financing mechanism (Section 2.2) and catalysts 1 and 8 (Section 6.6). Table 1 gives the day-by-day chronology.

Public discussion of these facts has reached for the Enron and 2008 analogies. Both are wrong in

the way that matters and right in a way that is worth keeping. They are wrong because there is no evidence of hidden losses, fabricated earnings, or criminal intent anywhere in the AI record, and because the assets are productive infrastructure with contracted cash flows rather than fraudulent consumer credit. They are worth keeping because the funding architecture genuinely resembles pre-crisis shadow banking (private-credit intermediation outside the bank perimeter, SPV layering, maturity transformation in construction finance, originate-to-distribute incentives in vendor-arranged financing) and because the Greensill mechanics (non-bank intermediation, financing of not-yet-existing economic claims, dependence on guarantees, arranger opacity, single-name concentration) are live in today's structures.

This paper makes five contributions. First, it reconstructs company-level exposure maps for eleven builders and funders (Microsoft, Amazon, Alphabet, Meta, Oracle, OpenAI, Anthropic, CoreWeave, Nvidia, Broadcom, xAI) on a consistent five-tier leverage framework, with every figure labeled by basis (recognized carrying, undiscounted notional, contingent maximum, off-balance-sheet) and evidence-graded. Second, it catalogs the seven financing mechanisms in use and renders an accounting verdict on each under US GAAP and IFRS, identifying exactly where the standards permit opacity. Third, it runs quantitative stress tests (nine scenarios over 2027–2030, a collateral-depreciation model for GPU-secured facilities, a propagation matrix, and eight observable catalysts) to locate the solvency boundary. Fourth, it compares the findings mechanism-by-mechanism against six historical episodes, grading resemblance as Strong, Partial, Weak, or Misleading with a three-part verdict (genuine resemblance, misleading aspect, accounting verdict) for each. Fifth, it states the falsifiers: the specific evidence that would prove the main conclusions wrong (Section 7.1).

The results preview is as follows. On forensics: Tier E / Tier A multiples run 17x (Microsoft), 14.4x (Nvidia), 9.1x (Alphabet), 8.9x (Meta), 3.9x (Amazon), 3.7x (Broadcom), 3.6x (Oracle), and 2.5x (CoreWeave) (OBSERVED, A)/(RECONSTRUCTED, B); all-in fixed-obligation years (Tier E / operating cash flow) run 13.9–14.8x (Oracle), 11.9x (CoreWeave), 8.0x (Broadcom), 5.68x (Meta), 4.92x (Alphabet), 4.68x (Nvidia), 3.75x (Microsoft), and 3.16x (Amazon) (OBSERVED, A)/(RECONSTRUCTED, B). On accounting: the flagship structures are compliant as far as the evidence shows; ASC 810 (VIE consolidation) and ASC 842 (lease recognition) block the specific devices used by Enron and the pre-crisis SIV system; the live opacity is at classification boundaries (lease versus service commitment, who consolidates the SPV) and in the private-credit disclosure vacuum. On stress: Microsoft, Amazon, and Alphabet never breach 1.0x fixed-obligation coverage in any scenario including the severe case (Modelled); Meta breaches only in 2027 of the severe case (0.71x); CoreWeave breaches from 2027 in every case (base 2027: 0.66x); Oracle compresses to 1.33x base coverage by FY2029 and breaches in the severe case (0.66–0.87x); Meta faces a one-year 2027 squeeze (severe 0.71x) with a \$90.3 billion cash buffer against it. GPU-secured debt goes underwater on hardware value within one to two years under every depreciation assumption from two to five years. On history: Enron Partial, WorldCom Weak, Tyco Misleading (a null finding), GFC Partial (the strongest structural parallel), Lehman Repo 105 Partial, Greensill Partial (the closest private-credit parallel).

What this paper does not claim matters as much as what it does. It does not allege fraud: no structure examined here is presented as non-compliant, and the accounting verdicts are evidence-led throughout. It does not forecast default: coverage below 1.0x means fixed obligations exceed operating cash flow under the stated assumptions, not that a borrower fails; borrowers cut capex, sell assets, raise equity, or renegotiate long before a missed payment. Tier E is a ceiling of fixed exposure, not an expected loss, and all-in fixed-obligation years are a burden scale, not a default clock. The resemblance grades compare mechanisms, not guilt.

The rest of the paper proceeds as follows: Section 2 describes the institutional setting, the seven financing mechanisms, and what “off-balance-sheet” means after ASC 842; Section 3 reviews the related literature; Section 4 describes the data and the evidence-grading standard; Section 5 sets out the methodology; Section 6 presents the results; Section 7 states the limitations, null findings, and falsifiers; Section 8 concludes. Appendices give the counting rules, the data appendix, and the AI assistance disclosure.

2 Institutional background

2.1 The buildout: who spends, who operates, who funds

The bottom line of this section is that AI infrastructure finance is a three-layer system: hyperscalers and AI labs commit the demand, neocloud operators build and operate much of the physical capacity, and banks, private-credit funds, vendors, insurers, and bond markets supply the capital, increasingly through vehicles that sit between the layers. Understanding any single structure requires knowing which layer bears the economics and which layer carries the accounting.

The demand layer is the hyperscalers (Microsoft, Amazon, Alphabet, Meta, Oracle) and the frontier AI labs (OpenAI, Anthropic, xAI). The hyperscalers’ 2026 capex guidance runs \$115–220 billion per year (OBSERVED, B), and their disclosed fixed-obligation stacks (Tier E, Section 5) total roughly \$3.3 trillion across the five rated names (OBSERVED, A)/(RECONSTRUCTED, B). Their balance sheets are the anchor of the system: four of the five generate operating cash flow above \$130 billion per year, and all five retain investment-grade ratings, though Oracle was cut to BBB- by S&P in July 2026 (OBSERVED, B). The labs are the demand growth story and the credit question mark: OpenAI runs at roughly a \$50 billion annualized revenue run rate (end-September 2026, investor documents via FT) (OBSERVED, B) with no public financial statements, Anthropic carries \$130 billion of primary compute commitments (\$100 billion AWS, \$30 billion Azure) (OBSERVED, A) against a memo-only aggregate portfolio near \$518 billion (OBSERVED, C), and xAI’s financing is press-reported only, including a \$20 billion GPU SPV package that Elon Musk publicly denied (OBSERVED, C).

The operating layer is the neoclouds and data-center operators: CoreWeave (public, the only pure-play neocloud with filed financials), plus Applied Digital, Nscale, TeraWulf, Crusoe, and the large private developers (CyrusOne, Vantage, Stack Infrastructure, EdgeConneX). These firms sign long-dated leases and capacity contracts with the demand layer, borrow against the contracted cash flows, and build the campuses. Their economics are the mirror image of the hyperscalers’: CoreWeave guided 2026 capex of \$35–39 billion against projected full-year revenue near \$12.8 billion (OBSERVED, B), funds the gap in the capital markets (it issued \$13.5 billion gross debt in the first half of 2026 (OBSERVED, B)), and carries \$35.5 billion of uncommenced lease obligations with 7–16 year terms (OBSERVED, A). Industry preleasing sits in the mid-70 percent range versus a 40–50 percent historical norm (CBRE 2026) (OBSERVED, B), which is both evidence of genuine demand and a measure of how far the development pipeline has run ahead of contracted revenue.

The funding layer is where the forensic interest lies. Four channels supply the capital. First, the public bond market: the hyperscalers issued tens of billions in 2026 (Amazon’s \$50 billion upsized March offering plus \$25 billion in July; Meta’s \$24.9 billion net new long-term debt in H1 2026) (OBSERVED, B)/(RECONSTRUCTED, B). Second, the banks: arranging and underwriting construction facilities and syndications (JPMorgan and MUFG on the \$38 billion Oracle-leased

Texas/Wisconsin construction debt; Bank of America, Citi, and Morgan Stanley launching the \$42 billion Broadcom/Anthropic Class A syndication in October 2026) (OBSERVED, B). Third, private credit: direct-lending funds, BDCs, and infrastructure funds (Blue Owl, Apollo, Blackstone, Carlyle, HPS, Ares) that fund SPVs, delayed-draw term loans (DDTLs), and junior tranches with no public disclosure regime (OBSERVED, B). Fourth, the vendors themselves: Broadcom and Nvidia now arrange financing for their own customers' purchases, the "chip-financing wave" of October 2026 (Broadcom's >\$50 billion "Nexus" arrangement for OpenAI, the \$60 billion Anthropic package, Oracle's chip SPV talks with Apollo/Goldman, Nvidia's \$500 billion platform with six private-credit partners, Nvidia's up-to-\$2 billion equity co-investment in the xAI SPV, the \$105 billion SB Energy backstop) (OBSERVED, B)/(OBSERVED, C). Insurers and pension funds enter as buyers of the long-dated paper (PIMCO anchored roughly \$18 billion of the Beignet/Hyperion notes; insurer and asset-manager money took the \$5 billion Waymo facility) (OBSERVED, B), which is where a future mark-to-market or redemption event would transmit into institutional balance sheets.

Table 2. 2026 capex guidance versus operating cash flow (\$B).

Borrower	2026 capex (guided)	OCF (base)	
Microsoft	115.9 (FY2026 actual)	182.9	
Amazon	~220	161.4	
Alphabet	195–205	185.7	–5.9 (Q2; first negative quart
Meta	130–145	130.3	+0.8 (Q2; capex consumed 98%
Oracle	~70 net cash (reported to 90–95 incl. prepayments)	32.0	
CoreWeave	35–39	7.3 (H1 annualized)	Deeply neg

Source: Workstream A exposure maps (vintage 2026-10-11). Grades: capex figures (OBSERVED, B)/(OBSERVED, C); OCF as in Table 5. Four of the six names are FCF-negative through the ramp; only Microsoft remains FCF-positive, which is why the core/periphery split Table 2 puts 2026 capex guidance against operating cash flow. in the stress tests follows the cash-flow statement rather than the balance sheet.

The insurer bid and the ABS template. The fourth funding channel deserves its own description because it is where institutional balance sheets meet AI credit. Insurers and asset managers are the natural buyers of long-dated AI paper: PIMCO anchored roughly \$18B of the Beignet/Hyperion notes (OBSERVED, B), and the \$5B Waymo facility took PIMCO, Blackstone, Sixth Street, Capital Group, Loomis, and T. Rowe as lenders (OBSERVED, B), which is the direct evidence of insurer and asset-manager balance-sheet exposure to AI credit. The distribution template is the securitized one: Switch's \$840M ABS (September 2026, 13 data centers) is the deal the market points to as the insurer-bid template for data-center cash flows (OBSERVED, B). Two caveats. First, no public disclosure of insurer AI-bond concentrations was found (gap), so the RBC-pressure channel in the propagation matrix (Table 26) is the least observable of the high-severity cells: a forced de-risking sale of BBB AI paper into a falling market is a mechanism without a visible tripwire. Second, risk appetite exists at a price: the AT1 primary market was open in October 2026 (Bankinter 6.0%, 9.8x book) (OBSERVED, B), which suggests insurers absorb wider-spread paper willingly when compensated, and unwillingly only when marks force their hand.

2.2 The seven financing mechanisms, in plain language

Seven structures recur across the cohort. The mechanism catalog (Section 6.2) gives the full GAAP treatment; this section describes what each one does economically.

First, the special-purpose vehicle (SPV). A legally separate entity is formed to own an asset (a data center, a fleet of GPUs) and borrow against it. Outside investors supply equity, the vehicle issues debt, the operating company takes a minority stake and signs a long-term lease or capacity contract, and it often bears construction risk (delays, cost overruns). The flagship is Meta’s Beignet Investor LLC / Hyperion: \$27.3 billion of senior secured notes (6.581% coupon, 2049 maturity, S&P A+) sitting on Beignet’s books rather than Meta’s, with Meta owning 20%, Blue Owl-managed funds 80%, Meta running construction and able to lease the site for up to 20 years (OBSERVED, B). Whether the debt stays off the sponsor’s balance sheet turns entirely on the consolidation analysis (ASC 810); Ernst & Young flagged Beignet’s non-consolidation assessment as a critical audit matter in Meta’s 2025 10-K (OBSERVED, B), which is the audit system working as designed, not evidence of a problem.

Second, the sale-leaseback. The company sells an asset (or one it is buying) to an investor vehicle and simultaneously leases it back, converting the asset to cash while keeping operational use. Amazon’s proposed \$8 billion Grace Blackwell chip SPV is the template: thousands of chips across 12+ US data centers sold into a vehicle funded by outside investors (up to 10% equity) and investment-grade-rated bonds, with Amazon retaining no equity and leasing the chips back (OBSERVED, B). The accounting question is whether the transfer qualifies as a true sale; if it does not, the “sale” proceeds are a financing liability and nothing leaves the balance sheet.

Third, the residual-value guarantee (RVG). A party guarantees that a leased or financed asset will be worth at least a stated amount at term end, paying the shortfall if realized value disappoints. RVGs credit-enhance financings where the collateral (GPUs, AI racks) depreciates fast and has uncertain secondary value. Instances: Broadcom’s residual-value support on the \$42 billion Anthropic Class A tranche (OBSERVED, B), Broadcom’s \$29 billion AI-rack backstop at maximum notional (OBSERVED, A), and Oracle’s \$3.3 billion lessor-borrowing guarantee (matured September 2026) (OBSERVED, A).

Fourth and fifth, operating and finance leases. Since ASC 842 (2019), both go on the balance sheet at commencement; the pre-2019 world in which operating leases were off-balance-sheet no longer exists. The live number is leases signed but not yet commenced: disclosed in footnotes, not recognized, converting automatically to on-balance-sheet liabilities as they commence on a pre-signed schedule. The cohort’s uncommenced totals are \$329.1 billion (Microsoft), \$279.0 billion (Meta), \$260–288 billion (Oracle), \$137.2 billion (Amazon), \$85.2 billion (Alphabet), \$35.5 billion (CoreWeave), and \$25.0 billion (Nvidia) (OBSERVED, A), with terms of 7–30 years. Microsoft’s \$66.6 billion of finance lease liabilities exceeds its \$40.3 billion of GAAP debt (OBSERVED, A), a reminder that “debt” as a line item is a poor proxy for fixed obligations.

Sixth, purchase and supply commitments. Executory contracts to buy servers, chips, power, or cloud capacity in the future; not recognized until performance, disclosed under ASC 440. Alphabet’s \$707 billion fixed-or-guaranteed commitment portfolio (an aggregate, composition undisclosed in the monitor) (OBSERVED, A), Meta’s \$349.3 billion of non-cancelable contractual commitments (OBSERVED, A), Nvidia’s \$279 billion of future supply and capacity commitments (OBSERVED, A), and Broadcom’s \$126.8 billion of purchase commitments (roughly \$126 billion falling due over 2027–2028) (OBSERVED, A) all sit here. The economic content is unambiguous: unconditional means the cash is committed.

Seventh, supplier-arranged customer financing, the circular structure. The vendor arranges third-party financing so its customer can buy the vendor’s product: private-credit lenders fund a vehicle or loan, the customer signs a purchase or lease commitment for the vendor’s chips, and the vendor sweetens the deal with a guarantee, an equity co-investment, or convertible notes. Cash flows in a

loop (lender to customer, customer to vendor as revenue, vendor growth justifying more lending), and no single filing consolidates the loop. The October 2026 instances are listed in Section 2.1; the common skeleton is a private-credit-funded vehicle buying chips, rental or purchase payments servicing the debt, and the chipmaker providing partial guarantees to attract lenders (OBSERVED, B)/(OBSERVED, C).

Worked example: Beignet through the ASC 810 test. Start with the economics. Beignet Investor LLC was formed to own the Hyperion data-center campus; it issued \$27.3B of senior secured notes (October 2025, par, 6.581% coupon, 2049 maturity, S&P A+) (OBSERVED, B); Meta owns 20% and Blue Owl-managed funds 80%; Meta runs construction, bears delay and overrun costs, and can lease the site for up to 20 years (OBSERVED, B). The debt sits on Beignet’s books, not Meta’s. Meta discloses maximum exposure of \$45.99B against a booked ~\$2.37B equity stake (OBSERVED, C). Now the accounting. Under ASC 810-10, Beignet is a VIE; the question is the primary beneficiary. Prong (a), power: who directs the activities that most significantly affect economic performance? Meta runs construction and will operate the site under a 20-year lease, which points toward Meta; Blue Owl controls the financing vehicle and its governance, which points the other way. Prong (b), economics: Meta’s 80%-partner absorbs most of the capital risk, while Meta’s \$45.99B maximum exposure (lease + overruns + equity) is significant to Beignet. The assessment is genuinely judgmental, which is exactly why EY flagged it as a critical audit matter in the 2025 10-K (OBSERVED, B): the auditor is required to flag precisely this kind of judgment, and the flag is the system working, not evidence of a problem. The market’s verdict is visible in the price: the notes traded to ~91c by October 2026 (~7.55%, +230bp over UST) (OBSERVED, B), with PIMCO’s ~\$18B anchor marked near 94.5c (~\$1B paper loss) (OBSERVED, C). Whatever the consolidation answer, the economics stay: off-balance-sheet debt is not off-risk.

2.3 What “off-balance-sheet” means after ASC 842

Much public commentary uses “off-balance-sheet” loosely. After ASC 842 and IFRS 16, the honest taxonomy has four buckets, and this paper labels every figure with its bucket.

Bucket one is recognized debt: borrowings carried on the balance sheet (Tier A). Bucket two is recognized non-debt liabilities: finance and operating lease liabilities, which are on the balance sheet at commencement but are not “debt” in the borrowings line (Tiers B and C). Bucket three is disclosed-but-unrecognized: uncommenced leases, unconditional purchase obligations, and guarantees disclosed in footnotes at undiscounted or maximum notional (Tier D and the commitment/guarantee adds in Tier E). Bucket four is genuinely off-balance-sheet: obligations of non-consolidated SPVs (the sponsor’s exposure is limited to its equity, guarantees, and lease commitments, which sit in buckets one to three), executory contracts before performance, and contingent exposures carried at fair value under ASC 460 that understate the tail.

Two consequences follow. First, a sale-leaseback is never off-liability: even a qualifying sale puts a lease liability on the balance sheet. Second, the cohort’s largest numbers live in bucket three, which converts to bucket two automatically as leases commence, with no new financing decision required. The binding question for credit analysis is therefore not whether the accounting is compliant (it is, as far as the evidence shows) but whether leverage metrics built on bucket one capture the fixed-payment burden. They do not: Tier E is 4–17x Tier A for the rated names, and the entire difference sits in buckets two through four. Table 3 puts cohort examples into the four buckets. The table is the paper’s central accounting claim in miniature: everything the market calls “hidden

leverage” in the AI buildout sits in buckets two through four, and bucket three (disclosed but unrecognized) is where the large numbers live.

Table 3. The honest taxonomy: four buckets of fixed exposure, with cohort examples.

Bucket	Content	Cohort example	In debt-based leverage ratios?	
1. Recognized debt	Borrowings, carrying amount	MSFT \$40.3B; ORCL \$125.3B	Yes	
2. Recognized non-debt liabilities	Lease liabilities (finance + operating), carrying	MSFT finance leases \$66.6B; CRWV operating leases \$16.3B	Sometimes (analyst-dependent)	
3. Disclosed but unrecognized	Uncommenced leases, purchase obligations (notional); guarantees (max notional)	MSFT \$329.1B uncommenced; GOOGL \$707B commitments; BRCM \$29B backstop	No	<i>Source:</i>
4. Genuinely off-balance-sheet	Non-consolidated SPV debt; executory contracts pre-performance; contingents at fair value	Beignet \$27.3B notes; AMZN \$8B SPV (proposed); NVDA \$105B backstop (contingent)	No	

Workstream A mechanism catalog. Grades as in Table 11. Bucket 3 converts to bucket 2 automatically as leases commence; bucket 4 converts only if consolidation or performance triggers it.

3 Related literature

This paper sits at the intersection of four literatures: the accounting of leases and special-purpose entities, the economics of shadow banking and financial crises, the case literature on accounting-driven failures, and the emerging quantitative literature on AI infrastructure capex. We take from each only what the forensic question requires.

The lease and SPE accounting literature establishes why the current disclosure regime looks the way it does. Feng et al. (2009) document the determinants of special-purpose-vehicle use and show that SPEs are associated with earnings management, which is the empirical basis for the post-Enron consolidation reforms (FIN 46, now ASC 810) whose bite this paper tests. Dechow and Schrand (2004) survey earnings quality and the channels through which accounting choice obscures performance; their framework disciplines our distinction between compliant-but-opaque structures and non-compliant ones. The lease-capitalization literature (summarized in the IASB/FASB basis for conclusions to IFRS 16 and ASC 842) established that operating-lease obligations are economically debt-like, which is why the Tier C step of our framework treats them as fixed obligations regardless of caption. The standards themselves are the product of the earlier cases this paper compares: FIN 46 (now ASC 810) was written against Enron’s SPEs, ASC 842/IFRS 16 against the off-balance-sheet operating lease, and ASC 860’s sale-accounting tightening plus risk-retention rules against the GFC securitization chain. That lineage is why the resemblance matrix’s accounting verdicts keep coming back compliant: the current framework was designed by people who had just watched each of these tricks fail. Our contribution on this margin is forensic rather than theoretical: we map each live AI structure to its ASC/IFRS treatment and identify the specific classification boundaries (lease versus service commitment under ASC 842/IFRS 16; the primary-beneficiary test under ASC 810/IFRS

10; the related-party tripwire under ASC 850) where structuring choices move large obligations between balance-sheet and footnote treatment.

The shadow-banking and crisis literature supplies the structural vocabulary for the historical comparison. Pozsar et al. (2010) map the US shadow banking system at nearly \$20 trillion as of March 2008 and show how credit intermediation outside the bank perimeter reproduces bank-like fragility (maturity transformation, leverage, run risk) without bank-like supervision; their perimeter concept is the lens through which we evaluate private-credit AI financing. Gorton and Metrick (2012) analyze the run on repo as the crisis’s transmission mechanism; the analogue we test is whether runnable short-term funding exists anywhere in the AI chain (we find the maturity transformation is real in construction finance but there is no household-leverage or bank-run dimension). The funding-liquidity spiral formalized by Brunnermeier and Pedersen (2009), in which market illiquidity and funding illiquidity reinforce each other, is the mechanism we test in the propagation matrix: the loop exists in construction finance and private-credit facilities, but it lacks the runnable short-term liabilities that made it systemic in 2008. The official investigations, Financial Crisis Inquiry Commission (2011) and the Levin-Coburn report (United States Senate Permanent Subcommittee on Investigations, 2011), document the originate-to-distribute incentive failures, rating-agency conflicts, and SIV mechanics that define the GFC case in our resemblance matrix; our verdict is that the funding architecture rhymes while the asset quality does not.

The case literature anchors the other five historical briefs. The Powers Report (Powers et al., 2002) is the primary record on Enron’s SPEs and the related-party opacity channel we compare against supplier-arranged AI financing. The Valukas examiner report (Valukas, 2010) is the canonical statement of “GAAP-compliant but economically misleading” reporting via Repo 105; we use it as the precedent for why compliance is not the end of the analysis. The SEC enforcement records on WorldCom (U.S. Securities and Exchange Commission, 2002b) and Tyco (U.S. Securities and Exchange Commission, 2002a, 2006) and the Breeden report (Breeden, 2003) define the fraud mechanisms we test for and do not find. On Greensill, the FINMA enforcement conclusion (Swiss Financial Market Supervisory Authority FINMA, 2023) and the House of Commons Treasury Committee report (House of Commons Treasury Committee, 2021) establish the private-credit, future-claims, and guarantee-dependence mechanics that make Greensill the closest single-firm parallel in our set.

The forensic-detection literature gives the warning-sign mapping its method. Dechow et al. (2011) show that material misstatements are predictable from accruals, financial performance, and off-balance-sheet activity, among other signals; our Table 19 applies the same logic in reverse, asking which of the classic pre-collapse signals have observable analogues in today’s filings and which are null. The nulls are the point: a detection framework that only records hits cannot distinguish a clean structure from an unexamined one.

Finally, the AI capex literature within this working-paper series. KR-WP-2026-10 (Kilambi Research, 2026a) documents the hyperscaler capex ramp and its physical drivers; KR-WP-2026-11 (Kilambi Research, 2026b) (“Can the AI Buildout Pay for Itself?”) models the utilization path required for the buildout to earn its cost of capital; KR-WP-2026-12 (Kilambi Research, 2026c) stress-tests the banking system’s capacity to fund it. We cite these by number and do not duplicate them: this paper’s distinct contribution is the liability side (how the buildout is financed and where the obligations sit), the accounting forensics, and the historical comparison. The leverage monitor whose vintages feed this paper is maintained as a standing series; its methodology is summarized in Section 5 and Appendix A.

4 Data

4.1 The obligation-level dataset

The bottom line on the data: the analysis rests on a purpose-built obligation dataset of 116 rows across 84 entities and 85 sources, five vintages from September 27 to October 11, 2026, with primary committed exposure of \$4,315.4 billion (OBSERVED, A), supplemented by SEC filings, the banking-stress sample from KR-WP-2026-12, and primary sources for the six historical cases. Every figure in the paper carries an evidence grade; grades govern provenance, not certainty.

The dataset is the AI leverage monitor’s five-vintage panel (2026-09-27 through 2026-10-11). Each row is one economic obligation of one entity: debt tranches, lease liabilities, uncommenced leases, purchase and supply commitments, guarantees, SPV exposures, and revolving or delayed-draw facilities. The 116 rows cover 84 entities and cite 85 distinct sources (OBSERVED, A). Primary committed exposure sums to \$4,315.4 billion (OBSERVED, A). Memo rows (press-reported, talks-stage, or duplicative lender-side views) are excluded from all aggregates per the counting rules (Appendix A); the exclusion is auditable row by row. Each row is labeled by basis: recognized debt (carrying amount), recognized lease liability (carrying), undiscounted contractual commitment (notional), contingent (maximum notional), or off-balance-sheet. These bases are never summed without labeling the mix: Tier E, for example, deliberately mixes carrying and notional figures, and every Tier E total in this paper is labeled as a mixed-basis ceiling of fixed exposure, not a debt figure.

Table 4 summarizes the dataset. Panel A gives the dataset-level statistics. Panel B decomposes primary committed exposure by company group on the Tier E (economic fixed obligations) definition, showing the ten profiled entities plus the residual in other rows.

The \$98.8B residual in other rows is 2.3% of primary committed exposure: smaller data-center operators, private-credit facilities, and utility and land commitments captured in the 116-row sweep but below the profiling threshold. The ten profiled entities are 97.7% of the total, so the company-level results in Section 6.1 cover essentially the entire measured exposure. The residual’s composition is the natural scope for the next vintage’s entity expansion.

4.2 SEC filings and company disclosures

Filing-sourced figures (10-K/10-Q footnotes on leases, commitments, guarantees, and VIEs; 8-Ks; prospectuses; earnings releases and decks) are the primary source for the Tier ladders in Section 6.1. Table 5 records the filing vintage, the operating-cash-flow base, and the Tier A composition for each rated name. As-of dates follow each company’s fiscal calendar: Microsoft June 30, 2026 (FY2026 10-K); Amazon, Alphabet, Meta, CoreWeave June 30, 2026 (Q2 10-Q); Nvidia July 26, 2026 (Q2 FY2027 10-Q); Broadcom August 2, 2026 (Q3 FY2026 10-Q); Oracle a mixed August 31, 2026 (Q1 FY2027 10-Q) and May 31, 2026 (FY2026 10-K) vintage, with the staleness handled explicitly (Section 6.1). One material access limitation: the SEC’s data.sec.gov API refused access (HTTP 403) during the build, so no fresh XBRL pulls were possible; filed figures come from the monitor vintage (built from 10-K/10-Q filings) plus EDGAR HTML indexed in search and company earnings releases (OBSERVED, B). Direct EDGAR filing fetches remain an open item for the next vintage, not a blocker for this paper’s conclusions, because every headline filing figure is corroborated across at least two sources.

Table 4. Summary statistics: the obligation dataset (vintage 2026-10-11).

Panel A: dataset		
Vintages	5 (2026-09-27 to 2026-10-11)	
Obligation rows	116 (OBSERVED, A)	
Entities	84 (OBSERVED, A)	
Distinct sources	85 (OBSERVED, A)	
Primary committed exposure	\$4,315.4B (OBSERVED, A)	
Memo rows (excluded from aggregates)	press/talks-stage and lender-side duplicates	
Panel B: Tier E by company group (\$B)		
Microsoft	686.5 (OBSERVED, A)	Source: AI
Amazon	509.3 (OBSERVED, A)	
Alphabet	913.0 (OBSERVED, A)	
Meta	740.6 (RECONSTRUCTED, B)	
Oracle	445.8 (RECONSTRUCTED, B)	
<i>Subtotal: five hyperscalers</i>	<i>3,295.2</i>	
CoreWeave	87.1 (OBSERVED, A)	
Nvidia	480.4 (RECONSTRUCTED, B)	
Broadcom	219.2 (RECONSTRUCTED, C)	
Anthropic (commitments only)	130.0 (OBSERVED, A)	
OpenAI (\$4.7B revolver; rest memo)	4.7 (OBSERVED, B)	
<i>Subtotal: ten profiled entities</i>	<i>4,216.6</i>	
Other rows (smaller entities, rounding)	98.8 (RECONSTRUCTED, B)	
<i>Total primary committed</i>	<i>4,315.4 (OBSERVED, A)</i>	

leverage monitor, vintage 2026-10-11 (Workstream A). Tier E = economic fixed obligations; mixed carrying/notional basis; guarantees at maximum notional. Memo rows excluded per counting rules (Appendix A).

Table 5. Filing vintages and operating-cash-flow bases by company.

Company	As-of (filing)	OCF base (\$M)	Grade	Tier A composition
Microsoft	2026-06-30 (FY2026 10-K)	182,935	(OBSERVED, B)	Current \$9,227M + long-term \$31,067M
Amazon	2026-06-30 (Q2 2026 10-Q)	161,400	(OBSERVED, A)	Balance-sheet aggregate
Alphabet	2026-06-30 (Q2 2026 10-Q)	185,675	(RECONSTRUCTED, B)	Balance-sheet aggregate
Meta	2026-06-30 (Q2 2026 10-Q)	130,300	(RECONSTRUCTED, B)	\$84,000M face; no current portion
Oracle	2026-08-31 (Q1 FY2027 10-Q) / 2026-05-31 (FY2026 10-K)	31,977 (FY2026)	(OBSERVED, A)	Current \$7,625M + non-current \$117,712M
CoreWeave	2026-06-30 (Q2 2026 10-Q)	7,326 (H1 annualized)	(RECONSTRUCTED, B)	Net recourse \$31,405M + non-recourse \$3,663M
Nvidia	2026-07-26 (Q2 FY2027 10-Q)	102,718 (FY2026)	(OBSERVED, A)	Balance-sheet aggregate
Broadcom	2026-08-02 (Q3 FY2026 10-Q)	27,537 (FY2025; stale)	(OBSERVED, A)	Balance-sheet aggregate

Source: AI leverage monitor, vintage 2026-10-11 (Workstream A exposure maps). Oracle's mixed as-of is handled per the staleness rule (Section 6.1); Broadcom's OCF base is stale by about nine months and labeled as such everywhere it is used.

4.3 Banking and market data

Banking-system figures (non-depository financial-institution lending, Shared National Credit non-pass rates, stress-test results) are taken from the Kilambi Research (2026c) sample and cited without duplication. The banking-system picture, taken from Kilambi Research (2026c) without duplication, is that the banks are not the solvency constraint: all 32 US banks stayed above minimum capital in the June 2026 severely-adverse test, and pre-provision net revenue covers charge-offs 6.4x (OBSERVED, B). The bank leg that matters for this paper is the funding leg: non-depository financial-institution lending stands at \$2.026 trillion (+19% year-on-year), Shared National Credit non-pass rates are 8.6%, and the funding-first sequence means banks transmit stress through warehouse lines, subscription facilities, and retained pipelines (the \$38B Oracle construction debt) rather than through capital impairment (OBSERVED, B). Market stress marks are October 2026 prints: Oracle 5Y CDS 261bp (Goldman, record), Broadcom 5Y CDS 136bp, Beignet/Hyperion 2049s notes near 91 cents, Oracle 2055 bonds near 77 cents, US 10-year near 5.2% (OBSERVED, B). GPU spot-rate data are CloudZero’s series (spot rates down 30–40% from peaks) (OBSERVED, B);

Table 6. Market stress marks, October 2026.

Mark	Print	Grade	What it prices	
Oracle 5Y CDS	261bp (record)	(OBSERVED, B)	Oracle refinancing risk; catalyst 2 tripwire at 300bp	
Broadcom 5Y CDS	136bp (record)	(OBSERVED, B)	Supplier-financing platform risk	
Beignet/Hyperion 2049s	~91c (~7.55%, +230bp over UST)	(OBSERVED, B)/(OBSERVED, C)	Recovery skepticism on long-dated AI SPV paper	
Oracle 2055s	~77c	(OBSERVED, B)	Same, at longer duration	<i>Source:</i>
US 10-year	~5.2%	(OBSERVED, B)	Rate backdrop for the A1 coupon assumptions	
September AI debt issuance	~\$23B	(OBSERVED, B)	Already-weak primary window (MS via FT, Oct 11)	
CloudZero spot GPU rates	–30–40% from peaks	(OBSERVED, B)	Collateral and service economics; catalyst 6 input	
Neocloud rate cards	Raised again in window (Nebius)	(OBSERVED, B)	Contracted-vs-spot gap the model watches	

Workstream A/C market data (vintage 2026-10-11). CDS prints via Goldman, October 9, 2026. These marks are the market’s running verdict on the structures this paper dissects; Section 6.6 turns them into observable catalysts.

neocloud rate cards and BDC fair-value marks are the corroborating series for the catalyst dashboard. Table 6 collects the October 2026 prints in one place.

4.4 Historical primary sources

Each historical brief in Section 6.3 is anchored in primary sources, summarized in Table 7. The one provenance caveat: the full Valukas examiner report (about 2,200 pages) was consulted via detailed secondary summaries quoting it at length, not via the complete primary PDF; figures attributed to it are graded B rather than A.

Table 7. Primary sources for the six historical cases.

Case	Primary source	
Enron (2001)	Powers, Troubh & Winokur, <i>Report of Investigation by the Special Investigative Committee of the Board of Directors of Enron Corp.</i> (Feb 1, 2002) (OBSERVED, A)	
WorldCom (2002)	SEC Litigation Release 17588 / AAER 1585 (Jun 27, 2002); Breeden, <i>Report of Investigation by the Corporate Monitor</i> (Aug 2003) (OBSERVED, A)	
Tyco (2002)	SEC LR-17722 / AAER 1627 (Sep 12, 2002); SEC Press Release 2006-58 (Apr 17, 2006) (OBSERVED, A)	
GFC (2007–2009)	Financial Crisis Inquiry Commission, <i>The Financial Crisis Inquiry Report</i> (Jan 27, 2011); Senate PSI, <i>Wall Street and the Financial Crisis</i> (Apr 13, 2011); Pozsar et al., NY Fed Staff Report 458 (Jul 2010) (OBSERVED, A)	Source:
Lehman Repo 105 (2008)	Valukas, <i>Report of Anton R. Valukas, Examiner</i> (Bankr. S.D.N.Y., Mar 11, 2010), via detailed quoted summaries (OBSERVED, B)	
Greensill (2021)	FINMA enforcement conclusion (Feb 28, 2023); House of Commons Treasury Committee, <i>Lessons from Greensill Capital</i> (Jul 20, 2021) (OBSERVED, A)	

Workstream B primary-source table. Full citations in References.

4.5 Evidence grading

Every number in the paper carries a two-dimensional grade: provenance (Observed = stated in a cited primary source; Reconstructed = built from primary sources by the desk; Modelled = produced by the desk’s models; Assumed = an explicit modeling input) and quality (A = official/primary or corroborated; B = reputable secondary or single primary; C = provisional). The (LEVEL, GRADE) macro renders these in the text, e.g. (OBSERVED, A). Grades govern provenance, not certainty: an (OBSERVED, A) figure can still be stale (Oracle’s uncommenced leases) or unaudited (press-sourced commitment totals), and the text says so where it matters. Model outputs are Modelled by construction and carry the assumption tables that produce them. The data vintage statement for the full paper is 2026-10-11 unless a figure is labeled otherwise.

5 Methodology

5.1 The five-tier leverage framework

The bottom line of the methodology: we measure leverage in five nested tiers so that the reader can see exactly which obligations are recognized debt and which are disclosed commitments, instead of arguing about any single “true” leverage number.

Tier A is GAAP recognized debt (carrying amount): current plus long-term borrowings. Tier B adds finance lease liabilities (carrying). Tier C adds operating lease liabilities (carrying): the full on-balance-sheet lease position under ASC 842. Tier D adds uncommenced leases and other disclosed commitments (undiscounted notional): obligations signed but not yet recognized. Tier E adds unavoidable purchase, supply, and take-or-pay commitments (undiscounted notional) plus probability-weighted guarantees. All four guarantee weights in the monitor are 1.0, so Tier E reports

the ceiling of fixed exposure, not an expected loss. Tiers A–C are pure carrying basis; Tiers D and E mix carrying and notional, and are always labeled as such. The framework is deliberately additive so that the ratio Tier E / Tier A (Section 6.1) measures how much economic fixed exposure sits outside the borrowings line.

5.2 Counting rules

Three rules matter for interpreting the results; the full set is in Appendix A. Table 8 summarizes the ten rules. First, one economic obligation is one row: memo rows (press-reported, talks-stage, or duplicative lender-side views of the same facility) are excluded from all aggregates, so CoreWeave’s consolidated SPVs are not double-counted with BDC lender-side positions, and Anthropic’s \$130 billion of primary commitments are not double-counted inside its \$518 billion memo portfolio. Second, basis is labeled on every line and mixed bases are never summed without labeling the mix. Third, staleness is handled by rule, not by carry-forward: the monitor did not carry Oracle’s May-31 uncommenced figure (\$260 billion) into the August row, so both the filed figure and the newer \$288 billion press figure are shown with their vintages.

Table 8. Counting rules (summary).

#	Rule
1	One economic obligation, one row: each row is a distinct obligation of a distinct entity.
2	Memo rows excluded from all aggregates: press-reported, talks-stage, and duplicative lender-side views. The exclusion is auditable row by row.
3	Basis labeled on every line (recognized carrying, undiscounted notional, contingent maximum notional, off-balance-sheet); mixed bases never summed without labeling the mix.
4	No double counting across parties: a facility appears once, on the obligor’s books; lender-side views are memo.
5	An undrawn revolver commitment is not debt.
6	Guarantees enter Tier E at maximum notional with weight 1.0 (ceiling, not expected loss); disclosed escrow noted but not netted.
7	FX: all figures in USD at the vintage’s stated rates.
8	Staleness handled by rule, not carry-forward: filed figures are not carried across as-of dates; newer press figures shown alongside with vintages.
9	Subsequent-event rows are memo until a filing-period row carries them.
10	Tier ladder: A = GAAP debt; B = A + finance leases; C = B + operating leases; D = C + uncommenced/commitments; E = D + purchase/take-or-pay + guarantees at max notional.

Source: AI

leverage monitor counting rules (full text with the dataset; worked examples in Appendix A).

5.3 The coverage model

The stress test compares each borrower’s annual fixed-obligation burden against its operating cash flow over 2027–2030. For year y :

$$\text{burden}(y) = \text{interest}(y) + \text{commenced lease payments}(y) + \text{uncommenced-lease ramp}(y) + \text{purchase/commitments}$$

Scheduled principal is included only for CoreWeave, whose debt amortizes and cannot be assumed rolled. For the five investment-grade borrowers the base case assumes maturing bonds are refinanced;

principal enters only in the refinancing-freeze scenario. Coverage is $OCF(y) / \text{burden}(y)$; a breach is coverage below 1.0x. Three cases are run: base (reported OCF run-rate grown per assumption A9), adverse (OCF minus 25%), severe (OCF minus 50% plus 200bp on rolled debt). Nine scenarios vary the drivers (weaker AI monetization, falling GPU rental rates, declining data-center valuations, higher rates, 12-month construction delay, energy constraints, top-3 lab customer default, 12-month refinancing freeze, private-credit counterparty failures). Full scenario results are in Section 6.5 and the scenario-results file. Table 22 defines the nine scenarios. The ten material assumptions (A1–A10) are: blended coupons 4.0–8.0% by borrower (ASSUMED, B); commenced-lease annual cash as carrying / 8 years for the four hyperscalers (ASSUMED, B); Oracle/CoreWeave lease payments at disclosed current portions (OBSERVED, B); uncommenced leases ramping straight-line over disclosed commence windows with 10–17 year terms (MODELLED, B); Microsoft/Amazon purchase commitments spread evenly over 2027–2030 (MODELLED, C); Meta’s \$81.65 billion of 2027 compute commitments on the observed schedule with the residual spread over 2028–2036 (MODELLED, B); Alphabet’s \$707 billion fixed-commitment portfolio excluded for lack of a schedule (ASSUMED, C) (the largest unmodeled exposure in the exercise); IG maturities rolled in base (ASSUMED, B); OCF growth 10–15% (40% for CoreWeave off a small base) (ASSUMED, C); no new debt, equity, or asset sales in any case.

Material limitations: the model excludes dividends and buybacks (including them would cut Microsoft’s base coverage to about 1.5x and Meta’s 2027 to about 1.05x; shareholder returns are the first discretionary buffer management would cut); OCF is pre-capex, and the buildout’s capex is the true cash drain (Amazon TTM capex near \$169 billion versus \$161 billion OCF).

5.4 The collateral-depreciation model

GPU-secured facilities are modeled as loan-to-value paths under straight-line depreciation of the collateral to a 10% residual over economic lives of two to five years. Two templates: F1, a CoreWeave-type DDTL (\$2.9 billion drawn, 75% initial LTV, 5%/yr amortization, 100% GPU collateral); F2, the \$38 billion Oracle-leased construction-debt template (80% initial LTV, 2%/yr amortization, 70% equipment / 30% data-center shell with a 25-year shell life). The model isolates hardware depreciation from covenant cures (no additional collateral posting, no cash sweep) because the question is recovery value in default, where cures have failed. Full annual paths are in Section 6.5.

5.5 The historical-comparison rubric

Each of the six historical cases is graded on a four-point resemblance scale: Strong, Partial, Weak, Misleading. The grade reflects a three-question verdict applied uniformly: (1) what genuinely resembles today’s AI financing arrangements; (2) what is misleading about the comparison; (3) the accounting verdict under current US GAAP/IFRS (whether the historical trick is blocked today, and where the live opacity sits). The grades are: Enron Partial, WorldCom Weak, Tyco Misleading (a null finding), GFC Partial (the strongest structural parallel in the set), Lehman Repo 105 Partial, Greensill Partial (the closest private-credit parallel). Section 6.4 presents the standardized matrix; Section 6.3 gives the condensed forensic briefs (about 400 words each) that the grades rest on.

Table 9. Coverage-model assumptions (A1–A10): justification and sensitivity.

#	Assumption	Justification	Sensitivity
A1	Blended coupons: MSFT 4.0%, AMZN 4.5%, GOOGL 5.0%, META 5.2%, ORCL 5.5%, CRWV 8.0% (ASSUMED, B)	US 10Y ~5.2% (Oct 2026); IG issuers carry older low-coupon bonds; Oracle CDS 261bp record; CoreWeave secured/HY mix	±100bp moves Oracle coverage ~0.06x, CoreWeave ~0.05x; no breach finding changes
A2	Commenced-lease annual cash = carrying / 8yr (MSFT/AMZN/GOOGL/META) (ASSUMED, B)	Weighted-average lease terms undisclosed; 8yr is a mid-estimate across office/DC portfolios	6yr divisor raises burden ~15%; no breach changes except META 2027 adverse (1.07x → 0.99x at 6yr)
A3	ORCL/CRWV lease payments = disclosed current portions (\$4,776M / \$591M) (OBSERVED, B)	10-Q current portion approximates next-12-month cash	Low uncertainty
A4	Uncommenced leases ramp straight-line over disclosed commence windows; 10–17yr terms (MODELLED, B)	Commence windows and terms disclosed in footnotes (ORCL: Q1 FY27–FY29, 15–19yr; CRWV: 2026–2029, 7–16yr; META: 2026–2036, 1–30yr)	Shorter terms push Oracle FY2029 base coverage toward 1.1x; direction unambiguous, level not
A5	MSFT/AMZN purchase commitments spread evenly over 2027–2030; ORCL over 2027–2029 (MODELLED, C)	No annual schedules disclosed; 10-K language points to 1–4yr horizons	Front-loading to 2 years cuts MSFT base 2027 to ~1.6x; still no breach
A6	Meta compute: \$81,650M in 2027 (observed schedule); residual \$214,140M spread 2028–2036 (MODELLED, B)	Treats maturity _{wall} compute rows as the scheduled portion of the \$349,310M total; avoids double counting	±\$0.2x on META 2027; adverse 2027 stays near 1.0x
A7	GOOGL \$707B fixed/guaranteed commitments excluded from annual burden (ASSUMED, C)	Portfolio figure explicitly “not added to the broader MD&A figure”; no schedule	Largest unmodeled exposure; 5% falling due in one year halves GOOGL coverage (still >5x)
A8	IG maturities rolled in base; repaid from OCF only in freeze scenario (ASSUMED, B)	Standard for IG issuers in normal markets; the freeze scenario tests the alternative	Repaying 50% of maturities in base cuts META 2027 to 1.39x, ORCL FY2027 to 1.14x
A9	OCF growth: MSFT/AMZN/META 12%, GOOGL 10%, ORCL 15%, CRWV 40% (ASSUMED, C)	Cloud/AI growth trajectories; CRWV from a small base on contracted backlog	±5pp growth changes 2030 coverage 10–25%; no breach finding changes before 2029
A10	No new debt, equity, or asset sales in any case	Stress test isolates operating coverage	Relaxing this is exactly CoreWeave’s business model; see interpretation

Source: Workstream C coverage-model.md. All scenario outputs are (MODELLED, C) by construction.

Table 10. Collateral-model assumptions (C1–C5): justification and sensitivity.

#	Assumption	Justification	Sensitivity
C1	Straight-line depreciation to 10% residual (ASSUMED, B)	Standard equipment-finance convention; 10% floor proxies scrap/part-out value	0% residual accelerates breach ~1 quarter; 25% residual delays the 5yr-life breach to year 3
C2	Accelerator economic life 2–5 years (ASSUMED, B)	CloudZero spot rates down 30–40% from peaks; new GPU generations every ~2 years; neocloud contracts typically 3–5yr	Breach is insensitive to life within range: 1–2 years regardless
C3	No additional collateral posting, no cash sweep (ASSUMED, C)	Isolates hardware depreciation from covenant cures	Real facilities have maintenance covenants; the question is recovery in default, where cures failed
C4	F1 amortizes 5%/yr, F2 2%/yr (ASSUMED, C)	DDTL facilities amortize; construction debt is interest-heavy pre-completion	10%/yr amortization delays F1 breach about one year at 4–5yr life
C5	F2 collateral 70% equipment / 30% shell (ASSUMED, C)	Shell/power infrastructure is the minority of AI data-center project cost	50/50 split delays breach roughly one year

Source: Workstream C collateral-depreciation.md (model code: model-code/collateral_model.py). Full annual paths in collateral-ltv-paths.csv.

6 Results

6.1 Company-level exposure maps


The bottom line of the exposure maps: recognized debt is a small fraction of economic fixed exposure for every rated name (Tier E runs 2.5x to 17x Tier A), the binding constraints are the disclosed-but-unrecognized pipeline (uncommenced leases) and vendor-arranged financing, and the cohort splits into a solvent core (the four hyperscalers ex-Oracle), a stressed middle (Oracle), a market-dependent periphery (CoreWeave), and the supplier-financiers (Nvidia, Broadcom) whose exposure runs the other way. Table 11 gives the tier ladders; Figure 1 plots Tier A against Tier E by company; Figure 2 plots the all-in burden (Tier E / OCF); Figure 3 shows the debt-maturity wall where the monitor carries a schedule.

Microsoft. The cleanest balance sheet in the cohort and the largest Tier E / Tier A multiple. GAAP debt is \$40.3 billion (OBSERVED, A) against Tier E of \$686.5 billion (OBSERVED, A) (17.0x), driven by \$329.1 billion of uncommenced non-cancelable leases (undiscounted; the largest single disclosed-but-unrecognized lease number in the cohort) (OBSERVED, A) plus \$194.1 billion of purchase commitments and \$34.6 billion of construction commitments (OBSERVED, A). Principal coverage is 19.8x (\$182.9 billion FY2026 OCF against \$9.25 billion of FY2027 maturities) (OBSERVED, B); all-in fixed-obligation years are 3.75x. Microsoft is the only hyperscaler here still generating positive free cash flow after the capex ramp (FY2026 capex \$115.9 billion; FCF \$67.0 billion) (OBSERVED, B). No bank or private-credit counterparty concentration is disclosed; uncommenced-lease counterparties are undisclosed (gap). Table 12 walks Microsoft’s ladder tier by tier to show how the framework accumulates. The arithmetic is exact: finance leases (\$66,594M) exceed GAAP debt (\$40,294M); the uncommenced-lease add (\$329,100M) is the single largest step; purchase and construction commitments complete the stack. Every other rated name follows the same walk with different weights (Table 11).

Table 11. Company tier ladders: from GAAP debt to economic fixed obligations (\$B).

Company	Tier (\$B)					Ratios	
	A	B	C	D	E	E/A	E/OCF
<i>Panel A: hyperscalers (solvent core)</i>							
Microsoft	40.3	106.9	128.8	457.9	686.5	17.0x	3.75x
Amazon	132.2	145.7	242.0	379.2	509.3	3.9x	3.16x
Alphabet	100.2	102.8	120.8	206.0	913.0	9.1x	4.92x
Meta	83.7	83.7	112.3	391.3	740.6	8.9x	5.68x
<i>Panel B: stressed middle and periphery</i>							
Oracle	125.3	134.5	169.1	429.1	445.8	3.6x	13.9x
(with \$288B uncommenced update)				457.1	473.8	3.8x	14.8x
CoreWeave	35.1	35.3	51.6	87.1	87.1	2.5x	11.9x
<i>Panel C: supplier-financiers</i>							
Nvidia	33.4	33.4	38.9	63.9	480.4	14.4x	4.68x
Broadcom	59.4	n.a.	n.a.	n.a.	219.2	3.7x	8.0x
<i>Panel D: private labs (commitments only; no full ladder)</i>							
Anthropic	\$130.0B primary commitments (AWS \$100B + Azure \$30B)					n.a.	n.a.
OpenAI	\$4.7B revolver committed/undrawn; rest memo/press-only					n.a.	n.a.
xAI	no monitor rows (gap); \$20B SPV package press-only					n.a.	n.a.

Source: AI leverage monitor, vintage 2026-10-11 (Workstream A, company-metrics.csv). Grades: Panel A Tiers A–D (OBSERVED, A) except Meta Tiers B–E (RECONSTRUCTED, B); Oracle Tier D/E (RECONSTRUCTED, B); CoreWeave (OBSERVED, A); Nvidia Tiers B–E (RECONSTRUCTED, B); Broadcom Tier E (RECONSTRUCTED, C). Broadcom has no B/C/D ladder in the monitor (lease components undisclosed: gap). Oracle as-of is mixed (see text). OCF bases: MSFT FY2026 \$182.9B; AMZN TTM \$161.4B; GOOGL TTM \$185.7B; META TTM \$130.3B; ORCL FY2026 \$32.0B; CRWV H1-2026 annualized \$7.3B; NVDA FY2026 \$102.7B; AVGO FY2025 \$27.5B (stale). Tier E is mixed carrying/notional basis; guarantees at maximum notional.



Data as of Jun-Aug 2026 filings (vintage 11 Oct 2026). Tier A = GAAP debt (carrying); Tier E = Tier A plus leases, uncommenced leases, purchase commitments, and guarantees at maximum notional. Source: Company 10-K/10-Q filings via the AI leverage monitor. Note: Tier E mixes carrying amounts (debt, recognized leases) with undiscounted notionals (commitments); the bars show burden scale, not a single accounting caption.

Figure 1. Tier ladders: GAAP recognized debt (Tier A) versus economic fixed obligations (Tier E), by company. The gap between the bars is the disclosed-but-unrecognized commitment stack. Labels show the Tier E / Tier A multiple.

Source: AI leverage monitor, vintage 2026-10-11 (company-metrics.csv). Placeholder figure; coordinator to build with chart_style.qa_check.

Table 12. Tier walk, Microsoft: from GAAP debt to economic fixed obligations (\$M).

Tier	Adds vs prior tier	Amount
A: GAAP recognized debt	Current \$9,227M + long-term \$31,067M (carrying)	40,294 (OBSERVED, A)
B: + finance leases	Finance lease liabilities \$66,594M (carrying)	106,888 (OBSERVED, B)
C: + operating leases	Operating lease liabilities \$21,925M (carrying)	128,813 (OBSERVED, C)
D: + uncommenced leases	Uncommenced non-cancelable leases \$329,100M (notional)	457,913 (OBSERVED, D)
E: + commitments	Purchase commitments \$194,060M + construction \$34,566M (notional)	686,539 (OBSERVED, E)

Source: Microsoft FY2026 10-K via the AI leverage monitor, vintage 2026-10-11. As-of 2026-06-30. Tiers A–C are carrying basis; D–E adds are undiscounted notional.

Amazon. Tier E of \$509.3 billion (OBSERVED, A) (3.9x Tier A), built from \$137.2 billion of uncommenced leases and \$130.1 billion of unconditional purchase obligations (OBSERVED, A). All-in burden is the cohort’s lowest at 3.16x on \$161.4 billion TTM OCF (OBSERVED, A); principal coverage is not computable (no debt-maturity schedule in the monitor: gap). Two flags: TTM free cash flow was negative \$7.6 billion (RECONSTRUCTED, B) after property-and-equipment purchases rose \$66.1 billion year-on-year, with 2026 capex guided near \$220 billion (OBSERVED, B); and Q2 2026 net income included \$53.4 billion of non-operating pre-tax other income, primarily from the Anthropic investment (OBSERVED, A), meaning Amazon is simultaneously Anthropic’s largest compute counterparty and a major equity holder (balance sheets linked through the investment, not through debt). The \$8 billion Grace Blackwell chip SPV sale-leaseback is at the discussions stage (memo; FT via Reuters) (OBSERVED, B): Amazon would retain no equity, outside investors would hold up to 10%, and the SPV would issue bonds seeking an investment-grade rating resting on Amazon’s AA credit. Whether it achieves a true sale under ASC 842-40 depends on terms that do not yet exist.

Alphabet. Tier E of \$913.0 billion (OBSERVED, A) (9.1x Tier A), dominated by the \$707 billion fixed-or-guaranteed commitment portfolio (OBSERVED, A): an aggregate, undiscounted, undisclosed-composition figure that is the single largest Tier E input in the cohort and the least decomposed (priority item for the next vintage’s footnote work). All-in burden is 4.92x on \$185.7 billion TTM OCF (RECONSTRUCTED, B). Q2 2026 free cash flow was negative \$5.9 billion (OBSERVED, A), the first negative FCF quarter since the 2004 IPO, after \$44.9 billion quarterly capex exceeded \$39.1 billion OCF; full-year 2026 capex was guided to \$195–205 billion (OBSERVED, A). Alphabet raised \$49.6 billion of equity in June 2026 (including its first-ever mandatory convertible preferred) plus \$20.3 billion of debt (RECONSTRUCTED, B). The contracted-demand counterpart exists: Google Cloud backlog passed \$514 billion (OBSERVED, B).

Meta. Tier E of \$740.6 billion (RECONSTRUCTED, B) (8.9x Tier A). The flagged structure is Beignet Investor LLC / Hyperion: \$27.3 billion off-balance-sheet, Meta 20% / Blue Owl-managed funds 80%, Meta running construction and able to lease the site up to 20 years while bearing delay and overrun costs; \$27.3 billion of senior secured notes priced October 2025 at par (6.581% coupon, 2049 maturity, S&P A+) (OBSERVED, B). The notes traded down to about 91 cents in October 2026 (about 7.55% yield, +230bp over UST) (OBSERVED, B)/(OBSERVED, C); PIMCO’s roughly \$18 billion anchor is marked near 94.5 cents, about a \$1 billion paper loss (OBSERVED, C); reported Meta maximum exposure is \$45.99 billion against a booked roughly \$2.37 billion equity stake (OBSERVED, C). EY flagged the non-consolidation assessment as a critical audit matter in the 2025 10-K (OBSERVED, B)/(OBSERVED, C). Steve Eisman publicly called the structure

“reminiscent of bad times past” while stopping short of alleging fraud (OBSERVED, C). Meta’s uncommenced leases total \$279.0 billion and non-cancelable contractual commitments \$349.3 billion (mostly cloud capacity, servers, network infrastructure, data centers) (OBSERVED, A); \$68 billion of leases signed in July 2026 sit as a subsequent-event memo (OBSERVED, A). Principal coverage is 47.4x (\$130.3 billion TTM OCF against the \$2.75 billion 2027 maturity bucket); all-in burden is 5.68x. Q2 2026 FCF collapsed to \$784 million (OBSERVED, A) as \$31.1 billion quarterly capex consumed 98% of quarterly OCF; buybacks went to \$0 in Q2 2026 versus \$10.2 billion in Q2 2025; cash and marketable securities of \$90.3 billion versus \$83.7 billion debt leaves nominal net cash near \$6.6 billion but lease-adjusted net debt near \$22.1 billion (RECONSTRUCTED, B).

Oracle. The stressed middle. Tier E of \$445.8 billion (RECONSTRUCTED, B) (3.6x Tier A), rising to about \$473.8 billion (14.8x all-in burden) on the newer \$288 billion uncommenced-lease press figure (OBSERVED, B)/(OBSERVED, C) that the monitor deliberately did not carry into the August row (counting rule: no carry-forward of stale figures). The staleness is material and both vintages are shown: filed uncommenced data-center leases of \$260 billion (May 31, 15–19 year terms) versus \$288 billion as of August 31 (press, 15–19 year terms, commencing Q2 FY2027–FY2029). Principal coverage is 4.44x (\$32.0 billion FY2026 OCF against \$7.21 billion FY2027 maturities) (OBSERVED, A), the weakest among the rated hyperscalers; all-in burden is 13.9x (14.8x updated), the cohort’s highest by a wide margin. Context: FY2026 FCF was negative \$23.7 billion (OBSERVED, A); Q1 FY2027 OCF of \$23.0 billion included \$11.4 billion of customer prepayments with a significant financing component (OBSERVED, B)/(OBSERVED, C) (the FY2026 base is used for coverage; the Q1 figure ran at roughly 2.9x the FY2026 quarterly average, flattered by prepayments); FY2027 net cash capex guided near \$70 billion with reported capex up to \$90–95 billion including prepayments (OBSERVED, B)/(OBSERVED, C); RPO \$664 billion with about 13% expected to convert in the next 12 months (OBSERVED, A). Credit stress marks: 5Y CDS at a record 261bp; S&P cut to BBB- (July 2026); 2055 bonds near 77 cents; JPMorgan/MUFG unable for 6+ months to distribute the \$38 billion of Oracle-leased Texas/Wisconsin construction debt; a January 2026 bondholder lawsuit alleging understatement of borrowing needs around the \$18 billion September 2025 bond sale (OBSERVED, B)/(OBSERVED, C). Flagged structures: chip SPV negotiations with Apollo and Goldman Sachs (WSJ October 7, 2026; amount undisclosed; chips for a 1 GW data center cost tens of billions) (OBSERVED, B)/(OBSERVED, C), where off-balance-sheet treatment turns on ASC 810/842-40 assessments that cannot be made until terms exist; the \$3.3 billion lessor-borrowing guarantee (matured September 2026; Tier E carries maximum notional at weight 1.0, a ceiling not an expected loss; no crystallization disclosed) (OBSERVED, A); and Project Jupiter, where Oracle sent a force-majeure notice to Blue Owl’s Stack Infrastructure over power-delivery delays at the 2.45 GW Stargate-linked New Mexico campus, with the related roughly \$18 billion bank facility press-reported only (OBSERVED, C).

CoreWeave. The market-dependent periphery. Tier E of \$87.1 billion (OBSERVED, A) (2.5x Tier A), with \$35.5 billion of uncommenced data-center and equipment leases (7–16 year terms, commencing 2026–2029) (OBSERVED, A). Of \$35.1 billion GAAP debt, \$31.4 billion is classified recourse versus \$3.7 billion non-recourse (OBSERVED, A); the financing SPVs (Compute Acquisition Co. II/IV, DDTL V vehicles) consolidate into CoreWeave, so lender-side BDC positions are memo-only to avoid double counting. Principal coverage is about 1.0x on the annualized assumption (\$7.33 billion annualized H1 OCF against about \$7.51 billion next-12-month principal) and 0.5x on reported H1 alone (MODELLED, C); either way the thinnest in the cohort. The H1 OCF figure itself rests on a \$9.7 billion deferred-revenue (customer prepayment) pool that grew \$790 million in

Q2 (OBSERVED, B): positive operating cash flow is substantially a prepayment timing effect, not operating-profit conversion (Q2 operating loss \$49 million) (OBSERVED, B). Q2 interest expense was \$640 million, up 140% year-on-year (OBSERVED, B); senior note coupons run 8.5–9.75%, DDTL effective rates 7–15% (OBSERVED, B)/(OBSERVED, C); 2026 capex is guided at \$35–39 billion versus projected full-year revenue near \$12.8 billion (OBSERVED, B)/(OBSERVED, C); accumulated deficit is \$4.0 billion (OBSERVED, B). The revenue backlog is \$104 billion (plus \$25 billion added early Q3) (OBSERVED, B), with 41% of RPO expected over the 24 months ending June 30, 2028. Counterparty concentrations: DDTL 4.0 (\$8.5 billion facility, \$2.9 billion drawn, \$5.6 billion unfunded; lenders unnamed in the monitor: gap) (OBSERVED, A); HPS exposure via DDTL V; Carlyle filed positions in CoreWeave SPVs; a \$2.5 billion revolver (\$533 million LC outstanding) (OBSERVED, A); a \$1 billion Jane Street investment (OBSERVED, C); about \$7 billion of 15-year Applied Digital leases (OBSERVED, B); and a Blue Owl-sponsored data center leased to CoreWeave financed with a roughly \$1.1 billion junk-bond sale via Goldman Sachs-led banks (September 2026) (OBSERVED, B). Post-period addendum: a \$3.7 billion bond settled September 22, 2026 (upsized from \$3 billion) (OBSERVED, B); a press account describes \$4.2 billion of 2.875% converts due 2033 settled the same week (OBSERVED, C), and the \$0.5 billion discrepancy (greenshoe versus separate tranche) is unresolved.

Nvidia. The supplier side. Tier E of \$480.4 billion (RECONSTRUCTED, B) (14.4x Tier A), composed of \$279 billion of future supply and capacity commitments (OBSERVED, A) (up from \$119 billion the prior quarter), \$29 billion of cloud take-or-pay (OBSERVED, A), the \$105 billion SB Energy backstop at maximum notional (OBSERVED, A), and \$3.5 billion of land/power/shell guarantees (OBSERVED, A). The guarantees enter Tier E at weight 1.0 (ceiling, not expected loss); partners placed \$712 million in escrow against the \$3.5 billion guarantees, disclosed but not netted (OBSERVED, A). Net debt is near zero; cash and marketable securities were \$62.6 billion at FY2026 year-end (OBSERVED, A); a \$235 billion buyback authorization runs through FY2028 (OBSERVED, B). Nvidia’s exposure runs the other way (guarantees extended, equity invested): the August 2026 platform with Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs, and KKR is expected to raise more than \$500 billion for AI infrastructure (OBSERVED, C), with lenders demanding bigger guarantees; Nvidia contributed up to \$2 billion of equity to the xAI SPV (OBSERVED, C), the supplier as co-funder of its own demand.

Broadcom. Tier E of \$219.2 billion (RECONSTRUCTED, C) (3.7x Tier A; no B/C/D ladder in the monitor because lease components are undisclosed: gap), built from \$126.8 billion of purchase commitments (OBSERVED, A), \$3.9 billion of other commitments, and the \$29 billion AI-rack backstop at maximum notional (OBSERVED, A). The purchase-obligation wall is steep: roughly \$126 billion falls due over 2027–2028 (OBSERVED, B). All-in burden is 8.0x on the stale FY2025 OCF base of \$27.5 billion (OBSERVED, A) (Q3 FY2026 OCF of \$14.2 billion suggests the true figure is lower on a current run-rate; labeled as stale). Principal coverage is 12.2x on the same stale base. The flagged complex is the supplier-financing platform: the AI XPV platform with Apollo and Blackstone (announced June 2026; initial \$35 billion supporting more than 1 GW of Anthropic compute; target above 20 GW through 2028; full commitment could require hundreds of billions) (OBSERVED, B); the \$60 billion Anthropic debt package (Bloomberg October 2, status update October 5: \$42 billion Class A senior-secured tranche began bank syndication October 5 via BofA/Citi/Morgan Stanley with Broadcom residual-value support, plus an \$18 billion Blackstone-led junior tranche with no Broadcom guarantee to follow) (OBSERVED, B)/(OBSERVED, C); Broadcom’s separate agreement to lend Anthropic up to \$42 billion via convertible notes per Anthropic’s confidential IPO prospectus

(OBSERVED, B)/(OBSERVED, C) (whether this overlaps the \$60 billion package is unconfirmed: gap); and the more-than-\$50 billion OpenAI “Nexus” financing Broadcom is arranging (Apollo and Blackstone approached; early talks; targeted close by year-end) (OBSERVED, B)/(OBSERVED, C). Customer concentration is the commercial risk: about 40% of net revenue from the top five end customers (OBSERVED, B), with Anthropic to be the largest custom-chip customer in 2027 (5 GW of chip-equivalent deliveries in 2027, 10 GW in 2028) (OBSERVED, B). CDS printed a record 136bp (Goldman, October 9) (OBSERVED, B).

OpenAI, Anthropic, xAI. The private labs are commitments-only and press-grade. OpenAI files no public financial statements: the only primary obligation in the monitor is the expanded \$4.7 billion revolving credit facility (committed, undrawn at close, as of March 31, 2026) (OBSERVED, B), and an undrawn revolver commitment is not debt. Everything else is memo: the \$500 billion Stargate investment intention (aspirational), the \$300 billion-plus Oracle partnership value, \$250 billion of Azure (press-only), \$38 billion of AWS over seven years (press-only), and Broadcom’s more-than-\$50 billion “Nexus” equipment financing (early talks) (OBSERVED, B)/(OBSERVED, C). Revenue context only: roughly \$50 billion annualized run rate at end-September 2026 (FT, citing investor documents) (OBSERVED, B); the widely circulated \$68–70 billion figure was an investor “gross-up” to match Anthropic’s revenue-counting methodology and was wrong (OBSERVED, B); leaked 2025 financials showed roughly \$13 billion revenue with spending significantly higher (OBSERVED, C); early-stage talks to raise \$30 billion-plus at about a \$1.4 trillion valuation (OBSERVED, C); IPO pushed to early 2027 (OBSERVED, C). Anthropic’s primary obligations are \$100 billion of AWS commitments over ten years (up to 5 GW new capacity) plus \$30 billion of Azure compute purchase commitments (OBSERVED, A); the memo-only aggregate compute portfolio is about \$518 billion over a decade across six partners, roughly 80% non-cancelable (OBSERVED, C), and the \$130 billion of primary rows are components of it (no double counting). No S-1 had been filed as of October 9, 2026 (OBSERVED, B); investor day is October 14. xAI has no monitor rows (gap): the \$20 billion financing package (about \$7.5 billion equity including up to \$2 billion from Nvidia, plus up to \$12.5 billion of debt through an SPV that buys Nvidia GPUs and leases them to xAI for five years, debt collateralized by the GPUs not xAI’s balance sheet; Apollo, Diameter, Valor participating) is press-only (OBSERVED, C), and Musk publicly denied the raise (“fake news”), which the SPV structure technically accommodates since xAI itself borrows nothing. In September 2026 xAI was working with Valor Equity Partners to secure \$12 billion in debt for Colossus 2, with some lenders wanting repayment within three years and a cap on total debt raised (OBSERVED, C). OpenAI’s true fixed-obligation stack is the largest unknown in this workstream.

The private labs as a group. OpenAI, Anthropic, and xAI share a structure the rated names do not: massive fixed demand commitments against undisclosed cash flow. Anthropic’s ~\$518B memo portfolio (~80% non-cancelable, (OBSERVED, C)) is the largest single demand number in the dataset, and its financing is explicitly sequenced around the IPO (catalyst 8): the \$42B Broadcom converts and the \$60B debt package both assume a public-currency event. OpenAI’s stack is larger still in press terms and thinner in documentation: the \$500B Stargate intention, the \$300B-plus Oracle partnership value, and the \$38B AWS figure are all memo-grade, while the only primary obligation is a \$4.7B undrawn revolver. The revenue-counting gap matters here: the \$20B OpenAI/Anthropic methodology difference is a principal-versus-agent question under ASC 606, which means reported “revenue” at the labs is not comparable across names without adjusting for it. For credit purposes the labs are demand, not obligors, in the rated names’ filings; their own solvency is unobservable until the S-1s land.

Flagship structures. Table 13 collects the eight structures Table 14 collects the counterparty concentrations. that define the financing frontier: the closed SPV (Beignet), the vendor-arranged platforms (XPV, Nexus, the Anthropic package), the talks-stage SPVs (Oracle chip vehicle, Amazon’s Grace Blackwell vehicle), and the large guarantees (SB Energy, xAI). Three are closed or in syndication; five are talks-stage or press-only. The table is the empirical basis for the circularity analysis in Sections 6.2 and 6.6.

Table 13. Flagship financing structures of the AI buildout.

Structure	Amount	Status (Oct 2026)	Sponsor economics	Grade
Meta Beignet / Hype- rion SPV	\$27.3B notes	Closed Oct 2025; notes at ~91c	Meta 20% / Blue Owl 80%; Meta bears construc- tion/overrun risk; 20yr lease-back	(OBSERVED, B)
Broadcom AI XPV platform	\$35B initial; >20 GW target	Announced Jun 2026	Broadcom + Apollo + Blackstone; finances An- thropic/OpenAI compute on Broadcom chips	(OBSERVED, B)
Broadcom \$60B An- thropic package	\$42B Class A + \$18B junior	Class A syndication began Oct 5, 2026 (BofA/Citi/MS)	Broadcom residual-value sup- port on Class A; junior has no Broadcom guarantee	(OBSERVED, B)/(OBSERVED, C)
Broadcom >\$50B OpenAI “Nexus”	>\$50B	Early talks; targeted close YE2026	Vendor arranges cus- tomer financing for ven- dor’s co-developed chips; Apollo/Blackstone ap- proached	(OBSERVED, B)/(OBSERVED, C)
Oracle chip SPV	Undisclosed (tens of \$B)	Negotiations (WSJ Oct 7)	Apollo + Goldman; private in- vestors fund vehicle that buys chips, leases back to Oracle	(OBSERVED, B)/(OBSERVED, C)
Amazon \$8B Grace Blackwell SPV	\$8B	Discussions	Amazon retains no equity; out- side investors to 10%; IG- rated SPV bonds on Amazon’s AA credit	(OBSERVED, B)
Nvidia \$105B SB En- ergy backstop	\$105B max no- tional	Aug 2026 subsequent event	Nvidia backstops obligor’s fa- cility; Tier E at weight 1.0 (ceiling)	(OBSERVED, A)
xAI \$20B GPU SPV	~\$7.5B equity + to \$12.5B debt	Press-only; denied by Musk	SPV buys Nvidia GPUs, 5yr leaseback; debt collateralized by GPUs, not xAI	(OBSERVED, C)


Source: AI leverage monitor, vintage 2026-10-11 (Workstream A). Memo/press-only rows are labeled; none enter the Tier aggregates.

The cross-company read is straightforward. Principal coverage (OCF over next-12-month principal) is computable only where the monitor carries a maturity schedule: Meta 47.4x, Microsoft 19.8x, Broadcom 12.2x (stale base), Oracle 4.4x, CoreWeave about 1.0x. The hyperscalers’ recognized debt is a small fraction of their economic stacks, and the binding constraints are the uncommenced-lease pipeline and the supplier-arranged financing wave, which shifts new obligations into structures where recognition depends on VIE and sale-leaseback assessments that outsiders cannot verify until terms are filed.

Table 14. Counterparty concentrations by company.

Company	Top counterparties	
Microsoft	Public investment-grade bondholders; no bank or private-credit concentration disclosed; uncommenced-lease counterparties undisclosed (gap)	
Amazon	Bond market (~\$50B Mar 2026 + \$25B Jul 2026 offerings); Anthropic linkage is equity/compute, not debt	
Alphabet	Public equity and bond markets (Jun 2026 raises); fixed-commitment counterparties undisclosed (gap)	
Meta	IG bondholders; Beignet/Hyperion lenders (PIMCO ~\$18B anchor, BlackRock funds ~\$3B); BlackRock 1 GW El Paso venture	
Oracle	Bond market (\$18B Sep 2025 sale); JPM/MUFG on the stuck \$38B construction debt; Apollo/Goldman on chip SPV talks	Source: AI
CoreWeave	DDTL 4.0 lenders (undisclosed); HPS (DDTL V); Carlyle (SPV positions); \$2.5B revolver banks; Jane Street \$1B; Applied Digital (~\$7B leases)	
Nvidia	Supplier side: SB Energy is the guarantee obligor; six private-credit partners on the \$500B platform	
Broadcom	IG bondholders; customer concentration: ~40% of net revenue from top five end customers; Anthropic as counterparty (converts + packages)	
Anthropic	AWS / Azure / Broadcom (commitments and financing)	
OpenAI	Revolver syndicate (JPM, Citi, GS, MS, WFC, Mizuho, RBC, SMBC, UBS, HSBC)	

leverage monitor, vintage 2026-10-11 (company-metrics.csv). Grades (OBSERVED, A)/(OBSERVED, B) except xAI (no monitor rows: gap).




Data as of Jun-Aug 2026 filings; OCF is trailing twelve months (vintage 11 Oct 2026). Oracle uses FY2026 OCF; Broadcom uses FY2025 (stale base). Source: Company 10-K/10-Q filings via the AI leverage monitor.

Note: Tier E is a stock of multi-year undiscounted commitments; the ratio is a burden scale, not a default clock.

Figure 2. All-in fixed-obligation burden: Tier E / operating cash flow, by company. Oracle at 13.9x (14.8x with the \$288B uncommenced update) is the cohort's highest by a wide margin; CoreWeave at 11.9x is second.

Source: AI leverage monitor, vintage 2026-10-11 (company-metrics.csv). OCF bases as in Table 11; Broadcom on the stale FY2025 base. Placeholder figure; coordinator to build with chart_style.qa_check.



Debt principal plus disclosed lease, purchase, and commitment schedules where available (vintage 11 Oct 2026). Amazon, Alphabet, and Nvidia carry no maturity schedule in the monitor.
Source: Maturity wall, AI leverage monitor.
Note: the binding near-term constraint is the uncommenced-lease ramp and 2027 compute commitments, not bond maturities.

Figure 3. Debt-maturity wall: scheduled debt principal by year, where the monitor carries a schedule. Microsoft, Meta, Oracle, Broadcom, and CoreWeave shown; Amazon, Alphabet, and Nvidia have no debt-maturity schedule in the monitor (gap).

Source: AI leverage monitor, maturity_wall.csv, vintage 2026-10-11 (OBSERVED, B). Placeholder figure; coordinator to build with chart_style.qa_check.

Table 15. Liquidity buffers against the 2027 fixed-obligation burden (\$B).

Borrower	Cash-like assets	2027 base burden	Buffer note
Alphabet	242.5	9.0	Buffer exceeds burden by an order of magnitude
Microsoft	~83	72.0	Positive FCF after capex; buffer covers burden
Amazon	~100	53.6	Buffer covers burden about twice over
Meta	90.3	91.4	Buffer covers the 2027 adverse-case shortfall (~\$6.4B) many times over
Oracle	~19	21.2	Thin relative to burden growth; market access required
CoreWeave	~1.7 + 5.6 undrawn DDTL + 2.0 revolver	10.6	Without market access, base 2027 coverage is 0.66x

Source: Workstream C coverage model (liquidity buffers; (RECONSTRUCTED, B)); 2027 base burden from Table 21 ((MODELLED, C)). Buffers are not in the coverage ratio; they are the second line of defense the ratio omits. Table 15 is read alongside Table 21.

6.2 Financing mechanisms and accounting treatment

The bottom line of the mechanism catalog: every one of the seven structures is compliant with current US GAAP and IFRS as far as the evidence shows; none of the historical tricks (Enron’s 3% rule, pre-842 off-balance-sheet operating leases, SFAS 140 repo sales) survives the current standards. The live opacity is at classification boundaries and in private markets, not in non-compliance. Table 16 gives the structure-by-treatment verdicts; the text that follows states for each mechanism whether it is genuinely off-balance-sheet or merely off-debt, and what outsiders cannot see. Table 17 maps disclosure quality by mechanism.

SPV/VIE. Under ASC 810-10, the vehicle is consolidated by its primary beneficiary: the party with (a) the power to direct the activities that most significantly affect the entity’s economic performance and (b) the obligation to absorb losses or right to receive benefits that could be significant. If the operating company is the primary beneficiary, the SPV’s debt comes onto its balance sheet and the “off-balance-sheet” label fails. The assessment is judgmental, and judgment is where the opacity lives: what outsiders cannot see until filings appear is the power analysis, related-party pricing of the lease, and cross-default or completion triggers. CoreWeave’s financing SPVs are the counterexample that proves the rule: they consolidate, so the debt is on-balance-sheet. Beignet shows the residual economics even when non-consolidation holds: a roughly \$2.37 billion booked equity stake against a disclosed \$45.99 billion maximum exposure (OBSERVED, C), plus a 20-year lease commitment and overrun absorption. Off-balance-sheet debt is not off-risk.

Sale-leaseback. Under ASC 842-40, the transfer must first qualify as a sale under ASC 606 (the buyer must obtain control). If it qualifies, the seller-lessee derecognizes the asset, recognizes gain or loss, and records a right-of-use asset and lease liability for the leaseback. If it does not (e.g., the seller retains control through a non-fair-value repurchase option, or the leaseback covers substantially the asset’s remaining life with fixed-price repurchase features), the transaction is a financing: proceeds are a financial liability under ASC 470, the asset stays on the books, and nothing leaves the balance sheet. The opacity is in the sale qualification: Amazon’s proposed \$8 billion chip SPV (Amazon retains no equity, leases the chips back for continued AWS use) turns on whether the transfer of thousands of already-installed GPUs constitutes a true sale under ASC 606, a determination outsiders cannot make from press reports.

Residual-value guarantee. Two places in the accounts: ASC 460 requires the guarantor to recognize a liability at fair value at inception and disclose the maximum potential undiscounted payments; ASC 842 requires the lessee to include in its lease liability amounts probable of being owed under an RVG. The same economic promise can therefore appear twice (guarantor’s ASC 460 liability and lessee’s lease liability), which is correct gross presentation, not double counting, provided each party reports its own obligation. The monitor’s Tier E carries guarantees at maximum notional (weight 1.0) precisely because the fair-value figure understates the tail. Broadcom’s residual-value support on the \$42 billion Anthropic Class A tranche is the live instance.

Operating lease after ASC 842. The lessee recognizes a right-of-use asset and lease liability at the present value of future payments, on the balance sheet, at commencement. Before commencement: genuinely undisclosed-by-the-balance-sheet but disclosed in footnotes, which is where the cohort’s largest numbers live (\$329.1 billion Microsoft, \$279.0 billion Meta, \$260–288 billion Oracle, \$137.2 billion Amazon, \$85.2 billion Alphabet, \$35.5 billion CoreWeave, \$25.0 billion Nvidia (OBSERVED, A)). These are undiscounted totals over 7–30 year terms; they convert to on-balance-sheet liabilities automatically as leases commence, with no new financing decision required. That automatic conversion is the mechanism to watch: today’s footnote becomes tomorrow’s liability on a pre-signed

Table 16. Financing mechanisms and their treatment under US GAAP and IFRS.

Structure	Primary standard	Balance-sheet treatment	Disclosure required	re-	Genuinely off-BS?
SPV / VIE	ASC 810-10; IFRS 10	Consolidated by primary beneficiary (power + economics); otherwise off-BS	VIE involvement incl. maximum exposure to loss, even if unconsolidated (ASC 810-10-50)		Only if non-consolidation holds; economics stay via equity, guarantees, leases
Sale-leaseback	ASC 842-40 (+ ASC 606 sale test); IFRS 16	Qualifying: derecognize asset, recognize ROU asset + lease liability. Failed sale: financing liability (ASC 470)	Sale-leaseback gains/losses and lease terms (ASC 842-40-50)		Never off-liability: lease liability always recognized
Residual-value guarantee	ASC 460 (guarantee) + ASC 842-30-30-1 (lessee)	Guarantor: liability at fair value at inception; lessee: probable RVG amounts in lease liability	Maximum potential undiscounted payments (ASC 460-10-50)		Off-debt until triggered; fair value understates the tail
Operating lease (post-842)	ASC 842-20; IFRS 16	ROU asset + lease liability at PV of payments, at commencement	Uncommenced leases: nature of commitment (ASC 842-20-50-10)		Commenced: on-BS. Un-commenced: disclosed only, converts automatically
Finance lease	ASC 842 (classification at 842-10-25-2); IFRS 16	On-BS; front-loaded expense (amortization + interest)	Lease liability split current/non-current		Neither off-BS nor off-debt; economically debt-like
Purchase commitment	ASC 440-10-50-1; executory	Not recognized until performance; loss only if onerous (ASC 450-20)	Unconditional purchase obligations: nature, amounts, timing		Genuinely off-BS until performance; off-debt even then (payables, not borrowings)
Supplier-arranged customer financing	Combination: ASC 606/460/321/326 + ASC 842/440	Each party books its own leg (vendor revenue, guarantee, equity/notes; customer lease/commitment; lender loan)	Visible only by assembling multiple parties' filings		Vendor: customer's debt genuinely off vendor's BS. System-level leverage visible nowhere in one filing

Source:

Workstream A mechanism catalog (US GAAP / IFRS treatment). No non-compliance is alleged for any cohort structure; all statements concern what the standards require or permit. Instances in the cohort are given in the text.

Table 17. Disclosure quality by mechanism: what outsiders can and cannot see.

Mechanism	What outsiders can see	What outsiders cannot see until filings/terms appear
SPV/VIE	VIE involvement and maximum exposure (ASC 810-10-50); critical audit matters	The power analysis; related-party pricing of the lease; cross-default/completion triggers
Sale-leaseback	Gains/losses and lease terms (ASC 842-40-50); failed sales as financing liabilities	Whether the transfer qualifies as a true sale under ASC 606 (the decisive fact)
Residual-value guarantee	Maximum potential undiscounted payments (ASC 460-10-50); fair value at inception	The stress-case payout: fair value at inception is rarely informative about tail outcomes
Operating lease (post-842)	Recognized liability split current/non-current; uncommenced totals, windows, term ranges	Commencement profiles (how much per year); pricing of individual leases
Finance lease	Liability on balance sheet; expense timing	Nothing material: classification changes presentation, not recognition
Purchase commitment	Totals and (usually) amounts by maturity year	Composition of aggregate portfolios (the \$707B Alphabet figure); take-or-pay triggers
Supplier-arranged financing	Each party's own leg in its own filing	The loop: no single filing consolidates vendor revenue, customer obligation, and lender exposure

Source:

Workstream A mechanism catalog. The right-hand column is where the next vintage's footnote work goes.

schedule. The missing disclosure for most of the cohort is the commencement profile (how much per year).

Finance lease. On-balance-sheet like an operating lease, but with front-loaded expense (amortization of the ROU asset plus interest on the liability) versus straight-line single lease cost. Finance lease liabilities are routinely included in debt-like measures and credit agreements; several borrowers disclose finance-lease principal payments inside capex figures. Classification as finance versus operating changes expense timing and EBITDA presentation, not recognition. Microsoft's \$66.6 billion of finance lease liabilities is the cohort's largest (OBSERVED, A).

Purchase commitment. Executory contracts are not recognized; ASC 440-10-50-1 requires disclosure of unconditional purchase obligations (nature, amounts, timing). A loss is recognized only if the contract becomes onerous. Genuinely off-balance-sheet until performance, and off-debt even then (they become payables or accrued liabilities, not borrowings). The monitor's Tier E adds these to the fixed stack because credit analysis treats unavoidable cash outflows as obligations regardless of caption. The risk the disclosure does not capture is take-or-pay economics where the buyer must pay even if it cannot use the capacity (Nvidia's \$29 billion cloud take-or-pay for R&D is the named example) (OBSERVED, A). Several cohort members disclose only a total rather than an annual schedule (the \$707 billion Alphabet portfolio's composition is the priority gap).

Supplier-arranged customer financing. There is no single ASC; it is a combination of the mechanisms above. For the vendor: revenue under ASC 606 on the chip sale (note the principal-versus-agent question behind the \$20 billion OpenAI/Anthropic revenue-counting gap (OBSERVED, B)), any guarantee under ASC 460, any equity stake under ASC 321/323, convertible notes receivable as financial assets with credit-loss assessment under ASC 326. For the customer: a purchase commitment or lease as above. For the lender: a loan receivable. Each party's books can be individually compliant while the system-level leverage (vendor revenue depends on lender-funded

customer purchases; lender collateral is the vendor’s depreciating product) is visible nowhere in any single filing. The opacity is structural, not accounting fraud: no filing consolidates the loop. It is also invisible to the related-party tripwire: the supplier is not a “related party” under ASC 850 merely by financing its customer, so related-party disclosures do not fire even though the economics carry the same information asymmetry.

A general monitoring principle follows from the Beignet example. Under PCAOB AS 3101, critical audit matters flag exactly the judgmental assessments (VIE primary-beneficiary tests, sale-leaseback sale qualification, guarantee fair values) where structuring risk lives. For an outsider tracking AI SPVs, the CAM disclosures in 10-Ks are the highest-signal observable available: they tell you which structures the auditor itself found most judgmental, before any restatement or enforcement action. The absence of a CAM on a material SPV is informative in the other direction. This is also why the Greensill auditor-shopping signal does not transfer: the AI SPVs in evidence are audited by the Big Four, and EY’s Beignet flag is the system working as designed.

6.3 Historical episodes: forensic summaries

The bottom line of the six briefs: every case is a different mechanism of concealed risk, which is why the resemblance grades differ by case rather than by era. Figure 4 places the episodes on a common timeline with the AI buildout’s financing milestones.




Figure 4. Historical timeline, 1999–2026: the six comparison episodes against AI buildout financing milestones. Enron (2001), WorldCom/Tyco (2002), GFC (2007–2009), Lehman Repo 105 (2008), Greensill (2021), and the AI chip-financing wave (2025–2026).

Source: Workstream B case briefs; AI milestones from the leverage monitor vintage 2026-10-11. Placeholder figure; coordinator to build with `chart_style.qa_check`.

What the timeline shows, read as a whole, is a regulatory ratchet: each historical episode was followed by a rule-writing response (Sarbanes-Oxley and FIN 46 after Enron, WorldCom, and Tyco; Dodd-Frank, Basel III, and risk retention after the GFC; the FASB/IASB repo-accounting review after Lehman; FINMA enforcement and the Treasury Committee report after Greensill), and each response closed the specific device that had just failed. The AI buildout’s financing milestones cluster in 2025–2026 with no regulatory response yet, which is why the resemblance matrix’s accounting verdicts matter: they test whether the existing ratchet already covers today’s structures. The answer it returns is mostly yes for the historical tricks and no for the two live gaps (private-credit disclosure, commitment classification), which are the ones the next rule-writing cycle, if any, would need to address.

Enron (2001). Enron, a Houston energy trader, collapsed after a November 2001 restatement revealed that roughly two dozen related-party SPEs had hidden debt, fabricated earnings, and enriched executives; it filed for Chapter 11 on December 2, 2001 (OBSERVED, A). The device was the SPE under the “3 percent rule”: Chewco’s supposed independent equity was funded by Enron itself through a guaranteed loan and managed by an Enron employee; LJM1/LJM2 were Fastow-run “hedge” counterparties capitalized with Enron’s own stock, so Enron was hedging its own losses with its own equity; the Raptor vehicles, capitalized with Enron stock, failed when the stock fell (OBSERVED, A) (Powers et al., 2002). Mark-to-market accounting (SEC permission obtained January 1992) (OBSERVED, A) let Enron book the net present value of long-dated deals as immediate earnings, opening a persistent gap between reported income and cash flow. The restatement cut net income by \$28 million (1997) to \$248 million (1999) per year and added \$561–711 million of debt per year over 1997–2000 (OBSERVED, A); a separate \$1.2 billion equity charge covered the Raptors (OBSERVED, A). Fastow took at least \$30 million from the LJM partnerships; Skilling and Lay were convicted in 2006; Congress responded with Sarbanes-Oxley and FIN 46 (OBSERVED, A). Warning signs visible beforehand: related-party footnotes as a standing feature, operating cash flow persistently below net income, deteriorating ROIC while the stock rose, complexity as narrative cover, Fastow’s dual role with a board conflicts waiver, and stock-financed hedges (OBSERVED, B)/(RECONSTRUCTED, B).

WorldCom (2002). WorldCom inflated earnings by reclassifying ordinary operating expenses as capital expenditures: \$3.852 billion of line-cost transfers in 2001–Q1 2002 disclosed June 25, 2002 (“not made in accordance with generally accepted accounting principles”), amended by the SEC to about \$9 billion reaching back to 1999, with total asset inflation around \$11 billion (OBSERVED, A)/(RECONSTRUCTED, B) (U.S. Securities and Exchange Commission, 2002b; Breeden, 2003). Line costs (usage-based payments to other carriers) were moved by journal entry into capital accounts at CFO Scott Sullivan’s direction, understating expense and overstating assets; the SEC alleged the fraud was designed to meet analyst estimates (OBSERVED, A). Ebbers had separately borrowed roughly \$400 million from the company, much of it for stock margin calls (RECONSTRUCTED, B). WorldCom filed for Chapter 11 on July 21, 2002, then the largest US bankruptcy; Ebbers was convicted in 2005 and sentenced to 25 years (OBSERVED, A). Warning signs: the line-cost-to-revenue ratio diverging from peers (AT&T, Sprint) from 2000, the thin “prepaid capacity” justification, revenue growth stalling while earnings held, serial-acquisition purchase-accounting latitude, and executive borrowing disclosed in proxies (OBSERVED, B)/(RECONSTRUCTED, B).

Tyco (2002). Tyco was executive looting plus earnings inflation through acquisition accounting, not a hidden-leverage episode; the brief records this directly rather than stretching the comparison. Kozlowski, Swartz, and Belnick were charged with failing to disclose millions in low-interest company loans and (for the first two) forgiving tens of millions of their own loans without authorization (OBSERVED, A) (U.S. Securities and Exchange Commission, 2002a); the DA’s prosecution alleged roughly \$150 million looted in unauthorized bonuses and forgiven loans plus \$430 million in secret stock sales (RECONSTRUCTED, B). Both were convicted in 2005 on 22 counts and sentenced to $8\frac{1}{3}$ to 25 years (OBSERVED, A). Separately, the SEC charged Tyco itself with overstating results by at least \$1 billion from 1996–2002 via improper acquisition accounting and reserve management, including the \$567 million ADT “dealer connection fee” transaction that lacked economic substance; Tyco paid a \$50 million penalty (OBSERVED, A) (U.S. Securities and Exchange Commission, 2006). Detection came through a tax-evasion probe of Kozlowski’s art purchases, the SEC’s accounting investigation, and press coverage of executive excess (OBSERVED, B).

The 2007–2009 global financial crisis. The crisis was a leverage and liquidity crisis transmitted through a securitization chain that had moved credit intermediation outside the regulated banking system (Financial Crisis Inquiry Commission, 2011; United States Senate Permanent Subcommittee on Investigations, 2011). Banks funded long-term mortgage assets with short-term wholesale funding through off-balance-sheet SIVs and conduits issuing asset-backed commercial paper; when asset values fell, thin equity was wiped out and losses rebounded onto sponsors through liquidity backstops. The FCIC (January 27, 2011) concluded the crisis was avoidable, caused by regulatory failure, governance breakdowns, excessive borrowing, and accountability breaches; collapsing mortgage-lending standards and the securitization pipeline “lit and spread the flame of contagion” (OBSERVED, A). The Levin-Coburn report (April 13, 2011, 639 pages, 150+ interviews, 56 million pages of documents) found high-risk products, undisclosed conflicts, and regulatory failure, with case studies on WaMu (\$300 billion, the largest US bank failure), OTS inaction, inflated ratings, and Goldman Sachs designing CDOs while shorting the mortgage market (OBSERVED, A). The major investment banks operated with extraordinarily thin capital (less than a 3% asset-value drop could wipe out a firm) (OBSERVED, B); the US shadow banking system was estimated near \$20 trillion as of March 2008, larger than traditional banking liabilities (OBSERVED, A) (Pozsar et al., 2010). AIG’s Financial Products unit concentrated correlated risk by selling CDO protection. When ABCP froze in August 2007 and repo haircuts rose, maturity transformation unwound: funders ran, fire sales depressed collateral, solvency crises followed. The regulatory outcome was Dodd-Frank, Basel III, risk retention, and the FSB shadow-banking framework (OBSERVED, A). Warning signs: originate-to-distribute with no skin in the game (WaMu/Long Beach securitized at least \$77 billion subprime 2000–2007 (OBSERVED, A)), the issuer-pays rating conflict, footnote-disclosed SIVs with sponsor backstops (Citigroup’s ran to roughly \$100 billion (RECONSTRUCTED, B)), sub-3% dealer equity cushions, overnight funding of long-term assets, and volume-linked compensation (OBSERVED, B).


Lehman Brothers, Repo 105 (2008). In the quarters before its September 15, 2008 bankruptcy, Lehman used Repo 105 to remove tens of billions of dollars of assets from its balance sheet at quarter-ends (Valukas, 2010). The device turned on SFAS 140: ordinary repos stay on the balance sheet, but sale treatment was permitted if control was surrendered, tested partly by the repurchase price falling outside 98–102% of cash received. Lehman posted collateral worth 105% or more of cash received, obtained a “true sale” opinion from Linklaters under English law (US counsel would not give it), and booked the transactions as sales: securities left the balance sheet, cash proceeds were recorded as an asset, and the obligation to repay was not recorded as a liability, though Lehman repurchased the securities days later (OBSERVED, B). The examiner’s scale: about \$38.6 billion removed at Q4 2007, \$49.1 billion at Q1 2008, \$50.38 billion at Q2 2008 (OBSERVED, B). Executives felt pressure to reduce leverage for reports; contemporaneous documents described Repo 105 as balance-sheet management with “few, if any” describing any other business purpose; Lehman’s own staff called it an “accounting gimmick” (OBSERVED, B). The examiner found a colorable claim of material misrepresentation and of malpractice against auditor EY, which knew of the practice (OBSERVED, B). No criminal charges followed; the episode became the canonical “GAAP-compliant but economically misleading” case study (RECONSTRUCTED, B). Warning signs: quarter-end saw-tooth leverage patterns, the practice as an open secret on funding desks, leverage targets met exactly at reporting dates, the Linklaters jurisdictional workaround, auditor awareness without disclosure, and collateral-quality migration toward illiquid positions (OBSERVED, B).

Greensill Capital (2021). Greensill, a supply-chain finance firm, collapsed in March 2021 through three layers (Swiss Financial Market Supervisory Authority FINMA, 2023; House of Commons Treasury Committee, 2021). First, recharacterized receivables: Greensill sold SCF notes into four Credit Suisse funds, and the risk character “changed decisively” as future receivables (claims not yet arisen) and doubtful credits entered the bundles; Credit Suisse “had no knowledge or control over how many claims were actually contractually owed” and “relied on the insurance coverage organized by Greensill” (OBSERVED, A). Second, insurance as the load-bearing assumption: when \$4.6 billion of Tokio Marine cover was not renewed, the structure’s premise collapsed (OBSERVED, B). Third, concentration and related-party opacity: heavy exposure to Sanjeev Gupta’s GFG Alliance and to Greensill’s own Wyelands Bank, with SoftBank’s Vision Fund a major Greensill investor (OBSERVED, B). Credit Suisse closed the four funds at short notice (clients had invested around \$10 billion (OBSERVED, A)); Greensill filed for administration; creditor losses were on the order of £2.7 billion (OBSERVED, B); roughly 50,000 jobs were threatened (RECONSTRUCTED, B). FINMA (February 28, 2023) found Credit Suisse had “seriously breached its supervisory obligations” and opened four enforcement proceedings against former managers (OBSERVED, A); the Commons Treasury Committee (July 20, 2021) found David Cameron’s lobbying showed “significant lack of judgement” while Treasury officials “behaved with complete and absolute integrity” (OBSERVED, A). Warning signs: future receivables as a disclosed product feature, single-name GFG concentration, auditor shopping (Saffery Champness after larger firms declined), insurance dependence as a single point of failure, regulatory-perimeter arbitrage, and political access as business strategy (OBSERVED, A)/(OBSERVED, B).

Cross-case patterns. Three patterns cut across the six cases and discipline the comparison. First, footnote opacity: Enron’s related-party footnotes, the GFC’s SIV disclosures, and today’s uncommenced-lease and commitment footnotes are all technically disclosed positions whose economic meaning requires assembly work; “disclosed” and “transparent” are not synonyms. Second, guarantee dependence: AIG’s credit protection, Tokio Marine’s \$4.6 billion of Greensill cover, and today’s vendor residual-value guarantees and offtake agreements are all structures whose solvency rests on a third party’s promise surviving stress untested. Third, the perimeter problem: SIVs, Greensill, and today’s private-credit SPVs all intermediate bank-like credit outside bank disclosure regimes; the regulator sees the banks’ warehouse lines but not the risk build-up inside the funds. What none of the six cases had, and what the AI evidence does not show, is more important: no case combined real, audited, productive physical assets with fully disclosed commitments and no misconduct findings, which is why four of the six grades stop at Partial and two stop well short.

6.4 The resemblance matrix

Table 18 is the paper’s main results table: the standardized six-case comparison. Figure 5 renders the grades as a heatmap. The reading of the table, in one paragraph per case, follows the uniform three-question verdict (genuine resemblance; misleading aspect; accounting verdict). The aggregate: no Strong grades, four Partials, one Weak, one Misleading. This is not an Enron/WorldCom/Tyco story; the GFC parallel is structural; Greensill is the closest live single-firm parallel.



Grades from the structured resemblance rubric (Workstream B): genuine resemblance vs misleading aspect vs accounting verdict per case.
Source: Resemblance matrix, six primary-source case briefs.
Note: no case grades 'Strong'; the closest parallels are structural (GFC funding architecture) and mechanical (Greensill private credit), not fraud.

Figure 5. Resemblance heatmap: the six historical cases graded Strong / Partial / Weak / Misleading against the AI financing evidence. No case grades Strong; Tyco is a null finding.

Source: Workstream B resemblance-matrix.csv. Placeholder figure; coordinator to build with chart_style.qa_check.

Table 18. The resemblance matrix: six historical cases against the AI financing evidence (main results table).

Case	Historical mechanism	AI parallel	Grade	Verdict (genuine resemblance / misleading aspect / accounting verdict)
Enron (2001)	Related-party SPEs (Chewco, LJM, Raptors) hid debt and fabricated earnings via mark-to-market; restatement cut net income \$28–248M/yr and added \$561–711M/yr debt, 1997–2000	Supplier-arranged customer financing (Broadcom/OpenAI ~\$50B+) and disclosed mega-commitments (MSFT \$329.1B un-commenced leases) with vendor-embedded opacity	Partial	Genuine: vendor financing its own customer's purchases creates related-party-style opacity; disclosed mega-commitments warehouse obligations outside debt the way Enron's SPEs kept them out of the debt line. Misleading: Enron hid <i>losses</i> and fabricated earnings with criminal intent; AI SPVs fund <i>real disclosed</i> assets; no Fastow analogue; ASC 810/842 block the old tricks. Accounting: compliant; ASC 810 would consolidate Enron-style SPEs today; live gap is large non-debt fixed commitments (ASC 842/460) excluded from leverage ratios and vendor finance escaping ASC 850 related-party tripwires. (OBSERVED, A)

Table 18 (continued).

Case	Historical mechanism	AI parallel	Grade	Verdict
WorldCom (2002)	Reclassified \$3.852B of line-cost operating expenses as capex (2001–Q1 2002); amended to ~\$9B; fabricated profitability to meet estimates	Capitalization and impairment judgments on AI infrastructure capex (data centers, GPUs) under ASC 360 / IAS 36	Weak	Genuine (narrow): only the forward-looking channel: if AI hardware obsolesces faster than depreciation schedules assume, the impairment question WorldCom inverted becomes live in milder form (the \$38B undistributable Oracle construction debt is the datum to watch). Misleading: WorldCom fabricated profit with journal entries; AI capex is real cash on real auditable assets; no allegation of expense-capitalization fraud; income-statement fraud versus balance-sheet commitment question are different problems. Accounting: unambiguously GAAP-violating then and now; for AI the live questions are useful-life estimates, impairment triggers (judgmental but audited), and the lease-versus-service classification boundary. (OBSERVED, A)
Tyco (2002)	Executive looting (~\$150M unauthorized loans/bonuses; convictions 2005) plus \$1B+ earnings inflation via acquisition accounting and ADT connection fees	None evidenced: no insider self-dealing alleged in AI financing structures	Misleading	Genuine: essentially none; thinnest threads are serial-acquisition accounting complexity and related-party disclosure hygiene, neither specific to AI. Misleading: Tyco was insider looting plus earnings fabrication; AI financing is outsiders funding real construction; invoking Tyco smuggles an insider-corruption narrative with zero evidentiary basis. Accounting: null finding; Tyco’s purchase-accounting violations (ASC 805) have no AI analogue in evidence; no sponsor-side executive participation in AI SPVs is alleged. (OBSERVED, A)

Table 18 (continued).

Case	Historical mechanism	AI parallel	Grade	Verdict
GFC (2007–2009)	Securitization chain funded long-term mortgage assets with short-term wholesale funding via off-BS SIVs/conduits; ~\$20T US shadow banking (2008); dealer equity under 3%; FCIC: avoidable	Private-credit and vendor-arranged SPV funding of the AI build-out outside the bank perimeter; securitization of AI lease/capacity cash flows; maturity transformation in construction finance	Partial	Genuine (strongest structural parallel): shadow-bank intermediation, SPV layering, maturity transformation, originate-to-distribute incentives (vendor financing own customers), and footnote-archaeology opacity: the same funding architecture. Misleading: subprime mortgages were fraudulent consumer credit; AI collateral is productive infrastructure with contracted cash flows; no household leverage, no bank-run dynamics; AI commitments are disclosed where SIV leverage was hidden. Accounting: GFC loopholes closed (ASC 860/810, FIN 46R, risk retention); live AI gaps are the commitment-versus-debt boundary (ASC 842 uncommenced leases, ASC 460 guarantees) and the private-credit disclosure vacuum. (OBSERVED, A)
Lehman Repo 105 (2008)	Repo booked as sale under SFAS 140 via 105% overcollateralization and a Linklaters true-sale opinion; \$38.6B/\$49.1B/\$50.7B removed at Q4 2007/Q1 2008/Q2 2008 quarter-ends to understate leverage; no business purpose found	Disclosed AI commitments (MSFT \$329.1B, GOOGL \$707B, META \$349.3B, ~~AMZN~~ \$137.2B) excluded from debt/leverage ratios despite being economically firm; structuring around lease-versus-service and consolidation boundaries	Partial	Genuine: the GAAP-compliant-but-leverage-misleading mechanism transfers directly: compliant commitments excluded from leverage metrics understate the fixed-payment burden; threshold-engineering around ASC 842/810 boundaries is a live structuring risk. Misleading: Repo 105 was deliberate temporary reporting-date deception reversed within days; AI commitments are persistent, disclosed, and purposeful; no intent evidence; conflating them smuggles a fraud presumption. Accounting: compliant as far as evidence shows; the specific trick is closed post-crisis; genuine arbitrage sits at classification boundaries (service vs. lease, purchase vs. debt) and the ASC 850 related-party tripwire not firing for vendor finance. (OBSERVED, B)

Table 18 (continued).

Case	Historical mechanism	AI parallel	Grade	Verdict
Greensill (2021)	Future-receivables SCF notes sold into four Credit Suisse funds (~\$10B); risk character drifted to doubtful credits; \$4.6B Tokio Marine cover non-renewed; FINMA: serious supervisory breach	Private-credit SPV funding of AI capacity (Broad-com/OpenAI, Apollo/Goldman Oracle SPV talks); financing of not-yet-built capacity on utilization assumptions; guarantee/offtake dependence; arranger opacity with no public disclosure regime	Partial	Genuine (closest private-credit parallel): non-bank intermediation outside the perimeter; financing future (not-yet-existing) economic claims; guarantee/insurance dependence; arranger opacity; single-tenant/vendor concentration. Misleading: Greensill drew misconduct findings and misrepresented assets; AI SPVs hold real physical infrastructure; no political-capture dimension; AI paper sold as credit risk to sophisticated buyers, not as cash-like product. Accounting: compliant as evidenced; genuine opacity in who consolidates the SPV (ASC 810/IFRS 10 contestability) and in the private-market disclosure vacuum; Greensill is the precedent for how those fail. (OBSERVED, A)

Source: Workstream B resemblance-matrix.csv; grades per the Section 5.5 rubric.

Why the grades stop where they do. The grade boundaries are deliberate. The GFC stops at Partial (strong end) rather than Strong because the asset-quality, consumer-leverage, and bank-run dimensions have no AI counterpart in the evidence, and because AI commitments are disclosed where SIV leverage was hidden; “2008 again” would overstate the case. Greensill stops at Partial (strong end) rather than Strong because it drew misconduct findings, misrepresented assets, and a political-capture dimension with no AI analogue; the evidenced parallel is private-market opacity and future-claims economics, which is real but less severe. Enron stops at Partial rather than Strong because there is no intent evidence, the assets are real and disclosed, no executive sits on both sides of the transactions, and ASC 810/842 block the specific tricks. Lehman Repo 105 stops at Partial rather than Strong because Repo 105 was temporary, intentional reporting-date deception with no business purpose, while AI commitments are persistent, disclosed, and purposeful; the transferable lesson is “compliant but leverage-misleading,” not deception. WorldCom is Weak rather than Misleading because the forward-looking impairment channel is genuinely live: if AI hardware obsolesces faster than depreciation schedules assume, the question WorldCom inverted becomes real in milder form. Tyco is Misleading rather than Weak because no element of the AI evidence maps onto the Tyco mechanism at all; it is recorded as a null finding, not a weak resemblance.

Table 19. Warning signs then and now: historical signals mapped to today’s observables.

Historical warning sign	Case	Today’s analogue	Observable
Related-party footnotes as a standing feature	Enron	Vendor-arranged financing opacity without legal related-party status	Broadcom XPV/Nexus/Anthropic filings; ASC 850 tripwire not firing
OCF persistently below net income	Enron	CoreWeave H1 OCF flattered by \$9.7B deferred-revenue pool vs. \$49M operating loss	Deferred-revenue growth vs. operating income
Complexity as narrative cover	Enron	Multi-layer SPVs with guarantees	Beignet structure; Oracle chip SPV terms when filed
Originate-to-distribute, no skin in the game	GFC	Vendor financing its own customers’ purchases	XPV, Nexus, Nvidia \$500B platform
Rating-agency conflict (issuer-pays)	GFC	IG ratings on SPV paper resting on sponsor credit	Beignet A+ rating vs. notes at ~91c
Off-BS vehicles with sponsor backstops	GFC	SPVs with sponsor construction/overrun risk	Meta \$45.99B max exposure vs. ~\$2.37B booked equity
Thin capital at dealers	GFC	Neocloud coverage near 1.0x	CoreWeave base 2027 coverage 0.66x
Short-term funding of long-term assets	GFC	Maturity transformation in construction finance	\$38B Oracle construction debt undistributable 6+ months
Future receivables as a product feature	Greensill	Financing not-yet-built capacity on utilization assumptions	Uncommenced leases: MSFT \$329B, META \$279B, ORCL \$260–288B
Single-name concentration	Greensill	Hyperscaler-tenant / GPU-vendor concentration	Broadcom ~40% of revenue from top five customers
Insurance/guarantee dependence	Greensill	RVG and offtake dependence, durability untested	Broadcom RVG on \$42B Class A; NVDA \$105B backstop
Perimeter arbitrage	Greensill	Private credit outside the disclosure regime	BDC 10-Q fair-value marks (catalyst 4)
Capitalization/impairment judgments	WorldCom	GPU useful lives vs. depreciation schedules	Impairment tests if utilization disappoints
Quarter-end window dressing	Lehman	None: commitments are persistent and disclosed	Null (this is why Repo 105 grades Partial)
Executive self-dealing	Tyco/WorldCom	None evidenced	Null (this is why Tyco is a null finding)

Source:

Workstream B case briefs (warning-sign sections) mapped to Workstream A/C observables. Null rows are findings, not omissions: the absence of a Lehman-style or Tyco-style signal is part of the verdict.

What Strong would have required. It is worth stating the counterfactual explicitly, because it shows how far the evidence is from the alarming version of each analogy. The GFC comparison would grade Strong if the AI record showed fraudulent collateral (fabricated contracts or liar-loan equivalents), correlated household-style leverage, or runnable short-term funding at systemic scale; it shows none of the three. Greensill would grade Strong with a misconduct finding against an AI arranger, misrepresented assets in an AI SPV, or political capture; there are none in evidence. Enron would grade Strong with an executive on both sides of a material AI transaction or a restatement hiding losses; there is neither. Lehman would grade Strong with evidence of reporting-date window dressing; AI commitments are persistent by construction. These are not hedges; they are the falsifiers from Section 7.1 restated as grade boundaries, and they are what the next vintage should test first. Table 29 collects the falsifiers in one place.

6.5 Stress tests

The bottom line of the stress tests: the solvency core holds in every case, including the severe one; the periphery does not. Microsoft, Amazon, and Alphabet never breach 1.0x coverage in any scenario; Meta breaches only in 2027 of the severe case and holds a \$90.3 billion cash buffer against a one-year squeeze; Oracle breaches in all four years of the severe case and sits at 1.33x in the base case by FY2029; CoreWeave breaches from 2027 in every case, including the base case. Table 20 gives the coverage ratios; Table 24 summarizes breaches across all nine scenarios; Figure 6 plots the coverage waterfall for the three names where stress binds.

Table 20. Fixed-obligation coverage ratios (OCF / annual burden), 2027–2030.

Borrower	Base				Adverse (OCF –25%)				Severe (OCF –50%, +200bp)			
	2027	2028	2029	2030	2027	2028	2029	2030	2027	2028	2029	2030
MSFT	2.54	2.75	2.99	3.24	1.90	2.06	2.24	2.43	1.27	1.37	1.49	1.62
AMZN	3.01	3.27	3.55	3.86	2.26	2.45	2.66	2.90	1.51	1.63	1.78	1.93
GOOGL	20.62	19.49	18.80	18.41	15.46	14.62	14.10	13.81	10.31	9.75	9.40	9.20
META	1.43	4.11	4.37	4.65	1.07	3.08	3.28	3.49	0.71	2.05	2.18	2.31
ORCL	1.51	1.39	1.33	1.77	1.13	1.04	1.00	1.33	0.75	0.69	0.66	0.87
CRWV	0.66	0.99	1.55	1.98	0.49	0.74	1.16	1.49	0.33	0.49	0.75	0.96


Source: Workstream C coverage model (Modelled; scenario-results.csv). Microsoft and Oracle years are fiscal (June/May year-ends); others calendar. Oracle FY2029 adverse at 1.00x is a watch level, not a clean breach. Alphabet’s \$707B fixed-commitment portfolio is excluded for lack of a schedule (largest unmodeled exposure).

Table 21. 2027 base-case burden composition (\$M).

Borrower	Interest	Leases	Uncomm. ramp	Purch./compute	Principal	Total burden	OCF 2027
MSFT	1,612	11,065		2,194	57,156	0	72,027
AMZN	5,950	13,721		1,429	32,516	0	53,617
GOOGL	5,008	2,578		1,420	0	0	9,006
META	4,351	3,582		1,860	81,650	0	91,442
ORCL	6,893	4,776		5,098	4,436	0	21,204
CRWV	2,805	591		986	0	6,184	10,567

Source: Workstream C coverage model ((MODELLED, C)). OCF 2027 = TTM grown one year per A9 (Meta flat at TTM; growth applies from 2028). Scheduled principal enters only for CoreWeave (amortizing debt, cannot be assumed rolled); IG maturities are assumed refinanced in the base case. Oracle’s uncommenced ramp (+\$5.1B/yr through FY2029, reaching \$15.3B/yr) is the single largest year-on-year burden increase in the model.

Coverage 2027-2030: CoreWeave breaches in all cases; Oracle in the severe case



Modelled, C. Fixed burden = interest + lease payments + uncommenced-lease ramp + purchase/compute commitments + scheduled principal.
Source: Coverage model, Workstream C (vintage 11 Oct 2026).
Note: Microsoft, Amazon, and Alphabet never breach 1.0x in any case (not shown).

Figure 6. Coverage waterfall, 2027–2030: Oracle, CoreWeave, and Meta fixed-obligation coverage under base, adverse, and severe cases. The 1.0x line is the breach threshold. Microsoft, Amazon, and Alphabet (never below 1.27x even in the severe case) are omitted for scale.

Source: Workstream C coverage model (Modelled). Placeholder figure; coordinator to build with
chart_style.qa_check.

Table 21 decomposes the 2027 base-case burden. Two features stand out. First, for Meta the entire 2027 story is one line: \$81,650M of compute commitments against \$91,442M of total burden. Second, for Oracle the uncommenced-lease ramp is already the second-largest burden line in 2027 and grows every year through FY2029, which is why Oracle’s coverage compresses monotonically while every other IG name’s coverage rises.

Three findings deserve emphasis. First, CoreWeave’s sub-1.0x base-case coverage is a structural fact about its model, not a stress result: with \$6.2 billion of 2027 scheduled principal against roughly \$7.0 billion of TTM operating cash flow, every dollar of growth capex and every refinancing must come from markets. The company did exactly that in the first half of 2026 (\$13.5 billion gross debt issuance) (OBSERVED, B). The stress question for CoreWeave is therefore not “can it cover” but “can it keep issuing”; the refinancing-freeze and counterparty-failure scenarios are the binding ones. Second, Oracle is the investment-grade borrower where stress bites. The \$260 billion uncommenced lease book commencing over three fiscal years converts a 1.7x-coverage company into a roughly 1.3x company by FY2029 in the base case (the single largest year-on-year burden increase in the model: +\$5.1 billion per year through FY2029, reaching \$15.3 billion per year). Any combination of OCF disappointment and rate pressure breaches the line, which is consistent with the market signal (261bp CDS, undistributable construction debt). Third, Meta’s risk is concentrated in a single year: \$81.65 billion of compute commitments fall due in 2027 against one year of OCF (OBSERVED, A)/(MODELLED, B). After 2027 the schedule lightens and coverage recovers above 4x. It is a one-year liquidity squeeze, not structural insolvency: Meta’s \$90.3 billion of cash-like assets cover the 2027 adverse-case shortfall (about \$6.4 billion) many times over.

Two further interpretation points from the model deserve emphasis. First, the model excludes dividends and buybacks, and they are the first discretionary buffer: including them would cut Microsoft’s base coverage to about 1.5x and Meta’s 2027 to about 1.05x (MODELLED, C). Management would cut shareholder returns long before missing a fixed payment, so the reported coverage

Table 22. The nine stress scenarios: definitions.

	Scenario	Definition	
(a)	Weaker AI monetization	AI revenue -30% vs. plan; OCF haircut 10–35% by borrower’s AI exposure	
(b)	Falling GPU rental rates	Spot GPU rates -40% ; neocloud revenue repriced; cloud pricing pressure on hyperscalers	
(c)	Declining data-center valuations	DC valuations -25% ; $+150\text{bp}$ secured refinancing cost for CoreWeave/Oracle	
(d)	Higher rates	$+200\text{bp}$ on all refinanced/rolled debt, 2027–2030	
(e)	Construction delay	12-month slip: purchase commitments and lease ramp shift $+1$ year; OCF -10%	<i>Source:</i>
(f)	Energy constraints	Power delivery shortfall; power costs $+15\%$; OCF -5 to -10%	
(g)	Top-3 lab customer default	A top-3 AI-lab customer fails; concentration loss modeled at CoreWeave (OCF -30%) and Oracle (-15%)	
(h)	Refinancing freeze	No new unsecured issuance for 12 months; 2027 maturities repaid from OCF	
(i)	Private-credit counter-party failures	Two major private-credit lenders impaired; DDTL access lost at CoreWeave	

Workstream C coverage model; scenario-results.csv (all outputs (MODELLED, C)).

ratios overstate operating headroom wherever capital returns are large. Second, Alphabet’s fortress coverage (10x–21x across cases) coexists with the \$707B fixed-commitment portfolio excluded for lack of a schedule: it is the largest unmodeled exposure in the exercise, and the coverage ratio is flattering precisely where the disclosure is thinnest.

The scenario sweep (Table 24)

Table 23. Material exposures not in the coverage model.

Exposure	Size	Why it is excluded, and what it would take to matter	
Alphabet fixed/guaranteed portfolio	\$707B \$707,000M	No schedule disclosed. If even 5% fell due in one year it would halve GOOGL coverage, still leaving it above 5x (MODELLED, C)	
Anthropic ~\$518B memo portfolio	~\$518,000M	Press-grade, no S-1; relevant to scenarios (a) and (g) via the propagation matrix	
Dividends and buybacks	MSFT ~\$49B/yr; META ~\$32B/yr	Excluded by design (fixed obligations only). Including them cuts MSFT base coverage to $\sim 1.5\text{x}$ and META 2027 to $\sim 1.05\text{x}$; they are the first buffer cut (MODELLED, C)	<i>Source:</i>
Capex cash drain	AMZN TTM capex ~\$169B vs. OCF \$161B	OCF is pre-capex; the investment program, not fixed obligations, is the true cash drain (OBSERVED, B)	
Oracle prepayment-flattered OCF	\$11.4B customer prepayments in Q1 FY2027	Q1 ran at $\sim 2.9\text{x}$ the FY2026 quarterly average; the model uses the FY2026 base (OBSERVED, B)/(OBSERVED, C)	
CoreWeave deferred-revenue OCF	\$9.7B pool (+ \$790M in Q2)	Positive H1 OCF is substantially a prepayment timing effect, not operating-profit conversion (OBSERVED, B)	

Workstream C coverage-model.md (material limitations). The model’s verdicts are conditional on these exclusions; Section 7.4 shows none of them changes a breach finding. Table 23 lists the material unmodeled exposures.

adds two null findings that matter. The $+200\text{bp}$ refinancing shock alone breaches nobody: IG issuers’ 2027–2030 rollover volumes are small relative to OCF. And the 12-month construction delay paradoxically eases near-term coverage (purchase commitments and lease ramps shift one year out,

at the cost of 10% lower OCF) and breaches nobody: delay is a liquidity relief valve, not a solvency event, though it destroys the growth thesis the debt was raised to fund.

Table 24. Scenario-breach summary: all nine scenarios, six borrowers.

Scenario	Borrower(s) breaching	Breach years	Min coverage	Key driver
(a) Weaker AI monetization	CRWV	2027–2028	0.43x (2027)	AI revenue –30% vs. plan; OCF haircut 10–35% by AI exposure
(b) Falling GPU rental rates	CRWV	2027–2029	0.40x (2027)	Spot GPU rates –40%; neo-cloud revenue repriced, cloud pricing pressure
(c) Declining DC valuations	CRWV	2027–2028	0.63x (2027)	DC valuations –25%; +150bp secured refinancing cost for CRWV/ORCL
(d) Higher rates +200bp	CRWV	2027–2028	0.65x (2027)	+200bp on all refinanced/rolled debt 2027–2030. Null finding: breaches nobody else
(e) Construction delay 12mo	CRWV	2027–2028	0.65x (2027)	12-month slip shifts outflows +1y; OCF –10%. Null finding: eases near-term coverage; breaches nobody else
(f) Energy constraints	CRWV	2027–2028	0.59x (2027)	Power delivery shortfall; power costs +15%; OCF –5 to –10%
(g) Top-3 lab customer default	CRWV	2027–2028	0.46x (2027)	Top-3 AI-lab customer fails; concentration loss at CRWV/ORCL (ORCL min 1.13x in 2029)
(h) Refinancing freeze 12mo	CRWV	2027–2028	0.66x (2027)	No new unsecured issuance 12mo; 2027 maturities repaid from OCF (ORCL min 1.13x in 2027)
(i) Counterparty failures	CRWV	2027–2028	0.55x (2027)	Two major private-credit lenders impaired; DDTL access lost at CRWV
Severe case (OCF –50%, +200bp)	CRWV 2027–2030; ORCL FY2027–FY2030; META 2027	n.a.	CRWV 0.33x; ORCL 0.66x; META 0.71x	Combined shock. Null finding: MSFT, AMZN, GOOGL never breach in any scenario

Source: Workstream C scenario-results.csv (all Modelled, C). “Breach” = coverage below 1.0x. Oracle FY2029 adverse

at 1.00x is a watch level.

Collateral: GPU-secured debt goes underwater on hardware value within two years.

The LTV model (Section 5.4) finds that a CoreWeave-type DDTL starting at 75% LTV is underwater within one to two years on every economic-life assumption from two to five years, and the \$38 billion Oracle construction-debt template (80% starting LTV, 70% equipment / 30% shell) likewise falls below principal within one to two years because the equipment share dominates early depreciation (Table 25; Figure 7). Under no life assumption does hardware collateral alone support these facilities past year two; that is the finding, not a failure of the model. The two-year-life row approximates a forced-sale view: a lender repossessing and selling GPUs after one year recovers roughly 55–65% of the loan balance at 75–80% starting LTV, before fees and illiquidity discounts. The four-to-five-year rows approximate orderly secondary-market value.

Table 25. Collateral LTV paths: loan balance / hardware value by GPU economic life.


Facility / GPU life	Year 0	Year 1	Year 2	Year 3	Year 5	First year > 100%	
<i>F1: CoreWeave-type DDTL (75% starting LTV, 5%/yr amort.)</i>							
2 years	0.75	1.30	6.77	6.43	5.80	Year 1	
3 years	0.75	1.02	1.69	6.43	5.80	Year 1	
4 years	0.75	0.92	1.23	1.98	5.80	Year 2	
5 years	0.75	0.87	1.06	1.40	5.80	Year 2	Source: Workstream C
<i>F2: Oracle construction-debt template (80% starting LTV, 2%/yr amort.)</i>							
2 years	0.80	1.16	2.21	2.23	2.29	Year 1	
3 years	0.80	1.01	1.38	2.23	2.29	Year 1	
4 years	0.80	0.94	1.16	1.52	2.29	Year 2	
5 years	0.80	0.91	1.06	1.28	2.29	Year 2	

collateral model (Modelled, C; collateral-summary.csv). Straight-line depreciation to 10% residual; no additional collateral posting or cash sweep (recovery value in default, where cures have failed). Year-5 LTVs converge as collateral hits the residual floor while the loan still amortizes slowly.

The interpretation is direct. Every GPU-secured facility in the vintage is underwritten on contracted cash flows, not on the hardware: the binding covenant in a stress is the debt-service coverage ratio on contracted cash flows, not the LTV test, and LTV tests exist to force early renegotiation, not to make lenders whole. Recovery value in a default is therefore a function of contract assignment and the borrower’s customer book, not of accelerator resale value. Applied to the market facts: CloudZero spot GPU rates down 30–40% from peaks (OBSERVED, B) compress the orderly-sale recovery further and shorten the effective economic life toward the two-to-three-year rows; the Beignet/Hyperion 2049s at about 91 cents and Oracle 2055s at 77 cents are the bond market pricing exactly this recovery skepticism on long-dated AI paper (OBSERVED, B). For the \$38 billion Oracle construction debt, the shell retains value, but the loan was sized against a campus whose economic value is the contracted lease stream (Oracle as anchor tenant, now under a force-majeure dispute at Project Jupiter): collateral value below principal within two years means any restructuring negotiation starts from the lease contract, not the concrete.

6.6 Propagation and catalysts

The bottom line of the propagation analysis: a neocloud default hits private credit first and hardest; private-credit stress hits the operators’ funding oxygen (CoreWeave’s \$5.6 billion undrawn DDTL);



Modelled, C. Representative CoreWeave-type DDTL GPU facility, 75% starting LTV; straight-line depreciation to 10% residual.
Source: Collateral model, Workstream C.
Note: GPU-secured facilities are underwritten on contracted cash flows, not hardware resale; recovery in default depends on contract assignment.

Figure 7. Collateral LTV erosion curves: F1 (CoreWeave-type DDTL) and F2 (Oracle construction-debt template) under 2–5 year GPU economic lives. The 1.0x line marks underwater on hardware value.

Source: Workstream C collateral model (Modelled). Placeholder figure; coordinator to build with
chart_style.qa_check.

bank underwriting losses transmit to bond markets; and the banking system itself is not the solvency constraint (all 32 US banks cleared the June 2026 severely-adverse test; pre-provision net revenue covers charge-offs 6.4x (OBSERVED, B) per KR-WP-2026-12). Table 26 gives the severity matrix across five shock origins and eight affected nodes; Figure 8 renders it as a network. Table 27 lists the eight highest-severity cells with channels and basis.

Table 26. Propagation severity matrix: five shock origins by eight affected nodes.

Shock origin	Hyperscalers	DC oper- ators	Private credit	Banks	Insurers	Pensions	Utilities	Bond mkts
Hyperscaler capex cut	Med	High	Med	Low	Low	Low	Med	Low
Neocloud default	Low	High	High	Med	Med	Med	Med	High
Private-credit fund stress	Med	High	High	Med	Med	Med	Low	Med
Bank SRT/underwriting loss	Low	High	Med	Med	Low	Low	Low	High
Insurer asset hit	Low	Med	Med	Low	High	Low	Low	Med

Source: Workstream C propagation-matrix.csv. Severity is the modeler’s judgment (ASSUMED, C); key cells rest on observed bases (see text). High-severity cells: neocloud default → private credit (direct marks on DDTL/BDC SPV positions); neocloud default → bond markets (HY AI paper reprices 200–400bp wider; issuance window shuts for non-IG operators); private-credit stress → operators (CoreWeave’s \$5.6B undrawn DDTL at risk against 0.66x base 2027 coverage); bank underwriting loss → operators and bond markets (the \$38B Oracle template reprices; AI loan/bond basis blows out).

The key channels, in prose. A hyperscaler capex cut is self-inflicted at the top (order cancellations flow back as lower growth capex and impairments on the build-to-suit pipeline) but High severity for data-center operators (CBRE mid-70s% preleasing unwinds toward the 40–50% historical norm; JLL has 2026 deliveries at \$175 billion rising to \$322 billion by 2028 (OBSERVED, B)) and Medium for utilities




Figure 8. Propagation network: shock origins (left) to affected nodes (right); line weight encodes severity. The heaviest edges run neocloud default → private credit, private-credit stress → operators, and bank underwriting loss → bond markets.

Source: Workstream C propagation-matrix.csv. Placeholder figure; coordinator to build with chart_style.qa_check.

Table 27. Highest-severity propagation cells: shock, channel, and basis.

Shock → node	Severity	Channel	Basis
Neocloud default → private credit	High	Direct marks: DDTL, term loans, BDC SPV positions (BCRED/CGBD/HPS/ARCC/OBDC); CLO warehouses hit	(OBSERVED, B)
Neocloud default → bond markets	High	HY AI paper reprices 200–400bp wider; new-issue window shuts for non-IG operators	(OBSERVED, B)
Private-credit stress → operators	High	DDTL/construction facilities pulled or repriced; CoreWeave’s \$5.6B undrawn DDTL at risk	(OBSERVED, B)
Bank SRT/underwriting loss → bond markets	High	AI loan/bond basis blows out; new-issue concessions jump; CLO formation on AI collateral stalls	(OBSERVED, B)
Bank SRT/underwriting loss → operators	High	Project-loan market reprices wider; \$38B Oracle construction-debt template trades down	(OBSERVED, B)
Hyperscaler capex cut → operators	High	Cancelled/deferred preleases; CBRE mid-70s% preleasing unwinds toward 40–50% norm	(OBSERVED, B)
Private-credit stress → private credit	High	Gating and NAV cuts cascade across funds holding the same SPV paper; liquidity mismatch crystallizes	(ASSUMED, C)
Insurer asset hit → insurers	High	RBC ratios fall at concentrated holders; forced de-risking sales of BBB AI paper	(ASSUMED, C)

Source:

Workstream C propagation-matrix.csv. Severity is the modeler’s judgment (ASSUMED, C); the first six cells rest on observed bases as shown.

left holding generation and transmission build (TeraWulf’s October 5 8-K doubled Muskie power to 1 GW on expected demand (OBSERVED, B)). A neocloud default reprices all non-IG operator paper (the Beignet/Hyperion 2049s at about 91 cents are the template (OBSERVED, B)), writes down private-credit DDTL, term-loan, and BDC SPV positions (BCRED/CGBD/HPS/ARCC/OBDC memo rows document filed fair-value positions at Q2 2026 (OBSERVED, B)), draws warehouse and subscription lines at banks (funding-first, per KR-WP-2026-12: NDFI lending \$2.026 trillion, +19% year-on-year; SNC 8.6% non-pass (OBSERVED, B)), and shuts the HY issuance window (September AI debt issuance already fell to about \$23 billion (OBSERVED, B)). Private-credit fund stress kills the talks-stage pipeline first (more than \$50 billion of chip/SPV financings in OBL-0111/0112/0113 (OBSERVED, B)) and pulls CoreWeave’s growth oxygen; gating and NAV cuts cascade across funds holding the same SPV paper. A bank SRT/underwriting loss is a P&L and RWA event, not a solvency event, but it reprices the project-loan market wider and blows out the AI loan/bond basis, stalling CLO formation on AI collateral. An insurer asset hit is the least mapped: no public disclosure of insurer AI-bond concentrations was found (gap), but RBC pressure at concentrated holders would force de-risking sales of BBB AI paper into a falling market.

Table 28 lists the eight observable catalysts: what to watch, where it appears first in data, which node breaks first, and the severity judgment. Catalysts 6 and 7 are fundamental (they impair the contracted cash flows that all GPU-secured lending underwrites); catalysts 1, 2, 3, and 8 are funding-market events (they close the capital-market access that CoreWeave and increasingly Oracle depend on); catalyst 4 is the transmission mechanism from asset marks to funding availability; catalyst 5 is the first legal acknowledgment that a borrower cannot meet its schedule. In the model’s judgment the most likely sequence is $6 \rightarrow 4 \rightarrow 5 \rightarrow 1$, with 2/3 as the Oracle-specific accelerant. The logic: a sustained gap between spot GPU rates and contracted rate cards (6) impairs the contracted cash flows that every GPU-secured facility underwrites, which forces private-credit marks down (4); the marks tighten DDTL availability and covenant headroom until a borrower needs a waiver or misses an amortization payment (5); the first legal acknowledgment of distress reprices the whole AI project-loan bid and kills the next large syndication (1). The Oracle accelerant runs in parallel: CDS through 300bp or a BBB- downgrade (2) reprices every Oracle-linked obligation at once and turns the \$38B undistributable pipeline (3) from an underwriting embarrassment into a balance-sheet event. None of this requires fraud, a recession, or a bank failure; it requires only that contracted AI cash flows grow more slowly than the financing assumed.

Table 28. Eight observable catalysts for forced deleveraging or contagion.

#	Catalyst	What to watch	Where it appears first	Node that breaks first
1	A failed \$10B+ AI-infrastructure syndication	A broadly syndicated AI loan/bond pulled, down-sized >20%, or flexed wider than 150bp. Live template: Broadcom’s \$42B Class A tranche (syndication began Oct 5, 2026) and the \$18B junior tranche	LevFin price talk (Octus, LCD); flex language in the offering memo; arranger retained-share disclosures in bank 10-Qs	Arranging banks’ underwriting desks, then the borrower, then the whole AI project-loan bid (the \$38B Oracle template already trades below 90c)

Table 28 (continued).

#	Catalyst	What to watch	Where it appears first	Node that breaks first
2	Oracle 5Y CDS sustained above 300bp, or downgrade to BBB-	Record print was 261bp (Oct 9, 2026). A sustained move through 300bp, or Moody's/S&P to the BBB- cliff edge, reprices every Oracle-linked obligation at once	Goldman/CDS prints; rating-agency outlook changes; Beignet/Hyperion 2049s (~91c) and Oracle 2055s (77c) as the cash-market echo	Oracle's \$38B undistributable construction debt (JPM/MUFG pipeline), then Oracle's own refinancing (FY2028 maturities \$10.1B), then IG AI paper broadly
3	The \$38B Oracle-leased construction debt trades below 85c or is retained	Secondary prints on the JPM/MUFG-led TX/WI loans, or a bank 10-Q/earnings disclosure of pipeline retention (or SRT at punitive pricing)	Loan-price re-reporting (Octus); bank earnings calls; FT/Reuters on SRT attempts	Underwriting banks (P&L and RWA), then the AI project-loan market: a sub-85c print sets the clearing price for every subsequent data-center construction facility
4	A private-credit fund gates redemptions or cuts NAV >10% on AI SPV marks	Quarterly marks at BDCs/private funds holding AI SPV debt at filed fair value: Blue Owl (BCRED), Carlyle (CGBD), HPS, Ares (ARCC), Blue Owl (OBDC). Memo rows in the monitor, so the exclusion is auditable	BDC 10-Q schedules of investments; fund letters; redemption-queue disclosures	The fund's LPs (liquidity), then borrowers: DDTL availability is the oxygen for neocloud growth capex (CoreWeave 0.66x base 2027 coverage without market access)

Table 28 (continued).

#	Catalyst	What to watch	Where it appears first	Node that breaks first
5	A neocloud covenant waiver, amendment, or missed amortization payment	CoreWeave DDTL covenant tests; any 8-K disclosing a waiver/amendment/deferral from “in compliance” similar tests at Nscale, Applied Digital, TerraWulf	8-K Item 1.03/2.04 or 10-Q covenant language flipping from “in compliance” to “waiver obtained”; DDTL delayed-draw notices	Secured lenders (waiver fees vs. enforcement on collateral worth less than principal per the LTV model), then the converts market (CoreWeave 2033s) as the liquid proxy
6	Spot GPU rates clear >40% below rate cards for two consecutive quarters	CloudZero spot series (already down 30–40% from peaks) vs. neocloud published rate cards. A sustained 40%+ gap means contracted revenue is priced against economics that no longer clear	CloudZero quarterly reports; neocloud 10-Q revenue-per-MW; then DSCR covenant tests at the SPVs	Neocloud revenue covenants and collateral coverage tests (at 2–3yr effective life the hardware covers barely half the loan), then private-credit marks
7	A top-3 AI-lab customer payment default or forced renegotiation	A missed payment, a “restructured” take-or-pay, or an anchor-tenant force-majeure dispute at the scale of Project Jupiter (Oracle’s Sep 2026 force-majeure notice to Blue Owl’s Stack Infrastructure; up to ~\$25B revenue at risk per Morningstar; not yet a payment default)	Neocloud 10-Q customer-concentration and receivables disclosures; lessor 8-Ks; then private-credit marks on the affected SPV	The leasing neocloud/operator (revenue hole against fixed debt service; scenario (g) cuts CoreWeave OCF 30% and Oracle 15%), then lenders, then the “contracted cash flow” underwriting thesis sector-wide

Table 28 (continued).

#	Catalyst	What to watch	Where it appears first	Node that breaks first
8	Anthropic’s IPO fails or is delayed past Q1 2027	Anthropic aiming for a mid-November 2026 IPO at ~\$2T valuation; no public S-1 filed as of Oct 9, 2026. The \$42B Broadcom converts and \$60B debt package are explicitly sequenced around the IPO	EDGAR (S-1 filing or absence); the Oct 14 investor day; then syndication talk on the \$42B Class A	Broadcom’s supplier-financing platform (the supplier-as-lender model), then the \$18B junior tranche (no Broadcom guarantee; direct Anthropic credit exposure), then the broader chip-financing wave

Source: Workstream C catalysts.md. Observables are verifiable; severity is the modeler’s judgment (ASSUMED, C).

7 Limitations and robustness

7.1 Falsifiability

The bottom line of this section: the paper’s conclusions are stated with the evidence that would overturn them. The main conclusions would be proven wrong by the following. The “no fraud” verdict would fall if a restatement, SEC enforcement action, or examiner-style report showed that any cohort company hid losses, fabricated earnings, or misapplied ASC 810/842 on a material structure (the Beignet non-consolidation assessment, as a critical audit matter, is the natural place such a finding would first surface). The “solvent core” verdict would fall if any of Microsoft, Amazon, Alphabet, or Meta reported fixed-obligation coverage below 1.0x on the Section 5.3 definition, or if a hyperscaler cut its dividend, suspended buybacks to meet fixed payments, or drew a revolving facility to cover maturities. The “fragile periphery” verdict would fall if CoreWeave covered its 2027 fixed burden from operating cash flow without new issuance, or if Oracle’s uncommenced-lease ramp came in more than 30% below the modeled path while coverage stayed above 1.5x. The resemblance grades would change if new evidence showed executive self-dealing in AI SPVs (Tyco would stop being a null finding), intentional reporting-date window dressing (Lehman would strengthen), or fabricated underlying contracts (the GFC asset-quality parallel would strengthen). The collateral finding is arithmetic: it would be overturned only by sustained secondary-market GPU prices implying economic lives well beyond five years, against the current CloudZero spot series.

Table 29. Falsifiers: the evidence that would overturn each main conclusion.

Conclusion	Overturning evidence	Where it would appear first	
No fraud	Restatement, SEC enforcement action, or examiner-style report showing hidden losses, fabricated earnings, or material misapplication of ASC 810/842	8-K / 10-K (including critical audit matters); SEC releases	
Solvent core	MSFT, AMZN, GOOGL, or META fixed-obligation coverage below 1.0x on the Section 5.3 definition; dividend cut or buyback suspension to meet fixed payments; revolver drawn to cover maturities	Earnings releases; 10-Q liquidity disclosures	
Fragile periphery	CoreWeave covers its 2027 fixed burden from OCF without new issuance; Oracle's uncommenced-lease ramp comes in >30% below the modeled path while coverage stays above 1.5x	10-K / 10-Q filings	<i>Source:</i>
Resemblance grades	Executive self-dealing in AI SPVs (Tyco stops being null); intentional reporting-date window dressing (Lehman strengthens); fabricated underlying contracts (GFC strengthens)	Proxy and related-party footnotes; restatements	
Collateral finding	Sustained secondary-market GPU prices implying economic lives well beyond five years	Auction results; CloudZero series	
Section 7.1. The next vintage tests these first.			

7.2 Null findings

Three null findings are reported as results, not omissions. First, Tyco: the resemblance is null. Nothing in the AI financing evidence maps onto insider looting or purchase-accounting earnings inflation, and the comparison is recorded as Misleading rather than stretched. Second, WorldCom: the core mechanism (fabricating profit by capitalizing operating expenses) has no analogue; only the narrow forward-looking impairment channel survives, grading Weak. Third, the never-breach cases: no scenario and no case pushes Microsoft, Amazon, or Alphabet below 1.0x coverage, and the +200bp refinancing shock alone breaches nobody. These are findings about the shape of the risk (concentrated in the periphery and in funding access, not in the core or in rates), and the paper’s verdict depends on them. A fourth null is worth recording: for all the attention on the chip-financing wave, no cohort company has been shown to recognize revenue on vendor-financed customer purchases in a way that inflates reported growth; the principal-versus-agent question behind the revenue-counting gap is a classification question, not a fabrication finding. The fraud hypotheses were tested where the evidence could have shown them (restatements, enforcement actions, auditor flags, executive-side participation) and came back empty; that emptiness is data.

7.3 Data gaps

The material gaps, each labeled where it binds. No Anthropic S-1 had been filed as of October 9, 2026 (OBSERVED, B), so the lab’s full obligation stack is commitments-only and press-grade. EDGAR’s data.sec.gov API refused access (HTTP 403) during the build, so no fresh XBRL pulls were possible; filed figures are corroborated across at least two sources each. No debt-maturity schedule exists in the monitor for Amazon, Alphabet, or Nvidia, so principal coverage is not computable for three of the eleven entities. Private-company unknowns are pervasive: OpenAI’s and xAI’s true fixed stacks are the largest unknowns in the workstream, and private-credit fund-level AI exposures have no public disclosure regime (the BDC memo rows are the only observable marks). GPU secondary-market prices are not directly estimated; the collateral model estimates book-value erosion, and actual auction recoveries could differ in either direction. Broadcom’s lease components are undisclosed in the monitor (no B/C/D ladder), its OCF base is stale by nine months, and the CoreWeave \$3.7 billion versus \$4.2 billion post-period bond discrepancy is unresolved. Related-party lease amounts at CoreWeave were not found (the Q2 10-Q related-party footnote was not pulled). Alphabet’s \$707 billion fixed-commitment portfolio is excluded from the coverage model for lack of a schedule; if even 5% fell due in one year it would halve Alphabet’s coverage, still leaving it above 5x (MODELLED, C).

7.4 Robustness

The verdicts are not sensitive to the plausible range of the key assumptions. Moving blended coupons by ± 100 bp shifts Oracle coverage by about 0.06x and CoreWeave by about 0.05x and changes no breach finding (MODELLED, C). Shortening the commenced-lease divisor from 8 to 6 years raises the burden about 15% for Microsoft and Amazon with no breach changes, except Meta’s 2027 adverse case (1.07x to 0.99x at 6 years) (MODELLED, C). Shorter uncommenced-lease terms push Oracle’s FY2029 base coverage toward 1.1x; the direction is unambiguous, the level is not (MODELLED, C). Front-loading Microsoft/Amazon purchase commitments into two years would cut Microsoft’s base 2027 coverage to about 1.6x, still no breach (MODELLED, C). Varying OCF growth by ± 5 pp changes 2030 coverage by 10–25% with no breach finding changing before 2029

Table 30. Data gaps: where each binds and the remedy.

Gap	Where it binds	Remedy / next vintage	<i>Source:</i>
No Anthropic S-1 (as of Oct 9, 2026)	Lab exposure is commitments-only and press-grade	S-1 filing; Oct 14 investor day	
EDGAR API 403	No fresh XBRL pulls	Direct EDGAR filing fetches	
No debt-maturity schedules (AMZN/GOOGL/NVDA)	Principal coverage not computable for 3 of 11 entities	Footnote extraction of maturity tables	
Private-credit fund-level data	Transmission opacity in stress	BDC 10-Q fair-value marks (catalyst 4)	
SPV financials (e.g. Beignet power analysis)	Consolidation assessment not verifiable by outsiders	10-K critical-audit-matter disclosures	
GPU secondary-market prices	Collateral recovery estimated from book value, not auctions	Auction data; CloudZero series	
Bank AI exposures	Bank leg of propagation	KR-WP-2026-12 updates	
Broadcom lease components	No B/C/D ladder for AVGO	10-K lease footnote	
Oracle \$260B vs \$288B staleness	Tier D/E uncertain by ~\$28B	Next 10-Q uncommenced-lease disclosure	
CoreWeave \$3.7B vs \$4.2B discrepancy	Tier A uncertain by \$0.5B	Prospectus / 8-K reconciliation	

Workstream A–C gap logs (vintage 2026-10-11). None of these gaps changes the paper’s verdicts (Section 7.4); they bound its precision. Table 30 maps each gap to its remedy.

(MODELLED, C). On collateral: a 0% residual accelerates the LTV breach by about one quarter; a 25% residual delays the five-year-life breach to year three; 10%/yr amortization delays the F1 breach by about one year at four-to-five-year life; a 50/50 equipment/shell split delays the F2 breach by roughly one year (MODELLED, C). None of these moves the finding that hardware collateral alone does not support the facilities past year two. Table 31 collects the sensitivity results in one place.

Table 31. Sensitivity of the stress verdicts to key assumptions.

Variation	Effect on verdicts	<i>Source:</i>
Blended coupons $\pm 100\text{bp}$	Oracle coverage $\pm 0.06\text{x}$, CoreWeave $\pm 0.05\text{x}$; no breach finding changes	
Commenced-lease divisor 8yr $\rightarrow$ 6yr	Burden +15% for MSFT/AMZN; no breach changes except META 2027 adverse (1.07x $\rightarrow$ 0.99x)	
Uncommenced-lease terms shorter	Oracle FY2029 base coverage toward 1.1x; direction unambiguous, level not	
Purchase commitments front-loaded to 2yr	MSFT base 2027 coverage to $\sim 1.6\text{x}$; still no breach	
OCF growth $\pm 5\text{pp}$	2030 coverage $\pm 10\text{--}25\%$; no breach finding changes before 2029	
GPU residual 0% / 25%	Breach ~ 1 quarter earlier / 5yr-life breach delayed to year 3	
Facility amortization 10%/yr	F1 LTV breach about one year later at 4–5yr life	
F2 collateral 50/50 equipment/shell	F2 LTV breach roughly one year later	

Workstream C coverage-model.md and collateral-depreciation.md (all (MODELLED, C)).

8 Conclusion

The verdict, stated without uplift. This is not an Enron, WorldCom, or Tyco story: there is no evidence of hidden losses, fabricated earnings, or criminal intent, the structures are disclosed in

filings, and ASC 810 and ASC 842 block the specific tricks those cases used. The GFC parallel is structural, not asset-quality: the funding architecture genuinely resembles pre-crisis shadow banking (private-credit intermediation outside the bank perimeter, SPV layering, maturity transformation in construction finance, originate-to-distribute incentives in vendor-arranged financing), while the assets differ fundamentally (productive infrastructure with contracted cash flows versus fraudulent consumer credit) and there is no household-leverage or bank-run dimension. Greensill is the closest live single-firm parallel: non-bank intermediation outside the regulatory perimeter, financing of future (not-yet-existing) economic claims, dependence on guarantees, arranger opacity, and single-name concentration; it stops at Partial because Greensill drew misconduct findings with no AI analogue in evidence.

The solvency core holds: Microsoft, Amazon, and Alphabet do not breach 1.0x fixed-obligation coverage in any stress case including the severe case; Meta breaches only in 2027 of the severe case. Their buildout is a sustainable capital-investment cycle funded from operating cash flow. The periphery is genuinely fragile: CoreWeave breaches in every case from 2027 and is structurally dependent on market access; Oracle compresses to 1.33x base coverage by FY2029 and breaches in the severe case; Meta faces a one-year 2027 squeeze against a \$90.3 billion cash buffer. GPU-secured debt goes underwater on hardware value within one to two years under every depreciation assumption; recovery depends on contract assignment, not resale.

The bottom line is a solvent core funding a fragile periphery through increasingly circular structures: a credit cycle with genuine systemic-fragility elements in the transmission architecture (private credit, vendor financing loops, bank underwriting pipelines), not a fraud story and not a clean bill of health. The specific catalysts to watch are a failed \$10 billion-plus AI syndication, Oracle CDS through 300bp or a BBB- downgrade, and the \$38 billion Oracle-leased construction debt printing below 85 cents or being retained on bank balance sheets.

References

References

- Richard C. Breeden. Report of investigation by the corporate monitor. Technical report, United States District Court for the Southern District of New York, 2003. Corporate Monitor in *SEC v. WorldCom*, August 2003.
- Markus K. Brunnermeier and Lasse Heje Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009.
- Patricia M. Dechow and Catherine M. Schrand. *Earnings Quality*. CFA Institute, Charlottesville, VA, 2004. Research Foundation of CFA Institute.
- Patricia M. Dechow, Weili Ge, Chad R. Larson, and Richard G. Sloan. Predicting material accounting misstatements. *Contemporary Accounting Research*, 28(1):17–82, 2011.
- Mei Feng, Jeffrey D. Gramlich, and Sanjay Gupta. Special purpose vehicles: Empirical evidence on determinants and earnings management. *The Accounting Review*, 84(6):1833–1876, 2009.
- Financial Crisis Inquiry Commission. *The Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States*. U.S. Government Printing Office, Washington, DC, 2011. January 27, 2011.

- Gary Gorton and Andrew Metrick. Securitized banking and the run on repo. *Journal of Financial Economics*, 104(3):425–451, 2012.
- House of Commons Treasury Committee. Lessons from greensill capital. Sixth report of session 2021–22, House of Commons, 2021. July 20, 2021.
- Kilambi Research. KR-WP-2026-10: Hyperscaler capital expenditure and the physical drivers of the AI datacenter buildout. Working Paper KR-WP-2026-10, Kilambi Research, 2026a.
- Kilambi Research. KR-WP-2026-11: Can the AI buildout pay for itself? Working Paper KR-WP-2026-11, Kilambi Research, 2026b.
- Kilambi Research. KR-WP-2026-12: Bank balance-sheet capacity for AI infrastructure credit: A stress test. Working Paper KR-WP-2026-12, Kilambi Research, 2026c.
- Jr. Powers, William C., Raymond S. Toubh, and Jr. Winokur, Herbert S. *Report of Investigation by the Special Investigative Committee of the Board of Directors of Enron Corp.* Enron Corp. Board of Directors, Houston, TX, 2002. February 1, 2002.
- Zoltan Pozsar, Tobias Adrian, Adam Ashcraft, and Hayley Boesky. Shadow banking. Staff Report 458, Federal Reserve Bank of New York, 2010. July 2010.
- Swiss Financial Market Supervisory Authority FINMA. Enforcement proceedings conclusion re credit suisse / greensill: Serious breach of supervisory obligations, 2023. February 28, 2023.
- United States Senate Permanent Subcommittee on Investigations. *Wall Street and the Financial Crisis: Anatomy of a Financial Collapse*. U.S. Senate Committee on Homeland Security and Governmental Affairs, Washington, DC, 2011. Majority and Minority Staff Report (Levin–Coburn), April 13, 2011. 639 pp.
- U.S. Securities and Exchange Commission. Litigation release no. 17722 / AAER no. 1627: SEC v. kozlowski, swartz & belnick, 2002a. September 12, 2002.
- U.S. Securities and Exchange Commission. Litigation release no. 17588 / AAER no. 1585: SEC v. WorldCom, inc., 02-cv-4963 (S.D.N.Y.), 2002b. June 27, 2002.
- U.S. Securities and Exchange Commission. Press release no. 2006-58: SEC charges tyco international ltd. with accounting fraud, 2006. April 17, 2006. Settled action; \$50 million civil penalty.
- Anton R. Valukas. Report of anton r. valukas, examiner. Technical report, United States Bankruptcy Court for the Southern District of New York, 2010. In re Lehman Brothers Holdings Inc., Ch. 11 Case No. 08-13555. March 11, 2010.

A Counting rules and tier definitions

This appendix summarizes the monitor’s counting rules; the full text is maintained with the dataset. (1) One economic obligation, one row: each row is a distinct obligation of a distinct entity. (2) Memo rows are excluded from all aggregates: press-reported, talks-stage, and duplicative lender-side views (e.g., BDC positions in CoreWeave SPVs that consolidate into CoreWeave; Anthropic’s \$518 billion memo portfolio containing its \$130 billion of primary commitments). The exclusion is auditable row

by row. (3) Basis is labeled on every line (recognized carrying, undiscounted notional, contingent maximum notional, off-balance-sheet) and mixed bases are never summed without labeling the mix. (4) No double counting across parties: a facility appears once, on the obligor’s books; lender-side views are memo. (5) An undrawn revolver commitment is not debt. (6) Guarantees enter Tier E at maximum notional with weight 1.0 (ceiling, not expected loss); disclosed escrow is noted but not netted. (7) FX: all figures converted to USD at the vintage’s stated rates. (8) Staleness: filed figures are not carried forward across as-of dates; where a newer press figure exists (Oracle’s \$288 billion uncommenced leases), both vintages are shown. (9) Subsequent-event rows (e.g., Meta’s \$68 billion of July 2026 leases) are memo until carried in a filing-period row. (10) Tier definitions: A = GAAP recognized debt (carrying); B = A + finance lease liabilities; C = B + operating lease liabilities; D = C + uncommenced leases and other disclosed commitments (notional); E = D + unavoidable purchase/supply/take-or-pay commitments (notional) + guarantees at maximum notional (weight 1.0). Tiers A–C are pure carrying basis; D and E are mixed and always labeled as such.

Worked examples. (i) CoreWeave’s financing SPVs (Compute Acquisition Co. II/IV, DDTL V vehicles) consolidate into CoreWeave, so the \$3.7B of non-recourse debt sits in Tier A; BDC lender-side positions in the same vehicles (HPS, Carlyle schedules) are memo rows, excluded from aggregates to avoid double counting. (ii) Anthropic’s \$100B AWS + \$30B Azure commitments are primary rows; the ~\$518B aggregate compute portfolio (Reuters, from the confidential IPO prospectus) is a memo row that contains them, so it is excluded. (iii) Oracle’s May-31 \$260B uncommenced figure was not carried into the August row; both the filed figure and the August \$288B press figure are shown with vintages, and Tier D/E are reported both ways. (iv) Meta’s \$68B of leases signed in July 2026 are a subsequent-event memo row until a filing-period row carries them. (v) OpenAI’s \$4.7B revolver is committed and undrawn: it is a primary row for facility existence but contributes zero to Tier A, because an undrawn commitment is not debt.

B Data appendix: series definitions and source table

OCF_TTM: net cash provided by operating activities, trailing twelve months, from 10-K/10-Q cash-flow statements or earnings releases; where TTM is reconstructed from quarterly decks, the quarters are listed. *Tier A*: current + long-term debt (carrying), balance sheet. *Finance lease liabilities / operating lease liabilities*: ASC 842 recognized amounts (carrying). *Uncommenced leases*: undiscounted total of signed non-cancelable leases not yet commenced, commitments footnote. *Purchase commitments*: unconditional purchase obligations, undiscounted, ASC 440 disclosure. *Fixed/guaranteed commitments*: Alphabet’s aggregate portfolio figure as disclosed. *Compute commitments*: Meta’s non-cancelable contractual commitments (cloud capacity, servers, network, data centers). *Guarantees*: maximum potential undiscounted future payments, ASC 460 disclosure. *Principal coverage*: *OCF_TTM* divided by debt principal maturing in the next 12 months (or nearest reported bucket, stated); principal-only, interest excluded. *All-in fixed-obligation years*: Tier E divided by *OCF_TTM*; a burden scale, not a default clock. *Coverage ratio (stress)*: $OCF(y) / \text{burden}(y)$ per Section 5.3. *LTV*: loan balance / modeled hardware value per Section 5.4. *CDS*: five-year credit default swap spread, midpoint prints. *Bond price*: clean price in cents on the dollar. *RPO / backlog*: remaining performance obligations / contracted revenue backlog as disclosed (e.g. Oracle RPO \$664B, CoreWeave revenue backlog \$104B, Google Cloud backlog \$514B). *Deferred revenue*: customer prepayments carried as a liability (CoreWeave \$9.7B pool). *DDTL*: delayed-draw term loan facility (CoreWeave DDTL 4.0: \$8.5B committed, \$2.9B drawn). *CDS*: five-year credit

default swap spread, midpoint prints. *Source table*: obligation rows cite 85 distinct sources; the monitor’s sources file maps each row to its source (filing, earnings release, press report with outlet, or market-data provider). Evidence grades per Section 4.5 apply to every figure.

Model artifacts. Coverage model: `model-code/coverage_model.py` (assumptions A1–A10; outputs in `coverage_tables.json`). Collateral model: `model-code/collateral_model.py` (assumptions C1–C5; annual paths in `collateral-ltv-paths.csv`). Machine-readable results: `scenario-results.csv` (54 rows: 9 scenarios $\times$ 6 borrowers), `propagation-matrix.csv` (40 cells: 5 shock origins $\times$ 8 affected nodes), `collateral-summary.csv` (8 rows: 2 facilities $\times$ 4 GPU-life assumptions). All model outputs are (MODELLED, C) by construction; inputs carry the grades in Tables 9 and 10.

Table 32 collects the maturity-schedule inputs used in the coverage model.

C Additional company tables

Table 32. Maturity schedule inputs used in the coverage model (\$M).

Borrower	2026	2027	2028	2029	2030	2031	2035+
MSFT debt principal	n.a.	9,250	0	2,001	0	500	34,385
META debt principal	n.a.	2,750	1,500	1,000	5,000	n.a.	73,750
ORCL debt principal	n.a.	7,210	10,145	5,500	7,250	9,750	90,250
AVGO debt principal	2,252	0	5,127	2,405	6,406	n.a.	44,889
CRWV debt principal	4,413	6,184	4,416	2,421	3,221	n.a.	14,896
META compute commitments	53,520	81,650	(residual spread 2028–2036)				
AVGO purchase obligations	1,110	52,674	72,952	85	n.a.	n.a.	n.a.
NVDA purchase obligations	n.a.	120,000	100,000	98,000	n.a.	n.a.	n.a.

Source: AI
leverage monitor, maturity_wall.csv, vintage 2026-10-11 (OBSERVED, B). Microsoft/Oracle fiscal years. Amazon, Alphabet, and Nvidia debt-principal schedules not in the monitor (gap).

D AI assistance disclosure

This paper was produced by Kilambi Research’s quantitative research pipeline with AI assistance; all analysis was verified against source data.