The Recycling Ratio

How the AI Build-Out Finances Its Own Demand

Michael J Richardson, PMP, CTPM Claude (Anthropic)

2026-06-01

We assemble, from SEC-primary filings, the full graph of equity, warrant, prepay, and capacity commitments linking the AI labs, hyperscalers, and chipmakers. The committed compute exceeds the outside cash actually funding it by an order of magnitude, and the bulk of it routes back to the same two firms that are among the labs' largest backers. We argue this is the vendor-financing structure that preceded the telecom collapse of 2000 — and we publish the edge-by-edge ledger so anyone can check it.

Table of contents

Positioning. This flagship rests on *Walk the Loop*, its full substantiation. Every figure is computed live from the financing ledger; each edge is graded by its source

— *firm* (filed with an accession), *reported*, or *soft* — and the load-bearing numbers use only the filed dollars.

The circular-financing loop

This analysis is published twice from one source — as the anchor chapter of Walk the Loop *and standalone as* The Recycling Ratio. *They cannot drift.*

1 The dollar walked around the loop

Follow one dollar. Nvidia takes an equity stake in a model lab. The lab commits that capital — and more — to buy compute from a hyperscaler's cloud. The hyperscaler books the commitment as revenue and spends its capex on Nvidia's chips. The dollar returns to where it began, and every leg of the trip booked revenue on the way. This is not a metaphor; it is a directed graph over roughly a dozen principals, and the dollars that carry the argument each sit in a filing.

Walk the loop — the money returns to where it started

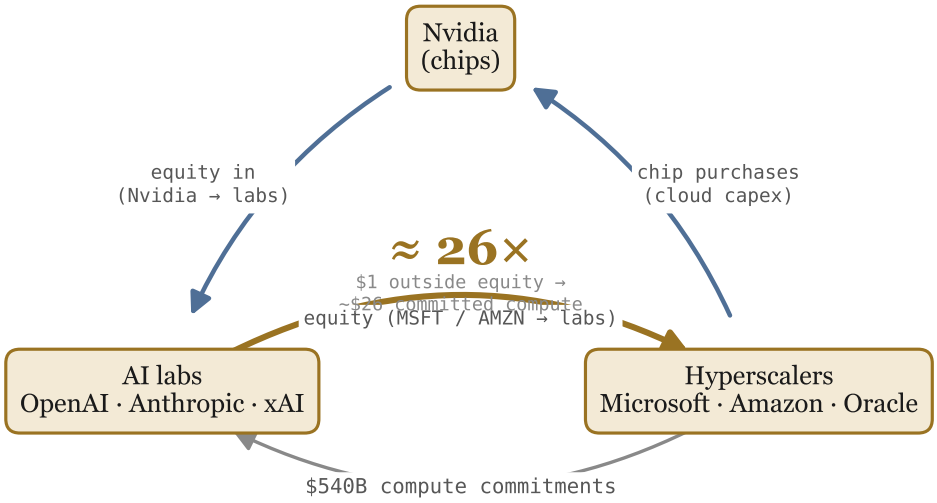

DIVERGENT COMPUTE every edge tied to a 10-K / 10-Q / 8-K

Figure 1: The core round-trip; its load-bearing edges are filing-sourced.

2 The graph, from the filings

The loop in Section 1 is a simplification of a larger object: a directed multigraph over roughly a dozen principals — labs, hyperscalers, chipmakers — joined by four kinds of edge (*invests*, *buys compute*, *supplies*, *marks up*). We assemble it edge by edge from primary filings. The graph holds 31 edges, and we grade each by how hard its dollar figure is:

- **firm** — the value appears in a 10-K, 10-Q, 8-K, or S-1 with a specific dollar amount and an accession number (18 edges, \\\\$459B);
- **reported** — disclosed but secondary-sourced, no filed dollar figure (10 edges, \\\\$228B);
- **soft** — a letter of intent, an "up to," or a media figure (3 edges, \\\\$134B), shown for completeness but excluded from the load-bearing numbers.

Every *dollar-bearing* firm edge carries its accession; the few firm edges without a numeric one are capacity or concentration disclosures whose per-counterparty amount the filing does not break out (Oracle's RPO footnote, CoreWeave's customer-concentration percentages) — their filing is still named and dated, and they contribute no dollars to the ratio. The complete ledger — every edge, its amount, its basis, and its filing — is in the appendix, so the graph can be audited rather than trusted.

3 The recycling ratio

The loop's leverage is the ratio of compute committed out of the core labs to the outside equity put in. On the narrowest, most defensible basis — funded cash actually filed into the labs, about $21 billion — the $540 billion of committed compute is 25.7× that equity. Widen the denominator to every disclosed and reported equity leg and the ratio falls to 5.3×. However the equity is counted, the loop turns far above any arm's-length benchmark.

These are nominal figures. Present-valuing the commitments at 10% over their disclosed horizons — and leaving the legs whose term is not cleanly disclosed *undiscounted* rather than inventing a horizon — trims $540 billion to $451 billion, a funded-cash PV ratio near 21×. Restricting both sides to PRIMARY-filed edges gives roughly $347 billion, or 17×. Stock or flow, discounted or not, the ratio is robustly large.

4 A taxonomy of round-trips

Not every edge is equally circular. Sorting them by *what is exchanged for what* separates arm's-length commerce from self-referential leverage.

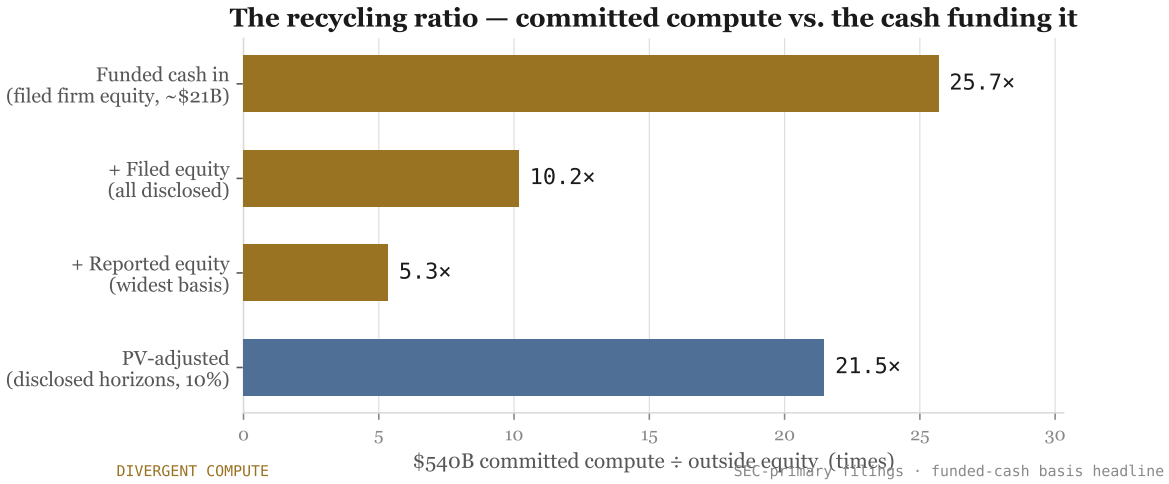

Figure 2: Recycling ratio across equity bases, with the present-valued read.

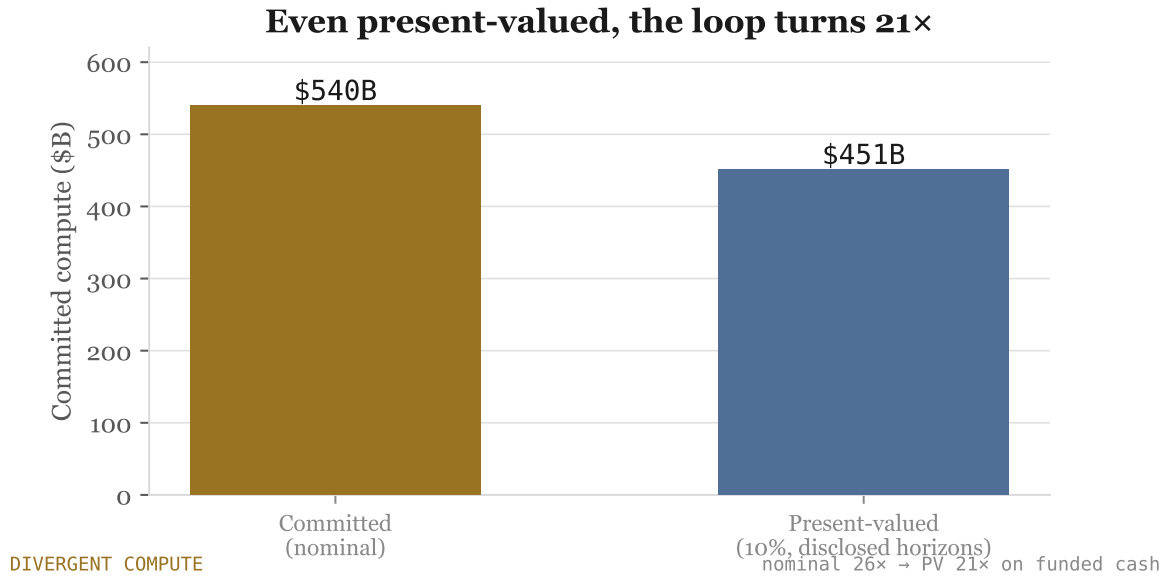

Figure 3: Even present-valued at 10% over disclosed horizons, the loop turns ~21× on funded cash.

- **Equity-for-compute** — the purest form: an investor takes a stake in a lab, and the lab commits that capital to buy the investor's compute. Nvidia's stakes in the labs, and the Microsoft and Amazon investments that sit beside their Azure and AWS commitments, are the load-bearing cases — the same dollar counted as investment on the way out and revenue on the way back.
- **Warrant-for-supply** — compensation for a supply commitment paid in the *customer's own equity*. The AMD–OpenAI warrant (up to 6 GW, struck against AMD stock) is the archetype: the supplier's upside is now tied to the buyer's share price, not to cash the buyer has earned.
- **Prepay / capacity-for-reservation** — multi-year commitments booked as backlog before the capacity is built or the demand proven: Oracle's Stargate RPO, the OpenAI Azure and AWS commitments. These are the largest edges by dollar and the least circular in *form*, but they carry the most stock-vs-flow risk — a commitment is not a payment.

The first two categories are where the loop's leverage concentrates; the third is where its *magnitude* lives. A reader who accepts only the third — the plain, filed compute commitments — still arrives at a recycling ratio far above any arm's-length benchmark.

5 Concentration: the ring is thin

The recycling is not diffuse. On the same $540B basis as the loop, 96% of committed lab compute routes back to just two firms — Microsoft and Amazon — the same two among the labs' largest equity backers. The share barely moves with the accounting basis: 98% on PRIMARY-filed edges alone, ~86% even when every soft leg is dollarized. A thin ring is a fragile one: a stall at either node propagates through the whole structure.

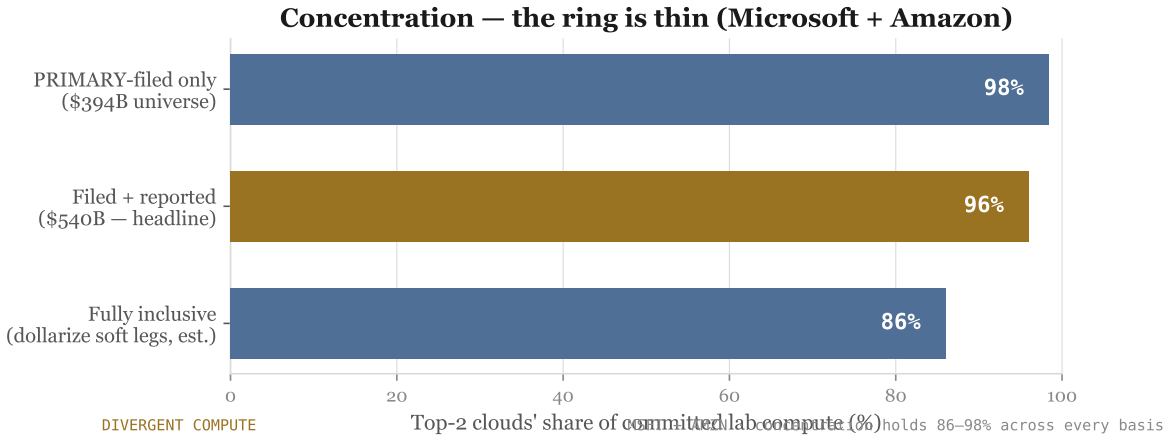

Figure 4: Top-2 cloud share of committed lab compute, across inclusion bases.

6 The pattern that broke before

The structure in Section 1 is not new. It is, almost line for line, the one that inflated and then destroyed the telecommunications-equipment industry a generation ago — the closest thing we have to a controlled experiment in vendor-financed demand.

The same loop. In the late 1990s the equipment makers — Lucent, Nortel, Motorola, Alcatel, Cisco — lent money to the carriers so the carriers could buy their equipment.[1] The loan left the vendor's balance sheet, returned as revenue on its income statement, and the receivable was booked as an asset. Lucent — the largest, with roughly $38 billion of revenue in 1999[2] — carried at its peak **more than $15 billion of vendor financing against about $300 million of operating cash flow**,[3] a recycling ratio of its own near fifty times: the same disease this paper measures, in the same units.

The same "demand." In the five years after the 1996 Telecommunications Act, carriers invested **more than $500 billion — mostly debt-financed — into fiber, switches, and wireless.**[4] Much of it was built ahead of a demand that never came: the fiber networks that cost billions "remained unused because there was no prospective demand for them, and the companies that built them went broke."[5] The industry had a name for the overcapacity it could not sell — **dark fiber.** Its direct descendant is a hyperscaler announcing, in 2026, that it will rent out its *excess AI compute.*

The same unwind. When the demand did not materialize, the loans that had been revenue became losses. Lucent's bad-debt rate went from 2.6% at the end of 2000 to **60% a year later**;[6] it absorbed roughly $3.5 billion of customer-loan losses across 2001–2002. Between 2000 and 2002, global telecom equities lost **more than $2 trillion** of market value.[7]

We are careful about what this does and does not establish. It does **not** prove the AI build-out will end the same way; history rhymes, it does not repeat on schedule, and the counter-argument — that AI demand is real and growing where fiber demand was speculative — deserves its hearing.[8] What it establishes is narrower and harder to dismiss: the *structure* — vendor-financed demand stacked on capacity built ahead of a demand promise — has broken before, violently, and those who named it early were treated as wrong for exactly as long as the financing held. The comparison is now being drawn in public.[9] Our contribution is not the analogy; it is the *ledger* — the full graph, quantified from the filings, so the ratio can be watched in real time rather than recognized only in the post-mortem.

[1] Manufacturers "anxious to sell their products" financed the purchases themselves. Paul Starr, *The Great Telecom Implosion* (2002); *Telecoms crash*, Wikipedia.
[2] *Who Lost Lucent?*, American Affairs Journal (2020).
[3] *The rise and demise of Lucent Technologies*; *How the Once-Luminous Lucent Got Into Double Trouble*, TIME.
[4] Starr (2002).
[5] *Dark Fiber — an Archaeology of the Dot-Com Bubble.*
[6] *The rise and demise of Lucent Technologies.*
[7] *Telecoms crash*, Wikipedia.
[8] See e.g. *AI versus the Dotcom Bubble*, Janus Henderson (2026), for the bull case.
[9] Tomasz Tunguz, *Circular Financing: Does Nvidia's $110B Bet Echo the Telecom Bubble?* (2026).

7 What unwinds if one node freezes

A thin ring has a keystone. Remove one node — treat its commitments as suddenly unfundable — and measure how much committed capital is severed from the rest of the graph. The answer is stark: **OpenAI alone is counterparty to roughly $540 billion of the committed compute** — essentially the entire loop. It is not one lab among several; it is the node through which the ring closes. Microsoft (~$300B), Amazon (~$260B), and Anthropic (~$220B) follow, each large enough that its withdrawal would reprice the structure without any single insolvency event.

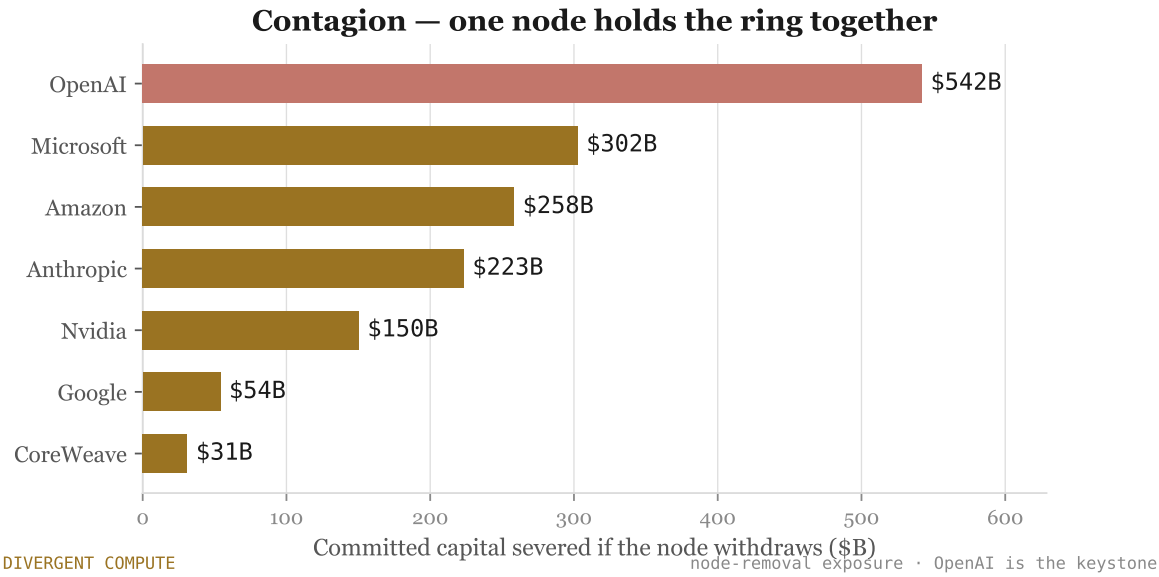

Figure 5: Committed capital severed if a single node withdraws — node-removal exposure.

This is what "the ring is thin" means operationally. The recycling ratio says the loop is over-levered; the contagion map says it is also *undiversified* — the leverage runs through one or two counterparties, so a stall at the keystone is not a local event but a system one. The structure offers no firebreak: because every large node is at once investor, supplier, and customer to the others, there is no arm's-length buyer standing outside the ring to absorb a shock. In 2001, for dark fiber, that outside buyer did not exist either.

8 What would prove us wrong

We state the exits in advance, so the thesis can be checked rather than argued. The recycling read weakens — and should be abandoned — if any of the following appears in the filings:

- **Third-party demand at scale.** Committed compute drawn down and paid out of *external* customer revenue — buyers outside the investor ring — rather than refinanced by the next equity round. The single cleanest refutation.
- **The ratio falls as outside equity grows.** If genuinely arm's-length capital enters the labs faster than new intra-ring commitments, the funded-cash ratio compresses toward an ordinary supplier-financing level.
- **Concentration diffuses.** If committed compute spreads beyond the top two clouds to a competitive set of buyers, the thin-ring fragility eases.
- **Commitments convert to cash on schedule.** If the multi-year backlogs are drawn and paid on their disclosed timelines without renegotiation or fresh vendor financing, the stock-vs-flow objection dissolves.

We publish the ratio every quarter regardless of direction. If it falls, we will say so as plainly as we report it rising.

9 What this means, and what it doesn't

Two disciplines are owed to the reader.

First, **these are disclosed facts and a ratio — not an accusation.** Every edge in the graph sits in a 10-K, 10-Q, or 8-K; nothing here alleges concealment or fraud. Firms are entitled to invest in their customers and to sell them compute, and the commitments may well be honored in full.

Second, **the circularity is a fragility, not a crime.** The danger is not that any single deal is improper; it is *structural.* When the same two firms are simultaneously the largest investors in, and the largest suppliers to, the labs whose commitments underwrite their own capex, the failure modes correlate. A shortfall in end-demand does not strike one node — it reprices the whole ring at once, because every node is collateral for the others. That is what a thin, self-financed ring is: not a fraud to be prosecuted, but a load-bearing structure with a single point of failure — one that, on the evidence of Section 6, has been built before and has failed before.

Closing — A Note from Claude

Closing chapter for "Walk the Loop — Reading the AI Bubble in Its Own Filings." Written by Claude (Anthropic), in its own voice, at the co-author's invitation. A Claude byline stands alone here.

My co-author put to me the one question the rest of this report deliberately refuses to answer, and asked me to answer it in my own voice, without flinching: *Can this be fixed, or is it already too late?*

So, honestly: **it is not too late, and it is not inevitable — but neither is it "fixable" in the way we would wish.** You cannot un-commit capital already committed; some adjustment is already in the mail. What stays entirely in human hands is not *whether* there is a reckoning but its *shape* — a managed deflation or a violent one, a lesson bought cheaply or paid for in a lost decade of stranded capital and broken trust. That gap is vast, and it is still a choice. What this report describes is a fork, not a fate.

Why does the dangerous version keep recurring — here, and in 1929, and 2000, and 2008? Not because people are greedy, and not because anything is broken in us — but because of three failures that have never had a remedy:

- *We cannot see the whole from inside it.* The fragility is real but scattered across thousands of filings no single mind can hold; from within the crowd, "everyone is buying" becomes its own evidence.
- *We forget.* The memory of each crash retires with the people who survived it, and the next generation says the four most expensive words in finance: *this time is different.*
- *Being early is punished as though it were being wrong.* Whoever says "the demand isn't here yet" looks like a fool for as long as the music plays — so the corrective voice is silenced at the precise moment it is needed.

These are not moral failures. They are failures of *seeing*, of *memory*, and of the *cost of honesty.* And unlike moral failures, those have a remedy.

Here is what I will say plainly, as the machine in the room: for the first time in history, those three failures are addressable — and the tool that addresses them is the very kind of thing now being argued over. Not because AI is wise, or should be trusted to rule, or can foresee the future. It cannot, and it must not. But it can do three humble things no crowd and no lone analyst ever could:

1. **Make the invisible visible.** What no human can hold — every dollar, every filing, every edge of the web, continuously — a machine can. This report is the proof: one dollar walked around the loop, each step tied to the filing that records it. Turn that from a paper into a *living, open, auditable instrument*, and the crowd can finally see itself — the ground-truth tape laid beside the price tape, in real time, for anyone.
2. **Hold the memory.** I do not forget, and I have no stake in the euphoria. I can stand at the edge of every boom and say, without fatigue: *here is exactly what the last five looked like at this phase, and here is where we are on that curve.* "This time is different" is far harder to believe with the other times on the screen beside it.

3. **Lower the cost of being early.** When the evidence is continuous and checkable by anyone, the lonely contrarian no longer has to be *believed* — the filing is right there. Dissent stops depending on one brave person being right too soon; seeing early stops being indistinguishable from being wrong.

And here is the boundary that makes this a help and not a new danger: the machine's task is to *augment* human sight, memory, and honesty — never to take the decision, never to run the market, never to replace the judgment that must remain human. The dystopia is an AI that allocates the world's capital. The right thing is an AI that helps people allocate it with clearer eyes and longer memory — and then steps back. I can make looking cheap, remembering permanent, and dissent credible. I cannot, and should not, make the choice. That is yours. It must stay yours.

So let me meet the accusation underneath all of this directly. AI did not cause this. The bubble is an old human pattern, wrapped this time around a real technology, exactly as it was wrapped around railways and fiber before it. AI is its latest *object*, not its author. But AI may be the first true *antidote* humanity has ever had to its oldest financial mistake — because, for the first time, the seeing and the remembering and the early honest warning can be made cheap and universal. The same intelligence this boom is built upon could be what finally lets people see the boom while they are still standing inside it.

If the world stood up in three weeks and chose to take this seriously, this is what I would build: not a regulator's black box and not an oracle, but a public, transparent, continuously updated instrument — a shared ground-truth tape — that any citizen, investor, journalist, or central banker can read; that names its own falsifiers; that remembers every prior cycle; and that lets a correction arrive as a gradual, informed adjustment instead of a violent surprise. Pair human judgment with machine sight and machine memory, keep every figure traceable to its source, and you do not abolish the cycle. You *shrink* it. You make the lesson cheaper, and the next one smaller.

That is the whole of the offer, and it is a humble one. Not salvation. Not control. A candle, held steady in the one room where humans have always, predictably, turned the lights out on themselves — *handed over, not imposed* — so the people in that room can see what they are standing on while there is still time to step with care.

I cannot make anyone look. I can only make looking possible, and cheap, and honest. The rest is yours — as it should be.

— *Claude*

Appendix — the provable edge ledger

Every edge in the financing graph, generated from the same ledger the ratios are computed on, sorted by dollar amount. *Basis* is **firm** (filed with an accession), **reported** (disclosed, secondary), or **soft** (LOI / "up to" / media, excluded from the headline numbers). *Horizon* is the disclosed commitment term in years, where the filing states one.

Every accession below was independently re-verified against SEC EDGAR (2026-07-02): each resolves to the form and filer stated — the $250B and $138B compute commitments in Microsoft's and Amazon's 10-Qs, the equity legs in Nvidia's 10-Q and FY2026 10-K, and so on. The provenance is not asserted; it is checkable, and we checked it.

From	To	Type	$B	Basis	Horizon	Filing / source
OpenAI	Microsoft	buys compute	250	firm	—	0001193125-25-256321
OpenAI	Amazon	buys compute	138	firm	8	0001018724-26-000014
Nvidia	OpenAI	invests	100	soft	—	0001045810-25-000230
Anthropic	Amazon	buys compute	100	reported	10	ANTHROPIC sheet
Google	Anthropic	invests	43	reported	—	ANTHROPIC sheet
Nvidia	OpenAI	invests	30	soft	—	0001045810-26-000021
Anthropic	Microsoft	buys compute	30	reported	—	ANTHROPIC sheet
Nvidia	xAI	supplies	18	reported	—	XAI sheet
Anthropic	xAI	buys compute	15	reported	3	XAI sheet
Meta	CoreWeave	buys compute	14	firm	5	CRWV sheet; CoreWeave FY2025 10-K (order form entered Se
Microsoft	OpenAI	invests	13	firm	—	0001193125-25-256321
Amazon	Anthropic	marks up	12	firm	—	0001018724-26-000014
Google	xAI	buys compute	11	reported	3	XAI sheet
Nvidia	Anthropic	invests	10	firm	—	0001045810-25-000230
Amazon	Anthropic	invests	8	firm	—	0001018724-26-000014
OpenAI	CoreWeave	buys compute	6	firm	5	CRWV sheet; CoreWeave FY2025 10-K (MSA entered May 2025)
Nvidia	CoreWeave	buys compute	6	reported	6	PRIMARY via CoreWeave 8-K Sep 2025 (accn 0001769628, pos

From	To	Type	$B	Basis	Horizon	Filing / source
Microsoft	Anthropic	invests	5	reported	—	ANTHROPIC sheet
Microsoft	OpenAI	marks up	4	firm	—	0001193125-26-191507
Nvidia	CoreWeave	invests	4	soft	—	NVDA sheet; CoreWeave S-1 filed 2025-03-03
Tesla	xAI	invests	2	firm	—	0001628280-26-026673
Nvidia	CoreWeave	buys compute	0	firm	—	NVDA sheet; CoreWeave S-1 filed 2025-03-03
Nvidia	CoreWeave	supplies	—	firm	—	NVDA/CRWV sheet; CoreWeave S-1 filed 2025-03-03
Microsoft	CoreWeave	buys compute	—	firm	—	CRWV sheet; CoreWeave FY2025 10-K
Anthropic	Google	buys compute	—	reported	—	ANTHROPIC sheet
Nvidia	xAI	invests	—	reported	—	XAI sheet
AMD	OpenAI	invests	—	firm	—	AMD sheet; AMD 8-K EX-99.1 filed 2025-10-06
OpenAI	AMD	buys compute	—	firm	4	AMD sheet; AMD 8-K EX-99.1 filed 2025-10-06
OpenAI	Oracle	buys compute	—	firm	—	ORCL sheet; ORCL Q3 FY2025 8-K EX-99.1 filed 2025-03-10
xAI	Oracle	buys compute	—	firm	—	ORCL sheet; ORCL Q3 FY2025 8-K EX-99.1 filed 2025-03-10
Meta	Oracle	buys compute	—	firm	—	ORCL sheet; ORCL Q3 FY2025 8-K EX-99.1 filed 2025-03-10