COMMENTARY · AI INFRASTRUCTURE

Who Wins When AI Can't Plug In

Electricity, not money, is becoming the limit on how fast America builds AI. That doesn't end the AI trade. It changes who gets paid.

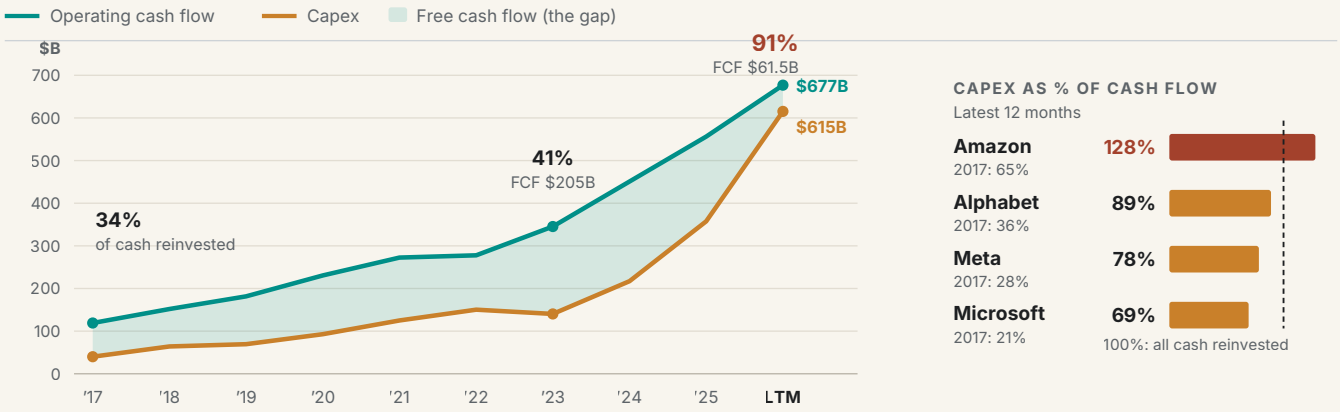
By **Leandro Perez** · November 12, 2026 · Filings through November 10, 2026; market data as of November 6–9

The AI buildout has a new problem

In the 12 months to September, Amazon, Microsoft, Alphabet and Meta spent $615 billion on capital projects. That's 91% of all the cash their businesses generated. In 2017, the same four companies reinvested about a third.[1]

EXHIBIT 1 **The AI capital spending problem**

Amazon, Microsoft, Alphabet and Meta combined, $ billions, 2017 to Sept. 2026; right, capex as a share of each company's cash flow.



Source: company filings via finboard.net and StockAnalysis, checked against September-quarter 2026 releases (note 1). Microsoft fiscal years end in June. Capex = purchases of property and equipment, Amazon gross of incentives.

For two years, investors have asked when that spending will pay off. A different question is becoming more urgent: what happens when a company has the capital, the chips and the customers, but can't get the electricity to switch the capacity on?

It's no longer hypothetical. In August, Texas ordered an audit of the data centers waiting to connect to its grid, whose requests total about 474 gigawatts, more than five times the state's record peak demand.[2] Much of that is speculative, and the audit is a gate rather than a wall.[3] New York has paused permits for large data centers for up to a year.[4] Most other states are simply making data centers pay more for power.[5]

The intuitive reading is that a power shortage is bad for everyone in AI. I don't think that's right. A limit on new capacity treats two kinds of company very differently. Some own capacity that is already powered and earning. Others have committed capital, chips or contracts to capacity that doesn't exist yet. If the limit binds, the first group's assets become scarcer and the second group's promises become harder to keep.

This isn't a forecast that the buildout will slow; queues are inflated, and capital can move to places with spare power. It's an attempt to work out who wins and who pays if it does, and whether the market has noticed.

Why building less doesn't create value

My first hypothesis was simple: if the hyperscalers can't build as fast, they keep more cash. That's true. It's also the least interesting part of the story.

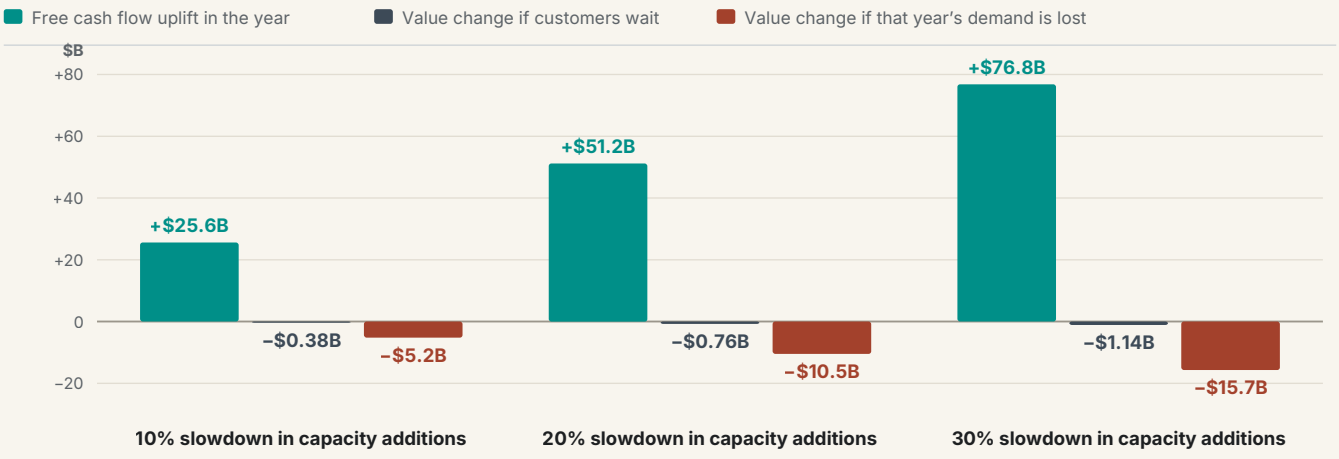
In my model, a 20% slowdown in new capacity would add about $51 billion to the four companies' combined free cash flow over a year, close to 8% of their operating cash flow.[6]

But delaying a profitable project doesn't make it more valuable. It moves cash from the future into the present. As I estimated in *The Central Bank of AI*, capacity financed at 6% over six years needs to return about 20 cents per dollar a year just to repay its cost.[7] My model assumes new capacity earns 25 cents, so every year of delay gives up returns above that hurdle.

How much that costs depends on what customers do. If they wait, because every provider faces the same constraint, the four companies lose less than $1 billion of value in my scenario. If they take the work to a rival, the loss is about $10.5 billion.[6] Either way, reported free cash flow goes up while the business becomes worth less.

EXHIBIT 2 **More cash doesn't always mean more value**

Illustrative effect of a one-year slowdown in new capacity on Amazon, Microsoft, Alphabet and Meta combined, $ billions. A scenario, not a forecast.



LP Research model. Assumptions: 85% of capex is tied to new AI capacity (Amazon 70%); 60% of that can be deferred within 12 months; new capacity earns 25¢ per dollar a year; 6% cost of capital; six-year life. Inputs, formulas and company detail in the research appendix. Note 6.

That's the counterintuitive part: a delay that hits the whole industry costs the incumbents little, because customers have nowhere else to go. For the largest cloud providers, a statewide power constraint is far less damaging than a problem at their own sites.

The remaining route to profit is price, and scarcity already carries one. AWS raised GPU reservation prices twice this year, by about 15% and then about 20%,[8] and one-year rental rates for Nvidia's older H100 chips rose almost 40% between October and March.[9]

But scarcity pays slowly, because most cloud revenue sits in multi-year contracts. Microsoft has a $677 billion commercial backlog, and only about 30% of it will be recognized in the next 12 months.[10] To offset a 20% slowdown at Azure's current growth rate, I estimate prices on existing capacity would need to rise about 5% if every contract could reprice immediately, and nearly 14% if only about a third can.[11] Pricing power arrives as contracts roll over. In the first year, the fastest-growing businesses give up the most.

Not every hyperscaler is on the selling side. Meta reported $349 billion of non-cancelable commitments at the end of June, mostly for third-party cloud capacity and infrastructure, on top of $279 billion of leases that hadn't started.[12] In a scarcity market, Meta is a buyer. It gets the capex relief and pays the premium.

When contracts move faster than buildings

The more important question in a slowdown is who has already promised to pay for capacity that doesn't exist yet. Anthropic is the best-documented case. Its amended IPO filing, published on November 10, shows about $900 billion of commitments, counting those signed since the end of September, almost all of it for compute and infrastructure. They include at least $121.3 billion with Google, $110.0 billion with AWS and $47.2 billion with Microsoft, each payable "regardless of actual usage," plus $451 billion of leases that have been signed but haven't started.[13]

The headline is enormous, but the structure is mostly protective. Leases start only when buildings are delivered. About $186 billion of the total carries termination rights or other conditions, and only $70.7 billion falls due in the 12 months after September.[14] If buildings are late, Anthropic mostly loses time, not cash.

There are two exceptions. The first is the cost of filling the gap. When contracted capacity falls short, Anthropic buys short-term compute that "currently costs significantly more than comparable compute procured under our multi-year commitments." It spent $5.0 billion on that kind of compute in the third quarter, a quarter in which it reported a $567 million operating loss after a $974 million profit in the second.[15] That's what a capacity shortage looks like in a lab's income statement.

The second is a set of five-year equipment leases with Broadcom, carrying about $161 billion of rent. They can't be canceled, and the filing adds: "Rent payments may commence earlier in certain circumstances, such as when the readiness of a data center facility is delayed." Here, the cost of a late building lands on the lab.[16]

Elsewhere, the line between commitment and ambition is blurrier. OpenAI and Broadcom announced 10 gigawatts of custom accelerators last year without disclosing terms,[17] and of seven planned U.S. Stargate sites, only Abilene was running in April.[18] Nvidia's guarantee of up to $105 billion of leases at one OpenAI campus is contingent: it phases in only as each lease begins and pays out only if the tenant defaults.[19] I'd treat the 10 gigawatts as a target, not an order book.

The parties with bargaining power, the large labs and hyperscalers, have mostly pushed timing risk down the chain to developers, neoclouds and their lenders. They're the ones left holding hardware that can't earn anything yet. A year ago, Satya Nadella said Microsoft might have "a bunch of chips sitting in inventory that I can't plug in."[20]

Take $10 billion of AI hardware that waits a year for power. If it would have earned 45 cents per dollar in its first year, with its earning power falling by a quarter each year, it gives up about $3.8 billion of cash in today's money. Because each chip generation erodes the value of the last whether or not the old one was running, most of that isn't postponed. It's gone. The roughly $0.8 billion of interest at 8% is paid either way, so it strains liquidity rather than adding to the loss.[21]

The accounting can hide this for a while. Depreciation usually starts only when equipment goes into service, so idle chips don't touch reported earnings. In my scenario, about a year of waiting leaves the hardware expected to earn back less than it cost, which is when impairment becomes a live question. Whether a write-down follows depends on accounting judgments and on rental prices, which this year have been rising.[22]

Winners, losers and unintended consequences

One company's capex is another's revenue, so the obvious conclusion is that slower building hurts chipmakers. I think that's only half right.

When power is scarce, a buyer with a fixed number of megawatts wants the most computing from each one. That favors the newest, most efficient chips, and rules that make electricity more expensive make performance per watt more valuable. But scarcity also raises the value of every chip already plugged in, which slows replacement. My view is that a power shortage changes the mix of chip demand more than its level.

The bigger difference is between merchant and custom silicon. An Nvidia GPU waiting on one customer's late building can usually be sold to someone else. A custom accelerator is designed around one customer's software, network and racks. Broadcom's AI business, about 56% of its revenue and mostly custom chips, is growing quickly, but its schedule depends on a handful of buildings it doesn't control.[23]

The most exposed companies are those that borrowed against future capacity. CoreWeave shows how a business can hold enormous contracted revenue and still be financially fragile. It reports a $132.5 billion backlog and more than 4.6 gigawatts of contracted power, against more than 1.7 gigawatts running. It also carries about $47 billion of debt, held $4.1 billion of unrestricted cash at the end of September and spent about $7.5 billion more than it generated in the third quarter. Interest absorbed 23% of revenue.[24]

Suppose 500 megawatts of its pipeline arrives a year late. Using its own run-rate guidance and margins, I estimate that would push about $5.5 billion of revenue and $3.4 billion of adjusted EBITDA out of the next 12 months, roughly as much operating cash as the company now holds.[25] The backlog wouldn't change at all. About $9 billion of expected customer prepayments and the possibility of contract extensions soften the blow. But CoreWeave's constraint is liquidity, not demand, and with equity worth only about a third of enterprise value, small changes in timing move the stock a lot.[26]

Oracle, far larger but rated one notch above junk, has $288 billion of leases still to start.[27] Both stocks have fallen by more than half in a year.[28]

The least obvious consequence is geopolitical. America and China are constrained by opposite inputs. The U.S. has the best chips and most of the capital and is running short of power. China added more than 430 gigawatts of wind and solar in 2025 alone, and Beijing estimates its data centers run at 20% to 30% utilization.[29] Its shortage is advanced chips.

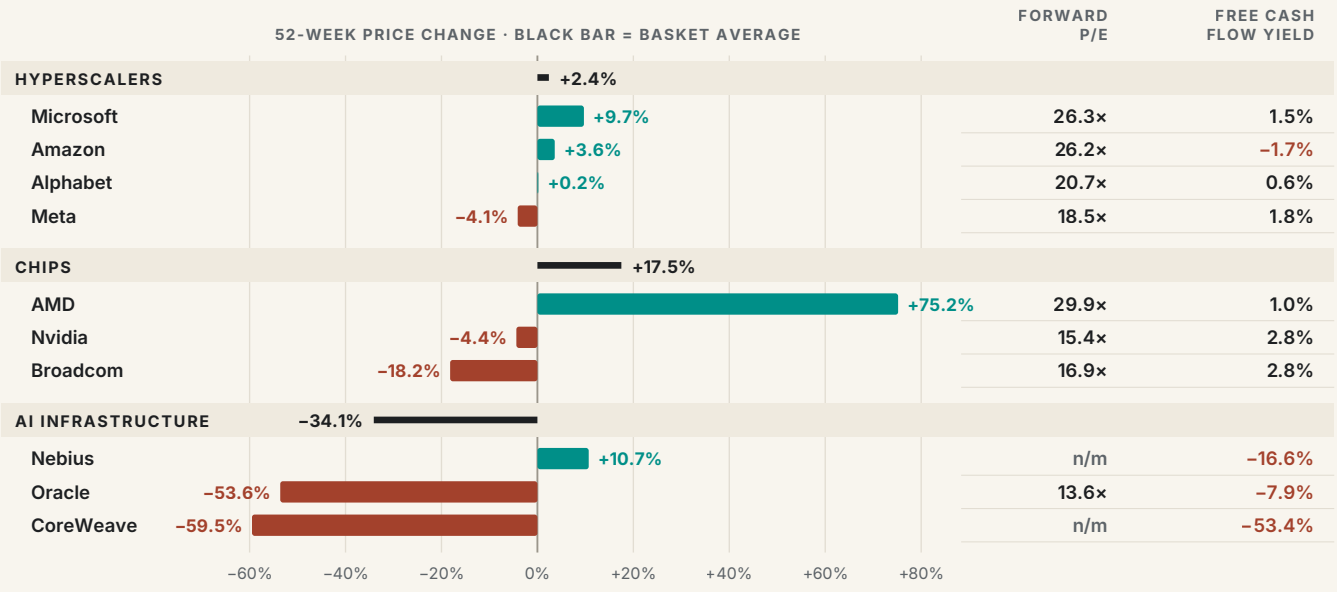
So a policy meant to protect American electricity customers could slow the U.S. in exactly the input where it leads. I won't put a date on any Chinese catch-up; nothing in the data supports one. But national security is the strongest argument for building, it appeals to both parties, and Executive Order 14318 already steers federal land, faster permits and loan support toward large projects.[30] My guess is two tiers: states make ordinary data centers pay their own way, while projects with a national-security label get federal help.

What is the market missing?

If the economics split, the stocks should too. They have.

EXHIBIT 3 **The diverging AI trade**

Share price change over the past 52 weeks, with forward P/E and trailing free cash flow yield.



Source: StockAnalysis, Nov. 6–9, 2026. Price change excludes dividends; basket averages are equal-weighted; forward P/E uses consensus estimates. n/m = not meaningful. Note 28.

The hyperscalers have barely moved, the infrastructure names have fallen by a third and the chipmakers have split.[28] I'm not attributing these moves to power policy; leverage, market share and customer concentration explain most of them. But the market clearly no longer treats AI as one trade.

The valuations raise a more uncomfortable point. The companies best placed to benefit from scarcity, the cloud sellers, carry some of the richest valuations: Microsoft and Amazon trade at about 26 times forward earnings, with free cash flow yields of 1.5% or less. The companies most exposed to a slowdown are already priced for one, with Nvidia at about 15 times and Oracle under 14.[28] A company that benefits from a trend isn't necessarily a good investment if its stock already reflects the benefit.

That makes the scenario the market seems to underweight not a worse shortage but no shortage at all. If the Texas audit clears most credible projects and other states raise prices rather than block construction, capacity arrives on time and the cheapest stocks in the AI trade are the most likely to rerate. This piece would then become a footnote to a simpler story about demand, and I'd rather find that out early than defend a framework the data have left behind. What I'm watching:

- ERCOT's December 10 report and how many verified projects connect on time.[3]
- The gap between running and contracted capacity at CoreWeave and the Stargate sites.
- GPU rental and reservation prices as new capacity comes online.
- Whether hyperscalers blame future capex cuts on sites and deliveries, or on demand.

In the last phase of this cycle, the scarce input was chips, and the company that made them captured most of the profits. In the next phase, the scarce input may be a powered building. The market may stop rewarding companies for how much capacity they promise to build and start rewarding them for how profitably they run the capacity they already own.

LP Research Investment Scorecard

What the analysis shows, for a reader with a minute.

THE NUMBERS THAT MATTER

REPORTED

$615B

Hyperscaler capex, 12 months to September

91% of operating cash flow, up from about a third in 2017. Free cash flow fell from ~$205B in 2023 to $61.5B.

SCENARIO ESTIMATE

~$51B

Free cash flow gained from a 20% slowdown

Value still falls: under $1B if customers wait, about $10.5B if they go elsewhere.

DISCLOSED, MIXED CATEGORIES

~$900B

Anthropic's aggregate commitments

Spans obligations, unstarted leases and newer contracts running to 2045; $70.7B falls due in the 12 months after September.

REPORTED

$132.5B vs **~$47B**

CoreWeave backlog against debt

The backlog pays as capacity is delivered; the debt is due on schedule. Unrestricted cash: $4.1B.

WHERE THE ECONOMICS DIVERGE

Group	Could benefit from	Exposed to	What prices show
Operational hyperscalers Microsoft, Amazon, Alphabet	Running capacity reprices as contracts roll	Slow repricing; more cash, not more value. Meta is a buyer and pays the premium	Flat in a year; Microsoft and Amazon at ~26× forward
Chip manufacturers Nvidia, Broadcom, AMD	Scarce power rewards performance per watt; GPUs can be resold	Late buildings delay installs; custom chips are tied to one customer	Split: AMD +75%, Broadcom −18%; Nvidia at ~15×
AI infrastructure developers CoreWeave, Oracle, Nebius	Big backlogs; scarce capacity is worth more once delivered	Debt and interest run on time while revenue waits for power	Down 34% on average; CoreWeave and Oracle down more than half
AI labs Anthropic, OpenAI	Most commitments start on delivery, so delays cost time, not cash	Costlier short-term compute fills gaps; some rent can start before a site is ready	Not listed; the cost shows up in filings

THE INVESTMENT CONCLUSION

> The companies that gain most from scarce power aren't automatically the best stocks. Microsoft and Amazon own the capacity that becomes more valuable, but they already trade at about 26 times forward earnings, and their pricing power arrives only as contracts roll. The sharper risk sits with businesses that borrowed against capacity they can't yet switch on: their debt comes due on schedule while their revenue waits for the grid. Much of that is already in prices. The outcome the market may be underpricing is a constraint that never binds.

WHAT COULD CHANGE THE THESIS

Grid connection approvals

SUPPORTS

ERCOT's Dec. 10 report clears few credible projects; New York extends its pause

UNDERMINES

Most verified projects connect on schedule

Compute pricing

SUPPORTS

GPU rental and reservation prices hold or rise as capacity arrives slowly

UNDERMINES

Prices fall as new capacity comes online

Contracted vs. running capacity

SUPPORTS

The gap widens (CoreWeave: more than 4.6 GW contracted vs. more than 1.7 GW running)

UNDERMINES

The gap closes on schedule

Sources: $615B, note 1; ~$51B, note 6; ~$900B, notes 13–14; CoreWeave, note 24; market data, note 28; indicators, notes 3, 8–9 and 24. Scenario figures come from LP Research models and are illustrations under stated assumptions, not forecasts.

Notes and sources

1. Operating cash flow and purchases of property and equipment for the 12 months to Sept. 30, 2026, from StockAnalysis cash-flow statements compiled from company filings (Microsoft, Amazon, Alphabet, Meta) and checked against September-quarter releases (Microsoft via sec-api, Alphabet, Amazon, Meta). Microsoft: OCF $195,878M, capex $135,921M (fiscal years end in June). Amazon: OCF $168,300M, capex $215.2B gross ($210,657M net of proceeds and incentives). Alphabet: OCF $179,551M, capex $159,804M. Meta: OCF $133,087M, capex $104,438M. Combined: OCF $676.8B, capex $615.4B, or 91%; free cash flow $61.5B. In 2017 the four reported combined OCF of $119.2B and capex of $40.0B (34%), and in 2023 free cash flow of about $205B, from 10-K data via finboard.net (Microsoft, Amazon, Alphabet, Meta). Ratios and totals are LP Research calculations.
2. Akin Gump, alert on Governor Greg Abbott's Aug. 3, 2026 directive requiring an audit of large loads seeking to connect to ERCOT before further approvals, with a request queue of about 474 GW against ERCOT's record peak demand (akingump.com).
3. Pillsbury, Aug. 24, 2026: on Aug. 20 the Public Utility Commission let ERCOT keep processing applications on a conditional basis while it verifies roughly 250 to 300 projects totaling about 200 GW (pillsburylaw.com). ERCOT, update to the PUC, Nov. 5, 2026: information requests reached 465 loads totaling about 177 GW, with a final report targeted for Dec. 10 (ercot.com).
4. Route Fifty, July 2026: Governor Kathy Hochul's July 14 order pausing state environmental permits for data centers of 50 MW or more for up to a year (route-fifty.com).
5. Debtwire (ION Analytics), review of state legislation to June 14, 2026: 16 states had enacted 30 data center laws, mainly on cost allocation, ratepayer protection and incentives; none of the moratorium bills had survived, and Maine's and Vermont's were vetoed (ionanalytics.com). Examples by type are in appendix table B1.
6. LP Research model (research appendix, sections C2–C3 and table B3); illustrative, not a forecast. Capex deferred = slowdown (20%) × share of capex tied to new AI capacity (85%; Amazon 70%) × share deferrable within 12 months (60%) × LTM capex (Amazon net of incentives) = $58.5B. Free cash flow rises by the deferred amount less half a year of forgone return at 25¢ per dollar: $51.2B in total (Amazon $15.5B, Alphabet $14.3B, Microsoft $12.1B, Meta $9.3B), or 7.6% of combined LTM operating cash flow (note 1). Value change from a one-year delay at a 6% cost of capital over a six-year life: $-\text{NPV} \times r/(1+r) = -\$0.76\text{B}$ if demand waits, and $C \times (r-y)/(1+r) = -\$10.5\text{B}$ if that year's demand is lost. Every input other than reported capex and cash flow is an assumption.
7. LP Research, *The Central Bank of AI*, Oct. 28, 2026, note 12: capital recovery factor $0.06/[1-(1.06)^{-6}] = 0.203$, or 0.223 at 9%, a simplified financing requirement rather than a full return hurdle. The 25¢ annual return on new capacity is an LP Research assumption; at 25¢ the net present value is about 23¢ per dollar invested.
8. Computing (press report), 2026: AWS raised prices on its GPU capacity reservations by about 15% at the start of the year and about 20% in July, saying prices are "updated periodically based on supply and demand" (computing.co.uk).
9. SemiAnalysis data, as reported by Tiger Brokers (secondary): one-year H100 rental contract prices rose almost 40% between October 2025 and March 2026 (itiger.com).
10. Microsoft Form 10-Q for the quarter ended Sept. 30, 2026 (fiscal Q1 2027), as summarized by sec-api.io (secondary): commercial remaining performance obligation of $677B, about 30% of it expected to be recognized over the next 12 months; Azure revenue growth of 48% (sec-api.io).
11. LP Research estimate (appendix section C4). Price increase on existing capacity needed to offset the lost growth = $s \times g \times m / \varphi$, with a 20% slowdown (s), 48% growth (g, note 10), a 50% cash margin (m, assumption) and φ, the share of revenue that can reprice within a year: 4.8% if $\varphi = 100\%$, 13.7% if $\varphi = 35\%$. The 35% is an assumption anchored on Microsoft's roughly 30% of backlog recognized within 12 months.
12. Arkolith, citing Meta's Form 10-Q for the quarter ended June 30, 2026 (secondary): $349B of non-cancelable contractual commitments, mostly for third-party cloud capacity and infrastructure, plus $279B of leases signed but not yet commenced (arkolith.com).
13. Anthropic Form S-1/A, Amendment No. 2, filed Nov. 10, 2026 (sec.gov). Aggregate commitments of about $900.9B: $320.0B of contractual obligations at Sept. 30 (including $313.8B of purchase commitments), $451.2B of leases executed but not commenced ("expected to commence between 2026 and 2028," weighted-average term of about 10 years), $17.4B of short-term commitments, and about $105.7B of longer-term and $6.6B of short-term commitments signed after Sept. 30. Named minimums, each payable "regardless of actual usage": Google at least $121.3B (April 2026 to July 2033), AWS at least $110.0B (May 2026 to April 2036) and Microsoft at least $47.2B and up to $49.2B (November 2026 to June 2033). Outside the total: an estimated $125B to $150B of additional Google TPU compute, $150B to $175B of arrangements with no fixed amount, and AMD purchases expected to exceed $20B. June 30 comparisons are from the original Form S-1 (sec.gov).
14. Same filing: about $186.2B of the total carries termination rights or contingent conditions; $70.7B is expected to be paid in the 12 months after Sept. 30, of which $4.2B is conditional; $451.6B is expected before Dec. 31, 2030, and the rest from 2031 to 2045. Capacity from Nscale and Lambda accrues fees only as it is delivered and accepted. The filing warns that "because much of the related capacity has not yet been delivered, the timing of our payments, and in some cases the amounts, may differ from those shown."
15. Same filing: Q3 2026 results include "$5.0 billion of compute under short-term commitments," which "currently costs significantly more than comparable compute procured under our multi-year commitments"; Q3 operating loss of $567M. Q2 2026 operating income of $974M is from the original Form S-1. The company says it generally procures compute "12 to 18 months in advance of availability."

16. Same filing: "five-year master equipment leases" with Broadcom, "non-cancelable by either party during the applicable lease period," with about $161.2B of rent and payments "typically commencing upon AI infrastructure equipment becoming 'ready for service.'" The filing adds: "Rent payments may commence earlier in certain circumstances, such as when the readiness of a data center facility is delayed." Its risk factors note that "utility providers and regulators may also reallocate previously committed capacity that our data center facilities may require."
17. Reuters, October 2025 (via TradingView): OpenAI and Broadcom agreed to deploy 10 GW of OpenAI-designed accelerators from the second half of 2026 through 2029, without disclosing financial terms; Reuters noted it was unclear how OpenAI would pay (tradingview.com).
18. Epoch AI, April 2026: of seven U.S. Stargate sites with more than 9 GW planned, only Abilene was operating, at about 0.3 GW, and five were scheduled for late 2028 (epoch.ai).
19. NVIDIA Form 10-Q for the quarter ended July 26, 2026, Note 10 (sec.gov): "In August 2026, we entered into guarantees, capped at a total of $105 billion," for SB Energy leases at a campus OpenAI will use; "Each guarantee generally becomes effective upon commencement of the applicable lease, the first of which is expected in fiscal year 2029," and payment is "triggered upon certain tenant defaults." See also *The Central Bank of AI*, notes 3 and 7.
20. Tom's Hardware, November 2025, reporting Satya Nadella's remarks on Microsoft's power constraints (tomshardware.com).
21. LP Research model (appendix section C5 and Exhibit A3); illustrative. $10B of hardware already owned; first-year cash return of 45¢ per dollar (35¢ and 55¢ also tested); earning power declining 25% a year from purchase; six-year life; 8% financing cost. Present value of cash not earned during a 12-month delay: $3.77B (range $2.93B to $4.61B). Interest: 8% × $10B = $0.8B, paid whether or not the hardware runs, so not added to the economic loss. Under these assumptions hardware needs a first-year return of about 41¢ to earn back an 8% cost of capital.
22. Under U.S. GAAP (FASB ASC 360-10), depreciation generally begins when an asset is placed in service, and a long-lived asset is tested for recoverability, by comparing its carrying amount with its undiscounted future cash flows, when events suggest the carrying amount may not be recoverable; any impairment is then measured against fair value. In the LP Research scenario, remaining undiscounted cash after a 12-month delay at 45¢ is $8.95B against a $10B carrying value (appendix table B7). Real outcomes depend on how assets are grouped, management's assumptions and market rental prices (notes 8–9).
23. TIKR, summary of Broadcom's fiscal Q3 2026 results and call (secondary): AI revenue about 56% of total revenue, about 73% of it from custom accelerators; customers include Google, Anthropic and OpenAI (tikr.com).
24. CoreWeave Q3 2026 earnings release, Form 8-K (sec.gov): revenue $3.66B; interest expense $846M; revenue backlog $132.5B; more than 1.7 GW of active power and more than 4.6 GW contracted; cash and equivalents $4.1B ($6.4B including restricted cash); free cash flow of about −$7.5B for the quarter. Debt of about $46.9B is the sum of the balance-sheet debt lines, and interest at 23% of revenue is an LP Research calculation.
25. LP Research estimate (appendix section C6 and table B9), not company figures. CoreWeave's guidance of a $20B to $21B revenue run rate exiting 2026 on about 1.85 GW (note 24) implies about $11.1B of revenue per GW-year. A 12-month delay on 500 MW pushes out about $5.5B of revenue and, at the Q3 adjusted EBITDA margin of 62%, about $3.4B of adjusted EBITDA.
26. CoreWeave Q3 2026 release (note 24): about $9B of customer prepayments and deposits expected over the next three months. Market capitalization of $25.6B and enterprise value of $71.6B per StockAnalysis, Nov. 6, 2026 (equity 36% of enterprise value). In 2025 a CoreWeave customer extended its contract after a third-party developer's delay (Reuters via TradingView).
27. Oracle Form 10-Q for the quarter ended Aug. 31, 2026: "$288 billion of additional lease commitments, substantially all related to data center arrangements," expected to commence between fiscal Q2 2027 and fiscal 2029 (sec.gov). S&P's downgrade to BBB− on July 9, 2026, as reported by PPC Land and cited in *The Central Bank of AI*, note 14.
28. StockAnalysis statistics pages dated Nov. 6–9, 2026 (stockanalysis.com): 52-week price change, forward P/E on consensus estimates, and trailing free cash flow yield. Microsoft +9.7%, 26.3×, 1.5%; Amazon +3.6%, 26.2×, −1.7%; Alphabet +0.2%, 20.7×, 0.6%; Meta −4.1%, 18.5×, 1.8%; AMD +75.2%, 29.9×, 1.0%; Nvidia −4.4%, 15.4×, 2.8%; Broadcom −18.2%, 16.9×, 2.8%; Nebius +10.7%, n/m, −16.6%; Oracle −53.6%, 13.6×, −7.9%; CoreWeave −59.5%, n/m, −53.4%. Equal-weighted basket averages (LP Research calculation): hyperscalers +2.4%, chips +17.5%, AI infrastructure −34.1%. Price changes exclude dividends.
29. Al Jazeera, May 28, 2026, citing BloombergNEF, the IEA and Chinese government estimates: China added more than 430 GW of wind and solar capacity in 2025, and Beijing estimates its data centers run at 20% to 30% utilization (aljazeera.com).
30. White & Case on Executive Order 14318, "Accelerating Federal Permitting of Data Center Infrastructure," July 2025: directs agencies to speed permitting, make federal land available and offer financial support, including loans and loan guarantees, for qualifying data center projects, including those of 100 MW or more (whitecase.com).

Figures are as reported in company filings and releases available on Nov. 10, 2026, unless noted; market data are as of Nov. 6–9, 2026. Secondary summaries and press reports are labeled as such. Figures described as LP Research estimates or models are scenario results under stated assumptions, not forecasts; the models, inputs and supporting tables are in the separate research appendix.

This commentary is for information only and is not investment advice. It isn't a recommendation to buy or sell any security.