The AI Bubble Is Running Out of Air

How the gold rush, the dot-com crash and the 2028 election point to the same ending for AI money



Written by Lior Alexander

The technology is real. The money is the bubble. It takes 31 minutes on average to read.

You can see the end of a boom first at a county planning meeting.

That is where the costs stop being theoretical. In Loudoun County, Virginia, and Hill County, Texas, residents have been debating the industrial footprint of a technology they may never use directly. The argument has escaped the folding chairs. It now reaches governor's races and Congress. Texas's governor, once a champion of data center expansion, has paused approvals. This is what a financial story looks like before financial markets recognize it.

I think the AI frenzy has about two years left.

I don't mean the technology. I build with AI every day, and I think we're still underestimating what it will do. I mean the financing: valuations that assume permanent acceleration, debt that depends on decades of demand, and infrastructure ordered long before its customers have earned enough to pay for it. Those commitments are real. The future used to justify them is not yet here.

I want this boom to last. I don't think it will.

A gold rush begins with a discovery, not a mistake. The trouble starts when people stop paying for what is there and start paying for what the next investor will believe. New money pushes prices up. Higher prices summon more money. Eventually the crowd becomes part of the product.

Railroads did it. The internet did it. Crypto did it. None needed its underlying technology to be worthless.

AI has an additional vulnerability: the people who will decide what the boom costs don't all own a piece of it. And they have votes.

The 2028 presidential election may turn AI into a rare bipartisan target. Ratepayers see larger bills. Towns see new industrial campuses. Workers see software advertised as a substitute for their labor. The companies financing the expansion see a glorious future. Many voters see an invoice arriving first.

For now, the White House has provided cover. A new president could change that. And a politician need not promise to stop AI to make the economics of new data centers less attractive. 'Build it, but pay for your own power' may be enough. Industry leaders are already trying to shape the rules. Voters may ask a different set of questions.

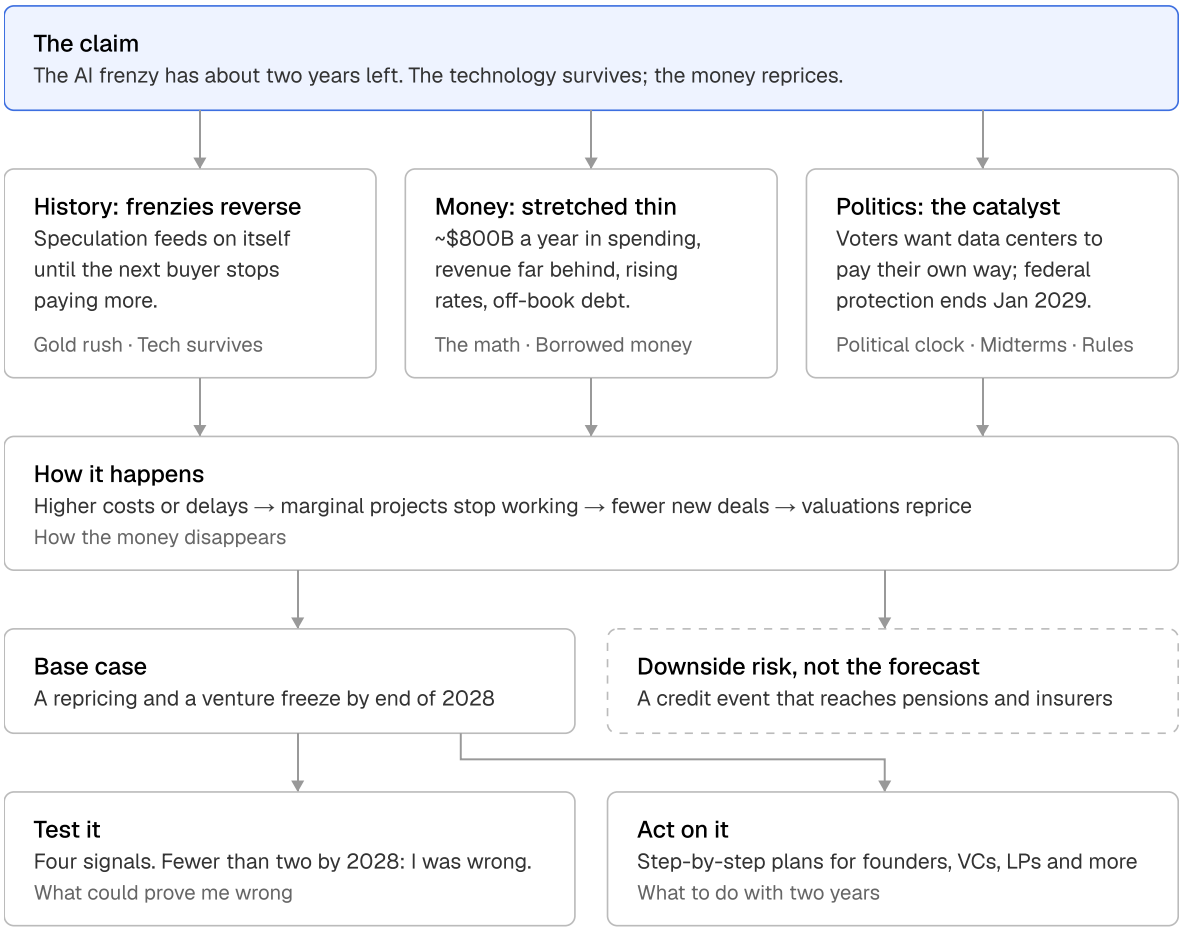
Investors cannot wait for the law to change. They have to price what might change. If power becomes more expensive, permits slower, and subsidies less secure, the value of a project falls today—even if it will not open for years. The first casualties need not be companies that fail. They may be projects that nobody decides to finance.

That is my forecast: AI keeps working, but the money behind its expansion becomes harder to obtain at today's prices. The reversal may already be underway. I expect the 2028 campaign to turn scattered doubts into a broader retreat. The interesting question is not whether AI wins. It's who pays, and who is still willing to pay, while it does.

How this essay is structured

The argument, from claim to action

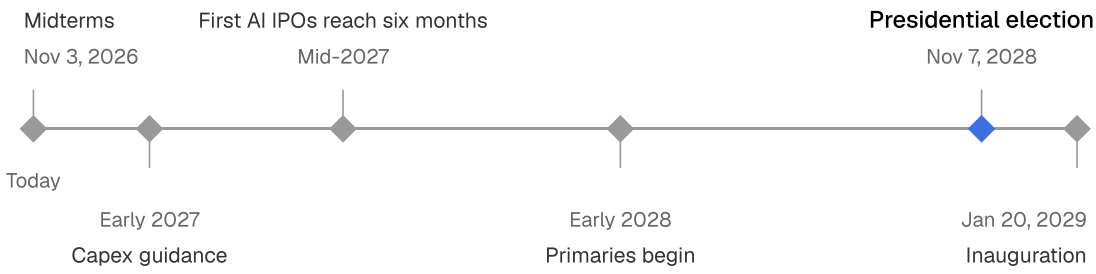
Each step names the sections of the essay that make it



THE ARGUMENT OF THIS ESSAY · CLAIM, THREE REASONS, MECHANISM, OUTCOME, TEST AND ACTION

The next two years run on a calendar

Checkpoints between now and the next inauguration, to scale



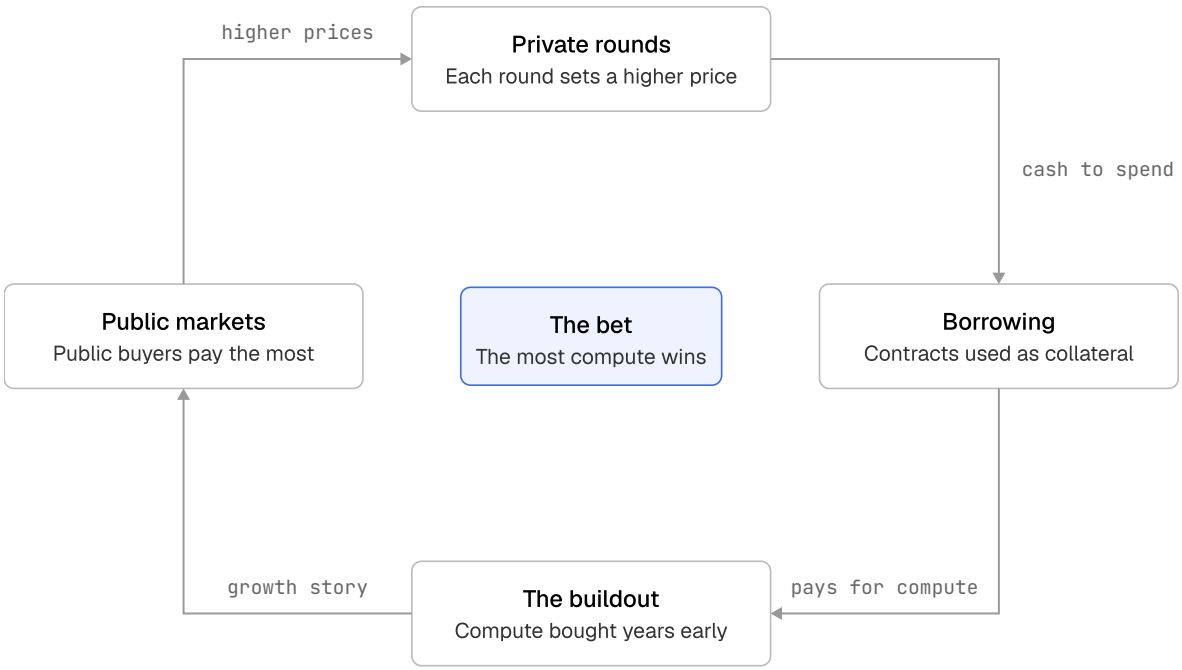
KEY DATES FOR THE FORECAST · MIDTERMS, EARNINGS SEASON, IPO SIX-MONTH MARKS, PRIMARIES, ELECTION, INAUGURATION

How the AI money machine works

AI's financing problem is a problem of time. Computing power is bought today. The revenue that must justify it is expected tomorrow. Between those dates lies an enormous bill, and someone must carry it.

How the AI money machine feeds itself

Illustrates reinforcing incentives, not required cash flows: each source can also act on its own



HOW THE BIGGEST AI COMPANIES FUND THE BUILDOUT · THREE INTERCONNECTED SOURCES OF CAPITAL

The money comes from three directions. Private investors fund labs in exchange for a claim on future profits. Lenders fund chips and buildings against promised customer payments. Public markets set prices for the suppliers, customers and increasingly the AI companies themselves. They aren't one pipeline. But they are exposed to one shared assumption: demand will arrive fast enough.

Consider the first link. OpenAI's $122 billion financing round and Anthropic's $965 billion valuation aren't just remarkable prices. They're expectations about future earnings, brought into the present as spendable capital. The labs use that capital to rent computing power and buy the capacity they hope will make those earnings possible. Some investors also supply the chips and cloud services being purchased. The same dollars can help finance a company and then show up as revenue at its supplier.

Now follow the promise one step further. A developer can borrow against a long-term contract and use the proceeds to build. The tenant gets compute without writing the entire check today; the lender gets a schedule of payments. If the tenant is Meta, that may be conventional infrastructure finance. If the tenant must raise another round just to pay its bills, the contract rests on a fundraising forecast. A promise supports a loan that supports another promise.

The public market adds a price tag to those expectations. An IPO can replenish capital and let earlier investors sell, but not every lab needs one and not every data center depends on one. Profitable hyperscalers can pay from operating cash flow. My concern is narrower: the places where optimistic forecasts, expensive financing and the need for another check all meet.

No one has to be secretly coordinating this. Each participant can make a reasonable decision. The problem is that many of those decisions depend on the same future. AI already earns real revenue. The question is whether it earns enough, soon enough, to pay for the future everyone is building in advance.

Every gold rush ends

A frenzy is not defined by stupid investors. It is defined by a price that needs the next buyer more than the underlying business. Early profits invite new money; new money raises the price; the price becomes evidence of its own wisdom. This can continue long after participants begin to suspect the ending.

During the California gold rush, the merchant Sam Brannan made his fortune supplying miners, not finding gold. The modern shovel sellers are the companies supplying chips, electricity and data centers. They can prosper even when their customers struggle to earn adequate returns. That doesn't make their revenue imaginary. It makes the distinction between *selling into a boom* and *profiting from its ultimate use* essential.

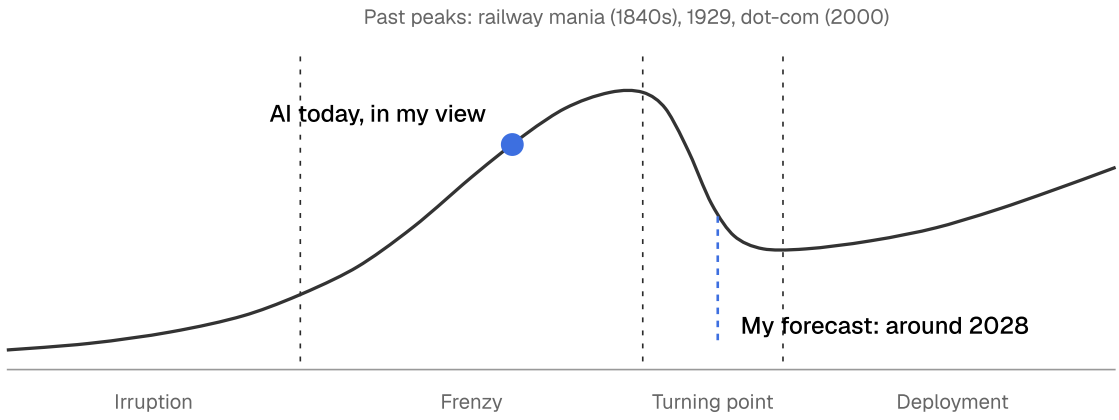
The pattern has a name

Economist Carlota Perez describes technological revolutions as a long installation, a speculative frenzy and a turning point before the technology becomes broadly useful. That is a framework, not a countdown timer. Its value here is a distinction: a technology can change the world while the investors who financed its first expansion lose money.

Here's my departure from her pattern. Government often writes the new rules after the frenzy breaks. In AI, the expectation of those rules may help break the frenzy first.

My reading: where AI sits in Perez's pattern

My interpretation mapped onto Carlota Perez's framework. A schematic, not a forecast from her model.



SCHEMATIC AFTER CARLOTA PEREZ, TECHNOLOGICAL REVOLUTIONS AND FINANCIAL CAPITAL (2002)

What the market is asking investors to believe

Public markets are also being asked to pay for future success today. SpaceX, which absorbed xAI, listed in June at around $1.75 trillion. Anthropic is reportedly seeking a valuation above $2 trillion in its own planned listing. These aren't proof of a top, but they show how much future success is already priced into the present.

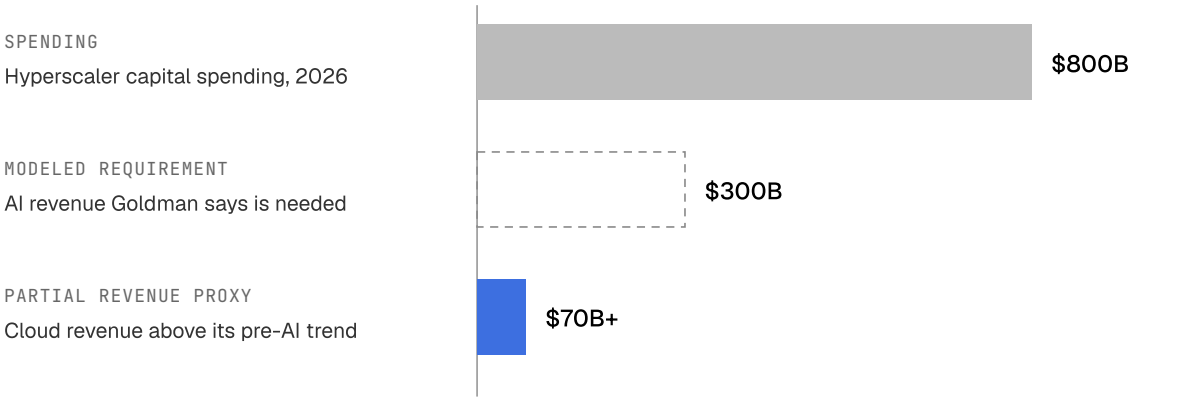
The math the frenzy has to clear

The question that matters more than any valuation is whether the money comes back. Goldman Sachs estimates the five largest US tech companies will spend about $800 billion this year, much of it tied to AI, and as much as $1.2 trillion next year. Its model puts the revenue needed to support that spending at roughly $300 billion annually.

There is no clean, industry-wide AI revenue line. One imperfect proxy—cloud revenue above its pre-AI trend—is closer to $70 billion a year. That excludes returns embedded in software, advertising and other businesses; comparing the two figures doesn't prove a $230 billion hole. It does show how much confidence today's spending requires about revenue that isn't yet directly visible. Goldman itself says the companies have a long way to go.

Three different numbers, not one measured gap

Goldman Sachs estimates, per year · each bar measures something different



The proxy misses AI revenue earned through ads, software and devices. The bars show scale, not a shortfall.

GOLDMAN SACHS ESTIMATES, SEPTEMBER 2026 · CLOUD REVENUE ABOVE ITS PRE-AI TREND USED AS A PROXY FOR AI REVENUE

The bullish case is not imaginary. Cloud demand is growing, backlogs are substantial, and some skepticism is already in market prices. But this is a race against a financial clock. Customers have to arrive, pay enough, and keep paying while the infrastructure depreciates and the debt accrues interest. Being right about the destination is not enough when you have borrowed against the arrival date.

The technology survives the bubble

History does not tell us when AI's financing boom will end. It tells us why a technology's eventual importance is a poor defense of what investors paid for it. Three precedents matter, for three different reasons.

Railroads: the infrastructure can outlive the capital

Railroads transformed America, but building them was ruinously expensive. Jay Cooke's firm collapsed in 1873 after overextending itself on a second transcontinental line. The financial panic damaged the economy; the tracks remained. The lasting value of infrastructure didn't guarantee a return for the people who financed its construction. AI data centers may have the same split between social usefulness and investor profit.

The internet: a financing loop can look like demand

In 2000, the internet was not the thing that failed. Financing was. As rates rose, companies counting on one more round found that the next investor had disappeared. The Nasdaq ultimately lost 78% of its value. The websites did not vanish. The prices attached to them changed.

Telecom equipment makers added another vulnerability. Lucent and Nortel financed customers who then bought more equipment from them. When those customers failed, the suppliers lost both sales and the loans supporting those sales. Nortel's shares fell from a peak of $86.75 to pennies. The fiber was used later. That didn't rescue the original balance sheets.

A modern echo is Nvidia's proposed investment of up to $100 billion in OpenAI, much of which was expected to return as chip purchases. The eventual Nvidia commitment was scaled back, but the relationship illustrates the point: an investment by a supplier can finance demand for that supplier. Such arrangements aren't necessarily improper or doomed. They do mean reported orders and independent end-customer demand are not quite the same thing.

Crypto: regulation isn't enough by itself

Crypto offers a caution against easy political storytelling. Funding plunged as rates rose and failures accumulated. Yet the SEC's 2023 cases against Coinbase and Binance did not permanently crush Coinbase's shares. Regulation by itself isn't a magic pin. It becomes powerful when investors are already reconsidering leverage, growth and cash flow.

What's different this time

AI has pieces of all three stories: durable assets, interdependent financing and a sector sensitive to confidence. It also faces higher financing costs. The Fed raised rates in September, and the 10-year Treasury yield has moved above 5%. That raises the return a new project must earn before investors will build it.

The financial conditions are already tightening. My argument is that politics could give investors the reason, and the deadline, to stop treating today's assumptions as permanent.

The political clock is ticking

AI's expansion has an unusual political weakness: its costs are visible before its benefits are widely shared. Communities see construction, households see power bills, and workers see a threat to jobs. That gives a presidential candidate an opening even without opposing the technology.

The current White House has largely favored the industry's freedom to expand. In December 2025, Trump ordered a federal effort to challenge state AI laws. Congress had declined broad state-law preemption, including a 99–1 Senate vote. That makes federal protection real but politically vulnerable: it depends heavily on an administration leaving office in January 2029. Even Republicans such as Florida governor Ron DeSantis have opposed the executive approach.

The coalition is already forming

These concerns come from different places, but they can converge at the ballot box. Anxiety about jobs broadens the audience; electricity bills and local projects give candidates concrete policies to propose.

The electricity issue is particularly tangible. In the PJM grid, an analysis attributed 63% of a capacity-price increase to data centers. A Virginia poll found widespread concern about the effect on power bills. And a Reuters/Ipsos poll found majorities in both parties favoring stricter AI regulation.

General support for AI regulation isn't a vote for any particular law. The distinction that matters here is between regulating model safety and changing the economics of expansion.

The policy that changes the numbers

Here the evidence becomes more relevant to investment. A national poll found 66% want data centers to pay for the grid upgrades they require. An industry-commissioned survey likewise found strong support for having data centers cover their own electricity costs. The House voted overwhelmingly for regulators to consider allocating those costs to large users.

A candidate can promise to win the AI race against China while insisting that its builders cover more of their infrastructure costs. That's an electoral message with a financial consequence: higher expected costs can lower the value of future projects.

Not every project would face a new bill. Some already pay for grid upgrades or pass electricity costs to tenants. The risk is concentrated in projects priced around favorable rates, incentives, and fast permits. The forecast depends on enough future investment relying on those assumptions.

Potential 2028 candidates are already exploring that ground. Josh Shapiro has moved from promoting AI investment in Pennsylvania toward restricting data center expansion. Alexandria Ocasio-Cortez has advocated a moratorium. Josh Hawley has been sharply critical of AI despite his hawkish stance on China. They disagree about almost everything else. That is precisely why a shared critique of who bears AI's costs matters.

A presidential candidate doesn't have to persuade voters that AI itself is bad. They can argue that AI companies should pay for the disruption they create. My bet is that a serious 2028 candidate will find that argument useful—and that markets will react to the prospect of changed rules before any election or law settles it.

Why 2028? Presidential campaigns turn diffuse unease into national policy proposals that investors can assign probabilities to. Markets reprice anticipated costs; they don't wait for a law to pass.

The timing is conditional. If the issue stays local, projects may simply move. For 2028 to become the catalyst I predict, national candidates must make the change credible and broad enough to affect growth assumptions across the industry.

The midterms: a dress rehearsal

The November 3 midterms provide an early test of my thesis: can a local backlash against AI infrastructure become a campaign message that candidates carry into office? The results deserve scrutiny, but a victory alone cannot establish that data centers determined a race.

The backlash is already visible across states. In at least 16 governor's races, candidates have called for temporary moratoriums. The point isn't that sixteen states will stop building. It's that an issue once confined to zoning hearings has become something statewide campaigns must answer.

What would actually count as a political signal?

The test isn't simply who wins. It's whether candidates campaign on infrastructure costs, act on those promises, and change the expectations of investors exposed to new construction.

RESULT	WHAT I'D INFER
Critics win in both parties and follow through	Stronger evidence that cost allocation and tighter permits can become a national 2028 platform. A victory alone wouldn't prove data centers caused it.
Mixed results and uneven policy	A serious but regional cost problem; developers may move projects rather than stop them. The national forecast remains open.
Critics lose and campaign pauses disappear	A substantial challenge to the claim that backlash will drive the 2028 campaign. I would revise the political timetable.

The sequence to watch now is campaign message, election result, policy follow-through, then market reaction. Only by distinguishing those steps can we tell whether a local grievance is becoming a national financing risk.

The industry wants to write the rules

The industry's regulatory posture looks contradictory only until you ask which rules it wants. Some AI executives are calling for stronger oversight while their companies spend heavily to shape it. Those are not necessarily opposing positions. Rules can constrain a market and still favor the firms best equipped to comply.

AI executives have advocated safety standards and international coordination, and leading companies signed a voluntary White House accord that leaves much of the oversight in their own hands. The accord even anticipates that its principles might eventually become law. This is a bid to shape the terms on which regulation arrives, not simply to avoid it.

The industry is divided on the details. Anthropic and supporters of more stringent safeguards are funding political groups, while other AI investors are supporting efforts against state-by-state rules. Collectively, AI interests have committed large sums to the midterms. But the arguments largely revolve around safety, audits, standards and which level of government writes them. Those are rules the largest firms can plan for—and sometimes absorb more easily than smaller competitors.

That creates a political risk the industry cannot simply lobby away. Campaign contributions may help companies promote the safety rules they prefer, but they cannot dictate what voters demand or what an elected government ultimately does. A candidate who welcomes AI-industry support for oversight could also embrace tougher requirements on electricity costs, subsidies and permits. The companies may help make regulation politically viable, only to find that the coalition supporting it wants a different kind of regulation. Buying influence over the first draft is not the same as owning the second.

But there is another argument building outside the industry's preferred vocabulary. Voters aren't necessarily asking who should audit a model. They're asking why a company promising vast profits should pass part of its electricity bill to everyone else. Safety regulation affects how a firm operates. Cost allocation affects whether its next project is worth building.

The railroads tried this

The railroad precedent is especially revealing. Historian Gabriel Kolko argued that some railroads welcomed federal regulation in the nineteenth century because it offered an alternative to a patchwork of hostile state rules and unpredictable price competition. They got a relatively weak federal regulator in 1887.

But the institution they helped bring into being didn't remain under their influence. Farmers, merchants and workers pushed for stronger controls, and the 1906 Hepburn Act gave the commission greater rate-setting power. The industry had helped write the first version of the rules. It couldn't control the coalition that demanded the second.

The parallel has limits. AI doesn't have a settled federal regulator of the kind railroads eventually faced. What it has is executive protection in some areas and voluntary promises in others. Those are easier to change. The next set of rules may be written by people who never accepted the industry's first draft.

The companies may help write the first draft. They cannot assume they'll write the second.

The boom is built on borrowed money

If investors change their minds about AI, why should that affect companies with real buildings and signed contracts? Because the industry's expansion has put a large amount of borrowed money between today's spending and tomorrow's hoped-for income. Morgan Stanley expects AI-related borrowing to approach $570 billion this year. But the interesting question isn't the total. It's whose promise makes the debt safe—and what happens if that promise becomes less attractive.

A strong borrower can make a risky project financeable

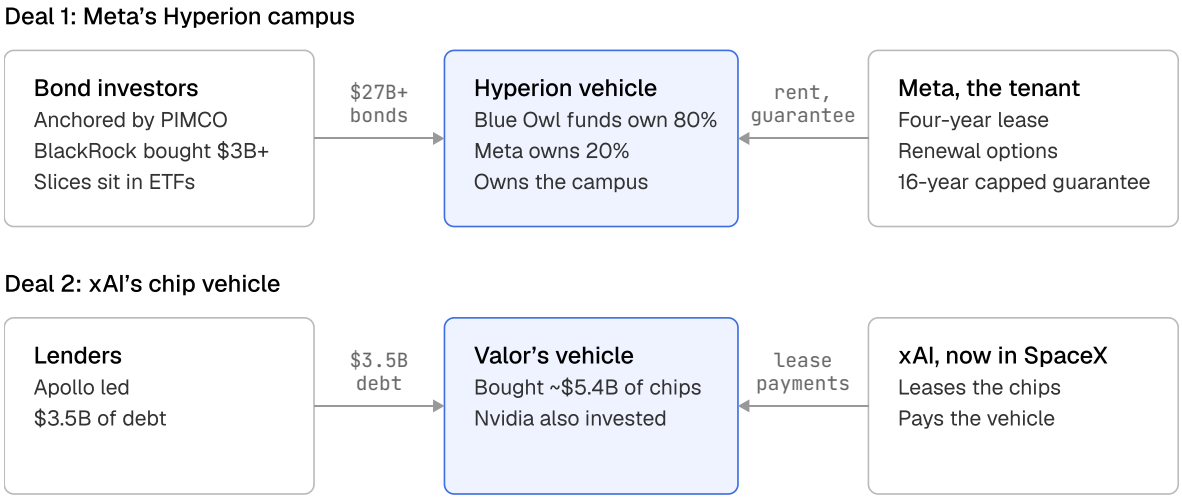
Meta's Hyperion campus in Louisiana shows the safest version of an elaborate arrangement. Rather than simply borrowing against its own balance sheet, Meta formed a venture with funds managed by Blue Owl. The venture raised more than $27 billion, outside Meta's consolidated debt; Meta kept a minority stake, arranged to lease and operate the campus, and received a $3 billion distribution.

That might sound like Meta transferred the risk away. The contracts show the limits of that interpretation. Its initial lease is only four years, so it also provided a residual-value guarantee that Moody's has put at up to $28 billion under specified circumstances. The guarantee helps make long-lived bonds attractive despite a short initial lease. In effect, the project's lenders are relying substantially on Meta's ability to honor its contract.

Hyperion isn't where I'd expect widespread defaults to begin: Meta is a strong borrower. Its lesson is that moving debt into a separate vehicle doesn't eliminate exposure. Investors can mark down even well-protected bonds when expected returns change, without any missed payment.

In both deals, the debt sits in a separate vehicle

Arrows show money flowing in. The two deals involve different parties and different risks.



Across AI private credit generally, the ECB says banks hold the safer senior loans, while pension funds and insurers hold the riskier layers. That is not specific to either deal.

TWO SEPARATE DEALS, FROM META'S HYPERION ANNOUNCEMENT AND REPORTING ON THE VALOR-XAI VEHICLE; ECB NOTE ON AI PRIVATE CREDIT

When the promise rests on the next customer payment

Now move down a level, to cloud providers that must finance new chips to fulfill contracts. CoreWeave began borrowing against Nvidia GPUs with a $2.3 billion deal; its debt has since grown dramatically. An $8.5 billion loan illustrates what lenders actually care about: contracts with Meta worth at least $19 billion. The chips are collateral, but the customer's payments are what make the loan credible.

That's reassuring when the customer is Meta. The structure is more fragile when the customer also depends on continuous outside financing. The borrower needs the contract to remain valuable, the customer needs to keep paying, and the hardware loses economic value as newer chips become available. In the xAI ecosystem, a vehicle bought $5.4 billion of chips to lease to the lab, much of it financed with debt. Different entities own different pieces of the same bet on continued AI demand.

Buildings can last decades; chips lose economic value much faster. Often different companies own them, so the building's debt need not match the chips' lifespan. But each chip replacement requires another financing decision at whatever rates and prices prevail then.

The risk doesn't disappear when it changes hands

The point isn't that the transactions are fraudulent. It's that pensions and insurers can be exposed to AI through stocks, private credit and bonds that all depend on similar growth assumptions. Different instruments don't always mean different underlying risks.

Confidence can retreat before anyone defaults: lenders demand more yield, developers postpone borrowing, and the next deal simply fails to close.

How the money disappears

Imagine a developer considering a new AI facility in 2028. Customers still want the computing power. But borrowing is expensive, and campaign proposals raise uncertainty about grid costs, permits, and subsidies. A thinly profitable project may no longer clear its investors' hurdle. As the crypto example showed, politics has more force when financing is fragile already.

The facility could still be profitable. It may simply no longer be profitable enough. The developer doesn't default. It walks away from the project.

Multiply that decision across enough projects and the effects spread. Suppliers lose expected orders; growth forecasts fall; investors pay less for businesses that assumed permanent expansion. Venture firms, facing lower valuations and less confidence, tighten funding elsewhere in AI. The financial frenzy can end while the technology keeps improving.

This spread is the critical bet, not a certainty. Some projects will move to friendlier states; existing capacity may become more valuable. For a broad retreat, delayed construction must be large enough to undermine the growth already priced into suppliers, infrastructure companies and private AI valuations.

Three ways the politics could play out

Regulation won't affect everyone equally. Safety rules may favor incumbents. Sweeping bans are a more severe scenario, not my base case. My prediction rests on a narrower mix of higher costs, slower permitting and less generous incentives that makes enough new projects unattractive.

SCENARIO	MAIN EFFECT	MARKET IMPLICATION
Limited safety regulation	Compliance obligations	Often manageable for incumbents
Infrastructure cost backlash — my base case	Higher grid costs, permit delays, fewer incentives	Marginal projects deferred; financing and growth assumptions repriced
Broad populist intervention	Moratoriums, taxes or sweeping liabilities	Risk of a much wider reset

Existing facilities and well-capitalized operators may even benefit. The pressure is greatest where valuations assume an uninterrupted building boom.

Some of the reversal is already visible

Markets have started asking whether the industry's spending will earn enough. Meta's shares fell sharply after a strong earnings report accompanied by higher capital spending and negative free cash flow. The message wasn't that Meta had stopped growing. It was that growth alone no longer answered the cost question.

At the local level, Heatmap Pro estimates at least $260 billion in planned data center investments have been canceled following opposition this year. That figure deserves careful interpretation—it describes announced projects, not necessarily sunk dollars or losses—but it shows how local decisions can disrupt national expansion plans. National politics could make those decisions harder to dismiss as isolated events.

Falling prices aren't defaults

A falling bond price is not a default. An investor can lose money on a bond while the borrower continues paying every coupon. This distinction matters because a repricing is my central forecast; a credit crisis is not. Defaults require something

worse: insufficient cash flow or a refinancing that cannot be completed. The pension funds and insurers holding this credit need not see a borrower default to suffer losses: a change in required yields can mark down the bonds and loans they own.

Actual distress is a separate downside: tenants can fail, equipment can lose collateral value, or refinancing can become impossible. None of that requires a large hyperscaler to fail, and none is necessary for my base-case repricing.

Why venture capital can react so quickly

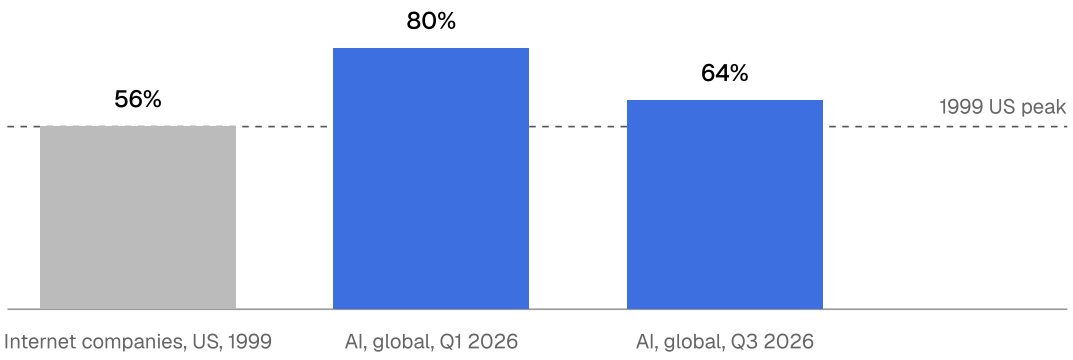
Venture capital can change direction without anyone changing their mind about AI's usefulness. An investor may still love the product and refuse yesterday's valuation. Funds preserve reserves; boards demand a longer runway; founders planning to raise in six months find the market has already left. Belief doesn't disappear. Its price changes.

We've seen this behavior before. Venture investment contracted sharply after 2000; Sequoia's "R.I.P. Good Times" memo became a shorthand for the cuts that followed the 2008 shock; global startup funding fell substantially in 2022–23. These episodes differed in cause, but each showed how quickly risk tolerance can change across otherwise unrelated companies.

AI is unusually exposed because it has captured a remarkable share of venture funding. Crunchbase puts AI at 64% of global venture investment in the third quarter of 2026. That's not directly comparable to the US internet-company share in 1999, but the dependence is unmistakable: if the dominant category revalues, the whole venture ecosystem feels it.

AI's share of venture capital rivals the internet's at its peak

Share of venture dollars. 1999 is US only (PwC; the NVCA counted about two-thirds). 2026 is global (Crunchbase). An indicative comparison, not a like-for-like one.



1999: US VENTURE DOLLARS (PWC MONEYTREE) · 2026: GLOBAL VENTURE DOLLARS (CRUNCHBASE) · INDICATIVE, NOT LIKE-FOR-LIKE

The chain reaction

The likely sequence isn't one cinematic crash. It's a change in the economics of saying yes:

1. High valuations, debt, interest rates and immense spending leave little room for disappointment.
2. Investors begin accounting for power charges, permitting and subsidies that may not survive the election.
3. Projects are delayed; the expected growth of their suppliers is repriced.
4. Lenders demand more protection, investors pay lower prices and venture firms preserve cash.
5. The technology continues spreading, but not every company or project financed at peak assumptions survives.

Nothing in that chain requires a major default. A more severe branch opens if weak tenants begin missing payments and losses pass through credit vehicles. That's a possibility worth watching. It isn't the event on which this essay's forecast

depends.

What could prove me wrong

A forecast worth publishing should have a way to fail. I'm not predicting a particular crash date or a universal collapse. By the end of 2028, I expect at least two of the following four things to have happened:

- The big hyperscalers' combined capex falls year over year in actual dollars, not just grows more slowly, or most of them cut their forward guidance in the same quarter.
- AI's share of global venture funding drops below 40%, from 64% in the third quarter of 2026.
- Most of the AI companies that go public before mid-2028 at valuations above $10 billion trade below their listing prices six months later. If none list, this signal counts as unmet.
- Down rounds spread: at least three of the ten most valuable private AI companies as of October 2026 (listed in the appendix) raise money at lower valuations than their last round within a twelve-month span.

If fewer than two happen, the market forecast was wrong. And I won't wait until 2028 to look for contrary evidence:

- AI valuations stay near today's levels through the 2028 campaign, even as regulatory risk becomes concrete.
- The big AI IPOs trade well after listing, and venture funding for AI stays near today's levels.
- New data centers keep getting financed even as states make them pay their full electricity costs.
- AI revenue and operating cash flow grow faster than the debt and lease obligations tied to them.
- The 2028 campaigns don't produce a serious bipartisan push on data centers, power bills or AI jobs.
- Spreads on data center and chip-backed debt stay stable as political scrutiny rises. On its own, this wouldn't disprove a valuation correction, but it would mean the downside branch isn't coming.

The point of setting these tests now is to resist the temptation to explain away every result afterward. If the numbers hold, I was early or wrong. If they break, the financial thesis was worth preparing for. But another question remains: what caused the break?

The four tests measure the outcome, not its cause. Higher rates or disappointing revenue could trigger them without politics playing a part. That would support the market call, not the electoral argument. For the latter, I'll look at reactions around specific political events: permit freezes, cost-allocation bills, election results. If the most exposed developers and lenders consistently react more sharply than less-exposed AI businesses, politics becomes a more plausible accelerator. It still won't be proof. If that pattern never emerges, I may have called the money right and the reason wrong.

A closer, earlier check on the mechanism is the market for new infrastructure itself: whether financing spreads widen, equity requirements rise, or projects fail to close specifically after consequential cost-allocation and permitting decisions. Those signals wouldn't replace the four outcome tests. They would help distinguish a politically driven financing retreat from a general selloff caused by rates or weak AI revenue.

There are also objections I take seriously.

This is the strongest one. But the 2028 coalition is aiming at electric bills, permits, jobs and who gets rich, and leaving frontier research alone. A candidate can promise to beat China and still promise to make data centers pay for their own power. Josh Hawley, one of the loudest anti-AI voices in the Senate, is also one of its most aggressive China hawks.

True. Meta, Google, Microsoft and Amazon can absorb losses that sank a 2001 telecom. That's why so much of the risk has moved off their balance sheets and into special purpose vehicles, private credit and insurance portfolios. The giants are protected. The chain underneath them isn't.

It might. This is the best case against my prediction, and I take it seriously. If AI revenue grows fast enough, the data centers fill up, the contracts cover the debt, and today's valuations look reasonable in hindsight. But that requires more than AI becoming useful. It requires the profits earned from AI to justify the extraordinary amount of capital

committed to building it, on the schedule today's investors are paying for. The technology can succeed without those profits arriving fast enough. Meta just posted negative free cash flow while beating its revenue forecast. That gap, between a technology that works and profits that arrive in time, is what I'm betting on.

The railroads were the economy in 1906. Congress regulated their prices anyway, because voters were angry enough. Size only makes an industry a bigger target.

Fair. Sequoia's 2008 warnings were partly wrong too: the downturn was real, but tech recovered faster than they predicted. So I'm naming a window for preparation, not a date for a crash. I'd rather be early and prepared than right and broke.

What to do with two years

Predictions are cheap until someone has to act on them. If I'm right, the next two years are less a countdown to catastrophe than a chance to buy optionality. Your best move depends on where you sit.

Some precautions are sound even if I'm completely wrong: shorten the distance to profitability, examine debt covenants, avoid one-customer dependence, and match financing to asset life. Advice that depends on my forecast should be weighed against your own circumstances, not taken as a universal rule.

If you're a founder

1. Don't wait for the perfect metrics. Every past freeze closed the funding window faster than founders expected, and the companies that survived 2001 were the ones that raised in 1999 and hadn't spent it. Size the round for 36 months of runway at your current burn, not 18. If you're profitable or already have three years of runway, you don't need to raise just because the window is open. Dilution is a cost too.
2. If you're choosing between a bigger valuation with heavy terms (participating preferences, multiples above 1x, ratchets) and a lower valuation with clean ones, take the clean round. Those terms decide who gets paid when the next round is flat or down.
3. Put on paper exactly what you'd cut, and when, if you can't raise again for two years. Know the month you'd pull the trigger. Founders who draft this in a downturn draft it too late.
4. Count only revenue from signed, paid annual contracts. Pilots and proofs of concept are the first budgets customers cut when their own investors get nervous. Make converting them your top commercial priority.
5. If any one customer is more than a quarter of your revenue, start fixing it. Concentration is a risk even the biggest AI labs are flagging to their own investors.
6. Avoid long take-or-pay contracts sized for your best-case growth. Model what happens to your margins if compute and power costs rise 25%, and price your product so you can survive it.
7. If you have venture debt with covenants tied to revenue or growth milestones, renegotiate now, while lenders are still friendly. Covenants written for a boom get enforced in a bust.

If you're a VC

1. Sort every company into three groups: profitable or default alive, one more round from safety, and dependent on today's prices to survive. Tell the third group now, while they can still act.
2. A flat round in 2027 beats no round in 2028. Make it the default advice at every board meeting.
3. In every past freeze, the best returns came from backing your winners when nobody else would. Hold more capital back for follow-ons than you think you need.
4. Discount pilot revenue, revenue from a single customer, and revenue that depends on the customer's own fundraising. Ask what the business looks like if its customers' budgets are cut by a third.
5. If most of your fund is in AI, and much of that is in companies that need constant new capital to buy compute, you're carrying the same bet many times over. Balance new deployments accordingly.
6. Don't wait to write your own "R.I.P. Good Times." Send a measured version this year: plan for 36 months, cut what you can't defend, raise while the window is open. Early warnings are cheaper than emergency ones.
7. Tell your LPs where your marks depend on frenzy-era prices. Markdowns that come early cost you less trust than markdowns that come all at once.

If you're an LP or allocator

1. The same bet can sit in your venture funds, your private credit funds, your public equities and your insurance portfolios at once. Count it in one place.
2. Ask how much they hold in data center and chip-backed debt, how concentrated it is by tenant, how long the leases run compared to the loans, and who guarantees what if a tenant leaves.
3. Don't concentrate commitments in funds deploying at peak prices. Keep capital available for the funds that will be buying in 2028 and 2029, when prices are lower.

If you build or finance data centers

1. Write leases so that changes in electricity rates and grid charges pass through to the tenant. The policy voters want most is data centers paying their own power costs. Don't let that land on your balance sheet.
2. If one customer's contract is the only thing repaying the loan, get a guarantee, diversify the tenants, or don't build.
3. About half the projects that met sustained local opposition this year were canceled. Treat community resistance as a permitting risk you underwrite, not a PR problem you manage later.
4. Finance buildings with long-term debt and chips with loans that are paid off well within the chips' useful life. Never finance a five-year asset with a fifteen-year loan.
5. Fund grid upgrades and community benefits up front. The industry's best defense against a cost-allocation law is making it unnecessary.

If you work in AI

1. Paper wealth at frenzy prices isn't money. Don't make life decisions that assume today's valuation. If you share my forecast, sell what you can when you can.
2. The teams closest to revenue are the last ones cut in a downturn.

The calendar to watch

Don't mistake a headline for a signal. These are the checkpoints I'd put on a calendar:

-
- the hyperscalers' capex guidance in their fourth-quarter earnings, and the first moves by new governors on power costs and permits.
- how the big AI IPOs trade six months after listing.
- the primary platforms on data centers, power bills and AI.

If two of my four end-of-frenzy signals show up early and persist, shift from preparation to defense: cut burn to reach default alive, secure the financing you can get, and stop committing capital at valuations that depend on the frenzy continuing. For the signals that move every quarter, capex guidance and AI's share of venture funding, "persist" means two quarters in a row, not one bad print. The other two count once they're broad, not the moment one company stumbles: most of the $10 billion-plus AI IPOs below their listing prices at six months, or down rounds at three of the ten leading companies within a year. One company's bad news is a company problem. A pattern across companies is a market problem.

The turning point

The first draft of AI's rules asked how companies would control their technology. The next will ask who pays for the power, who absorbs the disruption and who keeps the profits. Those questions sound political. They are also financing terms. The industry is building as though it will write the next draft too. The 2028 election may take the pen away.

We have two years. Let's not spend them pretending we'll be the ones holding the pen.

Appendix: the reference list

The down-round test in “What could prove me wrong” uses a fixed group of companies, recorded here so nobody has to argue in 2028 about who counts. These are the ten most valuable private AI companies as of October 28, 2026, ranked by their latest observable valuation, whether from a funding round or a secondary sale such as an employee tender. Three rules decide who’s included: the company is independent and private (no subsidiaries, and none already acquired or pending acquisition); its core business is AI models, AI software or AI infrastructure; and it’s headquartered in the US or Europe.

Two rules govern the test itself. First, a down round means a new primary financing round priced below the company’s latest primary round, not below a tender or secondary price. Second, this cohort is frozen. If a company goes public or is acquired, it stops being eligible for the down-round test, and IPOs get judged by the IPO test instead, but nobody is added to replace it, and the threshold stays at three. That makes the test harder to meet over time, but it means the same thing in 2028 as it does today.

#	COMPANY	LAST VALUATION	TRANSACTION	DATE	SOURCE
1	Anthropic	$965B	Series H	May 2026	Anthropic
2	OpenAI	$852B	$122B primary round	Mar 2026	Bloomberg
3	Databricks	$190B	$5B strategic round	Aug 2026	Yahoo Finance
4	Cognition	$48B	$2B round	Sep 2026	TechCrunch
5	Safe Superintelligence	$32B	Series B	Apr 2025	Reuters
6	Crusoe	$30.9B	Series F	Sep 2026	Business Insider
7	Scale AI	$29B	Meta investment (49% stake)	Jun 2025	Reuters
8	Mistral AI	$24.4B	Series D	Sep 2026	TechCrunch
9	ElevenLabs	$22B	Employee tender	Sep 2026	ElevenLabs
10	Etched	$21B	$700M round	Aug 2026	TechCrunch

Compiled from the Prompt20 AI valuation leaderboard (valuations as of October 11, 2026) and the original funding reports linked above. Two entries are stale: Safe Superintelligence and Scale AI last raised at a published price in 2025. Nvidia invested in Safe Superintelligence in July 2026, but no valuation was disclosed.