

Impact of rising financing costs on the AI CapEx cycle

State Street · Ninghui Liu;Aditya Patel;Krishna Bhimavarapu · 2026-08-26 · English original

The AI story is shifting. As capital spending accelerates, investors are asking a new question: who funds the buildout, and at what cost?

The question the market has stopped asking

For much of the past year, the AI debate has been about who captures the largest share of AI-related capital spending. However, now the focus is shifting to more nuanced and uncomfortable questions: who funds the buildout, and at what cost? What happens if end-user demand validation and productivity payoff arrives slower than the financing clock allows?

Credit markets are already signaling concern about AI-related capital expenditures (CapEx). Since December 2025, the spreads on the JP Morgan Liquid Hyperscaler basket have widened significantly, while the broader investment-grade (IG) spreads have stayed relatively tight. That is not a default warning for the world's strongest balance sheets—but it may be a repricing of the capital intensity, uncertainty around return on invested capital (ROIC) on the incremental AI CapEx, and duration risk now embedded in the AI trade.¹

Figure 1: Bond market starts to question AI buildout

Our argument is deliberately not a bubble call. We remain constructive on equities, particularly US equities. The point is narrower, and we think, more useful: the AI cycle is evolving from one driven primarily by growth and earnings to one increasingly shaped by balance sheet considerations. The goal is not to predict market direction, but to provide a framework for understanding the transition of the AI cycle from an earnings-driven story to one increasingly shaped by capital availability, financing costs, and macroeconomic risks.

A useful lens: Installation before deployment

Carlota Perez's work on technological revolutions offers a helpful frame. Before settling into a longer deployment phase of broad diffusion and productivity gains, major technologies tend to move through a speculative installation phase. This phase is often characterized as financial capital chasing a new opportunity, aggressive overbuilding, and valuations running ahead of realized economics. The transition between the two often occurs when returns begin to disappoint many of the investors who funded the installation phase.²

The paradox is that the frenzy is socially productive, even when it is financially painful: society may inherit useful infrastructure, lower access costs, technical know-how, and new business capabilities, even if the investors who fund the buildout earn poor returns.

The overbuilding of canals, railways, and fiber left behind cheap, ubiquitous infrastructure that society later used to great effect. The technology usually delivers on its promise in the longer term; investor outcomes during the buildout are far more mixed. In other words, a shakeout after an installation boom is closer to a feature than a bug of most technological revolutions.

Today's simultaneous, unprecedented buildout of data centers, semiconductors, networking, cooling, and power looks like an installation phase—which is precisely why the financing question deserves just as much attention as the demand question.

Why AI financing costs matter now

The old narrative was clean—Microsoft, Google, Amazon and Meta spend on CapEx, then NVIDIA, Broadcom, Taiwan Semiconductor Manufacturing Company, and Micron collect that spending, and investors buy the supply chain.

Markets have long valued hyperscalers as capital-light cash generators with high ROIC and growth profiles, supported by durable moats. Until recently, investors believed these platforms could always deploy as much cash as necessary to invest to win in the AI-era. However, now the proposition has changed to how much do these companies need to spend to avoid falling behind the pack; a 'red queen race' of sorts: every company has their own reason to keep investing.³ The likely outcome is rising hardware costs, rising financing costs, rising power costs, and declining prices for AI services.

Those selling chips upstream have been thriving, while those buying them downstream are faced with uncertain economics on these large investments. Not only does it create unprecedented financing needs for these companies, but also new execution risk which increases the uncertainty around ROI on these spends. Bonds are increasingly being issued for ten or twenty years; graphics processing units (GPUs) can be obsolete in three to five. Cloud providers are progressively using more expensive capital to buy progressively shorter-lived equipment.⁴

The macro-financial evidence also corroborates this shift. In the most exposed economies, the Bank of International Settlements (BIS) estimates AI-related investment is approaching 1% of GDP—and it's increasingly debt-financed. According to Dealogic data, global tech companies issued \$428.3 billion of bonds in 2025, up from about \$150 billion in 2023, while private credit lending to AI-related firms grew to about \$200 billion in 2025, with potentially incremental scope.⁵

A material and growing share of the buildout is being funded off the balance sheet entirely—through leases, offtake agreements, and special-purpose vehicles that the BIS labels "shadow borrowing." Public balance sheets no longer show the full risk map.⁶

Not one AI trade, but six

Treating AI as a single exposure also obscures where the risk sits. It is more useful to think in layers, each with a distinct business model, moat, and vulnerability. The framework also makes the circularity and contagion channels far easier to see.

Figure 3: The AI value chain: Competitive strengths and key vulnerabilities by layer

NVIDIA, Broadcom, TSMC, Micron, power/networking

Scarcity pricing, high current demand, strong margins/FCF

Cyclicality, future overcapacity, new-entrant threat

Bridge capacity the ecosystem needs today. Critical to the open-source ecosystem

Most acute financial risk: leverage, under-utilization, GPU obsolescence

Distribution moats, existing cash flows, high switching costs

Rising CapEx intensity, uncertain incremental ROIC

L4 – Large language model (LLM) developers

Strategic value, brand, capability frontier

Monetization uncertainty, open-source pressure, commoditization

Salesforce, SAP, ServiceNow, vertical software-as-a-service

Systems of record, trust, workflow control

Must prove AI pricing power / outcome-based pricing

Source: State Street Investment Management, as of August 5, 2026.

Layer 1 is genuinely flush: NVIDIA reported FY2026 revenue of \$215.9 billion, data center revenue of \$193.5 billion, and free cash flow of \$96.7 billion with property, equipment and intangible purchases of just \$6.0 billion.⁷

Layer 3 platforms are durable franchises: Microsoft alone reported FY2026 revenue of \$331.8 billion and operating income of \$155.2 billion, alongside commercial remaining performance obligations of \$678 billion and additions to property and equipment, including finance-lease-related additions, of \$116 billion.⁸ Falling free cash flow or lighter buybacks aren't the issue for these names, but uncertainty around incremental ROIC is.

The genuine financing and unit-economics stress is concentrated in Layers 2 and 4. Layer 2 is squeezed from both sides: uncertain demand and GPU pricing on the revenue line, obsolescence risk, and rising interest costs on the cost line. Debt-financed growth is always like riding a bicycle fast downhill: exhilarating while balance holds, but a small bump in utilization, refinancing cost, or decline in collateral value can produce an outsized crash.

Layer 4 faces an unresolved path to profitable monetization at scale: open-source competition raises the risk that “intelligence” itself proves hard to monetize. Crucially, most LLM developers are still private companies, so their financial strain doesn't show up in daily market prices. Yet, a large part of the revenue backlog booked by Layers 1, 2, and 3 is ultimately contingent on these frontier labs continuing to raise capital through initial public offerings or private rounds. That is a meaningful dependency at a time when macro volatility is elevated, unit economics are unresolved, and end-user demand for proprietary models is uncertain.

Circularity vs. contagion: Read in a balanced light

A rising share of AI demand is also ecosystem-financed rather than end-user-validated. That doesn't by itself prove a bubble—circularity is a normal feature of installation phases. It is, however, worth monitoring: investors should keep asking how much demand reflects genuine end-user willingness to pay versus capital being recycled within the ecosystem.

When Layer 1 and Layer 3 surpluses indirectly finance Layer 2 capacity and Layer 4 model development—while Layer 4 is simultaneously among the largest sources of demand for that infrastructure—some of the “strong AI demand” being reported is partly self-referential until enterprise ROI is fully validated.

The contagion channel follows from the same map. If broad diffusion ultimately requires lower inference prices, cheaper tokens, or bundled AI features, the margin pressure falls disproportionately on Layers 2 and 4.

Layers 5 and 6 may benefit from cheaper AI, but compute operators and/or model developers can face compressing margins and rising solvency risk, and even Layers 1 and 3 could see lower ROIC and a valuation de-rating. A chain is only as strong as its weakest link—and genuine refinancing risk is concentrated in Layer 2 (and, to a lesser and largely equity-funded extent, Layer 4). For the ecosystem, the key risk exposure is not refinancing but contagion.

The cyclical aspect: Merited by changing fundamentals

Wider hyperscaler spreads are, in large part, warranted: a deteriorating technical backdrop of heavier bond supply is colliding with a reappraisal of a more capital-intensive growth profile and greater ROIC uncertainty.

Order-book coverage for hyperscaler issuance reportedly fell from nearly five times in February to under two times by July,⁹ even as 2026 global AI-related issuance heads toward the high hundreds of billions based on Morgan Stanley estimates.¹⁰ That flow lifts credit default swaps (CDS) premia without equities collapsing, and helps explain why CoreWeave and Oracle 5Y CDS contracts trade above ~700bps and ~200bps respectively as of July 23rd while the mega-caps sit near 50–85bps.¹¹

One nuance deserves emphasis: some Layer 2 CDS may be trading wider than fundamentals alone would justify because equity holders of these highly levered names might be buying CDS protection—cheaper than puts—while retaining their equity exposure.

The macro backdrop: less supportive now, not permanently hostile

With bond and private-credit based borrowings increasingly financing the AI investment cycle, the macro backdrop matters more than markets appreciate. Higher interest rates can raise uncomfortable questions around project returns, balance sheet capacity, and the sustainability of ever-rising capital expenditure. Our base case remains that US monetary policy ultimately stays accommodative and we do not expect another meaningful tightening cycle from the Fed in the near-term.

However, markets remain reluctant to fully price rate cuts given above-target inflation, elevated energy prices, lingering geopolitical uncertainty in the Middle East, and a labor market that remains resilient at the margin. As a result, policy rates may not move materially higher, but they may also remain restrictive for longer than many growth-oriented sectors would prefer.

The more interesting headwind may not come from Washington at all. For decades, Japan has been one of the world's largest exporters of savings, helping anchor global borrowing costs through ultra-low interest rates, substantial overseas asset holdings, and persistent demand for foreign fixed income. That backdrop is gradually evolving, as the Bank of Japan (BoJ) is normalizing monetary policy as inflation becomes more durable, wage growth broadens, and Japan re-emerges as an increasingly investable destination for domestic and global capital alike.¹²

However, the persistence of a weak yen has, at times, prompted questions about whether policy normalization has lagged the underlying improvement in economic fundamentals. While the yen's weakness culminated in an unprecedented coordinated intervention with the US recently, markets are now increasingly focused on the prospect of further BoJ rate hikes. Our view is that the next hike could come as early as September, driven primarily by domestic fundamentals rather than external pressure.¹³ At the same time, maintaining a steady and independent normalization path will be important to limit the risk of unnecessary market volatility that could accompany a materially faster tightening cycle while preserving policy credibility.

The implications extend well beyond Japan. Although the financing of the AI buildout is overwhelmingly

concentrated in US capital markets, those markets ultimately draw on a broad global investor base that includes patient, long-term capital from Japanese banks, insurers, pension funds, and asset managers. As Japanese yields become increasingly competitive and domestic investment opportunities improve, the incentive to allocate marginal capital abroad may gradually diminish. This does not necessarily mean a large-scale repatriation of assets, but it does suggest that the exceptionally supportive global capital conditions of the past decade may become somewhat less abundant.

The key risk may not be Japan's normalization alone, but central banks losing the luxury of moving at their preferred pace. Rising bond yields, ongoing fiscal expansion, geopolitical uncertainties, and the evolving financing needs of AI itself have all made markets more sensitive to interest rate expectations. While we do not expect a renewed Fed tightening cycle and see a sustained normalization for the BoJ, bond and currency market volatility could still tighten financial conditions more than what the macro data warrant, which could further lift the cost of capital for AI. The irony is that AI may ultimately alleviate some of these constraints through timely delivery of productivity gains, but for now, it is also contributing to them through its rising capital intensity and financing. Resultingly, AI is becoming an important dynamic reshaping the macro-financial environment. While we do not view this as an imminent risk, it is a risk investors would be prudent to monitor.

Implications for investors

The AI cycle is moving from installation to the more demanding phase of deployment. That transition typically warrants greater investor discipline. Rather than treating AI as a single trade, investors should separate the value chain into distinct risk pools, each with its own economics, financing profile, and vulnerabilities.

For hyperscalers, the implications are especially important. Microsoft, Alphabet, Amazon, and Meta remain exceptional franchises, supported by deep customer relationships, powerful distribution, and formidable cash generation. Yet they are also funding a highly capital-intensive AI buildout at a time when the global macro backdrop is less supportive. Incremental capital is being committed to assets with uncertain useful lives, residual values, and payback periods. The AI debate has therefore moved beyond earnings momentum; it is increasingly a balance-sheet and capital-allocation story.

Lower buybacks are not the central issue. They are great when reinvestment opportunities are scarce; today, the concern is not the strategic logic behind AI CapEx, but the uncertainty of returns on these unusually large commitments. The cleaner implication is that company-level ROIC certainty has declined. Before AI, these firms could be valued as asset-light, consistent compounders. Now, incremental returns depend on a broader set of variables: adoption speed, technological diffusion, power availability, and costs, financing conditions, and obsolescence risk. All else equal, a wider distribution of outcomes could lead markets to assign lower valuation multiples, especially if US risk-free rates remain elevated.

The key offset is captive demand. These platforms are not building speculative capacity in search of customers; they sit at the center of cloud distribution, sticky enterprise workflows, lucrative advertising ecosystems, and consumer platforms with durable network effects. Many also have direct or indirect exposure to leading frontier labs. Even if some LLM providers struggle to monetize or finance themselves, integrated platforms can still earn attractive returns by embedding AI into existing products, defending core franchises, improving internal productivity, and deepening customer lock-in. That makes the "cyclical industrial" framing directionally useful, but premature as a central thesis—consistent with our constructive house view on US equities.

The ultimate test, then, is not how much capital is spent, but whether that capital can be converted into

durable returns after financing costs. AI CapEx should therefore be viewed neither as inherently bullish nor bearish, but as the bridge between promise and proof.

1 “Big Tech AI bond binge shatters ‘unspoken contract’ with investors” February 23, 2026, CNBC. 2 Perez, Carlota. Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages, 2002. 3 Perez, Carlota. Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages, 2002. 4 “The question everyone in AI asking: How long before a GPU depreciates?” November 14, 2025, CNBC. 5 “AI spending spree drives global tech debt issuance to record high” Reuters, December 22, 2025. 6 Bank of International Settlements, “AI and the global Economy: implications for central banks” July 28, 2026; “Financing the AI boom: from cash flows to debt” January 7, 2026; “Financing the AI infrastructure boom: on- and off-balance sheet borrowing” March 16, 2026; “AI disruption in private credit: exposure to software firms in BDCs” July 14, 2026; “Private credit’s software lending meets AI disruption” March 16, 2026. 7 FactSet Standardized Financials, as of August 11, 2026. 8 FactSet Standardized Financials, as of August 11, 2026. 9 “Bond Investors Push Back As AI Debt Heads Toward \$570 Billion” Forbes, July 17, 2026. 10 “Cover Ratios for Hyperscaler Bonds Declining | The Daily Spark” Forbes, July 17, 2026. 11 CNBC, “Big Tech AI bond binge shatters ‘unspoken contract’ with investors” February 23, 2026 and Forbes, “Bond Investors Push Back As AI Debt Heads Toward \$570 Billion” July 17, 2026. 12 State Street Investment Management, “Japan is back, and how!” May 19, 2026. 13 State Street Investment Management, “Labor weakness argues against Fed hikes” August 10, 2026.

For institutional / professional investors use only.

Investing involves risk including the risk of loss of principal.

The whole or any part of this work may not be reproduced, copied or transmitted or any of its contents disclosed to third parties without SSGA’s express written consent.

Source: https://www.ssga.com/nz/en_gb/institutional/insights/impact-of-rising-financing-costs-on-the-ai-capex-cycle