Data center financing

Does growth belie fragility?

Advait Arun

Advait Arun is Senior Associate, Capital Markets at the Center for Public Enterprise (CPE), where his research focuses on the barriers to mobilizing finance for infrastructure. Before CPE, he worked at the US Treasury on international and sustainable infrastructure investment. He is the author of "Bubble or Nothing: Data Center Project Finance," published in November 2025.

In 2025, Advait Arun of the Center for Public Enterprise argued data center financing was fragile at best, rife with asset-liability mismatches that should raise flags among institutional investors. Six months on, PREA CEO Zoe Hughes spoke with Arun to discuss if his views hold — where they have softened, where they have hardened, and where investors should be asking more questions.

▶ Aerial view of a data center being constructed among existing centers


iStock.com/Gerville

KEY TAKEAWAYS

GROWTH STILL BELIES FRAGILITY

Asset-liability mismatches still run through the AI and data center capital stack: cash flows that don't service debt, GPUs depreciating faster than their leases, and circular financing where growth collateralizes more growth.

THE HYPERSCALER CASH CANNON

Mega-deals like Meta-Blue Owl's $27 billion Hyperion bond have pushed back the maturity wall, giving the largest hyperscalers more flexibility. Investors price it closer to Treasuries than AI.

PRICE THE WHOLE STACK

Below the hyperscalers, investors should test tenant concentration, liquidity, and refinancing risk. Neoclouds like CoreWeave, borrowing at 14% to 17% with two customers driving 70% of revenue, demand a different lens, including GPU depreciation and refinancing cliffs.

GUIDE TO THE LINGO

Here are five definitions you need to know about AI:

- **Hyperscaler:** The handful of cloud platform giants, such as Microsoft, Google, Amazon, Meta, and Oracle, that own and operate data centers at globally significant scale and sit at the top of the AI capital structure.
- **Neocloud:** Pure play AI infrastructure providers, such as CoreWeave, Lambda, Crusoe, and Nebius, that rent GPU capacity to hyperscalers and AI developers. Not the same as traditional cloud.
- **GPU (graphics processing unit):** The specialized chip (overwhelmingly produced by Nvidia) that powers AI workloads. The unit of account in data center economics.
- **Compute** (noun)**:** Shorthand for GPU hours, as in buying compute or compute capacity.
- **Inference:** Running a trained AI model to generate output (i.e., the answer to your ChatGPT prompt). Distinct from training, the upfront, capex-heavy work of building the model in the first place. Inference is the recurring, revenue-generating workload.

For institutional investors, data centers are unavoidable. They sit at the intersection of real estate, infrastructure, and equity-market concentration risk, and the capital flowing into data centers is a chief motor of US economic growth.

In November 2025, Advait Arun of the Center for Public Enterprise published a paper, "Bubble or Nothing."[1] The paper maps the capital structure of the AI and data center sector from hyperscaler down to GPU producer, and traces the asset-liability mismatches and potential flags for institutional investors and investment managers to watch.

In the paper, Arun argued that the sector's explosive growth belies a deep fragility: cash flows that don't service liabilities, collateral whose value is eroding, and a circular financing structure that rests firmly on other companies' growth.

Six months on, Arun's view has evolved. The maturity wall he was most worried about has receded, thanks to a new model of bond sales financing data centers. But concentration risk and the industry's asset-liability mismatches continue.

It's worth starting the conversation with a reminder of the numbers behind today's AI buildout.

- The Magnificent Seven account for over half of the S&P 500 by market cap.
- The amount of debt tied to AI ballooned to $1.2 trillion as of October 2025, making it the largest segment in the investment grade market, surpassing US banks.
- According to the Federal Reserve, AI-related investment added almost 1% to GDP growth in the first three-quarters of 2025, representing 39% of all GDP growth and surpassing the level of the dot-com boom of the 2000s.[2]

For institutional investors, whose mandates run 20 to 30 years, the question is no longer whether to have a view on AI infrastructure. It's whether the view is correctly priced. Arun's paper is one of the few attempts to map every actor in that capital structure. The following conversation is what he's seeing now.

Walk us through how "Bubble or Nothing" came together and its core argument.

The paper comes from my own realization over last summer, as attention on AI continued to rise, that a lot of folks in the energy space and policymaker world did not have a good map of all the different actors in the system. People were conflating the hyperscalers with the data centers, with some of the debt providers. They're all contractually related in different ways, but nobody had drawn that map between which actors were doing what in the system.

Given that I work a lot on project finance and infrastructure finance, seeing how similar this sector was — entirely abstract from the technology — to real estate, to infrastructure, existing transactions we had already worked with, I thought it would be helpful to provide a clearer view of the capital structure. Who are the actors? What assets and liabilities are on their balance sheets? And importantly, how do they connect? (see 1)

These balance sheets are interlinked. If we can study those interlinkages, we can dig deeper into the sector and where it's going.

The phrase that runs through the paper is 'growth belies fragility.' Where does that fragility actually live?

It lives in the debt financing structures. It lives in revenue not matching debt, GPU-backed loans, the

1

Market structure of the AI sector

Each player in the data center and AI universe is contractually related in different ways. Having a clearer view of the capital structure can help with understanding the sector and where it is going.

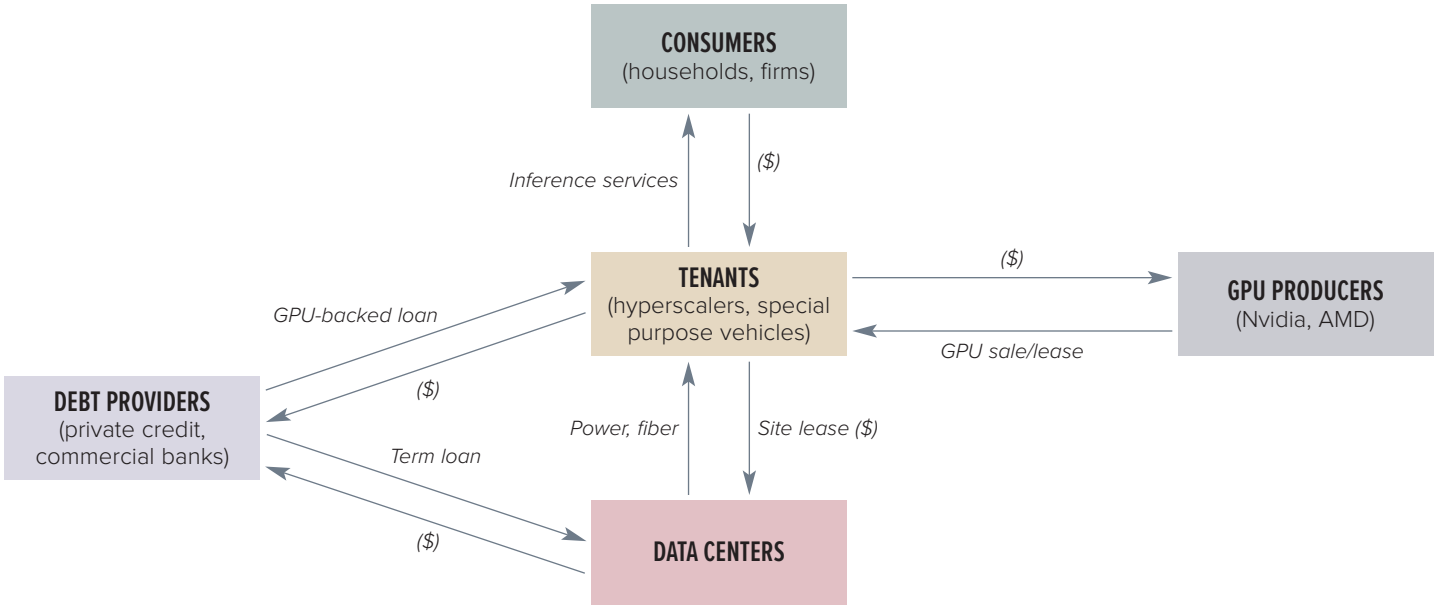

Source: Center for Public Enterprise.

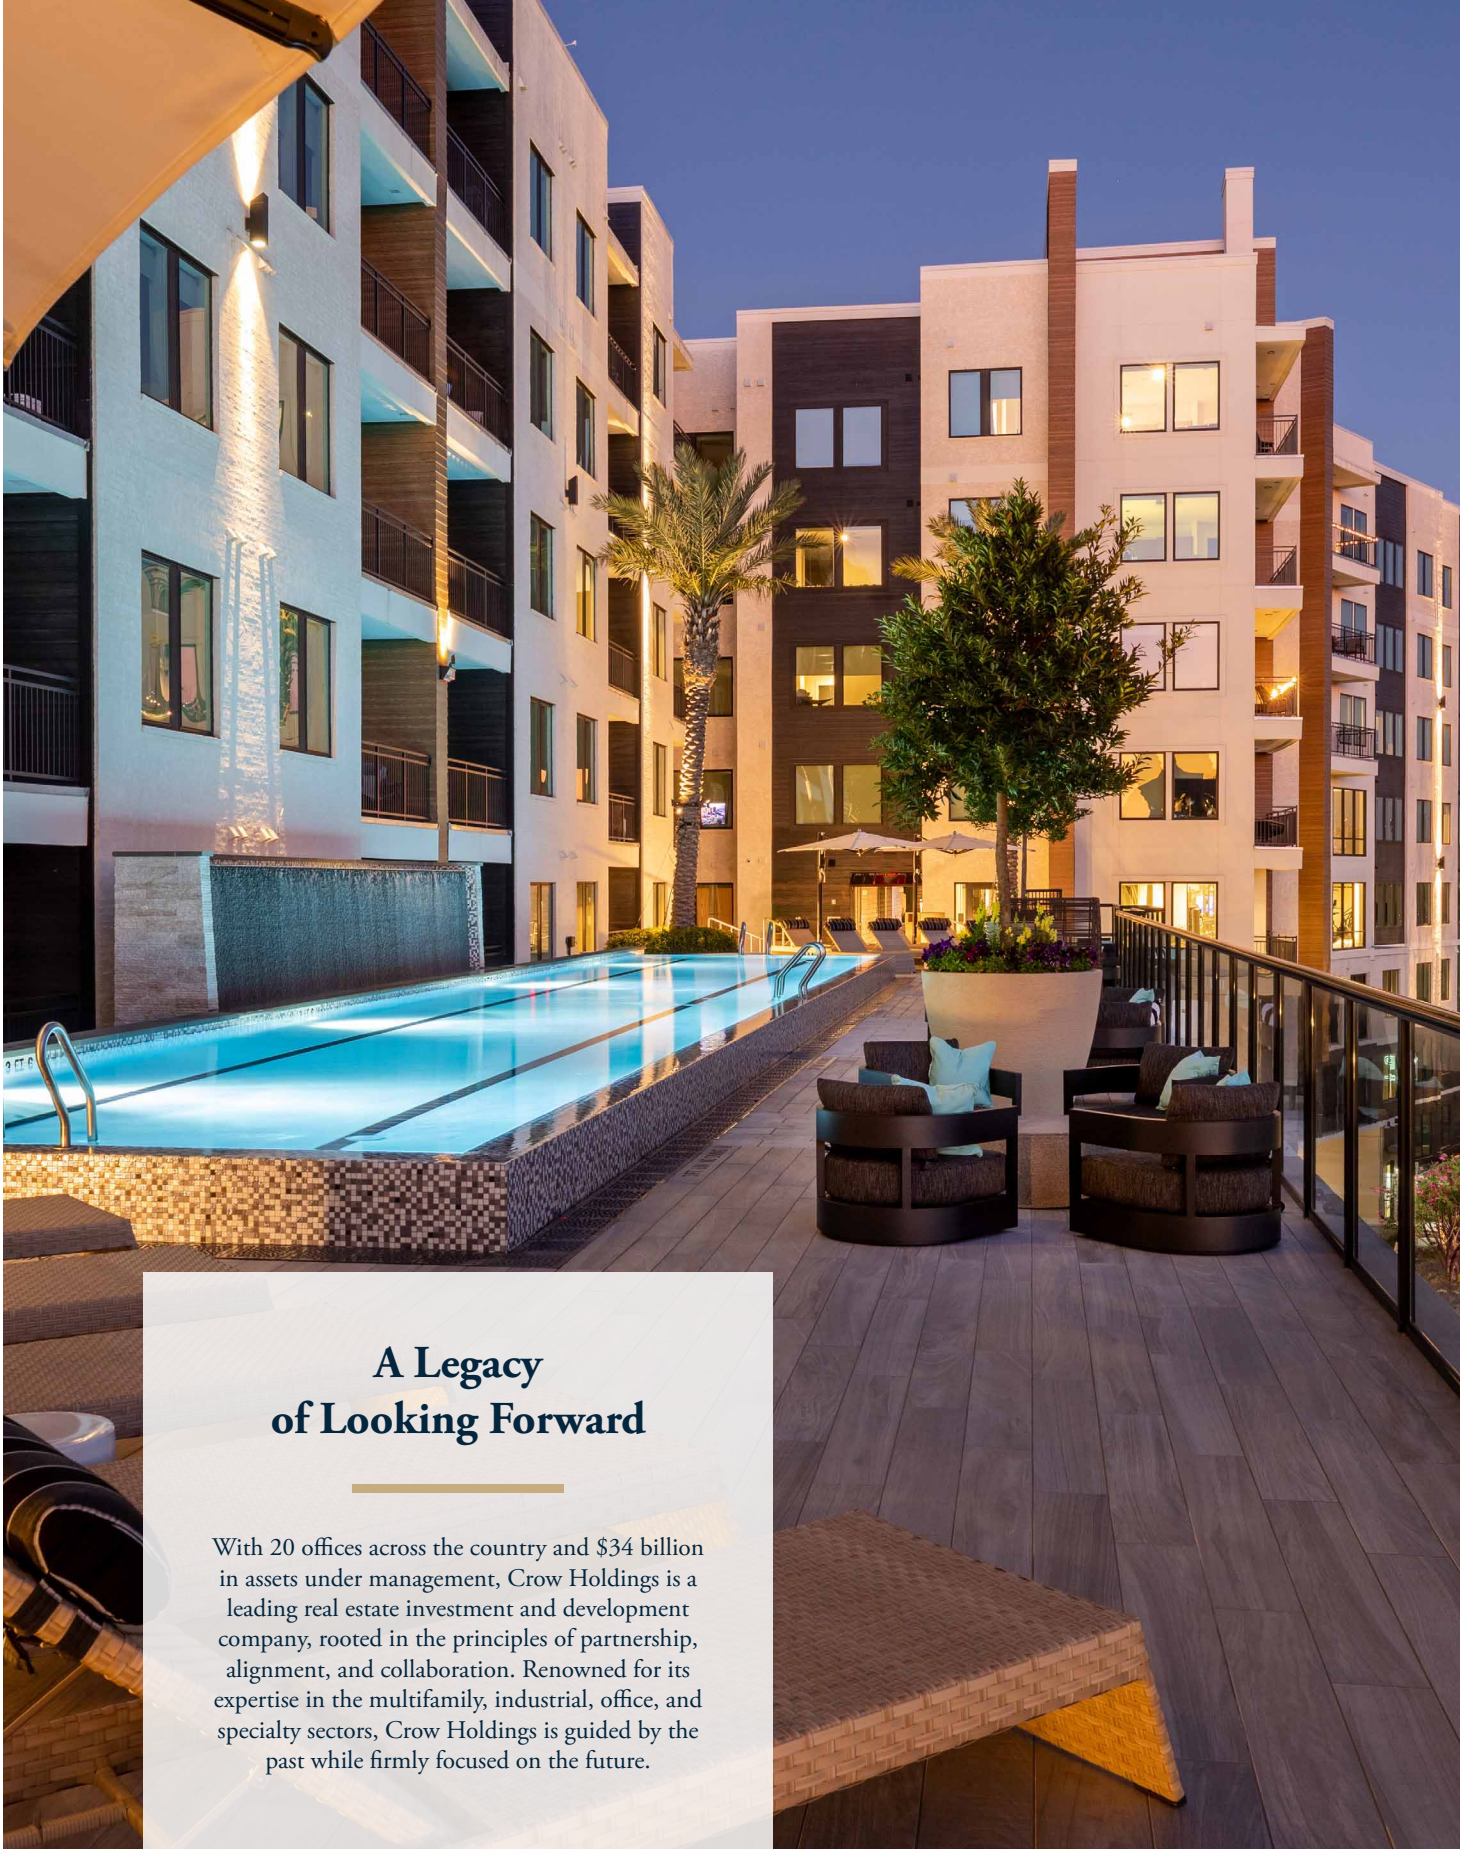

A Legacy of Looking Forward

With 20 offices across the country and $34 billion in assets under management, Crow Holdings is a leading real estate investment and development company, rooted in the principles of partnership, alignment, and collaboration. Renowned for its expertise in the multifamily, industrial, office, and specialty sectors, Crow Holdings is guided by the past while firmly focused on the future.

CROW HOLDINGS

CROWHOLDINGS.COM

A 30 second read of "Bubble or Nothing" by Advait Arun

The paper argues there are four themes that investors and policymakers should pay attention to:

1. Cash flow uncertainty. Today, AI firms are incentivized to earn market share by offering flat-rate pricing and unlimited use, locking in users for the future. Is this a race to the bottom?
2. GPU depreciation. The GPU is the keystone asset for hyperscalers and officially depreciates over a period of five to six years. But the pressure to stay competitive by deploying the latest, most powerful chips is constant. In 2024, Nvidia CEO Jensen Huang announced the firm would design new chip architecture every year[3] — further devaluing older GPUs.
3. Roundabouting. Hyperscalers are cross-investing, cross-selling, and cross-guaranteeing each other's expansion. The collateral justifying the risk-taking is, frequently, another company's future growth.
4. Funding future growth with debt. Hyperscalers are issuing record amounts of corporate debt. Below them, neoclouds and developers are raising debt on the back of hyperscaler tenant credit, at much higher rates and with much thinner cushions.

relationship between neoclouds and hyperscalers, who are competing for the same space but collaborating with each other for the same service. And it lives in what it means for the economy to be taking on so much more debt and spending so much more on capex for this one sector. The paper intends to provide a sense of where people should be paying attention.

A race to the bottom for market share?

Let's turn to the first of the four issues you flag in the 2025 paper. Cash flow uncertainty, or as you frame it a potential race to the bottom, a race for market share over margins. What's actually happening right now?

First and foremost, there's a psychological component we can't ignore. A lot of the people working on tech and software over the last 20 years have actually had a lot of success with this playbook, which is the realization that if I can be a monopolist of this service provision, if the network effects that I create can be the moat, I will have achieved the market share needed to have pricing power over the long term. That playbook means, eventually, *eventually*, I'll recoup the capital expenditure I had to take along the way. Examples of loss leaders are Uber and Amazon Web Services. But Uber is very capital light. And Amazon Web Services was built on top of revenue and infrastructure from Amazon.

The data center and AI sector is different in that the capex is real. And the hyperscalers and providers are all competing for something that users treat as a fundamentally identical product. When Google released Gemini 3, users flocked away from Claude and OpenAI to jump into Gemini 3. Then Spring 2026 came, and Claude Code took the world by storm. The users of AI services are really fickle. This idea that network effects create a moat hasn't shown up just yet.

Isn't this capture now, monetize later? Netflix priced cheap, killed cable, and locked in users. Once Claude is helping me run a business, I can't easily walk away either.

It's a really good challenge question, because I don't think we're out of the woods in terms of a world where that could happen. But there's a difference.

The hyperscalers building these inference services already have an existing tech service model they're bringing to bear. They're working in a new market but trying to integrate it into an existing older service that they're already making a lot of money from. Google has Google Workspace. Microsoft has Azure. Meta has its apps, advertisements, and social media.

If they all can't find an easy way to displace their competitor for inference services, they're going to fold that back into their overall suite of tech services.

That trend back towards enterprise revenues matters in the short term. It's easier to get that stable revenue stream out of enterprise services, even at a lower price point, from a market you already definitively have built a moat around. No one else is replacing Google Docs anytime soon. No one else is doing WhatsApp. That's where short-term stability will come from for hyperscalers and for investors underwriting data centers where these hyperscalers are the tenants.

The hyperscalers can lean on enterprise and the existing revenues generated from those services. What about the firms in the middle — the pure play AI service providers who aren't Google or Meta or Microsoft and don't have a WhatsApp to fall back on? How do we look at those firms, against that backdrop of unprecedented capex into data center build out?

It's really hard to track some of the revenues. Anthropic is posting some really good ARR (annualized run rate) figures, allegedly overtaking OpenAI. We still have to be skeptical, not just because it's a number multiplied by 12,

but because all the companies calculate ARR in slightly different ways, and Anthropic just released Claude Code where everyone just flocked to it. We're still a little bit unsure where the revenue trend will even go.

I'm not going to close the door on saying the revenues could circle back to justify this expenditure, where the market power of an individual model gets to that point. Maybe Claude is the model of the future. But I don't know if I'm going to rubber stamp that just yet.

When the accounting really matters to credit and value

I would argue the paper's most quietly explosive section is actually looking at accounting and depreciation because at the heart of a tenant's creditworthiness is its GPUs, and the value of those assets. What's been happening in terms of GPU depreciation that raised flags for you?

There are three things going on.

First, physical deterioration. We have no confirmed sense of the useful life of GPUs. It could be fast, around three years. A lot of companies say five years. But what we do know is that if demand really picks up, you have to replace the key capital asset rapidly.

Second, there's consistent downward pressure on resale value, because Nvidia and AMD are making new GPUs every year. The leading tech companies are going to get them first, and if you want to be on the cutting edge, you're going to need the new one. In reality, everyone is forcibly contributing to the deterioration of their own capital stock.

There's consistent downward pressure on resale value, because Nvidia and AMD are making new GPUs every year.

Third, the accounting. Hyperscalers are shifting their depreciation timelines to about five or six years. This helps to smooth earnings so as not to report sharp losses each year. But in the interaction of all of these factors, hyperscalers and neoclouds have taken out billions of dollars of GPU-backed loans to generate working capital. It's inventory finance. But if the value of the inventory is unknown or unstable, I see that as a collateral risk problem.

Do you believe there's a fundamental mismatch for data center landlords and owners in terms of their tenants, especially when a data center lease typically runs 10 years and GPUs depreciate, at best, every five to six years?

For me, refinancing deadlines are a big risk management flag. Data center leases will be 10 to 15 years, loans 15 to 20 to 30 years on the data center, and GPUs need to be replaced much quicker. Those various refinancing deadlines were really on top of my mind because that's what threatens the sector's fragile capital structure with the possibility of a Minsky moment.

AI financing: Round and round we go

You talk about roundabouting in your paper. That's the dot-com-era term for circular financing or the web of cross-investments, cross-purchases, and equity stakes which bind together companies in a sector. You define it as 'risk mutualism' and I was struck by this phrase: "where the collateral justifying the risk-taking is another company's future growth." Is this the flag that worries you most?

It's a flag. It's not a red one yet. Vendor financing is not abnormal in any real sense — suppliers need healthy off-takers, off-takers need suppliers. It is in everyone's best interest to ensure that the partners they're contracting with are creditworthy. But what's striking is not just the volume of roundabouting or the connections between the actors; it's how much capital expenditure is being pulled out of these cross-commitments.

Our team calls it a kind of risk mutualism. What worries me is that none of these commitments can be backed up with cash on hand on the balance sheets of these AI companies. A lot of the commitments are purchase commitments over the next five years. Some of them are repurchase commitments. A lot of it is equity warrants, which are claims on future growth that don't need to be exercised currently. A lot of market actors seem to be trusting that these purchase commitments — these options, these warrants — mean something in the capital structure. But if there's no cash on hand to pay them, they might not.

At the very top level, hyperscalers, as the ultimate buyers, will be fine given their liquidity position. But for

everyone under that layer, dependent on that buyer's market, that's a different analysis and that's where the impacts of roundabouting look more worrying.

And that web of cross-investments and commitments highlights your fourth concern on debt-fuelled growth.

Borrowing the future

Perhaps the best way to illustrate what you're talking about — and where all four flags of your paper meet together — is your example of CoreWeave, a neocloud or GPU-as-a-service provider. CoreWeave has raised more than $12 billion in loans, much of it priced at 10% to 14% and one tranche at 17%, which is triple the cost of investment-grade debt. That debt sits on an SPV, collateralized by CoreWeave's GPUs and service contracts. Two customers, Microsoft and Nvidia, generate over 70% of CoreWeave's annual revenue. And the company carries $56 billion in lease commitments running about a decade. When you stack all of that up, what's most exposed?

At worst, CoreWeave is a GPU warehouse for the hyperscalers. At best, it's the infrastructure for a free-wheeling wildcat market of AI products that anyone else can use.

What worries me is that the company is taking out more debt to roll over its old debt. It's not at the point where it's borrowing to pay interest, but it's running really close to that limit. We can assume that, below the hyperscalers, everything is a little more fragile than it seems.

There's a new liquidity document that came out after I wrote the report.[4] CoreWeave's lenders are allowing it to keep only $100 million of cash on hand at any one time, which is ridiculously low given its investments. If there's a collateral call on CoreWeave, does CoreWeave have the cash in the right account to pay that back? I don't think so. That's just one canary in the data center coal mine.

What's changed?

You wrote about Hyman Minsky in "Bubble or Nothing." Why?

Hyman Minsky was a financial economist who wrote about the financial sector's ability to make the macroeconomy more fragile by piling leverage onto ordinarily safe assets. The economy shifts from being composed of what he calls hedged units, which take a lot of equity and some debt financing and can always pay back their liabilities, into a world of speculative finance, and then Ponzi finance, where firms have to increase their leverage just to roll over their existing liabilities. So the whole economy becomes more vulnerable to instability.

The linkages between all of these actors, to fund each other's debts, is what creates this web of liabilities, where one shock, one collateral call, one margin call could send the structure tumbling. A Minsky moment is when that Ponzi structure starts collapsing like a house of cards. That's when people start saying, 'Is this a Minsky moment? Are we almost there?'

Has your view stayed the same or changed since releasing the report?

I saw one news piece that made me really think something is actually going on that's different here, that doesn't involve loan refinancing or rolling over debt.[5] There's at least one data center transaction in the pipeline where the lenders have dropped the idea of having one mini-permanent loan and one permanent loan after the mini-permanent loan. They've folded it all into one long-term loan over the life of the data center, which also includes construction with capitalized interest. The hyperscaler has committed to starting to pay its lease even if the power interconnection hasn't happened at the end of construction. If there are construction overruns, the hyperscaler will pay for that too.

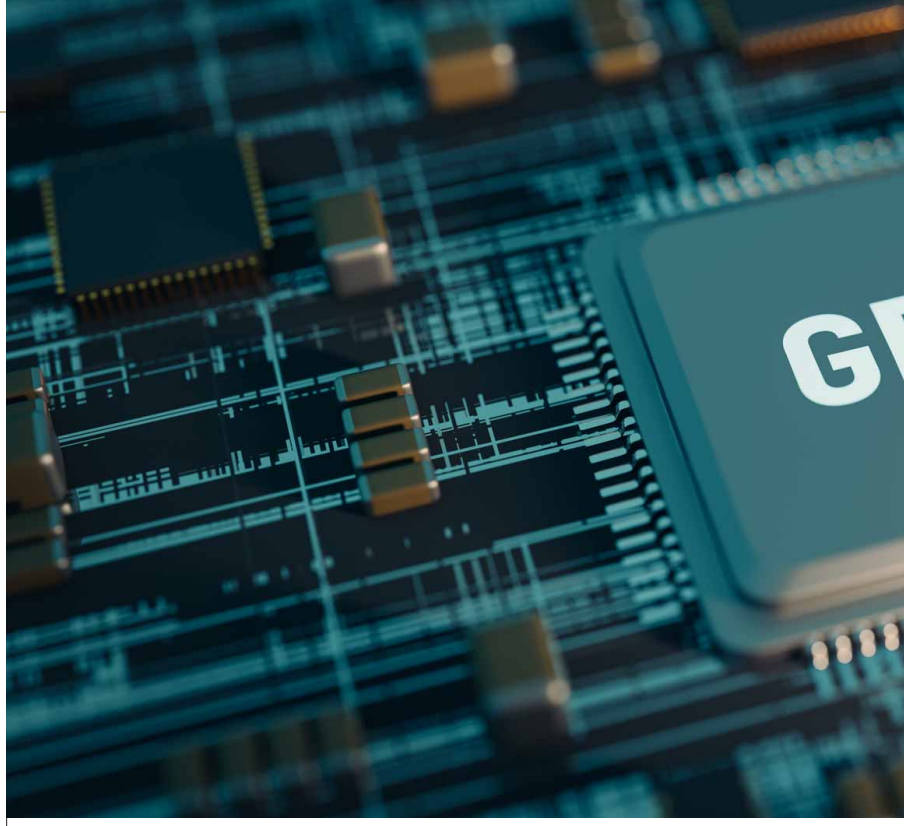

▲ At the heart of a tenant's creditworthiness is its GPUs, and the value of those assets

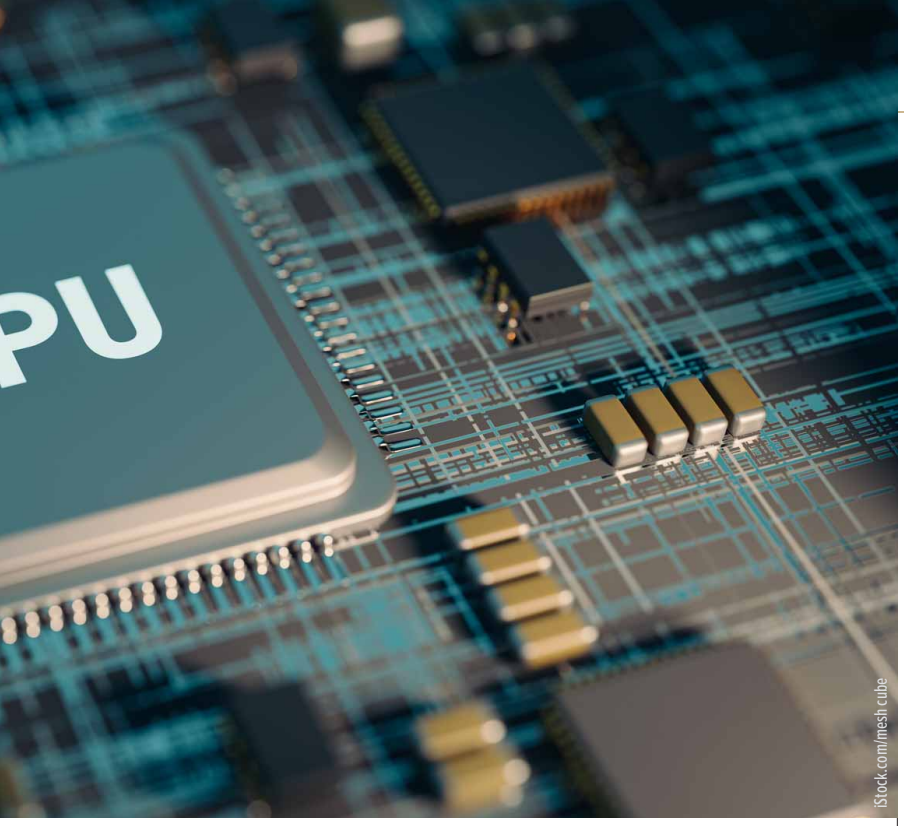

iStock.com/mesh cube

The deal actually makes a lot of sense. The hyperscaler is committing to getting this data center built, hell or high water, with its cash, and the lender doesn't need to worry about tenant churn or refinancing risk if all of a sudden the market moves against you in the middle of the loan. There isn't a maturity wall.

In October 2025, Meta and Blue Owl closed a $27 billion bond issuance to finance the Louisiana Hyperion data center they're developing.

Project finance is never in public markets. PIMCO was the largest investor in the Hyperion deal and they structured this project finance transaction with Blue Owl as a corporate bond. First, this structure is unprecedented[6] to the best of my knowledge. Second, this is the biggest corporate bond issuance in history. Third, by the time February [2026] came around, the bond was trading at $107, which implies a really low yield. Investors were treating this as so near Treasuries in terms of how secure it is.

Since then, we've seen the guarantee terms trickle out.[7] There's a residual value guarantee, where Meta is paying the lease up to a certain point, but also guaranteeing the broader asset value. It's also committing to deal with all of the construction overruns. Meta will also guarantee against tenant churn, promising not to vacate for a certain amount of time — or, if it does, it will sell at an equal price.

Investors love this deal. They're in essence getting direct exposure to Meta's balance sheet as corporate debt. Even though it's not direct exposure, investors seem to think so. Similar deals could definitely happen in the future.

Only the very largest hyperscalers can do this. Does that change the Minsky question and the flags you raised in your paper?

The hyperscaler cash cannon seems to be something that lets everyone blunt their understanding of what the sector's actual risks are. Having so much cash around — the fact that PIMCO's investment-grade bond is something that allows you to get direct exposure to the hyperscaler balance sheet and *not* to AI — sort of allows everyone to not worry about pricing the risk here fully. Maybe the Minsky moment is submerged or delayed. Examples of the financial instability hypothesis will still be playing out underneath the hyperscaler corporate level, like with CoreWeave.

But the nature of the market correction is changing. The hyperscalers have more runway, even as they're issuing debt. Combine that with the energy shortage and grid capacity — they'll get control over the market simply because no one else can play in it. That slows growth in the sector and concentrates the debt issuances. The market impact looks more concentrated. It's a risk, but it's less of a danger than a Minsky-like moment.

Is this 2008 all over again?

I don't believe what is happening with data center spending is similar to the global financial crisis. Our houses aren't on the line. The ABS and CMBS markets for data centers and AI seem fairly small and concentrated in a way that the CDO market for 2008 wasn't. CDS is not a relevant factor here either. But what's different about this means we have to create different metrics to take a look at this sector.

[1] Arun, A. 2025. "Bubble or Nothing: Data Center Project Finance." Center for Public Enterprise. November 12.
[2] Rubinton, H. and B.A. Patro. 2026. "Tracking AI's Contribution to GDP Growth." Federal Reserve Bank of St Louis. January 12.
[3] Bradshaw, Tim. 2024. "Nvidia Chief Vows Flagship AI Chips Each Year Despite Delays." *Financial Times*. August 29.
[4] Investing.com. 2026. "CoreWeave Amends Credit Agreement to Adjust Financial Covenants and Liquidity." January 2.
[5] Obando, S. 2026. "Moody's Sees $3T in Data Center Spending by 2030." *ConstructionDive*. January 20.
[6] Ji, C. 2025. "Meta, Blue Owl, and AI: Here are the Details of Wall Street's Biggest Private-Credit Deal Ever." DowJones MarketWatch. October 23; JPMorgan North America Credit Research. 2025. "Bridging the Gap Between Public and Private HG Markets." October 28.
[7] Rajgopal, S. 2025. "The Strange Case of Meta." *Forbes*. November 16.