



January 2026

Silicon, Steel & Stress

Part I: An Overview

www.unityinvestments.com

About Unity Investments

Our mission at Unity is simple yet profound: to create better access. Specifically, Unity identifies, catalyzes, and capitalizes on the most compelling alternative investment opportunities and shares them with our investors. In the long run, we aim to level the playing field.

Today, our investments business predominantly focuses on high-yielding private credit opportunities in the U.S. By building our origination network, underwriting team, and fund administration capabilities in-house, we control the entire investment process end-to-end to ensure quality and consistency. Selected for asymmetrical risk-reward and downside protection, our private credit deals are bespoke, privately negotiated, rigorously underwritten, and creatively structured. To learn more, please visit www.unityinvestments.com.





This two-part paper series asks a focused question: Can the credit-fueled infrastructure buildout generate sufficient cash flow to repay its debt? Part I sets the market and structural context; Part II quantifies repayment economics and stress scenarios.

Scale and Asset Mix

AI infrastructure capex is unprecedented: US\$2.9 trillion by 2028 and US\$5.2 trillion by 2030. Importantly, the spend is concentrated in compute hardware (servers and storage), which carries faster depreciation and higher obsolescence risk than long-lived real assets such as shells, land, and grid infrastructure. For credit, the key issue is that assets with very different lifespans are often financed together in one capital stack.

Sponsors and the Financing Gap

Large projects span hyperscalers, sovereign investors, and alternative-asset sponsors. While hyperscalers can fund a significant portion internally, the 2025–2028 cycle still implies an estimated US\$1.5 trillion of external financing gap. Within that gap, private credit is the dominant marginal supplier (~US\$800 billion), making its leverage, underwriting standards, and vehicle structures central to system risk. Because private credit often uses fund-level leverage, true exposure may be higher than headline figures, but remains difficult to measure given limited transparency.

Where Risk Actually Sits: Three Credit Archetypes

The paper's central analytic point is that credit risk is determined by structure and repayment engine, not by headline capex. AI infrastructure financing falls into three recurring archetypes: project finance, GPU platform credit, and hyperscaler balance sheet. Together, these archetypes illustrate how risk accumulates and migrates across a layered capital stack, as different pools of capital with varying durations, return targets, and loss tolerance finance the same underlying assets.

Policy Adds Momentum, Not Repayment

AI infrastructure is increasingly viewed as a strategic national asset. Policy support may reduce the risk of abandonment, but does not immunize private capital from loss and may encourage overbuilding by dampening market discipline, in some ways echoing lessons from the Dot-Com telecom cycle.

Ultimately, the sustainability of the AI capex supercycle depends on whether the cash flows supporting layered credit structures prove durable. Part II stress-tests these structures against obsolescence, concentration, and systemic leverage.

1. An Overview



1.1 Preface

While we believe AI will be transformative and significantly boost global productivity, the basic laws of economics remain unchanged. The danger lies in our natural tendency to over-extrapolate. We are here to separate the two and purely focus on analyzing economic fundamentals.

To ground the discussion, we map the buildout (scale, sponsors, and asset allocation) and then dissect how credit is being deployed across the capital stack, from construction loans and permanent takeouts to GPU platform credit and hyperscaler balance-sheet funding. We conclude by organizing these exposures into

three archetypes that clarify where risk truly resides.

Ultimately, the goal is to equip readers with a clear, fact-based understanding of AI infrastructure finance and to break down the key risks that concern lenders. The private credit market is inherently opaque; however, by meticulously combing through hundreds of pages of industry and company research, we hope to provide a detailed and grounded analysis of how credit flows into the sector today, based on observable precedents and identifiable financial dynamics, to assess what that implies for the sector's future stability.



2. How Big is the Market & Who are the Sponsors?

2.1 Size of the AI Infrastructure Buildout

2.1.1 US\$5.2 Trillion: Global AI Infrastructure Projection Through 2030

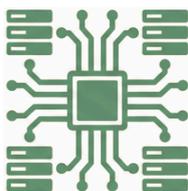
Our analysis begins with an observable fact: the investment quantum for AI is unprecedented. Specifically, we estimate US\$2.9 trillion of AI-related capex by 2028 and US\$5.2 trillion by 2030.^{*1,2} The breakdown of this expenditure, detailed in Figure 1, reveals the capital allocation and underlying assets that credit markets are being asked to finance by 2030.

Although the investment quantum is relatively clear, the path is not. The steep 34% CAGR implied between

our 2028 and 2030 anchor points represents a consensus buildout, assuming today's severe constraints are resolved. The true trajectory, however, can be dictated by forces such as (1) downside technological shock that comes with the shift to inference and breakthroughs in silicon efficiency, which could suppress per-unit hardware demand, and (2) upside cost shock resulting in persistent constraints and fractured supply chains that could instead inflate capital intensity per watt.

For credit analysis, the precise CAGR is less critical than this range of outcomes. Our anchor points establish the scale of capital committed under the current paradigm, allowing us to stress-test financing structures against the disruptions that could render them most vulnerable.

DESIGNERS



Computational Hardware:

Dominating the spend, an estimated US\$3.3 trillion (64%) is allocated to servers (primarily GPUs/CPUs) and storage. This is the core productive asset whose utilization and technological obsolescence are primary credit risks.

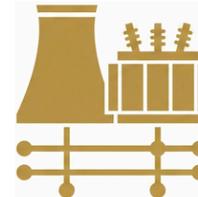
BUILDERS



Facility Construction:

Approximately US\$0.8 trillion (16%) is allocated to labor, shell, site development, and land, or assets with longer physical lifespans that are dependent on tenant demand.

ENERGIZERS



Power & Network Infrastructure:

Roughly US\$1.0 trillion (20%) funds electrical systems, power generation, and networking, representing critical enabling infrastructure with its own distinct value and risk profile.

Breakdown of 2028 & 2030 AI Capex Estimates

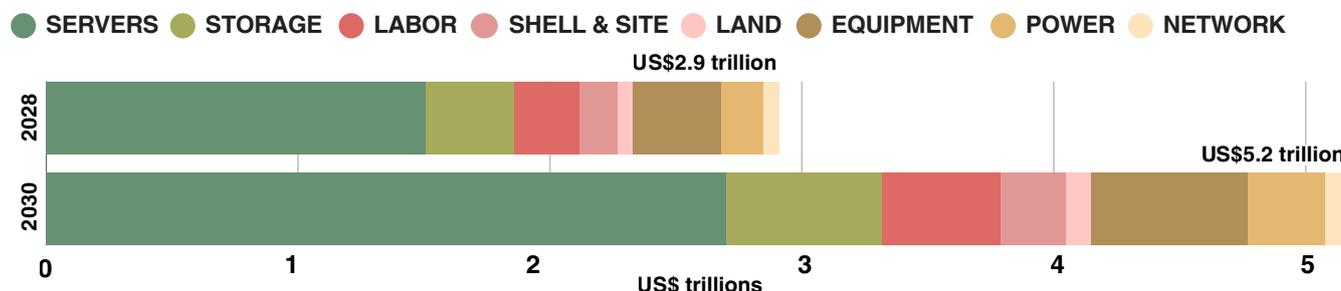


Figure 1.³ Segment allocations of US\$2.9 trillion and US\$5.2 trillion total AI capex estimate. See Appendix 2.1 for specific dollar amounts and percentage breakdowns.

* Projections for AI infrastructure investment vary significantly, ranging from ~US\$3 trillion (JLL, Morgan Stanley) to US\$5.2 trillion (McKinsey). This variance stems from differing methodological lenses, but reveals three convergent truths critical for credit analysis: unprecedented scale, dominant cost drivers by compute hardware and power infrastructure, and structural imperative—at this magnitude, sustainability depends less on the headline capex figure and more on the financing structures deployed to fund it. We use McKinsey's US\$5.2 trillion projection as a conservatively ambitious anchor that best captures the full capital intensity of the AI stack and illuminates the systemic credit risks under examination.



2.1.2 Major AI Infrastructure Announcements

Spanning North America, EMEA, and Asia, massive gigawatt-level data center undertakings sponsored by listed companies and sovereign wealth funds alike are sprouting up globally.

Though not exhaustive, these projects alone account for a meaningful US\$2.5 trillion of infrastructure spend.

While these announcements underscore the sheer size of the investment, they also raise a pivotal question: who will finance this buildout, and what are the implications for credit markets?

Major Announcements & Estimates

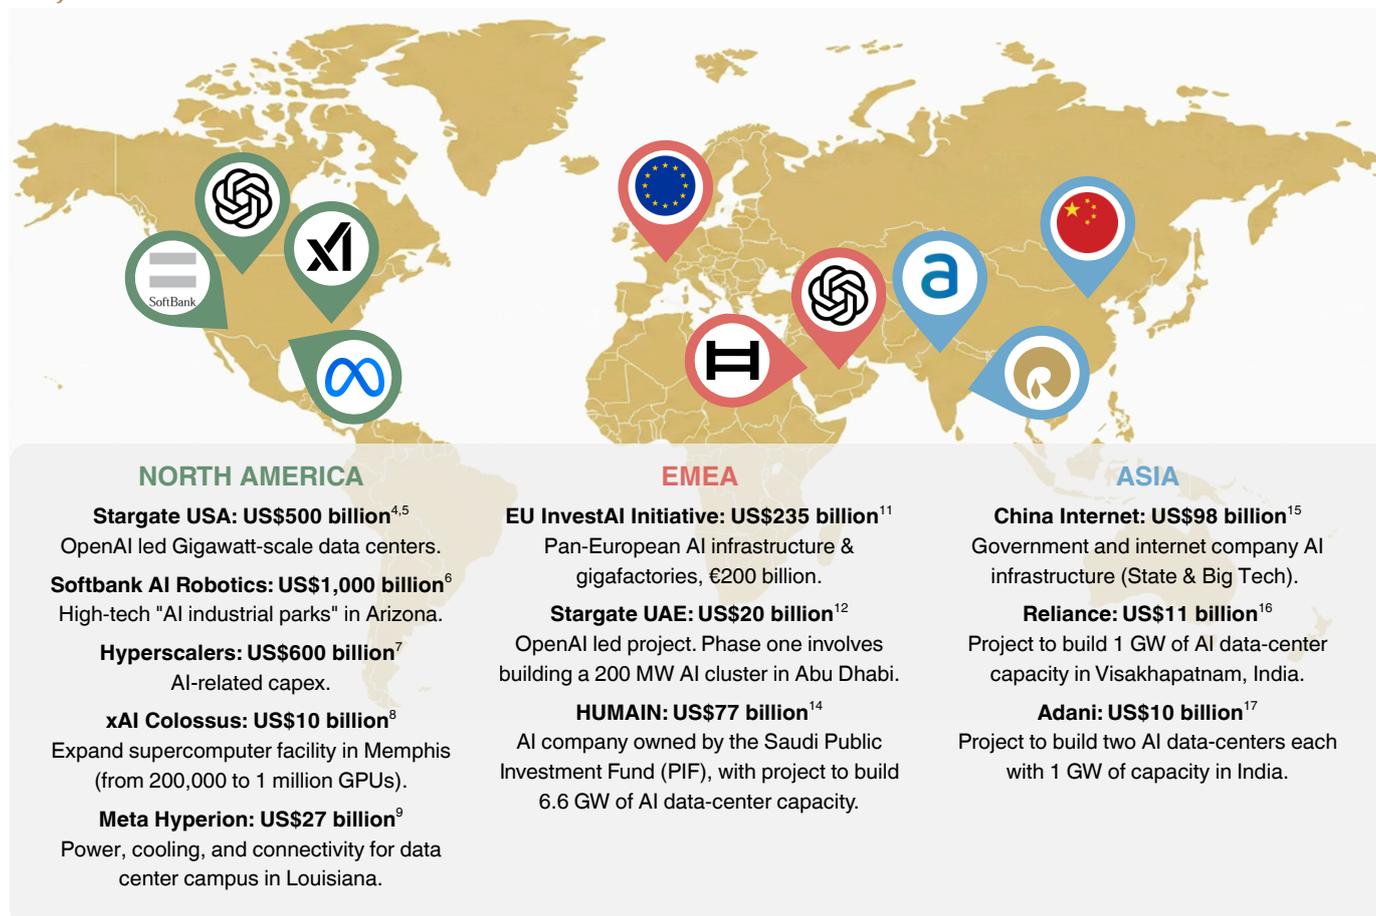


Figure 2.¹⁸ Major AI infrastructure announcements and estimates. Figure 2 reflects major publicly disclosed, market-visible projects rather than an exhaustive measure of total global capacity expansion. See Appendix 2.2 for further details, due dates, debt lenders, and equity sponsors.

OpenAI's "Stargate"

Stargate's scale is real, but committed equity covers only a small fraction of projected capex, leaving execution, refinancing, and technology risk squarely with lenders over long construction timelines. At the same time, the enormous funding gap also highlights the uncertainty surrounding the project's scale and execution timeline.

Overview and Scope of the Program

In January 2025, OpenAI, SoftBank, and a consortium of partners announced Stargate, a long-term initiative to build up to 10 GW of AI data center capacity.¹⁹ By September 2025, the partners disclosed five U.S. sites (three Oracle-led and two SoftBank-led campuses). These five sites account for almost 7GW and over US\$400 billion of committed investment, out of the program's ultimate target of 10GW and US\$500 billion.⁵

The initiative is structured as a network of multiple campuses rather than a single "mega-facility." Its scale dwarfs typical data-center projects.

To understand where such a vast sum is allocated, the planned US\$500 billion investment can be broken down across key technical domains (see Appendix 2.3 for further details):

- **Compute Hardware:** US\$200–225 billion
- **Data Center Construction:** US\$125–150 billion
- **Power Infrastructure:** US\$50–75 billion
- **Networking Infrastructure:** US\$25–40 billion
- **Operational Expenses:** US\$25–40 billion
- **Contingency:** US\$10–20 billion
- **Software Platform & Tools:** US\$15–25 billion
- **Research & Development:** US\$10–15 billion



Continued...

Per-gigawatt costs vary widely due to uncertainty around GPU architecture, power design, cooling density, and redundancy requirements. Small design choices can shift total capex by tens of billions of dollars, materially amplifying lender exposure.

Funding Commitments and the Financing Gap

Initial press coverage emphasized that Stargate would be funded through a combination of equity from sponsors and massive debt issuance; however, the actual equity committed has been modest relative to the headline budget. Specifically, OpenAI and SoftBank reportedly committed US\$19 billion each toward Stargate’s initial funding (40% stakes apiece); Oracle and Abu Dhabi-backed MGX were reported to contribute US\$7 billion each. ^{*.20,21}

These contributions amount to about 10.4% of the US\$500 billion target, leaving a funding gap of roughly US\$450 billion. The intention is to finance the remainder through public and private capital markets. Six months after launch, press reports noted that site selection and power procurement were still unresolved for several campuses, contributing to delays in project execution. In short, the project’s headline is real, but the path to completion is fraught with funding, technology, and execution risks that mirror the broader uncertainties of the AI infrastructure cycle.

2.2 The Hyperscalers

The Hyperscalers—namely, Amazon Web Services (AWS), Google Cloud, Microsoft Azure, Meta, and Oracle—are the primary drivers of demand, but not the sole source of capital.** While they can fund a significant portion of capex from robust operating cash flows—projected at over US\$500 billion in 2025—they still tap debt markets at scale to preserve liquidity and fund long-duration infrastructure.

The Hyperscalers are projected to spend ~US\$320 billion on AI-related capex in 2025.²² In 3Q 2025 alone, they spent US\$142 billion.²³ In particular, Amazon leads this cohort with a US\$100 billion allocation to scale AWS’s AI footprint. Signaling a growing appetite and conviction, Hyperscaler investments are concentrated on expanding data center capacity, strengthening backbone networks, and enhancing computational capabilities.

Despite their cash generation, the Hyperscalers have become large debt issuers. In 2025, the group issued

about US\$121 billion in bonds—four times their five-year average—to help fund infrastructure while preserving internal cash for other uses.²⁹

2025 Hyperscalers Capex vs. Cash Flow

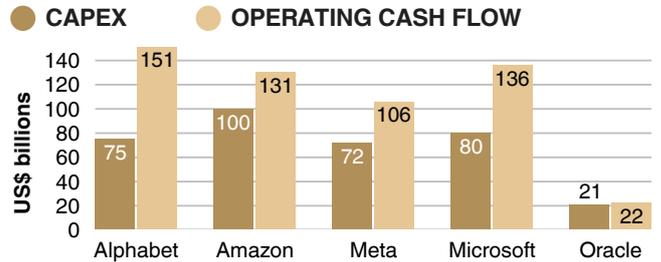


Figure 3. ^{24,25,26,27,28} AI infrastructure capex and operating cash flow of the Hyperscalers in 2025.

2.3 Anatomy of the US\$1.5 Trillion Financing Gap

Projected over the full 2025-2028 cycle, the total AI infrastructure capital requirement reaches US\$2.9 trillion. As shown in Table 5, roughly half, or around US\$1.4 trillion, is expected to be funded internally from Hyperscaler operating cash flows. The remaining US\$1.5 trillion must be sourced externally, comprising US\$1.15 trillion in credit and US\$350 billion in other capital, creating a massive demand for credit and alternative capital that will profoundly stress-test the capacity and risk appetite of global financial markets. Below is a detailed breakdown of the funding sources.

2025 Hyperscalers Capex vs. Cash Flow

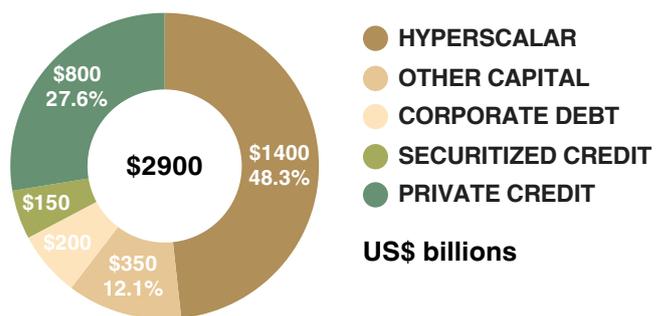


Figure 4.¹ Breakdown of financing paths. Includes 1) Hyperscaler Cash Flow, 2) Other Capital (PE³⁰, VC³¹, SWFs³²), 3) Corporate Debt³³, 4) Securitized Credit (ABS³⁴, CMBS³⁵), and 5) Private Credit³⁶ (via ABF and debt funding of JVs). See Appendix 2.4 for further details.

* Separately, OpenAI has reportedly agreed to spend US\$300 billion with Oracle for AI compute capacity (4.5GW), a commercial hosting/offtake arrangement distinct from Stargate’s equity funding.

** Capitalized “Hyperscalers” refers specifically to Amazon Web Services, Google Cloud, Microsoft Azure, Meta, and Oracle, while the lower case “hyperscalers” refers to all companies that build and operate data centers to provide on-demand cloud computing services. We define a sub-category because 1) they are the most influential, and 2) it is easier to quantify their spendings and cash flows versus those of the industry.



Fund Level Leverage

Public filings and ratings commentary show that large, public direct-lending vehicles—commonly called Business Development Companies (BDCs)—operate with debt-to-equity ratios above 1x and in some cases approaching regulatory caps of 2x, reflecting material balance-sheet leverage.^{37,38}

The implication is that private-credit lending often enters asset-level structures already levered at the platform level, before any additional project-level debt is applied. Effective system leverage is therefore likely higher than commonly measured, and widely cited aggregate figures—such as the US\$800 billion private-credit estimate—most plausibly reflect gross loan origination or exposure rather than unlevered equity capital, though opacity around fund-level leverage limits precise system-wide measurement.

A second, more aggressive interpretation is that the US\$800 billion figure reflects equity capital committed to private credit funds, rather than the total capital deployed. In that case, applying standard fund-level leverage of 1.0-1.2x would imply US\$1.4-1.6 trillion of actual lending capacity to AI infrastructure. This would introduce a second layer of leverage on top of already levered project finance structures, materially amplifying systemic exposure and increasing the potential for pro-cyclical stress.

In short, the first interpretation is more consistent with industry reporting conventions and is therefore more likely; however, the lack of transparency around how private credit exposure is measured means the second scenario cannot be ruled out. The true level of system leverage likely lies between these bounds, as most private credit funds are not transparent with their actual levels of borrowing, potentially obscuring the true scale of leverage embedded in the financing of AI infrastructure.

Fundraising Continuity and Capital Recycling

A further structural weakness lies in the system's reliance on continuous capital recycling. Large sponsors increasingly operate multi-vehicle platforms—spanning closed-end private funds, non-traded REITs, and BDCs—to rotate assets, access different investor bases, sustain AUM growth, and extend maturities. While this structure can improve capital

efficiency and provide liquidity pathways for early investors, it also amplifies the risks of layered leverage, related-party transactions, and growing exposure to less sophisticated private-wealth capital.* These dynamics raise questions about arm's-length valuation, transparency, and investor understanding of complex, illiquid structures originally designed for institutional balance sheets.³⁹

Blue Owl's Capital Cycle:

Part 1: Two Separate Transaction Facts

The US\$27 billion Meta Hyperion Joint Venture^{40,41}
Meta partnered with Blue Owl funds in an 80/20 JV to develop the US\$27 billion Hyperion Louisiana campus. For its equity share, Blue Owl contributed an initial US\$7 billion in cash. Market reports indicate the project debt was structured US\$27 billion private placement, with PIMCO subscribing to US\$18 billion and BlackRock to over US\$3 billion, and with an interest rate of 6.6%. Upon completion, Meta will enter into operating lease agreements with the JV and provide a residual value guarantee for the first 16 years of operations.

The US\$3.3 Billion BODIT Portfolio Acquisition⁴²
In a separate transaction, Blue Owl launched Blue Owl Digital Infrastructure Trust (BODIT), a REIT. BODIT raised US\$1.7 billion and used it to acquire a US\$3.3 billion portfolio of 11 stabilized data centers from Blue Owl's Fund I. The transaction was structured 'net of existing debt assumed,' resulting in a net equity purchase price of approximately US\$1.5 billion, and implies roughly US\$1.8 billion in pre-existing leverage.

Part 2: The Capital Lifecycle Model

While the two transactions involve entirely separate assets and are not directly linked, they exemplify the recurring, sequential phases that characterize Blue Owl's integrated platform strategy.

- **Origination:** Blue Owl-managed funds deploy equity (e.g., the US\$7 billion commitment) to secure controlling stakes in hyperscaler-anchored development projects, which are financed with massive project-level debt (e.g., the reported ~US\$27 billion private placement). This creates new, institution-grade assets within dedicated vehicles.
- **Transition – Stabilization:** Post-construction, once the asset's risk profile shifts and lease cash flows are established (as with Meta's long-term lease), the project transitions from a development bet to stabilized infrastructure with predictable returns, ready for long-duration ownership.
- **Recycling:** The BODIT acquisition is an example of this phase. Vehicles like BODIT, designed as permanent capital, exist to absorb mature assets. By purchasing the US\$3.3 billion stabilized portfolio from its own earlier Fund I, Blue Owl recycles capital, provides a liquidity event for Fund I investors, and seeds a new perpetual fee-generating vehicle.

* For example, Brookfield's non-traded REIT annual report risk factors explicitly warn that acquiring properties from affiliates may create conflicts and may lack arm's-length negotiations.



Continued...

Importantly, the BODIT transaction does not constitute a committed or announced exit pathway for the Hyperion project. Rather, it serves as an observable precedent that informs investor expectations regarding Blue Owl's capacity to manage asset transitions across vehicles as assets mature.

The Strategic Takeaway: Why the Connection Matters

Taken together, the Meta JV and the BODIT transaction reveal a continuous financing loop rather than isolated deals. Development-stage, hyperscaler-anchored assets can be originated within closed-end funds—while vehicles like BODIT

function as absorbers of mature assets, extending platform AUM and fundraising capacity.

This model enhances capital efficiency and supports continuity of fundraising. At the same time, it creates layered leverage and potential conflict of interest. The fragility, therefore, resides in the strategy's dependence on two continuous flows: the uninterrupted origination of new, levered projects by closed-end funds, and the uninterrupted appetite from permanent capital vehicles (and their underlying investors) to absorb the resulting stabilized assets. A failure in either flow risks stalling the entire engine.

Blue Owl's Data Center Playbook

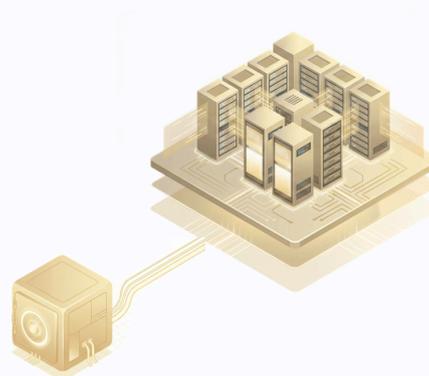
THE PRECEDENT

Blue Owl's Fund I acquired 11 data centers

- Investors seek high, absolute returns over a defined period (e.g., 7-10 years) and a clear exit.
- Fears: The project fails or exits are not realized.

US\$1.7 billion fundraise by BODIT to acquire 11 data centers from Fund I

- The vehicle allows for exit and offtake arrangements both from and with Blue Owl.
- Investors prioritize reliable dividend yield and capital preservation, accepting illiquidity for higher income.



META HYPERION

Funded with US\$7 billion in equity from Blue Owl

- US\$7 billion from a new group of investors.
- The structure set in the precedent signals a tested exit plan—potential for investors to rely on Blue Owl to replicate the BODIT model for Hyperion despite high leverage.
- Fears: Project delays and financing frictions could strand investor capital.

Figure 5. Blue Owl's data center playbook that could facilitate the Meta Hyperion joint venture.

Having established the scale of deployment and principal sources of capital, we should note that, importantly, credit risk is based on how that capital is structured, what assets or cash flows sit beneath it, and which assumptions lenders are implicitly underwriting. Two AI infrastructure projects of similar size may carry different risk profiles depending on whether repayment relies on enterprise-level cash flow, stabilized lease income, contracted compute revenues, or future asset monetization.

This distinction becomes especially salient in a market characterized by rapid technological iteration, long

construction and manufacturing timelines, and heterogeneous asset lives. GPUs, power infrastructure, shells, and land do not depreciate—or retain value—on the same schedule, yet they are often financed together within a single capital stack. As a result, credit exposure is shaped less by headline capex figures and more by where lenders sit in the stack and what ultimately services the debt. With this context established, the analysis turns to how credit is structured in practice. The following archetypes are not theoretical categories but observable financing patterns that determine where risk ultimately resides.

3. The AI Infrastructure & Credit Archetypes



3.1 Lay of the Land

Private credit financing of AI infrastructure is heterogeneous, segmented by borrower type, collateral, and repayment engine.

We have identified three main AI infrastructure credit archetypes: 1) the Project Finance archetype, 2) the GPU Platform archetype, and 3) the Hyperscaler Build-to-Operate (BtO) archetype. Together, these archetypes form a stacked and interdependent financing ecosystem rather than a linear spectrum, with capital flowing through different structures at different stages of the AI infrastructure lifecycle. This section elucidates how risk is engineered, where cash flows originate, and why certain layers attract private credit while others remain the domain of investment-grade capital markets.

3.2 Archetype 1: Project Finance

3.2.1 Archetype 1A: Construction Financing

- **WHAT LENDERS UNDERWRITE:** construction execution, grid interconnection, and lease-up timing.
- **WHAT BREAKS FIRST:** refinancing access if stabilization is delayed or rates rise.

Spending on data centers is expected to outpace office-building construction as soon as next year, marking a major shift in US commercial real estate.⁴³ In project finance, credit is broken down into the construction phase and the completed phase. During construction, lenders provide short-term construction and bridge financing to a developer or project SPV before a data center is fully built and leased. Since the asset generates no operating cash flow during this phase, lenders depend on reserves, advance payments, or sponsor support to cover interest. Repayment is ultimately tied to the asset's economic viability after completion.

Once the project is delivered and leased, the borrower is expected to refinance the project with long-term permanent debt—namely, Commercial Mortgage-Backed Securities (CMBS), Asset Backed Securities (ABS)—or sell the project to repay the loan.^{44,45}

Construction financing structures like mini-permanent loans, which bridge construction and early operations, are common. These typically have a total tenor of three to four years, after which refinancing at stabilization is anticipated.

In this archetype, the borrower is typically a project-level SPV with equity contributed by the sponsors and debt raised on a limited-recourse basis against the project's cash flows.⁹ The SPV owns the data center and isolates the sponsor from direct balance-sheet exposure, while long-term lease commitments from hyperscaler tenants underpin the revenue used to repay the financing.⁴⁶

Financing is underwritten on construction execution risk, grid interconnection timing, and future lease commitments. Lenders often require robust Engineering, Procurement, and Construction (EPC) contracts, performance bonds, and completion guarantees to mitigate delivery risk. Grid connection and tenant pre-leasing milestones are often requirements for drawdowns, reflecting the importance of minimizing execution risk before stabilization.⁴⁷

A high-visibility, real-world example is the large Oracle-linked data center campus in New Mexico, where a consortium of 20 banks provided a US\$18 billion facility to support development tied to Oracle as an anchor tenant. The reported pricing was Secured Overnight Financing Rate (SOFR) + 250bps, with a four-year maturity and extension options. The loan will be repaid upon completion and lease-up of the facility.⁴⁸

Principal risks are mainly execution risk (namely, delays, cost overruns, and grid interconnection) and lease-up risk (such as failure to secure anchor tenants on favorable terms). If interest rates rise or market conditions deteriorate during the construction period, developers may struggle to refinance their bridge facilities into permanent debt, potentially leading to distress or lender takeovers to complete and stabilize the asset. Further, utilities now frequently demand cash or letters of credit covering a project's share of grid-upgrade costs—often US\$50–150 million or more for a single facility.⁴⁹ This multi-million-dollar-per-megawatt liquidity burden must be posted before



construction, acting as a severe filter that favors only the deepest-pocketed sponsors and adding a material, pre-revenue credit risk for lenders.⁵⁰

With 57% of AI data center construction projects experiencing holdups of three months or more in 2025, tenants should anticipate later completion dates and incorporate flexibility into their timelines as delays in construction and supply chains will persist, impacting project schedules.⁵¹

3.2.2 Archetype 1B: Permanent Financing

- **WHAT LENDERS UNDERWRITE:** stabilized lease cash flows and tenant credit quality.
- **WHAT BREAKS FIRST:** asset values under oversupply or pricing pressure.

Permanent Financing refinances stabilized assets into long-term debt backed by lease cash flows. In the U.S., this transition is increasingly supported by structured finance markets, in particular by CMBS and ABS.^{*,34,35} The public ABS market provides a critical outlet for long-term lending on stabilized, leased facilities, with strong tenant covenants and predictable cash flows underpinning the permanent financings.^{52,53} In 2026, the market could witness data center ABS and CMBS originations surging to US\$50 billion.⁵⁴

A clear illustration of a stabilized takeout is Blackstone's US\$3.5 billion CMBS refinancing of ten QTS data centers, backed by the data centers themselves. Reported terms include a floating coupon around 6.3%, interest-only payments for an initial two-year period, and up to three one-year extension options.⁵⁵ The deal reflects how mature data center assets can access long-dated, structured capital at scale from institutional investors.

Demand for the leased, stabilized assets backing permanent financing remains strong. We expect lease rates to grow at a 5-7% CAGR through 2030. At the same time, 77% of capacity under construction is already pre-committed.⁵⁶ Part of this demand comes from hyperscalers who are opting to lease space from third-party owners instead of building all their own facilities. In 2024, 40% of hyperscaler data center capacity in the US was leased from third-party owners, up from 35% in 2023, while the remaining 60% was developed and operated directly by hyperscalers

themselves. At the same time, public and private funds are increasingly prioritizing the sector, with 95% of major institutional investors planning to raise their allocations to data centers;⁴³ however, this aggressive pre-committing itself could mask future oversupply if demand forecasts are overly optimistic, a risk for permanent lenders relying on long-term cash flow stability.

Monthly Office Construction Spending

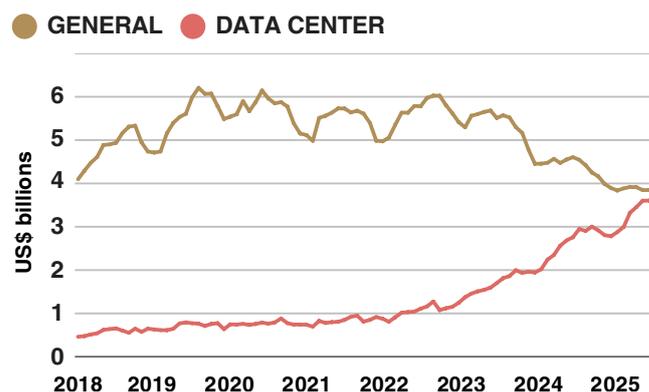


Figure 6.⁵⁷ Monthly construction spending on private general office buildings and data centers. General includes administration buildings, computer centers, office buildings, and professional buildings. Data centers refers to buildings that cater to the hardware needed for storing, processing, and transmitting digital information.

3.3 Archetype 2: GPU Platform Credit

- **WHAT LENDERS UNDERWRITE:** platform execution and the contracted compute revenues, and the assumed economic durability of GPU collateral over the debt tenor.
- **WHAT BREAKS FIRST:** cash flow if platform utilization falters, deployment is delayed, customer concentration becomes a risk, or rapid GPU value changes.

Secured lending to a GPU-cloud ("neocloud") platform is often used to acquire GPU computing infrastructure and the relevant equipment ("compute/equipment SPVs"). A typical borrower is an acquisition SPV within a broader operating platform. The SPV owns/holds the GPU fleet and related infrastructure, which it pledges to the lender as collateral. On top, the parent provides additional credit support through guarantees and equity pledges.

* Over the trailing twelve months ("TTM"), CMBS and ABS issuance combined totaled US\$557 billion (November 2024 to November 2025).



Lenders monitor performance using contract-linked tests like a Contract Realization Ratio, which measures actual billings/receipts versus projected contracted cash flows.⁵⁸ The repayment engine is built around the expected cash flow to be “guaranteed” by multi-year committed compute contracts.

That said, a fundamental risk underlies this collateral model: GPUs are subject to rapid technological obsolescence, as new architectures will likely deliver performance leaps and steadily erode the economic

value of prior generations. For lenders, this creates a structural vulnerability, as multi-year debt amortization must outpace the depreciation of the assets securing the loan. This risk may be compounded by a potential workload transition around 2027, when AI inference workloads may overtake training, a transformation that can undermine the utilization and pricing assumptions embedded in today’s GPU cloud contracts as inference is often more distributed and cost-sensitive. These risks will be examined and quantified further in Paper II.^{59,60}

CoreWeave: A GPU-Compute “Neocloud” Scaling via Structured Credit

CoreWeave exemplifies the GPU platform credit model: heavy leverage today in exchange for future contract cash flows. The company has US\$50 billion in remaining performance obligations backed by multi-year, take-or-pay contracts with customers like OpenAI and Meta, typically supported by 15–25% prepayments. Nevertheless, the near-term financial picture is stark, as 1H 2025 operating cash flow was negative US\$190 million, while capex reached US\$3.9 billion, highlighting the gap between massive upfront hardware spend and delayed revenue realization.⁶¹

To bridge this gap between future promises and current funding needs, CoreWeave employs a layered capital structure:

Secured, Asset-Backed SPV Debt:

Fulfilling specific long-term contracts (e.g., with OpenAI).

- **August 2025:** US\$2.6 billion secured delayed-draw term loan (DDTL 3.0) raised through specific acquisition entities (e.g., Coreweave Compute Acquisition Co. VII, LLC)
- **Pricing:** SOFR +400 bps
- **Collateral:** GPU^{62,63}

Unsecured Senior Notes (Non-Convertible):

General corporate purposes and debt repayment.

- **July 2025:** US\$1.75 billion of senior unsecured notes
- **Pricing:** 9.00% coupon, issued at par
- **Maturity:** February 1, 2031⁶⁴

Unsecured Senior Notes (Convertible):

Diversify funding sources and provide equity conversion optionality.

- **December 2025:** US\$2.25 billion of 1.75% convertible senior notes due 2031, upsized from an initial US\$2.0 billion
- **Maturity:** 2035
- **Pricing:** 1.75% coupon⁶⁵
- **Rating:** Fitch rating ‘BB-/RR4’⁶⁶

CoreWeave’s capital structure—secured, asset-backed SPV facilities alongside unsecured corporate notes—has produced

very high leverage. It has a total indebtedness of US\$14.2 billion, while EBITDA for the nine months ended September 30, 2025, totaled US\$1.7 billion, implying an 8.5x Debt/EBITDA ratio and 4.0x Debt/Revenue ratio as CoreWeave generated US\$3.6 billion of revenue over the same period.⁶⁷

For creditors and the system at large, it reflects several core risks:

- **The Cash Flow Mismatch:** Negative free cash flow occurs despite a US\$50 billion+ contract backlog (RPO).
- **Layered & Concentrated Leverage:** US\$18.8 billion total debt stacking secured SPV loans with unsecured/corporate debt. Historic revenue concentration where 77% of revenue came from the top 2 customers.
- **Execution as the Collateral:** Repayment depends on flawlessly deploying capital to fulfill massive, multi-year take-or-pay contracts. The asset (GPU) value is secondary to platform performance.
- **The “Sponsor” is a Contract:** Credit support comes not from a balance sheet but from contract prepayments (15–25% TCV) and the credibility of the counterparty (e.g., Meta, OpenAI). Lender underwriting is on the future platform’s viability.
- **Collateral Obsolescence vs. Debt Tenor:** The secured SPV debt being collateralized by GPUs creates a structural mismatch between multi-year debt amortization schedules and the sub-two-year technological replacement cycle, potentially exposing lenders to collateral decay even if contracts perform as expected.

The market is already partially pricing in this elevated risk profile, as CoreWeave’s share price dropped by over 50% from its peak in the second half of 2025.⁶⁸ Compared to established hyperscalers, CoreWeave pays a significant premium for debt. For example, its recent secured term loan is priced at SOFR +400bps, a spread approximately 355 basis points wider than a benchmark hyperscaler investment-grade credit facility from Amazon priced at SOFR +45bps.⁶⁹ This differential reflects how lenders underwrite the higher platform and execution risk inherent in the GPU neocloud model.

* EBITDA: US\$1,677,046,000 (for the nine months ended September 30, 2025) = Operating Income + Depreciation and Amortization = US\$43,589,000 + US\$1,633,457,000 = US\$1,677,046,000

** Both cases are senior secured credit facilities, and the ~355 bps spread is the debt market’s direct pricing of the incremental platform and execution risk of a GPU neocloud versus a cash-generative hyperscaler.



3.4 Archetype 3: Hyperscalers' Balance Sheet

- **WHAT LENDERS UNDERWRITE:** diversified enterprise cash flows.
- **WHAT BREAKS FIRST:** systemic or regulatory shocks to the core business, and the accumulation of asset-linked contingent obligations.

Hyperscalers are the most well-capitalized backers of AI infrastructure. Funded by public markets and internal cash flows, hyperscalers build and operate data centers primarily for their own cloud and AI services. In this model, Hyperscalers can raise financing at the parent-company level through investment-grade senior unsecured bonds and bank facilities. They can also raise through leased capacity, JVs, and SPVs (for instance, some large SPV arrangements may exist separately, such as Meta/Blue Owl for specific campuses). The use of proceeds is typically stated as general corporate purposes, including large data center capex, GPU deployments, working capital needs, and refinancing existing debt.

A typical borrower in this archetype issues investment-grade senior unsecured notes and may also use corporate revolvers/term loans. These instruments are generally unsecured and carry terms that emphasize broad corporate credit rather than specific collateral.

The repayment engine is the Hyperscaler's consolidated operating cash flow across its diversified businesses—including cloud services (AWS/Azure/Google Cloud Platform ("GCP")), advertising, subscriptions, and e-commerce—rather than revenue tied to a specific data center's lease or rent roll. Debt service relies on enterprise-wide free cash flow and balance-sheet liquidity, not project-specific DSCR, enabling these investment-grade credits to sustain very large capex plans with unsecured financing.⁷⁰

For the time being, the primary vulnerability in this archetype is largely company-level risk because debt is unsecured and serviced by broad enterprise earnings; major threats would come from extreme systemic disruption to the core business model (such as regulatory shocks, macro downturns, or structural declines in cloud/advertising revenue). Individual data center underperformance or hardware obsolescence is

unlikely to threaten near-term debt service.

That said, this insulation has to be maintained through explicit risk absorption at the corporate level. In certain hybrid structures, hyperscalers have provided asset-specific support—such as residual value guarantees or termination payments—to preserve operational flexibility (e.g., short lease tenors) while enabling third-party capital formation. For example, in Meta's joint venture with Blue Owl (see pages 8 & 9), Meta entered into four-year operating leases but provided a capped residual value guarantee covering the first 16 years of operations.⁴⁰ These mechanisms do not change how lenders underwrite hyperscaler credit, but they reintroduce asset-level obsolescence and utilization risk as contingent corporate obligations, effectively using investment-grade balance sheets to underwrite assets whose long-term economics may be uncertain.

Overall, hyperscaler corporate credit is currently among the lowest-risk layers within the credit stack; however, this archetype, at times, internalizes risk subconsciously, shifting uncertainty from asset-level lenders to corporate balance sheets in a manner that is stable until incentives, liquidity, or political tolerance to honor such support are tested.

3.5 Co-Mingling of Archetypes and Government Involvement Adds Complexity

In practice, these archetypes often blend within single financing stacks, creating hybrid structures. This comingling of risks—each with different durations, collateral, and repayment priorities—makes it difficult to disentangle exposures. In a downturn, complex cross-collateralization could complicate creditor recoveries, echoing structural opacity seen in past crises.

Additionally, this structural complexity is compounded by a new, powerful variable: the treatment of AI infrastructure as a strategic national imperative. This framing alters incentives, timelines, and risk perception, but not repayment mechanics.

Specifically, U.S. executive actions and policy guidance have articulated a clear federal objective to expand domestic AI compute and data-center capacity as a matter of national economic competitiveness and



security. Successive administrations under both Biden and Trump have acted favorably toward the expansion of AI infrastructure, signaling durable bipartisan support and continuity in policy direction.^{71,72} Strategic framing lowers the likelihood of outright abandonment of core infrastructure, but does not protect private capital from loss.

Importantly, strategic policy framing can also inadvertently amplify overbuilding pressures. In the White House's 2025 AI Action Plan, by framing AI as a national priority and removing barriers to infrastructure expansion, policymakers can reinforce optimistic expectations about near-term demand by providing an implicit government "put".⁷³ This dynamic lowers perceived project risk and can encourage developers, investors, and lenders to pursue larger, faster buildouts, even where monetization paths remain uncertain.

Thus, the national strategic imperative surrounding AI infrastructure complicates the direct analogy to the Dot-Com bubble's purely market-driven cycle. While government policy and regulation shaped the economic landscape of the late 1990s—such as deregulation under the Telecommunications Act of 1996 that helped unleash massive private investment in fiber networks—there was no broad federal program that directly sustained or "saved" dark fiber capacity after the telecom crash. In the five years after 1996, telecom operators invested more than US\$500 billion, largely funded through leverage, into fiber deployment, switching infrastructure, and wireless networks.⁷⁴ Much of this capital was deployed ahead of realized demand, leaving utilization well below capacity in the near term.

After the Internet Bubble burst, wholesale bandwidth and long-haul fiber capacity experienced a severe price collapse as overcapacity emerged, with transmission costs falling by as much as ~90% in the early 2000s.⁷⁵ Much of the excess fiber laid during that period remained unused and was largely a consequence of private overinvestment rather than strategic public intervention.

This highlights a key distinction: governments today may be more willing to support AI infrastructure deployments explicitly as a strategic priority, whereas in the telecom build-out of the early 2000s, the overcapacity was resolved through market forces,

often resulting in sharply lower prices and financial distress for private capital without a government backstop. However, even with a more likely policy backdrop today from the federal government, it does not abolish the fundamental risk of overinvestment relative to near-term monetization. In fact, by accelerating deployment and dampening market signals, it may prolong overcapacity and deepen eventual financial distress for private capital stacks, even as it secures the long-term strategic asset. This reinforces our core argument: the eventual indispensability of the infrastructure does not immunize the financing built upon it from the pressures of a boom-bust cycle, a lesson firmly rooted in the Dot-Com experience, but now operating under a powerful new political-economic influence.

2026 Policy Signals Shift AI Infrastructure Costs

This dynamic is already observable in practice. In early 2026, amid growing community backlash over electricity prices, water usage, and local environmental impacts, President Trump publicly emphasized that large technology firms building AI data centers should "pay their own way."⁷⁶ In response, Microsoft announced a set of voluntary commitments: agreeing to pay electricity rates sufficient to cover incremental grid and infrastructure costs, forgoing local tax abatements, and structuring utility agreements to ensure that nearby communities would not face higher power bills.⁷⁷

This illustrates how strategic policy framing can sustain deployment momentum even as underlying project economics become more demanding. Importantly, this mechanism does not shield private capital from loss. Instead, it preserves buildout continuity through public acceptance by shifting social, environmental, and infrastructure externalities onto corporate balance sheets, layering additional financial pressure atop already complex and levered capital structures. In this way, strategic framing may reduce the likelihood of outright project abandonment, while simultaneously amplifying balance-sheet and credit stress within the private financing stack—reinforcing the broader boom-bust risks outlined above.



4. Closing



This paper establishes a framework for understanding the AI infrastructure boom through a credit lens. Three conclusions follow.

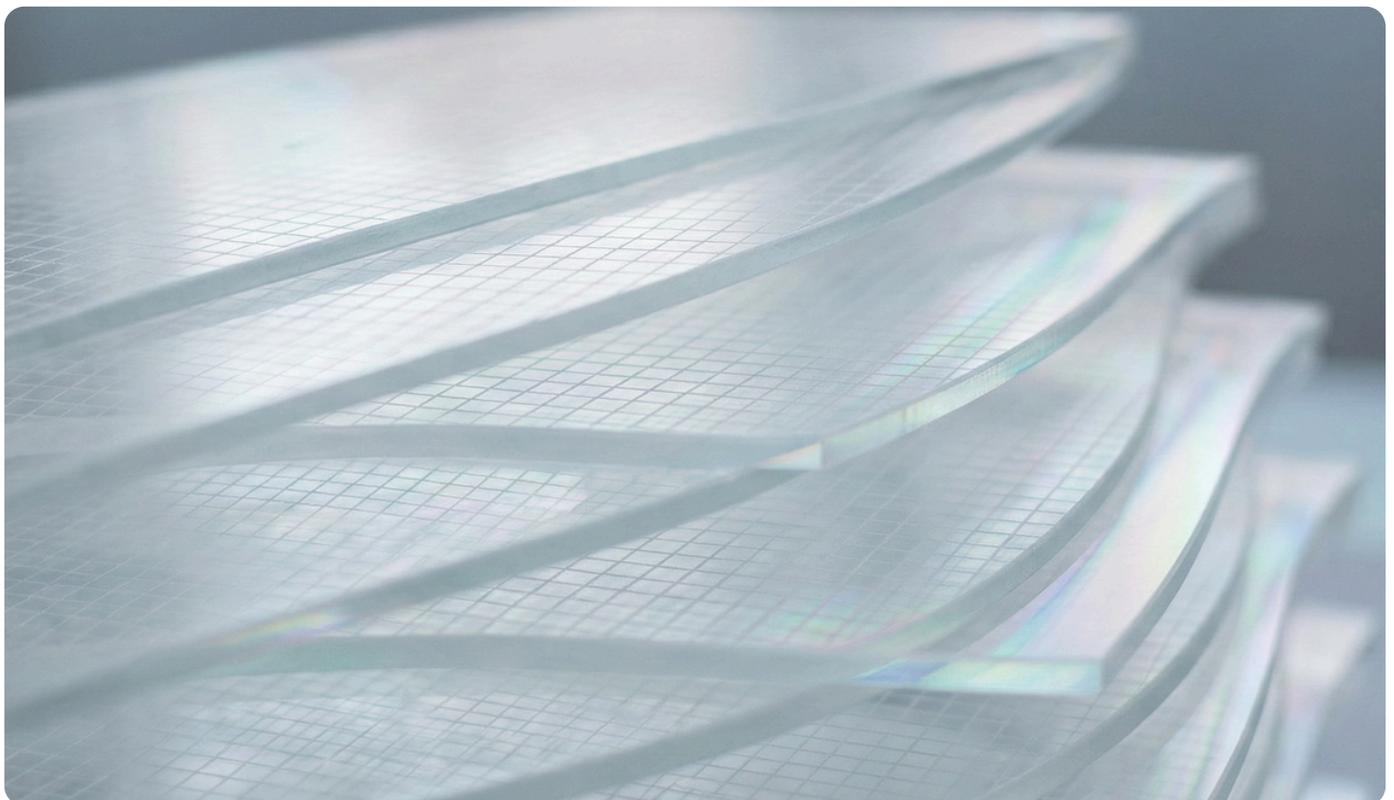
- **Private credit is the dominant engine.**
Of the US\$1.5 trillion external financing gap, approximately US\$800 billion is expected to come from private credit, making its structures, leverage, and capital-recycling models central to systemic risk.
- **Risk is stratified across three archetypes.**
Financing is not uniform but split into distinct layers: Project Finance (construction/lease-up risk), GPU Platform Credit (contract fulfillment risk against large deficits), and Hyperscaler Balance Sheet (lowest-risk, enterprise-funded). These layers are interdependent.
- **A strategic imperative might distort the cycle.**
Government support for AI as a national priority accelerates deployment and may prevent asset abandonment, but it does not eliminate financial risk. It can instead prolong overcapacity and amplify distress for private capital stacks.

The sustainability of this buildout hinges not only on the scale of investment but also on the cash flows supporting these layered debt structures. Part II will zoom in on key risks—namely, obsolescence, concentration, and systemic leverage—to stress-test the durability of these credit structures.

At the end of the day, we are merely trying to comprehend and draw insights from the enormous amount of data, narratives, and announcements that are available in the market. If our analysis allows us to better gauge the future direction by even 1%, then we would have done our investors justice. If you have any further questions regarding AI, credit, and our thoughts, please feel free to reach out to IR@unityinvestments.com.

$$A = P \left(1 + \frac{r}{n}\right)^{nt}$$

This is Life, Compounded.



References / Appendix



References

- [1] Morgan Stanley Research. (n.d.). *Bridging a \$1.5 trillion data center gap*. Morgan Stanley. <https://www.scribd.com/document/903369274/Bridging-a-1-5tr-Data-Center>
- [2] McKinsey & Company. (n.d.). The cost of compute: A \$7 trillion race to scale data centers. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers>
- [3] McKinsey & Company. (n.d.). The data center dividend. <https://www.mckinsey.com/featured-insights/week-in-charts/the-data-center-dividend>
- [4] Reuters. (2025, October 1). Key stakeholders in the \$500 billion Stargate AI project. <https://www.reuters.com/business/media-telecom/key-stakeholders-500-billion-stargate-ai-project-2025-10-01/>
- [5] OpenAI. (n.d.). Five new Stargate sites. <https://openai.com/index/five-new-stargate-sites/>
- [6] Reuters. (2025, June 20). SoftBank's Son pitches \$1 trillion US AI hub. <https://www.reuters.com/business/media-telecom/softbanks-son-pitches-1-trillion-us-ai-hub-tsmc-trump-team-bloomberg-news-2025-06-20/>
- [7] CreditSights. (n.d.). Technology hyperscaler capex: 2026 estimates. <https://know.creditsights.com/insights/technology-hyperscaler-capex-2026-estimates/>
- [8] Data Center Dynamics. (n.d.). xAI's Memphis Colossus GPU deployment. <https://www.datacenterdynamics.com/en/news/xai-elon-musk-memphis-colossus-gpu/>
- [9] Global Data Center Hub. (n.d.). Meta & Blue Owl's \$27 billion data center bet. <https://www.globaldatacenterhub.com/p/meta-blue-owls-27b-bet-is-this-the>
- [10] Bloomberg. (2025, October 16). Blue Owl seals largest private capital deal for Meta's AI growth. <https://www.bloomberg.com/news/articles/2025-10-16/blue-owl-seals-largest-private-capital-deal-for-meta-s-ai-growth>
- [11] European Commission. (2025, February 11). EU launches InvestAI initiative to mobilise €200 billion. https://luxembourg.representation.ec.europa.eu/actualites-et-evenements/actualites/eu-launches-investai-initiative-mobilise-eu200-billion-investment-artificial-intelligence-2025-02-11_en
- [12] The National. (2025, December 5). Stargate UAE's first phase to be completed in Q3 2026. <https://www.thenationalnews.com/business/2025/12/05/stargate-uaes-first-phase-to-be-completed-in-third-quarter-of-2026/>
- [13] AlInvest. (2025). Abu Dhabi Stargate: A \$20B AI powerhouse. <https://www.ainvest.com/news/abu-dhabi-stargate-uae-20b-ai-powerhouse-future-global-cloud-infrastructure-2507/>
- [14] Data Center Dynamics. (n.d.). Saudi Arabia's Humain seeks US data center equity partner. <https://www.datacenterdynamics.com/en/news/saudi-arabias-ai-co-humain-looking-for-us-data-center-equity-partner-targets-66gw-by-2034-with-subsidized-electricity/>
- [15] South China Morning Post. (2025). China's AI capital spending set to reach \$98 billion. <https://www.scmp.com/tech/tech-war/article/3315805/chinas-ai-capital-spending-set-reach-us98-billion-2025-amid-rivalry-us>
- [16] Reuters. (2025, November 26). Reliance Industries JV to invest \$11 billion in data centers. <https://www.reuters.com/world/india/indias-reliance-industries-jv-invest-11-billion-data-center-2025-11-26/>
- [17] Data Center Dynamics. (n.d.). Adani could spend \$10 billion on Indian data centers. <https://www.datacenterdynamics.com/en/news/adani-could-spend-10bn-on-two-data-centers-in-india-report>
- [18] Allianz Global Investors. (2025). Stephen Luncheon presentation. <https://hk.allianzgi.com/-/media/AllianzGI/AP/HongKong/Documents/Presentation/eventpresentation/2025/09/25/01/48/20250930-Stephen-Luncheon>
- [19] OpenAI. (n.d.). Announcing the Stargate project. <https://openai.com/index/announcing-the-stargate-project/>
- [20] Reuters. (2025, January 23). OpenAI and SoftBank each commit \$19 billion to Stargate. <https://www.reuters.com/technology/openai-softbank-each-commit-19-bln-stargate-data-center-venture-information-2025-01-23>
- [21] GSN Online. (n.d.). MGX invests alongside tech giants in Stargate. <https://www.gsn-online.com/news-center/article/abu-dhabis-mgx-invests-along-tech-giants-us-500bn-stargate-ai>
- [22] CNBC. (2025, February 8). Tech megacaps to spend more than \$300 billion in 2025. <https://www.cnbc.com/2025/02/08/tech-megacaps-to-spend-more-than-300-billion-in-2025-to-win-in-ai.html>
- [23] Goldman Sachs. (2026). Why AI companies may invest more than \$500 billion. <https://www.goldmansachs.com/insights/articles/why-ai-companies-may-invest-more-than-500-billion-in-2026>
- [24] Data Center Dynamics. (2025). Google expects 2025 capex to surge to \$75 billion. <https://www.datacenterdynamics.com/en/news/google-expects-2025-capex-to-surge-to-75bn-on-ai-data-center-buildout/>
- [25] Data Center Dynamics. (2025). Amazon 2025 capex to reach \$100 billion. <https://www.datacenterdynamics.com/en/news/amazon-2025-capex-to-reach-100bn-aws-revenue-hit-100bn-in-2024>
- [26] Meta Platforms, Inc. (2025). Meta reports third quarter 2025 results. <https://investor.atmeta.com/investor-news/press-release-details/2025/Meta-Reports-Third-Quarter-2025-Results/>
- [27] Reuters. (2025, January 3). Microsoft plans to spend \$80 billion on AI-enabled data centers. <https://www.reuters.com/technology/artificial-intelligence/microsoft-plans-spend-80-bln-ai-enabled-data-centers-fiscal-2025-cnbc-reports-2025-01-03/>



- [28] Nasdaq. (n.d.). Oracle, Microsoft and Alphabet earnings preview. <https://www.nasdaq.com/articles/oracle-microsoft-and-alphabet-are-part-zacks-earnings-preview>
- [29] Investing.com. (n.d.). Hyperscaler bond spreads widen as tech giants issue \$121 billion in debt. <https://www.investing.com/news/stock-market-news/hyperscaler-bond-spreads-widen-as-tech-giants-issue-121bn-in-debt-93CH-4367010>
- [30] PE Insights. (2024). Global private equity assets reach \$10.8 trillion. <https://pe-insights.com/global-private-equity-assets-reach-record-10-8tn-in-2024/>
- [31] National Venture Capital Association. (2025). 2025 NVCA yearbook. https://nvca.org/press_releases/nvca-releases-2025-yearbook-showcasing-2024-vc-trends/
- [32] IQ-EQ. (n.d.). Sovereign wealth funds are positioning portfolios for growth. <https://iqeq.com/insights/sovereign-wealth-funds-are-positioning-portfolios-for-growth/>
- [33] SIFMA. (n.d.). US corporate bond statistics. <https://www.sifma.org/research/statistics/us-corporate-bonds-statistics>
- [34] SIFMA. (n.d.). US asset-backed securities statistics. <https://www.sifma.org/research/statistics/us-asset-backed-securities-statistics>
- [35] SIFMA. (n.d.). US mortgage-backed securities statistics. <https://www.sifma.org/research/statistics/us-mortgage-backed-securities-statistics>
- [36] Unity Investments. (n.d.). Private credit primer: Part I. <https://www.unityinvestments.com/post/private-credit-primer-part-1-origin-strategies-players-and-bubbles>
- [37] Houlihan Lokey. (2018). Direct lending update. <https://cdn.hl.com/pdf/2018/direct-lending-update-2018.pdf>
- [38] BDC Investor. (n.d.). Most leveraged BDCs. <https://www.bdcinvestor.com/screens/most-leveraged-bdcs/>
- [39] Eversheds Sutherland. (n.d.). Structuring publicly traded private equity vehicles. <https://www.eversheds-sutherland.com/en/poland/client-tools/publicly-traded-private-equity/structuring-investment-vehicles>
- [39] Eversheds Sutherland. (n.d.). Structuring publicly traded private equity vehicles. <https://www.eversheds-sutherland.com/en/poland/client-tools/publicly-traded-private-equity/structuring-investment-vehicles>
- [40] Meta Platforms, Inc. (2025). Meta announces JV with funds managed by Blue Owl Capital. <https://investor.atmeta.com/investor-news/press-release-details/2025/Meta-Announces-Joint-Venture-with-Funds-Managed-by-Blue-Owl-Capital-to-Develop-Hyperion-Data-Center/>
- [41] Longbridge. (n.d.). Data center investment news. <https://longbridge.com/en/news/262131565>
- [42] U.S. Securities and Exchange Commission. (2025). OSNL Form 8-K. <https://www.sec.gov/Archives/edgar/data/2069692/000206969225000025/osnl-20251201.htm>
- [43] CRE Daily. (n.d.). Data center investment surges amid AI boom. <https://www.credaily.com/briefs/data-centers-investment-surges-amid-ai-boom/>
- [44] Project Finance Law. (2025, June). Data center financing structures. <https://www.projectfinance.law/publications/2025/june/data-center-financing-structures>
- [45] Investopedia. (n.d.). Mini-perm financing. <https://www.investopedia.com/terms/m/miniperm.asp>
- [46] Osborne Clarke. (n.d.). EU and UK data center financing trends. <https://www.osborneclarke.com/insights/eu-and-uk-financing-trends-are-reshaping-data-center-projects>
- [47] Skadden. (2025). Hyperscaler data centers. <https://www.skadden.com/insights/publications/2025/09/insights-september-2025/corporate/hyperscaler-data-centers>
- [48] Reuters. (2025, November 7). Banks lend \$18 billion to Oracle-tied data center project. <https://www.reuters.com/business/finance/banks-lend-18-billion-oracle-tied-data-center-project-bloomberg-news-reports-2025-11-07>
- [49] EUCI. (n.d.). Utilities adopt large-load tariffs for data centers. <https://www.euci.com/utilities-adopt-large-load-tariffs-to-cope-with-the-costs-and-power-demands-of-data-centers/>
- [50] ISO New England. (n.d.). FERC order to show cause. https://www.iso-ne.com/static-assets/documents/100012/ferc_order_to_show_cause_on_to_initial_to_funding.pdf
- [51] JLL. (n.d.). Global data center market outlook. <https://www.jll.com/en-de/insights/market-outlook/global-data-centers>
- [52] Ion Analytics. (n.d.). Data center refs to spur ABS issuance. <https://ionanalytics.com/insights/infralogic/analysis-data-center-refs-to-spur-abs-issuance>
- [53] Hogan Lovells. (n.d.). Financing data centers in the US and Europe. <https://www.hoganlovells.com/en/publications/financing-data-centers-in-the-us-and-europe>
- [54] JLL. (n.d.). Global data center sector to nearly double to 200GW. <https://www.jll.com/en-ca/newsroom/global-data-center-sector-to-nearly-double-to-200gw-amid-ai-infrastructure-boom>
- [55] Commercial Property Executive. (n.d.). Blackstone's QTS to close \$3.5B data center refi. <https://www.commercialsearch.com/news/blackstones-qts-to-close-3-5b-data-center-refi/>
- [56] Capacity Global. (n.d.). Global data center capacity set to double. <https://capacityglobal.com/news/data-centre-capacity-set-to-double-jll/>
- [57] U.S. Census Bureau. (n.d.). Value of construction put in place. https://www.census.gov/construction/c30/historical_data.html
- [58] U.S. Securities and Exchange Commission. (2025). CRWV Form 8-K. <https://www.sec.gov/Archives/edgar/data/1769628/000176962825000033/crwv-20250728.htm>
- [59] Rev. (2026). NVIDIA CES 2026 keynote transcript. <https://www.rev.com/transcripts/nvidia-at-ces-2026>



- [60] JLL. (n.d.). Data center outlook. <https://www.jll.com/en-us/insights/market-outlook/data-center-outlook>
- [61] U.S. Securities and Exchange Commission. (2025). CRWV Form 8-K. <https://www.sec.gov/Archives/edgar/data/1769628/000176962825000062/crwv-20250930.htm>
- [62] CoreWeave. (2025). CoreWeave closes \$2.6B secured debt financing. <https://investors.coreweave.com/news/news-details/2025/CoreWeave-Closes-2-6-Billion-Secured-Debt-Financing-Facility-Strengthening-Market-Position-as-AI-Cloud-Leader>
- [63] MarketScreener. (2025). CoreWeave closes \$2.6B secured debt facility. <https://www.marketscreener.com/news/coreweave-inc-closes-a-2-6-billion-secured-debt-financing-facility-ce7c5fdd8bf022>
- [64] U.S. Securities and Exchange Commission. (2025). CoreWeave press release. <https://www.sec.gov/Archives/edgar/data/1769628/000176962825000030/coreweave-pressreleasepric.htm>
- [65] Latham & Watkins. (2025). Advises on CoreWeave convertible senior notes offering. <https://www.lw.com/en/news/2025/12/latham-watkins-advises-on-coreweave-convertible-senior-notes-offering>
- [66] Fitch Ratings. (2025). Fitch rates CoreWeave senior unsecured convertible notes. <https://www.fitchratings.com/research/corporate-finance/fitch-rates-coreweave-new-senior-unsecured-convertible-notes-bb-rr4-10-12-2025>
- [67] CoreWeave. (2025). Investor presentation. <https://d18rn0p25nwr6d.cloudfront.net/CIK-0001769628/e660873e-365c-4e39-9c88-923d050a4d55.pdf>
- [68] The Globe and Mail. (n.d.). Is CoreWeave stock yesterday's news? <https://www.theglobeandmail.com/investing/markets/markets-news/Motley%20Fool/36981486/is-coreweave-stock-yesterday-s-news/>
- [69] U.S. Securities and Exchange Commission. (2024). Amazon Form 10-Q. <https://www.sec.gov/Archives/edgar/data/1018724/000101872424000161/amzn-20240930xex101.htm>
- [70] Janus Henderson Investors. (n.d.). Mega issuance and the AI arms race. <https://www.janushenderson.com/en-sg/investor/article/mega-issuance-and-the-ai-arms-race-big-techs-impact-on-credit-spread>
- [71] The White House. (2025, January 14). Executive order on advancing US AI infrastructure. <https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2025/01/14/executive-order-on-advancing-united-states-leadership-in-artificial-intelligence-infrastructure/>
- [72] The White House. (2025, July). Accelerating federal permitting of data center infrastructure. <https://www.whitehouse.gov/presidential-actions/2025/07/accelerating-federal-permitting-of-data-center-infrastructure/>
- [73] The White House. (2025). America's AI action plan. <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>
- [74] Runkevicius, D. (2025, March 24). Is the AI boom headed for its dark fiber moment? Forbes. <https://www.forbes.com/sites/danrunkevicius/2025/03/24/is-the-ai-boom-headed-for-its-dark-fiber-moment/>
- [75] Bandwidth glut lives on. (2004). Wired. <https://www.wired.com/2004/09/bandwidth-glut-lives-on>
- [76] The Guardian. (2026, January 13). Trump, Microsoft and the electricity cost of data centers. <https://www.theguardian.com/technology/2026/jan/13/trump-microsoft-data-centers-electricity>
- [77] CNBC. (2026, January 13). Microsoft data centers drive up energy prices. <https://www.cnbc.com/2026/01/13/microsoft-data-centers-energy-ai-prices.html>

Appendix

Appendix 2.1: Table 1.³ Segment allocations as a percentage of US\$2.9 trillion and US\$5.2 trillion total AI capex estimate. Figures may not sum to totals because of rounding.

	Segment	% of US\$5.2 billion total	2028 Capex (back-cast, US\$ billions)	2030 Capex (US\$ billions)
Designers	Servers (GPU, CPU)	52%	\$1,510	\$2,700
	Storage	12%	\$350	\$620
Builders	Labor	9%	\$260	\$470
	Shell and site	5%	\$150	\$260
	Land acquisition	2%	\$60	\$100
Energizers	Electrical & mechanical equipment	12%	\$350	\$620
	Power generation	6%	\$170	\$310
	Network infrastructure	2%	\$60	\$100
	Totals	100%	\$2,900	\$5,200



Appendix 2.2: Table 2.¹⁸ Major AI infrastructure announcements and estimates. Table 2 reflects major publicly disclosed, market-visible AI infrastructure projects rather than an exhaustive measure of total global capacity expansion.

Region	Borrower	Amount (US\$ billions)	Purpose	Due	Debt Lenders	Equity Sponsors
North America	Stargate USA ^{4,5}	\$500	Build gigawatt-scale data centers	2029	Mizuho, JPMorgan Chase (US\$2.3 bn for Abilene)	SoftBank, OpenAI, Oracle, MGX
	SoftBank AI Robotics / Project Crystal Land ⁶	\$1,000 (concept)	High-tech "AI industrial parks" in Arizona	TBD	Not yet financed	SoftBank, TSMC (approached)
	Hyperscalers (aggregate) ⁷	\$600	AI-related capex	2026	Investment-grade bonds (>US\$150 bn issued 2025)	Alphabet, Amazon, Meta, Microsoft, Oracle
	xAI Colossus supercomputer ⁸	\$10	Expand Memphis supercomputer (200k to 1m GPUs)	2026	Morgan Stanley (US\$5 bn loans)	xAI, SpaceX (US\$2 bn), Saudi PIF (US\$0.8 bn)
	Meta Hyperion data-center campus (Louisiana) ⁹	\$27	Data centers, power, cooling, connectivity	2049 ¹⁰	PIMCO (US\$18 bn), BlackRock (US\$3 bn)	Blue Owl (80%), Meta (20%)
EMEA	EU InvestAI Initiative ¹¹	\$235	Pan-European AI infrastructure & gigafactories	Mid-2030s	European Investment Bank	Public-private partnerships, member states
	Stargate UAE (Abu Dhabi) ¹²	\$20 ¹³	Build 200 MW AI cluster (first phase)	Q3 2026	Not disclosed	SoftBank, G42, OpenAI, Oracle, Nvidia, Cisco
	HUMAIN (Saudi PIF) ¹⁴	≈\$77	Build 6.6 GW AI data-center capacity	2034	US\$10 bn venture fund (planned)	Saudi Public Investment Fund, AWS, AMD
Asia	China Internet (State & Big Tech) ¹⁵	\$98	Government and internet company AI infrastructure	2025	Government funding (US\$55bn)	Alibaba, Tencent, and major internet companies
	India Tech – Reliance ¹⁶	≈\$11	Build 1 GW AI data-center capacity (Visakhapatnam)	2025–2030	Project finance (details not public)	Reliance, Brookfield, Digital Realty
	India Tech – Adani ¹⁷	≈\$10	Build two 1 GW AI data centers	2026	Not disclosed	Adani Group, Google partnership (US\$5 bn)
Totals		\$2,588				

Appendix 2.3: Table 3. Breakdown of US\$500 billion Stargate investment into key technical domains

Investment Category	Percentage of Total	Estimated Cost (US\$ billions)	Key Components & Notes
Compute Hardware	40-45%	\$200-225	AI accelerators (GPUs), CPUs, high-bandwidth memory, fast storage, and high-speed interconnects.
Data Center Construction & Land	25-30%	\$125-150	Land acquisition for 20 sites and buildout of Tier 3/4 facilities, including power, cooling, and security.
Power Infrastructure	10-15%	\$50-75	High-voltage grid connections, backup power systems, renewable energy integration, and power management.
Networking Infrastructure	5-8%	\$25-40	High-bandwidth, low-latency interconnects between sites, internet bandwidth, and specialized hardware.
Operational Expenses	5-8%	\$25-40	Initial 5-year costs for personnel, power, utilities, maintenance, and security operations.
Contingency	2-4%	\$10-20	Fund for unforeseen challenges, cost overruns, and risk mitigation.
Software Platform & Tools	3-5%	\$15-25	AI software stacks, orchestration platforms, data management, and security software.
Research & Development	2-3%	\$10-15	Next-gen data center tech, AI model optimization, and sustainable energy research.

Appendix 2.4: Table 4.¹ Breakdown of financing paths for US\$2.9 trillion capex projection

Financing Paths	Amount (US\$ billions)	Industry Size Today (US\$ billions)	% of Industry
Hyperscaler Cash Flow	\$1,400	(see Figure 3, operating cash flow total)	N/A
Other Capital (PE, VC, SWFs)	\$350	PE: 10,800 ²⁰ VC: 1,250 ²¹ SWFs: ≥12,000 ²²	N/A
Non-Credit Total	\$1,750		
Corporate Debt Issuance	\$200	2326 TTM, 11381 outstanding (2Q25) ²³	8.6%
Securitized Credit Issuance (ABS, CMBS)	\$150	ABS: 489 Issued TTM ²⁴ CMBS: 68 Issued TTM ²⁵	26.9%
Private Credit (via ABF and Debt Funding of JVs)	\$800	1,800 AUM ²⁶	44.4%
Credit Total	\$1,150	4,683	24.6%

The information contained herein is for informational purposes only and is not intended as, and may not be relied on in any manner as legal, tax or investment advice, a recommendation, or as an offer to sell, a solicitation of an offer to purchase or a recommendation of any interest in any security offered by Unity Angels Management Inc. or its affiliates ('Unity'). The information is subject to change and may also be incomplete.

Unity believes that information to be reliable but does not assume any responsibility for the independent verification of such information and makes no representation, warranty or guarantee as to its accuracy or completeness. Therefore, Unity assumes no liability for the information provided.

Unity may issue, and may in the future issue, materials that are inconsistent with, and reach different conclusions from, the information presented herein. Those materials may reflect different assumptions, views, and analytical methods and Unity is under no obligation to ensure that such other materials are brought to the attention of any user.

