



What a Wonderful World

Financing the AI revolution and finding where the risks land

*"I see trees of green
Red roses too
I see them bloom
For me and you
And I think to myself
What a wonderful world"*

– Louis Armstrong

Executive summary

The artificial-intelligence investment cycle has entered a new phase.

Artificial intelligence (AI) has given rise to one of history's largest investment cycles. The most valuable technology companies are racing to secure chips, power and land to support their growing ambitions. Investors are rushing to finance this exciting revolution, focusing on the promise of blooming red roses and the wonderful world of abundance and productivity rising over the horizon.

Early in the cycle, the largest firms could finance almost any plausible data-centre budget from cash. But that is changing: Capital spending is catching up with operating cash flow, while debt issuance, long-term leases, project vehicles, private credit and securitization are filling the gap.

At first glance, the system looks financeable. The largest borrowers are still highly rated, public bond markets are deep and private capital is abundant. More importantly, demand for compute still exceeds deliverable supply, meaning long-term contracts can turn uncertain future usage into cash flows that lenders can underwrite. **That is the forest. The trees tell a more complicated story.**

The same dollar of expected AI demand can support several claims: A model developer commits to cloud capacity, the cloud provider signs a lease, a developer finances the facility. The facility owner issues debt, the equipment is financed again, and a guarantee may then make the structure look investment-grade.

Nothing improper is required for risk to become difficult to see.

The core risk is not a simple shortage of capital, but a mismatch among five clocks: the time needed to build the infrastructure, the time needed to connect reliable power, the time needed for customers to adopt AI, the economic life of the equipment and the maturity of the debt.

If the clocks stay aligned, the system can work well. If they do not, risk will migrate to project companies, private funds, insurers, banks, utilities and, in some cases, ratepayers or taxpayers. Some of it may still return to the original sponsor through guarantees, lease obligations or reputational support.

That is where a comparison with housing finance before 2008 becomes useful.

Data centres are clearly not subprime mortgages. The borrowers, the collateral and the banking regime are different. Yet the financing mechanism rhymes: Risk is being sliced, transferred, enhanced and distributed. Ratings can make complex exposures look familiar, and off-balance-sheet treatment can obscure economic leverage. Correlation can stay hidden until many projects turn out to be relying on the same tenants, chips, power markets and refinancing window.

The lesson from 2008 is not that securitization is inherently dangerous but rather that **moving risk is not the same as removing it**. In other words, the financial steps taken toward a wonderful world ahead might just set the stage for a systematization of risks throughout the economic and financial system.



1. Still in the buildout phase

AI investment remains infrastructure-led. Building data centres requires vast amounts of land, equipment and energy. More importantly, it requires enormous amounts of capital. The figures are unusually unstable because the pricing of key components evolves, and demand grows exponentially, leading to frequent upward revisions of capital spending estimates.

In early October 2026, Goldman Sachs estimated that AI-related capital spending by hyperscalers will go from \$800 billion in 2026 to \$1.2 trillion in 2027 and \$1.4 trillion in 2028.

If we add spending by the rest of the ecosystem, it's easy to conclude that the buildout has reached such a scale that it will have an outsized impact on the global economy.


The evidence still looks like a buildout phase. Capabilities are improving, adoption is rising and investment is booming, but broad labour displacement and economy-wide productivity gains have not yet arrived. That timing gap matters, because it means today's infrastructure is being financed partly against benefits that are plausible but not yet fully visible in aggregate data.

EXHIBIT 1 | The buildout from the macro perspective

Private investment ex-housing is surging in the United States, particularly in the tech sector. It is becoming an increasingly important share of GDP, comparable to other large infrastructure investment cycles. Despite this buildout, corporate balance sheets don't appear to be under stress at the macro level. In fact, liabilities as a share of GDP are edging lower while profits are surging.

Surging profits deserve special attention. Are they artificial because of the circular investment nature of the AI buildout? Are they an early indicator of improving productivity that, for now, accrues to shareholders? Or are they simply the result of a historically low household savings rate and a historically large budget deficit? The answer matters, because it could determine whether the cycle is on a sustainable footing or is already living on borrowed time.


United States: Nominal Non-Residential Private Investment as a % of GDP
U.S. Bureau of Economic Analysis (BEA), as at 2026 Q2




United States: Corporate Liabilities as a Share of GDP
Macrobond, as at 2026 Q2



United States: Nominal Tech Investment as a % of GDP
U.S. Bureau of Economic Analysis (BEA), as at 2026 Q2



United States: Corporate Profit as a % of GDP
Macrobond, as at 2026 Q2





2. From cash flow to debt: Is capital a binding constraint?

The largest hyperscalers entered this cycle in excellent shape. In fact, they may have prepared for it. They didn't know exactly when the next big technological boom would come, but they knew it would come eventually – AI research was promising before the large language model came into public consciousness circa 2023 – so it made a lot of sense to keep their balance sheets pristine to be ready for it.

And ready they were. Alphabet, Amazon, Meta and Microsoft had strong cash generation, low leverage and ample liquidity. Early AI spending could be funded by operating cash flow, cash reserves, lower share repurchases and modest equity issuance. But that model is reaching its practical limit.

In fact, in 2026 operating cash flow from hyperscalers is about the same as capital spending – both \$800 billion. Because a dollar used for a data centre cannot also fund an acquisition, a dividend, a buyback or another investment, debt financing is becoming vital.

The change is already visible. An analysis by Bloomberg Intelligence shows AI-related issuance of corporate bonds reached an all-time high in 2025, approaching \$300 billion. At the third-quarter mark of 2026, tech corporate bond issuance had already surpassed \$400 billion. Given the trajectory of capex spending, records should be shattered again in 2027, and probably in 2028, if the broader macro environment remains constructive.

So basically, AI has become a recognizable credit sector. Its index weight depends on how the sector is defined. A narrow measure counts only large technology issuers, but a broader measure adding data-centre operators, utilities, chipmakers, neoclouds and project bonds provides a range from about 10% to more than 16% of the U.S. corporate index by mid-2026.

As for the supply of credit, capital is certainly not binding yet; but "capital is available" isn't the same as "capital will remain cheap."

For now, institutional demand is strong. Investment-grade technology debt offers scale, liquidity and familiar issuers, project bonds offer duration, data-centre asset-backed securities offer a rated spread, and private credit offers a complexity premium.

Demand absorbed extraordinary issuance in the first half of 2026, and markets remained functional.

But even though public markets are deep, absorption has a price.


Heavy supply can widen sector spreads, raise borrowing costs elsewhere and consume investor balance-sheet capacity. Private-market dry powder is also substantial, but it isn't cash sitting in a vault. It is committed capital subject to mandates, concentration limits, return hurdles, draw schedules and investor confidence. One dollar cannot fund every strategy at once.

The broader constraint comes from governments competing for the same dollars. Governments are borrowing heavily to fund defence, energy security and persistent deficits – at a time when central banks are no longer the automatic buyers they were throughout the 2010s. Rising interest rates are natural in this environment.

So, can global markets finance the AI buildout? Probably, yes. But the more useful credit question is harder: What cash flow repays each layer of the capital stack?

EXHIBIT 2 | The competition for duration

Gross bond issuance has been rising in the United States in recent years. Bond issuance by hyperscalers has been growing the fastest, from \$14 billion in 2024, to \$68 billion in 2025 and \$132 billion so far in 2026 yet remains small relative to other types of debt. Meanwhile, gross Treasury issuance and corporate investment-grade issuance are closing in on last year's number after only three quarters. Both are set to reach record issuance if the year-to-date trend continues.





3. Don't lose the trees for the forest

In the world of finance, details matter quite a bit. And excitement around high-velocity themes often leads to losing them over the big picture.

A Microsoft bond is a claim on a diversified global business. A project bond is a claim on a lease. A GPU loan is a claim on customer contracts and rapidly depreciating equipment. A private fund interest is a claim on a portfolio whose marks may move slowly. A utility investment may be supported by a dedicated customer, or by a wider rate base. These exposures are not the same.

Investors in the infrastructure sector are familiar with those distinctions: Upfront costs are large, assets take years to build, and contracts can support long-dated debt. But AI adds a technology-cycle problem: A pipeline doesn't become obsolete because a better pipeline launches the following year, as a CPU or a GPU could.

When we venture into the forest, we find that the current system financing the AI buildout is quite complicated and layered. While its layers are not entirely novel, their number – we have identified seven – point to a complex, multifaceted structure mixing and spreading economic, technological and financial risks throughout the financial system. And we've seen that movie before; a complex systemic system can have outsized impacts on the real and the financial worlds.

Layer 0: Final demand, the source of repayment

Layer 0 is simply the economic engine driving a business.

It consists of payments made to suppliers of services by end users, and the money runs up the chain to the application provider, the model developer, the cloud or neocloud, the data-centre operator, the equipment and power providers, and finally the creditors. Each link gets paid along the way.

There is already an inherent risk at this ground-level layer because it still isn't clear where the money will flow (our previous [paper](#) attempts to answer this question).

A model provider can report fast revenue growth while its infrastructure partners earn weak returns.

This distinction is crucial: AI can transform the economy but still destroy capital for some investors. Railways changed commerce, yet many railway securities failed. Fibre networks became essential, yet many telecom investors lost money. Social value and investor return are not the same thing.

A durable monitoring framework separates four variables: AI usage, customer spending, provider revenue and cash available for debt service after maintenance capital spending. They will not move one for one.

Layer 1: Hyperscaler balance sheets

This is the cleanest layer.

Microsoft, Alphabet, Amazon and Meta have diversified revenue, substantial cash generation and deep market access. They can absorb a weak return on a specific campus without threatening debt service. For bondholders, the main risk is deterioration rather than default: Spreads can widen as more demand for financing arises, ratings can drift lower as free cash flow weakens and investors may absorb leverage and duration risk without sharing fully in the equity upside.

As of this writing, Oracle is the outlier and the first real test case for the new financing model.

It entered 2026 at BBB with a negative S&P outlook, below the larger cloud providers.

Oracle had \$55.7 billion of capital expenditures in fiscal 2026 against \$32.0 billion of operating cash flow. It issued \$42.7 billion of senior notes and another \$5 billion of mandatory convertible preferred stock. Its cloud growth is genuine, as is its reliance on external financing.

Calling Oracle the weak link would be too simple. It is better understood as the first test of whether an investment-grade issuer with rapidly growing demand, negative free cash flow, large customer commitments and a smaller legacy cash-flow cushion can finance the buildout without losing market confidence.



EXHIBIT 3 | The credit hierarchy

Corporate spreads remain tight across the board, and hyperscalers have usually benefited from some of the tightest spreads in the market. That is changing, however, with recent large issuance. Hyperscalers' credit spreads are now in line with investment-grade bonds but still well below high-yield bonds.

Corporate Spreads: Investment Grade, High Yield and AI
Bloomberg, as at 9/29/2026



Layer 2: Credit enhancement and risk transfer

This is where things start getting a bit more complicated. Credit enhancement lets a weaker borrower borrow on a stronger company's credit.

A developer building an AI campus may have a modest balance sheet and only one or two customers. On its own, it would find long-term debt expensive, if it could raise any. A contract with a stronger party changes that, whether it's in the form of a long-term lease, a take-or-pay commitment, a financial backstop, a residual-value support, a minimum-revenue arrangement or a completion guarantee.

Each one attaches part of a hyperscaler's creditworthiness to someone else's debt.

Hut 8 is an example of how this works. The publicly traded energy and digital-infrastructure company develops and operates large-scale computing facilities. Its River Bend campus, under construction in Louisiana, is designed to provide the power, cooling and servers needed for AI and other high-performance computing workloads. A 15-year lease with AI cloud provider Fluidstack has a base contract value of \$7 billion, and Google provides a financial backstop for rent and certain obligations. That support helped a Hut 8 subsidiary issue \$3.25 billion of 16.5-year investment-grade senior secured notes. The debt is nonrecourse to Hut 8, but the economic anchor is Google.

In plain terms, bondholders are lending to a project company, not directly to Hut 8 or Google. They accept investment-grade terms because Google has agreed to support key payments. Credit enhancement moves risk toward hyperscalers, but it stays at arm's length. Google does not carry the Hut 8 notes on its balance sheet, but it has taken on a contingent obligation that could become real if the tenant falls short. So, while not direct, part of the exposure lands on Google anyway.

These tools are not limited to project vehicles. A lease or an offtake contract with an investment-grade tenant can also support a developer that borrows on its own balance sheet. What defines this layer is not the legal structure, but where the credit comes from: the stronger party's promise.

Box A: Nvidia, the central bank of AI

Nvidia, while not a hyperscaler, has been an enthusiastic proponent of credit enhancement. In fact, it increasingly acts like the central bank of the AI investment cycle. It supplies the chips at the centre of the system, but it also uses its balance sheet to support the neoclouds that buy them.

Through equity investments and long-term commitments to purchase cloud capacity, Nvidia gives these younger and more highly leveraged operators more predictable revenue. That strengthens their credit profile and makes lenders more willing to finance their data centres and equipment.

CoreWeave illustrates the mechanism. Nvidia invested \$2 billion in the company and committed to buy up to \$6.3 billion of residual cloud capacity through 2032. Such commitments

reduce the risk that newly built servers will sit unused. They also create future cash flows that can support more borrowing, more data-centre construction and more purchases of Nvidia chips. Nvidia, therefore, helps keep credit flowing through the ecosystem while reinforcing demand for its own products.

The role extends beyond individual borrowers. Nvidia has also partnered with major asset managers and banks to mobilize third-party capital for AI infrastructure. The central bank analogy has limits, however. Nvidia is a commercial company, not a public institution: Its support can be reduced or repriced, and its financial strength depends on the same AI cycle it is helping to sustain. Credit enhancement can lower the risk of a particular borrower, but it does not remove risk from the system.



Layer 3: Data-centre project finance

Project finance adds further complexity; credit moves toward data-centre projects, and away from their sponsors.

In project finance, a separate company is set up to run a single project, such as a power plant, a pipeline, an airport or a toll road. This company owns or builds the facility and raises its own equity and debt. Lenders are repaid mainly from the project's contracts and assets, and they have little or no claim on the parent company if things go wrong.

So, project finance explicitly keeps the project separate from its sponsor's balance sheet.

Meta's Hyperion campus, a massive AI data-centre project under construction in Louisiana, is a clear example. Instead of paying for and owning the whole campus, Meta placed it in a joint venture.

Funds managed by Blue Owl own 80%, while Meta owns 20% and continues to manage construction and the property. Blue Owl and its investors provide most of the outside capital, and the project vehicle raises debt secured by the campus and the payments tied to it. Once the facilities are ready, Meta leases them from the joint venture. The lease payments are the recurring cash flow that covers operations and pays interest and principal to lenders.


Meta trades a large upfront construction bill and direct responsibility for the debt for a 20% stake, future lease payments and specific guarantees. Blue Owl's funds and the bondholders take on more of the financing, construction and asset-value risk. But the risk has not disappeared; it has been split among Meta, the joint venture's investors and its lenders, each exposed through a different contract.

Such financial arrangements have surged in popularity. Five years ago, the value of future lease payments represented less than \$100 billion in off-balance commitments by hyperscalers. They now amount to more than \$1 trillion. This is perfectly rational – hyperscalers have all the incentive in the world to spread risk around.

EXHIBIT 4 | Off-balance-sheet commitments

The five hyperscalers increasingly rely on off-balance-sheet commitments, including future data-centre lease payments. Once a niche practice, these obligations are now economically and financially significant. Commitments from highly rated companies provide the cash-flow support needed to finance data-centre investment.

Hyperscalers Off-Balance Sheet Commitments - Uncommenced Lease
Alphabet, Amazon, Meta, Microsoft and Oracle SEC filings, as at 2026 Q2



If it is sensible at the micro level, there are many questions when we look at it from the macro perspective.

One such question regards the potential mismatch between the life cycle of the assets used as collateral and the maturity of the debt that is issued. This is one of the key differences from traditional infrastructure projects.

Technology can become obsolete faster. So far, older chips have remained remarkably useful despite the arrival of new generations, allaying some fears that the speed of depreciation of some assets is underestimated. But it might still become a problem. A building designed around one generation of computing equipment is weaker collateral than a toll road or a pipeline.

Most large AI deals combine layers 2 and 3. Hyperscalers use their financial muscle to enhance the credit of weaker partners and move risks toward data-centre projects, while keeping these commitments off their balance sheets. But when the same small group of tenants anchors many projects, lenders across all the vehicles are exposed to the same credit and the same demand assumptions. Ringfencing isolates each project from its sponsor, but all projects face the same underlying risks.



In other words, the risk is spread across many partners and investors, but it does not disappear.

Layer 4: Asset-backed AI financing

This is the cycle's most important financial innovation, and the layer most likely to confuse investors. This is the layer of securitization of property and lease cash flows.

Data-centre asset-backed securities allow an owner to borrow against a group of operating facilities and the rent paid by their tenants. The buildings, equipment, leases and reserve accounts are placed in a separate legal entity, which issues bonds to investors. Tenant payments flow into that entity and are used first to operate the facilities, then to pay interest and repay the bonds.

For example, a 2026 Stack Infrastructure transaction issued \$695 million of A-rated notes backed by 11 U.S. data centres and their lease income. This structure gives the owner cash upfront, while investors are repaid over time from tenant rent. For investors, the key question is simple: Will the facilities remain occupied and generate enough cash? Losses can arise if a major tenant leaves, leases aren't renewed, a site suffers downtime, costly upgrades are needed or the debt cannot be refinanced on favourable terms.

GPU financing looks more like aircraft leasing or equipment finance. A special-purpose borrower buys the equipment. Loans are advanced against eligible assets. Customer contracts generate cash flow, and the hardware secures the debt. Advance rates, amortization, maintenance, insurance and residual-value assumptions determine lender protection.

The aircraft analogy is helpful but incomplete. As discussed above, questions persist about the economic life of GPUs, even if so far it hasn't been a problem. As is always the case with financial innovation, promises come along with potential pitfalls.

These financial innovations help spread risk across investors, and in turn credit can flow to keep the AI cycle going. The flipside is a lack of transparency, as it becomes less clear where risk lives, who is exposed and by how much.

Layer 5: Utilities, governments and the public balance sheet

A data centre cannot operate without power, water and roads.

Those assets often sit outside the project vehicle. Utilities may build generating capacity and grid connections, municipalities may issue industrial-revenue bonds, states and provinces may offer tax incentives, and public agencies may fund roads, water or workforce programs.

So, the credit question is not only whether the data centre can pay its own debt. It is also who pays if the expected power or water doesn't arrive.

Power may turn out to be the truly scarce asset. Rapid load growth can push up local electricity prices when generation and transmission fail to keep pace. Power-market design, interconnection rules and connection timing can decide a project's economics. And the public side can hand out more than money; it also grants permits, which can be withheld.

Oracle's Project Jupiter in New Mexico shows how that risk travels back into private finance. The 1,400-acre campus, part of the Stargate initiative, is designed to handle 2.45 gigawatts. Oracle is the main tenant, and the developer is Stack Infrastructure, a unit of Blue Owl. According to media reports, as of September 2026, the project had faced delays affecting a planned natural-gas pipeline, denial of a permit tied to its energy plans and legal challenges over water and air-quality permits.

We then heard that Oracle had sent the developer a force majeure notice. According to initial reports, Oracle isn't trying to walk away as a tenant but is rather trying to delay its payments if the campus fails to come online in 2028 as planned. Oracle says the project remains on schedule, that such notices are common in developments of this size and that it is rethinking the project's power plan. Blue Owl says the notice doesn't change the financial commitments. Whether the move would actually release Oracle from any obligation is unclear.

The notice matters because of what it signals. In layer 2, a strong tenant's lease turned a project into financeable debt.



Jupiter shows that the promise is only as firm as its exceptions. If power fails to arrive for reasons outside the tenant's control, force majeure may let the tenant pay later, or less. The cost of a delay caused by public infrastructure could then shift from the tenant to the developer's equity and the project's lenders.

Politics around data centres, which are unpopular in the United States, are worth watching. When permits, pipelines and water rights depend on public decisions, political risk becomes credit risk.

The lesson of this layer is that the public balance sheet doesn't just fund what data centres need. Its permits and approvals decide whether the private contracts built on top of it pay on time.

Layer 6: Financing the financiers

The final layer determines whether a project problem stays local.

Private capital sits at the centre of the financing ecosystem. Blue Owl, Apollo, Blackstone, Brookfield, KKR, Pimco and other managers provide equity, construction debt, private placements and long-term asset finance. But private doesn't mean isolated.

Private-credit funds are themselves funded by life insurers, pension plans, sovereign wealth funds, banks, business development companies, semi-liquid retail funds, mutual funds and foreign institutions. Banks may finance the project, lend to the fund, provide the hedge and hold part of the securitization.

Meanwhile, some off-balance-sheet AI structures can be described as "shadow borrowing." Debt sits in project vehicles, but the ecosystem remains connected to hyperscalers through leases and guarantees, to private funds through ownership and to banks through funding lines. Stress can travel through refinancing pressure, falling private-credit appetite, guarantee activation or forced balance-sheet support.

The upshot is that it is hard to find corners of international finance that are not exposed in one way or another to the AI buildout. This does not make the system fragile by definition. Long-dated institutional capital in particular is well suited to infrastructure such as data centres.

Closed-end funds do not promise daily liquidity; insurance liabilities can be stable and fully amortizing project debt can avoid a maturity wall.

Beyond the obvious risks of defaults reverberating throughout the financial system, danger could appear when asset liquidity and investor liquidity don't match. A private loan can be fundamentally sound yet difficult to sell. If a semi-liquid fund promises periodic redemptions, liquidity pressure may arrive before credit losses do.

Credit agencies have an important role to play. By being watchdogs of credit risks, they could monitor shadow borrowing, increase transparency and allow financial institutions to build prudent exposure to the AI buildout.

The good news is that they remain constructive on the theme, and most of their rating actions have been positive. The more concerning news is that rating agencies significantly underestimated the risk embedded in mortgage-related securities in the run-up to the Great Financial Crisis, partly because of the complex structure of the financing ecosystem. Let's hope lessons have been learned.

4. So, where are the risks?

This conversation brings us to the inevitable question: In this tangled web of demand for and supply of capital, where do the actual risks reside? It's one thing to equate "AI will change the world" with "AI credit is a safe bet," but that is losing the trees for looking at the forest.

As we learned in the credit buildup preceding the 2008 financial crisis, risk doesn't simply move in one direction, but rather through contracts. In the case of AI, the first loss may sit with a project company, but the final loss can migrate through guarantees, leases, fund leverage, utility tariffs or reputational support.

As with any other infrastructure project, a lot of things can go wrong with the AI buildout, each with a potentially material impact on different layers of the financing system. While we don't pretend to be exhaustive, the list below does provide an efficient visual on how risks could make their way through the multiple layers depending on the shock.



Shock	First balance sheet affected	Where risk may migrate
Construction delay	Project vehicle and developer	Equity sponsor, completion guarantor, lender, anchor tenant
Power connection delay	Project vehicle	Utility, developer, equipment owner, tenant under delay clauses
Lower AI utilization	Model provider or cloud operator	Neocloud, facility landlord, equipment lender, equity sponsor
Price compression	Cloud and model providers	Neocloud margins, project coverage, GPU residual values
Faster chip obsolescence	Equipment owner	GPU lender, lessor, guarantor, private fund
Tenant default	Facility vehicle	Lease guarantor, securitization noteholders, insurer
Higher long-term rates	Refinancing borrower	Equity holder, property valuation, fund marks
Private-fund redemptions	Fund manager	Secondary market, bank facility, remaining investors
Hyperscaler capex cut	Entire supply chain	Data centres, utilities, equipment makers, local governments

What to watch: a checklist for the prudent investor

After that conversation, what should investors track as canaries in the coal mine as they participate in the financing efforts of the technological revolution of our time?

The best indicators sit below the headline capex number. No single measure is decisive, but together they show whether leverage is supporting productive infrastructure or compensating for weak cash conversion.

- **Cash conversion:** Is hyperscaler operating cash flow keeping pace with capital spending, and how much extra yield must issuers offer to sell new debt?
- **Hidden commitments:** Look at leases signed but not yet started, guarantees, backstops, residual-value support and capacity commitments – the obligations that sit outside reported debt.
- **Delivery:** Track contracted, connected and operating megawatts separately. Watch for delay or force majeure notices, and check whether utility grid costs are charged to data-centre customers or passed on to the wider rate base.

- **Revenue quality:** Watch revenue per unit of compute rather than unit growth alone, customer concentration at neoclouds and project vehicles, and project debt-service coverage once construction and acceptance tests are complete.
- **Collateral value:** Follow GPU resale prices, utilization and how long the chips actually stay useful compared with the depreciation assumptions.
- **Liquidity in private credit:** Monitor redemption queues and secondary-market discounts at private funds, plus bank lending to private-credit funds and project vehicles.

Conclusion: Look more closely

The good news is that the AI buildout can be financed. The largest sponsors remain strong, investor demand is deep, the infrastructure is real, and long-term contracts can match long-term capital. The economic gains from AI may ultimately justify today's spending.

But the financing system now reaches far beyond hyperscaler bonds. It includes project companies, private funds, securitizations, equipment loans, insurers, banks, utilities and public balance sheets. That makes the system more flexible but harder to see.

Thus, investors run the risk of misidentifying what they own.



A highly rated bond may depend on one tenant, a diversified fund may depend on one technology cycle, a guaranteed project may rely on a guarantor whose earnings share the same shock and a long-lived building may house short-lived equipment.

That is the lesson from 2008 worth keeping: **Financial innovation can distribute risk efficiently, but it can also distribute risk so widely that no one sees the concentration.** AI does not need to disappoint for every lender to lose money. It only needs to monetize more slowly than the debt was structured to assume.

The forest still looks healthy. The time to inspect the financing architecture is before the market comes under pressure. The next phase will be decided borrower by borrower and contract by contract.

Welcome to the wonderful world of AI financing.

Sébastien Mc Mahon, MA, PRM, CFA
Chief Economist, iA Financial Group

Maxime Houde, CFA
Vice-President, Portfolio Manager, Thematic Investing

Étienne Bergeron, M.Sc
Associate Director, Macro Strategy

iA Global Asset Management™ is one of Canada's largest asset managers, overseeing, as of June 30, 2026, \$156 billion across a range of institutional and retail mandates. We help investors achieve their long-term wealth creation goals through innovative investment solutions designed for today's complex markets. We are building upon our historic success, supporting the growth of our core strengths, and exploring innovative ways to meet investor needs. We are rooted in history and innovating for the future. Our experienced portfolio managers use a proprietary investment methodology, rooted in iAGAM's unifying commitment to strong risk management, analytical rigour, and a disciplined, process-driven approach to asset allocation and security selection.

General Disclosures The information and opinions contained in this report were prepared by iA Global Asset Management ("iAGAM"). The opinions, estimates and projections contained in this report are those of iAGAM as of the date of this report and are subject to change without notice. iAGAM endeavours to ensure that the contents have been compiled or derived from sources that we believe to be reliable and contain information and opinions that are accurate and complete. However, iAGAM makes no representations or warranty, express or implied, in respect thereof, takes no responsibility for any errors and omissions contained herein and accepts no liability whatsoever for any loss arising from any use of, or reliance on, this report or its contents. There is no representation, warranty, or other assurance that any projections contained in this report will be realized. The pro forma and estimated financial information contained in this report, if any, is based on certain assumptions and analysis of information available at the time that this information was prepared, which assumptions and analysis may or may not be correct. This report is not to be construed as an offer or solicitation to buy or sell any security. The reader should not rely solely on this report in evaluating whether or not to buy or sell securities of the subject company. The reader should consider whether it is suitable for your particular circumstances and talk to your financial advisor. iA Global Asset Management, iAGAM and the iA Global Asset Management logo are trademarks of Industrial Alliance Insurance and Financial Services Inc. and are used under license. iA Global Asset Management and iAGAM are tradenames under which Industrial Alliance Investment Management Inc. and its subsidiary iA Global Asset Management Inc. operate.