Agincourt Capital Management, LLC

FEBRUARY 2026

INVESTMENT UPDATE

Every so often, an investment cycle emerges that fundamentally reshapes markets. The AI investment cycle appears to be one of those moments, requiring a massive commitment of capital to fund the compute hardware, power, and infrastructure needed to support AI data centers. These investments will ripple through capital markets over the coming years, with meaningful implications for fixed income investors. This *Investment Update*—the first in a two-part series—provides an overview of the AI ecosystem, the scale of the emerging AI capex cycle, and its implications for the broader bond market. Next month's update will examine the impact in greater detail across key fixed income subsegments.

Estimates for global AI infrastructure investment range from more than $5T (trillion) through 2030, according to J.P. Morgan, to upwards of $7T, based on estimates from McKinsey. At the midpoint, this level of spending equates to roughly 3% of U.S. GDP annually. As the chart on this page shows, J.P. Morgan projects ~$1.5T to be funded through operating cash flow, with the remaining ~$4T requiring external capital. Traditional fixed income markets are projected to absorb half that amount. The scale of this financing requirement raises critical questions around funding sources, market capacity, and long-term returns.

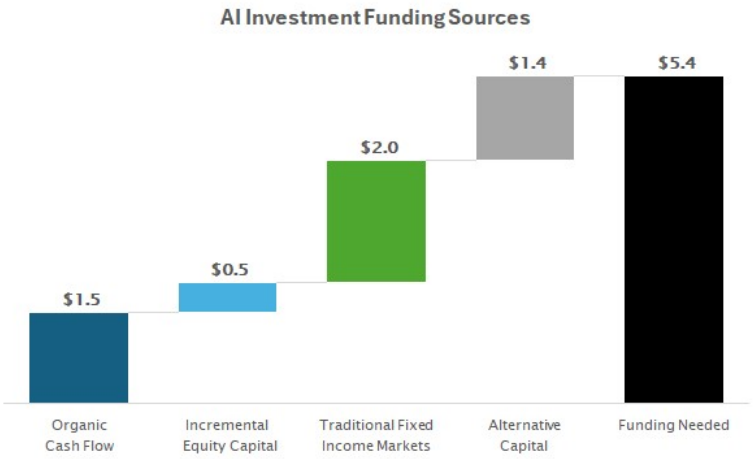

Before we start to unpack these questions and the implications for bond markets, we will first provide an overview of the data center and AI ecosystem. A data center is a large building (or group of buildings) purposely built to house compute hardware (microprocessors, servers, networking) and infrastructure (racks, connections, cooling). The compute hardware and infrastructure within the data center is used to develop AI systems (large language models, agents) that underpin AI applications (ChatGPT, Gemini), or to provide traditional non-AI functions, such as web hosting, data storage, and cloud-based software delivery.

The AI ecosystem spans multiple layers, beginning with compute providers and extending through physical infrastructure, power generation, and AI model development. At the top of the ecosystem are the hyperscalers—Microsoft, Amazon, Alphabet, Meta, and Oracle—which supply a large share of global cloud-based compute capacity. This capacity is either rented to third parties or consumed internally within vertically integrated technology stacks. Alongside the hyperscalers are neoclouds, such as CoreWeave, which are AI-native providers offering specialized compute to third parties. While neoclouds compete with hyperscalers for AI workloads, they also count the hyperscalers as core customers and financial backers. A key distinction between the two groups is their financial profiles: hyperscalers generate substantial free cash flow from diversified, non-AI businesses, while neoclouds generally operate with negative free cash flow due to significant upfront investment requirements against relatively modest revenues and profits.

Supporting this compute layer is the physical infrastructure required to house and power AI workloads. Data center developers and operators build and own the facilities that house the compute infrastructure, provide power and fiber connectivity, and operate the non-compute infrastructure. These facilities are leased to hyperscalers, neoclouds, and enterprise customers.

Power providers—including regulated utilities and independent power producers—are also critical participants, as AI data centers are highly energy-intensive with power availability representing one of the key bottlenecks to continued growth. Estimates suggest that meeting projected data center power demand will require a material expansion of generation capacity over the coming years, adding cost, complexity, and long lead times to the AI build out.

At the application layer, foundational AI model developers and application providers—such as OpenAI, Anthropic, and certain hyperscalers—use leased compute capacity to train and deploy large language models and AI agents. These models underpin a rapidly expanding range of consumer and enterprise applications, driving demand for additional compute and infrastructure.

Enabling this entire ecosystem are technology hardware suppliers, including microprocessor designers such as Nvidia, server and networking manufacturers, and providers of electrical and industrial equipment. Together, these participants form a tightly linked ecosystem in which capital investment, demand growth, and financing needs reinforce one another.

This degree of interconnectedness has raised questions about the sustainability of the investment cycle and whether current dynamics reflect durable fundamentals or speculative excess (a bubble). The graphic on this page is a good depiction of the overlap between suppliers, customers, and financing. While the unconventional nature of this financing and supply chain adds risks to the ecosystem, it isn't an indication of a bubble in and of itself. Rather, bubbles are typically associated with excessive optimism, aggressive financing structures, and rapid increases in market prices that can't be justified by the underlying fundamentals and prospective financial returns.

Perhaps the more important question is, do the underlying fundamentals and prospective financial returns justify the investment cycle? Demand for AI infrastructure and compute power is exploding and far exceeds available and under-construction capacity. At the same time, supply is being severely limited by power availability. As such, supply-demand appears to be supportive of investment.

Questions about financial returns are uncertain. J.P. Morgan estimates that ~$650B (billion) in annual revenue would be needed to achieve an internal rate of return of 10%, based on its >$5T capex projection through 2030. To put that in context, OpenAI expected to end 2025 at an annualized revenue run-rate of ~$20B, up from ~$30MM (million) in 2022. Revenue should scale rapidly as OpenAI adds paying subscribers and likely raises prices. Similar dynamics should play out across the ecosystem. Revenues and prospective financial returns are trending in the right direction. The question going forward is: will the growth curve extend out long enough to generate enough revenue (or cost savings) to sufficiently compensate for the inherent risks?

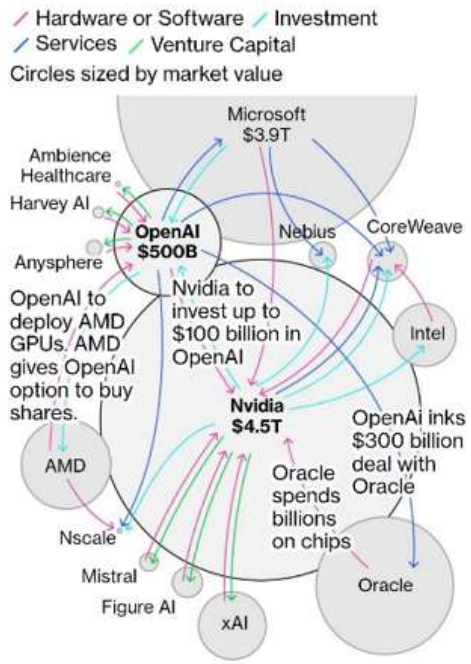

Source: Bloomberg News

The answer to that question depends, in part, on who funds the capex cycle and how it is funded. The short answer is everyone involved in the ecosystem will help fund the cycle, with the largest financers being the companies that are buying the compute infrastructure and those that are building the physical data centers. Estimates vary depending on the intended use of a data center; however, a leading-edge, 1-gigawatt (GW) AI data center can cost $30-60B, with the compute and networking infrastructure accounting for >50%, followed by power, non-compute infrastructure, land, and construction costs.

Prior to 2025, most funding for data centers was from internal operating cash flow, incremental equity investments, and a modest amount of securitized debt. The securitized market has played a growing role in financing, with issuance from data center developers and operators doubling from ~$10B in 2021 to ~$20B in 2025. More is on the way, with J.P. Morgan projecting $300B of data center securitized issuance over the 2025-2030 period.

Until recently, the high-grade corporate bond market hasn't played an explicit role in financing data center build outs. That changed in fall 2025 when the floodgates opened and hyperscalers brought ~$115B in new debt over a two-month period. That was just the tip of the iceberg. As of February 2026, ~$50B has already been issued at the hyperscaler level, with more expected throughout the year.

Given this backdrop, we can conclude that the staggering capex needs of the AI investment cycle will require most subsegments of the fixed income markets to absorb massive amounts of new debt over the next several years. The high-grade corporate bond market will continue to lead the way with the hyperscalers being the largest AI issuers. Similarly, utilities should see rising issuance needs as power generation and grid upgrade capex grows.

Elsewhere, the leveraged finance markets will see growing issuance from neoclouds and foundational model and application developers. Data center developers and operators will also utilize leveraged finance markets while continuing to lean into securitized markets to diversify funding sources.

The bubble question remains an uncertainty, and it should be acknowledged that, in real time, the line between sustainable, strong growth and an unsustainable bubble is blurry. The AI investment cycle certainly has some characteristics of a bubble: circularity, some questionable financing structures, excessive optimism, and rapid increases in market prices. But it also contains characteristics that suggest it is not a bubble: demand far outpacing supply, strong profitability and cash flows at the hyperscaler level, revenue growth rates that suggest a long demand curve, and the real value of the underlying AI compute infrastructure and the land and buildings themselves.

Considering all of this, it is no surprise that debt and equity markets for AI-linked companies have exhibited volatility with fits and starts of optimism and fear. A dose of caution and patience is warranted. After all, this is only the beginning, particularly for bond markets, and more supply is on the way. Next month we will dig deeper into the implications for the key subsegments of the bond market. Stay tuned!