



THE HOTTEST IPO OF 2025

The Hottest IPO of 2025

By Jeff Brown, Founder & Chief Investment Analyst, Brownstone Research

Just last month, whispers surfaced about one of my favorite private companies – an industry-leading data center company – preparing to go public in 2025.

It's not a household name... yet. But it's quietly become indispensable to the next wave of breakthroughs in artificial intelligence (AI), machine learning (ML), and 3D rendering.

This isn't your typical cloud service provider. It's a company focused on delivering cutting-edge GPU infrastructure designed to handle the world's most demanding computational workloads, specifically designed for machine learning and artificial intelligence.

Wall Street is already buzzing with speculation. Some believe this company could become the next Amazon Web Services, potentially becoming a juggernaut in the fast-growing AI cloud computing space.

But what makes this company so special? Why are the world's top AI firms clamoring for access to its state-of-the-art data centers? And why are investors gearing up to buy shares of this business when it finally hits the public markets?

To understand why this company is positioned to dominate, we first need to explore the driving forces behind the multitrillion-dollar AI

hyperscale data center buildout. This buildout will be one of the biggest infrastructure buildouts of our generation. And this company is at the heart of it all.

Demand for Compute Is Insatiable

To grasp why the world's appetite for computational power is seemingly limitless, we need to examine the demands of large language models (LLMs) and the race toward artificial general intelligence (AGI).

AGI refers to a level of artificial intelligence that can understand, learn, and apply knowledge across a wide range of tasks, much like a human.

Unlike narrow AI, which is designed for specific tasks (like playing chess or recommending movies), AGI will possess the cognitive abilities to solve any problem and perform any task that a human can, at an expert level regardless of the subject matter.

This means AGI would not be limited to pre-defined functions but would have the capacity to reason, plan, learn, and communicate effectively.

AGI models will help revolutionize many industries. It will accelerate scientific discoveries by processing and analyzing vast amounts of data. Amounts far beyond human capacity.

It will lead to insights in fields like physics, biology, and material and environmental sciences. AGI will even be able to conduct self-directed research.

And those companies that succeed in creating AGI will have an advantage in winning the ultimate prize – developing artificial super intelligence (ASI). An ASI will not only be self-aware, its capacity will exceed that of all human intelligence in all fields of study.

This is why many of the world’s most complex problems will be solved. New drug discoveries will cure nearly every disease. We’ll be able to predict weather patterns far in advance. And we’ll have limitless clean energy in the form of nuclear fusion.


Trillions of dollars of monetary incentives are in place, which is why the largest tech companies in the world, as well as a handful of venture capital-backed startups, are racing to develop not just LLMs, but AGI.

Major players in the race include Microsoft (via its investment in OpenAI and its GPT models), Meta’s Llama, Google’s Gemini, xAI (currently used on X as Grok), and Apple’s soon-to-be-launched Apple Intelligence. Additionally, private companies like Anthropic and its Claude AI, Cohere, and Mistral AI are all leaders in this space.

The Race to AGI

November 30, 2022, was the day the AI industry changed forever. This is when the private startup, OpenAI, released its ChatGPT chatbot. In a few years, this day will be remembered

Time It Took to Reach 100 Million Users Worldwide



BROWNSTONE
RESEARCH

Source: PwC Analysis

as the inflection point in generative AI and a moment in history just as significant as the day Steve Jobs announced the iPhone in 2007.

ChatGPT was the first time most people could interact at a massive scale with a general-purpose AI without any special technical knowledge.

ChatGPT’s simplicity and utility resulted in ChatGPT becoming the fastest app to reach 100 million users, a feat accomplished in just two months, dramatically besting TikTok’s record at nine months. (See chart above.)

Adoption at this scale and this speed demonstrated the clear demand for a generative AI for the mass market.

Despite the technology not being perfect in 2023 and 2024, generative AI continued to improve month by month and found instant utility in students, professionals, consumers, researchers, programmers, and just about anyone with a question they need answered.

Adoption has been rapid. For example, we already see 90% of software developers using an LLM to help speed along their coding.

Marketers use it to help with marketing campaigns, artists use generative AI to conceptualize ideas, bloggers use it to write content, lawyers use it to both summarize legal agreements as well as write them... the list goes on and on.

LLMs improve the quality and time to completion of all these things. The tech is very good at predicting correct outcomes like what good software code looks like. Or profitable patterns in cryptocurrency markets.


Despite this adoption, many people deride ChatGPT as just a chatbot. A tool that answers only basic questions. And yes, that is how it started, but that's not what this technology is today... that first step was just a stepping stone to what's next.

In the chart above we can see how with each new iteration of OpenAI's GPT model, it gets a little smarter. When ChatGPT came out, it was built on the GPT-3 model. This model has the intelligence of someone in elementary school. That's why it could only answer basic questions.

But as the models get bigger and ingest more data and have more compute resources, they'll get smarter. GPT-4 was a step further. It can handle more complex tasks that take more steps to get to a solution.

Today, OpenAI's current LLM technology is known as GPT-4o, the "o" stands for "omni." GPT-4o is what's known as a multi-modal LLM which is capable of ingesting audio, images, video, text, and even software code. And it is already capable of performing many tasks better than the majority of humans, and able to do so in just seconds.

Base Scaleup of Effective Compute



GPT-5, or whatever it will be called, is expected to be released in early 2025 and will take another major step forward. I don't expect that it will be an AGI, but perhaps just one step away.

While not capable of self-directed research, it will likely be more capable than most humans at most tasks. The next version will demonstrate far more advanced reasoning capabilities, which will make the technology that much more useful.

And what comes after that will likely be AGI. I'm on record predicting at least one company will achieve AGI before the end of 2026.

And to get there we will need a lot more computational resources (compute). And the workhorses of this compute will be GPUs (graphics processing units). GPUs are used for AI applications because they allow the parallel processing needed to both train and run AI – meaning GPUs can run many separate calculations at once.

In contrast, serial processing is performed by central processing units (CPUs) enabling just one calculation before moving to the next. It's sequential.

While multi-core CPUs can run a limited number of parallel processes, it's not designed to handle thousands of tasks at once like GPUs.

The training and running of these foundational LLMs isn't possible on a desktop or laptop computer. It requires stringing together thousands, tens of thousands, and even hundreds of thousands of GPUs.

To do that, we need structures like the hyperscale data centers being built today.

The capital is readily available to build it, and the hyperscale data center infrastructure is being built as quickly as possible. Sometimes I've referred to these simply as "AI factories" – an appropriate name as they are designed precisely to train AIs and ultimately develop an artificial general intelligence.

Data Center Growth: The New Infrastructure Race

At the moment, America has more than 5,000 data centers. But here's the problem – almost all of them were built before the AI boom kicked off with the release of ChatGPT in November 2022.

These data centers simply weren't designed to handle the massive computational power and storage required to train and operate artificial intelligence (AI) applications.

That means the vast majority of today's infrastructure is obsolete when it comes to AI workloads. They don't have the size or power density to process the enormous volumes of data AI requires.

So, the industry is now racing to build entirely new facilities – data centers engineered specifically for AI. A report from Synergy Research Group highlights that at least 130 AI-focused hyperscale data centers are set to come online annually for the next decade.

For instance, Elon Musk's xAI has built a giant AI data center outside Memphis comprised of 100,000 NVIDIA GPUs. He calls it Colossus. And by early 2025 Colossus will have expanded its computational power to 200,000 GPUs.

Amazon spent \$10 billion on just one hyperscale data center in rural Mississippi.

Meta has pledged \$800 million to build just one AI data center in Indiana.

Microsoft and OpenAI are spending \$100 billion on a hyperscale data center they call "Stargate." I've seen it under construction. I've actually been there. And the scale is difficult to comprehend.

And that's just the start... you could say the AI data center buildout is like a space race. Each of these companies wants to dominate the potentially lucrative field of generative AI. And they're willing to pay any price to do it.

In 2024 alone, hyperscale AI data center spending will surpass \$1.4 trillion. And McKinsey projects that global demand for data center capacity will triple by 2030.

This is where my #1 IPO for 2025, CoreWeave, comes into play.

Its business model – renting out GPU power for AI and machine learning workloads – places it at the heart of this gold rush. The demand for compute is so overwhelming that even the largest companies with the deepest pockets are outsourcing capacity to CoreWeave.

On a recent earnings call, Microsoft CEO Satya Nadella admitted they simply can't get enough GPU servers online to meet customer demand. And in November, Microsoft expanded its contract with CoreWeave – buying close to \$10 billion worth of compute resources from CoreWeave by the end of the decade.

This isn't a short-term trend. As the AI revolution accelerates, the demand for CoreWeave's high-performance GPU cloud services will only increase.

CoreWeave Is Built for the AI Era

CoreWeave has carved out a unique niche in cloud computing. Unlike traditional cloud providers, CoreWeave focuses on cutting-edge GPU infrastructure specifically for AI and machine learning. That's all. And they can tailor the servers to the specific needs of their clients.

The company came from humble beginnings. Founded in 2017 under the name Atlantic Crypto, CoreWeave began as a cryptocurrency mining venture. Its three co-founders – Michael Intrator, Brian Venturo, and Brannin McBee – had no prior experience in the cloud computing sector.

In fact, all three were trading energy futures at their hedge fund when they discovered an emerging opportunity: leasing GPUs.

During the crypto bull market of 2017, they realized that a single GPU could pay for itself in mere days by mining Ethereum. Starting in Venturo's grandfather's New Jersey garage, they began scaling operations.

But the cryptocurrency crash of 2018 forced a pivot. They raised \$1.2 million in seed funding, diversified their focus, and rebranded as CoreWeave, shifting to cloud infrastructure for high-performance computing (HPC).

Fast-forward to today: CoreWeave has raised \$12 billion in equity and \$7.5 billion in debt financing, enabling it to operate 28 data centers by the end of 2024. Its focus on GPU resources for high-performance computing (HPC) has made it indispensable to companies with the most demanding computational needs.

By specializing in GPU-focused infrastructure, CoreWeave sets itself apart from legacy cloud providers like Amazon Web Services (AWS), Microsoft Azure, Oracle, IBM, and Google Cloud, which were built around more generalized CPU-based computing for more typical cloud software workloads.

While there's a huge need for these kinds of data centers, the growth in demand comes from the new hyperscale AI data centers built around GPUs.

And CoreWeave's purpose-built GPU servers deliver a performance-adjusted cost that the claim is up to 80% less expensive than large, generalized clouds.

The release of ChatGPT in November 2022 amplified the demand for CoreWeave's GPU-centric services, propelling the company to new heights. But what truly cements CoreWeave's competitive advantage is its deep partnership with Nvidia (NVDA), the undisputed leader in GPU technology.

The Nvidia Advantage

Nvidia dominates the AI GPU market, with over 90% of all AI GPUs being Nvidia-built. These GPUs are the most sought-after computing hardware in the world today, and CoreWeave has become a preferred partner of Nvidia's for cloud-based GPU solutions.

Their partnership goes back to July 2021 when Nvidia named CoreWeave, its first elite cloud services provider for compute. And shortly after that, in May 2023, Nvidia invested \$100 million in CoreWeave's Series B funding round.

Why would Nvidia favor a relatively small upstart like CoreWeave over giants like Microsoft or Amazon? The answer lies in strategy.

Microsoft, Amazon, and Google are all developing their own custom silicon to compete directly with Nvidia, vying for a bigger piece of the AI hardware market. These companies have the scale and ambition to eventually challenge Nvidia's leadership.

CoreWeave, by contrast, has no such competing agenda. Its sole focus is on providing the most efficient and optimized Nvidia GPU cloud solutions. For Nvidia, this makes CoreWeave a trusted ally – one that ensures Nvidia's latest and greatest GPUs reach the market quickly and efficiently.

This partnership gives CoreWeave a unique edge. In August 2023, CoreWeave became the first cloud provider to bring Nvidia's H200 Tensor Core GPUs to market. These GPUs are an upgrade to the highly coveted H100 series, which CoreWeave also received before many competitors.


Early access to Nvidia's cutting-edge GPUs amplifies CoreWeave's value proposition. It ensures that its clients – many of the world's top AI developers – have access to the best hardware available. This advantage keeps CoreWeave's customers coming back and is reflected in the company's rapidly growing revenue.

Valuation

As with any private company, piecing together financial details can feel like detective work. But based on available data, the trajectory of CoreWeave's growth is staggering.

We anticipate even more clarity when the company files its S-1 with the SEC ahead of

CoreWeave Revenue



its IPO. These filings often reveal key metrics and strategies that were previously unknown, offering invaluable insights for investors.

For now, here's what we do know: CoreWeave's revenue growth has been explosive.

The company is projecting its revenue to more than triple in 2025, climbing from \$2.4 billion to a jaw-dropping \$8 billion. This kind of growth at scale is rare and underscores the insatiable demand for cloud compute power, especially GPU-based solutions. Quite simply, CoreWeave can sell as much GPU-based computational power as it can build.

And this isn't just growth for growth's sake – it's highly profitable growth. Research from Sacra estimates that CoreWeave operates with an 85% gross profit margin despite its highly competitive pricing in cloud services. For comparison, Microsoft Cloud's gross margin is just over 70%.

High gross margins signal a lean, efficient business model that can translate into robust net profitability. While we'll need to see the S-1 filing to confirm operating expenses and bottom-line margins, such an impressive gross margin positions CoreWeave as one of the most profitable companies in its class.

CoreWeave's most recent valuation pegged the company at \$23 billion. For a business expected to generate \$8 billion in revenue next year, that seems like a bargain.

Given CoreWeave's high gross margins, I anticipate the company will IPO with a valuation between \$48 billion and \$64 billion, roughly 6–8x its 2025 projected sales. However, once it hits the public markets, I expect CoreWeave will quickly jump to trade at more than 10x sales, which would push its valuation well above \$100 billion by the end of 2025.

CoreWeave will be one of the single best pure plays on this explosion in artificial intelligence, which is why it will draw so much institutional capital.

This is an IPO to watch closely. If CoreWeave's valuation aligns with our expectations, it could represent an incredible, fast-moving investment opportunity for investors. The Brownstone Research team will be watching the filings closely and will notify subscribers when it's time to act.

To contact us, call toll free Domestic/International: 1-888-493-3156, Mon-Fri: 9am-5pm ET or email memberservices@brownstoneresearch.com.

© 2025 Brownstone Research, 1125 N Charles St, Baltimore, MD 21201. All rights reserved. Any reproduction, copying, or redistribution, in whole or in part, is prohibited without written permission from the publisher.

Information contained herein is obtained from sources believed to be reliable, but its accuracy cannot be guaranteed. It is not designed to meet your personal situation—we are not financial advisors nor do we give personalized advice. The opinions expressed herein are those of the publisher and are subject to change without notice. It may become outdated and there is no obligation to update any such information.

Recommendations in Brownstone Research publications should be made only after consulting with your advisor and only after reviewing the prospectus or financial statements of the company in question. You shouldn't make any decision based solely on what you read here.

Brownstone Research writers and publications do not take compensation in any form for covering those securities or commodities.

Brownstone Research expressly forbids its writers from owning or having an interest in any security that they recommend to their readers. Furthermore, all other employees and agents of Brownstone Research and its affiliate companies must wait 24 hours before following an initial recommendation published on the Internet, or 72 hours after a printed publication is mailed.