

THE AI BOTTLENECK SERIES: Waiting for the Watts

Artificial intelligence (AI) has emerged as one of the most powerful secular growth themes of our time, transforming how businesses operate and how people live and work. Yet AI is not a single technology. It is a complex ecosystem supported by a long and interconnected value chain, spanning semiconductors, networking equipment, data centers, power infrastructure, and software. As AI adoption accelerates, strain can emerge at key points along this supply chain, creating bottlenecks where demand outpaces available resources. While these constraints can create uncertainty, we believe they can also reveal compelling investment opportunities. In this series, we explore the most significant AI bottlenecks, the forces driving them, and the investment opportunities they may create.


Within the current AI ecosystem, electricity may be the most limiting constraint today. A large AI data center campus can require a gigawatt or more of continuous power, roughly equivalent to the aggregate consumption of a mid-sized American city. While that need may seem massive, the challenge is less about generating that power than connecting to it. An AI data center can typically be built in 18 to 24 months, but connecting a facility of that size to the public grid can take five to seven years, due to required engineering studies, regulatory approvals and physical upgrades.¹ In Texas alone, ERCOT, the state grid operator, is reviewing requests from large new customers, roughly 90% of them data centers, for about 474 gigawatts, which is more than five times the most power the state has ever used at once.²

Why Is This Happening?

The cause of today's bottleneck lies in the two decades leading up to the introduction of ChatGPT in 2022. As Figure 1 shows, U.S. electricity use grew steadily from 1975 to 2005, then stopped. More efficient lighting, appliances and factories, plus the movement of heavy industry overseas, kept electricity consumption essentially flat for nearly twenty years.

With electricity demand flat, utilities and independent power producers commissioned new plants only sparingly, and the manufacturers of turbines, transformers and other grid equipment calibrated their production to replacing aging assets rather than enlarging the fleet. The launch of ChatGPT changed the demand picture abruptly. Data centers consumed roughly 4% of U.S. electricity in 2023, and we estimate that share could reach approximately 30% by 2030, depending on the pace of AI adoption. We project total U.S. electricity demand growth could exceed 6% a year through 2040, a pace not seen since the postwar era. We recognize this is a bold projection, well above most official forecasts and a sharp break from the recent past. In our view, however, it reflects the scale of data center capacity already planned or under construction, and

Figure 1: U.S. Electricity Demand is Projected to Surge



Source: U.S. Energy Information Administration and Alger estimates for 2026 onward.

that growth is why the power supply chain is straining: the industry spent twenty years sized for stagnation and is now being asked to grow rapidly. Consider the equipment that new capacity requires: a gas turbine ordered today may not arrive until 2031, and a large transformer takes up to four years for delivery.³

The Solution: Bring Your Own Generation

As we wrote back in April in The AI Power Problem commentary, data center operators have responded by taking power into their own hands. Bring Your Own Generation (BYOG), the practice of installing gas turbines, engines or fuel cells on site to power the data center directly, has moved from a niche solution to a central strategy. Binding orders for on-site power generation equipment now total about 75 gigawatts, and OpenAI's flagship Stargate campus in Texas alone will run on more than 500 gas engines.⁴

A one-gigawatt campus can produce tens of billions of dollars in revenue each year. Multiply that by a five-to-seven-year wait for a grid connection, and the opportunity cost of delay can exceed \$100 billion.⁵ Paying a premium

for on-site power that can be running in under two years is a rational decision. BYOG also sidesteps a growing political problem. Concerns about higher household electricity bills, driven in part by the cost of grid upgrades to serve new data centers, have made data centers unpopular neighbors; in August, Texas paused approvals of new grid-connected data centers pending a statewide audit. A campus that brings its own power is not asking ratepayers to fund new transmission lines on its behalf.

Looking Ahead

For data center operators, we believe on-site power is a necessity rather than a choice, and the companies that supply it are booked years in advance, which gives them strong pricing power. In our view, select gas turbine and engine manufacturers, makers of the switchgear and transformers that connect generators to servers, and the specialized contractors that install them remain primary beneficiaries of this bottleneck. The grid will catch up eventually. Until it does, the companies that help data center operators bypass the wait for a grid connection will continue to command a premium.

¹ Teplin, C., Swisher, J., & Strowbridge, D. (November 2025). PJM's speed to power problem and how to fix it. RMI. Construction timeline is an Alger estimate..

² Office of the Governor of Texas. (August 2026). Letter to the Public Utility Commission of Texas and ERCOT.

³ GE Vernova Inc. (July 2026). Transformer lead times per Wood Mackenzie, as reported by POWER Magazine (January 2026).

⁴ SemiAnalysis. (September 2026).

⁵ JLL and Alger (September 2026). 2026 data center outlook. Assumes approximately \$20 million in annual cloud services revenue per megawatt of operational capacity.

The views expressed are the views of Fred Alger Management, LLC ("FAM") and its affiliates as of October 2026. These views are subject to change at any time and may not represent the views of all portfolio management teams. These views should not be interpreted as a guarantee of the future performance of the markets, any security or any funds managed by FAM. These views are not meant to provide investment advice and should not be considered a recommendation to purchase or sell securities.

Risk Disclosures: Investing in the stock market involves risks, including the potential loss of principal. Growth stocks may be more volatile than other stocks as their prices tend to be higher in relation to their companies' earnings and may be more sensitive to market, political, and economic developments. **Companies involved in, or exposed to, AI-related businesses may have limited product lines, markets, financial resources or personnel as they face intense competition and potentially rapid product obsolescence, and many depend significantly on retaining and growing their consumer base.** These companies may be substantially exposed to the market and business risks of other industries or sectors, and may be adversely affected by negative developments impacting those companies, industries or sectors, as well as by loss or impairment of intellectual property rights or misappropriation of their technology. Companies that utilize AI could face reputational harm, competitive harm, and legal liability, and/or an adverse effect on business operations as content, analyses, or recommendations that AI applications produce may be deficient, inaccurate, biased, misleading or incomplete, may lead to errors, and may be used in negligent or criminal ways. AI companies, especially smaller companies, tend to be more volatile than companies that do not rely heavily on technology.

Important Information for US Investors: This material must be accompanied by the most recent fund fact sheet(s) if used in connection with the sale of mutual fund and ETF shares. Fred Alger & Company, LLC serves as distributor of the Alger mutual funds and ETFs.

Alger pays compensation to third party marketers to sell various strategies to prospective investors.