A&M INFRA & CAPITAL PROJECTS BY ALVAREZ & MARSAL

Powering AI: Who Builds, Who Delivers, Who Leads

1. THE NEW ROI: TIME-TO-ELECTRON

As the world joined forces intending to decarbonize, artificial intelligence (AI) showed up with an insatiable hunger for the same energy we were so desperately trying to conserve.

> **AI is not just reshaping industries; it is prompting new approaches to development, risk management and planning in the energy infrastructure intended to meet AI's energy demand.**

Hyperscale data centers, fueled by AI's exponential growth, now demand levels of speed, scalability and reliability that traditional utility models simply cannot match. What was once a standard utility risk is now a front-line strategic imperative. Power delivery timelines are no longer measured in years; they are counted in months, and slower timelines risk leaving significant AI-driven value unrealized. For developers, investors and operators alike, the race to secure electrons has become as critical as the race to build and train the AI models. In response, many in the sector are turning to the speed and certainty offered by deploying known and available technologies with gas as the key energy in areas where it is abundant. This model, often underwritten by investment grade PPAs, still carries longer-term asset value risks. Today's accelerated solutions could become tomorrow's stranded liabilities.

The traditional metrics for evaluating long-term power generation projects such as cost per megawatt-hour, emissions profiles and long-term asset life are taking a back seat. Speed-to-power delivery has become a leading metric of strategic advantage. Projects that can deliver electrons within 12 to 24 months are considered bankable. Those that cannot are increasingly seen as less competitive, no matter how low-cost or efficient they may be.

In this new reality, time-to-electron defines strategic advantage, and energy infrastructure planning is adapting accordingly.

2. SOLVING FOR SPEED: WHY GAS IS BECOMING A DEFAULT CHOICE

This surge in power demand is attributed to the increasing computational requirements of AI applications and the expansion of cloud services. According to Berkley Lab, 2024 United States Data Center Energy Usage Report, data center energy consumption held steady at approximately 60 terawatt-hours (TWh) between 2014 and 2016. However, as the deployment of accelerated AI servers began gaining momentum in 2017, energy use began to climb. By 2018, data centers were consuming around 76 TWh, accounting for 1.9 percent of the total U.S. electricity demand. This upward trajectory continued, with U.S. data center energy consumption surging to 176 TWh in 2023, a number now representing 4.4 percent of the nation's total electricity demand.[1]

In terms of energy, short-term forecasts show a wide range as to how much is going to be consumed by data centers ranging from 325 TWh to 580 TWh. This translates to an installed capacity of 60 GW to 110 GW.
As data center energy demand accelerates, developers and operators face pressure to deliver reliable capacity quickly. While renewable and nuclear remain critical to long-term decarbonization, they face near-term constraints. Solar and wind projects are limited by permitting delays, transmission bottlenecks and intermittency

[1] Lawrence Berkeley National Laboratory, 2024 United States Data Center Energy Usage Report, Arman Shehabi et al., December 2024, https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf

challenges. Meanwhile, small modular nuclear reactors (SMRs) remain early-stage technologies with regulatory and commercialization timeline challenges. In contrast, natural gas offers a proven, dispatchable option that can be deployed quickly at scale, using established technologies and financing structures.

Assuming an average load of 450 TWh by 2028, met entirely by large, efficient, natural-gas-fired combined cycle gas turbines, this load will require approximately 2.5 trillion standard cubic feet (TCF) per year of natural gas (or 6.8 BCF per day). For context, in 2023, the U.S. electric power sector consumed approximately 12.93 TCF of natural gas, accounting for about 40 percent of the U.S. total natural gas consumption, which was at approximately 32.50 TCF for 2023.[2]

This growth in consumption is reshaping the data center landscape across the United States, with new developments emerging in both traditional hubs and frontier regions. The strategic placement of emerging data centers, particularly those focused on AI and hyperscale operations, is increasingly influenced by proximity to natural gas infrastructure. This trend reflects the need for reliable, dispatchable power to meet the substantial energy demands of these facilities. Emerging hyperscale datacenter locations include Texas, Louisiana, Ohio, Arizona, Nevada and Utah.

Location	Project
Abilene, TX	The Stargate Project, a $500-billion AI infrastructure initiative led by OpenAI, SoftBank and Oracle, is constructing 10 massive data centers in Abilene. This development aims to position Texas as a central hub for AI innovation.[3]
Richland Parish, LA	Meta is investing $10 billion in a hyperscale data center in Richland Parish, marking Louisiana's entry into the data center market. The facility is expected to be transformative for the state's economy.[4]
Phoenix, AZ	Phoenix is projected to experience a 553 percent increase in data center capacity, reaching 5,340 MW, driven by its affordable energy and proximity to major tech markets.[5]

In short, access to natural gas remains a key enabler of power availability in many of the emerging hyperscale clusters, even as longer-term decarbonization policies cast uncertainties over future infrastructure planning.

> **As traditional grid-connected utility-scale solutions fall behind hyperscale demand curves, a new class of power developers is stepping into the gap. Data center operators, once dependent on traditional utilities, are increasingly turning to private power developers who can deliver firm, scalable electrons on compressed timelines.**

In today's landscape, the ability to secure power has become a prerequisite for growth; those who can guarantee electrons within 12 to 24 months are emerging as critical partners in the AI economy. In this same world, transmission bottlenecks, interconnection delays and permitting challenges are no longer background risks. Instead, they are strategic blockers shaping investment decisions.

While AI algorithms surge in their complexity, power developers do not seem eager to reinvent the wheel. All eyes are turning to established and reliable technology for power generation. With the eagerness to be independent and deliver on the new ROI, one solution seems to have emerged: on-site/behind-the-meter combined cycle gas turbine (CCGT) power plants, sited in areas where access to natural gas poses no significant risk. In doing so, we are powering the intelligence of the future using the fuels of the past — an irony that underscores the compromises shaping this next wave of development.

---

[2] "The Breakdown of the Merchant Generation Business Model," WBK-PRG Merchant Generation White Paper, June 2017.
[3] "Announcing The Stargate Project," OpenAI, January 21, 2025.
[4] "Landry Announces Meta Selects North Louisiana as Site of $10 Billion Artificial Intelligence Optimized Data Center," Office of Governor Jeff Landry, December 4, 2024.
[5] "U.S. Data Center Powerhouses: The 5 Fastest-Growing Hubs," Upwind, November 12, 2024.

One case demonstrating the developer logic is Homer City Generating Station in Indiana County, Pennsylvania. The $10-billion project is expected to deliver a 4.5 GW natural gas powerplant. This translates to $2,222 per kW installed almost 2.5 times higher than what was a standard CCGT CapEx five years ago. This facility is designed to operate behind the meter, supplying electricity directly to on-site data centers with a strategic location atop the Marcellus Shale natural gas resources. Construction is slated to begin in 2025, with power generation expected by 2027. This accelerated timeline is made possible by leveraging existing infrastructure from a retired coal plant.

While the instinct to question natural-gas-fired turbines is valid, the decisions taken in today's reality seem reactive rather than fully premeditated. Several drivers are propelling these solutions:

1. Speed and Deployment Certainty
Avoids backlogged interconnection queues and cumbersome system upgrades for GW scale generation; is faster to permit especially if in attainment areas with no New Source Review (NSR) requirements; creates manageable supply chain challenges especially if turbines are available.

2. Baseload and Relatively Flexible Power
Provides 24/7 uptime, redundancy and operational flexibility for AI workloads.

3. Familiar Technology and Financing Models
Leverages proven models that can simplify the processes to underwrite, scale and maintain versus emerging energy alternatives (i.e., small modular reactors or renewables plus storage).

4. Design Flexibility
Supports modular plant GW scale designs with phased construction. Can be built for potential future carbon capture add-on, dual fuel capability or hydrogen blending.

5. Offtake Direct Control Over Power Procurement
Allows data center operators to bypass utility interconnection delays, manage procurement timelines directly and secure onsite power for their hyperscale operations.

3. STRATEGIC RISKS IN THE RUSH TO BUILD

While looking more closely at this trend with a pragmatic lens, history reminds us that today's necessities can become tomorrow's liabilities.

> **As developers accelerate investments in CCGTs to meet near-term AI-driven demand, disciplined capital planning is essential to avoid stranded assets in a shifting energy landscape.**

While CCGTs offer familiarity and speed, they are far from plug-and-play. Power developers face real technical challenges, among them the complexity of integrating firm generation behind the meter. Global supply chain strains, along with rising demands for the larger class turbines and other already in-high-demand specialized equipment, are driving up project costs. More importantly, these strains are shifting project schedules and planned commercial operations dates. Under traditional models, these pressures would have made investors hesitate.

To inform long-term investment decisions, we have identified six strategic risks tied to the rapid rise of behind-the-meter natural gas generation, each with historical parallels and practical implications for today's developers:

6 Ethan Howland, "Largest US gas-fired power plant planned for data centers in Pennsylvania ," Utility Dive, April 3, 2025.

	Risk	Reality	Historical Parallel	Mitigations
Regulatory and Policy	New regulations on carbon, methane or local air pollutants (NOx, CO) could heavily impact operating costs or force early retirement.[7]	CCGTs today may be favored by the "current policy" and is "relatively" clean, but "current" and "relatively" are moving targets. ESG pressures and federal/state carbon policies could turn "advantage today" into "liability tomorrow."	Coal plants built in the 1970s and 1980s seemed like permanent baseload infrastructure until environmental regulations and gas prices flipped the economics.[8]	Implement modular and flexible design features that allow for future retrofits or upgrades to meet evolving environmental standards and regulations.
Fuel Supply and Price Volatility	Natural gas access today does not mean secure, affordable gas tomorrow.	Local bottlenecks, seasonal shortages can drastically increase costs or reduce availability.	In the mid-2000s, New England CCGTs faced crippling winter gas shortages, spiking spark spreads and forcing operational curtailments.	Include backup fuel oil. Lock in supply through firm transportation agreements and financial hedges. Choose site projects near pipeline interconnects or storage facilities to reduce basis risk.
Demand Overestimation and Load Shifts	AI–driven demand forecasts could prove overly aggressive, leading to underutilized infrastructure: slowdown or geographical shifts of growth.	Basis risk (local gas price deviations from Henry Hub) is often underestimated. AI compute workloads may shift geographically (e.g., cheaper land, new energy hubs) or technologically (e.g., more efficient chips, lower power demand per compute unit).	In the early 2000s, merchant gas plants were built across the U.S. in anticipation of soaring deregulated power demand. Many became stranded when the market failed to materialize.[9]	Build scalable CCGT blocks that can be phased in or out based on actual load realization. Structure PPAs with staged capacity commitments or termination options. Choose locations that offer repurposing potential for grid interconnection or alternative off-takers.
Regulatory and Policy	Building long-life assets in a space where energy technologies are evolving faster than the financing cycles.	Breakthroughs in long-duration storage, nuclear microreactors or renewable baseload hybrids could dramatically alter competitiveness of islanded CCGT plants within 10–15 years.	Peaker plants lost value rapidly once grid-scale batteries became financially viable for short-duration dispatch services.[10]	Design the balance of plant (BOP) to allow for reuse if the core prime mover technology is replaced. Allow flexibility for future integration of battery storage, hydrogen and/or renewables. Avoid 20-year fixed structures. Instead, structure flexible financing with triggers to re-evaluate after 5–7 years.
Fuel Supply and Price Volatility	Building a high-availability islanded system is not just "buy a CCGT and switch it on." Integration, redundancy, operability and maintenance, all must be flawlessly executed.	Downtime penalties in SLA-driven data center contracts represent a substantial operational risk. System integration risk (not just generation risk) such as ability to load follow could destroy economics if not properly managed.	The U.S. coal-to-gas conversion wave in the 2000s–2010s revealed the risks of retrofitting assets without fully adapting operational systems. Many plants faced fuel supply constraints, inefficiencies and reliability issues, ultimately leading to financial underperformance and early retirements.	Embed Tier IV-level redundancy, proven operation and maintenance partners, and robust testing protocols from design through operations.
Demand Overestimation and Load Shifts	Selling or refinancing a standalone, islanded CCGT tied to a single data center or AI cluster may be difficult in secondary markets.	These assets are bespoke, not fungible like utility-scale plants tied into the grid.	Dedicated industrial cogeneration plants in the 1990s often struggled to find buyers once the industrial host shut down or relocated.	Engineer the plant for potential future interconnection to wholesale markets. Include investor exit windows, change-of-use triggers and take-out refinancing options.

[7] "Power Sector Evolution," U.S. EPA, accessed May 14, 2025.

[8] Ira Shavel, "Coal Plant Retirements and Power Markets," Brattle Group, October 2015.

[9] Raymond L. Gifford et al., "The Breakdown of the Merchant Generation Business Model," WBK-PRG Merchant Generation White Paper, June 2017.

[10] Brian Martucci, "4-hour batteries best replacements for aging Maine peaker plants: study," Utility Dive, May 3, 2024.

4. CONCLUSION

Data center growth, driven by AI, is accelerating demand for fast, firm power. Developers are responding with behind-the-meter gas projects that bypass grid bottlenecks, provide for easier permitting and deliver capacity on compressed timelines. This response solves for speed, but introduces exposure to capital intensity, regulatory volatility and structural rigidity.

Several of these projects are showing CapEx profiles well above historical benchmarks. Whether those costs are recoverable depends on future utilization, market evolution and optionality built into the asset. Without design flexibility or integration pathways, these assets risk being stranded as technology and policy advance. Speed will remain a defining advantage. But in this environment, it is the developers who can execute with precision — while preserving adaptability — who will lead.

A&M's Infrastructure and Capital Projects team brings hands-on experience working with both power producers and data center owners and operators to navigate this challenge. From capital strategy to execution, we help our clients deliver fast without compromising long-term value.

We acknowledge the significant contributions of Iman Aziz, SPURS/ Hubert H. Humphrey Research Fellow at MIT Department of Urban Studies and Planning.

This article contains contributions from George Barrenechea.

CONTACT US

Hassan Hamdan
Managing Director
hhamdan@alvarezandmarsal.com

Dennis Bocuzzi
Senior Director
dbocuzzi@alvarezandmarsal.com

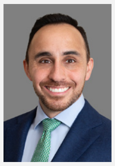

Jorge Macedo
Director
jorge.macedo@alvarezandmarsal.com

ABOUT ALVAREZ & MARSAL

Founded in 1983, Alvarez & Marsal is a leading global professional services firm under the parent company Alvarez & Marsal Inc. Renowned for its leadership, action and results, Alvarez & Marsal provides advisory, business performance improvement and turnaround management services, delivering practical solutions to address clients' unique challenges. With an international network of experienced operators, world-class consultants, former regulators and industry authorities, Alvarez & Marsal helps corporates, boards, private equity firms, law firms and government agencies drive transformation, mitigate risk and unlock value at every stage of growth.

To learn more, visit: **AlvarezandMarsal.com**

Follow A&M on:

ALVAREZ & MARSAL
LEADERSHIP. ACTION. RESULTS.℠