Originally published as Carbon Direct, which is now Relae. Learn more at relae.co

Carbon Direct

The Navigator

Special Power Edition | July 2026



Contents

Letter from the Editor 3

Credits 5

Acronyms and abbreviations 5

AI meets the grid: Interconnection queue analysis in PJM and ERCOT 6

Unlocking clean, firm power with enhanced geothermal systems 11

The billion-dollar case for enabling data center load flexibility 12

Inside NERC's Level 3 Alert on data center loads 14

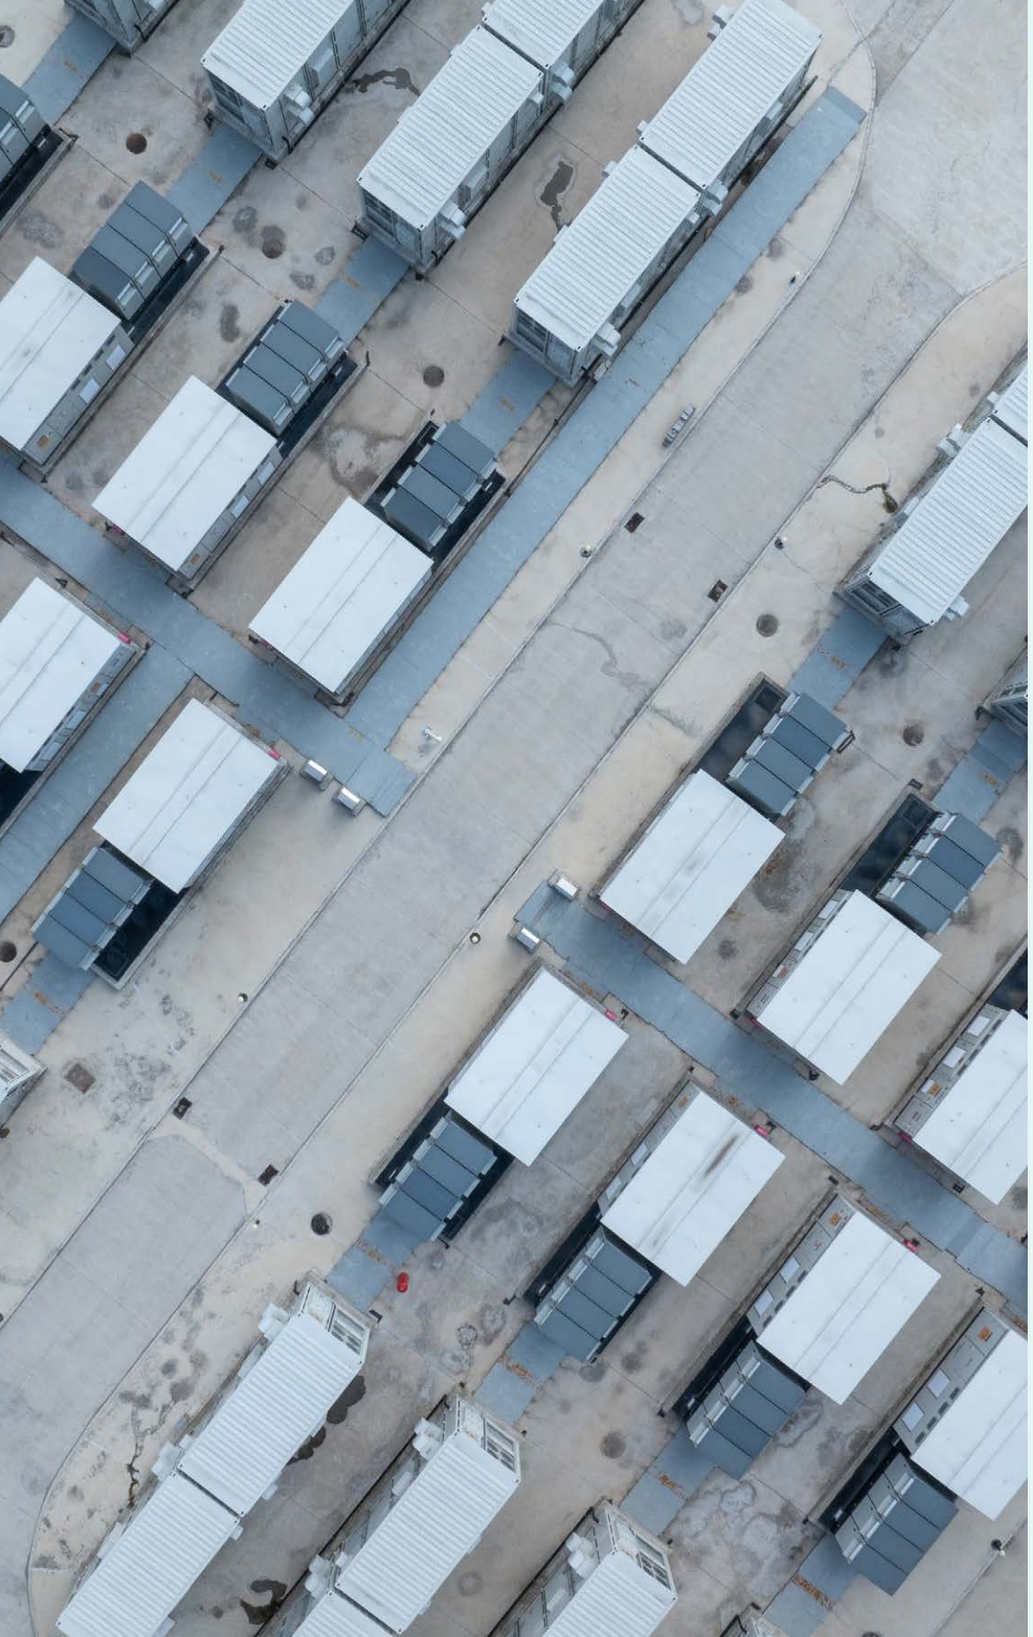
The shifting playbook for corporate power procurement 15

Why behind-the-meter power emissions belong in scope 2. 17

Developer responses to data center opposition 18

Disclaimer 19





Letter from the Editor

There's a version of the energy transition story that's been told a thousand times: renewables get cheaper, fossil fuels get squeezed out, and grids slowly green. It's a satisfying narrative arc. It's also, at the moment, incomplete. AI changed the timeline.

Data centers are arriving on grids at a pace and scale these systems weren't designed for, driving the fastest load growth the US has seen in a generation. The infrastructure that kept the lights on through a century of gradual change is now being asked to absorb a step change in consumption in a matter of years, not decades. The old story isn't wrong. It's just no longer sufficient. This issue is about what comes next.

On the supply side, we find that the interconnection queue (i.e., the regulatory pipeline connecting new power generation to grids) is a more serious bottleneck than the headline numbers suggest (**pg. 6**). We also look at why enhanced geothermal systems are finally breaking out of their resource-constrained niche, pulled forward by hyperscalers willing to act as anchor customers for clean, firm power (**pg. 11**).

On the demand side, Carbon Direct's power system modeling puts a dollar figure on what grid inflexibility costs: up to US$5.5 billion a year at the high end of data center load growth forecasts in ERCOT (40 GW). It also shows that 25% load flexibility capability, provided by data centers, eliminates that exposure entirely (**pg. 13**). As Texas SB6 is turning that flexibility into a regulatory requirement, other regulators may follow, and data center operators are looking to on-site battery storage, fuel cells, and other emerging technologies as part of the solution. Separately, NERC's first-ever Level 3 Alert on data center loads makes clear that instantaneous disconnections are now an active reliability issue, not a coming one (**pg. 14**).



Carbon accounting is evolving alongside grids. The GHG Protocol's proposed scope 2 revisions would expose an emissions gap many large buyers haven't actively tracked. This gap is largest in the same markets that are absorbing the most data center demand (**pg. 15**). Behind-the-meter power generation is accelerating as a workaround for grid constraints, and a widespread misclassification is emerging that companies need to correct before contracts close (**pg. 17**).

No discussion of data centers would be complete without a careful look at community opposition, which is rapidly becoming a key determinant of successful project completion. As we show, there are three response postures that separate developers who advance projects from those who don't (**pg. 18**).

Grids are being asked to do something genuinely new. Our goal, as always, is to help you find clarity in what that means for your work.

Colin McCormick, PhD
Editor and Chief Innovation Officer

Credits

Report authors and contributors

Given the highly collaborative nature of Carbon Direct, many people contributed insights to this report beyond the list below. The following team members contributed substantially to developing this report.

Editor: Colin McCormick, PhD, Chief Innovation Officer
Deputy Editor: Adrianna Sutton, Managing Editor of Science Writing
Managing Editor: Jen Taveras, Senior Business Manager

Douglas Bryan, Senior Manager of Power and Energy Systems
Derya Eryilmaz, PhD, Vice President of Power Commercialization
Daniel Garcia, PhD, Senior Manager of Life Cycle Assessment
Jonathan Goldberg, Chief Executive Officer and Founder
Grant Gutierrez, PhD, Head of Community Impacts
Liam Kilroy, Energy Data Scientist
Julia Millot, Senior Power Decarbonization Manager
Shirin Mavandad, Senior Research Associate
Peter Psarras, PhD, Director of Low-Carbon Energy
Britt Warthen, Senior Brand Designer

Acronyms and abbreviations

AI: artificial intelligence
BESS: battery energy storage systems
BTM: behind-the-meter
CCS: carbon capture and storage
CFE: carbon-free energy
CILR: customer-initiated load reduction
DOE: Department of Energy (United States)
EGS: enhanced geothermal systems
EPRI: Electric Power Research Institute
ERCOT: Electric Reliability Council of Texas
FERC: Federal Energy Regulatory Commission
GHG: greenhouse gas
GPU: graphics processing unit
GW: gigawatt
GWh: gigawatt-hour
IA: interconnection agreement
MW: megawatt
MWh: megawatt-hour
NERC: North American Electric Reliability Corporation
OBBB: One Big Beautiful Bill Act (United States)
PJM: regional transmission organization covering Pennsylvania, New Jersey, and Maryland
PPA: power purchase agreement
REC: renewable energy certificate
SB6: Senate Bill 6 (Texas)
US: United States

AI meets the grid: Interconnection queue analysis in PJM and ERCOT

The supply perspective part 1: How fast is new power generation being added to the grid to meet rising data center demand?

Derya Eryilmaz, PhD and Liam Kilroy

THE BOTTOM LINE: The interconnection queue, the regulatory process by which new power generation connects to grids, has become a critical bottleneck for meeting AI-driven data center power demand in two of the largest US electricity markets, ERCOT (covering Texas) and PJM (covering Pennsylvania, New Jersey, and Maryland). Despite holding over 500 GW of planned generation capacity, actual queue timelines in these markets run two to five times longer than FERC's targets, and most projects are ultimately withdrawn before reaching commercial operation. The One Big Beautiful Bill Act (OBBB) has accelerated this mismatch by simultaneously driving withdrawals of renewable energy projects and a surge in new natural gas entries. Data centers cannot rely on existing grids and traditional interconnection processes to predict the availability of clean, firm power; they need rigorous, project-level analysis to identify when and where this power is likely to materialize.

An unprecedented surge in energy demand

AI-driven data center growth is creating a geographically concentrated surge in US electricity demand, unlike anything grids have managed before. Since October 2025, more than 10 GW of new data center projects have been announced in ERCOT and PJM alone. Nationally, data centers are projected to grow from their current 4% of total US electricity consumption to 7–12% by 2028. The scale and speed of this demand has exposed a structural problem: the regulatory process for connecting new power generation to grids wasn't built for this.

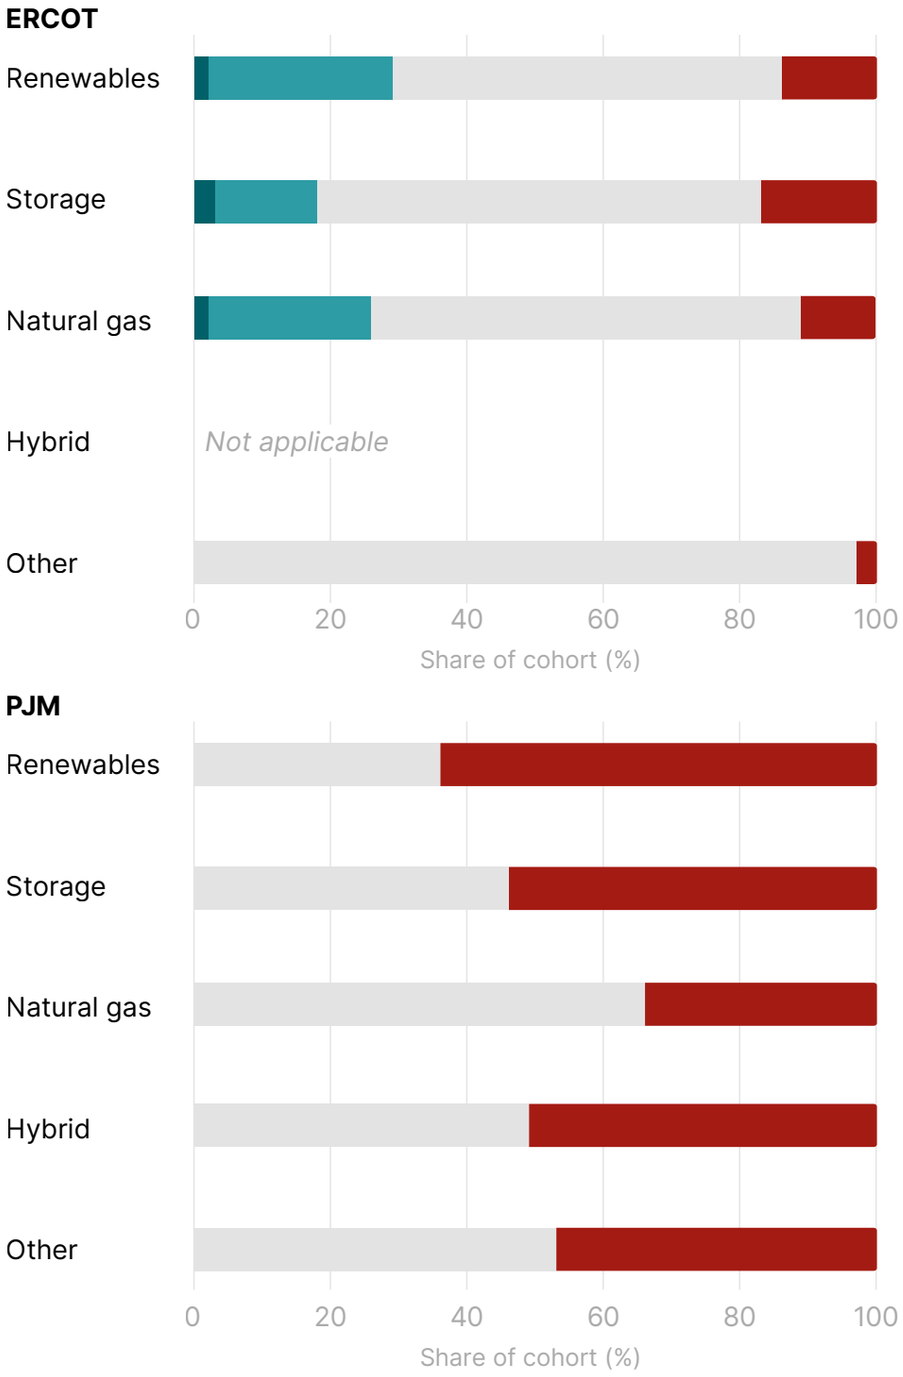
What is the interconnection queue?

Before any new power plant can deliver electricity, it must clear the interconnection queue, a multi-stage regulatory process that determines when and how a power generator connects to a grid. In theory, the queue is no more of a hindrance than most other regulatory and engineering processes that precede new power generation coming online. In practice, however, it has become the critical bottleneck between planned capacity and actual power delivery.

The numbers look reassuring at first glance. There are approximately 1,500 active projects in the ERCOT and PJM queues, representing over 500 GW of potential generation and storage capacity. This far exceeds projected data center load growth through 2030. But this figure is misleading. While FERC's 2023 interconnection reform rule targets a queue timeline of no more than 8–11 months from interconnection request to agreement, actual timelines average 20 months in ERCOT and 40 months in PJM. In fact, most projects never make it: in ERCOT, only 2.8% of potential projects have reached commercial operation; in PJM, more than half of renewable projects in the queue have been withdrawn.

The queue holds a lot of capacity on paper. It delivers far less in practice.

> In ERCOT, only 2.8% of potential projects have reached commercial operation; in PJM, more than half of renewable projects in the queue have been withdrawn.



The One Big Beautiful Bill reshapes the pipeline

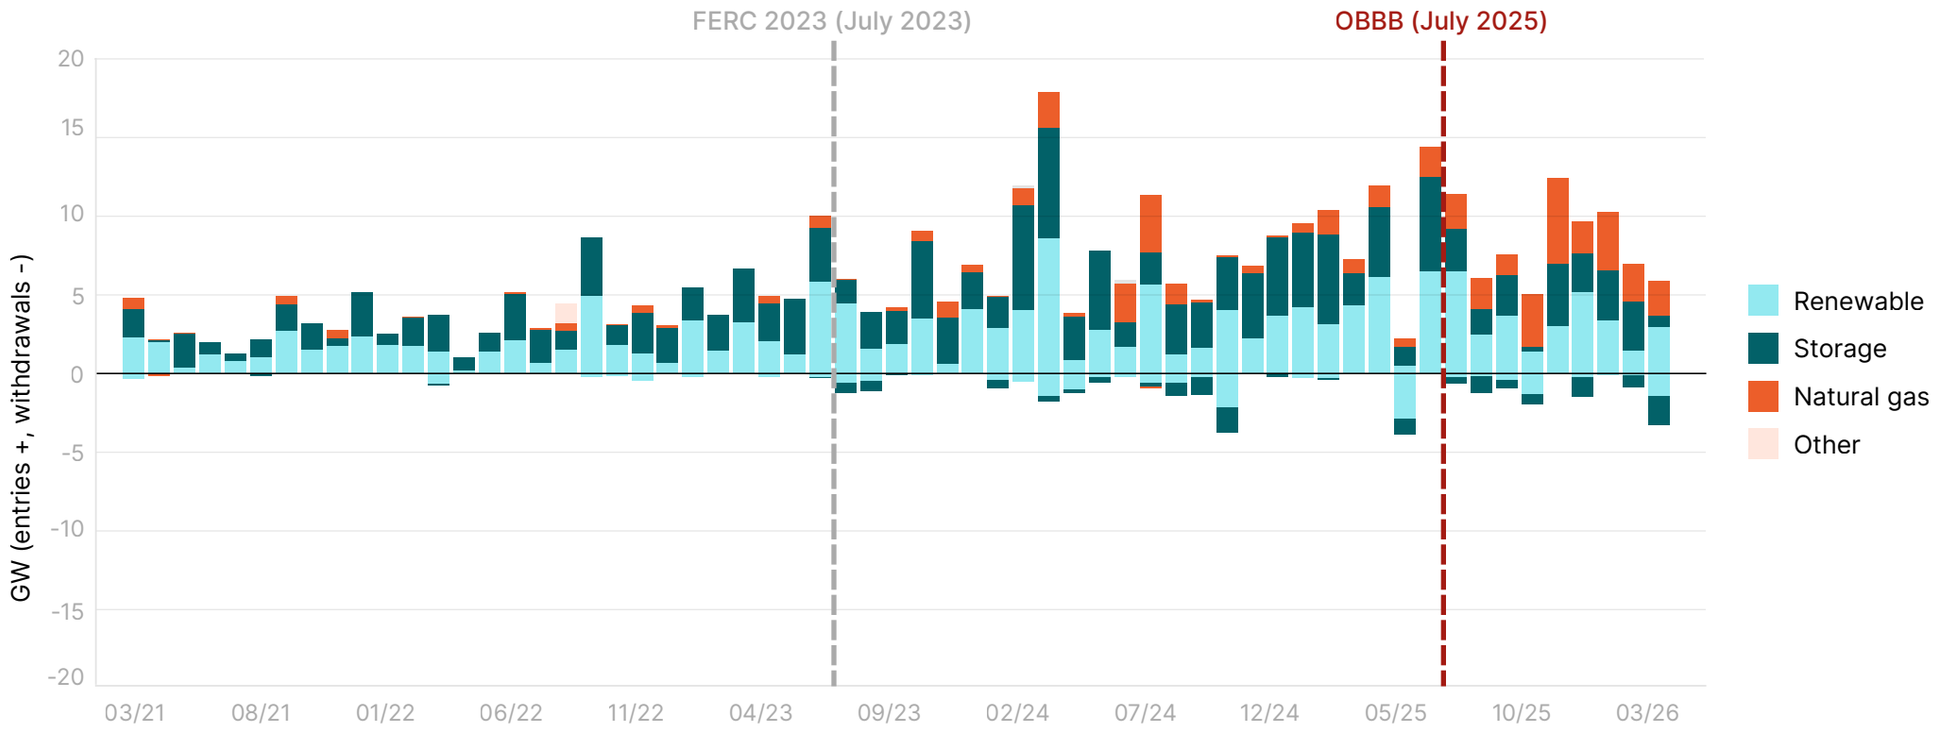
The OBBB has fundamentally altered the technology mix entering the queue, accelerating natural gas while triggering a retreat of renewables. Post-signing of the OBBB, gas interconnection requests in ERCOT accelerated by approximately 150% on a monthly basis, rising from 1.1 GW/month to 2.7 GW/month. Over the same period, storage entries in ERCOT fell 42% month-over-month.

Renewables tell a similar story. Withdrawal rates increased significantly in both markets following the bill's passage. PJM showed a more complex picture: nine gas projects totaling 3.5 GW were also withdrawn, reflecting broader market uncertainty rather than a simple pivot toward gas. However, the net effect across both markets is a queue that looks increasingly oriented toward dispatchable fossil generation, with consequences for anyone counting on clean capacity to materialize on schedule.

In service Past IA Under study Withdrawn

Success rate by technology (March 2021–March 2026). Notes: IA = interconnection agreement, a key milestone in project progression toward grid connection. Cohort window uses earliest available entry data from each Independent System Operator (i.e., PJM's interconnection request submitted date, and ERCOT's screening study start date). Renewable projects include wind and solar; storage projects include battery energy storage systems (BESS); natural gas projects include any gas technology such as combined cycle or simple cycle; hybrid projects include any fuels with compound labeling (e.g., solar + BESS, wind + BESS), and other projects include any technology listed outside of renewable, natural gas, storage and hybrid. Source: Carbon Direct analysis of Independent System Operator generation interconnection queue.

ERCOT monthly queue entries vs. withdrawals by fuel group
March 2021–March 2026



Monthly capacity changes in the ERCOT interconnection queue between March 2021 and March 2026, FERC 2023 is indicated by the gray dotted line and OBBB is indicated by the red dotted line. Source: Carbon Direct analysis of PJM and ERCOT interconnection queue.

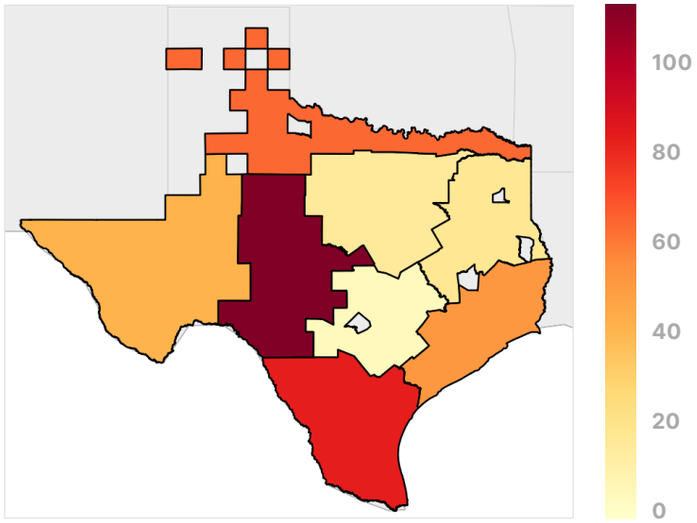
Where the load is, and where the capacity isn't

A geographic mismatch between where data center load is growing and where queue capacity is sited creates a near-term reliability risk that the raw queue numbers obscure.

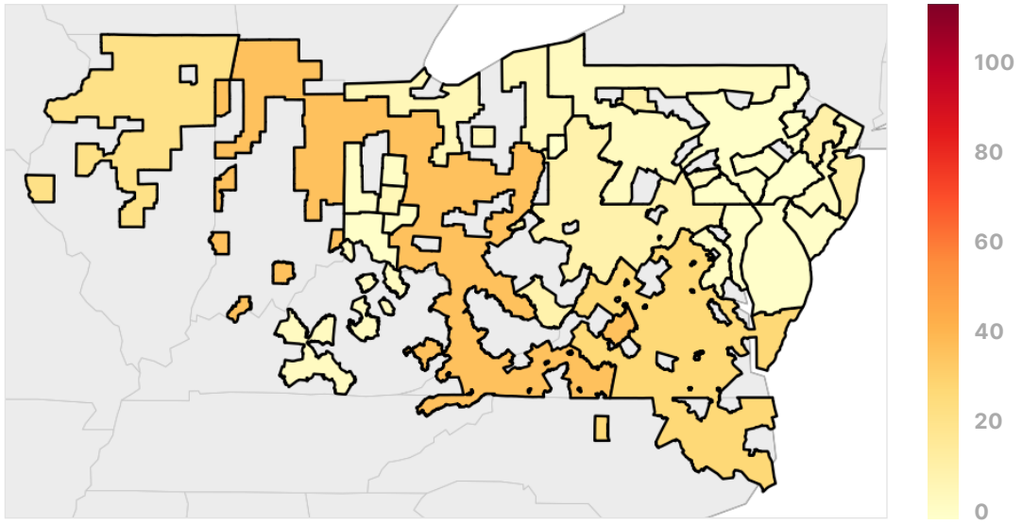
ERCOT's two largest load zones account for 53% of regional load but hold only 26% of queue capacity. This means the highest-demand areas face the greatest exposure to supply gaps. That imbalance is compounded by technology-specific queue dynamics: gas projects move through ERCOT's process significantly faster than renewables or storage, averaging 17 months compared to 27 for renewables. The post-OBBB surge in gas entries is therefore directly relevant to near-term reliability. Dispatchable gas capacity is advancing faster through the very regions under the most pressure.

PJM shows better geographic alignment between data center load and queue capacity. But low project success rates mean this current alignment doesn't guarantee adequate supply in the future. The legacy of PJM's 2022 queue freeze, combined with the OBBB's reduction of renewable incentives, leaves the region facing compounding risks as data center demand accelerates.

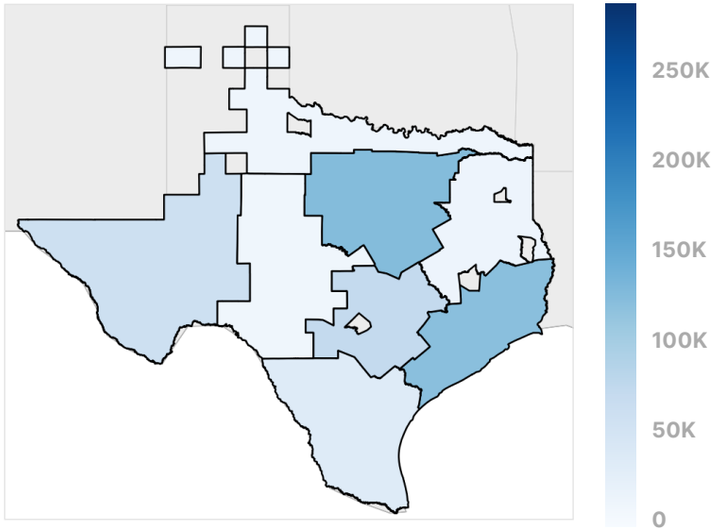
ERCOT: Total active projects (GW) by hub zone



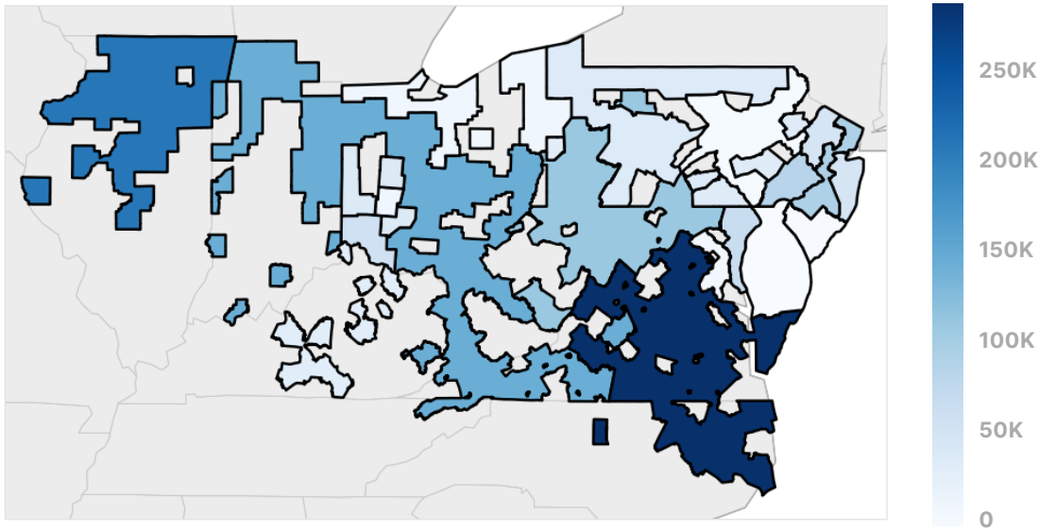
PJM: Total active projects (GW) by hub zone



ERCOT: Total zonal load (GWh), 2025 data



PJM: Total zonal load (GWh), 2025 data



Active project capacity (i.e., projects with interconnection agreement signed, under construction) versus total zonal load for ERCOT (top) and PJM (bottom). Note: We treat ERCOT weather zones and PJM transmission operator zones as roughly equivalent for this analysis. This requires making a few assumptions. For ERCOT, projects are assigned to weather zones based on a two-tier lookup: project county is mapped to ERCOT's county-to-zone profile, and projects without county data fall back to ERCOT's capacity, demand, reserve reporting zone field as mapped to weather zones. For PJM, each project is assigned to a zone using a three-tier lookup: project transmission owner is mapped to load zone, with county-level overrides for ambiguous cases, and projects with no transmission owner fall back to the state level. Source: Carbon Direct analysis using Independent System Operator/Regional Transmission Organization, Energy Information Administration, and US Census datasets.

What this means for data centers

ERCOT is better positioned than PJM to close near-term reliability gaps: with shorter timelines, higher execution rates, and a growing pipeline of dispatchable gas capacity that can bridge supply gaps while clean energy scales. However, being better positioned is not the same as well-positioned. Neither ERCOT nor PJM offers a reliable off-the-shelf solution for data centers that need firm, clean power on a compressed timeline.

Meeting the scale and pace of AI-driven demand for sustainable power requires more than queue participation. It requires rigorous analysis of project viability, technology siting, and dispatch optimization to determine where firm, clean capacity will actually materialize, and when. Data centers that treat the interconnection queue as a proxy for power availability will find the gap between what's planned and what actually gets built to be an expensive surprise.



Read more on interconnection queues for data center power generation here →



Unlocking clean, firm power with enhanced geothermal systems

The supply perspective, part 2: How can new geothermal power technology help bring clean, firm power to the grid fast?

Peter Psarras, PhD

WHY IT MATTERS: Enhanced geothermal systems (EGS) are rewriting the rules on where geothermal energy is possible, with the potential to scale beyond a handful of states with compatible geology. Solar and wind can't reliably fill the 24/7 baseload gap hyperscalers need, and nuclear timelines stretch well past 2030. EGS is arriving at precisely the moment that grids most need dependable, low-carbon power supply—and hyperscaler demand is pulling it forward.

Conventional geothermal energy has always had a location problem. It requires a rare coincidence of high subsurface temperatures, naturally occurring fluid, and sufficient rock permeability. These are conditions that, in the US, can only be found in California, Nevada, and a handful of other states. That geographic constraint has kept geothermal energy in a niche despite its obvious appeal: a source of firm, renewable, low-carbon power with no purchased fuel costs.

Enhanced geothermal systems (EGS) dissolve that constraint by engineering the conditions geothermal energy needs. Instead of locating natural permeability, EGS creates it to make subsurface heat, which exists at sufficient depth almost everywhere, readily accessible. The DOE estimates this approach could unlock more than 5,500 GW of potential energy resources in the US.

One of the technology's long-standing obstacles has been induced seismicity. High-pressure fluid injection can trigger fault slip. This happened in Basel in 2006 and in Pohang in 2017, killing both projects and setting the industry back years. Next-generation approaches, like Sage Geosystems' gravity-assisted fracturing, address this directly. By avoiding overpressures and directing fractures away from critically stressed fault zones, Sage's method is designed to work with the subsurface rather than against it.

Recent deals signal that EGS has crossed from the demonstration phase into commercial deployment. Sage's US$97 million Series B, co-led by Ormat Technologies and Carbon Direct Capital, will fund the first commercial EGS facility at an existing Ormat plant—pairing novel subsurface engineering with six decades of conventional geothermal operational expertise. Meanwhile, Fervo Energy's deal with Google and Meta's 150 MW power purchase agreement with Sage reflect hyperscalers' growing role as anchor customers for clean, firm technologies, providing the revenue certainty needed to unlock project financing.



BLOG POWER CLIMATE STRATEGY

The new geothermal energy: How EGS unlocks clean, firm power at scale



Peter Psarras, PhD
Senior Decarbonization Engineer

Read more on EGS here →

The billion-dollar case for enabling data center load flexibility

The demand perspective: Can new data center load become more flexible in order to ease pressure on the grid?

Jonathan Goldberg, Douglas Bryan, and Liam Kilroy

THE BOTTOM LINE: As AI-driven data center demand pushes new load in ERCOT (up to 40 GW by 2028), Carbon Direct's power system modeling reveals that data center flexibility is critical for avoiding billions in consumer welfare impacts. Without load flexibility, forced outages appear at 30 GW of data center load and escalate sharply beyond that threshold. With emergency curtailment capabilities, all forced outages are eliminated across every scenario modeled, requiring an average of only 5% curtailment for less than 1% of annual hours. Texas SB6 is turning this flexibility into a regulatory requirement, and on-site battery storage is emerging as the technology that enables both compliance and continuous uptime.

Grids under pressure

AI data center growth is straining US grids at a pace and scale that traditional grid management tools were not designed to handle. NERC's 2025 Long-Term Reliability Assessment projects summer peak demand growth of 224 GW nationally over the next decade—more than double the prior year's projection. Texas is among the most acute examples: industry surveys project data center load in ERCOT could reach or exceed 40 GW by 2028.

For decades, utilities and grid operators managed relatively flat load growth; AI data centers represent a step change in what planners must now respond to. Historical demand growth gave utilities lead time to procure new power generation. The pace of data center deployment doesn't. Supply-side planning cycles that once had years to respond now have months, forcing regulators to develop novel approaches to grid management in real time.

Texas SB6: Flexibility becomes a requirement

Texas SB6 represents the most significant restructuring of large load interconnection rules in ERCOT's history, and it signals a broader shift. Load flexibility is now a regulatory prerequisite to interconnection.

SB6 requires all new loads above 75 MW connecting to ERCOT to install remote curtailment equipment and demonstrate the ability to reduce power draw on demand during declared grid emergencies. The mandate exists because voluntary programs haven't worked. Demand response programs have struggled to attract data center participation for a straightforward reason: in an AI arms race, the value of uninterrupted compute time far exceeds available curtailment payments. That valuation asymmetry is precisely what a mandate is designed to overcome.

The technical case for compliance is also strengthening. EPRI's DCFlex program is developing protocols and standards for large-load demand response. Emerald AI has demonstrated that 25% load curtailment is achievable without GPU service interruption. Google has already contracted 1 GW in demand response agreements with multiple US utilities—a signal that hyperscalers are moving from skepticism to execution.

Our solution

Carbon Direct's power system modeling framework quantifies what data center load flexibility is worth (and what inflexibility costs) across a range of ERCOT growth scenarios.

The analysis modeled data center loads of 5–40 GW, in 5-GW increments, paired with co-located gas generation covering approximately 70% of data center energy needs, under two conditions: no load flexibility and 25% emergency curtailment capability. The results are stark.

Without flexible load management, forced load shedding first appears at 30 GW of data center

load. At 40 GW, shedding reaches 158 GWh across 81 hours, equivalent to nearly three times ERCOT's average daily energy consumption. At ERCOT's standard value for lost load benchmark, US$35,000 per MWh, that translates to approximately US$5.5 billion in annual consumer welfare losses.

What flexibility actually requires

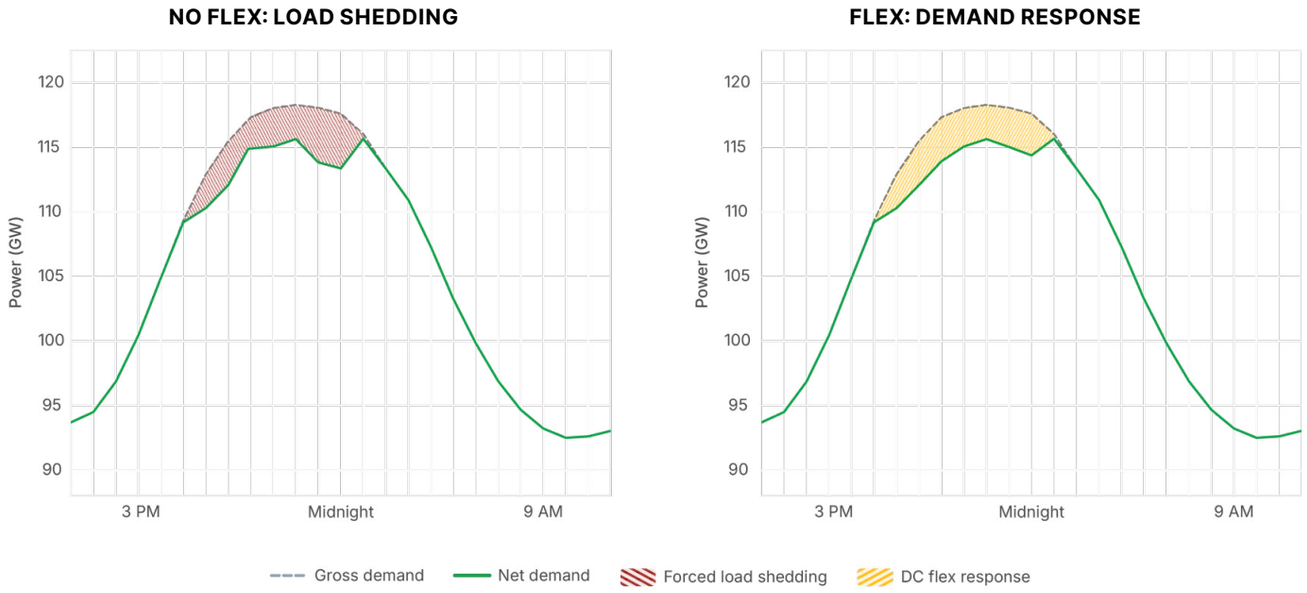
Enabling 25% emergency curtailment capability eliminates forced load shedding entirely in every scenario modeled. In the flexible 40 GW scenario, 165 GWh of controlled, short-duration curtailment is spread across 86 hours—less than 1% of annual operating hours—with an average demand response of just 5% of nameplate load and a maximum single event of 14%.

In the model, demand response functions as a "ghost battery" at the data center's grid node, absorbing grid stress exactly as a physical battery would. That analogy has a direct real-world counterpart: on-site battery energy storage systems.

The battery storage opportunity

Co-located battery storage is the highest-value deployment opportunity to emerge from this analysis. It converts a compliance obligation into an uptime guarantee.

For data center operators, co-located storage allows curtailment events to become battery discharges rather than workload interruptions,



Hourly ERCOT load with 40 GW data center demand. Forced load shedding events (left) and data center (DC) flex demand response deployed to mitigate shedding events (right). Source: Carbon Direct, modeled using CD-PyPSA-USA.

maintaining compute continuity while delivering the grid relief SB6 requires. For storage developers, the commercial case is now quantified: each hour of demand response in the 40 GW scenario delivers approximately US$64 million in avoided consumer welfare losses. A well-positioned co-located battery asset can credibly claim to preserve that value.

Developers and investors need to design for flexibility

The data centers being built today are large enough to reshape grid reliability across entire regions. The regulatory environment, led by Texas SB6, is beginning to reflect that scale by treating load flexibility as a prerequisite for interconnection, not a voluntary option. Designing curtailment capability from the start converts a compliance requirement into a durable grid asset and an uptime advantage. The modeling evidence is clear: doing so costs far less in compute disruption than the alternative costs regional grids.

Read more on the case for data center load flexibility here →

Inside NERC's Level 3 Alert on data center loads

The grid perspective: What new power reliability challenges are large data centers creating for grid operators?

Douglas Bryan and Colin McCormick, PhD

WHY IT MATTERS: NERC's Level 3 Alert signals that data center load behavior is now an active grid reliability issue. The modeling requirements it introduces will raise the cost and complexity of interconnection for any large load that can't demonstrate predictable, controllable operation and will reward those that can.

On May 4, 2026, NERC issued a Level 3 Essential Actions Alert, its highest-urgency classification, targeting a pattern of large-scale data center disconnections that existing grid planning frameworks were not designed to handle.

This was triggered by repeated customer-initiated load reductions (CILRs), in which data centers' automatic protection systems dropped more than 1,000 MW from a grid in a matter of seconds. These events are a self-preservation reflex. Facilities trip offline when grid frequency deviates to protect sensitive equipment. However, at hyperscale that reflex can amplify the very instability it's designed to mitigate, removing massive blocks of load so rapidly that the grid struggles to rebalance.

NERC's companion Reliability Guideline extends the concern into long-term planning, requiring utilities and ISOs to model large-load behavior across three timescales: sub-second electromagnetic transients, seconds-to-minutes stability, and hours-to-years resource adequacy. Together, the alert and the guidelines establish controllable, predictable load behavior as a new baseline for interconnection.

The stakes are concrete. Data center operators that cannot demonstrate compliance across all three modeling timescales face longer interconnection queues and potentially restricted access to constrained transmission nodes. For storage developers, the alert introduces a newly quantified reliability value: co-located batteries can buffer CILRs, preventing automatic disconnections and absorbing the grid stress that the protection systems were designed to offload. This strengthens the commercial case for pairing storage with large data center loads.

Learn more with Carbon Direct's Colin McCormick and Douglas Bryan on this episode of the Volts Podcast, hosted by David Roberts.

Read more on NERC's Level 3 Alert here →

VOLTS PODCAST | JUNE 10, 2026

Why is NERC so worried about data centers?



BLOG POWER ENERGY + ELECTRICITY

Inside NERC's Level 3 Alert on data center loads



The shifting playbook for corporate power procurement

The emissions perspective part 1: How is corporate clean energy procurement evolving as data center power demand accelerates?

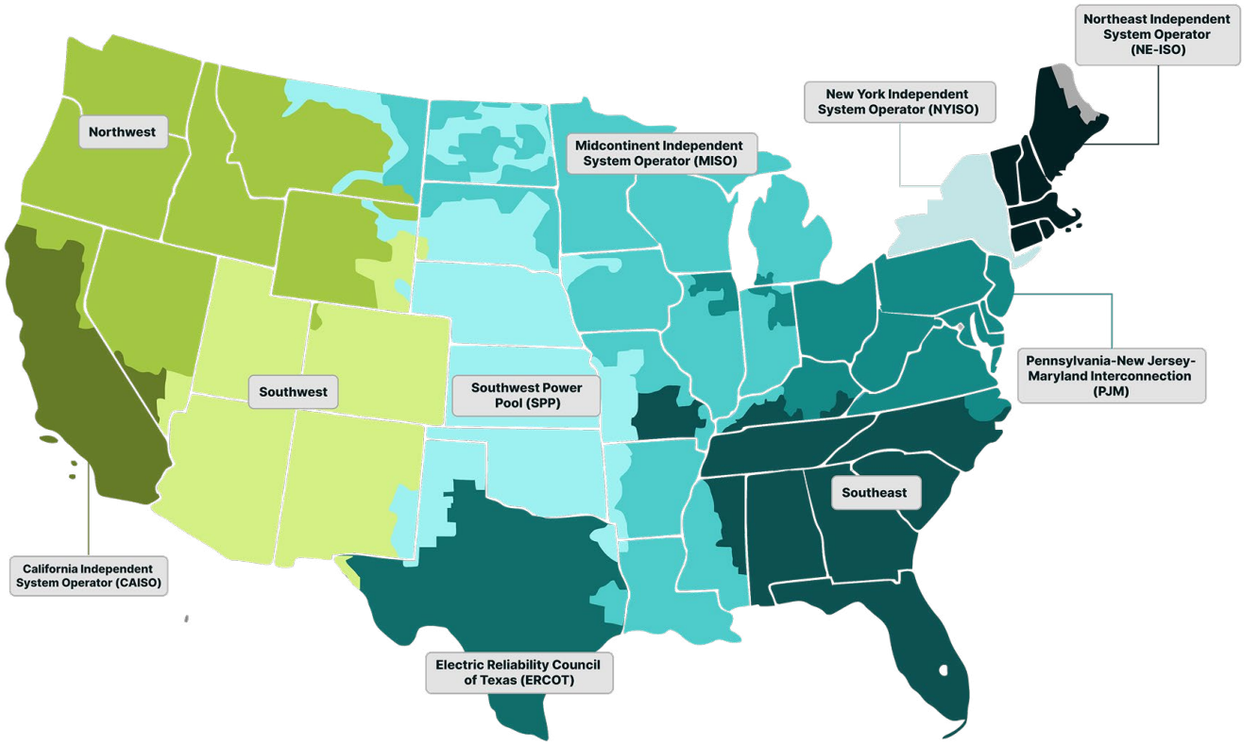
Julia Millot and Douglas Bryan

WHY IT MATTERS: Large electricity consumers who relied on annual RECs face significant and unquantified exposure if proposed GHG Protocol scope 2 revisions take effect. That exposure is greatest in the same markets where data center demand is growing fastest.

The GHG Protocol's proposed scope 2 revisions would place tighter temporal and geographic constraints on how corporate energy buyers reduce market-based emissions with RECs. This exposes an emissions reduction gap that most buyers haven't yet measured.

Under current practice, companies can retire RECs sourced from any grid within the country, at any time of the year, to reduce their reported market-based emissions. The proposed revisions tighten that ability significantly: certificates must come from the same grid, and in the same hour, as consumption. That shift makes annual, location-agnostic REC retirement, which is currently the dominant approach, insufficient.

The gap is largest in PJM and ERCOT, the same markets that are absorbing the fastest data center load growth. Both grids rely heavily on fossil generation during evening peaks when renewable output is low. Buyers with large footprints in these markets may look clean on paper under an annual matching approach, but considerably less so under an hourly approach.

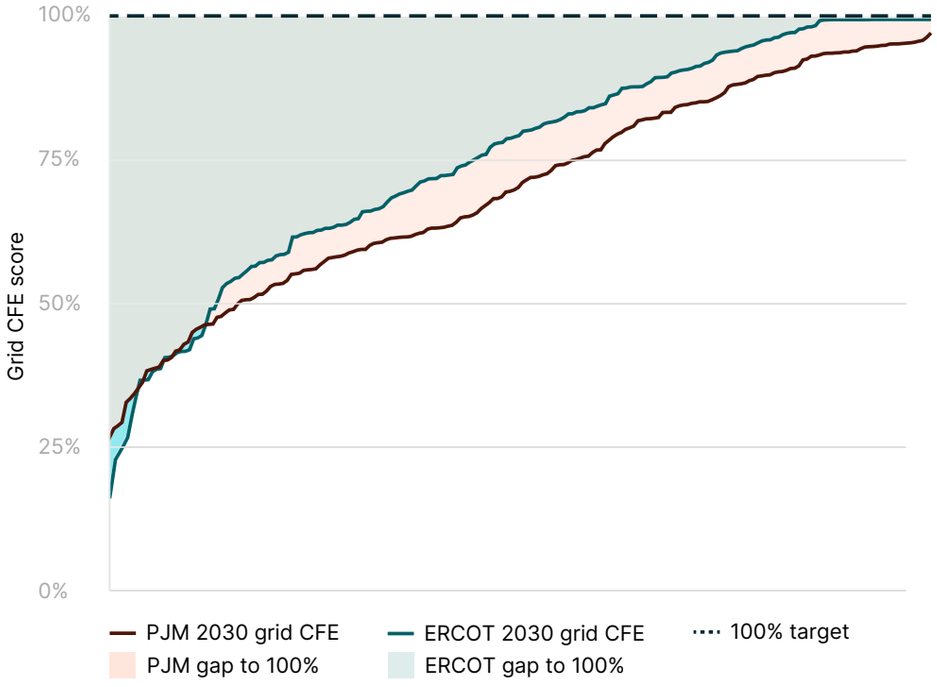


Map of power grids in the contiguous United States. Source: Carbon Direct, adapted from Midwest Energy Efficiency Alliance.

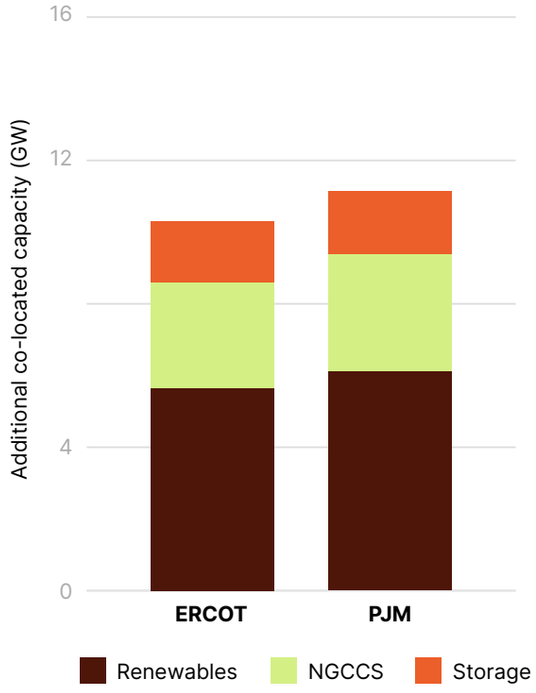
Closing that gap requires a fundamentally different procurement mix. Carbon Direct's modeling shows that achieving 100% hourly, carbon-free energy (CFE) for a representative 4 GW load requires 9.6 GW of additional clean capacity in ERCOT and 10.5 GW in PJM. Natural gas-fired power generation with CCS, long-duration storage, and expanded renewables all play significant roles in reaching that goal. For either grid to achieve 100% CFE in every hour, this mix is required. Companies that engage now can help shape the project pipeline available in their target markets by 2030 and lock in offtake terms, before competition for the most valuable projects tightens. Buyers who wait until the proposed scope 2 revisions are final will find shorter lead times, fewer partners, and less leverage.

Read more on the shifting playbook for corporate power procurement here →

2030 grid CFE score, hours ranked dirtiest to cleanest
Representative 2030 hours, sorted by CFE score (lowest to highest)



Additional capacity for 4 GW load at 100% CFE
Hourly CFE target 100%



Carbon-free energy (CFE) gap and additional capacity required to meet hypothetical CFE demand. Note: Natural gas with CCS is included in the 100% CFE stack, though it represents a ~95% (rather than fully zero) scope 2 emissions reduction. Modeling assumes technology costs as per the 2024 NLR Annual Technology Baseline-Conservative scenario.

Why behind-the-meter power emissions belong in scope 2

WHY IT MATTERS: Behind-the-meter procurement is accelerating as a solution to grid constraints. Companies that misclassify it now will face restatement risk and credibility exposure as voluntary and regulatory disclosure standards converge on stricter scope 2 treatment.

The emissions perspective, part 2: How does accurate carbon accounting apply to new, behind-the-meter power generation arrangements at data centers?

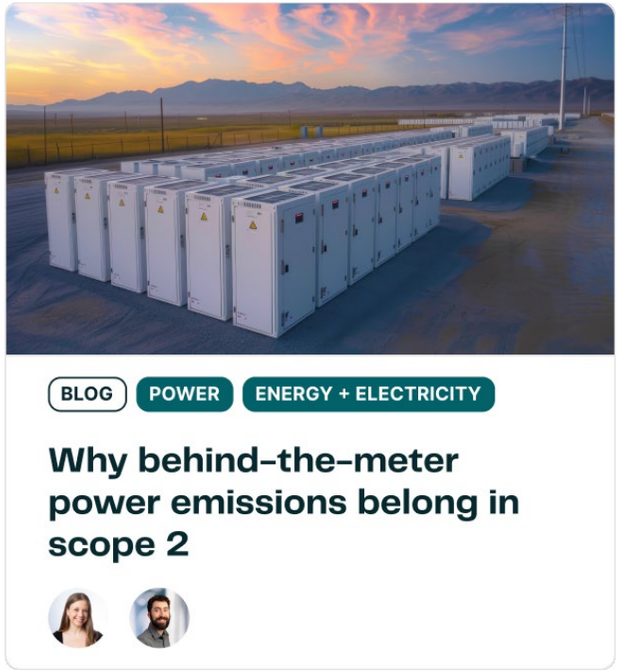
Julia Millot and Daniel Garcia, PhD

Large power users are turning to behind-the-meter (BTM) generation to bypass grid constraints. This involves sourcing electricity through a private line from a third-party-owned asset rather than the utility meter. Widespread confusion around this practice is causing companies to debate: Where does BTM electricity belong on GHG inventories?

The GHG Protocol's *Corporate Accounting and Reporting Standard* is unambiguous on this point. Section 5.4 of the *Scope 2 Guidance* explicitly classifies consumed BTM electricity as scope 2 emissions, regardless of ownership or contractual structure. Scope 3, Category 8 applies only to energy not already captured in scope 1 or 2, which rules it out here. The confusion stems from the physical pathway: because BTM power arrives via private line rather than a utility meter, some buyers and advisors have treated it as a purchased service rather than an energy input, thereby producing the wrong reporting.

The stakes of misclassification are significant. Companies that lock in the wrong reporting before a BTM contract closes risk compounding the error across multi-year agreements and facing restatement obligations as disclosure standards tighten. Voluntary and regulatory frameworks are converging on stricter scope 2 treatment. Inventories built on a misclassification today will need to be unwound later at greater cost and with greater reputational risk.

However, correcting this is straightforward. BTM electricity volume, and its associated emissions factor, belong within scope 2 reporting. Getting classification right requires legal, sustainability, and procurement teams to align before BTM agreements are signed, not after.



Read more on behind-the-meter accounting here →

Developer responses to data center opposition

The community perspective: How are local perceptions of data centers impacting success rates, and how can developers best respond?

Grant Gutierrez, PhD

WHY IT MATTERS: How data center developers respond to community opposition over the next 12–18 months will influence the scale of AI infrastructure buildout. Developing projects in a silo is unlikely to yield desired results. Retreating and repositioning preserves capital, but does not resolve the political conditions that drive community opposition. Making commitments to community benefits and clean energy raises the floor of what communities can negotiate for, but does not address rate impacts at the tariff level. Structural ratepayer protection addresses that gap but cannot substitute for community engagement. Developers that integrate meaningful community engagement and benefits planning with clean energy and rate protection will likely avoid intense public opposition.

Community opposition to AI data center development is no longer a peripheral concern. State legislatures are establishing comprehensive policy frameworks. Court rulings are setting precedents for community advocates. Project cancellations continued through the first half of 2026. The question now isn't whether opposition matters, but how the industry responds.

Three distinct response postures are emerging across pending and contested projects:

01 **Retreat and reposition**

02 **Engage through benefits and clean energy commitments**

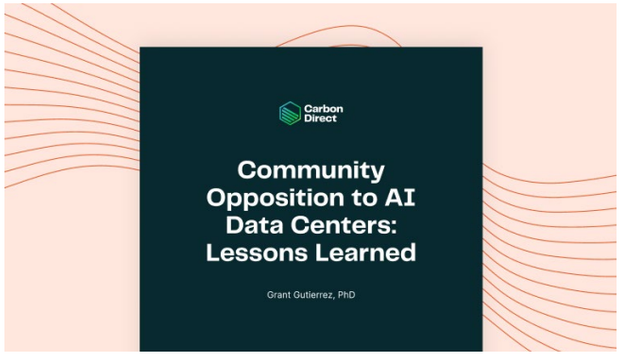
03 **Reposition as a grid asset through structural rate protection**

Retreat and reposition: A developer withdraws the application, then re-emerges under a different entity name, a different jurisdiction, or both. The pattern preserves capital and optionality, but does not resolve the underlying political conditions that drove the original opposition. It's a tactical move rather than a strategic one.

Engage through benefits and clean energy commitments: Community benefits agreements, voluntary infrastructure cost coverage, and paired clean generation are all part of the toolkit developers draw on when assuming this response posture. Google's Pine Island project in Minnesota set a visible standard, with 1,600 MW of new, clean power generation and full coverage of grid infrastructure costs. Other developers are following suit. However, this response does not address rate impacts at the tariff level.

Reposition as a grid asset through structural rate protection: Microsoft filed a Ratepayer Protection Tariff with the Nevada Public Utilities Commission earlier this year that would require large-load customers to pay their own infrastructure costs and contribute to existing system upkeep. Google announced a three-year, 100 MW virtual power plant deal with Voltus on June 2 of this year. The White House Ratepayer Protection Pledge has gathered hyperscaler signatories. These moves address residential rate impact at the structural level rather than offsetting it through benefits.

We are keeping an eye on what combination of responses best enables projects to advance through public processes rather than circumvent them.



Read more on Community Opposition to Data Centers here →

Disclaimer

This report is prepared by Carbon Direct for general informational purposes only. It is intended solely for the private use of its recipients, who are under no obligation to act upon the information provided herein. While we endeavor to keep the information up to date and correct, we make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability with respect to the publication or the information, products, services, or related graphics contained in the publication for any purpose. Any reliance you place on such information is therefore strictly at your own risk. The contents of this report, including but not limited to interpretations of policy, market trends, and any forward-looking statements, are based on the best judgment of Carbon Direct at the time of preparation and are subject to change without notice.

This report does not constitute financial, investment, legal, or tax advice, and should not be construed as such. Recipients are encouraged to consult with their own advisors before making any decisions based on the information contained herein. Neither Carbon Direct nor any of its affiliates, officers, directors, employees, or agents make any representation or warranty, express or implied, regarding the accuracy, completeness, or reliability of the information contained in this report, nor shall they be liable for any errors or omissions in the content or for any actions taken in reliance thereon. By accessing or using this report, recipients agree to release and hold harmless Carbon Direct and its affiliates, officers, directors, employees, and agents from any and all claims, damages, losses, liabilities, costs, and expenses arising out of or in connection with the use of or reliance on the information contained herein.

Carbon Direct reserves the right to amend, modify, or update this report at any time without prior notice. This disclaimer and any use of the report shall be governed by and construed in accordance with the laws of New York, and any disputes arising hereunder shall be subject to the exclusive jurisdiction of the courts of New York.

Carbon Direct may use AI-enabled tools to assist with research, data analysis, and drafting. Such tools are used to support personnel and do not replace professional judgment. Final analyses, conclusions, and recommendations are subject to human review prior to issuance.