Four Bets on the Grid

How Data Center Builders are Solving the Power Problem

May 2026

noreva

Data Centers and the Power Problem

The bottleneck in data center development has changed. When the AI boom started, we naively assumed it would be the construction or compute limitations. But now it's obvious that getting the data center synced to the grid, or otherwise connected to a power source, is the most challenging logistical aspect of building data centers.

Large-load interconnection queues have gone from a niche concern to the defining constraint. In ERCOT alone, large-load request volume grew by 4x just last year. And in PJM, the most recent long-term forecast concluded that the system may face unmet demand by 2030 even if every single project in the interconnection queue gets built on time.

In other words, load is now straining the system as much as the generator queue it's been chasing for years, if not more.

Against that backdrop, data center developers are keeping their foot on the pedal, coming up with creative solutions for these limitations.

Four strategies have emerged on the demand side. Each one is a different bet on how the power system will evolve to accommodate data center-driven load growth. Watching which market participants choose which strategy tells us a lot about how they think this all plays out.

1. Relocate – East to West

Instead of fighting for interconnection capacity in saturated markets like Northern Virginia, Santa Clara, Hillsboro, etc. find a place where power is available, or at least where the queue is shorter, and the utility is more willing to plan around you. In practice this has meant a decisive shift toward places in what Noreva is dubbing "America's Capacity Belt" along the of untapped natural gas in the Midwest and Southern Texas.

JLL's year-end 2025 market report forecasts that Texas is on track to overtake Northern Virginia as the world's largest data center market by 2030.

Load relocation is the cheapest and least technically ambitious of the four strategies, which makes it the default for anyone with the flexibility. There's no need to invent new commercial structures, procure equipment, or lobby regulators about cost allocation. Instead, developers are just moving to where the electrons are.

These builders anticipate that existing grids in primary markets will stay constrained into the long-term; they don't believe transmission is going to get built fast enough to change that, and they're willing to take on latency, workforce, and customer-proximity costs to get power *now*.

It also says something about their view of ERCOT specifically – that a market with relatively fast interconnection, abundant land, permissive siting, and a lot of potential for gas-fired capacity queued up is a better bet than a vertically integrated utility market where a co-operative relationship with the incumbent is table stakes.

2. Re-sequence, or Move the Goalposts

Keep the site but phase the compute ramp-up to match the pace at which power actually shows up. Builders and utilities enable re-sequencing through three tactics that often show up together: 1) phased interconnection agreements that start with partial firm service and upsize as upgrades finish; 2) higher queue priority in exchange for an obligation to reduce load during system stress; and 3) co-location with a specific generator, which physically anchors the load ramp to that plant's construction schedule.

FERC's December 18, 2025 order to PJM is the clearest institutional expression of the strategy: the Commission directed PJM to create new transmission pathways for co-located large loads, explicitly rewarding data centers willing to trade flexibility for speed. The DOE's parallel ANOPR, which FERC is working to finalize by June 2026, pushes the same idea nationally.

Re-sequencing is the strategy for developers who think the grid will deliver, just not on the timeline that a naive reading of their capex plan would require. The bet is that the regulated system works, and that a phased agreement preserves optionality on eventually getting cheap, reliable firm service, which is a better long-run position than owning your own power plant.

This strategy is taken by firms that have already committed billions to chips, land, and shell construction. Unfortunately for them, they will bear meaningful carrying costs when compute release is deferred.
Big tech appears to be tilting this way on most of their earlier announced projects because it’s the most pragmatic approach. This suggests a view that the timing gap is bridgeable through phased agreements and patience, rather than by building a parallel power system.

3. Load Islanding – Fully Behind the Meter

Disconnecting, or nearly so, from the existing power infrastructure is another strategy that gets flashy headlines. This path requires the firm to own generation, storage, and run the whole system independently, or at least designing it so that it could. In the purest version, the data center never touches the utility grid. But in more common proposals, there's a small grid tie for redundancy. The primary architecture is behind-the-meter generation (often natural gas, sometimes renewable) backed by very large battery banks.

Meta's Project Walleye is the clearest current expression of the strategy and the premium it commands. Walleye is the $3 billion construction loan Meta is raising to finance its 1 GW Prometheus campus in New Albany, Ohio – a first-of-its-kind deal in which lenders fund both the data center buildings and the Socrates behind-the-meter gas plant that powers them, in a single syndication.

Some back-of-envelope math on the project finance tells the story of the premium Meta is paying for this project. Williams is building Socrates for roughly $1.6 billion to deliver 400 MW, or about $4,000/kW of installed capacity. For comparison, Blackstone's recent acquisition of the 620 MW Hill Top plant in Pennsylvania (itself called out by industry analysts as richly priced) came in at about $1,600/kW. So, on any traditional $/kW basis, Meta is paying roughly 2x what a sophisticated PE acquirer pays for comparable operating capacity. Meta knows this – the deal is a statement that the premium is worth paying to skip the queue. Getting their data center energized on their own schedule, rather than PJM's, is worth more than the $/kW delta.
While the industry remains very skeptical of such projects actually getting built, let alone financed, one recent report identified 46 data centers, or roughly 30% of the planned U.S. pipeline, of planned behind-the-meter data center capacity.
This strategy is expensive and comes with a ton of risks that otherwise are handled by grids and utilities such as reliability and fuel-supply. If your business case relies on speed to market, but the utility tells you

it's a five-year queue, you've got to get creative, reach deep in your pockets, or do both – and a behind-the-meter architecture is where many are landing.

This is the strategy most tightly linked to the vibe-investing ecosystem around AGI. It requires a capital partner who will fund a several-hundred-million-dollar power plant on the theory that the compute behind it will be worth an order of magnitude more. Outside that ecosystem, and for all but the most aggressive projects, islanding is a pressure valve rather than a playbook.

4. Load Matching

Pairing an interconnection request with new capacity so that when load shows up on the system, so does the capacity and energy to serve it. The generation can be utility-owned and bundled into an Integrated Resource Plan (IRP, the long-term capacity plan utilities file with state regulators), IPP-owned and contracted to the data center, or co-developed. The defining feature is that the generation is visible to the system operator, goes into rate base or wholesale markets, and is shared by all grid customers. The system gets bigger to accommodate the new load, and the regulator can see and price both sides of the transaction. If you've heard of the BYOG (Bring Your Own Generation) mantra, its typically referencing this strategy.

The regulated version looks like Dominion Energy's 2024 IRP: 5,934 MW of new gas between 2030 and 2036, with its 2025 update increasing expected gas needs another 40%. The merchant version looks like Vistra's Permian Basin Power Plant expansion, announced September 2025, adding 860 MW of new ERCOT capacity sold into wholesale markets against data-center-driven demand.

Load matching is the strategy the regulated power system can actually price and approve. New gas added to the served grid solves for both capacity and energy in a form a PSC or ISO can sign off on, without forcing ratepayer advocates into unwinnable cost-causation fights. What it says about the future grid is the most conservative view of the four: that the regulated system, properly stressed and properly priced, will respond. Gas is the fuel of choice because it's the only dispatchable resource with a short enough lead time to matter – and because at the right price, it solves the system operator's problem at the same time it solves the data center's. That alignment is why, of the four strategies, it's the one regulators and grid operators most want to see more of.

What the Four Strategies Tell Us

No data center developer uses just one of these. The most sophisticated are running two or three in parallel – relocating to Texas while matching load to a new gas plant, for example, seems quite common. So is accepting a phased interconnection from the ISO while building behind-the-meter as a bridge.

Relocation says the geography of U.S. compute will shift permanently.

Re-sequencing says the existing system catches up, just slower than AI business models want.

Islanding says the system won't catch up fast enough, and speed is worth almost any premium.

Load matching says the regulated system works when priced right.

Contact:

Yaniv Yaffe – VP of Product

yaniv@noreva.ai

Commercial

sales@noreva.ai

Noreva

675 Third Ave, 31st Floor

NY, NY 10017