datacenters.builtfor.ai
AI INFRASTRUCTURE INTELLIGENCE

SECTOR RESEARCH · POWER & GRID Q3 2026 · Infrastructure Research

U.S. AI DATA-CENTER POWER · 2027 OUTLOOK

The Power Wall

The constraint on the AI buildout has moved from chips to power — and specifically to the speed of interconnection. The queue is enormous, but most of it will not be built on time. We favor markets and strategies that route around the wall.

EXECUTIVE DASHBOARD · U.S. AI DATA-CENTER POWER, 2027 OUTLOOK

OUR CALL

Constrained

Long behind-the-meter & firm power (gas, geothermal, nuclear) and the spillover developers that control energized, interconnected sites. Underweight greenfield capacity whose timeline depends on clearing the queue.

BY THE NUMBERS

In the interconnection queue	1,745 GW
Completion-adjusted (~11%)	198 GW
Operational today	562 GW
Largest queue	Texas
Of queue is solar / wind / storage	83%
States with moratorium bills	18
ISO/RTO markets tracked	7

Our view

The bottleneck in the U.S. AI data-center buildout is no longer the supply of chips or the willingness of capital — it is electricity, delivered firm and on time. The headline numbers make the constraint look soluble: our tracking of the national interconnection queue shows **1,745 GW** of generation and storage seeking to connect, roughly **3×** the **562 GW** operating today. That framing is misleading. Interconnection queues have historically converted to operating capacity at only about **11%**, which puts the realistic near-term addition closer to **198 GW** — a fraction of the wall of announcements. The gap between what is queued and what actually energizes is the single most important variable for siting, financing, and timing decisions through 2027.

- **Power, not silicon, sets the pace.** Independent infrastructure analysts now rank power availability as the first, second, and third site-selection criterion (Dell'Oro Group). Compute can be ordered faster than it can be energized.
- **The queue overstates supply.** At a ~11% historical completion rate, the queue is a measure of ambition, not delivery. Transmission, transformers, and turbines — not interest — are the gating items.
- **Tier-1 hubs are hitting the wall first.** The densest markets face transmission saturation and rising political resistance; demand is already migrating to spillover metros with available headroom.
- **What's queued isn't firm.** About **83%** of queued capacity is solar, wind, and batteries — clean and fast, but not 24/7. The firm-power gap is being bridged by behind-the-meter gas now and, later, by nuclear and geothermal.
- **Policy is the new swing factor.** Our policy tracker shows moratorium-type bills active in **18** states and large-load-tariff / ratepayer measures in **6**—a rising tax on the simplest siting plans.

The call: Constrained. Through 2027 we expect demand to outrun firm, on-time supply in the marquee markets, and we would position accordingly. We are **long behind-the-meter and firm power** — the gas, geothermal, and nuclear strategies that bypass the interconnection queue rather than wait in it — and **long the developers and spillover markets that already control energized, interconnected sites**, with the miner-pivot names chief among them. We are **underweight greenfield capacity whose timeline depends on clearing the queue**, which is where the wall bites hardest. The scarce asset, throughout, is power that can be energized on schedule; the rest of this report is a map of where to find it — and where not to look.

From chips to power: the demand case

The demand signal is unambiguous and still accelerating. Frontier labs are contracting power and compute at a scale that was inconceivable two years ago: Anthropic has reportedly been pursuing on the order of **10 GW** of capacity over the coming years (The Information), and in April 2026 announced a **3.5 GW** next-generation TPU agreement with Google and Broadcom, with deployments beginning in 2027 (company announcement). Independent forecasters put data-center capital spending on a path toward roughly **$1 trillion** a year — en route to about **$1.7 trillion by 2030** (Dell'Oro Group). The adoption data underneath these commitments continues to compound — third-party web-traffic measures show total generative-AI usage rising double digits month-over-month into early 2026, with newer entrants growing fastest.

For a power analyst, the precise share splits matter less than the aggregate: every incremental gigawatt of frontier compute is a gigawatt of round-the-clock electricity demand that did not exist in the utility planning base a year earlier. That is why the binding question has inverted. The industry spent 2023–24 asking whether it could get enough accelerators; it is spending 2026 asking whether it can get enough power, in the right place, in time. The rest of this report is about the supply side of that question.

The queue illusion: a wall that won't all clear

The interconnection queue is where the demand story meets physical reality — and where it breaks down. Across the seven ISO/RTOs and the non-ISO utilities we track, **1,745 GW** of generation and storage is actively seeking to connect. Taken at face value, that is several times the operating fleet and far more than the buildout needs. But queues are a register of *requests*, not of delivery. Historically, only about **11%** of requested capacity has ever reached commercial operation; the rest is withdrawn as projects fail to clear studies, secure financing, or survive cost re-allocations. Apply that completion rate and the queue's 1,745 GW resolves to roughly **198 GW** of credible near-term supply. The wall is real; most of it is a mirage.

That ~11% is a blended, historical figure, and it cuts in one direction worth flagging: the queue has been dominated by speculative, under-capitalized projects — predominantly renewables — that withdraw the moment interconnection-upgrade costs surface. Hyperscaler- and lab-backed projects, which can absorb those costs to lock in power, should clear at materially higher rates, so the effective completion rate at well-funded nodes is likely above the average — plausibly two to three times the blended figure, on the order of 25–35% at the best-capitalized nodes (our analytic estimate, not an observed rate). That advantage is widening as hyperscalers increasingly fund their own network upgrades — paying for the substations, gen-ties, and even transmission segments that a developer would normally wait on — which can pull their nodes through faster than the queue average. The catch is that capital cannot manufacture a transformer on demand: the physical supply chain enforces an absolute ceiling on deployment velocity regardless of balance-sheet scale, so even self-funded nodes hit the equipment wall. The landscape is consequently bifurcating into a capitalized, high-velocity core and a broader queue trapped behind systemic equipment bottlenecks.

The constraint, in other words, is not developer interest — it is throughput. The physical inputs to interconnection are on multi-year lead times that the demand cycle cannot wait out: high-voltage transformers run roughly three to four years, and large-frame gas turbines are effectively sold out into 2029–2030, on roughly five-year backlogs (GE Vernova, Siemens Energy, Mitsubishi). A queue position is not power; it is a place in line behind those constraints.

Exhibit 1 Active interconnection queue by grid region (GW)

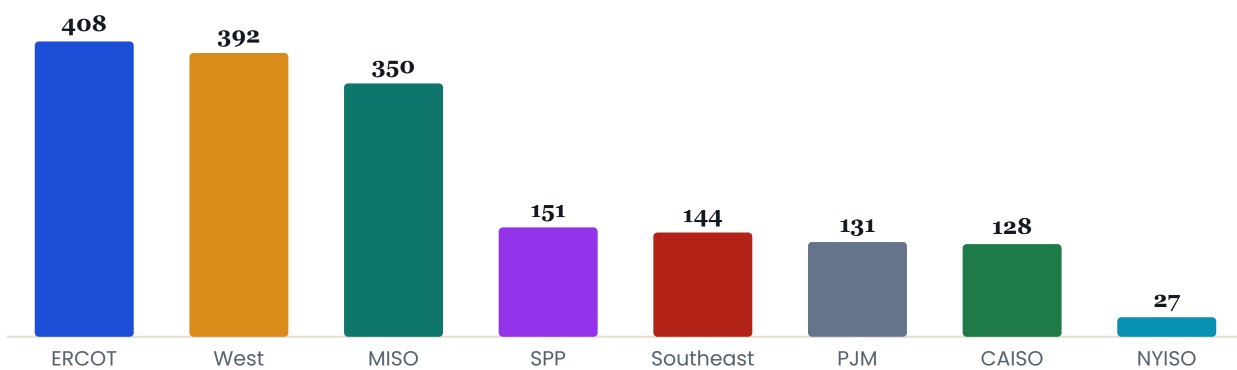

Source: Lawrence Berkeley National Laboratory (Berkeley Lab), Electricity Markets & Policy, "U.S. Interconnection Queue Data through 2025." (CC BY 4.0), as of 2026-06-18. datacenters.builtfor.ai analysis.

Exhibit 2 Queue by region — active vs. completion-adjusted

ISO / RTO	ACTIVE QUEUE (GW)	COMPLETION-ADJ. (~11%)
ERCOT	408	46
West	392	44
MISO	350	40
SPP	151	17
Southeast	144	16
PJM	131	15
CAISO	128	14
NYISO	27	3

Source: Lawrence Berkeley National Laboratory (Berkeley Lab), Electricity Markets & Policy, "U.S. Interconnection Queue Data through 2025." (CC BY 4.0), as of 2026-06-18. datacenters.builtfor.ai analysis.

Where the power is — and isn't

Queued capacity is highly concentrated. **Texas** alone accounts for roughly **437 GW** of active requests — more than the next several states combined — and the queue clusters in ERCOT, the West, and MISO. (The 437 GW is the state of Texas; the ERCOT *region* figure in Exhibit 1 differs because parts of the state sit outside ERCOT.) That concentration is not the same as availability. The densest existing data-center markets are precisely where transmission headroom is scarcest: hosting capacity in much of the West is effectively exhausted, and one regional assessment put the transmission investment required to serve even a modest load increase at roughly **$45 billion** over a decade (Energy Strategies). The Tier-1 hubs that defined the last cycle are the ones now hitting the wall.

The investable consequence is geographic. As marquee markets saturate — on the grid and in the statehouse — incremental demand is migrating to *spillover* metros that still have interconnection headroom and a willing political environment. We expect that rotation to be one of the defining features of the 2026–27 siting map, and it is the core of our preference for spillover markets over Tier-1 hubs.

Concretely, our spillover screen rewards markets that pair genuine interconnection headroom and shorter study queues with an accommodating policy posture and access to land, water, and gas — not the cheapest power, but power that can actually be energized on the buildout's timeline. That points to the central-Ohio corridor (the Columbus/AEP market), the Dakotas and the broader SPP/MISO interior where miner-pivot developers already hold energized sites, West and South Texas inside ERCOT, and selectively the Mid-South and Mountain West (Atlanta; Phoenix, the latter with a water asterisk — a permitting and reputational constraint more than a physical one, since modern closed-loop liquid cooling uses far less water than the legacy evaporative systems, but in a desert the optics and siting approvals still bite). The common disqualifier is the inverse of cheap power: a multi-year interconnection queue with no behind-the-meter path around it.

The generation gap: what's queued isn't firm

Even where capacity does clear, it does not necessarily solve the data center's problem, because AI load wants firm, 24/7 power and the queue is dominated by intermittent resources. About **83%** of queued capacity is solar, wind, and batteries; gas — the most straightforward source of round-the-clock generation — is roughly **240 GW**. Storage shifts solar into the evening but does not, at current durations, underwrite a continuous gigawatt-scale load. The result is a firm-power gap that the headline queue figure conceals.

The market's near-term answer is to stop waiting for the grid. Behind-the-meter generation — predominantly on-site natural gas — has moved from a fringe approach to something on the order of a third to a half of new developers in the West in roughly a year (Energy Strategies), precisely because it sidesteps the interconnection queue. It is not free: on-site build can require materially more capital than a grid connection, it still depends on firm natural-gas pipeline capacity — itself permitting- and backlog-constrained in some basins — and it raises emissions and gas-demand questions. The firmer, cleaner answers — advanced nuclear and geothermal — are real but late. Utility plans now include gigawatt-scale nuclear and geothermal commitments, and early data-center PPAs exist (for example, a 13 MW, 20-year geothermal contract between Ormat and Switch, with deliveries beginning around 2030), but scale nuclear is generally not expected before the mid-2030s.

For the 2027 window, then, the firm-power swing factor is gas — but not the heavy-frame turbines whose multi-year backlogs, sold out into 2029–2030, we flagged earlier. The gas that actually energizes by 2027 is the fast-deploying kind: reciprocating-engine gensets and aeroderivative units that can be sited in roughly twelve to twenty-four months, plus turbine capacity ordered well ahead of need. That distinction — fast modular gas now, heavy-frame and firm-clean capacity later — is what lets gas bridge the 2027 gap rather than miss it, and it is why behind-the-meter is a near-term lever and not a 2030 one.

Exhibit 3 What's in the queue: firm vs. intermittent

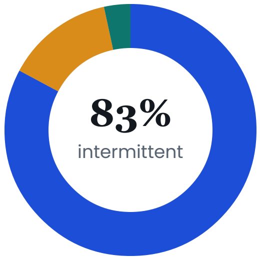

Solar / wind / storage	83%
Gas (firm, dispatchable)	14%
Other (nuclear, geothermal, etc.)	4%

Source: Lawrence Berkeley National Laboratory (Berkeley Lab), Electricity Markets & Policy, "U.S. Interconnection Queue Data through 2025." (CC BY 4.0). datacenters.builtfor.ai analysis. Shares are of the 1,745 GW active queue. "Intermittent" = solar, wind & batteries — note that storage MW is not continuous firm generation. Gas is the queue's main firm, dispatchable resource.

The players: who builds the power — and who wins

The buildout has organized into a stack, and the economics differ sharply by layer because each owns a different scarce asset. At the top sit the demand-side players — the frontier labs and hyperscalers (OpenAI, Google, Microsoft, Meta, Anthropic, xAI). Their investment-grade balance sheets are what underwrite the entire chain: a 15-year lease is only bankable because the offtaker is creditworthy, which is why developers compete hardest for hyperscaler- and lab-backed contracts.

The most distinctive dynamic is one layer down. A cohort that spent the last cycle mining bitcoin — Cipher Mining, Hut 8, TeraWulf, Riot, CleanSpark, Bitdeer — is converting its scarcest possession into AI capacity: *energized, already-interconnected sites.* That is precisely the asset the queue cannot supply on demand, so these developers can deliver power-ready capacity years faster than a greenfield project waiting in line. The market is bifurcating them into a "Tier 1" that has signed investment-grade leases — Applied Digital's CoreWeave and hyperscaler campuses, and the roughly one gigawatt (on a critical-IT-load basis) of Fluidstack-intermediated, Anthropic-backed capacity spread across Hut 8 (245 MW at River Bend, a 15-year, $7 billion lease), Cipher (168 MW critical IT load, up to 244 MW gross, at Barber Lake, Texas), and TeraWulf (168 MW in Texas plus 360 MW at Lake Mariner), each with Google financial backing — and a "Tier 2" that controls energized power and land but has not yet contracted it, leaving it exposed to merchant-power and lease-up risk until it does. The structural edge is the same one the report keeps returning to: control of connected power beats a place in the queue.

The equipment layer is where scarcity turns into pricing power. Power-distribution and thermal-management suppliers — Vertiv, Eaton, Schneider Electric, GE Vernova — sit on the critical path, with order backlogs up sharply year-over-year and delivery times stretching beyond twelve months (company earnings). Their advantage is a near-oligopoly on lead-time-constrained gear, amplified by the technology transition toward liquid cooling and 800-volt DC architectures that rewards whoever leads it (Dell'Oro Group). Power-electronics and storage entrants — Enphase, Fluence, Energy Vault — are pushing into the same medium-voltage and behind-the-meter opportunity from the edges.

Two layers decide whether any of it gets built on time. Firm-generation providers win because firm, 24/7 power is the scarce product: gas developers today, and geothermal (Ormat, Fervo) and nuclear later. And construction is itself a bottleneck — only a handful of U.S. EPCs (on the order of seven to ten) can execute 500 MW-plus projects, which hands firms like Quanta, MasTec, and Primoris pricing power and makes them a gating constraint. **The throughline: advantage accrues to whoever controls the scarcest link** — energized sites, firm generation, lead-time-critical equipment, or execution capacity. The losers are greenfield projects waiting in the queue and intermittent-only generators that cannot serve a continuous load.

Exhibit 4 The AI data-center power value chain

LAYER	REPRESENTATIVE PLAYERS	SOURCE OF ADVANTAGE
Frontier labs & hyperscalers	OpenAI, Google, Microsoft, Meta, Anthropic, xAI	Investment-grade demand underwrites the chain
Capacity developers (miner-pivot)	Applied Digital, Cipher, Hut 8, TeraWulf, Riot, CleanSpark, Bitdeer	Control energized, interconnected sites → power-ready fastest
Neoclouds	CoreWeave, Fluidstack	Aggregate GPU compute between developers and labs
Power & cooling infrastructure	Vertiv, Eaton, Schneider, GE Vernova	Oligopoly on lead-time-critical gear; cooling / 800VDC shift
Power electronics & storage	Enphase, Fluence, Energy Vault	Medium-voltage & behind-the-meter entrants
Firm generation & PPAs	Gas developers; Ormat, Fervo (geothermal); nuclear	Firm 24/7 power is the scarce product
EPC & construction	Quanta, MasTec, Primoris	~7–10 firms can build 500 MW+; scarcity = pricing power

datacenters.builtfor.ai analysis. Representative names, not recommendations.

Policy & ratepayer risk

The final constraint is political, and it is tightening. As data-center load growth shows up on residential utility bills and in local land-use fights, legislatures are responding. Our policy tracker, which classifies data-center-specific legislation across all 50 states, finds moratorium-type measures active in **18** states and large-load-tariff or ratepayer-protection bills in **6** — measures designed to ensure that hyperscale loads, rather than ordinary ratepayers, bear the cost of the grid upgrades they require. For developers, the effect is to raise the cost and lengthen the timeline of the simplest siting plans, and to reward markets whose policy posture remains accommodative. Policy is no longer a background variable; it is part of the underwriting.

Scenarios

We frame the 2027 outlook around the single most important uncertainty: how much of the queue actually energizes, and how fast firm power fills the gap. Our base case assumes a partial queue clear plus accelerating behind-the-meter gas; the bull case adds interconnection reform and earlier firm-capacity delivery — and its most likely catalyst is now visible, in FERC's June 18, 2026 Section 206 show-cause orders pressing all six RTOs to speed large-load tariffs; if those bite quickly, the wall erodes faster than 2027. The bear case is a stalled queue and slipping gas timelines.

Exhibit 5 2027 data-center power scenarios

SCENARIO	DRIVER	OUTCOME
Bear	Queue stalls; gas + transmission slip	Supply-short
Base	Partial queue clear + behind-the-meter gas	Constrained
Bull	Interconnection reform + early firm capacity	Balanced

datacenters.builtfor.ai analysis. Illustrative; see methodology.

Risks to the view

- **Faster reform.** Large-load interconnection fast-tracks are accelerating — on June 18, 2026 FERC issued Section 206 show-cause orders requiring all six RTOs/ISOs to justify their current large-load tariff treatment or file reforms — and could clear the queue faster than the historical completion rate implies.

- **Demand air-pocket.** A slowdown in frontier-model capex or an efficiency step-change would relieve the power constraint outright.
- **Gas execution.** The behind-the-meter bridge depends on turbine and pipeline timelines that are themselves stretched.
- **Policy whipsaw.** Federal incentives or pre-emption could offset state-level restriction — or tighten it further.

Methodology & disclosures

Forecasts, scenarios, and interpretation are the analyst's own. The interconnection-queue figures are generated from Lawrence Berkeley National Laboratory's U.S. interconnection-queue dataset (CC BY 4.0). **Our queue figure is a filtered subset**: we count only requests with an active status and their primary nameplate capacity, which is narrower than LBNL's headline total of more than 2,000 GW of generation and storage (that figure also includes additional in-progress statuses and the storage leg of hybrid projects). The completion rate is computed from the dataset's own operational-versus-withdrawn history. Resource shares (Exhibit 3) are taken against this active total, so the intermittent share, the gas figure, and the slice percentages reconcile. The generation-mix and demand figures are from the U.S. EIA; the policy classification is from our own all-50-state tracker (built on LegiScan data, CC BY 4.0), which is published and auditable on our policy page. External market facts are attributed to their primary sources; figures presented as estimates — including the 25–35% well-funded-node completion range — are labeled as such. This report is informational only and is not investment advice.