



April 5, 2026

Data Centers and Power: Speed, Reliability, and Tradeoffs

Data center-driven load growth is accelerating across U.S. power markets, raising a seemingly simple question: how will these facilities quickly secure reliable power at scale?

In reality, the answer is anything but straightforward. Uncertainty around the magnitude and timing of demand complicates long-term planning, while independent system operators (ISOs) face real limits when it comes to rapidly bringing multi-gigawatt levels of new capacity online. At the same time, emerging winter reliability risks in numerous power markets are challenging traditional assumptions about which resources can reliably serve load and how often critical conditions will arise, while soaring capacity prices are driving affordability concerns and increasing political scrutiny. Against this backdrop, data center operators are being forced to confront a fundamental tradeoff between what they want – always-on, dedicated power – and what they realistically need to move quickly.

In a recent webinar hosted by Hannam & Partners, Dr. Brent Nelson, Senior Managing Director of Markets and Strategy at Ascend Analytics, outlined how developers, operators, corporates, and investors can navigate these constraints while shifting toward more flexible and pragmatic approaches to power procurement that better align with the realities of today's energy markets.

Key Takeaways

- Data center-driven demand forecasts range widely, making it difficult for ISOs and developers to plan effectively. Supply is lagging demand growth, with most markets unable to add capacity fast enough to meet the appetite for demand. This mismatch will likely create persistently tight conditions and shift where data centers are built, as developers prioritize locations where power is easiest to access.
- Winter poses increasingly serious reliability challenges as cold weather events create long-duration overnight peaks that renewables and short-duration storage cannot easily serve. But winter risks also raise questions about the reliability of both grid-connected and behind-the-meter (BTM) gas. However, because extreme winter events are relatively rare, many data centers may be overestimating curtailment risk and overbuilding for it.
- While 24/7 power availability is a central goal for data centers and hyperscalers, the top priority remains speed to market. In considering strategies such as 'bring your own power,' obstacles related to EPC capacity, permitting, equipment availability, and community support may force a shift away from the ideal of always-on, dedicated gas power toward more pragmatic solutions.
- Flexible approaches to securing power, such as combining grid power with BTM storage or paying other loads to curtail during peak events, can improve speed to market while maintaining reliability.

- Storage has become faster to build than gas, is supported by tax incentives, and can effectively handle summer peaks and short-duration reliability events. While it does not fully solve long-duration winter challenges, storage should be considered as an alternative to gas as a central part of any practical, speed-focused data center power strategy.

How Large is the Scale of Data Center-Driven Demand Growth?

After 20 years of mostly flat or declining demand, load growth is accelerating in power markets across the US. As shown in **Figure 1**, every ISO is forecasting significant load growth: the only uncertainties are how much and how fast.




Figure 1. Forecasted peak demand growth by 2030 (GW). Source: Ascend Analytics market forecasts

With data centers specifically, the exact magnitude of load growth is almost a complete unknown, with forecasts ranging from bullish to ludicrous. Currently, the spectrum of projected outcomes is too wide for energy market stakeholders to effectively plan for managing data center-driven load growth.

What is certain in this context, though, is that energy markets have historically been incapable of adding capacity at a rate that would be sufficient to meet current demand projections. For example, ERCOT's most recent adjusted load growth scenario called for an 8 GW per year growth in peak demand, as shown in **Figure 2**.

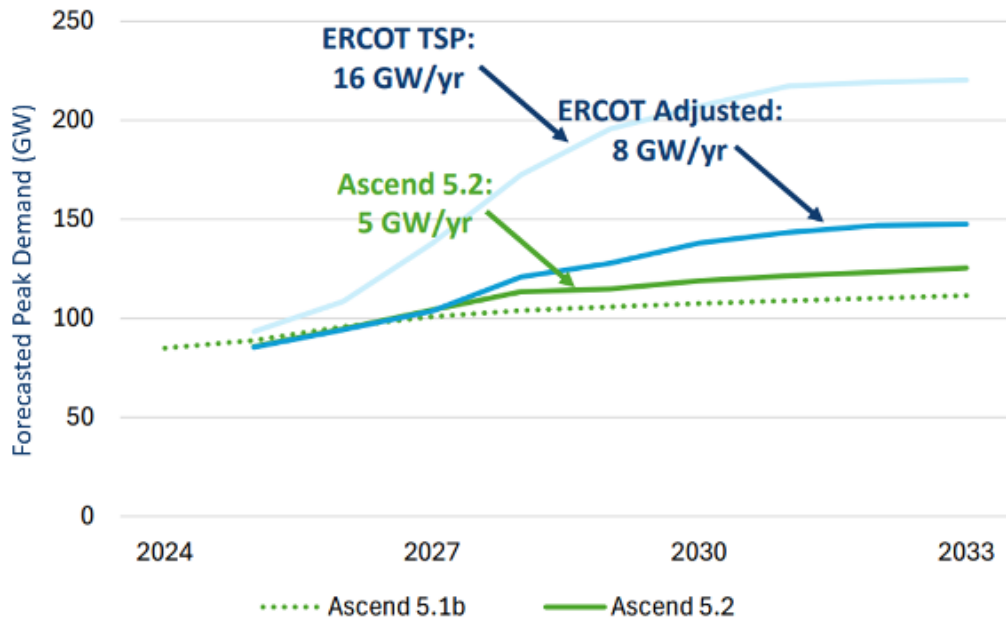


Figure 2. ERCOT forecasted peak demand (GW)

However, as shown in **Figure 3**, ERCOT has historically only been able to add a maximum of 2–3 GW per year of winter peak capacity (ELCC-adjusted). In most markets, the pace of supply additions will restrict demand growth and lead to perpetually tight conditions.




Figure 3. ERCOT historical annual capacity additions (GW)

Ultimately, demand and supply will both go where supply is easiest to add. Historically, industry siting practices have followed sources of inexpensive power. With the current concept of cheap power coupled with the notion of zero-carbon power, data centers are moving toward locations that offer low-cost, low-emissions resources. Additionally, the movement of some data center supply to behind the meter (BTM) will decouple the impacts of load growth from what eventually materializes in wholesale markets.

How Should Data Centers Account for Winter Reliability Risks?

Data centers evaluating power reliability often focus on summer peaks, but winter conditions can pose a greater risk to grid reliability. As illustrated in **Figure 4**, summer demand spikes are typically short, occur during daylight hours in sunny conditions, and as a result can be met with solar generation and short-duration batteries. Winter peaks, however, often last through long, cold nights when solar output is unavailable and wind production may be limited, leading to prolonged high demand periods, as illustrated in **Figure 5**. These long-duration winter peak events require dispatchable resources capable of delivering power for many consecutive hours, which makes reliability planning more complex.

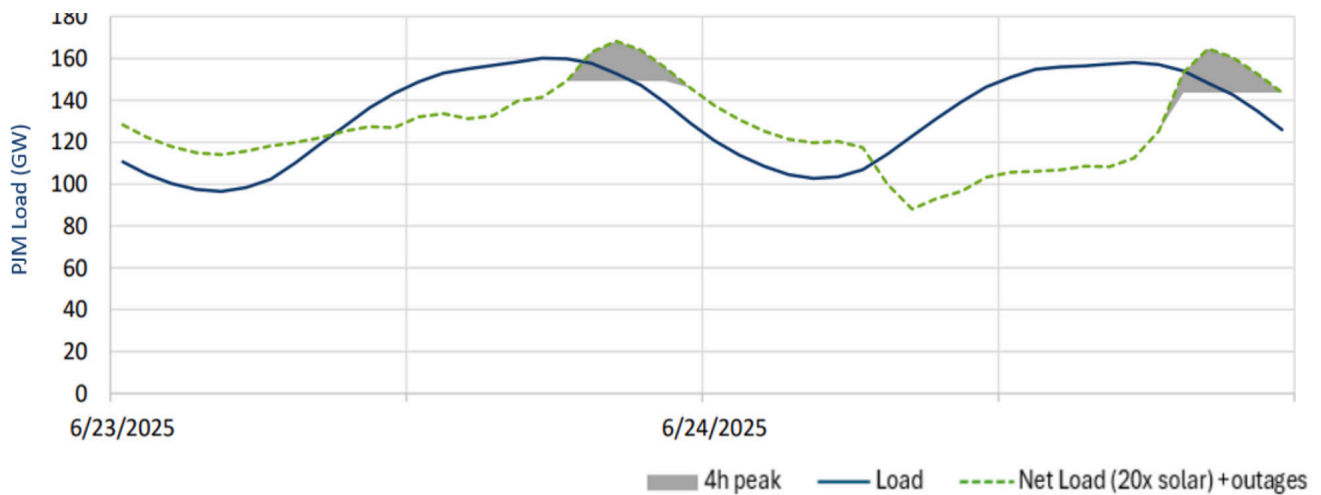


Figure 4. PJM summer load during 2025 heat wave (GW)

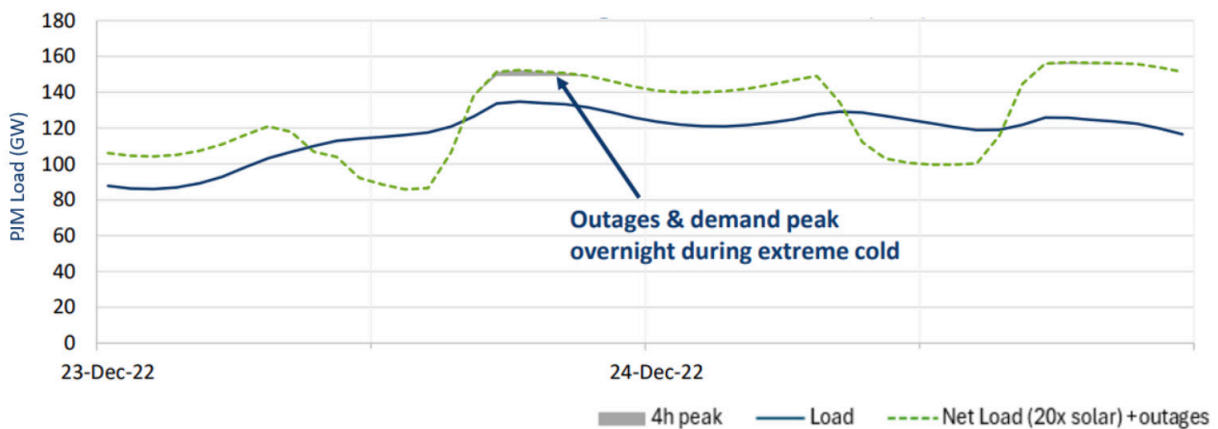


Figure 5. PJM winter load during winter storm Elliot (GW)

At the same time, cold weather increases the risk of forced outages at thermal plants and can disrupt fuel supply for natural gas generators. For data center operators, this dynamic should raise questions about the reliability of both grid-connected and BTM gas resources during the severe winter events when load curtailment will be most likely.

As performance during winter conditions becomes increasingly important, preventative maintenance, weather hardening, and secure fuel supply will be critical. As data centers rush to secure capacity, it is also important to recognize that extreme winter tightness tends to occur only during rare events: many data centers may be overestimating the impacts of curtailment risk under interruptible load tariffs.

What Strategies Should Data Centers Use to Manage Reliability Risks?

Data center developers often approach power procurement with a primary objective to secure 24/7 electricity with no interruptions. In practice, however, the fastest and most cost-effective path to reliable power may look different from this ideal. As data center demand accelerates, developers are encountering real-world constraints across the energy supply chain, from limited engineering, procurement, and construction (EPC) capacity to growing uncertainty around supply availability for new natural gas generation. Natural gas projects can also face permitting delays, air quality standards, and community opposition.

These constraints create a gap between what data centers want – guaranteed constant power – and what they may actually need to deploy quickly (with speed to market being the top priority) and reliably. In many cases, the opportunity cost of waiting for dedicated infrastructure can be significant.

Those operators who can leverage flexible reliability strategies stand to move far more quickly than those who don't. One potential approach for data centers involves combining a grid connection with interruptible load tariffs and BTM battery energy storage systems (BESS). Under this model, data centers would draw most of their power from the grid but use on-site batteries to cover short curtailment events during the summer, and sacrifice some operations during any curtailment hours that last beyond the battery duration during rare severe winter events where the grid runs short on supply.

More flexible load management strategies may also help. In certain conditions, data centers could effectively buy reliability by compensating other loads to ramp down during tight periods. This kind of system coordination would allow more data centers to connect without dramatically increasing peak demand, potentially improving both affordability and speed to market while still delivering a high level of operational reliability.

Does Clean Energy Still Matter for Data Center Customers?

Clean energy remains important to data center customers, though the dominant driver is speed – both in terms of bringing new facilities online and delivering compute. With demand for AI and cloud services accelerating rapidly, the cost of falling behind can be enormous. As a result, developers are often willing to prioritize power availability and deployment timelines ahead of sustainability goals in the near term.

That shift does not mean clean energy has fallen off the agenda for hyperscalers and data center operators: many continue to maintain long-term commitments to carbon reduction and 24/7 clean energy procurement. However, renewable resources alone do not resolve the most immediate reliability challenges, especially during winter.

In practice, this means that many data center operators and hyperscalers are likely to utilize any reliable power source that enables faster deployment, while continuing to pursue carbon-free goals over the long term. This consideration may tilt the scales further toward storage as limitations on gas generation become more apparent.

Is 'Bring Your Own Power' a Plausible Option for Data Centers?

The idea of data centers 'bringing their own power' has gained attention as electricity demand from hyperscale facilities grows. In theory, building BTM generation such as dedicated gas plants could give operators greater control over reliability and reduce dependence on the grid.

In practice, however, making this approach work at scale is far more complicated. Operating generation assets requires competencies that most technology companies do not traditionally possess. Developing gigawatt-scale power infrastructure means navigating complex permitting requirements, securing (increasingly scarce) EPC resources, managing ongoing operations and maintenance, and complying with environmental regulations.

Fuel availability is another challenge. Even if a data center builds dedicated gas generation, it still depends on reliable fuel delivery, especially during those extreme weather events when gas supply disruptions are most likely.

Ultimately, 'bring your own power' could be possible for data center developers and hyperscalers. However, it is not as straightforward as is often assumed.

Why Should Data Centers Consider Battery Energy Storage Systems?

Historically, gas turbines have been viewed by data center developers as the go-to solution for firm, on-site generation. That assumption is changing rapidly. Turbine supply shortages have driven capital costs sharply higher while delaying project timelines, and reliability expectations for new gas turbines have been downgraded as ISOs worry more about winter peaks.

As shown in **Figure 6** for PJM, storage can easily outcompete newly built simple-cycle gas turbines and is increasingly competitive with combined-cycle gas plants in many markets. Even longer-duration storage options are becoming viable alternatives as turbine prices remain elevated. Storage also benefits from investment tax credits and faster deployment timelines, giving it an additional advantage for developers racing to bring new data center capacity online.




Figure 6. Capacity resource Net CONE waterfall in Ascend's PJM 5.3 forecast release (\$/kW-yr)

Additionally, the resources most likely to be built in the near term are already sitting in interconnection queues, and storage projects far outpace gas projects in that regard. With constrained turbine supply slowing new gas builds, batteries and other flexible technologies can be better positioned to meet rapidly growing load.

From a reliability perspective, batteries can solve a large portion of the near-term challenge. A few hours of storage can effectively manage summer peak conditions and short-duration curtailments, especially when paired with grid power. The more difficult reliability challenge involves winter events, when long-duration overnight demand can stretch beyond typical battery discharge windows, though these events are relatively rare.

Even so, storage presents a strong option for ensuring data center reliability. In a world of expensive turbines, uncertain fuel supply, and pressure to move quickly, batteries are far more competitive than many data center developers may realize.

Interested in Learning More About How Ascend Analytics Supports and Informs Data Center Investment Decisions?

AscendMI™ (Ascend Market Intelligence) delivers proprietary power market forecasts that have been trusted in hundreds of projects and resource planning activities, supporting over 300 gigawatts and \$25 billion in project assessments. Please contact us to learn more.

About Ascend Analytics

Ascend Analytics is the leading provider of market intelligence and analytics solutions for the power industry.

The company's offerings enable decision makers in power development and supply procurement to maximize the value of planning, operating, and managing risk for renewable, storage, and other assets. From real-time to 30-year horizons, their forecasts and insights are at the foundation of over \$50 billion in project financing assessments.

Ascend provides energy market stakeholders with the clarity and confidence to successfully navigate the rapidly shifting energy landscape. Visit us at ascendanalytics.com

