



The Race for Power:

How Behind-the-Meter
Solutions Are Accelerating
Data Center Deployment

As hyperscalers, neoclouds and AI-native companies race to deploy increasingly large graphics processing unit (“GPU”) clusters, reliable power continues to represent the defining constraint on data center development. In response, Independent Power Producers (“IPPs”) have expanded behind-the-meter (“BTM”) solutions to accelerate deployment timelines, transforming BTM resources from a contingency measure into a strategic imperative for AI infrastructure growth. Drawing on recent experiences across data center, power and capital markets, FTI Consulting experts explore the business models, financing structures, value creation opportunities and risk-allocation challenges driving BTM development.

AI’s Growing Power Challenge

Driven by unprecedented demand for AI computing, the North American data center market has expanded at an estimated 37% compound annual growth rate since 2021.¹ By mid-2026, the region is expected to support approximately 70 GW of operational and pre-leased capacity, with another 17 GW under construction and more than 200 GW advancing through various stages of planning and development.²

Growth in demand has significantly outpaced the industry’s ability to build new capacity. As a result, vacancy rates have fallen from approximately 6% in 2021 to just 1% in 2026.³ At the same time, development timelines have stretched to five to ten years, as projects encounter permitting hurdles, grid interconnection backlogs, transmission constraints and persistent supply chain bottlenecks.

concentrations of data center capacity, benefiting from extensive fiber networks, cloud connectivity and decades of infrastructure investment. However, increasingly constrained access to power is limiting the pace of growth in these regions. As securing power has become the primary gating factor for new development, hyperscalers and AI infrastructure providers are expanding into emerging markets where large-scale power resources can be deployed more quickly and reliably. New approaches to power generation, transmission and procurement are emerging as critical enablers of data center growth.

Why Is the Market Turning to BTM Solutions?


The U.S. electric grid is confronting one of the most consequential demand growth cycles in modern history, with mounting resource adequacy challenges across key markets including PJM Interconnection LLC (“PJM”), Electric Reliability Council of Texas (“ERCOT”), Midcontinent Independent System Operator (“MISO”) and parts of the Western Interconnection.

In its latest Long-Term Reliability Assessment, North American Electric Reliability Corporation (“NERC”) forecasts 224 GW of summer peak load growth and more than 245 GW of winter peak load growth over the next decade.⁵ This represents a more than 69% increase in projected summer demand growth and a 65% increase in projected winter demand growth compared with NERC’s

prior forecast,⁶ underscoring the speed and scale at which electricity demand expectations are evolving.

At the same time, the expansion of generation, transmission and distribution (“GT&D”) infrastructure struggles to keep pace. As a result, developers, hyperscalers, and other large-load customers are seeking alternatives that can deliver power on timelines aligned with project development schedules.

Figure 1: North America Data Center Capacity⁴



Source: DatacenterHawk, FTI analysis


Regulatory developments are further reinforcing this trend. In PJM, where electricity demand is projected to increase by approximately 85 GW over the next 15 years,⁷ FERC recently determined that PJM's existing tariff lacked clear and consistent rules governing co-located generation and large-load facilities and was unjust and unreasonable. FERC has directed PJM to formalize rules for co-located generation and large-load facilities, helping facilitate new approaches to serving growing load.

Similarly, system operators are increasingly encouraging large-load customers to provide operational flexibility through onsite generation, energy storage, demand response and flexible load management. In Texas, ERCOT's Batch Zero process was established in response to more than 438 GW of large-load interconnection requests.⁸ The framework requires developers to demonstrate project viability and financial commitment earlier in the interconnection process.

Against this backdrop, BTM power has emerged as an actively deployed near-term solution.

BTM solutions have evolved into a diverse set of technologies differentiated by speed-to-power, scalability, fuel availability, operational flexibility, emissions performance and lifecycle economics. Today, reciprocating engines, aeroderivative gas turbines, simple-cycle turbines and generators account for most near-term AI data center deployments due to their modular designs, proven reliability and relatively short development timelines. FTI notes that, while aeroderivative gas turbines continue to face supply chain constraints due to Original Equipment Manufacturer ("OEM") backlogs, acceptance of reciprocating engine solutions has gained meaningful traction and market acceptance as a faster speed-to-power solution. Fuel cells, meanwhile, are emerging as a viable option for edge data centers. Over time, other technologies, including advanced nuclear, geothermal and gas with carbon capture, are expected to play an increasingly important role as developers seek more scalable, reliable and low-carbon solutions to meet the growing power demands of AI infrastructure and other energy-intensive applications.

Figure 2: NERC North America Power System Assessment⁹



The Business Case for BTM Power

Historically used in mission-critical industries, BTM power is increasingly adopted by hyperscale and AI data centers seeking greater control over power. This shift is changing the relationship between data centers and the grid. Rather than relying on the grid for primary power and maintaining backup generators for emergencies, developers are integrating onsite generation, storage and microgrid controls into campus infrastructure.

The appeal of BTM extends beyond access to megawatts; onsite generation can be deployed incrementally alongside computing demand, often years before transmission upgrades or grid interconnections become available. The economics of BTM are driven not only by the cost of electricity, but also by the value of accelerating development timelines. In markets where utility interconnections can take five to ten years, the ability to bring computing capacity online sooner may create significantly higher value than the incremental cost of onsite generation.

BTM also alters the cost structure of power delivery. By shifting the generation onsite, developers can reduce exposure to T&D charges, system upgrade costs and ISO and utility fees. These potential benefits must be weighed against the additional capital investment and operational responsibilities associated with generation assets, fuel supply and electrical infrastructure. Procurement itself is also changing. Earlier data center agreements were energy transactions priced to hedge electricity cost; AI-era agreements are capacity transactions, gigawatt-scale, twenty years or longer and priced for guaranteed availability. Consequently, BTM is best evaluated through a broader lens that incorporates speed-to-power, reliability, risk, operational resilience and contract structure.

Ultimately, the strategic value of BTM lies in the certainty and speed-to-market it provides. The key question is not whether BTM power is cheaper than the grid. Rather, it is whether securing reliable power sooner can justify the additional investment required to obtain it. For a growing number of AI developers, the answer is yes, reflecting a broader shift from optimizing for the lowest cost of power to optimizing for assured access to it.

Figure 3: Components of FTM and BTM PPAs Prices¹⁰

	FTM	BTM
Energy	✓	✓
+		
Capacity	✓	✓
+		
Environmental Attributes	✓	✓
+		
ISO charges	✓	✗
+		
Utility Charges	✓	✗
+		
Retail Margins	✓	✗
+		
System Upgrade Cost	✓	✗
+		
Onsite Electrical Infrastructures	✗	✓
=		
All-in PPA	✓	✓

Rethinking Resilience and Redundancy in AI Data Centers

Traditional data centers rely on the utility grid as their primary source of power, making grid outages the largest reliability risk. While the U.S. grid is highly reliable, delivering approximately 99.95% availability on average,¹¹ even a few hours of annual interruption are unacceptable for mission-critical digital infrastructure. To achieve four-nines (99.99%) and five-nines (99.999%) availability, operators typically layer UPS systems, battery storage and diesel generators on top of a single utility supply. These systems are designed to reduce potential downtime from hours per year to only minutes.

BTM generation redefines this reliability model. Rather than depending on a single utility supply supported by standby assets, BTM campuses are designed as integrated power systems that combine onsite generation, energy storage and advanced control platforms. Reliability is achieved through the coordinated operation of multiple resources capable of responding to the rapid load fluctuations and stringent power-quality requirements characteristic of AI workloads.

Importantly, most hyperscalers continue to view the utility grid as the preferred long-term source of power due to its scale, efficiency and potential access to low-carbon electricity. BTM generation is generally considered a complementary solution that can enhance resilience, accelerate deployment or bridge periods when grid capacity and interconnection timelines cannot keep pace with AI infrastructure demand.

In a BTM campus, because the generation capacity is distributed across multiple modular assets, the loss of an individual unit has lower impact on overall operations. Additional technologies, including synchronous condensers, flywheel systems and advanced control platforms, can further enhance voltage stability, frequency response and system performance. Properly designed BTM microgrids can therefore achieve reliability levels comparable to traditional data center architectures through resource diversification.

This architecture also provides greater operational flexibility and can reduce capital investment in backup infrastructure. Because redundancy is embedded within the generation fleet itself, operators may require fewer dedicated standby generators and associated electrical systems than in traditional grid-centric designs. Generation assets can be maintained without materially disrupting operations and, in certain configurations, onsite generation, storage and intelligent controls can significantly reduce, or even eliminate, reliance on conventional diesel backup systems.

At the same time, the industry's experience operating large-scale BTM generation systems for AI data centers remains relatively limited. While individual generation, storage and control technologies have established operating histories, the application of these resources as integrated, utility-scale power systems supporting hyperscale AI workloads is still evolving. Achieving high levels of availability, including three-nines (99.9%) reliability and beyond, will require continued operational experience, refinement of system architectures and the refinement of best practices. As additional deployments come online, operators, equipment providers and regulators will gain valuable insights that are likely to further improve the reliability, resilience and economics of BTM-powered data center campuses.

Figure 4: FTI Analysis of FTM and BTM Architectures¹²

	Data Center with FTM Power Supply	BTM Microgrid Data Center
Availability Goal	99.99%-99.999%	99.99%-99.999%
Power Architecture	Utility Grid + Standby Backup Systems	Integrated Generation, Storage & Controls
Single Point of Contingency	Higher Dependency on Grid Supply	Lower Through Onsite Resources Diversity
Primary Reliability Risk	Utility Outage or Grid Disturbance	Fuel Supply, Controls Integration and Asset Coordination
Frequency Stability	Primarily Supported by Bulk Power Grid	Supported by Generation Controls, Storage & Microgrid EMS
Voltage Regulation	Utility Service + UPS Systems	Generation, Storage & Equipment
Transient AI Load Management	UPS and Batteries Absorb Short-Duration Events	Generation, Storage and Controls Manage Dynamic Load Ramping
Inertia & Grid-Forming Capability	Primarily Supplied by Bulk Power System	Rotating Assets, Advanced Inverters and Controls
Maintenance Flexibility	Reliant on Backup Architecture During Maintenance	Fleet-Based Redundancy Allows Maintenance While Operating
Scalability	Dependent on Utility Infrastructure Expansion	Modular Expansion of Generation and Storage Assets

The BTM Challenge: Allocating Risk to Unlock Scale

Despite the strong momentum behind BTM, scaling the model introduces a different set of challenges than traditional FTM power infrastructure. For many projects, the primary obstacle is aligning capital commitments and allocating risk among hyperscalers, developers, power providers and investors.

Many BTM solutions involve generation assets that are purpose-built for a specific data center campus under long-term contractual arrangements. Unlike FTM power plants, which can sell electricity into wholesale markets or serve multiple customers, these assets often have limited alternative uses. This introduces material stranded-asset risk: delays in customer expansion, changing technology requirements, workload consolidation or shifts in site selection can significantly impair utilization and project economics. The more specialized the infrastructure, the greater the potential exposure.

Development timelines further complicate BTM deployment, as developers and power providers often must commit substantial capital to create credible powered-land offerings before securing long-term customer commitments. This creates a classic coordination challenge between customers seeking certainty of power and investors seeking certainty of demand, placing a premium on commercial structures that provide both power certainty and revenue visibility.

Increasingly, these risks are addressed contractually through delivery and availability commitments, liquidated damages and capped remedies. As these mechanisms become standardized, BTM is viewed as a transparent and allocable risk that can be efficiently priced, underwritten and financed by the market.

Looking Ahead: From Bespoke Projects to Scalable Platforms

As power becomes an increasingly strategic input to AI infrastructure, market participants are pursuing vertical integration and consolidation to gain greater control over the value chain and accelerate growth. IPPs, fuel suppliers, developers, technology providers and infrastructure investors are becoming critical players in the ecosystem. What began as a response to grid constraints is evolving into a broader investment thesis, one in which access to firm, scalable power is becoming as important as access to fiber, land and capital.

That shift has reshaped how AI infrastructure is financed. Continued growth in tech sector capital investment has sustained demand across the energy and digital infrastructure stack. As long-term customer commitments and contracted cash flows become more prevalent, BTM projects are increasingly viewed through an infrastructure lens, creating opportunities for repeatable financing structures and broader access to institutional capital.

As the market matures, competitive advantage will shift from executing individual projects to building scalable infrastructure platforms that combine power, computing, capital and commercial execution. Organizations that can replicate these capabilities across multiple sites while effectively managing risk will be best positioned to capture long-term value.

How FTI Consulting Can Help

FTI Consulting's Power, Renewables & Energy Transition Practice ("PRET") offers a team of highly experienced economists, industry specialists, former utility and energy company executives, regulators and accountants to serve the regulatory and strategic needs of our power, utility and infrastructure clients. For data centers, investor-owned utilities, municipalities, cooperatives, developers or regulators, the PRET team provides our clients with holistic and actionable strategies, pertinent analysis and approaches to compete across the energy value chain.

FTI Consulting's Data Center Practice is an industry leader offering an expansive suite of services to investors across the globe, including due diligence, M&A advisory, market assessments, strategy, merger integration and carve-out support, performance improvement, business transformation, and turnaround and restructuring. Our data center team has worked on dozens of recent North American M&A transactions, providing commercial, operational, technical, ESG and sustainability, Workplace Health Safety and Environmental (WHSE), financial, tax and cyber due diligence services with expertise across the colocation spectrum (retail, wholesale, hyperscale, edge, crypto mining, GPU (graphics processing unit) cloud and carrier hotel segments).

Our services include:

- Growth Strategy
- Financial and Operational Strategy and Business Transformation
- Turnaround and Restructuring
- Utility Rate Case Advisory and Rate Design
- Capital Project Planning
- Buy- and Sell-Side M&A
- Power Market Price Forecasts
- Commercial and Financial Due Diligence
- Safety and Reliability Compliance
- Financial and Operational Compliance
- Dispute and Expert Witness

Fengrong Li, CFA, CIRA

Senior Managing Director
Power, Renewables & Energy Transition

Chris Post, CFA

Senior Managing Director
firstname.lastname@fticonsulting.com

Henrique Souza

Managing Director
Telecom, Media & Technology

The views expressed herein are those of the author(s) and not necessarily the views of FTI Consulting, Inc., its management, its subsidiaries, its affiliates, or its other professionals. FTI Consulting, Inc., including its subsidiaries and affiliates, is a consulting firm and is not a certified public accounting firm or a law firm.

FTI Consulting is the leading global expert firm for organizations facing crisis and transformation.

FTI Consulting is dedicated to helping organizations manage change, mitigate risk and resolve disputes: financial, legal, operational, political and regulatory, reputational and transactional. FTI Consulting professionals, located in all major business centers throughout the world, work closely with clients to anticipate, illuminate and overcome complex business challenges and opportunities.

© 2026 FTI Consulting, Inc. All rights reserved. [fticonsulting.com](https://www.fticonsulting.com)



Endnotes

¹ FTI analysis in connection with this article.

² FTI analysis.

³ FTI analysis.

⁴ FTI analysis.

⁵ Long-Term Reliability Assessment: North American Electric Reliability Corporation (NERC) (January 2026), https://www.nerc.com/globalassets/our-work/assessments/nerc_ltra_2025.pdf.

⁶ Ibid. at 25.

⁷ PJM's Updated 20-Year Forecast Continues To See Significant Long-Term Load Growth, January 14, 2026. <https://insidelines.pjm.com/pjms-updated-20-year-forecast-continues-to-see-significant-long-term-load-growth/>.

⁸ Public Utility Commission of Texas ("PUCT") Approves ERCOT's Batch Zero Process for Connecting Large Electricity Users While Protecting System Reliability for Texans, June 18, 2026, <https://www.ercot.com/news/release/06182026-puct-approves-ercots>.

⁹ North American Electric Reliability Corporation, Long-Term Reliability Assessment, January 2026.

¹⁰ FTI analysis.

¹¹ According to recent Institute of Electrical and Electronics Engineers ("IEEE") benchmarking data, the median System Average Interruption Duration Index (SAIDI) was 246 minutes per year, corresponding to roughly 99.95% availability. IEEE Power & Energy Society Distribution Reliability Working Group, IEEE Benchmark Year 2026 Results for 2025 Data, presented at the Distribution Reliability Working Group Meeting, Montreal, Canada, July 21, 2026. <https://cmte.ieee.org/pes-drwg/wp-content/uploads/sites/61/2026-DRWG-Benchmarking-Results-July-21-2026.pdf>.

¹² FTI analysis.