



CAMUS encord

ZERO LAB
PRINCETON UNIVERSITY

Flexible Data Centers: **A Faster, More Affordable Path to Power**

How flexible grid connections and bring-your-own capacity speed up the path to grid power and ensure data centers cover incremental costs

December 2025



Authors & Acknowledgments

Modeling and analysis for this study was conducted in summer and early fall of 2025. Inputs and assumptions reflect the best available data from public sources at the time. Camus, Princeton and encoord are independent organizations, and all paper conclusions and analyses are their own.

Authors

Carlo Brancucci, Dylan Cutler, Emma Elgqvist, Jesse Jenkins, Nathan Case, Scott Burger, Steven Brisley, Wallace Kenyon, Will Frazier, Wilson Ricks

The authors wish to acknowledge members of the Google Energy team for their valuable feedback and input on earlier drafts of this report.

Contacts

Dylan Cutler, dylan@camus.energy
Steven Brisley, steven@camus.energy

Suggested Citation

Carlo Brancucci, Dylan Cutler, and Jesse Jenkins. Flexible Data Centers: A Faster, More Affordable Path to Power. Camus, encoord, and Princeton ZERO Lab, December 2025.

Images courtesy of Adobe Stock unless otherwise noted.

Acknowledgements

This report benefited greatly from input and review by contributors throughout the industry, including:

Doug Middleton, Tesla Energy
Duncan Campbell, Scale Microgrids
Ethan Tiao, Emerald AI
Fiona Dearth, Generate Capital
Garretson Oester
Kate R. Sherwood, ZeroEnergy
Logan Goldie-Scot, Generate Capital
Pam Maines, ZeroEnergy
Sean Jones, Tesla Energy
Varun Sivaram, Emerald AI

The conclusions reflected in this study are not intended to represent the views of the contributors. We appreciated their help in providing useful context and guidance, which the authors used to shape the research.

Research Sponsors

This research was funded by Google LLC (Google), who reviewed the analysis prior to publication.

Flexible Data Centers: A Faster, More Affordable Path to Power

Across the United States, data centers face mounting delays in connecting to the electric grid¹ – with timelines of 3 to 7 years², far longer than the 18-24 month construction timeline for a new data center. This study demonstrates a practical path to accelerate access to grid power. By combining **flexible grid connections** with **bring-your-own capacity (BYOC)** arrangements, data centers can reach full operation years sooner while maintaining reliability and improving affordability for all customers.

The challenge

Two bottlenecks dominate the data center interconnection process. **Transmission constraints** occur when lines cannot carry additional power without upgrades, while **generation constraints** arise when the system lacks sufficient accredited generation capacity to provide firm service for new load. Both constraints are often binding. In PJM, the most recent capacity auction reached its price cap³ while transmission upgrades often face multi-year delays. Traditionally, the only remedies are to wait for new infrastructure or shift development to a different location—neither of which meets the speed-to-power needs of today’s AI-driven data center buildout.

The opportunity

Flexible grid connections and BYOC programs provide a two-part solution that directly addresses these bottlenecks. Under a **flexible grid connection**, a data center receives both firm (uninterruptible) service and conditional firm service⁴, where a portion of the load uses grid power in normal conditions and relies on on-site or co-located resources - including demand-side flexibility - during limited periods of system stress. Through **BYOC** programs, the data center directly procures the accredited capacity needed to meet firm-service requirements— through clean energy PPAs⁵, VPPs, or on-site resources—rather than waiting years for utility procurement or ISO queue processes.

Together, these mechanisms replace the traditional “build first, connect later” model with a different approach: **connect now, operate flexibly during the hours the grid is constrained**. This approach aligns the data center’s need for rapid access to power with the utility’s or ISO’s obligations to maintain reliability and ensure affordability. Flexible connections address transmission bottlenecks, while BYOC addresses generation bottlenecks.

¹ Source: [403 Large Loads Letter \(October 2025, Department of Energy\)](#)

² Sources: [Virginia Data Centers Face Seven-Year Wait for Power Hookups - Bloomberg](#); [The United States Needs Data Centers, and Data Centers Need Energy | ITIE](#); [Can US infrastructure keep up with the AI economy? | Deloitte](#)

³ Source: [PJM capacity prices set another record with 22% jump | Utility Dive](#)


⁴ We use “conditional firm service” to describe a contractual arrangement that provides conditional access to the grid, where a portion of a large load may be curtailed during limited periods of system stress to maintain system reliability. Similar arrangements are referred to as “non-firm,” “flexible connections,” or “flexibility commitments.” Our usage most closely aligns with published definitions of non-firm or flexible connection agreements as described in [Non-Firm Grid Connections](#)

⁵ A detailed definition and explanation of bring-your-own capacity is included on page 12



Flexible Connection

A flexible connection provides a combination of **firm** service (uninterrupted grid power) and **conditional firm** service, which delivers grid power when available and uses local generation or demand flexibility to serve remaining demand during brief grid constraints.



Bring-Your-Own Capacity (BYOC)

Bring-your-own capacity (BYOC) lets a data center supply its required accredited generation capacity through its **own off-site PPAs or on-site resources**, enabling capacity delivery in 1-2 years instead of waiting for multi-year utility planning, procurement and construction timelines.




Figure 1: Summary of Flexible Connection and Bring-Your-Own Capacity

The demonstration

To test this integrated approach, Camus led a collaborative analysis with Princeton University's ZERO Lab and encoord. The team used a three-tier methodology—system, utility, and site-level modeling—applied to six real candidate sites within one PJM utility's territory. **Camus** combined the transmission constraints identified by encoord and the system-level capacity findings from Princeton with site-level modeling to determine how an optimal mix of on-site and off-site resources would meet each site's flexibility requirements.

The study combined:

- **encoord's SAInt platform** to simulate available firm transmission capacity and flexibility requirements across hourly time horizons (8760 hours per year) at each site
- **NREL's REopt model** to determine cost-optimal portfolios of on-site flexibility resources
- **Princeton's GenX model** to assess generation capacity requirements and system-level cost and emissions impacts

This is the **first publicly available study** to combine real utility transmission system data, system-level capacity expansion modeling, and site-level capacity optimization to evaluate how flexibility can accelerate data center interconnections.


Just as importantly, **the methodology demonstrated here provides a repeatable blueprint that any utility can follow**, using the tools and data already available for new load studies.

Key finding #1: Flexible data centers can connect 3-5 years faster

A 500-MW data center using flexible grid connection + BYOC can reach full operation in roughly two years; **three to five years faster** than traditional interconnection processes. Across constrained sites:

- Grid power was available for **more than 99% of all hours**
- On-site resources (e.g. batteries, generators, load flex) were dispatched **40–70 hours per year**
- Transmission constraints led to **7–35 curtailed hours annually**, with events lasting 4–16 hours
- Generation shortfalls added **~32 hours per year**, concentrated in extreme weather events

Flexible connection and bring your own capacity (BYOC) shorten the wait for a data center grid connection by 3-5 years





Normal and high cases assume major transmission upgrades (e.g. new 230 kV or higher lines, significant reconductoring of lines) are required. Low case assumes extension of MV transmission and/or development of a new or retrofitted substation. Normal case for generation buildout assumes a new generation resource must navigate the PJM interconnection queue (start to finish).

Sources: U.S. Department of Energy Transmission Impact Assessment, Bloomberg, primary industry interviews

Figure 2: Speed to Power Gains (above) and Figure 3: Grid Availability (below)

Grid power remains available >99% of the year; on-site or co-located resources are dispatched 40-70 hours annually to manage grid constraints

Daily minimum deliverable power (% of nameplate) – based on transmission capacity (Hare)



Dispatch of on-site resources during longest transmission grid constraint (Hare)



Key finding #2: Flexible grid connections and BYOC significantly reduce and internalize incremental supply costs

A central question for utilities and regulators is whether new data centers increase costs for other customers. This analysis finds that the same tools that accelerate access to power—flexible grid connections and bring-your-own capacity (BYOC)—also provide a clear pathway to mitigate or avoid those costs. Used together, these mechanisms **mitigate new system buildout** and **shift remaining costs onto the data center**, substantially reducing the risk of cost shifts for other customers.

- Each gigawatt of new data center demand adds **\$764 million in system supply costs**⁶ under a traditional firm-only interconnection—driven by 2.17 GW of required nameplate generation additions across natural gas, storage, solar, and wind
- **Flexible grid connections** with 20% conditional firm avoid **273 MW of new build**, primarily battery storage and natural gas, eliminating **\$78 million** in incremental system costs per GW
- **BYOC** internalizes **\$326 million** in capacity costs per GW, with the data center procuring accredited resources directly and offering them into the market to increase supply
- Data center **payments for the energy portion** of their bill cover an additional **\$329 million** per GW of new demand

Across these components, a **data center contributes ~\$733 million per GW** toward the costs associated with its incremental load, reducing the net system cost increase by **nearly 100%**. Flexible grid connections reduce the amount of new capacity the system must build, while BYOC ensures the data center, not other customers, pays for the capacity required to serve its firm load.

Flexible data centers also **increase utilization** of transmission and generation assets, spreading fixed costs more broadly and creating **new opportunities to ease rate pressure for all customers**.




Figure 4: Incremental Cost Coverage by Flexible Connection and BYOC

⁶ Incremental system supply costs estimated by the capacity expansion modeling (via GenX)

The implications

Flexible connections and BYOC offer a faster, more affordable path for integrating large data centers:

- **Speed:** three to five years faster access to full power for the data centers
- **Affordability:** data centers directly cover incremental costs while increasing grid utilization, enabling utilities to spread fixed costs over more electricity sales
- **Reliability:** data centers gain reliable grid supply for >99% of the year, while utilities gain additional demand side resources to alleviate system stress

The foundations for rapid progress are already in place. Utilities have the data needed to evaluate flexible grid connections, and the planning tools required to do this work are readily available. Early demonstrations show this approach can work. What's needed now is for utilities, regulators, and data center developers to build on that foundation by adopting advanced planning tools, defining clear service agreements, and launching near-term, large-scale demonstrations that make flexible grid connections and BYOC available to data centers that are willing to operate flexibly.

While the grid will still require major transmission and generation investments to meet long-term AI-driven demand, flexibility offers **the fastest, most practical path forward today**. It allows data centers to connect faster, protects reliability, and gives planners the time and breathing room needed to build the infrastructure of the future—a strategy that strengthens the grid instead of waiting on it.

Table of Contents

Grid constraints: generation and transmission	9
New path: flexible grid connections + bring-your-own capacity	11
Flexible grid connection	11
Bring-your-own capacity (BYOC)	12
Speeding up the path to grid power	12
How we studied it: real data, 3 perspectives	13
Site Selection	14
Finding #1: Flexible data centers can connect 3-5 years faster	16
A faster path to full power	16
Transmission constraints: 4 of 6 sites faced limited grid capacity	17
On-site flexibility: enabling full power within grid limits	19
Generation constraints: securing firm service via BYOC	21
Generation constraints: Managing curtailment events for the conditional firm service	24
Solving the two bottlenecks	25
Why speed matters to data centers	26
Finding #2: Flexible connections avoid \$78 million and 273 MW of capacity buildout per GW of data center demand; BYOC internalizes \$326 million per GW of capacity costs	29
Each GW of new data center demand adds \$764 million in baseline system supply costs	29
Flexible connections (20% conditional firm) avoid \$78 million and 273 MW of new capacity per GW of new demand	30
BYOC internalizes \$326 million of capacity costs alongside the data center's payment of \$329 million in energy costs per GW	30
With flexibility and BYOC, data centers contribute ~\$733 million per GW of demand, covering nearly all of the incremental supply costs	32
Flexible Data Centers Boost Utilization & Grid Efficiency	32
Detailed methodology	35
Tier 1 – System Planner	35
Tier 2 – Grid Planner	36
Tier 3 – Site Planner	37
Making flexibility work in practice	38
Technology readiness: proven tools for flexibility	38
Expanding transmission capacity with Dynamic Line Rating	40
Regulatory progress: overcoming barriers	41
Research limitations and future topics	44
Recommendations for immediate action	45
For Utilities	45
For Regulators and Policymakers	46
For ISO/RTOs and Vertically Integrated Utilities	46
For Developers and Utilities	46
Appendix: Assumptions, Methodology, & Additional Analyses	48

Grid constraints: generation and transmission

Connecting new data centers to the electric grid has become one of the biggest challenges facing utilities, as the scale and speed of demand outpace available capacity. Data centers face two distinct yet often overlapping bottlenecks on the path to grid power: **generation constraints** and **transmission constraints**. These constraints stem from a diverse set of root causes: a lack of transmission infrastructure capacity, outdated planning tools and processes, speculative requests flooding generation queues, conservative risk evaluation, etc. The result in most regions of the U.S. is the same: more costs, constraints, and delays.

Understanding both constraints is essential to explaining why many large load projects are delayed and how flexible solutions can unlock faster, more affordable connections.




Figure 5: Two Bottlenecks to Connecting Data Centers

Generation constraints stem from a lack of accredited capacity available to reliably meet existing peak demand, capacity reserve margins, and the incremental demand of a new data center. In most markets, a load serving entity (LSE) cannot grant firm service to a large new load unless it can demonstrate sufficient accredited generation capacity to cover that load during system peaks.

Historically, LSEs maintained enough capacity headroom to enable new loads to connect without delay. Today, multi-year generation interconnection queues and accelerating load growth have made generation capacity a binding constraint in many regions. The most recent PJM capacity auction reached its price cap—a clear signal that supply is tight⁷.

⁷ Source: [PJM capacity prices set another record with 22% jump | Utility Dive](#)

Transmission constraints occur when lines and substations cannot safely carry additional power without exceeding thermal, voltage, or contingency limits. In many regions, these bottlenecks have become the primary cause of interconnection delays.

Transmission expansion is complex and capital-intensive: new transmission lines often require 7 to 10 years⁸ from planning to energization, while reconductoring an existing line or simply extending service can be completed in as few as 2-3 years. However, even as utilities move quickly to expand capacity, the surge of large load requests is straining staff, equipment, and supply chains, further extending timelines⁹. Siting, permitting, and cost allocation of new facilities introduce additional challenges.

While a new data center may face only a transmission or generation constraint, the two often coincide. A site near an existing substation may still encounter both local transmission bottlenecks and regional generation shortfalls that limit access to 24/7 power.

Our analysis of six potential data center sites in one utility's territory found that **four faced transmission constraints** and that PJM overall is generation constrained – meaning that **all six sites must overcome generation constraints**. In practice, addressing one type of constraint may not speed interconnection if the other remains, as the slower of the two becomes the critical path.

For this study, we assume data center construction and the extension of standard transmission service require two years to complete while **major transmission upgrades require five to seven years to alleviate related constraints**¹⁰. During the period before the site is energized, the utility incurs planning and construction costs, while the data center outlays cash for grid upgrade expenses and bears opportunity costs from idle infrastructure. In this study, we also assume **no surplus generation capacity is available in PJM** and that procuring new accredited capacity by **building new generation resources requires roughly five years** to navigate the interconnection queue¹¹.

By quantifying how often each constraint occurs, this study establishes the foundation for a new approach: using **flexible grid connections** to navigate transmission constraints and **bring-your-own capacity** (BYOC) to overcome generation constraints. Combined, these strategies remove key obstacles on the critical path that delay large load interconnections.

⁸ [Transmission Impact Assessment \(U.S. Department of Energy\)](#)

⁹ [Can US infrastructure keep up with the AI economy?](#) (Deloitte)

¹⁰ This assumption is inclusive of broader system upgrades as well as local connection to the new data center site. It reflects conversations with utilities and data center developers across the U.S. and is not representative of a single utility. Timelines for different projects and utilities vary from 2-10 years.

¹¹ Interconnection.FYI ([2024 Queue Changes](#), published June 1, 2025)

New path: flexible grid connections + bring-your-own capacity

Rather than waiting for new grid infrastructure or capacity to be built, data centers can **combine flexible grid connections and bring-your-own capacity arrangements to access power faster** while maintaining grid reliability. Together, these mechanisms replace the ‘build first, connect later’ model with one that connects new loads sooner—provided they can flex when the grid requires.

Flexible grid connection

A **flexible grid connection** allows a data center to receive a mix of firm and conditional firm service¹². Firm service guarantees grid power at all times; conditional firm service provides power except during specific hours when the grid is constrained. During these constrained hours, the data center uses on-site or co-located resources, such as batteries, solar, gas generators, or compute flexibility, to stay within grid limits.




Figure 6: Explanation of Flexible Connection


This approach effectively turns a static interconnection into a dynamic one, allowing new loads to connect within existing infrastructure limits. Success requires clear parameters for flexibility: how often curtailment might occur, how much load would be limited, and for how long. Utilities can define these through deterministic limits (e.g., maximum curtailed hours per year) or probabilistic data that lets data centers assess their own risk.

¹²We use “conditional firm service” to describe a contractual arrangement that provides conditional access to the grid, where a portion of a large load may be curtailed during limited periods of system stress to maintain reliability. Similar arrangements are referred to as “non-firm,” “flexible connections,” or “flexibility commitments.” Our usage most closely aligns with published definitions of non-firm or flexible connection agreements as described in [Non-Firm Grid Connections](#)

Bring-your-own capacity (BYOC)

Bring-your-own capacity addresses generation capacity constraints. In this model, the data center directly procures accredited capacity through bilateral contracts with net new or in-queue resources such as solar, wind, storage, natural gas, nuclear, virtual power plants, or onsite flexibility, rather than waiting for the load serving entity to expand its portfolio.

A bring-your-own capacity (BYOC) tariff addresses generation capacity constraints more quickly by procuring or building new capacity directly



This represents a simplified, illustrative model of bring-your-own capacity. Eligible PPA resources are limited to those that can be accredited by the ISO/RTO or relevant utility – including locational requirements based on regional supply needs. Note: Bring Your Own Generation (BYOG) constructs offer a parallel structure for data centers to bilaterally procure deliverable energy supply to match their demand, and can be paired with BYOC contracts to procure accredited capacity and delivered, time-matched energy.

Figure 7: Explanation of Bring Your Own Capacity (BYOC)

By doing so, the data center satisfies its share of PJM’s resource adequacy requirement and can be granted firm service for that portion of its demand sooner. BYOC effectively accelerates the addition of new capacity to the system by allowing large customers to either contract with projects already in the interconnection queue or develop co-located or behind-the-meter generation, helping relieve the generation constraint while maintaining reliability standards.

Speeding up the path to grid power

When used together, flexible grid connection and BYOC reduce or eliminate the obstacles that delay large load interconnections. The flexible grid connection removes the need to wait for transmission upgrades; BYOC removes the need to wait for new generation procurement. Combined, they allow utilities to provide earlier access to grid power without compromising reliability.

This study set out to answer one key question: **how can utilities and data centers best address transmission and generation constraints to speed up the path to grid power?** Our analysis demonstrates that combining flexible grid connections and bring-your-own capacity provides a practical, repeatable way to do so; one that uses existing utility data, physics-based modeling, and system-level planning to bridge today’s infrastructure gap.

How we studied it: real data, 3 perspectives

This study evaluates how flexibility and capacity can accelerate data center interconnections while reducing grid impacts, using three tiers of analysis: **system planner**, **grid planner**, and **site planner**. An optimal pathway minimizes a combination of time to connect and lifecycle cost of meeting data center demand. All analyses focus on the PJM region, with detailed modeling of one utility's service territory within the PJM footprint.




Figure 8: Summary of Three-Tiered Modeling Approach

A key feature of this approach is using results from one tier to impact the other tiers. There are three primary places where exchange of data between tiers supports our findings:

1. Transmission constraints and hourly curtailment requirements from the grid planner inform technology selection by the site planner
2. Generation constraints from the system planner inform accredited capacity procurements by the site planner
3. The conditional firm service level from the site planner informs the data center flexibility evaluated by the system planner, which informs findings on system costs and affordability

These **create a feedback loop**, where the system- and grid-planner analyses inform the site-planner requirements, which are used to evaluate system-level cost and emission impacts. Expanding on prior studies¹³, **this work incorporates real transmission system models and time-series SCADA data**, adding a rarely captured transmission planning perspective to the analysis of data center flexibility at six specific sites.

¹³ Prior studies include [Rethinking Load Growth - Nicholas Institute](#), [Fast, Flexible Solutions for Data Centers - RMI](#), and [Practical Guidance and Considerations for Large Load Interconnections - GridLab](#)

Site Selection

The sites selected for analysis were chosen to represent a broad range of interconnection characteristics, including proximity to major substations, proximity to major thermal generation stations, quantity of intersecting high voltage transmission lines, and peak utilization in the base case. All sites were located within approximately 125 miles of each other, with two being as close as 10 miles from each other. Unless clearly stated, all analysis evaluates sites one-at-a-time.

Table 1: Data Center Sites

Site ¹⁴	Interconnection Voltage Level ¹⁵	Line Connections	Peak Utilization ¹⁶	Additional Notes
Hare	230 kV	Substation with 230 kV / 69 kV voltage levels Junction of 4 - 230 kV transmission lines	58% maximum and 42% minimum loading of the set of 230 kV class lines at peak system demand	Approximately 200 MW delivered to 69 kV system at peak demand
Koala	230 kV	Substation with 230 kV / 69 kV voltage levels Junction of 3 - 230 kV transmission lines	58% maximum and 9% minimum loading of the set of 230 kV class lines at peak system demand	Approximately 70 MW delivered to 69 kV system at peak demand
Pony	230 kV	Substation with 230 kV / 69 kV voltage levels Junction of 2 - 230 kV transmission lines	22% maximum and 10% minimum loading of the set of 230 kV class lines at peak system demand	Approximately 80 MW delivered to 69 kV system at peak demand
Shark	230 kV / 500 kV	Substation with 500 kV / 230 kV Junction of both 500 kV and 230 kV transmission lines	58% maximum and 15% minimum loading of the set of 230 kV class lines at peak system demand	POI for thermal generation
Snake	230 kV / 500 kV	Substation with 500 kV / 230 kV / 69 kV voltage levels Junction of both 500 kV and 230 kV transmission lines	39% maximum and 15% minimum loading of the set of 230 kV class lines at peak system demand	POI for thermal generation Approximately 210 MW delivered to 69 kV system at peak demand
Whale	230 kV	Substation with 230 kV / 69 kV voltage levels Junction of 3 - 230 kV transmission lines	49% maximum and 6% minimum loading of the set of 230 kV class lines at peak system demand	Approximately 95 MW delivered to 69 kV system at peak

¹⁴ We have affectionately named the six sites by animal names to obfuscate the actual sites within the utility transmission network. These will be used throughout the paper to refer to the site distinct locations analyzed.

¹⁵ Two of the sites are located at major substations near large thermal generation facilities with numerous high voltage transmission line interconnections. These two sites had ample transformer capacity between 500 kV and 230 kV, and the interconnection voltage level is therefore marked as both levels. The rest of the sites are located at the 230 kV level with a variety of transmission line count configurations.

¹⁶ The peak utilization represents the maximum and minimum loading of the 230 kV lines at peak system demand, regardless of power flow direction; this should not be interpreted as a hosting capacity metric. Where applicable, the consumption of the 69 kV system (from the 230 kV system) is noted.



Finding #1: Flexible data centers can connect 3-5 years faster

A faster path to full power

Our analysis of six individual 500-megawatt data center sites found that combining **flexible grid connections** with a **bring-your-own capacity (BYOC)** approach can **reduce the time to full power by three to five years** compared to the traditional interconnection process.




Figure 9: Speed to Power Benefits of Flexible Connections and BYOC

Across all the sites we assessed, **grid power remained available for more than 99% of all hours** in the year¹⁷, with on-site or co-located resources dispatched for only 40 to 70 hours annually to stay within transmission or generation limits. In total, this combined approach enables utilities to **connect up to 3x as much data center capacity within two years**¹⁸ as they could in 5-7 years under conventional methods, while maintaining system reliability.

The subsections that follow detail how each constraint—transmission and generation—was evaluated for each of the data center sites, how they were addressed through flexibility and capacity procurement, and how these solutions create a pathway to faster grid power.

¹⁷ We note that while grid power is available >99% of the hours in the year, the optimized dispatch that includes onsite resources results in 95-96% of total site energy being served from the grid

¹⁸ The Hare site would be limited to 154 MW until the transmission constraint is resolved. A flexible grid connection would support 500 MW of service (firm + conditional firm), which is 3.2x larger

Transmission constraints: 4 of 6 sites faced limited grid capacity

Out of the six data center sites evaluated, four faced transmission constraints that would limit their ability to draw full power under a traditional, static interconnection. Across those constrained sites, each modeled individually as an incremental 500-megawatt data center, the available firm capacity ranged from **154 to 326 megawatts**, with the remainder dependent on the transmission infrastructure upgrades. These results confirm that even within the same utility service territory, available grid capacity can vary significantly across locations, making the interconnection process highly site-specific.




Figure 10: Available Grid Capacity and Service Levels by Site

Although the severity of transmission constraints varied, all four sites showed similar patterns in frequency. **Transmission constraints occurred only during a small fraction of the year: between 7 and 35 total hours.**

The key difference was the duration and magnitude of those events, which determined how much flexibility each site required to operate within grid limits. At the three moderately constrained sites (Koala, Whale, and Pony), curtailment events were brief and infrequent, lasting no more than four to five hours. The most constrained site (Hare) faced longer events—up to sixteen hours—necessitating a larger portfolio of on-site resources to maintain full operations.

Table 2: Transmission Constraints by Site

Transmission Constraints by Site								
Site	Firm Limit	Annual curtailment	% of Year	Peak curtailment	Longest Tx curtailment event	Energy req'd during longest event	Number of events	Average event duration (hours)
Koala	326 MW	7 hours	0.08%	109 MW	4 hours	247 MWh	3	2.3 hours
Whale	322 MW	11 hours	0.13%	113 MW	4 hours	354 MWh	4	2.8 hours
Pony	314 MW	13 hours	0.15%	121 MW	5 hours	392 MWh	4	3.3 hours
Hare	154 MW	35 hours	0.40%	281 MW	16 hours	2,299 MWh	4	8.8 hours

Two sites (Shark and Snake) did not face any transmission constraints for the 500 MW of additional demand.

Even at the most constrained site, the total number of curtailed hours remained **well below 1% of the year**. This finding highlights a critical insight for utilities and developers: waiting several years for new infrastructure may be unnecessary when temporary, well-defined constraints can be managed through flexibility at the site level.

These site-level findings do not necessarily indicate that other sites in the utility's service territory (or in PJM more broadly) would require less than 1% curtailment. But they do suggest that conducting similar analyses for candidate sites could be worthwhile.

The two remaining sites in our analysis were able to accommodate the 500-megawatt load without triggering transmission violations, indicating that sufficient transmission capacity does exist in portions of the network, particularly those located proximate to the 500kV transmission backbone. Together, these results illustrate that transmission constraints are both **common and unevenly distributed**, and that flexibility offers a practical alternative to time-consuming grid upgrades where constraints do occur.

The differences across sites, each located within roughly 125 miles of one another, underscore the importance of making location-specific, time-varying transmission capacity data available to developers as they evaluate potential interconnection points. Today, that information is largely opaque, leaving developers to site projects without a clear view of where grid headroom exists. **Utilities could help unlock faster, better-informed siting decisions by sharing data that reflects how available transmission capacity fluctuates over time**, whether through dynamic hosting capacity maps or other mechanisms.



The following section examines how the four transmission-constrained sites could operate at full capacity more quickly by combining firm and conditional firm service, supported by a mix of on-site and co-located resources that enable data center operations to continue unaffected.

On-site flexibility: enabling full power within grid limits

For the four sites with transmission constraints, we evaluated how a **flexible grid connection**, combining firm and conditional firm service, could allow each data center to connect and operate its full 500 megawatt load from the first day of energization.

The scale of each site's transmission constraint determined the amount of flexibility required. Across all constrained sites, the curtailed hours were brief and infrequent, allowing each data center to maintain reliability with a targeted combination of **battery storage, on-site generation, and compute flexibility** tailored to its specific grid limits.

Among the six sites studied, **Koala and Hare represented the smallest and largest transmission constraints** respectively, and were modeled in greater detail to identify cost-optimal portfolios of on-site resources. Both sites required a combination of battery storage, small-scale gas generation, and compute flexibility, with Koala also incorporating a modest amount of co-located solar generation¹⁹ to reduce energy costs. All sites were assumed to have the same data center load profile and benefited from approximately 65 megawatts of headroom between their nameplate capacity (500 MW) and the load profile's actual peak demand (~435 MW), providing a small additional operating margin during constrained periods²⁰.

This analysis sought a **cost-optimal mix** of on-site resources – used to manage transmission grid constraints and, when not being used for that primary use case, deliver bill savings for the data center²¹. As a result, the recommended portfolio of on-site resources is likely to differ from the minimum-viable portfolio required to manage transmission grid constraints.

At Koala, the less constrained site, the optimal mix included **26 MW of solar PV, 49 MW of 4-hour battery storage (196 MWh), 11 MW of on-site gas generation, and compute flexibility capable of reducing or shifting up to 25% of the load up to 20 hours per year** to meet a peak of 109 MW curtailment and total curtailment of 7 hours per year. Notably, Koala's curtailment could be served with compute flexibility alone; however, doing so would require dispatching the full 25% compute flex capability, leaving no operational margin for error.

¹⁹Solar is not used to alleviate the transmission constraint; it's deployed as part of a cost-optimal portfolio to manage site-level costs. Because of this, it may be possible to site this solar at a nearby location and utilize virtual net metering

²⁰ The selected load profile includes peak load at 87% of nameplate capacity while the utilization rate (average load / nameplate capacity) is ~80%. This is considered to be “conservative” as most data centers operate at lower utilization rates and have lower peak-to-nameplate ratios.

²¹ Because compute flexibility is a newer technology with uncertain operational costs, this modeling assumes that compute flexibility *is not* be used for bill management and is only applied to the transmission constraint management.

Hare, which faced the largest transmission constraint, required a larger portfolio: approximately **155 MW of 4-hour battery storage (620 MWh)**, **48 MW of on-site gas generation**, and the **same level of compute flexibility** to meet a peak of 281 MW curtailment and total curtailment of 35 hours per year.




Figure 11: Cost-Optimal Portfolios of On-Site Resources By Site

Both configurations enabled the sites to serve their full 500-megawatt nameplate demand through a combination of firm and conditional firm service, with on-site resources used only during the limited hours when transmission capacity was constrained.

This analysis demonstrates that **on-site resources can overcome local transmission constraints entirely**, provided that the utility or transmission operator defines clear parameters for conditional firm service—either as a deterministic limit on curtailment frequency and duration or as probabilistic data that allows the utility and data center to assess operational risk.

With those parameters in place, the data center can identify an optimal resource mix to remain within grid limits while maintaining full operational capability.

While this flexible grid connection strategy mitigates the transmission constraint, achieving firm service also depends on access to accredited generation capacity. The next section examines this second constraint and how a bring-your-own capacity (BYOC) approach can accelerate access to firm power by allowing data centers to procure or provide accredited capacity directly.

Generation constraints: securing firm service via BYOC

Flexible grid connections resolve transmission constraints, but large loads must still address another barrier: ensuring enough accredited generation capacity to support firm service.

As described in the methodology, we modeled a bring-your-own capacity (BYOC) approach in which the data center meets its firm capacity requirement by procuring accredited resources directly rather than waiting for the load-serving entity to add new generation. This allowed the data center to satisfy PJM's capacity obligations sooner and secure firm service for a larger portion of its demand.

Available accredited capacity: limited off-site supply

Based on PJM queue data, we estimated that roughly **158 megawatts of accredited capacity per 500-megawatt data center** could realistically be procured from off-site resources within a two-year development window. This comprised roughly **118 MW of solar, 11 MW of wind, and 29 MW of VPP accredited capacity**. Beyond that, additional capacity would need to be met with on-site or co-located resources.

Table 3: Available PPA Capacity

Off-Site Resources (PJM)	Available per 500 MW data center ²²	PJM Average ELCC*	Available Accredited Capacity
Solar PV (Tracking)	2,000 MW	6%	118 MW
Wind (Onshore)	50 MW	23%	11 MW
Virtual Power Plant	50 MW	59%	29 MW
Total			158 MW

*ELCC figures are averages based on PJM forecasts for 2027 - 2034.

This estimate was derived by dividing the total accredited capacity of eligible projects in PJM's active interconnection queue²³ by the 24 gigawatts²⁴ of forecasted data center additions between 2025 and 2030, then adjusted for local resource mix within the modeled utility's territory. A detailed explanation of this methodology, along with high- and low-availability sensitivity cases, is provided in Appendix C.

²² The quantity of available off-site PPA capacity per data center may vary significantly based on how quickly PJM processes its queue, changes to the growth rate of data center demand, changes to ELCC ratings, and the pace at which new generation is added to the queue. We have provided a sensitivity analysis in the Appendix C that analyzes the impacts of High PPA availability and Low PPA availability.

²³ Source: <https://www.interconnection.fyi>

²⁴ Source: [PJM Board Letter | Large Load Additions | August 2025](#)

Portfolio results: Koala and Hare

Using these constraints, we modeled our two bookend sites (Koala and Hare) to evaluate how each could reach sufficient firm capacity through a **mix of off-site and on-site accredited resources**. At Koala, the optimal BYOC mix included 158 MW of accredited off-site PPAs for solar, wind, and VPP capacity complemented by 168 MW of accredited capacity from on-site natural gas generation (with 202 MW of nameplate capacity). At Hare, by contrast, the full 154MW of accredited capacity requirements could be met with off-site PPAs alone.

Table 4: Resource Portfolio by Site

Power Portfolio (Cost + Speed Optimal)			Koala	Hare
Firm Service	Off Site BYOC (Accredited)	Solar	118 MW	114 MW
		Wind	11 MW	11 MW
		VPP	29 MW	29 MW
	On Site BYOC (Accredited)	BESS	–	–
		Generator	168 MW ²⁵	–
Firm Service Level		326 MW	154 MW	
Conditional Firm Service	On Site (Not Accredited)	Solar	26 MW	–
		BESS	49 MW / 196 MWh	143 MW / 572 MWh
		Generator	11 MW	60 MW
		Compute Flex (20 hrs)	25% of Load (Max 109 MW)	25% of Load (Max 109 MW)
		Headroom ²⁶	65 MW	65 MW
	Cond. Firm Service Level ²⁷		174 MW	346 MW
Grand total			500 MW	500 MW

²⁵ The 168 MW of accredited generator capacity requires 202 MW of nameplate generator installed on site due to the 83% ELCC required by PJM

²⁶ Headroom refers to the difference between nameplate site capacity (500 MW) and site-level peak demand (435 MW) based on peak utilization rates.

²⁷ Conditional Firm Service Level includes resources whose cumulative nameplate capacities may exceed the service level because not all resources are dispatched at the same time. This mix finds the cost-optimal way to serve the data center demand during periods of grid constraint.



These portfolios reflect tradeoffs between speed, cost, and carbon intensity. Off-site PPAs are often preferred by developers due to their lower operational complexity and alignment with corporate sustainability goals, but they are also more expensive per accredited megawatt than on-site generation. By contrast, on-site gas generation offers a lower-cost pathway to accredited capacity but comes with higher emissions and additional permitting²⁸ and operational considerations. Our modeling treated off-site PPAs as the strictly preferred source of capacity when available, with on-site options filling the gap when queue-limited supply prevented sufficient procurement.

Capacity value and cost sensitivity

The cost of accredited capacity depends on each resource's **effective load carrying capacity (ELCC)**, a measure of how much a resource can reliably contribute to meeting peak demand. ELCC values vary by season and system conditions. In our base calculations, we used the **average of PJM's forecasted ELCCs for 2027–2034**, representing a near-term view of capacity accreditation across the region.

The GenX model used in our system-level analysis also generates its own dynamic ELCC values, which we present in Appendix E. In a “Summer Peak Risk” scenario from that modeling, the solar ELCC increased from 6% to 27% and wind from 23% to 35%, reducing the cost of accredited capacity by more than half—solar from \$1,123 to \$244 per kW-year and wind from \$710 to \$454 per kW-year. These results highlight how more granular, seasonally weighted ELCC methods could improve the economics of low-carbon capacity procurement for data centers and other large loads.

²⁸ There are operational limits to onsite resources with emissions that vary by region. Example: [Emissions and Air Permitting Requirements for Standby Generator Sets](#)



Generation constraints: Managing curtailment events for the conditional firm service

In addition to managing transmission constraints and securing firm capacity through BYOC, data centers taking conditional firm service must be prepared for **generation curtailment events**—hours when the balancing authority requires data centers with conditional firm service to reduce net demand due to a lack of available generation. These events are distinct from the broader need to provide accredited capacity for firm service.

Across the modeled sites, we estimate approximately **32 hours per year** when conditional firm customers could be asked to reduce grid demand. These events are infrequent, with two to three occurrences annually, generally concentrated during periods of extreme weather—though they can last up to ~16 hours. While infrequent, they are long enough to require meaningful flexibility.

Table 5: Generation Constraints

Generation Constraints					
Site	Annual curtailment	% of Year	Longest generation curtailment event	Number of events (per year)	Average event duration
All Sites	32 hours	0.37%	~16 hours	2-3	~8-16 hours

The same on-site or co-located resources used to manage transmission constraints can serve these generation curtailment needs. In our modeling, batteries, generators, and compute flexibility were rarely used at full capacity during transmission events, leaving substantial headroom to cover generation curtailments as well.

For example, at sites where compute flexibility is available for 100 hours per year, that compute flexibility would only be dispatched to manage transmission constraints between 7 and 35 hours in a year, leaving significant excess capacity to alleviate generation constraints during other periods.

Although exceptionally long or severe events could require upsizing on-site resources, our cost-optimal portfolios were sufficient to meet both types of curtailment obligations under typical conditions.²⁹

²⁹ The GenX model provides seasonal (rather than hourly) dispatch insights. As a result, we did not quantify how often generation and transmission constraints occur simultaneously. Future research using joint hourly load and generation forecasting would strengthen understanding of this overlap.

Solving the two bottlenecks

When applied together, flexible grid connections and bring-your-own capacity (BYOC) resolve the two bottlenecks that most often delay large-load interconnections: transmission capacity and accredited generation capacity. A two-pronged approach allows data centers to reach **full 500-megawatt operation in two years**, rather than waiting five to seven years for new infrastructure to be built.

At each modeled site, we found that grid power remains available more than **99 percent of the year**, with on-site or co-located resources dispatched during only **40 to 70 hours annually** to manage short-term transmission or generation constraints. The resulting energy mix³⁰ is almost entirely grid-supplied, with batteries, compute flexibility, or generators used intermittently for flexibility and to optimize energy costs through peak shaving.




Figure 12: Grid Availability and Dispatch During Period With Largest Transmission Constraint

By resolving both bottlenecks simultaneously, this approach reduces interconnection timelines without compromising reliability. These findings demonstrate that **flexibility does not materially limit access to power for the data center's compute loads**; data centers can continue to access grid power the vast majority of the time *and* maintain normal operations when the grid is constrained. Taken together, the flexible grid connection + BYOC model provides a faster, lower-risk pathway to grid power—achieving full operational capacity in roughly two years.

³⁰ The energy mix for each site included 95-96% grid supply, with the remaining coming from on-site resources. This should not be confused with grid availability, which was >99%; during peak cost periods the sites chose to procure energy from on-site resources for bill management, despite the availability of grid power

Why speed matters to data centers

Accelerating the path to power creates tangible value for AI data centers. Each megawatt of capacity represents between **\$4 million and \$12 million in annual revenue**³¹, depending on workload type and utilization. Accessing capacity three to five years sooner transforms project economics, allowing data centers to generate revenue years earlier and accelerating business growth and market share.

Across the two modeled sites, the present value (PV) of **additional EBITDA** generated by accessing grid power five years earlier ranges from **\$4.7 billion to \$5.5 billion per site**. Gaining access three years earlier generates between \$2.3 billion and \$3.2 billion per site. These values assume median revenue of \$8 million per MW, EBITDA margin of 45 percent³² and an 8 percent discount rate.

Implementing flexible grid connection + BYOC adds **\$1.2–\$1.4 billion in lifecycle costs**³³, including on-site flexibility and higher-priced accredited capacity via PPAs. Even with these additional costs, the EBITDA gains exceed incremental expenses across +3, +4, and +5-year scenarios, delivering **net returns of \$0.9 to \$4.3 billion per site**.




Figure 13: Incremental EBITDA vs. Lifecycle Costs by Site

³¹ Sources include [CBRE](#), [TNI](#), and primary research interviews with AI data center operators

³² Estimated based on industry examples, including [Iron Mountain](#) (45.6%) and [Equinix](#) (49.1%)

³³ These costs are incremental to standard data center backup systems (typically diesel generators), which continue to be available for grid outages and are not used for flexible connections nor BYOC accreditation

Importantly, the site-level returns depend on two primary factors: 1) how many years faster the data center can connect, and 2) revenue per MW. Across both sites, our sensitivity analysis shows that developers begin seeing net positive returns with **at least two years** of speed-to-power gains—even when accounting for on-site resource investments and higher-priced clean energy PPAs.³⁴




Figure 14: Sensitivity Analysis for Average Net Returns (EBITDA minus Lifecycle Costs)

³⁴ This finding is highly dependent upon the estimated revenue and EBITDA gains from faster speed to power and is intended as a simple heuristic, not investment advice



Finding #2: Flexible connections avoid \$78 million and 273 MW of capacity buildout per GW of data center demand; BYOC internalizes \$326 million per GW of capacity costs

A central question for utilities and regulators is whether new data centers increase costs for other customers. Our system-level analysis of adding incremental data center demand across PJM by 2030 finds that flexible grid connections and bring-your-own capacity (BYOC) provide a clear pathway to significantly reduce system cost increases. Used together, these tools allow data centers to internalize nearly the full cost of their incremental demand—while reducing required capacity buildout and increasing utilization for the system.

Each GW of new data center demand adds \$764 million in baseline system supply costs

Because PJM is capacity-constrained, serving large new loads under a traditional firm-only service model requires significant new accredited generation. In the 2030 base case³⁵, GenX modeling shows that each gigawatt of new data center demand drives **2.17 GW of nameplate generation capacity additions**, including 1.1 GW of natural gas, 775 MW of battery storage, and the remainder from solar and wind. These additions increase annual system supply costs (capacity + energy) by **\$764 million per GW**. This value forms the baseline against which flexibility and BYOC are measured.




Figure 15: Nameplate Capacity Buildout per GW of New Demand (2030, GenX Model)

³⁵Our base case uses GenX modeling in 2030 and is incremental to PJM’s planned load growth with a “low data center” case, including ~12 GW of data center demand, ~11GW of other load growth, and ~20 GW of plant retirements

Flexible connections (20% conditional firm) avoid \$78 million and 273 MW of new capacity per GW of new demand

Allowing 20% of incremental demand to take conditional firm service significantly reduces the accredited capacity PJM must build to maintain reliability. For each gigawatt of new data center demand, this level of flexible grid connection **avoids 273 MW of nameplate capacity buildout** (~13% of the baseline additions), and **\$78 million** in incremental system supply costs (not including any potential local transmission network savings).

Most of the avoided capacity comes from battery storage, with smaller reductions in gas and wind. These avoided additions represent true savings: capacity the system no longer needs to procure.




Figure 16: Reduction in New Nameplate Capacity Buildout (20% Conditional Firm Service)

BYOC internalizes \$326 million of capacity costs alongside the data center’s payment of \$329 million in energy costs per GW

After flexible grid connections reduce new system buildout, PJM must still secure accredited generation to serve the firm portion of new load. Under a BYOC framework, the data center procures this accredited capacity directly—via clean energy PPAs, VPPs, and/or on-site resources—rather than relying on the load-serving entity to build or procure additional generation.

For every gigawatt of new data center demand, BYOC has the potential to internalize **~\$326 million** in capacity costs. This reflects the data center paying for its firm capacity requirement (800 MW per 1 GW of nameplate demand) at the market clearing price under the 20% conditional firm scenario

(\$1,116 per MW-day³⁶). Notably, the market clearing price is set at the minimum price required to incent enough new entry of firm capacity to serve the incremental demand. Because generators don't all bear the same cost profile, this market clearing price, if applied to all new generation and storage as is assumed in the analysis, is higher than the incremental cost to add capacity to the system. The difference manifests as additional profit for generators (or "rent" in economic terms). Instead, by purchasing BYOC capacity at marginal cost bilaterally, **the data center effectively pays the marginal cost** for the incremental capacity required and **internalizes the full incremental capacity costs**.

If the procured capacity is supplied to PJM to offset the data center's demand at inframarginal prices, the clearing price in PJM's capacity auction would remain the same as if the data center did not exist: both the demand (from the data center with flexibility) and the supply (from incremental firm capacity) would shift by equal amounts, cancelling out their impact on the wider capacity market.




Figure 17: Diagram Explaining How Flexible Connection & BYOC Negate Capacity Price Impacts

These costs are not avoided, instead the BYOC model allows them to be shifted off the system and onto the data center, ensuring other customers do not bear the burden. Notably, BYOC is not the only way to ensure data centers pay for their full share of additional infrastructure build out without shifting costs to other customers. Tariff and rate design can also ensure costs are fairly allocated.

BYOC does not address energy costs. Instead, energy costs are covered by the data center's rates and tariff. This analysis estimates energy payments from the data center by multiplying the total annual demand for a 1 GW data center by the PJM marginal energy cost from the GenX model outputs. For the scenario where incremental data centers connect using 20% conditional firm grid connections, the data center contributes **\$329 million per GW** to cover the costs of its energy consumption³⁷.

³⁶ The capacity clearing price is higher than PJM's current cap and assumes the cap would be increased

³⁷ Energy price effects will vary across systems. In systems where energy costs increase as a result of the additional demand, mechanisms like the [Clean Transition Tariff](#) offer pathways to prevent cost shifts. In addition, data centers may opt to pay premiums for clean energy or hourly-matched products.

With flexibility and BYOC, data centers contribute ~\$733 million per GW of demand, covering nearly all of the incremental supply costs

Including avoided costs from the flexible connection, capacity contributions via BYOC, and energy costs paid by the data center, the data center contributes \$733 million per GW of demand in the case of 20% conditional firm service—equal to 96% of incremental system costs. Across a range of cases with varying conditional service levels, we find that data centers with flexible interconnection and BYOC arrangements cover between 91% and 101% of incremental system costs³⁸—effectively **avoiding any significant cost shift** to other electricity consumers




Figure 18: Incremental Cost Coverage by Flexible Connection and BYOC

Flexible Data Centers Boost Utilization & Grid Efficiency

Eliminating the incremental system costs of new data centers is only part of the affordability story. Flexible interconnections also help utilities **increase utilization** of transmission infrastructure and generation assets, creating opportunities to ease rate pressure for all customers.

First, flexible data centers increase utilization of existing transmission and distribution infrastructure. When a large load connects via flexible service, the utility can deliver significantly more energy across existing grid assets, while longer-lead upgrades are planned and built. For a

³⁸ See Appendix D for details; values include cases that range from 0% to 60% conditional firm service. The residual costs arise from non-linearities in supply and demand curves, where GenX prices don't perfectly align with system cost changes across cases. Additional information regarding model output prices and how flexible connections and BYOC negate changes to other customers' capacity and energy prices is available in Appendix D.

single 500-MW facility, this amounts to roughly **3.3 terawatt-hours (TWh) of additional annual electricity sales**³⁹ delivered through the existing transmission system. This spreads fixed transmission costs across additional electricity sales and leads to higher revenue for the utility providing opportunities to reduce rates for all customers.

Second, flexible data centers may increase utilization of the generation fleet. Data centers generally operate at **higher load factors than the broader system**. This analysis included a data center load profile with a load factor of 80%. When this demand is added to the system, overall load factor of the generation fleet rises: GenX modeling shows PJM’s system-wide load factor increasing from 57.1% to 60.1% in 2030 with new data center load⁴⁰. Incorporating flexibility, such as serving 20% of all new data center load with conditional firm service, raises this load factor further. The firm portion of the flexible data center’s demand operates at nearly 100% load factor, allowing each gigawatt of accredited generation capacity to serve more demand over the year.




Figure 19: Load Factor Impacts of Incremental Data Center Demand (Without, With Flexibility)

As these effects scale, utilities can move more megawatt-hours through the same infrastructure and derive more value from their existing generation portfolios—**spreading fixed costs over more energy sales** and creating opportunities to ease rate pressure for all customers. This dynamic persists when data centers maintain a load factor above the system average, a likely but not guaranteed outcome.

Flexible interconnections and BYOC do more than prevent cost shifts. They also help make the entire system more efficient: improving utilization and bolstering long-term affordability without compromising grid reliability.

³⁹ TWh sold to the data center is based on the assumed load profile for the 500 MW data center; includes ~80% average load (~400 MW) for 8,760 hours with 95-96% of GWh provided by the grid

⁴⁰ Based on GenX modeling for the base case, calculated as [total MWh] divided by [peak MW] * [8760 hours]



Detailed methodology

This section expands on the three-tier framework, detailing the data sources, modeling assumptions, and analytical steps used to evaluate flexibility and capacity across system, grid, and site levels.⁴¹

Tier 1 – System Planner

Tier 1 evaluates the **system-level generation and transmission buildout** required in PJM by 2030 using the [GenX model](#),⁴² populated with inputs from [PowerGenome](#). The capacity expansion model **assesses the expected costs and resource mix of the grid build out over the next five years** as PJM's grid evolves to accommodate substantial load growth.

Because the study focuses on near-term data center growth, the modeling captures the PJM portion of the Eastern Interconnect, reflects current market conditions that limit new capacity availability for certain resource types, and simulates system expansion through 2030. This simulation forms the “baseline scenario”, which incorporates PJM's current load forecasts⁴³—including ~24 GW of data center demand—and uses effective load carrying capacities (ELCCs)⁴⁴ for all included technologies.

To quantify the effects of (1) incremental data center demand and (2) load flexibility on system capacity needs, the study conducts two comparative analyses⁴⁵:

1. **Low demand case:** The 24 GW of data center demand in the baseline scenario is reduced to 12.2 GW while preserving ~11 GW of non-data center load growth and ~20 GW of planned retirements. Comparing this case with the baseline isolates the cost and capacity impacts of increased data center demand.
2. **Conditional firm cases:** The system impacts of conditional firm service are evaluated by modeling the capacity and cost implications of serving the incremental 11.8 GW of new data center load (baseline minus low-demand case) under varying flexibility levels. While results for 20%, 40%, and 60% conditional firm are modeled, the main findings presented in this report focus on the 20% case, with additional results provided in Appendix E. The conditional-firm portion may be curtailed during constrained periods but is assigned a marginal cost high enough to prevent curtailment during typical market operations, reflecting a data center's preference to draw from the grid whenever power is available.

⁴¹ Additional information regarding assumptions and methodology is available in Appendices A and B

⁴² GenX, a least-cost optimization model developed at Princeton and MIT, identifies cost-optimal generation, storage, and transmission investments needed to meet a defined system demand while observing physical and policy constraints. PowerGenome aggregates data from public sources (FERC, EIA, EPA, and NREL) to characterize generation technologies and costs, forecasted loads, existing generation fleets, and other aspects of capacity expansion modeling

⁴³ Current PJM forecast as of July 15, 2025

⁴⁴ It is important to note that GenX can (and by default does) calculate ELCCs endogenous to the model, rather than taking the exact PJM ELCC calculations as inputs. In the baseline case, these ELCCs are quite similar to the ones established by PJM, but this allows us to evaluate additional scenarios where peak capacity risk assumptions result in alternate ELCCs.

⁴⁵ An additional sensitivity evaluates how alternative summer peak risk periods affect ELCC values across technologies. Details on this scenario and its findings are provided in the Appendix E.



Tier 2 – Grid Planner

Tier 2 assesses the **transmission impacts and flexibility requirements for each of the data center sites** through detailed annual simulations of the utility network. Leveraging [encoord's SAInt integrated modeling platform](#), this study characterized how the generation fleet could be securely dispatched to reliably meet demand, and identified hours when load could not be fully served under the current system configuration and reliability standards.

The proposed and applied approach iterated between a nodal production cost model (PCM, formulated as a unit commitment DC optimal power flow model) and contingency analyses based on AC power flow (ACPF) simulations. Any unacceptable system contingency responses are addressed with subsequent PCM runs including mitigating security constraints. The results from this integrated planning approach **defined the amount of local load flexibility or generation** (which could be on-site or co-located downstream of the transmission substation) **required to serve demand when the grid is constrained**.

The PCM and ACPF models were implemented in SAInt starting from a transmission planning model of the eastern interconnection provided by the utility. To make annual simulations computationally feasible, we extracted the subset of the network covering the target utility's service territory and adjacent transmission interties. The interties were characterized with approximating impedances and historical time-series SCADA data from the utility to capture expected operational behaviors based on local generation and demand changes. This reduced network model was paired with additional SCADA data for system demand at step-down transformers and assumed generator performance characteristics for local units.

With a single integrated model containing a partitioned and fully validated transmission network and calibrated production cost models, we **added a 500 MW data center to six locations** and **evaluated the ability of the system to meet the additional demand**⁴⁶ in an operationally secure manner.




Figure 20: Diagram of Grid-Level Modeling (Transmission Constraints)

⁴⁶ While we primarily focus on results on adding each of the new data center sites one at a time, there is a scenario in the Appendix E that looks at impacts of all six sites added simultaneously

Tier 3 – Site Planner

Once **transmission constraints** were identified, we created an **8760-hour flexibility profile** (where non-zero hours indicate times the full data center load cannot be served) as an input to the site-level capacity expansion model. In simpler terms, this profile **mapped out when and for how long the grid may be unable to meet the data center's demand**, pinpointing the hours when on-site resources would be needed to keep the data center running normally.

Using NREL's [REopt model](#), we identified the cost-optimal mix of technology types, sizes, and dispatch that both addressed flexibility requirements and minimized site costs under the utility's tariff. The optimization included solar PV, battery storage (BESS), natural gas generators, and compute flexibility.

After optimizing behind-the-meter technologies, we evaluated accredited capacity options that would address **generation constraints** and support the firm portion of the load. We assumed the data center must procure accredited capacity to offset its firm load obligation, an assumption supported by Tier 1 results showing a capacity constrained system. Rather than relying on the load-serving entity to add new generation, our modeling reflected a bring-your-own capacity (BYOC) approach in which the data center directly procures or develops accredited resources.

In this BYOC framework, the data center satisfied its capacity requirement through a prioritized strategy that balances speed, cost, and carbon:

- **First, maximize low-carbon, off-site accredited capacity** that can be deployed within roughly two years, selecting from solar, wind, and virtual power plant (VPP) resources⁴⁷. We calculated the cost of a MW of each accredited resource (incorporating the energy value of the wind and solar PPAs), and selected in least-cost order until available capacity was exhausted.
- **Then, fill any remaining gap with on-site or co-located resources** that can be accredited while also supporting transmission flexibility, such as battery storage, natural gas generation, or compute flexibility.

Clean energy PPAs were prioritized for their availability in interconnection queues and alignment with common procurement practices, while on-site and co-located resources were modeled as secondary options when off-site accredited capacity is insufficient. The resulting site-level portfolios captured trade-offs between cost, carbon, and interconnection speed and fed back into the grid- and system-level analyses, linking site-level choices to regional reliability and affordability outcomes.

⁴⁷ This analysis did not consider energy storage, off-shore wind, hydropower, natural gas, coal, or nuclear generation within the off-site accredited capacity options. Gas and coal were excluded due to their carbon emissions. Off-shore wind, hydropower, and nuclear were excluded due to development timelines (5+ years). For future analysis, we recommend inclusion of storage (standalone or hybrid).

Making flexibility work in practice

Implementing flexibility at scale requires collaboration between utilities, data center developers, technology providers, and regulators. The following sections highlight the readiness of key enabling technologies, the growing regulatory momentum behind flexibility, and the immediate actions available to utilities, ISO/RTOs, regulators, and data center developers.

Technology readiness: proven tools for flexibility

The technologies that make flexible grid connections possible are already deployed across the power system today. Batteries, generators, compute flexibility, VPPs, and clean energy PPAs provide mature, scalable ways to support firm + conditional firm service.

Flexible Interconnection Software

Utilities increasingly rely on software to plan for and operate flexible interconnections. These platforms are in use across the industry and provide the core functions needed to make flexibility reliable: integrating grid and telemetry data, running hourly constraint analyses, forecasting available capacity, and issuing clear operating limits to large loads.

→ Real-world examples: [Camus](#) and [encoord](#)

Battery Storage

Battery energy storage is a proven reliability asset for utilities and data centers alike. Systems of hundreds of megawatts are deployed across PJM, ERCOT, and CAISO, regularly supporting grid stability and responding to dispatch signals within seconds. For data centers, batteries provide a clean, fast-acting resource to manage generation curtailment events and operate through short-term transmission constraints.⁴⁸

→ Real-world example: [Calibrant and Aligned](#)

Gas Generation

On-site natural gas generation is a mature and widely used backup power source. Generators can provide continuous power during curtailment events or outages, with start-up times often measured in seconds. While emissions are a consideration, modern gas units with advanced controls can operate in low-duty cycles consistent with flexible service, ensuring reliability without continuous operation.

→ Real-world example: [Enchanted Rock](#)

⁴⁸ In addition, batteries are increasingly being considered as tools to solve transience and ramping challenges, though these were not evaluated in this study.

Compute Flexibility

Compute flexibility—the ability to adjust non-time-sensitive computing workloads in response to grid conditions—is emerging as a powerful new tool for data center flexibility. Managed through software rather than hardware, it allows partial load reductions or task shifting with little to no impact on service delivery.

→ Real-world examples: [Google](#), [Emerald.AI](#)

Virtual Power Plants (VPPs)

VPPs aggregate distributed energy resources such as batteries, EV chargers, and flexible loads to provide dispatchable capacity at scale. In PJM and other markets, VPPs are accredited to provide capacity and ancillary services. As part of a BYOC strategy, VPPs could enable data centers to source accredited capacity faster and closer to load centers while providing reliability and resilience benefits to local community members.

→ Real-world examples: [Voltus](#), [Tesla](#), [Sunrun](#)

Off-Site Clean Energy PPAs

Most large data centers are already accustomed to contracting for renewable power through power purchase agreements (PPAs). In 2023 alone, corporations signed more than 30 gigawatts of renewable PPAs in the United States, with hyperscale data centers accounting for a large share. These contracts provide both accredited capacity and emissions reductions, directly supporting bring-your-own capacity (BYOC) approaches for firm service.

Clean Firm Portfolio PPAs

Clean firm portfolio PPAs are an emerging solution to BYOC and energy, in which optimized portfolios of resources — including batteries, solar, wind and VPPs — are bundled and operated together to supply firm, deliverable capacity and hourly-matched/shaped energy to meet data center demands and contracted through a single aggregated PPA and/or clean capacity tariff. Individually, wind and solar have lower ELCC values, VPPs and batteries cannot meet energy needs, and the scale of capacity offered by VPPs today is inadequate to meet data center requirements on their own. Constructing a portfolio of resources helps overcome the limitations of any one asset, delivering a combined package that is fast to build, scalable and relatively cost competitive. With terawatts of wind, solar and storage in interconnection queues today, this approach could help meet near-term, large-scale demand from data centers and minimize reliance on new and existing fossil generators.

→ Real-world example: [Firma Power](#)

SIDEBAR

Expanding transmission capacity with Dynamic Line Rating

While our primary analysis focused on on-site solutions to manage transmission constraints, we also evaluated **Dynamic Line Rating (DLR)** as a complementary utility-side option. DLR is a grid-enhancing technology that increases transmission capacity by adjusting line ratings in real time based on weather conditions such as temperature, wind, and solar irradiance. By replacing static assumptions or ambient-adjusted ratings⁴⁹ with data on real-world temperature and wind speeds, DLR can reveal significant unused capacity on existing lines.

Using findings from an NREL study (2007–2013 weather data), our analysis applied time-varying conductor ratings for 138 kV, 230 kV, and 500 kV lines across the utility’s portion of the PJM system⁵⁰. The results show that DLR can meaningfully boost available transmission capacity during most hours—especially in winter, when PJM constraints are most common.

At the modeled data center sites, the impact was substantial: with DLR, **Koala required no curtailment even at 700 MW of demand**, while **Hare’s curtailment dropped to just three hours per year**, with a 50 MW curtailment peak. These results demonstrate that DLR has potential to alleviate transmission constraints, yet we note that studies⁵¹ have shown the need for detailed analysis.

While data centers cannot directly deploy DLR, they can partner with utilities through cost-sharing or demonstration projects to accelerate adoption. Doing so can expand firm service levels for both the participating site and surrounding customers, using existing grid assets more efficiently while long-term upgrades progress.




Figure 21: Ratio of Dynamic Line Rating to Standard Line Rating for One-Year Period

⁴⁹ Source: [FERC Rule to Improve Transmission Line Ratings Will Help Lower Transmission Costs](#)

⁵⁰ These time-varying ratings are adjusted from the static, normal ratings assigned to each line and taken from the transmission planning case data provided by the utility.

⁵¹ Source: [Lessons from first deployment of Dynamic Line Ratings](#)



Regulatory progress: overcoming barriers

Most current regulations were designed for an era of firm, always-available electric service—an era when utilities built enough transmission and accredited generation to ensure any customer could draw maximum power at any time. Large, flexible loads require a different approach. Regulators increasingly recognize the need for flexibility, but often lack familiarity with the specific barriers that prevent utilities from offering firm + conditional firm service today. The following section outlines the key regulatory gaps and the changes needed to enable faster, more affordable interconnections.

Barriers to Implementation




Figure 22: Four Barriers to Implementing Flexible Connections & BYOC

1) Planning frameworks assume every load is always at its maximum

Integrated Resource Planning (IRP) and Resource Adequacy (RA) processes were designed around firm, always-available demand. These frameworks do not yet provide an adequate way to represent firm + conditional firm service, where a portion of a customer's load can reduce or self-supply during constrained hours. As a result, **planners generally assume all new data center load requires firm accredited capacity**, even when flexibility can reduce capacity needs and accelerate timelines.

What needs to change: Regulators should incorporate limited large load flexibility where voluntarily offered as an explicit input in IRP and RA processes, allowing utility planners to quantify how flexibility can reduce near-term firm capacity requirements and defer new capacity buildout.

2) Accreditation methods don't consistently define and value load-modifying resources

While some markets (such as PJM) accredit demand response, **current accreditation methods do not consistently quantify or recognize the value that flexible grid connections could provide to the system**. E3's analysis for SPP⁵² and GridLab plus Telos Energy's analysis for NV Energy⁵³ show that

⁵² [Demand Response in SPP - E3](#)

⁵³ [Bringing Data Center Flexibility into Resource Adequacy Planning: A case study of NV Energy - GridLab and Telos Energy](#)

accreditation frameworks, like effective load-carrying capability (ELCC), demonstrate high capacity value of demand response resources that can help address emergency system conditions for a limited number of hours per year. BYOC, including the use of on-site flexibility resources, relies on a clear method for valuing the capacity contribution of these flexible loads.

What needs to change: Regulators should extend accreditation methods, where necessary, to recognize the reliability contribution of emergency load modifying resources in RA planning under predetermined bounds of duration and annual availability.

3) Tariffs only allow firm or non-firm service—and often don't offer non-firm options at all.

Most tariffs treat service as binary: firm or non-firm.⁵⁴ This prevents utilities from offering firm + conditional firm service, where the customer receives firm service during most hours but agrees to reduce demand or self-supply during a small number of constrained hours. In addition, most large-load tariffs do not offer *any* non-firm service option, leaving no formal pathway for a data center to be studied as a curtailable customer. Without this structure, utilities cannot operationalize flexible grid connections at scale—even when doing so avoids significant upgrades or capacity procurement.

What needs to change: FERC and state commissions should encourage transmission providers to reform their processes to better utilize voluntary flexible loads, including establishing clear operating parameters (such as frequency, duration, and curtailment notice) and implementing reasonable telemetry and verification requirements to ensure reliable system operation⁵⁵.

4) Transmission and resource adequacy commitments must be recognized as independent

Flexibility commitments to the **transmission provider** that help bridge the gap until infrastructure upgrades are completed are distinct from flexibility commitments to the **balancing authority** to reduce capacity procurement needs even when these roles are fulfilled by the same entity (e.g., in an ISO/RTO). However, these are often conflated. FERC guidance is needed to make clear that these are independent⁵⁶.

What needs to change: Regulators, likely FERC, should clarify through rulemaking or guidance that transmission flexibility commitments and resource-adequacy flexibility commitments serve different purposes and must be treated separately unless explicitly coordinated.

⁵⁴ Source: [How DOE's Proposed Large Load Interconnection Process Could Unlock the Benefits of Load Flexibility](#)

⁵⁵ Often today, load modifying resources are subject to the same telemetry, verification, and testing requirements as supply side resources, which is overly burdensome. Although it's important to consider what is necessary to ensure reliable operation, supply side resource requirements should not be the default starting point.

⁵⁶ As noted in the [Nicholas Institute policy briefing](#), where the transmission provider and the balancing authority are the same entity, these two flexibility commitments could be rationalized into a single contractual or tariff framework

Regulatory momentum is already underway

Although regulatory frameworks are still evolving, momentum is building across federal, regional, and state levels.

At the **federal level**, DOE's *Advanced Notice of Proposed Rulemaking* (ANOPR)⁵⁷ calls on FERC to reform large-load interconnection processes and to enable accelerated study tracks for curtailable and hybrid loads. This is one of the first explicit federal steps toward codifying flexibility as a standard service option.

Across regional markets, progress is accelerating:

- **SPP** is studying co-located generation and conditional service models through its HILLGA and CHILLS initiatives⁵⁸
- **PJM** initiated a Critical Issue Fast Path (CIFP)⁵⁹ process for Large Load Additions, exploring bring-your-own-capacity (BYOC) options
- **ERCOT** has adopted protocol revisions (NPRR1234⁶⁰, NPRR1188⁶¹) to operationalize flexibility and is studying how to embed it in long-term planning

At the **state and utility level**, progress is tangible. Utilities including PG&E (CA), Avangrid (NY), and ComEd (IL) are already using flexible grid connection models to accelerate both generation and large-load interconnections⁶². These efforts demonstrate that flexibility can be implemented under existing constructs while broader tariff and market reforms advance. However, replicating these early results at scale - and fully realizing the speed and affordability benefits modeled in this study - depends on formalizing these approaches through updates to planning frameworks, accreditation methods, and tariff design.

⁵⁷ [403 Large Loads Letter](#)

⁵⁸ [High Impact Large Load \(HILL\) Integration - Southwest Power Pool](#)

⁵⁹ [CIFP Large Load Additions - PJM](#)

⁶⁰ <https://www.ercot.com/mktrules/issues/NPRR1234>

⁶¹ <https://www.ercot.com/mktrules/issues/NPRR1188>

⁶² Sources: [PG&E](#) (CA), [Avangrid](#) (NY), and [ComEd](#) (IL)

Research limitations and future topics

This study demonstrates that combining flexible grid connections with a bring-your-own-capacity approach can enable data centers to access power far sooner—while improving system utilization and managing long-term costs. Our analysis applies this model to real transmission data from a single utility within the PJM footprint, paired with system-level modeling of PJM’s generation fleet.

While the findings show that the combined firm + conditional firm approach can be cost-effective for the grid *and* the data center developer, our work represents a demonstration of the methodology for a select number of sites and system configurations, rather than a comprehensive national assessment. Several areas merit further study to validate, refine, and extend the insights presented here.

1. Evaluate the system-level impact of BYOC portfolios, especially with more clean generation

Our BYOC analysis focuses on identifying cost-optimal accredited capacity for firm service but does not assess how these portfolios interact with broader clean procurement strategies. Existing research—including studies on RE100 and 24/7⁶³ carbon-free energy—shows that large-scale clean power procurement can significantly reduce system emissions. Future work could build on these findings to examine how flexible interconnection and BYOC procurement together influence renewable buildout, emissions trajectories, and long-term market dynamics.

2. Assess operational interactions when multiple flexible grid connections exist in the same region

This analysis evaluates curtailment parameters for a single flexible grid connection—such as frequency, duration, and magnitude. But as multiple flexible data centers operate within the same area, their flexibility commitments will interact with one another and with the underlying physics of the grid. Key questions, such as how curtailment is ordered across sites, how flexibility obligations are allocated, and how portfolios behave when several flexible loads respond simultaneously, require further study. Modeling alternative program designs (e.g., zonal limits, LIFO vs. pro-rata allocation, or auction-based approaches) would help utilities and developers understand how scaling flexible interconnections across a region affects operational feasibility and financial outcomes.

3. Expand the analysis across additional utility and ISO/RTO regions

This demonstration focuses on PJM and one anonymized transmission operator within the PJM footprint. Transmission and generation constraints vary significantly across regions—shaped by weather patterns, system topology, planning practices, and the pace and location of load growth. Applying this framework in additional utility and ISO/RTO regions would help quantify how flexible grid connections and BYOC perform under different grid conditions and provide a more comprehensive view of where this approach offers the greatest value.

⁶³ Example studies include [On the means, costs, and system-level impacts of 24/7 carbon-free energy procurement - ScienceDirect](#) and [Advancing Decarbonisation through Clean Electricity Procurement - IEA](#)

Recommendations for immediate action

Accelerating large-load interconnections will require utilities, ISO/RTOs, regulators, and developers to take coordinated steps that enable firm + conditional firm service. The following recommendations translate this study's findings into clear, actionable priorities for each stakeholder group. Together, these actions would help connect new data centers and other large loads faster, at lower cost, and with greater transparency.




Figure 23: Recommended Actions to Accelerate Progress

For Utilities

1. Invest in advanced planning tools to quantify site-specific flexibility requirements.

Use time-varying (e.g. hourly) simulations to identify how often and by how much new loads or generation would exceed local grid limits. These analyses enable utilities to define curtailment parameters (including frequency, duration, and magnitude) that reflect real system constraints and can be transparently shared with developers.

2. Publicly signal willingness to engage on flexible grid connection and BYOC models.

Developers lack visibility into which utilities are open to exploring flexible interconnections. Publicly communicating an openness to evaluate such approaches, without committing to specific outcomes,

would help developers prioritize engagement and identify early opportunities for demonstration projects.

3. Develop and standardize parameters for firm + conditional firm service.

Establish a clear, repeatable framework for defining the conditions under which curtailment may occur. These parameters (e.g. number and duration of events, peak curtailment magnitude, and advance-notice requirements) should be clearly defined and consistent in structure.

For Regulators and Policymakers

4. Create standard requirements for and encourage development of firm + conditional firm tariffs.

Encourage utilities and RTOs to offer hybrid service combining firm and conditional firm delivery, and define the minimum requirements for such tariffs: clear curtailment limits, telemetry and verification standards, and consistent procedures for how flexibility is initiated and enforced.

5. Modernize planning and accreditation frameworks to recognize flexibility.

Encourage utilities and ISO/RTOs to incorporate firm + conditional firm service into Integrated Resource Planning and Resource Adequacy modeling. Where needed, expand existing accreditation methods to properly value emergency load-modifying resources and clean energy resources.

For ISO/RTOs and Vertically Integrated Utilities

6. Expand access to conditional firm service and BYOC models for flexible loads.

Continue or initiate efforts to create pathways for large flexible loads to access conditional firm service and bring-your-own capacity models. Ensure that study processes can evaluate flexibility at the time of interconnection and that results are integrated with system planning and capacity accreditation frameworks.

For Developers and Utilities

7. Launch large-scale demonstration projects and scale-up programs.

Collaboratively develop large-scale demonstrations and phased scale-up programs that test firm + conditional firm service models under real-world operating conditions. These demonstrations can provide the insights regulators and utility planners need to establish tariffs, planning procedures, and programs that safely scale flexible interconnections.





CAMUS

encord

ZERO LAB
PRINCETON UNIVERSITY