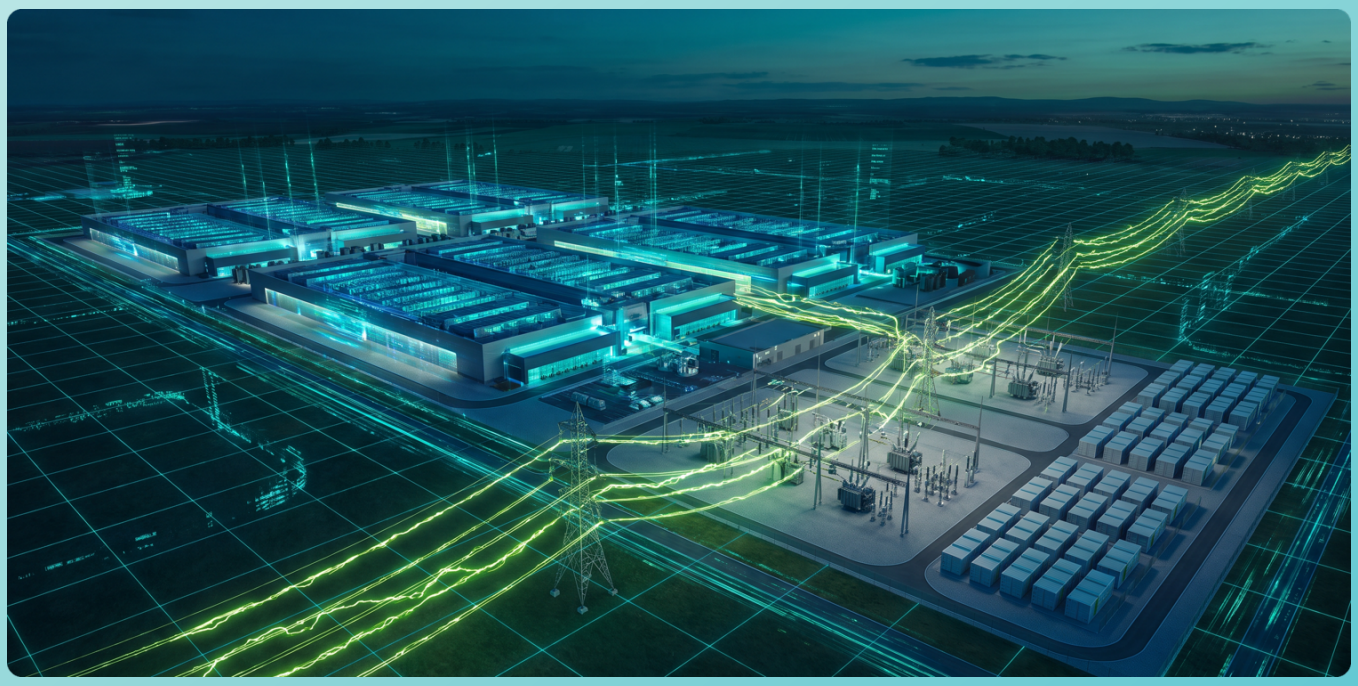
VedeniEnergy
Plugged Into The Energy Industry

OCTOBER 2026 - II

Flexible Large Loads: Turning Data Centers into Grid Assets

Curtailable interconnection, on-site resources, and workload orchestration as a bridge across the interconnection bottleneck

WHITEPAPER



SCAN TO EXPLORE

VedeniEnergy

- +1 463-266-4496
- www.vedeni.energy
- info@vedeni.energy
- North Vernon, IN 47265, US.

Abstract

The arrival of artificial intelligence data centers has turned large-load interconnection into one of the most pressing problems in the North American power sector. Requests for new service now exceed the physical capacity of many systems, and the conventional answer of building generation and transmission first takes far longer than a data center takes to construct. This paper examines **flexible large loads** as a near-term response: arrangements in which data centers and other large consumers accept curtailment or shift a portion of demand during peak and emergency conditions in exchange for faster and less costly connection. It reviews the evidence that **data center demand flexibility** can unlock tens of gigawatts of headroom, the technical enablers that make a **flexible data center** possible, the tariff and contract structures emerging for **curtailable interconnection**, the reliability and measurement challenges that remain unresolved, and the role of **load flexibility** on the **grid** in deferring build-out and supporting resource adequacy.

About this Paper

Vedeni Energy prepared this white paper as part of its Vedeni.Insights series of free resources on North American wholesale electricity markets. The analysis draws on regulatory filings, grid operator materials, national laboratory and university research, and published field demonstrations available through October 2026. It aims to describe the state of practice and the open questions evenly, without advocating for any single market design.

Audience and Use Cases

The paper is written for utility planners and regulatory staff, grid operator stakeholders, data center developers and their energy procurement teams, independent power producers, and investors who need a concise, current view of how flexible large-load service is taking shape and what it implies for interconnection timelines, cost allocation, and reliability.

Executive Summary

The interconnection queue for large loads has become the binding constraint on data center growth, and flexibility is quickly becoming the price of admission to a faster lane.

Data centers consumed roughly 4.4 percent of U.S. electricity in 2023, and national laboratory projections place that share between 6.7 and 12 percent by 2028. Utilities and grid operators are receiving requests far larger than their existing systems, while generation and transmission take years longer to build than the facilities they would serve. The result is a widening gap between when large loads want power and when the grid can firmly deliver it.

A growing body of research shows that this gap can be narrowed by asking new loads to be modestly flexible. Duke University's Nicholas Institute estimated in 2025 that the existing U.S. system could absorb 76 GW of new load if that load accepted curtailment averaging 0.25 percent of its annual uptime, rising to 98 GW at 0.5 percent and 126 GW at 1 percent. Follow-on work by Camus Energy, encoord, and Princeton's ZERO Lab found that a 500 MW data center pairing a flexible connection with its own capacity could reach full power in about two years, three to five years sooner than the traditional path, while drawing grid power in more than 99 percent of hours.

Regulators have moved quickly. FERC accepted the Southwest Power Pool's curtailable Conditional High Impact Large Load Service in June 2026, issued show cause orders to all six FERC-jurisdictional grid operators that same month directing them to consider transmission services for flexible large loads, and in October rejected most of PJM's interim curtailment framework on jurisdictional grounds. Texas Senate Bill 6 has written curtailment obligations into statute, and CAISO has proposed flexible interim and permanent interconnection services. The reliability questions are real, however: voltage-sensitive data center load has tripped offline in Virginia in blocks of 1,500 MW and, in July 2026, roughly 3,800 MW, and NERC has concluded that existing standards are not adequate for emerging large loads. Flexibility is therefore not a substitute for new infrastructure, but a credible bridge to it, provided the telemetry, verification, and contractual discipline keep pace.

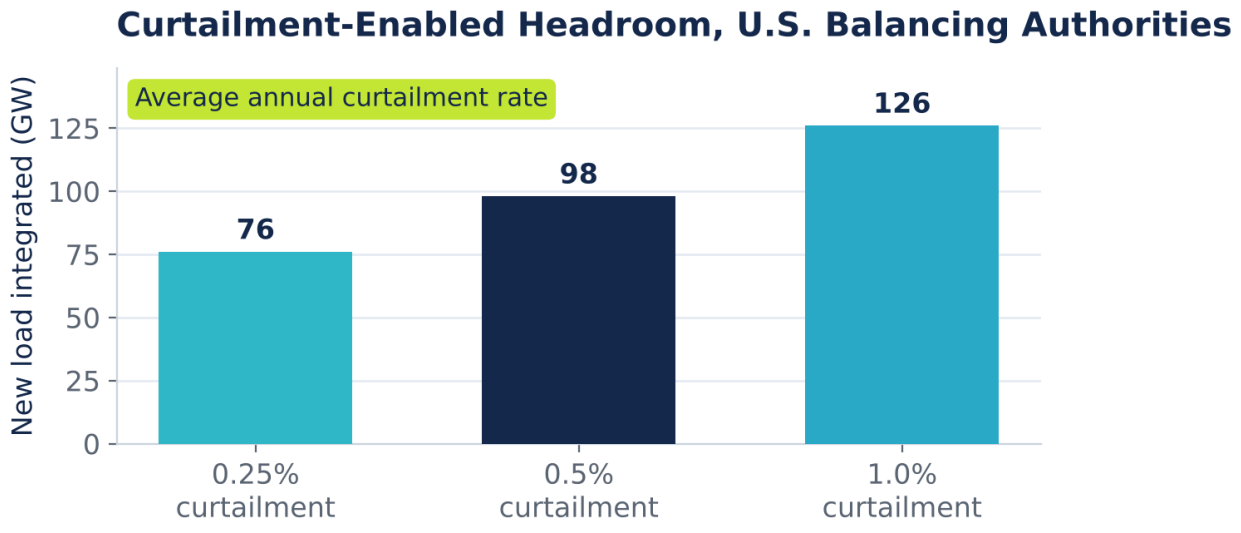


Figure 1. Estimated new load the existing U.S. system could integrate at different average annual curtailment rates (Duke University Nicholas Institute, 2025).

01 The Interconnection Bottleneck

For most of the past two decades, U.S. electricity demand was nearly flat. Planning processes, interconnection procedures, and tariffs were built for that world, and they are now straining under a wave of requests from facilities that each draw as much power as a small city.

Lawrence Berkeley National Laboratory found that data center electricity use rose from 58 TWh in 2014 to 176 TWh in 2023 and could reach between 325 and 580 TWh by 2028. Growth on that scale is not evenly distributed. It arrives in large, concentrated blocks at a limited number of substations, and it arrives quickly: a hyperscale facility can typically be built in eighteen to twenty-four months, while the network upgrades and new generation needed to serve it on a firm basis commonly take three to seven years.

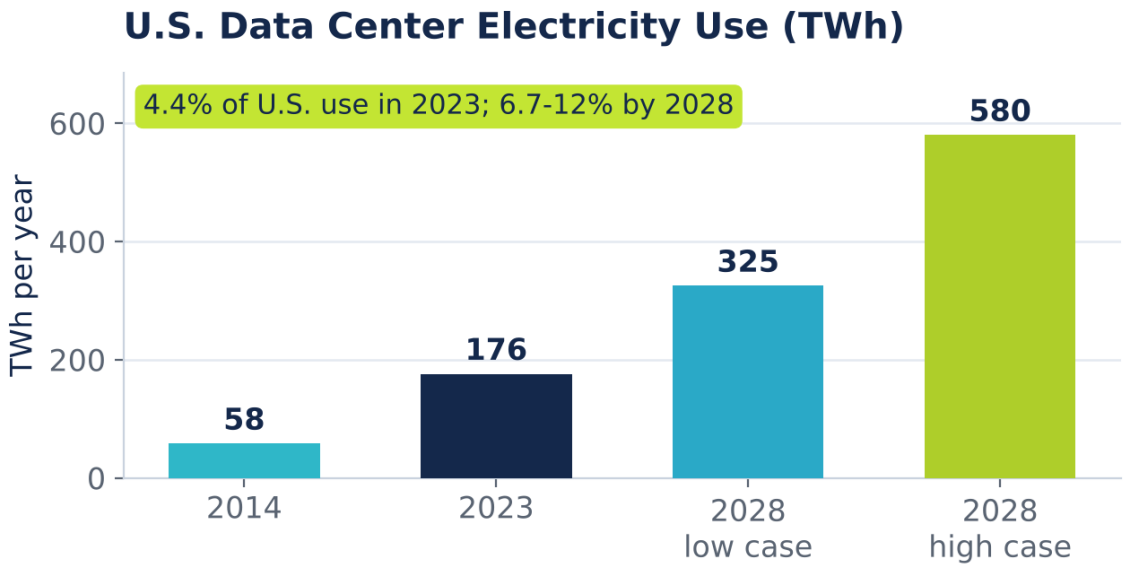


Figure 2. Historical and projected U.S. data center electricity consumption (Lawrence Berkeley National Laboratory, 2024).

The queues reflect that mismatch. ERCOT reported large-load requests in the hundreds of gigawatts during 2026, well above its record hourly peak of 91.1 GW set on July 22, 2026, with only a small fraction approved to energize. In the West, NV Energy's 2026 Integrated Resource Plan reported roughly 22,000 MW of data center interest against a 2025 combined system peak of about 8,250 MW. Many of these requests are speculative or duplicative, but even after heavy discounting the volume is large enough to overwhelm study processes designed for a few large customers at a time.

The economics of waiting are severe for both sides. A developer that cannot energize forgoes revenue on a multibillion-dollar investment, while utilities and their existing customers face the risk of building generation and transmission for load that may never materialize. Capacity markets have already felt the pressure; PJM's recent base residual auctions cleared at or near their price caps, and PJM has proposed a backstop procurement to close a projected shortfall. Firm service on traditional terms is becoming both slower and more expensive to provide.

Against this backdrop, a question that once seemed academic has moved to the center of regulatory debate: must every new large load be served as if it were inflexible every hour of the year? If a data center could agree to reduce grid withdrawals during the relatively few hours when the system is genuinely stressed, the grid capacity required to serve it could be dramatically smaller, and the connection could come much sooner.

★

KEY TAKEAWAY

The bottleneck is a timing mismatch: data centers can be built in about two years, while firm generation and transmission take three to seven. Flexibility is valuable because it closes that gap.

02 Curtailment-Enabled Headroom

Power systems are built for peak demand, and peaks are rare. That simple fact is the foundation of the case for flexible large loads.

Most bulk power systems operate well below their planned capacity for the great majority of hours. The Duke University Nicholas Institute formalized this observation in its February 2025 study, Rethinking Load Growth, introducing the concept of curtailment-enabled headroom: the amount of new load the existing system can absorb if that load accepts brief, infrequent reductions during the highest-stress hours. Across 22 balancing authorities representing most U.S. load, the study found that an average annual curtailment rate of 0.25 percent would accommodate about 76 GW of new load, roughly 10 percent of national aggregate peak demand. At 0.5 percent the figure rose to 98 GW, and at 1 percent to 126 GW.

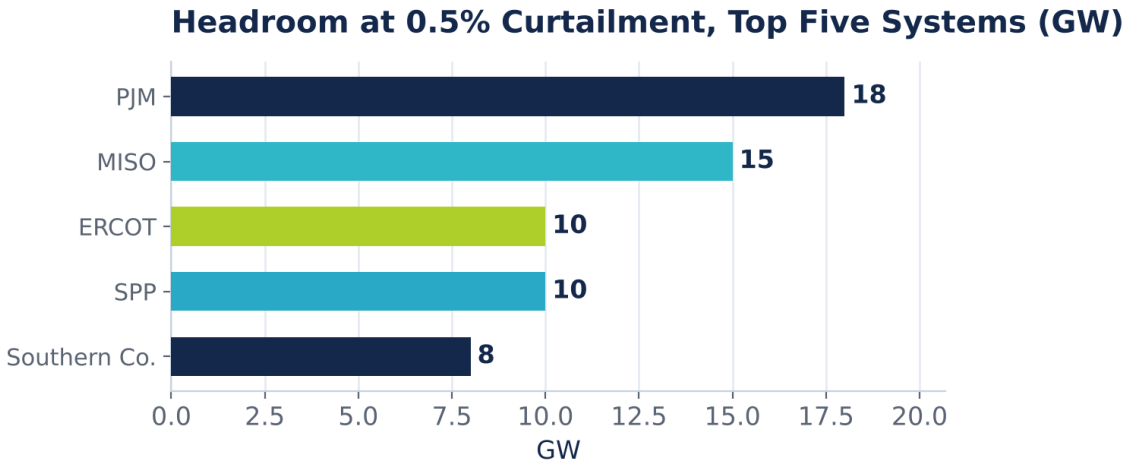


Figure 3. Systems with the largest curtailment-enabled headroom at a 0.5 percent annual curtailment rate (Duke University Nicholas Institute, 2025).

Two features of the results matter as much as the headline numbers. First, curtailment events were typically short, averaging a few hours. Second, most were partial: across the modeled events, at least half of the new load continued to run for the large majority of curtailment hours. A flexible data center, in other words, would rarely have to go dark; it would more often trim a portion of its grid draw while continuing to serve critical workloads.

A second Duke analysis, published in February 2026, extended the question from interconnection to economics. Modeling the decade ahead, it found that data center flexibility reduced the need for new peaking capacity and lowered projected retail electricity prices by roughly 0.5 to 2.8 percent, with data centers themselves seeing average prices $2.5 to $8.0 per MWh lower. Because flexibility shaves the most expensive peak hours, its value is concentrated precisely where system costs are

highest.

These results are system-level estimates, and they have been criticized when stretched beyond that scope. Headroom is not uniformly available at every substation, and local transmission constraints may bind in different hours than system-wide resource constraints. A curtailment rate that looks modest in aggregate may be much higher at a congested node. The findings are best read as evidence that flexibility is a material resource worth designing for, not as a promise that any individual site can connect without upgrades.

76 GW
headroom at 0.25% curtailment

126 GW
headroom at 1% curtailment

0.5-2.8%
projected retail price reduction

★

KEY TAKEAWAY

Modest, infrequent curtailment unlocks tens of gigawatts of system headroom, but the value is site-specific and must be validated against local transmission constraints.



Brief, partial curtailment during rare peak hours frees capacity the rest of the year.



03 Technical Enablers of the Flexible Data Center

A data center can reduce its grid draw in three broad ways: by supplying part of its own load on site, by storing energy, and by changing what its computers do and when.

The first enabler is on-site generation. Most large data centers already have backup generation sized to carry the facility through outages, but those assets have traditionally been reserved for emergencies. Regulatory clarity has begun to widen their use. In May 2025 the U.S. Environmental Protection Agency confirmed that stationary emergency engines may run up to 50 hours per year in non-emergency demand response arrangements that help prevent local load shed, under defined conditions. Developers are also adding dedicated bridge generation, fuel cells, and co-located plants, and EPRI's DCFlex initiative is testing renewable diesel substitutes to reduce the emissions trade-off of running backup engines more often.

The second enabler is battery storage. Batteries respond in milliseconds, can provide reserves and voltage support, and can bridge the minutes or hours needed to ramp workloads down gracefully. In Hillsboro, Oregon, Aligned Data Centers and Calibrant Energy are deploying a 31 MW, 62 MWh battery that Portland General Electric can call on during critical periods, allowing the campus to expand years before grid upgrades would otherwise permit. The Camus, encoord, and Princeton study found that the least-cost portfolio for a 500 MW site ranged from about 49 MW of storage and 11 MW of on-site generation at a lightly constrained location to about 155 MW of storage and 48 MW of generation at a moderately constrained one.

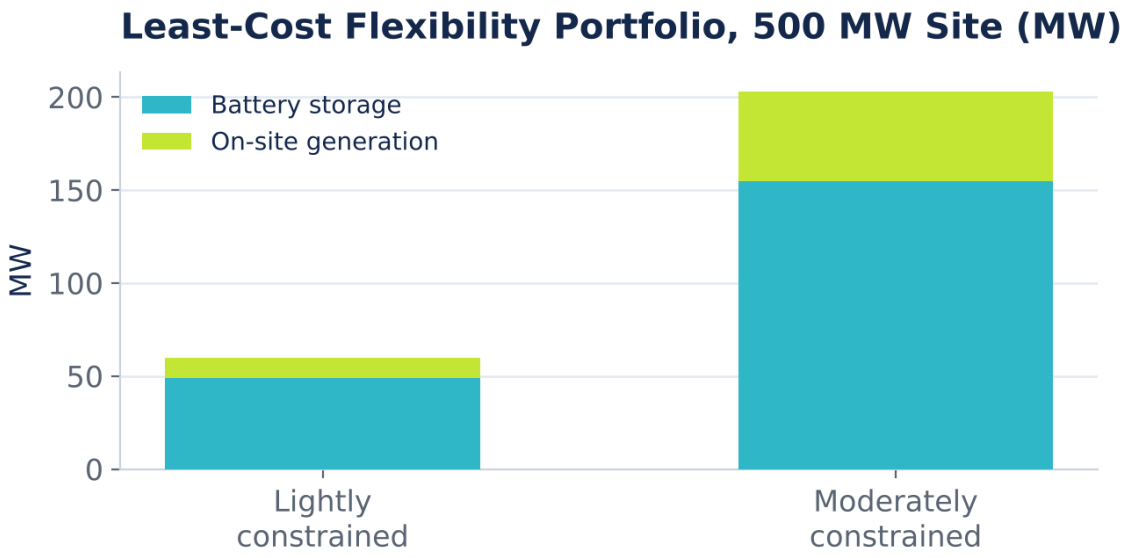


Figure 4. Modeled least-cost on-site flexibility portfolios for a 500 MW data center (Camus Energy, encoord, and Princeton ZERO Lab, 2025).

The third enabler, and the most novel, is workload orchestration. Not all computing is equally urgent. Artificial intelligence training runs, batch analytics, and media processing can often be slowed, paused, or moved to another region, while latency-sensitive inference and customer-facing services cannot. Google began shifting non-urgent compute across time and location in 2023, reduced machine-learning load during three grid events with the Omaha Public Power District in 2024, and in August 2025 signed agreements with Indiana Michigan Power and the Tennessee Valley Authority to deliver demand response specifically from machine-learning workloads.

The clearest public evidence of software-only flexibility comes from Phoenix. In a May 2025 demonstration under EPRI's DCFlex initiative, Emerald AI's Conductor platform steered a 256-GPU cluster at an Oracle Cloud Infrastructure data center to cut power by 25 percent for three hours during a Salt River Project peak event while meeting predefined performance commitments for each job. The control levers included GPU power capping, frequency scaling, job pausing, and resource reallocation, applied according to tiered service-level agreements that allowed different jobs to tolerate different slowdowns. DCFlex expanded to nine demonstration sites in February 2026 and plans to bring a 96 MW facility in Manassas, Virginia, designed for flexibility from the outset, online in late 2026.

Workload flexibility does not require a data center to stop computing. It requires the operator to know which jobs can bend, by how much, and for how long, and to commit to that knowledge in advance.

★

KEY TAKEAWAY

Generation, storage, and orchestration are complementary. The most robust flexible data centers will layer all three so that no single lever must carry the full curtailment.

04 Tariffs and Contracts for Curtailable Interconnection

Flexibility becomes a grid asset only when it is written into an enforceable product. In 2026 those products began to appear in tariffs, statutes, and bilateral agreements across the country.

The Southwest Power Pool has moved furthest at the regional level. FERC accepted SPP's High Impact Large Load study process in January 2026, offering a path to service within roughly 90 days for qualifying loads, and accepted the Conditional High Impact Large Load Service, known as CHILLS, on June 5, 2026. CHILLS provides non-firm transmission service using available capacity for a term of one to seven years while a customer secures designated resources or completes the upgrades needed for firm service. The CHILLS load is separately metered, is not treated as critical load, and can be interrupted when the system is constrained or in emergency. SPP is separately developing a Price Adaptive Load service for customers willing to respond to prices rather than directives.

Texas has taken a statutory route. Senate Bill 6, enacted in 2025, applies to loads of 75 MW and above. It authorizes ERCOT to direct new large loads with substantial backup generation to deploy that generation or curtail during emergencies, requires utilities to be able to curtail new large loads during firm load shed, and directs the Public Utility Commission to establish a large-load demand management service. ERCOT brought the first version of its Large Load Curtailment Manager into service in mid-2026, held an implementation workshop in October, and has created a Provisional Controllable Load Resource path for loads seeking more grid capacity ahead of firm service.

At the federal level, FERC chose show cause orders over a national rule. On June 18, 2026 it directed PJM, MISO, SPP, CAISO, ISO New England, and NYISO to justify or reform their tariffs for loads above 50 MW connecting above 69 kV, with an explicit request to consider new transmission services for loads able to limit withdrawals during stress. CAISO responded with a straw proposal for two interconnection services, a Flexible Interim Load Interconnection that converts to firm service when upgrades are complete and a Flexible Load Interconnection Permanent option, with a FERC filing targeted for November 16, 2026. PJM's Interim Resource Adequacy Service, which would have let large loads connect without matching capacity if they accepted pre-emergency curtailment, was largely rejected by FERC on October 9, 2026 as intruding on state authority over retail service and potentially discriminating against large loads, although FERC approved PJM's large-load definition and registry.

Mechanism	Jurisdiction	Curtailment basis	Status (Oct. 2026)
CHILLS	SPP	Non-firm service, 1-7 year bridge to firm	Accepted by FERC, June 2026
SB 6 curtailment and LLCM	ERCOT / Texas	Emergency and firm load shed directives	Rules due Dec. 2026
FILI and FLIP	CAISO	Interim or permanent flexible interconnection	Filing due Nov. 2026
Interim Resource Adequacy Service	PJM	Pre-emergency curtailment absent capacity	Largely rejected, Oct. 2026
Machine-learning demand response	I and M, TVA	Bilateral utility agreements	Signed Aug. 2025
Utility-dispatched battery	PGE (Oregon)	Dispatchable standby program	Under development

Table 1. Selected flexible large-load mechanisms and their status.

Retail tariffs are evolving in parallel. Lawrence Berkeley National Laboratory counted more than fifteen utilities with large-load tariff proposals by early 2026, with contract terms coalescing around ten to fifteen years and features such as minimum demand charges, exit fees, collateral, and ramp schedules. A growing number pair those protections with bring-your-own-capacity options, as in NV Energy's Clean Transition Tariff and a large-load tariff Microsoft proposed in Nevada, and with explicit consideration of flexibility. The common thread is a trade: a customer that accepts more operational or financial responsibility receives earlier or cheaper access to the grid.



KEY TAKEAWAY

Curtailable interconnection is shifting from bespoke deals to standardized products, but the federal-state jurisdictional line, as FERC's PJM decision shows, still shapes what any grid operator can offer.

05

Reliability, Telemetry, and Measurement

A flexible load is only as valuable as the confidence a system operator can place in it. That confidence depends on how the load behaves both when it is asked to curtail and when it is not.

The reliability concern is not hypothetical. In July 2024 a fault on a 230 kV line in Northern Virginia caused about 1,500 MW of data center load to transfer abruptly to backup power, an event NERC documented as an early case study of voltage-sensitive load. A similar event occurred in 2025, and on July 22, 2026 approximately 3,800 MW of data center load disconnected in two waves, reducing PJM's total load by about 3.8 percent. The underlying cause is protective: uninterruptible power supplies and transfer logic are designed to shield servers from grid disturbances, but when many facilities react to the same voltage dip, the grid experiences a sudden loss of load that can push voltage and frequency outside normal bounds.

NERC has responded on several fronts. A March 2026 gap assessment concluded that existing standards and interconnection practices are inadequate for emerging large loads, a voluntary reliability guideline followed in May, and a Level 3 alert issued on May 4, 2026 asked registered entities to take specific actions. NERC has told FERC it intends to file registry criteria and initial standards for large loads by the end of 2026, although industry-wide ride-through thresholds may follow later. Grid operators are filling the gap with interconnection requirements of their own, such as real-time telemetry, operator-initiated disconnection capability, ramp limits, and ride-through performance.

1,500 MW
Virginia load loss, July 2024

3,800 MW
PJM data center trip, July 2026

Dec. 2026
NERC large-load standards target

Flexibility introduces a second set of measurement problems. Traditional demand response settles against a baseline, an estimate of what the customer would have consumed absent the event. For a data center whose load is driven by changing compute jobs, constructing a credible baseline is difficult, and an operator could in principle schedule discretionary work so that a curtailment appears larger than it is. Curtailable interconnection products avoid part of this problem by defining performance as a ceiling on withdrawals rather than a reduction from a baseline. SPP's requirement that conditional load be separately metered is an example: the question becomes whether the load stayed under its limit, which revenue-grade metering can answer directly.

The harder question is enforcement during a true emergency. Contracts can specify penalties, automatic disconnection schemes, and audits, but a curtailment that fails during a capacity shortfall

cannot be remedied after the fact. Operators are therefore likely to require physical assurance, through protection schemes, telemetry, and periodic testing, before they count flexible load in planning studies. EPRI's Flex MOSAIC framework, released in March 2026, is an early attempt to give utilities and data centers a common vocabulary for describing and verifying flexibility capabilities.

> *The test of a flexible load is not whether it can curtail in a demonstration, but whether an operator will stake reliability on it during the worst hour of the year.*

KEY TAKEAWAY

Ceiling-based performance, revenue-grade metering, operator control, and ride-through requirements are becoming the practical foundation for counting flexibility as a reliability resource.



Revenue-grade metering and real-time telemetry turn flexibility into a verifiable service.

06 Deferring Build-Out and Supporting Resource Adequacy

The economic case for flexible large loads rests on a simple idea: infrastructure sized for a few hours of peak demand is expensive, and avoiding or delaying it creates value for everyone who pays for the grid.

The Camus, encoord, and Princeton ZERO Lab study, released in December 2025 and funded in part by Google, tested this idea using real transmission data from a PJM utility and six candidate 500 MW sites. Combining a flexible grid connection, which addresses transmission limits, with a bring-your-own-capacity arrangement, which addresses generation limits, brought each site to full operation in roughly two years. Grid power remained available in more than 99 percent of hours, and on-site resources were dispatched only about 40 to 70 hours per year.

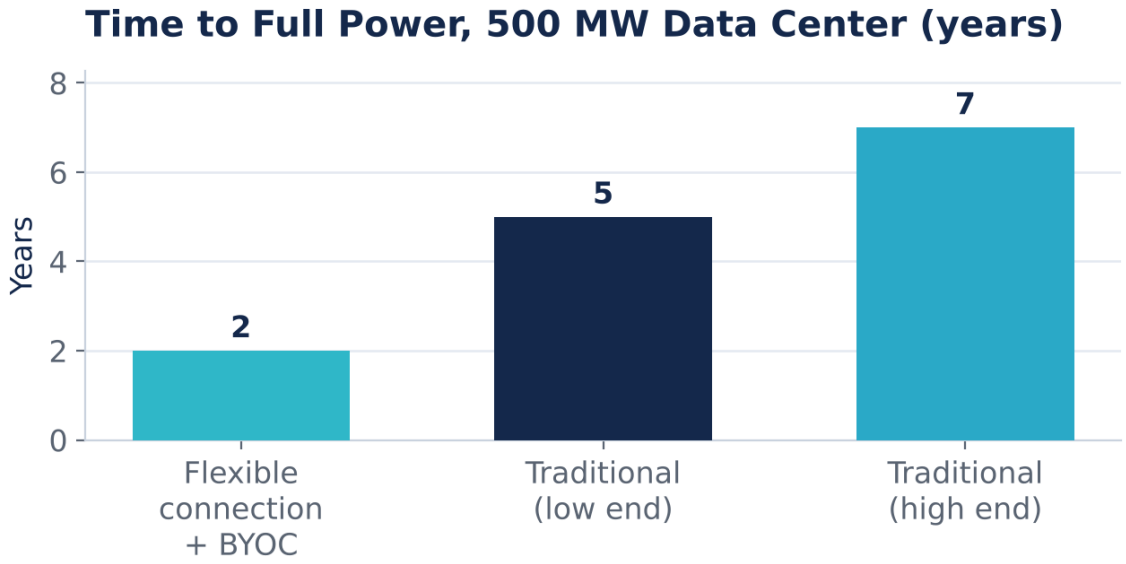


Figure 5. Modeled time to full power for a 500 MW data center under flexible and traditional interconnection (Camus Energy, encoord, and Princeton ZERO Lab, 2025).

99%+	40-70 h	$733M/GW
hours served by grid power	annual on-site dispatch	contribution to incremental costs

The system-cost findings are equally significant. Flexible data centers in the study contributed about $733 million per gigawatt toward the costs of their incremental load and reduced the net system cost increase by roughly 96 percent compared with the same volume of inflexible load. For

the developer, connecting three years earlier was worth an estimated $2.3 to $3.2 billion in additional earnings per site, easily covering the lifecycle cost of the batteries and generators needed to make flexibility dependable.

For resource adequacy, flexible load changes the planning problem rather than eliminating it. A load that reliably curtails during the system's tightest hours does not add to the peak that reserve margins are built around, so its accredited contribution to capacity need is smaller than its nameplate demand. ERCOT has begun reflecting SB 6 curtailment assumptions in its reserve margin scenarios for the same reason. The risk runs in both directions: counting flexibility that fails to materialize understates the true capacity requirement, while ignoring dependable flexibility leads to overbuilding that existing customers ultimately fund.

Flexibility also interacts with markets in ways still being worked out. A large block of load that steps off the grid when prices spike will dampen scarcity pricing, which lowers costs for consumers but may weaken the signals that attract new generation. Several of PJM's stakeholders raised this concern about capacity price formation, and FERC Commissioner Rosner cautioned that excluding interim large loads from PJM's capacity demand curve could distort the market. Designing flexible service so that it complements, rather than undermines, investment signals is one of the central tasks for the next round of filings.

> *The choice facing a large load is rarely between revenue and curtailment. It is between revenue now under flexible terms and revenue years later under firm ones.*

★

KEY TAKEAWAY

Flexibility is a bridge to new infrastructure, not a replacement for it. Its value lies in connecting load years earlier and right-sizing what must ultimately be built.

07 Outlook and Implications for the West

The West faces the same pressures as the rest of the country, with the added complexity of a largely bilateral market structure that is now consolidating around two regional day-ahead markets.

Much of the Western Interconnection lies outside an RTO, so large-load interconnection is governed primarily by vertically integrated utilities and state commissions rather than by a single regional tariff. That structure places the design of flexible service in retail tariffs and bilateral contracts. Nevada illustrates the stakes: data center interest reported in NV Energy's 2026 resource plan was nearly three times the utilities' combined system peak. Oregon regulators approved a large-load rate class for Portland General Electric in 2026, and PGE's study process already catalogs each new data center's flexibility capabilities. The Hillsboro battery project shows how a utility program for dispatchable customer resources can become a de facto flexible interconnection product.

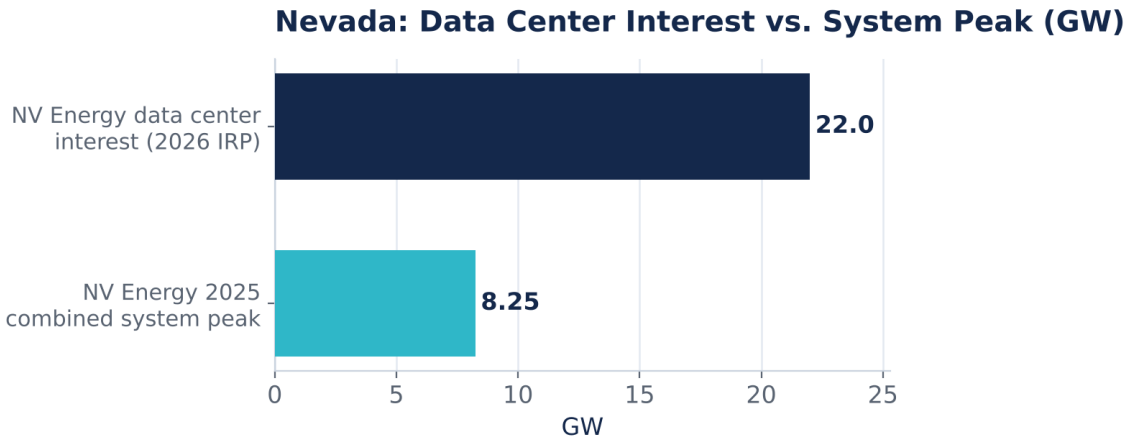


Figure 6. Data center interest relative to system peak in NV Energy's service territory (NV Energy 2026 Integrated Resource Plan, as reported).

CAISO's flexible interconnection proposals will matter well beyond California once filed, because they will be the first Western ISO framework for curtailable load service. They also expose a practical limit: participating transmission owners, not CAISO, own the retail and wholesale provisions that govern study and cost recovery for new load, so any product must be mirrored in utility tariffs before customers can use it. Across the broader West, the start of CAISO's Extended Day-Ahead Market in May 2026 and SPP's Markets+ planned launch in 2027 will change how flexible load is scheduled and priced, even though neither market directly governs interconnection.

Several developments will shape the next twelve months. FERC will rule on the six show cause proceedings and CAISO's filing; NERC is expected to file its first large-load standards; the Texas

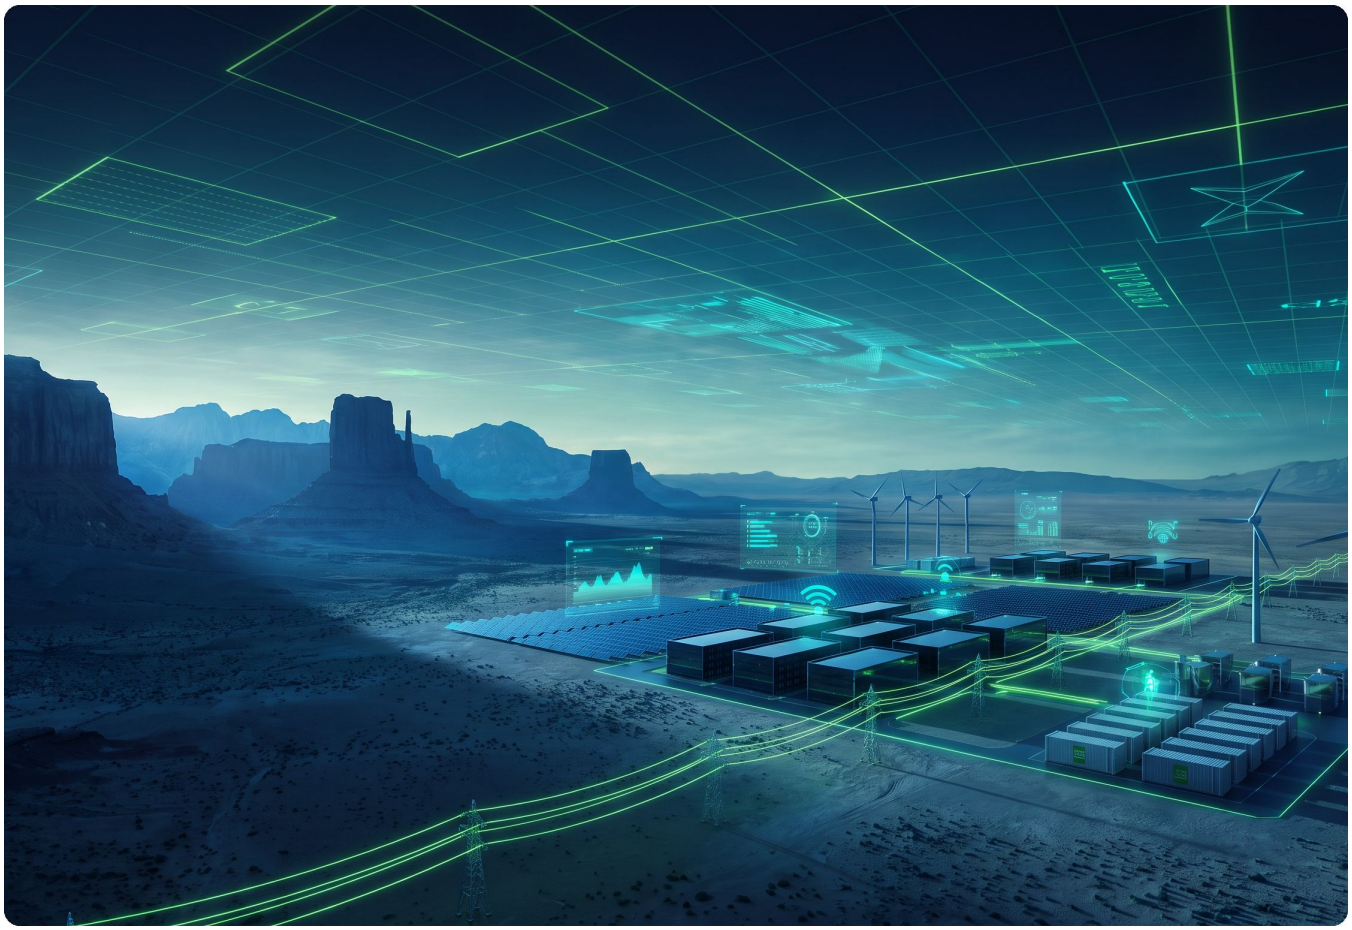
commission must complete its SB 6 rules; EPRI's flexibility-by-design demonstration in Virginia should begin producing results; and more utilities will test whether flexibility provisions in retail tariffs attract customers. The likely end state is a menu of service tiers, from fully firm service at full cost to conditional and price-responsive service at lower cost and faster timing, each with clear performance obligations.

For Western market participants, the practical implications are concrete. Developers should design facilities with flexibility in mind from the outset, including metering, controls, on-site resources, and internal service tiers that define which workloads can bend. Utilities and regulators should define flexibility as a measurable withdrawal limit with physical assurance, and integrate it into resource plans and transmission studies. Market designers should ensure that flexible load participates in a way that preserves price signals for the generation and transmission the region will still need.

★

KEY TAKEAWAY

In the West, flexible large-load service will be built mostly through utility tariffs and bilateral agreements, making state commissions and transmission owners the decisive actors.



A menu of firm, conditional, and price-responsive service is emerging across Western systems.

Key Findings

The evidence reviewed in this paper supports the following findings.

- → Large-load interconnection has become a timing problem: data centers can be built in about two years, while the firm generation and transmission to serve them typically take three to seven.
- → Modest flexibility unlocks substantial capacity. Average annual curtailment of 0.25 to 1 percent could allow the existing U.S. system to integrate 76 to 126 GW of new load.
- → On-site generation, batteries, and workload orchestration are complementary enablers; field demonstrations have shown 25 percent cluster power reductions sustained for three hours without breaching performance commitments.
- → Curtailable interconnection products now exist or are pending in SPP, ERCOT, and CAISO, while FERC has directed all six jurisdictional grid operators to consider transmission services for flexible loads.
- → Reliability risks from voltage-sensitive data center load are real and documented, and NERC standards for large loads are still being developed.
- → Flexibility can reduce system costs and connect load years sooner, but it complements rather than replaces new infrastructure and must be designed to preserve investment signals.

Method and Scope

This paper synthesizes public sources available through October 21, 2026, including FERC orders and dockets, grid operator stakeholder materials, NERC assessments, state legislation, national laboratory and university studies, and published field demonstrations. Quantitative results are reported as published by their authors; Vedeni Energy did not independently model curtailment headroom or system costs.

Conclusion

Flexibility is the fastest bridge between load growth and the grid that must serve it

The electricity system is confronting load growth on a scale and at a speed it was not designed to absorb. Building generation and transmission remains essential, but it cannot happen fast enough to meet every request on firm terms. Flexible large loads offer a way to use the capacity that already exists during the many hours when it is idle, and to reserve new construction for the needs that flexibility cannot meet.

The building blocks are now visible: research quantifying system headroom, demonstrations proving that compute can bend, tariffs that convert flexibility into enforceable service, and a federal regulator pressing every major grid operator to act. What remains is the harder work of making flexibility dependable enough to plan around, through clear performance definitions, physical assurance, sound measurement, and market rules that reward reliability rather than eroding it. Utilities, developers, and regulators that solve those problems first will connect load faster and at lower cost, and will turn data centers from a strain on the grid into one of its most valuable assets.

References

01 **T. H. Norris, T. Profeta, D. Patino-Echeverri, and A. Cowie-Haskell**
Rethinking Load Growth: Assessing the Potential for Integration of Large Flexible Loads in US Power Systems. Durham, NC: Nicholas Institute for Energy, Environment and Sustainability, Duke University, Feb. 2025.

02 **M. Ross and J. Ewing**
Data Centers and Generation Capacity over the Next Decade: Potential Benefits of Flexibility. Durham, NC: Nicholas Institute, Duke University, Feb. 17, 2026.

03 **A. Shehabi et al.**
2024 United States Data Center Energy Usage Report. Berkeley, CA: Lawrence Berkeley National Laboratory, Dec. 2024.

04 **C. Brancucci, D. Cutler, and J. Jenkins**
Flexible Data Centers: A Faster, More Affordable Path to Power. Camus Energy, encoord, and Princeton University ZERO Lab, Dec. 2025.

05 **P. Colangelo et al.**
Turning AI Data Centers into Grid-Interactive Assets: Results from a Field Demonstration in Phoenix, Arizona. arXiv:2507.00909, Jul. 2025.

06 **American Public Power Association**
SRP participates in artificial intelligence data center demonstration. Public Power Current, Jul. 2025.

07 **Electric Power Research Institute**
EPRI's DCFlex Initiative Expands to Nine Demonstration Sites Across U.S., Europe. Press release, Feb. 2, 2026.

08 **Electric Power Research Institute**
Data Center Flexible Load Initiative (DCFlex): Flex MOSAIC and Demonstration. Briefing note, 2026.

09 **M. Terrell / Google**
Google expands data center demand response to machine learning workloads (Indiana Michigan Power and TVA agreements). Reported by Latitude Media and Data Center Dynamics, Aug. 2025.

10 **U.S. Environmental Protection Agency**
EPA Issues Clarification to Help Power Data Centers: Use of Backup Generators to Maintain Grid Reliability. Fact sheet and FAQ, May 1, 2025.

11 **Canary Media**
In a first, a data center is using a big battery to get online faster; Oregon utility tries out a faster, cheaper way to power data centers. 2025-2026.

References (continued)

12 **Federal Energy Regulatory Commission**
Interconnection of Large Loads to the Interstate Transmission System, Docket No. RM26-4-000. Advance notice of proposed rulemaking and docket materials, 2025-2026.

13 **White and Case LLP**
FERC orders grid operators to promptly revise or justify interconnection rules for data centers and large loads. Insight alert, Jun. 2026.

14 **McGuireWoods LLP**
FERC Issues Section 206 Show Cause Orders Directing All Six RTOs/ISOs to Justify or Reform Large Load Integration Rules. Client alert, Jun. 2026.

15 **Utility Dive**
FERC approves SPP non-firm, large-load transmission service (CHILLS). Jun. 2026.

16 **Southwest Power Pool**
Conditional High Impact Large Load Service Operating Agreement (Attachment BA, Addendum 1). SPP Tariff, 2026.

17 **Utility Dive**
FERC rejects most of PJM's large load interconnection plan. Oct. 12, 2026.

18 **PJM Interconnection**
Board Decisional Letter on Critical Issue Fast Path - Large Load Additions. Jan. 16, 2026.

19 **RTO Insider**
PJM Drops Non-capacity Backed Load, Shifts Focus to Resource Queue, PRD. Oct. 9, 2025.

20 **CAISO**
Large Load Considerations: Straw Proposal; Motion for Abeyance, Docket No. EL26-71-000. Aug. 2026.

21 **Baker Botts LLP**
Texas Senate Bill 6: Understanding the Impacts to Large Loads and Co-located Generation. Jul. 2025.

22 **ERCOT**
ERCOT Monthly, June 2026; Market Notice M-A090226-01 on SB 6 Large Load curtailment workshop, Sept. 2, 2026.

References (continued)

23 **U.S. Energy Information Administration**
Report on ERCOT record hourly load of 91.1 GW set July 22, 2026. Oct. 15, 2026.

24 **North American Electric Reliability Corporation**
Incident Review: Considering Simultaneous Voltage-Sensitive Load Reductions. Jan. 2025.

25 **Utility Dive**
PJM eyes data center, crypto reliability requirements after 3.8 GW of load trips offline. Aug. 2026.

26 **POWER Magazine**
Bringing a Power Plant's Worth of Load onto the Grid: NERC's Interconnection Reckoning. 2026.

27 **Lawrence Berkeley National Laboratory**
Electricity Rate Designs for Large Loads. Technical brief, Aug. 2026; Large Load Tariffs: Elements and Themes, presentation to the Washington UTC, Apr. 27, 2026.

28 **RCG Economics**
Nevada's Data-Center Boom Is Outrunning Its Power Grid (citing NV Energy 2026 Integrated Resource Plan). Jul. 21, 2026.

29 **Utility Dive**
Data centers are ready to negotiate flexibility for speed. Jun. 2026.

Disclaimer. The material presented in this paper is provided for informational purposes only and represents the authors' views and interpretations at the time of writing. While every effort has been made to ensure accuracy and completeness, neither the authors nor their affiliated organizations make any warranty, express or implied, regarding its correctness or suitability for any particular purpose. This document does not constitute legal, financial, or technical advice, and readers should independently verify all facts and seek professional counsel before acting on any information contained herein. Neither the authors nor their organizations accept liability for any loss or damage arising directly or indirectly from the use of this publication.

About Vedeni Energy



Vedeni Energy offers specialized services designed to help businesses navigate the complexities of the modern energy landscape, tailored to utilities, independent power producers, regulators, and other stakeholders through strategic insights, expert guidance, and innovative solutions.



Vedeni.Insights is Vedeni Energy's free library of white papers, reports, and in-depth technical analyses covering market trends, regulatory change, emerging technologies, and strategic investment opportunities.



Vedeni.IQ+ provides actionable wholesale electric power market intelligence, covering the stakeholder forums of North American wholesale markets and distilling complex market information into clear insights that help organizations elevate strategy, influence policy, and identify new opportunities.

TO LEARN MORE, VISIT US AT

W W W . V E D E N I . E N E R G Y

+1 463-266-4496 • info@vedeni.energy • North Vernon, IN 47265, US.