AI Data Center Growth: Challenges and Opportunities for Electric Power Systems

Oxford Energy Research Network
19 May 2026
University of Oxford

Pramod P. Khargonekar
University of California, Irvine

UCI University of California, Irvine

Source: FT, 17 April 2026

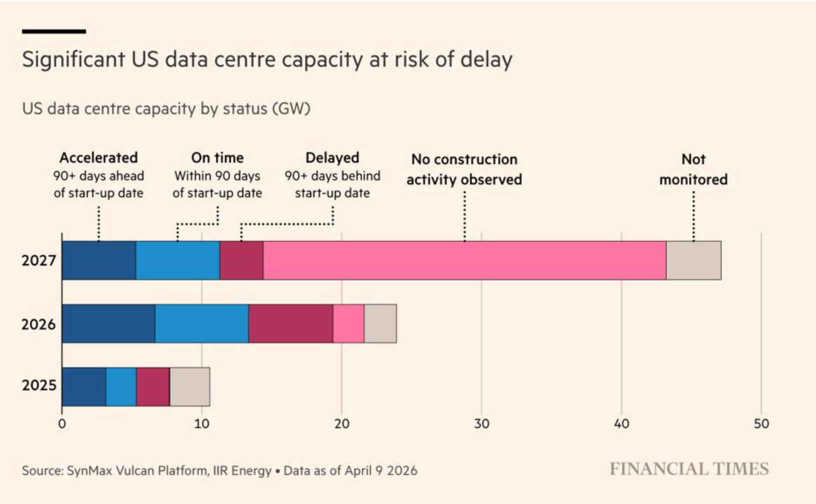

Source: WSJ, 7 April 2026

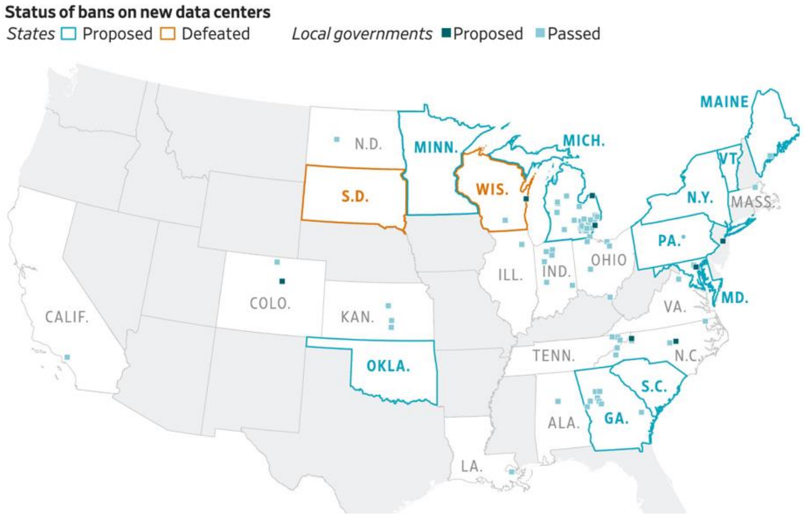

Source: Bloomberg, 16 April 2026

Texas Grid Issues Dramatic Power Demand Growth Projection

A power station in Corpus Christi, Texas. *Photographer: Eddie Seal/Bloomberg*

Source: Google, 19 March 2026

A new milestone for smart, affordable electricity growth

Mar 19, 2026
2 min read

We've now signed 1 GW of data center demand response with utility partners, harnessing our electricity growth for smarter utilization of U.S. electricity systems.

Michael Terrell
Head of Advanced Energy

Share

Outline

- AI data center growth in the US
- Major challenges in connecting to the power grid
- Data center flexibility
- Tariff design for data centers
- Conclusions

US Electric Energy Consumption Has Been Flattish

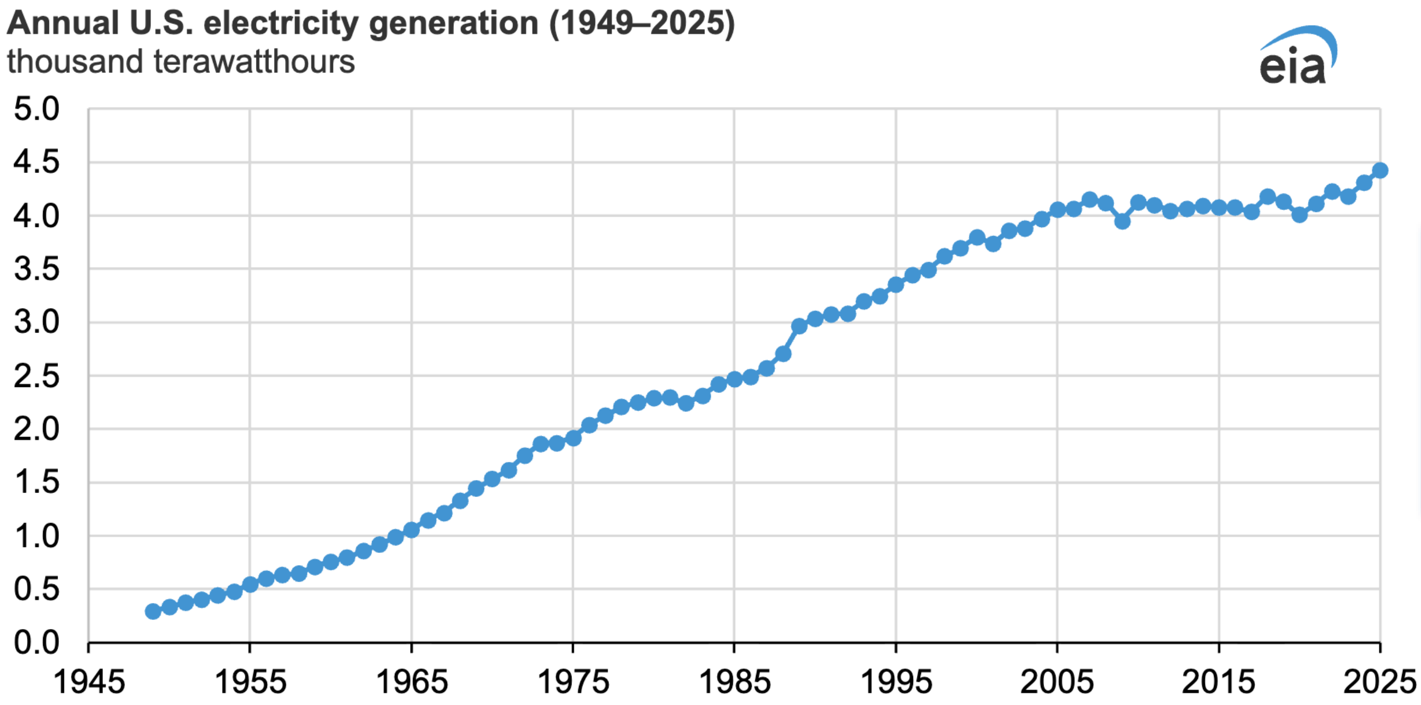

Data source: U.S. Department of Energy, *Monthly Energy Review* and *Electricity Data Browser*

4430 TWh - USA
322 TWh - UK
2625 TWh – EU
10370 TWh – China
2054 TWh - India

Source: US Department of Energy

AI Capex Boom

Alphabet, Amazon, Microsoft and Meta
$640B in 2026

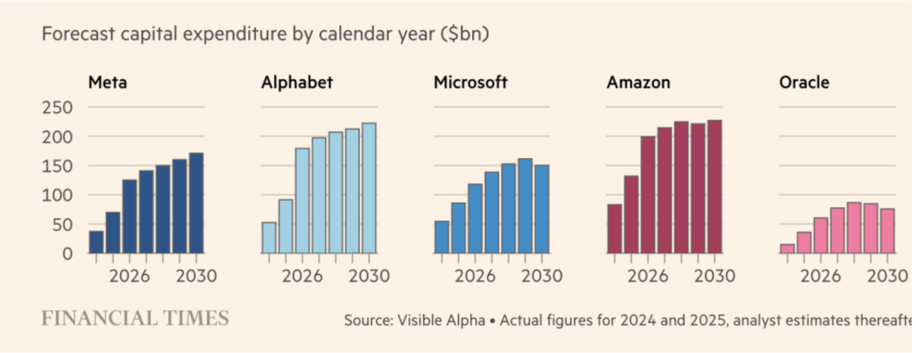

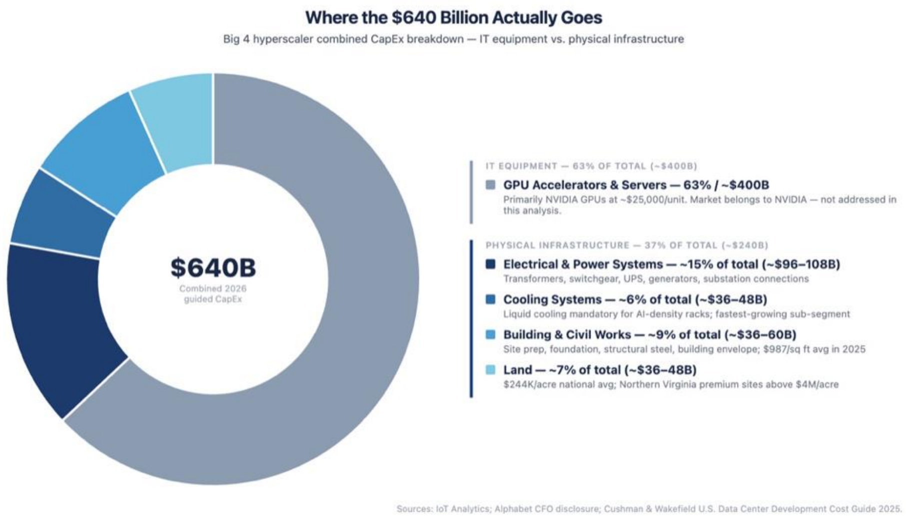

$96-108B
in Electrical and Power Infrastructure

Source: "Will the AI data centre boom become a $9tn bust"
Financial Times, March 28, 2026

Source: Avanza Energy, 2026

Current and Planned Data Centers

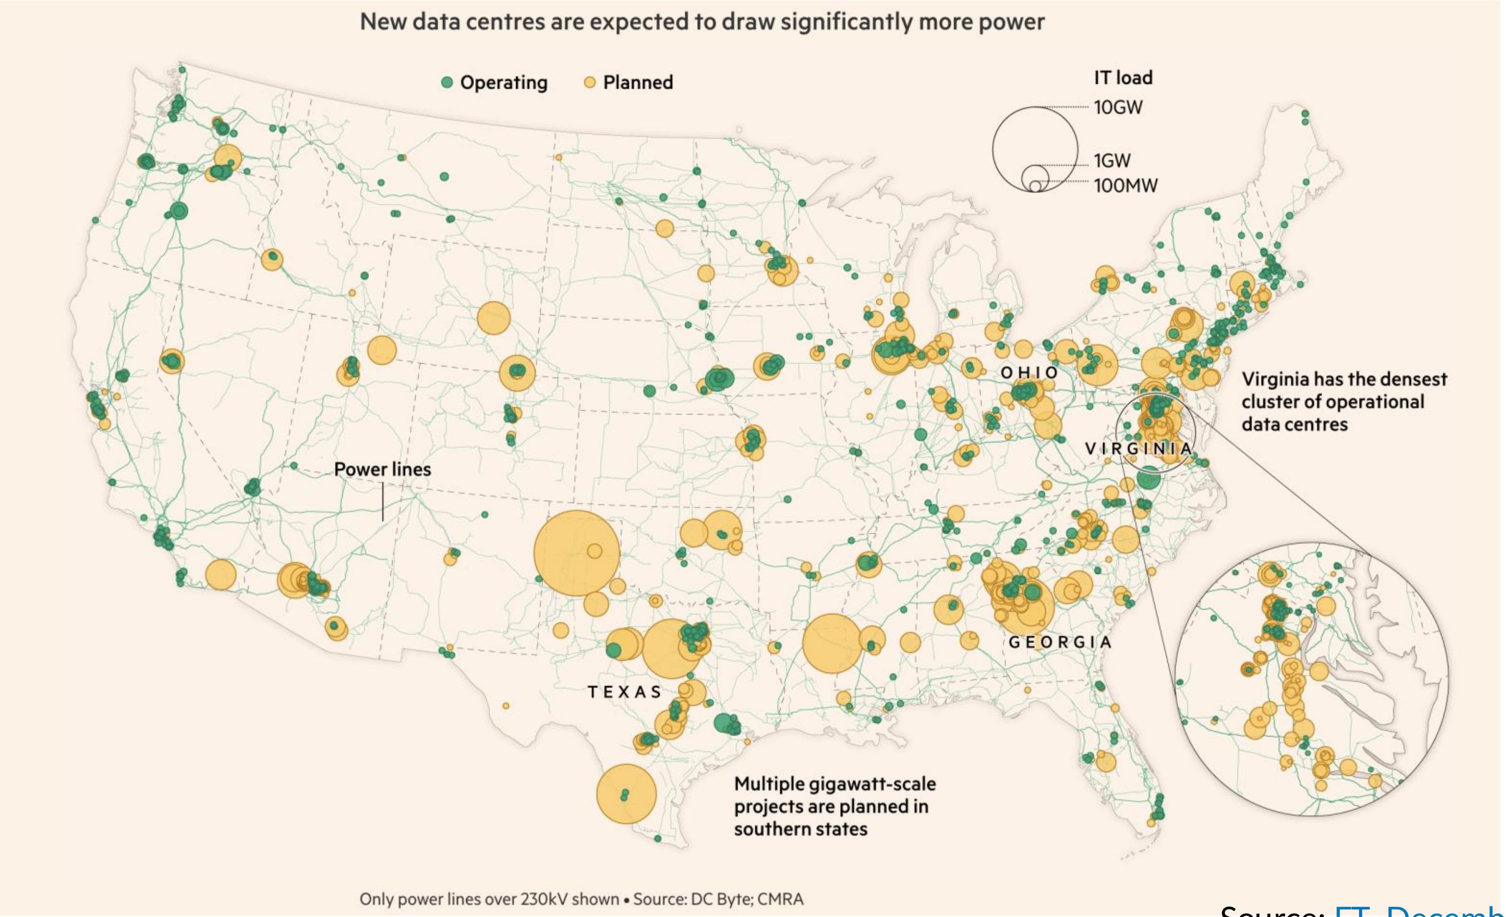

Source: [FT, December 8 2025]

Projections of Data Center Electricity Consumption

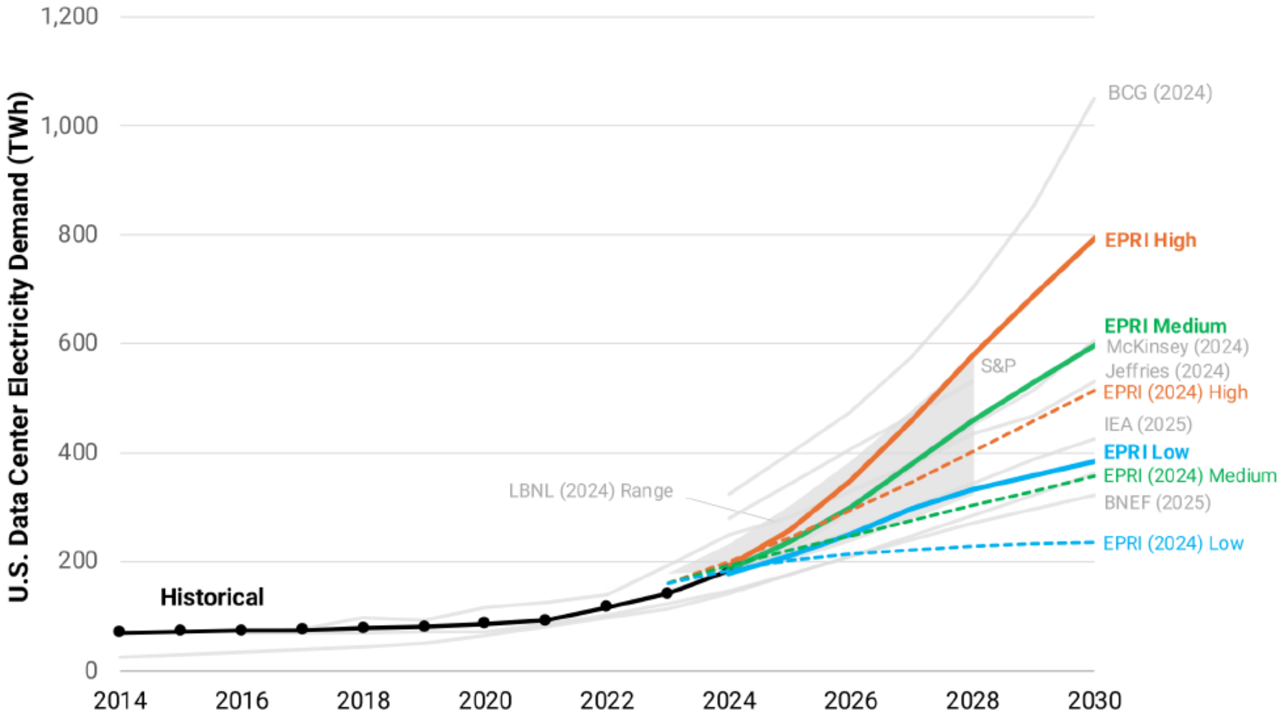

Figure ES-2. Comparison of U.S. data center annual electricity consumption projections. *Projections in this study span a similar range to the LBNL (2024) report. Shaded band shows scenarios from LBNL (2024); lines show recent external estimates, including BCG (2024), BloombergNEF (2025), EPRI (2024a), IEA (2025), Jefferies (2024), McKinsey (2024), S&P (2024). EPRI estimates include small- and large-scale data centers as well as cryptocurrency mining.*

Source: Electric Power Research Institute, 2026

Key Sources of Uncertainty

- Future market for AI technologies in the broader economy
- Competition among AI companies
- Improvements in energy efficiency of AI computing
- Shifts in AI data centers: training to inference
- Availability of chips: GPUs and HBM
- Provision of electric energy to AI data centers: generation, electrical infrastructure, permitting, …

Electric Power for AI: Constraints, Goals, and Aspirations

- Electric power system reliability
- Electric energy affordability for households
- Enabling AI for economic growth and competitiveness
- Deploying AI with environmental sustainability

North American Power Grid

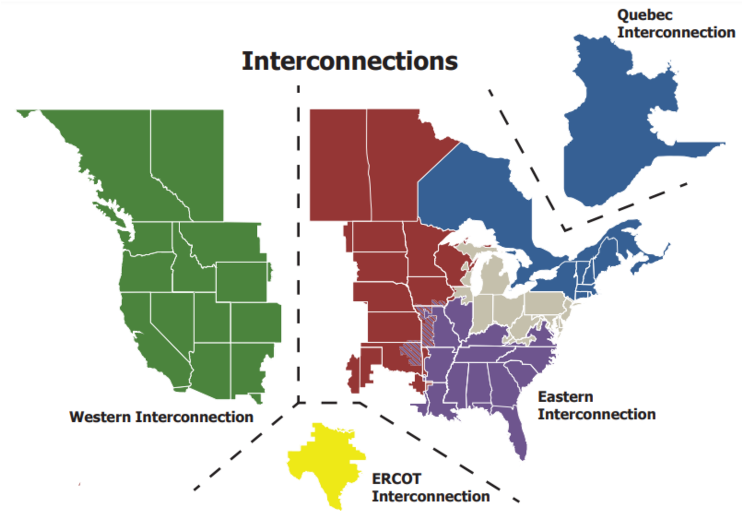

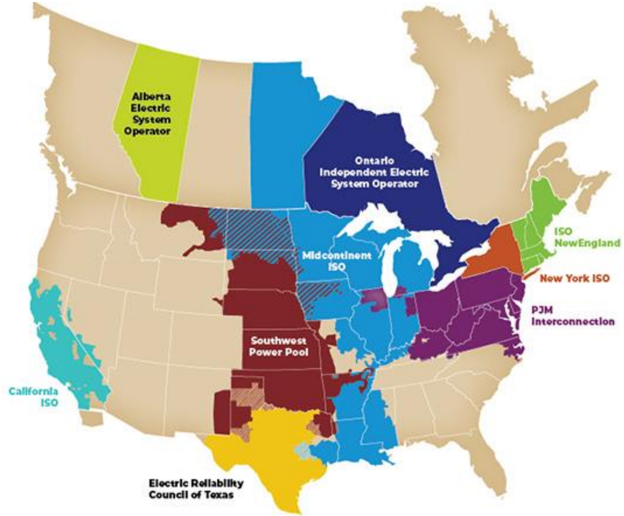

Key Regulatory Agencies:

- Federal Energy Regulatory Commission (FERC)
- North American Electric Reliability Corporation (NERC)
- States

- Regulated vertically integrated utilities
- Deregulated wholesale markets

Recent Data on Electric Power Needs

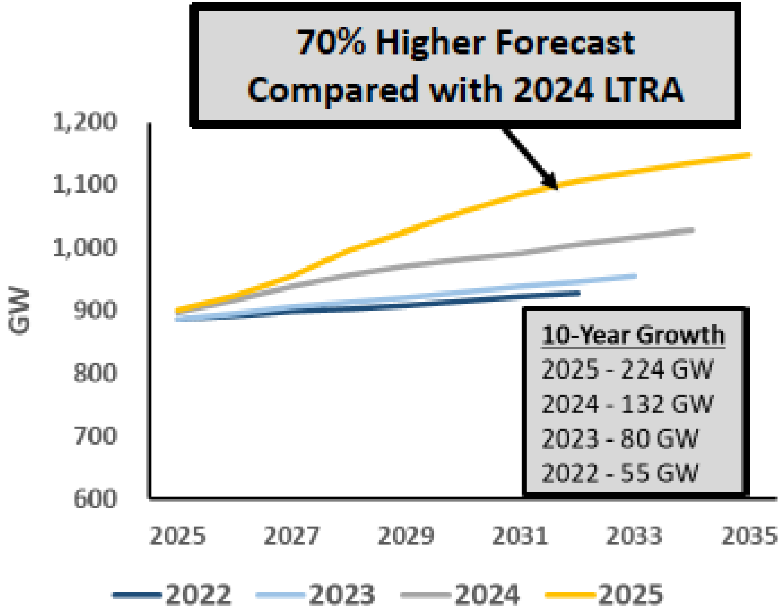

10-year BPS Summer Peak Demand Growth

Source: NERC 2026

Disclosed data center project pipeline reached 241GW

33% under active development

183 GW signed construction or electricity supply agreements

Texas and PJM account for 72% of large load capacity commitments

Source: Wood Mackenzie, 2026

PJM: Regional Transmission Organization

PJM

Key Statistics	
Member companies	1,116
Millions of people served	67+
Peak load in megawatts	**165,563**
Megawatts of generating capacity	**185,989**
Miles of transmission lines (BES)	88,333
Gigawatt hours of annual energy	828,161
Generation sources	1,673
Square miles of territory	368,906
States served	13 + DC

- 25.6% of generation in Eastern Interconnection
- 25.2% of load in Eastern Interconnection

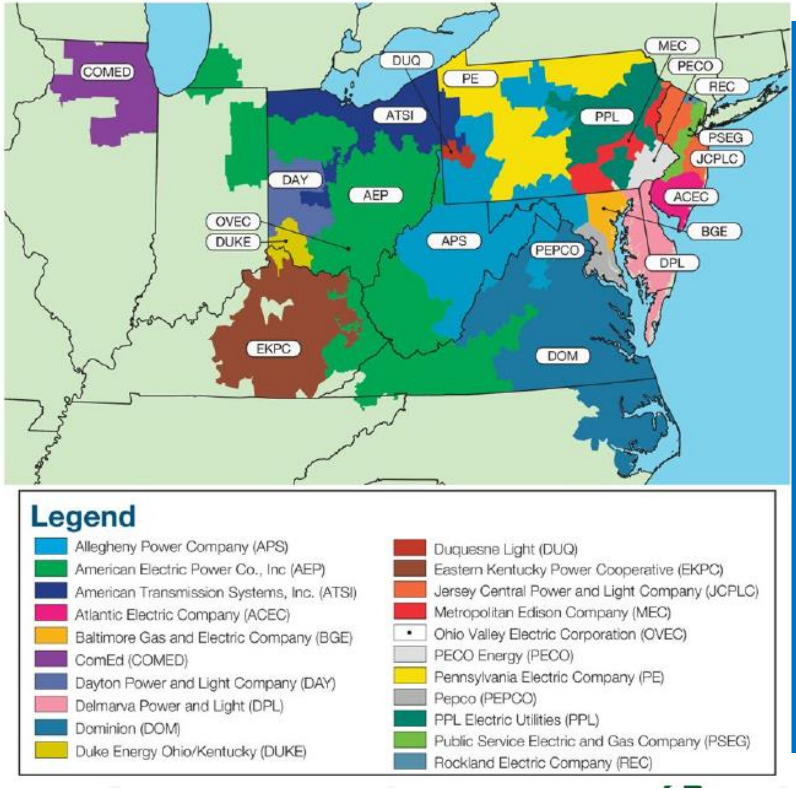

Key Responsibilities:

- Grid Reliability
- Transmission Planning
- System Operations
- Energy and Capacity Markets

Peak Load Growth in PJM

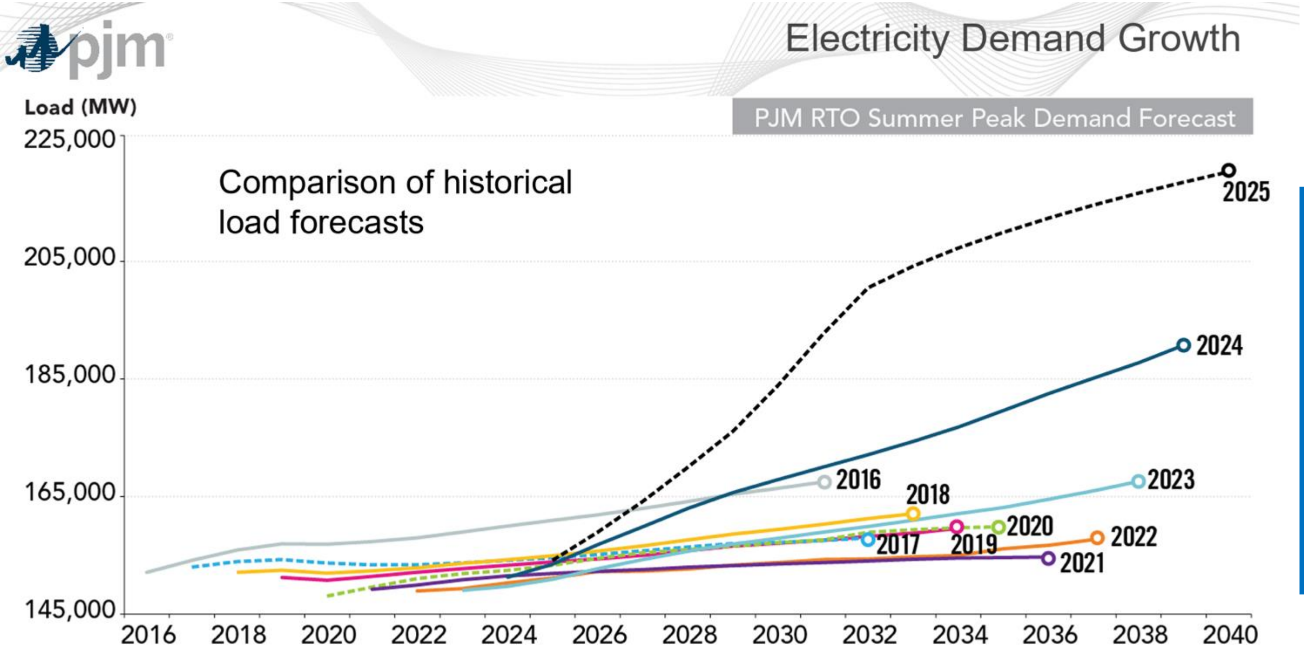

Wood Mackenzie:

- 55 GW of large load growth by 2030
- 100 GW by 2037

Source: PJM, February 2026

PJM Capacity Market Results

PJM capacity prices climb for a third straight year in the 2027/2028 auction

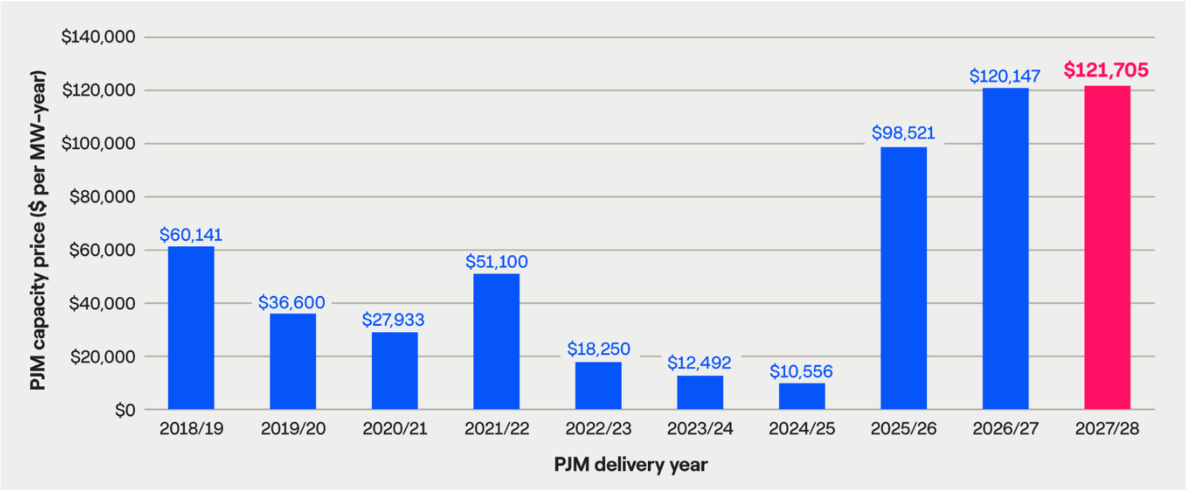

Source: PJM Interconnection

- Estimate of $194K/MW-year if there was no price cap.
- Auction fell short of target by 6.6 GW.
- Not much new supply came into the market
- Data center load growth is the primary reason for capacity market conditions

Source: Enel North America, 2025

PJM – Crucial Questions

Powering Reliability Through Market Design

Addressing Rising Demand and Constrained Supply, and Stimulating Investment To Support Durable Reliability

May 6, 2026

- "Should we preserve the concept of **resource adequacy as a common good** … and if so, who is responsible for making it financially durable?"
- "Or should we decide that **reliability**, in a period of scarcity, **must be explicitly rationed**?"
- "… should the primary long-term hedging instrument be the capacity product, or do we make a deliberate shift of **revenues supporting resource adequacy to the energy market** with long-term contracting?"

- PATH A: vast majority of load is required to be covered through long-term forward commitments
- PATH B: develop the operational and commercial framework to explicitly differentiate reliability
- PATH C: phased shift of revenue recovery from the capacity market to the Energy and Ancillary Services markets

Source: PJM, May 2026

Data Center Growth in Texas

ercot
PUBLIC

Large Load Interconnection Requests (as of March 26, 2026)

Actual and Projected Large Loads Growth 2022-2030

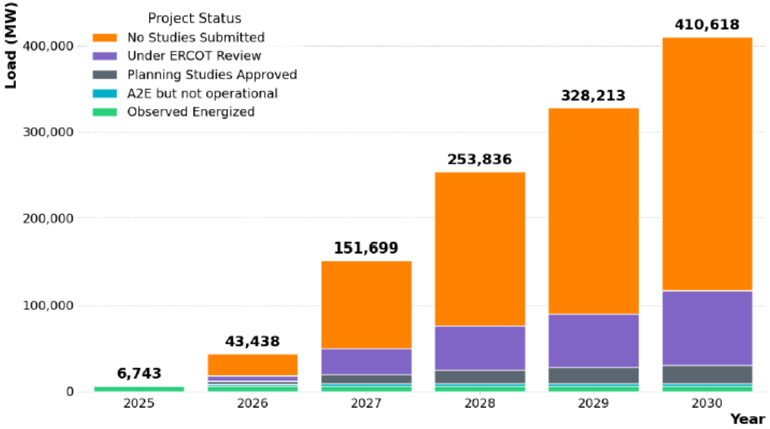

Project Status	2025	2026	2027	2028	2029	2030
No Studies Submitted	0	25,253	101,702	177,879	238,188	293,651
Under ERCOT Review	0	6,478	30,539	51,315	61,966	86,605
Planning Studies Approved	30	3,181	10,739	15,923	19,040	21,343
A2E but not operational	935	2,748	2,941	2,941	3,241	3,241
Observed Energized	5,778	5,778	5,778	5,778	5,778	5,778
Total (MW)	**6,743**	**43,438**	**151,699**	**253,836**	**328,213**	**410,618**

Large Loads by Project Type

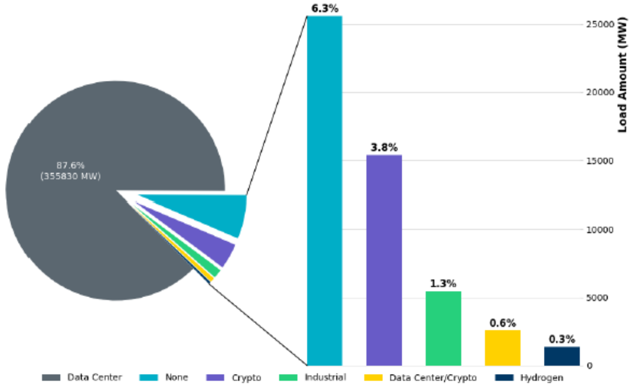

Key Takeaway: ERCOT is tracking approximately 410 GW of Large Loads seeking interconnection of which ~87% are data centers.

Source: [ERCOT, 2026]

State of Play in the Texas Grid

- 410 GW in the official queue as of April 2026 of which 87% are data centers; 99 projects > 1GW each
- 69 GW added in the last 12 months alone; ~198 GW filed since January 2026
- Texas Senate Bill 6 on interconnection standards for new loads of 75 MW or greater
- **Controllable Load Resource (CLR) pathway** and **Bring Your Own New Generation (BYONG) pathway**
- Key construct under development for co-location or bring your own generation: **Withdrawal-Limited Private Use Network (WLPUN)**

Major Challenges

- Large costs and long lead times for building and upgrading the transmission and distribution grids
- Long queues to connect generation and storage resources to the grid
- Long supply chain delays in key resources such as gas turbines, transformers, substations, reactive power components, …
- Maintaining grid reliability and resilience in the face of these expansions, inherent nature of data center loads, and extreme events such as storms, heat waves, …
- Rising cost of electricity leading to consumer worries and policy maker concerns

Timing Mismatch

FIGURE 1.4. Illustrative Time to Market for Various Grid Projects

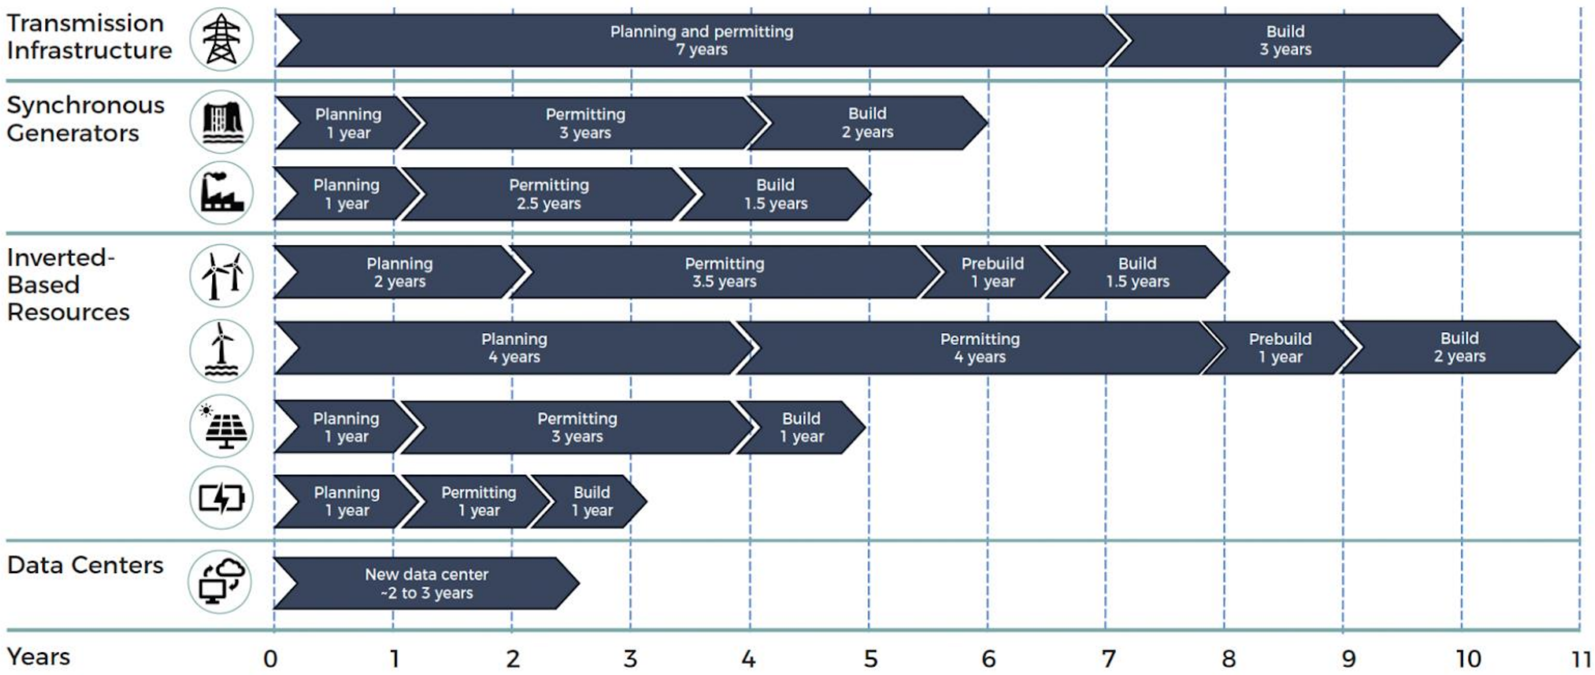

Timelines for grid infrastructure are not aligned with those for large load development, creating bottlenecks for grid supply of electricity. SOURCE: ADAPTED FROM S&P GLOBAL.

Source: Camus Energy

Many Pathways to the Future

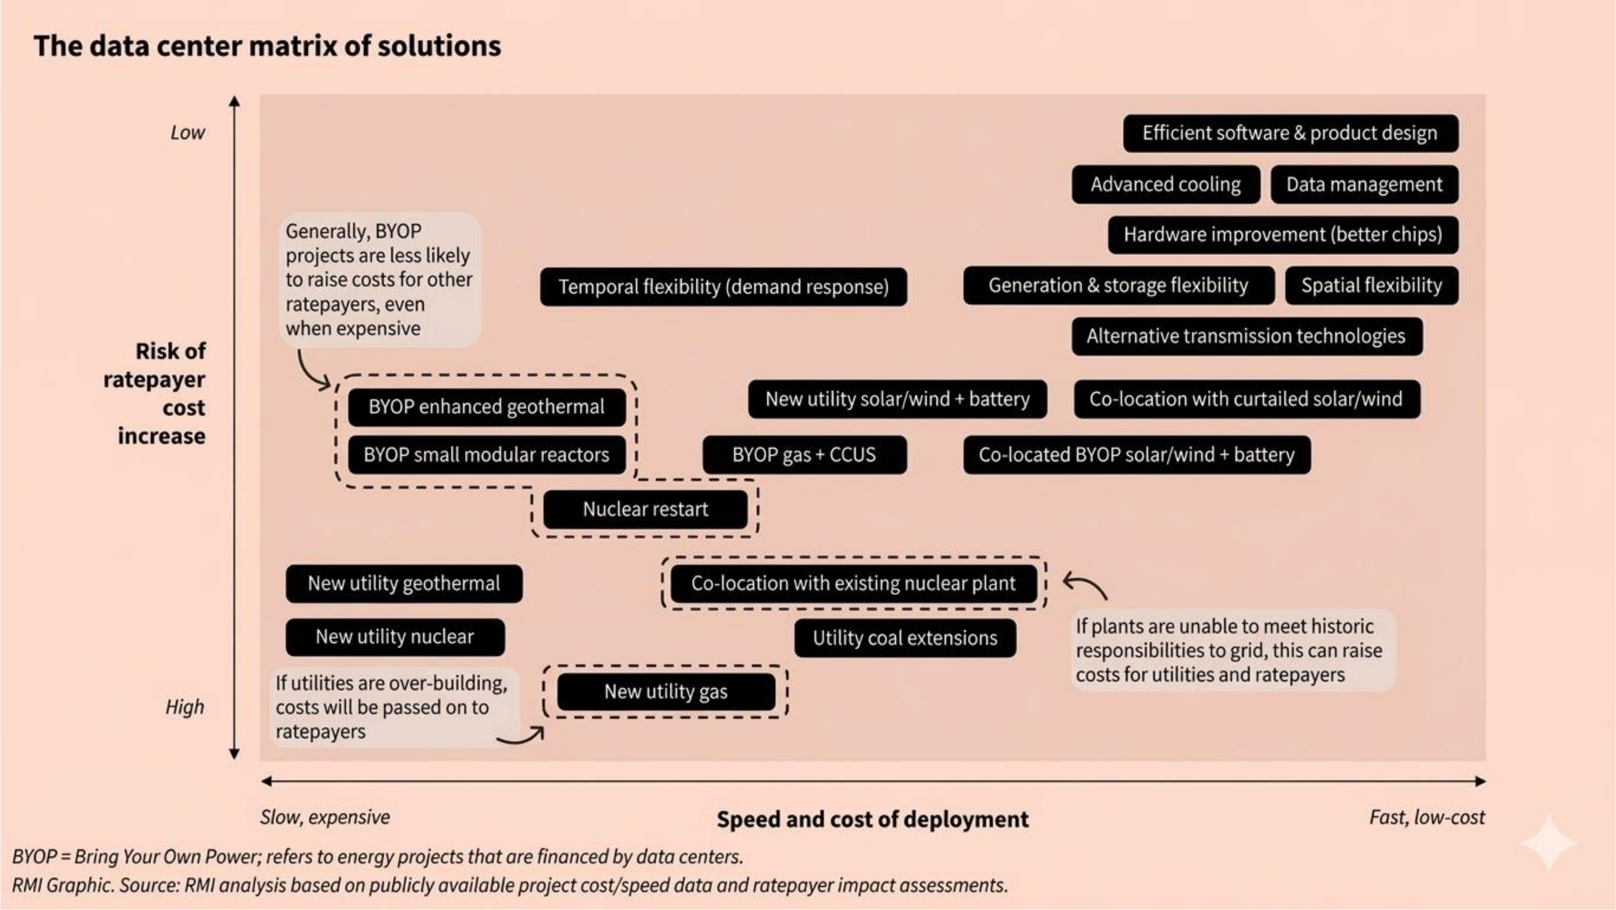

Source: Fast, Flexible Solutions for Data Centers, RMI, 2025

How can we get more out of the existing grid infrastructure and benefit ratepayers?

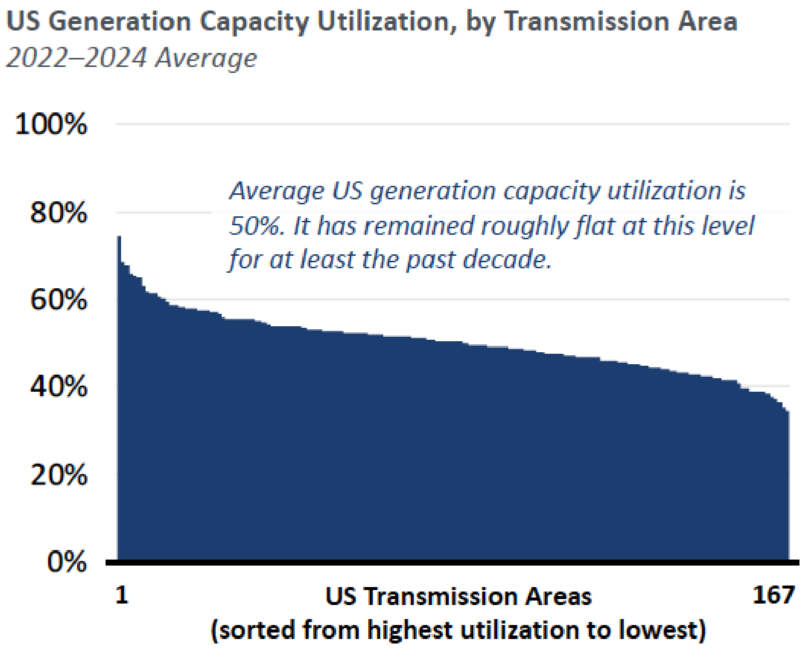

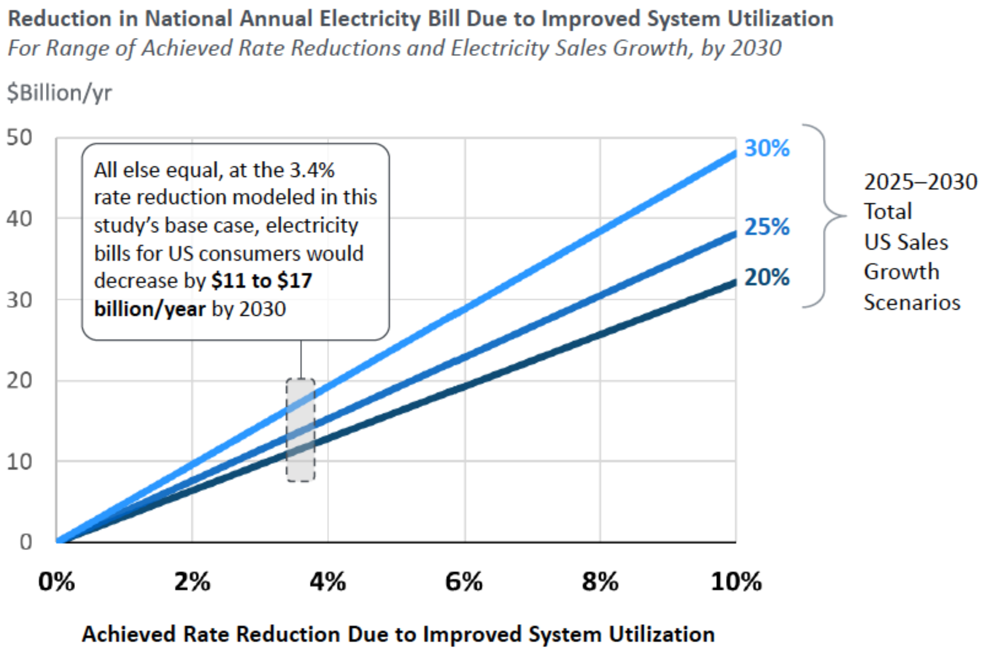

Source: [The Untapped Grid, Brattle Group, 2026]

Load Duration Curve

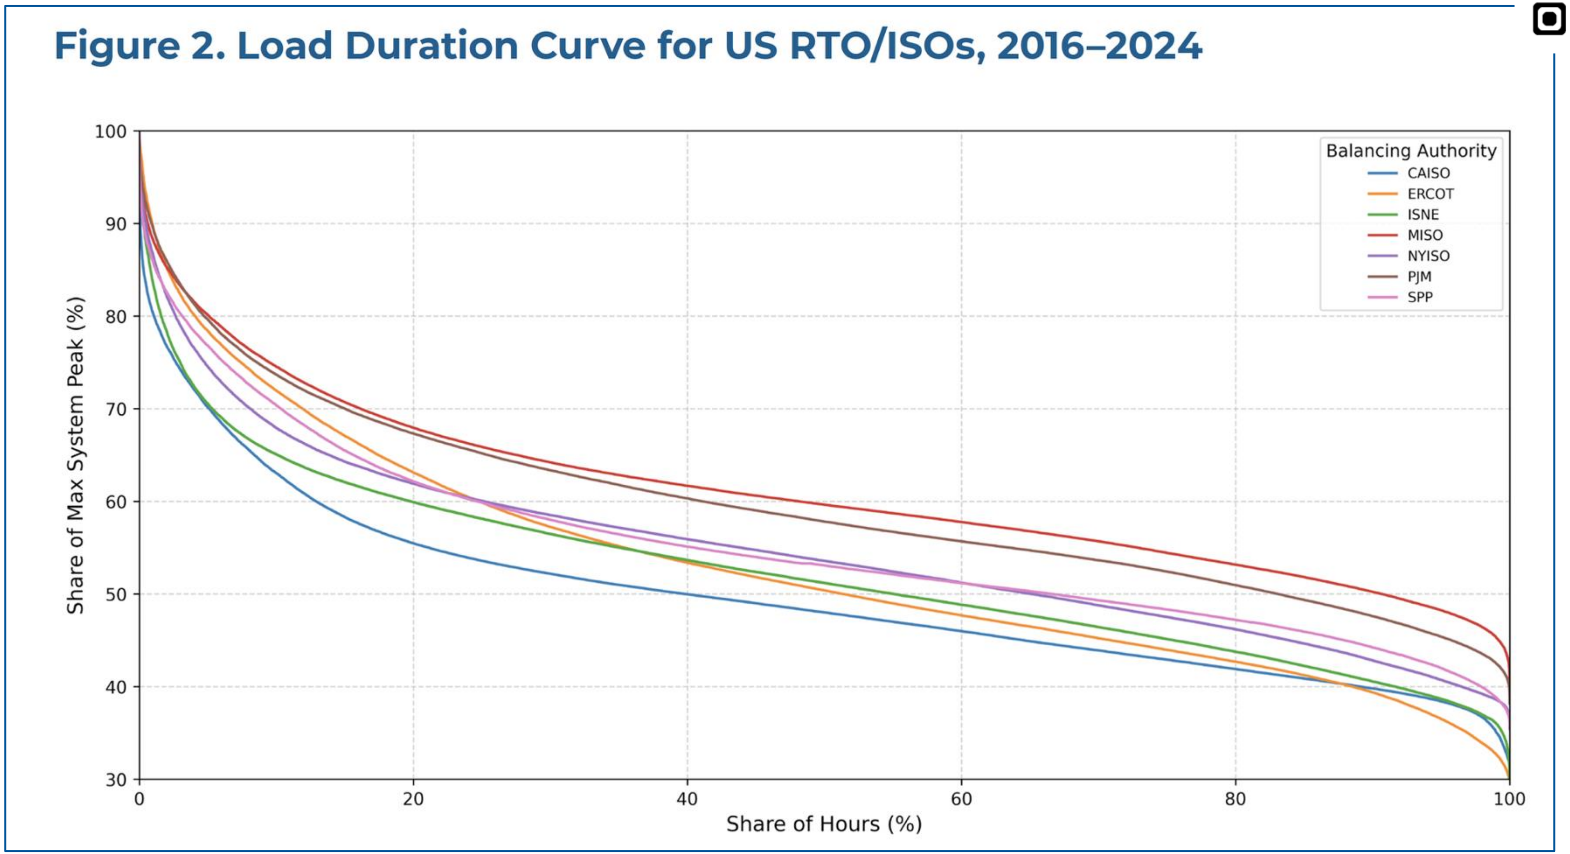

Source: Rethinking Load Growth, Norris et al, 2025

~100 GW of New Load Can be Served by Existing Grid

Figure 1. System Headroom Enabled by Load Curtailment of New Load by Balancing Authority, GW

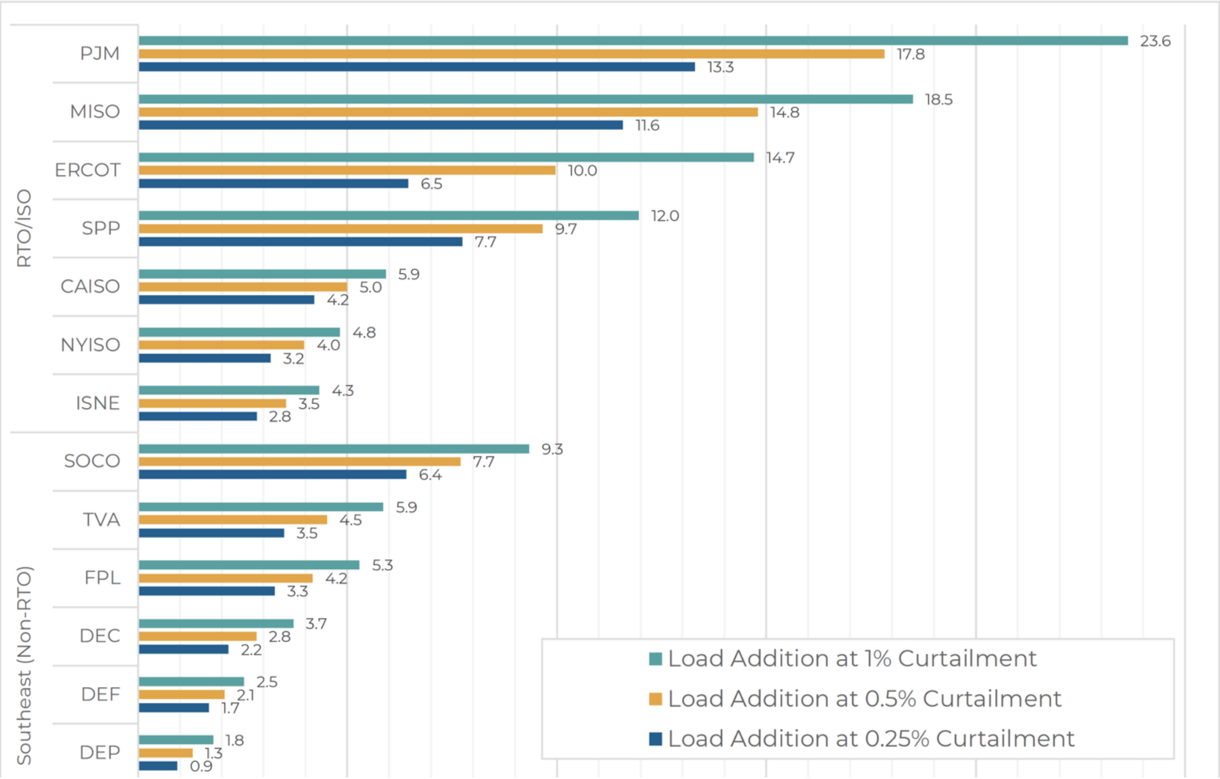

Source: Rethinking Load Growth, Norris et al, 2025

Potential of New Large Load Curtailment

- 76 GW of new load at an average annual load curtailment rate of 0.25%
- 98 GW of new load at an average annual load curtailment rate of 0.5%
- 126 GW of new load at an average annual load curtailment rate of 1.0%
- The average duration of load curtailment of 1.7 hours for 0.25%, 2.1 hours for 0.5%, and 2.5 hours for 1.0%
- Nearly 90% of hours during which load curtailment is required needs less than 50% curtailment

Source: Rethinking Load Growth, Norris et al, 2025

Power Network Constraints Will Reduce Potential Gains

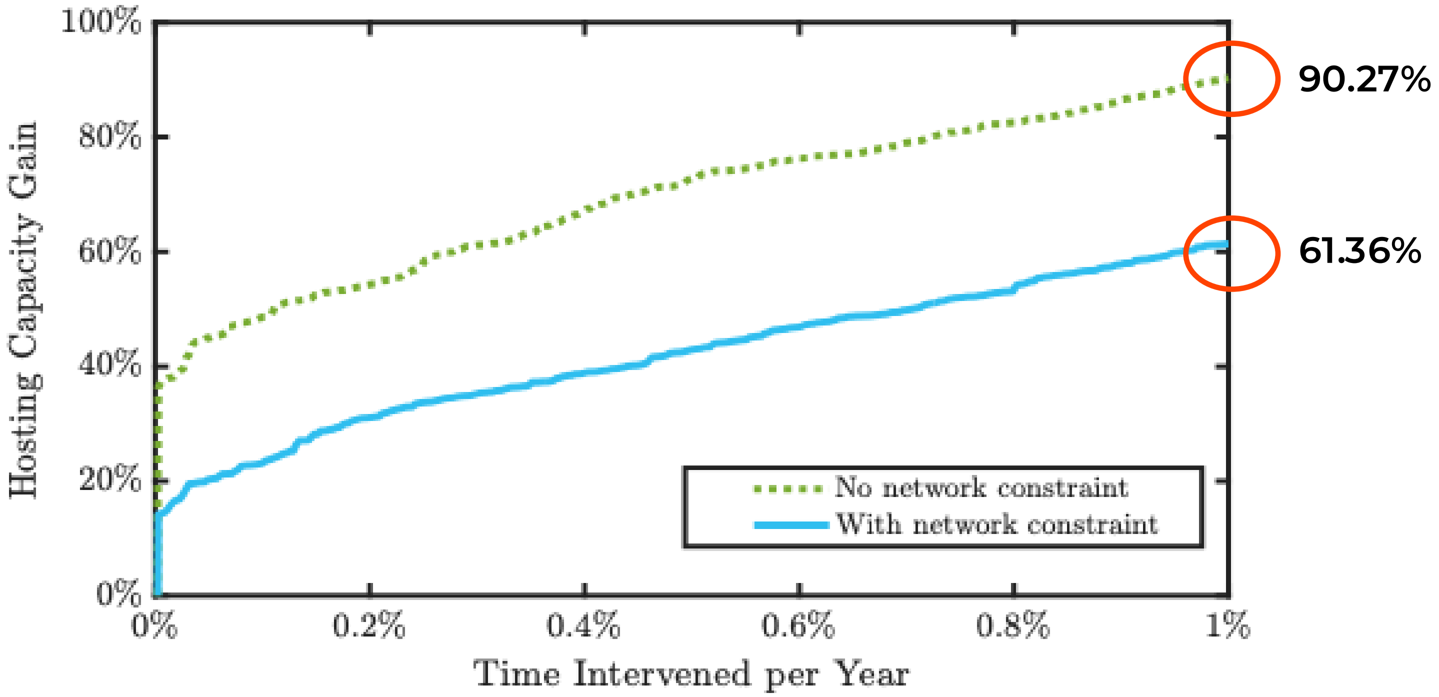

Source: Nan Gu, Ge Chen, Junjie Qin, " The Role of Flexible Connection in Accelerating Load Interconnection in Distribution Networks", accepted at HICSS 59'. https://arxiv.org/abs/2510.11476

Flexibility in Data Centers

Table 4. Flexibility Potential from Data Center Assets and Subsystems

Underlying Asset	Potential Capability	Response Time	Duration Sustained	Availability
On-site generation				
On-site battery storage				
On-site thermal storage				
Cooling systems				
Standby generation				
UPS				
Compute				

Source: Grid Flexibility Needs and Data Center Characteristics, EPRI, 2025

EPRI DCFlex Initiative

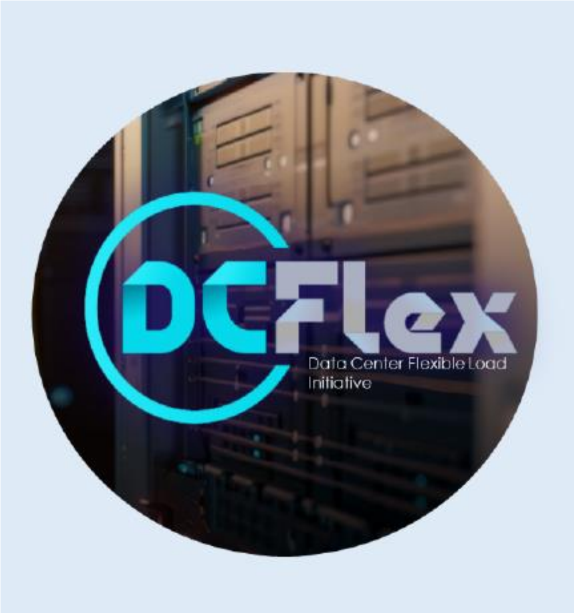

Source: https://dcflex.epri.com/

"DCFlex, will demonstrate how data centers can support and stabilize the electric grid while improving interconnection and efficiency"

ANALYSIS DATA + CLOUD GRID EDGE

Nvidia and Oracle tapped this startup to flex a Phoenix data center

The load flexibility pilot used Emerald AI's platform to reduce power use by 25% during peak demand — and maintain AI compute quality.

MAEVE ALLSUP | JULY 1, 2025

Source: Latitude Media, July 1, 2025

GPU Power Capping vs ML Performance

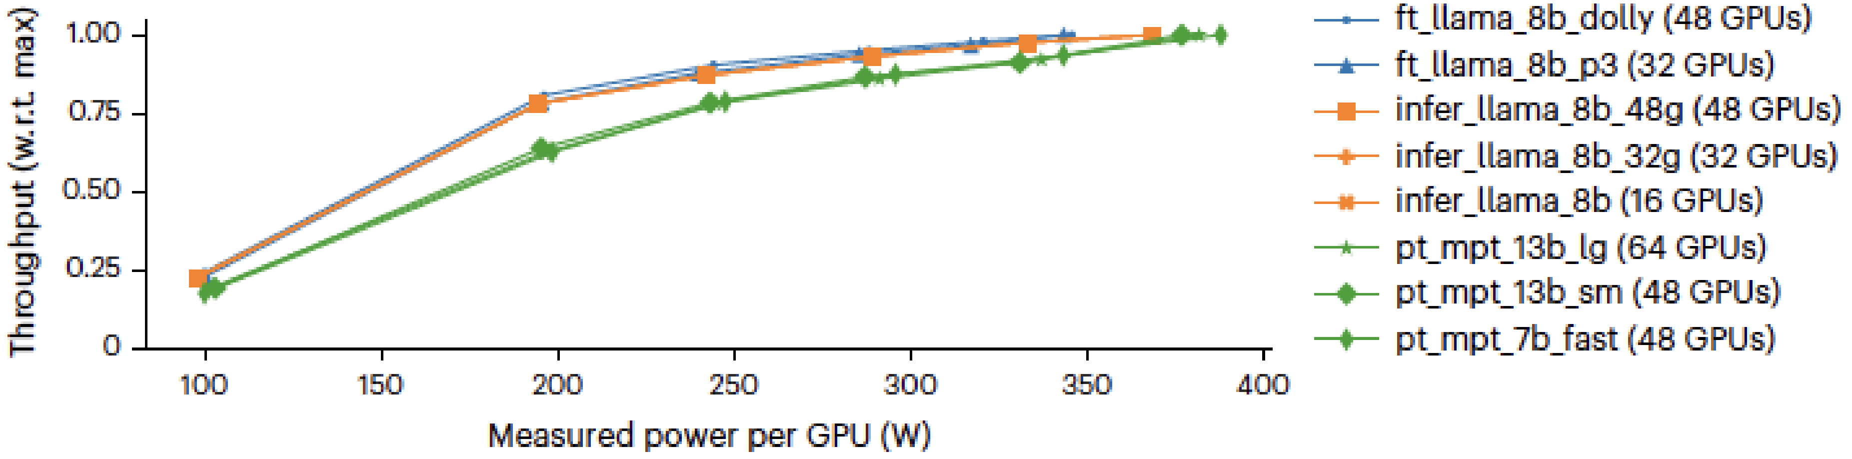

Source: [Colangelo et al, Nature Energy, 2026]

Data Center Flexibility Demonstration

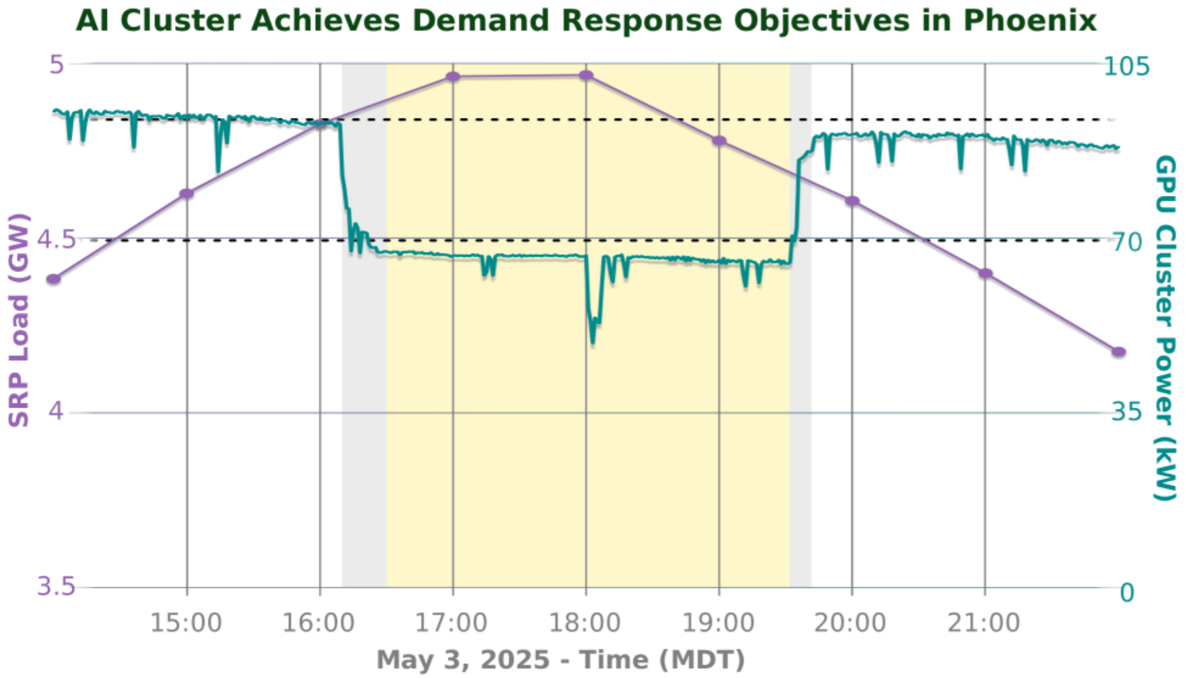

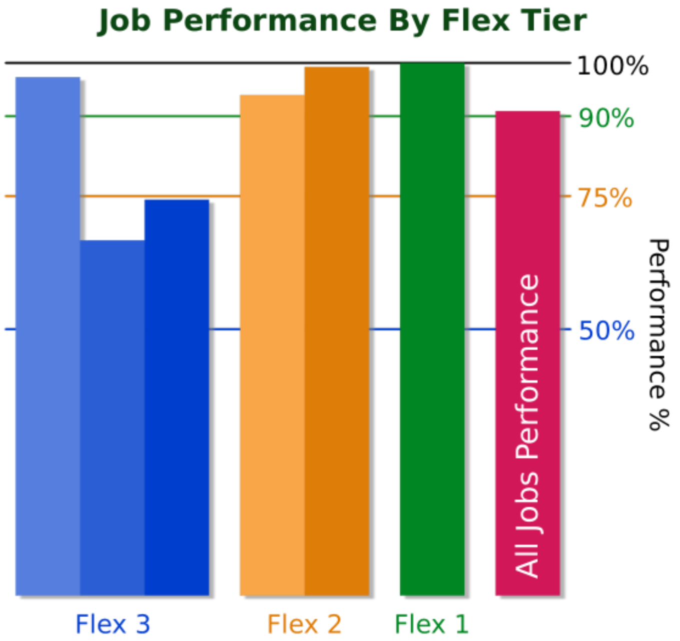

Source: [Colangelo et al, Nature Energy, 2026]

NVIDIA DSX Flex AI Data Center Architecture

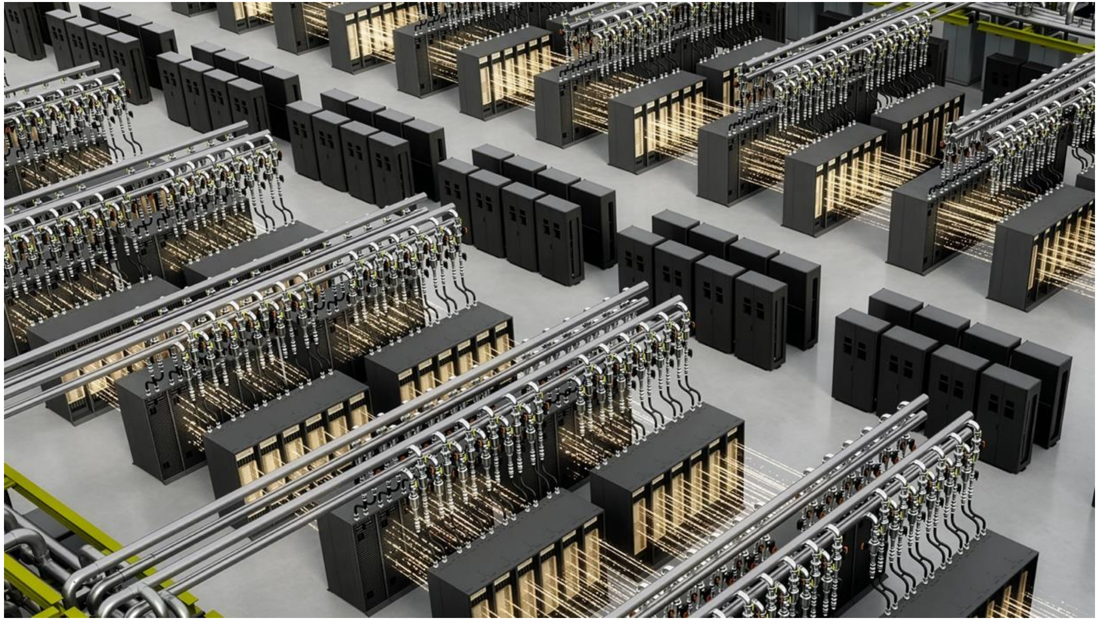

96MW AI Data Center in Virginia that aims to show:

- Peak load reduction: 20%-30% for multiple hours
- Long-duration curtailments – 10 hours
- Multiple events on one days
- Rapid response: 10-minute
- Carbon-aware dispatch.
- Tariff/market signals

Source: NVIDIA, March 2026

CAMUS encoord 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

A 500-MW data center using flexible grid connection + BYOC

- Grid power available for more than 99% of all hours
- On-site resources (e.g. batteries, generators, load flex) 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 data centers can connect 3-5 years faster

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.

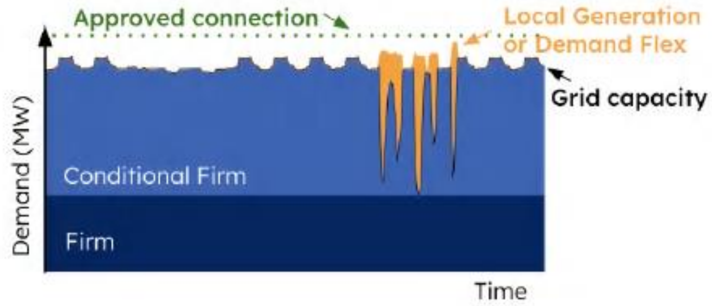

A bring-your-own capacity (BYOC) tariff addresses generation capacity constraints more quickly by procuring or building new capacity directly

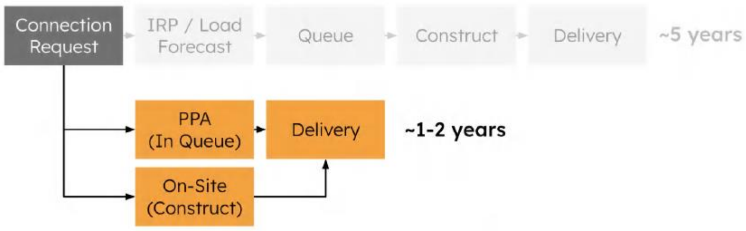

Source: Brancucci et al, December 2025, Research funded by Google

Compute Heat Rate and Flexibility

Compute heat rate = the maximum electricity price at which a given AI workload remains profitable

"At current PJM energy price cap levels (roughly $3,700/MWh), the economics of curtailing compute operations are generally unattractive for most AI workloads." PJM, May 2026.

Workload Type	R(w) $/MWh	C(non-elec) $/MWh	Raw CHR $/MWh	Effective CHR $/MWh	vs. Gas HR	Provenance
Frontier Inference (Opus/GPT-5)	$74,000	$4,250	$53,650	$53,650	~1,070x	Empirical
Mid-Tier Inference (Sonnet/GPT-4.1)	$14,800	$4,250	$8,120	$8,120	~162x	Empirical
Enterprise Agentic AI	$15,000	$4,500	$8,080	$8,080	~162x	Uncertain
Enterprise Contracted	$5,900	$4,250	$1,270	$1,270	~25x	Modeled
Commodity Inference (mini)	$1,850	$3,800*	-$1,500†	~$800†	~16x	Modeled
Frontier Model Training	$2,000	$5,200*	-$2,460†	~$500†	~10x	Modeled
Blended Average (2026)	**$12,500**	**$4,350**	**$6,270**	**$6,350**	**~127x**	**Modeled**

Source: The Compute Heat Rate, H. Royal, SSRN, March 2026

Cost Allocation Issues

- How to ensure additional costs of providing electric power and energy to AI data centers are fairly allocated?
- Complex processes with a mixture of federal and state regulations for allocating various costs and setting electricity rates for various classes of customers
- Mixed evidence so far on whether AI data center costs are showing up in residential customer bills
- Possibility of stranded costs of grid expansion falling on residential consumers
- Cost allocation for co-located generation with the AI data center is unclear and subject to dispute

Tariff and Contract Design for AI Data Centers

- AI Data Center developers negotiate with utilities for connecting to the grid and getting power
- Typical tariff: demand charge + energy charge
- Demand charge minimum to the data center: transmission charge + distribution charge + generation charge
- Energy charge related to the cost of energy delivered
- Multi-year binding contract

Pennsylvania regulators move to protect utility customers from data center costs

Posted: May 15, 2026 in Electricity Pricing, Energy Efficiency

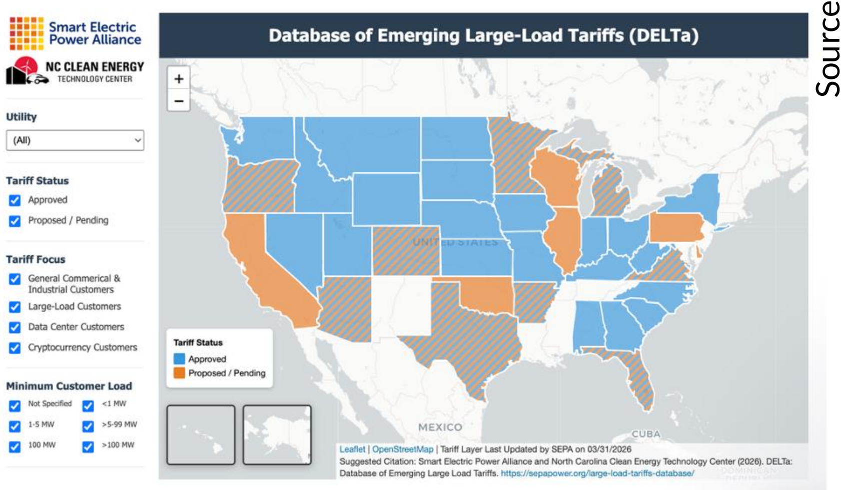

Source: Smart Electric Power Alliance

Source: Environmental Defense Fund, 2026

Billed Demand Should be Ratcheted

- Data centers request grid capacity before future demand is known
- Utilities make substantial investments to serve the requested capacity
- Realized demand may be much lower than requested capacity
- Billed demand = max(realized demand, **minimum billed demand**)
- Create a payment floor for data centers, limit cost shifting to ratepayers
- Minimum billed demand is subject to approval by the Public Utility Commission

Source: Zhijie Lai, Sen Li, Junjie Qin, Kameshwar Poolla, Pramod Khargonekar "Robust Tariff Design for Data Center Interconnection", Working Paper.

Higher Minimum Billed Demand Better Protects Ratepayers

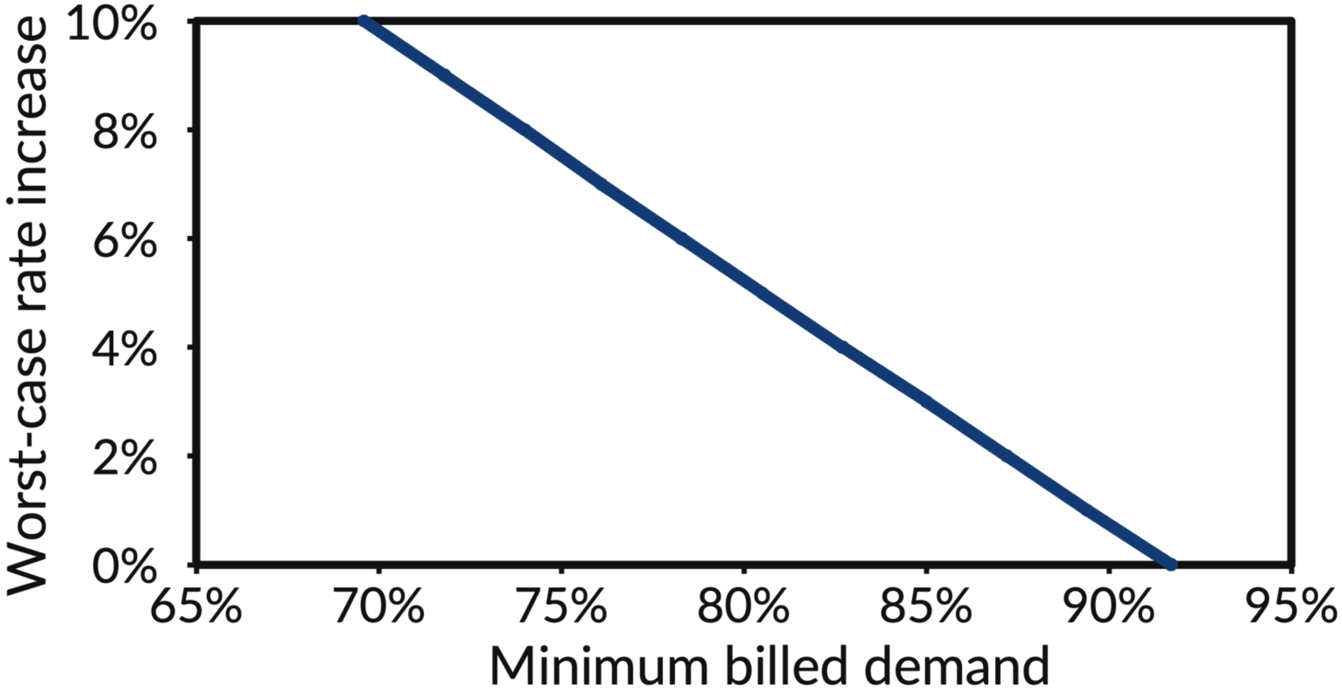

Source: Zhijie Lai, Sen Li, Junjie Qin, Kameshwar Poolla, Pramod Khargonekar “Robust Tariff Design for Data Center Interconnection”, Working Paper.

Implications for Data Center Tariff Design

- Higher minimum billed demand makes data centers financially responsible for reserved capacity
- Risk-free protection requires minimum billed demand above 90%.
- The 80–85% levels being proposed or adopted can reduce risk, but do not eliminate it
- This range may offer a more practical balance between ratepayer protection and data center participation
- Better demand forecasts can further reduce ratepayer risk

Source: Zhijie Lai, Sen Li, Junjie Qin, Kameshwar Poolla, Pramod Khargonekar "Robust Tariff Design for Data Center Interconnection", Working Paper.

Conclusions

- US electric power grid going through a generational transformation. AI Data Centers are a significant driver of this transformation
- Regulatory policy at the state level will play an important role
- Grid-data center planning, grid reliability, grid-data center operations and control likely to remain important topics for the foreseeable future
- Done right, transformed electric grid could serve the public for decades to come

Rising Opposition to Data Centers

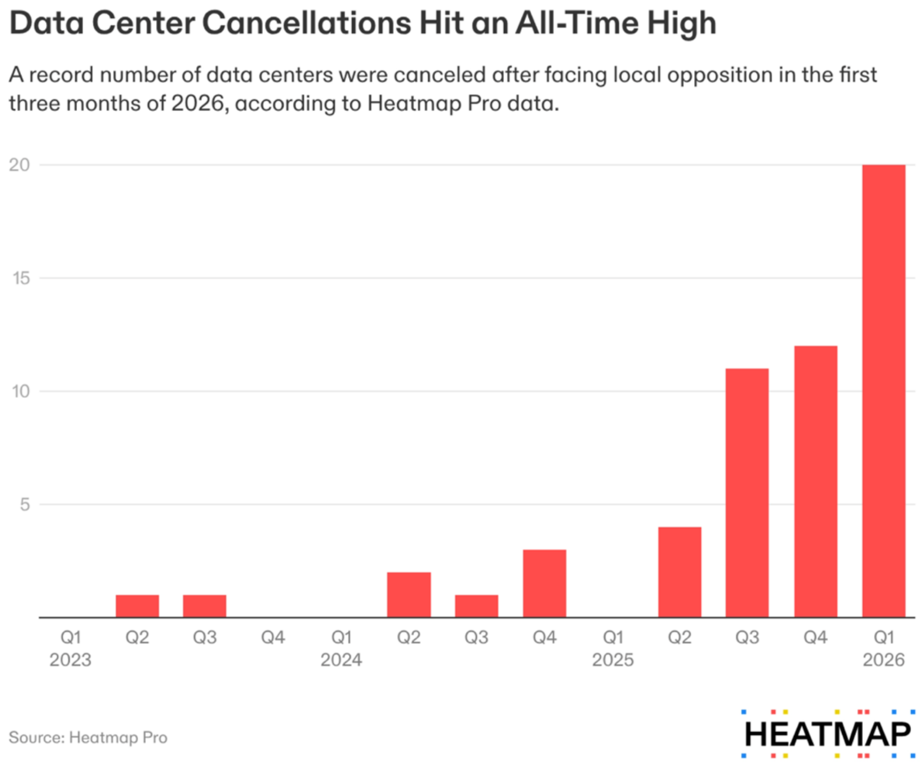

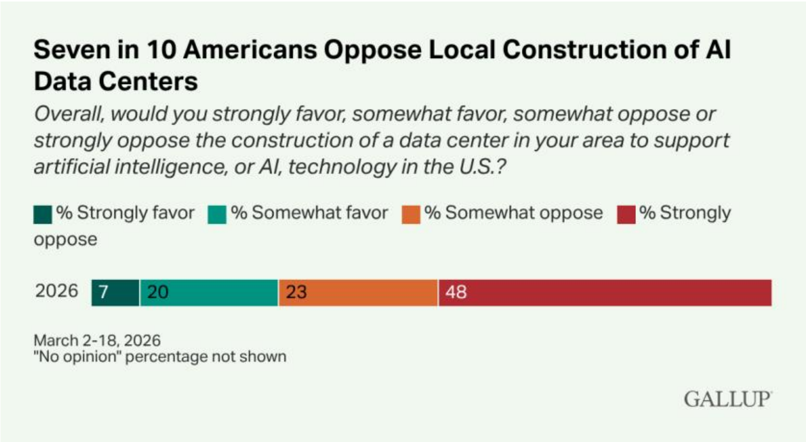

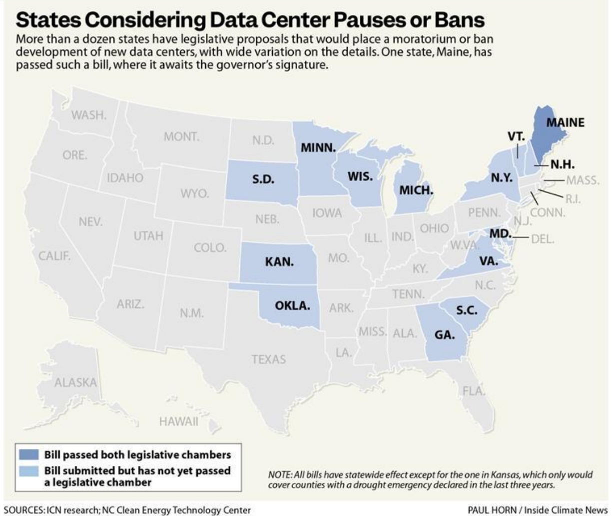

https://tinyurl.com/2d7v6a26

https://tinyurl.com/4zs4hj97

https://shorturl.at/YsAuC

Overarching Perspective

- Electric power systems are a critical infrastructure for all aspects of our society
- AI data center buildout puts great pressure on energy and water infrastructure
- Large sectors of the energy system can be electrified: transportation, residential heating, manufacturing and industry, …
- Electrification is also the furthest ahead in terms of decarbonization through solar, wind, batteries, …
- Electric energy system growth offers a major opportunity for positive change

To Build or Not to Build?

"Since we can't solve the grid crisis by blocking the data center buildout, we should focus on harnessing that buildout to improve the power grid overall. Some of the economy's richest and most innovative companies are willing to invest in new power plants and improve the grid's infrastructure. That is a generational opportunity."

Robinson Meyer, NYTimes, April 27, 2026

"The rapid, largely unregulated rise of data centers to fuel the AI and crypto frenzy is disrupting communities across the country and threatening Americans' economic, environmental, climate and water security. We urge you to join our call for a national moratorium on new data centers"

Food and Water Watch and 200+ organizations, December 8, 2025

Comments

Ideas

Questions?

pramod.Khargonekar@uci.edu
http://faculty.sites.uci.edu/khargonekar/