Data Center Flexibility

July 2026

Teneo

U.S. data center demand growth is driving a step-change in power demand

U.S. data center demand growth
(TWh)

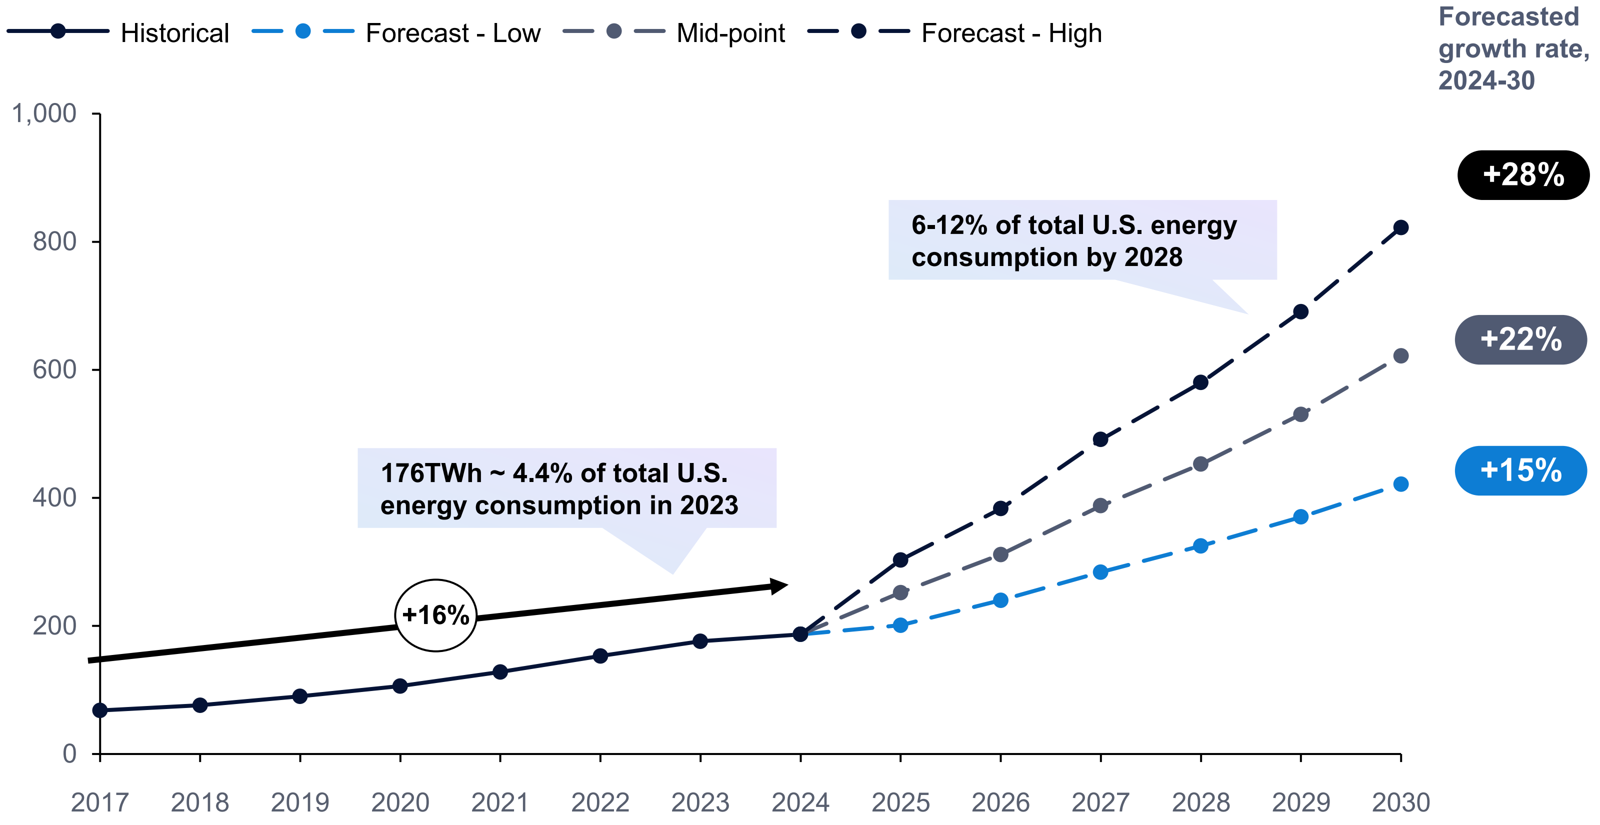

Observations

External estimates vary, but all point to very rapid growth:

Berkeley Lab estimates U.S. data center electricity demand will rise from 176 TWh in 2023 to 325–580 TWh by 2028, equivalent to roughly 74–132 GW of power demand

EPRI's updated 2030 scenarios imply 380–790 TWh and 56–132 GW of nominal IT capacity by 2030

Source(s): U.S. department of energy www.energy.gov, Lawrence Berkeley National Laboratory LBNL "2024 report on data center use", desk research

…and growing data center hubs are stressing regional infrastructure, contributing to interconnection (IX) delays

Top 10 US data center hubs, Q4 2025 (GW)

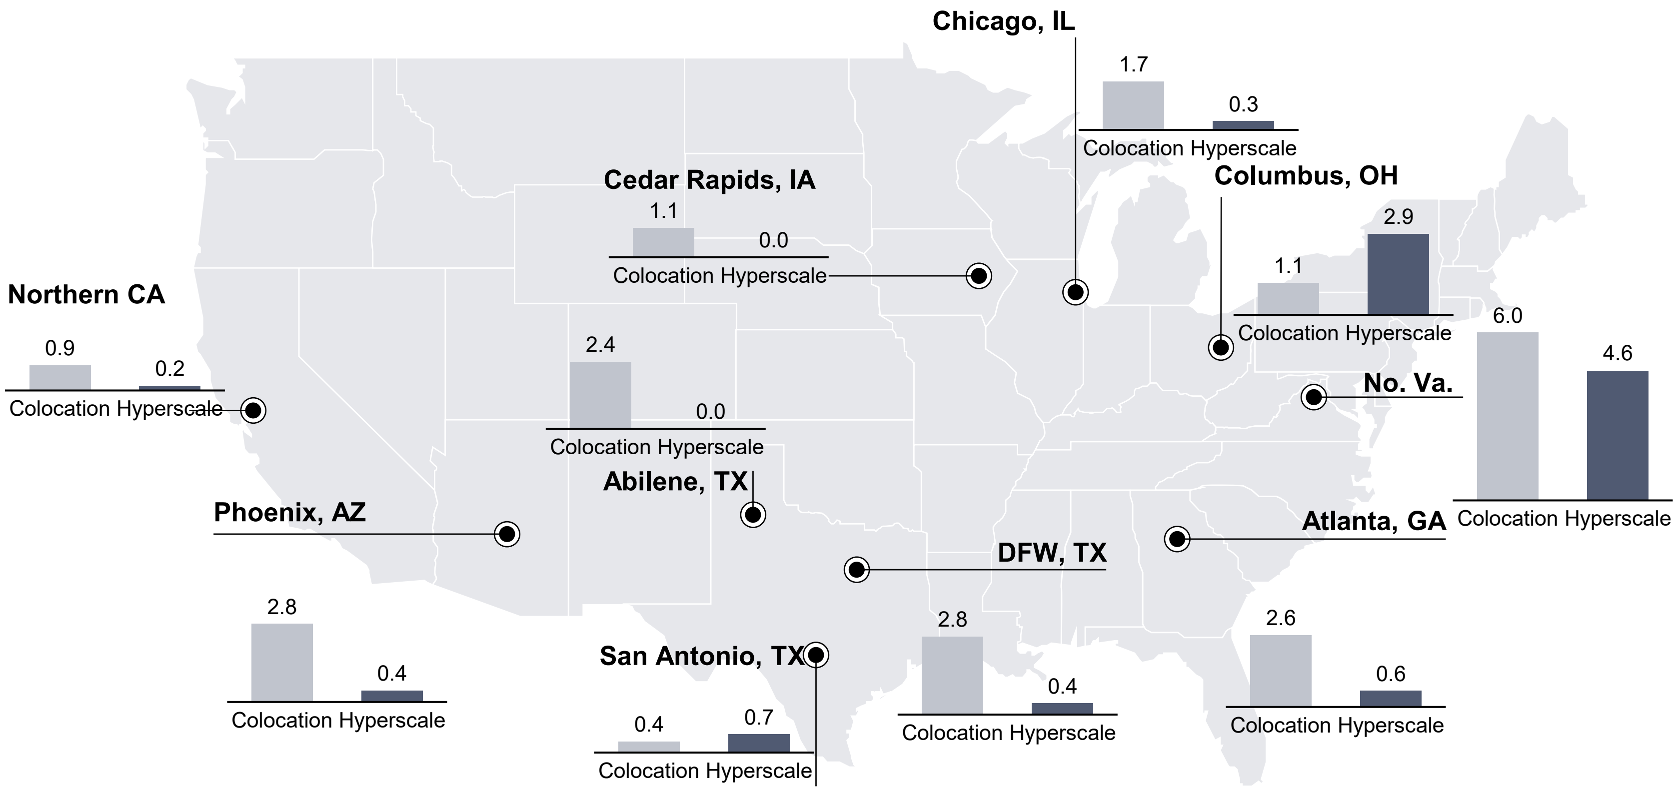

Transmission operator impacts

pjm	Emerging collocated generation, with new regulatory frameworks emerging to expedite IX
MISO	Constrained generation, with 50 GW of signed IX agreements remaining unbuilt
ercot	Emerging strategy of partial capacity commitments while awaiting transmission upgrades over the next 5-10 years
California ISO	Load-serving entities vote on projects to receive IX rights, requiring projects to execute PPAs in advance
SPP	High Impact Large Load process approved by FERC, reducing timelines to 90 days
New York ISO	Cluster study reforms see timelines reduced to c. 21 months from 5+ years
ISO new england	No large load track, with sufficient capacity anticipated through 2033

Source(s): Data Center Hawk, S&P, impactECI

Data center growth is outpacing the grid's ability to deliver new capacity, triggering costly system upgrades

ERCOT large load interconnection queue, by stage (2022-30P) *(GW)*

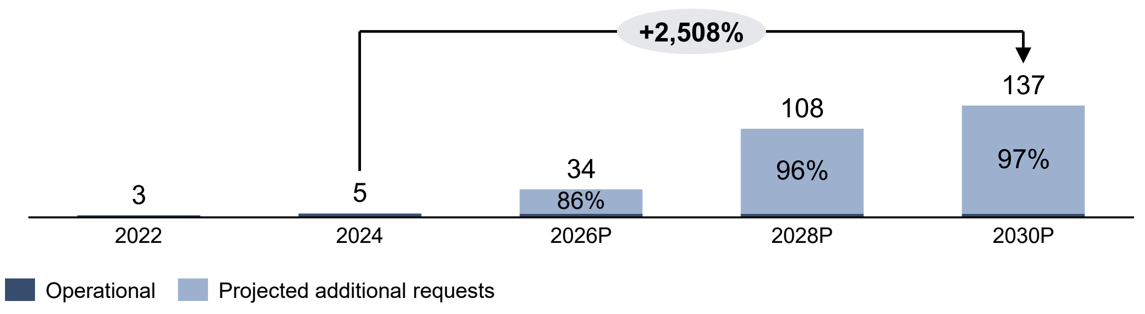

Estimated data center interconnection timeline in Texas *(# months)*

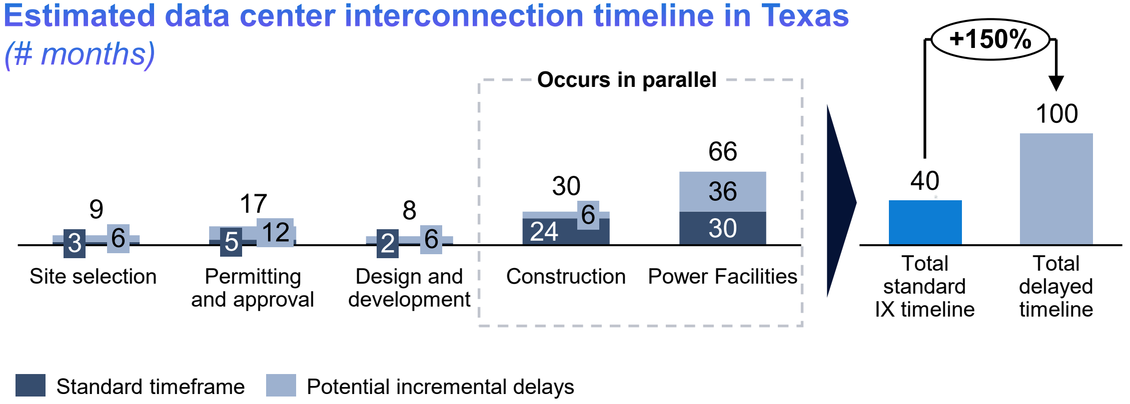

Data center related transmission costs approved in 2024, *[$M]*

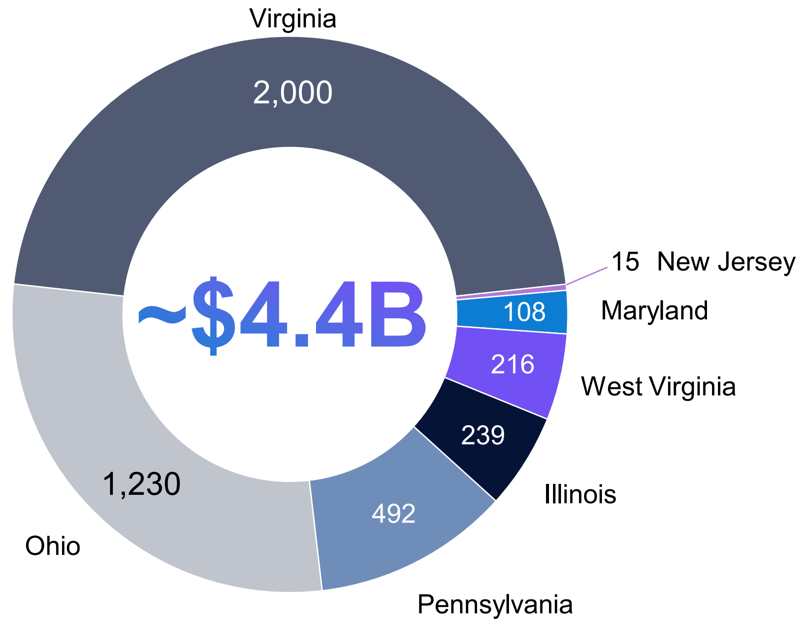

- In key markets, developers face multi-year interconnection timelines even before major upstream upgrades are complete
- As load clusters grow, the cost of network upgrades is increasingly being pushed into broader system costs and ratepayer debates. In 2024, customers across seven PJM states were billed $4.4 billion for data center-related transmission upgrades

Teneo **Source(s):** U.S. department of energy www.energy.gov, Lawrence Berkeley National Laboratory LBNL "2024 report on data center use", desk research

$64 billion of data center projects in the U.S. have been blocked or delayed in the last two years alone because of public opposition, negatively implicating economic growth and U.S. AI leadership.

$18B

worth of data center projects were blocked between May 2024 and March 2025.

$46B

worth of projects were delayed.

142

activist groups pushing back on data center development across 24 states in the U.S.

31,000

signatures on petitions against data center development since 2022.

Public awareness is dominated by perceived costs, and more informed audiences are more negative, signaling a clear failure by developers to communicate benefits effectively

Public sentiment skews negative, with many remaining uninformed

Views of data centers' impact in various areas, % of U.S. adults[1]

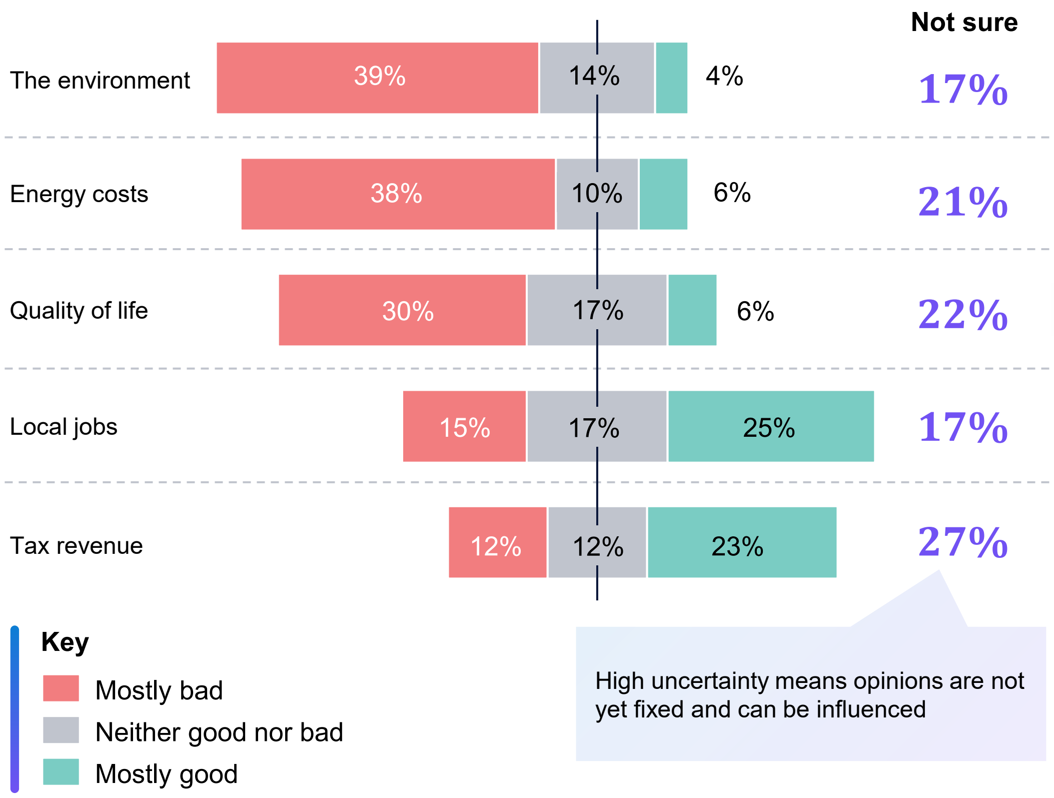

As awareness increases, perceptions become more negative across key impact areas

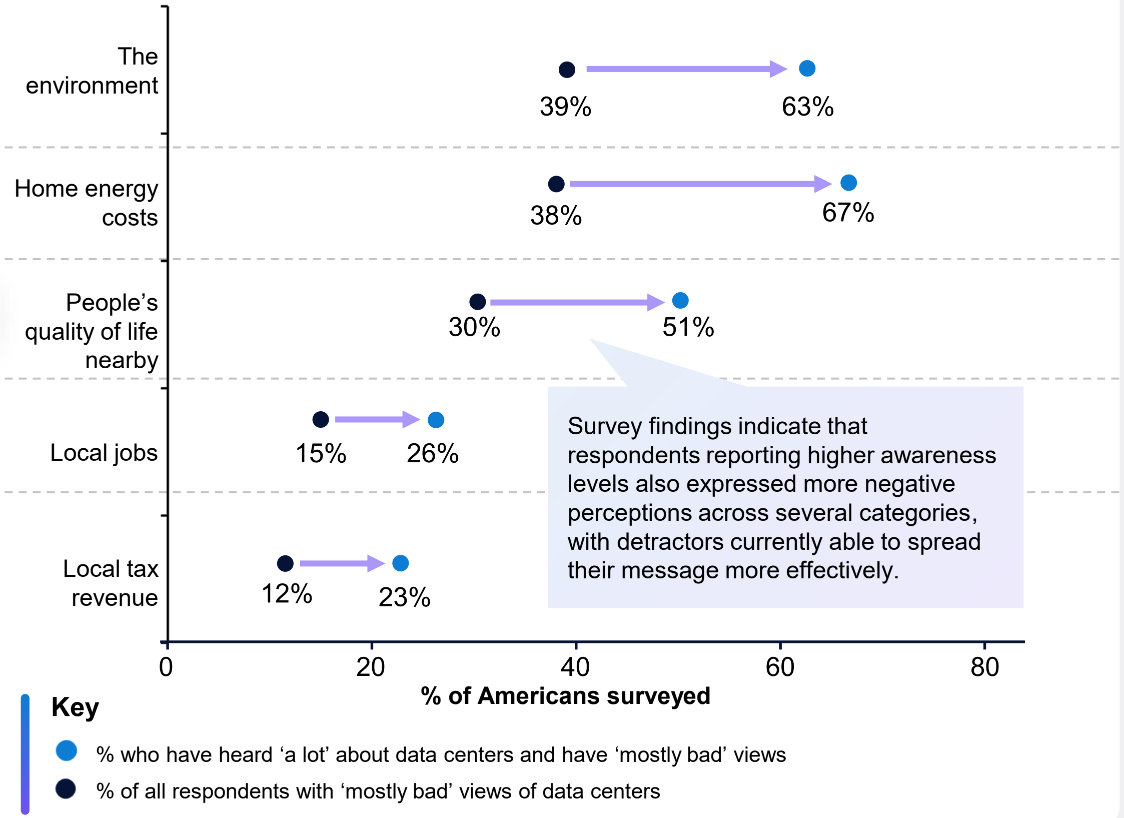

Teneo **Source(s):** Pew Research Center; Teneo research & analysis | **Note(s):** 1. Figures exclude those who did not answer and those who have not heard about data centers

Integrating flexibility into facility design can avoid delays and cancellations arising from public opposition during the early stages of project development

Percentage of projects which are delayed or canceled at each phase of the development process

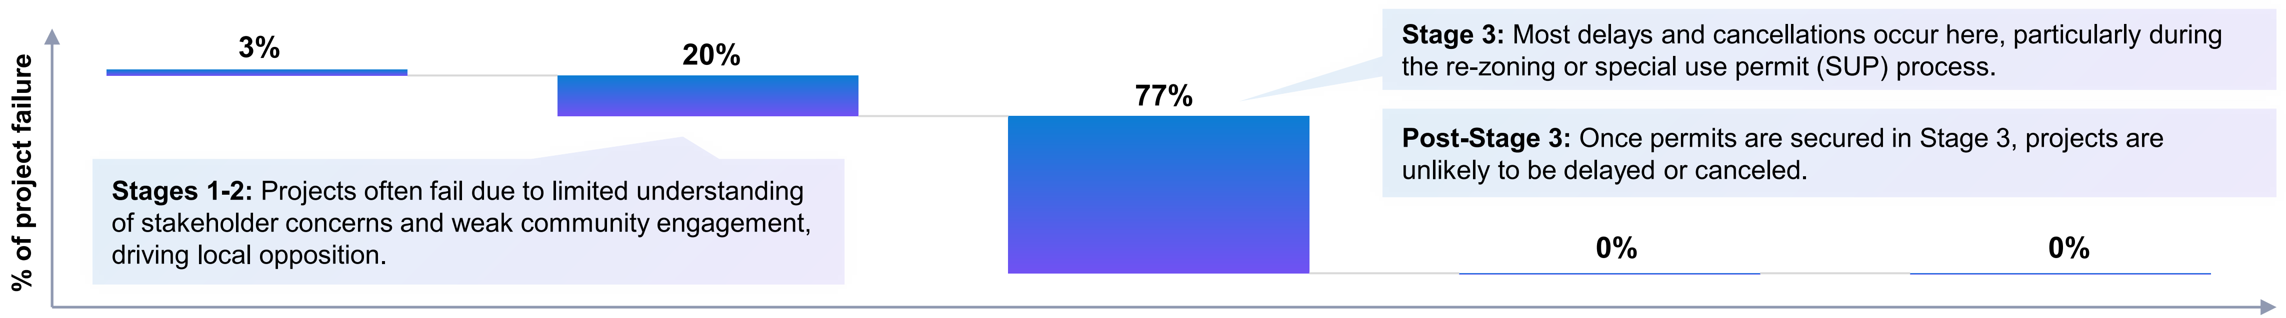

01 Site identification & land acquisition

02 Feasibility & due diligence

03 Permitting & regulatory approval

04 Design & construction

05 Commissioning & startup

Key stakeholders

01 – Stakeholders needed for early legitimacy and land use alignment:

- Local government officials
- Landowners

02 – Stakeholders needed for credibility on power, resource use and impacts:

- Utilities & infrastructure providers
- Community members and organizations

03 – Stakeholders needed for approval and public acceptance:

- Regulators
- Community members and organizations
- State and local government officials

04 – Stakeholders needed for delivery credibility during construction:

- Local workforce
- Businesses
- Residents / near neighbors

05 – Stakeholders needed for final validation before operations:

- Utilities
- Regulators
- Residents / near neighbors

Teneo

Source(s): Data Center Watch, Teneo research and analysis on ~30 data center projects that were delayed or canceled.

Current "build more infrastructure" responses are too slow, too costly and too risky for ratepayers

Key takeaways

- The traditional response to large-load growth - new transmission, utility-owned gas or new utility-owned firm generation - sits in the slow, expensive and high-ratepayer-risk corner of the solution set.
- Data center flexibility is placed in the opposite corner: fast to deploy, relatively low cost and low impact on customer bills.

Why flexibility matters?

- Operators gain a faster path to power.
- Utilities gain a lower-regret alternative to emergency overbuild.
- Investors gain exposure to a model that can monetize flexibility without waiting for decade-long infrastructure cycles

Risk of ratepayer cost increase

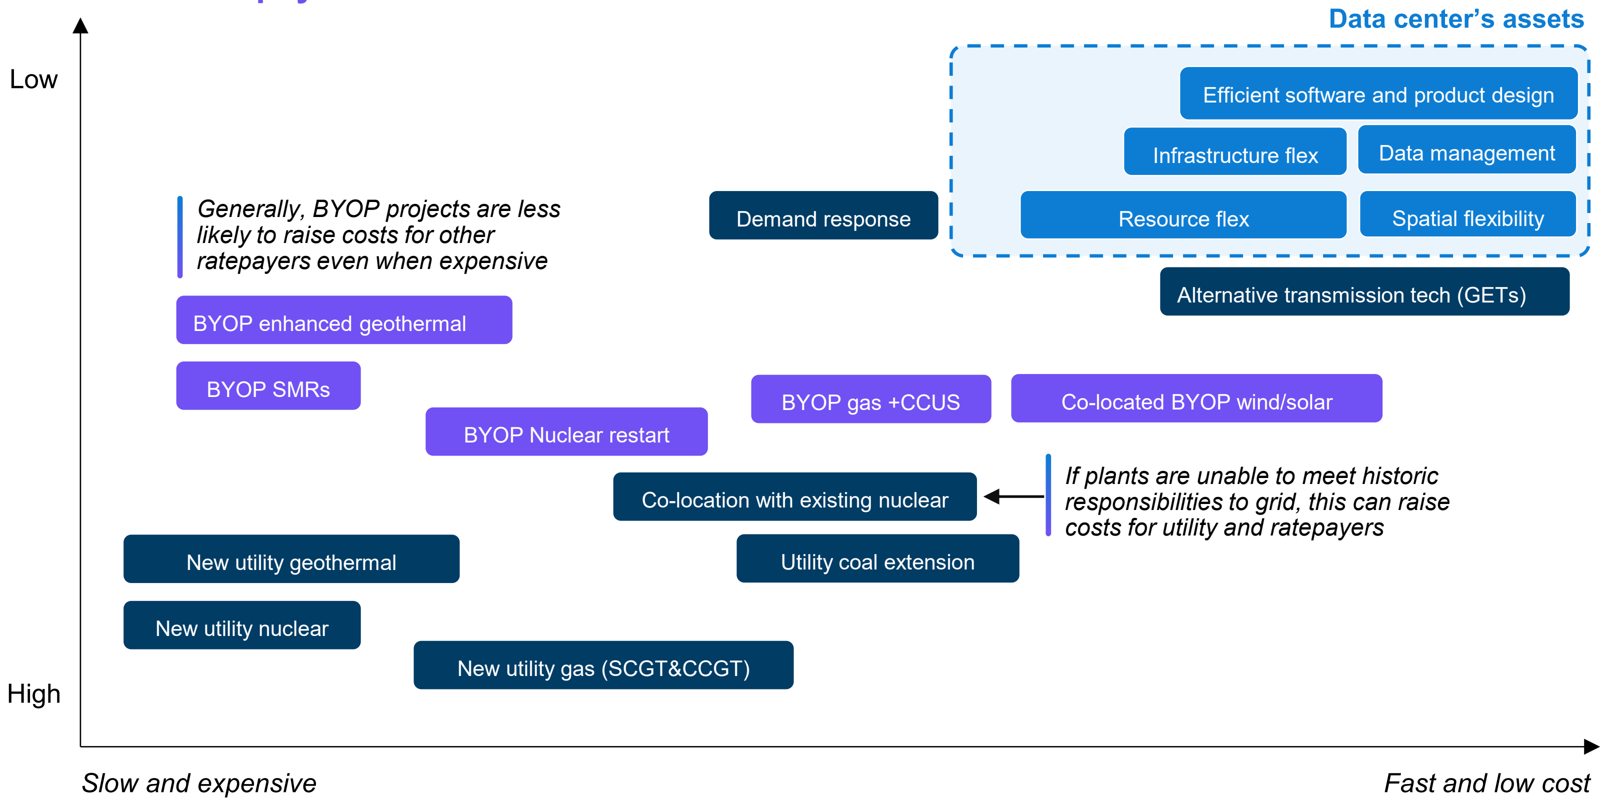

BYOP: Bring your own power, refers to energy projects that are financed by data centers

Most data center facilities can employ multiple capabilities to enable flexibility

Modes of data center flexibility

	Energy	Compute
Physical	1 Grid interactive UPS 2 Backup generation integration 3 BESS integration	6 Infrastructure flexibility
Contractual	4 Interruptible service models 5 Flexibility-contingent interconnection	7 Temporal shifting 8 Spatial load shifting

Definitions

1. Turns UPS units into fast-response ancillary services assets, providing FFR,1 FRT2 & PFAPR3
2. Uses natural gas and diesel gensets to deliver spinning reserves, capacity and reliability support within emissions limits
3. Deploys multi-hour batteries to shift load and shave peaks for both the facility and the grid
4. Provides faster interconnection and cost benefits in exchange for curtailment commitments during periods of grid stress
5. Accelerates energization timelines by verifying that facilities can shed or offset load when required
6. Uses thermal storage, dynamic setpoints and liquid or immersion cooling to shift cooling load without affecting IT performance
7. Reschedules delay-tolerant jobs to off-peak hours to cut cost and support grid stability
8. Routes workloads to data centers in other regions with lower grid stress

Source(s): [Sources text] | **Note(s):** 1. Fast frequency response; 2. Fault ride-through; 3. Post-fault active power recovery

Early pilots show that data centers can support the grid without compromising service reliability (SLAs)

Compute flexibility-Phoenix AI factory's load profile vs grid peak, job performance by flex tier

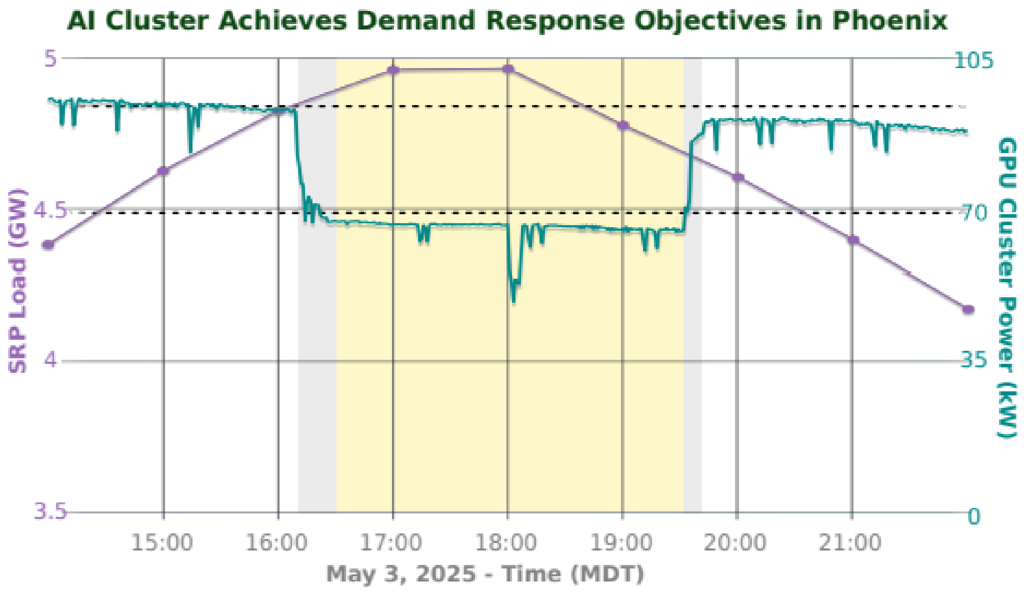

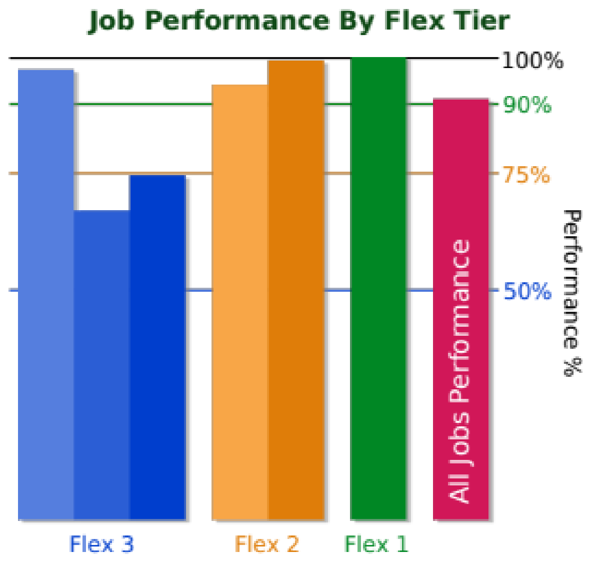

- **Emerald AI co-led with Oracle, NVIDIA,** Salt River Project and EPRI DCFlex in Phoenix AI factory pilot
- **Emerald AI cut power demand by 25% across 256 NVIDIA GPUs** while maintaining service levels for 3 hours, ramping down and up gracefully over 15 minutes. Its system continuously profiles workloads for flexibility, dynamically balancing live grid events and data center operations by rescheduling non-time-sensitive computing

The data center already has much of the physical and digital stack needed for flexibility; the challenge is orchestration, not invention

Teneo **Source(s):** Emerald AI, Google, Desk research

Early pilots show that data centers can support the grid without compromising service reliability (SLAs)

Temporal workload shifting pilot-Google compute load shifting illustration

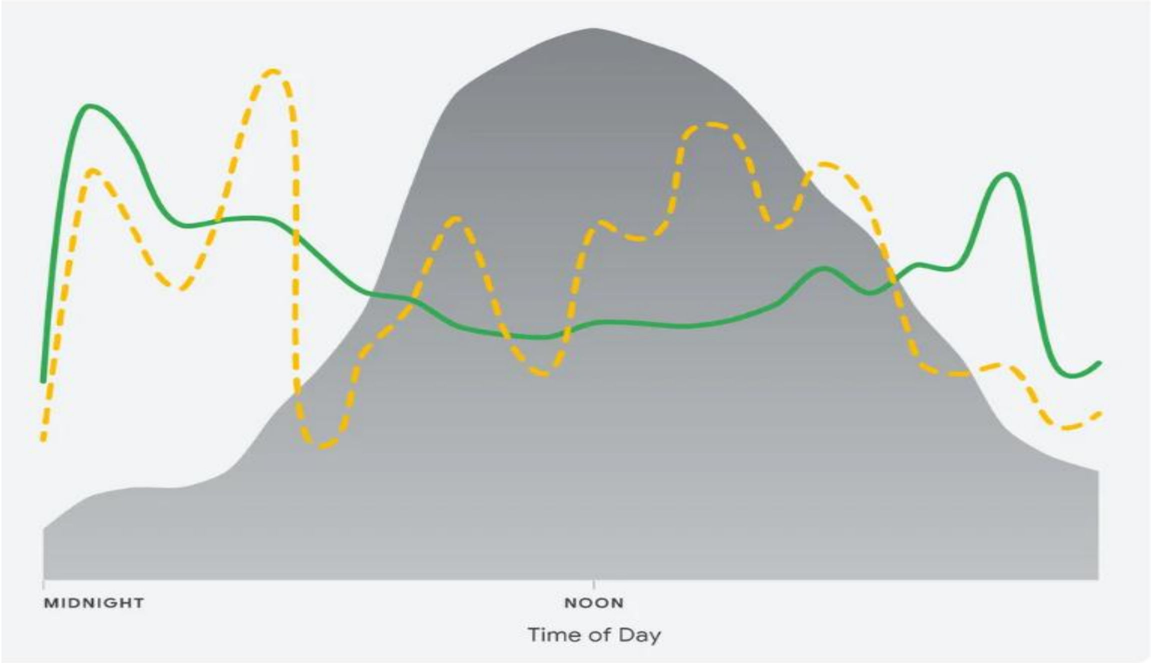

- Google shifts compute load from baseline (dashed line) to better align with less carbon-intensive times of the day, such as early morning and late evening (solid line)
- Gray shading represents times of day when more carbon-intensive energy is present on the grid

Key

- Adjustment load (solid line)
- Baseline load (dashed line)

The data center already has much of the physical and digital stack needed for flexibility; the challenge is orchestration, not invention

Source(s): Emerald AI, Google, Desk research

GWs of additional capacity can be unlocked with curtailing loads during peak demand hours or shifting energy using BESS

2024 Dominion Energy (VEPCO) hourly load intensity figures

Illustrative

Peak load (absolute peak: 23.1 GW)

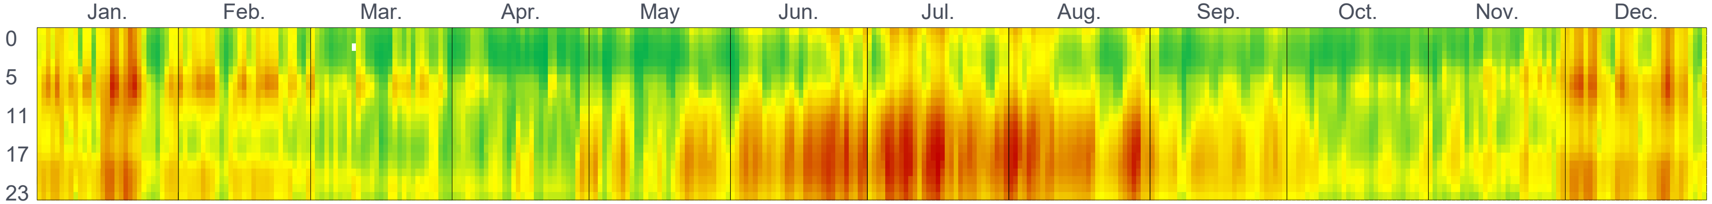

Dominion peak load is concentrated around extremely cold winter days and hot summer days and does not reflect most operating hours

Load with top 300 hours curtailed (scenario peak: 19.3 GW)

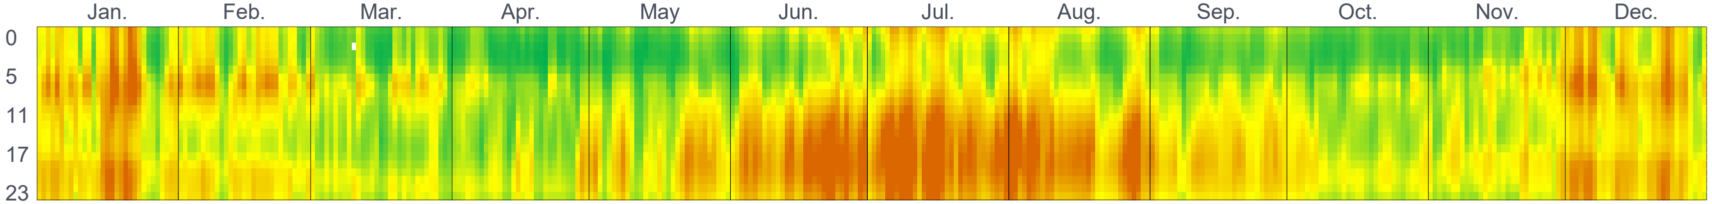

Curtailing 300 hours reduces absolute peak load by c. 15%, creating an additional c. 3.8 GW of available capacity, considering existing infrastructure

Load with aggregated 12-hour battery storage sized to 20% of peak (scenario peak: 18.7 GW)

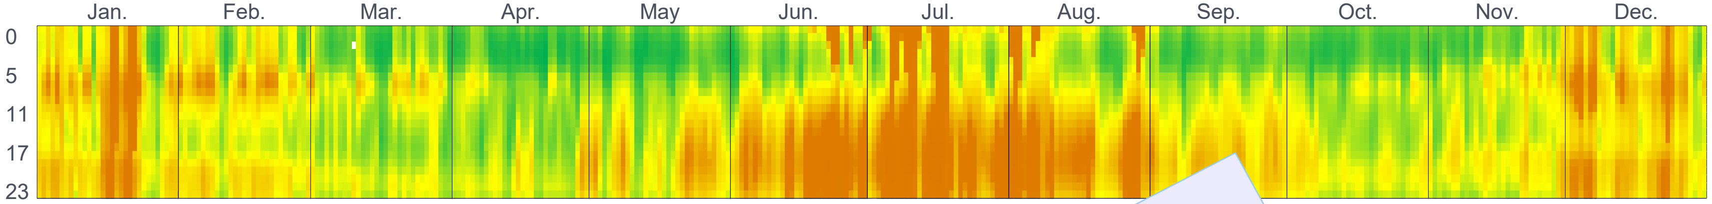

Backup energy storage shifts load to lower-peak hours and retains 100% data center operations, reducing absolute peak load by c. 19% and unlocking c. 4.5 GW

Battery storage increases load in some hours as they charge, but does not create a new peak load

Source(s): impactECI

Scaling flexibility requires coordination across a broad ecosystem of stakeholders

Six tier flexibility ecosystem, overview

1 Grid operators and utilities

- Coordinate overall system reliability and issue flexibility signals
- RTOs manage markets and large-scale coordination
- Local utilities handle distribution and customer interfaces

2 Aggregators and orchestration platforms

- Convert grid signals into actionable site level responses
- Enable real-time coordination and portfolio-wide flexibility at scale

3 Data center operators

- Build and operate facilities that embed flexibility into design and operations
- Operators must balance participation in grid programs with uptime and service commitments

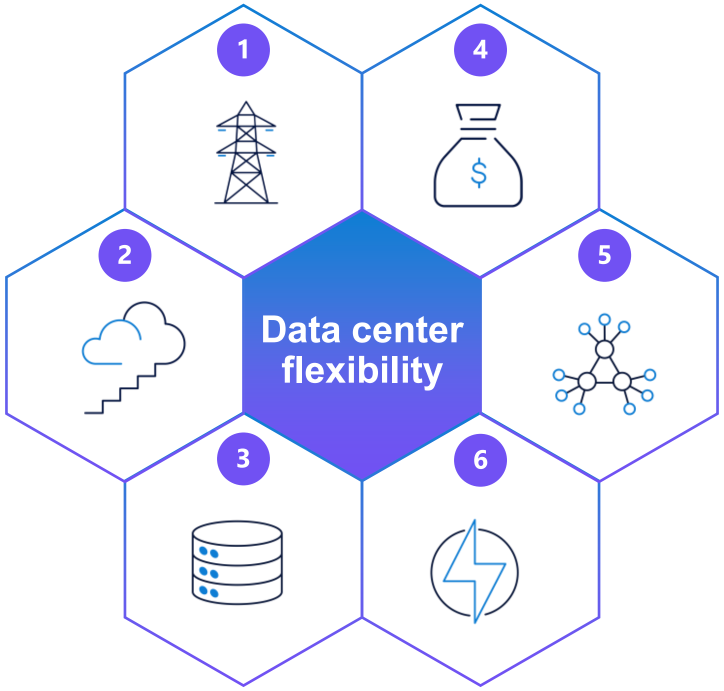

4 Financial stakeholders

Provide the funding, risk management and validation enabling data center ramp and flexible operations

5 Standards and community organizations

Establish common protocols, certifications and best practices to ensure interoperability

6 On-site flexibility assets

Provide physical means for flexible response, including backup generators, batteries, connected cooling and adaptive IT systems

Across markets, flexibility for large loads is advancing quickly, but through a patchwork of rules, pilots and participation pathways

Current RTOs/ISOs regulatory overview

Market	Key market mechanisms	Participation pathways	Recent developments
PJM	Capacity market, Large Load Additions (LLA) framework, flexible interconnection proposals	Dispatchable capacity, aggregated third-party resources, automated load response, emerging contractual flexibility for large loads	Hyperscalers and PJM are negotiating an LLA framework that would allow voluntary flexibility as an alternative to mandatory curtailment constructs, with new interconnection and tariff models under active stakeholder review
CAISO	Resource adequacy, flexible capacity, bilateral contracts, ancillary services, emergency programs (ELRP). PG&E's "Flex Connect" program has five data center/EV sites active, shifting from pilot to scalable model in California	Direct or hybrid participation, aggregator models, load modulation	Policy push for flexible, decarbonized capacity and hybrid data center-storage systems
ISO-NE	Forward capacity market (three-year-forward), frequency and reserve markets, transmission cost allocation. Specific flex load mechanisms for AI/data center clusters are still in process	Large centers (>20 MW) with self-funded interconnection, demand response programs, emerging pathways for flexible large loads	ISO-NE is revising tariffs and transmission cost structures for high-demand users and exploring mechanisms for AI/data center clusters to participate in capacity and flexibility products while maintaining reliability
ERCOT	Real-time energy markets, ancillary services (regulation, RRS, non-spin, ECRS), Controllable Load Resource (CLR) programs - Senate Bill 6 implementation	Direct participation as CLRs, telemetry-based curtailment, interruptible contracts, large-load interconnection tied to flexible, dispatchable load status	Senate Bill 06 (SB06) introduced stricter requirements on data centers, including cost sharing for grid upgrades, mandatory information sharing about BTM assets and load curtailment during emergency events. PGRR134 and associated proposals create a faster interconnection path for Interconnecting Large Load Entities that qualify as Controllable Load Resources
MISO	Resource adequacy, aggregation pilots, DR programs	Aggregator participation and direct flexible load bids	Early flexible interconnection proposals for hyperscale loads
SPP	Capacity and energy markets, HILL/CHILL, 90-day interconnection tariff and adopted system-wide	Flexible load and interruptible contracts, provisional service with paired generation	HILL/CHILL framework enables interconnection of large loads in 90 days when paired with new generation and subject to curtailment
NYISO	Capacity, reserves, ancillary services, advanced demand response, storage integration	Direct or aggregated flexible resource participation	NYISO has updated rules for flexible resource qualifications and is refining demand response and storage participation models that can accommodate large, controllable loads including data centers
Non-RTOs	Utility bilateral contracts, vertically integrated resource planning, state-run DR programs	Utility-specific pilots, expedited interconnection deals, bespoke flexibility or interruptible-service contracts	Utilities in high-growth states (AZ, UT, GA and throughout the Southeast) are piloting flexible data center demand response and bespoke large-load tariffs, using bilateral structures rather than organized market programs

Source(s): impactECI

Utility tariffs show early movement on flexibility, but most still lack a clear and standardized framework for large loads

Data center tariffs across 10 utilities

Utility	State	Contract length	Min demand	Min load factor	Power factor range	Requirement for investment	Cost assignment	Shed load required	Subject to interruptible service	Max hours of interruptible per year	Demand response
Cheyenne Light Fuel & Power / Black Hills Corporation	WY	Included	Included	Not applicable or not specified	Not applicable or not specified	Included	Included	Included	Included	Included	Not included
entergy Arkansas	AR	Not applicable or not specified	Not applicable or not specified	Not applicable or not specified	Not applicable or not specified	Included	Included	Included	Included	Included	Not applicable or not specified
Idaho Power	ID	Included	Included	Not applicable or not specified	Included	Included	Not applicable or not specified	Included	Included	Included	Not applicable or not specified
New York Power Authority	NY	Not applicable or not specified	Included	Not applicable or not specified	Not applicable or not specified	Included	Included	Not included	Not applicable or not specified	Not applicable or not specified	Not included
Montana-Dakota Utilities Co.	SD	Included	Included	Included	Included	Not included	Not included	Included	Included	Included	Included
Grant PUD	WA	Not included	Not included	Not applicable or not specified	Not applicable or not specified	Not included	Not included	Not included	Not applicable or not specified	Not applicable or not specified	Included
Indiana Michigan Power	IN	Not included	Included	Not applicable or not specified	Not applicable or not specified	Included	Not applicable or not specified	Not included	Not applicable or not specified	Not applicable or not specified	Not included
Kentucky Power	KT	Included	Included	Not applicable or not specified	Not applicable or not specified	Not applicable or not specified	Not applicable or not specified	Not included	Included	Included	Included
evergy	MO	Included	Included	Not applicable or not specified	Not applicable or not specified	Not applicable or not specified	Included	Not included	Not applicable or not specified	Not applicable or not specified	Included
Montana-Dakota Utilities Co.	ND	Included	Included	Included	Included	Included	Included	Included	Included	Included	Included
Dominion Energy	VA	Included	Included	Included	Not applicable or not specified	Included	Included	Not applicable or not specified	Not applicable or not specified	Not applicable or not specified	Not applicable or not specified

Included; Not included; Not applicable or not specified

Source(s): Stacy Sherwood - Jan 2025, Dominion Energy

For utilities, flexibility unlocks millions of dollars in value by deferring capacity, lowering operating risk and delivering steady ELCC value

Data center flexibility — Economics and value stack for utility

01 Transmission & capacity deferral

- Defers or avoids expensive transmission and peaking generation builds
- Gridlab-NV Energy: ~$300-400M+ NPV from 1–2 GW of curtailment, equivalent to 0.6-1.5 GW of firm capacity (67-87% ELCC)
- Acts as a "non-wires alternative"; value compounds in constrained regions (Northern Virginia, the Northeast)

02 Lower cost vs. grid capacity alternatives

- Marginal cost vs. $3,800/kW for BQDM-style DR, $2.5-6M/mile for new transmission
- Immediate deployment vs. 10-20 years for new lines
- Relieves stress at the source, but investment sits behind-the-meter, raising cost-sharing questions

03 Improved operations & reliability

- Smoother ramps and fewer hours hitting thermal and voltage limits
- Reduces reliance on the most expensive marginal resources
- Allows utilities to defer/resize projects justified only by worst-case hours
- Frees capital for higher-utilization investments (cleaner resources, network reinforcement)
- Avoids permitting, litigation and community-opposition risk on greenfield builds

04 Steady, biddable capacity value (ELCC)

- 1 GW of multi-hour flexibility can offset ~2/3 to >4/5 of new firm capacity needs
- Functions as a "virtual resource" in capacity auctions
- In vertically-integrated regions, captured in resource plans and reserve studies

05 Carbon emissions reduction

Regional impact varies (MIT 2025)

- ERCOT sees strong cost + emissions wins
- Mid-Atlantic gets cost wins but possible emissions uptick on fossil-heavy margin
- WECC modest on both

Source(s): impactECI; Gridlab.org; ITIF (Robin Gaster, Nov 2025); MIT CEEPR 2025 ("Flexible Data Centers and the Grid"); E3 (SPP demand response study); DOE National Transmission Needs Study 2023

Utilities capture large system value from a small number of peak-stress hours

NPV of system cost saving [$M] and ELCC contribution [MW] of 1-2 GW data center flexibility at different uptime requirement

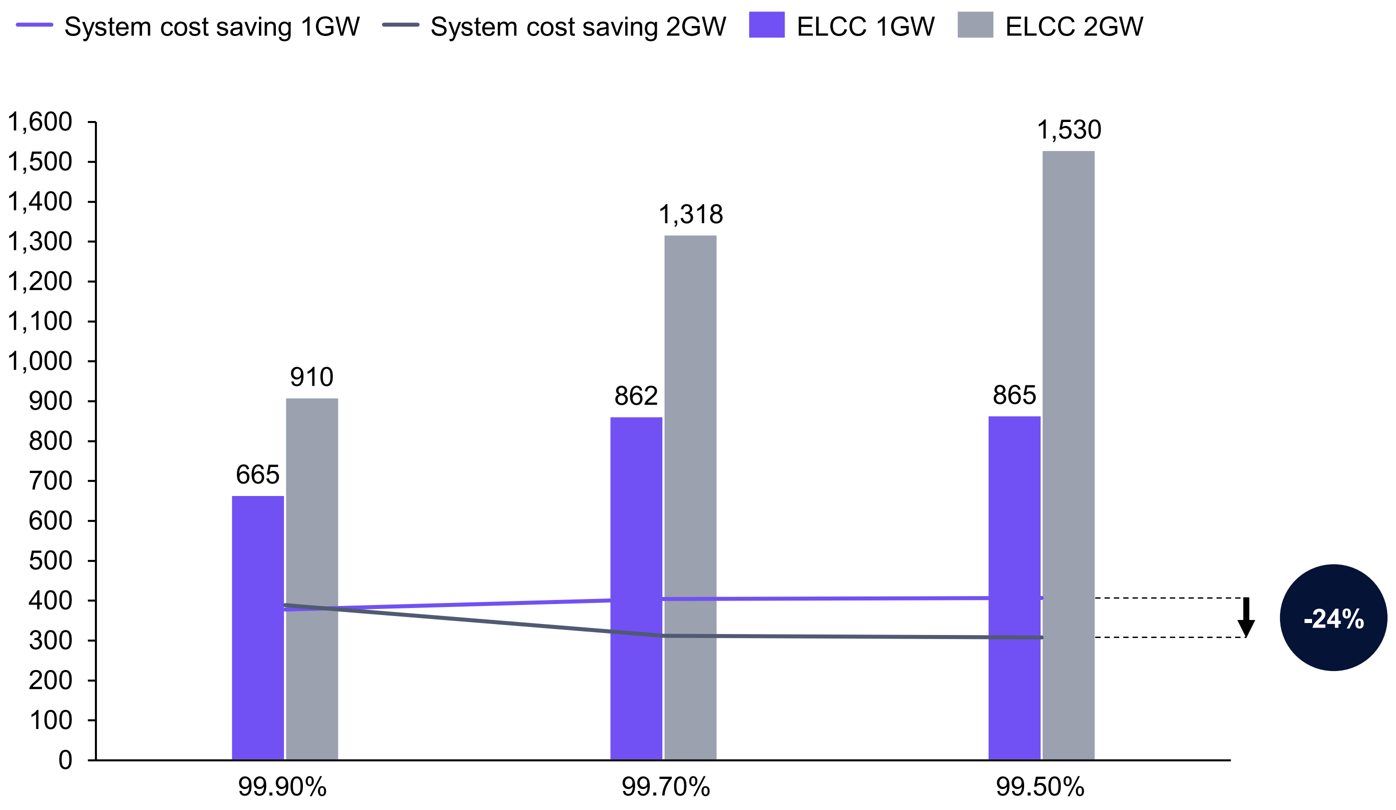

Key takeaways

1 GW of data center flexibility delivers ~$380–410M in NPV system benefits, driven by:

- Avoided gas peaker additions
- Deferred firm-capacity procurement
- Delayed battery and storage buildout
- Reduced reserve-margin requirements
- Deferred capacity investments over 2025–2050

Benefits largely plateau beyond 99.7% uptime, indicating that utilities capture most system value from **a limited number of annual peak-stress events**

Adding a second GW provides lower incremental benefit despite higher ELCC contribution, reflecting saturation effects within the NV-Energy system

- **Saturation effects:** The first GW of flexibility alleviates the most critical reserve-margin and firm-capacity constraints, avoiding high-cost peaker and firm-capacity additions
- **Shift from capacity constraints to energy constraints:** After the initial capacity bottleneck is relieved, the system increasingly requires energy delivery rather than additional firm capacity, reducing incremental system savings provided by additional flexibility

Aggregator and orchestration platforms are the translation layers that turn grid signals, through optimization and dispatch, into safe, reliable actions across an entire fleet

Optimize flexibility across multiple resource types

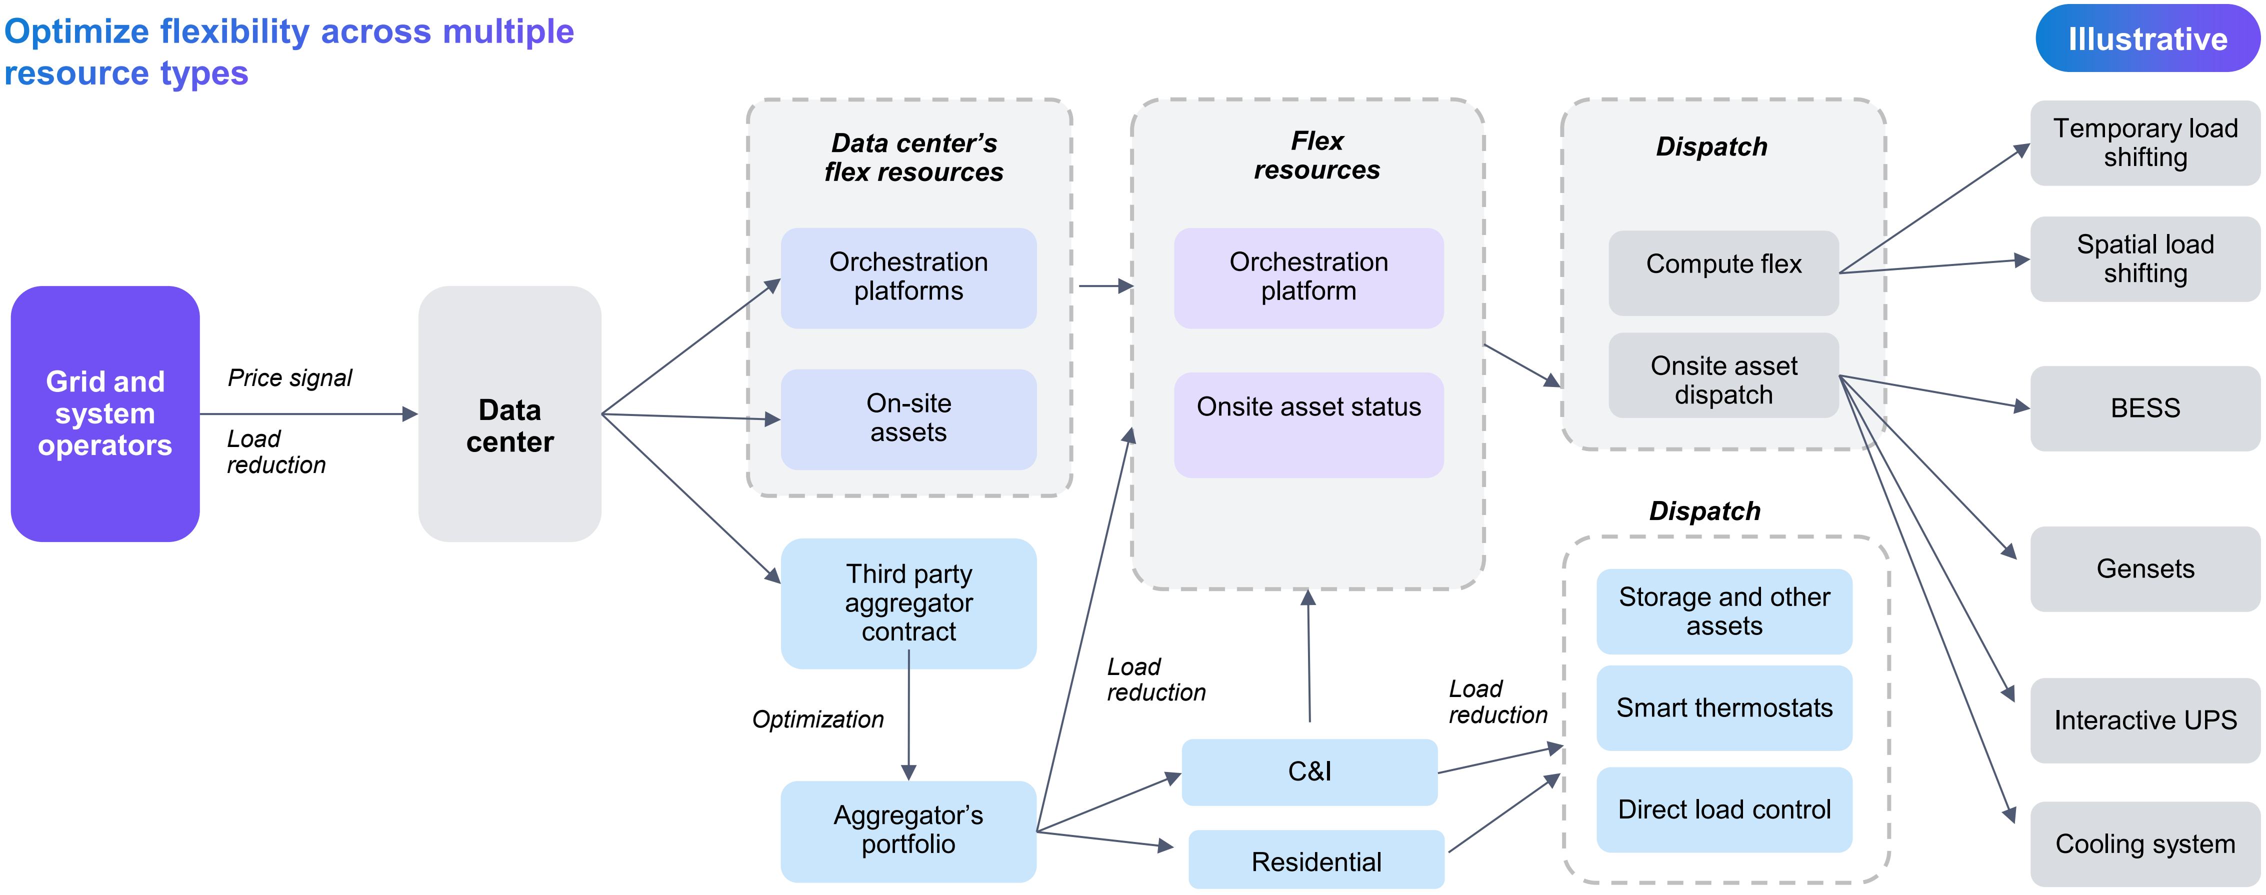

Teneo **Source(s):** Teneo analysis, impactECI

For data center operators and developers, flexibility pulls revenue forward by years and adds a recurring grid-services stack on top of core income

Data center flexibility — Economics and value stack for data centers

01 Accelerated grid connection

- Accelerate “speed to power”: federal guidance (DOE, FERC) and utility practice now prioritize adaptable facilities
- Trade a small amount of operational risk for major reductions in delay, redesign, stranded-asset, and lost-revenue risk

02 New grid-services revenue stack

Capacity, demand response, ancillary services and frequency regulation becomes recurring income

- PJM 2025/2026 capacity cleared at $98-170K/MW-year ($269–466/MW-day) - up from just $29/MW-day in 2024/2025
- Aggregators paying out $70-100K+/MW-year
- GridBeyond offers a 10-year guaranteed $73K/MW-year with 100% revenue share

03 Operational and sustainability gains

- Portfolio-level workload orchestration: treat the global fleet as one virtual asset, shifting AI/ML jobs by region based on price, marginal emissions and grid stress
- Aligns with hourly carbon accounting, GEB requirements and SEC climate disclosures, making 24/7 CFE claims credible
- Reduces fuel volatility risk, local pollutant exposure and O&M on gensets in markets with aging infrastructure

04 Competitive advantage in constrained markets

- Unlocks access to grid-constrained strategic regions where power blocks new development
- Differentiated tenant offering for co-location: flex-tier SLAs, dynamic pricing, verifiable low-carbon ops without tenants building their own controls
- Earlier energization=earlier sales of cloud, AI capacity, colocation or SaaS years ahead of queue-bound competitors

05 Goodwill and corporate reputation

- Repositions data centers from "grid burden" to "grid partner" as utilities increasingly require developers to fund triggered upgrades rather than spreading costs to ratepayers
- Builds credibility with utilities, regulators, communities and investors under heightened public scrutiny

Source(s): impactECI; PJM 2024/2025 capacity auction results; GridBeyond (Environment+Energy Leader); LBNL demand flexibility programs (2024); Company 10-Ks (AWS, NVIDIA, Digital Realty)

For data center operators and developers, "speed to power" can pull forward billions of dollars in revenue, alongside other economic benefits

Data center interconnection timeline, *[Years]*

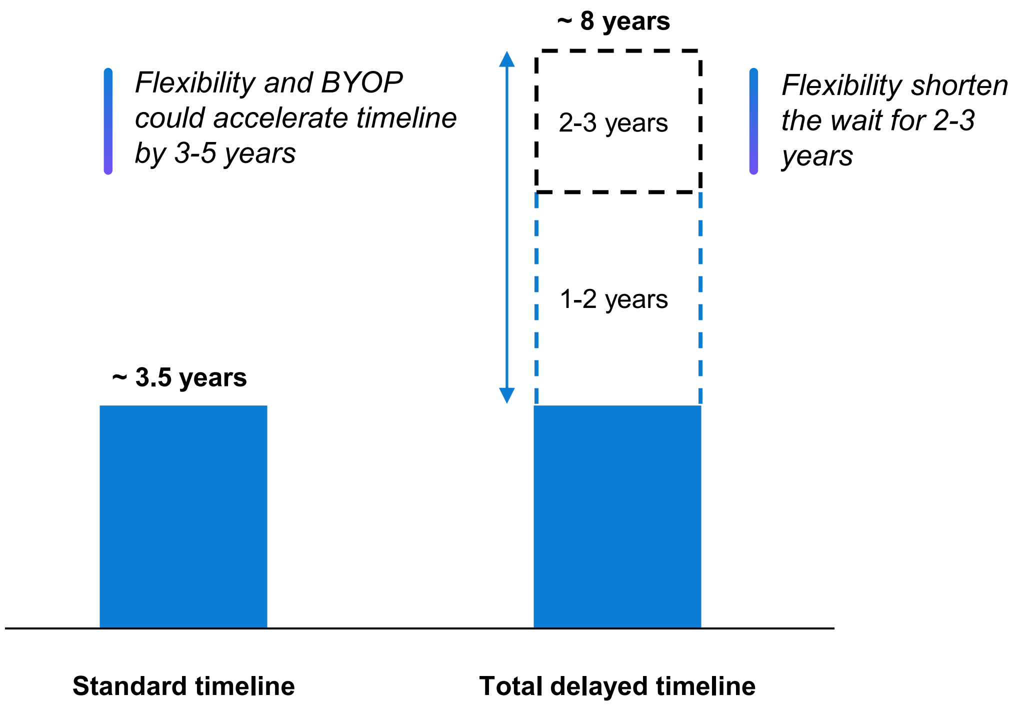

NPV of 100MW potential revenue stream of 2-year accelerated interconnection, *[$M]*

Illustrative

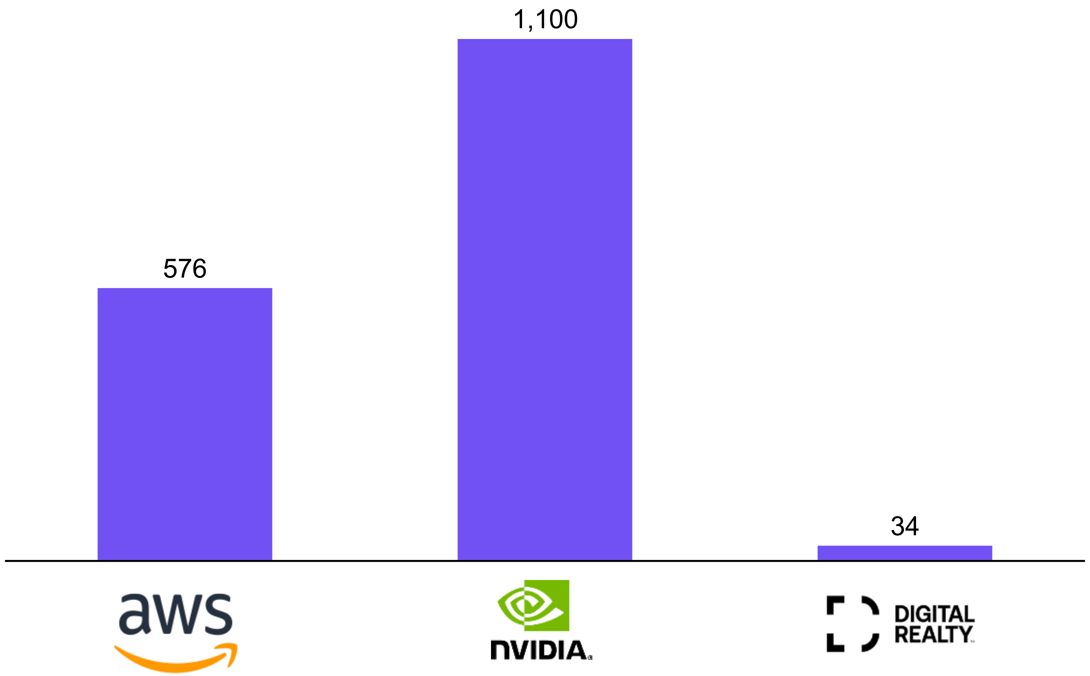

Source(s): impactECI, company 10k, industry news, desk research
Note(s): Hyperscale Amazon AWS 2024 10K North America profit $25B with est, 5.5GW in North America ~ $4.5m /MW-year, AI/HPC revenue based on NVIDIA ~ $9m/MW-year, Co-location based on Digital Realty 10K 2025 $0.25m/MW-year, discount rate 8.64%, standard COD date year 7, accelerated COD with thermal storage on year 5 of project timeline

Operational flexibility can increase project NPV by ~64% by accelerating time-to-power while avoiding grid electricity and interconnection costs

Indicative TX Data center project finance, NPV
($B)

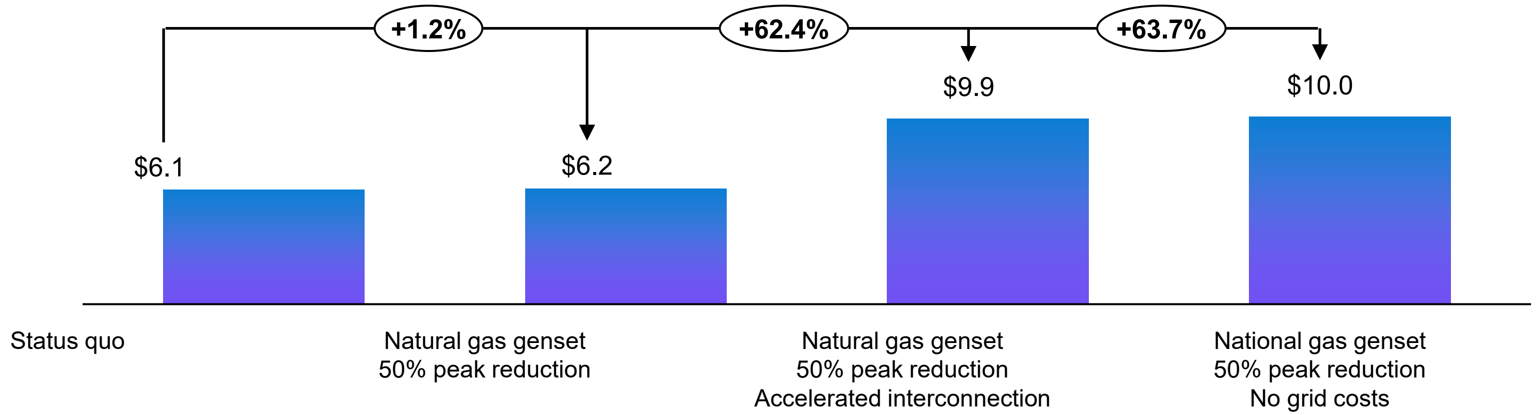

Selected inputs

Time to power	60 months	60 months	18 months	18 months
Revenue	$3,402/MWh	$3,402/MWh	$3,402/MWh	$3,402/MWh
Interconnection cost	$110M	$110M	$110M	$0
Generator hours	0	4,382	4,382	4,382
Gen var. O&M	$0.00/MWh	$31.21	$31.21	$31.21
Electricity expense	$46M	$9M	$9M	$9M

Indicative

Comments

- Value is concentrated in speed-to-power
- The benefits of avoiding electricity expenses through on-site generation are partially offset by natural gas fuel costs and variable O&M expenses
- Onsight generation can reduce 4CP load, lowering overall electricity costs in ERCOT
- Additional BTM generation and storage may enable facilities to participate in ancillary services markets, generating additional revenue from unused capacity

Source(s): Teneo, impactECI, JLL, CBRE, EIA AEO
Note(s): 1. Global assumptions include: 100 MW with 5.4 COP chiller in DFW area; Revenue based on the 3W/token and 0.3c/1M tokens; Cost of equity: 12%; Cost of construction debt : 8.5%, Cost of debt : 8.5%; GPU costs hit 2 months before first ramp; 2 ramp period; costs escalated by flat 3%, except electricity in Texas (4%); GPUs are replace every 5 years; Facility shell: $27,500/kW, GPU: 18,750/kW, Non fuel facility O&M: $11,000/MW-month; Non-fuel GPU O&M: $8,000/month; Free cooling estimated to begin around 50°F

Authors

Dan Gabaldon

Vice Chair of Management Consulting and Head of Energy

Robert Zabors

Vice Chairman, Management Consulting and Head of Infrastructure & Innovation

Omar Saadeh

Senior Director

Teneo®