EARTHJUSTICE
BECAUSE THE EARTH NEEDS A GOOD LAWYER

June 6, 2025

Via Electronic Filing

Matthew L. Homsher, Secretary
Pennsylvania Public Utility Commission
Commonwealth Keystone Building
400 North Street
Harrisburg, PA 17120

Re: En Banc Hearing on Interconnection and Tariffs for Large Load Customers; Docket No. M-2025-3054271

Dear Secretary Homsher,

Enclosed for filing in the above-referenced proceeding please find the Comments of Earthjustice as well as Attachments 1-4 associated with these comments. As evidenced by the attached Certificate of Service, all parties to the proceeding are being served with a copy of this document via e-Service or mail. Should you have any questions, please do not hesitate to contact me. Thank you.

Sincerely,

/s/ Caroline Weinberg
Senior Research and Policy Analyst
Earthjustice
1617 John F. Kennedy Blvd., Suite 2020
Philadelphia, PA 19103
cweinberg@earthjustice.org

cc:
Certificate of Service

CERTIFICATE OF SERVICE

I hereby certify that I have this day served a true copy of this electronically-filed document upon the parties, in accordance with the requirements of 52 Pa. Code § 1.54 (relating to service by a party).

SERVICE BY E-SERVICE AND MAIL

Celia Kosinski, Electrification Coalition
1111 19th Street NW #406
Washington, DC 20036
ckosinski@electrificationcoalition.org

Charles Harper, Evergreen Collaborative
1920 14th St NW
Apt 515
Washington, DC 20009
Charles@evergreenaction.com

Cindy Menhorn, VP
MCR Performance Solutions
2919 Seminary Drive
Greensburg, PA 15601
cmenhorn@mcr-group.com

Donna Kohut, PennFuture
610 North Third Street
Harrisburg, PA 17101
kohut@pennfuture.org

Eric Epstein
4100 Hillsdale Rd.
Harrisburg, PA 17112
epstein@efmr.org

Frank Lacey, President
Electric Advisors Consulting, LLC
3 Traylor Drive
West Chester, PA 19382
frank@eacpower.com

Gillian Graber, Executive Director
Protect PT
3344 Route 130
Harrison City, PA 15636

Joseph J Delmar Sr., Director
Cogentrix Energy Power Management
13860 Ballantyne Corporate Place
Suite 300
Charlotte, NC 28277
josephdelmar@cogentrix.com

Kevin Wright, President and CEO
ProtoGen, INC.
PO box 193
Quakertown, PA 18951
Contact@ProtoGenEnergy.com

Lauren Posey, Protect PT
3344 PA Route 130
Unit A
Harrison City, PA 15636
lauren@protectpt.org

Maureen Jorgensen, Calibrant Energy
4830 W. Kennedy Blvd.
Suite 600
Tampa, FL 33609
marketdev@calibrantenergy.com

Rebecca Lyttle, Office of Small Business Advocate
555 Walnut Street
1st Floor
Harrisburg, PA 17101

Richard C Culbertson
1430 Bower Hill Road
Pittsburgh, PA 15243
Richard.c.culbertson@gmail.com

Sara Majorsky
3344 Route 130
Harrison City, PA 15012
sara@protectpt.org

Stephen Parker
4100 Manor Oaks Ct.
Export, PA 15632

Travis Kavulla, VP of Regulatory Affiars
NRG Energy Inc.
1825 K Street Northwest Suite 1208
Washington, DC 20006

Dated: June 6, 2025

*_/s/ Caroline Weinberg_*________
Senior Research and Policy Analyst
Earthjustice
1617 John F. Kennedy Blvd., Suite 2020
Philadelphia, PA 19103
cweinberg@earthjustice.org

BEFORE THE
PENNSYLVANIA PUBLIC UTILITY COMMISSION

En Banc Hearing on Interconnection and Tariffs for Large Load Customers	Docket No. M-2025-3054271

COMMENTS OF EARTHJUSTICE

I. Introduction

Earthjustice appreciates the opportunity to submit comments on this timely and consequential hearing. Earthjustice is a nonprofit public interest environmental law organization. We represent concerned communities throughout the US in the face of energy affordability challenges, which includes protecting these communities from the ways emerging large digital loads may further exacerbate existing energy inequities.

We applaud the Commission for beginning the work needed to address the challenges that new large load customers pose to the electric grid and ratepayers in Pennsylvania. The potential impacts from improper planning for the connection of these loads to the grid are substantial but can be mitigated with appropriate structures in place. Each day, new research further illuminates the scale of the threat, especially within the PJM region.[1] Without protections, existing customers will foot the bill for large load-related capacity and infrastructure investments that existing customers do not benefit from, especially if the projected demand from large load customers fails to fully materialize.[2] To help the Commission as it seeks to minimize impacts on existing customers, we offer comments at a high level based on our involvement in similar proceedings throughout the country representing concerned communities. We look forward to continuing to engage with the Commission as a more detailed and specific model tariff is proposed and modified.

---

[1] Ivan Penn, The New York Times, *Data Centers' Hunger for Energy Could Raise All Electric Bills* (May 16, 2025), https://www.nytimes.com/2025/05/16/business/energy-environment/data-centers-utilities-electricity-bills.html?unlocked_article_code=1.Hk8.pezY.X1Y60lDEEWjG&smid=url-share; Cathy Kunkel, IEEFA, *West Virginia Ratepayers Footing the Bill for Infrastructure Build Out* (May, 2025), https://ieefa.org/sites/default/files/2025-05/UPDATED_West%20Virginia%20Ratepayers%20Footing%20the%20Bill%20for%20Infrastructure_May%202025.pdf.

[2] Ari Peskoe and Eliza Martin, *Extracting Profits from the Public: How Utility Ratepayers Are Paying for Big Tech's Power* (March 5, 2025), https://eelp.law.harvard.edu/extracting-profits-from-the-public-how-utility-ratepayers-are-paying-for-big-techs-power/ (attached as Attachment 1).

II. It Is Appropriate and Beneficial to Many Stakeholders for the Commission to Develop a Large Load Tariff

As noted in almost every testimony submitted to this hearing, all stakeholders stand to benefit from the increased standardization provided by a tariff structure for large load customers. A tariff will offer large load customers clarity on costs associated with developing these sorts of projects in Pennsylvania, will allow utilities to add such customers more efficiently and with more certainty, and will establish the urgently needed safeguards to protect ratepayers in all customer classes from stranded asset risk and potential cost shifting. It will be important that large load customers are required to adhere to this tariff and not opt out ad hoc via special contracts or private power purchase agreements (PPAs).[3] If permitted, PPAs must be fully disclosed to the public and subject to Commission review.

We are glad that the Commission has begun this necessary process, following in the footsteps of commissions throughout the country that have begun to institute similar protections. The Commission must follow common sense cost causation principles and mandate large load customers to truly bear the risk of their demand projections. These protections are appropriate and necessary.

A tariff structure that requires material up-front financial commitment from large load customers should offer the co-benefit of reducing uncertainty in utility load forecasts. Nevertheless, we urge the Commission to consider whether additional guidance is necessary to

[3] *See* Direct Testimony of Benjamin Inskeep, *In re Verified Petition of Indiana Michigan Power Company for Approval of Modifications to its Industrial Power Tariff I.P.*, Cause No. 46097 (Ind. Utility Reg. Comm'n Oct. 15, 2024), at 22, https://iurc.portal.in.gov/_entity/sharepointdocumentlocation/0c49aa8c-168b-ef11-ac21-001dd8067cf7/bb9c6bba-fd52-45ad-8e64-a444aef13c39?file=CN%2046097--%20CAC%20Exhibit%201--10-15-24FINAL_Redacted.pdf, "I also strongly agree with I&M's decision to serve new large load customers under a published tariff rather than through individually negotiated special contracts, which are often kept confidential. Such an approach is more transparent, administratively efficient, and fair to existing and potential new customers."

establish standard load forecasting methodologies across utilities. Past data reveal that PJM utilities have almost always over-forecasted demand over the past two decades,[4] and recent reporting explains how "phantom" data center loads can lead utilities to expect new load only for it not to materialize.[5] Commission efforts to help minimize this phenomenon moving forward would be appropriate and impactful.

Luckily, the Commission has the opportunity to build a structure informed by best practices in other jurisdictions.[6] Our comments focus on many such best practices that we hope will be informative to the Commission.

III. Data Centers Pose Unique Challenges to the Electric System

Throughout the hearing, the Commission heard from testifying large load customers who claimed that data centers constitute only a piece of a larger load growth puzzle. While it is true that there are other contributors to load growth across the United States, the unique characteristics of load growth from large loads like data centers make it appropriate for the Commission to take swift and decisive action to prevent harm from these potential new loads.

---

[4] The Bipartisan Policy Center and Koomey Analytics, *Electricity Demand Growth and Data Centers: Guide for the Perplexed* at 9 (February, 2025), https://bipartisanpolicy.org/download/?file=/wp-content/uploads/2025/02/BPC-Report-Electricity-Demand-Growth-and-Data-Centers-A-Guide-for-the-Perplexed.pdf (attached as Attachment 2).
[5] Brian Martucci, Utility Dive, *A fraction of proposed data centers will get built. Utilities are wising up* (May 15, 2025), https://www.utilitydive.com/news/a-fraction-of-proposed-data-centers-will-get-built-utilities-are-wising-up/748214/.
[6] Ann Collier and Justin Lindemann, Smart Electric Power Alliance, *Innovative Utility Tariffs Pave the Way for Flexible, Carbon-Free Data Centers* (February 21, 2025), https://sepapower.org/knowledge/innovative-utility-tariffs-pave-the-way-for-flexible-carbon-free-data-centers/.

A. Data Center Load Growth Is Larger and More Concentrated Than Other Types of Load Growth in PJM and Nationally

We appreciate the Data Center Coalition's (DCC) citation to Lawrence Berkeley National Lab's (LBNL) 2024 report on data center growth,[7] which establishes data center demand in the context of larger demand to come over the next decades. The DCC seemed to suggest that this broader context may be a reason to shy away from tariff systems targeting specific types of new loads, instead preferring more "flexible" "best practices."[8] The LBNL report, however, establishes that "if investments are made on the grid side but the expected load fails to show up, ratepayers could be unduly burdened by cost recovery,"[9] and that current data center demand can be used "as an opportunity to develop the leadership and a foundation for an economy-wide electricity infrastructure expansion."[10] In other words, LBNL is suggesting that, with considerable electric sector growth to come in the next decades, now is the time to institute structures and guidelines that ensure the grid can expand in sustainable and equitable ways. Such structures can and should ensure that new large load customers pay their fair share, protecting residential and small businesses customers in Pennsylvania.

Further, while it is true that data centers are not exclusively responsible for load growth, a recent study from the National Electrical Manufacturers Association estimates that they are the largest driver from 2023 through 2037, constituting 32% of growth nationally, while

---

[7] Lawrence Berkeley National Laboratory, *2024 United States Data Center Energy Usage Report* (December, 2024), https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf.
[8] Testimony of Lucas Fykes on Behalf of the Data Center Coalition at 3-4, PA PUC Docket No. M-2025-3054271 (April 23, 2025).
[9] Lawrence Berkeley National Laboratory, *2024 United States Data Center Energy Usage Report* at 69 (December, 2024), https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf.
[10] *Id.* at 7.

transportation is 24% and the residential, commercial, and industrial sectors combined are 44%.[11]

In PJM, according to a new report from ICF, data centers account for a clear majority of anticipated load growth.[12] Transportation load will be much more geographically dispersed than data center load growth, as charging infrastructure will need to be available at regular intervals on transportation corridors. As the Office of the Consumer Advocate (OCA) noted in testimony, because of large digital loads' concentrated nature:

> [large load customer's] size, transmission and distribution-related costs, potential for stranded investment, and interconnection safety and reliability issues are more similar to those of generators than the typical industrial customer. Furthermore, many large load customers will have primary or backup generation of commensurate size to its load, making them both a large load and a large generator.[13]

A glaring example of this concentrated load is the 4.5 GW proposal at Homer City, just east of Pittsburgh.[14] For these reasons, concentrated digital large loads are an appropriate candidate for the Commission's close scrutiny and guidance in order to protect Pennsylvanians.

---

[11] National Electrical Manufacturers Association, *A Reliable Grid for an Electric Future: NEMA's Grid Reliability Study* at 5 (April, 2025), https://www.makeitelectric.org/wp-content/uploads/2025/04/grid-reliability-study-nema-deck.pdf.
[12] ICF, *Rising current: America's growing electricity demand* at 7 (May, 2025), https://www.icf.com/-/media/files/icf/reports/2025/energy-demand-report-icf-2025_report.pdf?rev=c87f111ab97f481a8fe3d3148a372f7f.
[13] Testimony of Darryl Lawrence on Behalf of the Pennsylvania Office of Consumer Advocate at 2, PA PUC Docket No. M-2025-3054271 (April 23, 2025).
[14] Commonwealth of Pennsylvania, *Homer City Generation Site Redevelopment* (May 15, 2025), https://www.pa.gov/agencies/dep/about-dep/regional-office-locations/northwest-regional-office/northwest-community-information/homer-city-generation-redevelopment.html.

B. Data Center Grid Reliability and Economic Benefits Should Not Be Overestimated

Many of the large load customer testifying witnesses claimed that data center customers may increase grid stability or resiliency where they interconnect. For example, they claimed that their facilities would cause "invest[ment] in resources that reduce grid disruptions for all customers,"[15] or would "improv[e] system resilience, and reduc[e] net stress on the grid."[16] They also claimed that "data centers can support and stabilize the electric grid while improving interconnection and efficiency."[17]

In Pennsylvania, as in many states in the US, upgrades to transmission and distribution infrastructure are overdue. In fact, Pennsylvania was selected by the US Department of Energy in 2024 for a grant specifically geared towards advancing critical transmission projects throughout the country.[18] It is unclear, however, whether large load customers would contribute to meaningful system upgrades in a way that would benefit all customers, or whether they would at best ensure that their load additions do not further aggravate existing transmission and distribution challenges. In fact, even when data center customers spur investment in energy infrastructure to sustain their new demand, data center operations can still pose huge day-to-day reliability challenges.[19] The reliability risks from data center load behavior are so great that the North American Electric Reliability Corporation (NERC) has developed an entire research plan to address potential impacts.[20] Recent reporting has highlighted the ways that swings in data

---

[15] Testimony of Michael Fradette on Behalf of Amazon Data Services, Inc. at 6, PA PUC Docket No. M-2025-3054271 (April 23, 2025).
[16] Testimony of Shawn Smith on Behalf of Vantage Data Centers at 2, PA PUC Docket No. M-2025-3054271 (April 23, 2025).
[17] Testimony of Brendon Baatz on Behalf of Google, LLC. at 3, PA PUC Docket No. M-2025-3054271 (April 23, 2025).
[18] US DOE, Transmission Siting and Economic Development Grants Program, https://www.energy.gov/gdo/TSED.
[19] Tim McLaughlin, Reuters, *Big Tech's data center boom poses new risk to US grid operators* (March 19, 2025), https://www.reuters.com/technology/big-techs-data-center-boom-poses-new-risk-us-grid-operators-2025-03-19.
[20] FERC, *Presentation | NERC Seeks to Address Reliability Impacts from Large Load Integration* (April 17, 2025), https://ferc.gov/news-events/news/presentation-nerc-seeks-address-reliability-impacts-large-load-integration.

center electricity consumption might be distorting the normal flow of power for millions of Americans living nearby.[21] The Commission should thus be wary of vague talking points surrounding the potential grid stability or reliability benefits of large load customers and should ensure the model tariff establishes strong provisions pertaining to the investments these customers make in local grid upgrades and the ways that they operate once interconnected.

Separately, the large load customer testifying witnesses also touted the economic benefits they could bring to the region. Specifically, they claimed to offer "tremendous opportunity for Pennsylvania in terms of jobs, economic growth,"[22] and to "generate significant economic benefits for the communities in which they reside through job creation, tax revenue for schools and social services, and substantial local investment."[23] It is indisputable that data center taxes can generate revenues that can be reinvested in local area. However, data centers often receive substantial tax exemptions, causing state and local governments to miss out on potential additional revenue streams.[24] Any economic benefits from data centers should be accounted for in the larger context of this type of development's impact on a community. For example, in many places, including in Virginia, local residents are concerned about data center development's impact on water resources, open space, and historic sites.[25]

---

[21] Leonardo Nicoletti, Naureen Malik, and Andre Tartar, Bloomberg Technology, *AI Needs So Much Power, It's Making Yours Worse* (December 27, 2024), https://www.bloomberg.com/graphics/2024-ai-power-home-appliances.
[22] Testimony of Brendon Baatz on Behalf of Google, LLC. at 1, PA PUC Docket No. M-2025-3054271 (April 23, 2025).
[23] Testimony of Michael Fradette on Behalf of Amazon Data Services, Inc. at 3, PA PUC Docket No. M-2025-3054271 (April 23, 2025).
[24] Kasia Tarczynska and Greg LeRoy, Good Jobs First, *Cloudy with a Loss of Spending Control: How Data Centers Are Endangering State Budgets* (April, 2025), https://goodjobsfirst.org/cloudy-with-a-loss-of-spending-control-how-data-centers-are-endangering-state-budgets/.
[25] Charlie Paullin, Virginia Mercury, *Virginia Explained: Data center expansion, with all its challenges and benefits* (May 28, 2024), https://virginiamercury.com/2024/05/28/virginia-explained-data-center-expansion-with-all-its-challenges-and-benefits/.

Further, as far as job impact is concerned, often 80 percent or more jobs created by a given data center facility are temporary construction job rather than long term locally based positions, which is exemplified in Meta's own data presented on job creation of their US data center fleet.[26] Early research into job creation from data centers found that states and localities were giving data centers an average of two million dollars per job in tax exemptions and incentives.[27]

While it is appropriate for the Commission to weigh the benefits large load customers may bring to Pennsylvania, we urge the Commission to fully investigate such purported benefits to ensure they are not overestimated.

IV. The Model Tariff Must Incorporate Several Characteristics in Order to Meaningfully Protect Ratepayers

To address the above concerns with limited good jobs, the likelihood of over-forecasted demand, and potential impacts on other customer classes, we offer the following recommendations based on best practices established by several utilities in Ohio, West Virginia, Kentucky, and Indiana as well as by leaders in the energy sector like LBNL[28] and Energy Futures Group (EFG).[29]

---

[26] Meta Data Centers, US Data Center Fleet, https://datacenters.atmeta.com/us-locations/.

[27] Good Jobs First, *Study: State and Local Governments Pay $2 Million per Job to Tech Giants for Data Centers* (October, 2016), https://goodjobsfirst.org/study-state-and-local-governments-pay-2-million-job-tech-giants-data-centers/.

[28] Andrew Satchwell et. al., Lawrence Berkeley National Laboratory, *Electricity Rate Designs for Large Loads: Evolving Practices and Opportunities* (January, 2025), https://eta-publications.lbl.gov/sites/default/files/2025-01/electricity_rate_designs_for_large_loads_evolving_practices_and_opportunities_final.pdf.

[29] Stacey Sherwood, Energy Futures Group, *Review of Large Load Tariffs to Identify Safeguards and Protections for Existing Ratepayers* (January, 2025), https://energyfuturesgroup.com/2025/01/30/efg-teams-with-earthjustice-on-safeguards-and-protections-for-existing-ratepayers-related-to-large-load-tariffs/ (attached as Attachment 3).

A. The Tariff Should Establish a Customer Class Based on Load Profiles of Data Centers in the Region

The Commission should establish an appropriate lower-bound load (MW) threshold for customers to be eligible for service under the tariff and this bound should be developed in accordance with the size of interconnection requests that utilities in the state are receiving, likely in the range of 25 – 75 MW. In addition, as recommended in the testimony of the OCA,[30] a company's aggregate control of large load facilities should be accounted for, and this company should be subject to the tariff if its aggregated load surpasses a certain threshold (likely higher than the individual threshold) even if any of the company's individual facilities do not exceed the individual threshold. The Commission may reference Indiana Michigan Power Company's (I&M) tariff that includes language surrounding load aggregation across a company's individual facilities.[31]

As noted above and below, individual special contracts or PPAs evading these ratepayer protections should not be allowed, or if allowed, should be fully made public and subject to Commission review with sufficient public comment.[32]

---

[30] Testimony of Darryl Lawrence on Behalf of the Pennsylvania Office of Consumer Advocate at 4, PA PUC Docket No. M-2025-3054271 (April 23, 2025).

[31] Order of the Commission, *In re Verified Petition of Indiana Michigan Power Company for Approval of Modifications to its Industrial Power Tariff I.P.*, Cause No. 46097 (Ind. Utility Reg. Comm'n Feb. 19, 2025), at 52, https://iurc.portal.in.gov/_entity/sharepointdocumentlocation/2b48cf93-d9ee-ef11-be20-001dd80b8c52/bb9c6bba-fd52-45ad-8e64-a444aef13c39?file=ord_46097_021925.pdf ("I&M Settlement Order").

[32] Direct Testimony of Benjamin Inskeep, *In re Verified Petition of Indiana Michigan Power Company for Approval of Modifications to its Industrial Power Tariff I.P.*, Cause No. 46097 (Ind. Utility Reg. Comm'n Oct. 15, 2024), at 22, https://iurc.portal.in.gov/_entity/sharepointdocumentlocation/0c49aa8c-168b-ef11-ac21-001dd8067cf7/bb9c6bba-fd52-45ad-8e64-a444aef13c39?file=CN%2046097--%20CAC%20Exhibit%201--10-15-24FINAL_Redacted.pdf; Ari Peskoe and Eliza Martin, *Extracting Profits from the Public: How Utility Ratepayers Are Paying for Big Tech's Power* at 32, (March 5, 2025), https://eelp.law.harvard.edu/extracting-profits-from-the-public-how-utility-ratepayers-are-paying-for-big-techs-power/ (attached as Attachment 1).

B. The Tariff Should Include Collateral Requirements That Are Independent of Business Models of the Large Load Customers to Protect Other Ratepayers from Stranded Asset Risk

One of the DCC's three focuses of its written testimony pertained to flexibility in collateral requirements.[33] Generally, when it comes to the types of collateral the Commission should allow under this provision, we agree with the OCA in their spoken testimony when they stated that the Commission must not "simply expand the universe of potential options based on someone's business model."[34] Collateral has been an important pillar of several large load tariffs in other jurisdictions and we urge the Commission to use strong terms that do not offer large load customers undue flexibilities or exemptions. It may be appropriate for collateral requirements to be based on the creditworthiness of the customer, as has been established in the I&M tariff, but such creditworthiness thresholds should not be so flexible as to exempt substantial groups of customers, as the DCC has advocated for in other jurisdictions.[35] The collateral requirements themselves should correlate with minimum charges expected from the customer over the full contract term, in order to ensure that the collateral would protect other ratepayers if a large load customer unexpectedly discontinued service. The AEP Ohio proceeding has proposed language to this end, and has also proposed provisions that establish mechanisms for re-evaluation as the financial position of a given customer may change.[36] Collateral requirements, along with exit

---

[33] Testimony of Lucas Fykes on Behalf of the Data Center Coalition at 5, PA PUC Docket No. M-2025-3054271 (April 23, 2025).

[34] Spoken Testimony of Darryl Lawrence on Behalf of the Pennsylvania Office of Consumer Advocate at 31:30 in https://youtu.be/Z83c3m_9A-M?si=zhyu_-5mCBSmcRdB, PA PUC Docket No. M-2025-3054271 (April 23, 2025).

[35] Direct Testimony of Kevin C. Higgins, *In the Matter of the Application of Ohio Power Company for New Tariffs Related To Data Centers and Mobile Data Centers,* Case No. 24-0508-EL-ATA (The Public Utilities Commission of Ohio Oct. 18, 2024), at 6-7, https://dis.puc.state.oh.us/ViewImage.aspx?CMID=A1001001A24J18B60224E00611.

[36] Stacey Sherwood, Energy Futures Group, *Review of Large Load Tariffs to Identify Safeguards and Protections for Existing Ratepayers* at 18 (January, 2025), https://energyfuturesgroup.com/2025/01/30/efg-teams-with-earthjustice-on-safeguards-and-protections-for-existing-ratepayers-related-to-large-load-tariffs/ (attached as Attachment 3).

fees (discussed in the following section), provide necessary protections to all other customer classes who would otherwise end up bearing the stranded costs if large load customer demand fails to materialize or ends abruptly.

C. *The Tariff Should Establish Appropriate Exit Fees to Protect Other Ratepayers from Stranded Asset Risk*

If a large load customer elects to discontinue service before the end of its contract term, an exit fee is necessary to protect the other customer classes from potential stranded assets. The exit fee should be the value of the remaining minimum bill charge for the terminated or reduced capacity for the duration of the contract. This is the approach taken by I&M as well as Appalachian Power Company and Wheeling Power Company in their large load tariffs,[37] which the Commission could use as a model for these types of provisions. The Commission should also consider establishing a minimum notice period for early contract termination or service discontinuance. This will ensure that the utility can prepare for the change in load and this provision would act in conjunction with an appropriate exit fee.

D. *The Tariff Should Establish Incentives for Large Load Customers to Shed Discretionary High-Load Tasks While Disincentivizing the Use of Dirty Backup Generation as a Load Flexibility Mechanism*

The potential for large load customers to operate with some degree of flexibility was a recurring theme throughout the hearing. For example, in questions following the hearing, Vice Chair Kimberly Barrow asked data center developers, "To what extent do you intend to

---

[37] I&M Settlement Order at 55; *In re Application for Approval of Revision to Schedules LCP and IP*, Case No. 24-0611-E-T-PW (Public Service Commission of West Virginia April 7, 2025), at PDF 6, https://www.psc.state.wv.us/scripts/WebDocket/ViewDocument.cfm?CaseActivityID=639454&NotType=WebDocket.

collaborate with grid operators on the timing of discretionary high-load tasks like model training in order to avoid increasing coincident peaks?"[38] We believe the model tariff could help incentivize such behavior by including terms that mandate notice to the local operator of certain discretionary high-load tasks, allowing the large load customer and utility to coordinate these events in ways that minimize grid strain. Some Google data centers have already begun to operate with some flexibility, responsive to notice from the local grid operator.[39] A recent investigation modeled small changes in data center behavior during peak hours and found that these minimal shifts could create substantial headroom on the grid.[40] The Commission should include provisions in the model tariff that incentivize this type of behavior, while also protecting against the use of dirty backup generation as a mechanism for demand response. Many data centers have backup diesel generation on site, one of the most polluting types of energy production. Researchers have estimated that the public health costs of data center backup generation in Virginia alone could be hundreds of millions of dollars, even if these generators are operating at 10% of permitted levels.[41] Dirty backup generation is not the only option. In Wyoming, a data center developer and a battery company recently announced a partnership to deploy battery energy storage systems at data centers.[42] The Commission should include

[38] Barrow Questions at 1. Directed Questions of Vice Chair Kimberly Barrow at 1, PA PUC Docket No. M-2025-3054271 (May 1, 2025).
[39] Varun Mehra & Raiden Hasegawa, *Supporting power grids with demand response at Google data* (October 3, 2023), https://cloud.google.com/blog/products/infrastructure/using-demand-response-to-reduce-data-center-power-consumption.
[40] Tyler H. Norris, Tim Profeta, Dalia Patino-Echeverri, and Adam Cowie-Haskell, Nicholas Institute for Energy, Environment & Sustainability, *Rethinking Load Growth Assessing the Potential for Integration of Large Flexible Loads in US Power Systems* (February, 2025), https://nicholasinstitute.duke.edu/sites/default/files/publications/rethinking-load-growth.pdf (attached as Attachment 4).
[41] Yuelin Han, Zhifeng Wu, Pengfei Li, Adam Wierman, and Shaolei Ren, *The Unpaid Toll: Quantifying the Public Health Impact of AI* at 5 (December, 2024), https://arxiv.org/pdf/2412.06288.
[42] Zachary Skidmore, Data Center Dynamics, *Prometheus Hyperscale partners with XL Batteries to deploy long duration battery storage across US data centers* (May 16, 2025), https://www.datacenterdynamics.com/en/news/prometheus-hyperscale-partners-with-xl-batteries-to-deploy-long-duration-battery-storage-across-us-data-centers/.

protections in the model tariff such that data center customers are not incentivized to operate with "flexibility" if doing so necessitates emitting air pollutants from dirty generators.

When large load customers bring their own new generation to a development, it can offer substantial system benefits. NV Energy's Clean Transition Tariff offers a possible avenue for incentivizing large load customers to add new clean resources to the grid that support their operations.[43] The technologies are available for large load customers to add clean power to support their facility's operations, and the Commission should integrate language that promotes such activity into the model tariff.

E. The Tariff Should Establish a Minimum Contract Length to Protect Other Ratepayers from Stranded Asset Risk

Minimum contract lengths have been clearly established as a key piece of any large load tariff structure and were mentioned by almost all testifying witnesses during the hearing. This has been borne out by nearly all large load tariffs instituted to date.[44]

The large loads proposed in Pennsylvania will require substantial investments in the capacity and transmission system, and these investments carry significant risks. Therefore, a mandatory minimum contract length can help ensure that cost recovery on such investments are

---

[43] Emma Penrod, Utility Dive, *NV Energy seeks new tariff to supply Google with 24/7 power from Fervo geothermal plant* (June 21, 2024), https://www.utilitydive.com/news/google-fervo-nv-energy-nevada-puc-clean-energy-tariff/719472/.

[44] *See, e.g.,* I&M Settlement at 32; Order, *In re Appalachian Power Company and Wheeling Power Company Application for Approval for Revisions to Schedules LCP and IP.*, Case No. 24-0611-E-T-PW (Public Service Commission of West Virginia March 25, 2025), at 5, https://www.psc.state.wv.us/scripts/WebDocket/ViewDocument.cfm?CaseActivityID=638931&NotType=WebDocket ("West Virginia Tariff Order"); Order, *In re Electronic Tariff Filing of Kentucky Power Company to Revise its Industrial General Service Tariff*, Case No. 2024-00305 (Ky. P.S.C. March 18, 2025), at 3, https://psc.ky.gov/pscscf/2024%20Cases/2024-00305//20250318_PSC_ORDER.pdf ("Kentucky Power Company Order").

distributed over a longer period of time, in turn helping to ensure that large load customers adequately pay for necessary utility investments, borne by the cost causer.[45]

For example. the Kentucky Public Service Commission recently approved a 20-year minimum contract term for Kentucky Power's large load customers, which we believe may be appropriate in Pennsylvania as well.[46] We would encourage the Commission to also evaluate whether a minimum contract length of more than 20 years may be necessary to fully protect utilities and other ratepayers from bearing the cost of investments needed to serve new large load customers.

In addition, we recommend including a pricing adjustment mechanism to ensure that large load customers pay just and reasonable rates throughout the entire contract term. For instance, the Commission could include a tariff term stipulating that the monthly rates in the proposed tariffs could be reopened if the actual cost to serve large load customers increases such that the rates are no longer just and reasonable.[47] This would ensure that other customers would not have to pay increased rates if the actual cost to serve large load customers becomes greater than currently anticipated. Such a term would help protect against long-term uncertainty regarding the costs associated with serving large load customers.

---

[45] Ari Peskoe and Eliza Martin, *Extracting Profits from the Public: How Utility Ratepayers Are Paying for Big Tech's Power* at 31-32 (March 5, 2025), https://eelp.law.harvard.edu/extracting-profits-from-the-public-how-utility-ratepayers-are-paying-for-big-techs-power/ (attached as Attachment 1).
[46] Kentucky Power Company Order at 3.
[47] *See*, Evergy Metro, Inc., Special High-Load Factor Market Rate Schedule MKT at 3, https://www.evergy.com/-/media/documents/billing/missouri/detailed_tariffs_mo/special-high-load-factor-market-rate.pdf, where the price terms are necessarily shorter than the contract terms.

F. *The Tariff Should Establish Minimum Monthly Billing Demand and Restrictions on Reduction on Contracted Capacity to Protect Other Ratepayers from Stranded Asset Risk*

A high minimum billing demand requirement helps ensure that utilities are compensated for the cost of providing service sufficient to meet data centers' peak demand and that those costs are not shifted to residential rate payers and small businesses in Pennsylvania. We strongly recommend the adoption of a monthly billing demand of at least 90% of the highest of (a) the customer's on-peak contract capacity or (b) the customer's highest previously established monthly billing demand during the past 11 months or (c) the customer's maximum demand created during the billing month. As described above, there is still significant uncertainty as to how the increased demand associated with large load customers will materialize. A 90% minimum billing demand would align with tariffs recently approved by the Kentucky Public Service Commission and proposed before the Ohio Public Utilities Commission.[48] Other jurisdictions, such as the Indiana Utility Regulatory Commission and the West Virginia Public Service Commission, recently approved tariffs with 80% minimum billing demand.[49] A high minimum billing demand would help protect Pennsylvania utilities and their customers while still allowing large load customers sufficient flexibility.

---

[48] Kentucky Power Company Order at 3; Application for Approval of New Tariffs by Ohio Power Company, *In the Matter of the Application of Ohio Power Company for New Tariffs Related To Data Centers and Mobile Data Centers*, Case No. 24-508-EL-ATA (Public Utilities Commission of Ohio May 13, 2024), at 8, https://dis.puc.state.oh.us/ViewImage.aspx?CMID=A1001001A24E13B42822J00948.

[49] I&M Settlement Order at 32-33; West Virginia Tariff Order at 5-6.

G. *The Tariff Should Establish Restrictions on Reduction on Contracted Capacity to Protect Other Ratepayers from Stranded Asset Risk*

Along with minimum contract length and minimum monthly billing demand, limitations on reductions to contracted capacity are also important provisions for ensuring that large load customers pay their share of investments needed to serve them. For example, I&M's recent large load tariff establishes a 20% limit on reductions to a customer's contracted capacity. If the customer wishes to reduce its capacity under its long-term contract, it must pay a fee unless the reduction is agreed upon by the utility and submitted to the Indiana Utility Regulatory Commission for review and approval.[50] The I&M tariff also specifies that reductions in contracted capacity cannot occur during the first five years of a contract and requires that data centers provide 42 months' notice before reducing capacity.[51] In West Virginia, Appalachian Power Company and Wheeling Power Company's large load tariff likewise requires 42 months' notice for a contract capacity reduction and limits such reductions to 20% absent payment of an exit fee, unless there is mutual agreement that a reduction would be beneficial, or at least not detrimental, to the large load customer, utility, and all other customers.[52] Kentucky Power has a similar term establishing a five-year notice period for reductions to contract capacity and a reduction limit of 20% maximum contract capacity, with exceptions only in the case of mutual agreement.[53]

---

[50] I&M Settlement Order at 36-40.
[51] *Id*.
[52] West Virginia Tariff Order, Attach. A, Joint Stipulation and Agreement for Settlement at 6-7.
[53] Kentucky Power Company Tariff I.G.S. Revisions (Aug. 30, 2024), at PDF page 10 ("P.S.C. KY. NO. 13 ORIGINAL SHEET NO. 8-3"), https://psc.ky.gov/pscscf/2024%20cases/2024-00305/20240830_Kentucky%20Power%20Tariff%20Filing.pdf.

We strongly recommend that the Commission adopt similar restrictions on a large load customer's ability to reduce its contract capacity, in order to protect other ratepayers from having to bear the cost of investments already made to support the original contract capacity.

H. The Tariff Should Include a Cost Allocation Term to Protect Pennsylvanians

Direct assignment of certain costs to large load customers has become more common in tariffs and rates for large load customers and should be a part of the Commission's model tariff. The Commission may mirror recent language from Indiana or New York in developing these terms. I&M's recent large load tariff settlement assigns to each large load customer the cost of any "Full Planning Studies, including steady-state and dynamic studies, required because of the potential addition of a Large Load Customer."[54] Even with this provision, the settlement specifies that it does not limit the ability of the settling parties or Indiana Utility Regulatory Commission from addressing cost allocation in a subsequent proceeding.[55]

New York Municipal Power Agency's rider for rates and charges for high density load service includes an even stronger cost allocation provision. The rider mandates that, "[u]pon payment of security acceptable to the Utility, the Utility shall conduct, or cause to be conducted a feasibility study to evaluate whether the requested load can be safely served by the Utility."[56] The rider specifies that the customer is responsible for the reasonable costs of conducting this feasibility study.[57] Under the New York rider, the customer is also initially responsible for the entire cost of any new facilities necessary to supply the requested service, and the customer is

[54] I&M Settlement Order at 41.
[55] *Id.* at 9.
[56] New York Municipal Power Agency, Generic Tariff Rider A (Rates and Charges for Customers Requesting High Density Load ("HDL") Service), Leaf 95-96 (Mar. 23, 2018), https://ets.dps.ny.gov/ets_web/search/showPDF.cfm?%3B%3AIS%20%3B%2A%29LOUNWD%5CJ%5E8%2B%22%2B5%2F0MD%2F0%28%23%21%5E%2AS%3C%3F%5E%0A.
[57] *Id.* at Leaf 96.

required to pay the costs of any new facilities in cash before those facilities will be constructed.[58] However, at the end of each of the first ten years of service, the customer receives a "refund equal to the lesser of the annual non-supply related revenues from the customer, or one-tenth of the cost contribution paid by the customer."[59]

We recommend that the Commission's model tariff include a term directly assigning to large load customers certain costs associated with delivering service to those customers. At a minimum, this term should encompass any costs from conducting studies related to the potential addition of a large load customer. Like the New York rider, this term could also initially assign the cost of any new facilities needed to serve a customer to that customer, with a mechanism to refund a portion of that over a reasonable period of time that sufficiently protects against stranded cost risks.

I. The Tariff Should Include Provisions that Require Investment in Vulnerable Communities

In addition to appropriate cost allocation terms, the Commission should also consider terms that require investment in energy efficiency or other demand-side management initiatives that benefit low-income customers in Pennsylvania. This sort of additional investment was adopted in the I&M tariff, where, as part of the agreement, Amazon, Microsoft, and Google will each contribute $500,000 annually for five years to the Indiana Community Action Association (INCAA) to support low-income customer weatherization initiatives.[60] Stipulations like this can help compensate for some of the risks that large load customers add to the system. However, we appreciate the caution expressed by the OCA during the hearing about the ways these sorts of

[58] *Id.*
[59] *Id.*
[60] I&M Settlement Order at 60.

schemes could lead to geographic disparities if large load customers only make these investments hyper-locally, especially considering that the majority of Pennsylvania's lowest income customers live in urban areas where most of these large load customers are unlikely to site operations.[61] We urge the Commission to include these sorts of investment provisions with care towards ensuring equitable outcomes.

J. The Tariff Should Include Provisions that Require Investment in Grid-Hardening

New large loads will stress the grid in ways that many grid operators have not seen. Aside from requiring new large load customers to pay for the grid infrastructure to serve their own needs, these customers should also invest in upgrading the grid that serves all of the Commonwealth. These upgrades will also increase the amount of electricity that can serve data centers and other customers, in some places by up to 50%. Provisions that incentivize this sort of investment should be included in the model tariff. For example, the I&M tariff also included a stipulation that I&M "conduct a study, with input from interested Settling Parties as to the scope of the study and with opportunities for stakeholder feedback, to evaluate the potential of grid-enhancing technologies," which offer potential benefits to all customers.[62] These technologies offer near-term solutions that are proven to benefit all customers, and their benefits are especially important in this moment of load growth and uncertainty. We urge the Commission to creatively consider ways these investment provisions can be leveraged to benefit communities in an equitable manner.

[61] Spoken Testimony of Darryl Lawrence on Behalf of the Pennsylvania Office of Consumer Advocate at 42:46 in https://youtu.be/Z83c3m_9A-M?si=zhyu_-5mCBSmcRdB, PA PUC Docket No. M-2025-3054271 (April 23, 2025).
[62] I&M Settlement Order at 44.

V. Conclusion

We appreciate the opportunity to share insights that may be helpful to the Commission as they begin this necessary process based on our experience in other jurisdictions facing similar issues. We look forward to continued collaboration and public comment as the Commission proposes a detailed draft tariff that protects existing electricity customers in the Commonwealth and ensures that new large customers pay their fair share of their impacts on the grid.

Dated: June 6, 2025

Respectfully submitted,

/s/ Caroline Weinberg

Caroline Weinberg
Jacob Elkin, Esquire
Devin McDougall, Esquire
Charles McPhedran, Esquire
Mandy DeRoche, Esquire

Earthjustice
1617 John F. Kennedy Blvd., Suite 2020
Philadelphia, PA 19103
cweinberg@earthjustice.org

ATTACHMENT 1

Extracting Profits from the Public: How Utility Ratepayers Are Paying for Big Tech's Power

Eliza Martin
Ari Peskoe

March 2025

ENVIRONMENTAL & ENERGY LAW PROGRAM
HARVARD LAW SCHOOL

Extracting Profits from the Public:
How Utility Ratepayers Are Paying for Big Tech's Power

Eliza Martin and Ari Peskoe*

Executive Summary

Some of the largest companies in the world — including Amazon, Google, Meta, and Microsoft — are looking to secure electricity for their energy-intensive operations.[1] Their quests for power to supply their growing "data centers" are super-charging a growing national market for electricity service that pits regional utilities against each other. In this paper, we investigate one aspect of this competition: how utilities can fund discounts to Big Tech by socializing their costs through electricity prices charged to the public. Hiding subsidies for trillion-dollar companies in power prices increases utility profits by raising costs for American consumers.

Because for-profit utilities enjoy state-granted monopolies over electricity delivery, states must protect the public by closely regulating the prices utilities charge for service. Regulated utility rates reimburse utilities for their costs of providing service and provide an opportunity to profit on their investments in new infrastructure. This age-old formula was designed to motivate utility expansion so it would meet society's growing energy demands.

The sudden surge in electricity use by data centers — warehouses filled with power-hungry computer chips — is shifting utilities' attention away from societal needs and to the wishes of a few energy-intensive consumers. Utilities' narrow focus on expanding to serve a handful of Big Tech companies, and to a lesser extent cryptocurrency speculators, breaks the mold of traditional utility rates that are premised on spreading the costs of beneficial system expansion to all ratepayers. The very same rate structures that have socialized the costs of reliable power delivery are now forcing the public to pay for infrastructure designed to supply a handful of exceedingly wealthy corporations.

To provide data centers with power, utilities must offer rates that attract Big Tech customers and are approved by the state's public utility commission (PUC). Utilities tell PUCs what they want to hear: that the deals for Big Tech isolate data center energy costs from other ratepayers' bills and won't increase consumers' power prices. But verifying this claim is all but impossible. Attributing utility costs to a specific consumer is an imprecise exercise premised on debatable claims about utility accounting records. The subjectivity and complexity of ratemaking conceal utility attempts to funnel revenue to their competitive lines of business by overcharging captive ratepayers. While PUCs are supposed to prevent utilities

from extracting such undue profits from ratepayers, utilities' control over rate-setting processes provides them with opportunities to obscure their self-interested strategies.

Detecting wealth transfers from ratepayers to utility shareholders and Big Tech companies is particularly challenging because utilities ask PUCs for confidential treatment of their contracts with data centers, which limits scrutiny of utilities' proposed deals and narrows the scope of regulators' options when they consider utilities' prices and terms. Meanwhile, regulators face political pressure to approve major economic investments already touted by elected officials for their economic impacts. Rejecting new data center contracts could lead potential Big Tech customers to construct their facilities in other states. Indeed, Big Tech companies have repeatedly told utility regulators that unfavorable utility rates could lead them to invest elsewhere.[2]

In the following sections, we investigate how utilities are shifting the costs of data centers' electricity consumption to other ratepayers. Based on our review of nearly 50 regulatory proceedings about data centers' rates, and the long history of utilities exploiting their monopolies, we are skeptical of utility claims that data center energy costs are isolated from other consumers' bills. After describing the rate mechanisms that shift utility costs among ratepayers, we explain how both existing and new rate structures, as well as secret contracts, could be transferring Big Tech's energy costs to the public. Next, we provide recommendations to limit hidden subsidies in utility rates. Finally, we question whether utility regulators should be making policy decisions about whether to subsidize data centers and speculate on the long-term implications of utility systems dominated by trillion-dollar software and social media companies.

Table of Contents

I. Government-Set Rates Incentivize Utilities to Pursue Data Center Growth at the Expense of the Public

Data centers are large facilities packed with computer servers, networking hardware, and cooling equipment that support services like cloud computing and other data processing applications. While data centers have existed for decades, companies are now building much larger facilities. In 2023, companies began developing facilities that will consume hundreds of megawatts of power, as much as the city of Cleveland.[3] As several companies race to develop artificial intelligence (AI), the scale and energy-intensity of data center development is rapidly accelerating. By the end of 2024, companies started building gigawatt-scale data center campuses and are envisioning even larger facilities that will demand more energy than the nation's largest nuclear power plant could provide.[4]

The sudden and anticipated near-term growth of cloud computing infrastructure to accommodate the development of AI is driving a surge of utility proposals to profit from Big Tech's escalating demands. By 2030, data centers may consume as much as 12 percent of all U.S. electricity and could be largely responsible for *quintupling* the annual growth in electricity demand.[5] This growth is likely to be concentrated in regions with robust access to telecommunications infrastructure and where utilities pledge to quickly meet growing demand. Data centers could substantially expand utilities' size, both financial and physical, as they develop billions of dollars of new infrastructure for Big Tech.[6]

Data center growth is overwhelming long-standing approaches to approving utility rates. Nearly every consumer pays for electricity based on the utilities' average costs of providing service to similar ratepayers. A handful of special interests, particularly large industrial users, pay individualized rates that are negotiated with the utility and often require PUC approval. Data center growth could flip the current ratio of consumers paying general rates to special-interest customers paying unique contracts pursuant to special contracts. In this section, we summarize the potential for massive data center growth and then explore how this growth is challenging long-standing ratemaking practices and is causing the public to subsidize Big Tech's power bills.

A. Utilities Are Projecting Massive Data Center Energy Use

Industry experts and utilities are forecasting massive data center growth, and their projections keep going up. In January 2024, one industry consultancy projected 16 GW of new data center demand by 2030.[7] But by the end of the year, experts were anticipating data center growth to be as high as 65 GW by 2030.[8] Individual utilities are even more bullish. For example, Georgia Power anticipates its total energy sales will nearly double by the early 2030s, a trend it largely attributes to data centers.[9] In Texas, Oncor announced 82 gigawatts of potential data center load,[10] equivalent to the maximum demand of Texas' energy market in 2024.[11] Similarly, AEP, whose multi-state system peaks at 35 GW, expects at least 15 GW of new load from data center customers by 2030,[12] although AEP's Ohio utility added that "customers have expressed interest" in 30 GW of additional data centers in its footprint.[13]

There are reasons, however, to be skeptical of utilities' projections. Utilities have an incentive to provide optimistic projections about potential growth; these announcements are designed in part to grab investors' attention with the promise of new capital spending that will drive future profits.[14] When pressed on their projections, utilities are often reticent to disclose facility-specific details on grounds that a data center's forecasted load is proprietary information.[15] This secrecy can lead utilities and analysts to double-count a data center that requests service from multiple utilities.[16] To acquire power as quickly as possible, data center companies may be negotiating with several utilities to discover which utility can offer service first.

Technological uncertainty further complicates the forecasting challenge. Future innovation may increase or decrease data centers' electricity demand. The current surge in data center growth is traceable to the release of ChatGPT in 2022 and the subsequent burst of AI products and their associated computing needs.[17] Computational or hardware advancements might reduce AI's energy demand and diminish data center demand.[18] For instance, initial reports in January 2025 about the low energy consumption of DeepSeek, a ChatGPT competitor, fueled speculation that more efficient AI models might be just as useful while consuming far less energy. Even if more energy efficient AI models materialize, however, their lower cost could lead consumers to demand more AI services, which could drive power use even higher.[19]

Nonetheless, investment is pouring into data center growth. At a January 21, 2025 White House press conference, OpenAI headlined an announcement of $100 billion in data center investment with the possibility of an additional $400 billion over four years.[20] Earlier that month, Microsoft revealed that it would spend $80 billion on data centers in 2025, including more than $40 billion in the U.S.[21] Two weeks earlier, Amazon said it would spend $10 billion on expanding a data center in Ohio.[22] And two weeks before that, Meta announced its own $10 billion investment to build a new data center in Louisiana.[23]

While the scale and pace of data center growth is impossible to forecast precisely, we know that utilities are projecting and pursuing growth. In the next section, we explore the ratemaking and other regulatory processes that socialize utilities' costs and risks. Unlike companies that face ordinary business risks to their profitability, utilities rely on government regulators to approve their prices and can manipulate rate-setting processes to offer special deals to favored customers that shift the costs of those discounts to the public. This "hidden value transfer," a term coined by Aneil Kovvali and Joshua Macey, is a strategy employed by monopolist utilities to increase profits at the expense of their captive ratepayers.[24] Regulators are supposed to protect against hidden value transfers by aligning rates with the costs utilities incur to serve particular types of consumers. But this rate design strategy is rife with imprecision. In reality, ratepayers are paying for each other's electricity consumption, and data center growth could potentially exacerbate the cross-subsidies that are rampant in utility rates.

B. Utility Rates Socialize Power System Costs Using the "Cost Causation" Standard

The U.S. legal system bestows significant economic advantages on investor-owned utilities (IOUs), which are for-profit companies that enjoy state-granted monopolies to deliver electricity. Government-approved electricity prices reimburse utilities for their operational expenses and provide utilities an opportunity to earn a fixed rate of return on their capital investments. With a monopoly service territory and regulated prices designed to facilitate earnings growth, a utility is insulated from many ordinary business risks and shielded from competitive pressures.

Public utility regulators, or PUCs, must protect the public from a utility's monopoly power and, in the absence of competition, motivate the company to provide reliable and cost-effective service. To meet those goals, PUCs determine whether utility service is offered to all consumers within a utility's service territory at rates and conditions that are "just and reasonable."[25] This standard, enshrined in state law, requires PUCs to balance captive consumers' interests in low prices and fair terms of service against the utility's interest in maximizing returns to its shareholders. A utility rate case is the PUC's primary mechanism for weighing these competing interests by setting equitable prices for consumers that provide for the utilities' financial viability.

"Cost causation" is a guiding principle in ratemaking that dictates consumer prices should align with the costs the utility incurs to provide service to that customer or group of similar ratepayers. By approving rates that roughly meet the cost causation standard, PUCs prevent "undue discrimination" between utility ratepayers, a legal requirement that is typically specified in state law.

While the PUC makes the final decision to approve consumer prices, the utility drives the ratemaking process. In a rate case, the utility's primary goal is to collect enough money to cover its operating expenses and earn a profit on its capital investments. A utility proposes new rates by filing its accounting records and other data and analysis that form the basis of its preferred prices. Once it establishes its "revenue requirement," the utility then proposes to divide this amount among groups of consumers based on their usage patterns, infrastructure requirements, and other characteristics that the utility claims inform its costs of providing service to those consumers. Typical groups, also known as ratepayer classes, include residential, commercial, and industrial consumers. Finally, the utility proposes standardized contracts known as tariffs for each ratepayer class that include uniform charges and terms of service for each member of that ratepayer class.

Under this ratemaking process, residential ratepayers often pay the highest rates because they are distributed across wide areas, often in single-family homes that consume little energy.[26] The utility recovers the costs of building, operating, and maintaining its extensive distribution system to serve residential ratepayers by spreading those costs over the relatively small amount of energy consumed by households. By contrast, an industrial consumer uses far more energy than a household and is likely connected to the power system through higher voltage lines and needs less local infrastructure than residential ratepayers. The utility can distribute lower total infrastructure costs over far greater energy sales to generate a lower industrial rate. Properly designed rates should "produce revenues from each class of customers which match, as closely as practicable, the costs to serve each class or individual customer."[27]

But ratemaking is not "an exact science," and there is not a single correct result.[28] In a utility rate case, various parties advocate for their own self-interest by contesting the utility's filing. Consumer groups and other parties urge the PUC to reduce the utility's revenue requirement, which could potentially lower all rates. But once the revenue requirement is set, consumer groups are pitted against each other as they try to reduce their share of the total amount. Their arguments are based on competing approaches to cost causation, with each party claiming that lower rates for itself align with economic principles, fairness, and other subjective values. Well-resourced participants, such as industrial groups that have a significant incentive to argue for lower power costs, hire lawyers and analysts to comb through the utility's filings and argue that their rates should be lower.

But parties face an uphill battle challenging the utility's accounting records, engineering studies, and other evidence the utility files to justify its preferred rates. Because it initiates the rate case and generates the information needed for the PUC to approve a rate, the utility is inherently advantaged. The information asymmetry between utilities and other parties, as well as the imprecision and subjectivity of the cost causation standard, can facilitate subsidization across classes of ratepayers. We highlight three reasons that PUCs may purposefully or unwittingly approve rates that depart from the cost causation standard.

First, attributing the utilities' costs to various ratepayer classes depends on contested assumptions and disputed methodologies. Different approaches to cost allocation will yield different results. As a pioneer in public utility economics once explained, there are "notorious disagreements among the experts as to the choice of the most rational method of [ ] cost allocation — a disagreement which seems to defy resolution because of the absence of any objective standard of rationality."[29] Parties, including the utility, provide the PUC with competing analyses that are designed to meet their own objectives. For instance, industrial consumers will sponsor a study that concludes lower rates for the industrial rate class is consistent with the cost causation principle. Other parties favor their own interests in what can be a zero-sum game over how to divide the utility's revenue requirement.

Second, the PUC may have its own preferences. In most states, utility commissioners are appointed by the governor, but in ten states they are elected officials. Either commissioner may face political pressure to favor a particular ratepayer class. For instance, an elected commissioner may be inclined to provide lower rates to residential ratepayers who will vote on the commissioner's reelection. An appointed commissioner may choose to align utility rates with a governor's economic development agenda by providing lower rates to major employers, such as the commercial or industrial class. Other pressures may bias regulators in favor of other interests. As it weighs competing evidence about cost allocation provided by various parties in a rate case, the PUC has discretion to find a particular study more credible and may choose a rate structure that aligns with the sponsoring party's goals and the PUC's own preferences. While other parties may challenge a PUC's decision in court, courts are unlikely to overturn a PUC's judgment about cost allocation.[30]

Third, the utility may exploit its informational advantages and intentionally provide false information. A rate case is premised on detailed accounting records filed by the utility about the expenses it incurs to provide service. The spreadsheets and other information that the utility files are based on internal records not available to the PUC or rate-case parties. Even if the utility provides some of its records in response to a party's request, the information might be too voluminous for the PUC or other parties to verify. Ultimately, the PUC relies on the utility's good faith. However, recent cases show that utilities are filing fabricated or misleading records.[31]

A random audit of multi-state utility company FirstEnergy by the Federal Energy Regulatory Commission (FERC) found that the utility had hidden lobbying expenses tied to political corruption by mislabeling them as legitimate expenses in its accounting books. According to the audit, the utility's internal controls had been "possibly obfuscated or circumvented to conceal or mislead as to the actual amounts, nature and purpose of the lobbying expenditures."[32] The audit concluded that the utility's mislabeling allowed the inappropriate lobbying expenses to be included in rates.[33] Rate cases did not detect this deception. Only an audit, informed by an extensive federal sting operation, revealed the utility's deceit. Regulators have recently uncovered other utilities filing false or misleading information in regulated proceedings.[34]

Once the regulators approve utility rates, some consumers can shift costs to other ratepayers by fine-tuning their energy consumption. As we discuss in more detail in part II.B.3, rates for commercial and industrial ratepayers typically include demand charges that are tied to each consumer's energy consumption during the utility's or regional power system's moment of peak demand that year. By anticipating when that peak will happen and reducing consumption of utility-delivered power at that moment, a data center or other energy-intensive consumer can substantially reduce its bill. While this "peak shaving" can reduce power prices for other consumers, it also forces other ratepayers to pay part of the energy-intensive consumer's share of infrastructure costs.

Despite its flaws, ratemaking continues to be the dominant approach to financing power sector infrastructure. Uniform, stable prices provide predictable revenue that motivates investors to fund utility expansion. Rate regulation typically insulates investors from many ordinary business risks by putting ratepayers on the hook for the company's engineering, construction, or procurement mistakes. For instance, regulators often allow utilities to increase rates when their projects are over-budget. The utility rarely faces financial consequences for missteps that would cause businesses that rely on competitive markets to lose profits.

Some energy-intensive consumers can be exempted from this ratemaking process that socializes costs and shifts risks to the public. The special rates for these consumers are set in one-off agreements that can lock in long-term prices and shield it from risks faced by other ratepayers. These contracts, which typically require PUC approval, allow an individual consumer to take service under conditions and terms not otherwise available to anyone else. Special rates are, in essence, "a discriminatory action, but one that regulators can justify under certain conditions."[35]

To protect ratepayers, some state laws authorizing special contracts require PUCs to evaluate whether the contract meets the cost causation standard.[36] However, the "notorious disagreements" about how to measure whether a consumer is paying for its costs of service still plague the special-contract cost causation analysis. And, as we describe below, proceedings about special contracts present unique obstacles to evaluating cost causation.

In other states, however, laws authorizing special contracts do not prevent PUCs from approving below-cost contracts. For instance, Kansas law allows regulators to approve special rates if it determines that the rate is in the state's best interest based on multiple factors, including economic development, local employment, and tax revenues.[37] A recent law enacted in Mississippi strips utility regulators of any authority to review contracts between a utility and a data center.[38]

Regardless of the standard for reviewing special contracts, there is significant political pressure on regulators to approve these deals, even if such development results in higher electricity costs for other ratepayers. Regulators do not want to be seen as the veto point for an economic development opportunity, which may have already been publicized by the company and the governor. Because utilities may be competing for the profitable opportunity to serve a particular energy-intensive consumer, they have an incentive to offer low prices, even if that reduced rate results in higher costs for the utility's other ratepayers. As noted, despite their wealth, Big Tech companies seek low energy prices and make siting decisions based in part on price.[39] Regulatory scrutiny of special contracts is therefore a critical backstop for protecting ratepayers.

II. How Data Center Costs Creep into Ratepayers' Bills

When a utility expands its system in anticipation of growing consumer demand, it typically seeks to include the capital costs of new infrastructure in its rates. If approved, ratepayers share the costs of the utility's expansion pursuant to a cost allocation formula accepted by the PUC. This approach, while imperfect for the reasons described in the previous section, has facilitated population growth and economic development by forcing ratepayers to subsidize new infrastructure that will allow new residents and businesses to receive utility-delivered energy.

For many utilities, their expectations about growth are now dominated by new data centers. Rather than being dispersed across a utility's service territory like homes and businesses, these new data center consumers that are benefitting from utility expansion are identifiable and capable of paying for infrastructure that will directly serve their facilities. If PUCs allow utilities to follow the conventional approach of socializing system expansion, utilities will impose data centers' energy costs on the public. The easiest way for utilities to shift data centers' energy costs to the public is to simply follow long-standing practices in rate cases.

In our view, however, utilities are often using more subtle ratemaking methods to push data centers' energy costs onto consumers' bills.

In this section, we focus on three mechanisms that can force consumers to pay for data center's energy costs. First, special contracts between utilities and data centers, approved through opaque regulatory processes, are transferring data center costs to other consumers. Second, disconnected processes for setting federally regulated transmission and wholesale power rates and state-set consumer prices are: A) causing consumers to pay for interstate infrastructure needed to accommodate new data centers; B) putting consumers on the hook for new infrastructure built for data-center load that never materializes; and C) allowing data centers to strategically reduce energy usage during a few hours to reduce their bills and shift costs to other consumers. Third, data centers that bypass traditional utility ratemaking by contracting directly with power generators may also be raising electricity prices for the public. These co-location agreements between a data center and adjacent non-utility generator may trigger an increase in power market prices and distort regulated electricity delivery rates.

A. Shifting Costs through Secret Contracts

Special contracts are offered by utilities to energy-intensive consumers to attract their business. While regulators in many states are required to protect the public from such cutthroat practices that harm ratepayers, we explain in this section why we are skeptical about utility claims that special contracts for data centers do not force the public to pay for Big Tech's energy costs.

Our review of 40 state PUC proceedings about special contracts with data centers finds that regulators frequently approve special contracts in short and conclusory orders. While PUC rate case decisions are lengthy documents that engage with the evidence filed by the utilities and other parties, most PUC orders approving special contracts provide only cursory analysis of the utility's proposal. One challenge for PUCs is that few, if any, parties participate in these proceedings. As a result, the PUC has little or no evidence in the record to compete with the utility's claim that the contract isolates data center energy costs from other ratepayers' bills.

The PUC often deters parties from arguing against the utility's proposed special contract by reflexively granting utility requests to shield its proposal from public view.[40] The PUC's own grant of confidentiality adds a procedural barrier to greater participation and prevents the public from even attempting to calculate the potential costs of these deals.[41] But perhaps the greater impediment to third-party analysis of proposed special contracts is that ratepayers believe that they have little at stake in the proceedings. Unlike rate cases, which set the prices consumers pay, a special contract will only have indirect financial effects on other ratepayers if it shifts costs that the energy-intensive customer ought to pay on to other ratepayers' bills. Because meaningfully participating in a special contract case has a high cost and a generally low reward, otherwise interested parties have typically not bothered to contest them. But the scale of data center special contracts demands attention because the costs being shifted to the public could be staggering.

A special contract shifts costs to other ratepayers when the customer pays the utility a price lower than the utility's costs to serve that customer. To cover the shortfall, utilities will attempt to raise rates for other ratepayers in a subsequent rate case.[42] The amount of the shortfall, and whether there is any shortfall at all, depends on how the utility calculates its costs of providing service to the data center. As discussed above, there are "notorious disagreements" about appropriate methodologies, and even the term "cost" can itself be subject to dispute. Experts debate, for instance, when to use average or marginal costs and whether short- or long-term costs are suitable metrics. When utilities use one metric in a rate case and another metric in a special contract proceeding, they could be causing spillover effects that harm ratepayers.[43]

The disagreements about methodologies and complexities of the calculations underscore a foundational challenge to reviewing a special contract rate. As discussed above, PUC rate case decisions do not purport to assign utility costs to individual consumers but instead apportion cost responsibility among similar ratepayers grouped together as classes. But in a special contract proceeding, the utility makes the unusual claim that it can isolate its costs to serve a single consumer. Without contrary evidence filed by interested parties, the PUC may have little basis for rejecting the utility's analysis.

Even without the benefit of third-party analyses in special contract proceedings, PUC orders may summarize cross-subsidy concerns raised by their own staff. But challenging the utility's analysis is costly and time-intensive, and staff may not have the resources to provide robust analysis. Similarly, state ratepayer advocates occasionally participate in these proceedings and raise cross subsidy arguments, but they are also often stretched too thin to provide a detailed response to the utility's proposal. As a result, we find that many PUC orders approving special contracts simply conclude that the proposed contract is reasonable without meaningfully engaging with the proposal.[44]

Such PUC orders are therefore not persuasive in assuaging concerns that the public may be subsidizing Big Tech's energy costs. Moreover, as discussed, state regulators may face political pressure not to veto a significant construction project in the state. The utility's assertion that it is protecting other ratepayers may provide enough cover for regulators to approve a special contract. The obscurity and complexity of these proceedings provides utilities with opportunities to hide data center energy costs and force them onto other consumers' bills.

Recent litigation against Duke Energy, one of the largest utilities in the country, exposed that the company was acting on its incentive to shift costs of a special contract to its other ratepayers. Duke's scheme responded to a new power plant developer offering competitive contracts to supply small non-profit utilities that had been purchasing power from Duke.[45] Duke's internal documents disclosed through litigation revealed that the new company was far more efficient than Duke and the utility therefore could not compete for customers based on price. Nonetheless, Duke offered one of its larger customers a new contract that amounted to a $325 million discount compared to its existing deal with Duke.[46] Additional internal utility documents revealed that Duke developed a plan to "shift the cost of the discount" to its other ratepayers by raising their rates.[47] Duke's strategy to force its ratepayers to subsidize the special-contract customer's energy was discovered only because the power plant developer sued Duke in federal court under antitrust law.

While our paper focuses on how consumers are likely subsidizing Big Tech's energy costs through their utility rates, we acknowledge that the reverse is also theoretically possible. A data center taking service under special contracts could be *overpaying*. A utility proposing a special contract might prefer to overcharge one deep-pocketed customer through a special contract in order to reduce rates for the public. While this pricing strategy may seem politically attractive for the utility and PUC, it seems unlikely to attract new data centers.

Regardless of a utility's motivation, regulators are supposed to be skeptical of a sudden surge in utility spending. Superficial reviews of special contracts are insufficient when they are collectively committing utilities to billions of dollars for Big Tech customers. The recent Duke litigation illustrates how utilities take advantage of their monopolies to force ratepayers into subsidizing their competitive lines of businesses. Discounted rates can give a utility an edge in the data center market,[48] and hiding the costs of discounts in ratepayers' bills boosts utility profits. To prevent utilities from overcharging captive ratepayers for the benefit of their competitive businesses, both PUCs and FERC have developed regulatory mechanisms that attempt to prevent such subsidies.[49] For instance, FERC applies special scrutiny to contracts between utilities and power plants that are owned by the same corporate parent. FERC's concern is that because state regulators must let the utility recover its FERC-regulated costs in consumer's rates, "such sales could be made at a rate that is too high, which would give an undue profit to the affiliated [power plant] at the expense of the franchised public utility's captive customers."[50]

Special contracts with data centers are the latest iteration of a long-standing problem with monopolist utilities. Policing cost-shifts in this context is particularly challenging due to the opaque nature of the proceedings, the complexity and subjectivity of assessing the utility's costs of serving an a single consumer, and political pressure on PUCs to approve contracts.

B. Shifting Costs through the Gap Between Federal and State Regulation

When a PUC approves a utility's revenue requirement, it must allow the utility to include interstate transmission and wholesale power market costs that are regulated by FERC.[51] In much of the country, utilities procure power through markets administered by non-profit corporations called Regional Transmission Organizations (RTOs). Market prices are influenced by a host of factors, such as fuel and technology costs, and ultimately reflect generation supply and consumer demand. If supply is constrained by a data center demand surge, market prices would likely increase, at least in the short term. Consumers' utility bills will include these higher power market prices.

PUCs can protect ratepayers from market price increases by allocating the costs of higher prices to data centers. But PUCs rarely order utilities to adjust the formulae that spread FERC-regulated market and transmission costs to ratepayers. In this section, we illustrate how ratepayers can pay more for power due to data center demand by focusing on FERC-regulated transmission costs. Federal law provides FERC with exclusive authority to set utilities' transmission revenue requirements and allocate a utility's transmission revenue requirement to multiple utilities. Under FERC's rules, costs of a new transmission line can be paid entirely by a single utility or shared among utilities if there is agreement that the new line benefits multiple utilities. When costs are shared, a region-specific formula approved by FERC divides costs roughly in proportion to the power system benefits each utility receives, such as lower market prices and improved reliability.[52]

Under either the single-utility or multi-utility approach, PUCs apply their own formula for dividing FERC-allocated transmission costs among ratepayer classes. These separate cost allocation schemes can allow data center energy costs to creep into other consumers' bills when new data centers trigger a need for transmission upgrades. We illustrate by discussing examples of each type of transmission cost recovery and then explain how rate designs embedded in special contracts or tariffs can allow data centers to reduce their bills at the expense of ratepayers.

1. *Separate Federal and PUC Transmission Cost Allocation Methods Allow Data Center Infrastructure Costs to Infiltrate Ratepayers' Bills*

In December 2023, the PJM RTO, a utility alliance stretching from New Jersey to Chicago and south to North Carolina, approved $5 billion of transmission projects whose costs would be shared based among PJM's utility members.[53] PJM identified two factors driving the need for this transmission expansion: retirement of existing generation resources and "unprecedented data center load growth," primarily in Virginia.[54] Pursuant to its FERC-approved cost allocation method, PJM split half of the transmission costs across its footprint based on each utilities' share of regional power demand and allocated the remaining half using a computer simulation of the regional transmission network that estimates benefits each utility receives from the new transmission projects.[55] Under this approach, PJM assigned approximately half of the total cost to Virginia utilities, approximately 10% to Maryland utilities, and the remainder to utilities across the region.[56]

Each state's PUC then allocates the costs assigned by PJM to ratepayer classes of each utility it regulates. In Maryland, across the state's three IOUs assign, an average of 66 percent of transmission costs are assigned to residential ratepayers.[57] The larger of Virginia's two IOUs includes more than half of its transmission costs in residential rates.[58] Thus, in both states, residential ratepayers are paying the majority of regional transmission costs that are tied to data center growth. From the public's perspective, this result appears to violate the cost causation principle. After all, residential ratepayers are not causing PJM to plan new transmission.

PJM's approach, however, recognizes that new regional transmission benefits all ratepayers by improving reliability, allowing for more efficient delivery of power, and providing other power system improvements that are broadly shared. PJM developed its cost-sharing approach with the understanding that new transmission would be designed primarily to provide public benefits. New transmission designed for a few energy-intensive consumers, and not broad public benefits, is inconsistent with PJM's premise. That said, by increasing transmission capacity, new regional transmission lines for data centers may provide ancillary benefits to all ratepayers. PJM's power system simulation, which it uses to allocate half the costs of transmission expansion, demonstrates the shared benefits of this new infrastructure. Proponents of transmission expansion argue that such power flow models validate the current approach of allocating transmission costs to benefiting ratepayers because the models can calculate with reasonable accuracy who benefits from new transmission and therefore who should pay for it.

But even assuming that ancillary benefits for all ratepayers are adequate to justify current methods for regional transmission cost allocation, PJM only spreads costs among the region's utilities. Each utility then has its own methods, approved by PUCs, for allocating transmission investment to its ratepayers. The PUC-approved methods typically presume that ratepayers share in the benefits of new transmission in proportion to their total energy consumption. This approach causes residential ratepayers in Maryland, which consume more than half of the state's electricity, to pay for the lion's share of Maryland utilities' costs of new PJM-planned transmission. Without reforms, consumers will be paying billions of dollars for regional infrastructure that is designed to address the needs of just a few of the world's wealthiest corporations.[59]

Obsolete PUC cost allocation formulas can also cause ratepayers to pay for transmission costs that are not regionally shared. For instance, in July 2024, Virginia's largest utility applied to the PUC for permission to build infrastructure that would serve a new large data center. PUC staff reviewing the proposal found that but for the data center's request, the project "likely, if not certainly, would not be needed at this time."[60] In its application, the utility told state regulators that the $23 million project would be paid for through its FERC-approved transmission tariff.[61] Under the utility's existing state-approved tariff, about half of all costs assigned through the FERC-regulated tariff are billed to residential ratepayers, and the remaining half are billed to other existing ratepayers.[62] The bottom line is that existing tariffs force the public to foot the bill for the data center's transmission.

2. *Utilities May Be Saddling Ratepayers with Stranded Costs for Unneeded Transmission*

If a utility's data center growth projections fail to materialize, ratepayers could be left paying for transmission that the utility constructed in anticipation of data center development. Claiming that it was addressing this "stranded cost" issue, American Electric Power (AEP) of Ohio proposed a new state-regulated tariff that that would require data center customers to enter into long-term contracts with the utility before receiving service. AEP's proposed contract would require the data center to pay 90 percent of costs associated with its maximum demand for a ten-year period, including FERC-regulated transmission costs.[63] According to the utility, this upfront guarantee protects AEP's other ratepayers from the risk that the utility builds new infrastructure for a data center that never materializes and prevents the utility from offloading all of these "stranded" costs on other ratepayers.

While these long-term contracts would at least partially insulate AEP's ratepayers from data center transmission costs, neighboring utilities pointed out that they could still be left paying for stranded costs through PJM's allocation of transmission investments. Their protests explain that if AEP builds new transmission lines in anticipation of data center load growth, and those lines are paid for via PJM's regional cost allocation, then those costs would be split among all PJM-member utilities. As noted, PJM allocates half the costs of new transmission lines to its utility members based on their share of regional energy sales. If AEP's data center customers commence operations, AEP's own share of regional transmission costs would increase in proportion to its rising share of regional energy sales. In that scenario, other utilities in the region may not overpay for transmission needed for AEP's data center customers.

Protesting utilities in the Ohio PUC proceeding focus on the possibility that AEP's data center customers cancel their projects or consume less energy than anticipated after AEP has spent money developing new transmission to meet projected data center demand.[64] Under that scenario, total regional transmission costs would rise due to AEP's spending, but AEP's share of total costs would not increase proportionally. As a result, other regional utilities would face increasing costs to pay for infrastructure developed to meet AEP's unrealized data center energy demand. How much individual consumers pay for the new infrastructure would depend on how each utility allocates transmission costs to various ratepayer classes pursuant to a PUC rate case decision.

New transmission projects paid for by a single utility can also raise stranded cost concerns. In December 2024, FERC approved a contract that governed the construction of transmission facilities needed to provide service to a new data center.[65] Under the contract, the data center will immediately pay for new infrastructure needed to connect the facility to the existing transmission network but will not directly pay for necessary upgrades to existing transmission facilities. Instead, the utility AES pledged to include those upgrade costs in the transmission rates paid by all ratepayers through a subsequent regulatory process. A separate state-regulated tariff for energy-intensive consumers would require the data center, and not other consumers, to ultimately pay for the upgrades. In addition, the contract requires the data center to pay for the upgrades in the event it does not commence operations or uses less energy than would be required under the state-regulated tariff to pay for the upgrades over the time. Our understanding is that this approach to transmission cost recovery for new energy-intensive consumers is fairly common and not limited to data centers, but ratepayer advocates are concerned that data centers' commitments may be more uncertain than other types of energy-intensive consumers.

The Ohio ratepayer advocate therefore protested the contract, arguing that the language protecting other consumers from paying for the transmission upgrades was "unacceptably ambiguous."[66] The Ohio advocate urged FERC to require "specific language to preclude shifting data center costs" to other consumers.[67] FERC nonetheless approved the contract because it found that these concerns were premature and noted that they may be raised in future proceedings that directly address any proposed cost shifts.[68] In a short concurrence, FERC Commissioner Mark Christie questioned whether the rate treatment proposed by the utility that could burden consumers with stranded costs is justified.

3. *By Slightly Reducing Their Energy Use, Data Centers Can Increase Ratepayers' Transmission and Wholesale Market Charges*

Like other ratepayers, data centers pay an energy price for each unit of energy they consume as well as a monthly flat fee. Data centers, and many non-residential ratepayers, also face utility-imposed demand charges that are tied to their peak consumption during a specified month, year, or other time period. These charges are intended to reflect the costs of building power systems that have sufficient capacity to generate and deliver energy when consumer demand is unusually high. In RTO regions, PUC-regulated data center special contracts and tariffs likely reflect FERC-approved demand charges that incorporate regional transmission costs and may also include costs of procuring sufficient power plant capacity to meet peak demand. By reducing their energy use during just a few hours of the year, data centers may be able to reduce their share of regional costs that are allocated to demand charges and effectively force other ratepayers to pick up the tab.

Electricity use is constantly changing, and it peaks when consumers ramp up cooling and heating systems during exceptionally hot or cold days. Meeting these moments of peak demand is very expensive. Consumers pay for transmission and power plant infrastructure that is mostly unused but nonetheless necessary for providing power during a few peak hours each year. While utilities have employed several methods for assessing demand charges, many energy-intensive consumers are billed based on their own consumption at the moment the regional system reaches its peak demand.[69]

Data centers and other large energy users have significant incentives to forecast when this peak hour will occur and reduce their consumption of utility-delivered power during that hour. To avoid shutting down or reducing their production during hours when the system might hit its peak, energy-intensive consumers may install backup generators that displace utility-provided power. Large power users may already have their own power generators to protect against outages or improve the quality of utility-delivered power.[70] Needless to say, most consumers that face demand charges, such as small businesses, do not have a sufficient incentive to forecast the system peaks or install on-site generation. As data centers' share of regional energy consumption grows, Big Tech will be able to shift an increasingly large share of the region's costs to other ratepayers, particularly if their demand charges are easily manipulable.

PUCs can often prevent these cost shifts among consumers who take service from rate-regulated utilities in their states. Federal law requires only that the total costs allocated through FERC-approved tariffs must be passed on to utilities and then ultimately to consumers through PUC-regulated tariffs or special contracts. PUCs can choose their own methods for allocating those costs among ratepayers. Because data centers' special contracts are confidential, we often do not know whether utilities and PUCs are facilitating cost shifts through demand charges. Whether data centers are taking service under tariffs or special contracts, PUCs should ensure that rate structures are not allowing data centers to shift costs through manipulable demand charges.

That said, as we discuss below in part III.E, cutting peak consumption can reduce costs for everyone if utilities build their systems for a lower peak that accounts for a data center's ability to turn off or self-power. The problem is that utilities are expanding based on an assumption that data centers will operate at full power with utility-delivered power during peak periods. When a data center uses its own generation during peak periods to avoid demand charges, it is shifting the costs of an overbuilt system to the public.

C. Shifting Costs by "Co-Locating" Data Centers and Existing Power Plants

Power plant owners have developed their own scheme for attracting data centers that could shift energy costs from data centers to ratepayers. Under "co-location" arrangements, a data center connects directly to an existing power plant behind the plant's point of interconnection to the utility-owned transmission network. By delivering and taking power without using the transmission network, power plant owners and data centers argue that they ought to be exempt from paying utility-assessed energy delivery fees. Utilities have contested this arrangement because it denies them profitable opportunities to build new infrastructure to connect data centers to their networks.

In their haste to secure power as quickly as possible, data centers are looking to contract with existing generation, particularly nuclear power plants. By connecting directly to a power plant, data centers aim to avoid a potentially lengthy process administered by a utility to connect the data center to the utility's power delivery system. Locating load behind a power plant's point of delivery to the transmission network is not new. But the potential scale of data center growth and possibility that some significant share of that growth will co-locate has spawned disputes between power plant owners and utilities.

We highlight the key points about co-location by focusing on regulatory proceedings that involve Constellation, the largest owner of nuclear plants in the U.S., and Exelon, the largest utility in the U.S. that owns only delivery infrastructure and not power plants. Until 2022, Constellation and Exelon were housed under the same corporate parent. The company's restructuring into separate generation and delivery companies allows each of those businesses to independently pursue policies that best meet their financial interests. Data center growth began to rapidly escalate shortly thereafter and has revealed tensions between utilities and companies that compete in wholesale electricity markets for profits.

Co-location is a vague term. Because financial consequences will follow from any regulatory definition of co-location, utilities and power generators dispute how co-location technically functions. Constellation claims that because a data center co-located with one of its nuclear plants cannot receive power from the grid, it is therefore "fully isolated" from the transmission network.[71] Exelon counters that "as a matter of physics and engineering," the co-located data center is "fully integrated with the electric grid."[72] Utilities and other parties point out that a nuclear plant must operate in sync with the other plants connected to the transmission network and claim that the data center benefits from this arrangement even if the transmission system is not delivering power to it.[73]

This technical distinction could affect whether co-located entities are utility ratepayers that pay for delivery service. Constellation argues that because the utility is not delivering energy to the data center, the data center is not a utility customer, and it should not have to pay any FERC- or PUC-regulated delivery charges. Exelon opposes that result and has estimated that a single proposed co-location arrangement between a nuclear owner and a data center would shift between $58 million and $140 million of transmission and state-regulated distribution charges to other ratepayers.[74]

But Constellation and other generators dispute that calculation, claiming that this "phantom . . . 'cost shift' is, at best, merely a back-of-the-envelope estimate" of the revenue a utility would collect if the data center signed up as its customer.[75] Co-location, according to the nuclear plant owners, does not actually cause other ratepayers to pay higher transmission rates but instead precludes them from receiving lower delivery rates that they might pay when a new energy-intensive customer becomes a utility ratepayer and pays its proportional share of the utility's cost of service (a hypothetical that likely does not occur when the new customer receives a one-off price pursuant to a special contract).

But analysts are concerned that co-location can actually raise prices in interstate power markets. Across much of the country, generators are constantly competing through auction markets to supply power. In a few regions, market operators conduct separate annual, monthly, or seasonal auctions for capacity to procure sufficient resources for meeting peak consumer demand. Each power plant can offer capacity into the auction equivalent to its maximum potential for energy generation. In the PJM region, nuclear plants accounted for 21 percent of total capacity that cleared the most recent auction.[76]

PJM's independent market monitor, who fiercely promotes and defends PJM's markets, recently warned that colocation could "undermine" PJM's markets. He posited that if all nuclear plants in the region attracted co-located customers, "the impact on the PJM grid and markets would be extreme. Power flows on the grid that was built in significant part to deliver low-cost nuclear energy to load would change significantly. Energy prices would increase significantly as low-cost nuclear energy is displaced by higher cost energy . . . Capacity prices would increase as the supply of capacity to the market is reduced."[77] Should this scenario play out, the region's ratepayers could be forced to pay higher prices due to data centers' purchasing decisions. However, as noted, steep increases in demand due to data center growth could increase wholesale market prices regardless of whether data centers co-locate with existing power plants.

For utilities, opposing co-location is not purely about protecting their ratepayers or upholding the integrity of interstate markets. Co-location threatens their control over power delivery by allowing data centers to take energy directly from a large power producer. In some states, utilities might claim that state laws prohibit co-location because they provide the utility with a monopoly on retail sales.[78] Co-location would also reduce the profits that utilities would otherwise stand to gain from constructing new infrastructure to serve data centers.

In an ongoing FERC proceeding, Constellation claims that utilities' opposition to co-location is an anti-competitive ploy to capitalize on their state-granted monopolies.[79] The company alleges that co-location arrangements at two of its nuclear plants are "being held hostage by one or two monopoly utilities . . . [that] have taken the law into their own hands, and are unilaterally blocking co-location projects unless the future data center customers accede to utility demands to take [ ] transmission services . . . from the utility and sign up for retail distribution services."[80] Utilities may be trying to delay Constellation's projects until FERC provides clear guidance on co-location arrangements, including whether data centers and nuclear plants will pay any transmission charges.[81]

Even if FERC sets new rules the two sides are likely to continue squabbling about the details. With billions of dollars on the line, each side might have an incentive to litigate, which would add risk to co-location schemes.

III. Recommendations for State Regulators and Legislators: Strategies for Protecting Consumers from Big Tech's Power Costs

Without systematic changes to prevailing utility ratemaking practices, the public faces significant risks that utilities will take advantage of opportunities to profit from new data centers by making major investments and then shifting costs to their captive ratepayers. The industry's current approaches of luring data centers with discounted contracts or lopsided tariffs are unsustainable.

We outline five recommendations for PUCs to better protect consumers from subsidizing Big Tech's data centers: A) establishing guidelines for reviewing special contracts, B) shifting new data centers from special contracts to tariffs, C) facilitating competition and the development of "energy parks" that are not connected to any utility-owned network, D) requiring utilities to provide more frequent demand forecasts;, and E) allowing new data centers to take service only if they commit to flexible operations.

A. Establish Robust Guidelines for Reviewing Special Contracts

PUCs rarely reject proposed special contracts with data centers. As we discussed, many states' laws provide PUCs with broad discretion to approve special contracts, do not specify a particular standard of review, and even allow the PUC to approve a contract that shifts costs to other ratepayers. Given the unprecedented scale and pace of data center special contracts, PUCs should establish more rigorous guidelines for reviewing special contracts that are aimed at protecting consumers.

In Kentucky, the Public Service Commission must make several findings on the record before approving a special contract.[82] Under the PSC's self-imposed guidelines, special contracts that include discounts are allowed only when the utility has excess generation capacity. The guidelines limit discounts to five years and no more than half the duration of the contract. The PSC must also find that the contract rate exceeds the utility's marginal costs to serve that customer and that the contract requires the customer to pay any of the utility's fixed costs associated with providing service to that customer.

Applying its guidelines, the PSC recently rejected a utility's proposed special contract with a cryptocurrency speculator because it found the contract did not shield consumers from the crypto venture's power costs.[83] The PSC was critical of the utility's projections about regional market and transmission prices and therefore did not find credible the utility's claim that the contract would cover the utility's cost to provide energy to the crypto speculator. Industrial ratepayers, several environmental and local NGOs, and Kentucky's attorney general, acting on behalf of consumers, participated in the proceeding and criticized the proposed contract.

While the PSC's guidelines compel it to address vital consumer protection issues, the rule cannot force regulators to critically analyze the utilities' filing or prevent the PSC from merely rubber-stamping a utility's proposed special contract. Vigorous oversight cannot be mandated by law: it requires dedicated public servants. The effectiveness of any consumer protection guidelines depends on the people who implement it, including PUC staff that review utility proposals and the commissioners who make the ultimate decisions. Nonetheless, we believe that establishing guidelines that require regulators to make specific findings about a proposed special contract would improve upon the status quo.

B. Require New Data Centers to Take Service Under Tariffs

Special contracts are vehicles for shifting special interests' energy costs to consumers. Approved in confidential proceedings by PUCs facing political pressure to approve deals and often with no competing interests participating, special contracts allow utilities to take advantage of the subjectivity and complexity of their accounting practices to socialize energy-intensive customers' costs to the public. The existing guardrails that ostensibly allow regulators to police special contracts are not working to protect consumers.

Guided by their consumer-protection mandate, regulators should stop approving any special contracts and instead require utilities to serve data centers through tariffs that offer standard terms and conditions for all future data-center customers. Unlike a one-off special contract that provides each data center with unique terms and conditions, a tariff ensures that all data centers pay under the same terms and that the impact of new customers is addressed by considering the full picture of the utility's costs and revenue. This holistic and uniform approach ends the race-to-the-bottom competition that incentivizes utilities to attract customers by offering hidden discounts paid for by other ratepayers.

That said, standard tariffs are not a talisman for protecting consumers. As we have emphasized, cost allocation is an imprecise exercise that depends on myriad assumptions and projections. However, tariff proceedings and rate cases are more procedurally appropriate forums than a special contract case to consider and address cost-allocation issues. Unlike special contracts, tariffs are reviewed in open dockets that allow the public and interested parties to scrutinize proposals and understand long-term implications of proposed rates should they go into effect. Once approved, a data-center tariff can be revisited in subsequent rate cases where the utility proposes to increase rates and allocate its costs among ratepayers, including data centers. All ratepayers will have an incentive to participate in those cases and offer evidence that challenge data centers' interests.

Several utilities have already been moving away from special contracts to tariffs. Recent and ongoing proceedings are highlighting issues that demand careful scrutiny, including whether to create new data-center-only tariffs and how to protect existing ratepayers from costs of new infrastructure needed to meet data centers' demands. We briefly canvas these issues.

A threshold issue is whether an existing utility tariff for energy-intensive ratepayers is appropriate for data centers or whether a new tariff is necessary to address issues that are unique to data centers. Ratepayer classes are generally defined by the similar costs that the utility incurs to serve members of that class. Data centers may, of course, oppose new tariffs that impose more expensive prices than they would pay if they took service under existing tariffs for energy-intensive ratepayers.

In Ohio, for instance, AEP proposed to create classes for new data centers and cryptocurrency speculators and require ratepayers in those classes to commit to higher upfront charges and for a longer period of time than other energy-intensive consumers.[84] To justify the new data center class, AEP argued that data centers' unique size at individual locations and in the aggregate, as well as uncertainty about their energy use over the long-term and minimal employment opportunities, distinguish data centers from other energy-intensive consumers.[85] Data center companies responded that AEP had "failed to justify its approach to exclusively target data centers" and claimed that the utilities' costs to serve data centers was no different from other energy-intensive consumers that operate around the clock.[86] As of February 2025, the Ohio PUC has yet to rule on AEP's proposal.

FERC addressed similar issues in August 2024 when a utility proposed a new ratepayer class for energy-intensive cryptocurrency operations. Like AEP, the utility claimed that significant but uncertain demand growth justified approval of the new rate class, and therefore higher upfront payment commitments and longer terms for this new customer class were appropriate.[87] According to the utility, crypto speculators can more easily relocate their operations as compared to other energy-intensive consumers, and this mobility amplifies the risk of stranded assets built for new crypto customers that quickly set up shop elsewhere. FERC rejected the proposal because it found that the utility had provided insufficient evidence that new crypto operations "pose a greater stranded asset risk than other loads of similar size."[88] FERC's finding does not foreclose a utility from creating a crypto or data center ratepayer class, but instead signals that FERC will demand more persuasive evidence to justify approval of a new class.

State legislatures could remove any evidentiary hurdles by requiring large data centers to be in their own ratepayer class. With large data centers in their own class, regulators could more easily understand the effects data centers have on other ratepayers. For instance, parties might introduce evidence in a rate case showing how various cost allocation methods that raise costs for data centers would lower costs for other ratepayers. To avoid any claims of undue discrimination, the new rate class might include any new consumer above a specified capacity threshold that, as a practical matter, would likely capture only data centers.

Separating large data centers from other ratepayers could facilitate more protective cost allocation methods that better isolate data center costs from other ratepayers. Again, state legislatures might have a role to play. In Virginia, a bill proposed in January 2025 would require state regulators to determine whether cost allocation methods "unreasonably subsidize" data centers and to minimize or eliminate any such subsidies.[89] Such clear language would provide the PUC with guidance as it balances its obligations to protect ratepayers and facilitate growth in the state. In addition, it would force PUCs to revisit decades-old methods for dividing FERC-regulated transmission costs, as we discuss above.

As data centers shift to new tariffs, the largest potential cost shift in many states could be from the costs of new power plants built to meet data center growth. In most states, utilities are the dominant generation owners and can earn a PUC-set rate of return that they collect from ratepayers on their investments in new power plants. In general, utility expenses on new power plants are spread among ratepayer classes under the theory that all ratepayers benefit from the utility's power plants. But the staggering power demands of data centers defy this assumption. Recent tariff proceedings highlight that many utilities are proposing schemes that are not adequately shielding ratepayers from the costs of new generation for data center growth.

In Indiana, the utility Indiana Michigan Power expects new data centers to increase the peak demand on its system from 2,800 to 7,000 megawatts.[90] To facilitate this growth, the utility proposed to create special terms for new customers that demand at least 150 megawatts of power, a threshold that in practice limits their applicability to new data centers.[91] Like AEP Ohio's proposal, the updated tariff would require a new data center to commit to paying 90 percent of the utility's costs of new generation and transmission capacity needed to meet the data center's demand.[92] This 90 percent capacity payment and the tariff's twenty-year term, according to the utility, would "provide reasonable assurance" that data centers' payments to the utility "will reasonably align with the cost of the significant investments and financial commitments the Company will make to provide service."[93]

Consumer advocates generally supported the utility's efforts to insulate ratepayers from data centers' energy costs but argued that the proposed terms were "insufficient for protecting existing customers from large potential cost shifts in the event of the closure" of a large data center.[94] One of their solutions was to "firewall" the costs of new power plants built to meet data center growth from other ratepayers by requiring the utility to separately procure or build generation for data centers, and then allocating all costs solely to data centers.[95] Consumer advocates also urged regulators to require other modifications related to contract termination and other provisions to protect ratepayers from stranded costs if data center growth failed to materialize or decreased following an initial spike.[96]

Data center companies argued the other side, claiming that the terms were too onerous and benefited the utility shareholders who "would be shielded from business risk, while reaping regulated returns on large potentially more risky expansion of rate base" that would be backed by data centers.[97] Amazon observed that the utility's proposed twenty-year term is based on the ordinary approach to cost recovery of utility capital investments. But instead of the utility building its own plants and earning a return on them, Amazon claimed that the utility could more efficiently support data center growth through short-term contracts with non-utility generators or purchases via PJM's regional markets.[98] Amazon argued that rather than "imposing virtually all risks" associated with power plant development on data centers and reaping all of the profits for itself, the utility should instead share the risks of infrastructure development with new data centers.[99]

The Indiana proceeding highlights how utility ownership of generation can exacerbate cost shifts that benefit utility shareholders. The traditional utility business model of decades-long cost recovery of new utility-owned power plants through consumer rates is not designed to address a near-term tripling of a utility's demand due to just a few giant energy-guzzling warehouses. While "firewalling" data centers' power plant costs from other ratepayers is a viable approach, regulators must ensure that utility proposals actually protect consumers.

Under its "Clean Transition Tariff," Nevada Energy claims to insulate other ratepayers from data centers' energy generation costs by contracting with new clean energy resources and then passing those contract costs directly to a specific data center or other customer. In theory, this arrangement could isolate generation costs, but public utility staff and other intervenors concluded that the new tariff would not actually firewall data centers' generation costs from other ratepayers.[100] They found that complex interactions between the new tariff's proposed pricing structure and existing tariffs would shift costs to other ratepayers. For instance, PUC staff focused on the utility's proposal to account for the revenue it would have earned if the data center took service under a standard tariff and then charge other ratepayers for a portion of its “lost” revenue.[101] In February 2025, the utility agreed with intervenors to modify its proposal and defer consideration of some of these complicated cost allocation issues.[102]

A better option for protecting ratepayers from power plant costs would be to allow data centers to purchase energy directly from non-utility retailers but still pay the utility for delivery service. Several states allow for such retail competition for energy-intensive consumers. To even further isolate data center energy costs, regulators could cut the cord entirely between the utility and data centers. Off-the-grid energy parks or energy parks that only export energy to the utility could completely insulate ratepayers from data centers’ energy costs.

C. Amend State Law to Require Retail Competition and Allow for Energy Parks

Competition can protect consumers from utility market power and insulate ratepayers from cost shifts. Starting in the 1970s, a few states began to allow limited competition for electricity service to certain energy-intensive consumers.[103] In the 1990s, about a dozen states permitted all ratepayers to shop for power supply while continuing to require them to pay state-regulated rates for utility-provided delivery service. Additional states allowed energy-intensive consumers to similarly choose a power supplier. To protect ratepayers, states could require new data centers to procure power through competitive processes rather than confining them to utility-supplied power. States could go further and allow or require new data centers to isolate entirely from the utility-owned network by creating new energy parks.

A mandate that new data centers procure power from non-utility suppliers would protect ratepayers from short-term costs and long-term risks. Requiring the data center to contract with a competitive supplier rather than with the utility would ensure that all stranded costs associated with the generation are allocated between the data center and its supplier. In addition, isolating the utility from the deal would obviate the need for the type of complex energy price calculations, integral to Nevada Energy’s proposal, that link the data center’s power price to the costs of the utility’s legacy assets.

The costs of utility-built power plants for data centers could be astronomical. In the Indiana proceeding discussed in the previous section, the utility’s own estimates revealed that if it met data center demand with self-built plants it could spend as much as $17 billion on new power plants over the next several years.[104] The utility’s proposal to require data centers to commit to paying 90 percent of the infrastructure costs over a twenty-year period would improve upon the status quo but would not completely isolate those costs from other ratepayers, particularly if data center demand did not meet the utility's forecasts.

Even with a state prohibition on new utility power plants for meeting data center demand, ratepayers could still face higher bills from cost shifts. A data center procuring energy from the market would still pay utility-imposed delivery charges that could obscure discounts for data centers or include various other cost shifts. Islanding the data center and its power supply from the utility-owned system is a sure-fire approach for protecting ratepayers.

An energy park, according to a recent paper by Energy Innovation, "combines generation assets, complementary resources like storage, and connected customers."[105] Unlike typical behind-the-meter arrangements where a customer installs some on-site generation to complement utility-delivered power, an energy park would provide sufficient power for the connected customers' operations. This arrangement is "particularly compelling for large customers due to the cost advantages of sourcing electricity directly from the cheapest, cleanest sources and due to the challenges of connecting large capacities to the existing grid."[106] Avoiding the protracted utility-run interconnection processes would be a benefit for Big Tech companies who tend to move faster than the lumbering utility industry.[107]

A fool-proof way to insulate utility ratepayers from data center energy costs is to isolate a data center energy park from the utility-owned network. Isolation may be difficult, however, as an interconnected energy park could be more financially attractive to developers, even if it is only able to export power to the transmission system and unable to import utility-delivered power.[108] Connecting an energy park would require a utility-run interconnection process and would likely lead to the utility imposing transmission charges on the energy park. While transmission charges associated with an export-only energy park could facilitate cost shifts, they are likely to be much smaller than those embedded in special contracts and other arrangements for serving data centers with utility-delivered power that we have outlined in this paper.

Both competitive generation and energy park development face the same legal obstacle: state protection of utility monopolies. Under many states' laws, an entity that delivers or sells power to another entity is a "public utility." For instance, if a generation company owns the park's generation assets and Big Tech company owns the data center, the generation company would be regulated as a public utility. This designation could doom the project. States typically prohibit competition for electric service and regulators and courts might enforce the state's monopoly protections by prohibiting a multi-owner energy park located within the territory assigned to the incumbent utility.[109] Even if a state allows the energy park to move forward as a public utility, the PUC may be compelled to regulate its rates and terms of service in a way that render the project unviable.

One potential workaround is to locate an energy park outside a for-profit utility's service territory. But states' laws may nonetheless impose obstacles. In Georgia, for instance, state law allows a new energy-intensive consumer located outside existing utility service territories to choose a supplier but limits the premises to a single customer.[110] An energy park in Georgia could therefore include only one data center owner. Energy parks might also be able to locate within the service territory of a municipal or cooperative utility. The service territories of these non-profit entities may not be protected by state law, or they may not be financially motivated to defend their monopolies and might instead welcome an energy park's investment in their communities.[111] That said, some non-profit utilities may regard an energy park as an infringement on their monopolies.[112]

State legislatures could amend anachronistic laws that prevent energy park development and block data centers taking utility service from procuring non-utility generation. To avoid interminable utility complaints that competition harms consumers,[113] laws could be tailored to apply only to data centers or other energy-intensive consumers that would otherwise require a utility to incur significant costs to procure power or build new generation.

D. Require Utilities to Disclose Data Center Forecasts

For competition to be effective, market participants need information about potential data centers' location and power demands. When utilities withhold that information, they prevent generators and other infrastructure and technology developers from offering data centers solutions that compete with the utility's offering. PUCs could require utilities to file monthly or quarterly load forecasts, which would reduce utilities' informational advantages and better enable other companies to offer solutions that would protect ratepayers from a utility's ability to shift data centers' costs to other consumers.

In the AEP Ohio proceeding, a trade association representing non-utility companies that sell electricity to consumers uncovered that AEP was withholding information. It documented that the utility's demand forecasts it filed in prior proceedings were inconsistent with its projections about data center growth it revealed to justify its data center tariff proposal.[114] The trade association's analyst explained that by holding back information AEP "conferred a *de facto* competitive advantage to build transmission rather than allowing a market response from competitive merchant generation" to meet data center demand.[115] The analyst also conjectured that AEP's concealment might directly harm ratepayers if it delayed development of generation that might be needed to meet growing regional demand, which could lead to increased prices in PJM's capacity auction.[116]

PUCs can order utilities to provide demand projections more frequently and specify that utilities include new energy-intensive consumers at various stages of development. Utilities could also provide potential locations and demands of new energy-intensive consumers with enough specificity to be useful to market participants but sufficiently obscured to protect consumers' potentially confidential business information. Because many utilities have substantially increased their demand forecasts over the past year,[117] new reporting rules would be well justified as a means of protecting consumers, enabling competition, and ensuring reliability.

E. Allow New Data Centers to Take Service Only if They Commit to Flexible Operations that Can Reduce System Costs

State regulators could require utilities to condition service to new data centers on a commitment to flexible operations. This approach could benefit all ratepayers by avoiding or reducing the need for expensive infrastructure that would otherwise be needed when a new data center increases the utility's maximum demand. A study by researchers at the Nicholas Institute for Energy, Environment & Sustainability estimates that 76 GW of data centers could connect to the system if utilities curtail energy delivery for just a few hours per year.[118]

As discussed above, utilities and RTOs plan power system expansion to provide sufficient capacity for meeting consumers' maximum energy demand, which usually occurs on the hottest and coldest days of the year. Because the system is planned for these extreme weather days, a large portion of a power system's generation and delivery infrastructure is underutilized for most of the year. If a data center commits to reducing its consumption of utility-supplied power during peak demand periods, utilities could deliver power to the data center without building new infrastructure.

To implement a flexibility mandate, PUCs could order utilities to modify their tariffs and classify data center loads as interruptible customers whose power can be turned off under specified circumstances. Similarly, regulators could also require utilities to modify their interconnection procedures to designate data centers as controllable loads that must reduce their consumption under certain conditions.[119] These strategies could defer the immediate need for costly infrastructure upgrades to serve new data centers. Utilities, however, have historically been hostile to regulatory attempts to require measures that would defer or avoid the need for costly infrastructure upgrades that drive utilities' profits.

IV. Subsidies Hidden in Utility Rates Extract Value from the Public

Utility rates have always been a means of achieving economic and energy policy goals. By financing favored investments through utility rates, rather than through general government revenue, policymakers can avoid having to raise taxes and instead conceal public spending through complex utility rate increases. From the public's perspective, hiding subsidies in utility rates may be acceptable if the benefits of the favored investments exceed their costs. For data centers deals, however, utilities do not publicly demonstrate that ratepayers pay lower rates as a result of the contract. To the extent data center development offers other benefits, such as expanding the local economy or advancing national security interests, we argue that these secondary effects are either already accounted for through other policies or irrelevant to utility regulators.

The economic harm to ratepayers from data center discounts extends beyond the short-term bill increases that utilities are imposing on the public. We are concerned that meeting data center demand is delaying opportunities to initiate power sector reforms that would benefit all ratepayers. To power new data centers, utilities are proposing more of the same: spending capital on large central-station power plants and transmission reinforcements. These types of projects have been fueling utility profits for generations, but the power sector today can do so much more. Deploying advanced technologies and adopting new operational and planning practices could squeeze more value from existing utility systems, but these low-capital-cost solutions are not profitable for utilities and therefore not pursued.[120] By approving special contracts for data centers and tariffs that do protect ratepayers from Big Tech's energy costs, PUCs may be inadvertently fostering an alliance between utilities and Big Tech that could reinforce the industry's technological status quo.

A. Data Center Subsidies Fail Traditional Benefit-Cost Tests

When a utility spends money to supply a new data center, the data center should pay for those investments. However, if ratepayers ultimately benefit from new infrastructure needed for a data center, it may be reasonable for the utility to charge ratepayers a portion of the costs. The "beneficiary pays" principle, an analogue of the cost causation standard, justifies short-term bill increases when they are offset by longer term benefits that reduce ratepayers' bills. Just as consumers should pay costs that reflect a utility's cost to serve them, a utility may charge consumers for projects that ultimately lower their rates.

PUCs have applied the beneficiary pays approach in numerous contexts. For example, many states fund energy efficiency programs through utility rates. These programs directly benefit the ratepayers that make use of the program's discounts for energy audits, new appliances, and other interventions that can reduce power use. All ratepayers are billed for these subsidies that flow directly to a handful of individual consumers that take advantage of these benefits. PUCs approve of this spending when programs ultimately lower peak system demand or otherwise reduce power system costs more than the costs of funding the efficiency program. We acknowledge, however, that these calculations are premised on assumptions and judgments and can be as imprecise as the cost allocation exercises we critique in this paper. The best regulators can do is conduct these analyses transparently, which allows for judicial review, limits the potential for arbitrary regulatory decisions, and provides a basis for changing the policy in response to new evidence.

In special contract proceedings, utilities and PUCs offer no such transparency about data center deals. Instead, billion-dollar contracts are proposed and approved without public accounting of the costs and benefits. Given the stakes and the incentives of the parties, the burden ought to be on utilities to prove publicly that ratepayers are benefiting from these deals, or at worst are being held harmless.

Ratepayers should not be saddled with costs due to data centers' purported strategic national importance. In January 2025, the Biden administration declared that AI is "a defining technology of our era" that has a "growing relevance to national security."[121] "Building AI infrastructure in the United States on the time frame needed to ensure United States leadership over competitors," according to the Biden administration, will "prevent adversaries from gaining access to, and using, powerful future systems to the detriment of our military and national security."[122] If this frightening scenario proves true — that AI will be a privately owned global weapon — it's not clear what it has to do with utility rates.

Data center proponents also tout the economic benefits of new development, but the public is already paying for local job growth through their taxes. Apart from discounted utility rates, many data centers separately receive generous state and local subsidies that governments rationalize based on the supposed economic and employment benefits of permitting new development. Several states, for instance, offer sales tax exemptions that allow data center companies to purchase computers, cooling equipment, and other components without paying state tax. In Virginia, the exemption saved data center companies nearly a billion dollars in 2023 alone.[123] Data centers may also benefit from one-off incentive packages. Mississippi is providing an Amazon data center with nearly $300 million of workforce training and infrastructure upgrades.[124] Mississippi will also reimburse Amazon for 3.15 percent of the data center construction costs and provide tax exemptions that could be worth more than $500 million. In lieu of taxes, Amazon will pay approximately $200 million in fees to the county over five years.[125]

B. Data Center Subsidies Interfere with Needed Power Sector Reforms

The power sector needs major upgrades. Investment in new high-voltage transmission is historically low,[126] despite an acute need for new power lines that can connect consumers to cheaper and cleaner sources of energy and improve network reliability.[127] With low interconnectivity, the utility industry is siloed into regional alliances that make little engineering or economic sense. Meanwhile, utilities have been sluggishly slow to adopt monitoring, communications, and computing technologies that can improve the performance of existing high-voltage networks.[128] At the local level, utilities are failing to unlock the potential of distributed energy resources to lower prices.[129]

Data center growth provides utilities with an excuse to ignore these inefficiencies. Utilities don't have to innovate to supply Big Tech's warehouses and are instead offering to meet data center demand with transmission reinforcements and gas-fired power plants, which have been the industry's bread-and-butter for decades. Some utilities are even propping up their oldest and dirtiest power plants to meet data center demand.[130] Neither data centers nor regulators are challenging utilities to modernize their systems.

Power sector stagnation is the fault of utilities and the regulatory construct that incentivizes inefficient corporate decisions. Rate regulation enables excessive utility spending that crowds out cheaper alternative investments. Because they are monopolists, utilities do not face competition that might expose their inefficiencies. Regulated rates rarely punish utilities for inefficiencies or reward them for improving their operations through low-cost technologies. Ultimately, regulators must try to align utility performance with consumers' interests, but achieving this straightforward objective is dauntingly complex.

Data center growth now overwhelms many PUC agendas. By law, regulators must respond to utility proposals about rate increases, special contracts, infrastructure development, and other issues. Utilities' messaging to regulators and investors is that meeting data centers' growth targets is an urgent priority. The implication is that there's no time to act differently. With utilities' push for growth dominating their dockets, PUCs may find it even harder to reform inefficient utility practices and block unneeded investments. For ratepayers, beneficial projects will remain unfunded, and wasteful utility practices will persist.

As utilities wring profits from the public through special contract approvals, they may be developing a new alliance with Big Tech. Uniting utilities' influence-peddling experience with the deep pockets of Big Tech could further entrench utility control over the power sector. Utilities are already among the largest donors to state elected officials and have a century of experience navigating state legislatures and agencies to protect their monopoly control and otherwise advance their interests. A long-term partnership to push the common interests of utilities and data centers at statehouses, PUCs, and other forums could undermine reform efforts and harm ratepayers.

While energy-intensive consumers typically have a financial incentive to participate in PUC proceedings and argue for their own self-interest by opposing wasteful utility spending, we are concerned that a different scenario may play out for data centers. If utilities' growth predictions are realized, some utilities will have invested billions of dollars to serve data centers that will consume *a majority of all power* delivered by the utility. Under this scenario, the utility will be dependent on its data center customers for revenue and will need to retain them in order to justify its prior and future expansion. To prevent data center departures and attract new data center customers, utilities might continue to offer discounted rates. Rather than acting as watchdogs in PUC proceedings, data center companies may instead focus on securing more discounts. Insulated by special contract deals and favorable tariffs with friendly utilities, data center companies would focus on defending their discounts rather than disciplining the utility's spending in rate cases.

Outside of formal proceedings, utility-Big Tech alliances could amplify pro-utility political messages. Utilities have a pecuniary interest in the laws that govern PUC decisionmaking and push for changes that benefit their bottom lines. Utilities formally lobby state legislators and also pursue an array of public relations strategies to secure favorable legislative and regulatory outcomes. Big Tech has the financial capacity to significantly increase the amount of money supporting of pro-utility bills and regulatory actions.

An alternative approach — which requires data centers to power themselves outside of the utility system — sets up a formidable counterweight to utilities' monopoly power. If Big Tech is forced to power itself, it might defend against utility efforts to limit competition and return to the pro-market advocacy that characterized the Big Tech's power-sector lobbying efforts prior to the ChatGPT-inspired AI boom.

Appendix A

Big Tech Companies and Data Center Developers Testifying that Utility Prices Inform Where They Build New Facilities

- AEP Ohio Proposed Tariff Modifications, *supra* note 2, Motion to Intervene and Memorandum in Support of Sidecat, an Affiliate of Meta (Jun. 10, 2024) ("The applicable electricity rates and corresponding electric service tariffs for AEP Ohio will be a significant consideration for Meta when evaluating possible sites for new facilities, expansions at existing facilities, and otherwise operating its data center assets.").
- AEP Ohio Proposed Tariff Modifications, Direct Testimony of Brendon J. Baatz in Opposition of the Second Joint Stipulation and Recommendation, at 4 (Nov. 8, 2024) ("the terms and conditions in Schedule DCT are far more restrictive and burdensome than those imposed by investor-owned utilities in other states, which could prompt some data center customers to consider investing outside of Ohio").
- AEP Ohio Proposed Tariff Modifications, Second Supplemental Direct Testimony of Michael Fradette, on Behalf of Amazon Data Services, Inc., at 18 (Nov. 8, 2024) ("By rejecting a stipulation that unfairly discriminates against data centers, the Commission can help ensure that Ohio continues to be a leader in attracting investment from this vital industry.").
- AEP Ohio Proposed Tariff Modifications, Motion to Intervene of Data Center Coalition, at 4 (May 24, 2024) ("AEP Ohio's proposals, and potential proposals made by intervenors in the case, may have a significant impact on existing and planned data centers in AEP Ohio's service territory.").
- AEP Ohio Proposed Tariff Modifications, Direct Testimony of Brendon J. Baatz, at 11 (Oct. 18, 2024) ("If AEP Ohio's proposal is adopted, it would create an unfavorable environment for data center development in the state, potentially causing companies to reconsider their investment plans.").
- AEP Ohio Proposed Tariff Modifications, Direct Testimony of Kevin C. Higgins on behalf of The Data Center Coalition, at 7 (Oct. 18, 2024) ("If approved, the DCP tariff will adversely impact planned data center development in the Company's service territory."); *id.* at 11 ("At the same time, it is important that the Commission not take actions that would depress the growth of an important emerging industry by imposing unjust and discriminatory terms.").
- Indiana Michigan Power Proposed Tariff Modification, *supra* note 15, Direct Testimony of Kevin C. Higgins on behalf of The Data Center Coalition, at 6 (Oct. 15, 2024) ("If

approved, the IP Tariff changes could adversely impact planned data center development in the Company's service territory.").

- Indiana Michigan Power Proposed Tariff Modification, Direct Testimony of Justin B. Farr on behalf of Google, at 23 (Oct. 15, 2024) ("Modifications . . . have the potential to limit opportunities for . . . the development of shared solutions that can provide significant benefit to I&M's system by removing the financial incentive for I&M to collaborate with its customers to pursue innovative solutions to support their growth.").
- Indiana Michigan Power Proposed Tariff Modification, Direct Testimony of Michael Fradette on behalf of Amazon Data Services, Inc., at 37 (Oct. 15, 2024) ("The proposed [tariff] is not reasonable and in fact has a negative impact on Amazon's view for future investment actions within I&M's service territory. I&M has offered no reasonable justification for revising Tariff I.P. as proposed.").
- Contracts for Provision of Electric Service to a New Large Customer's Minnesota Data Center Project, Minn. Pub. Util. Comm'n Docket No. 22-572, Petition, at 28 ("The customer has made clear that the CRR Rate is critically important to its decision to select a site in Minnesota for its new data center. Without the CRR Rate, the economic feasibility of this new data center would be jeopardized.").
- In re Application of Pub. Serv. Co. of Colorado for Approval of a Non-Standard EDR Contract, Pub. Util. Comm'n of Colorado Proceeding No. 23A-0330E, Direct Testimony & Attachment of Travis Wright on behalf of Quality Technology Services, at 8 (Jun. 23, 2023) ("QTS selects its new locations extremely carefully. Electricity is one of the major costs to operating a data center, so the low EDR rate provided by Public Service, and the term of the EDR agreement, is a critical factor in determining to locate in Aurora."); *id*. at 10–11 ("Given that approximately 40 percent of the Aurora QTS Campus's operational expense will be attributable to utilities, with electric being the largest component, the cost per kWh can easily make or break a project, or drive QTS or its customers to invest resources elsewhere. The EDR ESA that we have negotiated with Public Service and are requesting approval of in this Proceeding, is a critical component of our business model for the Aurora QTS Campus."); *id*. at 16 ("Was the cost of electricity a critical consideration for QTS in deciding where to site its new operations? Yes. 40 percent of the operational cost of a data center is electricity, and this will usually be the largest line item on the budget. Additionally, this cost will continue for 40 years, and will scale the business. In contrast, real estate and development costs are one-time, up-front expenditures that are watered down as the

volume of business increases. The largest and fastest growing operations in our portfolio are in markets where electricity costs are competitive.”).

- In re Application of Ohio Power Company and New Albany Data Center, LLC for Approval of a Reasonable Arrangement, Pub. Util. Comm’n of Ohio Case No. 23-0891-EL-AEC, Joint Application, at 7 (Sep. 28, 2023) (“Without this reasonable arrangement, NADC could construct its own dedicated substation and take lower-cost service under AEP Ohio’s transmission voltage tariff – to the extent it would decide to develop its facilities in AEP Ohio’s service territory.”).
- Application of Nevada Power Company for Approval of an Energy Supply Agreement with Lumen Group, Pub. Util. Comm’n of Nev. Docket No. 19-12017, Application, Attachment A: Long Term Energy Supply Agreement White Paper, at 17 (Dec. 19, 2019) (“The ESA provides Google with important benefits . . . the blended rate provided for in the ESA is cost-effective and competitively priced compared to other available options, the fixed-price nature of the agreement provides Google with important cost-certainty into its energy expenditures . . .”).

Endnotes

---

* Eliza Martin is a Legal Fellow in the Environmental and Energy Law Program at Harvard Law School. Ari Peskoe is the Director of the Electricity Law Initiative. We thank Kent Chandler, Josh Macey, Abe Silverman, and Megan Wachspress for thoughtful feedback on our draft.

[1] *See, e.g.*, JOHN D. WILSON, ZACH ZIMMERMAN & ROB GRAMLICH, STRATEGIC INDUSTRIES SURGING: DRIVING US POWER DEMAND 8 (Grid Strategies, Dec. 2024) [hereinafter Grid Strategies Report]; Alastair Green et al., *How Data Centers and the Energy Sector Can Sate AI's Hunger for Power*, MCKINSEY & CO., ("Much of data center growth — about 70 percent — is expected to be fulfilled directed or indirectly (via cloud services, for instance) by hyperscalers by 2030"); EPRI, POWERING INTELLIGENCE: ANALYZING ARTIFICIAL INTELLIGENCE & DATA CENTER ENERGY CONSUMPTION 7 (May 2024) [hereinafter Powering Intelligence]; Jennifer Hiller & Katherine Blunt, *Inside the Audacious Plan to Reopen Three Mile Island's Nuclear Plant*, WALL ST. J. (Nov. 10, 2024), ("Analysts at Jefferies estimate Microsoft will pay between $110 and $115 per megawatt hour of electricity").

[2] *See, e.g.*, In re *Application of Ohio Power Company for New Tariffs Related to Data Centers*, Pub. Util. Comm'n of Ohio Case No. 24-508-EL-ATA [hereinafter AEP Ohio Proposed Tariff Modifications], Direct Testimony of Kevin C. Higgins on behalf of The Data Center Coalition, at 7 ("If approved, the [proposed] tariff will adversely impact planned data center development in the Company's service territory."); *id.* at 11 ("At the same time, it is important that the Commission not take actions that would depress the growth of an important emerging industry by imposing unjust and discriminatory terms."). *See* Appendix A for additional evidence.

[3] *See, e.g.*, Rich Miller, *Skybox Plans 300-Megawatt Campus South of Dallas*, DATA CENTER FRONTIER (Nov. 20, 2023); City of Cleveland, Office of Sustainability & Climate Justice (noting that the city has a 300-megawatt system).

[4] Palo Verde is the largest nuclear power station in the U.S. Its three reactors produce approximately 3.3 gigawatts. Meta announced a two-gigawatt data center development in December 2024. *See* Dan Swinhoe & Zachary Skidmore, *Meta Announces 4 Million Square Foot, 2 GW Louisiana Data Center Campus*, DATA CENTER DYNAMICS (Sep. 5, 2024).

[5] *See generally* Powering Intelligence; Alastair Green et al., *How Data Centers and the Energy Sector Can Sate AI's Hunger for Power*, MCKINSEY & CO.

[6] *See, e.g.*, Grid Strategies Report ("[A]nnual peak demand growth will average 3% per year over the next five years. While 3% growth may seem small to some, it would mean six times the planning and construction of new generation and transmission capacity.").

[7] *See* FED. ENERGY REG. COMM'N, SUMMER ENERGY MARKET & ELECTRIC RELIABILITY ASSESSMENT 46 (May 23, 2024) (showing 19 GW actual demand in 2023); Newmark, 2023 U.S. DATA CENTER MARKET OVERVIEW & MARKET CLUSTERS 7 (Jan. 2024) (projecting 35 GW in 2030); *AI is Poised to Drive 160% Increase in Data Center Power Demand*, Goldman Sachs (May 14, 2024).

[8] *See* Grid Strategies Report, at 12.

[9] *See* Georgia Power Company, Georgia Pub. Serv. Comm'n Docket No. 56002, Budget 2025: Load and Energy Forecast 2025 to 2044 (Jan. 31, 2025); Drew Kann and Zachary Hansen, *Data Centers Use Lots of Energy: Georgia Lawmakers Might Make Them Pay More*, THE ATLANTA JOURNAL CONSTITUTION (Feb. 13, 2025) (stating that Georgia Power executives stated that 80 percent of the company's forecasted electricity demand growth is due to data centers).

[10] Press Release, *Oncor Electric Delivery Company, Oncor Reports Third Quarter 2024 Results* (Nov. 6, 2024),.

[11] Robert Walton, *ERCOT Successfully Navigates Heat Wave, New Peak Demand Record*, UTILITY DIVE (Aug. 26, 2024).

[12] *See* Ethan Howland, *AEP Faces 15 GW of New Load, Driven by Amazon, Google, Other Data Centers: Interim CEO Fowke*, UTILITY DIVE (May 1, 2024); American Electric Power, 4th Quarter Earnings Presentation (Feb. 13, 2025).

[13] See, e.g., In re Application of Ohio Power Company for New Tariffs Related to Data Centers, Pub. Util. Comm'n of Ohio Case No. 24-508-EL-ATA [hereinafter AEP Ohio Proposed Tariff Modifications], Direct Testimony of Matthew S. McKenzie on Behalf of Ohio Power Company [hereinafter Ohio Power Company Testimony], at 2 (May 13, 2024)

[14] Indeed, investors are taking note. The authors have on file numerous reports from utility stock analysts that tout the potential of data center growth. Utilities' presentations to investors claim that data center growth will drive future earnings. *See, e.g.,* AEP 4th Quarter Earnings Presentation, *supra* note 13, at 13 (stating that "Load Growth Supports Financial Strength" and noting it is being driven by data centers).

[15] *See, e.g.*, In re *Verified Petition of Indiana Michigan Power Company for Approval of Modifications to its Industrial Tariff*, Indiana Util. Reg. Comm'n Cause No. 46097 [hereinafter Indiana Michigan Power Proposed Tariff Modifications], Testimony of Indiana Consumer Advocates, at 4 (Oct. 15, 2024) ("There has been a significant lack of transparency with these new loads . . . For example, with respect to new large loads coming to I&M's service territory, Google and Microsoft refused to answer CAC data requests about their anticipated load and electricity consumption, and Microsoft also refused to identify its forecasted load factor. CAC counsel reached out to counsel to these parties and requested to execute a non-disclosure agreement with each respective company so that CAC could obtain this pertinent information, but thus far, we have not received a proposed non-disclosure agreement or the confidential information."). Most of the figures in the Georgia Power filing cited at note 9 are redacted.
[16] *See, e.g.*, AEP Ohio Proposed Tariff Modifications, Ohio Power Company Testimony, *supra* note 13, at 2 ("Currently, AEP Ohio has limited ability to distinguish customers who are merely speculating on potential data center investments from customers who are willing to make long-term financial *commitments* to data center investments.") (original emphasis); *Large Loads Co-Located at General Facilities Technical Conference*, FERC Docket No. AD24-11-000, Transcript, at 26 (Aubrey Johnson, Vice-President, Systems & Resource Planning for the Midcontinent Independent System Operator explaining that "in many cases, these data centers are showing up in multiple places, so I have many members submitting loads that are all the same. So how do we have more clarity . . . to understand what the actual true load is?").
[17] *See generally* Powering Intelligence, at 7.
[18] *See, e.g.*, David Uberti, *AI Rout Sends Independent Power Stocks Stumbling*, WALL ST. J. (Jan. 27, 2025), ("DeepSeek's efficient approach have 'created panic among investors who question the sustainability of US data center and AI investments,' Guggenheim analysts wrote in a note"); JONATHAN KOOMEY, TANYA DAS & ZACHARY SCHMIDT, ELECTRICITY DEMAND GROWTH AND DATA CENTERS: A GUIDE FOR THE PERPLEXED (Bipartisan Policy Center & Koomey Analytics, Feb. 2025).
[19] The Grainger College of Engineering, *Why DeepSeek Could be Good News for Energy Consumption*, (Feb. 6, 2025); James O'Donnell, *DeepSeek Might Not be Such Good News for Energy After All*, MIT TECH. REVIEW (Jan. 31, 2025).
[20] *See* Deepa Seetharaman and Tom Dotan, *Tech Leaders Pledge Up to $500 Billion in AI Investment in the U.S.*, WALL ST. J. (Jan. 21, 2025).
[21] Jordan Novet, *Microsoft Expects to Spend $80 Billion on AI-Enabled Data Centers in Fiscal 2025*, CNBC (Jan. 3, 2025).
[22] Press Release, State of Ohio, *Governor DeWine Announces $10 Billion Investment Plan from Amazon Web Services in Greater Ohio* (Dec. 16, 2024).
[23] Dan Swinhoe & Zachary Skidmore, *Meta Announces 4 Million Sq Ft, 2 GW Louisiana Data Center*, DATA CENTER DYNAMICS (Dec. 5, 2024).
[24] *See generally* Aneil Kovvali & Joshua C. Macey, *Hidden Value Transfers in Public Utilities*, 171 PENN. L. REV. 2129 (2023).
[25] KEN COSTELLO, ALTERNATIVE RATE MECHANISMS & THEIR COMPATIBILITY WITH STATE UTILITY COMMISSION OBJECTIVES, NATIONAL REGULATORY RESEARCH INSTITUTE 2 (Apr. 2014).
[26] *See* U.S. Energy Information Administration, *Electric Power Monthly, Table 5.6.A. Average Price of Electricity to Ultimate Customers by End-Use Sector* (showing average residential, commercial, and industrial rates in each state).
[27] Alabama Elec. Co-op., Inc. v. FERC, 684 F.2d 20, 27 (D.C. Cir. 1982).
[28] Co. Interstate Gas Co. v. Fed. Power Comm'n, 324 U.S. 581, 590 (1945).
[29] JAMES C. BONBRIGHT, PRINCIPLES OF PUBLIC UTILITY RATES 338 (1961).
[30] *See, e.g.,* Off. of Consumer Counsel v. Dep't of Pub. Util. Control et al., 905 A.2d 1, 6 (Conn. 2006) ("In the specialized context of a rate case, the court may not substitute its own balance of the regulatory considerations for that of the agency, and must assure itself that the [department] has given consideration of the factors expressed in [the statute]."); Iowa-Ill. Gas & Elec. Co. v. Ill. Com. Comm'n, 19 Ill. 2d 436, 442 (Ill. 1960) (explaining that deference to the Commission is "especially appropriate in the area of fixing rates"); Farmland Ind., Inc. v. Kan. Corp. Comm'n, 37 P.3d 640, 650 (Kan. App. 2001) (providing that the Kansans Corporation Commission "has broad discretion in making decisions in rate design types of issues"); Ohio Consumers' Counsel v. Pub. Util. Comm'n, 926 N.E.2d 261, 266 (Ohio 2010) ("The lack of a governing statute telling the commission how it must design rates vests the commission with broad discretion in this area.").
[31] *See 2024 FERC Rep. on Enforcement*, FERC Docket No. AD07-13-018, at 51 (Nov. 21, 2024) ("Most audits find that public utilities recorded non-operating expenses and functional operating and maintenance expenses

in [Administrative and General] expense accounts, leading to inappropriate inclusion of such costs in revenue requirements produced by their formula rates"); *see also infra* note 34.

[32] *FirstEnergy Corp.,* FERC Docket No. FA19-1-000, Audit Report, at 48 (Feb. 4, 2022).

[33] *Id.* at 16.

[34] *See, e.g.*, *Application of Southern California Gas Company for Authority to Update its Gas Revenue Requirement and Bas Rates*, California Pub. Util. Comm'n Application 22-05-015, Decision 24-12-074, at 7 (Dec. 19, 2024) ("The decision [to use one-way balancing accounts] highlights a pattern of misclassification of costs at Sempra Utilities, where the company has charged ratepayers for lobbying, political activities, and expenses related to outside legal firms. These costs have been improperly booked as above-the-line expenses when forecasting future costs."); *Order Instituting Rulemaking*, California Pub. Util. Comm'n Rulemaking 13-11-005, Decision 22-04-034 (Apr. 7, 2022) ("As an experienced utility, SoCalGas should have known that its billing of lobbying against reach codes implicates several basic legal principles that are central to its duties to the Commission and to customers . . . Thus, aside from billing ratepayers for lobbying contrary to the intent of the Commission, SoCalGas appears on the face of the record to have misled staff about the direction of its lobbying…."). *See also 2024 FERC Rep. on Enforcement*, FERC Docket No. AD07-13-018, at 58 (Nov. 21, 2024) (summarizing that FERC audits revealed "improper application of merger-related costs; lobbying, charitable donation, membership dues, and employment discrimination settlement costs; improper labor overhead capitalization rates….").

[35] Costello, *supra* note 25, at 44. *See also Investigation into the Reasonableness of Rates & Charges of PacifiCorp*, Utah Pub. Serv. Comm'n Docket No. 99-035-10, 2000 WL 873337 (2000) ("[E]ach class of service does not pay precisely its 'share' of costs. This is true, for example, of the large customer groups, or special contract customers, according to some views of allocations.").

[36] *See, e.g.*, MINN. STAT. § 216B.162, subd.7 (2024); COLO. REV. STAT. ANN. § 40-3-104.3 (West 2018); MICH. COMP. LAWS § 460.6a(3).

[37] KAN. STAT. ANN. § 66-101i.

[38] *See* MISS. CODE ANN. § 77-3-271(3) ("A public utility may enter into a large customer supply and service agreement with a customer, which may include terms and pricing for electric service without reference to the rates or other conditions that may be established or fixed under Title 77, Chapter 3, Article 1, Mississippi Code of 1972. No approval by the commission of such agreement shall be required. With respect to such an agreement…the agreement, including any pricing or charges for electric service, shall not be subject to alteration or other modification or cancelation by the commission, for the entire term of the agreement….").

[39] *See* Appendix A.

[40] *See, e.g.*, *Application of El Paso Electric Company for an Economic Development Rate Rider for a New Data Center*, Pub. Util. Comm'n Texas Docket No. 56903, Order No. 1 (Aug. 2, 2024) (issuing standard protective order with no analysis); *Petition of Duke Energy Indiana for Approval of a Special Retail Electric Service Agreement*, Indiana Util. Reg. Comm'n Cause No. 45975, Order (Nov. 20, 2023) (granting Duke Energy's motion for confidential treatment); In re *Cheyenne Light, Fuel & Power Co. Petition for Confidential Treatment of a Contract with Mineone Wyoming Data Center LLC*, Wyoming Pub. Serv. Comm'n Docket No. 20003-238-EK-24 (Record No. 17600), Letter Order (Oct. 9, 2024) (authorizing confidential treatment); In re *Xcel Energy's Petition for Approval of Contracts for Provision of Service to a New Large Customer's Minnesota Data Center Project*, Minn. Pub. Util. Comm'n Docket No. E-002/M-22-572, Order (excising significant portions of the proposed service agreement and staff analysis because it is a "highly confidential trade secret"); *Tariff Filing of Kentucky Power Company for Approval of a Special Contract with Ebon International, LLC*, Kentucky Pub. Serv. Comm'n Case No. 2022-00387, Order (Dec. 4, 2024), at 3 (granting confidential treatment for utility filing and providing that the information "shall not be placed in the public record or made available for public inspection for five years or until further order[ed]").

[41]*See id; see also* Daniel Dassow, *University of Tennessee Professor Sues TVA for Records of Incentives to Bitcoin Miners*, KNOXVILLE NEWS SENTINEL (Oct. 29, 2024) (explaining how there was no information about the incentives that TVA gave a cryptocurrency company to build within its footprint, but that the company used 9.4 percent of all Knoxville Utilities Board electricity in 2023 while employing just thirty people).

[42] *See* Costello, *supra* note 25, at 21.

[43] *See* Peter Lazare, *Special Contracts and the Ratemaking Process*, 10 ELEC. J. 67, 68–70 (1997) (quoting a Commonwealth Edison filing that argues long-run costs are appropriate for rate cases and short-term costs are appropriate for special contract proceedings and explaining the implications of using different metrics).

[44]*See, e.g.*, In re *Application of Ohio Power Company and New Albany Data Center, LLC for Approval of a Reasonable Arrangement*, Pub. Util. Comm'n of Ohio Case No. 23-0891-EL-AEC, Order Approving the Application with Modification ("The proposed arrangement meets the burden of proof for obtaining a

reasonable arrangement under Ohio Adm. Code Chapter 4901:1-38. Furthermore, we find that the proposed arrangement, as modified by Staff, is reasonable and should be approved."). Occasionally, a state PUC applying its public interest standard will gesture at a utility's static marginal cost analysis or no-harm analysis for analytical support. *See, e.g.*, *Petition of Duke Energy Indiana for Approval of a Special Retail Electric Service Agreement*, Indiana Util. Reg. Comm'n Cause No. 45975, Order of the Commission (Apr. 24, 2024) ("In making such a determination [that the proposed agreement satisfies Indiana Code], two considerations are important: whether the rates negotiated between the utility and its customer are sufficient for the utility to cover the incremental cost of providing the service to the customer and still make some contribution to the utility's recovery of its fixed costs, and whether the utility has sufficient capacity to meet the customer's needs. As explained by [Duke Energy's Vice President of Rates and Regulatory Strategy], the Agreement requires that Customer cover the incremental costs of providing service to it, as well as contributing to Petitioner's recovery of fixed costs…Based on the evidence of record, we find and conclude that the terms and conditions contemplated in the Agreement are just and reasonable…Therefore, we find that the Agreement is in the public interest and is, therefore, approved…."); In re *Idaho Power Company's Application for Approval of a Special Contract and Tariff Schedule 33 to Provide Electric Service to Brisbie, LLC's Data Center Facility*, Idaho Pub. Util. Comm'n Case No. IPC-E-21-42, Order No. 35958 ("Commission Discussion and Findings: The Commission has jurisdiction over this matter under *Idaho Code* §§ 61-501, -502, and -503…We have reviewed the record in this case and find the Company's August 30, 2023, Filing including an amended ESA, revised Schedule 33, and additional modifications is consistent with the Commission's directive in Order No. 3577.").

[45] See Duke Energy Carolinas, LLC v. NTE Carolinas II, LLC, 111 F.4th 337, 344–46 (4th Cir. 2024).

[46] *Id.* at 347.

[47] *Id.* at 349.

[48] *See* Appendix A.

[49] *See generally* Kovvali & Macey, *supra* note 24.

[50] Cross-Subsidization Restrictions on Affiliate Transactions, 73 Fed. Reg. 11,013 (2008) (codified at 18 C.F.R. pt. 35).

[51] *See, e.g.,* Nantahala Power & Light Co. v. FERC, 476 U.S. 953 (1986).

[52] *See, e.g.*, Nat'l Ass'n of Reg. Util. Comm'rs v. FERC, 475 F.3d 1227, 1285 (D.C. Cir. 2007); Entergy Services, Inc. v. FERC, 319 F.3d 536 (D.C. Cir. 2003); South Carolina Pub. Serv. Auth. V. FERC, 762 F.3d 41 (D.C. Cir. 2014).

[53] PJM, *PJM Board of Managers Approves Critical Grid Upgrades*, PJM INSIDE LINES (Dec. 11, 2023).

[54] Sami Abdulsalam, Senior Manager, PJM Transmission Planning, Reliability Analysis Update at Transmission Expansion Advisory Committee Meeting (Dec. 5, 2023). *See also PJM Revisions to Incorporate Cost Responsibility Assignments for Regional Transmission Expansion Plan Baseline Upgrades*, FERC Docket No. ER24-843, Protest and Comments of Maryland Office of People's Counsel (Feb. 9, 2024) [hereinafter Maryland People's Counsel Protest].

[55] *See generally PJM Interconnection*, 187 FERC ¶ 61,012 at P 6 (2024); Maryland People's Counsel Protest, Affidavit of Ron Nelson, at 5.

[56] *See* Maryland People's Counsel Protest, Affidavit of Ron Nelson, at 5.

[57] *See Delmarva Power & Light Co. Modification of Retail Transmission Rates*, Maryland Pub. Serv. Comm'n Case No. 8890, Revised Tariff, Attachment E (Jul. 2, 2024) (allocating 68 percent of transmission costs to residential customers); *Potomac Electric Power Co. Modification of Retail Transmission Rates*, Maryland Pub. Serv. Comm'n Case No. 8890, Revised Tariff, Attachment F (Jul. 2, 2024) (allocating 53 percent of transmission costs to residential customers); *Baltimore Gas & Elec. Co. Updated Market-Priced Service Rates, Administrative Charges, and Retail Transmission Rates under Rider 1*, Maryland Pub. Serv. Comm'n Case Nos. 9056/9064, Attachment 2: Development of the Retail Transmission Rates (Apr. 30, 2024) (allocating 78 percent of transmission costs to residential customers).

[58] *Application of Virginia Electric and Power Co.*, Virginia Corp. Comm'n. Case No. PUR-2021-00102, Report of Chief Hearing Examiner Alexander F. Skirpan, Jr., at 9–10 (Jul. 14, 2021).

[59] The cost causation principle could require a shift from transmission rates based on average — or static marginal — costs, to dynamic marginal cost analyses. *See* In re *Application of Pub. Serv. Co. of Colorado for Approval of a Non-Standard EDR Contract*, Colorado Pub. Util. Comm'n Proceeding No. 23A-0330E, Commission Decision Denying Exceptions to Decision No. R24-0168 and Adopting Recommended Decision with Modifications, at 11–12 (May 15, 2024) ("[W]e emphasize that the Commission's review of future Non-Standard EDR contracts must entail detailed examination of how the addition of large loads to the Public Service's system may create a dynamic need for multi-billion new generation and transmission capacity investments that unpredictably show up with no meaningful notice to this Commission and may not be easily

captured in a static marginal cost analysis . . . To that end, the marginal cost analysis that Public Service applied to the EDR ESA with [the data center customer] may not be adequate in future proceedings where the Commission reviews a similar Non-Standard EDR contract especially in light of the rapidly evolving and dynamic interaction between rising demand and the potential costs of serving that growth.").
[60] *Application of Virginia Electric Power,* Virginia Corp. Comm'n. Case No. PUR-2024-00135, Report of Hearing Examiner Bryan D. Stogdale, at 47 (Feb. 14, 2025).
[61] *Application of Virginia Electric Power,* Virginia Corp. Comm'n. Case No. PUR-2024-00135, Report of Hearing Examiner Bryan D. Stogdale, at 23 (Feb. 14, 2025).
[62] *Supra* note 58.
[63] *See* AEP Ohio Proposed Tariff Modifications, Ohio Power Company Testimony, at 18–20 (May 13, 2024).
[64] *See* AEP Ohio Proposed Tariff Modifications, Prepared Direct Testimony of Dennis W. Bethel on Behalf of Buckeye Power, Inc. and American Municipal Power [hereinafter Buckeye Power Comments], at 18–19 (Aug. 29, 2024).
[65] *Dayton Power & Light Co.*, 189 FERC ¶ 61,220 (2024).
[66] *Dayton Power & Light Co.*, FERC Docket No. ER25-192, Protest of the Office of the Ohio Consumers' Counsel [hereinafter Protest of the Office of Ohio Consumers' Counsel], at 4 (Nov. 13, 2024); *Dayton Power & Light Co.*, FERC Docket No. ER25-192, Limited Comments of Buckeye Power (Nov. 21, 2024).
[67] Protest of the Office of the Ohio Consumers' Counsel at 5.
[68] *Dayton Power and Light Co.*, 189 FERC ¶ 61,220 at P 23 (2024).
[69] *PJM Interconnection and Virginia Electric and Power Company*, 169 FERC ¶ 61,041 (2019).
[70] *See, e.g.,* Walker Orenstein, *Amazon Wants to Limit Review of 250 Diesel Generators at Its Minnesota Data Center,* MINNESOTA STAR TRIBUNE (Feb. 17, 2025) (noting that Amazon wants to install 600 megawatts of on-site diesel-powered generators at its new data center).
[71] *Constellation Energy Generation v. PJM*, FERC Docket No. EL25-20, Complaint Requesting Fast Track Processing of Constellation Energy Generation, LLC [hereinafter Constellation Complaint], at 20–21 (Nov. 22, 2024).
[72] *Constellation Energy Generation v. PJM*, Docket No. EL25-20, Exelon Comments in Opposition to the Complaint, at 3 (Dec. 12, 2024) ("Constellation refers to Co-Located Load as being 'Fully Isolated' and repeats that term again and again, but it remains untrue. If the loads at issue were truly 'isolated,' the PJM Tariff would not apply to them; no FERC-jurisdictional tariff would. And there would be no reason for this proceeding. As further discussed . . . the loads — whether they are what Constellation labels 'fully isolated' or not — unavoidably rely upon and use grid facilities and grid services in multiple ways. As a matter of physics and engineering, the load is fully integrated with the electric grid — this is the opposite of 'Fully Isolated.'").
[73] *See, e.g.*, *Constellation Energy Generation v. PJM*, FERC Docket No. EL25-20, Comments of the Illinois Attorney General, at 12–13 (Dec. 12, 2024); *Large Loads Co-Located at General Facilities*, FERC Docket No. AD24-11-000, Post Technical Comments of the Organization of PJM States, Inc., at 4 (Dec. 9, 2024) (stating that "[t]ransmission customers have paid the costs of supporting the grid necessary to allow [ ] nuclear facilities to operate").
[74] *PJM Interconnection, LLC*, FERC Docket No. ER24-2172 [hereinafter Susquehanna Nuclear Interconnection Agreement], Protest of Exelon Corporation & American Electric Power Service Corporation, Declaration of John J. Reed & Danielle S. Powers, at 4 (Jun. 24, 2024).
[75] Susquehanna Nuclear Interconnection Agreement, Motion for Leave to Answer and Answer of Constellation Energy Generation and Vistra Corp., at 11 (Jul. 10, 2024).
[76] *See* PJM, *2025/2026 Base Residual Auction Report*, at 11 (2024).
[77] *See* *2024 Quarterly State of the Market Report for PJM: January Through September*, MONITORING ANALYTICS 3 (2024). *See also* Buckeye Power Comments, at 15 (Aug. 29, 2024) ("Co-location of data centers at existing multi-unit generators (nuclear plants are considered ideal) appears, at first blush, to be attractive as it can 'free-up' transmission capacity by reducing the net output of the generators that the transmission system must deliver. But co-location is a complicated scenario that can disrupt power markets and shift costs by removing large blocks of reliable base load power that will need to be replaced by other sources that will likely require transmission expansion elsewhere."); *Constellation Energy Generation v. PJM*, FERC Docket No. EL25-20, Comments of the Illinois Attorney General, at 3–4 (Dec. 12, 2024) ("The OAG's primary concern regarding co-location arrangements is the impact on resource adequacy and electricity energy and capacity prices . . . . The effect of removing the Illinois nuclear power plant capacity from the ComEd zone and from the PJM market generally can be expected to drive up prices . . . . In light of these multiple factors that are currently putting pressure on prices, co-location arrangements that reserve large blocks of power for discrete customers and prevent them from serving the grid as a whole can be expected to affect the 2027/2028 [capacity prices] . . .

. The OAG is concerned that co-location arrangements that abruptly remove large resources with high capacity values from the grid will cause further devastating price increases while the PJM markets struggle to respond.").

[78] *See infra* Section III.C.

[79] *See Constellation Energy Generation v. PJM*, FERC Docket No. EL25-20, Constellation Complaint, at 6–7 (Nov. 22, 2024) ("competition to serve data center loads [is] a threat to [utilities] bottom line").

[80] *Id.* ("Exelon's utilities already have taken the position that this Commission has decreed that Fully Isolated Co-Located Load is 'impossible' — and shut down any attempt by customers to co-locate data center load in their utility systems. As detailed in their petition for declaratory order filed in Docket No. EL24-149, Exelon is refusing to process necessary studies on these grounds, demanding expensive upgrades under their unified interconnection procedures, delaying agreed-upon work which will force a nuclear plant to take additional outages, and forcing additional services to be procured.").

[81] *See PJM Interconnection, LLC*, 190 FERC ¶ 61,115 (Feb. 20, 2025) (instituting a show cause proceeding pursuant to section 206 of the FPA, and directing PJM and the Transmission Owners to either (1) show cause as to why the Tariff "remains just and reasonable and not unduly discriminatory or preferential without provisions addressing the sufficient clarity or consistency the rates, terms, and conditions of service that apply to co-location arrangements; or (2) explain what changes to the Tariff would remedy the identified concerns if the Commission were to determine that the Tariff has in fact become unjust and unreasonable or unduly discriminatory or preferential, and therefore, proceeds to establish a replacement Tariff").

[82] *See In the Matter of: Electronic Tariff Filing of Kentucky Power Company for Approval of a Special Contract with Ebon International, LLC*, Kentucky Pub. Serv. Comm'n Case No. 2022-00387, at 2–4 (Aug. 28, 2023) (citing *Investigation into the Implementation of Economic Development Rates by Electric & Gas Utilities,* Kentucky Pub. Serv. Comm'n Admin. Case No. 327 (Sep. 24, 1990), *aff'd*, Kentucky Power Co. v. PSC of Kentucky, Franklin Circuit Court, Div. 1, Civil Action No. 23-CI-00899 (Dec. 30, 2024)).

[83] *Id.*

[84] *See* AEP Ohio Proposed Tariff Modifications, Ohio Power Company Testimony, at 2 (May 13, 2024). AEP Ohio requested PUC approval to create two new customer classifications: data centers with a monthly maximum demand of 25 MW or greater, and mobile data centers (cryptocurrency miners) with a monthly maximum demand of 1 MW or greater. AEP's proposed tariff would include new obligations for these customer classes, including a minimum demand charge of 90 percent for data centers, and 95 percent for cryptocurrency facilities, as opposed to the standard 60 percent minimum demand charge for other customers in the general service rate class. AEP Ohio would also require: the two customer classes enter into energy service agreements (ESAs) for an initial term of at least ten years, as opposed to the typical term of one to five years; requirements to pay an exit fee equal to three years of minimum charges should the customer cancel the ESA after five years; collateral requirements tied to the customer's credit ratings; requirements to reduce demand on AEP Ohio's system during an emergency event; and requirements to participate in a separate energy procurement auction than standard offer service customers

[85] *Id*. at 7–8.

[86] AEP Ohio Proposed Tariff Modifications, Initial Comments of Data Center Coalition, at 9–12 (Jun. 25. 2024).

[87] *Basin Electric Power Cooperative*, 188 FERC ¶ 61,132 at PP 15–16, 61 (2024).

[88] *Id.* at P 95.

[89] *See* H.B. 2101, 2025 Gen. Assemb., Reg. Sess. (Va. 2025).

[90] *See* Indiana Michigan Power Proposed Tariff Modifications, *supra* note 15, Direct Testimony of Andrew J. Williamson on Behalf of Indiana Michigan Power Company, at 5 (Jul. 19, 2024).

[91] *Id*. at 3, 6–7.

[92] *Id.* at 14.

[93] *Id.*; *id.* at 16 (tariff terms ensure data center provides "reasonable financial support for the significant transmission and generation infrastructure needed to serve large loads").

[94] Indiana Michigan Power Proposed Tariff Modifications, *supra* note 15, Direct Testimony of Benjamin Inskeep on Behalf of Citizens Action Coalition of Indiana, Inc. [hereinafter Citizens Action Coalition of Indiana Testimony], at 25 (Oct. 15, 2024).

[95] *Id.* at 36.

[96] *Id.* at 24–31.

[97] Indiana Michigan Power Proposed Tariff Modifications, *supra* note 15, Direct Testimony of Carolyn A. Berry on Behalf of Amazon Web Services, at 16 (Oct. 15, 2024).

[98] *Id.*

[99] *Id.*

[100] *See generally Application of Nevada Power Company to Implement Clean Transition Tariff Schedule*, Nevada Pub. Util. Comm'n Docket No. 24-05023 [Nevada Power Clean Transition Tariff], Direct Testimony of Manuel N. Lopez on Behalf of Regulatory Operations Staff (Jan. 16, 2025); Nevada Power Clean Transition Tariff, Direct Testimony of Jeremy I. Fisher on Behalf of Sierra Club, Docket No. PUCN 24-05023, at 10–20 (Jan. 16, 2025).
[101] *See generally* Nevada Power Clean Transition Tariff, Direct Testimony of Manuel N. Lopez on Behalf of Regulatory Operations Staff, at 7–8 (Jan. 16, 2025).
[102] Nevada Power Clean Transition Tariff, Stipulation (Feb. 7, 2025).
[103] *See, e.g.,* GA. CODE ANN. § 46-3-8 (allowing utilities to compete to provide service to certain new customers demanding at least 900 kilowatts).
[104] *See* Indiana Michigan Power Proposed Tariff Modifications, *supra* note 15, Citizens Action Coalition of Indiana Testimony, at 11 (Oct. 15, 2024) ("Using I&M witness Williamson's example portfolio that has an average resource cost of $2,000/kW and has an average accredited capacity of 50%, I&M will also need to make $17.6 billion in new generation investments to serve 4.4 GW of new hyperscaler load.").
[105] ERIC GIMON, MARK AHLSTROM & MIKE O'BOYLE, ENERGY PARKS: A NEW STRATEGY TO MEET RISING ELECTRICITY DEMAND 7 (Energy Innovation Policy & Technology, Dec. 2024).
[106] *Id*. at 8.
[107] *See id.* at 19.
[108] *See id.* at 8–21.
[109] *See, e.g.,* State ex rel. Utilities Commission v. North Carolina Waste Awareness and Reduction Network, 805 S.E.2d 712 (N.C. Ct. App. 2017), *aff'd per curiam*, 371 N.C. 109, 617 (2018).
[110] *See* Sawnee Electric Membership Corporation v. Public Service Comm'n, 371 Ga. App. 267, 270 (2024) (". . . [T]he text of the Act assigns each geographic area to an electric supplier but also includes the large load exception to allow customers to choose their electric supplier if certain conditions exist . . . the premises must be 'utilized by one consumer and have single-metered service").
[111] *See generally* David Roberts, *Assembling Diverse Resources Into Super-Powered "Energy Parks:" A Conversation with Eric Gimon of Energy Innovation*, VOLTS (Jan. 15, 2025) (featuring an Energy Innovation author describing energy parks in rural cooperative territory in Texas).
[112] *See, e.g.,* Paoli Mun. Light Dept. v. Orange County Rural Elec. Membership Corp., 904 N.E.2d 1280 (Ind. Ct. App. 2009) (ruling in favor of a cooperative utility that sued to prevent a municipal utility from providing electric service to a facility owned by that municipality but located within the cooperative's service territory).
[113] *See, e.g.,* Power for Tomorrow (last visited Jan. 29, 2025), which claims to be "the nation's leading resource" about the "regulated electric utility model" and generally opposes competition with utilities, in part by claiming that competition harms residential consumers. The effort is funded by utilities. *See* Energy and Policy Institute, *Power for Tomorrow* (last visited Jan. 29, 2025).
[114] AEP Ohio Proposed Tariff Modifications, Testimony of Paul Sotkiewicz on Behalf of the Retail Energy Supply Association, at 9–10 (Aug. 29, 2024).
[115] *Id.* at 15.
[116] *Id.* at 14–15.
[117] The trade group's analyst observed that in January 2023 AEP projected only 248 megawatts of data center growth through 2038, but one year later AEP projected 3,700 megawatts of data center growth by 2030. *Id.* at 10 (citing PJM reports).
[118] TYLER NORRIS ET AL., RETHINKING LOAD GROWTH: ASSESSING THE POTENTIAL FOR INTEGRATION OF LARGE FLEXIBLE LOADS IN U.S. POWER SYSTEMS 18 (Nicholas Institute for Energy, Environment & Sustainability, 2025).
[119] *Id*. at 5–6.
[120] *See* Ari Peskoe, *Replacing the Utility Transmission Syndicate's Control*, 44 ENERGY L. J. 547 (2023).
[121] Exec. Order No. 14,141, 90 FR 5469 (2025).
[122] *Id*.
[123] Va. J. Legis. Audit & Rev. Commission 2024-548, Report to the Governor & the General Assembly of Virginia: Data Centers in Virginia, at viii (2024).
[124] Brody Ford & Matt Day, *Price Tag Jumps for Amazon's Mississippi Data Centers Jump 60% to $16 Billion*, BLOOMBERG (Jan. 31, 2025).
[125] *Id.*
[126] *See generally* NATHAN SHREVE, ZACHARY ZIMMERMAN & ROB GRAMLICH, FEWER NEW MILES: THE US TRANSMISSION GRID IN THE 2020S, GRID STRATEGIES (Jul. 2024).
[127] U.S. Department of Energy, *National Transmission Needs Study* (Oct. 30, 2023).
[128] *See* Ari Peskoe, *Replacing the Utility Transmission Syndicate's Control*, 44 ENERGY L. J. 547 (2023)

[129] Sonali Razdan, Jennifer Downing & Louise White, *Pathways to Commercial Liftoff: Virtual Power Plants 2025 Update*, U.S. Department of Energy Loan Programs Office (Jan. 2025).
[130] *See, e.g,* Mississippi Power Company's Notice of IRP Cycle, Mississippi Public Service Comm'n Docket No. 2019-UA-231 (Jan. 9, 2025) (stating that because the utility has entered into two contracts with 600 MW of new load it will keep at least one coal plant open that had been slated for retirement); Mississippi Power Special Contract Filing, Mississippi Public Service Comm'n Docket No. 2025-UN-3 (Jan. 9, 2025) (showing that at least one of the two special contracts is with a data center).

ATTACHMENT 2

Bipartisan Policy Center

KOOMEY ANALYTICS

Electricity Demand Growth and Data Centers: A Guide for the Perplexed

February 2025

AUTHORS

Jonathan Koomey, Ph.D.
President
Koomey Analytics

Zachary Schmidt
Senior Researcher
Koomey Analytics

Tanya Das, Ph.D.
Director of AI and Energy Technology Policy
Bipartisan Policy Center

ACKNOWLEDGEMENTS

The authors extend their gratitude to NVIDIA for providing funding to write this report. BPC also thanks Lesley Jantarasami (BPC), Lori J. Pickford (BPC Action), and Kim Dean (BPC Action) for their contributions to the report. Dr. Koomey acknowledges ongoing funding on other projects related to understanding data center electricity use from the Lawrence Berkeley National Laboratory, AMD, Amazon, Cadence Design Systems, and Microsoft.

The Bipartisan Policy Center is a mission-focused organization helping policymakers work across party lines to craft bipartisan solutions. By connecting Republicans and Democrats, delivering data and context, negotiating public policy, and creating space for bipartisan collaboration, BPC helps turn legislators' best ideas into durable laws that improve lives. Since 2007, the Bipartisan Policy Center has helped shepherd countless bills across the finish line.

Executive Summary

Reports of unprecedented and “explosive” growth in electricity demand from data centers—facilities that host websites, store data, stream media, mine cryptocurrency, and train artificial intelligence models—have appeared in many major news outlets. These headlines encapsulate two widely expressed concerns:

- First, that rising energy demand from data centers could further overburden aging power infrastructure.
- Second, that this new source of demand could jeopardize efforts to mitigate climate change.

Focusing on empirical data, this report explores the accuracy of narratives that data centers are behind exploding demand for energy. It finds no evidence that this is true on the national level in recent years—indeed, the data show no rapid growth of energy demand—but we did find regional variations. It also shows that utility forecasts in 2023 were much higher than forecasts from prior years, demonstrating that utilities expect demand for electricity to grow in the years ahead.

This report also explains the key drivers of load growth for data centers, focusing on growth in computing services and improvements in efficiency. Both these drivers are subject to deep uncertainty. The report puts data centers’ projected growth into perspective, comparing it to load growth in a high electrification scenario. This comparison shows that even for an extremely high projection of data center electricity use, data centers are likely to be only one of several important contributors to load growth in the years ahead.

Authors' Note: During publication of this report, Chinese AI startup DeepSeek released a generative AI model that reportedly operates more efficiently and cost significantly less to develop than similar models from OpenAI, Google, and others. While the implications of this announcement are still being analyzed, this milestone underscores a key finding of this report: Innovations in AI systems including software, algorithms, and training methods could lead to substantial efficiency gains that reduce future electricity demand associated with AI technology compared to previous expectations.

Table of Contents

Introduction

Reports of unprecedented and "explosive" growth in electricity demand from data centers—facilities that host websites, store data, stream media, mine cryptocurrency, and train AI models—have appeared in many major news outlets [1, 2, 3, 4, 5, 6, 7, 8]. Sample headlines conveying a sense of alarm include "Amid explosive demand, America is running out of power" (*The Washington Post*); "A New Surge in Power Use Is Threatening U.S. Climate Goals" (*The New York Times*); and "Booming AI demand threatens global electricity supply" *(Financial Times)*.

These headlines encapsulate two widely expressed concerns:

- First, that rising energy demand from data centers could further overburden aging power infrastructure.
- Second, that this new source of demand could jeopardize efforts to mitigate climate change.

Such concerns have been amplified in recent reports by influential management consulting and investment advising firms [9, 10, 11, 12, 13, 14, 15]. Together, they suggest that a power crisis could be looming, and that AI's electricity consumption could be a national issue [16].

This report examines recent forecasts of demand growth in the U.S. electricity sector and claims about the potential impact of AI-related power consumption. How important will data-center-driven growth be compared with other changing sources of demand, such as the addition of new industries or the electrification of home heating and transportation? If growing demand is a problem, is it a national, regional, or local one?

The goal of this report is not to dismiss concerns about grid resilience, or to downplay the challenges of reconciling expected growth in electricity demand with simultaneous efforts to decarbonize the power sector, but to bring empirical data to bear on questions about the role of data centers (particularly AI data centers) in load growth. Such a critical analysis is long overdue, both to improve policymaking and to develop effective strategies for managing rapidly evolving demands on the U.S. electricity grid.

Understanding U.S. electricity sector load growth: A look at the data

Changes in the composition and geographic distribution of different industries and economic activities have always influenced U.S. electricity consumption, as have a host of other hard-to-predict factors, from technological developments and behavioral changes to short-term weather variations. This section explores the evidence for recent electricity load growth in the United States, starting with national-level data, followed by an assessment of trends in two states that are projecting higher than average load growth.

Analysis of national U.S. electricity demand shows that there is no evidence for explosive demand growth in recent years, in the annual, monthly, or sectoral data (see Appendix A for details), although some states have seen demand growth. For example, national electricity use was lower in 2023 than in 2022, as shown in **Figure 1**, as was the commercial sector's electricity demand.

Figure 1. Historical U.S. electricity use, primary energy use, and economic output (GDP)

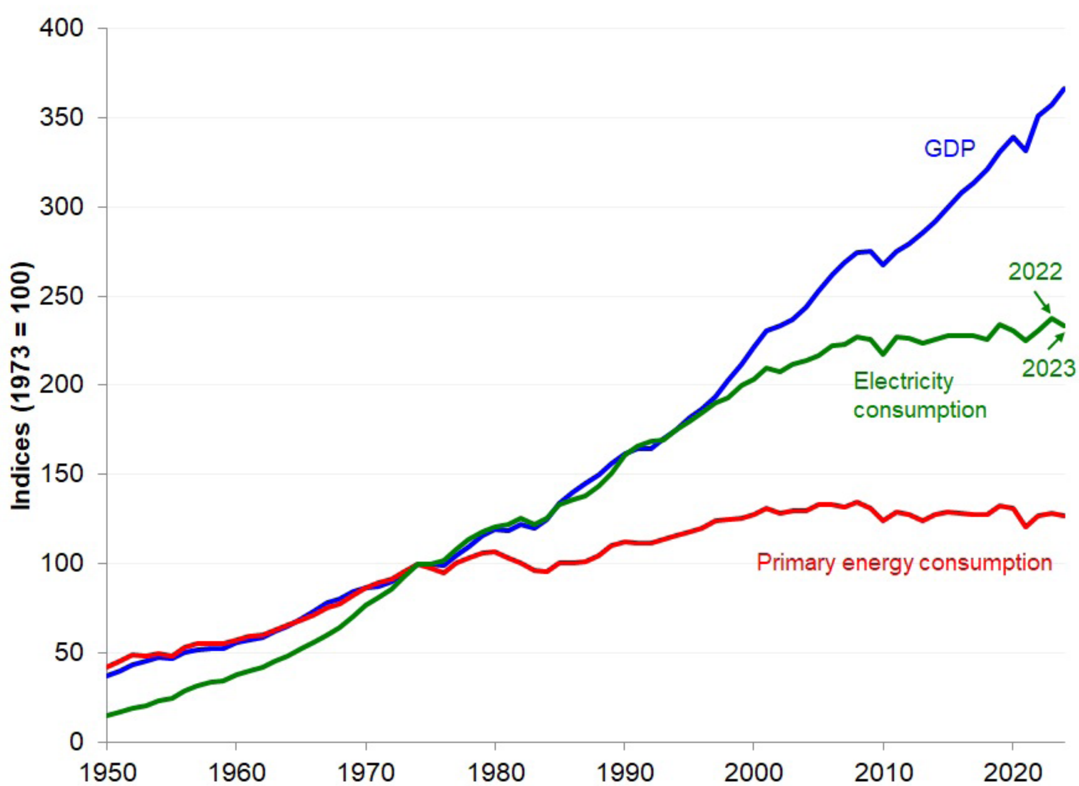

Source: https://www.eia.gov/opendata/

To understand the relationship between historical electricity use trends and projections, we used data from the Federal Energy Regulatory Commission's Form 714, which compiles utility sector forecasts and historical generation every year for many utilities and most balancing authorities. Although no exhaustive database of forecasts covers the whole U.S., we compiled data for 14 utilities and balancing authorities representing about three-quarters of U.S. generation in 2023 (see Appendix B for a complete list). **Figure 2** shows the results.

Figure 2. Aggregate electricity generation (actual and projected) for a large sample of U.S. utilities

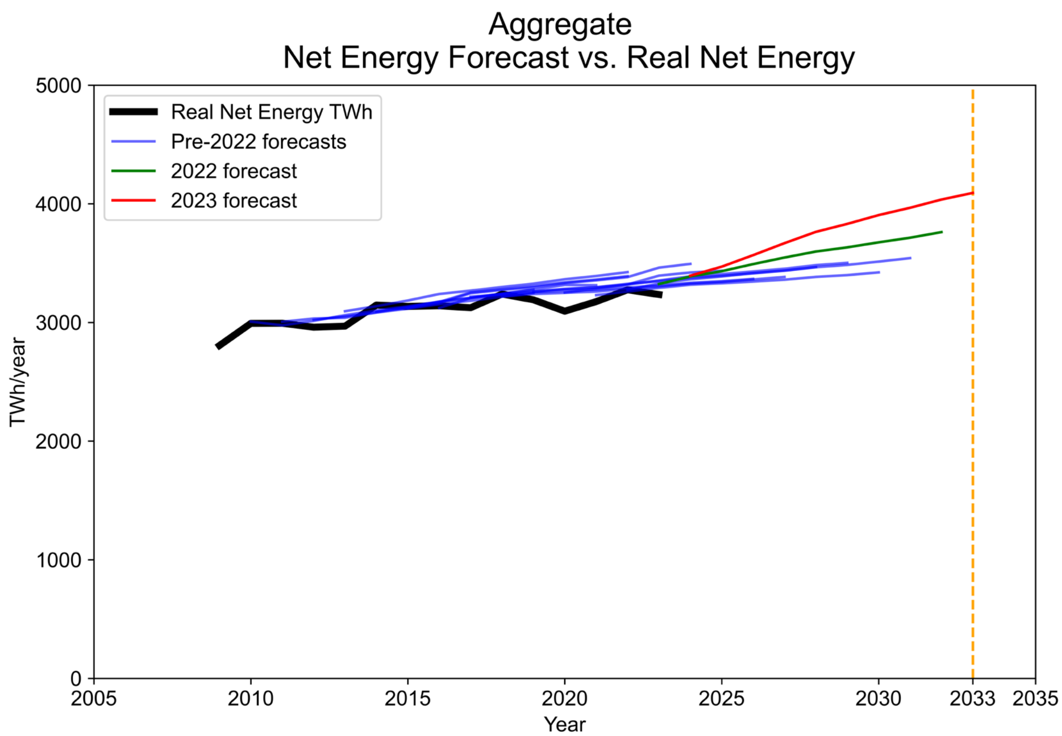

Source: https://www.ferc.gov/industries-data/electric/general-information/electric-industry-forms/form-no-714-annual-electric/overview

Figure 2 shows actual electricity generation rising slowly from 2009-2023, projections typically exceeding actual generation, and then significantly higher projected growth in the 2023 forecast—about 26% growth through 2033 (roughly 2.4% a year).

Utilities generate forecasts as part of their planning processes. In many states, the forecasts are incorporated into what are called "Integrated Resource Plans" in which options for meeting the projected demand are compared to each other. Those comparisons help the utility choose among those options to minimize costs, maintain or improve reliability, and achieve other goals, such as reducing emissions [17]. Uncertainty is almost always inherent in these forecasts [18, 19, 20, 21], but creating them is necessary so that utilities have a baseline against which to measure alternative scenarios.

Although there is no strong evidence for explosive electricity demand growth in the United States in recent years, these data may not tell the whole story. Looking beyond figures showing little load growth nationally, individual states and utilities are experiencing substantial load growth. For example, recent news stories about potential energy impacts of data centers have often focused on two states, Georgia and Virginia, which have seen growth in data center electricity use, Georgia more recently than Virginia.

Figure 3 shows Georgia Power's forecasts of future demand, compared with actual net generation going back to 2006, using data from the utility's Form 714 filings. The figure shows a large increase in Georgia Power's 2023 forecast (indicated by a red line) relative to its pre-2022 forecasts (indicated by blue lines) and its 2022 forecast (indicated by a green line). The 2023 projection shows growth of about 75% in total electricity generation by 2033.

Figure 3. Actual generation and 2006-2023 projections for Georgia Power

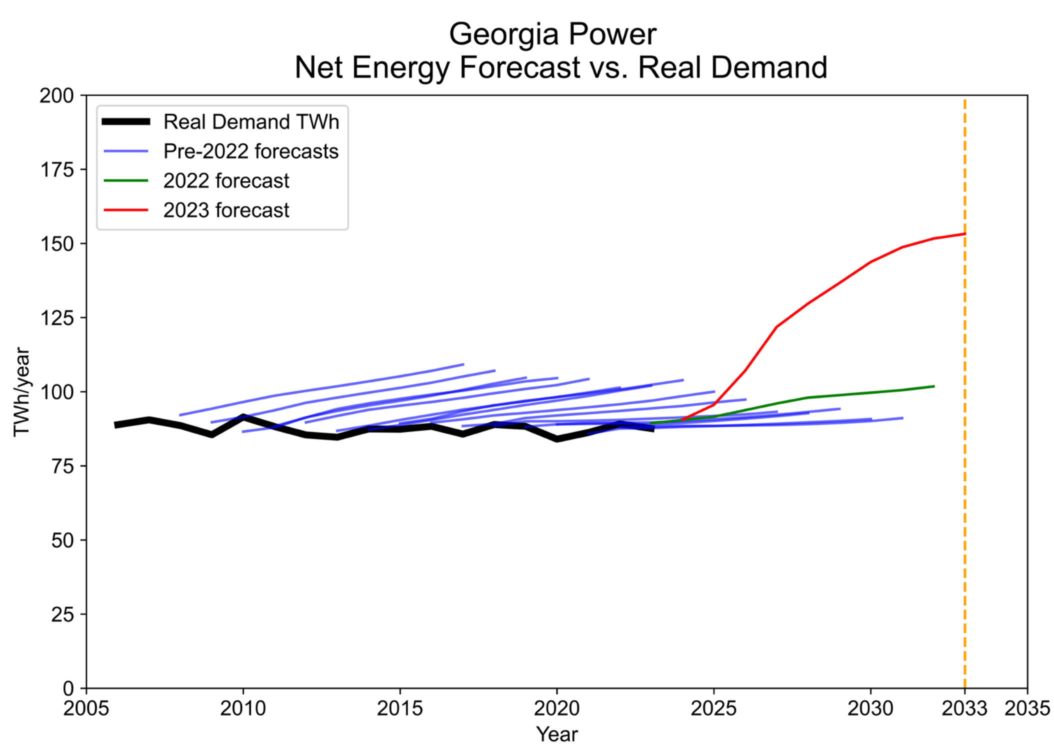

Source: https://www.ferc.gov/industries-data/electric/general-information/electric-industry-forms/form-no-714-annual-electric/overview

A similar pattern emerges from forecasts for the PJM Interconnection, which includes the state of Virginia, as shown in **Figure 4**. The 2023 forecast represents a substantial departure from past forecasts, showing a roughly 30% increase in demand by 2033 in the 2023 forecast.

Figure 4. Actual generation and 2006-2023 projections for PJM

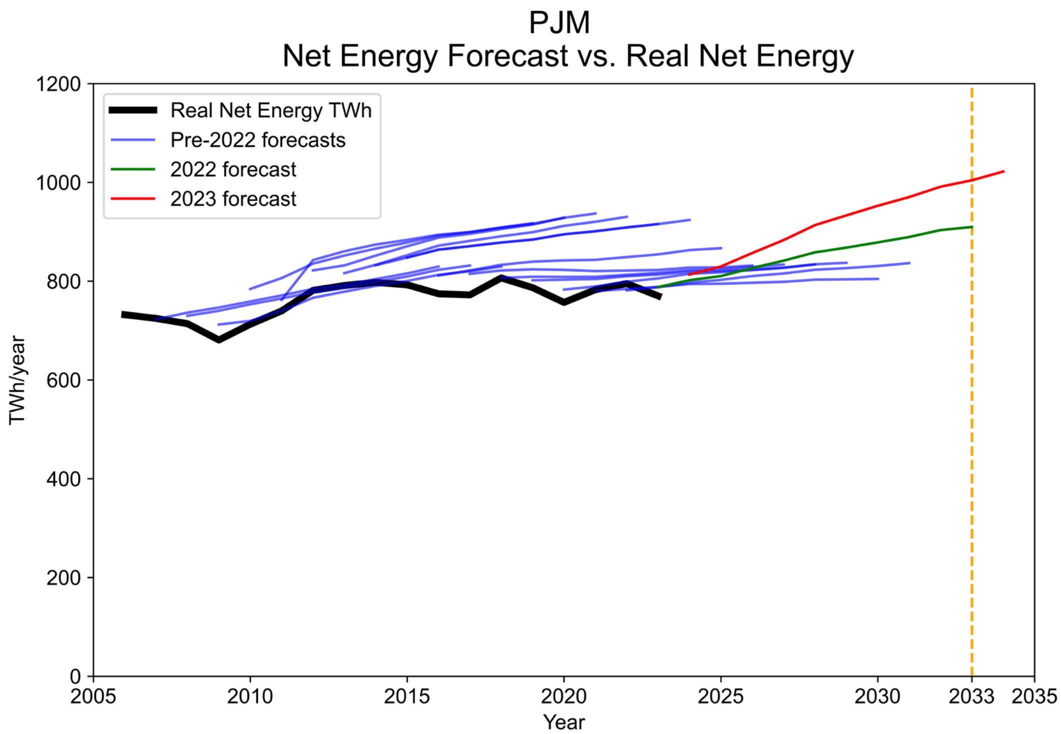

Source: https://www.ferc.gov/industries-data/electric/general-information/electric-industry-forms/form-no-714-annual-electric/overview

The 2023 forecast for the eastern Virginia portion of PJM anticipates one of the largest peak load increases for any utility in that region [22]. **Figure 5** shows the projected load increase for the DOM Zone in Virginia—the service territory of Dominion Energy, which covers much of the eastern part of the state—indicating a more than 50% increase in peak load projected to 2033 compared with 2023 [23].

Northern Virginia has the largest data center market in the world, with 70% of the world's internet traffic originating in or passing through Loudoun County [25].

Figure 5. Peak demand growth (MW) in DOM Zone

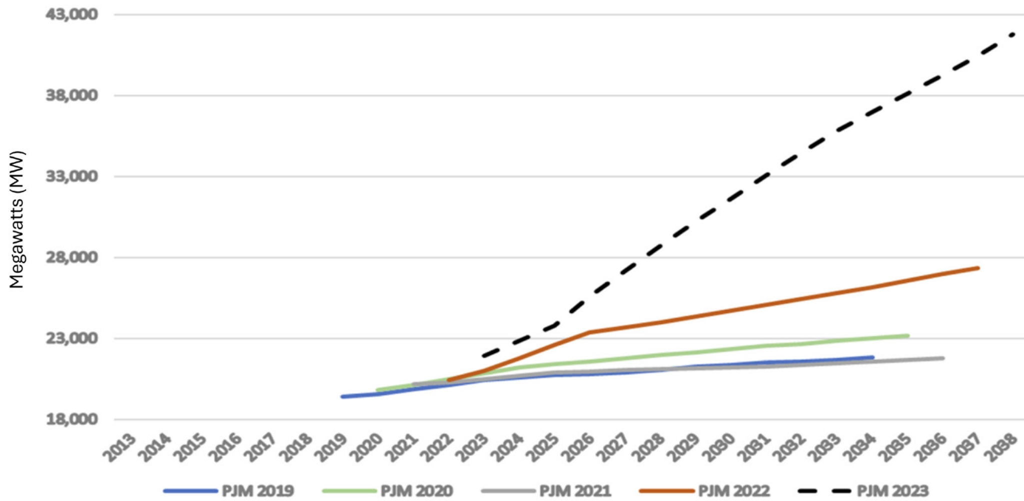

Source: https://rga.lis.virginia.gov/Published/2023/RD214

Dominion Energy has serviced and assessed data center load growth in its service territory for over a decade. In particular, Loudoun County, VA, contains 80% of Dominion's data center electricity demand [24]. Northern Virginia has the largest data center market in the world, with 70% of the world's internet traffic originating in or passing through Loudoun County [25].

Load growth from data centers in context

Load growth due to data centers in a specific region can be difficult to predict. Data center developers consider multiple states as possible locations for data centers, and they query multiple utilities simultaneously for electricity rates and incentives prior to making a final selection. Therefore, counting data center project proposals to forecast load growth can result in the overestimation of data centers likely to be built in a specific service territory. Only national or regional level tracking of these projects can give an accurate picture, but such tracking currently does not exist, at least in a publicly available form.

Figure 6 shows the sources of new load growth in Dominion Energy Virginia. In this case, the majority of projected load growth is driven by the commercial sector, and specifically by anticipated growth in data center demand [26]. Industrial and residential electricity sales, in contrast, are expected to be flat.

Figure 6. Electricity sales growth projected by sector for Dominion Energy Virginia

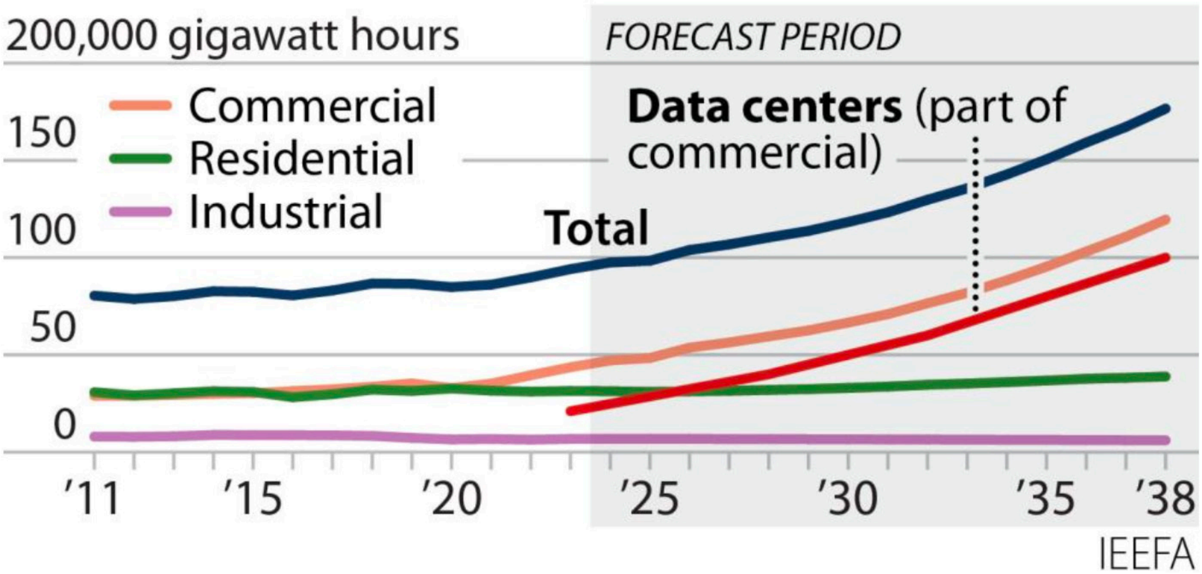

Source: https://ieefa.org/resources/dominion-virginias-improbable-integrated-resource-plan

Another contributing factor to the uncertainty of projecting data center demand growth is the availability of land and new transmission capacity needed to support new data centers. Easy expansion of data centers in Virginia has depended on cheap land, low power prices, and public support—conditions that prevailed until relatively recently. There is evidence that data center buildout is entering a new and more contentious phase with siting issues and growing public opposition [26]. Thus, expectations that growth will continue for the next 14 years as rapidly as it has in the past are uncertain.

Currently, utilities are collecting better data, tightening criteria about how to "count" projects in the pipeline, and assigning probabilities to projects at different stages of development [27, 28, 29, 30]. These changes are welcome and should help reduce uncertainty in forecasts going forward.

Electrification is also important to load growth

Data centers are only one potential source of load growth in coming decades. Others include shifts in industrial investments and electrification of vehicles, heat, and industrial processes [12, 31, 32]. Demand for air conditioning is also increasing as the world warms [33]. Due to the rapidly changing nature of the U.S. commercial sector, policymakers must consider these factors when deciding how to address electricity load growth.

Figure 7 shows U.S. electricity use in 2020 and projected use in 2030, highlighting projected load growth by sector for 2020-2030. The projections for buildings, industry, and transportation are part of a high electrification scenario compiled in the *Fifth National Climate Assessment* for the United States [34].

This scenario exercise was completed before recent concerns about electricity growth from data centers became salient—and before the onshoring effects from the Inflation Reduction Act became evident—so it is fair to infer that AI-driven scenarios of data centers' load growth were not included in the projections. We add on top of that bar (in red) the highest projections, by the Electric Power Research Institute (EPRI), of data center load growth to 2030 [35] for comparison.

This comparison shows that data centers would still only account for 25% of expected growth over 2020-2030. This comparison is admittedly a bounding one, given that the EPRI projections are at the high end of recent forecasts, but it does demonstrate that other sources of potential load growth are likely to be substantial compared with growth in data centers' electricity use. Total data center electricity use would be about 8% of U.S. electricity use in 2030 based on EPRI's projection, up from about 2% in 2020. The EPRI estimate for 2020 is comparable in percentage terms to estimates in [16] but higher than estimates for the world, which are around 1% [36].

Figure 7. U.S. electricity use (2020), compared with projected growth in electricity use by sector to 2030 in a high electrification scenario; top bar (red) is data center growth from EPRI's highest data center growth scenario

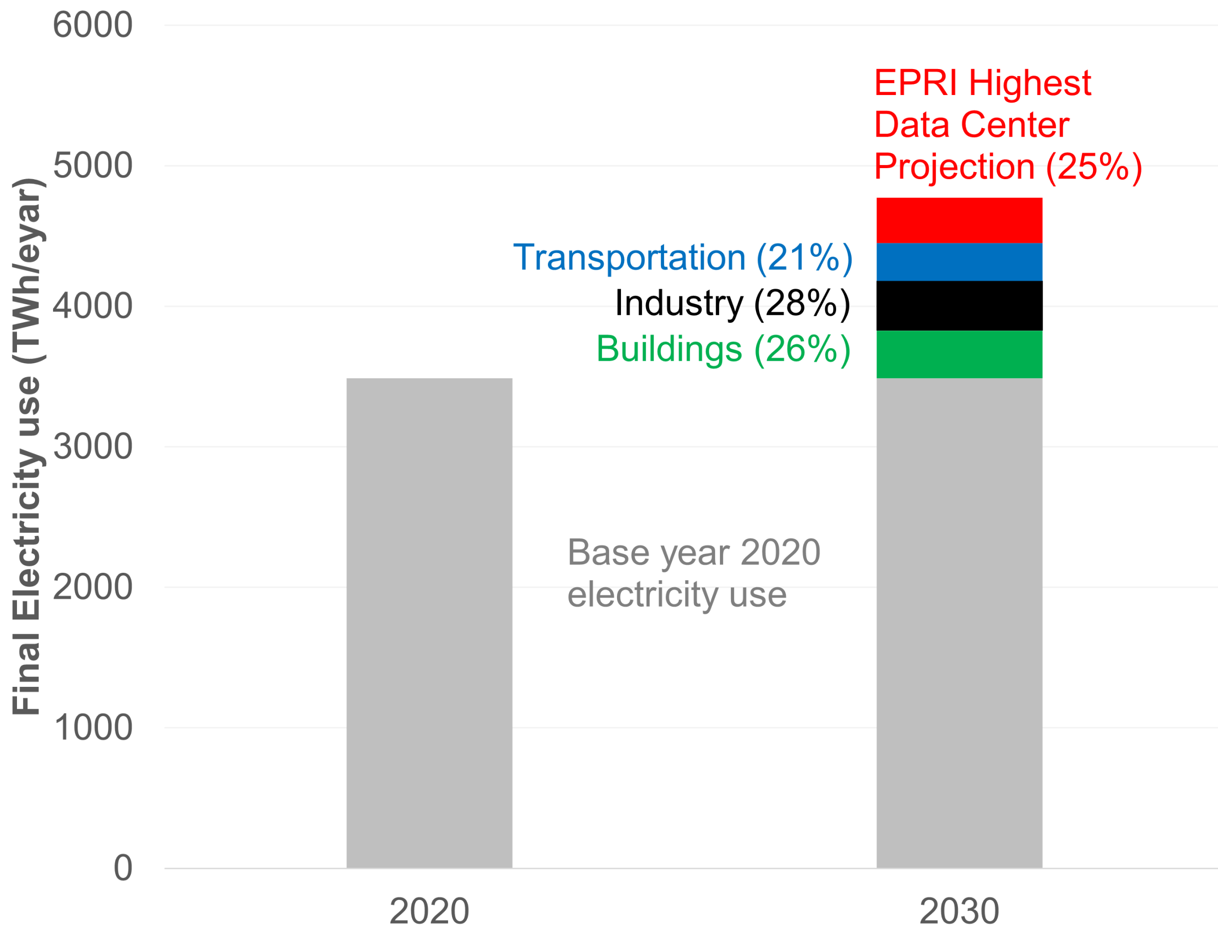

Understanding the potential for load growth from computing

The growth of two major factors will determine the impact of new computing applications on electricity demand:

- Service demand, as determined by the quantity of computations being performed.
- Computing efficiency, as determined by the amount of energy required to perform each computation.

As segments of the computing industry grow rapidly, an additional third factor may come into play in the form of *limits on industrial production capacity* to supply AI chips and servers, to expand the infrastructure needed to build new data centers (such as backup power generators), and to modernize the electricity grid. These three drivers are interdependent, and their trajectory in coming years is difficult to predict.

Consider future demand for AI compute: The industry's current growth projections are aggressive, but whether they materialize depends on businesses realizing positive economic returns from AI investments and on whether users' concerns about accuracy and reliability can be adequately addressed [37, 38, 39]. The industry's growth trajectory could also be affected by new technologies that can deliver similar services but without some of the risks and liabilities associated with current AI models.

Historically, as demand for computing services has increased, the efficiency of meeting that demand has also increased rapidly, offsetting some or all of the growth in demand for computing services [40]. In the early stage of the AI boom, efficiency was not top of mind, and companies bought all the AI equipment that was available, regardless of efficiency. As constraints in deploying AI manifested, the industry rightly began to focus on efficiency as one path to alleviating those constraints [41]. This pattern matches what happened from 2000-2005 when electricity used by data centers in the United States and globally roughly doubled [42], and the industry then focused on improving efficiency. This effort led to slower growth in data centers' electricity use in 2010 and little growth from 2010-2018 [43].

Growth in the ability to meet service demand can also be uncertain because of supply chain constraints in producing and deploying AI servers and supporting equipment. In the first phase of the recent AI boom, people bought as many AI servers as the industry could produce, causing shortages that could persist if demand continues to grow rapidly. There are also constraints on the physical systems—such as backup power generators for data centers—that may hamper the speed of AI deployment. If service demand growth moderates, these issues become less pressing. When growth in new technologies is rapid, it can affect the rate at which these technologies can be deployed.

Progress in computing efficiency in the future

How much potential is there to improve computing efficiency in the future? Most people know about progress in shrinking transistors [44, 45] in the form of "Moore's law." These changes resulted in rapid and consistent improvements in performance and efficiency from the first microprocessor in 1971 until the early 2000s, when real physical and economic limits started to slow things down [46]. During the modern era, these changes and others led to a doubling of computing energy efficiency at peak output every 1.6 years or so [40] through the year 2000. After 2000, this rate of improvement slowed to a doubling every 2.6 years [47].

Making smaller transistors is not the only way to improve efficiency. When it became harder to shrink transistors in the early 2000s, the industry had to rely on continued innovation in other areas. Charles E. Leiserson et al. [48] identified ways the industry could continue to push performance and efficiency forward at rapid rates, including changes to hardware architecture, using better software, improving algorithms, and adopting special purpose computing (including co-design of hardware and software).

The potential for these technologies to improve efficiency is vast. Leiserson and his co-authors cited an example of matrix multiplication that could be improved by a factor of 60,000 using these techniques. Other back-of-the-envelope calculations show that current technology is far from the ultimate physical limits for computing. Companies might need to rethink computing technology from first principles as they approach these limits in coming decades [49].

Some people have pointed to the "Jevons paradox" [50] to argue that improvements in the energy efficiency of computing always lead to a corresponding increase in energy use that swamps the efficiency gains. This claim reflects a misunderstanding about a complex phenomenon. As computing devices become more powerful, their costs per compute cycle drop rapidly. Lower costs can drive increased adoption, but whether that increased adoption fully or partially offsets efficiency improvements depends on the net benefits the new technology brings, as well as the offsetting effect of retiring the older technologies displaced by the new one. A technology that brings benefits greater than its costs will be adopted much faster than one whose benefits only slightly exceed its costs. But energy is only one cost of computing, albeit a significant one in some cases. It is not generally true that a change in energy efficiency must drive increases in energy or computing consumption.

Of course, new technology with significant net benefits will be adopted rapidly, and consumption of the inputs that allows that new technology to be adopted will also increase. This story, however, is not primarily about energy efficiency. It is about technological change driving down costs of new technology, which increases adoption, while displacing old ways of accomplishing the same task.

For AI systems, the question of efficiency is complicated. Historically, models have improved their performance with scale [51, 52], so that the bigger, more powerful models trained on more data outperformed earlier "best in class" models. AI researchers plowed all efficiency gains back into increasing aggregate performance, so electricity use increased, and that trend held for the most recent explosion of interest in AI [53].

For the most part, costs and benefits did not factor into these experiments, except in the narrow sense that the companies building the big models presumably kept to the budgets they set for those projects. However, business value (i.e., net benefits) will drive the next phase of AI deployment. Whether continuing to scale models to ever larger sizes will be the best approach to generate business value is an open question that will ultimately be answered in the marketplace, depending on people's willingness to pay for AI compute services and the cost to deliver those services. That open question is one of the main sources of deep uncertainty for AI service demand in the future.

Conclusion

It is incumbent on utilities, regulators, policymakers, and investors to investigate claims of rapid new electricity demand growth and to ensure that expectations are based on the latest and most accurate data and models. Although data centers' electricity use appears to be growing again, exactly how that growth will play out in coming years is deeply uncertain, both because growth in the use of AI is uncertain and because progress in efficiency is uncertain. It is likely, however, that other sources of demand growth, such as the onshoring of industry and the electrification of transportation, heating, and industry, will be bigger drivers of total demand growth than data centers in the medium to longer term.

Endnotes

1. Brian Bailey, "AI Power Consumption Exploding," *Semiconductor Engineering*, 2022. Available at: https://semiengineering.com/ai-power-consumption-exploding/.

2. Evan Halper, "Amid Explosive Demand, America Is Running Out of Power," *The Washington Post*, March 7, 2024. Available at: https://www.washingtonpost.com/business/2024/03/07/ai-data-centers-power/.

3. Camilla Hodgson, "Booming AI Demand Threatens Global Electricity Supply," *Financial Times*, April 17, 2024. Available at: https://www.ft.com/content/b7570359-f809-49ce-8cd5-9166d36a057b.

4. Laila Kearney, Seher Dareen, and Deep Kaushik Vakil, "US Electric Utilities Brace for Surge in Power Demand from Data Centers," Reuters, April 10, 2024. Available at: https://www.reuters.com/business/energy/us-electric-utilities-brace-surge-power-demand-data-centers-2024-04-10/.

5. Antonio Olivo, "Internet Data Centers Are Fueling Drive to Old Power Source: Coal," *The Washington Post*, April 17, 2024. Available at: https://www.washingtonpost.com/business/interactive/2024/data-centers-internet-power-source-coal/.

6. Brad Plumer and Nadja Popovich, "A New Surge in Power Use Is Threatening U.S. Climate Goals," The New York Times, March 14, 2024. Available at: https://www.nytimes.com/interactive/2024/03/13/climate/electric-power-climate-change.html.

7. Marietje Schaake, "Our Stressed-Out Power Grids Need Help," *Financial Times*, March 13, 2024. Available at: https://www.ft.com/content/d38ebd3e-84a2-4c22-b84c-6528a1e2dd3b.

8. Katie Tarasov, "Generative AI Requires Massive Amounts of Power and Water, and the Aging U.S. Grid Can't Handle the Load," CNBC, July 28, 2024. Available at: https://www.cnbc.com/2024/07/28/how-the-massive-power-draw-of-generative-ai-is-overtaxing-our-grid.html.

9. Srini Bangalore, Arjita Bhan, et al., *Investing in the Rising Data Center Economy*, Technology, Media, & Telecommunications Practice, McKinsey & Company, January 2023. Available at: https://www.mckinsey.com/~/media/mckinsey/industries/technology%20media%20and%20telecommunications/high%20tech/our%20insights/investing%20in%20the%20rising%20data%20center%20economy/investing-in-the-rising-data-center-economy_final.pdf.

10. Vivian Lee, *The Impact of GenAI on Electricity: How GenAI Is Fueling the Data Center Boom in the U.S.*, Boston Consulting Group, September 13, 2023. Available at: https://www.linkedin.com/pulse/impact-genai-electricity-how-fueling-data-center-boom-vivian-lee/.

11. Carly Davenport, Brian Singer, et al., Generational Growth: AI, *Data Centers*

and the Coming US Power Demand Surge, Goldman Sachs Equity Research, April 28, 2024. Available at: https://www.goldmansachs.com/pdfs/insights/pages/generational-growth-ai-data-centers-and-the-coming-us-power-surge/report.pdf.

12. T. Bruce Tsuchida, Long Lam, et al., *Electricity Demand Growth and Forecasting in a Time of Change*, The Brattle Group, May 2024. Available at: https://www.brattle.com/insights-events/publications/electricity-demand-growth-and-forecasting-in-a-time-of-change/.

13. Jeff Brown, *Data Center Electricity Consumption Is About to Quadruple*, Brownridge Research, June 3, 2024. Available at: https://www.brownridge.com/data-center-electricity-consumption-is-about-to-quadruple/.

14. Michael Cembalest, *Electravision: The Complicated Journey to an Electrified Future*, J.P. Morgan Asset & Wealth Management, 2024. Available at: https://privatebank.jpmorgan.com/nam/en/insights/latest-and-featured/eotm/annual-energy-paper.

15. John D. Wilson and Zach Zimmerman, *The Era of Flat Power Demand Is Over*, *Grid Strategies*, December 2023. Available at: https://gridstrategiesllc.com/wp-content/uploads/2023/12/National-Load-Growth-Report-2023.pdf.

16. Arman Shehabi, Sarah Josephine Smith, et al., *2024 United States Data Center Energy Usage Report*, Lawrence Berkeley National Laboratory, December 19, 2024. Available at: https://eta-publications.lbl.gov/publications/2024-lbnl-data-center-energy-usage-report.

17. Lauren Shwisberg, Katerina Stephan, and Mark Dyson, *Reimagining Resource Planning*, Rocky Mountain Institute, January 2023. Available at: https://rmi.org/insight/reimagining-resource-planning/.

18. Paul Craig, Ashok Gadgil, and Jonathan Koomey, "What Can History Teach Us? A Retrospective Analysis of Long-Term Energy Forecasts for the U.S.," *in Annual Review of Energy and the Environment 2002*, ed. R. H. Socolow, D. Anderson, and J. Harte (Palo Alto, CA: Annual Reviews Inc., 2002), 83–118. Available at: https://doi.org/10.1146/annurev.energy.27.122001.083425.

19. Jonathan G. Koomey, Paul Craig, et al., "Improving Long-Range Energy Modeling: A Plea for Historical Retrospectives," *Energy Journal*, 24(4): 75–92, October 2003. Available at: https://journals.sagepub.com/doi/abs/10.5547/ISSN0195-6574-EJ-Vol24-No4-4.

20. Irene Scher and Jonathan G. Koomey, "Is Accurate Forecasting of Economic Systems Possible?" *Climatic Change*, 104(3–4): 473–479, February 2011. Available at: http://link.springer.com/article/10.1007%2Fs10584-010-9945-z.

21. Jonathan Koomey, Karl Hausker, et al., "Abandon the Idea of an 'Optimal Economic Path' for Climate Policy," *WIREs Climate Change*, 14(6): e850, July 2023. Available at: http://doi.org/10.1002/wcc.850.

22. PJM Resource Adequacy Planning Department, *PJM Load Forecast Report*, PJM, January 2024. Available at: https://insidelines.pjm.com/pjm-publishes-2024-long-term-load-forecast/.

23. Dominion Energy, *Virginia Electric and Power Company's Report of Its 2023*

Integrated Resource Plan, Before the Virginia State Corporation Commission and North Carolina Utilities Commission, May 1, 2023. Available at: https://rga.lis.virginia.gov/Published/2023/RD214.

24. Dominion Energy, *Dominion Energy Service Territory: Data Center Forecasting*, June 26, 2023. Available at: https://www.pjm.com/-/media/committees-groups/subcommittees/las/2023/20230626/20230626-item-05---dominion-load-adjustment-method_results.ashx.

25. Virginia Economic Development Partnership, "The Dawn of Data: How Virginia Became the Data Center Capital of the World," *Virginia Economic Review*, Fourth Quarter, 2019. Available at: https://www.vedp.org/news/dawn-data.

26. Dennis Wamsted, *Dominion Virginia's Improbable IRP: Virginia Regulators Should Reject the Utility's Long-Term Resource Plan*, Institute for Energy Economics and Financial Analysis, November 2023. Available at: https://ieefa.org/resources/dominion-virginias-improbable-integrated-resource-plan.

27. Georgia Power, *2023 Integrated Resource Plan Update: Q2 2024 Large Load Economic Development Report*, Docket No. 55378, August 16, 2024. Available at: https://psc.ga.gov/search/facts-document/?documentId=219697.

28. Georgia Power, *2023 Integrated Resource Plan Update: Updated Load and Energy Forecast*, Docket No. 55378, 2024. Available at: https://www.georgiapower.com/content/dam/georgia-power/pdfs/company-pdfs/2023-irp-update-main-document.pdf.

29. Georgia Power, *2023 Integrated Resource Plan Update: Direct Testimony of Jeffrey R. Grubb, Francisco Valle, Lee Evans, and Michael A. Bush*, Georgia Public Service Commission, Docket No. 55378, December 4, 2023. Available at: https://psc.ga.gov/search/facts-document/?documentId=216591.

30. Georgia Power, *2023 Integrated Resource Plan Update: Rebuttal Testimony of Jeffrey R. Grubb, Francisco Valle, Lee Evans, Michael B. Robinson, and Michael A. Bush*, Georgia Public Service Commission, Docket No. 55378, March 18, 2024. Available at: https://psc.ga.gov/search/facts-document/?documentId=218033.

31. Electric Power Research Institute (EPRI), *US National Electrification Assessmen*t, April 3, 2018. Available at: https://www.epri.com/research/products/000000003002013582.

32. Pami Aalto, ed., *Electrification: Accelerating the Energy Transition* (Amsterdam: Academic Press, 2021). Available at: https://doi.org/10.1016/C2019-0-03794-6.

33. International Energy Agency, *Electricity Mid-Year Update*, July 2024. Available at: https://www.iea.org/reports/electricity-mid-year-update-july-2024.

34. A. R. Crimmins. C. W. Avery, et al., *Fifth National Climate Assessment*, U.S. Global Change Research Program, November 14, 2023. Available at: https://nca2023.globalchange.gov/.

35. Electric Power Research Institute, *Power Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption*, May 2024. Available at: https://www.epri.com/research/products/000000003002028905.

36. Jens Malmodin, Nina Lövehagen, et al., "ICT Sector Electricity Consumption

and Greenhouse Gas Emissions – 2020 Outcome," *Telecommunications Policy*, 48(3): 102701, April 2024. Available at: https://www.sciencedirect.com/science/article/pii/S0308596123002124.

37. Jason Koebler, "Goldman Sachs: AI Is Overhyped, Wildly Expensive, and Unreliable," *404 Media*, July 12, 2024. Available at: https://www.404media.co/goldman-sachs-ai-is-overhyped-wildly-expensive-and-unreliable/.

38. Allison Nathan, Jenny Grimberg, and Ashley Rhodes, *Gen AI: Too Much Spend, Too Little Benefit?* Goldman Sachs, June 25, 2024. Available at: https://www.goldmansachs.com/intelligence/pages/gs-research/gen-ai-too-much-spend-too-little-benefit/report.pdf.

39. Lexin Zhou, Wout Schellaert, et al., "Larger and More Instructable Language Models Become Less Reliable," *Nature*, September 25, 2024. Available at: https://doi.org/10.1038/s41586-024-07930-y.

40. Jonathan Koomey, Stephen Berard, et al., "Implications of Historical Trends in the Electrical Efficiency of Computing," *IEEE Annals of the History of Computing*, 33(3): 46–54, July–September 2011. Available at: http://doi.ieeecomputersociety.org/10.1109/MAHC.2010.28.

41. Don Nico Forbes, "Can Innovation Curb AI's Hunger for Power?" *The Wall Street Journal*, July 24, 2024. Available at: https://www.wsj.com/articles/can-innovation-curb-ais-hunger-for-power-e9c3d8bc.

42. Jonathan Koomey, "Worldwide Electricity Used in Data Centers," *Environmental Research Letter*s, 3(034008): September 23, 2008. Available at: http://stacks.iop.org/1748-9326/3/034008.

43. Eric Masanet, Arman Shehabi, et al., "*Recalibrating Global Data Center Energy-Use Estimates*," Science, 367(6481): 984, February 28, 2020. Available at: http://science.sciencemag.org/content/367/6481/984.abstract.

44. Gordon E. Moore, "Progress in Digital Integrated Electronics," *IEEE*, *IEDM Tech Digest*, 1975, 11–13. Available at: https://ieeexplore.ieee.org/document/4804410.

45. Robert H. Dennard, Fritz H. Gaensslen, et al., "Design of Ion-Implanted MOSFETs with Very Small Physical Dimensions," *IEEE Journal of Solid-State Circuits*, 9(5): 256–268, October 1974. Available at: https://ieeexplore.ieee.org/document/1050511.

46. Mark Bohr, "A 30-Year Retrospective on Dennard's MOSFET Scaling Paper," *IEEE SSCS Newsletter*, 12(1): 11–13, Winter 2007. Available at: https://ieeexplore.ieee.org/document/4785534.

47. Jonathan Koomey and Samuel Naffziger, "Energy Efficiency of Computing: What's Next?" *Electronic Design*, November 28, 2016. Available at: http://electronicdesign.com/microprocessors/energy-efficiency-computing-what-s-next.

48. Charles E. Leiserson, Neil C. Thompson, et al., "There's Plenty of Room at the Top: What Will Drive Computer Performance After Moore's Law?" *Science*, 368(6495): eaam9744, 2020. Available at: http://science.sciencemag.org/content/368/6495/eaam9744.abstract.

49. Jonathan G. Koomey, H. Scott Matthews, and Eric Williams, "Smart Everything: Will Intelligent Systems Reduce Resource Use?" *Annual Review of Environment and Resources*, 38: 311–343, October 2013. Available at: http://arjournals.annualreviews.org/eprint/wjniAGGzj2i9X7i3kqWx/full/10.1146/annurev-environ-021512-110549.

50. Richard York and Julius Alexander McGee, "Understanding the Jevons Paradox," *Environmental Sociology*, 2(1): 77–87, January 2, 2016. Available at: https://doi.org/10.1080/23251042.2015.1106060.

51. Jared Kaplan, Sam McCandish, et al., *Scaling Laws for Neural Language Models*, arXiv:2001.08361, January 23, 2020. Available at: https://doi.org/10.48550/arXiv.2001.08361.

52. Sebastian Sartor and Neil Thompson, *Neural Scaling Laws for Embodied AI*, arXiv:2405.14005, May 22, 2024. Available at: https://doi.org/10.48550/arXiv.2405.14005.

53. David Owen, "How Predictable Is Language Model Benchmark Performance?" *Epoch AI*. Available at: January 8, 2025. Available at: https://ourworldindata.org/grapher/ai-performance-knowledge-tests-vs-training-computation.

54. Dean Nelson, "Defining the Digital Infrastructure Industry," *Interglobix Magazine*, February 10, 2022. Available at: https://www.interglobixmagazine.com/defining-the-digital-infrastructure-industry/.

55. iMasons, *State of the Digital Infrastructure Industry: Annual Report*, 2024. Available at: https://imasons.org/blog/state-of-the-digital-infrastructure-industry-report-2024/.

56. International Energy Agency, Electricity 2024: *Analysis and Forecast to 2026*, January 2024. Available at: https://www.iea.org/reports/electricity-2024.

57. International Energy Agency, *World Energy Outlook 2024*, October 2024). Available at: https://www.iea.org/reports/world-energy-outlook-2024.

58. Bloomberg, "AI Is Already Wreaking Havoc on Global Power Systems," June 21, 2024. Available at: https://www.bloomberg.com/graphics/2024-ai-data-centers-power-grids/.

59. Victor Avelar, Patrick Donovan, et al., *The AI Disruption: Challenges and Guidance for Data Center Design*, Schneider Electric White Paper 110, version 1.1, September 7, 2023. Available at: https://www.se.com/eg/en/download/document/SPD_WP110_EN/.

Appendix A: Details on U.S. electricity demand growth

Figure A-1, which relies on data from the U.S. Energy Information Administration (EIA), shows that electricity consumption grew more rapidly than GDP (and more rapidly than primary energy consumption) between 1950-1973 (https://www.eia.gov/opendata/). Between 1973-2000, the growth in electricity demand tracked GDP growth, while primary energy consumption began to level off. Beginning in the early 2000s, electricity consumption, including primary energy consumption, was increasingly decoupled from GDP growth, and aggregate demand grew only slowly, increasing at a rate of about 0.3% per year from 2014-2023. In 2023, electricity consumption was down slightly compared with 2022.

Figure A-1. Historical U.S. electricity use, primary energy use, and economic output (GDP)

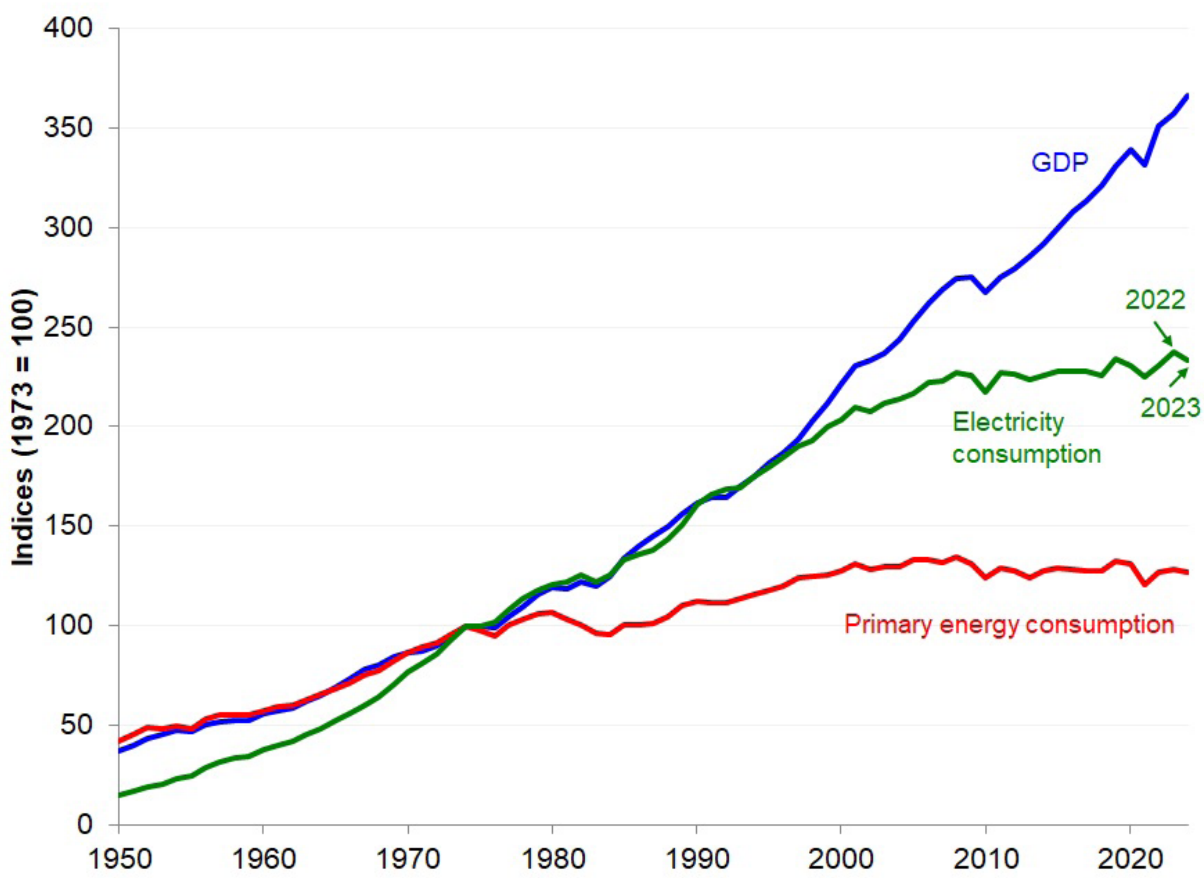

What about more recent trends? **Figure A-2** shows U.S. net electricity generation by month for 2018, 2019, 2022, 2023, and the first part of 2024, also from EIA (data from 2020 and 2021 are anomalous because of the COVID-19 pandemic and are not included). Monthly data are more subject to weather variations, but the figure does not show significant demand growth at the national level over the past few years.

Figure A-2. Monthly U.S. net generation for selected years

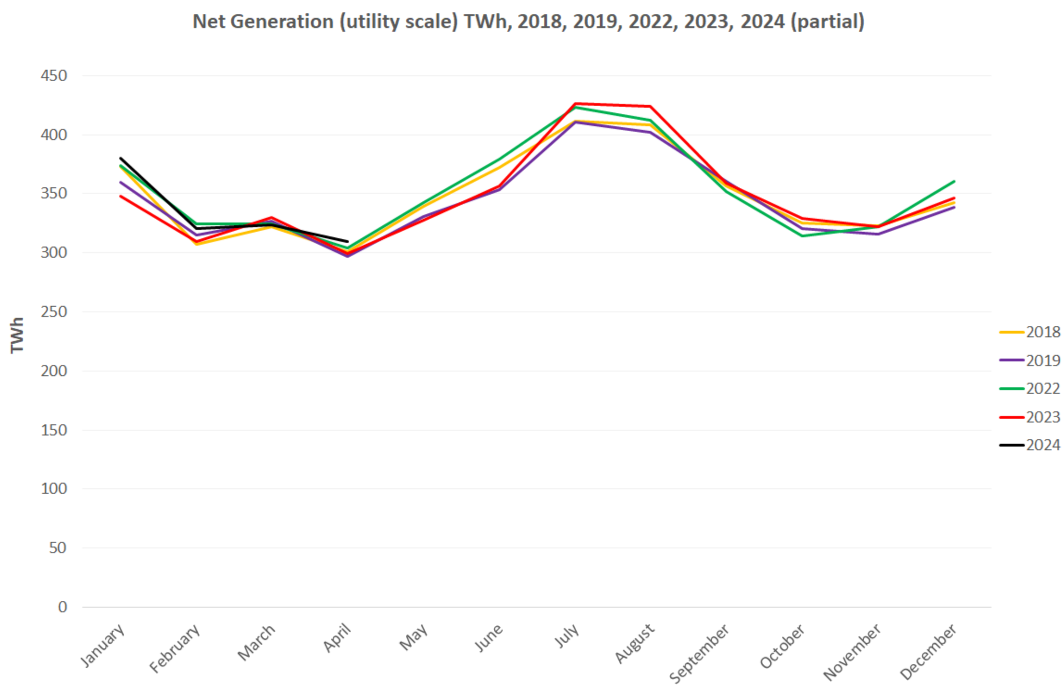

EIA also collects data on monthly electricity sales by sector. If data centers were adding significant new load, one would expect to see growth in the commercial and industrial sectors in recent years. Commercial sales were lower in 2023 compared with 2022, whereas industrial sales increased slightly between 2022-2023. Combined, commercial and industrial sector electricity sales declined slightly from 2022-2023.

Two additional factors could affect electricity sales data as they are currently reported: weather variations and "behind the meter" electricity generation from solar photovoltaics (this generation is not counted in either retail electricity sales or net generation statistics).

Heating and cooling loads in the residential and commercial sectors are temperature sensitive, and year-to-year (as well as month-to-month) weather variations can have a measurable effect on electricity use. **Figure A-3** shows EIA data for the sum of annual commercial and industrial sales from 2014-2023 (blue line) with commercial adjusted for the average weather in that period (orange line), which smooths out yearly weather variations. The results do not change much, but adjusting for weather produces a smaller reduction in electricity sales between 2023 and 2022.

Adding EIA estimates of behind-the-meter solar generation slightly increases total electricity consumption figures for the residential, commercial, and industrial sectors in 2014 and subsequent years (the green line in **Figure A-3** shows the combined effect of this adjustment for commercial and industrial customers). Total consumption increases slightly, but the shape of the curve does not change much.

Figure A-3. EIA electricity sales for the sum of commercial and industrial sectors, corrected for weather and behind-the-meter solar generation

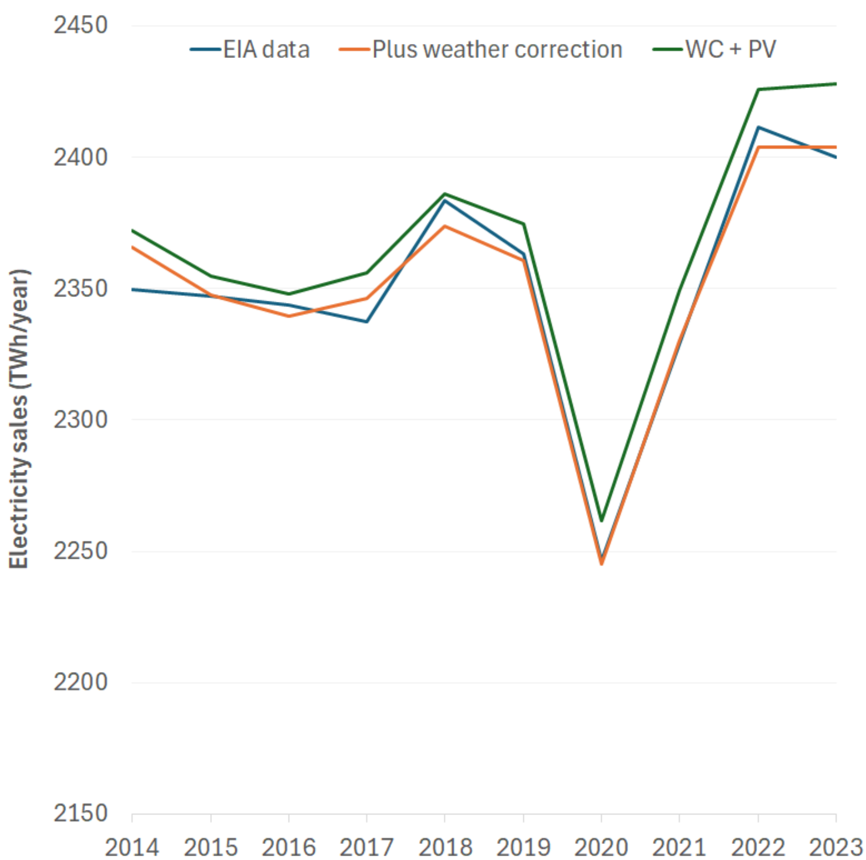

Appendix B: Utilities and balancing authorities included in aggregate forecast tallies in Figure 2

Arizona Public Service Company

Bonneville Power Administration

California Independent System Operator

Duke Energy Carolinas

ERCOT

Florida Power & Light Company

Georgia Power

ISO New England

MISO

NYISO

PJM

Puget Sound Energy

Southwest Power Pool

Tennessee Valley Authority

Appendix C: Understanding recent estimates of data center electricity use

Most estimates of data centers' electricity use have been conducted at the global level, with some regional disaggregation. **Figure C-1** summarizes the most important global estimates for data centers in recent years. These data only include compute data centers and exclude network providers and cryptocurrency. AI data centers are implicitly included within the category of compute data centers, and are sometimes explicitly identified, but have been relatively small until recently.

The peer reviewed estimates are those by Koomey [42], Masanet et al. [43], and Malmodin et al. [36]. These are shown as circles of different colors. From 2000-2005, data centers' electricity use doubled, then growth moderated to 2010 and significantly slowed down in the decade after that. By 2018, data centers used a little under 1% of the world's electricity.

The second group of data is from two organizations with some topic knowledge that produced nonpeer-reviewed estimates for 2021 from iMasons [54, 55] and for 2022 from the International Energy Agency (IEA) 2024a and 2024b [56, 57]. iMasons is an industry group of the world's top data center designers, builders, and operators. IEA is a government-sponsored institution that tracks energy consumption and production trends globally and regionally. These data points are colored squares.

The iMasons estimate for 2021 is close to the 2020 data point from Malmodin et al. The IEA 2024a estimate for 2022 is about 50% higher than Malmodin's 2020 data point. Unfortunately, the IEA report does not give enough information to determine how it arrived at this figure. It could indicate substantial recent growth, or it could be anomalous; nobody knows yet.

In October 2024, IEA released its World Energy Outlook [57], which gave a range of 240 to 340 TWh/year for total compute data center electricity use in 2022 (p. 186). The top end of that range corresponds to the earlier IEA estimate, while the bottom end is closer to iMasons' and Malmodin's estimates. This report also does not give enough detail to determine how it arrived at these figures.

The other three estimates are from Bloomberg [58], Goldman Sachs [11], and Brownridge [13]. These data points are colored triangles. The three firms are analysis shops that are ostensibly credible, but some of these estimates should give one pause. The Bloomberg numbers pre-2020 are particularly concerning, given how low these estimates are compared with the peer-reviewed analysis. The estimates for all three of these sources post-2020 are in the same ballpark as the IEA 2024a number, but given the lack of transparency for the data and methods for these estimates, it is impossible to know how accurate they are.

As time passes, more data will become available to determine whether the projected short-term growth post-2020 is a real phenomenon. Credible analysis always lags events in this space, which is often frustrating to policymakers.

Figure C-1. Estimates of global data center electricity use for historical years

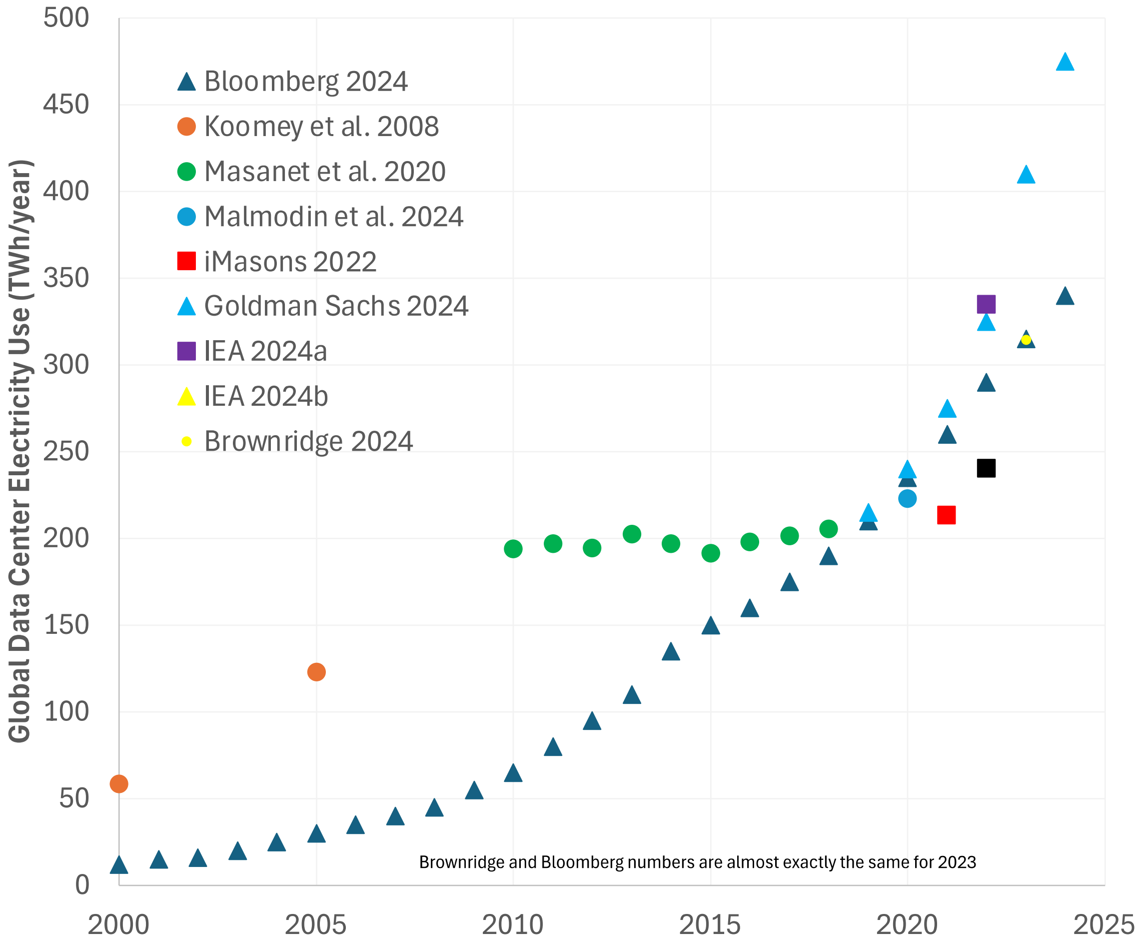

There are also a variety of projections that have been published over the past two years, and it is interesting to compare their projected growth rates to historical developments. **Figure C-2** shows average annual growth rates for a variety of time periods from different sources. At the top are growth rates for different periods from the peer-reviewed data. Growth from 2000-2005 was about 16% per year, and from 2005-2010 was about 9% per year. That growth slowed drastically after 2010.

The graph also shows growth rates from 2020-2022 or 2023 for IEA, Bloomberg, and Brownridge (the 2020 data point comes from Malmodin in all cases). The implied growth rate for the IEA 2024a data is over 20% per year, which would be remarkable if true. The Bloomberg and Brownridge growth rates are about 12% per year, which are rapid but not without historical precedent.

Figure C-2. Growth rates in projected global and U.S. data center electricity use

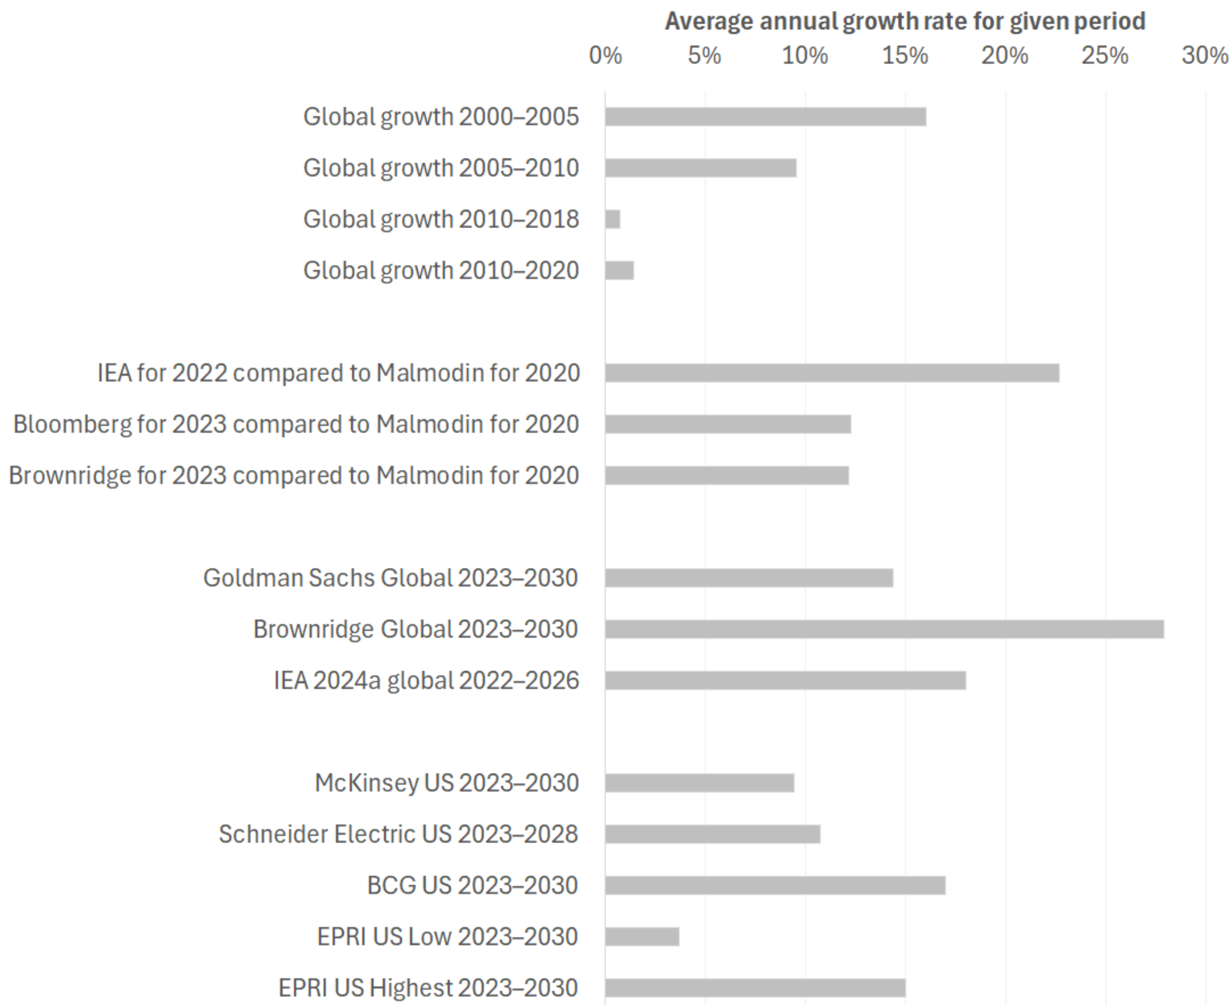

The Goldman Sachs global projections from 2023-2030 show growth of just under 15% per year; the Brownridge global projections over the same period show growth of 28% per year; and the IEA projection from 2022-2026 shows growth of 18% per year. The latter two projections exceed the growth rates in the 2000-2005 period, for an industry that is now much bigger than it was then.

For U.S. data centers' growth, figure C-2 also shows projections from McKinsey [9], Schneider Electric [59], Boston Consulting Group (BCG) [10], and the Electric Power Research institute (EPRI) [35], which gives a range of estimates. BCG and the EPRI high case show growth rates comparable to 2000-2005, while the McKinsey and Schneider electric cases show growth comparable to 2005-2010.

Which, if any, of these projections are credible? Comparisons to historical data can be suggestive, but they are not conclusive. If the benefits of AI are as compelling as many think, the growth in service demand for AI could indeed be unprecedented, driving absolute electricity use upward at rapid rates. At the same time, an initial period of rapid growth could lead the industry to invest heavily in improving efficiency, just as it did after the big growth spurt in the early 2000s.

There are some things we can say with confidence. Projections of computing's electricity use more than a few years out are no better than guesswork. Growth rates in electricity use equal to or greater than those prevailing in the 2000-2005 period (for an industry that was much smaller then) are unlikely to prevail for very long and are worthy of careful scrutiny. As the uncertainty range for these data is huge, caution and humility are the order of the day.

Bipartisan Policy Center

1225 Eye Street NW, Suite 1000
Washington, D.C. 20005

Where democracy gets to work

ATTACHMENT 3

Review of Large Load Tariffs to Identify Safeguards and Protections for Existing Ratepayers

By: Stacy Sherwood

On behalf of Earthjustice

Final Version

January 28, 2025

Table of Contents

Introduction

Cloud computing, artificial intelligence ("AI"), and cryptomining have resulted in an unprecedented projected growth in power demand throughout the nation, and many forecasts find that such demand will continue to grow significantly over the next decade. In its February 2024 analysis, EIA estimated that cryptocurrency mining in the U.S. may represent up to 2.3% of the annual total U.S. electricity demand.[1] Between May and August of 2024, there were predictions that data centers alone could reach as much as 7.5-9% of the United States' total electricity consumption by 2030.[2] [3] Due to the size and frequency of requests, forecasted load related to data centers and cryptomining are ever changing evolving and can change every few months.

The increase in power demand for data centers and other large consumption activities can negatively impact existing customers on the electric system and limit or eliminate progress on renewable energy and greenhouse gas emissions goals.[4] Negative impacts can include increased electricity demand that cannot be met with current capacity and increased congestion, a new customer's operations ceasing after a utility's significant investment in distribution and/or transmission infrastructure and procurement of new capacity. These translate into increased and abandoned costs left to be recovered from existing ratepayers.

For data centers, the full operating capacity does not typically occur for the first few years of a utility service contract, which impacts the timing of cost recovery and cash flow from servicing the load for the utility. Therefore, it's pertinent to include safeguard provisions in tariffs and special contracts to protect ratepayers and environmental goals, such as ensuring the facility is paying its fair share of transmission and distribution costs associated with service, requiring a certain number of jobs for economic development rates, and meeting decarbonization plans and goals of both the host jurisdiction and the host utility.

This report consists of four sections. The first section briefly considers why technology giants, such as Microsoft and Amazon, have an interest in designing their own contracts related to data centers and clean energy procurement. Second, this report summarizes a review of high-density tariffs and special contracts established for large load customers. Through this review, common provisions were identified, as well as details on how certain provisions can serve as

---

[1] *Tracking Electricity Consumption from U.S. Cryptocurrency Mining Operations*, U.S. Energy Information Administration, Feb. 1, 2024, https://www.eia.gov/todayinenergy/detail.php?id=61364.
[2] *How Data Centers Can Set the Stage for Larger Loads to Come*, Alexandra Gorin, Roberto Zanchi, and Mark Dyson, May 3, 2024, https://rmi.org/how-data-centers-can-set-the-stage-for-larger-loads-to-come/, accessed October 18, 2024.
[3] *Clean energy Resources to Meet Data Center Electricity Demand,* U.S. Department of Energy, August 12, 2024, https://www.energy.gov/policy/articles/clean-energy-resources-meet-data-center-electricity-demand#:~:text=Data%20center%20deployment%2C%20partly%20driven,of%20total%20load%20in%202023, accessed October 18, 2024.
[4] Although some may use the terms data center and cryptomining facility interchangeably, there is a distinction between the two, particularly when it comes to operation. Cryptomining facilities operate depending on the price signal from the crypto markets, with facilities operating up to 24 hours a day depending on the financials. Data centers have high load factors and operate on a 24/7 basis.

safeguards for ratepayers and/or environmental goals. The third section identifies ongoing proceedings and efforts to monitor as they could have a significant impact on the structure of high-density tariffs in the future. The final section of this report discusses certain safeguards more in-depth and identifies specific language for consideration in future tariffs and special contracts to serve as safeguards for ratepayers.

With the evolving market surrounding the electric service of data centers and large loads, it should be noted that this report was drafted based upon the information available throughout the latter half of 2024. The cases summarized in the third section of this report are based upon the information available at the time and will not include all details of the case, such as settlement proposals and commission orders. For clarity, in this document, a reference to a data center or cryptocurrency mining customer that the tariff would be applicable to will be identified as "customer," the utility will be referred to either as "utility" or "company," and those already on the power system will be referred to as "ratepayers."

Tech Giants' Interest

Technology giants, such as Amazon, Google, Microsoft, and Meta, all have significant stakes in locating and developing their data centers to support cloud computing and artificial intelligence. In addition to trying to develop a competitive edge in the data center world, each organization has corporate goals related to clean energy. Additionally, the technology giants may also have policies related to the implementation of their data centers. For example, requirements for onsite backup power. Price signals in the market help the companies determine which types of onsite power back up is procured (storage versus fossil fuel generators).

Corporations pursuing data centers may be proactively working with utilities on tariff development to find ways to reduce costs around onsite generation back up, energy costs, and achieving renewable energy goals. If a corporation is working with a utility to develop a tariff, the corporation can ensure the tariff supports its efforts to develop a competitive edge, while achieving corporate goals and requirements for siting data centers.

Review of Existing Tariffs and Special Contracts

A multitude of tariffs and special contracts were reviewed, from which a total of ten tariffs, each from a different state, were identified as being models for consideration based upon the safeguards included in the tariff language.[5] Regardless of the location, there are common rate structure elements, including:

- Contract length, requirements for investment by the new customer, and cost assignment.
- Demand, load factor, and power factor.
- Requirements to shed load and/or participate in demand response.

[5] A detailed summary of the reviewed tariffs and special contracts are provided in Appendix A of this report.

- Resource adequacy and requirements related to renewable or clean energy.

There is not one perfect tariff design that can adequately address the potential concerns related to large loads, and it is likely that large load tariffs will have to evolve over time, as loads and customers' requirements continue to change. However, there are elements of a rate structure that can serve as safeguards for existing ratepayers, ensure new customers pay their fair share of system costs, promote more efficient electricity usage, and minimize adverse impacts to clean energy and climate goals.

Figure 1 below provides the prevalence of safeguard provisions throughout the ten tariffs examined. A more detailed review of each of the requirements is provided in Appendix A, along with a link to the tariff or special contract. A green circle indicates that a safeguard is included as part of the tariff, while a red circle indicates that it is not a tariff requirement. If the circle is white, then it is considered not applicable, either because it was not mentioned, or in the case of demand response, it is not offered by the utility. As noted below, not one of the tariffs includes all the safeguard provisions discussed in this report. That is because safeguards are dependent upon a service territory's needs, which could pertain to ensuring the customer base does not suffer from stranded asset costs or to capacity and transmission constraints. For example, if there is excess capacity in a service territory, stakeholders may not be as concerned with having a robust demand response program or interruptible tariff.

Figure 8 Safeguards Included in Data Center and Cryptocurrency Tariffs

State	Utility	Document Type	Contract Length	Minimum Demand	Minimum Load Factors	Range for Power Factor	Requirements for Investment	Cost Assignment	Requirement to Shed Load	Load Subject to Interruptible Service	Maximum Hours of Interruptible Per Year	Demand Response
WY	Cheyenne Light, Fuel and Power Company d/b/a Black Hills Energy	SC	🟢	🟢	○	○	🟢	🟢	🟢	🟢	🟢	🟠
AR	Entergy Arkansas LLC	T	○	○	○	○	🟢	🟢	🟢	🟢	🟢	○
ID	Idaho Power Company	T	🟢	🟢	○	🟢	🟢	○	🟢	🟢	🟢	🟠
NY	New York Municipal Power Agency	T	○	🟢	○	○	🟢	🟢	🟠	○	○	○
SD	Montana-Dakota Utilities Company	T	🟢	🟢	🟢	🟢	🟠	🟠	🟢	🟢	🟢	🟢
WA	Grant County Public Utility District	T	🟠	🟠	○	○	🟠	🟠	🟠	○	○	🟢
IN	Indiana Michigan Power	T	🟠	🟢	○	○	🟢	○	🟠	○	○	🟠
KY	Kentucky Power	SC	🟢	🟢	○	○	○	○	🟠	🟢	🟢	🟢
MO	Evergy Missouri Metro	T	🟢	🟢	○	○	○	🟢	🟠	○	○	🟢
ND	Montana-Dakota Utilities Company	T	🟢	🟢	🟢	🟢	🟠	🟠	🟢	🟢	🟢	🟢

Note: For document type, “T” indicates a tariff and “SC” indicates a special contract.

Below is a more in-depth discussion of the safeguards in existing contracts and how they could be applied to future contracts for large loads.

Contract and Minimum Demand

The most prevalent safeguards include establishing a contract term length and minimum monthly demand to qualify for the tariff. The latter is a typical element of a commercial or industrial rate structure. This allows for targeting certain, or significant, energy loads. By establishing a monthly demand minimum for participation, the tariff can allow smaller load customers to receive service through another tariff, where the associated risks are not as significant. Minimum demand should be determined:

- in relation to the overall demand from the commercial and industrial customers and sector,
- in relation to the overall service territory's demand; and,
- through consideration of the available capacity in the system and the need for additional capacity builds.

> Caution: The tariff should indicate if the minimum demand is based upon the location, service point, or customer. There is potential for customers to find ways to avoid paying the tariff by structuring the demand in a manner that stays below the minimum demand threshold, such has having multiple meter points for a single customer

Not only can demand serve as a minimum requirement for a tariff, but there can also be a demand threshold that requires customers above a certain level of demand to have a special contract. This can be useful in large load scenarios as it will allow for the utility to ensure safeguards are in place for existing ratepayers, the Company, and the customer.
Idaho Power Company's Speculative High-Density Load tariff is offered to those with metered usage exceeding 2,000 kilowatt hours ("kWh") for at least three billing periods and requires customers with a minimum demand threshold of 1,000 kilowatts ("kW") to be served under this tariff. The tariff specifies that a special contract is required for loads over 20,000 kW.[6] The tariff language is provided below.

> SCHEDULE 20
> <u>SPECULATIVE HIGH-DENSITY LOAD</u>
>
> If the aggregate power requirement of a Customer who receives service at one or more Points of Delivery on the same Premises exceeds 20,000 kW, the Customer is ineligible for service under this schedule and is required to make special contract arrangements with the Company.
>
> Service under this schedule is applicable to electric service supplied to a Customer at one Point of Delivery and measured through one meter delivered at the primary or transmission service level. This schedule is applicable to Customers whose metered energy usage exceeds 2,000 kWh per Billing Period for a minimum of three Billing Periods during the most recent 12 consecutive Billing Periods. Where the

[6] Idaho Power Schedule 20 Speculative High-Density Load: https://docs.idahopower.com/pdfs/AboutUs/RatesRegulatory/Tariffs/20.pdf.

The contract term length is not related to the offering of the tariff; rather, this is a feature of the special customer service agreement. There are various lengths used by utilities and are likely dependent upon risk associated with the customer's service load. Of the arrangements reviewed, contract terms varied from two to ten years. In addition to the overall contract, some utilities required terms for renegotiation and/or pricing terms. Longer-term contracts, such as contracts of ten years or more, may have a shorter term related to pricing, as that is harder to accurately forecast over an extended period. Most of the contracts had contract length options within the three- to five-year span. This allows for limited forecasting on price and can accommodate ramp up in load, while also allowing for cost recovery of improvements to the system.

Some large load tariffs, such as those for facilities with a load greater than 50 MW, are proposing longer contract terms, such as 20-year minimums, with termination of the contract only if the facility ceases operation along with a penalty payment.[7] Large loads, like those more than 100 MW, will require significant investment in the electric system, both in capacity and the transmission system. Investments of that size are riskier given the level of cost recovery, depreciation of assets, the need for large capacity resource builds, and the fact that the significant load increased will be limited to one customer class rather than spread across multiple customer segments. The benefit of a longer contract term for this size of customer is that the cost recovery of the investment can be spread over the contract term. This will also allow for cost allocation that enables these customers to pay for their share of the utility investment needed to provide them with electric service. A negative of a long contract term, particularly if there is not much diversity in the customer class, is that an economic downturn or changes in the industry could significantly impact the load and need for service. For example, if the industrial customer class primarily consists of cryptocurrency mining customers, then a decrease in proof-of-work cryptocurrency value could limit the utility's revenue from that class. Therefore, it is important to develop a guardrail to alleviate the risk throughout the years of the contract. As noted in the Investment Requirement and Cost Assignment subsection below, the requirements for deposits throughout the life of the contract can offset some of this risk. A deposit can offset stranded costs if usage is below a minimum threshold or if the customer shuts down.

The contract itself can outline cost allocations to the customer, deposit terms, and credits to be returned to the customer for continued electric service and initial infrastructure investment to support the customer's load. Any known increases in load throughout the contract period can be addressed at the time of the contract being drafted, or through contract amendments, particularly if there is additional investment required to bring that load onto the system.

---

[7] Examples of these proposed tariffs include Kentucky Power Company's New Tarif Industrial General Service: https://psc.ky.gov/pscscf/2024%20cases/2024-00305/20240830_Kentucky%20Power%20Tariff%20Filing.pdf and Appalachian Power Company and Wheeling Power Company's Application for Approval of Revisions of Schedules LCP and IP https://www.psc.state.wv.us/scripts/WebDocket/ViewDocument.cfm?CaseActivityID=625853&NotType=WebDocket.

Load and Power Factors

In addition to contract and minimum demand levels, tariffs and special contracts also may establish a minimum load factor or a range for power factor to encourage consistent monthly energy usage. Encouraging consistent energy usage will ensure that utilities can cover the fixed cost to serve the load. Demand ratchets, discussed below, are another method of ensuring fixed costs are covered.

Load factor is the average power usage compared to peak power usage during the same period, measured as a percentage. The higher the percentage indicates the more efficient use of electricity. The desired effect of a minimum load factor is to smooth out demand peaks to lower the strain on the power infrastructure and increase reliability.

Power factor, also measured as a percentage, indicates the effectiveness of the use of incoming power by a specific load or equipment. The higher the power factor, the more efficient performance of the load/equipment. More efficient usage of power can reduce energy costs and system losses, which translates into savings for all customers.

Load factors are dependent upon the customer's usage. For example, an office building, which has low usage on weekends, can experience a load factor of 40-60%, whereas a cryptomining facility that is dependent on the value of the currency may have a lower load factor due to spikey monthly usage. A large load data center, since it is constantly active, will have a high load factor of 90-100%. Ultimately, the load factor is dependent on the type of customer/industry. The utility can include a load factor charge to penalize those customers that do not maintain a certain load factor, based on the type of customers being served under that tariff.

Demand Ratchet

While residential customers are billed on energy usage, commercial and industrial tariffs also include a demand charge component. A demand charge, which is used to cover fixed costs associated with a customer's load, is based upon the peak demand during the billing period.[8] The demand charge typically reflects a per kilowatt hour charge based upon the highest level of demand during a billing period. This charge allows the utility to recover the cost of providing a reliable service during those high peaks. Utilities must provide reliable service at those maximum demand levels; however, a customer may have significant shifts in demand by hour, day, or month.

[8] Peak demand is based on the level of demand over a 15-minute period.

Demand Ratchet Tariff Example

Here is an example of an 80% demand ratchet over an 11-month period. In this example, the demand charge is based upon the greater of the actual peak demand in the billing month or 80% of the highest peak demand recognized in the prior 11-month period.

Ex. In September, a facility's maximum peak demand was 400 kW and in the prior 11-months, the facility recognized its highest demand peak of 560 kW in July. The demand ratchet dictates that the demand charge for the month of September would be based on the greater of the 400 kW of actual usage or 448 kW (80% of 560 kW). Therefore, the facility would be charged a peak demand of 448 kW, since that is greater, resulting in the customer paying for 48 kW of demand it did not actually use.

One way that utilities reduce risks of serving customers that have large swings in demand is to assess demand charges using a demand ratchet.[9] The demand ratchet establishes the level of the demand charge based upon the actual peak demand, or a percentage of the highest demand recorded during the previous certain number of months, whichever is greater. The percentage of demand typically ranges from 80-85% of the previous period's demand, and the previous period can range from 9 to 11 months. Utilizing a demand ratchet encourages the customer to maintain a level of demand that is consistent as the customer would have to pay for demand not utilized if it does not.

Demand Shedding

Another safeguard that is often included or available is the opportunity to shed load, either through an interruptible tariff or through a demand response program. The availability of an interruptible tariff or a formal demand response program appears to be dependent upon the size of the service territory and utility type (investor-owned / cooperative / municipality). Even without a formal avenue to shed load, such as an interruptible tariff or demand response program, some tariffs included language for the utility to be able to enter into demand shedding agreements directly with customers. The highlighted language below identifies Black Hills Energy's Blockchain Interruptible Service requirements for interruptible service that is detailed in individual service contracts.[10]

[9] For more information on demand, please visit; https://www.santeecooper.com/rates/understanding-your-demand/#:~:text=Ratchet%20%E2%80%93%20A%20ratchet%20charge%20is,work%20and%20is%20being%20lost..
[10] Cheyenne Light Fuel and Power Company d/b/a Black Hills Energy, Electric Rates Blockchain Interruptible Service: https://ir.blackhillscorp.com/static-files/5c33d769-2d19-43f8-8898-a37af25481ef#:~:text=This%20tariff%20is%20applicable%20to,Agreement")%20with%20the%20Company.

ELECTRIC RATES

BLOCKCHAIN INTERRUPTIBLE SERVICE ("BCIS")

The Agreement shall be in accordance with the provisions of this BCIS tariff and at a minimum shall include:

1. Electric service is for new interruptible load expected to be 10,000 kW or greater;
2. A term of at least two (2) years;
3. Specific pricing for all electricity purchased, with the pricing terms being subject to renegotiation at least every three (3) years;
4. Identification of Customer and Company costs for any required new electric infrastructure;
5. Details specifying how service will be interrupted by the Company;
6. Negotiated service interruption provisions (size of interruptible load, notice of planned interruption, duration of interruption, and maximum hours of interruption per year) ;
7. BCIS customers that fail to interrupt service as required by the Agreement shall be responsible for all costs incurred by the Company due to such failure;
8. A release of liability of the Company for any losses or damages, including consequential damages, caused by or resulting from any interruption of service;

With the level of some proposed data centers' load being equivalent to 50% or more of an entire system's load, utilities and their systems would benefit from having a tariff that allows for interruptible service, either through a formalized tariff or on a case-by-case basis, which can be negotiated with or without a special contract. As these loads are large and unique compared to past loads, having a flexible interruptible tariff will likely allow a utility to accommodate customers while accounting for risk and available system capacity. Not one of the tariffs reviewed identified the maximum or minimum level of load that can be interruptible, rather the tariffs required the service agreement to identify the level of firm load, or the amount of demand that cannot be interrupted. Some contracts did include a maximum number of hours or interruption events; however, it is not necessary to establish a maximum number of hours or event durations within the tariff. This can be negotiated based upon the load and customer. For transparency and fairness purposes, the utilities may want to disclose in the tariff the compensation for interruptible service.

Commercial and industrial ("C&I") demand response and interruptible load programs are typically more cost-effective than residential demand response programs. Depending on program saturation, C&I can provide a more significant shed load than a residential program due to a higher level of load per customer.

It is important that pricing of interruptible and demand response efforts be done in moderation, with enough incentive to the ratepayer to offset the inconvenience of shedding load and reducing activity, but not too high as to incentivize high profitability from shedding load as it can be costly to other ratepayers. Pricing structure, limitations on overall hours of interruption, and having the utility determine when an interruptible or demand response event occurs can eliminate concerns related to profitability. Compensation for demand response efforts should be considered based upon the level of load that can be shed and how quickly the load can respond to a request. Commercial and industrial customers, depending on their industry, can typically shed higher amounts of load and in a short period of time (within 30 minutes to an hour). The ability to provide large amounts of load shedding quickly should be compensated appropriately to encourage customers to do so when necessary. Demand response or interruptible tariff compensation for load shedding should be compared among similar rate classes and rate design elements, such as number of hours and events and duration of the event. These factors, along with the need for capacity in a service territory, can influence the level of compensation offered for demand shedding.

Interruptible tariffs can have several elements to establish safeguards for the grid and to ensure that load reductions do occur. In Texas, there have been capacity issues when an interruptible service client does not respond to the request to reduce load. Some provisions that can be included in an interruptible service agreement include:

- Number of annual events and total hours. The number of events and overall hours for interruption per year should not be detrimental to the business.
- Event duration and seasonal requirements. There may be periods of time when demand reduction is more valuable than others, depending on the utility's peak season. This can influence the length of events, typically around two to four hours, and the timing of the events.
- Details of compensation that could be based on the level of demand or energy reduction, such as the dollar per megawatt, or could be offered through a discounted energy price throughout the year for participating.
- Penalty for not responding to an interruption event. The utility is relying on the reduction in load; however, if a customer does not respond, it can increase energy costs for others. Therefore, a penalty should be assessed to offset that increase in cost for not responding to the event and to encourage customer participation.

Investment Requirements and Cost Assignment

One way to limit risk to existing ratepayers from the addition of the customer's load is to assign costs to the customer, require contributions in aid of construction for system upgrades, and require surety bonds or minimum bills equivalent to a portion of the annual bill. These safeguards can lessen the risk to ratepayers by requiring the customer to be invested in the location. Assignment of costs for new or expanded electric service is not a new concept. Customers, both residential and commercial, can be responsible for line extensions and other identified costs to receive service. Cost assignments should be designated in the tariff, including guidelines on how to calculate the minimum bill.[11]

Depending on the size and characteristics of the load, there is potential for other customers throughout the service territory subsidizing the cost of service for a large load customer, particularly when discounted rates are provided to the large load customer. One way to avoid subsidization for a particular customer is to evaluate if the revenues received from the large load customer exceed the cost to serve the customer. An example of this is Evergy Missouri Metro's Special High-Load Factor Market Rate ("Schedule MKT"), noted in Table 1 below, which requires the utility to track all costs to serve each customer under this tariff and verify that the revenue collected is higher.[12] This provision is designed to ensure that non-Schedule MKT customers are not held liable for any deficiencies in revenues or from stranded investment or costs from serving the customer over the length of the contract. To track the costs and revenues associated with this, the tariff outlines the following:

- Utility must identify costs and revenues with each customer on the Schedule MKT in its books and records.
- During a rate proceeding, the portion of the revenue requirement associated with the costs to serve the customer shall be assigned to the customer and not the overall customer base.
- If the customer's rate revenues do not exceed the cost to serve the customer in the customer's revenue requirement, there must be an additional revenue adjustment to cover the shortfall in a true-up period.
- The customer served by Schedule MKT can argue whether a specific quantifiable societal or other benefit (e.g., added jobs or tax revenue) should be considered to offset the deficiency.

Concern: The cost assignment concerns are not only limited within a service territory but also across state lines for transmission infrastructure. In April, the Federal Energy Regulatory Commission ("FERC") approved a regional cost assignment for the PJM. The transmission upgrades are being implemented to support a cluster of data centers in northern Virginia. While the location of the data centers is in Virginia, ratepayers in Maryland have been assigned 10%

One example of a cost assigned could be for a feasibility study. As large new loads are requested on an electric system, a feasibility study is usually conducted to understand what system upgrades may be needed to accommodate the load safely, depending on size thresholds, including transmission and distribution upgrades.[13] Sometimes, the tariff includes

[11] Source for orange box: *Utilities poised for datacenter earnings boost, want clarity on cost recovery,* Allison Good, April 18, 2024, https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/utilities-poised-for-datacenter-earnings-boost-want-clarity-on-cost-recovery-81249390, accessed October 18, 2024.

[12] Evergy Missouri Metro's Special High-Load Factor Market Rate Schedule MKT can be found here: https://www.evergy.com/-/media/documents/billing/missouri/detailed_tariffs_mo/special-high-load-factor-market-rate.pdf

[13] Requirements for a feasibility study is dependent upon the service territory and the jurisdiction.

a provision that assigns the cost of the feasibility study on the customer, like in New York, which is shown below.[14]

RIDER A

RATES AND CHARGES FOR CUSTOMERS REQUESTING HIGH DENSITY LOAD ("HDL") SERVICE

B. APPLICATION FOR SERVICE:

b. Upon payment of security acceptable to the Utility, the Utility shall conduct, or cause to be conducted a feasibility study to evaluate whether the requested load can be safely served by the Utility.

a. The feasibility study will identify what, if any, upgrades to the Utility's facilities are required to serve the customer.

B. CUSTOMER COST CONTRIBUTION

A Customer requesting service under this Rider will be responsible for:

a. reasonable costs of conducting the feasibility study; and

If the system can accommodate the load with minimal system upgrades, the risk associated with the customer's electric service is likely limited. However, if significant upgrades are required, then those costs serve as potential risks to existing ratepayers. The cost for the feasibility study should be assessed to the customer seeking interconnection; sometimes this is done through a flat fee. Furthermore, the charges associated with upgrades, including the proportional cost of acquiring or building new generation to serve the customer, should be required to be funded by the customer and tied to a deposit or contribution in aid of construction, to limit risk exposure of stranded assets to the existing customer base.

Historically, a large load facility, like an Amazon warehouse or industrial process, is more permanent and will contribute towards cost recovery immediately, as the plant ramps up in its first year of operation and then will remain on the system for the foreseeable future. On the contrary, cryptocurrency mining facilities are seen as volatile as they are price sensitive and can be operated in non-permanent facilities, and traditional data centers can take years

[14] See Leaf 95-96 of Rider A Rates and Charges for Customers Requesting High Density Load ("HDL") Service, https://ets.dps.ny.gov/ets_web/search/showPDF.cfm?%3B%3AIS%20%3B%2A%29LOUNWD%5CJ%5E8%2B"%2B5%2F0MD%2F0%28%231V%28S<WX%0A, accessed November 11, 2024.

to get to full capacity, which can delay cost recovery and place the burden on existing ratepayers.

A definition and summary of how each requirement serves as a safeguard is provided in Table 1 below. In addition, each requirement has an example and is linked to one of the tariffs discussed in Appendix A.

Table 1 Common Tariffs Requirements

Requirement	Definition	Serves as a Safeguard?	Example
Contract Term Length	Length of the service agreement. It can be limited to a minimum and/or maximum number of years. In addition to a contract term, there could be a term length for pricing terms.	Yes. A limited term could limit potential risk to customers, as well as ensure that system upgrades or investment in new generation are paid for by the new customer rather than existing ratepayers.	Evergy Missouri Metro limits contract lengths to 10 years, with pricing terms no more than 5 years
Minimum Demand	Level of demand needed to qualify for the tariff	Yes. Provides a threshold for customers to qualify for the tariff and can be designed to target high demand users	Contracts varied significantly between 500 kW and 100,000 kWh per month. This will be dependent on the service territory's load compared to the new customer load.
Minimum Load Factor	Average power usage compared to peak power usage during the same period. The higher the percentage, the more efficiently the electricity is being used.	Yes. Establishing a penalty for not achieving a minimum load factor will encourage the customer to have energy usage consistent with its maximum peak. Smoothing out peaks can lower the strain on power infrastructure and reliability.	If required, the minimum load factor required was 85%. The reduces the opportunity for significant fluctuations in load and thus the reliability of service is more easily predictable by the utility.

Requirement	Definition	Serves as a Safeguard?	Example
Range for Power Factor	Effectiveness of incoming power by a specific load (or equipment) at a given time. The higher the power factor, the more efficient the load's performance.	Yes. Inefficient power usage can result in additional costs on the system. Establishing a power factor range can reduce energy costs, reduce system losses, and improve voltage regulation, which can limit outages and allow for additional loads to be added to the system from that customer.	If required, this would be 90% or greater. The Montana-Dakota Utilities Company requires a power actor between 97% lagging and 97% leading.
Requirements for Investment	Designated cost elements that are funded directly by the new customer, sometime viewed as a deposit in the form of Contributions in Aid of Construction ("CIAC"), bonds, or actual payments. This investment may be returned to the customer overtime.	Yes. Delineating expenses for the customer to pay or cover with a deposit eliminates concerns about discriminatory rates. Additionally, it encourages investment by the new customers, thus removing the risk from existing ratepayers, and ensures a term commitment to the service territory.	This requirement varied by utility, but could include new electric infrastructure, line extension or system upgrades, and feasibility studies. Other utilities require bonds for Value of Lost Load dependent upon the RTO requirements or a bond for the average bill for a time period.
Cost Assignment	Designation of which expenses related to providing service to the customer is the responsibility of the customer and not socialized to other ratepayers.	Yes. Eliminates the risk of a customer not paying their fair share of the investment in providing electric service. Some commissions have required utilities to track all costs related to the customer to ensure during rate cases that the revenues from the customer offset expenses to provide service to the customer.	Evergy Missouri Metro has a requirement to track all costs to serve the customer and verify that revenue collected is higher. The New York Municipal Power Agency requires costs associated with the purchased power adjustment and rate statement to be allocated to the customer.

Requirement	Definition	Serves as a Safeguard?	Example
Requirement to Shed Load	Utility requires the customer to drop a portion of its load during events with notice.	Yes. Increases system reliability and reduces capacity costs, depending on the type of event requiring load shedding. This could be done through an interruptible service rider, service agreement, or a formal demand response program.	Approximately half of the tariffs have a load shed requirement. The majority vary by contract. If there is an interruptible schedule, the customer is typically not subject to a demand response program. If there is not an interruptible program, then demand response programs were often, but not always available. Grant County Public Utility District does not offer an interruptible tariff or a demand response program through tariffs but does do arrangements on a customer-by-customer basis.
Load Subject to Interruptible	Can be a determined capacity subject to interruptible service (such as non-firm demand) or the amount of time when an interruption event may be announced.	Yes. While the tariff language can indicate a cap on the level of interruptible load to be included or excluded, it is recommended that the level of load be negotiated on a per customer basis.	For those requiring interruptible load, the amount of load subject is established in the contract with the customer. It is often limited to non-firm demand.
Maximum Hours of Interruptible per Year	A defined limitation on the number of hours that load can be interrupted per year. This is typically accompanied by penalty language in the event the customer does not respond to the interruptible load request.	Yes. Designating a maximum number of events or hours, or even length of events, can encourage participation from customers in an interruptible schedule.	There is a significant range in the number of hours, if any were specified in the tariff. Entergy Arkansas limits the maximum number of hours to 40 or 80 hours, depending on notice time, while other utilities such as Idaho Power Company set limits of 225 hours per year.

2024 Proposed Large Load Tariffs

Ohio

In Ohio, there are opposing opinions between the utility, AEP Ohio, and the technology giants like Amazon, Google, Meta, as well as the Data Center Coalition on the structure of large load tariffs. In July 2024, AEP Ohio, in its role as a distribution utility, proposed two new tariff designs as a result of an influx of data center load requests in its service territory in May 2024.[15] The initially-proposed tariff included two components, a Data Center Power designed for customers with a monthly demand of 25 MW or more, and a second Mobile Data Center component for cryptomining facilities with a monthly demand of 1 MW or greater.[16]

As of January 2025, there were two competing settlements that diverged substantially from the initial proposal, and the case is still pending before the Ohio Public Utilities Commission, with hearing dates in December 2024 and January 2025.[17] Depending on the decision in the case, it could set precedent and baseline safeguards throughout the nation as the filing's proposed terms have not been collectively included in any other utility tariffs for data centers.

The primary components of the initial proposal were changes to an existing rider, known as the Basic Transmission Cost Rider ("BTCR").[18] Currently the BTCR sets the minimum demand charge for a customer at 60% of the contracted capacity. AEP Ohio's initial proposal indicated that the amount was too low and sought to increase the minimum demand charge to 90-95% of the contracted demand. This is due to the significant difference for large load customers between the minimum and actual bill if all contracted load is utilized. In addition, AEP Ohio initially requested that data centers enter into 10-year service contracts to ensure funding for the significant investment that the utility will need to make over the next decade to accommodate the data center load interconnection requests. An exit fee was proposed for customers in the 10-year contract to pay to leave the contract after 5 years. As noted in the safeguard above, AEP Ohio is implementing elements to provide safeguards not only for ratepayers but also for the utility itself as it endeavors to grow the system. If the data centers are not located in the service territory after AEP Ohio builds out the transmission system, the unneeded capacity costs will be passed along to ratepayers located throughout PJM.

---

[15] Application for approval of New Tariffs By Ohio Power Company, *In the Matter of the Application of Ohio Power Company for New Tariffs Related to Data Centers and Mobile Data Centers,* Case No. 24-508-EL-ATA, https://dis.puc.state.oh.us/ViewImage.aspx?CMID=A1001001A24E13B42822J00948.

[16] Direct testimony of Matthew S McKenzie on behalf of Ohio Power Company, *In the Matter of the Application of Ohio Power Company for New Tariffs Related to Data Centers and Mobile Data Centers,* Case no. 24-508-EL-ATA, tariff pages begin on page 32, https://dis.puc.state.oh.us/ViewImage.aspx?CMID=A1001001A24E13B43247C00950.

[17] Full docket available at: https://dis.puc.state.oh.us/CaseRecord.aspx?CaseNo=24-0508

[18] Direct testimony of Matthew S McKenzie on behalf of Ohio Power Company, *In the Matter of the Application of Ohio Power Company for New Tariffs Related to Data Centers and Mobile Data Centers,* Case no. 24-508-EL-ATA, tariff pages begin on pages 15-16, https://dis.puc.state.oh.us/ViewImage.aspx?CMID=A1001001A24E13B43247C00950.

Provisions within the initially-proposed tariff that can serve as safeguards for ratepayers are summarized below:

- Minimum Load Eligibility
    - Tariff is applicable to customers requesting a minimum demand of 25 MW of service at a single location. The tariff would also be applicable to a parent company with multiple data centers that have an aggregate monthly maximum demand greater than 25 MW within a 24-month period.
        - By establishing aggregate demand for parent companies, this ensures that data centers locating around the service territory are not circumventing the eligibility requirements for the tariff.
- Minimum Billing Demand
    - Load ramp period which establishes monthly peak load requirements as the facility comes online and requires that the overall requested load of the facility commence service within three years. During the ramp up period, billing demand shall not be less than 90% of the customer's load ramp contract capacity.
        - This ensures that the fixed costs associated with serving this customer's level of load are paid for by the customer. Even if the customer has not reached that level of demand, the utility is already incurring the cost to provide services at the contracted demand levels.
    - Monthly billing demand once a customer is beyond the load ramp period shall not be less than 90% of the greater of (a) customer's contracted capacity or (b) customer highest previously established monthly billing demand during the past 11 months.
        - The inclusion of a demand ratchet ensures the customer is paying the fixed charges associated with this customer's demand.
- Range for Power Factor
    - Includes an excess reactive demand charge, assessed for each kVAR of reactive demand, leading or lagging, in excess of 50% of the metered demand.
        - This ensures that the customer is paying its fair share of the fixed charges to provide service, as it is based on the level of capacity contracted and not used.
- Retail Supplier Notice
    - If a customer wants to switch from standard offer service to a competitive supplier, then the customer must provide the utility with notice 60 days prior to the end of the supply period covered by the auction. The customer must remain on standard offer service for the six month period in which the customer has been receiving standard offer service.
        - This ensures that the utility does not over procure energy through the supply auctions.
- Contract Period
    - The initial contract period cannot be less than 10 years, including the load ramp period. There is an exit fee, equal to the minimum charges for 36 months after the notice of the termination, if the customer elects to leave after the completion of the 5th year of the contract.

        - The contract term is the average contract length and has an exit fee schedule that is designed to avoid stranded asset costs.
- Collateral Requirements
    - Customers must meet a credit and cash collateral requirement relative to 50% of the total minimum charges for the full contract term. The amount of collateral is reduced by one year's minimum charges for each year the customer is energized and makes on-time electric service payments. If the financial position of the customer changes over the term of the contract, the Company may ask for updated information and re-evaluate the collateral requirements.
        - This provision is unique compared to others reviewed, as the collateral is for the full contract term and the reduction of the collateral is based upon timely payments. Furthermore, the collateral provisions are typically calculated ahead of the contract signing and do not have re-evaluation requirements. This last provision would be useful as the industries related to cryptomining and data centers are ever evolving and dependent on a number of factors, such as contracts and price signals.
- Demand response
    - The initially proposed contract lacks a provision related to interruption outside of a requirement for the customer to reduce its demand during an RTO- or company-declared emergency event. There is a lack of detail related to the emergency events and no mention of voluntary interruptible events. While it is important to be able to react to emergency events, given the size of the loads anticipated, the ability to interrupt load for reliability purposes, particularly to address local reliability issues, would be of significant benefit to the system. While it may not be a standard provision, this tariff should have a special contract provision to determine interruptible load levels from large load facilities.

As noted above, as of this publication date, the case was ongoing with a multi-day hearing held on many of the issues covered above.

Indiana

On November 22, 2024, Indiana Michigan Power Company (I&M) introduced a settlement, involving all parties to the case including tech giants Amazon and Google and the Indiana Office of Utility Consumer Counselor, to amend their industrial power tariff.[19] This tariff is applicable to new or expanded facilities seeking to contract capacity of 70 MW or more or 150 MW of aggregated load across a company. Loads meeting this requirement are required to

[19] Before the Indiana Utility Regulatory Commission, *In the Matter of Verified Petition of Indiana Michigan Power Company for Approval of Modifications to its Industrial Power Tariff – Tariff I.P.*, Cause No. 46097, filed November 22, 2024, https://iurc.portal.in.gov/_entity/sharepointdocumentlocation/4aae5d78-18a9-ef11-8a6a-001dd80bd98a/bb9c6bba-fd52-45ad-8e64-a444aef13c39?file=46097_IndMich_Submission%20of%20Unopposed%20Settlement%20Agreement%20and%20Unopposed%20Motion%20for%20Acceptance%20of%20Out%20of%20Time%20Filing_112224.pdf.

have initial contracts of at least 12 years. The contract for the full load can start after a five-year ramp up period. Additionally, without incurring any fees, after the first five years of the contract, a customer can reduce its contract capacity by up to 20 percent, as long as the customer notifies I&M through written notice 42 months prior to the start of a PJM Interconnection delivery year. Contracts can be terminated, or contract capacity can be reduced beyond 20%, if an exit fee is paid and done so under the conditions listed above for reduced capacity.

In addition to these contract terms, the I&M settlement put forth several provisions related to I&M's integrated resource planning ("IRP"), interconnection, demand response, and clean tariffs. As part of its IRP, I&M has agreed to study grid enhancing technologies and tools to maximize the transmission grid efficiency and to relay the study's result in the next IRP. I&M also agreed to discuss any changes to its interconnection process with stakeholders, including large load entry requirements to the utility's queue, interconnection requirements, and load ramping requirements. To address emergency load reduction plans, I&M will meet with the parties to the settlement to discuss emergency response procedures and demand response opportunities for customers under this tariff. Finally, I&M agreed to collaborate with settling parties to develop a clean transition tariff proposal that will allow participants to support investment in carbon-free resources and ensure that all program costs are covered by participants and remain consistent with the five pillars in Indiana Code §8-1-2-0.6.

As part of the agreement, beginning six months after approval, I&M would provide semi-confidential reports to the Indiana Utility Regulatory Commission on new and pending large load customers. The settlement, which as of the publication of this report, has not been approved yet by the Commission,[20] also requires Amazon Web Services, Microsoft, and Google to each give $500,000 annually, for five years, to the Indiana Community Action Association, which supports low-income individuals in Indiana.

North and South Carolina

In North and South Carolina, Duke Energy has several initiatives they have proposed or adopted to address the growing demand from high energy users, including from data centers.

New rates for Data Centers and Industrial Customers

Duke Energy conducted a study which evaluated ways that high-volume users could pay their fair share into the system. The reason behind the focus has to do with the constrained power supply on their system compared to a few years ago. Duke is anticipating 18,000 gigawatt hours of additional load from new customers by 2028, with 25% of that load coming from data centers.[21] As a result of the study, Duke is adding electric supply contract terms for data centers and factories which require a minimum-take clause and upfront payments for infrastructure investments. The minimum-take clause requires qualifying customers to pay

[20] Full docket at https://iurc.portal.in.gov/docketed-case-details/?id=b8cd5780-0546-ef11-8409-001dd803817e

[21] *Duke Energy seeks take or pay power contracts for data centers,* Laila Kearney, May 7, 2024, https://www.reuters.com/business/energy/duke-energy-seeks-take-or-pay-power-contracts-data-centers-2024-05-07/ , accessed October 18, 2024.

for a certain amount of power regardless of actual use and requires upfront contributions for investment in system upgrades.

Clean Energy Tariff Options

In May 2024, Duke Energy signed memorandums of understanding with Amazon, Google, Microsoft, and Nucor to explore carbon-free energy generation and clean tariff options, called the Accelerating Clean Energy ("ACE") tariffs. The ACE framework includes a Clean Transition Tariff where Duke Energy would be able to provide commercial and industrial customers with new carbon-free energy options, while providing protection for non-participating customers and potentially lowering the long-term costs of investing in clean energy technologies.[22] The framework being proposed will occur in phases, with the purpose of helping customers meet their clean energy goals through tariff design and financing options.

One of those items that occurred outside of the framework included a green tariff proposal called the Green Source Advantage Choice Program, which was approved by the North Carolina Utilities Commission in July 2024. [23] The rider is offered to non-residential customers "who elect to direct the Company to procure renewable energy on behalf of the Customer's behalf" and who have a minimum maximum annual peak demand of 1 MW or an aggregated annual peak demand of 5 MW.[24] The tariff allows for large customers to increase Duke Energy's investment in solar energy by 150 MW per year, through a resource acceleration option in which customers can sponsor projects not selected in the company's annul competitive bidding process. The program limits procurement of renewables by the Duke Energy companies in North Carolina as follows:

- 4,000 MW of renewable energy from Duke Energy Carolinas ("DEC") and Duke Energy Progress ("DEP")
- DEP and DEC can only collectively own 2,200 MW of the capacity under this tariff
- The remaining 1,800 MW of renewable energy facilities must be developed by third parties that have entered into PPA's with one of the Companies or an eligible Green Source Advantage Choice customer.
- Annually, the Company must reserve 10% of the capacity for subscription by qualifying economic development customers. At the end of the third quarter each year, any unsubscribed economic development capacity can be released to all other qualified customers.

Some of the projections in place for the service territories customers include:

---

[22] *Responding to growing demand, Duke Energy, Amazon, Google, Microsoft, and Nucor execute agreements to accelerate clean energy options,* Duke Energy News Center, May 29, 2024, https://news.duke-energy.com/releases/responding-to-growing-demand-duke-energy-amazon-google-microsoft-and-nucor-execute-agreements-to-accelerate-clean-energy-options, accessed October 18, 2024.

[23] Docket Nos. E-2, SUB 1314 and E-7, SUB 1289, Before the North Carolina Utilities Commission, *In the Matter of Petition of Duke Energy Progress, LLC, and Duke Energy Carolinas, LLC, Requesting Approval of Green Source Advantage Choice Program and Rider GSAC,* Commission Order dated July 31, 2024, *https://starw1.ncuc.gov/NCUC/ViewFile.aspx?Id=acd1a9a3-9b00-4a3a-9700-4dae3a293cc2.*.

[24] Compliance tariff currently under review by the North Carolina Utilities Commission, Rider GSAC Green Source Advantage Choice, dated August 14, 2024, https://starw1.ncuc.gov/NCUC/ViewFile.aspx?Id=0d45934a-06ea-478d-8301-7a3b4377415a.

- Customers can pay for their portion of clean energy costs either through an up-front contribution in aid of construction payment or on their bill over time through a levelized demand charge payment.
- If a customer elects battery storage, the charging cost will be assessed as a charge to the customer and the discharging value will be assessed as a credit to the customer, effectively netting the amounts on the customer bill.

The docket for this item is ongoing and the tariff has not yet been approved by the Commission. Additionally, the overall ACE framework is an ongoing process that should continue to be monitored.

West Virginia and Kentucky

On July 18, 2024, Appalachian Power Company and Wheeling Power Company filed proposed revisions to its Schedules LCP and IP to include tariff terms related to the addition of customers with loads of 200 MW or greater in West Virginia.[25] On August 30, 2024, Kentucky Power Company filed revisions to its Tariff Industrial General Service ("Tariff I.G.S.") to address customers with loads of 150 MW or greater in Kentucky.[26] The initially-proposed changes to the tariffs were the same and include the following:

- Initial contract period of 20 years
- Either the customer or utility must provide at least five years' written notice to discontinue service of the terms of the schedule; however, this shall not reduce the 20-year initial contract term.
- If a permanent closure by the customers occurs in the first five years of the contract, the customer must pay a one-time exit fee equal to five years of minimum billing.
- A customer must provide written notice five years in advance to reduce the contract capacity by up to 20 percent of the contract capacity; however, mutual agreement can result in reduce contract capacity in less than five years.
- Demand ratchet requirement of no less than 90 percent of the greater of (a) the customer's on-peak contract capacity, or (b) the customer's highest previously established monthly billing demand during the past 11 months, or (c) the customer's maximum demand created during the billing month.
- Collateral is based upon creditworthiness of the customer. The collateral shall be equal to 24 times the customer's previous maximum monthly non-fuel bill.

---

[25] Before the West Virginia Public Service Commission, *In the Matter of Appalachian Power Company and Wheeling Power Company Application for Approval of Revisions to Schedules LCP and IP,* Case No. 24-0611-E-T-PW, https://www.psc.state.wv.us/scripts/WebDocket/ViewDocument.cfm?CaseActivityID=625853&NotType=WebDocket.

[26] Before the Kentucky Public Service Commission, *In the Matter of Kentucky Power Company's First Revised Tariff Sheet 1-1 (Index), First Revised Tariff Sheet 8-2 (Tariff I.G.S.), and Original Tariff Sheet 8-3 (Tariff I.G.S.),* Case No.2024-0830, https://psc.ky.gov/pscscf/2024%20cases/2024-00305//20240830_Kentucky%20Power%20Tariff%20Filing.pdf.

As of January 2025, this case is still pending before respective Commissions.[27] Notably, on January 22, 2025, the parties in the West Virginia proceeding filed a joint stipulation and settlement agreement signed by all parties. Under the terms of the settlement agreement, which is still pending approval, the large load tariff will apply to customers seeking to contract capacity of 100 MW or more or 150 MW of aggregated load across a company. Many of the settlement's terms mirror the terms of the Indiana settlement discussed above: for example, terms pertaining to minimum contract length, monthly billing demand, and reducing capacity during the contract period. The settlement also requires the utilities to track revenue and capital investments related to new large load customers, with the customers having the ability to seek confidentiality protections. The utilities, with input from the settling parties, must also conduct or utilize analyses to minimize transmission needs, but the cost of such analysis cannot exceed $50,000 pending further agreement.

Additional Considerations

Powering large loads from cryptocurrency mining and data centers is still evolving, which means there are changes announced monthly. In addition to reviewing the tariffs, several proceedings before public service commissions were reviewed to assess the fairness, reasonableness, and non-discriminatory elements of various contracts considered by public service commissions, in order to to better understand which safeguards have legal standing or precedent. Using the information from those proceedings and the tariffs discussed in the second section, there are additional rate provisions that should be considered when designing a large load tariff. These provisions will not only safeguard existing ratepayers, but also the efforts to achieve clean and renewable energy goals.

Avoid Discriminatory Rate Structures

As established by the Robinson-Patman Act, the Federal Trade Commission prohibits public service commissions from allowing unduly discriminatory rates. Public service commissions require approved rate structures to be just, reasonable, and non-preferential. While some commissions have approved tariffs that explicitly identify cryptomining and data centers, concerns regarding discriminatory rates and tariffs have been rising up throughout the states, as well at the federal level.

To avoid discriminating against certain industries, tariffs can include definitions and categories of service that can be related to the volatile and non-permanent nature of cryptomining and data centers.

Rather than explicitly naming cryptomining or data centers, utility tariffs have used the following definitions for high density tariffs:

- "Load that is portable and distributable"
- "High energy use density"
- "High variable load growth or load reduction"
- "permanency of service cannot be reasonable assured"
- "Evolving Industry"

[27] Joint Stipulation and Agreement for Settlement, Case No. 24-0611-E-T-PW, filed Jan. 22, 2025, https://www.psc.state.wv.us/scripts/WebDocket/ViewDocument.cfm?CaseActivityID=634939&NotType=WebDocket.

Black Hills Energy in Colorado offers a service tariff for "Indeterminate Service," which is defined below.[28]

> **BLACK HILLS COLORADO ELECTRIC, LLC**
> d/b/a BLACK HILLS ENERGY
>
> K. Indeterminate Service: Service that is of an indefinite or indeterminate nature where the amount and permanency of service cannot be reasonably assured in order to predict the revenue stream from applicant. For purposes of uniform application, "Indeterminate Service" may include such service as may be required for the speculative development of property, mobile buildings, mines, quarries, oil or gas wells, sand pits and other ventures that may reasonably be deemed to be speculative in nature.

In the Grant County Public Utility District ("PUD") service territory, in Washington, rather than adopting a tariff explicitly for cryptomining facilities and volatile users, the PUD adopted a new rate class, known as "evolving industries." Rather than explicitly call out specific users, it defined characteristics that those industries are known for. The definition of Evolving Industries rate class is based on three risk factors as shown below.[29] This rate class is charged a different rate than other C&I customers.

> To decide if an industry falls into the evolving industries class, the district used a test focused on certain risk factors presented by the industry in question. These risks are:
>
> - Regulatory risk — risk of detrimental changes to regulation with the potential to render the industry inviable within a foreseeable time horizon;
> - Business risk — potential for cessation or significant reduction of service due to a concentration of business risk in an evolving or unproven industry or in the value of the customer's primary output; and
> - Concentration risk — potential for significant load concentration within the district's service territory resulting in a meaningful aggregate impact and corresponding future risk to the district's revenue stream. Evaluation would begin to occur when industry concentration of existing and service request queue customer loads exceeds 5% of the district's total load.

[28] Black Hills Colorado Electric LLC d/b/a/ Black Hills Energy tariffs: https://www.blackhillsenergy.com/sites/blackhillsenergy.com/files/coe-rates-tariff.pdf, see PDF page 220.
[29] *A Blow to Crypto Miners Disputing Local Energy Rates,* James Gatto and Andrew Mina, April 10, 2020, https://www.sheppardmullin.com/media/publication/1859_A%20Blow%20To%20Crypto%20Miners%20Disputing%20Local%20Energy%20Rates.pdf, accessed October 18, 2024.

Renewable Energy Requirements

To date, most tariffs related to cryptomining and data centers do not have renewable energy or clean energy procurement requirements. Most efforts to have clean energy used to power these services are achieved through renewable energy credits pushed by a corporate goal rather than from a utility. Of the tariffs and proceedings reviewed, only one had an explicit renewable energy provision. Renewable energy requirements or clean energy tariffs should be designed in accordance with the "three pillars" of clean energy:

1. Incremental – energy is from a clean energy source that incremental to existing generation.
2. Temporality or being time-matched – power is generated in the same hour it is consumed.
3. Deliverable – power is deliverable in the same grid region.

In the Evergy Missouri Metro service territory, customers are subject to the Renewable Energy Standard Rate Adjustment Mechanism ("RESRAM") charge, which is an adjustable rate to allow for the utility to recover prudently-incurred costs related to procurement of renewable energy standard costs that are above and beyond the renewable energy costs already included in base rates. The provision included below states that a customer on Schedule MKT must pay future RESRAM charges unless they have renewable attributes that support its load which are greater than or equal to the existing Renewable Energy Standard.[30] As written, the provision rewards customers under this tariff if they are procuring renewable attributes on their own. Please note that the provision does not require actual investment in renewable energy resources to directly serve the load.

Special High-Load Factor Market Rate
Schedule MKT

6. A Schedule MKT Customer shall be subject to any future RESRAM charges imposed by Evergy Metro unless a Schedule MKT customer does have renewable attributes supporting its load greater than or equal to the then existing Renewable Energy Standard including any solar portfolio requirements. For Schedule MKT customers with renewable attributes supporting its load greater than or equal to the then existing Renewable Energy Standard, including any solar portfolio requirements, the MKT Customer's entire load will be subtracted from the calculation of total retail electric sales in in 20 CSR 4240-20.100. Renewable attributes means Renewable Energy Credits and solar Renewable Energy Credits that the MKT Customer has retired, or had retired on its behalf, documented annually from an established renewable registry.

While renewable energy credits are a step in the right direction, it is essential to include provisions to require data centers to invest in renewable energy in the surrounding community, either through investment in community solar, wind, roof top solar, and storage. Adding significant levels of load in communities, particularly those with clean energy targets,

[30] Evergy Missouri Metro Special High-Load Factor Market Rate Schedule MKT, https://www.evergy.com/-/media/documents/billing/missouri/detailed_tariffs_mo/special-high-load-factor-market-rate.pdf.

can derail clean energy achievements to date and could potentially result in increased environmental and health impacts due to increased generation needs. One of the three pillars of clean energy is incrementality. To achieve this, data centers must work to accelerate achievement of clean energy goals and/or offset any additional load powered by fossil fuel power plants. Utilities should work with potential customers to identify avenues to support the growth of renewable energy generation. For example, Meta worked with the Tennessee Valley Authority ("TVA") to develop a green tariff that supports the development of solar energy across the service territory to support Meta's corporate energy goals.[31] Depending on the economic development provisions, the green tariff is likely driving investment in the nearby community.

The clean transition tariff proposed by NV Energy in Nevada and Google and currently before the Public Utilities Commission of Nevada is another example of having clean energy serving large loads. The proposed tariff would allow for Google to power one of its data centers by purchasing power that NV Energy buys from the 115 MW Corsac Station Enhanced Geothermal Project at a price slightly higher than that paid by NV Energy. The tariff design prevents impacts to other ratepayers and allows Google to operate towards its 24/7 carbon free energy goal by 2030.

Power Purchase Agreements

Data center and cryptomining facilities are working with power plant operators and markets to establish power purchase agreements ("PPAs") to procure low-cost power options.[32] A power purchase agreement is between the buyer and seller, where a buyer commits to purchase an agreed amount of electricity over an established period. PPAs require approval from a utility commission if they involve a regulated utility.[33] There are two types of PPAs, physical and prepaid. A physical PPA is when the buyer takes physical delivery of the electricity generated either onsite in a behind-the-meter arrangement or offsite at a pre-determined point on the grid. A prepaid PPA is when the buyer pays the discounted cost of the PPA upfront. There is also something known as a virtual PPA, which is not a PPA but rather a financial instrument for a contract for difference.[34] Ultimately, state and local regulations on retail choice and electricity franchises establish the type of PPAs that are available by state.

As noted in Texas and by a case being considered by FERC, PPAs could have negative implications for other ratepayers. In Texas, a cryptocurrency company purchased low-cost electricity behind-the-meter through a PPA, which means that the energy utilized by the

---

[31] More information on the green tariff is provided here: *Meta Partners with Silicon Rand for Seven New Solar Projects in Georgia and Tennessee,* December 15, 2022, https://www.siliconranch.com/stories/meta-partners-silicon-ranch-walton-emc-tva, accessed October 18, 2024.
[32] For more information on power purchase agreements, please see: *Customer Power Purchase Agreements,* United States Environmental Protection Agency, https://www.epa.gov/statelocalenergy/customer-power-purchase-agreements, accessed October 18, 2024.
[33] Wholesale power sales, which do not involve an end user, are within the purview of the Federal Energy Regulatory Commission.
[34] Virtual PPAs are considered a financial instrument and are regulated by the Securities and Exchange Commission.

PPA customer is not offered in the ERCOT market. During a heat wave in summer 2023, ERCOT issued a request for curtailment of power. In response, the cryptomining company, through its wholesale agreements, sold its power into ERCOT, making over $24 million on energy savings, more than three times the revenue it made from cryptomining the prior month.[35] Due to the load flexibility and price sensitivity of cryptomining, the facilities are able to game the system to create additional profits at a significant cost to ratepayers, who are less flexible to respond to demand pressures and are not compensated for doing so, as ERCOT does not currently offer residential demand response programs.

Another case where ratepayers may not benefit is for the interconnection service agreement ("ISAs") change for a facility to provide power to a co-located data center or mine. Currently, the 2,228-MW Susquehanna nuclear facility in Pennsylvania provides power to PJM as a baseload resource.[36] However, in March, Talen Energy, which owns the nuclear plant and had a cryptomining facility and data center on site, sold the data center to Amazon and planned to sell up to 980 MW of nuclear power to Amazon through a behind-the-meter power purchase agreement. In late November 2024, FERC denied the application.[37]

Economic Development

The potential for economic development through increased tax revenues and potential jobs from large load projects is intriguing and viewed as a positive element of potential load growth by politicians and utilities. However, the opportunities of increased tax revenue are often offset by state and local government tax credits used to entice certain industries or large loads to locate in a specific area. Additionally, utilities often offer discounted rates to large loads, which means that there is potential for existing ratepayers subsidizing that customer and lower potential tax revenue from the electric service. These discounts do not have to come from an economic development tariff, rather they can be supported by existing laws and incentives which provide these to new loads and entities building in certain areas.

The issue with economic development for cryptomining facilities and data centers is that they typically do not produce a substantial number of full-time equivalent jobs compared to the level of load added to the system. Furthermore, with the tax credits, there is limited net tax revenue being provided to the area.[38] As a result, the economic development discounts provided to customers result in limited to no benefits to the area and can expose those living in the area to added risks and increased bills, as previously identified.

---

[35] "Texas Leaders worry that Bitcoin mines threaten to crash the state power grid," Keaton Peters, The Texas Tribune, July 10, 2024, https://www.texastribune.org/2024/07/10/texas-bitcoin-mine-noise-power-grid-cryptocurrency/, accessed October 18, 2024.

[36] *Talen-Amazon interconnection agreement needs extended FERC review: PJM Market Monitor,* Ethan Howland, July 11, 2024, https://www.utilitydive.com/news/talen-amazon-interconnection-agreement-ferc-constellation-vistra/721066/, accessed October 18, 2024.

[37] https://elibrary.ferc.gov/eLibrary/filelist?accession_number=20241101-3061&optimized=false; https://www.utilitydive.com/news/ferc-interconnection-isa-talen-amazon-data-center-susquehanna-exelon/731841/

[38] Reference for the orange box text: *Protect SC Consumers From Data Center Costs,* Frank Knapp, South Carolina Daily Gazette, September 12, 2024, https://scdailygazette.com/2024/09/12/protect-sc-consumers-from-data-center-costs/, accessed October 18, 2024.

With the focus from politicians on attracting new industries, utilities may want to consider reviewing and revising their economic development riders that allow for discounted rates. One AEP utility, Indiana Michigan Power in Indiana, sunset its Economic Development Rider tariff and adopted its Economic Development Rider 2 tariff, which increased the level of minimum demand and the minimum number of full-time equivalent jobs and capital investment guidelines. A summary of the differences to qualify for a discounted rate through the Economic Development Rider 2 is provided below.[39]

INDIANA
Economic Development Rider 2 (EDR 2)

To qualify, a new or expanding business must meet the following minimum criteria:

New Customer Criteria

- Add 500 kW or greater to one metered account
- Create at least 20 full-time equivalent (FTE) jobs or make a capital investment of $2,000,000 or more at the service location.

Existing Customer Criteria

- Increase billing demand by 250 kW or more above the Average Billing Demand during the 12 months prior to the date of application on one metered account
- Achieve a score of 100 or greater using the following calculations:
    - Base Score = New FTEs created X 10 + Capital Investment / 10,000
    - Load Multiplier = Estimated Load Increase (kW) / Base Average Billing Demand (Maximum of 1.0)
    - Final Score = Base Score X Load Multiplier.

Customer Account Status	Final Score	Discount Percentage on Total Non-Fuel Bill				
		Years 1 - 4	Year 5	Year 6	Year 7	Year 8
New	--	12.0%	9.0%	6.0%	3.0%	0.0%
Existing – Higher	> 200	6.0%	4.5%	3.0%	1.5%	0.0%
Existing – Lower	100 - 200	4.0%	3.0%	2.0%	1.0%	0.0%

Siting with Generation

As part of large load facilities procuring low energy costs, some are locating themselves near the power sources to ensure availability of low-cost energy. Not only are consuming companies looking to site near low-cost generation, but so are utilities. Several coal power plants have been revived or experienced increased run time in order to support new large loads.

[39]Indiana Michigan Power, Indiana Economic Development Rider 2, https://www.aep.com/assets/docs/economic-development/IN-EDR-2023-App.pdf.

While there is an option to build new generation, co-locating the data center or cryptocurrency facility with an existing coal or gas plant slated for retirement or transition to a gas-fired plant can be an attractive energy source for larger users. This can result in increased greenhouse gas emissions and local air and water pollution due to smaller, less efficient plants being built or from the proliferation of coal-fired plants that may have difficulty with emission compliance. Additionally, while some large loads are considering nuclear power sources, there are concerns about capacity limitations and increased wholesale market prices if such power plants dedicate power directly to a customer rather than to the open market.

Including Projected Loads in Forecasts

Prospective data load centers and cryptomining facilities are seeking the best electricity rates and terms. This can result in utilities over-forecasting new load additions and capacity needs. Inclusion of the loads into utility forecasting needs a level of certainty as to whether a project will move forward or not, and sensitivity analyses need to properly account for the level of load that may not come to fruition. A utility's capacity planning cycles will likely never match up with discussions of potential customers' loads. Therefore, utilities should assess the likelihood of the load addition using elements such as where the new load is in the interconnection process, whether a feasibility study has been conducted, and whether the location has been procured, such as through a land sale/lease contract or local zoning approval.

Providing reasonable estimates of large new loads is extremely important, as it can require investment in not only new generating capacity, but also the transmission and distribution systems. If utilities utilize their planning processes, such as integrated resource planning ("IRP"), or a regional transmission operator does long-term planning of new transmission infrastructure, those entities could invest in capacity and grid system upgrades that end up not being needed if the large loads do not come to fruition. This results in existing customers footing a bill for stranded assets and less load and fewer customers to share those stranded assets costs across.

Mitigating over- and under-building of assets ultimately resides with the utilities and their planning models.[40] The planning models themselves need to not only account for customer load growth requirements over a long-term, but they also need to assess transmission and distribution opportunities and investments in distributed energy resources, such as energy efficiency, demand response, renewable energy, and storage. With all that said, there does not seem to be a utility or transmission operator that has established a process that can properly account for large load additions. For example, in 2023, Georgia Power submitted a one-year update to its 2022 IRP filing, indicating that the utility's demand increased by 20% by 2030 compared to the prior year's filing. There was significant uncertainty among the added load, particularly as to where this projected increase in demand was in the process of

---

[40] *Demand Better: How growing demand for electricity can drive a cleaner grid*, Jeremy Fisher, Laurie Williams, Dori Jaffe, Megan Wachspress, Sierra Club, September 2024, https://www.sierraclub.org/sites/default/files/2024-09/demandingbetterreportfinal_sept2024.pdf, p. 24, accessed October 18, 2024.

being interconnected. Transparency regarding potential new loads in the planning process—including the timing of the interconnection process and feasibility studies and ramp up of load over time—can be beneficial in ensuring sufficient investment in capacity.

Adequate Available Capacity

Kentucky Power's Economic Development Rider ("EDR") tariff requires there to be sufficient capacity to accommodate the increased or new load proposed by the customer. If sufficient capacity is not available, the cost of capacity to serve the new load must be passed on to the customer, by decreasing the discounted rate received by the customer. This provision is made to ensure that if capacity is needed to serve the load, that those costs are not passed on to the existing ratepayers. Not limited to EDRs, tariffs can include limitations on the level of load served by a certain tariff, such as Idaho Power Company's Schedule 20 Speculative High-Density Load.[41]

Tariff E.D.R.
(Economic Development Rider)

Terms and Conditions

(1) The Company will offer the EDR to qualifying customers with new or increased load when the Company has sufficient generating capacity available. When sufficient generating capacity is not available, the Company will procure the additional capacity on the customer's behalf. The cost of capacity procured on behalf of the customer shall reduce on a dollar-for-dollar basis the customer's IBDD and SBDD. Such reduction shall be capped so that the customer's maximum demand charge shall be the non- discounted tariff demand charge. The reduction will be applied in reverse chronological order

Conclusion

An ideal tariff will limit risk based upon the load being added to the system. There are several ways to achieve this and therefore, there is not one uniform set of safeguards that should be established. However, at a minimum, tariffs or special contracts should include the following:

1. For large loads under 50 MW, contract terms are not longer than 10 years, and loads larger than 50 MW should consider longer contract terms such as 12-20 years. Either contract term should come with pricing and negotiation terms set intermittently throughout the overall contract term.
2. Minimum or tiered monthly load requirements to qualify for the tariff.
3. Penalties for not maintaining a good load factor (typically 85% or greater) or power factor (typically 90% or greater). Examples of this are provided in Table 1 above.
4. Establish minimum demand charges or a demand ratchet to ensure that a large customer's fixed charges for peak demand levels are recovered.
5. Identification of costs that should be assigned to the customer or the requirement for a bond or deposit to offset the cost risk to existing ratepayers. Requirement of

[41] Idaho Power Company Schedule 20 Speculative High-Density Load: https://docs.idahopower.com/pdfs/AboutUs/RatesRegulatory/Tariffs/20.pdf.

contributions in aid of construction for any grid upgrades related directly to providing service will offset potential for stranded assets costs.

6. To ensure that the large load customer is not being subsidized by the service territory's other customers, the utility should track costs and revenues from the large load customer and assess a true up mechanism if the revenues do not exceed the customer costs.
7. An interruptible service requirement that can be negotiated between the utility and the customer. An interruptible service agreement should include the number of events and total annual hours, length of events, load reduction requirement, and penalty payment for failure to respond. It should also have term limits to allow for renegotiation.
8. Adequate available system capacity, with a requirement for procuring new capacity to be backed by the customer or through the purchase of renewable energy.

While these elements can be considered as part of any tariff related to serving large loads that may be considered volatile or a significant impact to the system, these terms will vary based upon the service territory's characteristics and current ratepayers.

In addition to establishing safeguards in tariffs, utilities need to put forward reasonable forecasts which consider whether large loads will move forward to interconnection. As part of those forecasts, utilities and IRPs should take into consideration how large loads can be served by a variety of services including transmission and distribution upgrades and investments in distributed energy resources. Using distributed energy resources such as solar, storage, and energy efficiency can also assist utilities and states to meet their environmental goals.

State	Utility	Document Type	Link	Contract Length	Minimum Demand	Minimum Load Factors	Range for Power Factor	Requirements for Investment
Wyoming	Cheyenne Light, Fuel and Power Company d/b/a Black Hills Energy	Special Contract	https://ir.blackhillscorp.com/static-files/5c33d769-2d19-43f8-8898-a37af25481ef#:~:text=This%20tariff%20is%20applicable%20to,Agreement")%20with%20the%20Company.	Min 2 years; renogotiation at least every 3 years	10,000 kW	N/A	N/A	New electric infrastructure, line extension or system upgrades
Arkansas	Entergy Arkansas LLC	Tariff	https://cdn.entergy-arkansas.com/userfiles/content/price/tariffs/eal_lphlds.pdf	N/A	N/A	N/A	N/A	Security deposit equal to 3 months of average estimated bill. Contributions in Aid of Construction for all network upgrades. Security Bond equal to Value of Lost Load Per MISO Schedule 28
Idaho	Idaho Power Company	Tariff	https://docs.idahopower.com/pdfs/AboutUs/RatesRegulatory/Tariffs/20.pdf	Special Contract required for over 20,000 kW	1,000 kW	N/A	90% or greater	Upgrades for interconnection facilities
New York	New York Municipal Power Agency	Tariff	https://ets.dps.ny.gov/ets_web/search/showPDF.cfm?%3B%3AIS%20%3B%2A%29LOUNWD%5CJ%5E8%2B"%2B5%2F0MD%2F0%28%231V%28S<WX%0A	N/A	>300 kW or load density exceeds 250/kWh/ft^2/year	N/A	N/A	Feasibility study, entire cost of new facilities necessary to supply requested service, cash deposit or Letter of Credit
South Dakota	Montana-Dakota Utilities Company	Tariff	https://puc.sd.gov/commission/Tariffs/Electric/mdu/Section3/20.pdf	3-5 years	10,000 kW	85%	Between 97% lagging and 97% leading	No
Washington	Grant County Public Utility District	Tariff	https://www.grantpud.org/templates/galaxy/images/Rate_Schedule_No_17.pdf	N/A	No minimum- separatedby greater or less than 200 kW	N/A	N/A	No
Indiana	Indiana Michigan Power	Tariff	https://www.aep.com/assets/docs/economic-development/IN-EDR-2023-App.pdf	N/A	500 kW	N/A	N/A	Create at least 20 full-time equivalent jobs or make a capital investment of $2 million or more at the service location, must apply and receive economic development assitance from the state, local government, or other public agency
Kentucky	Kentucky Power	Special Contract	https://psc.ky.gov/tariffs/Electric/Kentucky%20Power%20Company/Tariff.pdf	10 years	500 kW	N/A	N/A	N/A
Missouri	Evergy Missouri Metro	Tariff	https://www.evergy.com/-/media/documents/billing/missouri/detailed_tariffs_mo/special-high-load-factor-market-rate.pdf	No more than 10 years, with pricing terms no more than 5 years	100,000 kW/month or projected to be 150,000 kW within 5 years of being a new customer	85% or greater	N/A	N/A
North Dakota	Montana-Dakota Utilities Company	Tariff	https://www.montana-dakota.com/wp-content/uploads/PDFs/Rates-Tariffs/NorthDakota/Electric/NDElectric38.pdf	3-5 years	10,000 kW	85%	Between 97% lagging and 97% leading	N/A

State	Utility	Cost Assignment	Requirement to Shed Load	Load Subject to Interruptible Service	Maximum Hours of Interruptible Per Year	Demand Response	Requirement for Renewables or Traditional Generation	Requires Adquate Available Capacity	Notes
Wyoming	Cheyenne Light, Fuel and Power Company d/b/a Black Hills Energy	N/A	As defined in contract	As specified in contract	As specified in contract	No	No	N/A	
Arkansas	Entergy Arkansas LLC	N/A	Yes	Non-firm demand	40 or 80 hours	N/A	N/A	N/A	
Idaho	Idaho Power Company	N/A	Yes	Unclear	225 hours	N/A	N/A	Yes	
New York	New York Municipal Power Agency	Purchased Power Adjustment and Rate Statement	No	N/A	N/A	Not Offered	N/A	N/A	
South Dakota	Montana-Dakota Utilities Company	No	Yes	Specified in electric service agreement	200 hours	N/A	N/A	N/A	
Washington	Grant County Public Utility District	No	No	N/A	N/A	Customer by Customer Basis	N/A	N/A	Classified as an "Evolving Industry"
Indiana	Indiana Michigan Power	N/A	N/A	N/A	N/A	N/A	N/A	N/A	Economic Development Rider. Requires that the customer provide to the Company's satisfaction that absent the availability of the ridre, the new or increased demand would be located out of the Company's service territory or not place into service.
Kentucky	Kentucky Power	N/A	Yes	Specified in electric service agreement	N/A	N/A	N/A	N/A	Economic Development Rider
Missouri	Evergy Missouri Metro	Revenues must exceed costs	No	N/A	N/A	Special Interruptible Contract	A Schedule MKT Customer shall be subject to any future RESRAM charges imposed by Evergy Metro unless a Schedule MKT customer does have renewable attributes supporting its load greater than or equal to the then existing Renewable Energy Standard including any solar portfolio requirements.	N/A	
North Dakota	Montana-Dakota Utilities Company	N/A	Yes	Specified in electric service agreement	200 hours	N/A	N/A	N/A	

ATTACHMENT 4

Duke NICHOLAS INSTITUTE *for* ENERGY, ENVIRONMENT & SUSTAINABILITY

Rethinking Load Growth

Assessing the Potential for Integration of Large Flexible Loads in US Power Systems

Tyler H. Norris, Tim Profeta, Dalia Patino-Echeverri, and Adam Cowie-Haskell

Authors and Affiliations

Tyler H. Norris, Nicholas School of the Environment, Duke University
Tim Profeta, Sanford School of Public Policy and Nicholas Institute for Energy, Environment & Sustainability, Duke University
Dalia Patino-Echeverri, Nicholas School of the Environment, Duke University
Adam Cowie-Haskell, Nicholas School of the Environment, Duke University

Acknowledgments

The authors would like to thank Jessalyn Chuang and Wendy Wen for their research assistance.

Citation

Norris, T. H., T. Profeta, D. Patino-Echeverri, and A. Cowie-Haskell. 2025. *Rethinking Load Growth: Assessing the Potential for Integration of Large Flexible Loads in US Power Systems*. NI R 25-01. Durham, NC: Nicholas Institute for Energy, Environment & Sustainability, Duke University.
https://nicholasinstitute.duke.edu/publications/rethinking-load-growth

Cover image courtesy Gerville via iStock

Nicholas Institute for Energy, Environment & Sustainability

The Nicholas Institute for Energy, Environment & Sustainability at Duke University accelerates solutions to critical energy and environmental challenges, advancing a more just, resilient, and sustainable world. The Nicholas Institute conducts and supports actionable research and undertakes sustained engagement with policymakers, businesses, and communities—in addition to delivering transformative educational experiences to empower future leaders. The Nicholas Institute's work is aligned with the Duke Climate Commitment, which unites the university's education, research, operations, and external engagement missions to address climate challenges.

Contact

Nicholas Institute | Duke University | P.O. Box 90467 | Durham, NC 27708
1201 Pennsylvania Avenue NW | Suite 500 | Washington, DC 20004
919.613.1305 | nicholasinstitute@duke.edu

CONTENTS

INTRODUCTION

A New Era of Electricity Demand

Rapid US load growth—driven by unprecedented electricity demand from data centers, industrial manufacturing, and electrification of transportation and heating—is colliding with barriers to timely resource expansion. Protracted interconnection queues, supply chain constraints, and extended permitting processes, among other obstacles, are limiting the development of new power generation and transmission infrastructure. Against this backdrop, there is increasing urgency to identify strategies that accommodate rising demand without compromising reliability, affordability, or progress on decarbonization.

Aggregated US winter peak load is forecasted to grow by 21.5% over the next decade, rising from approximately 694 GW in 2024 to 843 GW by 2034, according to the *2024 Long-Term Reliability Assessment* of the North American Electric Reliability Corporation. This represents a 10-year compound annual growth rate (CAGR) of 2.0%, higher than any period since the 1980s (NERC 2024). Meanwhile, the Federal Energy Regulatory Commission's (FERC) latest five-year outlook forecasts 128 GW in peak load growth as early as 2029—a CAGR of 3.0% (FERC 2024b).

The primary catalyst for these updated forecasts is the surge in electricity demand from large commercial customers. While significant uncertainty remains, particularly following the release of DeepSeek, data centers are expected to account for the single largest growth segment, adding as much as 65 GW through 2029 and up to 44% of US electricity load growth through 2028 (Wilson et al. 2024; Rouch et al. 2024). Artificial intelligence (AI) workloads are projected to represent 50% to 70% of data center demand by 2030—up from less than 3% at the start of this decade—with generative AI driving 40% to 60% of this growth (Srivathsan et al. 2024; Lee et al. 2025).

Analysts have drawn parallels to the 1950s through the 1970s, when the United States achieved comparable electric power sector growth rates (Wilson et al. 2024). Yet these comparisons arguably understate the nature of today's challenge in the face of stricter permitting obstacles, higher population density, less land availability, skilled labor shortages, persistent supply chain bottlenecks, and demand for decarbonization and greater power reliability. While historical growth rates offer a useful benchmark, the sheer volume of required new generation, transmission, and distribution capacity forecasted for the United States within a condensed timeframe appears unprecedented.

The immensity of the challenge underscores the importance of deploying every available tool, especially those that can more swiftly, affordably, and sustainably integrate large loads. The time-sensitivity for solutions is amplified by the market pressure for many of these loads to interconnect as quickly as possible. In recent months, the US Secretary of Energy Advisory Board (SEAB) and the Electrical Power Research Institute (EPRI) have highlighted a key solution: load flexibility (SEAB 2024, Walton 2024a). The promise is that the unique profile of AI data centers can facilitate more flexible operations, supported by ongoing advancements in distributed energy resources (DERs).

Flexibility, in this context, refers to the ability of end-use customers to temporarily reduce their electricity consumption from the grid during periods of system stress by using on-site generators, shifting workload to other facilities, or reducing operations.[1] When system planners can reliably anticipate the availability of this load flexibility, the immediate pressure to expand generation capacity and transmission infrastructure can potentially be alleviated, mitigating or deferring costly expenditures. By facilitating near-term load growth without prematurely committing to large-scale capacity expansion, this approach offers a hedge against mounting uncertainty in the US data center market in light of the release of DeepSeek and related developments (Kearney and Hampton 2025).

Summary of Analysis and Findings

To support evaluation of potential solutions, this study presents an analysis of the existing US electrical power system's ability to accommodate new flexible loads. The analysis, which encompasses 22 of the largest balancing authorities serving 95% of the country's peak load, provides a first-order estimate of the potential for accommodating such loads with minimal capacity expansion or impact on demand-supply balance.[2]

Specifically, we estimate the gigawatts of new load that could be added in each balancing authority (BA) before total load exceeds what system planners are prepared to serve, provided the new load can be temporarily curtailed as needed. This serves as a proxy for the system's ability to integrate new load, which we term *curtailment-enabled headroom*.

Key results include (see Figure 1):

- 76 GW of new load—equivalent to 10% of the nation's current aggregate peak demand—could be integrated with an average annual load curtailment rate of 0.25% (i.e., if new loads can be curtailed for 0.25% of their maximum uptime)
- 98 GW of new load could be integrated at an average annual load curtailment rate of 0.5%, and 126 GW at a rate of 1.0%
- The number of hours during which curtailment of new loads would be necessary per year, on average, is comparable to those of existing US demand response programs
- The average duration of load curtailment (i.e., the length of time the new load is curtailed during curtailment events) would be relatively short, at 1.7 hours when average annual load curtailment is limited to 0.25%, 2.1 hours at a 0.5% limit, and 2.5 hours at a 1.0% limit
- Nearly 90% of hours during which load curtailment is required retain at least half of the new load (i.e., less than 50% curtailment of the new load is required)
- The five balancing authorities with the largest potential load integration at 0.5% annual curtailment are PJM at 18 GW, MISO at 15 GW, ERCOT at 10 GW, SPP at 10 GW, and Southern Company at 8 GW[3]

1 Note that while *curtailment* and *flexibility* are used interchangeably in this paper, *flexibility* can refer to a broader range of capabilities and services, such as the provision of down-reserves and other ancillary services.

2 For further discussion on the nuances regarding generation versus transmission capacity, see the section on limitations.

3 A complete list of abbreviations and their definitions can be found at the end of the report.

Overall, these results suggest the US power system's existing headroom, resulting from intentional planning decisions to maintain sizable reserves during infrequent peak demand events, is sufficient to accommodate significant constant new loads, provided such loads can be safely scaled back during some hours of the year. In addition, they underscore the potential for leveraging flexible load as a complement to supply-side investments, enabling growth while mitigating the need for large expenditures on new capacity.

We further demonstrate that a system's potential to serve new electricity demand without capacity expansion is determined primarily by the system's load factor (i.e., a measure of the level of use of system capacity) and grows in proportion to the flexibility of such load (i.e., what percentage of its maximal potential annual consumption can be curtailed). For this reason, in this paper we assess the technical potential for a system to serve new load under different curtailment limit scenarios (i.e., varying curtailment tolerance levels for new loads).

The analysis does not consider the technical constraints of power plants that impose intertemporal constraints on their operations (e.g., minimum downtime, minimum uptime, startup time, ramping capability, etc.) and does not account for transmission constraints. However, it ensures that the estimate of load accommodation capacity is such that total demand does not exceed the peak demand already anticipated for each season by system planners, and it discounts existing installed reserve margins capable of accommodating load that exceeds historical peaks. It also assumes that new load is constant throughout all hours.

This analysis should not be interpreted to suggest the United States can fully meet its near- and medium-term electricity demands without building new peaking capacity or expanding the grid. Rather, it highlights that flexible load strategies can help tap existing headroom to more quickly integrate new loads, reduce the cost of capacity expansion, and enable greater focus on the highest-value investments in the electric power system.

This paper proceeds as follows: the following section provides background on the opportunities and challenges to integrating large new data centers onto the grid. It explores how load flexibility can accelerate interconnection, reduce ratepayer costs through higher system utilization, and expand the role of demand response, particularly for AI-specialized data centers. We then detail the methods and results for estimating curtailment-enabled headroom, highlighting key trends and variations in system headroom and its correlation with load factors across regions. The paper concludes with a brief overview of key findings, limitations, and near-term implications.

BACKGROUND

Load Flexibility Can Accelerate Grid Interconnection

The growing demand for grid access by new large loads has significantly increased interconnection wait times, with some utilities reporting delays up to 7 to 10 years (Li et al. 2024; Saul 2024; WECC 2024). These wait times are exacerbated by increasingly severe transmission equipment supply chain constraints. In June 2024, the President's National Infrastructure Advisory Council highlighted that transformer order lead times had ballooned to two to five years—up from less than one year in 2020—while costs surged by 80% (NIAC 2024). Circuit breakers have seen similar delays: last year, the Western Area Power Administration

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

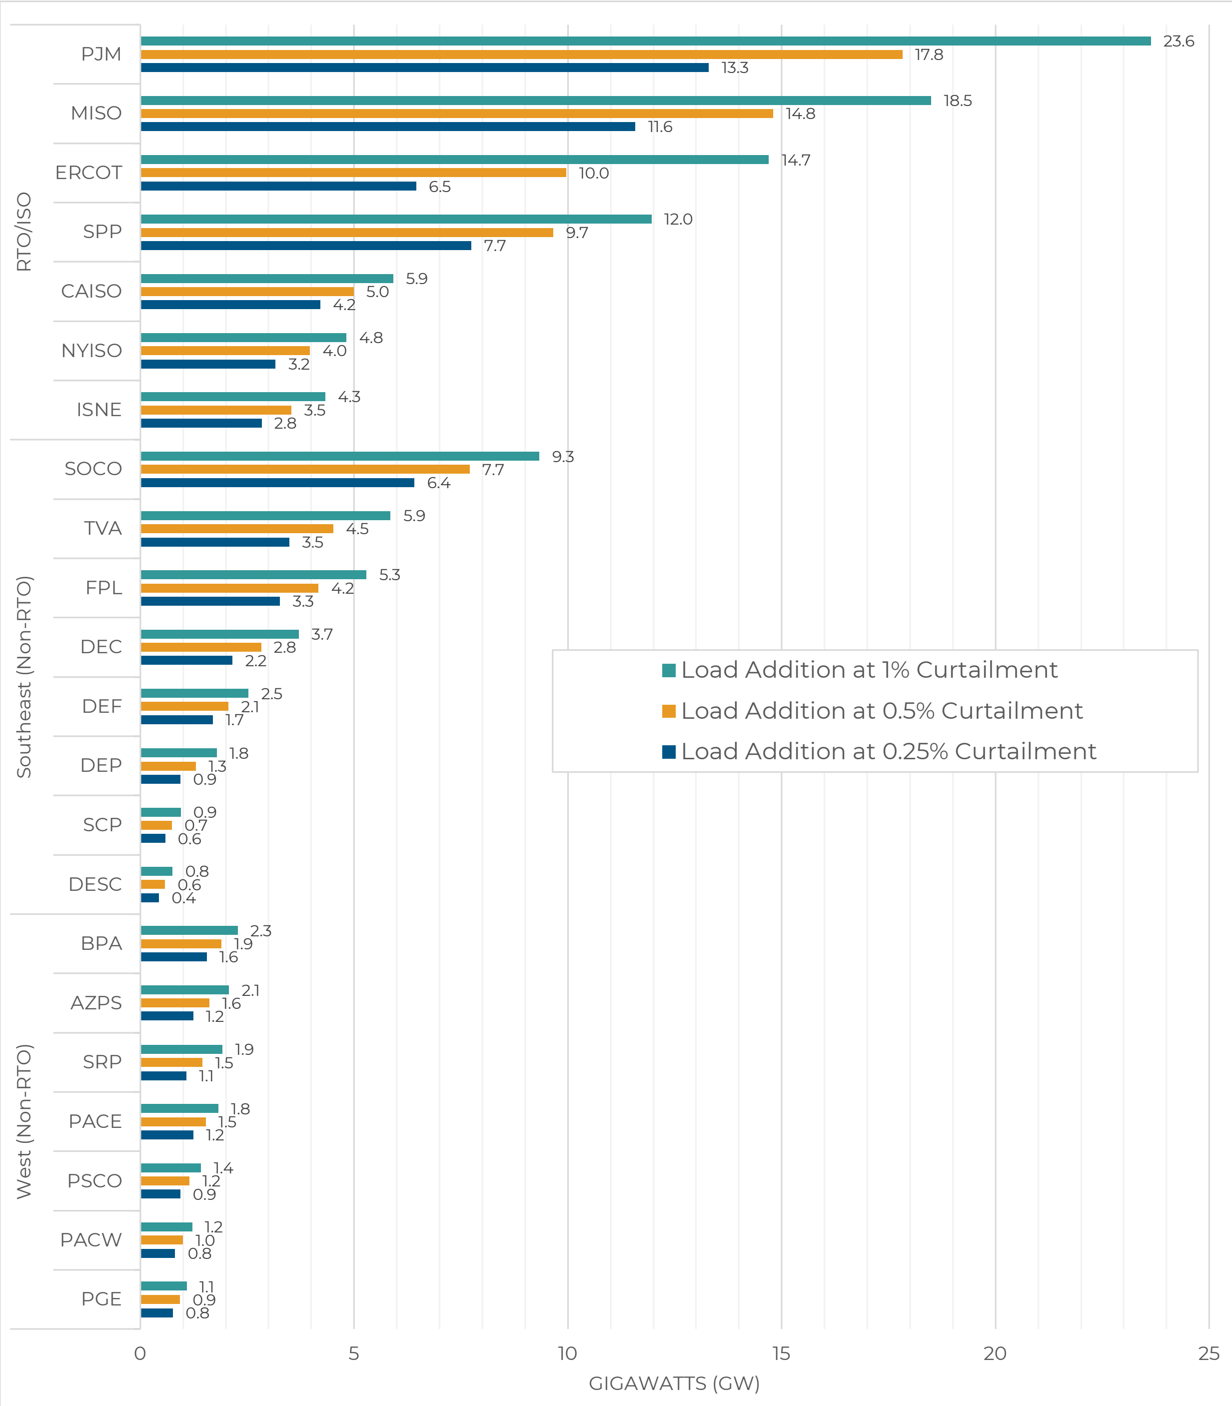

Note: *System headroom* refers to the amount of GW by which a BA's load can be augmented every hour in the absence of capacity expansion so that, provided a certain volume of curtailment of the new load, the total demand does not exceed the supply provisioned by system planners to withstand the expected highest peak. The headroom calculation assumes the new load is constant and hence increases the total load by the same GW hour-by-hour.

reported lead times of up to four and a half years for lower voltage classes and five and a half years for higher voltage classes, alongside a 140% price hike over two years (Rohrer 2024). Wood Mackenzie reported in May 2024 that lead times for high-voltage circuit breakers reached 151 weeks in late 2023, marking a 130% year-over-year increase (Boucher 2024).

Large load interconnection delays have recently led to growing interest among data centers in colocating with existing generation facilities. At a FERC technical conference on the subject in late 2024 (FERC 2024c), several participants highlighted the potential benefits of colocation for expedited interconnection,[4] a view echoed in recent grey literature (Schatzki et al. 2024). Colocation, however, represents only a portion of load interconnections and is not viewed as a long-term, system-wide solution.

Load flexibility similarly offers a practical solution to accelerating the interconnection of large demand loads (SIP 2024, Jabeck 2023). The most time-intensive and costly infrastructure upgrades required for new interconnections are often associated with expanding the transmission system to deliver electricity during the most stressed grid conditions (Gorman et al. 2024). If a new load is assumed to require firm interconnection service and operate at 100% of its maximum electricity draw at all times, including during system-wide peaks, it is far more likely to trigger the need for significant upgrades, such as new transformers, transmission line reconductoring, circuit breakers, or other substation equipment.

To the extent a new load can temporarily reduce (i.e., curtail) its electricity consumption from the grid during these peak stress periods, however, it may be able to connect while deferring—or even avoiding—the need for certain upgrades (ERCOT 2023b). A recent study on Virginia's data center electricity load growth noted, "Flexibility in load is generally expected to offset the need for capacity additions in a system, which could help mitigate the pressure of rapid resource and transmission expansion" (K. Patel et al. 2024). The extent and frequency of required curtailment would depend on the specific nature of the upgrades; in some cases, curtailment may only be necessary if a contingency event occurs, such as an unplanned transmission line or generator outage. For loads that pay for firm interconnection service, any period requiring occasional curtailment would be temporary, ending once necessary network upgrades are completed.[5] Such "partially firm," flexible service was also highlighted by participants in FERC's 2024 technical conference on colocation.[6]

Traditionally, such arrangements have been known as *interruptible* electric service. More recently, some utilities have pursued *flexible* load interconnection options. In March 2022, for example, ERCOT implemented an interim interconnection process for large loads seeking to connect in two years or less, proposing to allow loads seeking to qualify as controllable load resources (CLRs) "to be studied as flexible and potentially interconnect more MWs" (ERCOT 2023b). More recently, ERCOT stated that "the optimal solution for grid reliability is for

4 For example, the Clean Energy Buyers Association (2024) noted, "Flexibility of co-located demand is a key asset that can enable rapid, reliable interconnection."

5 Such an arrangement is analogous to provisional interconnection service available to large generators, as defined in Section 5.9.2 of FERC's Pro Forma *Large Generator Interconnection Agreement* (LGIA).

6 MISO's market monitor representative stated, "instead of being a network firm customer, could [large flexible loads] be a non-firm, or partial non-firm [customer], and that could come with certain configuration requirements that make them truly non-firm, or partially non-firm. But, all those things are the things that could enable some loads to get on the system quicker" (FERC 2024c).

more loads to participate in economic dispatch as CLRs" (Springer 2024). Similarly, Pacific Gas and Electric (PG&E) recently introduced a Flex Connect program to allow certain loads faster access to the grid (Allsup 2024).

These options resemble interconnection services available to large generators that forgo capacity compensation, and potentially higher curtailment risk, in exchange for expedited lower-cost grid access (Norris 2023). FERC codified this approach with Energy Resource Interconnection Service (ERIS) in Order 2003 and revisited the concept during a 2024 technical workshop to explore potential improvements (Norris 2024). Some market participants have since proposed modifying ERIS to facilitate the colocation of new generators with large loads (Intersect Power 2024).

Ratepayers Benefit from Higher System Utilization

The US electric power system is characterized by a relatively low utilization rate, often referred to as the *load factor*. The load factor is the ratio of average demand to peak demand over a given period and provides a measure of the utilization of system capacity (Cerna et al. 2023). A system with a high load factor operates closer to its peak system load for more hours throughout the year, while a system with a low load factor generally experiences demand spikes that are higher than its typical demand levels (Cerna et al. 2022). This discrepancy means that, for much of the year, a significant portion of a system's available generation and transmission infrastructure is underutilized (Cochran et al. 2015).

The power system is designed to handle the highest demand peaks, which in some cases may occur less than once per year, on average, due to extreme weather events. As a result, the bulk of the year sees demand levels well below that peak, leaving substantial headroom in installed capacity. Seasonal shifts add another layer of complexity: some balancing authorities may show higher load factors in summer, yet experience significantly lower utilization in winter, and vice versa.

The *load duration curve* (LDC) illustrates system utilization by ranking demand from highest to lowest over a given period. It provides a visual representation of how often certain demand levels occur, highlighting the frequency and magnitude of peak demand relative to average load. A steep LDC suggests high demand variability, with peaks significantly exceeding typical loads, while a flatter LDC indicates more consistent usage. Figure 2 presents LDCs for each US RTO/ISO based on hourly load between 2016 and 2024, standardized as a percentage of each system's maximum peak demand to allow cross-market comparisons.

A system utilization rate below 100% is expected for most large-scale infrastructure designed to withstand occasional surges in demand. Nevertheless, when the gap between average demand and peak demand is consistently large, it implies that substantial portions of the electric power system—generation assets, transmission infrastructure, and distribution networks—remain idle for much of the year (Riu et al. 2024). These assets are expensive to build and maintain, and ratepayers ultimately bear the cost.

Once the infrastructure is in place, however, there is a strong economic incentive to increase usage and spread these fixed costs over more kilowatt-hours of delivered electricity. An important consideration is therefore the potential for additional load to be added without significant new investment, provided the additional load does not raise the system's overall

Figure 2. Load Duration Curve for US RTO/ISOs, 2016–2024

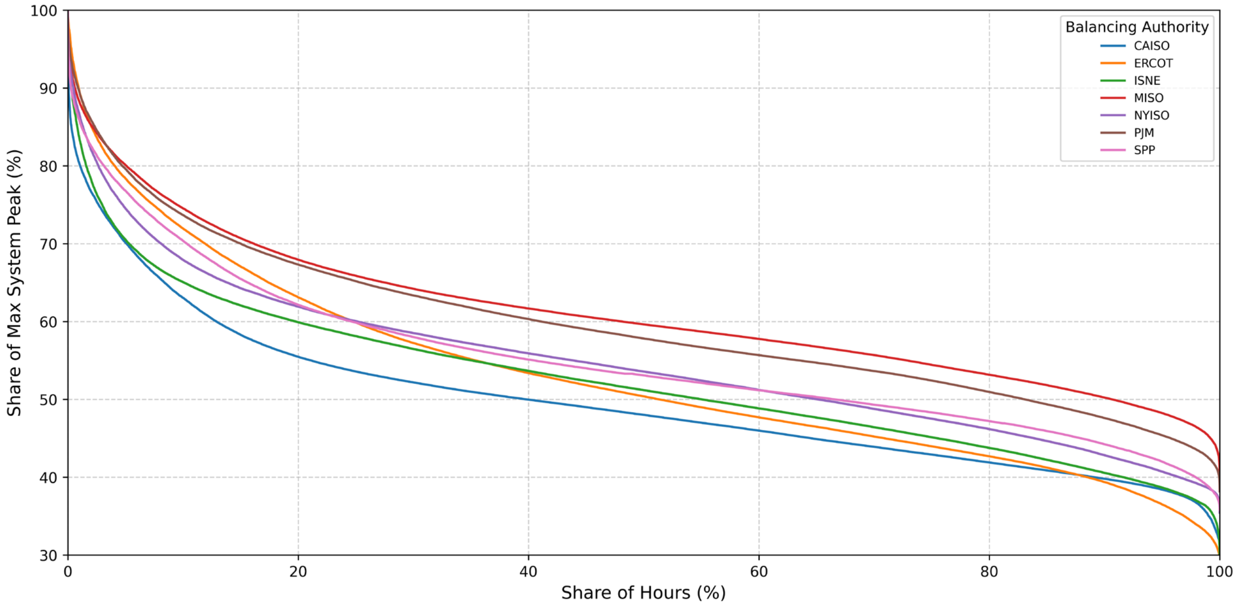

This figure is adapted from the analysis section of this paper, which contains additional detail on the data and method.

peak demand and thereby trigger system expansion.[7] When new loads are flexible enough to avoid a high coincident load factor, thereby mitigating contribution to the highest-demand hours, they fit within the existing grid's headroom.[8] By strategically timing or curtailing demand, these flexible loads can minimize their impact on peak periods. In doing so, they help existing customers by improving the overall utilization rate—thereby lowering the per-unit cost of electricity—and reduce the likelihood that expensive new peaking plants or network expansions may be needed.

In contrast, inflexible new loads that increase the system's absolute peak demand can drive substantial additional needs for generation and transmission capacity. Even a modest rise in peak demand may trigger capital investments in peaking plants, fuel supply infrastructure, and reliability enhancements. These cost implications have contributed to increasingly contentious disputes in which regulators or ratepayer advocates seek to create mechanisms to pass the costs of serving large loads directly to those loads and otherwise ensure data centers do not shift costs via longer contract commitments, billing minimums, and upfront investment (Howland 2024a; Riu et al. 2024). Some examples include:

- The **Georgia Public Service Commission (GPSC)**, citing "staggering" large load growth and the need to protect ratepayers from the costs of serving those customers, recently implemented changes to customer contract provisions if peak draw exceeds 100 MW, mandating a GPSC review and allowing the utility to seek longer contracts and minimum billing for cost recovery (GPSC 2025). This follows GPSC's approval

7 See the discussion on limitations and further analysis in the following section for additional nuance.

8 Demand charges are often based on coincident consumption (e.g., ERCOT's Four Coincident Peak charge uses the load's coincident consumption at the system's expected seasonal peak to determine an averaged demand charge that may account for >30% of a user's annual bill).

of 1.4 GW of gas capacity proposed by Georgia Power in response to load growth "approximately 17 times greater than previously forecasted" through 2030/2031, a forecast it revised upward in late 2024 (GPC 2023, 2024).

- **Ohio**, where American Electric Power issued a moratorium on data center service requests, followed by a settlement agreement with the Public Service Commission staff and consumer advocates that calls for longer contract terms, load ramping schedules, a minimum demand charge, and collateral for service from data centers exceeding 25 MW (Ohio Power Company 2024).
- **Indiana**, where 4.4 GW of interconnection requests from a "handful" of data centers represents a 157% increase in peak load for Indiana Michigan Power over the next six years. Stakeholders there have proposed "firewalling" the associated cost of service from the rest of the rate base, wherein the utility would procure a separate energy, capacity, and ancillary resource portfolio for large loads and recover that portfolio's costs from only the qualifying large loads (Inskeep 2024).
- **Illinois**, where Commonwealth Edison reported that large loads have paid 8.2% of their interconnection costs while the remaining 91.8% is socialized across general customers (ComEd 2024).

These examples underscore the significance of exploring how flexible loads can mitigate peak increases, optimize the utilization of existing infrastructure, and reduce the urgency for costly and time-consuming capacity expansions.

Demand Response and Data Centers

Demand response refers to changes in electricity usage by end-use customers to provide grid services in response to economic signals, reliability events, or other conditions. Originally developed to reduce peak loads (also called *peak shaving*), demand response programs have evolved to encompass a variety of grid services, including balancing services, ancillary services, targeted deferral of grid upgrades, and even variable renewable integration (Hurley et al. 2013; Ruggles et al. 2021). Demand response is often referred to as a form of *demand-side management* or *demand flexibility* (Nethercutt 2023).

Demand response is the largest and most established form of virtual power plant (Downing et al. 2023), with 33 GW of registered capacity in wholesale RTO/ISO programs and 31 GW in retail programs as of 2023 (FERC 2024a).[9] As a share of peak demand, participation in RTO/ISO programs ranges from a high of 10.1% in MISO to a low of 1.4% in SPP. A majority of enrolled capacity in demand response programs are industrial or commercial customers, representing nearly 70% of registered capacity in retail (EIA 2024).

Following a decade of expansion, growth in demand response program participation stalled in the mid-2010s partially because of depressed capacity prices, forecasted over-capacity, and increasingly restrictive wholesale market participation rules (Hledik et al. 2019). However, the resurgence of load growth and increasing capacity prices, coupled with ongoing advancements in DERs and grid information and communication technologies (ICT) appears likely to reverse this trend.

9 RTO/ISO and retail data may overlap.

Studies of national demand response potential have identified a range of potential scenarios (Becker et al. 2024), ranging as high as 200 GW by 2030 in a 2019 study, comprising 20% of the then-forecasted system peak and yielding $15 billion in annual benefits primarily via avoided generation and transmission and distribution (T&D) capacity (Hledik et al. 2019). Notably, this research was conducted before recent load growth forecasts.

The Participation Gap: Data Centers and Demand Response

For nearly two decades, computational loads—and data centers in particular—have been identified as a promising area for demand response. Early studies explored these capabilities, such as a two-phase Lawrence Berkeley National Laboratory study drawing on six years of research, which concluded in 2010 that "data centers, on the basis of their operational characteristics and energy use, have significant potential for demand response" (Ghatikar et al. 2010) and in 2012 that "[certain] data centers can participate in demand response programs with no impact to operations or service-level agreements" (Ghatikar et al. 2012). The 2012 study provided one of the earliest demonstrations of computational load responsiveness, finding that 10% load shed can typically occur within 6 to 15 minutes.

Despite this promise, data centers have historically exhibited low participation rates in demand response programs as a result of operational priorities and economic incentives (Basmadjian 2019; Clausen et al. 2019; Wierman et al. 2014). Data centers are designed to provide reliable, uninterrupted service and generally operate under service-level agreements (SLAs) that mandate specific performance benchmarks, including uptime, latency, and overall quality of service. Deviation from these standards can result in financial penalties and reputational harm, creating a high-stakes environment where operators are averse to operational changes that introduce uncertainty or risk (Basmadjian et al. 2018).

Compounding this challenge is the increasing prevalence of large-scale colocated data centers, which represent a significant share of the data center market (Shehabi et al. 2024). These facilities house multiple tenants, each with varying operational requirements. Coordinating demand response participation in such environments introduces layers of administrative and logistical complexity, as operators must mediate cost- and reward-sharing agreements among tenants. Further, while data centers possess significant technical capabilities, tapping these capabilities for demand response requires sophisticated planning and expertise, which some operators may not have needed to date (Silva et al. 2024).

Economic considerations have further compounded this reluctance. Implementing a demand response program requires investments in advanced energy management systems, staff training, and integration with utility platforms for which costs can be material, particularly for smaller or midsized facilities. At the same time, financial incentives provided by most demand response programs have historically been modest and insufficient to offset the expenses and opportunity costs associated with curtailed operations. For operators focused on maintaining high utilization rates and controlling costs, the economic proposition of demand response participation may be unattractive.

Existing demand response program designs may inadvertently discourage participation. Many programs were originally created with traditional industrial consumers in mind, with different incentives and operational specifications. Price-based programs may require high price variability to elicit meaningful responses, while direct control programs without sufficient guardrails may introduce unacceptable risks related to uptime and performance. The complexity of active participation in demand response markets, including bidding processes and navigating market mechanisms, adds another layer of difficulty. Without streamlined participation structures, tailored incentives, and metrics that reflect the scale and responsiveness of data centers, many existing demand response programs may be ill-suited to the operational realities of modern data centers.

Table 1. Key Data Center Terms

Term	Definition
AI workload	A broad category encompassing computational tasks related to machine learning, natural language processing, generative AI, deep learning, and other AI-driven applications.
AI-specialized data center	Typically developed by hyperscalers, this type of facility is optimized for AI workloads and relies heavily on high-performance graphics processing units (GPUs) and advanced central processing units (CPUs) to handle intensive computing demands.
Computational load	A category of electrical demand primarily driven by computing and data processing activities, ranging from general-purpose computing to specialized AI model training, cryptographic processing, and high-performance computing (HPC).
Conventional data center	A facility that could range from a small enterprise-run server room to a large-scale cloud data center that handles diverse non-AI workloads, including file sharing, transaction processing, and application hosting. These facilities are predominantly powered by CPUs.
Conventional workload	A diverse array of computing tasks typically handled by CPUs, including file sharing, transaction processing, application hosting, and similar operations.
Cryptomine	A dedicated server farm optimized for high-throughput operations on blockchain networks, typically focused on validating and generating cryptocurrency.
Hyperscalers/hyperscale data centers	Large, well-capitalized cloud service providers that build hyperscale data centers to achieve scalability and high performance at multihundred megawatt scale or larger (Howland 2024b, Miller 2024).
Inferencing	The ongoing application of an AI model, where users prompt the model to provide responses or outputs. According to EPRI, inferencing represents 60% of an AI model's annual energy consumption (Aljbour and Wilson 2024).
Model training	The process of developing and training AI models by processing vast amounts of data. Model training accounts for 30–40% of annual AI power consumption and can take weeks or months to complete (Aljbour and Wilson 2024).

Rethinking Data Centers with AI-Driven Flexibility

Limited documentation of commercial data center participation in demand response has reinforced a perception that these facilities' demands are inherently inflexible loads. A variety of recent developments in computational load profiles, operational capabilities, and broader market conditions, however, suggest that a new phase of opportunity and necessity is emerging.

In a July 2024 memo on data center electricity demand, the SEAB recommended the Department of Energy prioritize initiatives to characterize and advance data center load flexibility, including the development of a "flexibility taxonomy and framework that explores the financial incentives and policy changes needed to drive flexible operation" (SEAB 2024). Building on these recommendations, EPRI announced a multi-year Data Center Flexible Load Initiative (DCFlex) in October 2024 with an objective "to spark change through hands-on and experiential demonstrations that showcase the full potential of data center operational flexibility and facility asset utilization," in partnership with multiple tech companies, electric utilities, and independent system operators (Walton 2024a).[10]

The central hypothesis is that the evolving computational load profiles of AI-specialized data centers facilitate operational capabilities that are more amendable to load flexibility. Unlike the many real-time processing demands typical of conventional data center workloads, such as cloud services and enterprise applications, the training of neural networks that power large language models and other machine learning algorithms is deferrable. This flexibility in timing, often referred to as *temporal flexibility*, allows for the strategic scheduling of training as well as other delay-tolerant tasks, both AI and non-AI alike. These delay-tolerant tasks are also referred to as *batch processing* and are typically not user-prompted (AWS 2025).

This temporal flexibility complements the developing interest in *spatial flexibility*, the ability to dynamically distribute workloads across one or multiple data centers in different geographic locations, optimizing resource utilization and operational efficiency. As stated by EPRI in a May 2024 report, "optimizing data center computation and geographic location to respond to electricity supply conditions, electricity carbon intensity, and other factors in addition to minimizing latency enables data centers to actively adjust their electricity consumption … some could achieve significant cost savings—as much as 15%—by optimizing computation to capitalize on lower electric rates during off-peak hours, reducing strain on the grid during high-demand periods" (EPRI 2024). For instance, having already developed a temporal workload shifting system, Google is seeking to implement spatial flexibility as well (Radovanović 2020).

In addition to temporal and spatial flexibility, other temporary load reduction methods may also enable data center flexibility. One approach is dynamic voltage and frequency scaling, which reduces server power consumption by lowering voltage or frequency at the expense of processing speed (Moons et al. 2017; Basmadjian 2019; Basmadjian and de Meer 2018). Another is server optimization, which consolidates workloads onto fewer servers while idling or shutting down underutilized ones, thereby reducing energy waste (Basmajian 2019; Chaurasia et al. 2021). These load reduction methods are driven by advances in virtual workload management, made possible by the "virtualization" of hardware (Pantazoglou et al. 2016).

10 Pointing to EPRI's new DCFlex Initiative, Michael Liebreich noted in a recent essay, "For instance, when they see how much it costs to work 24/7 at full power, perhaps data-center owners will see a benefit to providing some demand response capacity… " (Liebreich 2024).

Finally, temperature flexibility leverages the fact that cooling systems account for 30% to 40% of data center energy consumption (EPRI 2024). For instance, operators can increase cooling during midday when solar energy is abundant and reduce cooling during peak evening demand.[11] While these methods may be perceived as uneconomic due to potential impacts on performance, hardware lifespan, or SLAs, they are not intended for continuous use. Instead, they are best suited for deployment during critical hours when grid demand reduction is most valuable.

Beyond peak shaving, data centers also hold potential to participate in ancillary services, particularly those requiring rapid response, such as frequency regulation. Studies have described how data centers can dynamically adjust workloads to provide real-time support to the grid, effectively acting as “virtual spinning reserves” that help stabilize grid frequency and integrate intermittent renewable resources (McClurg et al. 2016; Al Kez et al. 2021; Wang et al. 2019). This capability extends beyond traditional demand response by providing near-instantaneous balancing resources (Zhang et al. 2022).

Three overarching market trends create further opportunities for load flexibility now than in the past. The first is constrained supply-side market conditions that raise costs and lead times for the interconnecting large inflexible loads, when speed to market is paramount for AI developers. The second is advancements in on-site generation and storage technologies that have lowered costs and expanded the availability of cleaner and more commercially viable behind-the-meter solutions, increasing their appeal to data center operators (Baumann et al. 2020). The third is the growing concentration of computational load in colocated or hyper-scale data centers—accounting for roughly 80% of the market in 2023—which is lending scale and specialization to more sophisticated data center operators. These operators, seek-ing speed to market, may be more likely to adopt flexibility in return for faster interconnec-tion (Shehabi et al. 2024; Basmadjian et al. 2018). The overarching trends underpinning this thesis are summarized in Table 2.

An important consideration for future data center load profiles is the balance between AI-specialized data centers focused on model development and those oriented toward inferencing. If fewer AI models are developed, a larger proportion of computing resources will shift toward inferencing tasks, which is delay-intolerant and variable (Riu et al. 2024). According to EPRI, training an AI model accounts for 30% of its annual footprint, compared to 60% for inferencing the same model (EPRI 2024).

In the absence of regulatory guidance, most advancements in data center flexibility to date are being driven by voluntary private-sector initiatives. Some hyperscalers and data center developers are taking steps to mitigate grid constraints by prioritizing near-term solutions for load flexibility. For example, one such company, Verrus, has established its business model around the premise that flexible data center operations offer an effective solution for growth needs (SIP 2024). Table 3 highlights additional initiatives related to facilitating or demonstrating data center flexibility.

11 Cooling demand for servers is inherently dependent on server workloads. Therefore, reducing workloads saves on cooling needs as well.

Table 2. Trends Enabling Data Center Load Flexibility

Category	Legacy	Future
Computational load profile	• Conventional servers with CPU-dominated workloads (Shehabi et al. 2024) • Real-time, delay-intolerant, and unscheduled processing (e.g., cloud services, enterprise apps) • Low latency critical	• AI-specialized servers with GPU or tensor processing unit (TPU)-favored workloads (Shehabi et al. 2024) • Greater portion of delay-tolerant and scheduled machine learning workloads (model training, non-interactive services) • Higher share of model training affords greater demand predictability • Highly parallelized workloads (Shehabi et al. 2024)
Operational capabilities	• Minimal temporal load shifting • Minimal spatial load migration • High proximity to end users for latency-sensitive tasks • Reliance on Tier 2 diesel generators for backup • Limited utilization of on-site power resulting from pollution concerns and regulatory restrictions (Cary 2023)	• More robust and intelligent temporal workload shifting (Radovanović et al. 2022) • Advanced spatial load migration and multi–data center training (D. Patel et al. 2024) • Flexibility in location for model training • Backup power diversified (storage, renewables, natural gas, cleaner diesel) • Cleaner on-site power enables greater utilization
Market conditions	• Minimal electric load growth • High availability of T&D network headroom • Standard interconnection timelines and queue volumes • Low supply chain bottlenecks for T&D equipment • Low capacity prices and forecasted overcapacity • High cost of clean on-site power options • Small-scale “server room” model	• High electric load growth • Low availability of T&D network headroom • Long interconnection timelines and overloaded queues • High supply chain bottlenecks for T&D equipment • High capacity prices and forecasted undercapacity (Walton 2024b) • Lower cost of clean on-site power options (Baranko et al. 2024) • Data center operations concentrating in large-scale facilities and operators

Table 3. Implementations of Computational Load Flexibility

Category	Examples
Operational flexibility	• Google deployed a "carbon-aware" temporal workload–shifting algorithm and is now seeking to develop geographic distribution capabilities (Radovanović 2020). • Google data centers have participated in demand response by reducing non-urgent compute tasks during grid stress events in Oregon, Nebraska, the US Southeast, Europe, and Taiwan (Mehra and Hasegawa 2023). • Enel X has supported demand response participation by data centers in North America, Ireland, Australia, South Korea, and Japan, including use of on-site batteries and generators to enable islanding within minutes (Enel X 2024). • Startup companies like Emerald AI are developing software to enable large-scale demand response from data centers through recent advances in computational resource management to precisely deliver grid services while preserving acceptable quality of service for compute users
On-site power	• Enchanted Rock, an energy solutions provider that supported Microsoft in building a renewable natural gas plant for a data center in San Jose, CA, created a behind-the-meter solution called Bridge-to-Grid, which seeks to provide intermediate power until primary service can be switched to the utility. At that point, the on-site power transitions to flexible backup power (Enchanted Rock 2024, 2025).
Market design and utility programs	• ERCOT established the Large Flexible Load Task Force and began to require the registration of large, interruptible loads seeking to interconnect with ERCOT for better visibility into their energy demand over the next five years (Hodge 2024). • ERCOT's demand response program shows promise for data center flexibility, with 750+ MW of data mining load registered as CLRs, which are dispatched by ERCOT within preset conditions (ERCOT 2023a). • PG&E debuted Flex Connect, a pilot that provides quicker interconnection service to large loads in return for flexibility at the margin when the system is constrained (Allsup 2024, St. John 2024).
Cryptomining	• A company generated more revenue from its demand response participation in ERCOT than from Bitcoin mining in one month, at times accommodating a 95% load reduction during peak demands (Riot Platforms 2023).

ANALYSIS OF CURTAILMENT-ENABLED HEADROOM

In this section we describe the method for estimating the gigawatts of new load that could be added to existing US power system load before the total exceeds what system planers are prepared to serve, provided that load curtailment is applied as needed. This serves as a proxy for the system's ability to integrate new load, which we term *curtailment-enabled headroom*.[12] We first investigated the aggregate and seasonal load factor for each of the 22 investigated balancing authorities, which measures a system's average utilization rate. Second, we computed the curtailment-enabled headroom for different assumptions of ac-

12 SEAB proposed a similar term, *available flex capacity*, in its July 2024 report *Recommendations on Powering Artificial Intelligence and Data Center Infrastructure*.

ceptable new load curtailment rates. In this context, *curtailment* refers to instances where the new load temporarily reduces its electricity draw—such as by using on-site generation resources, shifting load temporally or spatially, or otherwise reducing operations—to ensure system demand does not exceed historical peak thresholds. Third, we quantified the magnitude, duration, and seasonal concentration of the load curtailment for each balancing authority. Finally, we examined the correlation between load factor, seasonal curtailment, and max potential load additions. This process is summarized in Figure 3.

Data and Method

Data

We considered nine years of hourly load data aggregated for each of the 22 balancing authorities, encompassing seven RTO/ISOs,[13] eight non-RTO Southeastern BAs,[14] and seven non-RTO Western BAs.[15] Together, these balancing authorities represent 744 of the approximate 777 GW of summer peak load (95%) across the continental United States. The dataset, sourced from the EIA Hourly Load Monitor (EIA-930), contains one demand value per hour

Figure 3. Steps for Calculating Headroom and Related Metrics

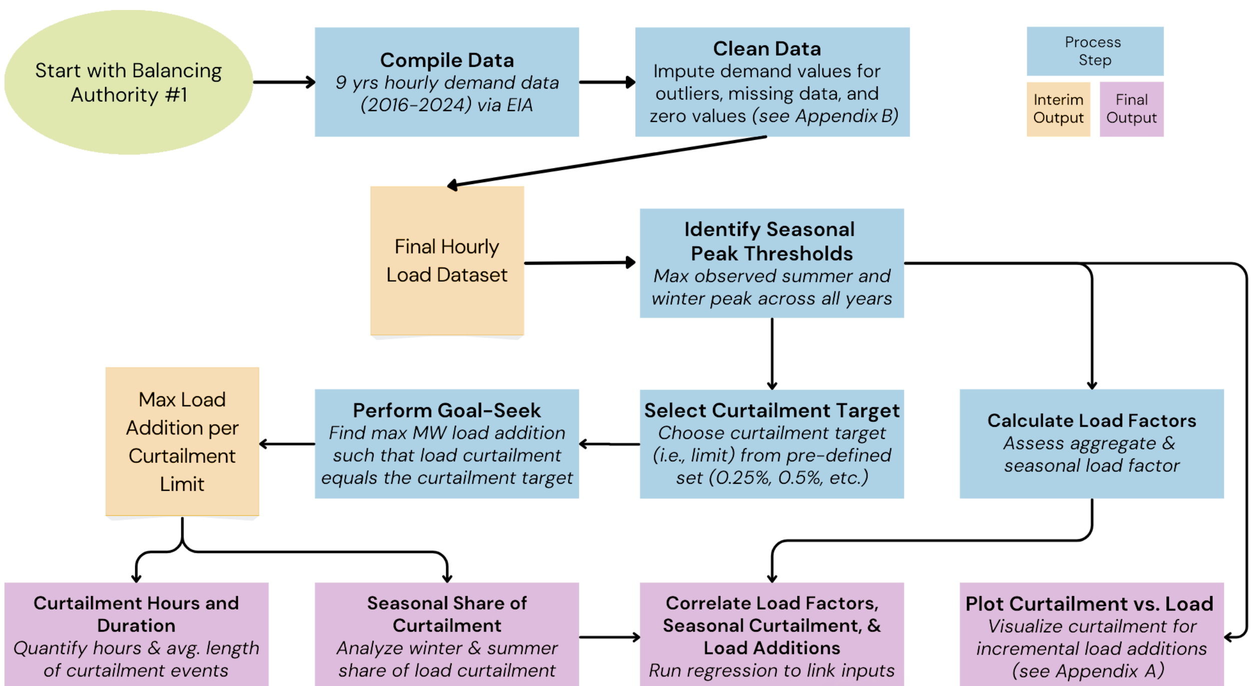

13 CAISO, ERCOT, ISO-NE, MISO, NYISO, PJM, and SPP.

14 DEC; DEP; DEF; DESC; FPL; Santee Cooper, SCP; Southern Company (SOCO); and TVA. Note the different BA codes used by EIA: DUK for DEC, CPLE for DEP, SCEG for DESC, FPC for DEF, and SC for SCP. Also note that Southern Company includes Georgia Power, Alabama Power, and Mississippi Power. A complete list of abbreviations and their definitions can be found at the end of the paper.

15 AZPS, BPA, PACE, PACW, PGE, PSCO, and SRP. Note that EIA uses the code BPAT for BPA. A complete list of abbreviations and their definitions can be found at the end of the paper.

and spans January 1, 2016, through December 31, 2024.[16] Data from 2015 were excluded because of incomplete reporting.[17] The dataset was cleaned to identify and impute values for samples with missing or outlier demand values (see details in Appendix B).

Determining Load Additions for Curtailment Limits

An analysis was conducted to determine the maximum load addition for each balancing authority that can be integrated while staying within predefined curtailment limits applied to the new load. The load curtailment limits (0.25%, 0.5%, 1.0%, and 5.0%) were selected within the range of maximum curtailment caps for existing interruptible demand response programs.[18] The analysis focused on finding the load addition volume in megawatts that results in an average annual load curtailment rate per balancing authority that matches the specified limit. To achieve this, a goal-seek technique was used to solve for the load addition that satisfies this condition,[19] for which the mathematical expression is presented in Appendix C. The calculation assumed the new load is constant and hence increases the total system load by the same gigawatt volume hour-by-hour. To complement this analysis and visualize the relationship between load addition volume and curtailment, curtailment rates were also calculated across small incremental load additions (i.e., 0.25% of the BA's peak load).

Load Curtailment Definition and Calculation

Load curtailment is defined as the megawatt-hour reduction of load required to prevent the augmented system demand (existing load + new load) from exceeding the maximum seasonal system peak threshold (e.g., see Figure 4). Curtailment was calculated hourly as the difference between the augmented demand and the seasonal peak threshold. These hourly curtailments in megawatt-hours were aggregated for all hours in a year to determine the total annual curtailment. The curtailment rate for each load increment was defined as the total annual curtailed megawatt-hours divided by the new load's maximum potential annual consumption, assuming continuous operation at full capacity.

Peak Thresholds and Seasonal Differentiation

Balancing authorities develop resource expansion plans to support different peak loads in winter and summer. To account for variation, we defined seasonal peak thresholds for each balancing authority. Specifically, we identified the maximum summer peak and the maximum winter peak observed from 2016 to 2024 for each balancing authority.[20] These thresholds serve as the upper limits for system demand during their respective seasons, and all

16 Additional detail on EIA's hourly load data collection is available at https://www.eia.gov/electricity/gridmonitor/about.

17 Fewer than half of the year's load hours were available, making the data unsuitable for inclusion.

18 For example, PG&E's and Southern California Edison's Base Interruptible Programs limit annual interruption for registered customers to a maximum of 180 hours (2.0% of all annual hours) or 10 events per month.

19 The goal-seek approach was implemented using Python's scipy.optimize.root_scalar function from the SciPy library. This tool is designed for solving one-dimensional root-finding problems, where the goal is to determine the input value that satisfies a specified equation within a defined range.

20 To identify the max seasonal peak load, summer was defined as June–August, while winter encompassed December–February. In a few cases, the BA's seasonal peak occurred within one month of these periods (AZPS winter, FPL winter, CAISO summer, CAISO winter), which were used as their max seasonal peak. To account for potential (albeit less likely) curtailment in shoulder months, the applicable summer peak was applied to April–May and September–October and the winter peak to November and March.

Figure 4. Illustrative Load Flexibility in PJM

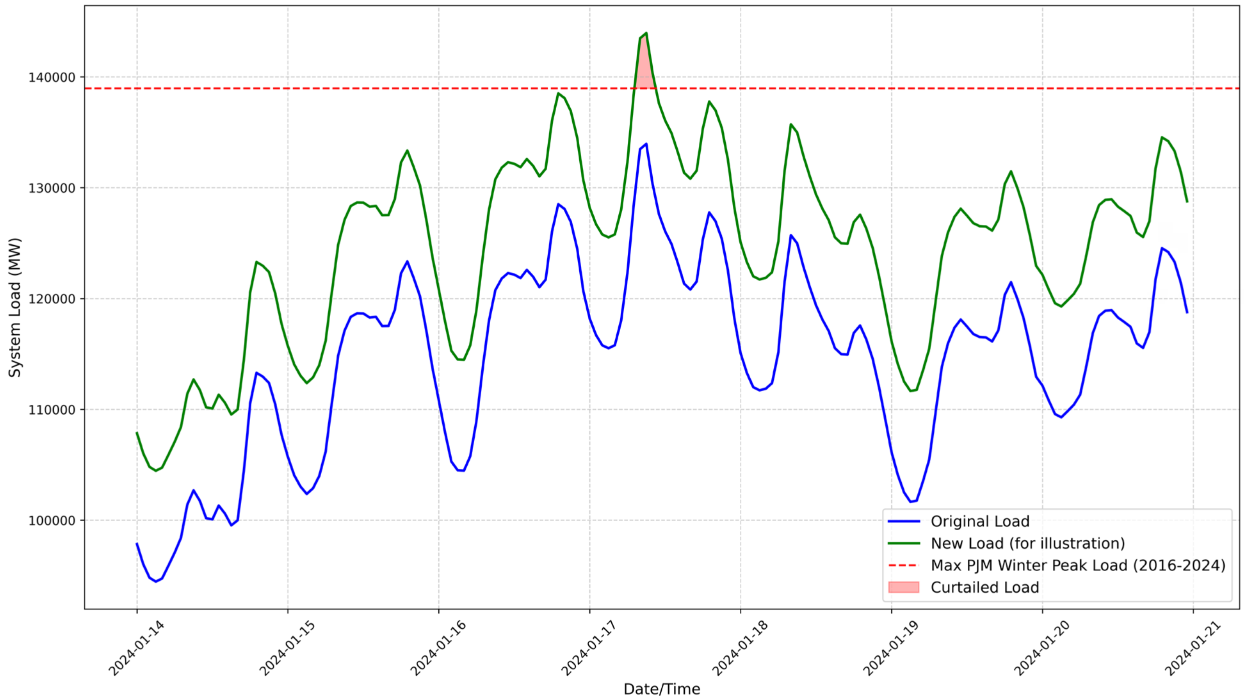

megawatt-hours that exceeded these thresholds was counted as curtailed energy. This seasonal differentiation captures the distinct demand characteristics of regions dominated by cooling loads (summer peaks) versus heating loads (winter peaks).

Year-by-Year Curtailment Analysis

Curtailment was analyzed independently for each year from 2016 to 2024. This year-by-year approach captures temporal variability in demand patterns, including the effects of extreme weather events and economic conditions. For each year, curtailment volumes were calculated across all load addition increments, resulting in a list of annual curtailment rates corresponding to each load increment. To synthesize results across years, we calculated the average curtailment rate for each load addition increment by averaging annual curtailment rates over the nine years. This averaging process smooths out year-specific anomalies and provides an estimate of the typical system response to additional load. This analysis was also used to calculate the average number of hours of curtailment for each curtailment limit and the seasonal allocation of curtailed generation.[21] We also assessed the magnitude of load curtailment required during these hours as a share of the new load's maximum potential draw to calculate the number of hours when 90%, 75%, and 50% or more of the load would still be available.

21 Consistent with the curtailment analysis, summer was defined as June–August and winter as December–February. For BAs located on the Pacific coast (BPA, CAISO, PGE, PACE, PACW), November was counted as winter given the region's unique seasonal load profile.

Figure 5. Load Factor by Balancing Authority and Season, 2016–2024

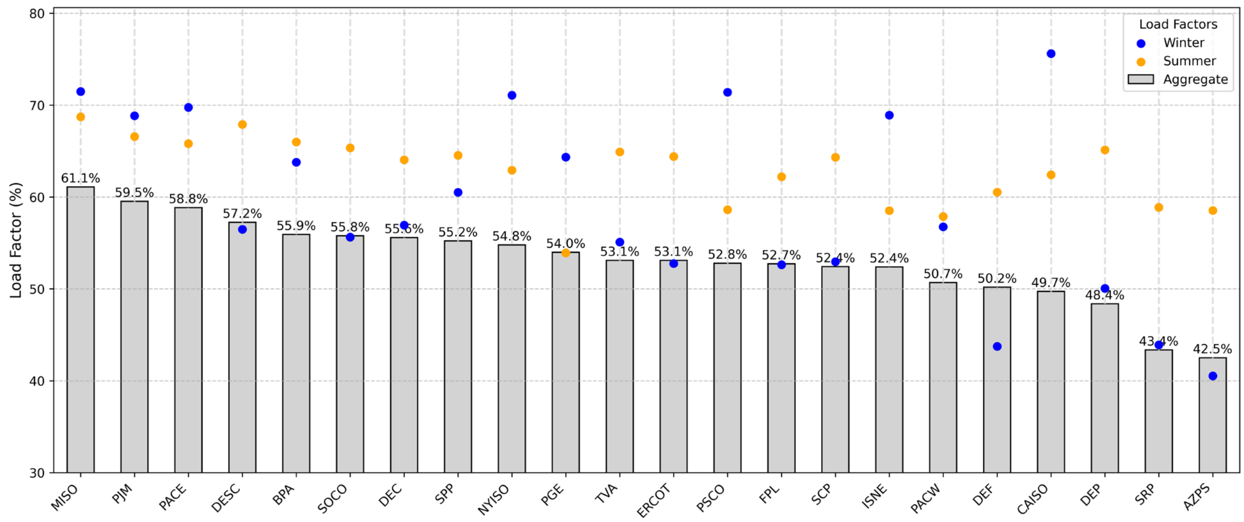

Figure 6. Load Duration Curves by Balancing Authority, 2016–2024

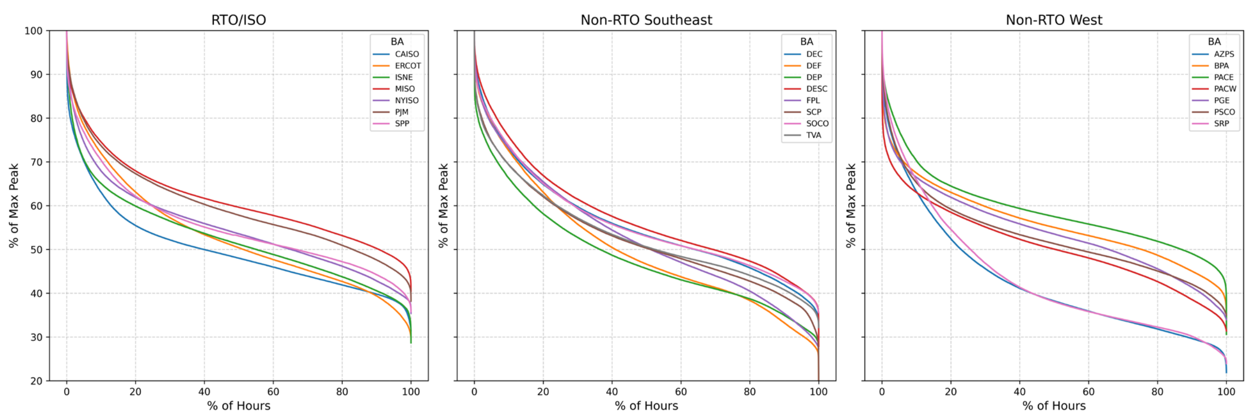

Results

Load Factor

In examining data for 22 balancing authorities, we found that aggregate load factors ranged between 43% to 61% (Figures 5 and 6), with an average and median value of 53%. The BAs with the lowest aggregate load factors were those in the desert southwest, Arizona Public Service Company (AZPS) and Salt River Project Agricultural Improvement and Power District (SRP). In terms of seasonal load factor, defined here as the average seasonal load as a share of seasonal maximum load (i.e., not as a share of the maximum all-time system load), winter load factors were notably lower than summer. The average and median winter load factor was 59% and 57% respectively, compared to 63% and 64% for summer. A majority of the balancing authorities had higher summer load factors (14) than winter (8).

Headroom Volume

Results show that the headroom across the 22 analyzed balancing authorities is between 76 to 215 GW, depending on the applicable load curtailment limit. This means that 76 to 215 GW of load could be added to the US power system and yet the total cumulative load would remain below the historical peak load, except for a limited number of hours per year

Figure 7. Headroom Enabled by Load Curtailment Thresholds, GW

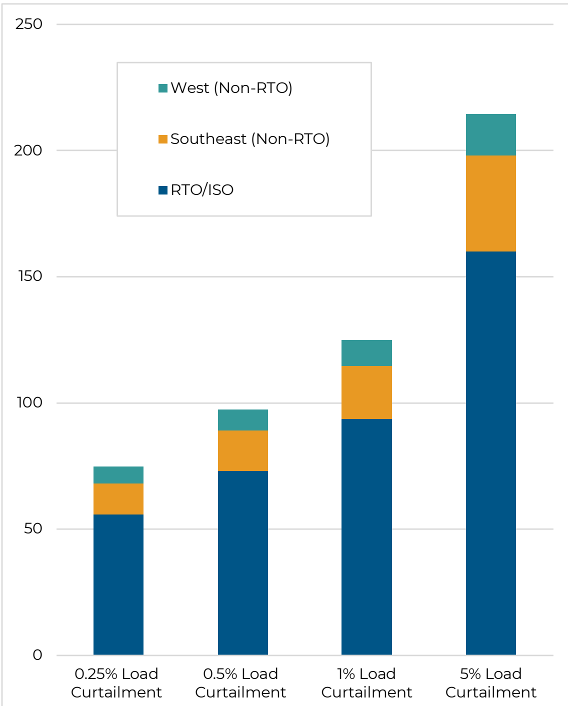

Figure 8. Headroom Enabled by 0.5% Load Curtailment by Balancing Authority, GW

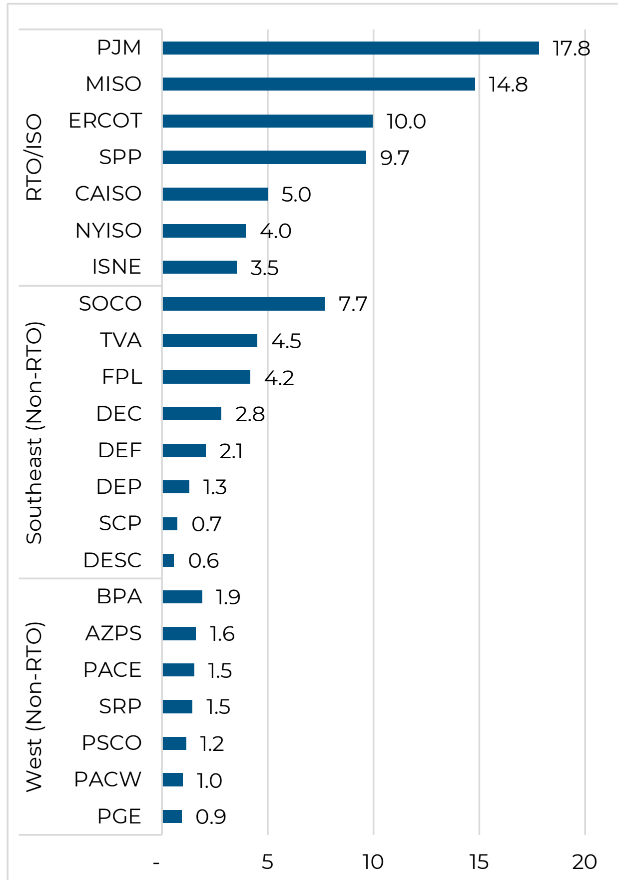

Figure 9. Load Curtailment Rate Due to Load Addition, % of System Peak

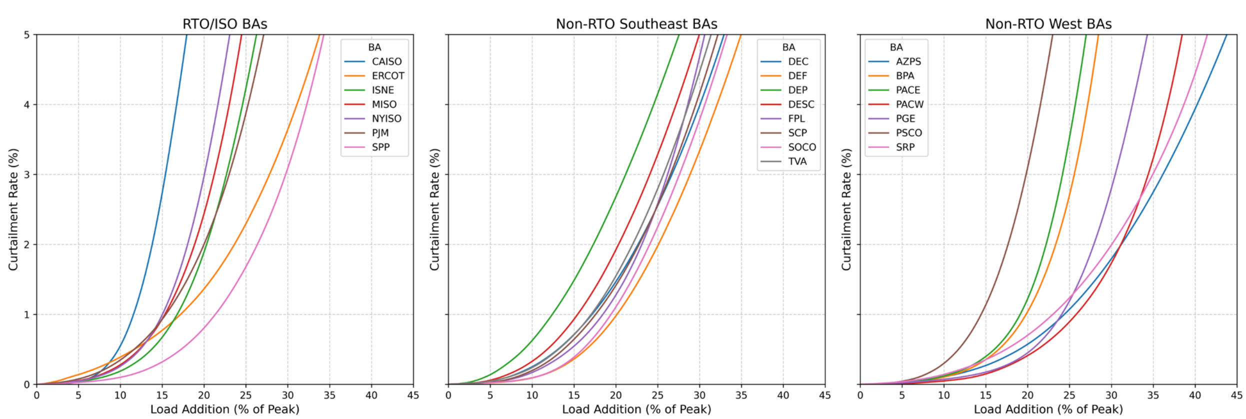

when the new load would be unserved. Specifically, 76 GW of headroom is available at an expected load curtailment rate of 0.25% (i.e., if 0.25% of the maximum potential annual energy consumption of the new load is curtailed during the highest load hours, or 1,643 out of 657,000 GWh). This headroom increases to 98 GW at 0.5% curtailment, 126 GW at 1.0% curtailment, and 215 GW at 5.0% curtailment (Figure 7). Headroom varies by balancing authority (Figure 8), including as a share of system peak (Figure 9). The five balancing authorities with the highest potential volume at 0.5% annual curtailment are PJM at 18 GW, MISO at 15 GW, ERCOT at 10 GW, SPP at 10 GW, and Southern Company at 8 GW. Detailed plots for each balancing authority, including results for each year, can be found in Appendix A.

Curtailment Hours

A large majority of curtailment hours retain most of the new load. Most hours during which load reduction is required entail a curtailment rate below 50% of the new load. Across all 22 BAs, the average required load curtailment times are 85 hours under the 0.25% curtailment rate (~1% of the hours in a year), 177 hours under the 0.5% curtailment rate, 366 hours under the 1.0% curtailment rate, and 1,848 hours under the 5.0% curtailment rate (i.e., ~21% of the hours). On average, 88% of these hours retain at least 50% of the new load (i.e., less than 50% curtailment of the load is required), 60% of the hours retain at least 75% of the load, and 29% retain at least 90% of the load (see Figure 10).

Curtailment Duration

The analysis calculated the average hourly duration of curtailment events (i.e., the length of time the new load is curtailed during curtailment events). All hours in which any curtailment occurred were included, regardless of magnitude. The results for each balancing authority and curtailment limit are presented in Figure 11. The average duration across BAs was 1.7 hours for the 0.25% limit, 2.1 hours for the 0.5% limit, 2.5 hours for the 1.0% limit, and 4.5 hours for the 5.0% limit.

Figure 10. Hours of Curtailment by Load Curtailment Limit

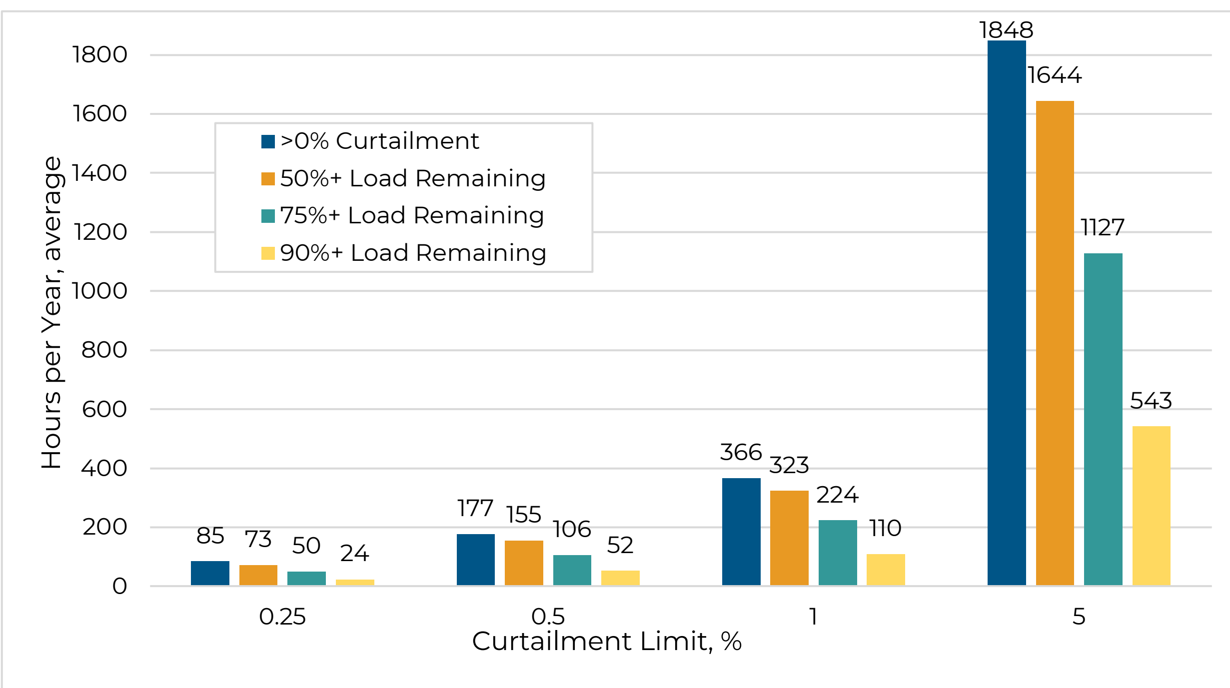

Seasonal Concentration of Curtailment

The analysis reveals significant variation in the seasonal concentration of curtailment hours across balancing authorities. The winter-summer split ranged from 92% to 1% for CAISO (California Independent System Operator), where curtailment is heavily winter-concentrated, to 0.2% to 92% for AZPS,[22] which exhibited a heavily summer-concentrated curtailment profile (Figure 12a).[23]

Figure 11. Average Curtailment Duration by Balancing Authority and Curtailment Limit, Hours

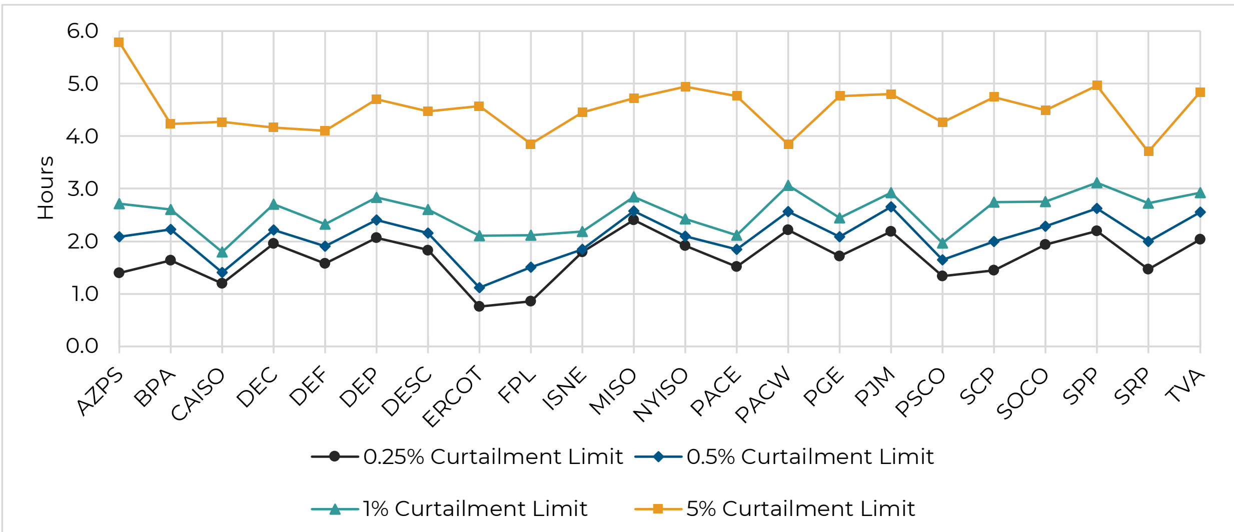

Figure 12. Seasonal Curtailment Analysis

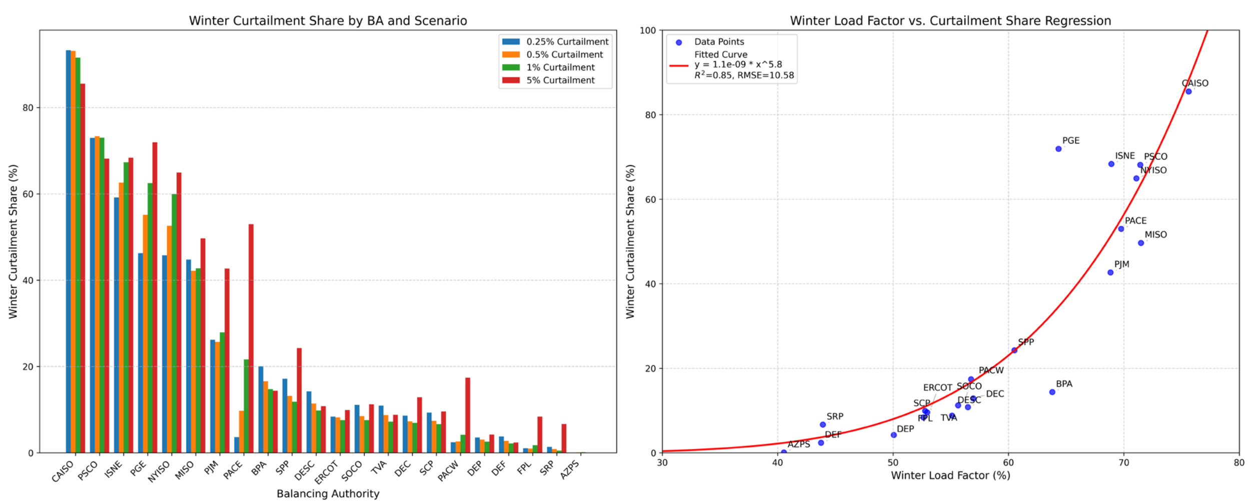

(a) (b)

22 Note the remainder of the curtailment occurred in these BAs in shoulder months (i.e., not summer, not winter).

23 These values correspond to the seasonal curtailment concentration for the 1% curtailment limit.

A key observation is the strong correlation between the winter load factor (system utilization during winter months) and the seasonal allocation of curtailment hours (Figure 12b). BAs with lower winter load factors—indicating reduced system utilization during winter—tend to have greater capacity to accommodate additional load in winter while experiencing a dispro-portionately higher share of curtailment during summer months. This trend is particularly pronounced in balancing authorities located in the Sun Belt region, resulting in a lower win-ter concentration of curtailment hours.

While most BAs exhibited relatively stable seasonal curtailment shares across increasing load addition thresholds, some demonstrated notable shifts in seasonal allocation as load additions increased (e.g., PACW, FPL, NYISO, ISO-NE, PACE, PGE). These shifts highlight the dynamic interplay between system demand patterns and the incremental addition of new load.

Figure 12a illustrates this variability, showcasing the relationship between winter load factor and winter curtailment share across curtailment scenarios.[24]

Discussion

The results highlight that the significant headroom in US power systems—stemming from their by-design low load factors—could be tapped to enable the integration of substantial load additions with relatively low rates of load curtailment. They also underscore substantial variation in flexibility across balancing authorities, driven by differences in seasonal and aggregate load patterns. This variation suggests that seasonal load factors may be strongly linked to how much additional load a balancing authority can integrate without requiring high curtailment rates.

To explore this relationship, we analyzed system load factors in relation to the additional load that each balancing authority could accommodate while limiting the load curtailment rate to 0.5%, 1.0%, and 5.0% (i.e., the load curtailment limit). To allow for meaningful comparison across BAs, the additional load was standardized as a percentage of the BA's historical peak load. To account for whether a balancing authority's curtailment was concentrated in the summer or winter, the seasonal load factor was selected corresponding to the season with the highest share of curtailment.

The analysis revealed that BAs with higher seasonal load factors tended to have less headroom for the load curtailment limits examined (Figure 13). In simpler terms, systems with higher utilization during their busiest season had less power generation capacity planned to be available that could serve new load without hitting curtailment limits. For example, CAISO, with a seasonal load factor of 76%, could accommodate less additional load compared to PacifiCorp West (PACW) and AZPS, which exhibited lower seasonal load factors and supported larger load additions as a share of peak system load. This relationship grew in statistical significance as the load curtailment limit increased, yielding an R^2 value of 0.48 and an RMSE of 3.04 at the 0.5% curtailment limit, and an R^2 value of 0.86 and an RMSE of 1.55 at the 5% curtailment limit (i.e., 86% of the variation in load addition capacity across balancing authorities can be explained by differences in load factor at a curtailment limit of 5.0%).

24 Note in Figure 12b that a high-degree polynomial function captures the nonlinear growth in the area under the load curve as curtailed load exceeds a fixed peak threshold. This fit generally aligns with expectations, demonstrating that higher-degree terms are necessary to capture the relationship between load factor and curtailed load.

Figure 13. Load Factor Versus Max Load Addition as Share of Peak Load

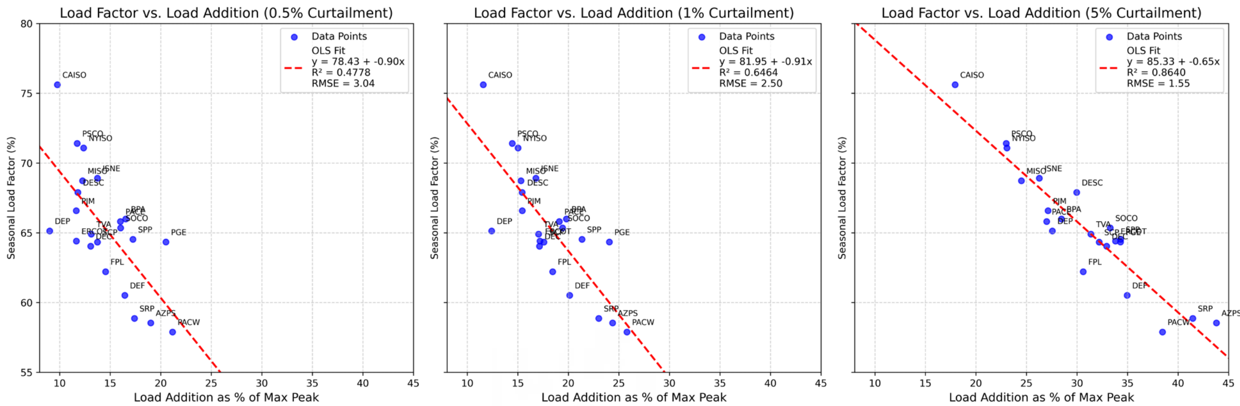

These findings emphasize the importance of load factor as a predictor of curtailment-enabled headroom. BAs with more uneven peak seasonal demand—characterized by relatively low system utilization in winter or summer—tend to have greater capacity to integrate new loads with limited curtailment. Conversely, systems with more consistent demand across the winter and summer face tighter limits, as their capacity to absorb additional load is already constrained by elevated baseline usage.

Limitations

This analysis provides a first-order assessment of power generation capacity available for serving new curtailable loads, and hence is an exploration of the market potential for large-scale demand response. The primary focus of the analysis is to ensure that total demand, subject to curtailment limits for new load, stays below the system peak for which system planners have prepared. Other considerations important for planning—such as ensuring adequate transmission capacity, ramping capability, and ramp-feasible reserves, among others—are beyond the scope of this study and therefore the results cannot be taken as an accurate estimate of the load that can be added to the system. Additionally, the analysis assumes the new loads do not change current demand patterns but rather shift the existing demand curves upward, and a more precise assessment of the potential for integration of new loads would require detailed characterization of the temporal patterns of the load. There is significant variation in how system operators forecast and plan for system peaks, accounting for potential demand response, and as a result there will be differences in the methods used to estimate potential to accommodate new load. Despite these limitations, the results presented here signal a vast potential that, even if overstated, warrants further research.

On the other hand, some aspects of this study may have contributed to an underestimation of available headroom. First, the analysis assumes that each BA's maximum servable load in the winter and summer is equivalent to the BA's highest realized seasonal peak demand based on the available historical data. However, the available generation capacity in each balancing authority should materially exceed this volume when accounting for the installed reserve margin. In other words, system operators have already planned their systems to accommodate load volume that exceeds their highest realized peak. Second, the analysis removed outlier demand values in some BAs to avoid using unreasonably high maximum peak thresholds, which would understate the curtailment rates. However, if some of the removed outliers properly represent a level of system load that the system is prepared to serve reliably, this analysis may have understated the curtailment-enabled headroom. Third, the analysis assumed all new load is constant and hence increases the total system load by the same gigawatt hour-by-hour, which would tend to overstate the absolute level of required gigawatt hour curtailment for a load that is not constant.

Future Analysis

Enhancing this analysis to more accurately assess the capacity to integrate large curtailable load would require addressing the following considerations:

Network Constraints

This analysis does not account for network constraints, which would require a power flow simulation to evaluate the ability of the transmission system to accommodate additional load under various conditions. As such, the results should not be interpreted as an indication that the identified load volumes could be interconnected and served without any expansions in network capacity. While the existing systems are planned to reliably serve their peak loads, this planning is based on the current load topology and the spatial distribution of generation and demand across the transmission network. A large new load could avoid exceeding aggregate peak system demand by employing flexibility, yet still cause localized grid overloads as a result of insufficient transmission capacity in specific areas. Such overloads could necessitate network upgrades, including the expansion of transmission lines, substations, or other grid infrastructure. Alternatively, in the absence of network upgrades, localized congestion could be addressed through the addition of nearby generation capacity, potentially limiting the flexibility and economic benefits of the new load. These factors underscore the importance of incorporating network-level analyses to fully understand the operational implications of large flexible load additions.

Intertemporal Constraints

This analysis does not account for intertemporal constraints related to load and generator operations. For load operations, response times affect system operations and management of operational reserves. Faster response times from flexible loads could alleviate system stress more effectively during peak demand periods, potentially reducing the reliance on reserve capacity. Conversely, slower response times may require additional reserves to bridge the gap between the onset of system imbalances and the load's eventual response. Moreover, the rapid ramp-down of large flexible loads could lead to localized stability or voltage issues, particularly in regions with weaker grid infrastructure. These effects may necessitate more localized network analyses to evaluate stability risks and operational impacts. On the generation side, intertemporal constraints such as ramping limits, minimum up and down times, and startup times can affect the system's ability to integrate fast-response demand. For instance, ramping constraints may restrict how quickly generators can adjust output to align with the curtailment of flexible loads, while minimum uptime and downtime requirements can limit generator flexibility.

Loss of Load Expectation

Peak load is a widely used proxy for resource adequacy and offers a reasonable indicative metric for high-level planning analyses. However, a more granular assessment would incorporate periods with the highest loss of load expectation (LOLE), which represent the times when the system is most likely to experience supply shortfalls. Historically, LOLE periods have aligned closely with peak load periods, making peak load a convenient and broadly applicable metric. However, in markets with increasing renewable energy penetration, LOLE periods are beginning to shift away from traditional peak load periods. This shift is driven by the variability and timing of renewable generation, particularly solar and wind, which can alter the temporal distribution of system stress. As a result, analyses focused solely on peak load may understate or misrepresent the operational challenges associated with integrating large new loads into these evolving systems.

CONCLUSION

This study highlights extensive potential for leveraging large load flexibility to address the challenges posed by rapid load growth in the US power system. By estimating the curtailment-enabled headroom across balancing authorities, the analysis demonstrates that existing system capacity—intentionally designed to accommodate the extreme swings of peak demand—could accommodate significant new load additions with relatively modest curtailment, as measured by the average number, magnitude, and duration of curtailment hours.

The findings further emphasize the relationship between load factors and headroom availability. Balancing authorities with lower seasonal load factors exhibit greater capacity to integrate flexible loads, highlighting the importance of regional load patterns in determining system-level opportunities. These results suggest that load flexibility can play a significant role in improving system utilization, mitigating the need for costly infrastructure expansion and complementing supply-side investments to support load growth and decarbonization objectives.

This analysis provides a first-order assessment of market potential, with estimates that can be refined through further evaluation. In particular, network constraints, intertemporal operational dynamics, and shifts in loss-of-load expectation periods represent opportunities for future analyses that can offer a deeper understanding of the practical and operational implications of integrating large flexible loads.

In conclusion, the integration of flexible loads offers a promising, near-term strategy for addressing structural transformations in the US electric power system. By utilizing existing system headroom, regulators and market participants can expedite the accommodation of new loads, optimize resource utilization, and support the broader goals of reliability, affordability, and sustainability.

REFERENCES

Al Kez, D., A. M. Foley, F. W. Ahmed, M. O'Malley, and S. M. Muyeen. 2021. "Potential of Data Centers for Fast Frequency Response Services in Synchronously Isolated Power Systems." *Renewable and Sustainable Energy Reviews* 151(November): 111547. https://doi.org/10.1016/j.rser.2021.111547.

Aljbour, J., and T. Wilson. 2024. *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption*. Palo Alto, CA: Electric Power Research Institute. https://www.wpr.org/wp-content/uploads/2024/06/3002028905_Powering-Intelligence_-Analyzing-Artificial-Intelligence-and-Data-Center-Energy-Consumption.pdf.

Allsup, M. 2024 "PG&E Is Laying the Groundwork for Flexible Data Center Interconnection." *Latitude Media*, November 15. www.latitudemedia.com/news/pg-e-is-laying-the-groundwork-for-flexible-data-center-interconnection/.

AWS. 2025. "What Is Batch Processing?" *Amazon Web Services*. aws.amazon.com/what-is/batch-processing/.

Baranko, K., D. Campbell, Z. Hausfather, J. McWalter, N. Ransohoff. 2024. "Fast, Scalable, Clean, and Cheap Enough: How Off-Grid Solar Microgrids Can Power the AI Race." *OffgridAI*, December. https://www.offgridai.us/.

Basmadjian, R. 2019. "Flexibility-Based Energy and Demand Management in Data Centers: A Case Study for Cloud Computing." *Energies* 12(17): 3301. https://doi.org/10.3390/en12173301.

Basmadjian, R., J. F. Botero, G. Giuliani, X. Hesselbach, S. Klingert, and H. De Meer. 2018. "Making Data Centers Fit for Demand Response: Introducing GreenSDA and GreenSLA Contracts." *IEEE Transactions on Smart Grid* 9(4): 3453–64. https://doi.org/10.1109/TSG.2016.2632526.

Basmadjian, R., and H. de Meer. 2018. "Modelling and Analysing Conservative Governor of DVFS-enabled Processors." In *Proceedings of the Ninth International Conference on Future Energy Systems (e-Energy '18)*. New York: Association for Computing Machinery. 519–25. https://doi.org/10.1145/3208903.3213778.

Baumann, C. 2020. *How Microgrids for Data Centers Increase Resilience, Optimize Costs, and Improve Sustainability*. Rueil-Malmaison, France: Schneider Electric. https://www.se.com/us/en/download/document/Microgrids_for_Data_Centers/.

Becker, J., K. Brehm, J. Cohen, T. Fitch, and L. Shwisberg. 2024. *Power Shift: How Virtual Power Plants Unlock Cleaner, More Affordable Electricity Systems*. RMI. https://rmi.org/wp-content/uploads/dlm_uploads/2024/09/power_shift_report.pdf.

Boucher, B. 2024. "The Challenge of Growing Electricity Demand in the US and the Shortage of Critical Electrical Equipment." *Wood Mackenzie*, May 2. https://www.woodmac.com/news/opinion/the-challenge-of-growing-electricity-demand-in-the-us-and-the-shortage-of-critical-electrical-equipment/.

Cary, P. 2023. "Virginia Environmental Regulators Drop Plan to Allow Data Centers to Rely on Diesel Generators." *Prince William Times*, April 12. www.princewilliamtimes.com/news/virginia-environmental-regulators-drop-plan-to-allow-data-centers-to-rely-on-diesel-generators/article_b337df48-d96a-11ed-8861-4b1de9b9963f.html.

Cerna, F., E. Naderi, M. Marzband, J. Contreras, J. Coelho, and M. Fantesia. 2022. "Load Factor Improvement of the Electricity Grid Considering Distributed Resources Operation and Regulation of Peak Load." *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4293004.

Cerna, F. V., J. K. Coêlho, M. P. Fantesia, E. Naderi, M. Marzband, and J. Contreras. 2023. "Load Factor Improvement of the Electricity Grid Considering Distributed Energy Resources Operation and Regulation of Peak Load." *Sustainable Cities and Society* 98(November): 104802. https://doi.org/10.1016/j.scs.2023.104802.

Chaurasia, N., M. Kumar, R. Chaudhry, and O. P. Verma. 2021. "Comprehensive Survey on Energy-Aware Server Consolidation Techniques in Cloud Computing." *The Journal of Supercomputing* 77(10): 11682–737. https://doi.org/10.1007/s11227-021-03760-1.

Clausen, A., G. Koenig, S. Klingert, G. Ghatikar, P. M. Schwartz, and N. Bates. 2019. "An Analysis of Contracts and Relationships between Supercomputing Centers and Electricity Service Providers." In *Workshop Proceedings of the 48th International Conference on Parallel Processing*, 1–8. Kyoto: ACM. https://doi.org/10.1145/3339186.3339209.

Clean Energy Buyers Association. 2024. *Post-Technical Conference Comments of the Clean Energy Buyers Association*. FERC Docket No. AD24-11-000. Washington, DC: Federal Energy Regulatory Commission. https://elibrary.ferc.gov/eLibrary/filelist?accession_number=20241209-5198.

Cochran, J., P. Denholm, B. Speer, and M. Miller. 2015. *Grid Integration and the Carrying Capacity of the U.S. Grid to Incorporate Variable Renewable Energy*. Golden, CO: National Renewable Energy Laboratory. https://www.nrel.gov/docs/fy15osti/62607.pdf.

ComEd. 2024. "Commonwealth Edison Company's Response to Constellation Energy Generation LLC's ('Constellation') Data Request Constellation-ComEd 5.01 RGP—Constellation-ComEd 5.04 RGP." *ICC Docket* Nos. 22-0486 / 23-0055 (Consol.) Refiled Grid Plan. Received May 30. https://www.icc.illinois.gov/docket/P2023-0055/documents/352947/files/617782.pdf.

Downing, J. N. Johnson, M. McNicholas, D. Nemtzow, R. Oueid, J. Paladino, and E. Bellis Wolfe. 2023. *Pathways to Commercial Liftoff: Virtual Power Plants*. Washington, DC: US Department of Energy. https://liftoff.energy.gov/wp-content/uploads/2023/10/LIFTOFF_DOE_VVP_10062023_v4.pdf.

EIA. 2024. "Electric Power Annual 2023." *US Energy Information Agency*. https://www.eia.gov/electricity/annual/.

Enchanted Rock. 2024. "Microsoft Is a Multinational Technology Company Producing Software, Electronics, PC's and Related Services like Data Centers." Case Study. December 12. https://enchantedrock.com/microsoft-is-a-multinational-technology-company-producing-software-electronics-pcs-and-related-services-like-data-centers/.

Enchanted Rock. 2025. "Bridge-To-Grid." enchantedrock.com/bridge-to-grid/.

Enel X. 2024. "How Data Centers Support the Power Grid with Ancillary Services?" July 10. www.enelx.com/tw/en/resources/how-data-centers-support-grids.

EPRI. 2024. *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption*. Palo Alto, CA: Electric Power Research Institute. https://www.epri.com/research/products/000000003002028905.

ERCOT. 2023a. *Ancillary Services*. Austin: Electric Reliability Council of Texas. https://www.ercot.com/files/docs/2023/06/06/Ancillary-Services-Handout-0524.pdf

ERCOT. 2023b. *Large Loads—Impact on Grid Reliability and Overview of Revision Request Package*. Presented at the NPRR1191 and Related Revision Requests Workshop, August 16. https://www.ercot.com/files/docs/2023/11/08/PUBLIC-Overview-of-Large-Load-Revision-Requests-for-8-16-23-Workshop.pptx.

FERC. 2024a. *2024 Assessment of Demand Response and Advanced Metering*. Washington, DC: Federal Energy Regulatory Commission. https://www.ferc.gov/news-events/news/ferc-staff-issues-2024-assessment-demand-response-and-advanced-metering.

FERC. 2024b. *Form No. 714—Annual Electric Balancing Authority Area and Planning Area Report*. Washington, DC: Federal Energy Regulatory Commission. https://www.ferc.gov/industries-data/electric/general-information/electric-industry-forms/form-no-714-annual-electric/data

FERC. 2024c. *Transcript of Technical Conference on Large Loads Co-Located at Generating Facilities, November.* Washington, DC: Federal Energy Regulatory Commission. https://www.ferc.gov/media/transcript-technical-conference-regarding-large-loads-co-located-generating-facilities.

FERC. 2024d. Pro Forma *Large Generator Interconnection Agreement (LGIA).* 18 C.F.R. § 35.28, Appendix C, § 5.9.2. Washington, DC: Federal Energy Regulatory Commission. https://www.ferc.gov/sites/default/files/2020-04/LGIA-agreement.pdf.

GPC. 2023. *2023 Integrated Resource Plan Update*. Atlanta: Georgia Power Company. https://www.georgiapower.com/content/dam/georgia-power/pdfs/company-pdfs/2023-irp-update-main-document.pdf.

GPC. 2024. *Large Load Economic Development Report for Q3 2024 PD.* Atlanta: Georgia Power Company. https://psc.ga.gov/search/facts-document/?documentId=220461.

GPSC. 2025. "PSC Approves Rule to Allow New Power Usage Terms for Data Centers," January 23. Atlanta: Georgia Public Service Commission. https://psc.ga.gov/site/assets/files/8617/media_advisory_data_centers_rule_1-23-2025.pdf.

Ghatikar, G., M. A. Piette, S. Fujita, A. T. McKane, J. Q Han, A. Radspieler, K. C. Mares, and D. Shroyer. 2010. *Demand Response and Open Automated Demand Response Opportunities for Data Centers*. Berkeley, CA: Lawrence Berkeley National Laboratory. https://eta.lbl.gov/publications/demand-response-and-open-automated.

Ghatikar, G., V. Ganti, N. Matson, and M. A. Piette. 2012. *Demand Response Opportunities and Enabling Technologies for Data Centers: Findings From Field Studies*. Berkeley, CA: Lawrence Berkeley National Laboratory. https://doi.org/10.2172/1174175.

Gorman, W., J. Mulvaney Kemp, J. Rand, J. Seel, R. Wiser, N. Manderlink, F. Kahrl, K. Porter, and W. Cotton. 2024. "Grid Connection Barriers to Renewable Energy Deployment in the United States." *Joule* (December): 101791. https://doi.org/10.1016/j.joule.2024.11.008.

Hledik, R., A. Faruqui, T. Lee, and J. Higham. 2019. *The National Potential for Load Flexibility: Value and Market Potential Through 2030*. Boston: The Brattle Group. https://www.brattle.com/wp-content/uploads/2021/05/16639_national_potential_for_load_flexibility_-_final.pdf.

Hodge, T. 2024. "Data Centers and Cryptocurrency Mining in Texas Drive Strong Power Demand Growth." *Today in Energy*, October 3. www.eia.gov/todayinenergy/detail.php?id=63344.

Howland, E. 2024a. "FERC Rejects Basin Electric's Cryptocurrency Mining Rate Proposal." *Utility Dive*, August 21. https://www.utilitydive.com/news/ferc-basin-electricscryptocurrency-bitcoin-mining-rate-proposal/724811/.

Howland, E. 2024b. "FERC Rejects Interconnection Pact for Talen-Amazon Data Center Deal at Nuclear Plant." *Utility Dive*, November 4. https://www.utilitydive.com/news/ferc-interconnection-isa-talen-amazon-data-center-susquehanna-exelon/731841/.

Hurley, D., P. Peterson, and M. Whited. 2013. *Demand Response as a Power System Resource: Program Designs, Performance, and Lessons Learned in the United States*. Cambridge, MA: Synapse Energy Economics. https://www.synapse-energy.com/sites/default/files/SynapseReport.2013-03.RAP_.US-Demand-Response.12-080.pdf.

Inskeep, B. 2024. *Testimony on Behalf of Citizens Action Coalition of Indiana, Inc.* "In the Matter of the Verified Petition of Indiana Michigan Power Co. for Approval of Modifications to Its Industrial Power Tariff." Before the Indiana Regulatory Commission. October 15, 2024. https://iurc.portal.in.gov/docketed-case-details/?id=b8cd5780-0546-ef11-8409-001dd803817e

Intersect Power. 2024. *Post-Technical Conference Comments of the Intersect Power LLC*. FERC Docket No. AD24-11-000. Washington, DC: Federal Energy Regulatory Commission. https://elibrary.ferc.gov/eLibrary/filelist?accession_number=20241209-5237.

Jabeck, B. 2023. "Flexible Capacity: The Secret Weapon for Securing Interconnection." *Data Center Frontier*, July 28. https://www.datacenterfrontier.com/sponsored/article/33008776/flexible-capacity-the-secret-weapon-for-securing-interconnection.

Kearney, L., and L. Hampton. 2025. "US Power Stocks Plummet as DeepSeek Raises Data Center Demand Doubts." *Reuters*, January 27. https://www.reuters.com/business/energy/us-power-stocks-plummet-deepseek-raises-data-center-demand-doubts-2025-01-27/.

Lee, V., P. Seshadri, C. O'Niell, A. Choudhary, B. Holstege, and S. A Deutscher. 2025. *Breaking Barriers to Data Center Growth*. Boston: Boston Consulting Group. https://www.bcg.com/publications/2025/breaking-barriers-data-center-growth.

Li, F., L. Braly, and C. Post. "Current Power Trends and Implications for the Data Center Industry." *FTI Consulting*, July. https://www.fticonsulting.com/insights/articles/current-power-trends-implications-data-center-industry.

Liebreich, M. 2024. "Generative AI—The Power and the Glory." *BloombergNEF*, December 24. https://about.bnef.com/blog/liebreich-generative-ai-the-power-and-the-glory/.

McClurg, J., R. Mudumbai, and J. Hall. 2016. “Fast Demand Response with Datacenter Loads.” In *2016 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, 1–5. Minneapolis: IEEE. https://doi.org/10.1109/ISGT.2016.7781219.

Mehra, V., and R. Hasegawa. 2023. “Supporting Power Grids with Demand Response at Google Data Centers.” *Google Cloud Blog*, October 3. cloud.google.com/blog/products/infrastructure/using-demand-response-to-reduce-data-center-power-consumption.

Miller, R. 2024. “The Gigawatt Data Center Campus Is Coming.” *Data Center Frontier*, April 29. https://www.datacenterfrontier.com/hyperscale/article/55021675/the-gigawatt-data-center-campus-is-coming.

Moons, B., W. Uytterhoeven, W. Dehaene, and M. Verhelst. 2017. “DVAFS: Trading Computational Accuracy for Energy Through Dynamic-Voltage-Accuracy-Frequency-Scaling.” *Design, Automation & Test in Europe Conference & Exhibition*. 488-93. https://doi.org/10.23919/DATE.2017.7927038.

NERC. 2024 *Long-Term Reliability Assessment*. Atlanta: North American Electric Reliability Corporation. https://www.nerc.com/pa/RAPA/ra/Reliability%20Assessments%20DL/NERC_Long%20Term%20Reliability%20Assessment_2024.pdf.

Nethercutt, E. J. 2023. *Demand Flexibility within a Performance-Based Regulatory Framework*. Washington, DC: National Association of Regulatory Utility Commissioners. https://pubs.naruc.org/pub/2A466862-1866-DAAC-99FB-E054E1C9AB13

NIAC. 2024. *Addressing the Critical Shortage of Power Transformers to Ensure Reliability of the U.S. Grid*. The National Infrastructure Advisory Council. Washington, DC: The President’s National Infrastructure Advisory Council. https://www.cisa.gov/sites/default/files/2024-09/NIAC_Addressing%20the%20Critical%20Shortage%20of%20Power%20Transformers%20to%20Ensure%20Reliability%20of%20the%20U.S.%20Grid_Report_06112024_508c_pdf_0.pdf.

Norris, T. 2023. *Beyond FERC Order 2023: Considerations on Deep Interconnection Reform*. NI PB 23-04. Durham, NC: Nicholas Institute for Energy, Environment & Sustainability, Duke University. https://hdl.handle.net/10161/31260.

Norris, T. 2024. *Pre-Workshop Comments for FERC Staff-led Workshop on Innovations and Efficiencies in Generator Interconnection*. Docket No. AD24-9-000. Washington, DC: Federal Energy Regulatory Commission. https://nicholasinstitute.duke.edu/publications/comments-ferc-workshop-innovations-efficiencies-generator-interconnection.

Ohio Power Company. 2024. *Joint Stipulation and Recommendation before the Public Service Commission of Ohio*. “In the Matter of the Application of Ohio Power Company for New Tariffs Related To Data Centers and Mobile Data Centers,” Case No. 24-508-EL-ATA, October 23. https://dis.puc.state.oh.us/ViewImage.aspx?CMID=A1001001A24J23B55758I01206.

Pantazoglou, M., G. Tzortzakis, and A. Delis. 2016. “Decentralized and Energy-Efficient Workload Management in Enterprise Clouds.” *IEEE Transactions on Cloud Computing* 4(2): 196–209. https://doi.org/10.1109/TCC.2015.2464817.

Patel, D., D. Nishball, J. Eliahou Ontiveros. "Multi-Datacenter Training: OpenAI's Ambitious Plan To Beat Google's Infrastructure." *SemiAnalysis*, September 4. https://semianalysis.com/2024/09/04/multi-datacenter-training-openais/.

Patel, K., K. Steinberger, A. DeBenedictis, M. Wu, J. Blair, P. Picciano, and P. Oporto, et al. 2024. *Virginia Data Center Study: Electric Infrastructure and Customer Rate Impact*. San Francisco: Energy and Environmental Economics, Inc. https://jlarc.virginia.gov/pdfs/presentations/JLARC%20Virginia%20Data%20Center%20Study_FINAL_12-09-2024.pdf.

Radovanović, A. 2020. "Our Data Centers Now Work Harder When the Sun Shines and Wind Blows." *The Keyword*, April 22. https://blog.google/inside-google/infrastructure/data-centers-work-harder-sun-shines-wind-blows/.

Radovanović, A., R. Koningstein, I. Schneider, B. Chen, A. Duarte, B. Roy, D. Xiao, et al. 2023. "Carbon-Aware Computing for Datacenters." *IEEE Transactions on Power Systems* 38(2): 1270–80. https://doi.org/10.1109/TPWRS.2022.3173250.

Riot Platforms. 2023. "Riot Announces August 2023 Production and Operations Updates." September 6. www.riotplatforms.com/riot-announces-august-2023-production-and-operations-updates/.

Riu, I., D. Smiley, S. Bessasparis, and K. Patel. 2024. *Load Growth Is Here to Stay, but Are Data Centers? Strategically Managing the Challenges and Opportunities of Load Growth*. San Francisco: Energy and Environmental Economics, Inc. https://www.ethree.com/data-center-load-growth/.

Rohrer, J. 2024. "Supply Chains Impact Power Transmission Systems." *Closed Circuit*, April 23. https://www.wapa.gov/supply-chains/.

Rouch, M. A. Denman, P. Hanbury, P. Renno, and E. Gray. 2024. *Utilities Must Reinvent Themselves to Harness the AI-Driven Data Center Boom*. Boston: Bain & Company. https://www.bain.com/insights/utilities-must-reinvent-themselves-to-harness-the-ai-driven-data-center-boom/.

Ruggles, T. H., J. A. Dowling, N. S. Lewis, and K. Caldeira. 2021. "Opportunities for Flexible Electricity Loads Such as Hydrogen Production from Curtailed Generation." *Advances in Applied Energy* 3(August): 100051. https://doi.org/10.1016/j.adapen.2021.100051.

Saul, J. 2024. "Data Centers Face Seven-Year Wait for Dominion Power Hookups." *Bloomberg*, August 29. https://www.bloomberg.com/news/articles/2024-08-29/data-centers-face-seven-year-wait-for-power-hookups-in-virginia.

Schatzki, T., J. Cavicchi, and M. Accordino. 2024. *Co-Located Load: Market, Economic, and Ratemaking Implications*. Analysis Group. https://www.analysisgroup.com/globalassets/insights/publishing/2024_co_located_load_market_economic_and_ratemaking_implications.pdf.

SEAB. 2024. *Recommendations on Powering Artificial Intelligence and Data Center Infrastructure*. Washington, DC: US Secretary of Energy Advisory Board. https://www.energy.gov/sites/default/files/2024-08/Powering%20AI%20and%20Data%20Center%20Infrastructure%20Recommendations%20July%202024.pdf.

Shehabi, A., S. J. Smith, A. Hubbard, A. Newkirk, N. Lei, M. A. B. Siddik, and B. Holecek, et al. 2024. *2024 United States Data Center Energy Usage Report*. Berkeley, CA: Lawrence Berkeley National Laboratory.

Silva, C. A., R. Vilaça, A. Pereira, and R. J. Bessa. 2024. "A Review on the Decarbonization of High-Performance Computing Centers." *Renewable and Sustainable Energy Reviews* 189(January): 114019. https://doi.org/10.1016/j.rser.2023.114019.

SIP. 2024. "Data Center Flexibility: A Call to Action Improving the Grid with a New Approach to Data Center Development." *Sidewalk Infrastructure Partners*, March. https://www.datacenterflexibility.com/.

Springer, A. 2024. *Large Loads in ERCOT—Observations and Risks to Reliability*. Presented at the NERC Large Load Task Force, October 8. https://www.nerc.com/comm/RSTC/LLTF/LLTF_Kickoff_Presentations.pdf.

Srivathsan, B., M. Sorel, P. Sachdeva., H. Batra, R. Sharma, R. Gupta, and S. Choudhary. 2024. "AI Power: Expanding Data Center Capacity to Meet Growing Demand." *McKinsey & Company*, October 29. https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand.

St. John, J. 2024 "A New Way to Fix Grid Bottlenecks for EV Charging: Flexible Connection." *Canary Media*, December 10. https://www.canarymedia.com/articles/transmission/a-new-way-to-fix-grid-bottlenecks-for-ev-charging-flexible-connection.

Walton, R. 2024a. "EPRI Launches Data Center Flexibility Initiative with Utilities, Google, Meta, NVIDIA." *Utility Dive*, October 30. https://www.utilitydive.com/news/epri-launches-data-center-flexibility-initiative-with-NVIDIA-google-meta/731490/.

Walton, R. 2024b. "'Explosive'" Demand Growth Puts More than Half of North America at Risk of Blackouts: NERC." *Utility Dive*, December 18. www.utilitydive.com/news/explosive-demand-growth-blackouts-NERC-LTRA-reliability/735866/.

Wang, W., A. Abdolrashidi, N. Yu, and D. Wong. 2019. "Frequency Regulation Service Provision in Data Center with Computational Flexibility." *Applied Energy* 251(October): 113304. https://doi.org/10.1016/j.apenergy.2019.05.107.

WECC. 2024. "State of the Interconnection." *Western Electricity Coordinating Council*, September. https://feature.wecc.org/soti/topic-sections/load/index.html.

Wierman, A., Z. Liu, I. Liu, and H. Mohsenian-Rad. 2014. "Opportunities and Challenges for Data Center Demand Response." In *International Green Computing Conference*, 1–10. Dallas, TX: IEEE. https://doi.org/10.1109/IGCC.2014.7039172.

Wilson, J. D., Z. Zimmerman, and R. Gramlich. 2024. *Strategic Industries Surging: Driving US Power Demand*. Bethesda, MD: GridStrategies. https://gridstrategiesllc.com/wp-content/uploads/National-Load-Growth-Report-2024.pdf.

Zhang, Y., D. C. Wilson, I. C. Paschalidis and A. K. Coskun. 2022. "HPC Data Center Participation in Demand Response: An Adaptive Policy With QoS Assurance." *IEEE Transactions on Sustainable Computing* 7(1): 157–71. http://doi.org/10.1109/TSUSC.2021.3077254.

ABBREVIATIONS

AI	Artificial intelligence
AZPS	Arizona Public Service Company
BA	balancing authority
BPA	Bonneville Power Administration
CAGR	compound annual growth rate
CAISO	California Independent System Operator
CLRs	controllable load resources
CPUs	central processing units
DEC	Duke Energy Carolinas
DEF	Duke Energy Florida
DEP	Duke Energy Progress East
DERs	distributed energy resources
DESC	Dominion Energy South Carolina
EIA	Energy Information Administration
EPRI	Electrical Power Research Institute
ERCOT	Electric Reliability Council of Texas
ERIS	Energy Resource Interconnection Service
FERC	Federal Energy Regulatory Commission’s
FPL	Florida Power & Light
GPUs	graphics processing units
ICT	information, and communication technology
ISO-NE	ISO New England
LGIA	Large Generator Interconnection Agreement
LOLE	loss of load expectation
MISO	Midcontinent Independent System Operator
NYISO	New York Independent System Operator
PACE	PacifiCorp East
PACW	PacifiCorp West
PG&E	Pacific Gas and Electric
PGE	Portland General Electric Company
PJM	PJM Interconnection
PSCO	Public Service Company of Colorado
RMSE	Root mean square error
RTO/ISO	Regional transmission organization/independent system operator
SCP	Santee Cooper, South Carolina Public Service Authority
SEAB	Secretary of Energy Advisory Board
SLAs	service-level agreements
SOCO	Southern Company
SPP	Southwest Power Pool
SRP	Salt River Project Agricultural Improvement and Power District
TPU	tensor processing unit
TVA	Tennessee Valley Authority

APPENDIX A: CURTAILMENT-ENABLED HEADROOM PER BALANCING AUTHORITY

Figure A.1. Curtailment Rate Versus Load Addition by RTO/ISO, MW

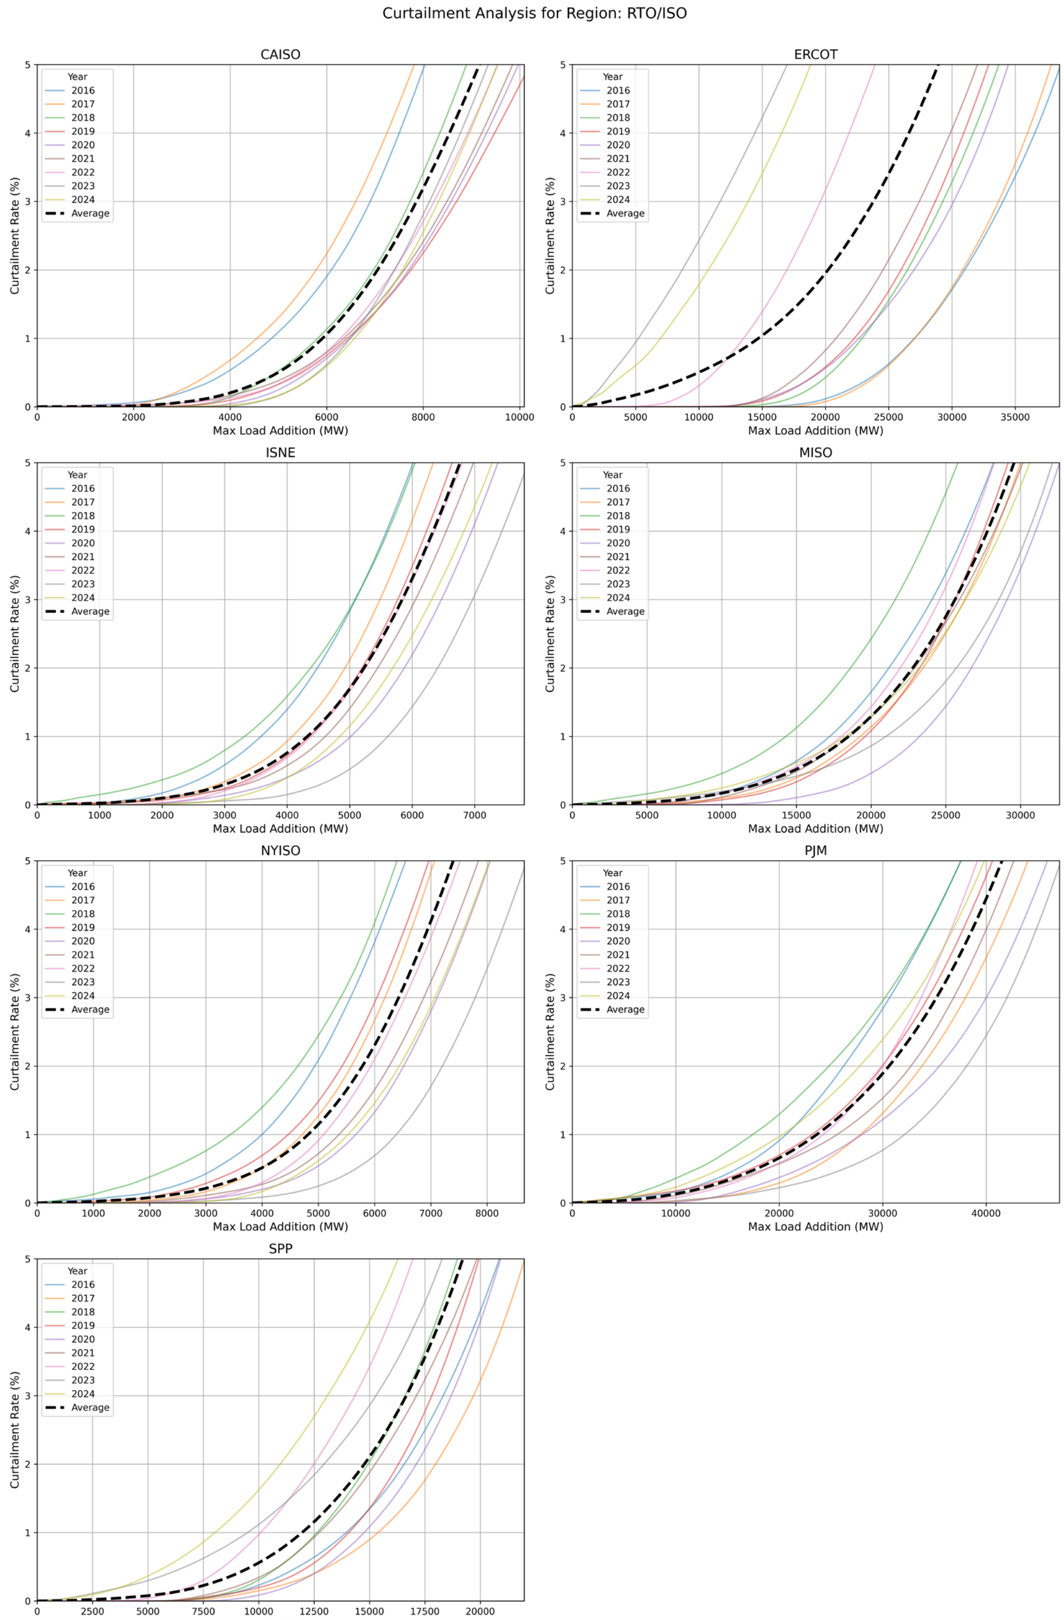

Figure A.2. Curtailment Rate Versus Load Addition by Non-RTO Southeastern Balancing Authority, MW

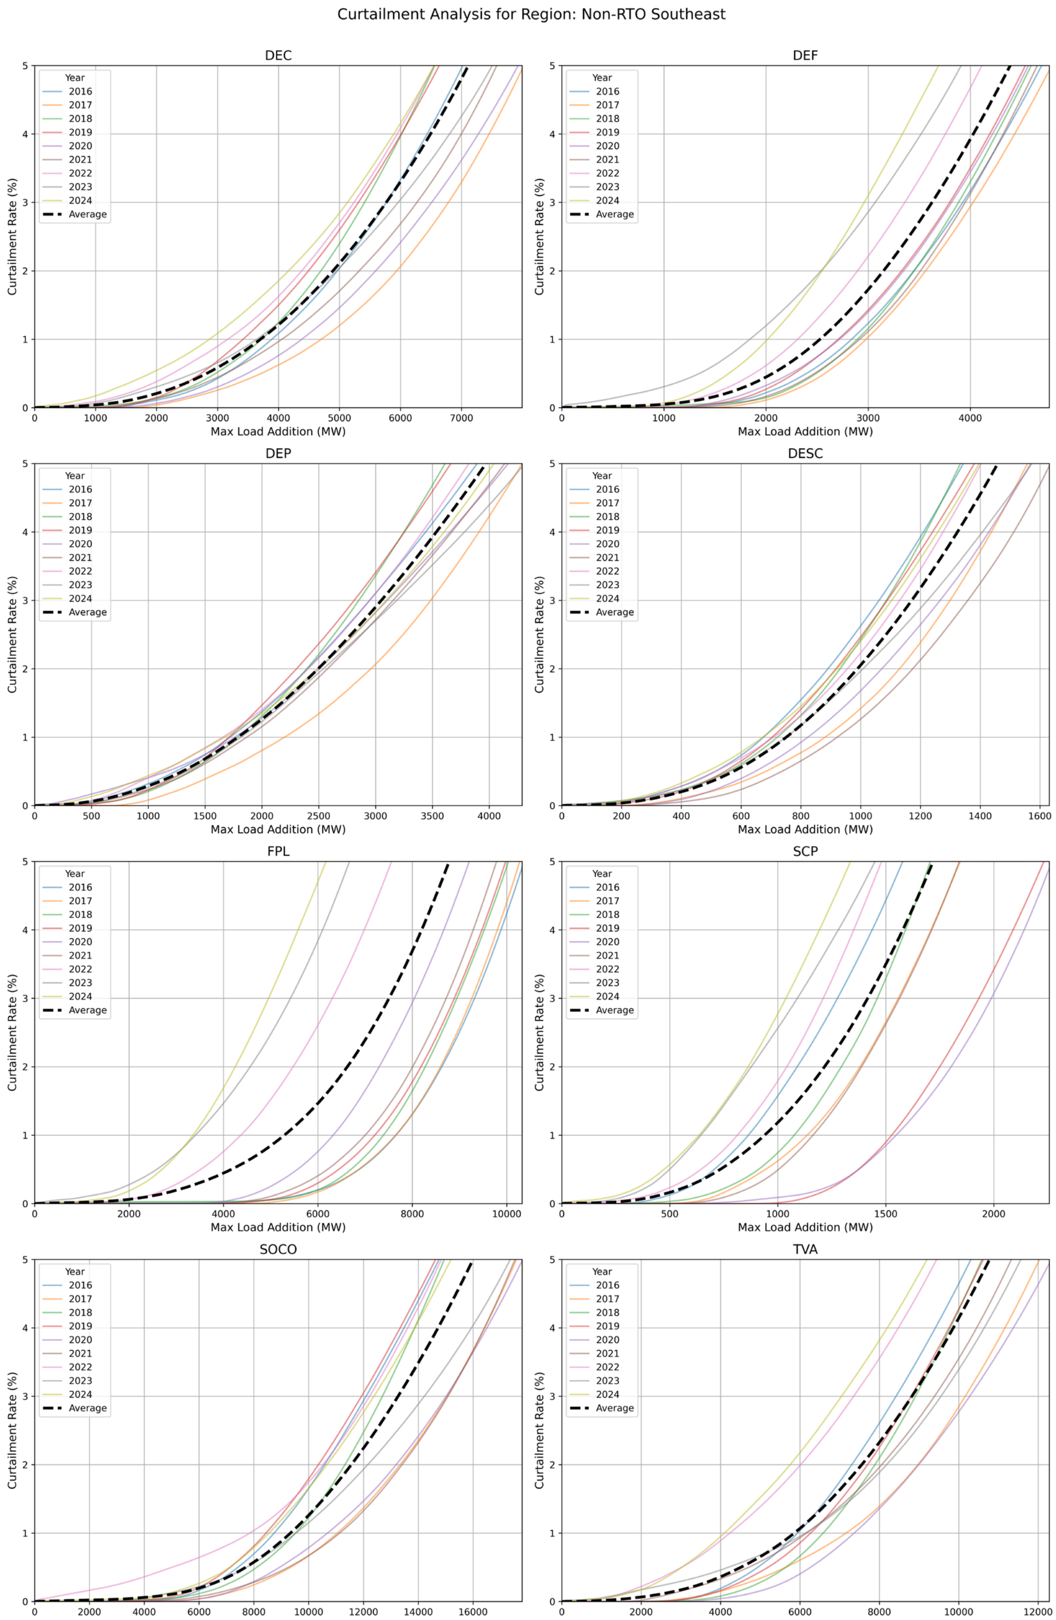

Figure A.3. Curtailment Rate Versus Load Addition by Non-RTO Western Balancing Authority, MW

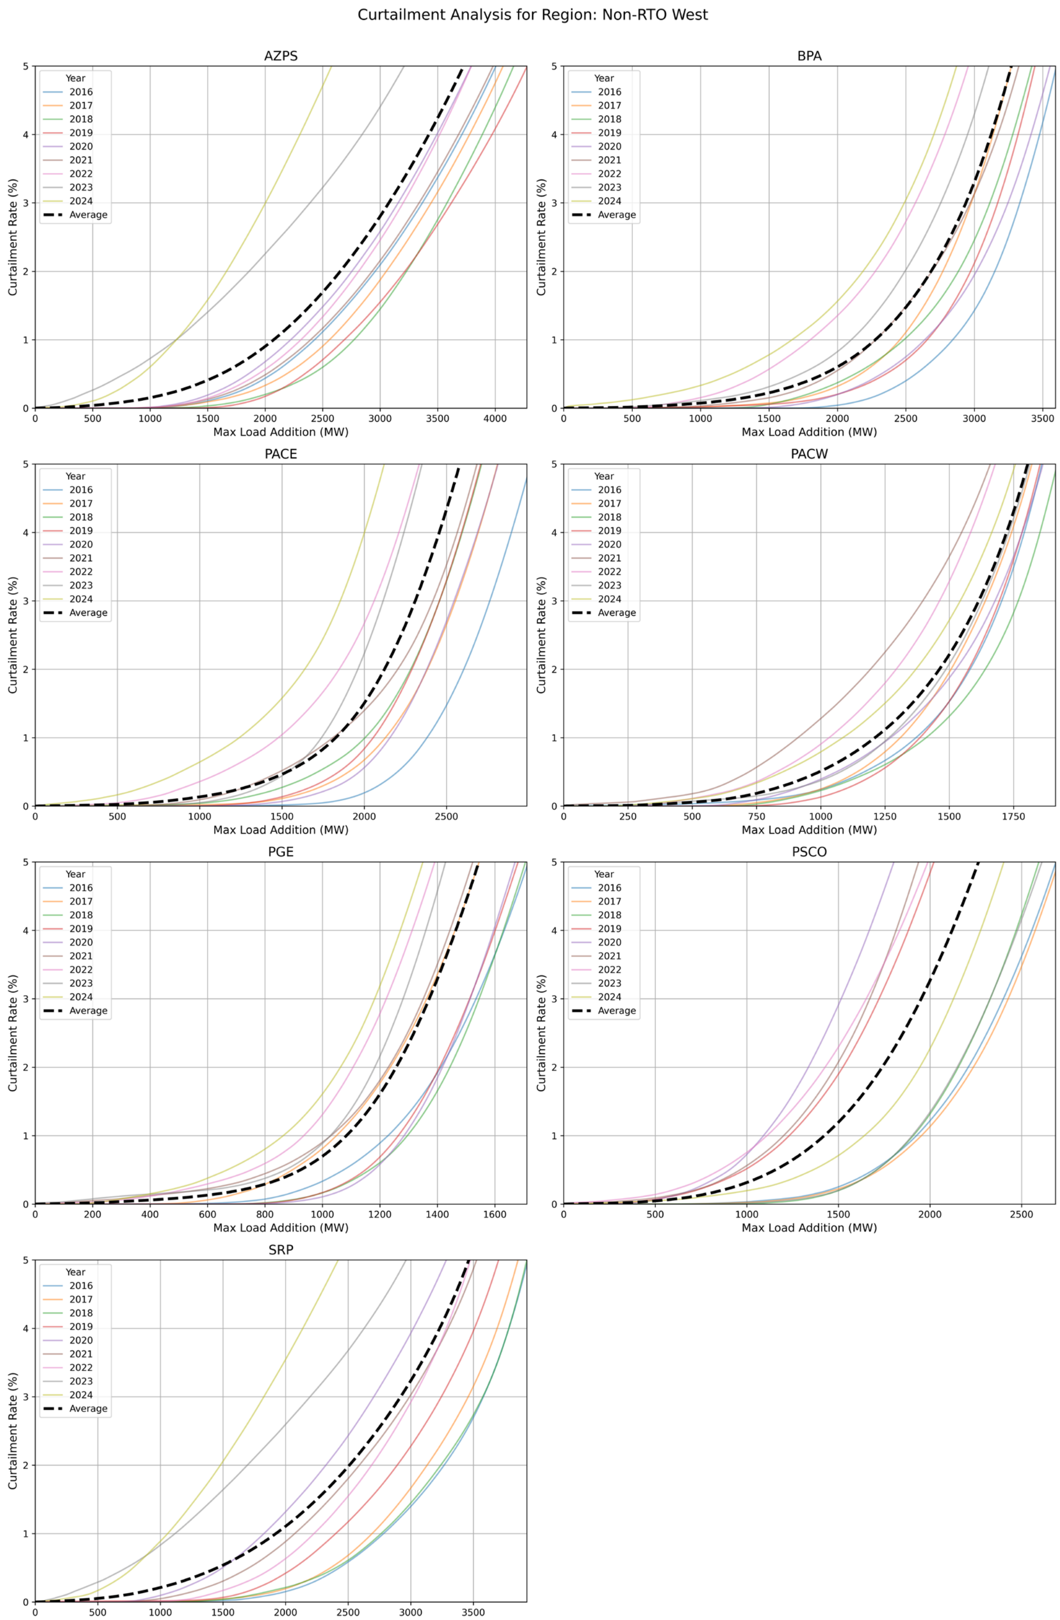

APPENDIX B: DATA CLEANING SUMMARY

The data cleaning process attempted to improve the accuracy of nine years of hourly load data across the 22 balancing authorities, including the following steps:

1. Data normalization
   - **Dates:** Date-time formats were verified to be uniform.
   - **Demand data:** Where the balancing authority had an "Adjusted demand" value for a given hour, this value was used, otherwise its "Demand" value was used. The final selected values were saved as "Demand" and a log was kept.
   - **BA labels:** Labels were mapped to align with widely used acronyms, including:
     - CPLE → DEP
     - DUK → DEC
     - SC → SCP
     - SWPP → SPP
     - SCEG → DESC
     - FPC → DEF
     - CISO → CAISO
     - BPAT → BPA
     - NYIS → NYISO
     - ERCO → ERCOT
2. Identifying and handling outliers
   - **Missing and zero values:** Filled using linear interpolation between adjacent data points to maintain temporal consistency.
   - **Low outliers:** Demand values below a predefined cutoff threshold (such as 0 or extremely low values inconsistent with historical data) were flagged. Imputation for flagged low outliers involved identifying the closest non-outlier value within the same balancing authority and time period and replacing the flagged value.
   - **Spikes:** Sudden demand spikes that deviated significantly from historical patterns were flagged. Corrections were applied based on nearby, consistent data.
   - **Erroneous peaks:** Specific known instances of demand peaks that are outliers (e.g., caused by reporting errors) are explicitly corrected or replaced with average values from adjacent time periods.
3. **Data validation**:
   - Seasonal and annual peak loads, load factors, and other summary statistics were computed and inspected to ensure no unexpected results. Max peaks were compared to forecasted peaks collected by FERC to ensure none were out of range.
   - Logs summarizing corrections, including the number of spikes or outliers addressed for each balancing authority, were saved as additional documentation.

APPENDIX C: CURTAILMENT GOAL-SEEK FUNCTION

Mathematically, the function can be expressed as

$$\frac{1}{N}\sum_{y=1}^{N}\left(\frac{Curtailment_y(L)}{L\,\cdot\,8{,}760}\cdot 100\right) = CurtailLimit$$

where

L	=	load addition in MW (constant load addition for all hours)
N	=	total number of years in the analysis (2016–2024)
$Curtailment_y(L)$	=	curtailed MWh for year y at load addition L
$L \cdot 8{,}760$	=	maximum potential energy consumption of the new load operating continuously at full capacity
$CurtailLimit$	=	predefined curtailment limit (e.g., 0.25%, 0.5%, 1.0%, or 5.0%).

For each hour t in year y, the curtailment is defined as

$$Curtailment_t(L) = \max(0, Demand_t + L - Threshold_t)$$

where

L	=	load addition being evaluated in MW
$Demand_t$	=	system demand at hour t in MW
$Threshold_t$	=	seasonal peak threshold applicable for hour t in MW (i.e., the maximum winter or summer peak across all years)

These hourly curtailments are aggregated to find the total annual curtailment

$$Curtailment_y(L) = \sum_{t\,\in\,T_y} Curtailment_t(L)$$

where

T_y	=	all hours in year y.

Replacing $Curtailment_y(L)$ in the original formula, the integrated formula becomes

$$\frac{1}{N}\sum_{y=1}^{N}\left(\frac{\sum_{t\,\in\,T_y}\max(0, Demand_t + L - Threshold_t)}{L\,\cdot\,8{,}760} * 100\right) = CurtailLimit$$