

California State Assembly

COMMITTEES ON UTILITIES AND ENERGY AND PRIVACY AND CONSUMER PROTECTION



COTTIE PETRIE-NORRIS AND
REBECCA BAUER-KAHAN
CHAIRS

Wednesday, January 28, 2026
1:30pm
1021 N Street, Room 1100

JOINT OVERSIGHT HEARING

AI's Energy Impacts

Findings

- *Artificial intelligence (AI)-driven data center growth is occurring at a scale and speed that outpaces most existing electricity planning frameworks. Forecasting methods, interconnection processes, and cost-allocation rules were not designed for large, uncertain, and highly concentrated loads arriving simultaneously across multiple regions.*
- *California has so far avoided rate impacts from data center growth, but that outcome is not guaranteed. High electricity prices, project-by-project review, and existing system capacity have limited impacts to date; however, growing transmission-level load heightens the risk of future cost shifts.*
- *New data center customers can reduce – or increase – rates depending on how infrastructure costs and revenue are allocated. Bill reductions only occur when new loads cover their full marginal and long-term costs; otherwise, ratepayers face higher bills, stranded assets, and “ghost load” risk.*
- *Uncertainty around data center demand materially increases financial and reliability risks for the grid. Overstated or duplicative load projections can drive unnecessary procurement and transmission investment, as evidenced in other regions, with costs borne by customers when projected load fails to materialize.*
- *Current regulatory pathways for backup generation and flexible demand may conflict with California’s environmental and equity goals. Expanded reliance on diesel backup generation, especially under streamlined siting and flexibility programs, could increase localized pollution and undermine emissions objectives if not carefully constrained.*

The increasing computational requirements of AI have fueled a wave of data center development, with an estimated 50 gigawatts (GW) of new data center capacity added to the global development pipeline in 2024 alone.¹ This is an unprecedented level of demand in an extraordinarily short amount of time. For context, 50 GW is roughly the entire electricity demand of California during its hottest, most stressed hours – essentially the load of the whole state during extreme heat events.² This rapid expansion of data centers is driving an unprecedented demand for energy. In 2024, data centers in the U.S. used approximately 200 terawatt-hours (TWh) of electricity, roughly what it takes to power Thailand for a year.³ As a result, access to electricity is a top concern shaping the development of digital infrastructure.

At the same time, access to *affordable* electricity is a top concern shaping the public discourse nationwide, particularly in areas with large data center concentrations. Across the country – from Georgia⁴ to Texas,⁵ Ohio⁶ to New York,⁷ to the hardest hit area in Northern Virginia⁸ – the impact of AI-fueled data center growth has become a kitchen-table issue, replete with yard signs,⁹ local moratoria,¹⁰ and high electricity bills fueling a frustrated public.^{11,12} According to the U.S. Energy Information Administration, residential utility bills rose 5% nationally by October 2025;¹³ however, this was much higher in data center-rich regions, with Virginia, Ohio, and Illinois experiencing double-digit increases.¹⁴

Yet these challenges have not yet surfaced in California. Despite being one of the top data center markets in the country,¹⁵ California has not seen an increase in electricity rates due to data center growth, according to the California Public Utilities Commission (CPUC).¹⁶ Although data

¹ iMasons. *State of the Digital Infrastructure Industry – Annual Report 2025*. 2025.

<https://imasons.org/publications/state-of-the-digital-infrastructure-industry-annual-report-2025/>

² September 6, 2022 reached a peak load of approximately 52 GWs, the highest on record. California's average annual demand is between 25-30 GWs.

³ James O'Donnell and Casey Crownhart, "We did the math on AI's energy footprint. Here's the story you haven't heard." *MIT Technology Review*, May 20, 2025; <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>

⁴ Kala Hunter, "How Georgia became the 'wild west' of data centers. Is transparency on the horizon?" *Georgia Public Broadcasting*, September 10, 2025; <https://www.gpb.org/news/2025/09/10/how-georgia-became-the-wild-west-of-data-centers-transparency-on-the-horizon>

⁵ Jason Plautz, "Trump, atoms, AI and the Texas data center gusher," *Politico*, January 4, 2026; https://www.politico.com/news/2025/12/23/fermi-america-data-center-amarillo-texas-00701800?_sp_pass_consent=true

⁶ Cliff Pinckard, "Ohio ranks fifth in the country for data centers. To power them, they're going nuclear." *Cleveland.com* (formerly *The Plain Dealer*), January 20, 2026. <https://www.cleveland.com/metro/2026/01/ohio-ranks-fifth-in-the-country-for-data-centers-to-power-them-theyre-going-nuclear-the-wake-up-for-tuesday-jan-20-2026.html>

⁷ Johan Sheridan, "Data center boom straining power grid as New York asks who should pay," *ABC News 10*, October 24, 2025; <https://www.news10.com/news/rising-energy-costs-new-york/>

⁸ Liam Bowman, "Concerns over data centers drive local election in Northern Virginia," *The Washington Post*, November 9, 2025. <https://www.washingtonpost.com/dc-md-va/2025/11/08/prince-william-county-gainseville-election/>

⁹ Ryan Murphy and Emily Feng, "Why more residents are saying 'No' to AI data centers in their backyard;" *NPR News*; July 17, 2025; <https://www.npr.org/2025/07/17/nx-s1-5469933/virginia-data-centers-residents-saying-no>

¹⁰ Anna Lynn Winfrey, "Another Columbus-area community passes 90-day moratorium on data centers," *Columbus Dispatch*, December 15, 2025; <https://www.dispatch.com/story/news/local/communities/dublin/2025/12/15/second-central-ohio-community-washington-twp-to-ban-data-centers-for-90-days-wants-dublin-to-join-in/87719185007/>

¹¹ Marc Levy and Jesse Bedayn, "Voters' anger at high electricity bills and data centers looms over 2026 midterms," *The Associated Press*, November 8, 2025; <https://apnews.com/article/2026-election-utility-bills-ai-data-centers-13703f61d1397612fd067e69b9093116>

¹² J. Saul, et al., "AI Data Centers are Sending Power Bills Soaring," *Bloomberg Technology*, September 29, 2025; <https://www.bloomberg.com/graphics/2025-ai-data-centers-electricity-prices/>

¹³ As compared to October 2024; EIA information released on December 23, 2025 with data for October 2025. <https://www.eia.gov/electricity/monthly/update/>

¹⁴ Kimball and Cortes, "Data centers are concentrated in these states. Here's what's happening to electricity prices." *CNBC*, November 14, 2025; <https://www.cnbc.com/2025/11/14/data-centers-are-concentrated-in-these-states-heres-whats-happening-to-electricity-prices-.html>

¹⁵ See Figure 2 below

¹⁶ CPUC, "California Data Center Development & Energy Needs: FAQs," 12.22.2025.

centers use significant amounts of electricity, they still represent a small share of California’s overall demand and are typically sited and managed in ways that have limited their impact on statewide electricity rates. In investor-owned utility (IOU) territories, the CPUC notes that commercial and industrial (C&I) rate structures often require project-by-project review, which helps identify data center-related system upgrades and assign those costs to the data center driving them. In addition, California’s above-average electricity prices can discourage electricity-hungry data center development, except where proximity to users is required to meet latency needs.

These trends may be changing. For instance, Pacific Gas & Electric (PG&E) residential electricity rates were 11% lower by January 2026 than in January 2024,¹⁷ in contrast to other California IOUs. But PG&E also reported a rise in future data center development during this time, from 5.5 GWs in February 2025¹⁸ to 10 GWs by July 2025.¹⁹ It will take years for this data center load to come online. But as that happens, the generation and new transmission lines to serve that load will also come online – the cost of which is traditionally shared amongst all customers.

The emergence of extensive data center electricity load, driven by the growth in AI computing, marks a rare opportunity in the energy industry to reassess existing rate design and customer protections. The promise of data center energy consumption is alluring: that increased electricity sales from data centers will cover all the new costs to serve those data centers and may even offset existing system costs. Such is the story for other large loads, like transportation electrification (TE), with the California Public Advocates Office (PAO) finding TE “may cause downward pressure on electric rates.”²⁰ The timing can feel almost too perfect: as California undertakes systemwide decarbonization and infrastructure renewal, a new, well-funded customer emerges, apparently willing²¹ to shoulder the cost of long-deferred investments. However, as the PAO cautions for TE, and as applies equally to data-center development, “to achieve this downward pressure on rates, effective management of multiple factors will be required, including efficient infrastructure buildout and cost constraints.”²²

California is better situated than many other parts of the country to accommodate data center driven load growth. California’s early adoption of resource planning, for both decarbonization and reliability, have provided a cushion, such that the California grid has significant capacity available in periods of low customer demand.²³ California also boasts a legacy of innovation,

¹⁷ PG&E press release, “PG&E to Lower Electric Prices on Jan. 1, Fourth Decrease in Two Years,” December 30, 2025; <https://investor.pgecorp.com/news-events/press-releases/press-release-details/2025/PGE-to-Lower-Electric-Prices-on-Jan--1-Fourth-Decrease-in-Two-Years/default.aspx>

¹⁸ PG&E press release, “PG&E Accelerating Connection of New Data Centers throughout Northern and Central California,” February 13, 2025; <https://investor.pgecorp.com/news-events/press-releases/press-release-details/2025/PGE-Accelerating-Connection-of-New-Data-Centers-throughout-Northern-and-Central-California/default.aspx>

¹⁹ PG&E press release, “PG&E Data Center Demand Pipeline Swells to 10 Gigawatts with Potential to Unlock Billions in Benefits for California,” July 31, 2025; <https://investor.pgecorp.com/news-events/press-releases/press-release-details/2025/PGE-Data-Center-Demand-Pipeline-Swells-to-10-Gigawatts-with-Potential-to-Unlock-Billions-in-Benefits-for-California/default.aspx>

²⁰ Pg. 14, Public Advocates Office, *Distribution Grid Electrification Model 2025 – Study and Report*, October 2025. <https://www.publicadvocates.cpuc.ca.gov/-/media/cal-advocates-website/files/press-room/reports-and-analyses/251030-public-advocates-office-distribution-grid-electrification-model-2025.pdf>

²¹ Amrith Ramkumar, “Microsoft Makes New Data-Center Pledges After Local Backlash,” *The Wall Street Journal*, January 13, 2025.

²² Pg. 14, PAO October 2025, *Ibid*.

²³ The transmission system uses less than 40% of its capacity on average. Pg. 8, L. Min, et al., *Powering AI at Speed in California*, Stanford Precourt Institute for Energy, September 2025; <https://drive.google.com/file/d/14NhXPKynQCjbBC99xVk6tA5pqCD9U052/view>

from early data center adoption in Santa Clara to energy efficiency requirements²⁴ to historic approaches to ensure adequate capacity.²⁵ The pace of AI deployment has strained most aspects of current energy governance and cost allocation; yet California’s high electricity costs may have provided the rare benefit of slowing recent in-state development, enabling California regulators the time to apply effective safeguards and lessons learned from other states.

The purpose of this hearing is to discuss the impact of AI and corresponding data center growth on the energy sector, both nationwide and in California. The hearing will detail impacts specific to California’s electricity system – from maintaining renewables goals, to ensuring system reliability, to protecting affordability. While data center electricity and water consumption are related, sometimes inversely,²⁶ this hearing will focus just on the electric side. Other committees may hold hearings later this session focused on other resource constraints, including water. Panelists will be asked to speak to needed safeguards to ensure data center development in California does not lead to runaway energy costs, stranded assets, or negative environmental impacts.

I. The Data Center Ecosystem.

Data centers pose an emerging challenge for California’s energy system. But not all data centers are built the same. Their power demands, latency needs, and location requirements vary significantly depending on their function.

Some key players include:

- Enterprise Data Centers – Private data centers that serve proprietary digital operations (e.g., banks, healthcare organizations, governments). The homeowners of the data center world.
- Colocation Data Centers – Facilities where multiple tenants rent space to host their own servers and technology equipment, without building or managing their own space. These are the landlords of the data center world.
- Hyperscale Cloud Providers – Large-scale operators that deliver rapid computing power and cloud-based storage to support AI and other data-intensive applications. Examples include Amazon Web Services (AWS), Microsoft Azure, and Google Cloud. These are like colleges or hospitals, where the entire campus is owned and managed by a single entity; they sell services rather than space.
- Edge Data Centers – Smaller, distributed facilities located closer to end-users or data sources. They support processes that rely on low latency. They are critical for services like autonomous vehicles and content delivery, often operating in conjunction with larger cloud or colocation providers to offload demand from centralized data centers. These are like the neighborhood corner store – built for speed and convenience, not scale.

²⁴ CEC, “Computer Rooms & Data Centers Fact Sheet,” August 2023. https://www.energy.ca.gov/sites/default/files/2023-09/2022_CEC-Computer_Room_and_Data_Centers_ADA.pdf

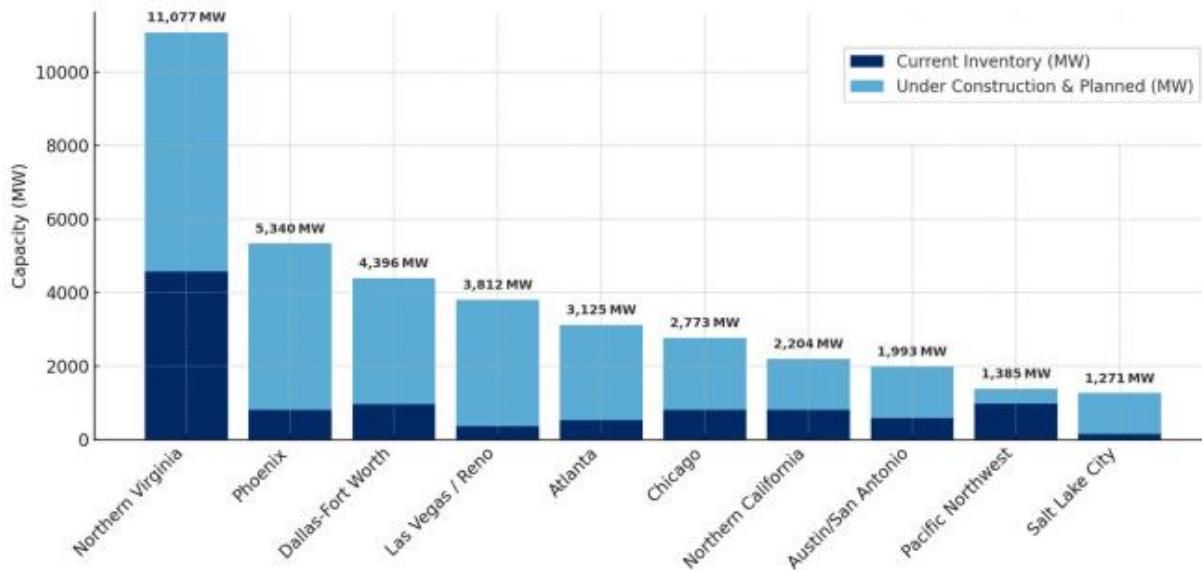
²⁵ Namely our Resource Adequacy program; <https://www.cpuc.ca.gov/industries-and-topics/electrical-energy/electric-power-procurement/resource-adequacy-homepage>

²⁶ For data center cooling, specifically, the less water you use (water towers), the more electricity you have to use (air chillers), and vice versa.

Data centers rely on long-term contracts and make money by leasing space and selling computing and storage services. Because building data centers requires very large upfront investments – often hundreds of millions or even billions of dollars – strong financing is necessary. Their profitability depends on factors such as energy costs, customer demand, lease terms, and how efficiently the facilities are operated.

The emergence of AI has led to a shift in data center architecture, from traditional centralized computing toward more decentralized networks. Traditionally, markets such as Loudoun County, Virginia and Silicon Valley, California have been key data center hubs due to their infrastructure and connectivity, as shown by the dark blue bars in Figure 1.²⁷ As demand for real-time processing increases, more decentralized approaches – such as edge computing – are changing where data centers are built. These trends are expanding data center development beyond traditional metro areas into regions with more favorable energy costs, permitting conditions, and transmission access, as shown by lighter blue bars in Figure 1. These trends have influenced the scale and location of data center development in California.

Figure 1. Top 10 U.S. data center markets by future power capacity (including current inventory and under construction & planned).



In California, data center development has been concentrated in Santa Clara and Los Angeles, due to proximity to fast fiber-optic connections and lower electricity rates served by those cities' municipal utilities. More recently, Pacific Gas & Electric has received a surge of requests for service to new data centers. Initially those requests were in and around the South Bay; however, the utility reports expanding development of large load customers in other regions of its service territory.²⁸

II. The Dawn of Data.

²⁷ L. Min, W. Chueh, and I. Ehrenpreis, *Powering AI at Speed in California*, Precourt Institute for Energy [Stanford University] (Sep. 2025), p. 7. <https://drive.google.com/file/d/14NhXPkynQCjbBC99xVk6tA5pqCD9U052/view>

²⁸ Michael Medieros, "PG&E Written Testimony," Letter to the Little Hoover Commission; December 8, 2025. <https://lhc.ca.gov/wp-content/uploads/PGE-LHC-Written-Testimony.pdf>

In 2024, global data generation surpassed 149 zettabytes (trillions of gigabytes).²⁹ As put into perspective by researchers at University of Texas at Austin: “if each gigabyte were a single-page document, the resulting stack of [these] papers could reach the Moon and back ~20 times.”³⁰ This scale of data reflects the rapid growth of connected devices, AI-driven queries, and cloud-based services that depend on data centers for real-time processing and storage.

It is the growth in AI-computing that has driven most of the growth in data center development and construction. However, the data and energy needs vary depending on the type of AI-computing. Traditionally, there are two types of AI workloads: training and inference. Training is when AI *learns* from large datasets; whereas inference is when that trained model is *used* by customers.

Data Hogs are Energy Hogs. As noted in reports by the Assembly Committee on Privacy and Consumer Protection,³¹ staggering quantities of data are required to train the most advanced AI models. For example, GPT- 4 – the large language model (LLM) embedded in ChatGPT 4 – is reported to have been trained on roughly 10 trillion words of text.^{32,33} Adjusting the model’s 1.8 trillion parameters continuously as it was exposed to this vast corpus required trillions upon trillions of computations, which were performed by running approximately 25,000 expensive, energy-consuming microchips for nearly 100 days nonstop, at an estimated cost of \$63 million.³⁴ It is estimated this training consumed 50 gigawatt-hours (GWhs) of

Box 1: A Note on Data Availability

Details about AI’s current and future energy impact are fairly murky. As noted by researchers at Lawrence Berkeley National Lab, “the lack of direct energy data available in a sector with rapidly evolving technologies limits this analysis...”^a

The researchers go on to note, “very few companies report actual data center electricity use and virtually none report it in context of IT characteristics...”^b

As such, values reported on AI or data center energy use in this document are *estimates* with large uncertainty.

Such vagueness may be fine for purposes of this discussion, but as highlighted below, can greatly limit state energy planners’ abilities to accurately forecast grid system needs. This dearth of information motivated both AB 222 (Bauer-Kahan, 2025) and the recently introduced AB 1577 (Bauer-Kahan, 2026).

a: pg. 7, Shehabi, Arman, et al. "2024 United States Data Center Energy Usage Report." December 2024, Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-2001637; https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report_1.pdf

b: pg. 68, Shehabi, *Ibid.*

²⁹ D. Ewim et al. “Impact of Data Centers on Climate Change: A Review of Energy Efficient Strategies”. *The Journal of Engineering and Exact Sciences* 9.6 (2023), 16397–01e. https://www.researchgate.net/figure/Global-annual-GHG-emissions-from-major-industries-Source-Kilgore-2023_fig1_373295068

³⁰ N. Ling, et. al, *Data Center Growth in Texas: Energy, Infrastructure, and Policy Pathways*; Bureau of Economic Geology, The University of Texas at Austin; December 2025; https://www.beg.utexas.edu/files/cee/Data_Center_White_Paper_BEG.pdf

³¹ Most recently in their background paper for their December 8, 2025, joint hearing on “Artificial Intelligence and Copyright” https://apcp.assembly.ca.gov/system/files/2025-12/background-paper-dec-8-stanford-ai-and-copyright-joint-info-hearing_updated.pdf

³² Schreiner, “GPT-4 architecture, datasets, costs and more leaked,” *The Decoder* (Jul. 11, 2023), <https://the-decoder.com/gpt-4-architecture-datasets-costs-and-more-leaked/>

³³ Begum, “OpenAI Releases GPT-4: A Smarter and Faster AI-Language Model with ‘Human-level Performance,’” *Vocal Media* (2023), <https://vocal.media/01/open-ai-releases-gpt-4-a-smarter-and-faster-ai-language-model-with-human-level-performance>.

³⁴ Ludvigsen, “The carbon footprint of GPT-4,” *Medium* (Jul. 18, 2023), available at <https://medium.com/data-science/the-carbon-footprint-of-gpt-4-d6c676eb21ae>.

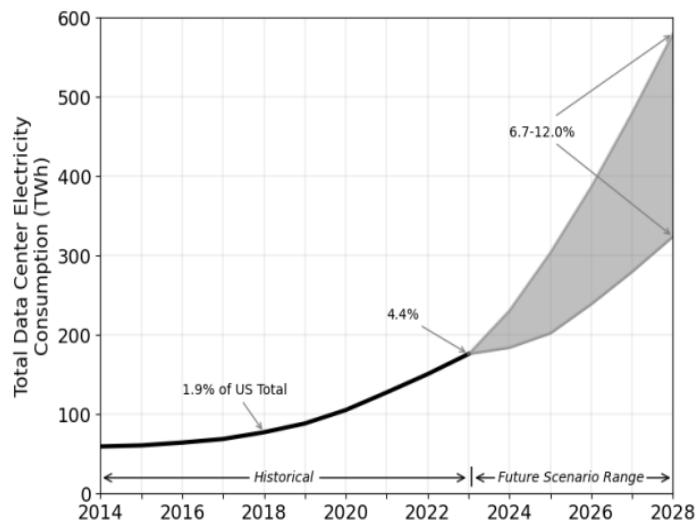
energy, enough to power San Francisco continuously for three days.³⁵

However, as noted in *MIT Technology Review*, “inference, not training, represents an increasing majority of AI’s energy demands. … It is now estimated that 80-90% of computing power for AI is used for inference.”³⁶ While rough estimates exist for how much energy each AI-query consumes – famously, OpenAI CEO Sam Altman has stated each ChatGPT query consumes approximately 0.34 watt-hours of electricity, enough to power a lightbulb for a few minutes³⁷ – in reality, the type and size of the model and your selected output (language, image, video) can make one query thousands of times more energy-intensive and emissions-producing than another. One billion of these inquiries every day for a year – well below the self-reported values of OpenAI³⁸ – would mean over 109 GWhs of electricity, enough to power 10,400 U.S. homes for a year, according to researchers writing in *MIT Technology Review*.³⁹

The type of AI workloads will also strongly shape where future data centers are built. Training large AI models requires enormous computing power but is relatively insensitive to latency, allowing these facilities to be sited in remote areas with access to abundant, low-cost electricity. In contrast, inference workloads support real-time applications and are highly latency-sensitive, which favors data centers located close to end users. As a result, energy-intensive training facilities are increasingly being built in rural or industrial regions, while smaller edge and micro data centers are expanding in urban areas to support inference-driven services such as autonomous driving and immersive media.

Exponential growth doesn’t feel exponential at first. As shown in Figure 2, U.S. data center annual energy use prior to 2016 was relatively stable at about 60 TWh. Despite the construction of new data centers to serve the rise of cloud-based online services, such as Netflix, increases in efficiency kept this growth consumption relatively flat. However, starting in 2017, electricity consumption accelerated nationwide primarily due to AI computing. By 2023, data centers were 4.4% of total U.S. energy consumption. By 2028, data centers are predicted to use between 6.7% and 12% of the country’s

Figure 2. Total U.S. data center electricity use from 2014-2028.¹⁰



³⁵ James O’Donnell and Casey Crownhart, “We did the math on AI’s energy footprint. Here’s the story you haven’t heard.” *MIT Technology Review*, May 20, 2025; <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>

³⁶ O’Donnell and Crownhart, *Ibid*.

³⁷ Kwan Wei Kevin Tan, “Sam Altman says the energy needed for an average ChatGPT query can power a lightbulb for a few minutes,” *Business Insider*, June 10, 2025; <https://www.businessinsider.com/how-much-energy-does-chatgpt-use-average-query-watts-altman-2025-6>

³⁸ Which in July 2025 estimated 2.5 billion prompts every day; Emma Roth, “OpenAI says ChatGPT users send over 2.5 billion prompts every day,” *The Verge*, July 21, 2025; <https://www.theverge.com/news/710867/openai-chatgpt-daily-prompts-2-billion>

³⁹ O’Donnell and Crownhart, *Ibid*.

electricity.⁴⁰ At that point, AI alone could consume as much electricity annually as 22% of all U.S. households.⁴¹

Just how utilities, regulators, and tech companies will meet this growing demand is uncertain.

Box 2: A Primer on CA's System Planning

California energy planning is conducted via layered programs administered by several state entities, chiefly: the CEC, the CPUC, and CAISO. These entities have a memorandum of understanding to provide clarity and specificity in how information and data are transmitted and utilized by each entity. The main planning processes at each organization include:

- The Integrated Energy Policy Report (IEPR) – *CEC* – every 2 years: forecasts all aspects of energy supply and demand. The demand forecast from the IEPR is a primary input for planning at the CPUC and CAISO.
- Integrated Resource Plans (IRP) – *CPUC* – every 2 years: process to ensure long-term resource procurement. Energy purchases are still conducted by individual energy suppliers, but the IRP ensures adequate resources are bought to meet the IEPR forecast and electricity sector greenhouse gas goals. Serves as a primary input for planning at the CAISO.
- The Transmission Planning Process (TPP) – *CAISO* – annually: estimates the transmission infrastructure needed to match the IRP portfolio provided annually by the CPUC with the demand forecast provided by the CEC. Once the final TPP is approved by the CAISO Board, the resulting approved infrastructure begins the development process, including permitting, licensing, and competitive solicitations as applicable.

New resources must be procured, and new transmission is likely needed to serve this load, all at cost to either ratepayers or the tech companies themselves. Yet given the large uncertainty projected (grey, shaded area in Figure 2), and the lack of transparent data on current and future data center energy needs (See Box 1), projections of future energy demands or estimates of resultant emissions are simply inadequate or inaccurate. This forces energy planners to assemble a puzzle with countless missing pieces.

III. Predicting the Unpredictable.

Last Wednesday, January 21, 2026, the California Energy Commission (CEC) adopted their energy demand forecast for 2025-2045.⁴² This forecast is foundational for resource procurement and system planning in the state, as the output of the forecast feeds into sequential planning streams at the CPUC and California Independent System Operator (CAISO),⁴³ as detailed in Box 2. In other words, the forecast is critical in determining how much energy generation to require utilities to purchase, where to upgrade power lines, and how to prevent future blackouts.

The 2025 IEPR provides electricity and gas demand forecasts which reflect expected impacts from economic projections, including data center growth, electric vehicle adoption, and other inputs.⁴⁴ The final 2025 adopted plan anticipates almost 20 GWs of load growth over the next 20 years in

⁴⁰ Shehabi, Arman, et al. "2024 United States Data Center Energy Usage Report." December 2024, Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-2001637; https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report_1.pdf

⁴¹ O'Donnell and Crownhart, *Ibid*.

⁴² Item 6, CEC Business Meeting Agenda, January 21, 2026;

<https://efiling.energy.ca.gov/GetDocument.aspx?tn=268217&DocumentContentId=105381>

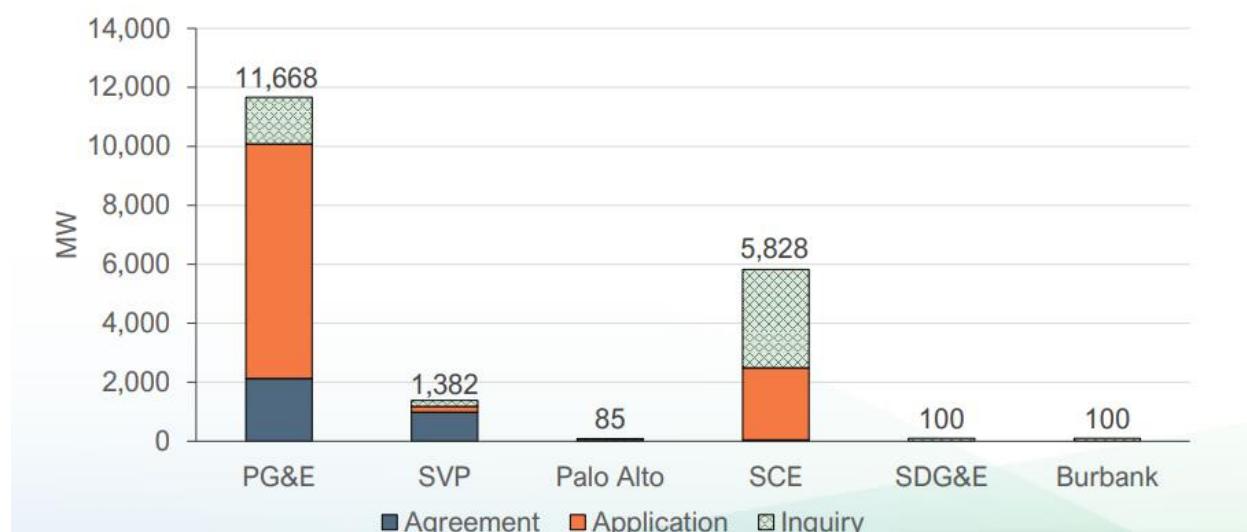
⁴³ a non-profit public benefit corporation regulated by the Federal Energy Regulatory Commission (FERC) who maintains operational control of ~80% of the state's transmission grid.

⁴⁴ CEC, "Resolution of the CEC Adopting the California Energy Demand Forecast, 2025-2045," <https://www.energy.ca.gov/filebrowser/download/9208?fid=9208>

CAISO's area, with close to 5 GWs arising from data centers. While significant, the adopted forecast shows a downward adjustment of almost 5 GWs by 2040 between the 2024⁴⁵ and 2025 IEPRs, largely due to the treatment of data centers' "known loads" in the 2025 models. If "known loads" are included, the peak forecast swells by almost 10 additional gigawatts by 2045 – a significant range of uncertainty.

All Models are Wrong, But Some are Useful.⁴⁶ As reported in *Politico*, Nick Fugate, the CEC's lead forecaster, has noted that uncertainty "has increased significantly" in the last few IEPRs.⁴⁷ "Known loads" are new inputs into the 2025 IEPR modeling, and are based on IOU-submitted project-level information, including energization requests and expected in-service dates. As shown in Figure 3, significant uncertainty exists around these known load requests.⁴⁸ Because these known loads lack a historical record within the CEC, the final 2025 forecast "Set Agreement" between the CEC, CPUC, and CAISO declines to advance the forecast with the known load values. Instead, the "Set Agreement" notes the CEC staff will review historical known loads throughout 2026 to "confirm assumptions informed by the IOUs" and "will continue to monitor data center applications for energization..."⁴⁹

Figure 3. Capacity Requests, as reported to the CEC from each utility as of August 2025.⁴⁸ (Key: PG&E = Pacific Gas & Electric; SVP = Silicon Valley Power; SCE = Southern California Edison; SDG&E = San Diego Gas & Electric)



Interestingly, the adopted 2025 agreement permits the use of known loads to inform local transmission planning at the CAISO but not reliability planning at the CPUC.⁵⁰ Parties representing large energy users, including data centers, have raised concern that this bifurcated inclusion of known loads will likely be "insufficient to serve projects."⁵¹ However, the "Set

⁴⁵ ~ 67 GWs anticipated by 2040 in the 2024 IEP (pg. 38, <https://efiling.energy.ca.gov/GetDocument.aspx?tn=266141>) versus ~62 in the 2025 IEP (slide 9, <https://www.energy.ca.gov/filebrowser/download/9328?fid=9328>)

⁴⁶ Quote largely attributed to statistician George Box.

⁴⁷ Noah Baustin, "California's new grid hog isn't who you think," *Politico*, January 21, 2026; <https://www.politico.com/newsletters/california-climate/2026/01/21/californias-new-grid-hog-isnt-who-you-think-00740515>

⁴⁸ Slide 5, presentation by CEC Manager of Demand Analysis Heidi Javanbakht to the Little Hoover Commission, December 11, 2025; <https://lhc.ca.gov/wp-content/uploads/5-Heidi-Javanbakht-CEC.pdf>

⁴⁹ Pg. 6, CEC, "Single Forecast Set Agreement," 2025 IEPR Forecast Supporting Documentation, filed January 23, 2026; <https://efiling.energy.ca.gov/GetDocument.aspx?tn=268288&DocumentContentId=105461>

⁵⁰ Pg. 7, CEC Set Agreement, *Ibid*.

⁵¹ Meredith Alexander, "Comments on the 2025 IEPR Draft Electricity Demand Forecast," January 9, 2026; <https://efiling.energy.ca.gov/GetDocument.aspx?tn=268213&DocumentContentId=105376>

Agreement” notes the CEC will closely monitor known load energizations throughout 2026, with the possibility of intervention in advance of 2027 should load growth outpace current forecast expectations.

While the demand forecast is independent of system cost, significant financial consequences do exist if the models prove grossly inaccurate. For instance, if the known loads were included in the planning forecast, at the current CPUC resource adequacy benchmark of \$11.53/kilowatt-month,⁵² the cost impact to serve that known load is in the range of \$500 million for 2026-2027.⁵³ This is a significant cost impact, especially given the uncertainty of whether this load will even materialize to pay for these costs.

The Ghost in the Model. Within the PJM Interconnection⁵⁴ – the largest power grid operator in the U.S., serving approximately 67 million customers from Chicago to New Jersey, including the “data center alley” of Northern Virginia – forecasted load growth has emerged as a key driver in utility cost surges over the last year. According to PJM’s Independent Market Monitor, load growth from new data centers was responsible for roughly \$9.3 billion of the \$14 billion regional capacity market bill for 2025-26, with costs escalating further over the next 2 years. PJM is also projecting an additional \$10+ billion in transmission expenditures, largely caused by data centers.⁵⁵

Yet many stakeholders are raising concern that these projections are inflated, with PJM recently adopting stricter vetting for its data center load forecast in response.⁵⁶ When projected energy demand doesn’t materialize, it is colloquially termed “ghost load.” Grid planners may see “ghost load” for several reasons. Data center developers often forum shop across multiple states and utility territories at the same time, looking for the best tax incentives, access to land, quick electrical hookups, and favorable regulatory conditions. Because utilities plan for and report this potential load individually, the same data center can be counted multiple times across different jurisdictions – even if it is never ultimately built. Utilities and transmission operators then make investment decisions based on these projections and report them to their grid planners. Those planners, in turn, use the inflated estimates to procure capacity and plan new transmission, which can drive higher electricity costs for customers.

In response to the emergence of “ghost loads,” various states have adopted strategies to make developers provide more upfront investments before utilities spend any ratepayer money. In Illinois, developers are required to provide a \$1 million deposit.⁵⁷ In Indiana and Ohio, long-term take-or-pay contracts are required where developers pay for at least 85% of the transmission service they request for at least 10 years, even if their actual usage is less.⁵⁸ And recently FERC ordered the PJM Interconnection to update its tariff to provide for interconnection of customers serving co-located load (i.e., data centers with onsite generation).⁵⁹ This FERC order largely

⁵² Table 2, CPUC, “Market Price Benchmark Calculations 2025,” October 1, 2025; <https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/energy-division/documents/community-choice-aggregation-and-direct-access/2025-mpbs.pdf>

⁵³ using an ~18% planning reserve margin, estimating ~\$465-\$565 million in additional cost

⁵⁴ “PJM” comes from the original name: the “Pennsylvania-New Jersey-Maryland Interconnection”

⁵⁵ IMM for PJM, “Analysis of the 2027/2028 RPM Base Residual Auction Part A,” January 5, 2026.

⁵⁶ Ethan Howland, “PJM trims near-term load forecast on stricter data center vetting, economic outlook,” *Utility Dive*, January 15, 2026; <https://www.utilitydive.com/news/pjm-interconnection-load-forecast-data-centers/809717/>

⁵⁷ John Pletz, “AI gold rush fuels ComEd crackdown on data center speculators,” *Chicago Business*, June 24, 2025; <https://www.chicagobusiness.com/utilities/comed-seeks-bigger-deposits-ai-era-data-centers>

⁵⁸ Nick Evans, “Ohio Manufacturers’ Association challenges new utility billing for data centers,” *Ohio Capital Journal*, November 13, 2025; <https://ohiocapitaljournal.com/2025/11/13/ohio-manufacturers-association-challenges-new-utility-billing-for-data-centers/>

⁵⁹ FERC Order to PJM, December 18, 2025, Docket EL25-49-000; <https://www.ferc.gov/media/e-1-el25-49-000-0#>

paves the way for the “Bring Your Own Generation” pathway under consideration at the PJM Board.⁶⁰

Implementing these policies at the state level has been shown to lead to immediate reduction in ghost load within a given state. In Ohio, for instance, AEP Ohio saw a 50% decrease in projected load upon implementing its large load tariff.⁶¹ In California, PG&E’s Rule 30 tariff application includes some of these reforms – such as minimum demand charges, early termination fees, and paying upfront some of the costs to interconnect.⁶² However, the PAO and others have recommended additional modifications to Rule 30 to ensure maximum ratepayer protections.⁶³

IV. Narrow Tolerances, Abrupt Drops.

On July 10, 2024, a device that protects electrical systems from damage during lightning strikes – a “lightning arrestor” – failed, resulting in a fault on a 230-kilovolt transmission line in the Eastern Interconnection. The transmission line was set up to automatically attempt restarting three times from each end of the line after a fault. Because both ends of the line were doing this, the line repeatedly shut off and restarted, causing six brief faults in just over a minute. The protection system worked as designed by detecting each fault and eventually (in 82 seconds) safely shut off the line. Nothing about the behavior of the transmission line’s protective equipment was unusual.

What was unusual was at the same time and near the same area, approximately 1.5 GWs of load dropped off the grid. As reported by the North American Electric Reliability Corporation’s (NERC) investigation of the event, none of the load was disconnected from the system by utility equipment; rather, the load was disconnected by the customer, specifically data center customers.⁶⁴ NERC concluded that a protective/control scheme on the data centers’ uninterruptible power supplies (UPS) – which are generally either power electronics that switch load to a battery bank or a flywheel that switches load to a diesel engine – was set to a particularly sensitive setting, causing over a gigawatt of load to drop off the grid and remain off for hours.

While grid operators had to act to stabilize the system following this load drop, the disturbance did not result in significant operational issues. However, NERC notes similar incidents have occurred in other interconnections with both cryptocurrency and oil and gas loads.

While CAISO, and many other grid operators, routinely plan for the sudden loss of large electric supply,⁶⁵ little planning has been done nationwide to manage the sudden loss of such significant load.⁶⁶ NERC recommends in its Incident Report that transmission operators enter into

⁶⁰ Joint Governor’s Joint Proposal for CIPF LLA, November 5, 2025; <https://www.pjm.com/-/media/DotCom/committees-groups/cipf-lla/2025/20251106/20251106-item-04h---joint-dcc-and-governors-proposal.pdf>

⁶¹ Zachary Jarrell, “AEP Ohio cuts its data center demand forecast in half,” *Biz Journal*, September 25, 2025. <https://www.bizjournals.com/columbus/news/2025/09/25/aep-ohio-data-center-demand-forecast.html>

⁶² A. 24-11-007; “Application of PG&E for Approval of Electric Rule No. 30 for Transmission-level Retail Electric Service,” filed November 21, 2024. <https://docs.cpuc.ca.gov/PublishedDocs/SupDoc/A2411007/7851/547330908.pdf>

⁶³ Akiya and McCormack,” Public Advocates Office Opening Brief,” A. 24-11-007, October 24, 2025; <https://docs.cpuc.ca.gov/PublishedDocs/Efile/G000/M584/K972/584972803.PDF>

⁶⁴ Pg. 2, NERC, “Incident Review: Considering Simultaneous Voltage-Sensitive Load Reductions;” January 8, 2025; https://www.nerc.com/globalassets/our-work/reports/event-reports/incident_review_large_load_loss.pdf

⁶⁵ For CAISO, this single contingency is the loss of the Diablo Canyon nuclear power plant. WECC rules require 6% contingency reserves for this loss. CAISO also requires electrical entities to maintain an additional 9% in reserves to account for other potential plant outages or higher-than-average peak demand, leading to a total 15% planning reserve margin.

⁶⁶ CAISO, “California ISO Planning Standards,” February 2, 2023; <https://www.caiso.com/Documents/ISO-Planning-Standards-Effective-Feb2023.pdf>

agreements with large load customers to ensure better coordination when reconnecting their systems. CAISO is already exploring various “ride-through” characteristics of various customer types and evaluating the potential impacts of on-site generation supporting these large loads on the transmission grid.⁶⁷

V. Sustainability Goals and Procurement Options.

Many hyperscalers have continued to uphold corporate sustainability commitments even as AI-related computing demand has surged, channeling investment not only into nuclear power but also emerging clean energy technologies and carbon capture.⁶⁸ Big tech firms seem to be single-handedly reviving the U.S. nuclear energy market, with nuclear project announcements ranging from Meta⁶⁹ to AWS⁷⁰ to Microsoft’s famous restart of Three Mile Island.⁷¹ As data center electricity demand grows, these commitments could help support further investment in underdeveloped carbon-free resources, though translating this potential into durable systemwide benefits is not guaranteed.

In 2024, Microsoft touted an agreement for 10.5 GWs of new renewable capacity in the U.S. and Europe, a scale Microsoft reported as “almost eight times larger than the largest corporate power purchase agreement (PPA) ever signed.”⁷² While this generation does not directly power its operations, Microsoft also reported an agreement with Powerex to match hourly demand at a Washington datacenter with direct deliveries of carbon-free power on a 24-hour basis throughout the year.⁷³

Microsoft is hardly alone in these sustainability achievements. Amazon has committed to matching all the electricity consumed by its operations with 100% renewable energy, a goal it achieved 7 years early in 2023. Amazon achieved this by investing billions in global solar and wind projects, as well as providing on-site solar directly on Amazon buildings.⁷⁴ While commendable, this investment does not mean Amazon’s operations are directly powered by renewable energy. Instead, an equivalent amount of renewable electricity is delivered to the grid and credited through accounting mechanisms.

Even with this approach, Amazon’s sustainability report shows that emissions from its direct operations increased by 7% in 2023, and its overall carbon footprint has grown by 34% since 2019.⁷⁵ Recent reporting from *S&P Global* noted “carbon-emitting sources have supplied 20%-

⁶⁷ Pgs. 4-5, CPUC, “California Data Center Development & Energy Needs: FAQs,” 12.22.2025.

⁶⁸ Spencer Kimball, “Exxon in advanced talks to power AI data centers with natural gas and carbon capture,” *CNBC*, October 31, 2025; <https://www.cnbc.com/2025/10/31/exxon-ai-data-center-natural-gas-carbon-capture.html?msocid=3f21a963703c6fbc1a35bfda71ff6eac>

⁶⁹ Meta newsroom, “Meta Announces Nuclear Energy Projects, Unlocking Up to 6.6 GW to Power American Leadership in AI Innovation,” January 9, 2026; <https://about.fb.com/news/2026/01/meta-nuclear-energy-projects-power-american-ai-leadership/>

⁷⁰ US EIA, “Data center owners turn to nuclear as potential electricity source,” *Today in Energy*, October 1, 2024; <https://www.eia.gov/todayinenergy/detail.php?id=63304>

⁷¹ Casey Crownhart, “Why Microsoft made a deal to help restart Three Mile Island,” *MIT Technology Review*, September 26, 2024; <https://www.technologyreview.com/2024/09/26/1104516/three-mile-island-microsoft/>

⁷² Bobby Hollis, “Accelerating the addition of carbon-free energy: An update on progress,” *Microsoft blog*, September 20, 2024; <https://www.microsoft.com/en-us/microsoft-cloud/blog/2024/09/20/accelerating-the-addition-of-carbon-free-energy-an-update-on-progress/>

⁷³ Powerex press release, “Powerex announces agreement to provide 24x7 hourly matching of carbon-free energy to Microsoft,” June 2023; <https://powerex.com/sites/default/files/2023-06/Powerex%20Announces%20Agreement%20with%20Microsoft%20for%202024x7%20Carbon-Free%20Energy.pdf>

⁷⁴ Amazon, “Amazon meets 100% renewable energy goal 7 years early;” *Amazon News*, last updated August 14, 2025; <https://www.aboutamazon.com/news/sustainability/amazon-renewable-energy-goal>

⁷⁵ 2023 *Amazon Sustainability Report*; <https://sustainability.aboutamazon.com/2023-amazon-sustainability-report.pdf>

35% of incremental power demand from Google and Microsoft.”⁷⁶ As a demonstration of this surge in carbon-emitting sources, current wait times for new gas-fired turbines can be as long as 7 years, with costs likewise escalating, due to high demand from data centers.⁷⁷

As the resource mix changes, hyperscalers are also shifting procurement strategies, moving from traditional PPAs to direct capacity investments, such as Google’s \$4.75 billion acquisition of Intersect.⁷⁸ Activity amongst the largest hyperscalers suggest a willingness to spend liberally, so long as the needed energy and capacity are readily available. There currently exist a wide spectrum of procurement options for data centers, ranging from fully off the grid to cost-of-service ratemaking under an existing – or updated – industrial tariff. A list of these various options was produced as part of a December 2025 report by E3 and is reproduced as Table 1 below.⁷⁹

Table 1. Sampling of Rate and Contract Design Options, with comparison of values. Key: tilde = partial utility service, check= full utility service; green = high, yellow = medium, and red = low score of the listed value.⁸⁰

Contracting Options	Definition	Utility service	Promotes Data center Growth	Risk Burden on Utility (U) or Developer (D)?	Protects Existing Customers	Shared Infrastructure Needed	Potential Ratepayer Benefits	Relative Ease of Implementation
Behind-the-Meter	Load co-located with generation, islanded/off-grid			D				
Wholesale Customer	Transmission -level customer with minimal utility service	~		D				
Physical PPA	Customer contracts with developer for physical power, utility provides T&D	✓		D				
Virtual PPA	Customer contracts for	✓		D				

⁷⁶ B. Brunettil, et al., “Hyperscaler procurement to shape US power investment,” *S&P Global*, December 19, 2025; <https://www.spglobal.com/sustainable1/en/insights/special-editorial/hyperscaler-procurement-to-shape-us-power-investment>

⁷⁷ Jared Anderson, “US gas-fired turbine wait times as much as seven years; costs up sharply;” *S&P Global*, May 20, 2025. <https://www.spglobal.com/energy/en/news-research/latest-news/electric-power/052025-us-gas-fired-turbine-wait-times-as-much-as-seven-years-costs-up-sharply>

⁷⁸ Alphabet Investor Relations, “Alphabet Announces Agreement to Acquire Intersect to Advance U.S. Energy Innovation,” December 22, 2025; <https://abc.xyz/investor/news/news-details/2025/Alphabet-Announces-Agreement-to-Acquire-Intersect-to-Advance-U-S-Energy-Innovation-2025-DVIuVDM9wW/default.aspx>

⁷⁹ E3, *Tailored for Scale: Designing Electric Rates and Tariffs for Large Loads*, December 2025; <https://www.ethree.com/wp-content/uploads/2025/12/RatepayerStudy.pdf>

⁸⁰ Figures 7 and 8 (pg 17-18) in E3 *Tailored*, *Ibid*.

	renewable energy credits with no physical delivery of power						
Semi-Islanded/Non-firm	Interruptible service, supplemented with on-site generation	~		D			
Rolled-in Incremental	New infrastructure costs are rolled into existing rate structure and thus spread across all customers.	✓		U			

As mentioned above, data center developers have been trying out these various procurement arrangements throughout the U.S. In California, activity has largely centered on individual project-by-project review for energization cost allocation, and traditional C&I tariffs for rate design. Other options listed in Table 1 are limited in California, such as the physical PPA design, which is subject to a statutory kWh limit and rarely open to new entrants.⁸¹

While data centers currently favor energization speed and build out, eventually the industry is likely to prioritize cost-effective energy arrangements to ensure continued growth; in which case alternate rate and contract designs will likely be explored more. For instance, under the right market conditions, one can envision data centers under existing IOUs energization agreements, exiting those agreements and turning to large amounts of on-site generation. The consequences of having large, heavily resourced firms seeking such alternate arrangements could be significant.

When the Backup Becomes the Main Act. It is also unclear, as hyperscalers are developing more fossil resources nationwide, whether that trend will also emerge in California. The CEC has exclusive authority to certify thermal power plants of 50 MW or more in California. The last natural gas plant the CEC permitted was in 2017,⁸² however, they have consistently authorized (via an exemption) diesel backup generating facilities co-located with data centers. Under the Small Power Plant Exemption (SPPE) process, applicants proposing facilities between 50 and 100 MW may seek an exemption from CEC certification. The CEC may grant an SPPE if it determines that the project would not result in significant adverse impacts to the environment or energy resources. If an exemption is granted, local land use authorities and other permitting agencies, such as the local air district, assume jurisdiction and act as responsible agencies under the California Environmental Quality Act (CEQA), conducting any additional environmental review required for their approvals.

Since the establishment of the SPPE process in the 1980s, the CEC has provided a total of 36 SPPEs, including 17 for data centers' backup power generation. All the SPPE applications filed since 2011 have been for backup generating facilities serving data centers, and all but one of the

⁸¹ Public Utilities Code §365.1

⁸² Huntington Beach Energy Project; <https://www.energy.ca.gov/powerplant/combined-cycle/huntington-beach-energy-project>

17 data center exemptions were provided for diesel generator systems.⁸³ Fuel costs, fuel availability, and energy capacity shape the selection of diesel generators over lower and zero-emitting backup power systems.

The CEC's power plant siting authority was designed to provide a transparent and predictable certification process, including meaningful opportunities for public participation and consideration of alternatives under CEQA. By contrast, projects approved through the SPPE offer limited formal opportunities for stakeholder engagement, and the program was not designed 40 years ago with today's large, data center–driven power systems in mind. Moreover, data centers in the state have concentrated within small, densely urban geographic areas, primarily in the Silicon Valley; as a result, these SPPEs may disproportionately impact certain communities near large population centers, increasing local air pollution concerns. *Given the scale and pace of data center expansion, the joint committee may wish to consider whether the SPPE remains the right pathway for authorizing these facilities.*

Moreover, as data center developers begin to embrace flexible demand in exchange for faster connectivity⁸⁴ – such as what is proposed in the recent FERC Advance Notice of Proposed Rulemaking (ANOPR)⁸⁵ – the potential for increased diesel usage could rise, adversely affecting local communities. Data centers have traditionally been considered inflexible loads, operating 24/7 to provide constant service such as cloud storage. However, certain AI workloads, such as AI training and machine learning are less time-sensitive than traditional data center workloads and can tolerate brief interruptions.⁸⁶ The promise of flexible data center demand is that users capable of reducing their grid usage, especially during periods of peak energy usage, could significantly reduce resource and infrastructure needs. However, just because a data center goes off the electric grid doesn't mean it goes dark. Depending on how rules for data center demand flexibility are written, data centers could still be considered flexible by switching their operations to their backup diesel units, greatly increasing their emissions. *The joint committee may wish to ask panelists what protections could be developed to ensure encouraging flexible demand at data centers does not inadvertently lead to overuse of onsite diesel backup generators.*

VI. Cost.

“Every houseguest brings you happiness – some when they arrive, and some when they are leaving.”⁸⁷

Imagine roommates sharing an apartment who split rent and utilities evenly. When a new roommate moves in, if the new roommate pays at least their share of the added costs – utilities, wear and tear, groceries – the original roommates aren't paying extra. If the new roommate pays less, the existing roommates subsidize them. If the new roommate covers not only their own costs but also contributes to fixed expenses like rent, everyone's share goes down. Fairness

⁸³ <https://www.energy.ca.gov/programs-and-topics/topics/power-plants/power-plant-compliance-and-siting>

⁸⁴ 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/sites/default/files/publications/rethinking-load-growth.pdf>

⁸⁵ Chris Wright, Letter to FERC to Initiate Rulemaking “...Regarding the Interconnection of Large Loads Pursuant to the Secretary's Authority Under Section 403...”, October 23, 2025. <https://www.energy.gov/sites/default/files/2025-10/403%20Large%20Loads%20Letter.pdf>

⁸⁶ Y. Numata, et al., “Fast, Flexible Solutions for Data Centers,” *RMI*, July 17, 2025; <https://rmi.org/fast-flexible-solutions-for-data-centers/>

⁸⁷ Aphorism widely attributed to Confucius

suggests a new roommate should at least cover what they add and ideally contribute to existing costs.

The emergence of extensive data center electricity load, driven largely by the growth in artificial intelligence computing, marks a rare opportunity in the energy industry to reassess this contract amongst roommates. A large, well-financed new roommate wants to move in. Ensuring new large loads provide sufficient revenue to meet or exceed the marginal cost of service⁸⁸ should protect existing utility customers.

It is the responsibility of the utility to provide electrical service to customers within their territory; known as their “obligation to serve.” Traditionally, the costs of providing that service, including potential grid upgrades, are socialized among all customers. However, existing rules for customers seeking to energize at the distribution-level⁸⁹ include provisions that reduce risks of stranded costs for ratepayers and require some customers to pay for part of their own energization.⁹⁰

Electric rate design generally rests on the principle that other ratepayers should not pay for upgrades from which they are not beneficiaries. Just how much an existing utility customer may or may not benefit from emergent data center energization is a topic of ongoing debate. Cost to energize new data centers are usually broken down by 1) the cost to connect them to the grid (their “energization” cost); and 2) the cost of continuing to provide them service (their rate schedule). For California’s IOUs, the traditional arrangement involves individual project-by-project review of the energization cost,⁹¹ and traditional C&I rate tariffs.

Electric Rule 30. In November 21, 2024, PG&E filed an application at the CPUC to establish a streamlined approach for energizing new transmission-level retail customers.⁹² While the application is broad in its definition of “large load customers” that would apply for service under this tariff, the PAO notes 75% of PG&E customer requests to interconnect at the transmission-level in the last two years are from data centers. That represents about 30 data center requests in 2 years, compared to 24 total customer requests for transmission-level service in the prior 8 years (2014-2022).⁹³ The Electric Rule 30 proposal is solely focused on the energization costs, not a unique large load rate.

Under the proposed Rule 30:

1. Large load customers are required to pay upfront the cost of new interconnection facilities, or the customer can build the facilities herself and transfer ownership to PG&E, if more cost-effective. Costs for upgrades on the existing transmission system would be paid by ratepayers. PG&E states this is due to these facility upgrades benefitting multiple customers as well as advancing California policy initiatives such as electrification.

⁸⁸ “Marginal cost of service” generally means the cost to the utility to serve an additional unit of electricity demand at a given time. These are the costs in power markets or the costs to build new infrastructure in serving new load.

⁸⁹ CPUC Electric Rules 15 and 16

⁹⁰ Essentially, customers are given a set “allowance” which is paid for by all ratepayers. Any costs in excess of the allowance are covered by the individual seeking to energize. The allowances are formula amounts set by the CPUC.

⁹¹ Via “Exceptional Case Filings” such as the recent approval of a 90MW Microsoft data center.

<https://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M590/K889/590889612.PDF>

⁹² A. 24-11-007

⁹³ Pg. 2, “PAO Opening Brief,” A. 24-11-007, filed October 24, 2025.

2. Large load customers are required to pay the actual costs of interconnection infrastructure, rather than the *estimated* costs, which are traditionally used for residential and commercial customer energization.
3. After the large load customer has received service and is providing revenues, the customer is eligible for a refund of their upfront payments over time. Refunds are based on the revenues produced by the large load customer. Thus, if the customer's load does not materialize, the customer does not receive back its entire upfront payment. A large load customer only receives refunds when it is producing revenues to offset the interconnection infrastructure costs. However, parties have raised that the timing of that repayment should be matched with the timing of the new revenue generated from the customer, such that annual refunds never exceed the net customer revenue.
4. Includes a minimum demand charge, minimum contract term, and early termination fee.
5. Information sharing provisions to facilitate community choice aggregators (CCAs) and PG&E having early notice of potential large load customers and being able to make cost-effective generation resource decisions.

While PG&E notes in its application that these provisions will protect existing ratepayers and ensure the revenues from large load customers result in bill reductions, others within the proceeding disagreed and proposed a number of modifications.⁹⁴ In July 2025, the CPUC adopted interim implementation of PG&E's Rule 30, requiring new transmission-level customers to be responsible for all initial costs, deferring decisions on rate recovery and cost allocation until a final decision, and denying PG&E's request for a memorandum account to record accrued interest, among other requirements.⁹⁵

Do Large Loads Mean Lower Rates? The promise of data center energy consumption is alluring: that increased electricity sales from data centers will cover all the new costs to serve those data centers and may even offset existing system costs. Such is the story for other large loads, like transportation electrification (TE), with the California Public Advocates Office (PAO) finding TE "may cause downward pressure on electric rates."⁹⁶ The timing can feel almost too perfect: as California undertakes systemwide decarbonization and infrastructure renewal, a new, well-funded customer emerges, apparently willing⁹⁷ to shoulder the cost of long-deferred investments. However, as the PAO cautions for TE, and as applies equally to data-center development, "to achieve this downward pressure on rates, effective management of multiple factors will be required, including efficient infrastructure buildout and cost constraints."⁹⁸

Rate design is an essential mitigation tool to ensure costs are not shifted onto other ratepayers. The treatment of large loads under retail ratemaking presents a challenge, as the pace of retail design has traditionally been aligned with smaller, incremental load growth.

⁹⁴ See "PAO Opening Brief," A. 24-11-007, filed October 24, 2025.

⁹⁵ D. 25-07-039. CPUC "Decision Partly Granting and Partly Denying Pacific Gas and Electric Company's Motion for Interim Implementation of Electric Rule Number 30," A. 24-11-007, July 24, 2025;

<https://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M574/K875/574875643.PDF>

⁹⁶ Pg. 14, Public Advocates Office, *Distribution Grid Electrification Model 2025 – Study and Report*, October 2025. <https://www.publicadvocates.cpuc.ca.gov/-/media/cal-advocates-website/files/press-room/reports-and-analyses/251030-public-advocates-office-distribution-grid-electrification-model-2025.pdf>

⁹⁷ Amrith Ramkumar, "Microsoft Makes New Data-Center Pledges After Local Backlash," *The Wall Street Journal*, January 13, 2025.

⁹⁸ Pg. 14, PAO October 2025, *Ibid*.

PG&E has provided “extensive analysis” in its Electric Rule 30 application showing that the substantial revenues from large load customers will result in bill reductions for existing ratepayers because fixed costs will be spread over a larger base of customers.⁹⁹ While the simplicity of the math is undeniable – larger denominator, smaller overall number – it is unclear if their analysis truly considers all known costs. PG&E’s analyses do not appear to be part of the public docket by which the committee might review.

Transmission-level customers, such as large data centers, typically pay rates about 14-17¢ less than distribution-level customers.¹⁰⁰ Distribution rates also contain statewide policy costs, such as wildfire mitigation, net energy metering, and public purpose program costs. Customers connecting at the transmission-level do pay certain distribution costs, including a portion of wildfire-related costs,¹⁰¹ but not nearly as significantly as distribution-level customers.

The primary way large loads reduce bills for other ratepayers is by (a) minimizing or avoiding infrastructure upgrades that would otherwise be socialized – either by limiting the need for new infrastructure or requiring the customer to pay those costs upfront – and/or (b) ensuring that large-load customers provide stable, long-term revenue over the life of any required upgrades, sufficient to cover not only fixed capital costs but also ongoing operations, maintenance, depreciation, and associated generation procurement costs.

The joint committee may wish to ask panelists how to structure tariffs to balance encouraging economic development while guarding against higher energy costs to existing ratepayers, stranded assets, and negative system impacts.

Conclusions. The ongoing surge of investment into AI data centers has been suggested to be bigger than the buildup of the interstate highway or the dot-com boom. AI-driven data center growth presents California with both a stress and a strategic choice. The scale, speed, and uncertainty of this new load are unlike prior waves of electrification, with the potential to challenge long-standing assumptions embedded in forecasting, planning, cost allocation, and rate design. California’s relative insulation to date reflects deliberate policy choices, higher baseline prices, and robust planning institutions; but those advantages are not guaranteed if safeguards do not evolve as rapidly as the load itself. The core question is not whether data centers should be served, but under what terms: who pays, when, and with what risks. If large loads are required to bear the full costs they impose, provide durable and verifiable revenue, and align procurement and operational practices with reliability and climate goals, they could support system investment without burdening existing customers. Absent those protections, the promise of “downward pressure on rates” risks becoming another case of privatized benefits and socialized costs. The task before policymakers is to ensure that California welcomes innovation without repeating the mistakes now playing out elsewhere; and that, when a powerful new roommate moves in, the household is better off, not worse.

#####

⁹⁹ Michael Medieros, “PG&E Written Testimony,” Letter to the Little Hoover Commission; December 8, 2025. <https://lhc.ca.gov/wp-content/uploads/PGE-LHC-Written-Testimony.pdf>

¹⁰⁰ Pg. 26, “Little Hoover Commission Study to Review Data Centers and California Energy Policy, Pt. II,” December 11, 2025. <https://lhc.ca.gov/wp-content/uploads/Revised-Commissioner-Information-12-11-25.pdf>

¹⁰¹ Medieros, *Ibid.*