


JULY 2026

AMERICANS *for*
RESPONSIBLE
INNOVATION

INFORMATION IS POWER

The Case for Data Center Transparency



WHITEPAPER BY
JESSICA MAKSIMOV

Executive Summary

Load growth from computational demand is outpacing the ability of grid planners, energy analysts, and policymakers to track, model, and manage its consequences. The Lawrence Berkeley National Laboratory (LBNL) projects that data center electricity consumption will rise from 4.4 percent to 6.7–12 percent of total U.S. electricity demand by 2028. On May 4, 2026, the North American Electric Reliability Corporation (NERC) issued its rarest grid alert level, citing data center load characteristics as a direct reliability threat. Yet the data collection infrastructure needed to respond to this growth remains nascent and insufficient to support U.S. energy policy.

This information gap has cascading consequences. Without granular, facility-level consumption data, energy analysts cannot accurately model load growth trajectories, and grid operators cannot characterize load behavior, run accurate simulations, or plan power delivery with confidence. Without certainty, grid operators are forced into conservative, costly planning assumptions, which can lead to inefficient resource dispatch, higher consumer costs from strained capacity markets, and degraded reliability over time. In addition, the heterogeneity of available data has also forced researchers to reconstruct consumption estimates with insufficient primary data. For example, studies estimating energy use per large language model (LLM) query have yielded figures ranging from 0.24 to three watt-hours depending on the methodology used.

To close this gap, Congress should direct the Energy Information Administration (EIA) to establish a standardized, mandatory reporting framework for data center electricity consumption. The framework must require disaggregated disclosures spanning the full facility consumption stack, including IT equipment; cooling systems; power delivery infrastructure; and AI workload types, differentiated between training and inference. In parallel, the Federal Energy Regulatory Commission (FERC) and NERC should subject large loads to NERC's Reliability Standards. This registration would compel developers to submit the simulation data, protection settings, and real-time consumption profiles that grid operators lack, while expanding real-time load measurements to validate facility models, close investigative gaps after grid disruptions, and align data center grid interconnection with long-term reliability planning.

Introduction

In 1973, the Organization of Arab Petroleum Exporting Countries (OAPEC) placed an oil embargo on the United States, triggering shortages across the country. Long gas lines formed.¹ President Nixon called for lowered thermostats and reduced driving speeds.² And the U.S. government found itself without credible energy statistics on which to base policy. In 1974, the General Accounting Office (now the Government Accountability Office) testified to Congress that the crisis “ma[de] us aware of the need for the best possible information upon which to base both public and private decisions.”³ Another witness remarked, “even now, there is no one, in or out of Government, who can give a credible, objective analysis of the extent or location of the real crisis.”⁴

The federal government’s reliance on voluntary data reporting by the energy industry sparked public mistrust. In a 1974 Senate committee hearing, Senator Glenn Beall noted, “Our citizens are not only questioning the extent of the shortage, but even asking: is there really a shortage of oil?”⁵ Another witness testified, “The energy problem is bad enough, but it is being made worse by a real crisis of confidence...People are afraid of being taken for suckers because they can’t seem to get a straight story. The government is woefully ignorant of basic energy facts.”⁶

In response, Congress created the Federal Energy Administration (FEA) to collect, analyze, and publicly disseminate energy information from energy-producing and consuming firms.⁷ In 1976, further solidifying federal information collection authorities, Congress established the Office of Energy Information and Analysis (now the Energy Information Administration (EIA)) to develop expertise in energy analysis and forecasting, provide congressional committees with energy information, and create reports on the U.S. energy situation and trends.⁸

Back then, the United States faced supply restrictions. Now, it faces heightened electricity demand and grid instability.⁹ The Lawrence Berkeley National Lab (LBNL) projects that data center loads will rise from 4.4 percent to 6.7–12 percent of total U.S. electricity consumption by 2028.¹⁰ On May 4, 2026, the North American Electric Reliability Corporation (NERC)—the bulk power system’s (BPS) grid reliability watchdog—issued a Level 3 Alert (its highest level), warning grid entities that data centers’ load type and level of demand are creating reliability concerns.¹¹ The very first “essential action” that NERC urges grid owners,

¹ Mark Holt et al., “The U.S. Energy Information Administration,” Congressional Research Service, September 9, 2020, <https://www.congress.gov/crs-product/R46524>.

² Roger Anders, “The Federal Energy Administration,” *U.S. Department of Energy*, November 1980, <https://www.energy.gov/management/articles/federal-energy-administration>.

³ Holt et al., “The U.S. Energy Information Administration.”

⁴ Ibid.

⁵ Ibid.

⁶ Ibid.

⁷ “Legislative Timeline,” Eia.Gov, U.S. Energy Information Administration, accessed May 12, 2026, https://www.eia.gov/about/legislative_timeline.php.

⁸ Ibid.

⁹ Arman Shehabi et al., *United States Data Center Energy Usage Report* (2024), <https://doi.org/10.2172/1372902>.

¹⁰ In the LBNL study, the load projection defines “data centers” broadly, encompassing telco and commercial edge, small and medium business, enterprise branch, internal, communications service provider, colocation, and hyperscale facilities. Shehabi et al., *United States Data Center Energy Usage Report*; Eric Masanet et al., “To Better Understand AI’s Growing Energy Use, Analysts Need a Data Revolution,” *Joule* 8, no. 9 (2024): 2427–36, <https://doi.org/10.1016/j.joule.2024.07.018>.

¹¹ Robert Walton, “NERC Issues Level 3 Alert, Mandates Action to Address Data Center Load Losses,” *Utility Dive*, May 5, 2026, <https://www.utilitydive.com/news/nerc-issues-rare-level-3-alert-over-data-center-load-losses/819295/>.

operators, and users to take is to collect “detailed” information on computational loads and ensure transmission owners reflect that information in their facility interconnection requirements.¹²

Data center interconnection requests are now growing from hundreds to thousands of megawatts.¹³ Assessing the full impact of this growth on the electric grid requires granular, facility-level data that analysts and planners largely do not have. Energy analysts need disaggregated consumption data to understand growth trajectories and computational demand characteristics. Grid operators and planners need precise models to characterize load behavior, evaluate risk, and safeguard reliability.¹⁴ Without this information, neither group can do its job effectively, and policymakers are left to make consequential decisions without the knowledge they need.¹⁵

Breaking It Down: How Data Centers Consume Energy

To accurately forecast and model data center load growth, profiles, and grid impacts, energy analysts need disaggregated information on the entire AI electricity-consuming stack. Factors that influence data centers’ electricity consumption include IT equipment characteristics, overhead energy use, AI workload behavior and efficiency, and data center facility size.

First, IT equipment comprises up to 60-95 percent of a data center’s total energy demand.¹⁶ This category consists of (1) the specialized chips that run AI workloads; (2) the general-purpose processors and memory chips that store and coordinate data as servers execute tasks;¹⁷ (3) the networking equipment that connects data center clusters; and (4) idle machine capacity, including servers provisioned to ensure low latency and service uptime and those consuming less power as they transition between tasks.¹⁸

Second, overhead energy is consumed by the data center’s supporting infrastructure. This equipment includes (1) the cooling systems that remove excess heat through pumps, fans, and water-cooled chillers; (2) the power delivery systems, which control and distribute electricity within the facility, such as emergency backups and uninterruptible power supplies (UPS); and (3) miscellaneous components, such as the lighting and security systems used across the data center.¹⁹

Third, AI data centers run two workload types: training and inference. Training develops and finetunes AI models, while inference uses those models to generate outputs in response to user prompts.²⁰ These

¹² Engineering Group, “Essential Action to Industry Computational Load Modeling, Studies, Instrumentation, Commissioning, Operations, Protection, and Control,” North American Electric Reliability Corporation, May 4, 2026, <https://www.nerc.com/globalassets/programs/bpsa/alerts/level-3-computational-load-alert.pdf>.

¹³ Jordan Aljbour et al., “Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption,” Electric Power Research Institute, May 28, 2024, <https://www.epri.com/research/products/3002028905>.

¹⁴ Large Loads Task Force, *Characteristics and Risks of Emerging Large Loads*, White Paper (North American Electric Reliability Corporation, 2025), <https://www.nerc.com/globalassets/who-we-are/standing-committees/rstc/whitepaper-characteristics-and-risks-of-emerging-large-loads.pdf>.

¹⁵ Masanet et al., “To Better Understand AI’s Growing Energy Use, Analysts Need a Data Revolution.”

¹⁶ Large Loads Task Force, *Characteristics and Risks of Emerging Large Loads*.

¹⁷ Cooper Elsworth et al., “Measuring the Environmental Impact of Delivering AI at Google Scale,” arXiv:2508.15734, preprint, arXiv, August 21, 2025, <https://doi.org/10.48550/arXiv.2508.15734>.

¹⁸ Elsworth et al., “Measuring the Environmental Impact of Delivering AI at Google Scale.”

¹⁹ Overhead energy is captured by the PUE metric. See Elsworth et al., “Measuring the Environmental Impact of Delivering AI at Google Scale”; Large Loads Task Force, *Characteristics and Risks of Emerging Large Loads*.

²⁰ Casey Crownhart, “AI Is Changing the Grid. Could It Help More than It Harms?,” *MIT Technology Review*, September 9, 2025, <https://www.technologyreview.com/2025/09/09/1123404/ai-grid-help/>.

workloads have distinct electrical load profiles. Training produces rapid demand fluctuations; during these runs, demand can spike or drop in under one second as the system alternates between training and progress saving.²¹ Inference, by contrast, supports interactive computations at lower, more stable power levels, though demand still fluctuates with user activity, rising during peak hours and falling during periods of lower use.²²

Since AI training and inference place fundamentally different demands on IT equipment and the grid, disaggregating consumption data by workload type is essential for accurate modeling. Tracking how much of a company's total energy supports training versus inference allows researchers to design targeted efficiency measures, such as tuning graphics processing units (GPUs) power limits during training runs to remove waste.²³ Disaggregated reporting by workload type (and the resulting ramp rates) would also inform grid planners about distinct AI load profiles, allowing them to accurately model AI data center behavior. Based on this information, operators can then procure appropriate reserves to balance the grid.²⁴ Without sufficient information, if operators underschedule power delivery, there may not be enough power to meet a sudden AI surge, risking blackouts. If they overschedule, they may waste money and fuel on power plants, and increase strain on capacity markets.²⁵

Fourth, a data center's physical size, including square footage and rack count,²⁶ can influence overall consumption. Facilities can range from closet-sized rooms to multi-acre hyperscale centers, influencing their overall demand on the grid.²⁷

Disaggregated disclosures across these components would give researchers the granular data needed to track efficiency gains, build more accurate consumption models, and identify opportunities for energy optimization. For grid planners and policymakers, this information is equally consequential; without it, they cannot reliably characterize how data centers behave on the grid, anticipate load growth, or make informed decisions about infrastructure investments and reliability planning. Ongoing, standardized data collection across the full electricity-consuming stack is, therefore, a prerequisite for sound energy policy.²⁸

²¹ Power spikes can occur when the workload is running due to alternating compute- and network-intensive phases.

²² McKinsey & Company projects that inference may constitute over 50 percent of all AI compute by 2030, but there is limited company data to make accurate assessments. Chhavi Arora et al., "The next Big Shifts in AI Workloads and Hyperscaler Strategies," McKinsey & Company, December 17, 2025, <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-next-big-shifts-in-ai-workloads-and-hyperscaler-strategies>.

²³ Nora Efram and Camron Assadi, *Future-Proof AI Data Centers, Grid Reliability, and Affordable Energy: Recommendations for States* (American Council for an Energy-Efficient Economy (ACEEE), 2025), <https://www.aceee.org/white-paper/2025/04/future-proof-ai-data-centers-grid-reliability-and-affordable-energy>.

²⁴ Shehabi et al., *United States Data Center Energy Usage Report*; Large Loads Task Force, *Characteristics and Risks of Emerging Large Loads*; Engineering Group, "Essential Action to Industry Computational Load Modeling, Studies, Instrumentation, Commissioning, Operations, Protection, and Control."

²⁵ North American Electric Reliability Corporation, "Resource Adequacy Risks Intensify Across North America as Demand Growth Surges," *Newsroom*, January 29, 2026, https://www.nerc.com/globalassets/our-work/assessments/nerc_ltra_2025.pdf; Rachel Reed, "How Data Centers May Lead to Higher Electricity Bills," *Harvard Law School*, September 3, 2025, <https://hls.harvard.edu/today/how-data-centers-may-lead-to-higher-electricity-bills/>.

²⁶ Rack count refers to the total number of standardized metal frames (racks) designed to hold rows of flat servers, storage drives, and network switches stacked on top of each other.

²⁷ Kyle Mason et al., "The Rise of Data Centers in the Grid," Regional Plan Association, June 30, 2025, <https://rpa.org/news/lab/the-rise-of-data-centers>; Shehabi et al., *United States Data Center Energy Usage Report*.

²⁸ Shehabi et al., *United States Data Center Energy Usage Report*.

The Measurement Gap: What Is(n't) Reported

Even as the need for information grows, data center operators are not subject to standardized metrics or reporting of their (1) electricity consumption, (2) load behavior, and (3) efficiency gains. This leaves researchers, grid entities, and policymakers with insufficient data for making informed demand projections and maintaining grid stability.²⁹

First, during the load interconnection process, data center operators report peak demand (the maximum power required at a single interconnection point), which allows utilities to build enough infrastructure to service that highest capacity point. However, this statistic does not capture data center load behavior, which impacts how these facilities use the grid.³⁰ Moreover, most hyperscalers publicly report electricity consumption as an organization-wide aggregate, meaning a company such as Amazon could lump its data center electricity consumption with its grocery stores and warehouses into one sum. A 2024 report found that of the thirteen examined major AI data center operators, only one (Google) publicly revealed its AI hardware electricity use, and only two (Meta and Apple) regularly reported the electricity use of their data center operations separately from all other company operations.³¹

Second, data center operators do not disaggregate reported information by function type, inhibiting electricity consumption modeling.³² A data center's electricity use varies with the functions it serves. Traditional small- to mid-size facilities supporting online services such as social media, banking, and healthcare share similar electrical characteristics with limited consumption variability.³³ AI data centers, however, are distinct: their servers run on high-power GPUs and tensor processing units (TPUs), driving greater power demands and, in turn, more intensive cooling needs.³⁴ Because cooling loads vary with the intensity of processing tasks, AI facilities' cooling energy profiles differ from those of traditional data centers.³⁵

²⁹ Mitul Jhaveri and Vijaykumar Palat, "Measuring AI's Energy/Environmental Footprint to Access Impacts," *Federation of American Scientists*, June 27, 2025, <https://fas.org/publication/measuring-and-standardizing-ais-energy-footprint/>; David Berreby, "As Use of A.I. Soars, So Does the Energy and Water It Requires," *Yale E360*, February 6, 2024, <https://e360.yale.edu/features/artificial-intelligence-climate-energy-emissions>.

³⁰ Christopher Tozzi, "Beyond Megawatts: Rethinking How We Measure Data Center Capacity," *DataCenterKnowledge*, accessed May 20, 2026, <https://www.datacenterknowledge.com/energy-power-supply/beyond-megawatts-rethinking-how-we-measure-data-center-capacity>.

³¹ Masanet et al., "To Better Understand AI's Growing Energy Use, Analysts Need a Data Revolution."

³² A load profile can be defined as the hourly or sub-hourly pattern of electricity use, typically represented as a graph of electrical load (power or energy consumption) versus time, showing how a facility's electricity demand varies over a specified period. See "Load Profile - An Overview," *ScienceDirect*, accessed May 20, 2026, <https://www.sciencedirect.com/topics/engineering/load-profile>; Monica Neukomm, "End-Use Load Profiles for the U.S. Building Stock," *Energy.gov*, December 4, 2018, <https://www.energy.gov/cmei/buildings/articles/end-use-load-profiles-us-building-stock>.

³³ Large Loads Task Force, *Characteristics and Risks of Emerging Large Loads*.

³⁴ SolarTech, "How Much Electricity Does a Data Center Use? Complete 2025 Analysis," *IAEI Magazine*, January 1, 2026, <https://iaeimagazine.org/electrical-fundamentals/how-much-electricity-does-a-data-center-use-complete-2025-analysis/>.

³⁵ Masanet et al., "To Better Understand AI's Growing Energy Use, Analysts Need a Data Revolution"; Walton, "NERC Issues Level 3 Alert, Mandates Action to Address Data Center Load Losses"; Efram and Assadi, *Future-Proof AI Data Centers, Grid Reliability, and Affordable Energy*.

Third, the scale of electricity demand growth from data centers, electrification, and reindustrialization is driving the need for efficiency gains to reduce strain on an aging and increasingly vulnerable U.S. grid.³⁶ Most data center facilities track their energy use efficiency with the power usage effectiveness (PUE) metric, which measures total facility power divided by IT equipment power.³⁷ However, PUE only reveals how efficient the cooling and power systems are, not how efficiently the IT equipment itself uses that power. Therefore, metrics such as IT usage effectiveness (ITUE) and IT work capacity (ITWC)³⁸ should be used to capture IT equipment efficiency, but they have seen little industry uptake and do not yet account for newer, accelerated hardware types such as GPUs and TPUs.³⁹

Why Reporting Matters

Grid Reliability

There is already evidence of data centers impacting BPS reliability. The Eastern Interconnection and the Electric Reliability Council of Texas (ERCOT) have both observed load-reduction events of approximately 1,500 megawatts (MW) as data centers transferred load to backup generation.⁴⁰ In Virginia, voltage fluctuations triggered a safety mechanism (used to protect electronic equipment), causing 60 data centers to trip offline and switch to on-site generators. This drop in load caused a surge in excess electricity and forced operators and the utility to urgently scale back power plants to prevent wider electricity outages.⁴¹

System operators and planners need consumption data and energy models to (1) characterize data center load behaviors and (2) study the risks they may pose to the bulk power system. However, large load (e.g., data center) owners and operators are not subject to NERC's Reliability Standards and, therefore, not required to disclose detailed, standardized load information to system operators.⁴² This information gap can lead to grid reliability issues.

Before a new data center load is interconnected, power system engineers build grid models and run simulations to assess how the new facility will affect voltages, currents, and power flows in the existing electrical infrastructure. These simulations allow planners and operators to understand new risks and design protections before loads connect to the real grid. Accurate simulations, however, require data on peak demand, load behavior, protection and control settings, and dynamic load models.⁴³

³⁶ Jonathan Koomey et al., *Electricity Demand Growth and Data Centers: A Guide for the Perplexed* (Bipartisan Policy Center, 2025), <https://bipartisanpolicy.org/report/electricity-demand-growth-and-data-centers/>; Morgan Butterfield and Tim Hess, "EIA Press Release: EIA Forecasts Strongest Four-Year Growth in U.S. Electricity Demand since 2000, Fueled by Data Centers," Washington, DC, January 13, 2026, <https://www.eia.gov/pressroom/releases/press582.php>; Aljbour et al., "Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption."

³⁷ Jhaveri and Palat, "Measuring AI's Energy/Environmental Footprint to Access Impacts."

³⁸ Steven Nadel, *Opportunities to Use Energy Efficiency and Demand Flexibility to Reduce Data Center Energy Use and Peak Demand*. (American Council for an Energy-Efficient Economy (ACEEE), 2025), <https://www.aceee.org/white-paper/2025/10/opportunities-use-energy-efficiency-and-demand-flexibility-reduce-data-center>.

³⁹ Ibid.

⁴⁰ Large Loads Task Force, *Characteristics and Risks of Emerging Large Loads*.

⁴¹ Tim McLaughlin, "Big Tech's Data Center Boom Poses New Risk to US Grid Operators," *Technology*, *Reuters*, March 19, 2025, <https://www.reuters.com/technology/big-techs-data-center-boom-poses-new-risk-us-grid-operators-2025-03-19/>.

⁴² Large Loads Task Force, *Characteristics and Risks of Emerging Large Loads*.

⁴³ Ibid.

In addition, the lack of reported data restricts system operators' ability to analyze events retroactively. Because most large loads lack real-time monitoring infrastructure, operators cannot validate their models against actual load profiles. As a result, notable grid events, such as sudden demand drops, can only be analyzed using recordings taken near these loads, not at them.⁴⁴

Finally, large loads such as data centers can shift demand across physical facilities in response to rising energy prices or emissions limits. Without visibility into the causes or timing of these fluctuations, utilities cannot accurately forecast demand, and operators cannot efficiently respond to real-time changes. Consequently, grid operators are left relying on excess power reserves to manage swings and maintain supply-demand balance, raising the risk of inefficiency and broader system strain.⁴⁵

Modeling Electricity Use

As loads come online, a lack of standardized disclosures means that researchers must reconstruct data center consumption using limited available data to understand and predict data center load growth trajectories and behaviors.

Researchers use bottom-up, top-down, and extrapolation models to estimate data centers' electricity use.⁴⁶ First, bottom-up modeling combines estimates of individual components' electricity use (e.g., servers' power draw and the data centers' overall efficiency) with the projected number of installed servers to calculate total consumption.⁴⁷ Though bottom-up models are the most granular approach, their predictions rely on the accuracy of the installed IT equipment base, sales, and shipment data.

Second, top-down modeling uses electricity consumption totals from national and regional surveys and statistics to interpolate the energy required by the facilities' physical infrastructure. This model's accuracy depends on data precision and availability. Because current reporting rarely breaks down data center components, analysts struggle to gauge how specific technology changes affect electricity demand.⁴⁸

Third, extrapolation modeling builds on existing bottom-up and top-down estimates to forecast future electricity use. This method also employs assumptions of overall market growth and energy efficiency improvements. Because extrapolation compounds input errors, the accuracy of underlying consumption data is paramount to producing reliable demand projections.⁴⁹

Improving these analyses depends on facility-level data that only data center operators possess. Key inputs—including AI hardware utilization, installed compute capacities, average system configurations, advanced PUE metrics, and AI workload types—would sharpen grid planners' and policymakers' understanding of projected load growth and behavior.⁵⁰ Yet as LBNL researchers note, few companies report comprehensive data center electricity use, as they consider these details proprietary.⁵¹

⁴⁴ Ibid.

⁴⁵ Ibid.

⁴⁶ Shehabi et al., *United States Data Center Energy Usage Report*.

⁴⁷ Masanet et al., "To Better Understand AI's Growing Energy Use, Analysts Need a Data Revolution."

⁴⁸ Shehabi et al., *United States Data Center Energy Usage Report*.

⁴⁹ Ibid.

⁵⁰ Masanet et al., "To Better Understand AI's Growing Energy Use, Analysts Need a Data Revolution."

⁵¹ Shehabi et al., *United States Data Center Energy Usage Report*.

Efficiency Proxies

Researchers have also tried to gauge AI efficiency through energy use per query, a metric that functions much like a car's miles-per-gallon rating.⁵² Currently, the absence of standardized data and methodologies produces widely varying results that are difficult to compare or validate.

A 2023 estimate by data scientist Alex de Vries put electricity consumption at approximately 3 watt-hours (Wh) per ChatGPT 3.5 query. His figure rested on several maximal assumptions: a “dense” model architecture that engaged every model parameter for every word generated, long 1,500-word combined queries and outputs, older semiconductor hardware, and chips running at their physical power limit.⁵³ In 2025, EpochAI reached a considerably lower conclusion, estimating that a typical GPT-4o query consumes just 0.3 Wh (with longer queries requiring 2.5-40 Wh), ten times lower than de Vries' figure. That estimate reflected updated assumptions: a ~400-word exchange, a model architecture activating only a fraction of its total parameters, newer and more efficient NVIDIA H100 semiconductors, and chips drawing roughly 70 percent of peak capacity.⁵⁴

Taking a broader view, Google researchers analyzed what they defined as the “full stack” of AI serving infrastructure, including active AI accelerator power, host system energy, idle machine capacity, and data center energy overhead. They estimated that Google's Gemini model uses 0.24 Wh per query. The researchers concluded that comprehensive measurements of AI serving metrics are “critical” for accurately comparing large language models and incentivizing efficiency gains across the full AI serving stack.⁵⁵

Each of these estimates is methodologically defensible, yet none are directly comparable. Without shared definitions of what to measure, which components to include, and how to account for efficiency improvements over time, researchers cannot accurately benchmark AI energy use or track whether models are becoming meaningfully more efficient.

What Steps Have Been Taken?

To fill the electricity consumption information vacuum, the EIA is already conducting a voluntary pilot survey of 196 data center operators across Texas, Washington state, and the Northern Virginia-D.C. region. The questionnaire covers the facilities' disaggregated consumption stack, including energy sources, electricity consumption characteristics, server metrics, and cooling systems.⁵⁶

On March 25, 2026, Senators Elizabeth Warren and Josh Hawley sent a bipartisan letter to the EIA Administrator, inquiring whether the EIA planned to collect more information on data center electricity consumption, with Warren noting that “Americans deserve to know how much energy data centers are sucking up and what that's doing to their utility bills.”⁵⁷ The Administrator responded that the EIA is planning to launch a second data center survey to target more facilities and collect their site characteristics (e.g., square

⁵² Efram and Assadi, *Future-Proof AI Data Centers, Grid Reliability, and Affordable Energy*.

⁵³ Alex de Vries, “The Growing Energy Footprint of Artificial Intelligence,” *Joule* 7, no. 10 (2023): 2191–94, <https://doi.org/10.1016/j.joule.2023.09.004>; Josh You, “How Much Energy Does ChatGPT Use?,” EpochAI, February 7, 2025, <https://epoch.ai/gradient-updates/how-much-energy-does-chatgpt-use>.

⁵⁴ You, “How Much Energy Does ChatGPT Use?”

⁵⁵ Elsworth et al., “Measuring the Environmental Impact of Delivering AI at Google Scale.”

⁵⁶ U.S. Energy Information Administration, “EIA Press Release: EIA Launches Pilot Survey on Energy Use at Data Centers,” Eia.Gov, March 25, 2026, <https://www.eia.gov/pressroom/releases/press585.php>; Molly Taft, “The US Government Will Ask Data Centers How Much Power They Use,” *Wired*, April 15, 2026, <https://www.wired.com/story/the-us-government-to-ask-data-centers-how-much-power-they-use/>.

⁵⁷ Taft, “The US Government Will Ask Data Centers How Much Power They Use.”

footage), behind-the-meter power generation, and IT specifications.⁵⁸ The Administrator also noted that the responses from these voluntary questionnaires will eventually help shape a mandatory national survey. However, this commitment exists only as a stated intent, not as a rulemaking, appropriation, or statutory requirement; its scope, timeline, and durability remain subject to shifting agency priorities, leadership changes, and budget constraints.⁵⁹

NERC has taken parallel action on grid reliability. As emerging large loads began connecting to the grid and creating new reliability risks, NERC created a Large Loads Working Group (LLWG) to identify mitigation strategies.⁶⁰ In 2025, NERC issued a Level 2 Alert after finding that large loads triggered “a series of disturbances” on the BPS.⁶¹ After analyzing the responses to that alert, NERC concluded that grid “entities generally did not have sufficient processes, procedures, or methods to address risks associated with computational loads.”⁶² In 2026, NERC escalated to a Level 3 Alert, citing new “critical risks” from computational loads including AI training, cryptocurrency mining, and traditional data center uses, and urging grid entities to implement immediate essential actions covering the modeling, study, commissioning, operation, protection, and control of these loads.⁶³

Policy Recommendations

Even though federal bodies are moving toward increased reporting, their efforts remain incomplete and reversible. The EIA’s pilot surveys are voluntary and geographically limited, and while the Administrator has signaled his intent to develop a mandatory national survey, that commitment carries no statutory force and no guaranteed timeline.⁶⁴ NERC’s Level 3 Alert, meanwhile, urges essential actions that large load owners are not obligated to follow. Without congressional action to codify mandatory, standardized, and system-wide reporting, the blind spots this paper documents will persist even as data center demand accelerates. ARI therefore proposes the following recommendations to convert today’s voluntary, fragmented efforts into a durable reporting infrastructure:

Expand EIA Surveys

Congress should expand the EIA’s pilot surveys to collect first-party data that would enable every downstream stakeholder to manage what they can (finally) measure. To that end, Congress should direct and fund the EIA to create a mandatory, standardized consumption reporting survey for data centers and facilities whose primary purpose is to house computing infrastructure. Congress must ensure that the EIA’s standard comprehensively captures the disaggregated elements of each facility’s full consumption stack, with AI data centers reporting training and inference metrics. At a minimum, the reporting standard should encompass IT

⁵⁸ Ibid.

⁵⁹ Tristan Abbey, “Letter to Senator Warren-EIA,” April 9, 2026, https://www.warren.senate.gov/imo/media/doc/sen_elizabeth_warren_-_eia.pdf.

⁶⁰ “Large Loads Working Group (LLWG),” North American Electric Reliability Corporation, May 11, 2026, <https://www.nerc.com/who-we-are/committees/reliability-and-security-technical-committee-rstc/subcommittees-working-groups-and-task-forces/large-loads-working-group-llwg>.

⁶¹ Engineering Group, “Large Load Interconnection, Study, Commissioning, and Operations,” North American Electric Reliability Corporation, September 9, 2025, <https://www.nerc.com/globalassets/programs/bpsa/alerts/2025/nerc-alert-level-2--large-loads.pdf>.

⁶² “Essential Action to Industry Computational Load Modeling, Studies, Instrumentation, Commissioning, Operations, Protection, and Control,” North American Electric Reliability Corporation, May 4, 2026, <https://www.nerc.com/globalassets/programs/bpsa/alerts/level-3-computational-load-alert.pdf>.

⁶³ “Essential Action to Industry Computational Load Modeling, Studies, Instrumentation, Commissioning, Operations, Protection, and Control.”

⁶⁴ Taft, “The US Government Will Ask Data Centers How Much Power They Use.”

and related equipment; power delivery systems; overhead energy; on-site generation; AI workload types; efficiency gains, including ITUE; miscellaneous loads; and site characteristics such as facility size. Since data protection is a key barrier to operator disclosures of disaggregated data, the EIA can anonymize and aggregate the reported information before public release.

Increase Reporting Through NERC Reliability Standards

Emerging large loads are currently not required to be registered entities under NERC's Registry Criteria Rules of Procedure, meaning they do not have to adhere to the watchdog's Reliability Standards. Even though large loads impact the BPS, NERC's Reliability Standards are currently explicitly designed for entities that own, operate, or use the BPS; these include generation and transmission entities, not end-use customers.⁶⁵

If the Federal Energy Regulatory Commission (FERC) and NERC applied these Standards to large loads, compliance could transform data collection. Large load owners would have to submit simulation data, such as peak demand, facility protection settings, and dynamic models, as well as real-time and day-ahead operational consumption profiles. This registration would eliminate hurdles in event investigations where operators struggle to obtain information from unregistered customers and increase the number of recorders that validate facility modeling and retroactively diagnose disruption causes. Furthermore, mandatory data sharing would allow grid planners to accurately forecast unprecedented load growth and align data center interconnection requests with long-term generation and transmission adequacy.

Conclusion

In 1974, a witness before the Senate Interior Committee warned that the energy crisis was worsened by a "real crisis of confidence," one rooted in the government's inability to obtain standardized reporting from the energy industry. Congress responded by building the institutions and authorities needed to collect, analyze, and publicly disseminate energy information. That information infrastructure became the foundation of U.S. energy policy for decades. Today, the United States faces a new kind of energy crisis: an explosion of demand that its institutions are not yet fully tracking or mitigating. The data centers driving this demand operate largely outside standardized reporting requirements, leaving grid planners, researchers, and policymakers in the dark.

To close the data gap, Congress must mandate nationwide EIA data center electricity consumption reporting, and NERC must collect granular load data. Together, these reforms would convert today's disclosures into the reporting infrastructure that energy analysts need to model demand, that grid operators need to safeguard grid reliability, and that policymakers need to make informed decisions about the country's energy future.

In the 1970s, Congress understood that markets, grids, and public trust all depend on a shared, credible picture of what the country is producing and consuming, and it built the institutions to provide one. Half a century later, Congress must put those same institutions to new work.

⁶⁵ Office of Public Participation (OPP), "Reliability Explainer," Federal Energy Regulatory Commission, accessed May 20, 2026, <https://www.ferc.gov/reliability-explainer>.

Bibliography

- Abbey, Tristan. "Letter to Senator Warren-EIA." April 9, 2026.
https://www.warren.senate.gov/imo/media/doc/sen_elizabeth_warren_-_eia.pdf.
- Aljbour, Jordan, Tom Wilson, and Poorvi Patel. "Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption." Electric Power Research Institute, May 28, 2024.
<https://www.epri.com/research/products/3002028905>.
- Anders, Roger. "The Federal Energy Administration." *U.S. Department of Energy*, November 1980.
<https://www.energy.gov/management/articles/federal-energy-administration>.
- Arora, Chhavi, Marc Sorel, Arjita Bhan, et al. "The next Big Shifts in AI Workloads and Hyperscaler Strategies." McKinsey & Company, December 17, 2025.
<https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-next-big-shifts-in-ai-workloads-and-hyperscaler-strategies>.
- Berreby, David. "As Use of A.I. Soars, So Does the Energy and Water It Requires." Yale E360, February 6, 2024.
<https://e360.yale.edu/features/artificial-intelligence-climate-energy-emissions>.
- Bothwell, Brian, and Kevin Walsh. *Artificial Intelligence: Generative AI's Environmental and Human Effects*. Technology Assessment. United States Government Accountability Office, 2025.
<https://www.gao.gov/products/gao-25-107172>.
- Butterfield, Morgan, and Tim Hess. "EIA Press Release: EIA Forecasts Strongest Four-Year Growth in U.S. Electricity Demand since 2000, Fueled by Data Centers." Washington, DC, January 13, 2026.
<https://www.eia.gov/pressroom/releases/press582.php>.
- Cornell Law School Legal Information Institute. "15 U.S. Code § 772 - Administrator's Information-Gathering Power." Accessed March 12, 2026. <https://www.law.cornell.edu/uscode/text/15/772>.
- Cornell Law School Legal Information Institute. "42 U.S. Code § 7135- Energy Information Administration." U.S. Code. Accessed March 12, 2026. <https://www.law.cornell.edu/uscode/text/42/7135>.
- Crownhart, Casey. "AI Is Changing the Grid. Could It Help More than It Harms?" *MIT Technology Review*, September 9, 2025. <https://www.technologyreview.com/2025/09/09/1123404/ai-grid-help/>.
- Elsworth, Cooper, Keguo Huang, David Patterson, et al. "Measuring the Environmental Impact of Delivering AI at Google Scale." arXiv:2508.15734. Preprint, arXiv, August 21, 2025. <https://doi.org/10.48550/arXiv.2508.15734>.
- Engineering Group. "Essential Action to Industry Computational Load Modeling, Studies, Instrumentation, Commissioning, Operations, Protection, and Control." North American Electric Reliability Corporation, May 4, 2026. <https://www.nerc.com/globalassets/programs/bpsa/alerts/level-3-computational-load-alert.pdf>.

- Engineering Group. “Large Load Interconnection, Study, Commissioning, and Operations.” North American Electric Reliability Corporation, September 9, 2025. <https://www.nerc.com/globalassets/programs/bpsa/alerts/2025/nerc-alert-level-2--large-loads.pdf>.
- Esram, Nora, and Camron Assadi. *Future-Proof AI Data Centers, Grid Reliability, and Affordable Energy: Recommendations for States*. American Council for an Energy-Efficient Economy (ACEEE), 2025. <https://www.aceee.org/white-paper/2025/04/future-proof-ai-data-centers-grid-reliability-and-affordable-energy>.
- Google. “Growing the Internet While Reducing Energy Consumption.” Data Centers. Google. Accessed May 12, 2026. <https://datacenters.google/efficiency>.
- Holt, Mark, Ashley Lawson, and Michael Ratner. “The U.S. Energy Information Administration.” Congressional Research Service, September 9, 2020. <https://www.congress.gov/crs-product/R46524>.
- Jhaveri, Mitul, and Vijaykumar Palat. “Measuring AI’s Energy/Environmental Footprint to Assess Impacts.” *Federation of American Scientists*, June 27, 2025. <https://fas.org/publication/measuring-and-standardizing-ais-energy-footprint/>.
- Koomey, Jonathan, Zachary Schmidt, and Tanya Das. *Electricity Demand Growth and Data Centers: A Guide for the Perplexed*. Bipartisan Policy Center, 2025. <https://bipartisanpolicy.org/report/electricity-demand-growth-and-data-centers/>.
- Large Loads Task Force. *Characteristics and Risks of Emerging Large Loads*. White Paper. North American Electric Reliability Corporation, 2025. <https://www.nerc.com/globalassets/who-we-are/standing-committees/rstc/whitepaper-characteristics-and-risks-of-emerging-large-loads.pdf>.
- Masanet, Eric, Nuoa Lei, and Jonathan Koomey. “To Better Understand AI’s Growing Energy Use, Analysts Need a Data Revolution.” *Joule* 8, no. 9 (2024): 2427–36. <https://doi.org/10.1016/j.joule.2024.07.018>.
- Mason, Kyle, Robert Freudenberg, Christine Garner, and Dave Zachin. “The Rise of Data Centers in the Grid.” Regional Plan Association, June 30, 2025. <https://rpa.org/news/lab/the-rise-of-data-centers>.
- McLaughlin, Tim. “Big Tech’s Data Center Boom Poses New Risk to US Grid Operators.” *Technology*. Reuters, March 19, 2025. <https://www.reuters.com/technology/big-techs-data-center-boom-poses-new-risk-us-grid-operators-2025-03-19/>.
- Nadel, Steven. *Opportunities to Use Energy Efficiency and Demand Flexibility to Reduce Data Center Energy Use and Peak Demand*. American Council for an Energy-Efficient Economy (ACEEE), 2025. <https://www.aceee.org/white-paper/2025/10/opportunities-use-energy-efficiency-and-demand-flexibility-reduce-data-center>.
- Neukomm, Monica. “End-Use Load Profiles for the U.S. Building Stock.” Energy.Gov, December 4, 2018. <https://www.energy.gov/cmei/buildings/articles/end-use-load-profiles-us-building-stock>.
- North American Electric Reliability Corporation. “Large Loads Working Group (LLWG).” May 11, 2026. <https://www.nerc.com/who-we-are/committees/reliability-and-security-technical-committee-rstc/subcommittees-working-groups-and-task-forces/large-loads-working-group-llwg>.

- North American Electric Reliability Corporation. "Resource Adequacy Risks Intensify Across North America as Demand Growth Surges." *Newsroom*, January 29, 2026.
<https://www.nerc.com/newsroom/resource-adequacy-risks-intensify-across-north-america-as-demand-growth-surges>.
- O'Donnell, James, and Casey Crownhart. "We Did the Math on AI's Energy Footprint. Here's the Story You Haven't Heard. | MIT Technology Review." MIT Technology Review, May 20, 2025.
https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/?gad_source=1&gad_campaignid=16701804390&gbraid=0AAAAADgO_mj7VQ9DMooJho-b9cGctnSkt&gclid=Cj0KCQiA5abIBhCaARIsAM3-zFX8NGkX7Biy9KwFWgdImGiQ8c01w1IhIG4DdFC_RObMK_7Now6IBBkaAiiXEALw_wcB.
- Office of Public Participation (OPP). "Reliability Explainer." Federal Energy Regulatory Commission. Accessed May 20, 2026. <https://www.ferc.gov/reliability-explainer>.
- Reed, Rachel. "How Data Centers May Lead to Higher Electricity Bills." *Harvard Law School*, September 3, 2025.
<https://hls.harvard.edu/today/how-data-centers-may-lead-to-higher-electricity-bills/>.
- Rep. Hastings, James F. "Text - H.R.14368 - 93rd Congress (1973-1974): Energy Supply and Environmental Coordination Act of 1974." June 22, 1974. 1974-04-24.
<https://www.congress.gov/bill/93rd-congress/house-bill/14368/text>.
- ScienceDirect. "Load Profile - An Overview." Accessed May 20, 2026.
<https://www.sciencedirect.com/topics/engineering/load-profile>.
- Shehabi, Arman, Sarah Josephine Smith, Alex Hubbard, Alexander Newkirk, Nuoa Lei, Md AbuBakar Siddik, Billie Holecek, Jonathan G. Koomey, Eric R. Masanet, and Dale A. Sartor. *United States Data Center Energy Usage Report*. 2024. <https://eta.lbl.gov/publications/2024-lbnl-data-center-energy-usage-report>.
- Smith, Addy, Alice Wu, Arjun Krishnaswami, Cole Donovan, and Zoe Brouns. "New DOE Re-Organization Raises Uncertainty for American Science, Energy Innovation, and Affordability." *Federation of American Scientists*, November 26, 2025.
<https://fas.org/publication/new-doe-re-organization-raises-uncertainty-for-american-science/>.
- SolarTech. "How Much Electricity Does a Data Center Use? Complete 2025 Analysis." *IAEI Magazine*, January 1, 2026.
<https://iaeimagazine.org/electrical-fundamentals/how-much-electricity-does-a-data-center-use-complete-2025-analysis/>.
- Taft, Molly. "The US Government Will Ask Data Centers How Much Power They Use." *Wired*, April 15, 2026.
<https://www.wired.com/story/the-us-government-to-ask-data-centers-how-much-power-they-use/>.
- Tozzi, Christopher. "Beyond Megawatts: Rethinking How We Measure Data Center Capacity." *DataCenterKnowledge*. Accessed May 20, 2026.
<https://www.datacenterknowledge.com/energy-power-supply/beyond-megawatts-rethinking-how-we-measure-data-center-capacity>.

- U.S. Energy Information Administration. “EIA Press Release: EIA Launches Pilot Survey on Energy Use at Data Centers.” Eia.Gov, March 25, 2026. <https://www.eia.gov/pressroom/releases/press585.php>.
- U.S. Energy Information Administration. “Legislative Timeline.” Eia.Gov. Accessed May 12, 2026. https://www.eia.gov/about/legislative_timeline.php.
- Vann, Adam. “The Legal Framework of the Federal Power Act.” Congressional Research Service, January 22, 2020. <https://www.congress.gov/crs-product/IF11411>.
- Vries, Alex de. “The Growing Energy Footprint of Artificial Intelligence.” *Joule* 7, no. 10 (2023): 2191–94. <https://doi.org/10.1016/j.joule.2023.09.004>.
- Walton, Robert. “NERC Issues Level 3 Alert, Mandates Action to Address Data Center Load Losses.” Utility Dive, May 5, 2026. <https://www.utilitydive.com/news/nerc-issues-rare-level-3-alert-over-data-center-load-losses/819295/>.
- Warren, Elizabeth, and Josh Hawley. “Letter to EIA: Data Center Transparency.” March 25, 2026. https://www.warren.senate.gov/imo/media/doc/letter_to_eia_re_data_center_transparency.pdf.
- You, Josh. “How Much Energy Does ChatGPT Use?” EpochAI, February 7, 2025. <https://epoch.ai/gradient-updates/how-much-energy-does-chatgpt-use>.