2026 White Paper

Powering Intelligence 2026

Updated Scenarios of U.S. Data Center Electricity Use and Power Strategies

Available online: powering-intelligence.epri.com

TABLE OF CONTENTS

EXECUTIVE SUMMARY 3

INTRODUCTION 10

UNDERSTANDING KEY DATA CENTER POWER METRICS 10

- Nominal Capacity 11
- Nameplate Capacity and Power Usage Effectiveness (PUE) 11
- Annual and Peak Electricity Use 11

NOMINAL CAPACITY: HISTORY AND FUTURE PROJECTIONS 12

- Future Scenarios 12

ANNUAL AND PEAK ELECTRICITY USE: HISTORY AND FUTURE PROJECTIONS 15

SUMMARY OF FUTURE PROJECTIONS 19

- Comparison to Other Studies and Forecasts 20

DATA CENTER LOAD GROWTH IN CONTEXT 22

- Data Center Shares of Total Electricity Demand 22
- Comparison to EV Charging Load 24

GENERATION AND CAPACITY IMPACTS OF DATA CENTER LOAD 25

CONCLUSIONS 27

REFERENCES 28

EXECUTIVE SUMMARY

Key Messages

- **Data centers, industrial onshoring, and transport electrification are driving renewed regional load growth.** In the United States and other nations, clusters of new, large loads are testing utilities' ability to keep pace and are spurring technical and business innovation.
- **Artificial Intelligence (AI) plays a growing, but uncertain role in data center load.** AI workloads are estimated to account for 15–25% of data center electricity today and that share is rising rapidly even as non-AI data center demand continues steady growth. However, the pace and scale of AI adoption, the power intensity of AI hardware and algorithms, and power system and supply chain constraints remain highly uncertain.
- **EPRI projects data centers to consume 9% to 17% of U.S. electricity by 2030, up from 4% to 5% today.** EPRI developed Low, Medium, and High growth scenarios of U.S. state-level data center power demand through 2030 using commercial project development data. The Low scenario assumes most projects under construction and one-fourth of those in advanced planning are fully operational by 2030. The High Scenario assumes all projects under construction or in advanced planning plus 30% of those in early planning quickly overcome supply chain and process constraints to be operational by 2030.
- **EPRI's revised projections are about 60% higher than its 2024 report estimates.** The increase is driven primarily by record levels of development over the past 18 months. The revised range is roughly consistent through 2028 with Lawrence Berkeley National Lab's projections in their December 2024 report to Congress, which was based upon projected chip and equipment shipments rather than commercial development data.
- **State-level load growth projections vary widely, creating localized challenges.** Today, Virginia is the only state where data centers consume over 20% of electricity; this share could increase to 39% to 57% by 2030. In the Medium scenario, seven additional states could exceed a 20% share by 2030.
- **Energy procurement strategies drive sharply different generation buildouts to meet this demand.** Assuming current U.S. state and federal energy policies, least-cost procurement strategies favor natural gas, with build rates in the High scenario more than double the recent average. In contrast, meeting all data center load with hourly-matched, carbon-free energy would favor renewable and battery additions, with new nuclear generation coming online where feasible. Supply bottlenecks as well as delays in permitting and siting could constrain the modeled additions of both generation and transmission capacity.
- **Accelerated collaboration is essential to support these levels of data center growth.** Collaboration between data center developers, energy and equipment providers, policymakers, and communities is key to data center buildout. Efforts like EPRI's DCFlex initiative, which is working with more than 60 companies to address grid reliability, flexibility, affordability, and other issues central to transforming data centers into grid assets, can help accelerate progress.

AI Power Demands Are Accelerating Data Center Load Growth

Data centers have become the fastest-growing source of U.S. electricity demand, and regional clusters of facilities are transforming local grid dynamics, testing utilities' ability to keep pace and spurring technical and business innovation. Forecasting future data center (DC) load growth is essential for power system planning but remains difficult because public reporting is limited, many announced projects are speculative, and there is fundamental uncertainty about the adoption of generative AI and successor technologies. AI applications are much more energy intensive than the streaming, communications, search, and other traditional data center workloads. While AI workloads are estimated to consume only 15–25% of data center electricity today (IEA (2025), JLL (2026)), that share is rising even as non-AI data center demands continue to grow steadily.

U.S. State-Level Data Center Load Projections

Drawing upon state-level data on operational capacity, construction in progress, and announced plans, EPRI developed three scenarios for U.S. data center capacity growth through 2030 (Figure ES-1):

- **Low growth.** Assumes that most projects under construction and one-fourth of those in advanced planning are fully operational by 2030.
- **Medium growth.** Assumes that all projects under construction, 75% of those in advanced planning, and 10% in early planning are fully operational by 2030.
- **High growth.** Assumes that all projects under construction or in advanced planning plus 30% of those in early planning are fully operational by 2030.

Figure ES-1. Data center nominal IT capacity by state. *Inner circles show capacity in 2021 (gray) and 2024 (blue). Outer band shows scenario range of projected capacity in 2030 (orange). Circle area is proportional to nominal IT capacity.*[1] *Estimates include small- and large-scale data centers as well as cryptocurrency mining. Results highlight continued concentration in established markets (e.g., Virginia, Texas) alongside emerging growth in new states as developers diversify geographically.*

---

[1] Nominal capacity of a data center typically refers to the maximum potential information technology (IT) load (i.e. servers and other IT equipment, exclusive of ancillary energy use for cooling and infrastructure).

These scenarios were designed to reflect uncertainty about the realization of planned projects and are not necessarily equally likely.[2] The Low growth scenario serves as a floor, assuming active construction sites are accurately estimated and come online, while the High growth scenario assumes that power, supply chain (materials and labor), and process constraints for many projects still in early planning can be quickly overcome. The Medium growth scenario falls roughly halfway between the other two.

In 2024, total U.S. data center nominal IT capacity is estimated to be 35 to 44 GW, a range that reflects uncertainty across accessed data sources. **By 2030, EPRI's scenarios for the realization of planned and announced projects result in a range of 56 to 132 GW U.S. total nominal capacity.** The results show wide variations across states in both the historic estimates and projected expansion (Figure ES-1). Capacity continues to accumulate in Virginia, Texas, and other primary data center markets, but the emergence of new capacity in other states such as Ohio, Indiana, Pennsylvania, Louisiana, and Mississippi, where there is currently little operating capacity, suggests increased prioritization of power access and land availability, particularly for large AI training centers.[3]

Translating nominal IT capacity estimates to peak electric load and annual electricity demand requires making assumptions about cooling and other on-site loads, data center annual load factors (i.e., variations in demand that differ by data center type and function), and the ramp rate for operationalizing nominal capacity. In 2024, total U.S. data center electricity use is estimated at 177–192 TWh, growing to **roughly 380 to 790 TWh by 2030**. The estimated 2030 range is around 60% higher than EPRI's prior estimates published in 2024, largely reflecting the accelerated pace of data center development over the past 18 months (Figure ES-2).

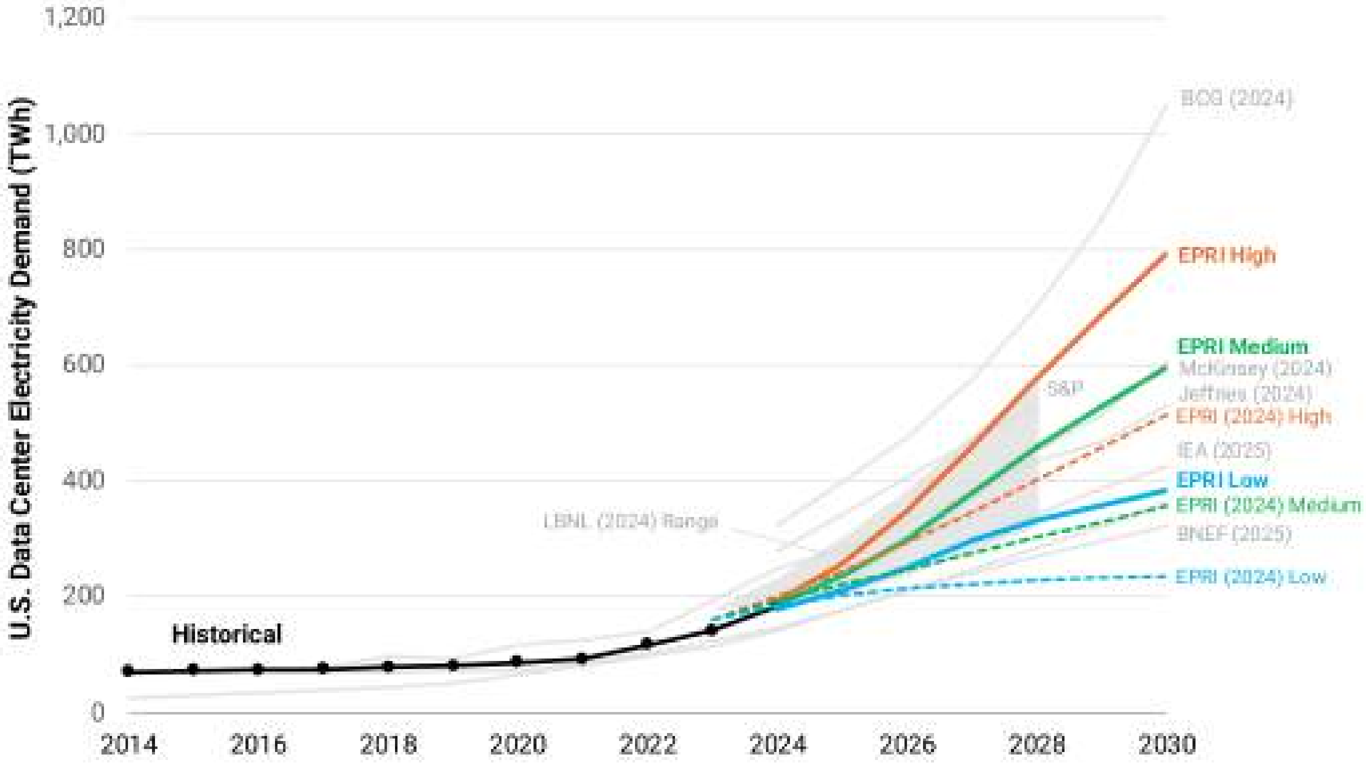

Figure ES-2. Comparison of U.S. data center annual electricity consumption projections. *Projections in this study span a similar range to the LBNL (2024) report. Shaded band shows scenarios from LBNL (2024); lines show recent external estimates, including BCG (2024), BloombergNEF (2025), EPRI (2024a),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

---

[2] See the Nominal Capacity section of the full report for more precise definitions of these scenarios and the data underlying them.

[3] Note that a shortcoming of the estimation approach used in this paper is that it relies heavily on public announcements of projects. For various reasons, the stage of project development where public announcements are made can differ by state and developer type.

Figure ES-2 shows projections of annual U.S. data center annual energy use through 2030 relative to other studies. EPRI's new scenarios span a range consistent with projections developed by LBNL (2024), despite fundamental differences in how they were created. The starting point for LBNL's estimates is shipments of processor chips and other IT equipment, versus the EPRI starting point of data center site development data.[4]

The aggregate peak load for operating U.S. data centers is estimated to be 21–22 GW in 2024, with projected growth by 2030 to between 45 GW (Low), 71 GW (Medium), and 94 GW (High), lower than the corresponding nominal IT capacity estimates.[5] The analysis shows that **announced nominal MW for data center projects should be treated as a pipeline indicator rather than a near-term peak forecast**, as non-IT loads, ramp-up, load shapes, onsite energy assets, and load flexibility materially impact peak effects. 2025 FERC forecasts[6] of peak demand growth align well with EPRI's High scenario in most regions.

Growing Data Center Share of Electricity Demand

Figure ES-3 shows the updated EPRI projections in the context of modeled scenarios for economy-wide electricity load growth. **The data center share of U.S. total electricity demand in 2030 ranges from 9% to 17%, an increase from 4–5% today.** At the state level, continued development of the largest DC market in Virginia implies a share increasing to between 39% and 57% by 2030, reflecting the many projects currently under construction or in advanced planning.

The figure also highlights wide variations across states in their pace of data center development. Today, Virginia is the only state where data centers consume over 20% of electricity. By 2030, seven additional states—Oregon, Iowa, Nebraska, Nevada, Wyoming, Arizona, and Indiana—could see data centers exceeding a 20% share (Medium scenario). Other states, such as Washington and New Jersey, are above the U.S. average today but have relatively little estimated capacity in construction or planning (based upon accessed public sources) and hence relatively low increases in their data center shares by 2030. On the other hand, several states with some existing capacity (e.g. New Mexico, Ohio, and Pennsylvania) and others with very little (e.g. Indiana, Louisiana, and Mississippi) are projected to emerge as new areas of concentrated development with 2030 shares exceeding 10%.

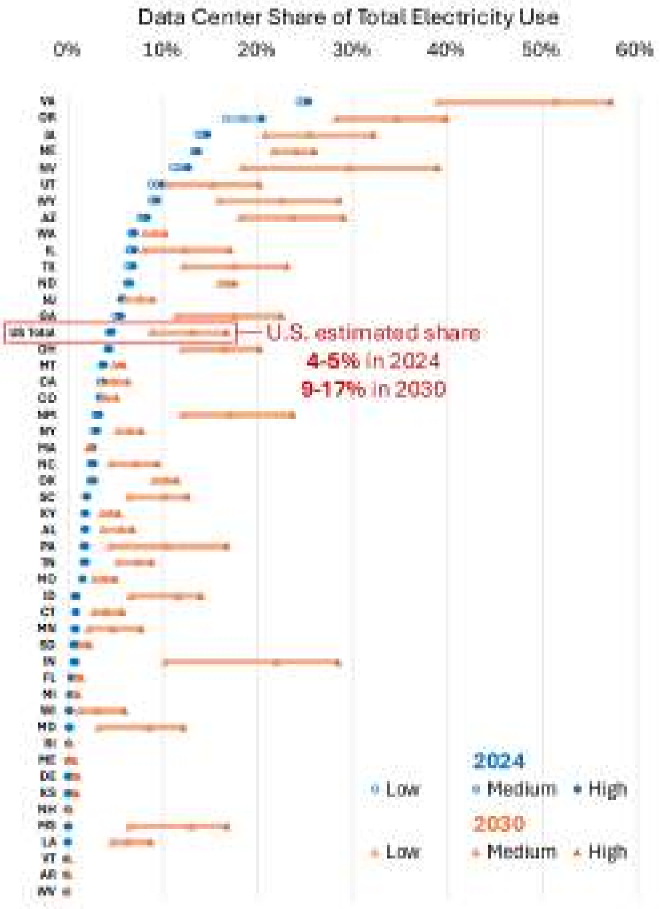

Figure ES-3. Data center share of total electricity demand by state. *Estimates include small- and large-scale DC as well as cryptocurrency mining. DC scenarios for 2030 are based on EPRI analysis of a range of industry sources. Non-DC electricity demand is based on US-REGEN model scenarios with current federal and state policies.*

---

[4] The LBNL study's projections also exclude cryptocurrency load, which is included in EPRI's estimates.

[5] These estimates reflect projected peak data center load before any potential flexible demand response. EPRI's DCFlex initiative, described later in this summary, is demonstrating data center demand flexibility as a strategy to speed access to power, delay or reduce grid buildout, and improve grid reliability.

[6] Utilities and balancing areas report peak load forecasts through FERC Form 714.

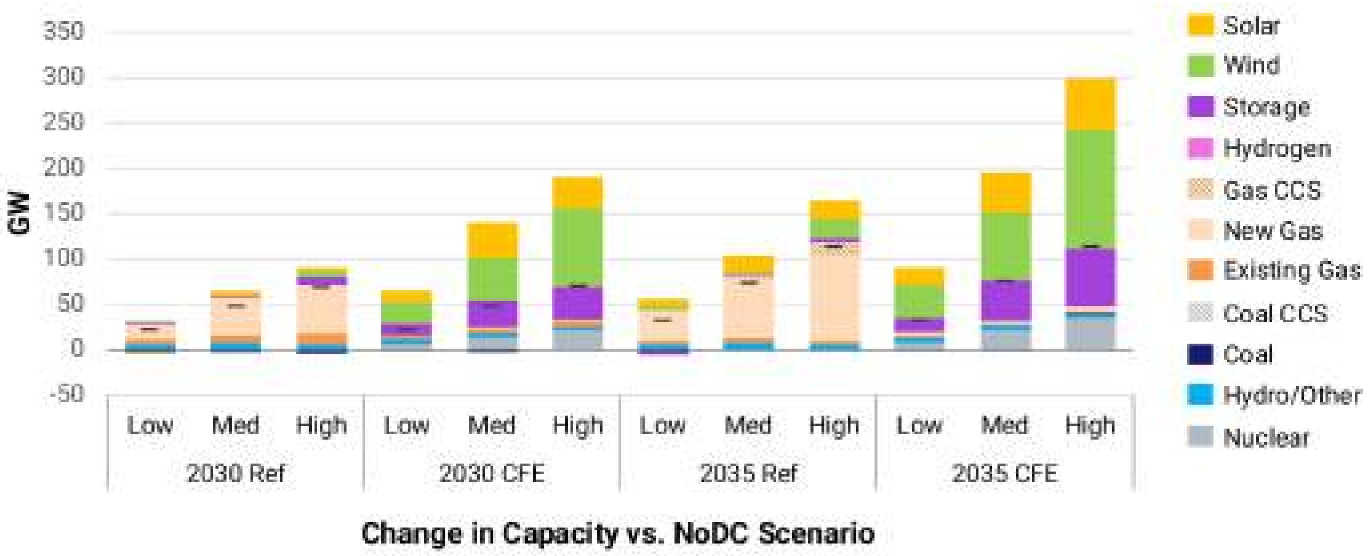

Figure ES-4. Change in U.S. capacity to meet new data center load. *Results from US-REGEN model scenarios with reference policies ("Ref") and 24/7 carbon-free energy targets ("CFE").*

Technologies for Powering Near-term Data Center Growth

Electric system responses to meet data center demand depend strongly on the policy environment and data center energy procurement objectives.[7] Figure ES-4 illustrates that under current state and federal policies, natural gas may dominate near-term incremental supply. If all data centers adopt and achieve 24/7 carbon-free energy (CFE) targets, portfolios shift towards low-emitting generation and energy storage.

Projected annual natural gas capacity builds from 2025 to 2030 range from 6.6 to 13.7 GW per year with reference policies, higher than the average of 5.7 GW per year over the past five years. These impacts are a departure from EPRI's 2024 analysis (EPRI 2024b), which included significantly higher wind and solar deployment to meet growing DC loads due to the Inflation Reduction Act's production and investment tax credits, which have been curtailed for many technologies under the 2025 budget bill (EPRI, 2026).

Under 24/7 CFE targets, generation and capacity responses to higher DC load come from wind, solar, nuclear, and energy storage. Nuclear and energy storage remain eligible for investment tax credits under current policy, increasing their value in CFE-constrained portfolios. Portfolios of incremental supply vary by region.

Modeled results provide an indication of the scale of new resources needed to meet projected data center load growth. However, supply chain bottlenecks from equipment manufacturing to permitting and siting processes could constrain additions of both generation and transmission capacity.

---

[7] This paper is generally agnostic about ownership and location of generation and storage to meet data center demand. For a discussion of the trade-offs of strategies ranging from traditional passive grid connections to off-grid power, see EPRI's report: Reconciling the Value of Grid Interconnection and Speed to Power: Strategies for Powering Data Centers in the AI Era.

Actions to Support Rapid, Reliable, Affordable Data Center Expansion

Data center expansion of the scale projected here faces many challenges. At the local and regional level, challenges arise from the scale of the centers themselves and mismatches in infrastructure timing. A typical new data center of 100 to 1000 megawatts represents a load equal to that of a new neighborhood of 80,000 to 800,000 average homes. While neighborhoods and the grid require many years to plan and build, data centers can be developed and connected in a few years. Added to the timing challenge are the supply chain issues associated with the scale of this growth. IT and power equipment and skilled labor are both regional and national-level challenges.

Essential strategies to support rapid, reliable, and affordable data center expansion include:

- **Improved short- and longer-term load forecasting for large loads.** Better tools are essential to enable effective responses to short-term system disruptions and to inform long-term investments in the grid. This report highlights both new approaches to understand data center load growth and the many data gaps and uncertainties that have hindered the effort. With grid development times reaching years to decades, rapid advances are needed to guide efficient investments.
- **More effective utilization of existing grid assets.** Identifying locations with underused assets, increased data center and grid load flexibility, generation capacity increases at existing sites, and grid-enhancing technologies that quickly reduce grid congestion are among the key strategies to better utilize the existing grid to accelerate cost-effective data center load growth.
- **Closer collaboration between data center developers, energy and equipment providers, policymakers and communities.** Collaboration is essential to maintain and enhance grid reliability and to address affordability and community impacts as data centers connect to the grid. EPRI-led collaborative efforts include:
  - **DCFlex.** EPRI is working with more than 60 companies through the DCFlex initiative to address grid reliability, flexibility, affordability and other issues central to transforming data centers into grid assets.
  - **GET SET.** EPRI's Grid-Enhancing Technologies for a Smart Energy Transition (GET SET) Initiative supports utility implementation of technologies and strategies to maximize existing grid assets.
  - **Mercury.** EPRI is partnering with Kraken on the Mercury Consortium to advance grid-edge technologies, including flexible load devices, and expand customer choice.
  - **Distributed Data Centers.** EPRI recently announced a new collaboration with Prologis, NVIDIA, and InfraPartners to study smaller-scale data centers designed for distributed inference.

INTRODUCTION

Data centers have become the fastest-growing source of U.S. electricity demand, and regional clusters of facilities are transforming local grid dynamics, testing utilities' ability to keep pace and spurring technical and business innovation. Computational service demands have increased rapidly, fueled by increased consumer demand for streaming and other data-intensive services, cryptocurrency, and artificial intelligence (AI). Data center (DC) capacity has grown rapidly to meet these surging demands. While U.S. data center development has historically been concentrated in several regions, the recent prioritization of power availability, clean power, and large tracts of land are causing development to spread.

A major uncertainty in projecting data center load growth is the broad emergence of AI technologies, highlighted by the rapid adoption of generative AI models since November 2022. AI applications are much more energy intensive than the streaming, communications, search, and other traditional data center workloads. An early rule of thumb held that an AI query required roughly ten times the power of a traditional search query; that comparison is increasingly outdated because of hardware and algorithmic advances and changing user behavior. Verbal requests (rather than text), contextual inputs (for example photos or reports), and content generation such as images and video all require additional processing and have no analog to earlier user requests. These interactions are becoming more common and more power intensive. While AI workloads are estimated to consume 15–25% of data center electricity today (IEA (2025), JLL (2026)), that share is rising even as non-AI data center demands continue to grow steadily.

Data center usage was once limited by human attention —how many movies a person can watch at one time or how many queries they can type. Advances such as deep research and agentic AI mean that usage is increasingly driven by hundreds or thousands of AI-generated requests that require only occasional human oversight. With 5.3 billion global internet users, widespread adoption of these tools could produce a step change in power requirements. On the other hand, history has shown that demand for increased processing has largely been offset by data center efficiency gains.

Although DC electricity demand is not publicly reported, DC-driven load growth is apparent in aggregate statistics in certain states, and utilities across the nation and around the world are experiencing sharp increases in service connection requests. Forecasting future data center (DC) load growth is essential for power system planning but remains difficult because public reporting is limited, many announced projects are speculative, and there is fundamental uncertainty about the adoption of generative AI and successor technologies. Several recent EPRI reports have discussed the key drivers of DC electricity demand and developed estimates of current and projected future data center load at the state level for the U.S.[8] This report updates EPRI's prior estimates, using new data on planned and announced projects rather than extrapolating historical trends. It also adds deeper modeling of key DC electricity metrics: nominal IT capacity, non-IT data center power use, capacity utilization, and annual and peak energy use.

UNDERSTANDING KEY DATA CENTER POWER METRICS

There is an important distinction and relationship between data center capacity, which refers to the amount of power a data center can supply to its equipment—or demand from the grid—in a given moment (measured in watts), and data center annual energy use, which refers to the actual consumption of power over the course of a year (measured in watt-hours). Both metrics describe the scale of data center activity and electricity demand, but they have different implications for electric system planning. Capacity affects peak supply and grid upgrades, while annual energy affects generation, fuel use, and emissions. Estimating both is uncertain because public data are limited, announced projects may not proceed at planned scale or pace, and projecting future state-level loads from available datasets and developer construction announcements is inherently imprecise.

---

[8] Powering Intelligence (EPRI, 2024a) presents data center load growth projections and discusses technology advances that could slow growth. Powering Data Centers (EPRI, 2024b) updates the load projections and analyzes energy and emissions impacts. Scaling Intelligence (EPRI, 2025a) explores the drivers of AI electricity demands and the impact on data center size, and Speed to Power (EPRI, 2025b) discusses strategies for powering data centers ranging from off-grid to traditional grid-connected power. EPRI's DCFlex initiative, a collaboration between utilities, developers, and other stakeholders, has produced several reports on regional load forecasting, interconnection requirements, system impacts, and opportunities/needs for data center flexibility.

Nominal Capacity

Nominal capacity of a data center typically refers to the maximum potential information technology (IT) load (i.e. servers and other IT equipment, exclusive of ancillary energy use for cooling and infrastructure). Because electric power is so essential to data center operation, this metric is used to characterize data center size and is often published by DC developers in project announcements. Although it is perhaps the most visible and publicly trackable metric available, nominal IT capacity is not always publicly reported. When nominal capacity data are not available either through public or private sources, it can be estimated based upon proxies such as square footage of buildings on a site combined with information on the data center's function, air permits for backup generation, and aerial assessments.

Nameplate Capacity and Power Usage Effectiveness (PUE)

Nameplate capacity of a data center typically refers to the maximum potential total load at the facility, which includes both IT and non-IT load (cooling, power conversion, lighting, etc.), and corresponds to what the data center requests for service interconnection. Non-IT load is often summarized by PUE, the ratio of total load to IT-only load. PUE can vary depending on cooling technology, regional climate, data center architecture and scale, chip design, computational use case, and other factors. Ancillary energy use for cooling and other infrastructure needs within a data center, while declining significantly over the past decades, still accounts for a significant share of total data center energy use. PUE and nameplate capacity are less frequently reported in public data than nominal capacity but can in some cases be estimated from site characteristics.

Annual and Peak Electricity Use

Translating nameplate capacity to annual electricity use and peak system load depends on additional assumptions about the ramp-up of active or operational capacity and patterns of capacity utilization. New facilities typically ramp up operational capacity over time as data halls, IT equipment, or (for colocation data centers) tenants are added. Developers anticipate future growth in computing power requirements, meaning that actual active operational load, especially in early years of operation, can be significantly lower than the stated nominal or nameplate "headline" capacity. The implication, aggregating across all data centers, is that fleet average capacity utilization dips during periods of rapid growth.

Annual utilization factors, i.e. how much active and operational equipment is run over the course of a year, vary across types of data centers and computational services. In general, data centers are operated with the goal of maximizing server uptime, leading to relatively high, constant load shapes when compared to other power demands. At the same time, the peak level of data center load is often significantly lower than the hypothetical maximum based on nominal capacity plus allowance for PUE.

The conceptual relationship between nominal, nameplate, and active capacity and realized peak load is illustrated in Figure 1.

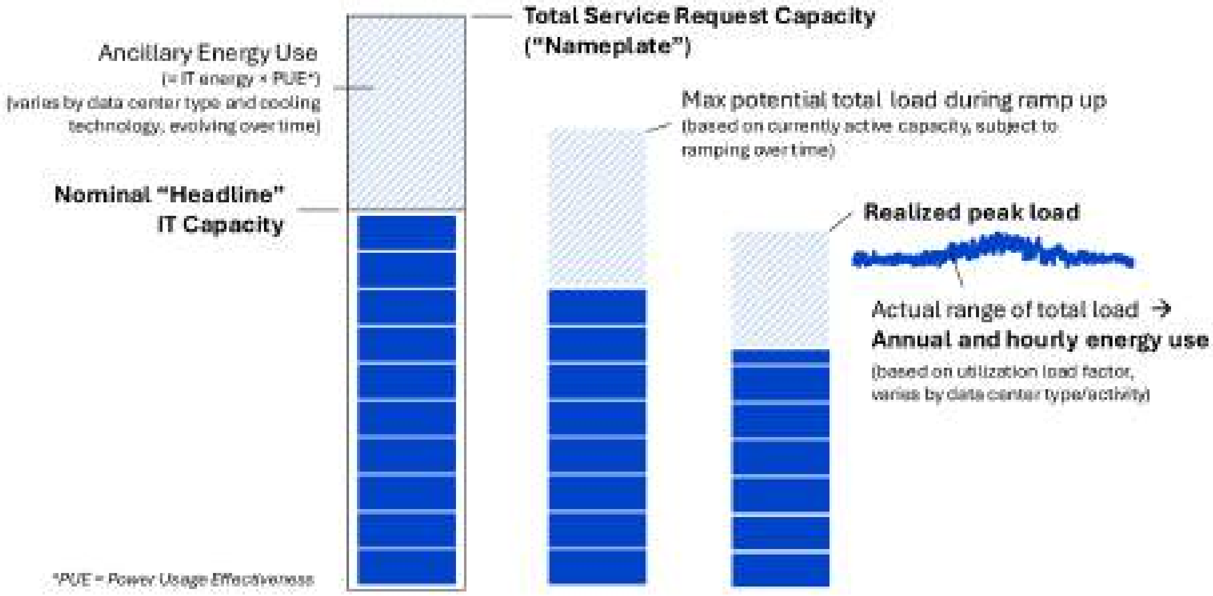

Figure 1. Illustrative diagram of data center capacity and utilization metrics. *Nominal IT capacity (what is typically reported in public announcements) differs from the total service request/nameplate load (which is higher once ancillary energy use is included). However, realized peak demand is lower due to ramp-up of active capacity and utilization.*

NOMINAL CAPACITY: HISTORY AND FUTURE PROJECTIONS

This report synthesizes a wide range of sources on current, historical, and announced future nominal IT capacity.[9] Many entities are attempting to track this information using various methods, with differences in focus and scope. Given the lack of formal reporting requirements or standards, assessments of individual sites can vary significantly, e.g. in how they estimate nominal IT power or how they define under construction or in advanced planning. There can also be discrepancies in the extent of coverage, with some sources including only data centers in major markets rather than providing comprehensive state-wide assessments or excluding data centers used for cryptocurrency mining.

Rapidly accelerating development of sites with increasing geographic diversity and the shift towards energy-intensive AI hardware with equipment secured inside buildings have exacerbated estimation challenges. In addition, single sources change their historical, construction, and planned estimates over time as new information becomes available, developer plans shift, or site visits identify discrepancies.

Despite these limitations, these sources provide a directional view of data center operations and development that serves as a starting point for state-level load analysis and can be validated with other observations. All estimates in this analysis include hyperscale, co-location, enterprise, and cryptocurrency mining unless otherwise noted.[10]

Future Scenarios

Scenarios for projected nominal capacity additions between 2024 and 2030 were developed based on state-level data across several sources that estimate planned data center projects in different stages of development:

- **Under Construction** – Definition can vary by source, based on criteria such as site preparation, boundary fence, building footings, vertical development, etc.
- **Announced (Advanced)** – Planned projects with at least some pre-construction milestones completed, such as track record (e.g., familiar vs new entrant), siting, permitting, power access, contracting, and design (criteria and their application differs by source).
- **Announced (Early)** – Planned projects in early stages of development.

---

[9] Sources include nominal capacity estimates published by Aterio, Avison Young, Baxtel, BloombergNEF, CBRE, Cushman & Wakefield, JLL, S&P Global, and Wood Mackenzie.

[10] Hyperscale data centers are large (or very large) facilities owned and operated by a single tenant, typically a large technology company, for proprietary use for large-scale computation and cloud services. Enterprise data centers are smaller facilities used by individual companies to support internal data services other than large-scale computation. Co-location facilities host data center services for multiple tenants for a variety of activities. Cryptocurrency mining facilities are designed validate blockchain transactions (e.g. Bitcoin), but their power access and land resources are increasingly being shifted to more general DC applications.

Table 1. Scenario assumptions for translating planned data center project pipeline data to realized nominal IT capacity by 2030. *Assumed completion rates vary by development stage for Low/Medium/High scenarios, applied to the maximum and average values across data sources.*[11] *Data was analyzed at the state level, with the same completion rates assumed for all states.*

	UNDER CONSTRUCTION		ANNOUNCED (ADVANCED)		ANNOUNCED (EARLY)	
	Max	**Average**	**Max**	**Average**	**Max**	**Average**
Nominal Capacity (GW)	17.2	12.9	36.7	36.7	114.0	67.4
Assumed Share Completed by 2030, by Scenario						
Low		90%		25%		
Medium	100%		75%		10%	
High	100%		100%		30%	

These estimates were combined into three scenarios (Low, Medium, and High) based on different assumptions about the share of projects in each stage that are completed by 2030. The scenarios generally assume that most or all projects currently under construction are completed, while a lower share for projects in earlier stages of development owing to their more speculative nature (Table 1).[12]

- **Low growth:** Assumes that most projects under construction and one-fourth of those in advanced planning are fully operational by 2030.
- **Medium growth:** Assumes that all projects under construction, 75% of those in advanced planning, and 10% in early planning are fully operational by 2030.
- **High growth:** Assumes that all projects under construction or in advanced planning plus 30% of those in early planning are fully operational by 2030.

These scenarios were designed to reflect uncertainty about the realization of planned projects and are not necessarily equally likely. The Low growth scenario serves as a floor, assuming active construction sites are accurately estimated and come online, while the High growth scenario assumes that power, supply chain (materials and labor), and process constraints for many projects still in early planning can be quickly overcome. The Medium growth scenario falls roughly halfway between the other two. The maximum was used as a basis for the Medium and High scenarios to represent an upper bound across datasets, recognizing that some provided less than comprehensive coverage. Total planned nominal IT capacity across all three categories (i.e., construction, planned (advanced), and planned (early)) amounts to 117 GW using the average across data sources and 168 GW using the maximum. See Figure 6 for a more detailed break-out by stage.

Based on our analysis of a wide range of sources, we estimate total nominal data center capacity in the U.S. in 2024 to be 35–44 GW.[13] This range reflects an increase of more than a factor of three relative to estimated nominal capacity of around 11 GW in 2021. Under the three future scenarios, total **U.S. nominal IT capacity in 2030 reaches 56 GW (Low), 96 GW (Medium), and 132 GW (High)**. Growth between 2024 and 2030 in each scenario is assumed to follow a linear pathway to the 2030 total. Nominal capacity estimates at the state level from 2021 to 2030 are shown in Table 2.

At the state level, projections reflect specific project announcements and indicate the emergence of several new data center clusters as developers increase geographic diversity. Figure 2 shows projected nominal capacity in each state compared to estimates of existing capacity. Capacity continues to accumulate in Virginia, Texas, and other primary data center markets, but the emergence of new capacity in other states such as Ohio, Indiana, Pennsylvania, Louisiana, and Mississippi, where there is currently little operating capacity, suggests increased prioritization of power access and land availability, particularly for large AI training centers. Other states, such as California and New York, have significant existing capacity but relatively little projected growth.

---

[11] EPRI removed clear anomalies from the source data before calculating averages and maximums, e.g., correcting data for facilities that were reported as 1000x their true size or removing facilities that have closed.

[12] Note that a shortcoming of the estimation approach used in this paper is that it relies heavily on public announcements of projects. For various reasons, the stage of project development where public announcements are made can differ by state.

[13] The low end of the range reflects the average of state-level capacity estimates across sources, while the high end reflects the maximum across sources. The medium scenario uses the average of the low and high values in 2024. Estimates of nominal capacity before 2024 also vary across sources, but only single values are shown.

Table 2. State-level data center nominal IT capacity (MW). *Estimates include small- and large-scale DC as well as cryptocurrency mining. 2024 ranges reflect variation across sources.*

STATE	2021	2022	2023	2024			2030		
				LOW	MEDIUM	HIGH	LOW	MEDIUM	HIGH
AL	17	167	170	315	353	390	454	780	981
AK	–	–	–	–	–	–	0	0	0
AR	–	–	–	8	11	14	8	11	14
AZ	472	652	740	1,657	1,864	2,072	2,537	3,853	5,288
CA	935	1,064	1,227	1,280	1,496	1,712	1,489	2,296	3,225
CO	92	248	258	170	241	313	333	404	503
CT	29	30	30	28	29	30	162	243	355
DE	–	–	–	10	14	17	10	14	17
FL	80	97	158	184	188	193	342	524	643
GA	297	789	1,174	1,505	1,774	2,043	3,014	5,358	7,466
HI	1	1	1	7	8	9	7	9	10
ID	15	17	17	69	89	109	290	554	689
IL	793	985	1,069	1,359	1,602	1,845	1,797	3,327	5,263
IN	19	22	160	236	292	348	2,079	5,235	7,423
IA	141	707	160	1,530	1,698	1.866	1,530	2,198	3,366
KS	1	1	1	37	44	51	37	45	54
KY	9	150	185	353	375	397	408	560	678
LA	2	7	9	9	9	9	809	1,130	1,591
ME	–	3	3	3	3	3	3	9	21
MD	9	11	11	60	63	67	547	1,576	2,350
MA	160	160	160	159	171	184	160	171	184
MI	19	21	60	90	97	103	118	179	214
MN	59	59	69	87	106	125	256	600	1,033
MS	–	–	–	31	46	61	602	1,316	1,770
MO	89	111	111	205	224	243	359	556	756
MT	–	66	66	96	106	115	96	106	115
NE	52	352	806	963	1,021	1,079	1,158	1,378	1,543
NV	487	606	613	674	866	1,057	1,268	2,529	3,922
NH	2	2	2	3	4	4	3	4	4
NJ	410	410	555	490	547	605	701	960	1,190
NM	36	46	46	353	393	434	385	676	1,123
NY	217	308	633	964	1,085	1,206	1,112	1,487	1,782
NC	344	491	506	605	746	887	1,096	1,856	2,533
ND	1	101	447	603	640	677	729	780	817
OH	183	396	1,170	1,987	2,104	2,222	3,248	5,087	6,772
OK	33	120	140	670	745	820	868	1,018	1,151
OR	203	741	1,602	2,299	2,982	3,664	2,641	3,635	4,522
PA	41	222	524	540	631	722	1,008	2,802	4,980
RI	1	2	2	2	3	3	2	3	3
SC	79	146	231	528	551	573	871	1,610	2,097
SD	8	8	8	23	30	38	23	31	38
TN	25	91	264	644	755	866	805	1,233	1,527
TX	1,915	2,722	5,159	6,547	7,315	8,083	8,271	14,053	20,881
UT	207	382	382	376	497	619	506	867	1,255
VT	–	–	–	1	2	2	1	2	2
VA	2,701	3,194	4,694	5,917	6,367	6,817	12,179	22,323	28,942
WA	561	605	1,027	1,014	1,163	1,285	1,119	1,298	1,467
WI	10	12	19	51	83	115	154	409	842
WY	212	212	212	305	341	376	463	773	1,096
U.S. Total	**10,967**	**16,537**	**25,547**	**35,074**	**39,774**	**44,473**	**56,057**	**95,862**	**132,499**

Figure 2. State-level data center nominal IT capacity (GW) today and in the planned 2030 pipeline. *Inner circles show capacity in 2021 (gray) and 2024 (blue). Outer band shows scenario range of projected capacity in 2030 (orange). Circle area is proportional to nominal IT capacity. Estimates include small- and large-scale DC as well as cryptocurrency mining. Results highlight continued concentration in established markets (e.g., Virginia, Texas) alongside emerging growth in new states as developers diversify geographically.*

ANNUAL AND PEAK ELECTRICITY USE: HISTORY AND FUTURE PROJECTIONS

While nominal data center capacity can be observed (subject to uncertainty in the data sources described above), there is currently very little public reporting of data center electricity use. To estimate annual and peak electricity use based on nominal DC capacity values, we must account for the factors described above: PUE, ramp-up of operational capacity, and peak and annual load factors.

We estimate that the average PUE across data centers in the U.S. was 1.32 in 2024, based on a capacity-weighted average of facility-level data. While earlier generations of smaller enterprise data centers had a PUE of 1.5 or higher, shifts toward better heat management, higher chip operating temperatures, increasing facility scale, and a recent shift towards liquid cooling have driven, and are expected to continue to drive, reductions in the fleet-average PUE over time. For example, large "hyperscale" data centers with liquid cooling currently under construction and coming online over the next few years could have a PUE as low as 1.1 (LBNL, 2024).

The extent to which nominal IT capacity is actually installed or operational at a given point in time is difficult to observe directly. We assume a time lag to reach this full nominal capacity based upon an observed lag between nominal IT capacity growth and maximum energy use inferred from state-level commercial and industrial energy data over the past several years. Specifically, we assume that when a facility comes online, only 20% of its full nominal capacity is active in the first year, ramping up by 20% each year until the facility is fully deployed.

Finally, we must estimate annual and hourly utilization of operational nameplate capacity. Public availability of temporally granular DC load data is limited (Masanet et al., 2024). For this study, EPRI collected facility-level data on load shapes from several anonymous facilities, which showed a peak utilization relative to "nameplate" (as defined by the facility) in the range of 62% to 80% (EPRI, 2025c). The collected profile for a large hyperscale facility (greater than 75 MW nominal capacity with a single tenant) exhibits higher utilization than respondents for smaller

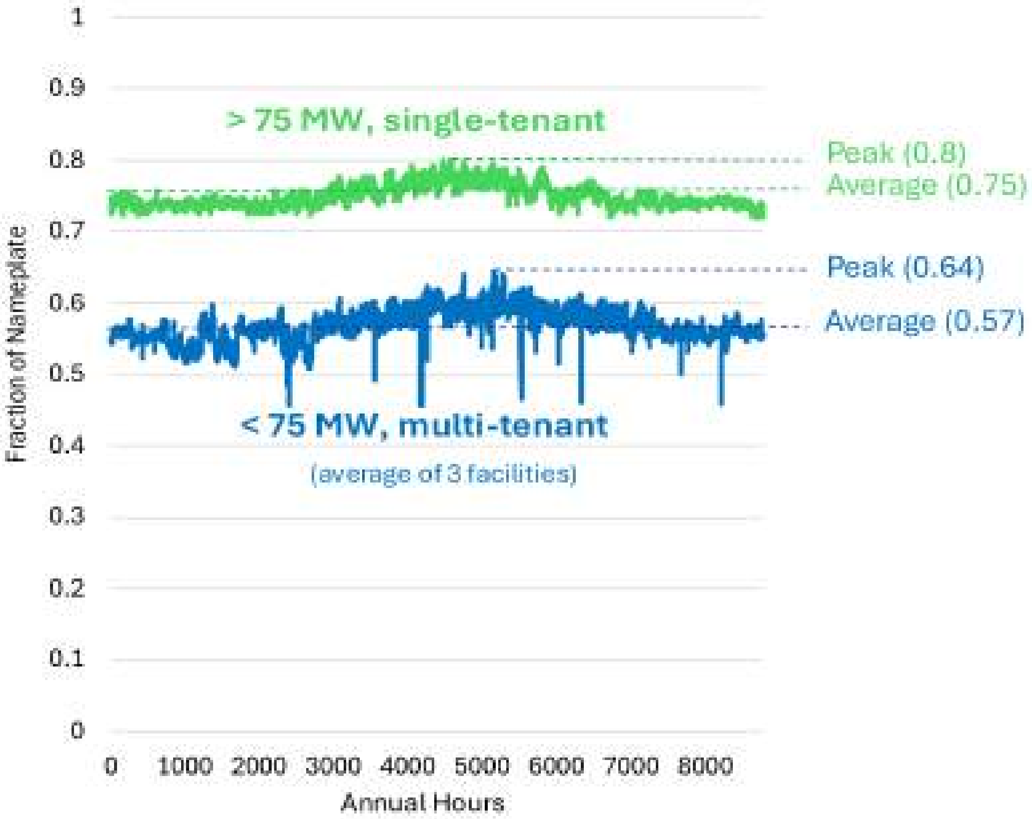

Figure 3. Hourly annual data center load shapes as fraction of nameplate capacity. *Observed profiles are based on EPRI facility data that show high, relatively flat utilization with higher sustained load for a large hyperscale facility than for smaller multi-tenant co-location facilities. Peak utilization is typically in the range of 62% to 80%, implying realized peak demand is lower than maximum/nameplate capacity.*

co-location facilities with multiple tenants (Figure 3). The average annual load factor relative to nameplate capacity was 75% for the hyperscale facility, while the average across the smaller co-location facilities was 57%. When measured against each facility's own realized peak demand (rather than nameplate), the load factor was higher, 94% for the hyperscale and 88% on average for co-location facilities.

Taken together, these data suggest that **aggregate DC load tends to be relatively constant through the year,**[14] **but that the realized peak is typically below the hypothetical maximum** or nameplate capacity values, a key consideration for power system planners translating announced MW into near-term peak impacts. Load shapes vary by workload and tenant mix, but the system planning implication is consistent: **peak load is often meaningfully below nameplate/ headline capacity.**

Applying assumptions about improving PUE, ramping of operational relative to nominal capacity, and utilization patterns based on observed load factors, we estimate that annual **electricity use over the past four years has grown more slowly than nominal capacity**. We estimate that total electricity use from data centers in the U.S. in 2024 was in the range of 177 to 192 TWh, or approximately twice the level of annual electricity use in 2021. Over the same period, nominal data center capacity is estimated to have increased by a factor of roughly 3–4. Although not directly observable, this estimated trend and the overall scale of data center electricity can be compared to observations of total electricity demand for validation, particularly in states where estimated data center power demands reflect a large share of the total.

---

[14] Despite relatively constant levels of output across minutes and hours, abrupt and large changes in load at second and sub-second timescales (that result from coordinated computing tasks stopping and starting) can have significant implications for operational reliability of the grid (for example, see EPRI, 2025d).

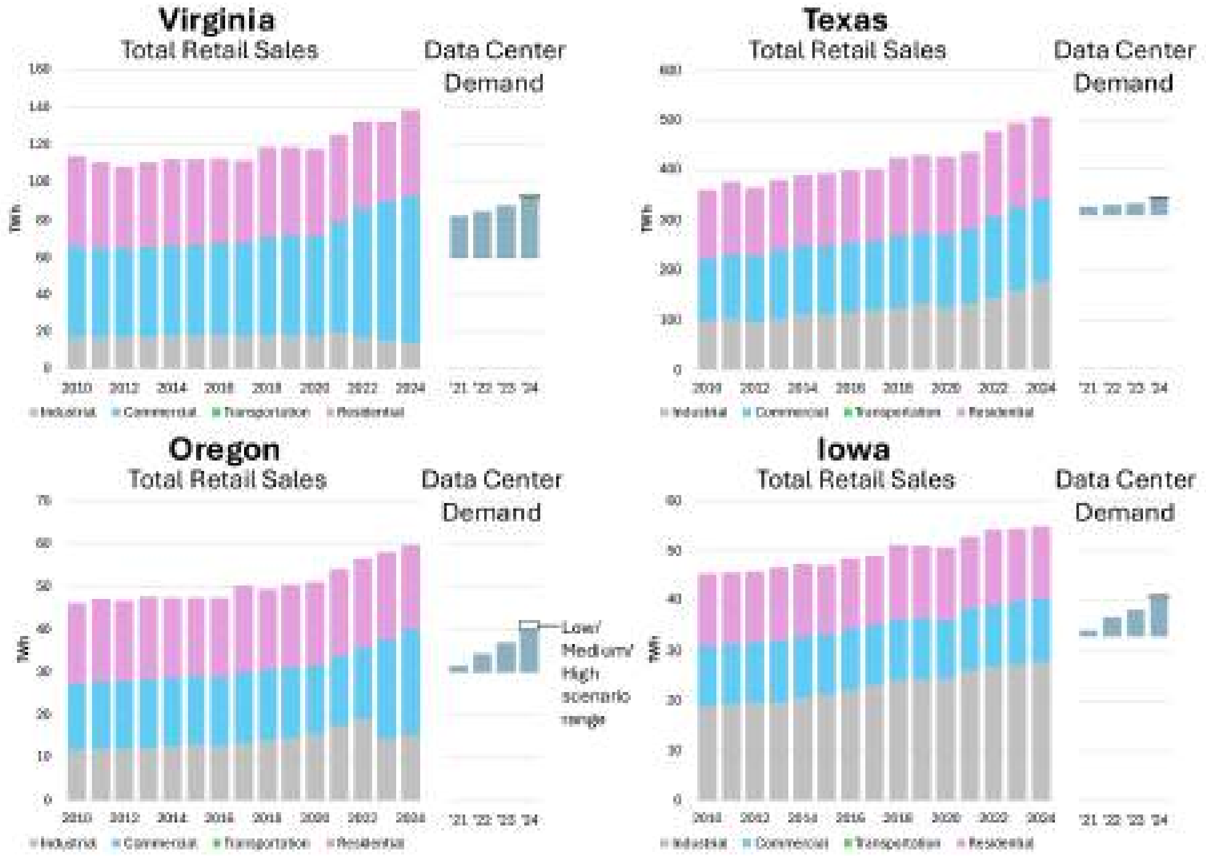

Figure 4. Comparison of modeled data center demand to total retail electricity sales by sector. *Retail sales by sector and state are based on EIA's State Energy Data System (SEDS) to show that the magnitude and growth of estimated DC load is broadly consistent with observed C&I totals, while highlighting ambiguity in whether DC load is reported as commercial versus industrial in some states. Data center demand is based on estimates in this analysis and shown aligned with top of industrial and commercial subtotal in 2024. Range in 2024 data center demand reflects data variation across sources, particularly in Oregon. Note axis scale varies by state.*

Comparing the modeled DC estimates to publicly reported commercial and industrial (C&I) sales data (from the U.S. Energy Information Administration's State Energy Data System dataset) provides an important cross-check. In many states, and at the national level, the scale of DC load is still too small to observe its signature separately from other trends (such as industrial migration or macro shocks such as the pandemic) and natural variation in weather-driven demand. But in states with high DC concentrations, the signature can be observed. Figure 4 shows our historical estimates of DC electricity use relative to total retail sales broken out by sector for four key states: Virginia, which has the highest concentration of DC electricity use in both absolute and relative terms (and has had a large DC market for over a decade); Texas, on par with Virginia in absolute scale of demand (when cryptocurrency mining is included) but lower in relative terms due to large industrial power demands in the state; and Oregon and Iowa, smaller in absolute terms but among the highest in terms of the DC share of electricity demand after Virginia.

DC load is sometimes reported by EIA as commercial demand and sometimes as industrial demand, creating ambiguity in interpreting the state data.[15] In addition, some re-classification over time is evident in Oregon and Virginia. In states like Texas, other drivers of industrial growth are certainly present, such as increased use of electricity in the upstream oil and gas sector, while in states like Iowa, DC growth may have been offset by declines in other industrial load. Still, the modeled DC demand trend is generally consistent with the trend in the C&I total sales in each state. **Indeed, this cross-check suggests that the modeled estimates of data center electricity use *could not* be *significantly higher* and still be consistent with observed total electricity sales.**

---

[15] The classification as commercial or industrial sometimes is dictated by the tariff class of the data center or terms of interconnection. Consequently, some states report both commercial and industrial data center loads.

Future data center electricity use was projected forward based on the nominal capacity projections and the assumed 20% operational ramp rate per year for new capacity. We also assumed continued improvements in PUE (which leads to lower energy use per nominal capacity) and a continued shift toward hyperscale and AI-driven workloads, which drive increased utilization and higher energy use per nominal capacity. Projected U.S. total data center energy use, shown in Figure 5, ranges from 383 TWh (Low) to 596 TWh (Medium) to 793 TWh (High) per year by 2030, an increase of between 2x and 4x from 2024. **This range is approximately 60% higher than the scenarios EPRI published in 2024.**[16]

Figure 5 also shows fleet-average utilization of nameplate data center capacity, which is estimated to have declined in recent years due to ramping of new additions. Going forward, recently added capacity begins to reach full utilization, while new additions continue to ramp gradually, with the net result that average utilization rises over time. In the Low growth scenario, average annual utilization of nominal capacity is higher than in the other scenarios because slower growth implies a lower share of new (and thus gradually ramping up) capacity in the total. However, utilization always remains below 100%.

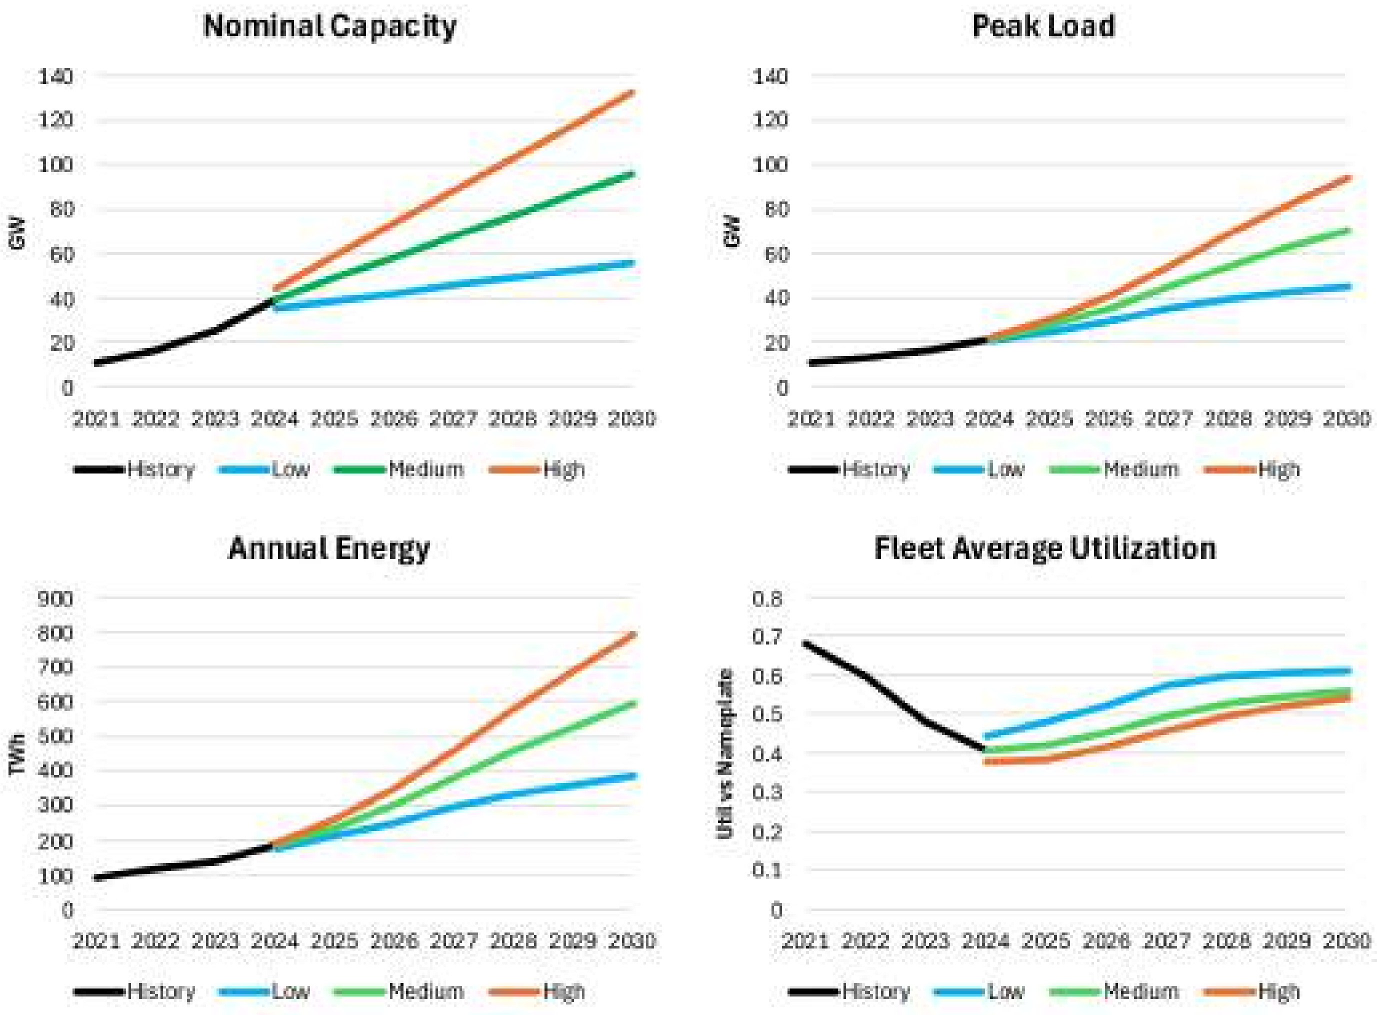

Figure 5. Estimated historical and projected energy and capacity metrics for U.S. data centers through 2030. *Historical estimates and Low/Medium/High projections show annual electricity consumption rising to 383-793 TWh by 2030, alongside implied capacity and fleet-average utilization. Results illustrate how ramp-up of new builds temporarily depresses average utilization and why realized peak impacts remain below nominal capacity. Estimates include small- and large-scale DC as well as cryptocurrency mining. Utilization refers to the average annual capacity factor relative to nameplate (or the product of nominal capacity and PUE).*

---

[16] See EPRI reports Powering Intelligence and Powering Data Centers.

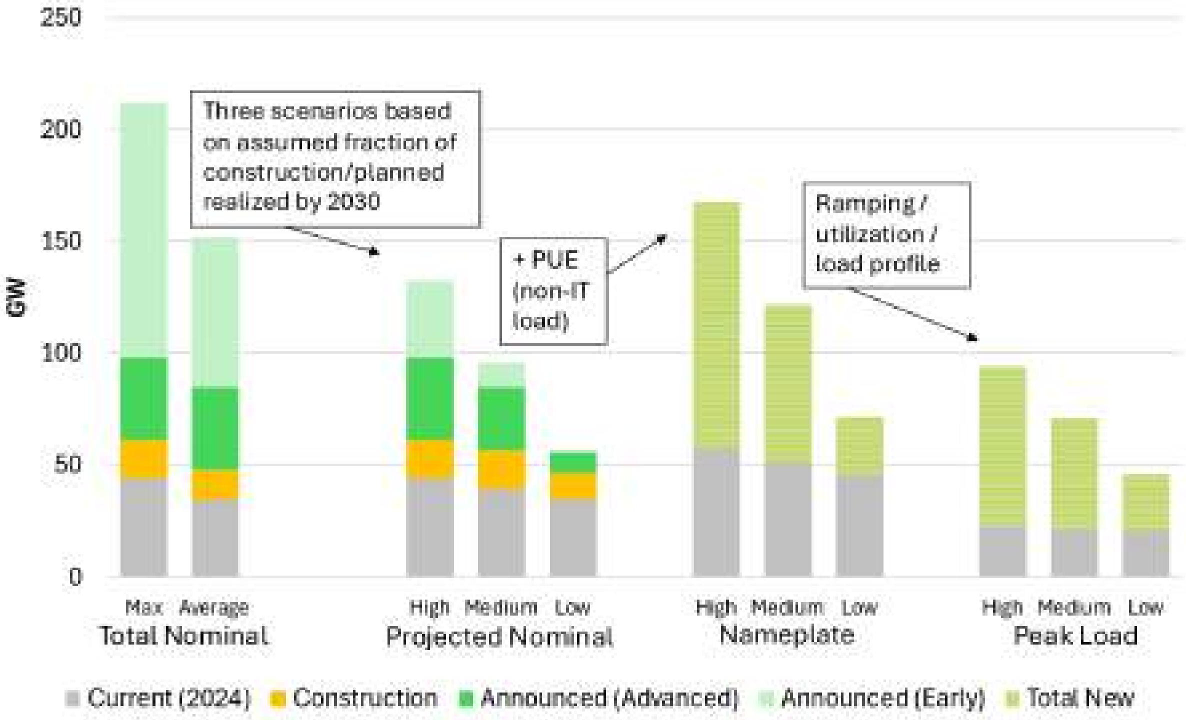

Figure 6. Planned project pipeline by stage translated into 2030 U.S. data center capacity. *Nominal IT capacity in Low/Medium/High scenarios broken out by stage, with total new implied nameplate capacity (IT x PUE) and realized peak load shown relative to current (2024) levels. Estimates include small- and large-scale DC as well as cryptocurrency mining. Scenarios based on EPRI analysis of a range of industry sources.*

Peak load impacts are also derived based on the application of data center load profiles across hours of the year described in EPRI (2025c) to the projections of nominal and active nameplate capacity. The non-coincident aggregate peak load for operating U.S. data centers (that is, the sum of calculated state-level data center peak loads, irrespective of whether those peaks occur at the same time) is estimated to be 21-22 GW in 2024, with projected growth to between 45 GW (Low), 71 GW (Medium), and 94 GW (High).[17] Because of the utilization factor less than one, the lag in utilization as new facilities ramp up, and declining PUE, peak data center load remains below nominal capacity values. State-level estimates of annual and peak electricity use are found in the online version of the report.

SUMMARY OF FUTURE PROJECTIONS

This section has described the relationship between three separate metrics of data center capacity in power terms (i.e. MW or GW):

- **Nominal IT capacity:** Reflects hypothetical maximum IT load (or data center computational capacity), announced projects subject to uncertain realization, and may be over-sized relative to actual live capacity as facilities ramp
- **Nameplate capacity:** Reflects hypothetical total load (or service requirement from utility), inclusive of cooling and other ancillary loads
- **Peak load:** Reflects actual peak power demand, lower than nameplate capacity due to ramping and utilization factors

Figure 6 summarizes these metrics at the U.S. level across the three scenarios, including the scenario assumptions about the completion rate of projects across the various planning stages. Understanding the differences and relative scaling between these metrics is important for translating the implications of DC project announcements into concrete implications for electric system planning.

---

[17] These estimates reflect projected peak data center load before any potential flexible demand response. EPRI's DCFlex initiative is demonstrating data center demand flexibility as a strategy to speed access to power, delay or reduce grid buildout, and improve grid reliability.

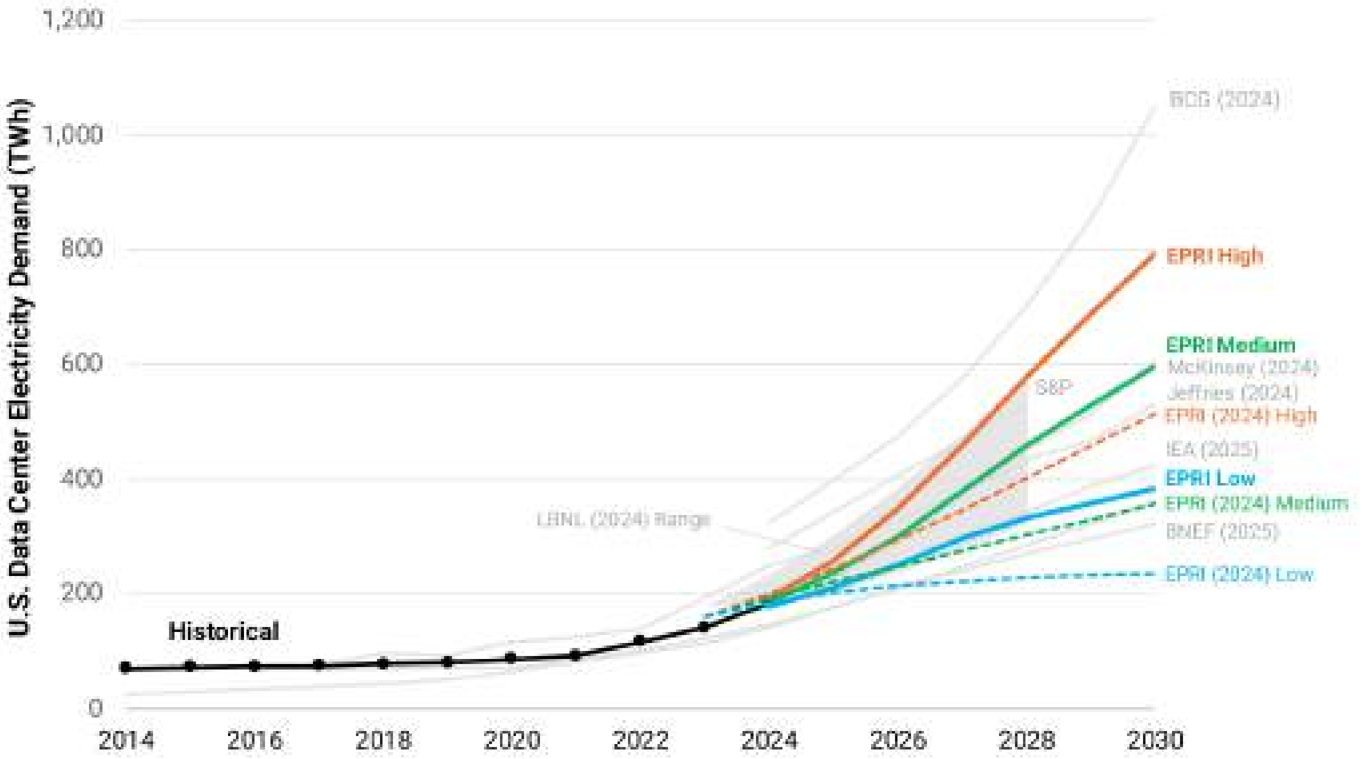

Figure 7. Comparison of U.S. data center annual electricity consumption projections. *Projections in this study span a similar range to the LBNL (2024) report. Shaded band shows scenarios from LBNL (2024); lines show recent external estimates, including BCG (2024), BloombergNEF (2025), EPRI (2024a), IEA (2025), Jefferies (2024), McKinsey (2024), S&P (2024). EPRI estimates include small- and large-scale data centers as well as cryptocurrency mining.*

Comparison to Other Studies and Forecasts

Several other studies have estimated total electricity demand from data centers (Figure 7). DC projections from this updated analysis generally span the range in the recent literature, including the LBNL (2024) report to Congress. This analysis utilizes state-specific project development data as its starting point because of EPRI's focus on electricity supply, where location is a primary issue. The LBNL report, in contrast, bases its estimates on chip and IT manufacturing projections, which provide useful bounds on US-level power demand that avoid the issues associated with using project announcements and estimating the power consumption of individual sites. The LBNL projections extend only to 2028, based on the time horizon of available chip forecasts when it was released.

These projections can also be compared to regional peak load forecasts reported by utilities and balancing areas via FERC Form 714. The projections in this analysis are state-based and reflect data center load only, while load forecasts reported to FERC are based on electricity planning areas and project changes in peak load across all categories. Figure 8 compares projected changes to peak demand between 2025 and 2030 for the five FERC regions with the largest growth forecasts with estimates from our analysis approximated by grouping states that align roughly with the regions. The FERC load growth forecasts, which aggregate electric company projections are generally consistent with the high end of the range analyzed here. They may also reflect different assumptions—potentially informed by project specific agreements—about the stage of projects, unannounced projects, and the translation of nominal or nameplate capacity to peak growth, for example they may assume higher utilization factors or faster ramping.

Additionally, because FERC and ISO forecasts are typically built up from individual utility forecasts, there is the possibility of double-counting data center projects that may have submitted multiple interconnection requests. In some regions, especially Texas/ERCOT and SWPP (which includes the Williston Basin oil and gas activity in North Dakota), the growth forecast also includes large new loads in other industrial segments.

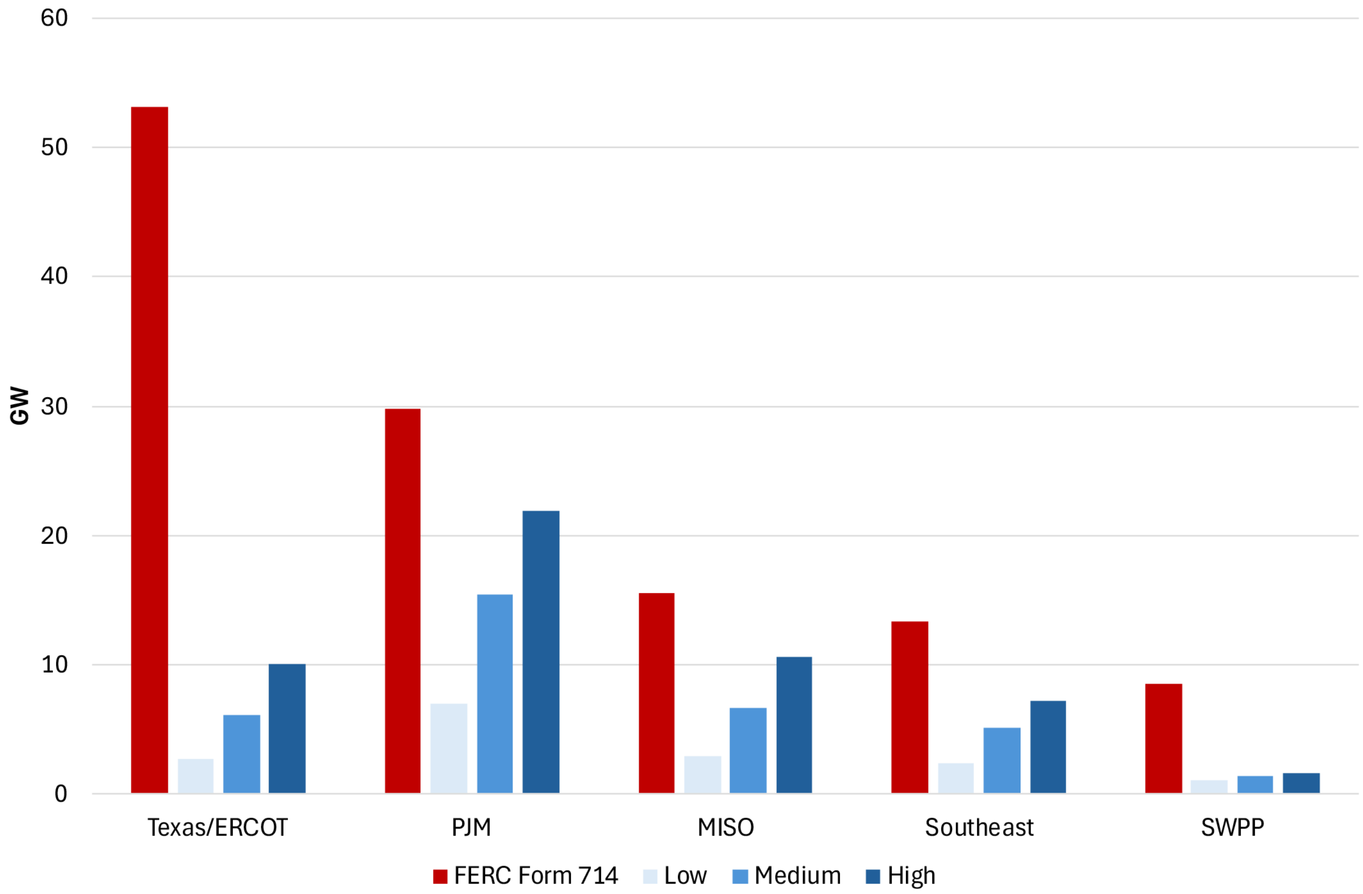

Figure 8. Peak demand growth (2025 to 2030) from planning forecasts vs. projected data center load growth scenarios in this analysis. *Total forecast peak load increases from FERC Form 714 are shown alongside DC peak increases under Low/Medium/High scenarios. Regional scenario results are approximated based on state groupings: PJM=VA,MD,DE,NJ,PA,WV,OH; MISO=IN,MI,IL,WI,MN,IA,MO,AR,LA; Southeast=NC,SC,AL,GA,MS; SWPP=OK,KS,NE,SD,ND.*

In the case of ERCOT, forecasts are required to consider load projections from individual transmission service providers (TSPs) per the HB 5066 legislation passed in 2023. In response, in the 2024 planning cycle, ERCOT accepted TSP "officer letter" loads without signed interconnection agreements when building its forecast. However, it also applied a downward adjustment to the TSP-aggregated forecast to account for potential double-counting and uncertain realization (the adjusted version is reflected in the FERC 714 data).[18]

ERCOT also provides a break-out of load growth into categories, including data centers and cryptocurrency. Even after the adjustment, forecasted load growth in these two categories between 2025 and 2030 accounts for nearly 30 GW of data center load growth, much higher than the scenarios modeled here. A possible explanation for this discrepancy is that ERCOT's forecasts assume much higher completion rates for announced projects than the scenarios described here; our analysis assumes uniform completion and ramp up rates of planned capacity across the country. Additionally, some developer load requests may not have been publicly announced and thus under-represented in the sources used here to estimate planned capacity.

---

[18] Information about ERCOT's load forecast can be found at the following links: https://www.ercot.com/files/docs/2025/03/03/ERCOT-Monthly-February2025.pdf; https://www.ercot.com/files/docs/2025/04/07/8.1-Long-Term-Load-Forecast-Update-2025-2031-and-Methodology-Changes.pdf.

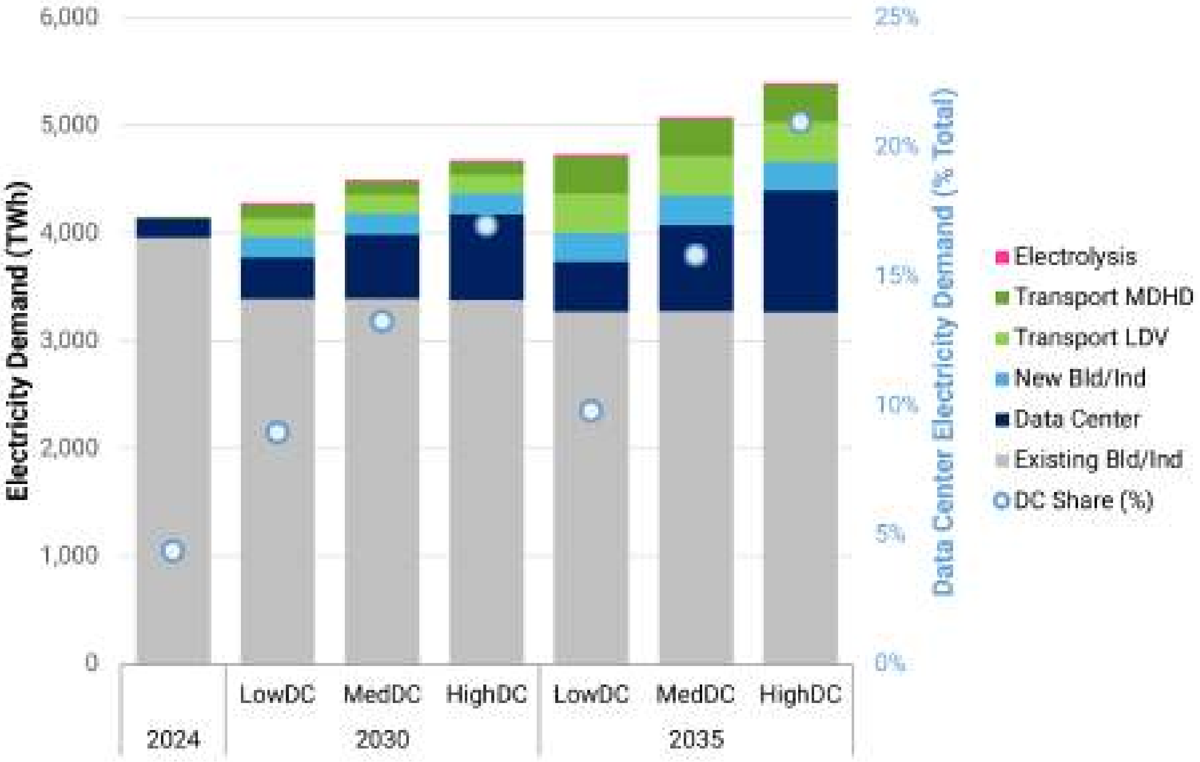

Figure 9. U.S. electricity demand by end use and data center share of total demand. *Values are shown for 2024, 2030, and 2035 across three data center scenarios. Outputs from EPRI's US-REGEN model under reference policies. The figure highlights how rapid DC growth can raise the DC share from today's 4-5% to 10-20% by 2035. MDHD = medium- and heavy-duty vehicles; LDV = light-duty vehicles.*

DATA CENTER LOAD GROWTH IN CONTEXT

Data Center Shares of Total Electricity Demand

Data centers reflect a large and growing share of total electricity demand. At the national level, **the current estimated share is 4–5%, but in several states, the share already exceeds 10%, with the highest concentration in Virginia, where the DC share exceeds 25%.** To understand how this share could evolve in the context of future projections, the DC load projections developed in this analysis were compared to total load projections from EPRI's US-REGEN energy systems model. After 2030, the DC load growth scenarios were extended by assuming a positive but declining growth rate through 2050.[19]

The results suggest **DC electricity use grows from about 4–5% of U.S. demand today to 9–17% by 2030 across scenarios, increasing further to 10–20% by 2035** (Figure 9). These DC demand shares exceed the projections in EPRI (2024b) for two reasons: higher estimated DC demand and lower projected non-data-center growth. The latter reflects updated policy assumptions, including reduced incentives for electrolytic hydrogen under revised Inflation Reduction Act guidance and slower transport electrification based on recent EV sales trends and more limited tax credits.[20] In the updated scenarios presented here, total U.S. electricity demand grows at an annual rate of 2% to 3.6% between 2025 and 2030, compared to a five-year average of 0.8% between 2019 and 2024.

---

[19] State-level annual growth rates modeled for 2025-2030 based on the project announcement methodology described here were assumed to decline by a factor of 4 for each subsequent 5-year period. With 3-5 year development times for data centers and a 3-year roadmap for server technologies, the current queue of announced projects provides a plausible basis for considering 2030 demand. However, the outlook beyond 2030 is highly uncertain given the pace of change in technology (hardware and algorithms) and the ultimate demands that society will place on data centers.

[20] See EPRI (2026) for analysis illustrating energy system impacts of changes to Inflation Reduction Act tax credits under the 2025 budget bill.

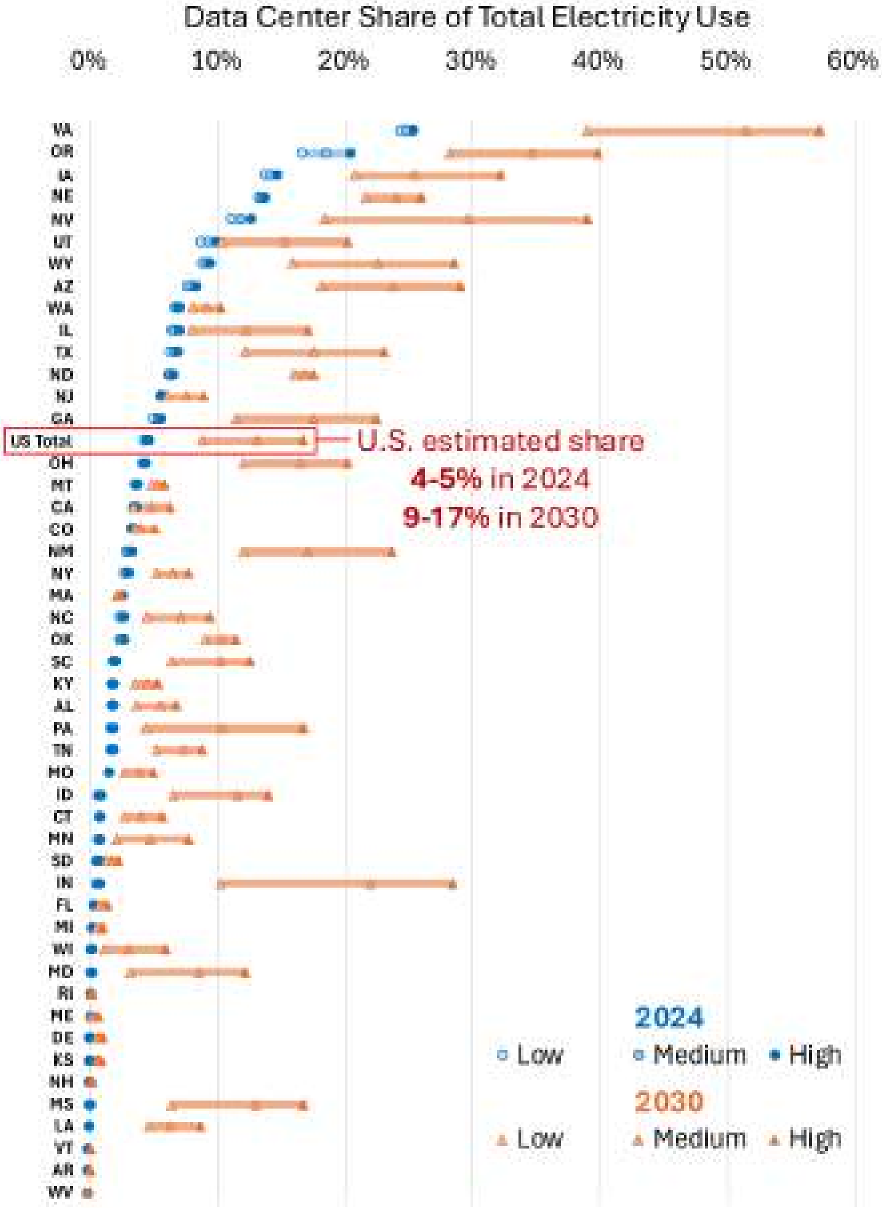

Figure 10. Data center share of total electricity demand by state. *Estimates include small- and large-scale DC as well as cryptocurrency mining. DC scenarios for 2030 are based on EPRI analysis of a range of industry sources. Non-DC electricity demand is based on US-REGEN model scenarios with current federal and state policies.*

At the state level, continued development of the largest DC market in Virginia implies a share increasing from around 25% to between 39 and 57% by 2030, reflecting the many projects currently under construction or advanced planning (Figure 10). Today, Virginia is the only state where data centers consume over 20% of electricity. By 2030, seven additional states—Oregon, Iowa, Nebraska, Nevada, Wyoming, Arizona, and Indiana—could see data centers exceeding a 20% share (Medium scenario). Other states, such as Washington and New Jersey, are above the U.S. average today but have relatively little estimated capacity in construction or planning (based upon accessed public sources) and hence relatively low increase in the data center share by 2030. On the other hand, several states with some existing capacity (e.g. New Mexico, Ohio, and Pennsylvania) and others with very little (e.g. Indiana, Louisiana, and Mississippi) are projected to emerge as new areas of concentrated development with 2030 shares exceeding 10%.

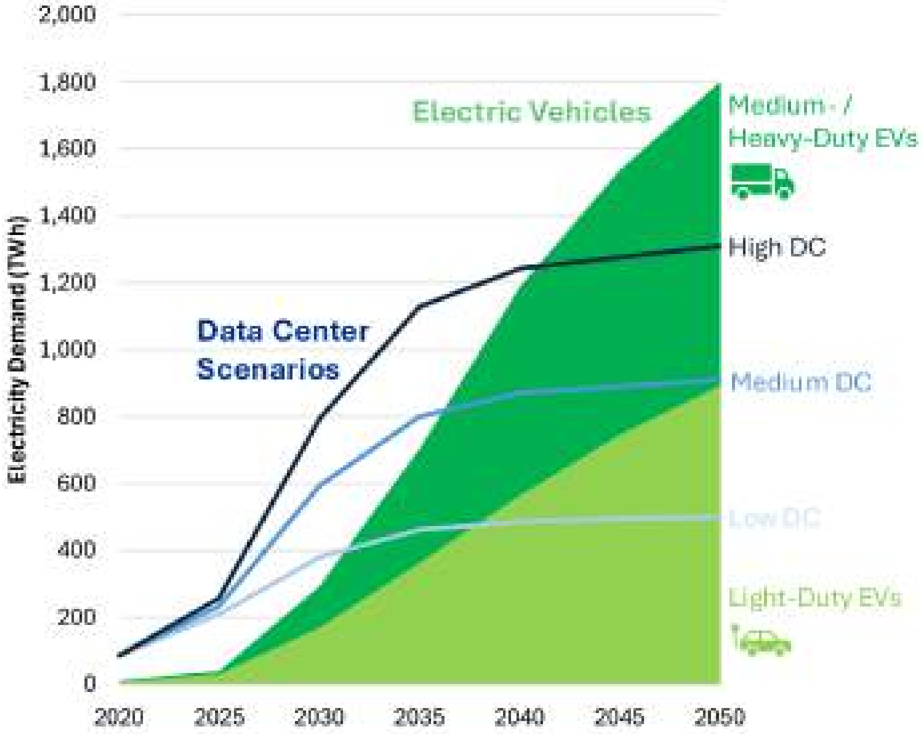

Figure 11. Reference electricity demand projections for data centers and transport electrification. *Electricity demand from light-duty vehicles and medium-/heavy-duty on-road vehicles are outputs from EPRI's US-REGEN model with reference policies. Under the reference trajectory, EV charging surpasses data center electricity use nationally in the mid- to late-2030s, though the cross-over can occur later (or not at all) in high data center states.*

Comparison to EV Charging Load

Data center load in 2024 is estimated to be 184 TWh nationally, while light-duty electric vehicle (EV) charging load in 2024 was around 12 TWh.[21] The updated DC projections in this analysis imply an annual growth rate in DC electricity use of 14–27% over six years. Meanwhile, EV adoption and charging electricity use is also growing, potentially at similar rates, but from a much lower starting point.

Both trends represent significant emerging loads, and modeling can help inform their relative magnitudes. However, substantial long-term uncertainty surrounds this comparison. The ultimate scale of AI and other data center service demands remains deeply uncertain, as well as the energy intensity of AI technologies (a function of both hardware and methodological aspects such as model scale and training approaches). Meanwhile, although the potential scope of future vehicle miles traveled in the future is better circumscribed, there is also uncertainty about the timing and extent EV adoption and efficiency improvements in vehicle technologies.[22]

In the US-REGEN modeling scenarios, 80–90% of on-road vehicle miles are projected to be electric by 2050 (including medium- and heavy-duty trucks), with reference assumptions for EV efficiency improvements. This high level of EV penetration reflects economic adoption by end-use consumers. Other projections include additional barriers to adoption, for example EPRI's eRoadMap scenario reaches 45% on-road electrification by 2050. Comparing US-REGEN's modeled EV load to the Medium data center scenario, EV charging represents a larger share of total electricity demand than data centers beginning around 2035, reaching around 1,800 TWh in 2050 vs. around 900 TWh for data centers (Figure 11). In the High DC growth scenario, EV load is still higher in the long run, but the crossing point occurs later, around 2040. With slower EV adoption (particularly for freight trucks) and/or more aggressive EV efficiency improvements, DC load could remain higher than EV load indefinitely.

---

[21] See EIA's Electric Power Monthly, Appendix D. https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=table_d_3.

[22] See EPRI's EVs2Scale 2030 Initiative as well as EPRI report *Valuing Improvements in EV Efficiency*.

The comparison between DC and EV load is perhaps more salient at the state and regional level. In 2035, U.S. total data center load in the Medium scenario is roughly equal to total EV load (based on economic adoption). However, the distribution across states is very different. Data center load, although diversifying geographically to some extent, has been shown to concentrate in certain regions. Vehicle charging load on the other hand is distributed more uniformly based on population and transportation corridors. The result is that in states with robust data center markets, including Virginia, Texas, Oregon, Iowa, Arizona, Georgia, as well emerging markets such as Indiana and Ohio, projected DC load in 2035 is significantly higher than EV load. In other states with relatively large population but more limited data center development, including California, Florida, New York, and New England, EV load significantly exceeds DC load by 2035. Moreover, EV charging load is distributed among many charge points ranging from a few kW to several MW, in contrast to a relatively small number of very large point-source data center loads on the scale of 100's of MW to GW.

GENERATION AND CAPACITY IMPACTS OF DATA CENTER LOAD

To assess potential electricity supply and emissions impacts of DC demand, we use EPRI's US-REGEN energy systems model to project generation and capacity responses under two policy scenarios.

- **Reference (Ref):** The first assumes a reference policy environment with on-the-books federal and state policies, including energy tax credit changes under the 2025 budget bill.
- **Carbon-Free Energy (CFE):** The second scenario assumes 24/7 carbon-free energy (CFE) targets for all DC loads by 2030.

These two cases are simulated across four scenarios of DC demand, including the Low, Medium, and High growth scenarios and a counterfactual that assumes no new DC demand after 2025.[23]

Under reference policies, higher DC demand leads to higher natural gas generation and capacity (Figure 12). Even with higher assumed capital costs for gas turbines, DC demand leads to increased investment in gas-fired resources over the next decade.[24] Projected annual natural gas capacity builds from 2025 to 2030 range from 6.6 to 13.7 GW per year with reference policies, compared to 3.3 GW per year with no new DC demand. The build rate in all scenarios is higher than the average of 5.7 GW per year over the past five years, although historical capacity additions have fluctuated within this range in recent decades.

These impacts represent large differences from the EPRI (2024b) analysis, in which the reference case included significantly higher wind and solar deployment to meet growing DC loads. Larger gas responses and more limited wind and solar deployment are primarily due to changes in Inflation Reduction Act production and investment tax credits, which lowered the competitiveness of wind and solar relative to gas-fired generation (EPRI, 2026; King, et al., 2025).

Under 24/7 CFE targets, generation and capacity responses to higher DC load come from wind, solar, nuclear, and energy storage. Nuclear and energy storage remain eligible for investment tax credits under current policy, increasing their value in CFE-constrained portfolios. Portfolios of incremental supply vary by region.

Under current policies, each additional MWh of data center demand carries an emissions intensity of 0.3 to 0.4 tCO_2, similar to the 2023 grid average. This reinforces the result that, without 24/7 CFE commitments, data center growth is met primarily by new and existing gas generation rather than through low-carbon additions.

Modeled results provide an indication of the scale of new resources needed to meet projected data center load growth. However, supply chain bottlenecks from equipment manufacturing to the permitting and siting process could constrain additions of both generation and transmission capacity. Although transmission within regions is not modeled explicitly in this analysis, interconnection of large new loads can require significant grid upgrades in some cases.

---

[23] For additional scenarios and discussion of supply-side responses, see "Supply-Side Technology Pathways to Meet Future U.S. Data Center Demand" (EPRI, 2025e).

[24] This paper is generally agnostic about ownership and location of generation and storage to meet data center demand. For a discussion of the trade-off of strategies ranging from a traditional passive grid connections to off-grid power, see EPRI's report: Reconciling the Value of Grid Interconnection and Speed to Power: Strategies for Powering Data Centers in the AI Era.

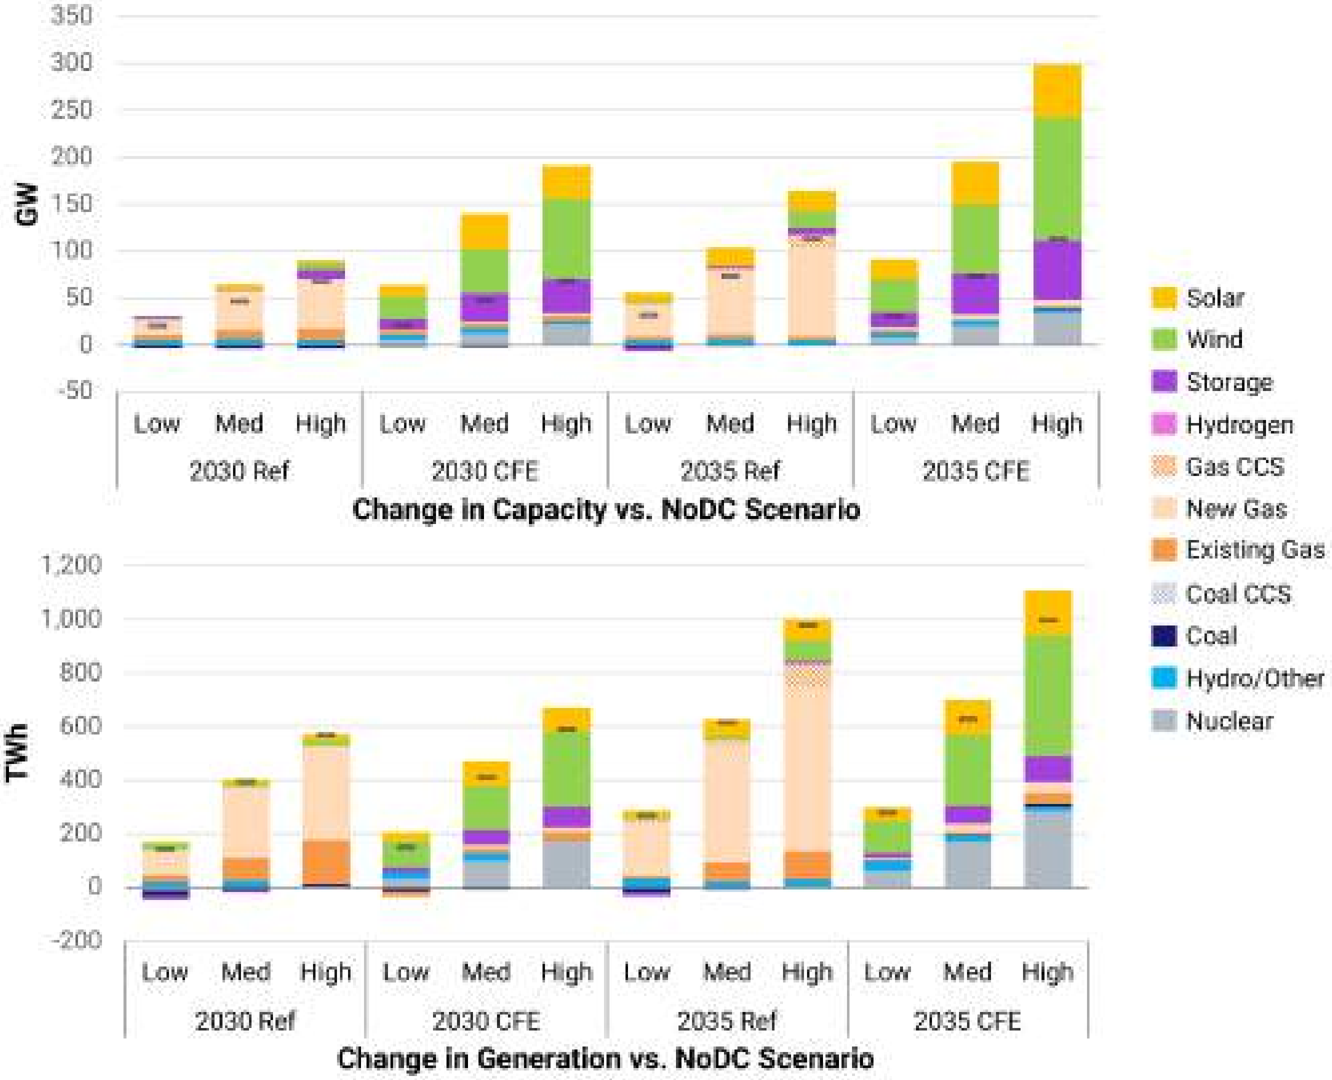

Figure 12. Change in U.S. capacity and generation (top and bottom panels, respectively) to meet new data center load. *Values are shown for 2030 and 2035 for scenarios with reference policies ("Ref") and 24/7 carbon-free energy targets ("CFE"). Under reference policies, incremental supply is dominated by natural gas, while 24/7 CFE targets shift additions toward low-emitting resources with the portfolio varying by region.*

CONCLUSIONS

Data center expansion of the scale projected here faces many challenges. At the local and regional level, challenges arise from the scale of the centers themselves and mismatches in infrastructure timing. A typical new data center of 100 to 1,000 megawatts represents a load equal to that of a new neighborhood of 80,000 to 800,000 average homes. While neighborhoods and the grid require many years to plan and build, data centers can be developed and connected in a few years. Added to the timing challenge are the supply chain issues associated with the scale of this growth. IT and power equipment and skilled labor are both regional and national-level challenges.

Better tools are essential to enable effective responses to short-term system disruptions and to inform long-term investments in the grid. This analysis replaces trend-based extrapolations with a project-pipeline approach and applies an explicit mapping from nominal IT capacity to annual electricity use and peak demand using assumptions about PUE, ramp-up, and utilization. The resulting 2030 range of data center electricity demand is higher than prior EPRI scenarios, which reflects the accelerated pace of data center development. Still, data gaps and fundamental uncertainties about future trends lead to a wide range of possible outcomes. With grid development times reaching years to decades, rapid advances in load forecasting are needed to guide efficient investments.

This work can inform planners and other stakeholders on several fronts. First, national and regional projections and integrated demand modeling can be used as a point of comparison with forecasts developed using other methodologies. Second, announced nominal capacity for data center projects should be treated as a pipeline indicator rather than a near-term peak forecast, as non-IT loads, ramp-up, and load shapes materially impact peak effects. Finally, system responses depend strongly on the policy environment and energy procurement objectives. Under reference policies, including updated federal tax credits and state-level policies, natural gas dominates incremental supply, while CFE commitments shift investment portfolios toward low-emitting generation and energy storage.

The analysis points to several areas for future work. First, these projections should be updated periodically as new DC project data become available using the repeatable methodology developed here, particularly for states such as Texas with large gaps between announced projects and utility forecasts. Second, modeling should be calibrated to DC load profiles and utilization patterns as more observations become available, particularly for emerging AI applications. Finally, modeling scenarios with alternate DC and supply assumptions should be conducted, including with alternate flexibility assumptions through EPRI’s DCFlex initiative.

Closer collaboration between data center developers, equipment providers, and power companies is essential to maintain and enhance grid reliability and to address affordability for all customers as data centers connect to the grid.

REFERENCES

BCG (2024), U.S. Data Center Power Outlook: Balancing Competing Power Consumption Needs (Vivian Lee, LinkedIn Pulse).

BloombergNEF (2025), U.S. Data Center Market Outlook.

EPRI (2026), *Energy System Implications of Market Drivers and State Policy.* EPRI, Palo Alto, CA: 2026. 3002034590.

EPRI (2025a), *Scaling Intelligence: The Exponential Growth of AI's Power Needs.* EPRI, Palo Alto, CA: 2025. 3002033669.

EPRI (2025b), *Speed to Power, Reconciling the Value of Grid Interconnection and Speed to Power: Strategies for Powering Data Centers in the AI Era.* EPRI, Palo Alto, CA: 2025. 3002034260.

EPRI (2025c), *Data Center Load Shape Library: 2025 Edition.* EPRI, Palo Alto, CA: 2025. 3002033424.

EPRI (2025d), *Regional Load Interconnection Requirements: United States and Europe.* EPRI, Palo Alto, CA: 2025. 3002033303.

EPRI (2025e), *Supply-Side Technology Pathways to Meet Future U.S. Data Center Demand.* EPRI, Palo Alto, CA: 2025. 3002034161.

EPRI (2024a), *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption.* EPRI, Palo Alto, CA: 2024. 3002028905.

EPRI (2024b), *Powering Data Centers: U.S. Energy System and Emissions Impacts of Growing Loads.* EPRI, Palo Alto, CA: 2024. 3002031198.

EPRI (2024c), *Valuing Improvements in Electric Vehicle Efficiency.* EPRI, Palo Alto, CA: 2024. 3002030215.

IEA (2025), Energy and AI. (Link)

Jefferies (2024), How Data Centers Are Shaping the Future of Energy Consumption. (Link)

JLL(2026). Global data center sector to nearly double to 200GW amid AI infrastructure boom. (Link)

King et al. (2025). What Passage of the "One Big Beautiful Bill" Means for US Energy and the Economy, Rhodium Group. (Link)

McKinsey (2024), How Data Centers and the Energy Sector Can Sate AI's Hunger for Power. (Link)

LBNL (2024), 2024 United States Data Center Energy Usage Report. (Link)

Masanet et al (2024), To Better Understand AI's Growing Energy Use, Analysts Need a Data Revolution, Joule. (Link)

S&P Global (2024), Data Centers: Rapid Growth Creates Opportunities and Issues. (Link)

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

THIS DOCUMENT WAS PREPARED BY THE ORGANIZATION(S) NAMED BELOW AS AN ACCOUNT OF WORK SPONSORED OR COSPONSORED BY THE ELECTRIC POWER RESEARCH INSTITUTE, INC. (EPRI). NEITHER EPRI, ANY MEMBER OF EPRI, ANY COSPONSOR, THE ORGANIZATION(S) BELOW, NOR ANY PERSON ACTING ON BEHALF OF ANY OF THEM:

(A) MAKES ANY WARRANTY OR REPRESENTATION WHATSOEVER, EXPRESS OR IMPLIED, (I) WITH RESPECT TO THE USE OF ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT, INCLUDING MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR (II) THAT SUCH USE DOES NOT INFRINGE ON OR INTERFERE WITH PRIVATELY OWNED RIGHTS, INCLUDING ANY PARTY'S INTELLECTUAL PROPERTY, OR (III) THAT THIS DOCUMENT IS SUITABLE TO ANY PARTICULAR USER'S CIRCUMSTANCE; OR

(B) ASSUMES RESPONSIBILITY FOR ANY DAMAGES OR OTHER LIABILITY WHATSOEVER (INCLUDING ANY CONSEQUENTIAL DAMAGES, EVEN IF EPRI OR ANY EPRI REPRESENTATIVE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES) RESULTING FROM YOUR SELECTION OR USE OF THIS DOCUMENT OR ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT.

REFERENCE HEREIN TO ANY SPECIFIC COMMERCIAL PRODUCT, PROCESS, OR SERVICE BY ITS TRADE NAME, TRADEMARK, MANUFACTURER, OR OTHERWISE, DOES NOT NECESSARILY CONSTITUTE OR IMPLY ITS ENDORSEMENT, RECOMMENDATION, OR FAVORING BY EPRI.

EPRI PREPARED THIS REPORT.

About EPRI

Founded in 1972, EPRI is the world's preeminent independent, nonprofit energy research and development organization, with offices around the world. EPRI's trusted experts collaborate with more than 450 companies in 45 countries, driving innovation to ensure the public has clean, safe, reliable, and affordable access to electricity across the globe. Together...Shaping the Future of Energy.®

PRINCIPAL INVESTIGATORS

GEOFF BLANFORD, *Principal Technical Executive*
650.855.2126, gblanford@epri.com

TOM WILSON, *Principal Technical Executive*
650.855.7928, twilson@epri.com

JOHN BISTLINE, *Senior Program Manager*
650.855.8517, jbistline@epri.com

NILS JOHNSON, *Principal Technical Leader*
650.855.7910, njohnson@epri.com

For more information, contact:

EPRI Customer Assistance Center
800.313.3774 • askepri@epri.com

3002034696 **February 2026**

EPRI
3420 Hillview Avenue, Palo Alto, California 94304-1338 USA • 650.855.2121 • www.epri.com

© 2026 Electric Power Research Institute (EPRI), Inc. All rights reserved. Electric Power Research Institute, EPRI, and TOGETHER...SHAPING THE FUTURE OF ENERGY are registered marks of the Electric Power Research Institute, Inc. in the U.S. and worldwide.