Data Centers in the AI Era

Course No: W03-002
Credit: 3 PDH

Ahmad Hammouz, P.Eng., LEED AP

CEDengineering.com
PDH FOR THE PROFESSIONAL

Continuing Education and Development, Inc.

P: (877) 322-5800
info@cedengineering.com

This course was adapted from the Global Efficiency Intelligence, LLC, " Data Centers in the AI Era: Energy and Emissions Impacts in the U.S. and Key States", which is in the public domain.

Executive Summary

The rising demand for data processing and storage, driven by advancements in AI and large-scale computing, has raised concerns about the energy use and emissions profiles of new data centers in the U.S. and globally. With data centers in the U.S. rapidly expanding, their energy and emissions footprint will largely depend on the technologies and pathways adopted. This report examines the geographical distribution, projected energy and emissions impacts, and state-level regulations affecting data center development, aiming to provide policymakers and stakeholders with insights to guide decisions on siting, construction, operations, and grid planning for low-carbon growth in computing demand.

The U.S. leads globally with over 3,000 data centers, representing 40% of the world's total, with top states like Virginia, California, and Texas housing over half of the nation's data centers. The trend in data center expansion is largely driven by factors such as land availability, access to affordable, reliable power, and proximity to end-users for low latency. Renewable energy access and incentives, including tax benefits, further influence data center siting, particularly as companies aim for 100% renewable energy.

We estimate data center electricity consumption through 2035 for the US and five key states where data center power demand is expected to grow: Virginia, Texas, Arizona, California, and Ohio. We estimate that by 2030, data center electricity demand could range from 345 to 490 TWh per year across efficiency and growth scenarios, reaching up to 655 TWh per year by 2035 under the highest growth and lowest efficiency scenario.

Much of the concern around data centers focuses on how much local electricity supply they will consume. We gathered data on projected electricity generation over time from the EIA for regional electricity markets and then used our estimates of electricity consumption by data centers over time to calculate the share of data center electricity use in total electricity generation. We find that data centers could make up 36-51% of Virginia's electricity consumption by 2030, or 28-54% by 2035, the largest share of total electricity demand out of the states studied. However, the other states and the US as a whole have lower overall shares of data center electricity demand, with Virginia representing a unique case where data centers are a dominant economic sector in the state. For the US as a whole, data center electricity use could increase from about 4% of the current total to 8-11% by 2030 and 8-15% by 2035, depending on growth rates and efficiency trajectories.

For the US as a whole, data center electricity use could increase from about 4% of the current total to 8-11% by 2030 and 8-15% by 2035, depending on growth rates and efficiency trajectories.

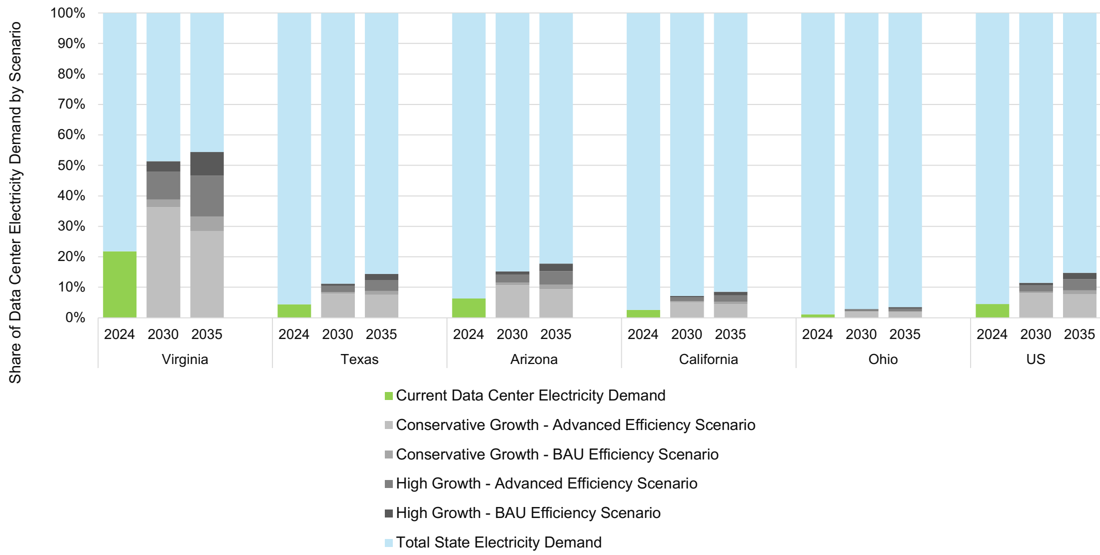

Figure ES1: Share of data center electricity demand out of total electricity demand, 2024-2035, all scenarios. Source: this study.

Note: the scenarios shown in gray are cumulative, i.e., the darkest gray bar for the High Growth-BAU Efficiency scenario represents the value from the top of that bar to the zero axis.

The emissions from data centers depend on a number of factors, modeled as scenarios in our study: computing demand growth, building efficiency (PUE), and procurement of renewable energy. Annual CO_2 emissions from data centers in the US are expected to rise through 2030. According to our analysis, emissions from data center electricity consumption could peak in 2030 at around between 63-83 Mt CO_2 per year. We estimate that about 25% of current electricity supply from data centers could come from directly procured renewable electricity. With a higher share and more rapid growth of RE for data centers, on top of ongoing grid decarbonization in the US, data centers could plateau and even begin declining in the near term, depending on growth scenarios, even as overall data center electricity consumption increases. Overall, our results show that increased RE procurement is a much larger driver of potential emissions reductions than increased efficiency (PUE).

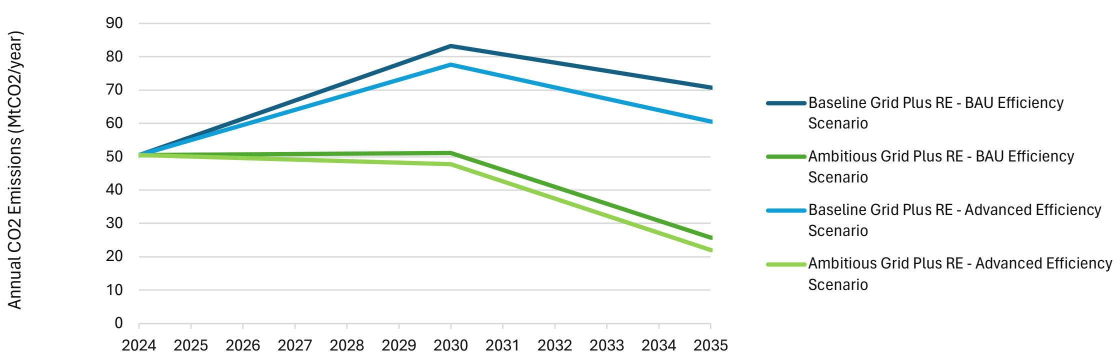

Figure ES2: Projected annual CO_2 emissions from data centers over time in the US, High Growth scenario

Based on our findings, we make policy recommendations for reducing emissions from data centers through energy system interventions across six key areas: renewable energy procurement, load management, grid optimization, enhanced state-level RE targets, improved energy management at data centers, and public data initiatives. Increasing renewable energy adoption can ensure that data center electricity consumption growth does not lead to an attendant increase in CO_2 emissions. Incentives for 24/7 clean energy tracking and long-term PPAs can help align data center energy needs with renewable energy supply. Load management improvements, including refined forecasting, demand response programs, and flexible scheduling, would help optimize grid resources. Grid infrastructure upgrades, especially expanded transmission networks and shared energy economy models, can help integrate renewables into data center operations. State-level ambitions should include data center-specific RE targets and incentives for high efficiency standards. Finally, public data initiatives can support transparency by tracking energy use and emissions benchmarks, which can help decision-makers understand and mitigate the emissions impacts of data centers.

Table of Contents

Introduction

As the demand for data processing and storage capacities increases, largely driven by the advancement of artificial intelligence and large-scale computing, there is a growing concern about the energy use and emissions profiles of new data centers in the United States and globally. As of 2021, data centers accounted for about 1% of global electricity demand (Digital Climate Alliance 2024), and up to 5% of electricity demand in the US, the world leader in data centers. The US has the largest number of data centers in the world, with nearly 3,000 data centers representing 40% of the world total number of data centers. The number of data centers in the US is growing rapidly, and thus the energy and emissions footprint of data centers hinge on pathways and technologies adopted in the US.

Overall electricity demand from data centers is projected by many experts to continue to grow significantly. The construction and operational practices of these data centers have significant implications for energy consumption and carbon dioxide (CO_2) emissions. This report explores the geographical distribution of these new facilities, their projected energy and emissions profiles, and state-level regulatory landscapes that could shape their development.

This study aims to equip policymakers and stakeholders with clear insights into the energy and emissions implications of new data centers driven by artificial intelligence and large-scale computing, facilitating informed decision-making regarding the location, construction, and operation of these facilities and the utility and grid planning.

1.1. Types of Data Centers

There are several key types of data centers as defined by their physical footprint and primary operator: hyperscale, colocation, enterprise, and edge data centers (in order of largest to smallest). Definitions vary, but we use these four broad categories to define data centers throughout this report. Note that there is some overlap between categories.

Hyperscale data centers are the largest-scale facilities designed to support scalable applications and are typically owned and operated by some of the major technology companies. They can support tens of thousands of servers, are extensively automated, and have high efficiency in power and cooling. They often provide cloud computing services and large-scale web applications. Examples of companies utilizing hyperscale data centers include Amazon Web Services (AWS), Google, Microsoft (Azure), and Meta.

Colocation data centers are facilities where businesses can rent space for their servers and other computing hardware. Customers share the data center infrastructure, such as power, cooling, and physical security. Colocation data centers vary in their scale, but supporting multiple businesses in a single data center can lead to demand for very large data centers. Major operators of colocation data centers include Equinix, Digital Realty, and CyrusOne.

Enterprise data centers are typically owned and operated by individual companies for their exclusive use. This gives companies full control over the infrastructure, security, and operations. These companies are typically large organizations with significant IT needs and the resources to manage their own facilities, although these facilities are typically smaller than colocation and hyperscale data centers.

Finally, **edge data centers** are smaller facilities located closer to end-users to provide faster data processing and reduced latency. They are typically near population centers or business hubs, and are often part of a larger network of edge locations. They can provide Internet of Things (IoT) applications, content delivery networks (CDNs), and faster data processing.

Given that these categories can have overlap, we present an estimate of the rough breakdown by total IT load for each data center type in the US. Colocation data centers make up the largest overall share of IT load, though hyperscale data centers (including some cloud providers and hyperscale operators that lease space within colocation facilities) have represented an increasing share of overall IT load in recent years.

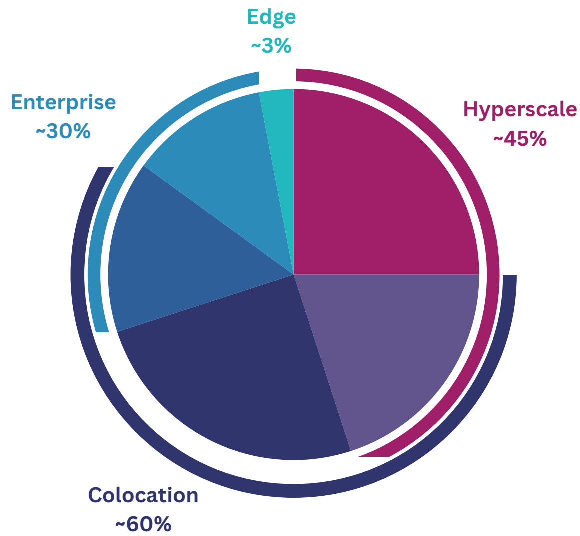

Figure 1: Estimated share of IT load by data center type.
Note: shares do not add up to 100% because there is overlap between data center types, and overlap is shown by the intersection of the outer concentric circles. (Sources: EPRI 2024, Cushman and Wakefield 2024)

1.2. Energy Use in a Data Center

The typical energy use breakdown of a data center can vary depending on the efficiency of the facility, the technologies used, and the operational practices in place. However, on average, the energy consumption of a data center is distributed across several key areas, with the IT load (i.e., computing processes) and cooling systems being the most significant contributors. As cooling systems grow more efficient, the IT load will take up a larger share of the total energy use.

Currently, the **IT load** consumes around 60% of energy at a data center, primarily to power servers and storage, but also networking equipment and other computing resources.

Cooling systems currently take up around 30% of energy at a data center. There are many possible types of cooling systems and cooling system components, which vary in their efficiency, including Computer Room Air Conditioners (CRAC), Computer Room Air Handlers (CRAH), chillers, cooling towers, fans, and liquid cooling systems. Cooling is essential to maintain the optimal operating temperature for IT equipment and prevent overheating, which can lead to failures and reduced lifespan of the equipment. For further description of cooling systems for data centers, see the Appendix.

Power infrastructure can consume up to 10% of total energy at a data center. This includes equipment to ensure Uninterruptible Power Supply (UPS), power distribution units (PDUs), transformers, and switchgear.

Finally, **lighting and other facility** systems can make up a small share of energy use at a data center. Given their large physical footprint, data centers need building lighting, security systems, and general building infrastructure.

The efficiency of a data center is often measured using the Power Usage Effectiveness (PUE) metric, which is calculated as:

PUE = **Total Facility Energy / IT Equipment Energy** (Equation 1)

The ideal theoretical PUE equals 1, indicating all energy is used by IT equipment with no additional overhead energy use. In practice, a PUE of 1.5 to 2.0 is fairly efficient for real-world data center. The ideal PUE can't be achieved in practice because some energy will always be required for essential non-IT functions, such as cooling, power conversion losses, and lighting, which are unavoidable even in the most optimized facilities. Data centers with a PUE above 2.0 are using a significant amount of energy for cooling and other non-IT infrastructure. The average PUE in data centers has declined from about 2.5 in 2007 to 1.58 in 2023 (EPRI 2024). Google has reported that it has achieved a PUE of 1.1 in its data centers as of 2024, down from 1.2 in 2010. Meta has reported a PUE of 1.08 as of 2023, down from 1.1 in 2019 (Meta 2024). Average PUE across multiple facilities is in part based on the type of facility, with data centers varying significantly based on their size, cooling system type, and the external climate.

Data centers rely on backup energy systems to maintain uninterrupted power in the case of an outage, with most facilities using diesel or natural gas generators. For computing equipment, UPS options may be installed that use batteries to provide immediate, short-term power during the transition to generators. Generators can supply more sustained power, typically ranging from 600 kW to 3,500 kW per generator (Bigelow 2024). While effective for emergency power, diesel generators can be significant sources of greenhouse gas emissions and local air pollution, especially when multiple generators are required for larger facilities.

This report assesses the energy use and emissions of future data centers that are likely to be built to meet computing demand in the US, largely driven by artificial intelligence and cloud computing. We first assess data center siting, exploring where data centers are currently located in the US, presenting data on planned data centers, and projecting where future data centers may be located. Using this information, we have developed a model for analyzing the energy and emissions profile of future data centers, projecting energy use and emissions through 2035. Next, we analyze the electricity system profiles of the states where future data centers may be located, focusing on the state of renewable energy development, case studies of data centers procuring renewable energy, and the regulatory landscape for data centers. Lastly, we discuss how technological innovation and energy efficiency may affect data center energy footprints, and provide policy recommendations for managing energy use and reducing emissions from data centers in the US.

Data Center Siting

2.1. Current Data Center Locations

The U.S. has the most data centers of any country in the world (DataCenterMap 2024), with approximately 3,000 data centers, or more than eight times the number of data centers than the country with the second-most data centers (the UK). Within the U.S., all 50 states have at least two data centers. The top ten states by number of data centers are displayed in Table 1 below. These ten states represent 63% of the total number of data centers in the U.S., and approximately over half of the total power demand. Virginia, California, and Texas are the top three states with the highest amount of data centers. In Virginia, data centers are primarily clustered in Northern Virginia. California data centers are largely located in Silicon Valley and Los Angeles. In Texas, data centers are sited near Dallas, Austin, Houston, and San Antonio. Some other city-level data center clusters within the top ten states include Columbus (Ohio), Chicago (Illinois), Miami (Florida), Portland and Hillsboro (Oregon), Seattle and Quincy (Washington), and Atlanta (Georgia) (Zhang 2024).

Table 1: Top ten states in the US by number of data centers

State	Number of Data Centers	Key Local Utilities	Major Data Center Operators
Virginia	473	Dominion Energy	Equinix, Digital Realty, AWS
California	278	PG&E, SCE	Equinix, CoreSite, Digital Realty, AWS, Google Cloud
Texas	278	Oncor Energy	Equinix, Digital Realty, IBM Cloud
Ohio	156	AEP	AWS, Google Cloud, Meta
Illinois	151	ComEd	Equinix, Digital Realty, Microsoft Azure
New York	128	ConEd, PSE&G	Equinix, DataBank, Digital Realty, QTS
Florida	123	FPL	Equinix, Digital Realty, AWS
Oregon	109	PacifiCorp	Digital Realty, NTT Global, AWS
Washington	89	Seattle City Light	Digital Realty, Sabey, Microsoft Azure
Georgia	84	Georgia Power	Equinix, QTS, Microsoft Azure

Sources: DataCenterMap 2024, Zhang 2024, Williams 2024, Cushman and Wakefield 2023

Note: Data is as of Q4 2024

2.2. Planned Data Center Locations

We surveyed several sources, including press releases, media articles, and analyst reports, in order to identify where planned data centers in the US will be built.[1] This is not an exhaustive list, but can indicate some trends on siting. We focused on the top states by current number of data centers. Table 2 presents summary data by state on the planned data centers we have identified.

1. The time frame of data collection was July through October 2024, recognizing that additional announcements have come forth since then. We used web searches and scientific literature database searchers, focusing on keywords of “data centers” and state names, plus “capacity”, “demand”, “growth”, “load”, “forecast”, “planned”, and other similar terms.

Table 2: Planned data centers by state

State	Number of Planned Data Centers	Estimated Total Land Area (million sq ft.)
California	19	5.5
Texas	19	12.8
Arizona	17	15.9
Georgia	17	18.1
Virginia	14	73.6
Ohio	13	6.2
Oregon	13	4.7
Illinois	12	6.2
New Jersey	4	0.6
New York	2	0.2
Idaho	1	1.0
Nevada	1	0.9
Total	132	146

Many of the data centers we identified had not yet announced their expected power demand or planned land area, so this could be an underestimate. At the same time, it is possible that many of these announced and planned data centers may not ultimately be constructed, or may represent multiple announcements for a single data center that is considering different sites.

Various sources have estimated that by 2030, data centers across the US could add an additional 20 GW of demand (Riu et al. 2024). The table above encompasses nearly 14 GW of demand across 132 planned data centers.

2.3. Projected Data Center Locations

In order to project where future, yet-to-be-announced data centers may be located, we first identify several key drivers of U.S. data center expansion, regional technological and growth patterns, and potential policy shifts that could affect data center siting.

With overall demand growth for data centers driven by cloud computing and AI, investors and data center operators seek locations based on three current key siting factors.

The first factor is the **availability of physical space**. Data centers can have a large physical footprint, and may need to plan for future expansions based on growth in demand. Gigawatt-scale data centers will require nearly a square kilometer of space, based on current technology. Data center operators are likely to site in areas with ample space and low land costs.

The second factor is the **availability of resources**, especially low-cost and reliable electric power. Data centers require significant amounts of electricity for their operations, and have historically operated at high capacity. Most data centers install redundant onsite backup generation, in addition to locating in areas with a reliable grid and affordable electricity. There are other types of resources that also constrain data center siting, including the availability of equipment and labor along the supply chain for data center construction and maintenance.

In addition, with the onset of climate-related extreme events, such as heat waves, fires, and storms, states with lower disaster risk or with strong infrastructure and resources for disaster prevention or relief may be preferred for future data center siting.

The third factor is **proximity to end users**. Locating close to customers can minimize latency and ensure rapid data transmission. In addition, locating near end users is related to the need to locate near existing fiber optic networks. Fiber optic cables are the primary medium for data transmission in and out of the data center.

Access to renewable energy resources will be a growing driver of data center siting. Renewable energy procurement is a major trend for hyperscale data center owners, with companies like Meta and Apple asserting that 100% of their energy requirements for data centers are met by renewable energy. Colocation data center operators, such as Equinix and Digital Realty, also procure renewable energy for their operations, with Digital Realty already meeting 100% of its energy requirements for data centers with RE, and others making commitments to do so. Data center companies are also investing in and developing RE generation projects. For example, Amazon Web Services (AWS) is planning to build two solar farms in Louisiana (Arizton 2023).

Alongside wind and solar energy, nuclear energy has emerged as a zero-carbon energy source of significant interest to tech companies and data center operators. Microsoft will procure nuclear energy from the Three Mile Island nuclear power plant in Pennsylvania via a 20-year PPA with Constellation Energy, aiming to recommission a unit that was retired in 2019 by 2028. AWS also signed a contract with Talen Energy for procuring a share of electricity from the Susquehanna nuclear power plant, also in Pennsylvania. Both arrangements will supply physically connected data centers in the PJM system (rather than a financial or virtual PPA purchasing electricity that goes into the larger grid) (U.S. EIA 2024). Data center operators are also showing interest in small modular reactors (SMRs), which are smaller and easier to deploy than traditional nuclear plants. For hyperscale data centers with constant electricity demand, nuclear power plants, which also operate continuously, can be an attractive source of stable zero-carbon electricity.

The demand for renewable energy will lead data center developers to locate in areas where electric utilities can commit to 100% RE provision, areas with robust RE PPA markets, and areas where companies can invest in their own onsite or nearby RE generation projects. In addition, data centers may also be sited in areas where there are renewable energy companies that can exclusively supply RE for data centers. Such companies include Dominion Energy, Avangrid, EDF Renewables, and ENEL Group.

The growth in demand for RE-powered data centers is also increasing demand for improved power infrastructure that can accommodate and supplement variable renewable energy, such as microgrids, smart grids, and energy storage. Companies like EdgeCore and Microsoft are adopting grid-interactive lithium ion batteries to provide uninterruptible power supply (UPS) systems, and other non-traditional types of batteries are also being explored (Arizton 2023). Thus, areas where suppliers of such technologies are established will also be attractive areas for building new data centers.

Potential shifts in policy landscapes could significantly influence data center site selection by emphasizing sustainability, providing incentives, and/or requiring regulatory compliance. A number of states have provided some form of tax incentives for data centers in order to attract investment. Of the top ten states by number of data centers, Florida, Georgia, Illinois, New York, Ohio, Texas, Virginia, and Washington have offered tax exemptions on equipment purchase for qualifying data centers (thresholds vary by state in terms of annual investment amount, required job creation,[2] and other criteria) (Stream Data Centers 2024). Florida, Illinois, and Texas also provide a sales and use tax exemption on electricity, while Georgia, Illinois, and Washington provide sales and use tax exemptions for energy-related infrastructure. Oregon has no sales tax, and there are local enterprise zones (such as Hillsborough, a major area for data centers) that have property tax exemptions; the state also provides property tax exemptions for energy efficiency infrastructure. These offer significant logistical and financial benefits for data centers.

Based on these drivers, current and planned data center locations that we have identified, and projected trends, we project large data center growth in five key regions: Virginia (Northern Virginia), Arizona (Phoenix area), Texas (DFW Area), California (Silicon Valley), and Ohio (Columbus area). Other regions will also experience growth, but we focus on these as some of the fastest growing areas for data centers, and that can have large impacts on the energy and emissions footprint of data centers depending on operational choices. Key drivers for these areas are summarized below.

Virginia, especially Loudoun and Prince William counties in Northern Virginia, is the leading location for data center development globally. This area already hosts a substantial portion of the world's data centers and is sometimes called "data center alley." Future data centers are likely to continue to site in Virginia due to its extensive fiber network and proximity to key government and business hubs. In addition, there are a variety of favorable tax incentives for data center operators.

Arizona, especially the Greater Phoenix Area, is rapidly emerging as a major data center hub in the U.S. Currently ranked eleventh among U.S. states for by number of data centers, Arizona is attractive to data center operators in part due to its climate, which is conducive to more energy efficiency free cooling, which can lower the PUE of a data center. Additionally, the state has a low natural disaster risk, minimizing potential operational disruptions. Arizona has strong RE potential and growing RE generation, an increasing priority for various companies operating data centers. The area's tax incentives and expanding infrastructure also make Phoenix an attractive location for data center growth.

Texas, particularly in the Dallas-Fort Worth metropolitan area, is centrally located to facilitate efficient data transfer across the U.S. DFW has strong connectivity, a business-friendly regulatory environment, and competitive tax incentives that appeal to data center operators. Texas also has a significant and growing renewable energy portfolio, especially wind energy. Finally, Texas has abundant and affordable land, an important resource for the large footprint of potential future data centers.

California's Silicon Valley is likely to remain a key region for data center growth due to its proximity to major tech companies and dense clusters of end users. Silicon Valley also provides direct access to advanced technology, innovation ecosystems, and a skilled workforce specialized in high-tech industries. Although costs are higher in California compared to other states, especially electricity prices, the demand for colocation and cloud services in this area is substantial.

2. Data centers are not highly labor intensive, with states requiring the creation of 5-50 jobs per data center (varying by state) based on $5-$150 million in investment in order to receive tax incentives (Wright et al. 2024)

Ohio, particularly around Columbus, is becoming a significant destination for data centers, offering access to major markets across the Midwest and beyond. Columbus has a well-developed fiber optic network and Ohio provides significant tax incentives and business-friendly policies specifically aimed at fostering data center development (see Section 4.5. for more recent developments). Columbus also has a growing tech ecosystem, supported by local universities and research institutions that contribute to a skilled workforce and collaborative environment for innovation.

Energy and Emissions Profiles of Emerging U.S. Data Centers

In this section, we project the energy use and emissions associated with US data centers over the next decade (though 2035), based on our assessment of the five key states listed above (Virginia, Arizona, California, Texas, and Ohio). We also estimated the energy use and emissions associated with data centers across the entire US.

3.1. Methodology

The magnitude of future growth in energy demand from data centers is highly uncertain. Growth estimates from before the advent of AI predicted that overall energy consumption from data centers in the U.S. would be flat or even decrease due to increased efficiency in computation, networks, infrastructure, and power scaling (Shehabi et al. 2016). However, with the rise of AI and cloud computing, current estimates project significant growth.

Our approach to projecting the energy and emissions impacts of data centers falls into the “extrapolation” approach, as defined by Shehabi et al. 2024. This approach entails using existing estimates to construct a base year, and then model projected growth (Shehabi et al. 2024). We link this approach to a state-level electricity mix and emissions factor model to estimate the emissions associated with data center electricity use for the states of interest. Contributions of our study include a longer time horizon to examine scenarios (through 2035), state-level estimates, a focus on emissions and electricity scenarios to understand decarbonization levers, and qualitative assessment of state-level policy and infrastructure landscapes. We recognize that our study makes necessary simplifications in order to examine the effects of different scenarios.

In Chapter 2, we estimated the number and capacity of current, planned, and projected data centers through online searches. The time frame of data collection was July through October 2024, recognizing that additional announcements have come forth since then. We used web searches and scientific literature database searchers, focusing on keywords of “data centers” and state names, plus “capacity,” “demand,” “growth,” “load,” “forecast,” “planned,” and other similar terms. Our results are not exhaustive, but combine estimates from the highest quality sources. The capacity of current data centers was used for the base year of our model, and we made growth projections to estimate future IT load and power demand.

A review of thirteen different studies found that projections of additional power demand through 2030 range from roughly 5 to 50 GW, with an average prediction of nearly 20 GW (Riu et al. 2024). Estimates of compound annual growth rates (CAGR) range from about 4% to 21% (Figure 2). This range of load growth corresponds to an additional electricity generation capacity of 9 to 100 GW, assuming demand is met with 75% renewables (Riu et al. 2024).

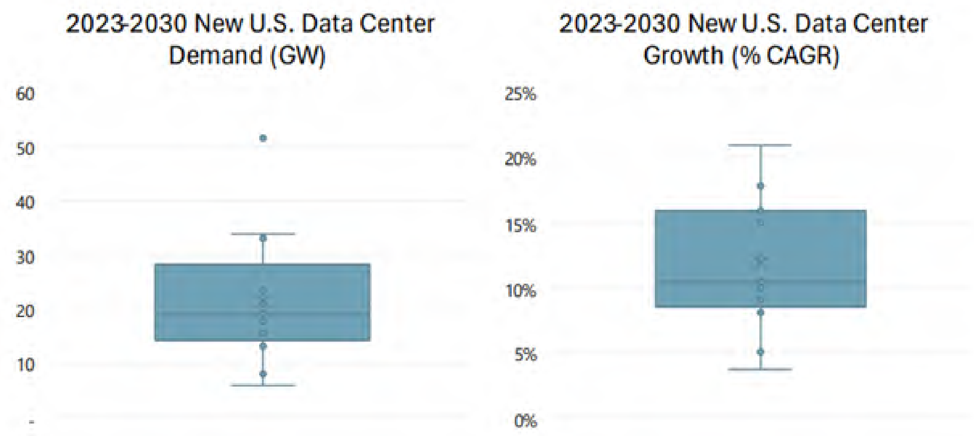

Figure 2: Review of study projections on total U.S. data center IT load growth from 2023 to 2030 (left) and projected compound annual growth rates (CAGR) (right) (Source: Riu et al. 2024)

In order to account for this uncertainty, we developed two growth scenarios: a Conservative Growth scenario and a High Growth scenario. These scenarios apply a CAGR to IT load in order to estimate the IT load over time. Table 3 shows the assumed values for these scenarios, with growth rates from 2030 onwards tapering off due to advances in computing efficiency and/or shifts in usage and demand.

Table 3: Growth scenarios used in this study

Scenario	Compound Annual Growth Rate in IT Load	
	2024-2030	2030-2035
Conservative Growth	5%	3%
High Growth	10%	7.5%

To estimate total electricity demand from data centers, we applied the PUE to IT load estimate the additional electricity demand from cooling systems and other equipment at data centers. We assumed that the average PUE of data centers in each state and across the US would improve over time, as advancements in cooling technology and other innovations lower the need for non-computing energy, and larger-scale and more efficient data centers make up a larger share of the national fleet. We developed two PUE scenarios, termed the Business-As-Usual (BAU) Efficiency scenario based on slower improvements, and the Advanced Efficiency scenario with more rapid improvement (Table 4), based on the PUE trends discussed in Section 1.2. We did not model the effects of backup power system, such as fuel generators, because this would require more granular facility-level data.

Table 4: Average PUE of data centers in this study across scenarios

	Average PUE		
	2024	2030	2035
BAU Efficiency	1.58	1.5	1.4
Advanced Efficiency	1.58	1.4	1.2

To estimate emissions from data center electricity use, we collected information on U.S. and state-level grid electricity emissions factors and projected their trajectories over time. We used the current emissions factor of each state's grid and projected linear decarbonization to either the state's stated net zero target, or to zero by 2050 if the state did not have an earlier net zero target. Virginia and California aim to achieve net zero by 2045, while Arizona, Texas, and Ohio do not have net zero targets. We also identified the electricity generation mix for each state, using this to estimate the share of fossil fuel-based generation (including coal, natural gas, and petroleum) vs. non-fossil fuel-based generation (including geothermal, hydropower, nuclear, solar thermal and PV, biomass, wind, and wood-based fuels).

Table 5: Electricity emissions profile of the analyzed states

State	Grid Electricity Emissions Factor (tCO_2/MWh)	Share of Non-Fossil Generation
Virginia	0.261	41%
California	0.179	52%
Texas	0.374	34%
Ohio	0.543	16%
Arizona	0.316	44%

Source: EIA 2024

A number of companies and data center operators aim to directly procure renewable electricity (RE) to power data centers via mechanisms such as Power Purchase Agreements (PPAs) or onsite generation of RE. This is a stated goal of many companies, with Apple and Meta having already achieved 100% RE, and several other companies aiming for this goal by 2025 (e.g., Microsoft, Amazon Web Services). While these high-profile companies have made significant advancements, there are many other data center operators that have later commitments or have not yet made commitments. To reflect these trends and possibilities, we developed two electricity emissions scenarios that we used to develop emissions factor trajectories over time.

The Baseline Grid Plus RE scenario that assumes that data centers will procure a growing share of their electricity needs from various renewable energy sources over time (with the rest of the electricity supply coming from the grid), and the Ambitious Grid Plus RE scenario assumes more rapid growth in procurement of RE. We estimate that currently about 25% of data center electricity consumption across the use comes from direct RE procurement, based on the share of data center electricity consumption occurring in hyperscale facilities run by large companies with ambitious RE procurement targets (see Section 1.1.).

Table 6: Electricity emissions scenarios in this study

	Share of RE in Electricity Consumption		
	2024	2030	2035
Baseline Grid Plus RE	25%	35%	45%
Ambitious Grid Plus RE	25%	60%	80%

3.2. Projected Electricity Consumption from Data Centers in the U.S. and Key States

We estimate that electricity consumption from data centers in the US will increase over time, driven by demand for computing from AI and cloud computing. After rapid growth through the end of the decade, we estimate that overall demand for computing and electricity consumption will grow more slowly, nearly plateauing as overall demand slows and efficiency increases. We estimate that by 2030, data center electricity demand could range from 345 to 490 TWh per year across efficiency and growth scenarios, reaching up to 655 TWh per year by 2035 under the highest growth and lowest efficiency scenario. Achievement of a very advanced fleet-wide PUE of 1.2 under the Advanced Efficiency scenario could lower electricity consumption by 14% relative to the BAU efficiency scenarios.

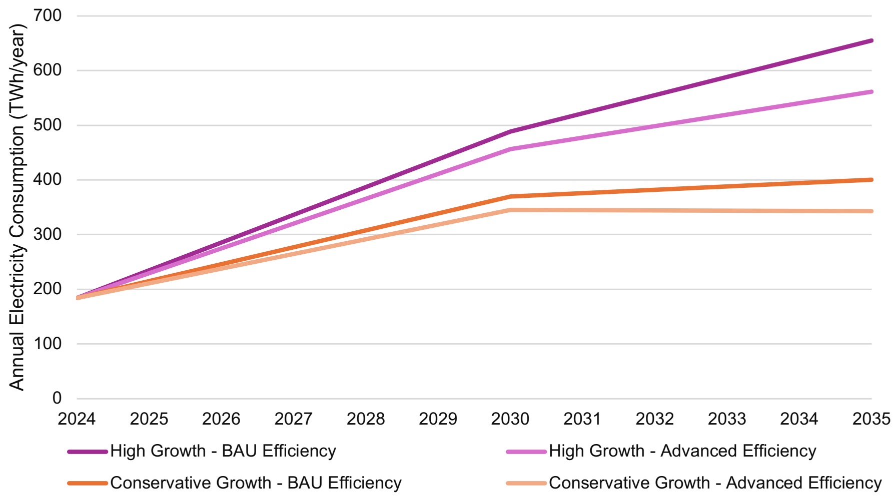

Figure 3: Projected annual electricity consumption from data centers in the US across scenarios, 2024-2035

Virginia has the highest expected electricity consumption of all states, followed by Texas. Figure 4 shows the range in estimated electricity consumption across growth and efficiency scenarios, with

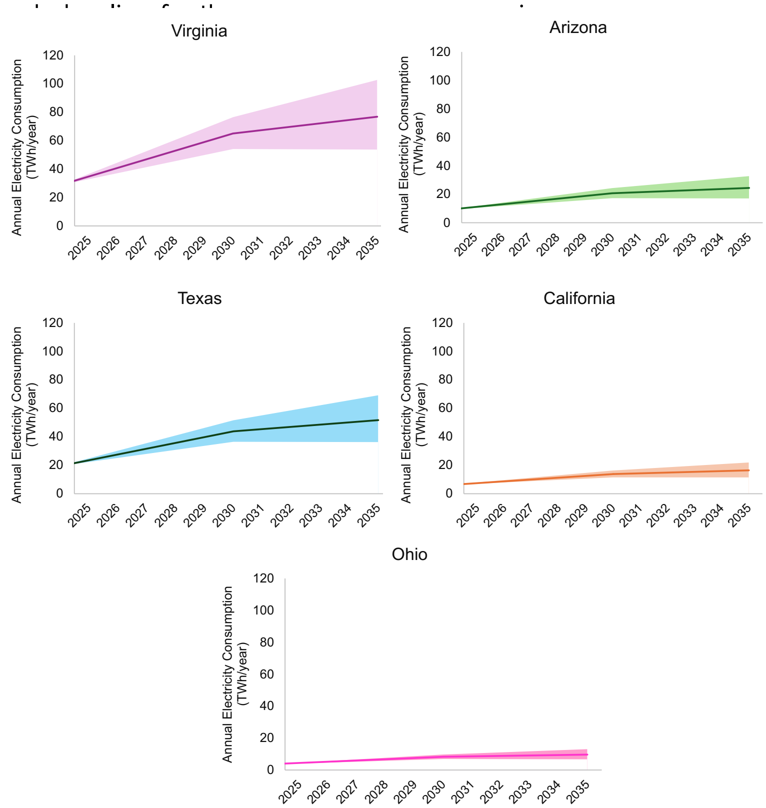

Figure 4: Estimated electricity consumption by data centers in five key states, 2024-2035.
Note: the top border of the shaded area represents the High Growth-BAU Efficiency scenario, while the bottom border represents the Conservative Growth-Advanced Efficiency scenario.

Much of the concern around data centers focuses on how much local electricity supply they will consume. Large estimates of data center energy consumption are taken into account in utility resource planning, which has the potential to drive overinvestment in fossil-fuel based power plants. We assessed the share of data center electricity use in total state electricity use over time. We gathered data on projected electricity generation from the EIA for regional electricity markets and then used our estimates of electricity consumption by data centers over time to calculate the share of data center electricity use in total electricity generation for the five states and the US. We find that data centers could make up 36-51% of Virginia's electricity consumption by 2030, or 28-54% by 2035, the largest share of total electricity demand out of the states studied. However, the other states and the US as a whole have lower overall shares of data center electricity demand, with Virginia representing a unique case where data centers are a dominant economic sector in the state. For the US as a whole, data center electricity use could increase from about 4% of the current total to 8-11% by 2030 and 8-15% by 2035, depending on growth rates and efficiency trajectories. Ohio has the lowest share of data center electricity demand out of the state total, of the states studied.

For the US as a whole, data center electricity use could increase from about 4% of the current total to 8-11% by 2030 and 8-15% by 2035, depending on growth rates and efficiency trajectories.

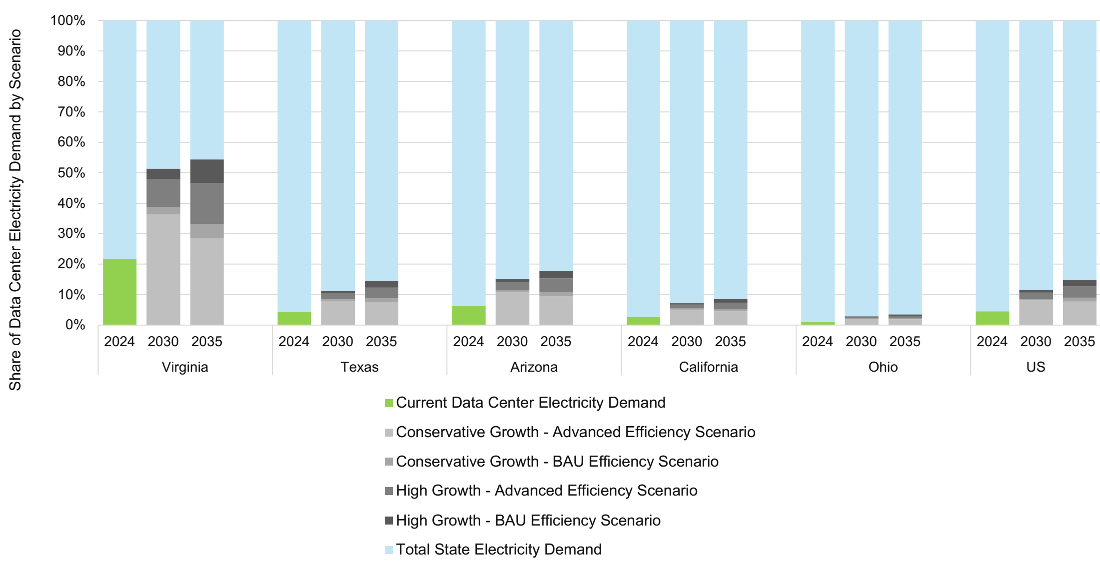

Figure 5: Share of data center electricity demand out of total electricity demand, 2024-2035, all scenarios.
Note: the scenarios shown in gray are cumulative, i.e., the darkest gray bar for the High Growth-BAU Efficiency scenario represents the value from the top of that bar to the zero axis.

3.3. Projected CO_2 Emissions from Data Centers in the U.S. and Key States

Annual CO_2 emissions from data centers in the US are expected to rise through 2030 under the Baseline Grid Plus RE scenario, driven by growth in computing demand across both growth scenarios. Under the High Growth scenario, emissions from data center electricity consumption peak in 2030 at around 83 Mt CO_2 per year (Baseline Grid Plus RE-BAU Efficiency Scenario) (Figure 7), while under the Conservative Growth scenario, emissions peak at 63 Mt CO_2 per year (Figure 6). However, under the Ambitious Grid Plus RE scenario, with more rapid integration of renewable electricity into data centers' electricity supply on top of ongoing grid decarbonization in the US, emissions would stay level under the High Growth scenario and even begin declining in the near term under the Conservative Growth scenario, even as overall data center electricity consumption increases. Overall, our results show that increased RE procurement is a much larger driver of potential emissions reductions than increased efficiency (PUE).

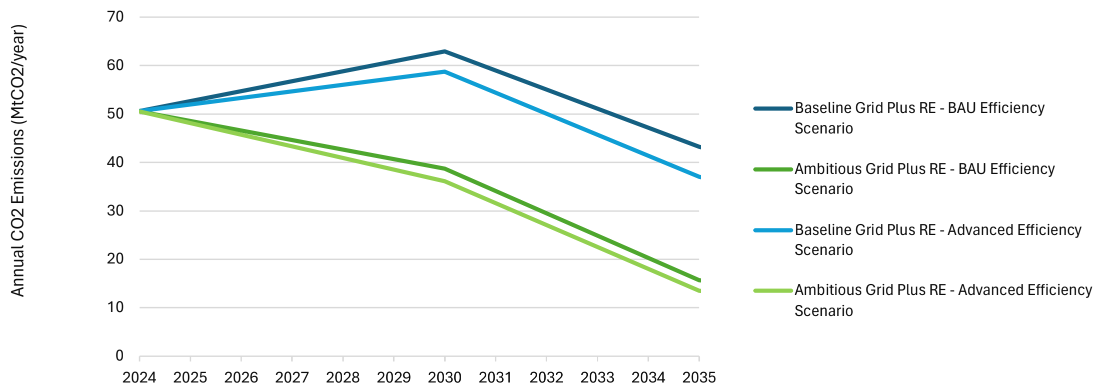

Figure 6: Projected annual CO_2 emissions from data centers over time in the US, Conservative Growth scenario

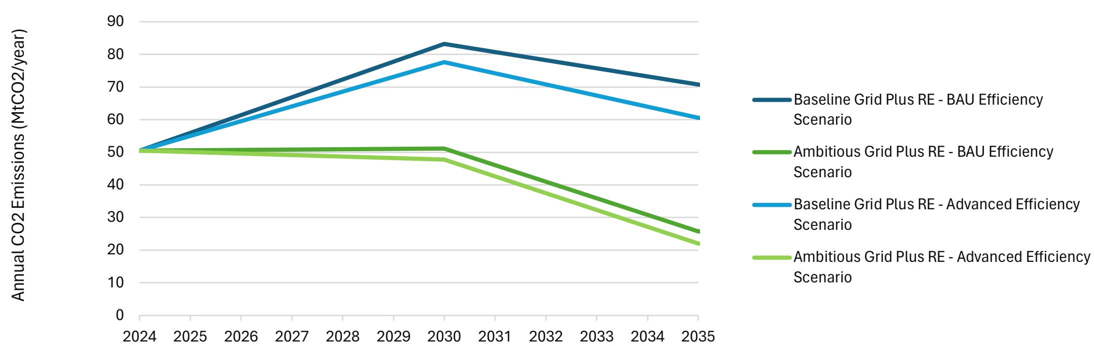

Figure 7: Projected annual CO_2 emissions from data centers over time in the US, High Growth scenario

Looking at results by state, even though Virginia currently has slightly higher emissions from data center electricity consumption than Texas, over time and across scenarios, emissions from data centers in Texas will surpass those of Virginia even though projected electricity consumption from data centers in Virginia will be higher than in Texas. This is because Virginia's grid emissions factor is lower than that of Texas, and Virginia's grid is projected to decarbonize faster based on a 2045 net zero goal for Dominion Energy. Likewise, although California is estimated to have nearly double the amount of data center electricity consumption relative to Ohio, the associated emissions are lower because California's grid emissions factor is three times lower than that of Ohio, which still significantly relies on coal for electricity generation. Across states and growth scenarios, it can be seen that with Ambitious levels of RE procurement, data center emissions will not increase, even as overall electricity consumption increases.

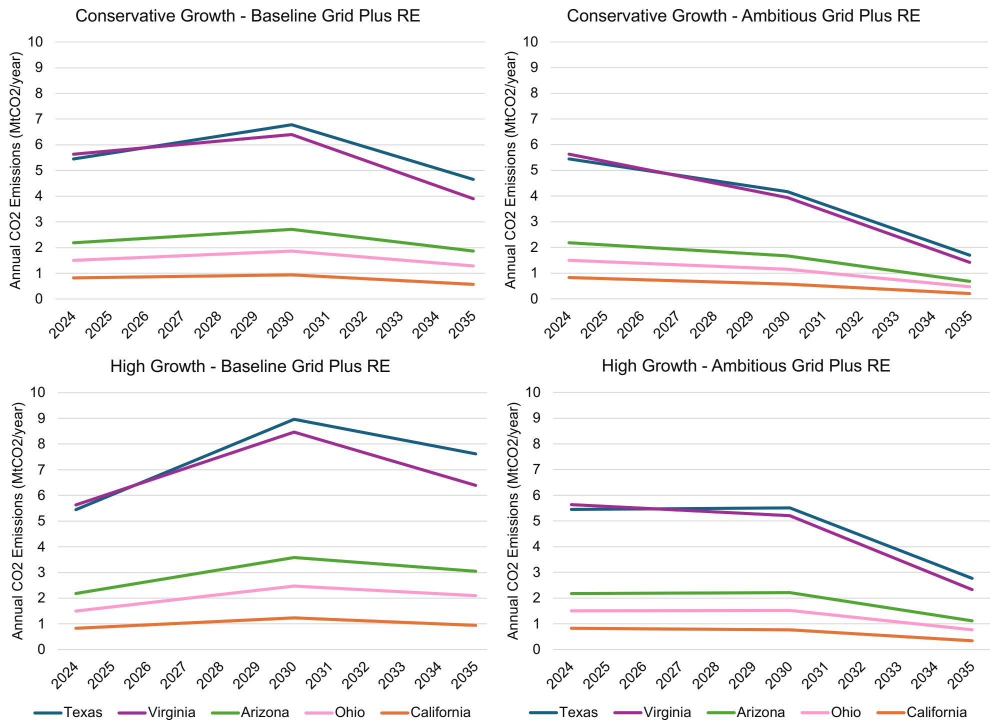

Figure 8: Emissions from data centers by state and growth/electricity scenario
Note: We present results here for the BAU Efficiency scenario. Under the Advanced Efficiency scenario, results would decrease slightly.

Another way to compare states is by the emissions intensity of computing across states in terms of the CO_2 emissions per unit of electricity consumed for data centers (tCO_2/MWh). Ohio has the highest emissions intensity, while California has the lowest. With more Ambitious levels of RE procurement, emissions intensity could be lowered more rapidly. Emissions intensity of computing could be a factor for decision-making about optimal siting of data centers from the perspective of minimizing CO_2 emissions.

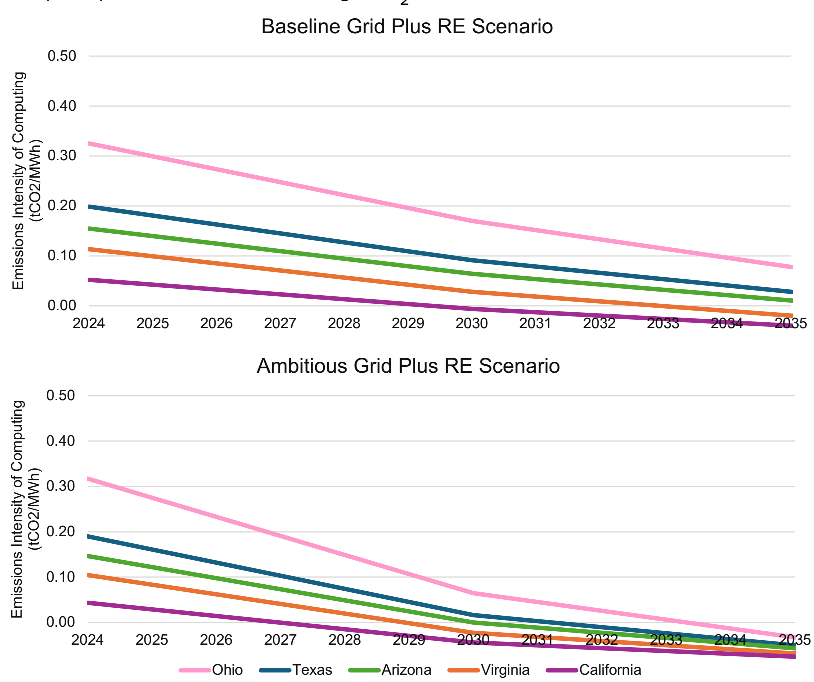

Figure 9: Emissions intensity of computing across the five studied states, 2024-2035, Baseline and Ambitious Grid Plus RE scenarios

Electricity System Profiles of Key Data Center Regions

This chapter analyzes the electricity system profiles of key data center states and regions. We assess key issues relevant to the electricity consumption and emissions impacts of data centers in each state/utility region: the renewable energy context, including the state of renewable energy development and available options for renewable electricity procurement for data centers; the electricity demand context; and the data center regulatory context. It is important to note that these factors change rapidly, and these profiles represent information as of late 2024.

4.1. Virginia

Dominion Energy is the primary utility serving Virginia, providing electricity generation, transmission, and distribution services to meet the state's energy needs. As the state's largest utility, Dominion is responsible for building new infrastructure, including power lines and substations, to support the heavy power demands of the data center sector in Northern Virginia. PJM Interconnection operates as the regional transmission organization (RTO) overseeing the electricity grid across Virginia and several other states in the Mid-Atlantic and Midwest. PJM ensures grid reliability, manages electricity flows across state lines, and coordinates the wholesale electricity market to balance supply and demand across the region.

As of 2022, renewable energy resources provided approximately 11% of Virginia's total electricity generation, with solar energy contributing the largest share. Solar power, primarily generated by utility-scale PV, accounted for nearly 6% of the state's net electricity generation. Solar energy production in Virginia more than doubled between 2020 and 2022, driven by a combination of favorable legislation, economic factors, and rising demand. Virginia has two nuclear power plants that make up about 31% of annual electricity generation.

A key legislative driver of renewable energy is the Virginia Clean Economy Act (VCEA), enacted in 2020, which mandates that Dominion Energy Virginia and Appalachian Electric Power achieve 100% renewable energy production (including nuclear) by 2045 and 2050, respectively. Additionally, Virginia has set a state-wide renewable energy capacity target: by 2028, the state aims to reach 5,500 MW of installed wind and solar capacity (Virginia DEQ 2024)

In addition to these grid decarbonization measures, there are also several programs offered to Dominion Energy commercial customers, such as data centers, for renewable energy procurement. There are multiple programs for renewable energy certificates (RECs), varying in their premium, certification level, which renewable energy sources are included, and where the RECs are sourced from (such as locally in Virginia or nationwide) (Table 7). Commercial and industrial purchasers of RECs may also be designated as Accelerated Renewable Energy Buyers and be exempted from certain utility charges.

Table 7: REC Programs offered by Dominion Energy for commercial electricity users

Program Name	Dominion Energy REC Select℠	Dominion Energy 100% Renewable Energy®	Dominion Energy Green Power®
Match a Portion	✓		✓
Match 100%	✓	✓[1]	✓
Cost	0.269¢/kWh $2.69/MWh	0.398¢/kWh $3.98/MWh	1.2¢/kWh $12/MWh
Energy & RECs[3]		✓	
Green-e® Energy Certified			✓
Renewable Resources			
Location	National RECs	VA and NC	VA and surrounding region[4]

Source: Dominion Energy 2024

In addition, the Renewable Energy Pilot Program enables Dominion Energy commercial customers to enter into Power Purchase Agreements (PPAs) with third-party renewable energy providers. Under this program, a third party, such as a solar or wind energy developer, can install renewable energy on a customer's premises and sell the generated electricity back to that customer, usually at a rate lower than utility-provided electricity. The Virginia State Corporation Commission (SCC) oversees this pilot program, with limitations set on the total renewable capacity that can participate. Dominion Energy's limit for these PPA arrangements is 500 MW for Virginia commercial and industrial users. Once this limit is reached, no additional projects can be accepted into the pilot. The initial phase of the pilot had a 50 MW cap, which was quickly reached (Vogelsong 2020).

Data centers in Virginia can also enter into offsite PPAs with renewable energy developers. For example, Microsoft and AES have entered into a significant renewable energy procurement agreement to power Microsoft's data centers in Virginia. Through this arrangement, Microsoft will receive renewable energy from a portfolio of 576 MW of wind, solar, and battery storage assets in the PJM interconnection, sourced by AES. The agreement includes both existing and new renewable assets. Similarly, Google has partnered with Virginia-based Energix Renewables through multiple PPAs to secure solar power for its data centers. Initially covering 1.5 GW of solar capacity with options for expansion, this agreement includes direct electricity provision and RECs.

Dominion Energy estimates that power demand will increase by 85% over the next 15 years, largely driven by data centers, necessitating substantial infrastructure upgrades. Power demand forecasts have been continually revised upwards. Dominion's recent Integrated Resource Plan (IRP) indicates that meeting this rising demand may require nearly doubling the current power supply through a mix of renewable and fossil fuel sources. This includes plans for a new natural gas facility. Out of nearly 21 GW of additional capacity proposed under the IRP by 2040, 28% will be natural gas under the stated "Plan A" (Dominion Energy 2023). Actual generation from natural gas facilities will represent a higher share of the electricity mix based on typical capacity factors, and this pathway would be more emissions-intensive than our Ambitious RE Procurement scenario, but possibly in line with our Baseline RE Procurement scenario if significant wind and solar energy are deployed. Since the most recent IRP, there has also been approval to build a nuclear fusion power plant in Virginia to meet growing electricity demand (Proctor 2024).

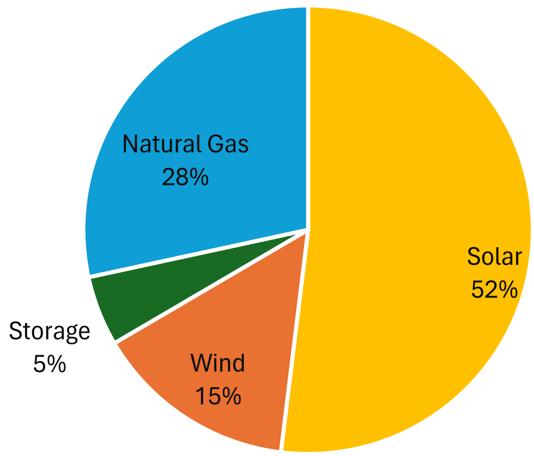

Figure 10: Dominion Energy share of energy source for planned additional capacity by 2040

Source: Dominion Energy Integrated Resource Plan 2024, "Plan A" (other plans have lower shares of proposed natural gas capacity, largely based on substitution with nuclear energy)

The economic incentives for localities to approve data centers are substantial, as seen in Loudoun County, where data centers contribute nearly a third of the county's budget. However, concerns over environmental and social impacts of data centers and the strain on the electrical grid have led to a large amount of proposed legislation. In previous state legislative sessions, numerous bills have been introduced to address the challenges of data centers. For example, proposed bills by Senators Russet Perry, Suhas Subramanyam, and others sought to take actions such as: provide tax credits to data centers that meet energy efficiency standards (i.e. 90% of energy from carbon-free sources), submit quarterly energy reports, classify data centers in a new category for utility rate structuring, etc. However, the bills have either been defeated, left in committee, or pushed to the 2025 session in order to allow the Joint Legislative Audit and Review Commission (JLARC) to complete a study on data center growth and its implications. The JLARC study, authorized in December 2023, will examine data center trends, assess economic impacts, and evaluate regulatory options. The defeat of various energy-related bills for data centers in Virginia indicates that there may be unmet need for legislative strategy and coordination during and after the JLARC study.

Direct regulation of data center siting in Virginia has occurred at the county level. In 2024, Fairfax Country (adjacent to Loudoun County) approved a zoning ordinance amendment establishing stricter standards for data centers, primarily addressing their size, location, and visual and noise impacts. The amendment mandates that data centers meet requirements for equipment enclosures, size limitations, setbacks from residential areas and Metro stations, noise mitigation, and architectural design (Schneider 2024).

4.2. Texas

The Texas electricity grid is managed primarily by the Electric Reliability Council of Texas (ERCOT), which operates as the state's regional transmission organization (RTO). ERCOT oversees about 90% of Texas's electric load, ensuring grid reliability and managing the state's wholesale electricity market. Unlike other parts of the U.S., ERCOT's grid has limited interconnections to neighboring states. The Texas electricity market is served by several major utilities and retail providers, including Oncor Electric Delivery, CenterPoint Energy, and Texas-New Mexico Power, as well as municipal utilities and electric cooperatives.

As of 2023, renewable sources accounted for nearly 30% of the state's electricity generation. Wind energy is the dominant renewable source in Texas, with nearly 41,000 MW of generating capacity, making up about two-thirds of the state's renewable capacity and one-fourth of its total electricity generation capacity. Solar energy in Texas is also experiencing rapid growth. By the end of 2023, Texas had nearly 18,500 MW of solar capacity. Solar power accounted for about 6% of Texas's electricity generation, with small-scale installations contributing roughly one-seventh of this capacity. In 2024 and 2025, nearly 24,000 MW of additional solar capacity is expected. Beyond wind and solar, Texas has a limited but diverse mix of other low-carbon energy sources. Biomass and biogas contribute a small fraction, around 0.2%, of the state's electricity, derived from sources like wood, landfill gas, and agricultural waste. Hydroelectric power remains minimal due to limited potential for development, and while there are no geothermal power plants, Texas is exploring opportunities to use geothermal resources from oil and gas operations.

Texas's deregulated electricity market, primarily managed by ERCOT, offers data centers multiple options for purchasing electricity compared to regulated markets like Virginia. In Texas, data centers can select from a range of retail electricity providers (REPs), allowing them to negotiate customized contracts, including options for renewable energy and the choice to pick fixed rates, time-of-use (TOU) pricing, long-term contracts, etc.

Several major tech companies have recently secured substantial renewable energy commitments in Texas to support their data center operations. Google has signed agreements with X-Elio and SB Energy to power its Dallas-area data centers and cloud region. The PPA with X-Elio involves offtaking from the 128 MW Bell Solar PV Plant, which will include a 100 MW battery storage system for 24/7 power support, starting operations in 2025. Meanwhile, SB Energy's "Orion Solar Belt" project, comprising three solar installations with a combined 900 MW capacity, is Google's largest solar investment globally. Google's ongoing renewable expansion in Texas also includes a $1 billion investment in its Dallas data centers, supported by 375 MW in solar PPAs. Microsoft has similarly advanced its renewable energy procurement in Texas, partnering with Leeward Renewable Energy (LRE) to secure 400 MW of solar capacity. Microsoft will source this energy from two LRE projects: the Morrow Lake Solar and Cradle Solar facilities, set to be operational by late 2024 and 2025, respectively (Swinhoe 2024c). In addition, Soluna Holdings has signed a PPA with EDF Renewables and Masdar for "Project Kati," a renewable computing data center in Texas designed for high-performance applications, including AI. Co-located with a wind facility, Project Kati will initially deliver 83 MW of renewable energy, with plans to double capacity in future phases (Decisions 2024).

Despite significant growth in renewable energy capacity, Texas is projecting a steep increase in electricity demand over the coming years, largely due to rising requests from data centers, hydrogen facilities, and electrified oil and gas operations, as well as a growing population. ERCOT has revised its forecast upwards, indicating that the state may need to nearly double its current electricity capacity by 2030, with data centers driving the largest share of growth (Figure 11). This updated outlook reflects recent legislation that allows ERCOT to count potential grid connection requests as part of its planning, even if agreements are not finalized. Because load estimates also inform transmission infrastructure planning, inflated demand projections can result in building costly and potentially underutilized infrastructure, increasing capital expenditures that may ultimately be passed on to consumers through higher rates. There is debate over whether the ERCOT projections may be overestimated if certain projects do not materialize, or if it is safer for ERCOT to plan for higher demand to avoid future shortages, given recent crises in grid reliability in Texas (e.g. the 2021 winter storm outages) (Foxhall and Guo 2024)

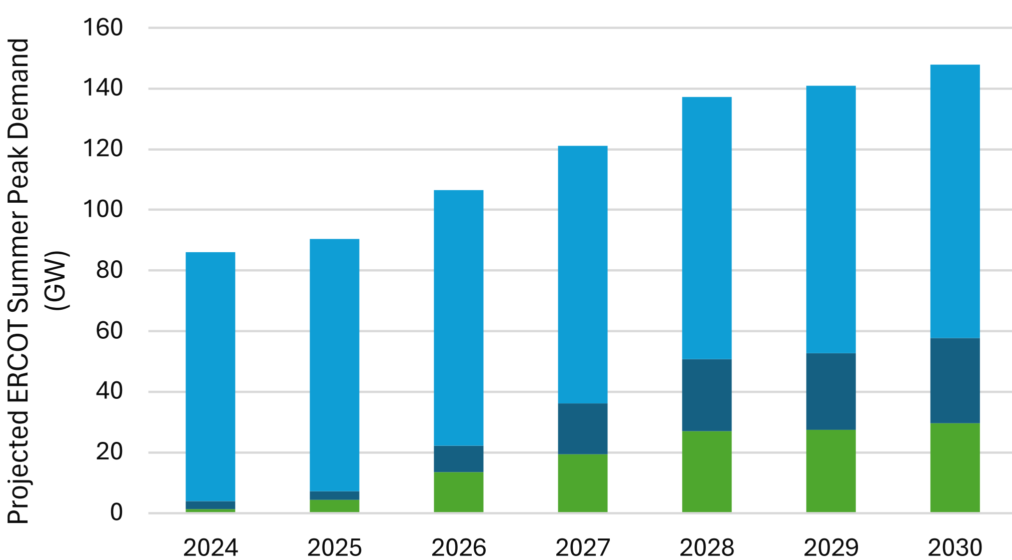

Figure 11: Share of projected data center load in ERCOT large load growth forecast. Other large load sources include cryptocurrency mining, hydrogen production, and oil and gas production. Data center load is based on loads with signed contracts or officer letters as of March 2024 (Source: adapted from ERCOT 2024).

Texas is exploring potential regulations on data centers. The state legislature and the Public Utility Commission of Texas (PUC) are examining ways to manage increased load, such as limiting data center locations, especially near residential neighborhoods, similar to recent restrictions in Virginia's Fairfax County. Another approach under discussion is to have data centers co-locate with power plants and potentially invest in their own power generation (Bloomberg 2024). This approach aligns with proposals in other states, such as Ohio, where utilities are shifting the cost of infrastructure expansion to new, high-demand customers like data centers.

4.3. Arizona

In Arizona, main electric utilities include Arizona Public Service (APS), Salt River Project (SRP), and Tucson Electric Power (TEP), both of which serve large portions of the state's population. APS is the largest electric utility in Arizona, serving the Phoenix metropolitan area and much of central Arizona, while TEP primarily serves Tucson and surrounding communities. These utilities operate within the broader framework of the Western Electricity Coordinating Council (WECC), which oversees grid reliability across the western United States, including Arizona.

Arizona has significant wind and solar potential. It is one of the top states in the US for solar-powered generating capacity, with over 6,100 MW of capacity from both utility-scale and small-scale installations, contributing to 10% of total electricity generation. Wind power is relatively untapped and only contributes about 1-2% of the state's electricity generation. Arizona also has the Palo Verde Nuclear Generating Station, the second-largest nuclear plant in the U.S., which provided 27% of Arizona's total net generation in 2023. Hydroelectric power accounts for about 5% of Arizona's electricity generation, however, this can significantly fluctuate due to ongoing drought conditions that impact water availability.

Arizona's electric grid is known for its high reliability, ranking 7th in the nation, even with the integration of renewable energy sources. The state's resilience can be attributed to several factors, including relatively new infrastructure and effective planning and maintenance. Arizona's electric utilities proactively manage tree trimming programs and rely on a significant portion of underground utilities, which reduces risks from power outages caused by falling trees or extreme weather events. Given heat levels in the state, utilities plan for a significant margin during peak summer periods when air conditioning usage surges.

Arizona previously had a renewable energy standard that targeted 15% RE for utilities by 2025, however, the standard was repealed in 2024. Nevertheless, utilities continue to set their own RE targets. APS aims to achieve 100% carbon-free electricity by 2050 and plans to reach 65% clean energy by 2030, with 45% of that coming specifically from wind and solar. APS also intends to eliminate coal from its energy mix by the end of 2031. Similarly, SRP has committed to a substantial increase in renewable capacity, with plans to add about 20,000 MW by 2035, 90% of which will be from renewable sources and 10% from natural gas. By 2050, SRP aims to reduce its fossil fuel usage by 90% and completely phase out coal.

Individual utilities have RE programs that consumers like data centers could elect to participate in. SRP offers two REC programs to help businesses meet their sustainability goals: REC Select and Solar Choice Select. REC Select allows businesses to purchase RECs, with SRP handling procurement and adding the cost to the customer's monthly bill. Businesses can choose from two REC options: "Solar and Wind," which includes newer solar and wind resources, or "Carbon-Free," which adds large hydro sources and includes older resources. Both options cover North American resources, with a maximum program capacity of 750,000 MWh for Solar and Wind and 1,000,000 MWh for Carbon-Free. Solar Choice Select allows businesses to offset up to 20% of their energy usage with utility-scale solar power generated directly on SRP's grid. Designed for large commercial accounts that consume over 750,000 kWh annually, businesses pay an additional $0.005 per kWh for the solar energy they use.

In addition to RECs, data center operators can also directly enter into PPAs for renewable energy. Google and Meta have both signed significant renewable energy PPAs in Arizona to support their data center operations and sustainability goals. Google's agreement with SRP includes more than 430 MW of wind, solar, and battery storage facilities managed by NextEra Energy Resources. Google's Mesa data center, set to open in 2025, aims to achieve at least 80% carbon-free energy usage on an hourly basis by 2026. Meta, partnering with SRP and Ørsted, has secured a PPA to procure the majority of the 300MW generated by Ørsted›s Eleven Mile Solar Center, which includes a 4-hour battery energy storage system. This energy will primarily power Meta's Mesa data center, with excess power distributed to SRP customers (Swinhoe 2024a) (Colthorpe 2023).

Arizona is experiencing significant load growth projections driven by multiple factors. APS's 2028 forecast increased from 8.6 GW to 9.5 GW in just one year, marking an 11% rise. This growth is fueled largely by an influx of large commercial, industrial, and data center loads. Additionally, Arizona's rising temperatures and peak summer heat are expected to push up peak demand levels, which is factored into load-growth projections (Wilson and Zimmerman 2023). As the state attracts data-intensive industries, APS anticipates that existing transmission capacity will be fully utilized before the end of the decade, requiring the Arizona grid to double in size over the next 15 years to keep pace with economic growth and electrification trends.

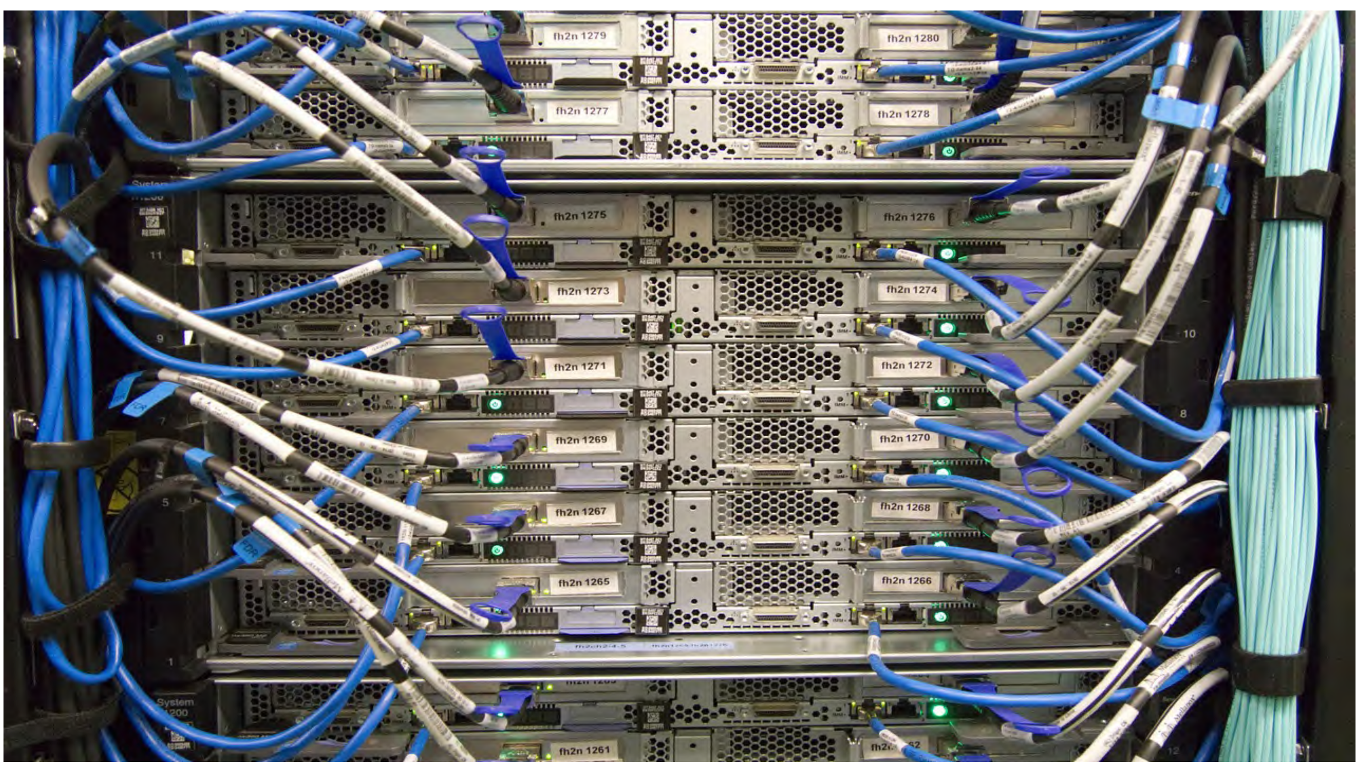

Arizona's approach to data centers is largely incentive-based, with limited regulation aside from local zoning ordinances like the one implemented by the city of Chandler. This ordinance restricts data centers to planning area development zones, sets specific noise limitations, and mandates pre- and post-construction sound studies as well as annual noise assessments for five years. These requirements are part of Chandler's effort to address recurring community concerns about noise from data centers, especially around backup generator maintenance and operational noise (Butler 2023). Statewide, however, Arizona's main focus is on incentives that attract data centers. These incentives include reductions in property tax and other financial benefits for companies that establish data centers in Arizona. In particular, data centers benefit from sales tax exemptions on equipment purchases, lowering their setup costs significantly.

4.4. California

California is primarily served by three major investor-owned utilities: Pacific Gas and Electric (PG&E), Southern California Edison (SCE), and San Diego Gas & Electric (SDG&E), each responsible for delivering power to different regions within California. The California Independent System Operator (CAISO) is the state's RTO, managing 80% of the state's electricity flow.

California first launched its Renewables Portfolio Standard (RPS) program in 2002. Currently, over 39% of the state's retail electricity sales come from RPS-certified renewables, with a total of 61% derived from zero-carbon sources, including large hydro and nuclear. Future targets aim to have California's retail electricity sales be supplied by 60% renewable energy by 2030, increasing to 100% by 2045. In 2024, clean energy supply matched demand for over 1,000 hours in 2024 within the CAISO service area (Lazo 2024). However, California still relies on natural gas as its largest electricity source, particularly to support the grid during peak demand or when renewable generation dips, such as during heat waves or wildfires, which are projected to be more common as the climate changes. To address blackouts and grid issues, California has added over 10,000 MW of battery storage.

Amazon has entered multiple PPAs with AES Corporation to secure renewable energy for its California operations. This includes energy from two solar-plus-storage projects in Adelanto, San Bernardino County, featuring the 150 MW Baldy Mesa project, with 75 MW of battery storage, and the smaller 50 MW Silver Peak project with 25 MW of storage. Microsoft also negotiated a PPA with AES Corporation to power its data centers with solar energy. This agreement includes 110 MW of solar power combined with energy storage. Equinix, a key data center operator, committed to 100% renewable energy for its global operations, including 16 data centers in California using both RECs and PPAs. Equinix signed a PPA with SunEdison in 2015 to purchase 105 MW of solar power from the Mount Signal Solar II project, a 150 MW solar farm near Calexico.

California plans to double its generation capacity by 2045, but unlike the other states profiled here, data centers are not projected to be the dominant driver of upward load growth revisions. CAISO's most recent forecast for 2028 increased from 49.3 GW to 49.8 GW, driven by the need for RE capacity to help meet decarbonization goals. In addition, California's load growth is largely expected to come from statewide efforts to electrify transportation and buildings. Data centers account for less than 0.3 GW of the anticipated load increase (Wilson and Zimmerman 2023).

California has implemented specific regulations related to data center efficiency and emissions. Since 2014, all state-owned and leased data centers over 200 square feet must comply with guidelines for temperature and humidity, as well as California's Title 24 Building Energy Efficiency Standards. Larger data centers, exceeding 1,000 square feet, are required to measure and report their PUE annually, aiming to achieve a PUE of 1.5 or lower through energy-saving measures. Data centers with a PUE above 1.5 are required to reduce their PUE 10% per year until they reach 1.5. Additionally, state agencies must prioritize energy-efficient equipment that meets certain standards (California Department of General Services 2014)

In 2023, California enacted the Climate Corporate Data Accountability Act (SB 253) and the Climate-Related Financial Risk Act (SB 261), mandating large businesses, including data centers, to disclose their direct and indirect greenhouse gas emissions. Under SB 253, businesses with annual revenue of more than $1 billion must report emissions from their facilities, including supply chain impacts, to the California Air Resources Board by 2026. SB 261 requires companies to disclose climate-related financial risks. Some experts have argued that this regulatory burden disadvantages California relative to other potential states for data center siting (Patrizio 2023).

4.5. Ohio

The primary electric utility serving Ohio is AEP Ohio, a subsidiary of American Electric Power. AEP Ohio operates within the jurisdiction of the PJM Interconnection, the regional transmission organization (RTO) that oversees the electricity grid and manages wholesale electricity markets across multiple states in the Midwest and Mid-Atlantic.

Ohio's renewable energy landscape is shaped by its Renewable Portfolio Standard (RPS), which mandates that 8.5% of the state's electricity come from renewable sources by 2026. In 2023, Ohio experienced significant growth in solar installations, although future growth is uncertain due to the ongoing effects of legislative changes. The 2021 passage of Senate Bill 52 (SB52) allowed individual counties to block renewable energy projects, creating regulatory uncertainty for developers. While some counties remain open to renewable development, others have used this authority to restrict projects, leading to a decline in renewable energy applications in 2023, even as installations continued based on applications from previous years.

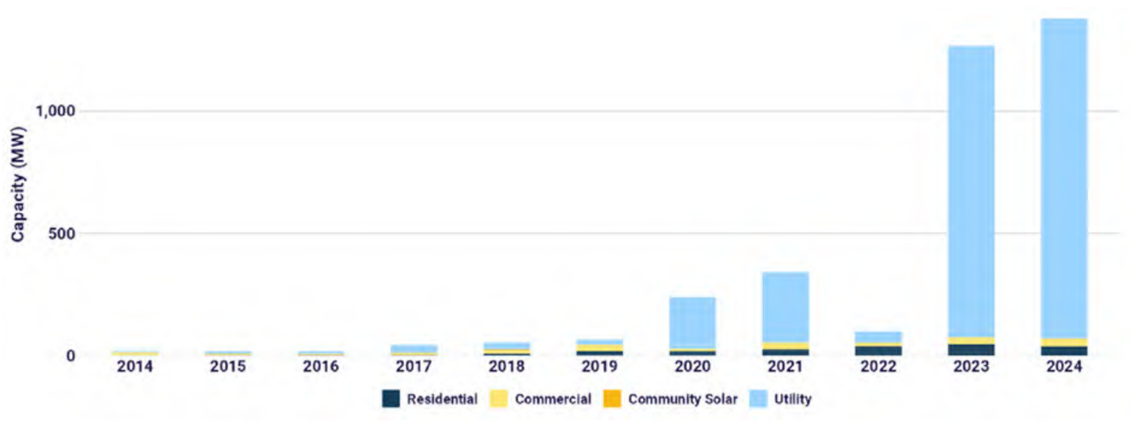

Figure 12: Annual solar installations in Ohio, 2014-2023. Source: (SEIA 2024)

Despite Ohio's relatively carbon-intensive electricity grid and constrained RE opportunities relative to the other studied states, there are still examples of RE PPAs by tech companies operating data centers in Ohio. Amazon secured a PPA for the entire generation capacity of the 150 MW Fox Squirrel Solar Phase 1 project, a part of the larger 577 MW Fox Squirrel Solar development by EDF Renewables North America, which is expected to be fully operational by the end of 2024. Microsoft has entered into a PPA with Apex Clean Energy to procure power from the 125 MW Wheatsborough Solar project in Erie County, developed under the Microsoft–Volt Energy Utility Environmental Justice Framework (Swinhoe 2024b). Although Microsoft currently lacks a cloud region in Ohio, it has acquired land near Columbus for potential future data center expansion.

AEP Ohio previously projected a steady annual load growth rate of around 1% to 2%. However, with the rapid influx of data center developments, AEP Ohio has now revised its forecast, expecting an average annual growth rate closer to 20%. The utility has signed multiple Electric Service Agreements (ESAs) and Letters of Agreement (LOAs) with various customers—primarily data center operators—aiming to add 4-5 GW of load in Central Ohio by 2030. This substantial load growth prompted AEP Ohio to present updated projections to PJM, which is also planning for significant additional transmission investments.

In response to the rapid expansion of data centers and its impact on Ohio's power infrastructure, AEP Ohio has implemented a series of regulatory measures. In March 2023, AEP imposed a temporary moratorium on new service requests from data center customers to allow its transmission planning team time to assess the strain on the electrical delivery system. Previously, the utility faced scrutiny over "gold-plating" transmission infrastructure—referring to over-investment in capital projects to increase profit margins—which has raised concerns among stakeholders about cost efficiency (Zullo 2023).

In 2024, AEP Ohio filed a proposal with the Public Utilities Commission of Ohio to introduce a new rate structure specifically for data centers, cryptocurrency mining, and mobile data center operations. Under this proposed rate, new data centers with loads exceeding 25 MW would be required to make a 10-year commitment, paying for a minimum of 90% of the electricity they request, regardless of actual usage. This rate structure has sparked pushback from tech companies and data center coalitions, who argue that it imposes unfair financial obligations.

Ohio's temporary moratorium on data centers indicates the need for more careful consideration of how planned growth will affect transmission infrastructure investments, and future discussions should take into account improved and more transparent load forecasting. Proposed adjustments to rate structures should also be coupled with incentives for efficiency and RE procurement in order to ensure that cost-sharing does not lead to data centers locking in high-carbon electricity, overestimating energy needs, or facing challenges in participating in load shifting programs. However, RE procurement may be a challenge for data centers unless Ohio's SB52 for RE approvals is adjusted. Future developments regarding RE installation and data center regulation in Ohio are thus highly uncertain.

Technological Innovation Trends

Technological innovation in computing is essential to meet the rising demands of AI and cloud computing at data centers. Our results have shown that even advanced efficiency improvement in the PUE of data centers will not significantly shift electricity consumption projections, especially relative to overall computing demand growth rate scenarios. Therefore, this section profiles some of the drivers of computing demand and emerging technologies and innovations that could significantly improve the efficiency of computing, referring to the amount of computational work (such as processing data or running applications) that can be accomplished per unit of energy consumed.

Advances in computing efficiency hold potential to counterbalance some of the demand growth, and could ostensibly meet larger demand for computing while using less electricity. We identify four key areas of improvement: advanced hardware, software efficiency, non-traditional computing, and cloud-enabled energy optimization.

The first area is **advanced hardware**, such as through the use of GPUs, TPUs, and specialized processors, which are designed to deliver higher performance with less energy. AI-specific hardware like tensor processing units (TPUs) and graphics processing units (GPUs) for machine learning applications provide significant energy savings by offering more processing power per watt compared to traditional CPUs. Examples include Nvidia's RAPIDS Accelerator for Apache Spark that allows GPUs to process data faster and with significantly less energy, reducing costs and carbon emissions by up to 80%; and Nvidia's Blackwell architecture and the Cerebras Wafer Scale Engine 3 (WSE-3) (L., Jennifer 2023).

This category also includes emerging alternatives such as 3D transistor stacking, which offers better power efficiency and faster operation compared to planar transistors, by effectively increasing transistor density without the limitations of two-dimensional scaling. High-speed switch technology and port utilization have reduced the power requirements of network switches by enabling a single switch to handle multiple devices, reducing the number of switches needed. Rollable ribbon fiber cables also increase fiber density (Commscope 2024).

Another critical area is **software efficiency**. By optimizing algorithms and software to better utilize hardware resources, significant energy savings can be achieved. Virtualization allows multiple virtual machines (VMs) to run on a single physical server, consolidating workloads and enhancing server utilization. This enables the dynamic allocation of resources like CPU, memory, and storage, improving scalability and reducing operational complexity. Containerization, as a lighter-weight alternative to virtualization, further enhances efficiency by isolating applications with a smaller footprint, allowing multiple containers to operate on a single server. Additionally, data reduction techniques such as compression and deduplication reduce the storage footprint and optimize network bandwidth, resulting in faster data transfers and lower latency (Whittington 2024). Techniques such as sparse modeling—where AI models are structured to focus on only the most relevant data—reduce computational requirements for model training and, consequently, power consumption (Patterson et al. 2022). Moreover, cloud-based orchestration tools and virtualization can dynamically manage computing loads, ensuring that resources are used efficiently across multiple applications and reducing energy waste. This is especially impactful in cloud environments where scaling policies can be optimized to prioritize energy-efficient operations. A recent example of breakthroughs in software efficiency is the release of a new generative AI model by the company DeepSeek.

DeepSeek achieved significant efficiency gains primarily through software optimizations that reduce computational overhead while maintaining high performance. DeepSeek's efficiency improvements primarily align with sparse modeling and data reduction techniques, optimizing both inference and training costs.

Non-traditional computing paradigms like quantum and neuromorphic computing offer longer-term potential for energy-efficient AI. Quantum computing, which utilizes qubits to perform calculations exponentially faster than classical computers, could transform high-complexity tasks like model training, provided the technology matures. Neuromorphic computing, which mimics the brain's neural architecture, also presents a pathway to achieve high efficiency in specific AI applications. Photonic computing, which relies on light rather than electrons for data transfer, could increase computing speed and bandwidth with reduced energy consumption. These alternatives are emerging pathways to handle computationally intensive tasks without relying solely on traditional silicon-based hardware (Balderas 2024).

Finally, **location and time-based energy optimization** using cloud computing is increasingly considered in data center operations. By leveraging cloud environments, companies can dynamically allocate computing resources based on demand. Container orchestration tools like Kubernetes allow cloud data centers to scale applications based on real-time demand, ensuring that energy is only used when necessary and that resources are optimized across workloads. In addition, shifting workloads to regions and times where renewable energy is abundant, which our analysis above has shown can vary significantly, can reduce the overall carbon footprint of AI workloads.

In addition to these innovation-driven improvements in efficiency that could mitigate computing demand growth, it is important to note that user demand for AI is one of the key uncertainties in future data center electricity demand. The business case for AI remains uncertain, and AI services will not necessarily see massive uptake by consumers and users due to a range of limiting factors, including concerns about privacy and data security, prices for AI services, learning curves for users, and the ethical, cultural, and regulatory landscape surrounding AI. These "demand side" factors could significantly mitigate growth in demand for AI.

Policy Recommendations

This study highlights the growing electricity consumptions from data centers across the US, showing that without an increase in RE procurement, CO_2 emissions from data center electricity use will increase along with electricity demand. In addition, our state-level quantitative and qualitative analysis has highlighted several challenges with renewable energy supply and procurement, policy support, and efficient management of data center load – challenges that are important to resolve in order to minimize the emissions impact of data centers. While the increases in computing efficiency profiled in Chapter 5 are also key to reducing the energy and emissions footprints of data centers, our report and the attendant recommendations are focused on energy system interventions.

To lessen the emissions impact of data centers in the U.S. and key states, several targeted policy areas should be addressed, with specific actions within each category to facilitate a cleaner and more efficient data center industry.

1. Increasing Renewable Energy Procurement

Encourage 24/7 Clean Energy Tracking: Support data centers in adopting 24/7 carbon-free energy models, which enable continuous tracking and verification of clean energy usage on an hourly basis and prioritize purchasing clean/carbon-free electricity on the data center's local grid rather than via RECs. Policies should incentivize data traceability and tracking systems to provide real-time data on RE generation and consumption. Implementing these systems may require partnerships with grid operators and utilities to ensure comprehensive data availability and accurate reporting across varied regions.

Incentivize Long-term Renewable Energy Agreements: By offering tax credits, grants, or streamlined permitting processes to data centers that commit to lengthy renewable energy contracts, policymakers can make it easier for data centers to access sustainable energy sources over the long term. Long-term PPAs not only provide stability for data centers but also promote the growth of renewable capacity, ensuring a steady supply of clean energy that meets their expanding energy demands.

2. Better Management of Data Center Load

Support Demand Response (DR) Program Participation: Demand response can significantly reduce data center electricity demand by allowing facilities to temporarily lower their energy use during local grid constraints, and companies like Google are already piloting DR programs for data centers (Vincent 2023). Utilities and policymakers should encourage data centers to participate in demand response programs. Incentives for flexible load management, such as rebates for DR participation, will help align data center operations with grid needs. DR programs may be implemented by the utility via direct load control, or they may be price-based, with the consumer or data center choosing when to consume electricity based on known price schedules. Price-based DR has large untapped potential in the US (Carvallo and Schwartz 2023).

Improve Load Forecasting for Data Centers: Federal and state agencies, utilities, and RTOs should collaborate on improved load forecasting tools tailored for data centers. This includes evaluating potential operational characteristics for AI, cloud, and other data-intensive services, as well as working with data center operators and other states to understand the true demand for computing and eliminate overestimation and double counting. Data center participation in DR programs should also be factored in to load forecasting and integrated resource planning processes, treating such programs as resources in utilities' capacity and resource adequacy analysis. More accurate forecasting will improve utilities' planning and investment in generation and transmission infrastructure, ensuring efficient use of resources.

Encourage Flexible Scheduling: Promote scheduling incentives that allow data centers to shift non-time-sensitive tasks to times when renewable energy is most available, thereby reducing emissions and grid strain during peak demand hours. For example, AI model training or large data processing tasks could be scheduled during off-peak hours when renewable energy is most abundant. Flexible scheduling can also use cloud computing to take advantage of location-specific renewable energy advantages.

3. Grid Infrastructure Optimization

Expand Transmission Networks to Support Renewable Integration: Public investments and policy support are necessary to modernize transmission lines and extend their reach, ensuring a reliable flow of renewable energy to major data center hubs. Federal authorities should play an active role in enhancing regional transmission planning, optimizing cost-sharing models, and lowering interconnection costs for renewable energy projects.

Integrate Data Centers into a Shared Energy Economy Model: Data centers could contribute to grid reliability by co-locating on-site power generation (such as solar, battery storage, or natural gas) to support both the data center's needs and the grid during peak demand periods or emergencies. This will need to be carefully managed and planned in coordination with local utilities.

4. Increasing State-Level Ambition for RE and Net Zero Targets

Mandate RE Targets Specific to Data Centers: Require states with significant data center activity to establish renewable energy procurement or carbon-neutral goals tailored for the data center industry, helping align the industry's growth with state decarbonization goals. State policies could further incentivize compliance by linking targets to benefits such as tax breaks, expediting the renewable energy transition for the sector. For example, bills have been introduced in Virginia that would provide incentives to data centers that achieve 90% carbon-free energy, and future legislation could restrict the numerous sales tax and property tax incentives for data centers in various states to only those whose operators have set e.g. 100% zero-carbon electricity targets.

Provide Incentives for High Efficiency Standards in Data Centers: States can establish specific energy efficiency targets for data centers, tied to incentives for facilities that meet benchmarks for Power Usage Effectiveness (PUE) and low-carbon cooling technologies. For example, California has a PUE standard that could be adopted in other states as a relatively achievable energy savings measure.

Support Establishment and Advancement of State Net Zero Goals: Some states have net zero targets that are earlier than the US 2050 target. Advocacy efforts should continue to encourage more states to adopt net zero targets, including for earlier than 2050, which will drive overall grid decarbonization and RE investments that will also assist with lowering data center emissions through lower-carbon electricity supply. Achieving earlier state-level targets would help make renewable energy more widely available for data centers.

5. Improve Data Center Energy Management

Incentivize Real-Time Energy Management Solutions: Promote real-time energy monitoring and AI-driven energy management to optimize data center energy consumption. Policies should provide grants or tax benefits for companies that implement energy-efficient management systems and transparent reporting, specifying, for example, a target PUE in order to receive incentives. Virginia previously introduced a bill requiring quarterly energy consumption reports to the state Department of Energy.

Expand Cloud Service Adoption to Reduce Redundant Loads: Encouraging the adoption of cloud services for the public and private sectors could help reduce redundant loads by consolidating small data centers into larger, more efficient facilities. Cloud providers typically operate at scale, enabling more energy-efficient management and resource allocation than individual data centers. By promoting cloud solutions, particularly in government and corporate sectors, policymakers could optimize energy use across various organizations, reducing overall demand.

6. Public Data and Reporting Initiatives

Enhance Data Collection and Transparency on Data Center Energy Use: The federal government should create a dedicated category within energy surveys to gather specific data on data center energy use and emissions. This data can guide future policy and identify areas of greatest energy use and improvement potential. For example, California recently enacted an emissions reporting standard that will apply to large data centers in the state. Such transparency policies could include detailed metrics on power usage effectiveness (PUE), hourly energy consumption profiles, renewable energy procurement levels, estimated emissions intensity, and the extent of participation in demand response programs. States can lead ahead of federal policy by mandating regular public reporting of these metrics and standardized the data collection protocol to enhance ease and uptake.

Establish an Open Repository for Data Center Emissions and Efficiency Benchmarks: Develop a publicly accessible repository for data center emissions, energy use, and efficiency benchmarks to support industry transparency and innovation. Doing so could foster competition in efficiency improvements and allow regulators to monitor progress.

These policy recommendations can help optimize energy use and reduce CO_2 emissions from data centers, while also supporting the integration of renewable energy. Figure 14 describes what various stakeholders would need to do to accomplish these policy actions.

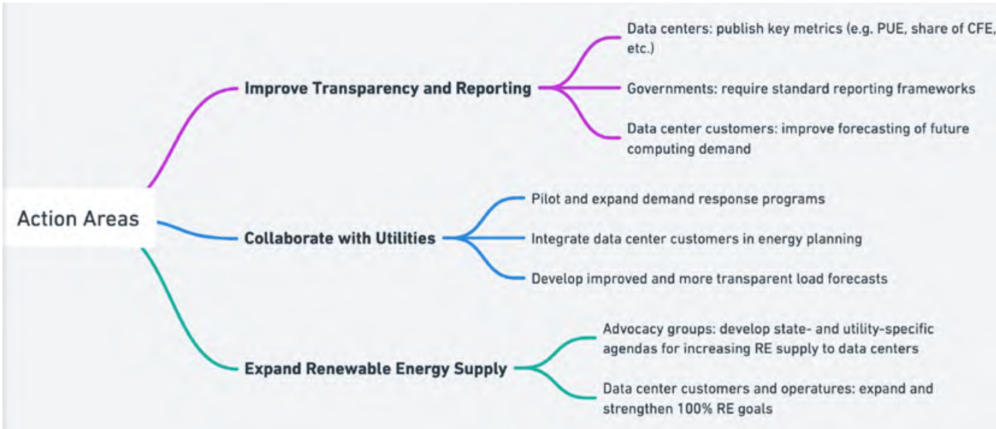

Figure 13: Policy actions by stakeholder for reducing emissions from data centers in the US

Conclusion

The siting of data centers across the US will be determined by a number of factors, including access to low-cost electricity, reliable power infrastructure, and proximity to end users, as well as emerging preference for renewable energy. States like Virginia and Texas lead in data center power demand due to robust infrastructure, large tech ecosystems, and competitive incentive structures. Emerging data center hubs such as Arizona and Ohio demonstrate that location decisions are also increasingly driven by favorable climate conditions, growing renewable energy resources, and supportive policy environments. These siting trends suggest that as cloud computing and AI continue to drive demand, regional and state-level factors will be critical in shaping the geographic distribution of data centers.

Computing demand is projected to grow significantly, with expected increases in data center electricity consumption alongside this. Our study projects that U.S. data center electricity demand could grow to 345 to 490 TWh per year by 2030 and potentially reach 655 TWh by 2035 under the High Growth-BAU Efficiency scenario. In Virginia, data centers are expected to account for a substantial portion of local electricity demand, representing 36-51% by 2030 and 28-54% by 2035, while nationally, data centers could rise from about 4% of total electricity use today to between 8-11% by 2030 and 8-15% by 2035, depending on growth and efficiency trends.

However, depending on levels of RE procurement for data centers, CO_2 emissions from data center electricity use need not grow as rapidly as electricity consumption, and may even decline in the near-term with ambitious levels of RE procurement. Annual CO_2 emissions from U.S. data centers may peak between 63-83 million metric tons by 2030, depending on the energy mix and growth projections.

As the energy demands of data centers continue to rise, especially with the influx of AI workloads, the alignment of data center operations with regional renewable energy goals and the development of infrastructure to support renewable generation will be necessary to reduce their emissions footprint. States and regions aiming to attract data center investments must balance incentives with regulations that support managed growth and emissions mitigation. Er recommend several policy actions to manage data center growth and minimize emissions. Renewable energy incentives, such as 24/7 clean energy tracking and long-term power purchase agreements, can help align data center needs with sustainable energy sources. Improved load management—through forecasting, demand response programs, and flexible scheduling—would better utilize grid resources and reduce emissions. Grid upgrades, including expanded transmission networks and shared energy models, would further facilitate renewable integration for data centers. Additionally, state-level policies could establish data center-specific renewable energy targets and efficiency incentives. Lastly, public data initiatives to increase transparency by tracking data center energy use and emissions can aid policymakers in decision-making for all of the above efforts.

References

Arizton. 2023. "U.S. Green Data Center Market Size, Industry Forecast 2028." https://www.arizton.comhttps://www.arizton.com/market-reports/us-green-data-center-market.

Balderas, Adolfo. 2024. "Transistors on the Edge: The Quest for Energy-Efficient Computing." *Harvard Technology Review* (blog). April 22, 2024. https://harvardtechnologyreview.com/2024/04/22/transistors-on-the-edge-the-quest-for-energy-efficient-computing/.

Bigelow, Stephen. 2024. "How to Design and Build a Data Center." TechTarget. 2024. https://www.techtarget.com/searchdatacenter/How-to-design-and-build-a-data-center.

Bloomberg. 2024. "Texas Regulator Wants Data Centers to Build Power Plants." Energy Connects. October 4, 2024. https://www.energyconnects.com/news/utilities/2024/october/texas-regulator-wants-data-centers-to-build-power-plants/.

Butler, Georgia. 2023. "Arizona City Adopts Data Center Zoning Ordinance." January 12, 2023. https://www.datacenterdynamics.com/en/news/arizona-city-adopts-data-center-zoning-ordinance/.

California Department of General Services. 2014. "Requirements of Data Centers and Server Rooms." California Department of General Services. 2014. https://www.dgs.ca.gov/Resources/SAM/TOC/1800/1820-3.

Carvallo, Juan, and Lisa Schwartz. 2023. "The Use of Price-Based Demand Response as a Resource in Electricity System Planning." None, 2205177, ark:/13030/qt3mp147d8. Lawrence Berkeley National Laboratory. https://doi.org/10.2172/2205177.

Colthorpe, Andy. 2023. "Meta Signs Solar-plus-Storage PPA with Ørsted for Arizona Data Centre Power." Energy-Storage.News. December 12, 2023. https://www.energy-storage.news/meta-signs-solar-plus-storage-ppa-with-orsted-for-arizona-data-centre-power/.

Commscope. 2024. "What's Next for the Data Center: 2024 Trends to Watch." https://www.anixter.com/content/dam/Suppliers/CommScope/Documents/commscope_ebook_2024-data-center-trends.pdf.

Cushman and Wakefield. 2023. "Americas Data Center Update." https://cw-gbl-gws-prod.azureedge.net/-/media/cw/global/insights/2024/americas-h2-2023-data-center-update.pdf?rev=a370c64425d145f6afc0d0aa4158bd89.

———. 2024. "2024 Global Data Center Market Comparison." https://cushwake.cld.bz/2024-Global-Data-Center-Market-Comparison.

DATA4. 2021. "A Closer Look at Free Cooling - DATA4." *DATA4 - Smart Data Centers at Scale* (blog). September 20, 2021. https://www.data4group.com/en/datacenter-dictionary/free-cooling/.

DataCenterMap. 2024. "USA Data Centers." 2024. https://www.datacentermap.com/usa/.

Decisions, Smart Energy. 2024. "Soluna Signs Wind PPA for Texas Data Center Project." 2024. https://www.smartenergydecisions.com/renewable-energy/2024/05/24/soluna-signs-wind-ppa-for-texas-data-center-project.

Digital Climate Alliance. 2024. "Promise and Peril: Sustainability and the Rise of Artificial Intelligence." https://www.digitalclimate.io/2024-ai-white-paper.

Dominion Energy. 2023. "Virginia Electric and Power Company's Report of Its 2023 Integrated Resource Plan." https://cdn-dominionenergy-prd-001.azureedge.net/-/media/pdfs/global/company/2023-va-integrated-resource-plan.pdf?la=en&rev=6b14e6ccd15342b480c8c7cc0d4e6593.

———. 2024. "100% Renewable Energy | Virginia | Dominion Energy." 2024. https://www.dominionenergy.com/virginia/renewable-energy-programs/100-percent-renewable-energy.

EPRI. 2024. “Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption.” https://www.epri.com/research/products/3002028905.

ERCOT. 2024. “Load Forecast.” 2024. https://www.ercot.com/gridinfo/load/forecast.

Foxhall, Emily, and Kayla Guo. 2024. “Texas Electricity Demand Could Nearly Double in Six Years, Grid Operator Predicts.” The Texas Tribune. June 20, 2024. https://www.texastribune.org/2024/06/20/texas-electricity-demand-forecast-ercot/.

L, Jennifer. 2023. “Nvidia’s Accelerated Analytics Can Cut Computing Cost and CO2 Footprint by 80%.” *Carbon Credits* (blog). August 29, 2023. https://carboncredits.com/nvidias-accelerated-analytics-can-cut-computing-cost-and-co2-footprint-by-80/.

Lazo, Alejandro. 2024. “California Hits Milestones toward 100% Clean Energy — but Has a Long Way to Go.” *CalMatters*, August 19, 2024, sec. Climate Change. http://calmatters.org/environment/climate-change/2024/08/california-clean-power-progress-grid/.

Meta. 2024. “2024 Sustainability Report.” https://sustainability.atmeta.com/wp-content/uploads/2024/08/Meta-2024-Sustainability-Report.pdf.

Patrizio, Andy. 2023. “How California Climate Disclosure Laws Could Impact Data Centers.” *Data Center Knowledge*, 2023. https://www.datacenterknowledge.com/sustainability/california-s-climate-disclosure-laws-could-have-adverse-effects-on-data-centers.

Patterson, David, Joseph Gonzalez, Urs Holzle, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David R. So, Maud Texier, and Jeff Dean. 2022. “The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink.” *Computer* 55 (7): 18–28. https://doi.org/10.1109/MC.2022.3148714.

Proctor, Darrell. 2024. “Virginia Officials Say State Will Be Home to World’s First Fusion Power Plant.” POWER Magazine. December 18, 2024. https://www.powermag.com/virginia-officials-say-state-will-be-home-to-worlds-first-fusion-power-plant/.

Riu, Isabelle, Dieter Smiley, Stephen Bessasparis, and Kushal Patel. 2024. “Load Growth Is Here to Stay, but Are Data Centers?” Energy and Environmental Economics, Inc.

Schneider, David. 2024. “Fairfax County, Virginia, Approves Updated Zoning Regulations Related to Data Centers | Insights | Holland & Knight.” 2024. https://www.hklaw.com/en/insights/publications/2024/09/fairfax-county-virginia-approves-updated-zoning-regulations.

SEIA. 2024. “Ohio.” SEIA. 2024. https://seia.org/state-solar-policy/ohio-solar/.

Shehabi, Arman, Sarah Smith, Alex Hubbard, Alex Newkirk, Nuoa Lei, Md Abu Bakar Siddik, Billie Holocek, Jonathan Koomey, Eric Masanet, and Dale Sartor. 2024. “2024 United States Data Center Energy Usage Report.” LBNL--1005775, 1372902. https://doi.org/10.2172/1372902.

Shehabi, Arman, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Inês Azevedo, and William Lintner. 2016. “United States Data Center Energy Usage Report.” LBNL--1005775, 1372902. https://doi.org/10.2172/1372902.

Stream Data Centers. 2024. “Tax Incentives for Data Centers.” Stream Data Centers. 2024. https://www.streamdatacenters.com/glossary/tax-incentives-for-data-centers/.

Swinhoe, Dan. 2024a. “Google Signs 480MW PPA in Arizona, Including Wind, Solar, and Battery Storage.” Data Center Dynamics. March 15, 2024. https://www.datacenterdynamics.com/en/news/google-signs-480mw-ppa-in-arizona/.

———. 2024b. “Microsoft Signs 125MW Solar PPA in Ohio.” January 15, 2024. https://www.datacenterdynamics.com/en/news/microsoft-signs-125mw-solar-ppa-in-ohio/.

———. 2024c. “Microsoft Signs 400MW Solar PPA in Texas.” March 28, 2024. https://www.datacenterdynamics.com/en/news/microsoft-signs-400mw-solar-ppa-in-texas/.

U.S. Energy Information Administration (EIA). 2024. “Data Center Owners Turn to Nuclear as Potential Electricity Source.” 2024. https://www.eia.gov/todayinenergy/detail.php?id=63304.

Vincent, Matt. 2023. "Google Is Now Reducing Data Center Energy Use During Local Power Emergencies." Data Center Frontier. 2023. https://www.datacenterfrontier.com/energy/article/33012588/google-is-now-reducing-data-center-energy-use-during-local-power-emergencies.

Virginia Department of Environmental Quality. 2024. "Renewable Energy | Virginia DEQ." 2024. https://www.deq.virginia.gov/our-programs/renewable-energy.

Vogelsong, Sarah. 2020. "With 'high Interest' from Customers, SCC Prepares to Restart Solar PPA Program • Virginia Mercury." *Virginia Mercury* (blog). 2020. https://virginiamercury.com/briefs/with-high-interest-from-customers-scc-prepares-to-restart-solar-ppa-program/.

Whittington, Brad. 2024. *Decarbonizing the Data Center For Dummies*. Wiley, Hitachi Vantara. https://www.hitachivantara.com/en-us/gated-forms/decarbonizing-data-center-dummies.

Williams, Mark. 2024. "AEP Ohio: Data Center Electricity Demands to Put Central Ohio on Same Tier as Manhattan." The Columbus Dispatch. 2024. https://www.dispatch.com/story/business/energy-resource/2024/05/14/aep-ohio-rolls-out-proposal-to-handle-surge-in-data-center-power-needs/73677235007/.

Wilson, John D, and Zach Zimmerman. 2023. "The Era of Flat Power Demand Is Over." Grid Strategies.

Wright, Scott, Alla Raykin, and Laurin McDonald. 2024. "Tricks and Traps of Data Center State Tax Incentives | Tax Notes." TaxNotes. 2024. https://www.taxnotes.com/special-reports/tax-technology/tricks-and-traps-data-center-state-tax-incentives/2023/12/28/7hmb7.

Zhang, Mary. 2024. "United States Data Centers: Top 10 Locations in the USA." Dgtl Infra. April 11, 2024. https://dgtlinfra.com/united-states-data-centers/.

Zullo, Robert. 2023. "Lack of Oversight on Transmission Spending Leads to Higher Electric Bills, Consumer Advocate Says • Virginia Mercury." *Virginia Mercury* (blog). October 11, 2023. https://virginiamercury.com/2023/10/11/lack-of-oversight-on-transmission-spending-leads-to-higher-electric-bills-consumer-advocate-says/.

Appendix

A1. Data Center Cooling Systems

This section summarizes the main types of cooling systems used in data centers. Some of these cooling systems can be used in combination.

Air-Based Cooling
Air cooling is the most traditional and widely used method for cooling data centers. There are different techniques within air-based cooling, but they all rely on moving cool air to absorb heat from the servers. Cold aisle / hot aisle containment entails data center racks arranged in alternating rows. Cold air is directed into the cold aisle in front of the servers, while hot air is exhausted into the hot aisle behind the servers. Computer Room Air Conditioning (CRAC) units are similar to standard air conditioners but designed for data center environments. They circulate cooled air into the server room to lower the temperature of equipment. Air-cooled chillers using condensers can also be used.

Water-Based Cooling
Evaporative cooling systems rely on the natural process of water evaporation to cool air. Air is passed through a water-saturated medium, where the heat evaporates the water and cools the air. The cooled air is then circulated through the data center. It is very energy-efficient and works well in dry, arid climates (like Arizona), although water requirements are significant. It can greatly reduce the need for mechanical cooling. Chilled water systems circulate chilled water through pipes in a data center. The chilled water absorbs heat from the air and then returns to chillers to be cooled again. Chilled water systems are often used in conjunction with air-based systems like CRACs. Cool water from nearby bodies of water or from cooling towers can also be used in chillers, however, water impacts are an issue.

Liquid IT Cooling
Liquid cooling is becoming increasingly adopted as data center power densities rise, and air cooling becomes less efficient in high-performance environments. Liquid coolants have a much higher heat capacity than air, making it more effective at removing heat. Liquid immersion cooling entails servers or components that are directly immersed in a non-conductive liquid. The liquid absorbs heat directly from the hardware, which is then cooled through external heat exchangers. Direct-to-Chip cooling circulates a liquid coolant directly to the hottest parts of the servers, such as the CPUs and GPUs, via a network of small pipes and heat sinks.

Free Cooling
Free cooling leverages outdoor air to cool data centers, using the natural temperature difference between the outside environment and the internal temperature of IT rooms. In locations where temperature differences are predictable, free cooling is easier to implement. Natural free cooling uses the building's thermal inertia and natural ventilation, often taking advantage of cooler nighttime air to pre-cool data centers. Mechanical ventilation may be used to manage air intake and exhaust to maintain optimal conditions (DATA4 2021).

Geothermal Cooling

Geothermal cooling leverages the stable temperature of the earth below the surface. Coolant or air is circulated through underground pipes, where the earth's natural coolness helps lower the temperature before returning it to the data center. Geothermal cooling is highly energy-efficient, although it entails high upfront costs for installation.