The Hidden Regulators:

Public Utility Commissions and AI Governance

Scott Babwah Brennan
& Lama Mohammad

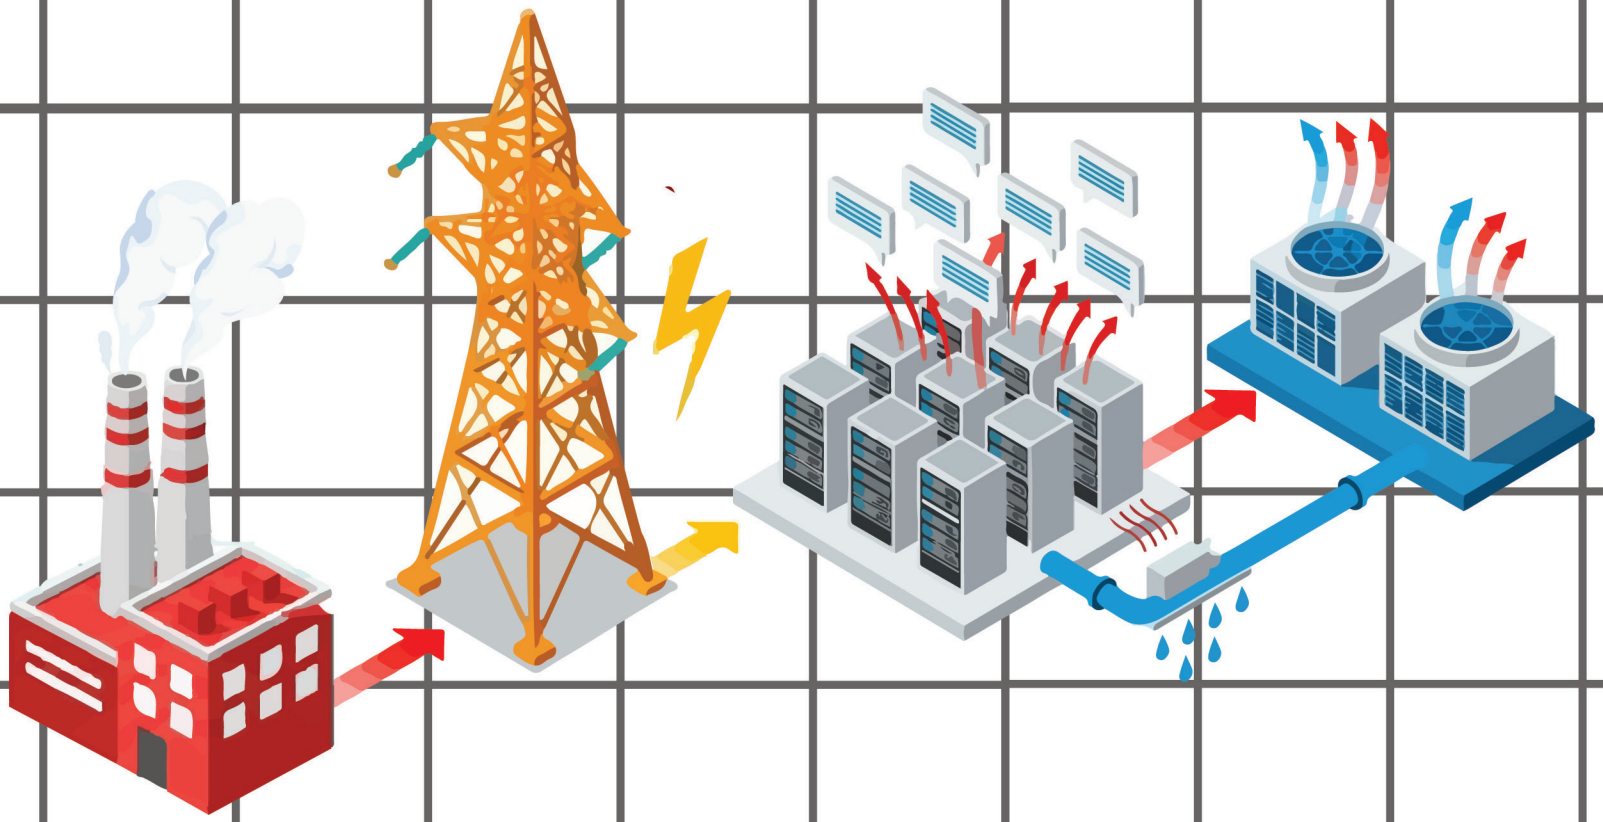

NYU | CENTER ON TECH POLICY

ABOUT THE CENTER ON TECHNOLOGY POLICY

NYU's Center on Technology Policy seeks to craft public policy for a better internet. Utilizing an interdisciplinary academic framework, CTP works to identify knowledge gaps and develop actionable policy frameworks that will enable us to realize the potential benefits of technology while minimizing its harms. By working closely with students and expanding the University's offerings in technology policy analysis, we seek to cultivate and train the field's future practitioners.

ABOUT THE AUTHORS

Scott Babwah Brennen is the director of the Center on Technology Policy at NYU and a consultant on technology policy issues.

Lama Mohammed is a Tech Policy Fellow/Researcher at NYU's Center for Social Media and Politics and a policy research affiliate at the Center on Technology Policy.

ACKNOWLEDGMENTS

We would like to thank Sarah Graham, Aizaz Bokhari, Michael Kleinman, Brian Chen, and Bonnie Brennen for their advice and feedback on the report.

This report was designed by Adrienne Testa.

Executive Summary

Suddenly, talk of data centers is everywhere. As artificial intelligence (AI) companies announce hundreds of billions of dollars in investments, many AI companies and industry analysts are claiming that AI innovation and development are now bottlenecked by a lack of infrastructure. Similarly, tech companies, academics, and policymakers argue that for the United States to maintain its AI dominance, it must significantly increase the number and capacity of data centers.

At the same time, recent research has examined the financial and environmental costs associated with new data center buildouts. Researchers and advocacy organizations are expressing concern that rapidly adding massive amounts of demand to the energy grid not only threatens grid stability but also risks increasing energy costs borne by other customers, harming the environment, and negatively impacting communities. Given these concerns, new data centers are finding growing opposition from advocacy groups, state policymakers, and even voters.

Within this chaotic and contested landscape, state Public Utility Commissions (PUCs) are emerging as influential players in AI regulation through their significant, albeit indirect, oversight of data centers.

PUCs are public bodies that oversee monopolistic utility companies that work to "ensure[] safe, reliable, affordable, and environmentally responsible utility service." They set the prices that regulated utilities can charge customers, approve infrastructure upgrades, and make rules surrounding how utilities operate in states.

PUCs do not directly regulate data centers; instead, they regulate the utilities that serve them. Yet, this role gives PUCs significant influence over how data centers operate and the costs they bear. But to understand and influence AI governance, we must consider not only the infrastructure — such as data centers — that enables AI, but also examine those who control and shape that infrastructure.

This report examines how PUCs are becoming important AI regulators while introducing PUCs and the wider energy regulator system for AI policy audiences.

Through interviews with commissioners and Commission employees across 12 states and a deep dive into how PUCs have been approaching data centers, we outline the specific powers that PUCs have to shape data center and AI policy, discuss key debates involving PUCs, and consider recent efforts to expand PUCs' regulatory authority over data centers.

We recognize five types of authority that PUCs have over data centers: rulemaking, rate setting, permitting, planning, and advising. Across these areas, PUCs have unique power to shape where, how, and if data centers are built.

PUC's indirect oversight of data centers provides an alternative route for the AI policy community to shape AI policy. AI policy professionals can bring a nuanced understanding of the trade-offs associated with data center development to PUC proceedings. That engagement, however, will require greater visibility into PUC processes and hearings, which, though often public, are hidden behind clunky websites, obtuse proceeding names, and arcane rules.

PUCs were first established to be public advocates, protecting consumers and communities from the excesses of powerful utilities. As we navigate a path forward in innovating smart, fair, and safe AI, PUCs can be valuable resources and important partners for the AI policy community.

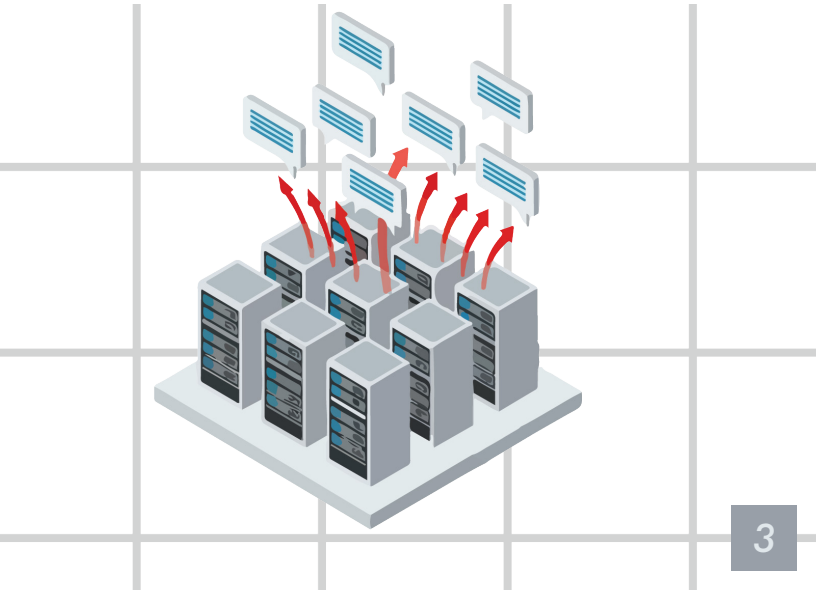

Background

Data Centers

For decades, data centers have served as the backbone of the internet. Housing large computer servers, data-storage systems, and network equipment, data centers store, process, and distribute digital data.

Given that AI companies in the United States compete with those abroad to develop the most advanced AI models and systems, the development of data centers has become both an economic and national security priority for the U.S. government. Just two days into his term, President Trump announced Project Stargate, a joint venture funded by OpenAI, SoftBank, Oracle, and MGX, to develop data centers in the United States.

As artificial intelligence models expand in size and become increasingly integrated into our lives, they require larger and larger amounts of computing capacity. For example, training OpenAI's GPT-4 cost more than $100 million and required approximately fifty gigawatt-hours of electricity, equivalent to the energy used by about 5,000 houses over a year. But asking queries of trained models, called inference, also uses large amounts of energy. One estimate suggested ChatGPT uses nearly 1 GWh each day. The increasing compute demands of both model training and inference are leading to the creation of hyperscale data centers that can have a power demand exceeding 100 megawatts (MW).

Given the power demand required to maintain hyperscale AI data centers, new policy challenges are emerging concerning its impact on states' electric grids, energy supply, job and workforce development opportunities, and the natural resources of local communities housed near data centers. For example, a report commissioned by the Department of Energy estimated that data centers accounted for about 4.4% of total U.S. electricity consumption (or about 176 million MWh) in 2023, which is "equivalent to the average annual consumption of 14 million households." The highest growth projections estimate that data centers could constitute up to 12% of total U.S. energy consumption by 2028.

Public Utility Commissions

Public Utility Commissions (PUC), which go by a range of different names, including Public Service Commissions and State Corporation Commissions, are quasi-judicial state regulators of utilities.[1] Because (aspects of) utility markets in states are run as monopolies, states have created Public Utilities Commissions to protect the interests of the public. These regulators are tasked with overseeing investor-owned utilities and, most importantly, setting the prices that utilities can charge consumers while ensuring that investments in new infrastructure are necessary and beneficial. While PUCs oversee a series of different utilities, including water and electricity systems, our focus here is narrowly on PUCs' regulation of energy utilities.

The State Energy Ecosystem

The scope and powers of a PUC depend on the state—and in particular, the structure of a state's energy system. Initially, electric utilities controlled all elements of the energy system, resulting in "vertically integrated" energy systems. Beginning in the progressive era during the early 20th century, PUCs were established to protect both consumers from monopolistic utilities and the wider grid. In the mid-to-late 1990s, about half of the states partially deregulated or "restructured" their energy systems, allowing for competition in energy generation. In abolishing the monopoly in energy generation, these states both permitted consumers to choose energy suppliers and removed generators from PUC regulation.

Broadly speaking, there are three main parts of a state's energy system: generating the energy, transporting it across the state, and finally, distributing it to consumers.

Generating facilities, or power plants, produce electricity through the burning of fossil fuels like natural

[1] We refer to them here as PUCs or Commissions.

gas or coal, through nuclear fission reactions, or through renewable sources, such as solar panels, wind turbines, or hydropower.

Every state in the U.S. has its own generation facilities; some states also import electricity from other states, and others export electricity to those states. For example, Maryland remains a net importer of energy, primarily from West Virginia, where large amounts of coal have historically been available to supply coal power plants, and Pennsylvania, where a substantial amount of natural gas is also available.

In vertically integrated states, utilities run power plants (along with transmission and distribution infrastructure) under the regulation of PUCs, which ultimately set the prices that utilities can charge. In restructured states, such as Pennsylvania, Maryland, or New Jersey, dedicated energy generation companies, called independent power producers (IPPs), non-utility generators (NUGs), or independent generation companies (IGCs), produce electricity and sell it within wholesale energy markets. In these states, the prices that IPPs can charge are not dictated by PUCs; instead, they are left to the market. The wholesale markets where IPPs sell the energy they generate are operated by Regional Transmission Organizations (RTOs) or, in some cases, Independent System Operators (ISOs). RTOs and ISOs are both regulated by the Federal Government through the Federal Energy Regulatory Commission (FERC), rather than by state PUCs. While the FERC sets basic regulations for how wholesale markets work, the prices vary according to supply and demand.

After energy is generated — and in deregulated states, sold within a wholesale market — it is transported into the grid to regional substations. This is called "transmission" and usually involves moving electricity through special high-voltage lines. Again, in vertically integrated states, the same utility companies that generate electricity also run the transmission lines and substations. In deregulated states, Transmission Owners (TOs) own the transmission lines — in most states, these are actually the same utility companies that handle distribution. The key difference is that other companies, known as RTOs, actually operate the transmission processes. When some states deregulated in the 1990s, utilities held on to their ownership of transmission infrastructure, but RTOs took over operation. For example, PJM Interconnection is the largest RTO, serving 65 million customers across 13 different states in the East Coast and Midwest regions of the country.

Because these systems often involve interstate transmission, the Federal Energy Regulatory Commission (FERC) is the primary economic regulator for RTOs, overseeing tariffs, market rules, and planning. An independent nonprofit organization called the North American Electric Reliability Corporation (NERC) oversees the reliability and security aspects of RTOs, ensuring there are no supply issues.

To complicate matters, RTOs operate in some vertically integrated states (such as Minnesota, Iowa, and Arkansas), and work with the utilities to transmit energy across state lines. In other vertically integrated states, utilities handle transmission themselves (including Florida, North Carolina, and Arizona).

Ultimately, electricity must be distributed from regional substations to consumers. In all states, this final distribution is run by utility companies. Because there is only one set of powerlines to a customer, these utilities are monopolies and are regulated by state Commissions.

For example, in Pennsylvania Talen Energy owns both nuclear and gas power plants. They generate energy and sell it within a wholesale market owned by PJM Interconnection. The energy is fed into large transmission lines owned by companies such as PPL Electric Utilities, but operated by PJM. Finally, a regulated utility company, such as PPL Electric Utilities delivers electricity from substations to customers.

Beyond these organizations, there are a series of additional regulators that play some role in state energy systems, including state departments of environmental protection, and federal regulators, including the Federal Energy Regulatory Commission (FERC), the Nuclear Regulatory Commission (NRC), and finally, the independent nonprofit regulator, the North American Electric Reliability Corporation (NERC). NERC, along with its eight regional sub-organizations, oversees the reliability of the bulk energy system. Less formally, there are also influential national associations or trade groups that do essential policy and guidance work, such as the National Association of Regulatory Utility Commissioners (NARUC) or regional NARUC groups,

such as the Mid-Atlantic Conference of Regulatory Utilities Commissioners (MACRUC), the Organization of PJM states (OPJM), or the Organization of MISO States (OMS).

Public Utility Commission Mandates

Whether in restructured or vertically integrated states, PUCs are primarily economic regulators. One chair of a Commission in an Eastern state said in an interview, "So typically if you talk to a regulator, they're going to describe themselves as an economic regulator, which means that they view themselves as calling balls and strikes and just regulating to the pure mission of rates."

Each time a utility wants to amend the rates it charges consumers, it must file a request with the Commission and go through a "rate case." To support its proposed rate change, a utility will submit a range of documentation. Other stakeholders can also file comments on the proposed rates.

Utility rate schedules include a series of broad "rate classes," such as "residential," "commercial," or "industrial." Each rate class includes a specific cost for each kWH of electricity used, along with other required provisions. In general, the cost for each kWh of electricity used will be lower for larger customers, both due to the advantages of buying in bulk and thanks to various efforts to incentivize commercial development.

Some Commissions, however, interpret their mandate in a broader way. The same Commission chair quoted above also noted that when Commissions see themselves only as "economic regulators:"

> In my opinion, that misses the larger mandate of things that are associated with climate change and the stress of that on the grid and the larger public interest. We're all operating in the same legal structure for the most part. We all have similar mandates. I think it's just a difference of opinion about how you're interpreting that mandate.

Several states have explicitly expanded the Commissions' mandate to include environmental protection and consideration. For example, New Jersey statute now expressly states that:

> The [Public Utilities] board may, after public hearing, upon notice, by order in writing, require any public utility to furnish safe, adequate and proper service, including furnishing and performance of service in a manner that tends to conserve and preserve the quality of the environment and prevent the pollution of the waters, land and air of this State, and including furnishing and performance of service in a manner which preserves and protects the water quality of a public water supply, and to maintain its property and equipment in such condition as to enable it to do so.

Public Utility Commission Structure

Depending on the state, PUCs are composed of between three and seven commissioners. In 10 states, citizens directly elect public utility commissioners. In the remaining 40 states, commissioners are appointed. In 38 of these, the governor appoints commissioners. In Virginia and South Carolina, the state legislatures appoint commissioners.

Republicans control the PUC in each of the 10 states that directly elect Commissions. Many states with appointed commissioners have partisan balance requirements, which require the presence of representatives from both major parties.

Over the last forty years, scholars have debated whether the method of selecting commissioners affects the work of PUCs. While some researchers have suggested that the direct election of commissioners may reduce the risks of regulatory capture by industry, others argue that it opens them up to capture through campaign donations, especially in the wake of the Citizens United decision.

Broadly, researchers have found that states with directly elected commissioners are more responsive to voter preferences, both for good and bad. On the one hand, states with elected commissioners tend to have lower

energy costs. On the other hand, those same states, on average, release as much as 40% more carbon dioxide per capita than states with appointed commissioners (Sautter and Twaite, 2009), as commissioners prioritize lowering prices over minimizing emissions to placate voters. At the same time, at least in the 1980s, states with elected commissioners had lower electric utility bond ratings, making it harder for utilities to raise capital.

PUCs and Data Centers: Central Issues

Public Utility Commissions do not directly regulate data centers. However, in regulating public utilities, PUCs do have significant — if indirect — influence on data centers, which has made PUCs powerful AI regulators. That influence, however, varies widely from state to state, just as the power of PUCs varies from state to state, both according to the structure of the energy market and the ways that states have interpreted the Commission's missions to serve ratepayers.

In the next section, we provide a detailed examination of the ways that Commissions impact data centers. But first, we describe some of the key contemporary debates that involve both Commissions and data centers.

Stressing the Grid

Data centers can draw incredible amounts of power. For example, one Data Center campus currently being built in Wyoming will draw 1.8 gigawatts of power, more than all the households in the state, and about the same as the entire city of Philadelphia.

For now, generation capacity has been increasing commensurate with demand growth. The US Energy Department forecasts that overall energy generation capacity will grow by 2.3% in 2025 and 3.0% in 2026, while overall demand in the U.S. is expected to increase by a similar amount over those two years. However, while data centers currently account for 4.4% of U.S. energy demand, some projections suggest that this share could grow to as much as 12% by 2028. As a result, some are projecting that data center energy demand may outstrip generation capacity.

Even if overall capacity continues to grow in line with demand, the large and consistent demand from data centers can stress the grid infrastructure. Building new generation facilities can take years — or, in the case of nuclear energy, even decades.

Broadly speaking, there are two primary ways energy producers can serve data centers. First, data centers can connect to the grid like any other rate payers. This is referred to as being "in front of the meter." Connecting to the grid ensures that a data center has reliable access to energy, without being wholly dependent on a single generator. Even if one power plant has issues, the data center will still be able to draw electricity from the rest of the grid. Connecting in this manner, however, means that data centers utilize the same infrastructure as the rest of the grid.

Grids are designed to have a greater capacity than is used at any normal time — called "planning reserve margin" or "operating reserve." This is often enough margin that the grid could survive the loss of the single largest generation facility, and still fully supply customers. This way, the grid can protect itself in the event of an unexpected shutdown of a generating facility or an unexpected surge in demand.

Unlike some other large load customers, data centers have very consistent or "flat" demand, requiring approximately the same amount of energy throughout the day. Some large manufacturers or other industrial customers not only have periods of high and low demand, but can also often shift their energy demand to different times of the day. Typically, tariffs for industrial customers include different rates for different times of the day, a mechanism that helps spread demand evenly throughout the day. However, data centers require a substantial amount of energy on a consistent basis. This requirement means that data centers will need backup generators in case there are issues with supply from the grid. These are typically diesel generators, which can be expensive, loud, and highly polluting to operate.

Instead of being in front of the meter, data centers can also connect "behind the meter," connecting directly to a generating facility. The generating station and the data center then negotiate a direct agreement (which may or may not be reviewed by the PUC, as discussed below), whereby a portion of the power plant's generating capacity will be allocated directly to the

data center, bypassing the grid. There are advantages and disadvantages to this. On one hand, data centers get direct, dedicated energy. On the other hand, if there are issues with the power plant or it needs to stop (or reduce) generation for routine maintenance, the data center will need to run backup generators.

However, having a data center connect directly to a generating facility essentially removes that facility — or some portion of it — from the grid, which threatens grid stability. To maintain sufficient reserve margin, utilities may need to build or acquire additional generating capacity. Depending on the market structure, this can also mean that more customers are competing for a limited amount of energy, resulting in higher rates.

Passing Costs on to Other Rate Payers

State legislatures, Commissions, activists, researchers, and community groups have all expressed concern that data centers will pass along added costs to other rate payers.

Importantly, there is mixed evidence that this has occurred in the last decade. On one hand, a recent study by researchers at Lawrence Berkeley National Laboratory, saw that at least on average between 2011 and 2024, residential energy prices did not increase faster than inflation. Moreover, they found that increasing energy demand tended to *decrease* prices paid by retail customers. They note that many have identified the largest contributor to increased residential energy costs "has been distribution and transmission expenditures — often devoted to refurbishment or replacement of existing infrastructure rather than to serve new loads." Increased demand makes it possible to spread these costs over more customers, lowering prices. The authors, however, note that it is not clear if continued demand increase will continue to lead to lower retail energy prices, or if it will ultimately increase prices.

On the other hand, a recent analysis of wholesale energy prices by Bloomberg found large increases in some locations near data centers. Similarly, an analysis by the Union of Concerned Scientists observed that in 2024 transmission upgrades in seven PJM states cost $4.3 billion, which they assume was passed onto consumers. A Carnegie Mellon study estimated there could be an 8% increase in average retail energy costs by 2030.

Indeed, this year, customers in several states have seen significant increases in energy costs. One analysis of rising prices in New Jersey, connects rising prices to simultaneous increase in demand from data centers and transportation electrification, and decrease in capacity due to decarbonization policies.

There are two primary ways that data center build-out can increase energy costs for other customers.

First, in unregulated states, where generating companies sell energy within a wholesale market, the energy costs will vary according to supply and demand. Data centers represent a large and rapid increase in energy demand. While supply as a whole has kept up with demand, that may not always be the case, especially in specific energy markets, given the time required to build new generation facilities. While it is possible in some states to import energy from other states, there are challenges and costs associated with this. Simply put, if data centers begin competing with other consumers for a limited energy supply, prices will increase.

Second, in both deregulated and vertically integrated energy systems, costs associated with building new energy generation capacity or other transmission and distribution infrastructure may be reflected in higher energy costs. In restructured markets, increased demand for energy in comparison to supply may motivate an increase in prices. In vertically integrated systems, some infrastructure upgrades may be passed on to consumers by being included in new rates approved by the Commission. As discussed below, some utilities are now requiring large customers to pay upfront for some of these infrastructure upgrades. However, in some cases, because the utility will retain ownership of the infrastructure, those costs are refunded through lower rates over long periods of times. For example, PG&E's interim Rule 30 in California (see below), includes a scheme where certain data centers pay upfront connection costs, but those costs are eventually refunded.

At the same time, in some cases, the utility is unable to recoup the full costs of infrastructure upgrades, leading to "stranded costs." These stranded costs will ultimately be passed on to other ratepayers. Although there are

now many planned new data centers, if AI demand and/or investment radically declines, many of these proposed data centers may not be completed, robbing utilities of the ability to recoup costs for necessary infrastructure upgrades. This could pass the cost along to other rate payers.

Notably, Commission-approved rates also include set returns on all infrastructure investments, often referred to as the Return on Equity (ROE). Part of the negotiation of new rate cases is setting the ROE for all capital investments, as well as a broader "Return on Rate Base," which is something like the overall profit a utility company can include in rates.

Endangering the Environment

Given the incredible amounts of energy and water used by data centers, there is a growing concern about their potential environmental impact. Advocates have highlighted various environmental concerns related to the impacts of data centers, including the construction of massive buildings that destroy local habitats, strains on municipal water systems, and pollution affecting the air and local ecosystems.

The energy that Data centers consume must come from somewhere. While the proportion of renewable energy in the overall energy mix in the U.S. is increasing, about 60% of the U.S. energy supply still comes from fossil fuels. This means that to meet this rising energy demand, some plants will increase their use of coal and natural gas, thereby increasing air pollution and carbon emissions. Building new nuclear power plants can take years if not decades. While there is much interest in small modular reactors (SMRs) for data centers, there remain significant technological and regulatory challenges to their widespread adoption. Data centers also generate significant heat, necessitating increased commercial ventilation and space cooling, which could lead to higher greenhouse gas emissions.

Since data centers require an uninterrupted energy supply, they typically have a large number of on-site backup generators. While they are usually only run in emergencies, they still release large amounts of environmental pollutants which could lead to more asthma cases.

Furthermore, in some cases, the toll of powering data centers disproportionately impacts communities of color and poorer communities. For example, Elon Musk's xAI data center was built near a predominantly Black community in Memphis that has already been plagued by industrial plants, and, as Politico reported, "leads the state in emergency department visits for asthma." The data center has not yet connected to the utility grid; instead it draws its power from a series of onsite generators that produce a large amount of pollution, and do not have the "pollution controls typically required by federal rules."

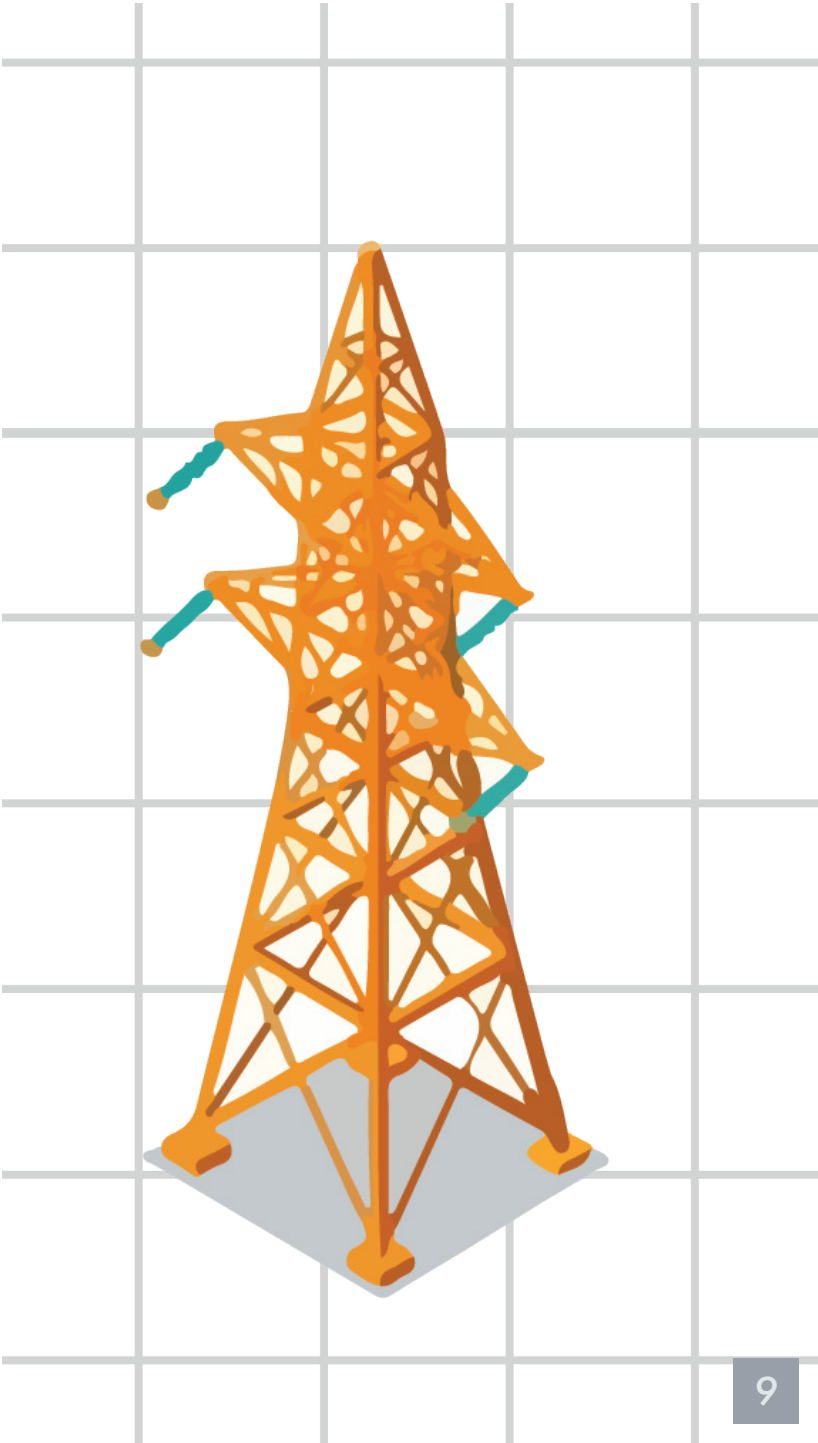

What Do Commissions Do?

Public Utility Commissions do not directly regulate data centers. However, they do have significant indirect authority over data centers. To better understand how PUCs influence data centers — and therefore AI — policy, we catalog the specific ways that PUCs can influence data centers across states.

Notice-and-Comment Rule Making

Commissions across the country create and enact regulations within the scope of the authority granted by state law. PUCs engage in "notice-and-comment rulemaking," where, after providing public notice of a proposed rulemaking, the PUC collects public feedback before implementing a regulatory change.

For many PUCs, the first step in any formal process is to open a "docket." PUCs have quasi-judicial powers, and opening a docket is akin to starting a court proceeding. Opening a docket grants PUCs a suite of capabilities, enabling them to hold hearings, request testimony, and solicit information from utilities and other companies.

For example, on April 3, 2025, Kevin Thompson, chairman of the Arizona Corporation Commission (ACC), opened a new docket:

> To review the existing rate classifications and the possible creation of more transparent rates for data center customers and the public. Other topics could include a review of utility mechanisms being implemented with data center customers, as well as behind-the-meter and in-front-of-the-meter solutions, and user-funded utility scale generation to help large customers such as data centers meet their power needs, while at the same time providing improvements to Arizona's power infrastructure that ultimately benefit the utilities' and other ratepayers through a more robust grid.

Since then, the ACC has collected comments from various stakeholders, including the Data Center Coalition, an industry trade group; the Interwest Energy Alliance, a utility trade group; and Underground Arizona, an advocacy group.

Similarly, on March 27, 2025, the PA PUC Chairman opened a docket related to data centers, holding an *en banc* hearing in late April "to seek information from industry leaders and interested parties concerning interconnection and tariffs for large load customers." It has since collected comments from a range of stakeholders.

While some dockets may be more focused on fact-finding, in many states, they are the first steps in PUCs' broader rulemaking authority. Below, we discuss in detail PUC's abilities to set new tariff rates and requirements; however, PUCs have broader powers to establish administrative rules. There is a wide variation across states in the scope of PUCs' rulemaking authority.

States like California, New York, and Illinois grant PUCs broad authority to craft new regulations.

Depending on the state, these regulations may establish the basic rules that utilities must follow regarding topics such as interconnection, outage reliability standards, and green energy supply options or tariffs.

For example, in August 2025, California PUC (CPUC) opened a new rulemaking to revise the rules around connecting distributed energy resources (DER) under Rule 21. Technically, Rule 21 is part of the master Tariff, but in California, it can be revised independently and, to some extent, functions as a regulatory statute. "Rule 21 is a tariff that describes the interconnection, operating, and metering requirements for certain generating and storage facilities seeking to connect to the electric distribution systems."

DERs are especially important for data centers, which must have significant backup generation capabilities on site in case of grid interruptions. Through rulemaking, the CPUC could influence the technical standards for generating facilities, necessary tests or assessments, and

cost/fee structures, all of which would materially impact data centers.

Once Commissions have instituted rules (or set rates, as described below), they have a series of procedural tools to help them implement new regulation. Most notably, many PUCs require that regulated utilities submit "compliance filings" or "tariff filings," to demonstrate that they are complying with regulation. For many Commissions, these letters may not have to receive full Commission approval.

A handful of states, however, treat these filings more formally, and require a full commenting and review process. Most notably, California uses an "advice letter," system to update or request approval or actions already authorized or required by Commission rules or decisions.

The California PUC distinguishes between three types of advice letters, tiers, 1, 2, and 3. The tier 1 process is reserved for the most minor matters; tier 3 is used for matters with more significant policy impact. For tiers 1 and 2, staff usually can rule without the full Commission. Tier 3 letters, however, must go through a full process.

During summer 2025, the California PUC approved Rule 30 for Pacific Gas & Electric (PG&E) , a major regulated utility. Rule 30 is an interim rule meant to serve while the Commission considered a permanent tariff. It establishes a temporary set of rules and an expedited process for PG&E to get approval to serve large load customers that operate at the transmission level (50-230 KeV) and that will fund 100% of infrastructure costs upfront. In addition to setting basic requirements, the rule allows requests to go through the simpler tier 2 approval process, rather than requiring a tier 3 process.

Rate Setting

Historically, the core power of PUCs has been to approve the prices that utilities charge customers for energy. The most common way they do so is by deciding periodic "rate cases," proceedings where PUCs review changes to the rates that an individual utility charges different types — or "classes" — of customers to protect both the interests of customers and the grid as a whole. However, there are two additional ways that PUCs shape the prices data centers pay for energy: approving special classes for data centers or large load customers, and reviewing bilateral contracts between utilities and data centers.

Bilateral Agreements/Contracts

Over the last decade, some utilities have negotiated individual bilateral agreements or contracts with data centers, rather than serving data centers only under existing rate classes. In some states, PUCs play a key role in overseeing and approving these bilateral agreements. These contracts typically include special energy rates, along with other provisions designed to incentivize data center development and/or mitigate the negative impact of serving data centers on the rest of the grid. In an interview, one policy leader at a Southern PUC said, "There's less of a chance of a disproportionate impact on residential customers from a data center if they operate under a special rate contract" because of the ability to build in additional protections.

While some consumer protections, such as investment recuperation, can be baked into the rate the customer pays for energy, special contracts can also include a wide range of other provisions to help protect utilities and customers from bearing additional costs. The scope of these provisions is often set by state statute. For example, in Arkansas, the ability of utilities to enter into special contracts is restricted by the "Commission Rules and Regulations Governing Promotional Practices of Electric and Gas Public Utilities." The rules not only articulate specific guidelines for promotional contracts, but they also stress that special contracts must still "benefit ratepayers in the aggregate."

These provisions are seen in both bilateral contracts and special rate classes. We discuss some of the most common provisions below.

Approving Special Rate Classes

In the last few years, some state utilities have begun to move away from bilateral contracts with large data centers, creating — or attempting to make — new rate classes that apply widely to all hyperscale data centers or other customers that use a vast amount of energy. Last year, AEP Ohio received approval from the Public Utilities Commission of Ohio to create the "Data Center Tariff," which is limited to data centers drawing more

than 25 MW of electricity. Utilities in a handful of other states are now trying to get data center rate classes approved. For example, two major utilities in Florida, Florida Power & Light and Duke Energy, are both in the process of creating a new rate class.

Additionally, over the past year, several state legislatures, such as those in Oregon, Minnesota, and Maryland, have enacted laws that require public utility Commissions to oversee the creation of new rate classes.

Creating new rate classes can take several years. One employee at a Southern PUC said "I guess theoretically it would be possible to create a separate rate class, and we have three commissioners that the governor appoints, so I don't want to speak for them at all, but I would highly doubt that a separate rate class would ever get established for, it's kind of a big deal to establish a separate rate class." Even still, utilities in 22 states have already approved new rate classes for large load customers.

In an interview, the chair of a PUC described how they recently began the process of developing a model tariff for large load customers. They stressed that the new tariff would not only apply to data centers, but to any large load customer. "What does it matter if it's a steel producer that's using 50 MW or if it's a data center using 50 MW, if it has the same load profiles, in terms of interruptability and all those things, and size and frequencies, is it running 24/7?...we need to be agnostic about a lot of things here and manage the grid. And that's what we're trying to do."

But special rate classes can take a range of forms. The Chair of a Commission in a state in New England said that while their state had not yet done special contract rates with data centers, instead, "we went with the approach of creating what we called an economic development tariff, which encourages existing businesses to expand, and it's more of an attempt to retain businesses rather than attract new ones. So it gives them sort of a discount on their rates."

Common Provisions in Bilateral Contracts and Special Rate Classes

Both special rate classes and bilateral contracts can contain a range of provisions beyond the base rate that a customer pays. A report from the Energy Futures Group analyzes new large-consumer contracts and special rate classes, identifying a set of distinct provisions. While these contractual terms are meant to address concerns about cost shifting — we do not yet have a comprehensive analysis analyzing their effect.

First, contracts can require data centers to directly pay for the infrastructure needed to connect them to the grid. Broadly speaking, there are two ways this can work. Contracts can require that large load customers pay a Contribution in Aid of Construction (CIAC) — a one-time upfront payment to cover all necessary infrastructure development, such as new substations or transmission lines. Alternatively, rather than paying all at once upfront, utilities can include an additional fee in each billing cycle specifically for equipment, often referred to as a "dedicated facilities charge." Like most loans, this charge will contain an interest payment.

Second, special agreements can set long contract lengths. Unlike residential customers, who can discontinue service at any time, data centers typically enter into long-term contracts for electricity that last 10, 15, or even 20 years. These long contracts help ensure that utilities can recoup building costs while making reliable plans for energy demand into the future.

Third, utilities can include a minimum monthly payment requirement — sometimes referred to as a "take-or-pay" clause. Unlike residential customers, who simply pay for the kilowatt-hours of electricity they consume, utilities need predictable revenue from large industrial loads to recover the costs of substations, feeders, and other dedicated facilities, and to plan system capacity reliably.

The most common approach ties the minimum payment to a percentage of the customer's contracted maximum demand. For example, if a data center contracts for 100 MW of capacity under a special tariff, the agreement may set a 90% minimum. In practice, that means if the facility peaks at 95 MW one month, it pays for 95 MW; if it only peaks at 85 MW the next month,

it must still pay for 90 MW. The Ohio Public Utilities Commission, for instance, recently approved this sort of agreement as part of a new rate class for AEP, where it applied a 90% ratchet.

Notably, new data centers may take months or even years to reach their full capacity. Contracts can include schedules, where the data center agrees to pay an increasing percentage each month over a set period.

Another way to handle monthly minimums is through a "demand ratchet." In this structure, the minimum payment will be a set percentage of the highest amount the data center used during a specific period — usually the past 12 months. For example, under a 90% ratchet, if a data center used 100 MW in January and then 80 MW in February, they'll still pay for 90 MW in February. But then, if they use 110MW in March and 80MW in April, they'll pay 99MW (90% of 110MW) for April. The key feature is that each new maximum sets an ever-changing floor.

Another provision included in contracts is "Time-of-Use (ToU) rates." In the U.S., most electricity customers pay the same rate for electricity throughout the day. However, in many other countries, charging customers more for energy during peak hours is very normal. In some contracts, utilities charge data centers a higher rate for the energy they use during peak hours. Of course, unlike many other large energy users, data centers have very regular electricity loads because they operate continuously throughout the day. This means that data centers have little ability to prioritize certain times of day.

Other contract terms include a set amount of time that service can be disrupted each year, or "demand shedding" — ways that data centers can reduce their maximum capacity over time, and "study fees", such as seen in AEP Ohio's new data center tariff, which requires new customers larger than 100MW to pay a study fee of as much as $100,000.

Permitting

In many states, before a utility builds a new facility or piece of infrastructure, the PUC must issue permits. Typically, new facilities require multiple permits from various regulators, including those for air, water, and other environmental purposes. In some states, such as Maryland or Arkansas, the PUC acts as a "clearing house," a central point for collecting all the necessary permits from other agencies. In other states, different agencies play that role. For example, in Florida, the environmental regulator serves in this capacity.

In most states, before a regulated utility can build a new facility or infrastructure upgrade, the PUC must issue a Certificate of Public Convenience and Necessity (CPCN) (or in some states a Certificate of Need (CoN)). One employee at a Southeastern PUC explained that the law in the state, "says it's illegal for any electric utility to construct anything without Commission approval other than ordinary extensions in the normal course of business, that is, running a new distribution line to a newly developed subdivision."

Importantly, data centers themselves do not require CPCNs: instead, in many states, the infrastructure necessary to support or serve the data center needs to be approved through a CPCN, including a generating facility, substation, or transmission line. For example, while the Montana PSC does not issue CPCNs for data centers, a new law passed this year now requires the Commission to issue CPCNs for any transmission line that is larger than 69 KV.

Although there are differences across states, CPCN approval typically hinges on whether the infrastructure is necessary, what impact it is likely to have on the grid and customer rates, and whether it aligns with both the state's energy policy and the utility's strategic plan (see the next section). Pennsylvania statutes state: "A certificate of public convenience shall be granted by order of the commission, only if the commission shall find or determine that the granting of such certificate is necessary or proper for the service, accommodation, convenience, or safety of the public."

Yet, in some states, the Commission must also consider the environmental impact when issuing a CPCN. For example, in Minnesota, the determination of whether to grant the CoN is made if the facility is in the public interest and "will provide benefits to society in a manner compatible with protecting the natural and socioeconomic environments, including human health." The Chair of one Commission in a MidAtlantic state said that after the state legislature included environmental review in the Commission's mandate, "So

when we analyze something that's brought before us, it's not just the dollars and cents as it was historically back in the good old days; it's now a multifaceted review that includes environmental impacts of whatever the utility or entity that's before us might do."

Over the past several years, a debate has arisen regarding the necessity for Commissions to issue CPCNs for backup generators at hyperscale data centers. Because data centers must run continuously, they must have backup generating capacity in case of either a scheduled or unscheduled outage. Due to the power requirements of data centers, these backup generators must have a large capacity. Often, they are diesel generators, which, when run, release a substantial amount of pollution.

Many regulated states require that Commissions grant CPCN or CoN approval for generating facilities over a specific size. Yet, because data centers have chosen to purchase many small generators — each smaller than the cutoff — rather than build a single large plant, there is disagreement over whether Commissions must issue permits.

Last year, a data center in Maryland requested an exemption from the CPCN process for its backup diesel generators. Maryland law requires CPCNs for generating facilities larger than 2MW. The facility was planned to have 168 generators, each with a capacity of less than 2MW. Therefore, the data center argued that because each generator was under the 2MW limit, it could be exempt from review. After the PSC ruled they did not qualify for an exemption, but had to go through the full CPCN process, the Maryland legislature passed a law amending the statute. The new law exempts generating facilities at data centers (or other large load customers) that are explicitly used for power outages and are not otherwise connected to the grid.

Similarly, this year, in March 2025, Amazon lost a bid to be exempt from the CoN process for 250 diesel generators at a proposed facility in Minnesota. Minnesota law exempts "small power production facilities… having a combined capacity at a single site of less than 80,000 kilowatts." While combined, Amazon's generating capacity exceeds 80,000 kW; they argued that because it wasn't connected to the grid, it should be exempt. Notably, Amazon could still go through the CoN process, but it will take time. In the meantime, Amazon has indicated that it will attempt to convince the legislature to change the law, just like what happened in Maryland.

Beyond CPCNs, at least one state requires a more direct data center permit process. While Montana does not require CPCNs for data centers, it has a unique provision in place. Existing law requires any customer larger than 5MW to:

> Come to the commission and show that service to this new customer will not harm the other existing retail customers. So I think that's a big stopgap that Montana has as far as just a check at the beginning of the process, how is this going to work? And we need to make sure it's not going to harm anybody else.

Similarly, as part of a new hyperscale rate class approved in early 2025, the Georgia PSC must now approve any data center larger than 100MW.

Planning

In many states, Commissions play a key role in strategic energy planning. These documents often set general frameworks for how utilities will approve data centers.

In most vertically integrated states, utilities are required to develop and submit regular strategic plans. These are commonly referred to as "Integrated Resource Plans" (IRPs) and are long-term (10-20 year) plans that model energy demand, describe how the utility will meet that demand, and explain how the utility will ensure it has adequate resources. In some states, these plans must also outline how the utility will meet other state energy priorities, such as reducing environmental impact or achieving clean energy goals.

IRPs are not binding documents, but instead forecasts or projections, or, as one employee at a Western state Commission described it, "a guesstimation as to what you're going to need to serve future loads." Moreover, "generally those guesstimations change. And here recently, generally they change every four years depending on the administration and who the president

is." However, "it does sort of give you sort of a glide path as to what you can expect from CPCN applications and what you might expect from general rate cases."

PUCs often have the authority to review and approve these strategic plans. The exact degree of oversight and involvement varies by state. While some Commissions are deeply involved in crafting these strategic plans, in other states, utilities are simply required to inform Commissions when they complete plans.

Given the significant increase in energy demand associated with serving hyperscale data centers, strategic plans can play a crucial role in helping utilities plan to meet that demand. Figure 1 shows a demand forecast from a Virginia Power and Electric 2024 IRP.

Figure 2.1.7: Customer Contracts Executed, as of July 2024

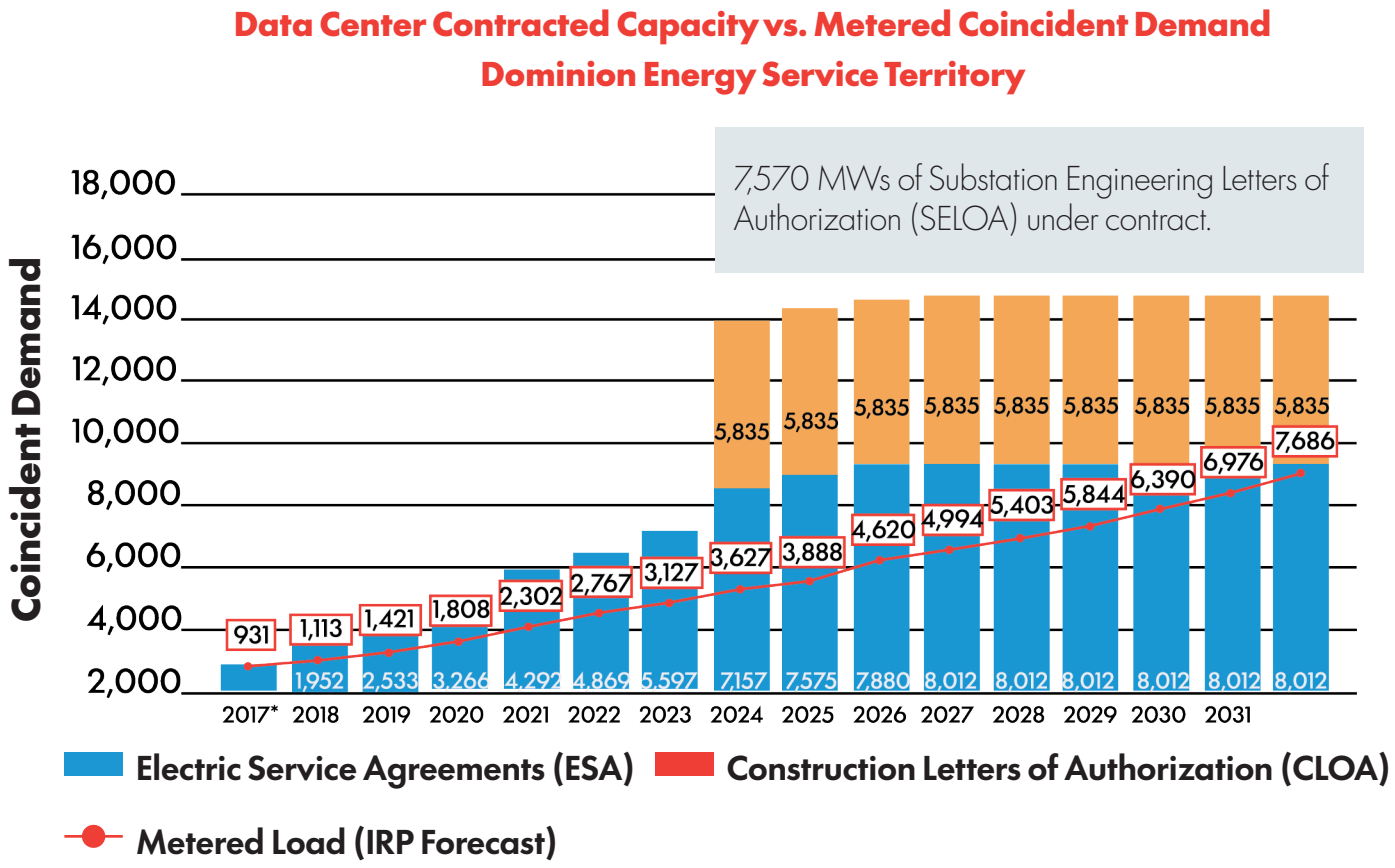

Note: The Company did not review ESAs prior to 2018 and assumed ESAs were equal to demand in 2017. Actual ESA totals will be higher than this assumption.

Figure 1: From the Virginia Power and Electric Plan IRP 2024 (VPE is a subsidiary of Dominion Energy)

Strategic plans often play a direct role in permitting approvals. Because CPCNs hinge on the "necessity" of a new development, being included in a strategic plan can support the case that a new facility or piece of infrastructure is essential to the health and upkeep of the grid.

In deregulated states, utilities must still file planning documents with Commissions that cover the regulated aspects

of energy systems. For example, in states where Commissions only oversee distribution, utilities must file a Distribution System Plan (DSP) or a Distribution System Implementation Plan (DSIP).

In some deregulated states, holistic strategic planning is left to the PUC. For example, NJ's BPU collaborates on an Energy Master Plan (EMP), which is legally required at least every 10 years. This guidance document often focuses on creating an energy roadmap to achieve climate goals and includes other key agencies, such as the Department of Environmental Protection.

Advising and Convening

Beyond their formal powers, many Commissions also play an important advisory role for other regulators. In many states, Commissions regularly provide expert commentary to legislators, executive agencies, or federal regulators about data centers. A senior staff member at one East Coast Commission described a specific office at the Commission that leads in this "critically important role" to "represent [the state's] interest, in our case, with PJM and FERC in a way that is consistent with the policy of our state government."

PUCs are one part of large and complicated energy ecosystems (see above), and Commissions frequently weigh in on matters before other parts of the energy system. While these comments do not carry the same force of law as official decisions, as recognized experts on state energy policy, Commissions can play a crucial and influential role in shaping data center policy.

For example, at the end of 2024, the Maryland PSC was asked by the legislature to advise on a proposal to permit a large data center to co-locate at Maryland's only nuclear power plant, Calvert Cliffs. The data center would have been "behind the meter," meaning it would not be connected to the grid, but draw power directly from the nuclear station. Due to its limited mandate in a deregulated state, the Maryland PSC would technically not oversee or approve this arrangement. However, the legislature still wanted the PSC's expert insight.

In their report to the legislature on the matter, the Maryland PSC concluded that a behind-the-meter large data center at Calvert Cliffs "could present significant risks to reliability in Maryland and the PJM region" and may result in "an increase in the marginal price of electricity." This is because the data center would essentially receive first dibs on the energy output from the nuclear station, effectively removing the station from the grid. The legislature appears to have listened to their advice, which aligned with FERC's comments, and the project has not proceeded.

Another example is that in 2024, FERC opened a docket to consider a PJM tariff for colocated loads. FERC held a technical conference at the end of 2024 and collected formal comments to assess the tariff. FERC received several hundred comments, including from a handful of state Commissions from Pennsylvania, Maryland, and New Jersey. In their comments, the Pennsylvania Commission argued that:

> PJM's Tariff must be amended to include provisions that ensure that participants in a co-located load arrangement are held financially responsible commensurate with the extent to which they receive benefits from the transmission grid according to the cost-causation and beneficiary pays principles.

While the comments of Pennsylvania's PSC have no force of law, FERC ruled in February that the PJM tariff did not state terms for co-location arrangements. PJM responded in March.

In addition to submitting or supplying commentary or expertise, some PUCs also serve as conveners for various members of the energy ecosystem. For example, a handful of states have recently held "technical conferences" to explore data centers, co-location, or related topics. For instance, in December 2024, the Virginia SSC held a technical conference "to explore current and anticipated future challenges presented by providing electric service to large-use, hyperscale customers," and then allowed additional public comments to be submitted for the next month. Similarly, the California PUC held a workshop on "Pending Loads."

Recent Legislative Expansions to PSC Power

The previous section reviews the powers that Commissions have across the country concerning data centers. These powers, however, are not static; rather, legislatures have been working to expand or contract the powers that PUCs have over data centers.

By our count, there have been more than a dozen new enacted laws that directly concern Commissions.

Broadly speaking, these laws fall into three categories: rate classes, process revisions, and study requirements.

New Rate Classification

This year, four states — Minnesota, Oregon, Maryland, and Texas — passed laws that explicitly direct the PUC to create a new tariff or rate class for large load users. Three of these four new laws offer few specifics as to what the new rate class would include, beyond ensuring that the tariffs are written in such a way as to "mitigate the risk of… shifting the costs" to other customers.

The new law in Texas, however, establishes a handful of specific provisions that must be included in the new rate class, including that large load customers pay for all interconnection costs directly, pay a study fee of at least $100,000, and disclose certain information about backup generation.

Several other states proposed similar legislation, but did not enact it this term.

Process Revisions

Beyond revising rate classifications, several states have made minor revisions to other processes through which the PUC exercises regulatory oversight over data centers. None of these laws introduces significant changes. A law in Montana requires that all transmission lines of a specific size be approved through a CPCN. A law in West Virginia establishes a process to expedite PUC oversight and exempt data centers from certain local regulations.

Indiana passed two laws. One creates a process where utilities can request approval from the Indiana Utility Regulatory Commission for an expedited generation resource (EGR) plan, helping to fast-track new generation facilities. But it sets conditions of approval, including a minimum 80% CIAC payment. Another law requires that any proposed data centers owned by a foreign company be approved by the Commission first.

Finally, Virginia Governor Glen Youngkin vetoed a bill that would have revised and expanded the IRP process to include discussion of reliability, affordability, decarbonization, and economic development.

New studies

Finally, two states now require that PUCs undertake new reviews or studies of data centers. Legislators in New Jersey enacted a law that requires the Board of Public Utilities to study the impact of data centers on the stability of the grid and on electricity costs. Other states proposed, but have not yet required, similar studies.

A law enacted in Virginia requires the SSC to conduct a biennial review to assess if the "Utility is using rates, tolls, charges, or schedules that contain reasonable classifications of customers. In making this determination, the Commission shall consider whether new or separate customer classifications are reasonable." Notably, the SSC must conduct a similar review of any newly proposed rate in a standard rate case.

Conclusions & Recommendations

As data centers become more central to AI innovation and development, the AI policy community should recognize PUCs as important state regulators of AI. Whether you believe data centers should be fast-tracked, slowed down, or just studied, Commissions represent an essential target for policy engagement.

At the same time, in a debate increasingly framed by AI accelerationists willing to pursue expansion at all costs and activists narrowly focused on either environmental protection or preventing data centers from being located in specific communities, AI policy professionals can offer a more nuanced perspective. In contrast to Manichean descriptions of data centers, we require a more rigorous accounting of the trade-offs inherent in building new infrastructure, along with specific recommendations for mitigating the negatives and supporting the positives.

It can be both true that we must build more data centers to continue the pace of AI innovation and development, and essential that we work carefully to protect ratepayers, the grid, and the environment from the costs of doing so. That serving massive data centers has harms does not mean we don't need more compute capacity; it means we have to be smarter about how we understand and regulate it.

However, increasing and diversifying engagement with utility Commissions will require more visibility into processes and hearings. Most Commission comments, hearings, and reports are publicly available. However, the data we need to track and understand what is happening across the country is often hidden in poorly designed public archives. While lobbyists can afford costly tracking services or dedicated personnel to monitor Commission proceedings, we need tools to make it easier for the broader AI policy community to follow the work of Commissions across states in real-time.

Utility Commissions were established during the Progressive Era as public advocates. They were designed to serve as bulwarks against public exploitation by utility monopolies granted the power to control the energy system that was rapidly becoming essential to both economic development and public well-being. As we negotiate a path forward on data center and AI regulation that maximizes the benefits and minimizes the harms, Commissions can be valuable partners.