State Policies to Mitigate Data Centers
January 2025

Introduction

The growth of data center electricity demand is quickly threatening to overwhelm climate and environmental goals, and potentially poses a risk to other ratepayers and electric utilities. This model state policy guide provides tangible options for state legislatures and/or public utility commission (PUC) rulemakings. These policies are not meant to stem the growth of data centers or any other specific end-use, but rather ensure that data centers do no harm to environmental or equity outcomes.

Policy Rationale

Data centers are one of the fastest growing segments of electricity consumption in the United States, with estimates that they may grow to 11 percent or more of US electricity consumption by 2030, on top of building, transportation, and industrial electrification. Data centers are able to come online relatively quickly and can be singularly large electricity consumers, with demands reaching hundreds of megawatts at a single facility, or the equivalent of hundreds of thousands of homes. In addition, data centers tend to cluster to take advantage of both infrastructure such as fiber and access to energy, and reduce data transmission costs. As a consequence of both clustering and sheer size, data centers can be highly impactful on local energy systems over short time frames, and strain utility decisions.

Traditional and AI-based data centers generally require uninterrupted power, meaning that not only do they seek low-cost reliable power, they also use substantial on-site backup generation, often in the form of dozens of large-scale diesel generators. These generators not only undergo testing, but may be required to run for extended periods under extreme weather conditions, creating local emissions hazards.

Today, the largest traditional and AI-based data center companies are the so-called “hyperscalers” of Amazon, Meta, Google, Microsoft, and Apple, who own and operate their own massive data centers. Other companies, like Digital Realty, CyrusOne, Vantage, and QTS sell “wholesale” data center space to hyperscalers and other technology companies.

Highly specialized data centers may also be used to “mine” proof-of-work cryptocurrencies, such as bitcoin. These facilities may also be hundreds of megawatts in scale, and can consume extraordinary amounts of energy, based on the prevailing cost of energy and the profitability of cryptocurrencies. These facilities differ from traditional and AI-based data centers in that (a) they are financially motivated to come online as quickly as possible because the value of mining bitcoin degrades over time, (b) their processors can be extremely noisy, to the detriment of localities, (c) their interest in staying online is entirely a function of profitability, and therefore they are prone to shutting down or moving. Finally, cryptocurrency mining operations are “curtailable” in that they may be able to shut down during periods of grid stress.

Sierra Club examined many of the issues associated with traditional data centers in the paper Demanding Better (2024), and associated principles document. In 2022, we collaborated with Earthjustice for a ground-up examination of cryptocurrency mining operations in The Energy Bomb.

For our purposes, as of late 2024, the hyperscaler technology companies have relatively strong climate and Environmental, Social, and Governance targets (ESG), potentially allowing us to work with those companies to improve local legislation on renewable energy, and siting. However, they are strongly motivated by growth, and may experience tension with those environmental goals as they expand. Cryptocurrency mining operations have only a passing or non-interest in environmental issues. The policy recommendations here are focused on traditional and AI-based data centers.

The fastest growing states for traditional and AI data centers are VA (up to 26 GW total potential), TX (growing to 7 GW), GA, AZ, NV, and CA (all growing to 4+ GW), OR, IA, OH, and IL (all growing to 3+ GW).

Policy Guidelines: Data Center Energy Use and Supply

The bulk of our recommendations are focused on energy issues, which have an outsized impact on Sierra Club's 2030 Goals. The basis of these guidelines can generally be found in Demanding Better, and are focused on **Renewable Energy Supply**, **Ratepayer Protection**, and **Transparency & Planning**.

Some of the policy options offered here are relatively novel, while others are extensions of existing practice. For ease of use, these policies are labeled by their relative difficulty as asks of the industry or regulators, all else held equal. **Fruit on the ground** 🥔 are straightforward, and may appeal to broad constituencies; **Low hanging fruit** 🍎 are known issue areas for regulators and utilities, and still represent interests; finally, **Stretch goals** 🖐 are novel approaches that may require iteration to get right, but have the potential to dramatically improve outcomes.

🥔 Fruit on the ground

- Abolish incentive rates
- Tariffs before infrastructure
- Cumulative impact air permits
- Utility growth transparency
- PUC reports to legislature

🍎 Low hanging fruit

- Large load tariffs
- Utility IRP
- State reporting
- PUE standards
- Demand response incentives

🖐 Stretch goals

- 24/7 procurement
- 24/7 tariffs
- Bolster RPS / CES
- Battery backup

Renewable Energy Supply

1. 🖐**24/7 Procurement: require that data centers procure or subscribe to locally-deliverable, additional, and zero-emissions renewable energy in all hours of the year ("24/7").** Data centers should be required to focus procurement efforts on zero emissions (preferably renewable) energy that serves their requirements in all hours of the year, without resorting to substantial new gas generation, or keeping existing coal generation online, to serve non-renewable hours.

    a. "Procure" may be applicable for data centers that operate in restructured markets (i.e. ISOs) or have retail or wholesale choice; "subscribe" is applicable to data centers that operate in a vertically-integrated utility structure

    b. "Locally-deliverable" may be defined electrically (i.e. balancing authority, regional transmission authority region, or based on transmission rights);

    c. "Additional" may be defined as new or recently built renewable energy that is not already identified for another use, such as procured for retail customers;

    d. "Zero-emissions energy" may be modified for local circumstances, but is meant to incentivize new renewable energy;

    e. "In all hours of the year" is meant to ensure that renewable energy is procured to serve grid requirements in all hours of the year, not just when wind and solar are readily available, providing an incentive to storage.

    Implementation: This policy may require some amount of flexibility in implementation, or a ramp-in period where an increasing fraction of energy is served by resources that meet the above criteria for existing data centers, or even phased in over time for new data centers. An additional form of flexibility may be to require a certain fraction of hours met through this standard, and/or flexibility in the resources that can meet the standard, including demand management, or non-renewable zero-emissions technologies.

    This policy option may require advanced energy consumption and generation tracking by the state or utility regulators. It may require data centers to submit load profiles and generation profiles for resources in their portfolio to demonstrate compliance. Alternatively, compliance with this requirement may require the development and/or deployment of time-based energy attribute credits (T-EACs), the tracking of energy consumption and generation on an hourly basis, at least for applicable entities and their generation portfolios.

    Compliance: For a data center-based compliance mechanism, enforcement and compliance may be considered through state public utility commissions (under an RPS-styled rule), state air agencies (under a climate rubric), or treasury (for tax-based incentives). Rules should include annual reporting requirements and/or audits.

    Proponents should consider if the policy is a compliance requirement with financial penalties or offered as an incentive with tax and/or utility rate treatment.

Additional considerations: (a) alternative compliance payments or offsets should be rigorously avoided; (b) proponents should be aware of the risk of "race-to-the bottom" leakage to other states.

This is a policy that may become increasingly more expensive for the entities subject to the rule as compliance stringency increases. It is critical that the costs of the policy are not unduly borne by other customers.

2. 🖐**24/7 Tariffs: require that utilities that serve data centers offer rigorous clean energy transition tariffs for locally-deliverable, additional, and zero-emissions energy in all hours of the year.** Utilities in non-restructured markets (i.e. have vertically integrated control over generation, transmission, and distribution) should be required to offer services to meet (1), above.

   Implementation: This policy can be implemented as a 24/7 carbon-free energy transition tariff, details and considerations of which are described in this paper from Regulatory Assistance Project. Key considerations include if the utility is charged with identifying and procuring sufficient resources to meet the requirements of the tariff, or if customers are able to identify resources that will work for their requirements (i.e. "bring your own capacity"), recognizing that more advanced customers may be able to identify resources, while smaller data center customers may need utilities to arrange all power supply contracts.

   The most efficient way to implement this policy may be to require utilities to create a subscription-based green tariff, where specific customer classes can (or are required to) subscribe to the tariff, which is organized on their behalf by the utility, as a class. In doing so, the utility can pull together a portfolio that works for the entire class, rather than boutique options for individual customers.

   Compliance: Enforcement of this provision can be required through state utility regulators. States with substantial presence of non-regulated utilities (i.e. municipal, public, or cooperative utilities) may need to look for mechanisms that ensure that utilities serving data centers are covered.

   Additional considerations: Numerous utilities have "green tariffs" which vary in effectiveness from those that allow utilities to purchase unbundled RECs to those that provide rigorous local, incremental renewable energy for purchasers. It is critical that energy resources used to serve a green tariff are not simply diverted from other customers, leaving other customers with a fossil-heavy portfolio but are additional. In addition, tariffs must be designed to protect other rate classes (see "ratepayer protection") below.

3. 🖐**Bolster RPS/CES: Revisit mandatory Renewable Portfolio Standards or Clean Energy Standards.** RPS and CES provide avenues to ensure that the entire electricity system of a state or region is becoming increasingly cleaner, and can be substantially more cost effective (and climate effective) than boutique tariffs for individual customer classes. As of 2023, 29 states had RPS, and 15 states have established a long-term CES, according to

LBL. Of those, however, just 1/3 of RPS require half or more of energy to be served by renewables by 2030.

A more rigorous RPS or CES may prevent utilities from building gas and retaining coal to serve new data centers, and ensures that all customers see the benefits of cleaner energy.

Implementation: Numerous resources provide background on RPS and CES and their implementation and considerations including from C2ES, NCSL, and NREL.

Ratepayer Protection

4. 🍎**Large load tariffs: require utilities to implement large-load tariffs with front-loaded costs or commitments.** Utilities are increasingly proposing to build (or actually building) large infrastructure to serve new large load customers, but the costs of those investments may be borne disproportionately by existing (i.e. residential) customers if those large loads do not transpire, or if unfair rate allocation principals are used. In 2024, several utilities (including AEP in OH and IN) extended proposals for new data center rate tariffs that would require firm, long-term commitments from new customers. Features of a large-load tariff may include:

    a. Minimum contract terms and premature exit or termination fees to ensure that new large customers are committed to a long-term relationship with the utility that will help pay down the costs of new infrastructure;

    b. Minimum charge based on anticipated or previously billed use that ensures that infrastructure to serve high-use baseload customers are paid for, even if those customers temporarily shut down;

    c. A requirement for collateral, escrow, or demonstration of financial viability to ensure that anticipated customers are backed by real demand and are able to pay costs.

    d. Incentives for demand management, renewable behind-the-meter generation, or utility use of behind-the-meter storage (see "backup generation and storage", below).

    Implementation: Tariffs are typically proposed by utilities and implemented before state utility regulators.

    Compliance: Legislation could require a utility commission to open proceedings and/or establish tariffs with specific guidelines.

    Additional considerations: Policy and regulators should work to ensure that rate structures do not bypass good (or low emissions) utility planning, or allow for utilities to "gold-plate" unnecessary projects. Policies meant to protect ratepayers should be considered in tandem with policies meant to increase renewable energy use by large load customers and/or utilities.

5. 🥔**Absolish incentive rates: Remove commercial or industrial incentive rates.** Some utilities have implemented incentive rate structures to drive new commercial or industrial customers to their service territories. Historically, incentive rates have been seen by some utilities as a way of expanding economic development for large factories, or dispersing fixed operational costs over a larger class of customers. However, in the face of rapid load growth and new infrastructure requirements, these discount or incentive rates may have the effect of shifting costs of new infrastructure to existing ratepayers.

   Implementation and Compliance: See (4), above.

6. 🥔**Tariffs before infrastructure: PUCs must require utilities to seek and obtain approval for tariffs for new large load customers *before* they incorporate load growth projections into new generation resource approval requests.** Commissions should not authorize utility investment in new resources that are needed primarily or in significant part to serve new large load customers until and unless they have approved a tariff that makes clear how the costs of this new generation will be allocated between existing ratepayers and the new large load customer. The alternative is major investment in resources on the basis of undifferentiated "load growth" without guarantees that those rate-based costs will fall primarily on the large load customers that cause the additional cost.

   This process is meant to reduce the risk of excess infrastructure, or rushed infrastructure, on smaller-scale customers. The establishment of a tariff structure provides clarity on the responsibility of large load customers to the utility, and protection for other ratepayers.

   Implementation: PUCs can establish a practice, or a formal finding, that utilities may not seek resource approval requests for large-scale new customers (i.e. above 50 MW) without having established an approved tariff for the customer. Legislation can be used to require the PUC to establish these rules.

Backup Generation and Storage

7. 🖐**Battery backup: require data centers to use battery storage for on-site backup generation, and maximize on-site solar.** Data centers deploy enormous amounts of backup generation to guarantee uptime even under adverse conditions. Typically, these backup generators are diesel or gas-fired turbines (or experimentally, fuel cells). Requiring a certain (large) fraction of the backup generation to be comprised of battery storage reduces local emissions of NOx, SO_2, and PM, and may allow the data center to accommodate a larger fraction of variable renewable energy if the system is configured appropriately (i.e. a 300 MW data center may be able to absorb up to 500 MW of energy during daylight hours, and store some of that energy into the evening). Data centers that allow some of their on-site storage to be used as a system resource can provide an increased benefit for their utility as it integrates increasing amounts of renewable energy.

   The duration of outages that are anticipated by on-site backup generation may exceed the duration of readily available battery storage technologies (*e.g.* a multi-day storm outage).

Therefore, data centers may be limited on the amount of backup that can be reasonably replaced by storage. Backup generation may, as a fallback, be served by a combination of batteries and a limited amount of backup generation fully equipped with controls, such as selective catalytic reduction, to mitigate emissions. In addition, data centers will need to develop mechanisms of ensuring that backup storage is fully charged under emergency conditions to maintain operations.

Implementation: This policy could either act as a requirement, incentive structure, or combination. A policy might require that a certain fraction of backup generation be composed of battery storage, and then offer an incentive for storage offered as a system resource (i.e. operated or partially operated by the utility).

Compliance: State or air district rules for air permits may be modified to preclude large concentrations of on-site generators, or require mitigations including substantial on-site battery storage.

8. 🥔**Cumulative impact air permits: require air permit assessments of simultaneous use of oil and gas backup generators**. Permits for backup generators are typically issued for a single site or operator, but adverse grid conditions that result in the use of backup diesel generators will trigger emissions from many generators at the same time, which could have health or environmental consequences far beyond that which could be realized from an individual generator.

   Air permitting agencies must be required to assess the cumulative impacts of multiple generators operating simultaneously under adverse conditions, and adjust permit conditions accordingly.

   Compliance: Agencies charged with issuing air permits.

Transparency & Planning

9. 🥔**Utility growth transparency: require utilities to disclose large load customer trajectories.** Some of the most problematic outcomes from data center load growth happen when utilities fail to disclose the growth that they anticipate from data centers and other large load customers until it's too late, triggering a planning emergency. To circumvent this information imbalance, regulators should require utilities to disclose large load customer trajectories with the appropriate level of certainty, well before submitting applications for new resources.

   Implementation: PUCs mandate compliance filings from regulated utilities serving over a certain threshold for all interconnection requests above a certain threshold. For example, PUC could require that all regulated utilities serving above a combined 1,000 MW of peak load, or 4 million MWh of commercial and industrial load, must submit an annual, or biennial accounting of all outstanding interconnection requests for customers with peak demands above 50 MW.

Compliance: Mandated by PUC, or required by legislation.

10. 🍎**Utility IRP: Require utilities to perform Integrated Resource Planning (IRP) with electrification and data center loads.** Utilities in 35 states are required to submit some form of IRP to regulators, assessing their projected customer demands, owned and contracted resources, and potential resource gaps. Better IRP regulations require utilities to assess least-cost methods of meeting customer requirements, including resource replacement (i.e. coal retirement) and least cost procurement with sensitivities. PUCS should require that utilities perform robust IRPs that take into account electrification loads (building, transportation, and industrial), new manufacturing, and data centers. Utilities must back up their assessments with robust projections of data center or other large loads, and assess the risks to their systems.

    IRPs can provide excellent opportunities to assess the costs and risks associated with serving large customers. IRP rules should be modified, or created, to assess ratepayer costs before and after data centers.

    In states that do not currently have robust IRP, the growth of data centers may provide an avenue to advocate for IRP rules.

    Implementation: PUCs can amend existing IRP rules, or be required to amend existing rules through a legislative process. For states without IRP, PUCs can either develop rules under existing framework or may be required to implement IRP rules.

    Compliance: Mandated by PUC, or required by legislation.

11. 🍎**State reporting: implement large load reporting to state energy offices or DEQ**. In the absence of a robust PUC, or in restructured states, state energy offices or departments of environmental quality may serve as a viable tracking agencies to track data center load growth through mandatory surveys.

    Implementation: May require enabling legislation.

12. 🥔**PUC reports to legislature: require periodic PUC reports to legislature on large load customer trajectories, electric affordability, infrastructure requirements, and economic development / jobs from data center growth.** In almost every state, PUCs are grappling with the impacts of data center growth, with narratives from data centers, utilities, and ratepayer advocates arguing about the value of new load, economic incentives and jobs, infrastructure, and rates. PUCs ultimately will make decisions about rate structures and utility requirements that seek to balance the public interest. Therefore, it is incumbent upon those regulators to report back to state representatives on the impacts of data centers and other large loads.

    Implementation: Legislative requirement to PUC (or state energy office, if applicable) for annual, or more often, reporting of data center growth, and the actual and projected impacts of electric affordability, infrastructure requirements, and economic development or jobs from data center growth.

Compliance: Mandated by legislation.

Energy Efficiency and Demand Response

13. 🍎**Establish power use efficiency ("PUE") standards**. Data centers use power both for computing and information technology ("IT" load) as well as cooling and other on-site uses. The fraction of total energy consumption to energy used for IT is termed "power use efficiency" ("PUE"). Most large data center companies have PUEs less than 1.5, and leading data centers have PUE of 1.2 or less. Google and Meta both report PUEs of around 1.1, and Microsoft around 1.2. Other large companies like Amazon, Digital Realty, and QTS do not publicly report PUEs. States could seek to establish PUE standards with resulting impacts in either tax incentives, permitting requirements, or rate treatment.

    Implementation: This would be a novel requirement in the US. Germany is looking at a federal PUE standard.

14. 🍎**Establish demand management / demand response incentives for traditional data centers.** A small amount of demand management or demand response at data centers could relieve congestion and utility strain, reduce capacity requirements, and allow new renewable energy to serve more requirements. Data centers generally use energy around the clock at a very high utilization, but with increasing AI loads, may be able to shift demand geographically (i.e. to be processed elsewhere) or temporally (i.e. to be processed at a different time of day).

    As an alternative to a demand management incentive, data centers could be called on to provide a minimum amount of demand response capability (through battery storage or reduced use) and charged higher utility rates otherwise.

    Implementation: An incentive structure could be implemented through utility rates.

    Additional considerations: Cryptocurrency mining operations have a high rate of demand response because their operations require effectively no continuity. As a consequence, when high demand-response incentives are offered, cryptocurrency mining operations have tended to capitalize on those incentives, earning substantial payment for a nearly free service. Therefore, demand response payments should be carefully calibrated to the difficulty entailed in curtailing operations. It is far harder for a traditional data center or large industrial facility to temporarily shut down than a cryptocurrency mining operation.