PFLN
PRO-FAMILY LEGISLATIVE NETWORK
BY WALLBUILDERS

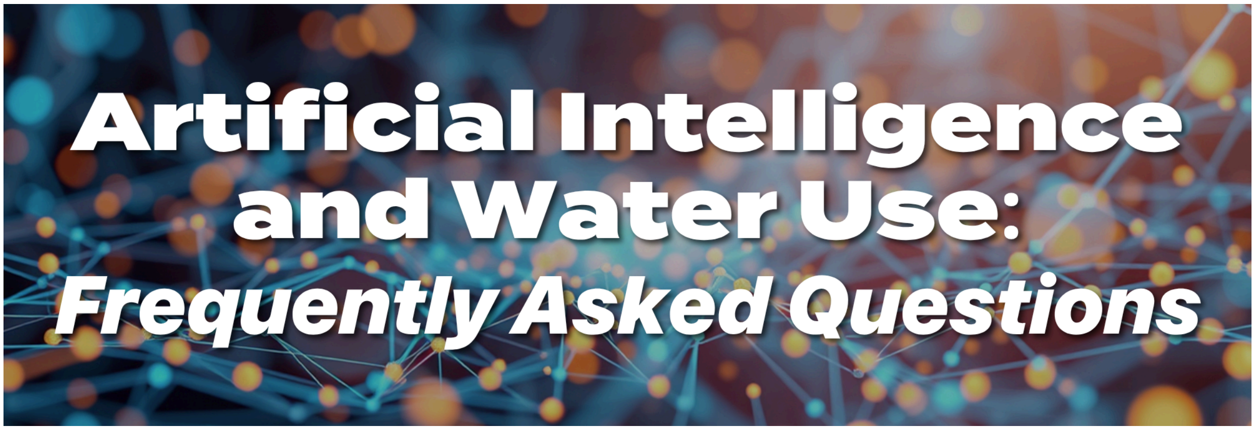

Artificial Intelligence and Water Use: *Frequently Asked Questions*

Introduction

This FAQ is designed to help state legislators understand the basic facts and policy considerations surrounding artificial intelligence infrastructure and water use. As AI systems expand, large data centers are being built across the country. In some communities, questions have arisen about how these facilities use water and how that demand fits within existing water systems and planning frameworks.

State lawmakers oversee water rights systems, public utilities, and land-use policy within their jurisdictions. Clear, accessible information about data centers, cooling technologies, water demand, and growth trends is therefore essential for informed oversight.

This document explains how data centers operate, why AI increases computing intensity, and how cooling systems affect water demand. It reviews industry growth and mitigation technologies, summarizes recent state-level disputes, and concludes with governance considerations for water planning and infrastructure oversight.

The FAQ is organized into five sections with thematically related questions. Sources for each section are provided at the end of the document.

Section I: Water-Related Policy Conflicts in the States

1. Why has data center water use become controversial in some states?

Controversy usually arises when rapid data center growth meets tight water supplies, limited local infrastructure, or incomplete public information. Several issues tend to drive disputes:

- **Use of drinking water for cooling**, especially in drought-prone or fast-growing areas.
- **High summer demand**, when cooling needs are greatest and communities are already under water pressure.
- **Whether local water and wastewater systems can handle additional demand** without raising rates or reducing reliability.
- **Questions about local benefits**, particularly when projects receive tax incentives but create relatively few permanent jobs.
- **Multiple facilities proposed in the same area**, raising concerns about how much growth is too much.
- **Lack of clear early disclosure** about projected water use, water sources, or conservation plans.

In many places, the debate is not just about one facility, but about how quickly development is expanding and whether existing rules were built for that scale.

2. *What are some recent examples of water-related disputes involving AI data centers?*

Tucson, Arizona

Tucson considered a proposed data center known publicly as “Project Blue,” a secretive annexation proposal that could have used hundreds of millions of gallons of water per year, including potable supplies during its early phases. Public outrage over limited disclosure and large projected water demand led the City Council to reject the annexation. In response, Tucson unanimously adopted a Large Quantity Water User ordinance requiring any customer expecting to use more than 7.4 million gallons per month to submit a public water-conservation plan and obtain council approval before connecting to the city’s system.

San Marcos, Texas

San Marcos reviewed a proposed $1.5 billion data center campus that required rezoning approval. Public debate focused heavily on projected water use, particularly amid ongoing drought conditions and concerns about the condition of the Edwards Aquifer. Residents also questioned whether the scale of the project matched the character of the community, even as developers highlighted construction jobs and tax revenue. In February 2026, the City Council voted 5–2 to deny the rezoning request, stopping the project. The decision was a local land-use vote rather than a change in statewide policy.

Monroe County, Georgia

In Monroe County, controversy emerged around proposals for multiple large data center developments, including a 1,632-acre “Forsyth Technology Campus” that could consume more than two million gallons of water per day at full buildout. Concerns focused on long-term water planning, infrastructure costs, legal exposure from inconsistent approvals, and the cumulative impact of clustered development. County officials adopted a temporary moratorium on new data center applications while revising zoning rules. During that period, local leaders began considering regulations modeled on Groton, Connecticut’s 2023 ordinance, including strict size caps, spacing requirements, and potential prohibitions on water-cooling systems.

West Des Moines, Iowa

In central Iowa, local reporting has documented significant water use associated with clustered data centers, including nearly 49 million gallons used by Microsoft’s West Des Moines facilities through September 2025. Data centers account for between 2 and 7 percent of the water pumped monthly by West Des Moines Water Works. During the summer, the utility imposed lawn-watering restrictions, though officials stated the ban would have occurred regardless of data center usage. The episode nevertheless intensified public scrutiny of overall industrial demand, long-term aquifer capacity, and whether alternative cooling technologies or revised permitting standards should be encouraged for future projects.

Northern Virginia

Northern Virginia hosts the nation's largest concentration of data centers, and sustained growth has pushed water and infrastructure issues into the General Assembly. In the 2026 session, lawmakers introduced legislation that would require water providers to report monthly volumes used by data centers and require certain high-load facilities to disclose anticipated water demand during development. These proposals emerged amid broader debates over grid strain, ratepayer impacts, and local oversight, following prior reform efforts that were vetoed. Rather than responding to a single project, Virginia lawmakers are considering how to manage cumulative growth across a region already saturated with data center development. These measures remain under consideration.

3. How are states and local governments responding to these conflicts?

In these states—and others facing similar debates—government responses have generally fallen into a few common categories:

- **Large water-user rules:** Some cities have created clear thresholds that trigger additional review for high-volume users. Tucson's Large Quantity Water User ordinance is one example, requiring conservation planning and public review once projected water use exceeds a defined level.
- **Temporary moratoria:** Local governments, including Monroe County in Georgia, have paused approvals while updating zoning rules, infrastructure standards, or cooling requirements.
- **Greater use of reclaimed or non-drinking water:** In areas where using potable water is politically sensitive, officials have encouraged or required the use of recycled wastewater or other non-potable supplies when available.
- **Clearer rules about who pays for infrastructure:** Many disputes have focused on whether developers must fund water and wastewater improvements tied to their projects rather than shifting costs to existing customers.
- **Statewide reporting and mitigation rules:** In states experiencing concentrated growth, such as Virginia, legislators are considering broader rules on reporting and mitigation so expectations are clearer before projects move forward.

4. What transparency or reporting concerns have contributed to these conflicts?

In many cases, disputes have escalated because key information was unclear or difficult to compare. Common concerns include:

- **Late disclosure:** Communities sometimes learn detailed water-use estimates or sourcing plans late in the approval process, limiting meaningful public review. The Tucson debate is a prominent example.
- **Inconsistent numbers:** Public discussions often mix peak-day demand with annual averages, or withdrawals with consumption, making it difficult to compare projects across locations.
- **Limited early project details:** Use of code names, nondisclosure agreements, or incomplete early filings can create distrust even if more information is provided later.
- **Difficulty gathering statewide data:** When reporting requirements differ across utilities and jurisdictions, policymakers may not have a clear picture of total data center water use in their state. For example, proposed reporting changes in Virginia reflect this concern.
- **Public blame during water shortages:** During droughts or other system challenges, large industrial users may receive attention even if they are not the main cause. Clear, routine reporting can help distinguish normal industrial demand from unrelated water problems.

Section II: Governance and Policy Considerations

5. How could federal preemption efforts affect state authority over data center water use?

Water allocation and land-use regulation have traditionally been state and local responsibilities. States control water rights systems, oversee public water utilities, and set zoning and development standards. Local governments typically handle permitting decisions tied to infrastructure capacity and land use.

That said, recent federal executive actions aimed at accelerating AI development have raised questions about how far federal influence might extend, especially where state policy touches AI infrastructure and related regulatory frameworks.

In December 2025, President Trump signed an executive order titled "Ensuring a National Policy Framework for Artificial Intelligence." At a high level, the order seeks to discourage or limit certain state AI laws by directing federal agencies to identify state measures viewed as burdensome and by coordinating potential legal challenges to those laws. The order also signals that access to certain federal funding streams could be reconsidered if states pursue policies the administration considers obstructive.

For lawmakers focused specifically on data centers and water use, an important detail is that the executive order contains an explicit carveout for AI computing and data center infrastructure. In practical terms, this suggests the order is aimed primarily at state laws governing AI systems and applications, rather than traditional state and local authority over physical infrastructure such as siting, utilities, or water management.

That does not mean boundary disputes are impossible. Federal agencies could still scrutinize whether a state action is framed as regulating "AI" itself rather than land use, utilities, or water allocation, and funding conditions could become a point of tension if federal grants are tied to broader AI policy goals. But unless Congress enacts new legislation clearly preempting state authority in this area—or courts interpret existing federal law to have that effect—core water rights systems, groundwater regulation, municipal utility governance, and local zoning decisions generally remain in state hands under current constitutional doctrine.

Some legal commentators have suggested that aspects of the executive order, particularly any use of funding pressure to influence state policy, could face legal challenges. Regardless of potential litigation, the practical takeaway is that federal policy may influence incentives and timelines, but day-to-day control over water sourcing, permitting, and infrastructure planning remains primarily with states and local governments.

6. What policy tools and guardrails should state legislators consider when addressing AI-related water demand?

State responses need not take the form of technology mandates. In most cases, the goal is clearer planning, predictable rules, and reduced conflict. Several broad approaches, many of which respond directly to the transparency and planning concerns described earlier in this document, may serve as practical options.

- **Disclosure and reporting:** Legislators may consider requiring standardized reporting of projected water demand and water sources early in the permitting process. This can include both typical daily use and peak summer demand so utilities can assess potential system strain. For very large facilities, periodic reporting may support long-term planning while still protecting sensitive business information.
- **Large-user thresholds and structured review:** States or cities can establish clear triggers for additional review once projected water use exceeds a defined level. Smaller projects would proceed under normal rules, while very large users would undergo additional planning steps or conservation review. Clear thresholds reduce uncertainty by setting expectations before major investments are made.
- **Water source and cooling expectations:** In water-stressed areas, policymakers may prioritize or require the use of reclaimed wastewater or other non-drinking supplies where feasible. Rather than prescribing a specific cooling technology, states may require a water

management plan that addresses sourcing, conservation practices, and how operations would adjust during peak-demand or drought conditions. This keeps the focus on measurable outcomes rather than engineering design choices.

- **Infrastructure cost allocation ("who pays"):** Many disputes arise over who funds necessary water and wastewater expansions tied to new development. Legislators may clarify when project-driven upgrades must be funded by the developer rather than existing ratepayers. Clear cost-allocation rules can reduce public backlash and avoid ad hoc negotiations.
- **Drought and contingency planning:** Large facilities may be required to submit contingency plans explaining how water use would be reduced during drought restrictions or other emergency conditions. Aligning these plans with existing local drought stages helps ensure that industrial users respond consistently with community-wide measures.
- **Conditioning incentives on measurable commitments:** When tax incentives, expedited permitting, or infrastructure support are offered, policymakers may choose to tie those benefits to clear commitments, such as non-potable sourcing where available, conservation targets, timely reporting, or funding of required upgrades. Linking incentives to performance can address public concerns without creating an overly complex regulatory structure.

Section III: The Basics

7. What is a data center?

A data center is a purpose-built facility that houses computer systems and related infrastructure used to store, process, and transmit digital information. The core "IT" equipment typically includes servers (the machines that run applications and perform calculations), storage systems (where data are kept), and networking equipment (routers and switches that move data within the facility and to the broader internet).

Unlike a typical office building, a data center is engineered around continuous operation. The equipment runs at high intensity and generates significant heat, so the facility must provide stable electrical power (often with backup systems), cooling, and environmental controls. Many modern facilities are designed with redundant power and cooling systems so that computing can continue even if a component fails or utility power is interrupted.

Data centers vary widely in scale. Some are relatively small and serve a single organization. Others are large "hyperscale" campuses operated by cloud providers to support many customers and extremely high computing workloads. Regardless of size, the defining feature is the same: they are built to deliver reliable computing capacity around the clock.

8. Why does AI require large data centers?

Modern artificial intelligence—particularly large-scale "generative" models—requires unusually high levels of computing power. Training a model involves running repeated mathematical calculations across very large datasets in order to adjust internal parameters. This process can take days or weeks and typically relies on specialized chips, such as graphics processing units (GPUs), that are designed to perform many computations in parallel.

These chips draw substantial electricity and produce significant heat. Concentrating them in large facilities makes it easier to provide high-capacity power connections, advanced cooling systems, and fast internal networking between machines working on the same tasks.

AI systems must also respond quickly to user requests and scale to high demand. Centralized data centers allow operators to pool computing resources, shift workloads between machines, and add capacity over time without rebuilding the system from scratch. For these reasons—computational intensity, specialized hardware, and scaling requirements—AI development tends to drive investment in larger and denser data center campuses.

Section IV: Industry Growth

9. Where are data centers most heavily concentrated?

Data centers cluster in a limited number of major hubs where power, fiber connectivity, and real estate-development conditions support large-scale projects. Northern Virginia remains the most prominent concentration in the United States and is frequently described as the world's leading hyperscale market. Other major U.S. hubs include Dallas–Fort Worth, Phoenix, Atlanta, and Chicago.

Growth is also expanding into additional states and smaller markets as land, transmission capacity, or interconnection timelines become constrained in the largest hubs. States such as Oregon, Iowa, and Ohio have attracted large cloud campuses in part because of available land and power infrastructure.

While national growth figures are substantial, the physical footprint of development tends to be concentrated in specific corridors and utility service areas. That geographic concentration shapes how impacts are experienced and evaluated at the local level.

10. Why are these locations often chosen for large AI facilities?

Large AI data centers tend to locate where multiple practical requirements can be met at once:

- **Power availability and delivery timelines:** Access to high-capacity electricity, and the ability to secure it within a predictable timeframe, is often the dominant factor in site selection.
- **Connectivity ecosystems:** Established hubs already have extensive high-speed fiber connections and multiple telecommunications providers. This allows data to move quickly and reliably between facilities and users, reducing delays and minimizing the risk of outages.
- **Large parcels and room for expansion:** Hyperscale projects often require land for phased buildouts, setbacks, substations, backup systems, and security.
- **Permitting and local experience:** Communities that have already permitted data centers often have clearer processes and fewer "unknowns," which can reduce uncertainty.
- **Climate and cooling strategy:** Local climate influences cooling needs. Cooler regions may allow more hours of air-based cooling, while hotter areas rely more heavily on evaporative systems. In some cases, access to reclaimed or non-potable water also factors into site decisions.

Taken together, these considerations explain why certain hubs continue to attract development and why expansion spreads to new regions when infrastructure or timelines become constrained in established markets.

11. How does data center water use compare to other industries?

Precise national totals for data center water use are not centrally tracked, but rough estimates can be derived from installed electrical capacity. As noted previously, major U.S. markets accounted for roughly 8,000 megawatts (MW) of data center capacity as of 2025. Using the aforementioned estimate of approximately 5,000 gallons per day per MW for evaporative-cooled facilities, total on-site water use nationally would likely fall on the order of 35 to 45 million gallons per day.

Even allowing for continued growth, this remains small relative to the largest national water-use categories. According to U.S. Geological Survey estimates, irrigation withdrawals are approximately 118 billion gallons per day, and thermoelectric power withdrawals about 133 billion gallons per day. Public water supply withdrawals total roughly 39 billion gallons per day.

In other words, national data center water use is measured in tens of millions of gallons per day, while major sectors operate in tens or hundreds of billions. The difference is measured in orders of magnitude.

However, comparisons at the local level can look very different. A single facility using hundreds of thousands—or in some cases millions—of gallons per day may represent a meaningful share of available capacity in a smaller municipality or water-stressed basin. For state-level review, the most relevant comparisons are therefore local: existing municipal supply, peak-day demand, other industrial users on the same system, and the cumulative impact of clustered development.

In short, data center water use is modest in national aggregate but can be significant in particular regional contexts.

Section V: Technology and Mitigation

12. How are data centers cooled today?

Understanding cooling systems is essential because cooling design largely determines on-site water demand. Data centers must continuously remove heat generated by servers and other computing equipment. As electricity is converted into computation, nearly all of it ultimately becomes heat. If that heat is not removed, equipment performance can decline and hardware can fail. Cooling systems therefore operate constantly to maintain reliability.

Most U.S. data centers rely on a combination of airflow management and mechanical refrigeration to keep equipment within required temperature and humidity ranges. Cooling approaches vary, but several common configurations are widely used:

- **Air-cooled systems:** These systems circulate conditioned air through server rooms. Warm air produced by the equipment is returned to cooling units, where refrigeration systems—either direct-expansion (DX) units or chilled-water systems—cool the air before it is recirculated. Depending on the configuration, air-cooled systems may operate with little or no on-site water use.
- **Water-cooled systems:** Many large facilities use chilled water to absorb heat from equipment. That heat is then transferred to cooling towers, where it is released to the atmosphere through evaporation. Because water evaporates during this process, these systems consume water on an ongoing basis.
- **Economizers ("free cooling"):** Where climate conditions permit, facilities can reduce mechanical refrigeration by using cool outside air (air-side economizers) or naturally cool water loops (water-side economizers). Air-side systems typically reduce both energy and water use. Water-side economizers may still rely on cooling towers, so overall water use depends on the specific configuration.
- **Hybrid configurations:** Many facilities combine these approaches and switch between modes depending on outside temperature, humidity, and system demand. A facility might rely primarily on dry cooling during cooler months and shift toward evaporative cooling during peak summer conditions.

As AI workloads increase computing density, some operators are adopting liquid-cooling technologies that circulate a coolant—often water-based or another non-conductive fluid—in a closed loop directly across or near the computer chips. Similar to the coolant in a car engine, this fluid absorbs heat and carries it away from the equipment before being cooled and reused. This approach can manage higher heat loads and improve cooling performance. However, even when heat is removed using liquid inside the building, it must ultimately be released to the outside environment through either evaporative systems (which consume water) or dry cooling systems.

From a policy perspective, the key question is whether heat removal relies primarily on evaporation—which consumes water—or on dry cooling, which reduces on-site water use but may increase electricity demand or capital cost.

13. What technologies can reduce or eliminate water use in data centers?

There is no single "water-free" solution that fits every site. Water use depends on climate, cooling design, and reliability requirements. However, several broad approaches can significantly reduce—and in some cases nearly eliminate—on-site water consumption:

- **Dry cooling instead of evaporative cooling:** Facilities can use large air-cooled systems to release heat without relying on cooling towers. This can substantially reduce or eliminate routine water consumption. The tradeoffs may include higher electricity use during hot weather and higher upfront construction costs.
- **Hybrid systems:** Some facilities rely primarily on dry cooling but use limited water during the hottest periods to improve efficiency. This can reduce total annual water use while maintaining reliability during peak conditions.
- **Improved internal cooling efficiency:** As computing equipment becomes more powerful and concentrated, some operators use liquid-based systems inside the building to move heat away from equipment more efficiently. While this does not automatically eliminate water use, it can reduce overall cooling demand and improve performance.
- **Use of non-potable water:** Where available, facilities may rely on reclaimed wastewater or other non-drinking sources rather than municipal drinking water, reducing pressure on potable supplies.
- **Operational water management:** Careful system monitoring, leak prevention, and improved cooling-tower management can reduce avoidable water losses without major design changes.

In practice, water reduction strategies involve tradeoffs among cost, energy use, land requirements, and reliability. For policymakers, the central questions are whether water use is minimized where feasible, whether potable supplies are protected where possible, and whether the facility's design choices are transparent during the approval process.

14. How does electricity generation for AI affect overall water demand?

The water used at a data center is only part of the picture. AI systems require large amounts of electricity, and producing that electricity can itself require water—depending on how it is generated.

The key distinction is between direct water use at the data center and indirect water use at power plants.

Many power plants that generate electricity using steam—such as coal, natural gas, or nuclear facilities—use cooling systems similar in principle to those described in Question 10. Plants that rely on cooling towers consume water through evaporation. Some older plants withdraw very large volumes of water and return most of it to the source, while others withdraw less but consume more through evaporation.

By contrast, electricity sources such as wind and solar use little or no water during normal operation. This observation concerns water use only; electricity sources differ significantly in reliability, cost, land use, and broader environmental or supply-chain impacts, which are beyond the scope of this water-focused discussion.

To illustrate scale, a 100-megawatt data center operating continuously requires roughly 2,400 megawatt-hours of electricity per day. The water impact of generating that electricity depends on which types of power plants are meeting the demand. If the additional load is supplied primarily by water-cooled thermoelectric plants, indirect water consumption can increase. If supply comes from low-water sources, the effect is much smaller.

For policymakers, this means a facility's total water footprint is not limited to what occurs on site. A data center with minimal on-site water use may still increase water demand elsewhere in the power system—particularly during peak summer periods when both electricity demand and water stress are highest.

Section VI Additional Information:

15. Why do data centers use water?

Data centers use water primarily for cooling. As servers and AI accelerators run, they convert electricity into computation and release heat as a byproduct. If that heat is not removed, performance can decline and equipment can be damaged.

Many facilities remove heat using systems that rely on water at some stage of the process, particularly evaporative cooling systems and cooling towers. In these systems, water absorbs heat and releases it to the atmosphere; some of that water is consumed through evaporation, and some may be discharged ("blowdown") to prevent mineral buildup inside the system.

Not all facilities use large amounts of water on site. Some rely on air-based ("dry") cooling or hybrid systems that reduce water demand. It is also important to distinguish between direct water use (water consumed in on-site cooling) and indirect water use (water used elsewhere to generate the electricity the facility consumes). The electricity-related effects are discussed in greater detail in Section III.

16. What types of water do data centers use?

Data centers generally rely on one of two broad categories of supply:

- **Potable (drinking-quality) water:** Some facilities use treated municipal drinking water for cooling because it is readily available and connected to existing utility systems. This can be controversial in water-stressed areas, especially when the same supply serves residents and businesses.
- **Non-potable water:** Where available, facilities may use reclaimed wastewater (treated effluent), recycled industrial water, or other non-drinking sources to reduce reliance on potable supplies.

In either case, the underlying source typically originates from surface water (rivers, lakes, reservoirs) or groundwater (aquifers), delivered through a municipal utility or, in some instances, withdrawn directly under state permit.

Cooling is by far the largest driver of on-site water demand. Smaller volumes may be used for humidification, maintenance, or other routine facility needs, but these are typically minor compared with cooling requirements.

The source and quality of water matter for policy because they affect drinking-water planning, wastewater capacity, and long-term infrastructure decisions. In practice, debates often center on whether potable supplies are being used for cooling and what alternatives are feasible.

17. How rapidly are AI data centers expanding in the United States?

U.S. data center growth has accelerated sharply since 2023–2024, with AI demand reflected in both operating capacity (measured in megawatts, or MW) and projects under construction. In commercial real estate firm CBRE’s tracking of major North American markets, total inventory reached 8,155 MW in the first half of 2025—up 17.6 percent from the prior half-year and 43.4 percent year-over-year. Vacancy fell to a record-low 1.6 percent. At the same time, 5,242.5 MW were under construction, and more than 74 percent of that pipeline was already pre-leased, largely to cloud and AI tenants securing future capacity. Northern Virginia alone accounted for more than 2,000 MW under construction during that period.

Federal energy forecasters are also explicitly linking near-term electricity demand growth to “large computing centers,” projecting U.S. electricity use to grow 1 percent in 2026 and 3 percent in 2027, with data centers a major driver.

In practical terms, expansion means not only more facilities, but larger and denser builds that require substantial power connections, land assembly, and cooling infrastructure on compressed timelines. Those pressures are often felt by state and local governments within a single planning or budget cycle.

18. How much water can a large data center use?

Water use varies widely depending on cooling design, climate, operating intensity, and whether the figure refers to withdrawals (water taken from a source) or consumption (water not returned because it evaporates or is discharged). Facilities that rely on evaporative cooling systems—often using cooling towers—tend to have higher on-site water consumption because water is lost through evaporation and periodic discharge ("blowdown").

Policy discussions commonly cite:

- **Approximately 300,000 gallons per day** as a representative daily figure for a large data center. This is often framed as roughly the demand of 1,000 households.
- **Up to roughly 5 million gallons per day** for very large facilities or campuses under certain conditions. This is sometimes compared to the daily use of a small town.

Another useful planning lens is water use per megawatt (MW) of electrical capacity. For facilities using evaporative cooling, the International Energy Agency estimates approximately 5,000 gallons per day per MW of IT load. Using that estimate, a 50 MW facility might use a few hundred thousand gallons per day, while a multi-hundred-megawatt campus can reach into the millions.

All of these figures require context. Peak summer demand can be significantly higher than annual averages. Actual use also depends on design choices such as dry cooling, hybrid systems, or reclaimed-water use, each of which may involve tradeoffs in cost, energy use, or reliability. Section III discusses those technological options in greater detail.

For Further Reading

Section I

- Congressional Research Service, "Data Centers and Their Energy Consumption," January 23, 2026, https://www.congress.gov/crs-product/R48646
- Congressional Research Service, "Data Centers and Cloud Computing: Information Infrastructure for Artificial Intelligence," February 5, 2025, https://www.congress.gov/crs-product/IF12899
- Lawrence Berkeley National Laboratory, Center of Expertise for Data Center Energy, "Water Efficiency," accessed February 10, 2026, https://datacenters.lbl.gov/water-efficiency

- U.S. Department of Energy, “Cooling Water Efficiency Opportunities for Federal Data Centers,” accessed February 10, 2026, https://www.energy.gov/femp/cooling-water-efficiency-opportunities-federal-data-centers
- U.S. Department of Energy, “Energy-Efficient Cooling Control Systems for Data Centers,” accessed February 10, 2026, https://www.energy.gov/eere/ito/energy-efficient-cooling-control-systems-data-centers
- National Laboratory of the Rockies, “High-Performance Computing Data Center Water Usage Efficiency,” February 17, 2026, https://www.nlr.gov/computational-science/reducing-water-usage
- Uptime Institute, “Tier Classification System,” accessed February 10, 2026, https://uptimeinstitute.com/tiers
- Environmental and Energy Study Institute, “Data Centers and Water Consumption,” June 25, 2025, https://www.eesi.org/articles/view/data-centers-and-water-consumption

Section II

- CBRE, “North America Data Center Trends H1 2025,” September 8, 2025, https://www.cbre.com/insights/reports/north-america-data-center-trends-h1-2025
- International Energy Agency, “Energy and AI,” April 10, 2025, https://www.iea.org/reports/energy-and-ai
- U.S. Energy Information Administration, “EIA forecasts strongest four-year growth in U.S. electricity demand since 2000, fueled by data centers,” *press release,* January 13, 2026, https://www.eia.gov/pressroom/releases/press582.php
- U.S. Energy Information Administration, “Electricity use for commercial computing could surpass space cooling, ventilation,” June 25, 2025, https://www.eia.gov/todayinenergy/detail.php?id=65564
- Synergy Research Group, “Synergy Updates Its Rankings of the World’s Top 20 Hyperscale Data Center Locations,” July 16, 2025, https://www.srgresearch.com/articles/synergy-updates-its-ranking-of-the-worlds-top-20-hyperscale-data-center-locations
- Brookings Institution, “AI, data centers, and water,” November 20, 2025, https://www.brookings.edu/articles/ai-data-centers-and-water/
- U.S. Geological Survey, “Total water use in the United States,” June 8, 2018, https://www.usgs.gov/water-science-school/science/total-water-use-united-states
- U.S. Geological Survey, “How much water is used by people in the United States?” June 17, 2025, https://www.usgs.gov/faqs/how-much-water-used-people-united-states

Section III

- U.S. Department of Energy, “Best Practices Guide for Energy-Efficient Data Center Design,” July 2024, https://www.energy.gov/sites/default/files/2024-07/best-practice-guide-data-center-design_0.pdf
- Lawrence Berkeley National Laboratory, Center of Expertise for Data Center Energy, “Water Efficiency,” accessed February 10, 2026, https://datacenters.lbl.gov/water-efficiency
- U.S. Department of Energy, “Cooling Water Efficiency Opportunities for Federal Data Centers,” accessed February 10, 2026, https://www.energy.gov/femp/cooling-water-efficiency-opportunities-federal-data-centers
- Nuoa Lei et al., “The water use of data center workloads: A review and assessment of key determinants,” *Resources, Conservation, and Recycling,* Volume 219, June 1, 2025, https://www.sciencedirect.com/science/article/abs/pii/S0921344925001892?via%3Dihub
- U.S. Energy Information Administration, “U.S. electric power sector’s use of water continued its downward trend in 2020,” December 17, 2021, https://www.eia.gov/todayinenergy/detail.php?id=50698
- U.S. Geological Survey, “Thermoelectric Power Water Use,” March 1, 20219, https://www.usgs.gov/mission-areas/water-resources/science/thermoelectric-power-water-use
- National Laboratory of the Rockies, “High-Performance Computing Data Center Water Usage Efficiency,” accessed February 10, 2026, https://www.nlr.gov/computational-science/reducing-water-usage

Section IV

- **Reuters**, “Desert storm: Can data centres slake their insatiable thirst for water?” December 17, 2025, https://www.reuters.com/sustainability/climate-energy/desert-storm-can-data-centres-slake-their-insatiable-thirst-water--ecmii-2025-12-17/

- **E&E News** by Politico, “States push to end secrecy over data center water use,” December 8, 2025, https://www.eenews.net/articles/states-push-to-end-secrecy-over-data-center-water-use/

- AP News, “Proposed data center prompts Tucson to regulate large water users, require conservation,” August 21, 2025, https://apnews.com/article/tucson-data-management-and-storage-arizona-general-news-environmental-conservation-42c1e554b02b4293685a08a4574db9f0

- Virginia Mercury, “These bills aim to regulate Virginia data center siting, generator use and SCC oversight,” January 19, 2026, https://virginiamercury.com/2026/01/19/these-bills-aim-to-regulate-virginia-data-center-siting-generator-use-and-scc-oversight/
- KUT News, “San Marcos City Council blocks proposed data center,” February 18, 2026, https://www.kut.org/energy-environment/2026-02-18/san-marcos-city-council-blocks-proposed-data-center
- 13MAZ.com, “Monroe County passes data center moratorium, says they will review potential limitations,” December 3, 2025, https://www.13wmaz.com/article/news/local/forsyth-monroe/monroe-county-passes-data-center-moratorium-will-review-potential-limitations/93-195cbb55-5a2a-4f29-879f-c0f3f859aef2
- Business Record, “Energy, water use under close watch as data centers expand in Iowa,” November 28, 2025, https://www.businessrecord.com/energy-water-use-under-close-watch-as-data-centers-expand-in-iowa/

Section V

- WhiteHouse.gov, “Ensuring a National Policy Framework for Artificial Intelligence,” *Executive Order,* December 11, 2025, https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy/
- Latham & Watkins, “AI Executive Order Targets State Laws and Seeks Uniform Federal Standards, ” December 17, 2025, https://www.lw.com/en/insights/ai-executive-order-targets-state-laws-and-seeks-uniform-federal-standards
- White & Case, “State AI laws under federal scrutiny: Key takeaways from the executive order establishing federal AI policy framework,” January 15, 2026, https://www.whitecase.com/insight-alert/state-ai-laws-under-federal-scrutiny-key-takeaways-executive-order-establishing
- Reuters, “Trump’s order targeting state AI laws faces political and legal hurdles,” December 12, 2025, https://www.reuters.com/legal/government/trumps-ai-order-faces-political-legal-hurdles-2025-12-12/