



**BUREAU OF
ECONOMIC
GEOLOGY**



COMPASS
Collaborative Optimization & Management of Power
Allocation – Surface & Subsurface strategies



The University of Texas at Austin
Energy Institute



The University of Texas at Austin
Jackson School of Geosciences



The University of Texas at Austin
**Walker Department
of Mechanical Engineering**
Cockrell School of Engineering



WATER USE REQUIREMENTS FOR DATA CENTERS IN TEXAS

A White Paper on the Evolving Demands of Water
Use in Data Center Infrastructure in Texas

Prepared by COMPASS Research Consortium

CONTACT INFORMATION

COMPASS

COMPASS@BEG.UTEXAS.EDU

10611 Exploration
Way, Austin, TX 78758



Scan QR Code for
More Information



**BUREAU OF
ECONOMIC
GEOLOGY**



COMPASS
Collaborative Optimization & Management of Power
Allocation - Surface & Subsurface strategies



The University of Texas at Austin
**Walker Department
of Mechanical Engineering**
Cockrell School of Engineering



The University of Texas at Austin
Energy Institute



The University of Texas at Austin
Jackson School of Geosciences

WATER USE REQUIREMENTS FOR DATA CENTERS IN TEXAS

A White Paper on the Evolving Demands of Water Use in Data Center Infrastructure in Texas



Cooling Modern Data Centers

AI computing and hyperscale cloud infrastructure pack enormous power into ever-smaller hardware, converting nearly all input energy into heat. As component densities rise, extreme heat fluxes demand sophisticated cooling solutions that are often heavily reliant on water.

Reframing Water in Data Center Strategy

Our white paper highlights the growing water footprint of data centers, arguing that water is no longer a secondary input. It is a central engineering, environmental, and policy challenge that will shape the future of sustainable digital infrastructure.

Inside the Report:

1

Brief Overview of the Impact of Data Centers on Energy and Water Demand

2

The Energy-Water-Data Nexus: Understanding Water Needs for Cooling

3

Understanding Indirect-Use Water Requirements for Data Centers in Texas

4

Water Quality and Treatment Considerations

5

Regional and Sectoral Comparison of Water Consumption


6

Transparent and Resilient Data Center Planning: Policy Recommendations

CONTACT INFORMATION

COMPASS | COMPASS@BEG.UTEXAS.EDU

10611 Exploration
Way, Austin, Texas



Scan QR Code for
More Information

Water Use Requirements for Data Centers in Texas

A White Paper on the Evolving Demands of Water Use in Data Center Infrastructure in Texas

About This Report – A Message from the COMPASS

At COMPASS, we are committed to advancing data-driven insights that support resilient infrastructure and sustainable resource management across Texas. This white paper reflects our growing focus on water use in large-load operations, particularly data centers, as part of a broader effort to understand the evolving demands on regional water systems.

The research spotlights a critical finding: where and when water is sourced can be just as consequential as how much is used. As data center development accelerates, these nuances will shape the future of water planning, permitting, and investment. We've included preliminary estimates, industry comparisons, engineering, and geospatial considerations to help stakeholders navigate this complex landscape.

Our goal is to inform public and private decision-makers about the opportunities and challenges associated with this fast-evolving sector and to help shape pathways that align economic competitiveness with community stewardship and regional resilience. This report reflects our commitment to supporting the growth of resilient digital infrastructure that meets the needs of both industry and communities in Texas.

Prepared by

COMPASS Research Affiliates Program at the University of Texas at Austin

Contact: compass@beg.utexas.edu



Executive Summary

Texas is rapidly emerging as one of the largest and fastest-growing data center markets in the United States, driven by energy availability, competitive energy pricing, abundant land, and robust infrastructure investments [1]. As of September 2025, Baxtel reports a total of 484 data centers in Texas, comprising 297 operational facilities, 109 under construction, and 124 planned [2]. The power demand from data centers in Texas has escalated sharply, with the sector accounting for a significant and growing share of regional electricity consumption. Texas data centers are concentrated in major hubs such as Dallas, Houston, Austin, and San Antonio, while expansive new campuses are being planned in the rest of the state [3, 4, 5]. With forecasted total capacity surpassing 70 GW by 2030, Texas data centers will consume electricity equivalent to tens of millions of homes, accentuating the immense scale of these operations [6].

Data centers are now recognized as the most energy-intensive elements of the digital economy [7, 8]. The total electrical power going into a single data center rack is expected to approach 1 MW by 2028 [9, 10]—equivalent to the peak hour energy use of roughly 250 U.S. homes flowing into a computing unit the size of a bookcase. Every unit of electrical energy flowing into microprocessors powering AI compute applications is eventually converted to heat, which needs to be removed. This is made more challenging by the fact that the electronic packages themselves have shrunk in size/volume over time, producing extremely high heat fluxes, sometimes comparable to heat fluxes on the surface of the sun ($\sim 60\text{--}70\text{MW}/\text{m}^2$) [11]. All this implies that at the hyperscale data center-level, thermal management to remove 100's MW of heat will be required. This magnitude of energy rivals numbers associated with power plants and is comparable to the energy requirements of small cities.

While energy consumption has dominated the recent discourses on data centers, water requirements — both direct and indirect — are now recognized as equally critical, particularly in the context of regional water scarcity and infrastructure resilience [12, 13]. As computational loads increase and thermal management becomes more complex, cooling technologies that rely heavily on water are being scaled across global networks of facilities [14]. This trend poses significant challenges for sustainable operations, particularly when economic incentives prioritize performance and cost efficiency over hydrological impact.

In addition to direct water use for cooling data centers, indirect water use occurs through electricity generation that powers data center operations. Since much of Texas's dispatchable electricity is produced by thermoelectric power plants fueled by natural gas, coal, or nuclear, which themselves consume significant water for cooling, the total water footprint of data centers extends beyond onsite consumption [15, 16]. This coupled water demand from both the operational and energy supply sides is estimated at 25 billion gallons in Texas in 2025, equivalent to about 0.4% of statewide water use [17, 18]. While the total statewide water use might appear low in percentage terms, this number is significantly higher in many water-stressed areas; additionally water-use will likely be much higher than the current estimates.

Efforts to address water requirements must emphasize the need for technological diversification and policy frameworks that consider local water availability, water treatment options, and the energy-water nexus. While data centers will drive economic growth and digital innovation, recognizing their water footprint will encourage more sustainable design and resource management practices critical to Texas's infrastructural resilience [19]. Left unaddressed, water issues have the potential to halt Texas's ambition and potential in AI and digital infrastructure.

This white paper addresses the urgent and growing need to understand and quantify the water footprint of data centers, alongside their escalating energy demands. Water has now emerged as a primary constraint in data center planning, particularly in regions vulnerable to drought, water stress, or infrastructure limitations. The adoption of water-intensive cooling systems, such as evaporative and hybrid technologies, while advantageous for energy efficiency, raises concerns over freshwater use and long-term sustainability. This white paper positions water not as a secondary input, but as a core engineering, environmental, and policy issue in the future of digital infrastructure.

Structured around three interrelated objectives, this white paper:

1. Quantifies the scale and growth of water and energy use in current and future data center configurations;
2. Highlights and assesses key water treatment technologies—thermal and membrane-based, that can reduce dependence on freshwater sources;
3. Proposes a regulatory and policy roadmap to support water stewardship and resilient infrastructure development in the data center sector.

Subsequent chapters build upon these goals and contribute to a holistic understanding of water in the data center context:

Chapter 1 – Brief Overview of the Impact of Data Centers on Energy and Water Demand

This chapter details the increasing infrastructure requirements of data centers, focusing on how the expansion of digital workloads and AI technologies is driving both energy and water consumption. It frames the foundational importance of resource management in data center design and planning.

Chapter 2 – The Energy-Water-Data Nexus: Understanding Water Needs for Cooling Data Centers

Chapter 2 establishes the conceptual framework of the energy-water nexus in digital infrastructure. It provides estimates of water use associated with various cooling systems and highlights the unique pressures of AI-related technologies on thermal management systems.

Chapter 3 – Understanding Indirect-Use Water Requirements for Data Centers in Texas

This chapter presents projected indirect water demand scenarios from 2025 to 2030 based on

empirical data and power capacity growth. It includes modeling outputs that compare different power generation technologies and their water use profiles.

Chapter 4 – Water Quality and Treatment Considerations

This chapter assesses the techno-economic aspects of sourcing and treating water for data center cooling. It analyzes treatment of non-conventional water sources (groundwater, seawater, produced water) using existing commercial solutions and concentrate disposal options, thereby highlighting new water sourcing options for data centers.

Chapter 5 – Regional and Sectoral Comparison of Water Consumption

Chapter 5 defines a simple accounting framework for total data center water use (direct plus indirect), then applies it to ERCOT-provided scenarios to compare data center-related withdrawals against other sectors over time.

Chapter 6 – Transparent and Resilient Data Center Planning: Policy Recommendations

The final chapter presents a strategic roadmap for sustainable water governance in the data center sector. It proposes policies related to water disclosures, usage caps, reuse incentives, and public-private collaboration for regional infrastructure resilience.

Texas's data center market is at the forefront of an unprecedented infrastructure expansion that will significantly influence the state's energy, water, and land resource dynamics. Strategic, forward-looking planning and regulatory coordination are essential to support this growth without compromising grid stability or aggravating regional water scarcity. The issues considered in this white paper will lead to a framework for balancing economic vitality with conscientious resource management in the digital age.

1. Brief Overview of the Impact of Data Centers on Energy and Water Demand

In recent years, the data center industry has undergone a fundamental transformation in both scale and complexity. Initially conceived as relatively compact facilities focused on supporting enterprise IT workloads, modern data centers, particularly those driven by AI, high-performance computing, and hyperscale cloud operations, have evolved into massive industrial infrastructure projects with growing and often conflicting demands on energy, water, land, and material resources. This convergence of pressures introduces profound challenges for engineering design, regulatory oversight, and corporate sustainability strategies [4, 20].

One of the most significant developments is the surge in energy demand, particularly due to the transition from CPU-based architectures to high-density GPU and AI accelerator systems [21]. Although these newer processors offer superior performance per watt, their absolute energy consumption per rack is significantly higher than that of traditional systems. For instance, a typical data center rack in 2010 drew approximately 4–5 kW of power, primarily from CPU-based servers. In contrast, AI-oriented racks equipped with NVIDIA's *Blackwell* or *Hopper* GPUs in 2025 can exceed 80–100 kW per rack, reflecting nearly a 15–20× increase in power density [22, 23]. Despite remarkable improvements in chip-level energy efficiency, modern GPUs deliver over 20× more performance per watt compared to the 2010 Fermi architecture, the overall scale of computation has vastly outpaced all efficiency gains [24]. The net result has been a dramatic rise in both total energy demand and thermal output across the data center sector, thereby increasing the dependency on advanced cooling systems and, by extension, on water. It is estimated that heat loads will approach 1 MW per rack by 2028, which portends a spike in water use by data centers, aggravated by hot, humid conditions in many parts of Texas.

The emergence of gigawatt-scale AI campuses further complicates these dynamics. These facilities, which can demand as much power as a small city, differ markedly from their predecessors in terms of design, cooling requirements, spatial footprint, and utility integration. Hyperscale and AI oriented data centers depend on high-capacity electrical infrastructure and specialized cooling systems whose design choices, such as evaporative towers, chillers, or reclaimed water loops, strongly shape their direct and indirect water footprints at the site and regional scale [25, 26]. Unlike traditional data centers, which aimed for redundancy and modularity, AI-focused campuses are optimized for throughput and density, often trading off flexibility for raw processing power [27, 28]. This evolution has transformed data centers into high-consumption industrial assets, with decisions around siting and design now affecting public infrastructure, regional water supplies, and critical materials supply chains.

Data center operators face a fundamental profit-maximizing tradeoff between investing in effective cooling technologies, minimizing energy and water costs, and meeting sustainability targets, especially when water quantity or quality are constraints. Operators prioritize cooling strategies that reduce both immediate operational expenses and long-term total cost of own-

ership. More effective cooling shrinks energy bills (which can account for over 40% of a data center's total spend) and makes the facility more competitive [21, 29].

At the heart of these developments lies a complex trade-off structure in which energy use, water use, land use, and material use efficiency cannot be simultaneously maximized. Optimizing for one resource often entails compromises elsewhere. An energy-efficient cooling system may be highly water-intensive; a water-conserving approach may raise electricity usage or capital costs; a dense AI configuration may reduce spatial footprint but increase thermal intensity and require the use of rare materials or challenging technologies. These trade-offs are often resolved not through sustainability metrics but through financial optimization: capital and operational expenditure, return on investment, and total cost of ownership considerations. In this context, profit-maximizing behavior frequently privileges efficiency gains that are measurable in utility bills or hardware performance.

Technology selection therefore, should reflect broader economic and infrastructural considerations. Operators prioritize reliability and cost, and in regions where water availability is not constrained by physical conditions or regulation, market incentives often favor options with higher water footprints. Rapid data center growth in such regions can outpace water planning frameworks, limiting transparency and policy oversight [30, 31]. Without standardized reporting of metrics such as water usage effectiveness (WUE), the scale and implications of data center water use will remain difficult for utilities, regulators, and communities to fully assess.

Despite the growing attention to data center-related water impacts, utilities and municipalities continue to face significant information gaps [18]. Most data centers do not readily report detailed facility-level data on water withdrawals, blowdown volumes, discharge quality, or system losses, limiting water suppliers' ability to forecast load growth, infrastructure stress, or treatment capacity needs. This lack of transparency limits municipal water suppliers and wastewater utilities from accurately forecasting demand, planning infrastructure, or managing treatment and regulatory compliance. As a result, discharged water, one of the least studied aspects of the data center footprint, remains poorly integrated into regional resource management and long-term financial planning, increasing uncertainty, operational risks, and the potential for stranded assets. This lack of standardized and comprehensive reporting complicates forecasting, increases uncertainty in utility-side capital planning, and raises the risk of stranded assets, regulatory noncompliance, and unanticipated operational burdens for both water supply and wastewater utilities.

To address the growing water requirements of data centers, a reorientation of design priorities and policy mechanisms is urgently needed. This includes the development of integrated metrics that capture both energy and water efficiency, the use of non-potable or reclaimed water where feasible, and the establishment of water reporting standards analogous to carbon disclosures. From an engineering perspective, investment in adaptive cooling systems, modular infrastructure, and site-specific hydrological planning will be essential to balance performance with water stewardship. Most importantly, the economic models that govern data center expansion must evolve to account for the true cost of water use.

2. The Energy-Water-Data Nexus: Understanding Water Needs for Cooling Data Centers

Data centers are rapidly becoming one of the fastest-growing large loads in the United States, intensifying scrutiny of water use from both onsite cooling and offsite electricity generation [32, 26, 33]. Their water footprint is important as freshwater supplies face pressure from population growth, agriculture, and water-intensive industries such as irrigation and manufacturing [34]. These needs are particularly significant in water-stressed regions of the southern and western U.S. Understanding water requirements, across quantity, quality, technology choice, and local system constraints, is essential for utilities and policymakers planning future development. This chapter establishes a framework for evaluating data center water demand and the associated infrastructure, policy, and resource challenges.

Data center direct water use is largely governed by the cooling technology choices. More importantly, there could be a tradeoff between energy use and water use [35]. This is because evaporative cooling is highly energy efficient but consumes significant amounts of water.

The cooling of a data center is often (but not always) done by two or more cooling technologies working together: the cooling technology employed to directly cool the chips and the cooling technology used to exhaust the heat from the data center. This is why the overall cooling system is divided into the following:

1. **Secondary cooling loop:** How the computing equipment and server racks are cooled.
2. **Primary cooling loop:** How the heat removed from the computing equipment is exhausted to the environment outside the data center.


Cooling of a data center is done by the flow of heat from the computational units to the secondary loop and from the secondary loop to the primary loop. The primary loop ultimately exhausts the heat to the atmosphere as illustrated in Fig 1 (a). Fig 1(b) shows a data center cooled with immersion liquid cooling (secondary loop) and a cooling tower (primary loop).

Secondary cooling loops

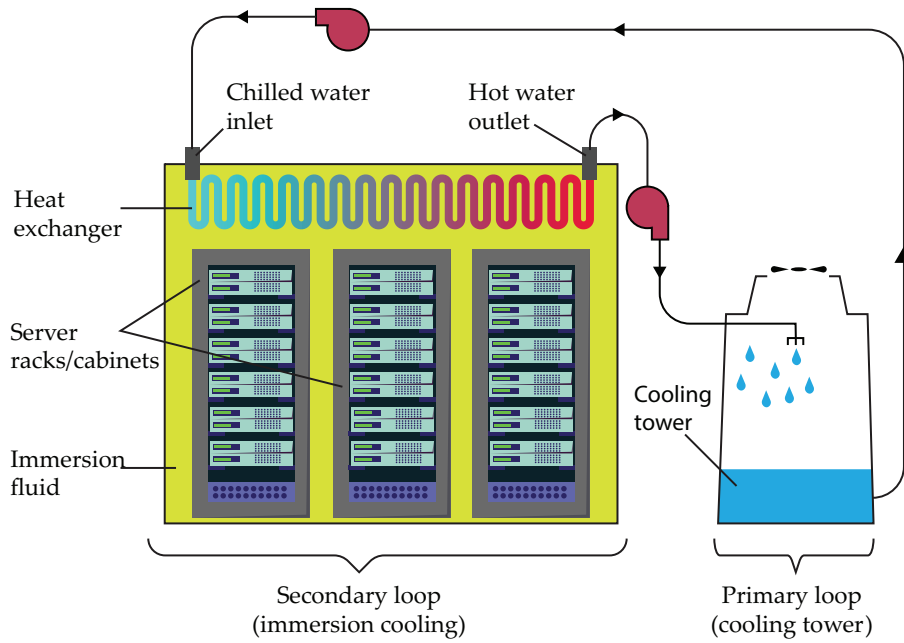
As data center power densities continue to increase, more advanced secondary cooling loops are required. Conventional approaches, such as direct air cooling, are increasingly unable to dissipate the high thermal loads generated by modern high-density server racks.

Interconnection between secondary and primary cooling loops

The total heat rejected by the primary cooling loop (e.g., cooling towers or heat pumps) approximately equals the total facility electrical load, of which IT power is the dominant component, regardless of the secondary cooling technology employed (e.g., cold plates, heat pipes, or liquid



(a)



(b)

Figure 1: (a) Flow of heat from heat sources to heat sink through secondary and primary loop, (b) An example data center in which the secondary cooling loop is liquid cooling and the primary loop is a cooling tower

cooling). However, the temperature level at which this heat is delivered, and the energy consumption of the secondary loop, can significantly affect the overall PUE and WUE. This is why more sophisticated secondary loops are being adopted by the industry, like liquid immersion cooling shown in Fig 1 (b).

Another design-decision aspect to note is that highly efficient secondary cooling loops allow operators to incorporate less efficient primary loops and avoid using evaporative cooling. For example, several data centers are operating with cold plates, heat pipes, and immersion cooling as secondary loops to avoid using cooling towers as a primary loop in order to drastically reduce direct water consumption. More sophisticated secondary cooling loops might allow data centers to have lower PUEs with non-evaporative primary cooling loops like direct air cooling and heat pump cooling.

Table 1 lists some choices for secondary cooling loops.

Table 1: Overview of key secondary cooling loops to cool data centers

Secondary loop cooling technology	PUE Range*	WUE Range* (gal/kWh)	CAPEX	Key Aspects
Direct air cooling (CRAC/CRAH, hot/cold aisle containment)	1.4-1.8	0-1.8	1X	Unable to work with high rack densities (>20-30 kW/rack). Airflow management complexity.
Rear door heat exchangers (RDHx)	1.3-1.6	0-1.5	2-3X	Requires chilled water loop integration. Throttles at very high rack densities (>70 kW/rack).
Direct-to-chip cold plate cooling (liquid cooling/phase change)	1.15-1.35	0-1.2	5-8X	Some components, like RAM, memory, and smaller hotspots, require airflow assistance due to no contact with the cold plate. Risk of leakage.
Single-phase liquid immersion cooling	1.05-1.25	0-1.0	10-15X	Requires specialized dielectric fluids. Maintenance requires fluid drainage and storage.
Two-phase liquid immersion cooling	1.03-1.20	0-1.0	15-20X	High fluid cost. Specialized tank requirement.
In-row liquid cooling units	1.25-1.5	0-1.5	2-4X	Requires chilled water supply. Some dependence on room-level airflow.
Microchannel direct liquid cooling (integrated cold plates)	1.10-1.30	0-1.2	8-12X	Higher pressure drop and pumping requirements. Microchannel damage (fouling, corrosion), complicated maintenance.
Spray cooling	1.05-1.20	0-1.0	15-25X	Control systems complexity. Limited commercial deployment maturity.

*The lower limit of PUE and the upper limit of WUE are based on cooling tower as the primary cooling loop. The upper limit of PUE and the lower limit of WUE (i.e., 0 gal/kWh) are based on a heat pump as the primary cooling loop

Primary cooling loops

Choice of evaporative systems for primary loops, such as cooling towers deliver very high performance and low energy use. However, they rely on continuous water evaporation and blowdown, driving up direct consumption [25]. Air-cooled and air-source heat-pump systems as primary loops can reduce or nearly eliminate direct evaporative water consumption and can achieve low WUE values (0.0053 gal/kWh), but they require much more electricity, often raising indirect water use (power generation) [36, 37]. Hybrid systems that pair cooling towers with heat pumps reduce water losses while maintaining high performance, but still require careful water treatment associated with cooling towers to control scaling, corrosion, and biofouling. There are other options for primary loops, such as geothermal and ocean-source heat sinks as well. Table 2 and Figure 2 summarize these technology-dependent water–energy tradeoffs.

Water remains the most effective medium for transferring heat from electronic equipment to the atmosphere, and most data centers rely on water-based cooling systems to maintain operational reliability. Although water-free alternatives exist, they typically require substantially higher capital investment. As shown in Table 2, widely deployed cooling technologies differ significantly in PUE and WUE, the primary metrics used to assess efficiency. While all deliver adequate thermal performance, their cost and resource footprints vary. Technologies that are still in early development or suitable only for niche applications are not considered in this analysis.

Table 2 highlights a consistent trade-off: reducing direct water use generally increases energy demand and operational cost. Conventional cooling towers dominate current practice due to their technological maturity and energy efficiency, though water use concerns are increasing in some communities. In contrast, air-cooled systems and heat pumps reduce onsite water withdrawals but impose higher energy requirements, especially in hot and humid regions, and higher capital costs, particularly for heat pumps. These additional electricity needs can shift water consumption upstream to power generation. Hybrid configurations seek to balance these effects by combining evaporative cooling with heat pumps. The primary cooling technologies differ not only in energy and water performance, but also in the physical mechanism by which heat is rejected—whether through direct evaporation into supply air, evaporation in a separate condenser loop, purely sensible heat rejection, or mechanical refrigeration. Table 3 distinguishes these mechanisms to clarify how engineering design choices translate into different water–energy tradeoffs. Figure 2 is a graphical summary of the four main cooling technology designs discussed above.

Water quality considerations (such as treatment to prevent scaling, corrosion, and biofouling) and water treatment technologies will add further operational and environmental constraints. Cooling towers require regular monitoring and maintenance, and downtime negatively affects data center reliability. These factors influence the volume and quality of water utilities must provide, and in regions with limited supply, this may increase pressure on local water resources and increase competition among users. Accordingly, the types of cooling systems deployed

Table 2: Overview of key primary cooling loops to cool data centers

Primary loop cooling technology	PUE Range	WUE Range (gal/kWh)	CAPEX	Key Aspects
Direct Evaporative Cooling (Adiabatic cooling)	1.35-1.5	0.3-0.5	X (baseline)	Increasing adoption. Direct evaporative cooling of server room's air reduces heat transfer inefficiencies.
Indirect Evaporative Cooling (Cooling tower)	1.1-1.35	0.46-0.66	X (baseline)	High water consumption due to evaporative losses. High energy efficiency. Being phased out due to water use concerns. Scaling and fouling are significant issues.
Air Cooling (Dry Cooling)	1.6-1.9	0-0.01	0.5X	Minimal water consumption but high PUE. Some water needed to maintain humidity. Not suited for high power density data centers or hot climates.
Heat Pump Cooling	1.3-1.6	0.08-0.13	5X	Energy-efficient cooling that can recover/reuse heat. Moderate water consumption depending on system design. High CAPEX. Not very common in Texas.
Hybrid Cooling (Heat Pump + Evaporative)	1.15-1.35	0.13-0.40	3X	Dynamic switching leverages evaporative cooling during favorable conditions. Lower PUE, WUE compared to direct evaporative cooling.
Alternative Environmental Sources (Geothermal / Ocean / Underground Storage) Cooling	1.2-1.5	0-0.01	10X	Dynamic switching required to switch to a different primary loop when underground or underwater temperatures exceed ambient temperatures. Very high CAPEX.

Table 3: Primary Cooling Loop Technologies Continued: Mechanisms and Tradeoffs

Technology	How Heat Is Rejected	Water Use Pathway	Mechanical Compression?	What Makes It Different
Direct Evaporative Cooling (Adiabatic)	Outside air is drawn into the data hall and cooled directly by evaporating water into the airstream.	Water evaporates directly into supply air.	No	Simplest evaporative method. Outside air mixes with IT air. Efficiency depends heavily on climate humidity.
Indirect Evaporative Cooling (Cooling Tower)	Heat from IT load is transferred to a water loop; heat is rejected outdoors via evaporation in a cooling tower.	Water evaporates in a separate condenser loop.	No (unless paired with chiller)	IT air does not mix with outside air. Higher control and reliability than direct evaporative systems. Most common large-scale approach.
Air Cooling (Dry Cooling)	Heat is transferred to ambient air through dry coolers or fin-fan heat exchangers (no evaporation).	No direct evaporative use.	No	Uses only sensible heat rejection. Performance drops in hot climates; larger equipment footprint required.
Heat Pump Cooling	Mechanical refrigeration cycle compresses refrigerant to move heat to an air- or water-cooled condenser.	Typically none onsite (unless paired with tower).	Yes	Uses electrically driven compressors. Can upgrade waste heat for reuse. Higher energy demand but reduced direct water dependence.
Hybrid (Evaporative + Heat Pump)	System switches between evaporative mode and mechanical compression based on ambient conditions.	Partial evaporative use depending on operating mode.	Yes (in part)	Balances water and energy tradeoffs seasonally. More complex controls.
Alternative Environmental Sources (Geothermal / Ocean / Underground Storage) Cooling	Heat exchanged with stable underground or large water-body temperatures.	Minimal evaporative use.	Sometimes	Highly site-specific. Relies on stable external heat sink; high upfront infrastructure cost.

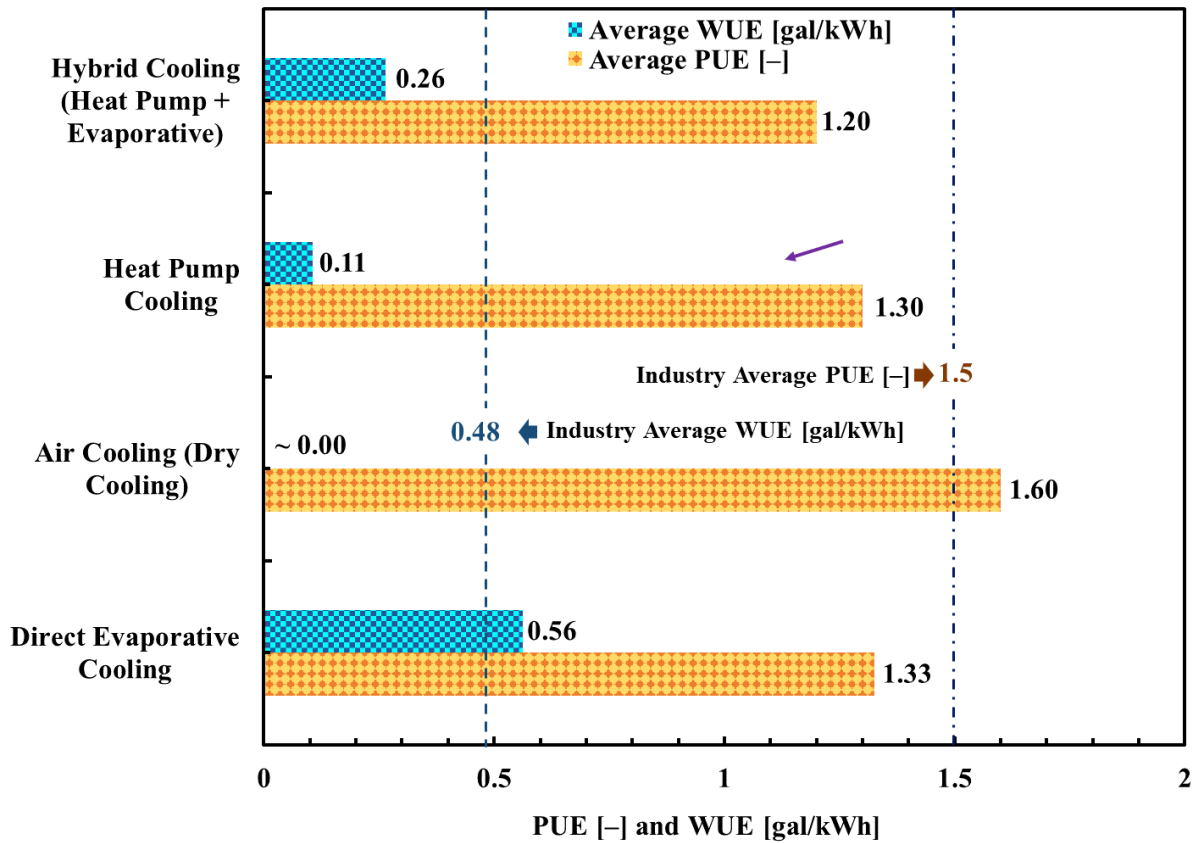


Figure 2: Graphical summary of power usage effectiveness (PUE) and water use effectiveness (WUE) associated with various cooling technologies for data centers.

and the availability of suitable water sources are likely to influence the shaping of the long-term sustainability and resilience of the global data centers.

Figure 3 shows different combinations of secondary and primary loops for data center cooling, resulting in different cooling configurations.

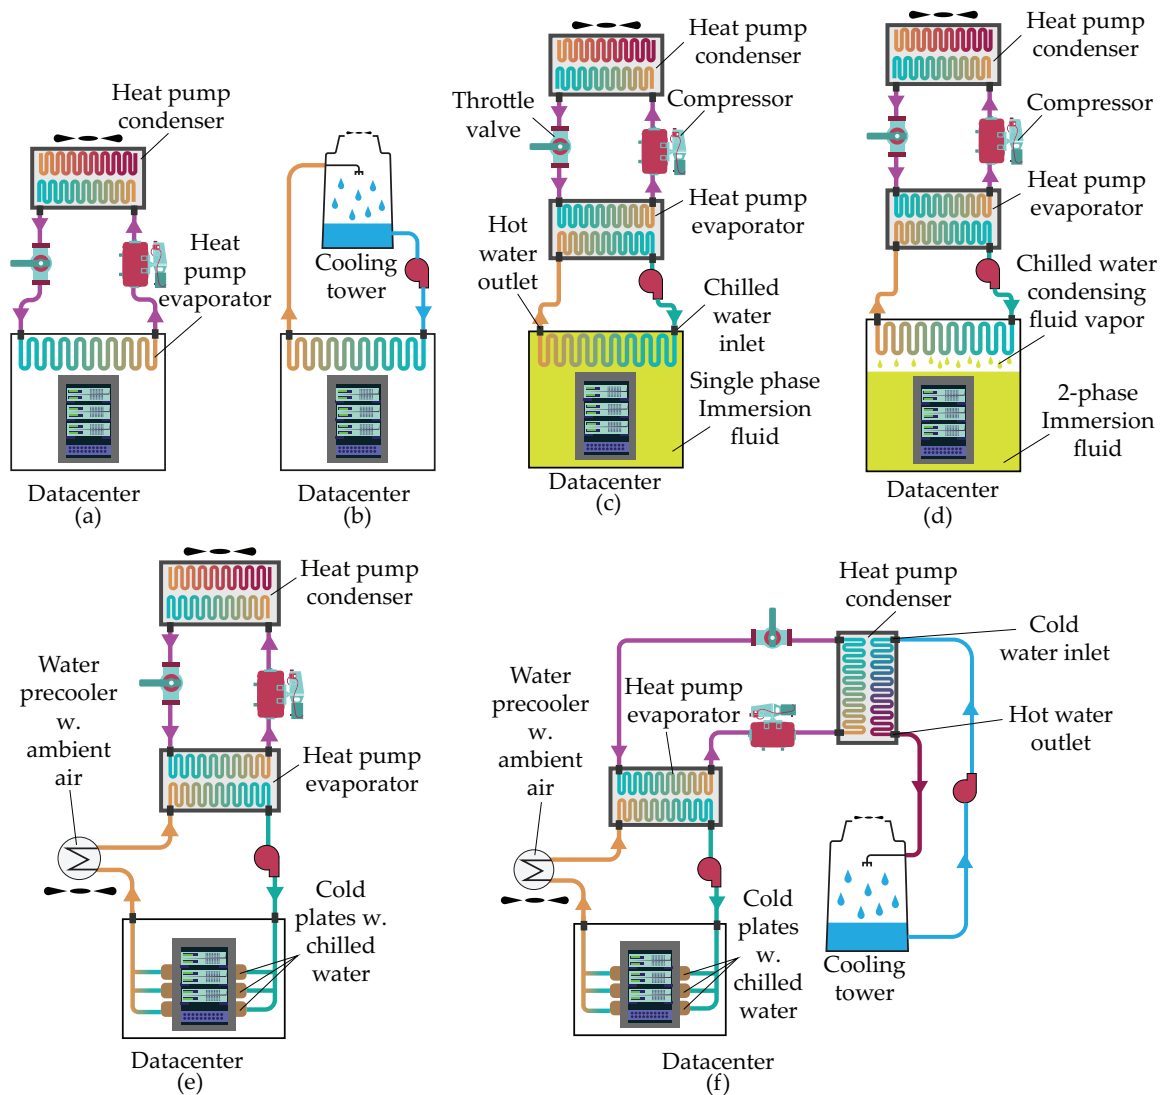


Figure 3: Different cooling configurations for data centers: (a) CRAC with heat pump, (b) CRAC with cooling tower, (c) Single phase immersion cooling with heat pump, (d) 2-phase immersion cooling with heat pump, (e) Cold plates with chilled water loop from heat pump, (f) Cold plates with chilled water loop from cooling tower-sourced heat pump.

3. Understanding Indirect-Use Water Requirements for Data Centers in Texas

An important, and at times unaccounted for, metric that data center operators and regulators need to consider is the secondary water requirements of data centers – specifically, water required to generate the electrical power to operate them.

Texas’s four leading power generation fuel sources (in decreasing percentage) are natural gas, renewables (wind and solar), coal, and nuclear [38]. With natural gas, power generation is either combined cycle-based or traditional gas turbine-based, with combined cycle accounting for a majority share. Water requirements for each of these generation schemes are reported in Figure 4 and Table 5. Metrics are reported in gallons per kilowatt-hour (gal/kWh) for both water *withdrawal intensity*—the total water taken from a source—and water *consumption intensity*. Water consumption is the amount of water withdrawn that is permanently removed from the water source due to process losses, evaporation, or other irreversible losses [36, 39]. While water consumption is undoubtedly a critical metric, water withdrawal is also an important consideration since the communities where these data centers are located will be impacted if power production withdraws large amounts of water.

Natural gas-based power generation will be the dominant technology for independent power infrastructure (termed “behind the meter”) for data centers. Data for power plant water consumption is available from the U.S. Energy Information Administration (EIA) [40]. Cooling systems for power plants are categorized into either open or closed systems. Between these two designations, there are five reported processes in use, three of which are commonly deployed across Texas for natural gas power plants (emphasized in orange in Table 4):




Figure 4: Schematic depictions of an open system once-through cooling with no cooling pond (ON), open system once-through cooling with a cooling pond (OC), and a closed loop recirculating air cooler (RI).

Table 4: Breakdown of cooling system types in use across Texas

Cooling System Type	Description
Closed Cooling System Types	
(DC) Dry Cooling	Air is used to cool the power facility with no water usage; functions similarly to a car radiator.
(RI) Recirculate: Induced Draft	Evaporative cooling via air pulled through the tower by fans mounted on top.
(RC) Recirculate: Cooling Pond	Warm water discharge is cooled in a pond through convection (evaporation) and conduction (mixing with pond volume) before reuse.
Open Cooling System Types	
(OC) Once Through with Cooling Pond	Warm water output is temporarily held in a cooling pond before discharge to the original water source; no reuse.
(ON) Once Through No Cooling Pond	Water is withdrawn from a natural source, used once for cooling, and discharged without reuse.

Water Intensity of Power Generation

Figure 5 shows the 2023 annual average water withdrawal and water consumption intensities for power generated by coal, natural gas, and nuclear.

Natural gas fuel type has the largest average water withdrawal intensity as a result of once-through cooling systems located on the Gulf Coast that withdraw vast quantities of seawater. Modern natural gas power plants use closed cooling systems which withdraw significantly less water. For recirculating induced draft cooling at natural gas power facilities (54% of the reporting facilities), the annual average water withdrawal intensity was 1.17 gal/kWh and water consumption was 0.86 gal/kWh (~75% of water withdrawal rate). For data centers connecting to the grid, their average indirect water requirements for withdrawal are 87 gal/kWh, and average water consumption is 0.96 gal/kWh. Overall annual water consumption for mid- and large-sized data centers ranges from 12 to 660 million gallons. Significant water use for power generation highlights the "hidden" water footprint of data centers and the need for thoughtful water considerations.




Figure 5: Average water intensities across all cooling types for coal, natural gas, and nuclear power production facilities in Texas for 2023.

Table 5: Estimated annual water consumption associated with power generation for data centers. Water consumption depends on fuel type and is estimated for mid-size (60 million kWh/year) and large-sized (600 million kWh/year) data centers.

Fuel Type	Water Consumption (gal/kWh)	Mid-Sized Data Center	Large-Sized Data Center
Coal	0.5 – 1	30 – 60 million gal	300 – 600 million gal
Nuclear	0.5 – 0.6	30 – 36 million gal	300 – 360 million gal
Natural Gas	0.2 – 1.1	12 – 66 million gal	120 – 660 million gal
Wind & Solar PV	~ 0	~ 0	~ 0

4. Water Quality and Treatment Considerations

Recent growth and general acceptance of desalination technologies open up possibilities for the use of saline water streams for data centers. Figure 6 provides a system-level description of a typical saline water desalination system, along with its associated inlet-outlet streams and available input power sources for its operation. As shown in the figure, potential saline streams can be broadly categorized as i) brackish water, ii) seawater, and iii) hypersaline industrial wastewaters, each distinguished by salinity level, geographic occurrence, and treatment complexity. Brackish water, typically containing 1,000–10,000 ppm of total dissolved solids (TDS), is widely distributed in inland aquifers and estuarine interfaces where freshwater and saline flows intermingle. These intermediate-salinity reserves often provide the most practical opportunities for decentralized or inland desalination due to their moderate energy requirements and potential proximity to new data centers. Seawater is a possibility for data centers in coastal regions. Seawater desalination (average TDS: ~35,000 ppm) is well-established and is the basis of large-scale desalination in certain energy-rich parts of the planet.

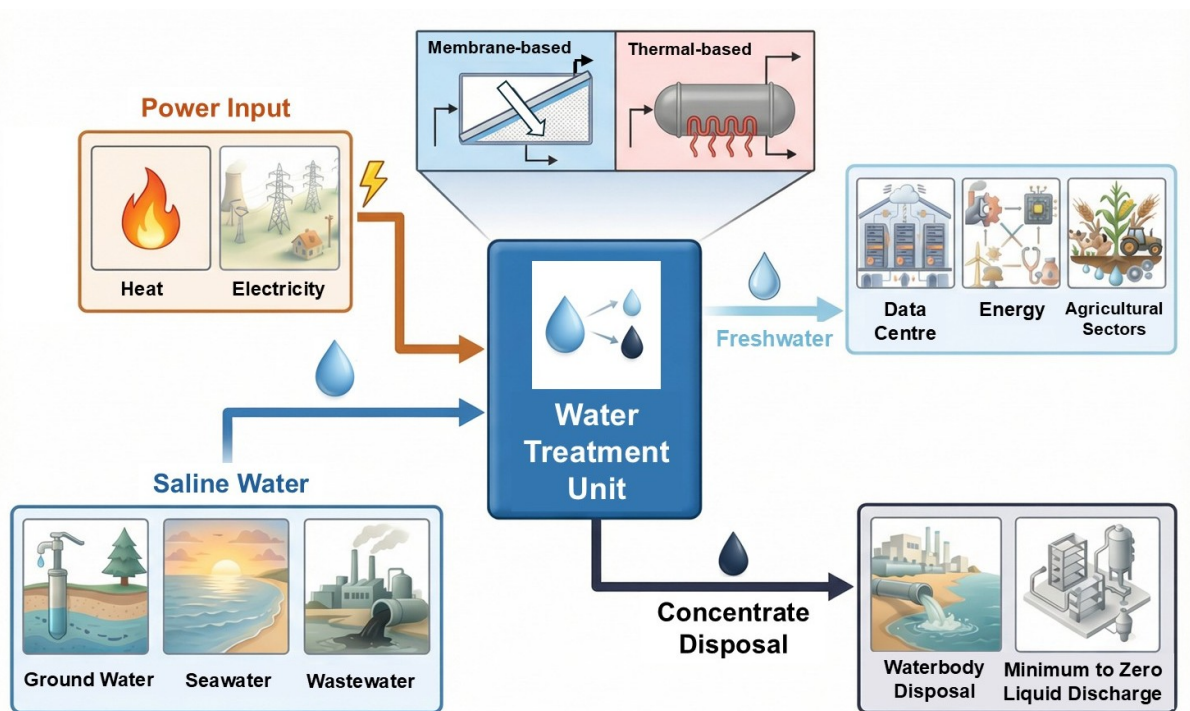


Figure 6: System-level schematic of a desalination system designed to produce freshwater suitable for applications such as data centers, agriculture, etc. The figure also describes the different sources of saline water, power input, and options for desalination-brine disposal.

At the high end of the salinity spectrum, hypersaline industrial wastewaters—notably oilfield-produced water (PW), concentrated mining effluents, and power-plant blowdown streams—can exhibit TDS levels exceeding 50,000 ppm, with some exceeding 200,000 ppm. PW, abundant across energy-producing basins such as the Permian and Eagle Ford in Texas, represents both a disposal challenge and a potential resource. Via advanced desalination processes,

these streams can serve as viable nontraditional water sources for data centers, reducing the stress on conventional freshwater sources while also supporting circular water-use strategies [41, 42].

As Figure 6 shows, desalination technologies broadly fall into two categories: thermal and membrane-based. Thermal techniques like mechanical vapor compression (MVC) and multi-stage flash (MSF) have been widely employed in the Middle East for decades for seawater desalination. It is noted that thermal techniques are highly energy- and CAPEX-intensive; their key advantage is that they are insensitive to feedwater salinity. Membrane-based techniques like reverse-osmosis (RO) are much less energy-intensive than thermal; however, they are subject to membrane degradation issues and can only treat streams while producing substantial freshwater recovery for feed TDS less than 50,000 ppm [42]. While PW can have much higher salinity, it should be noted that 20% of PW in the Permian Basin in Texas has salinity less than 50,000 ppm. Given that the Permian produces more than about 20 million barrels of produced water/day [43, 44], the sheer volume of water is sufficient to justify the use of either thermal or membrane-based water treatment. It is noted that extensive pre- and post-treatment of water will be required for use by data centers. As previously mentioned, this large volume of oilfield waste, when adequately treated, could generate significant freshwater output, while simultaneously mitigating environmentally hazardous brine management risks such as aquifer contamination, soil pollution, and increased seismic activity.

Following desalination, the residual concentrate stream requires careful management, as its disposal strongly determines environmental and sustainability aspects of the entire water treatment project. Any meaningful, holistic analysis of water treatment has to account for concentrate stream disposal, which is subject to regulation. The most conventional practice is discharge to natural water bodies, such as marine outfalls, rivers, or evaporation ponds, where the concentrate—typically containing up to 70,000 ppm of dissolved salts—is diluted and dispersed. The concentration limit for waterbody disposal strategy was determined based on typical desalination concentrate salinity in conventional seawater RO desalination applications (~70,000) and also based on the maximum possible brine concentration attainable by conventional RO process (65,000-75,000 ppm), dictated by its membrane burst pressure limits. For inland or water-stressed locations where direct disposal is not feasible, Minimal Liquid Discharge (MLD) approaches are increasingly adopted to maximize freshwater recovery and minimize the volume of liquid waste. MLD systems generally operate with concentrate salinities in the range up to 200,000 ppm [45]. Beyond improving water utilization, MLD offers a more environmentally benign disposal route, substantially reducing the potential for soil and groundwater salinization compared with conventional discharge. At the most advanced end of the treatment spectrum, Zero Liquid Discharge (ZLD) systems further concentrate the brine, beyond 200,000 ppm [42] up to near saturation limit (260,000 ppm for aqueous sodium chloride), crystallizing the remaining salts for solid waste management. Although ZLD entails greater energy intensity and CAPEX, it provides a closed-loop, zero-release solution that virtually eliminates liquid discharges and enables safer handling of hypersaline produced water and other industrial effluents. These attributes make ZLD attractive for data-center cooling frameworks

pursuing water-neutral and environmentally sustainable operations.

Figure 7 illustrates key results from recent analysis (by Prof. Bahadur's group) of water treatment of various streams, including groundwater, seawater, mild hypersaline PW from Permian Basin, and hypersaline PW from Permian Basin. Freshwater recovery, energy consumption, and costs are quantified as a function of the concentrate discharge technology (disposal to water bodies, MLD, and ZLD). Figure 7a shows that the freshwater yield is highest for brackish-water feed, which can be attributed to its relatively low salt content and high recoverable water fraction. Figure 7a also shows that freshwater yields reduce as the salinity of feedstream increases and the concentrate discharge requirements become less stringent. Clearly, significant quantity of freshwater can be obtained for various feedwater streams.

Figure 7b shows that brackish-water desalination needs the lowest energy; however, its energy demand doubles if the brine management strategy shifts from water-body disposal to ZLD operation, since additional energy is required for higher recovery. In general, MLD and ZLD systems, while being environmentally superior, inherently entail higher energy intensity compared to conventional brine disposal options.

The levelized cost of freshwater for different configurations is shown in Figure 7c. Brackish water desalination has the lowest cost among all cases analyzed. As feed salinity increases, the cost of freshwater production rises accordingly. For instance, for MLD upper-limit operation, the cost for seawater desalination is nearly 2.5 times higher than that of groundwater desalination. For a given feed salinity, increasing the discharge brine concentration from water-body discharge to MLD or ZLD limits raises costs for brackish and seawater cases, owing to the high energy and CAPEX of the thermal techniques required for extended recovery; this is consistent with the energy consumption trends.

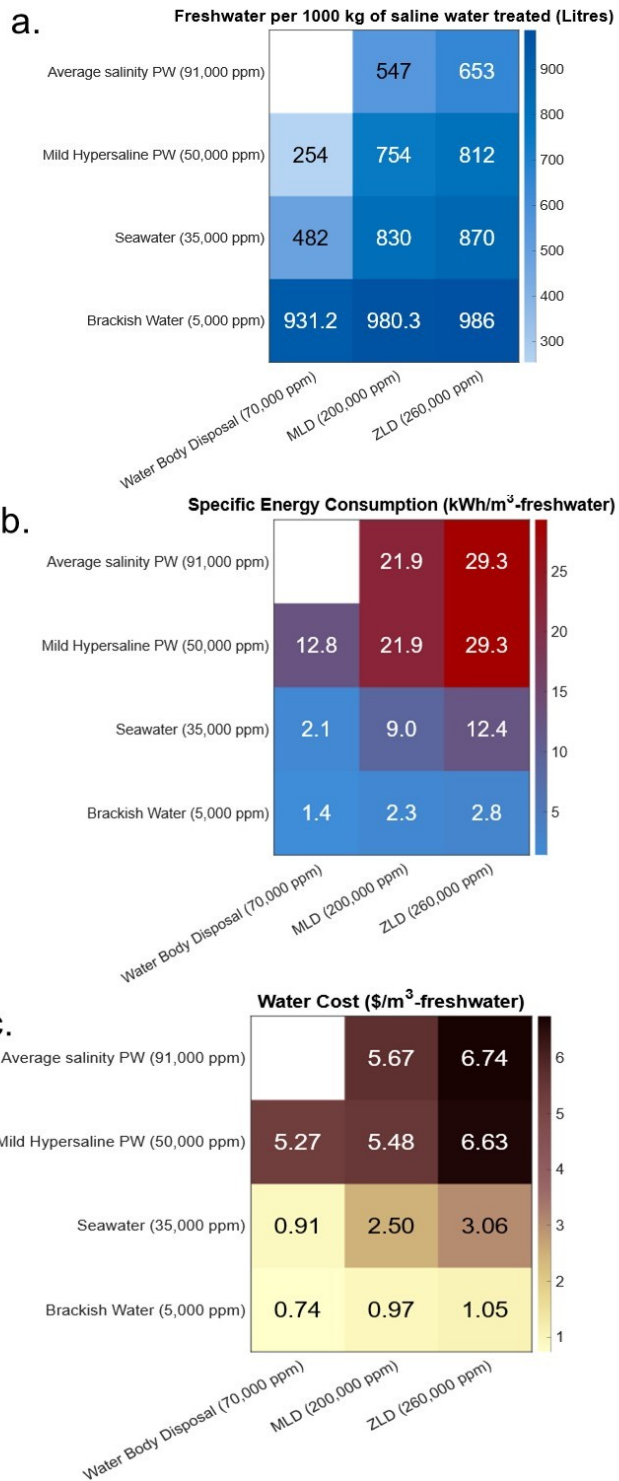


Figure 7: Key performance metrics of conventional desalination systems for various saline water streams and concentrate brine disposal options.

5. Regional and Sectoral Comparison of Water Consumption

Table 6: Projected Water Demand by Category in Texas (2025–2050), Billion Gallons [34]

Category	2025	2030	2040
Irrigation	3,078 (52.5%)	3,058 (50.9%)	2,834 (47.5%)
Livestock	108 (1.84%)	110 (1.83%)	116 (1.94%)
Manufacturing	435 (7.42%)	499 (8.31%)	499 (8.36%)
Mining	132 (2.25%)	132 (2.20%)	118 (1.98%)
Municipal	1,701 (29.0%)	1,897 (31.6%)	2,097 (35.1%)
Steam-electric	304 (5.19%)	304 (5.06%)	304 (5.09%)
Texas Total	5,861	6,003	5,971

Note: Data from 2022 Texas Water Development Board (TWDB) State Water Plan

Table 7: COMPASS Estimates for Data Centers Water Use in Texas (2025–2050), Billion Gallons

Category	2025	2030	2040
Texas Total (Est.)	5,861	6,003	5,971
Data Centers (Est.)	43.5 (0.75%)	223.6 (3.72%)	345.5 (5.70%)

Note: Data center water-use values represent the authors' average estimates across all modeled scenarios. These values are based on engineering calculations using publicly available assumptions on cooling systems, energy use, and representative facility designs. Actual water consumption will vary by site, technology, and operations. The authors do not have access to facility-level, operator-reported water-use data for Texas data centers.

We developed a scenario-based model to estimate Texas data center water use, expressing both direct and indirect consumption as a percentage of statewide demand.

Table 7 presents modeled estimates based on cooling requirements, server utilization, and expected deployment. These values are indicative; actual water use varies with location, cooling technology, climate, and operational strategy. The analysis considers multiple *capacity scenarios*, *facility operational cases*, and *electric grid mixes*. Detailed numbers and description in Appendix Table 8:

- **Capacity Scenarios:** Projected data center power capacities (in Gigawatts) for the years 2025, 2030, and 2040 under Low, Medium, and High growth cases.
- **Facility Cases:** Variations in data center operational efficiency, characterized by server utilization, PUE, and WUE (in gal/kWh). Three cases were considered: Base Case, High Efficiency, and Low Efficiency.
- **Grid Cases:** Different electric grid generation mixes within ERCOT were modeled to capture the variability of indirect water use. The mixes include ERCOT Low (renewables-heavy), ERCOT Moderate (balanced), and ERCOT High (thermal-heavy) scenarios.

Projected water use associated with data centers in Texas is expected to increase substantially under future capacity scenarios. Considering both direct water use for cooling and indirect use

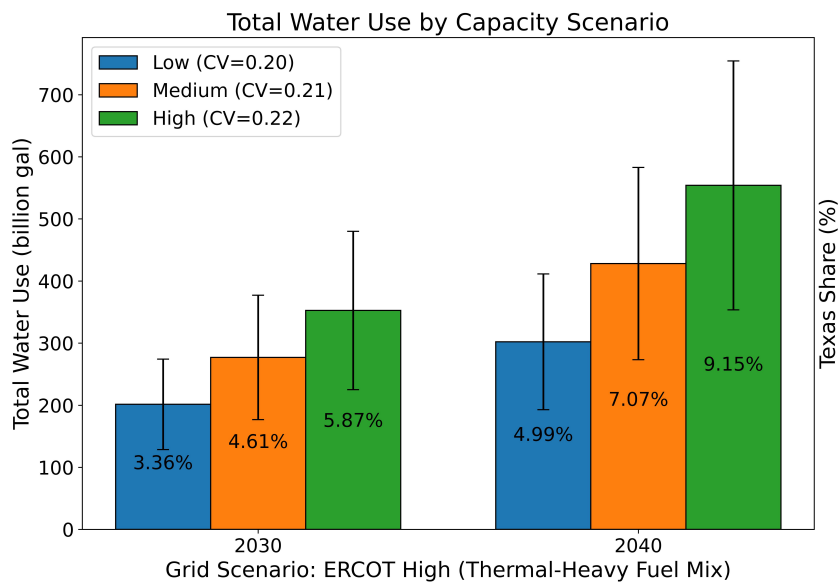
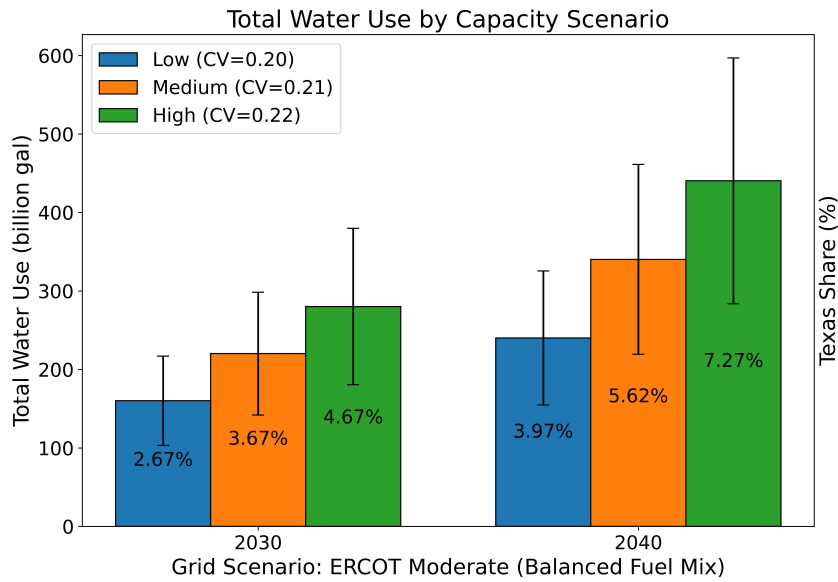



Figure 8: Total water use across capacity scenarios (Low, Medium, High) under three ERCOT grid mixes. Percent values on bars represent each scenario's share of statewide Texas water withdrawals. Error bars show the variation across min/max values. CV (coefficient of variation) is calculated as the ratio of the standard deviation to the mean of the total water-use average.

from power generation, annual withdrawals could rise from roughly 0.75% of statewide demand in 2025 to 3-5% under low-capacity scenarios (40–60 GW) or 5-9% under high-capacity scenarios (70–110 GW) by 2030-2040, approaching levels comparable to several major industrial sectors.

For context, the Texas Water Development Board (TWDB) projects total statewide water demand in 2040 to reach 5,970 billion gallons, with irrigation and municipal use accounting for 47.5% and 35.1%, respectively. Other significant water use sectors include manufacturing (8.36%), steam-electric generation (5.09%), mining (1.98%), and livestock (1.94%). Under the most water-intensive data center scenarios, withdrawals could exceed those of mining and livestock, and likely exceed those associated with steam-electric generation. This signals a potential shift in the composition of Texas’s water demand portfolio.

Although data centers are a minor contributor under low- and moderate-intensity scenarios, high-capacity deployments and thermally intensive grids could make them a significant source of consumptive water use, especially in water-stressed regions. The variability across scenarios also underscores the importance of incorporating energy grid mix, facility efficiency, and PUE/WUE metrics when projecting future consumptive water impacts of data centers.

Figure 8 collectively demonstrates how projected total water use associated with data center operations in Texas varies across three principal dimensions: capacity scenarios, grid fuel mixes, and temporal horizons (2030 versus 2040). Each chart corresponds to a distinct ERCOT grid scenario—Low (renewables-heavy), Moderate (balanced), and High (thermal-heavy)—and depicts the scaling of water withdrawals with increasing data center capacity levels (Low, Medium, High) over time. Coefficients of variation (CV) range from 0.20–0.22 across scenarios, reflecting moderate variability in projected water use due to differences in facility efficiency, capacity assumptions, and grid composition. Higher CV indicates that uncertainty in operational or grid conditions meaningfully affects potential water use estimates.

Table 7 situates these water demands in the context of statewide sectoral withdrawals, demonstrating that, although small relative to irrigation and municipal use, data center-related water demand is comparable to smaller industrial and thermal sectors and will grow in significance under high-demand futures.

Total water use rises with capacity across all grid scenarios. In 2040 under ERCOT High, withdrawals reach ~500 billion gallons for High capacity versus ~300 billion for Low. Thermal-heavy grids drive higher water use due to coal, nuclear, and natural gas dominance. Conversely, the ERCOT Low scenario, characterized by a high share of wind and solar, yields markedly lower water use at equivalent capacity levels. Over time, the proportion of statewide water use attributable to data centers rises, reflecting growing regional pressures. Under the ERCOT High mix, for example, Texas’s share increases from 5.87% in the High capacity scenario in 2030 to 9.15% by 2040. Sensitivity to scenario assumptions is substantial; differences in grid composition alone can produce swings exceeding 180 billion gallons in total water use by 2040 for the same capacity level. Collectively, these results underscore the compounding

effects of capacity expansion and grid composition on water resource demand, highlighting the importance of strategic siting, operational efficiency improvements, and low-water-intensity energy sourcing in future data center development.

6. Transparent and Resilient Data Center Planning: Policy Recommendations

Call for Action: Enhancing Transparency and Cross-Sector Communication

The accelerating expansion of data centers in Texas deepens the urgent need for greater transparency and systematic communication among key stakeholders, including data center operators, utilities, municipalities, state agencies, and private developers. Despite the sector's growing influence on water, energy, and land systems, data relevant to operational water use, grid impacts, and siting constraints remain fragmented or inaccessible. This opacity inhibits effective long-term planning, reduces resilience to resource stress, and increases the likelihood of siting decisions that exacerbate regional vulnerabilities.

Enhanced transparency will improve the fidelity of forecasting models, help identify emerging spatial and temporal resource conflicts before they materialize, and support equitable, resilient infrastructure siting. In doing so, Texas can align private investment with public resource stewardship and foster a more integrated, future-oriented planning environment.

Mapping Water Stress Against Infrastructure Expansion

2022 State of Texas Water Plan projections from the Texas Water Development Board [46] reveal substantial geographic variation in freshwater availability by 2040. When these projections are compared to the locations of operational and proposed data centers, a concerning pattern emerges: a number of large-load facilities are sited in regions projected to face significant water shortages (Figure 9). This spatial mismatch calls for rethinking siting paradigms not only in terms of short-term feasibility but also long-term resilience, system-level interactions, and equity implications for surrounding communities.

Red zones of projected water deficit in West and Central Texas overlap with several announced AI and hyperscale developments, while water-surplus regions in East Texas remain underutilized due to infrastructure gaps or permitting constraints. This distribution highlights a pressing need for proactive coordination between developers and public agencies to align capacity expansion with sustainable resource availability.

Temporal–Spatial Mismatches and Strategic Siting Trade-offs

Effective data center policy must account for the misalignment between rapidly growing load demand and the slower, uneven evolution of water and energy systems. Water availability fluctuates with hydrologic conditions and long-term aquifer trends, while grid and water-infrastructure expansions unfold over multi-year permitting and construction timelines. These temporal mismatches heighten the risk that facilities are approved under conditions that no longer hold when they become operational.

At the same time, spatial disparities create incentives to build in regions with strong grid access but chronic water constraints, while water-resilient regions remain underutilized due to limited transmission capacity or permitting bottlenecks. Without policy intervention, these misaligned incentives reinforce siting patterns that increase long-term vulnerability.

A strategic policy response requires integrated planning tools that combine hydrologic projections, grid development schedules, and permitting timelines into a unified siting framework. Targeted investment in resilient regions, paired with coordinated upgrades to transmission, water conveyance, and regulatory processes, can shift development toward areas better aligned with long-term resource security, reducing the likelihood of stranded assets and community-level impacts.

Infrastructure and Permitting Bottlenecks

Even in water-abundant regions, development is often hindered by regulatory fragmentation and infrastructure bottlenecks. Interconnection queues delay access to power; pipeline and conveyance limitations restrict water delivery; and overlapping jurisdictions slow down permitting. These challenges reflect not only logistical constraints but also deeper governance issues stemming from sectoral silos and misaligned planning timelines.

Addressing these barriers requires harmonization of permitting frameworks, clearer inter-agency coordination mechanisms, and alignment between public-sector timelines and private-sector investment cycles. Without such reforms, Texas risks underutilizing its water-resilient regions while overburdening its constrained ones.

Ripple Effects and Systems-Level Impacts

Data centers exert ripple effects across multiple interconnected systems. Large-scale groundwater withdrawals can lower water tables, affecting agricultural productivity and ecological habitats. Land-use conversions can modify runoff patterns, reduce flood resilience, and alter local microclimates. Housing markets and transportation systems may also experience indirect pressures as workforce and infrastructure demands grow.

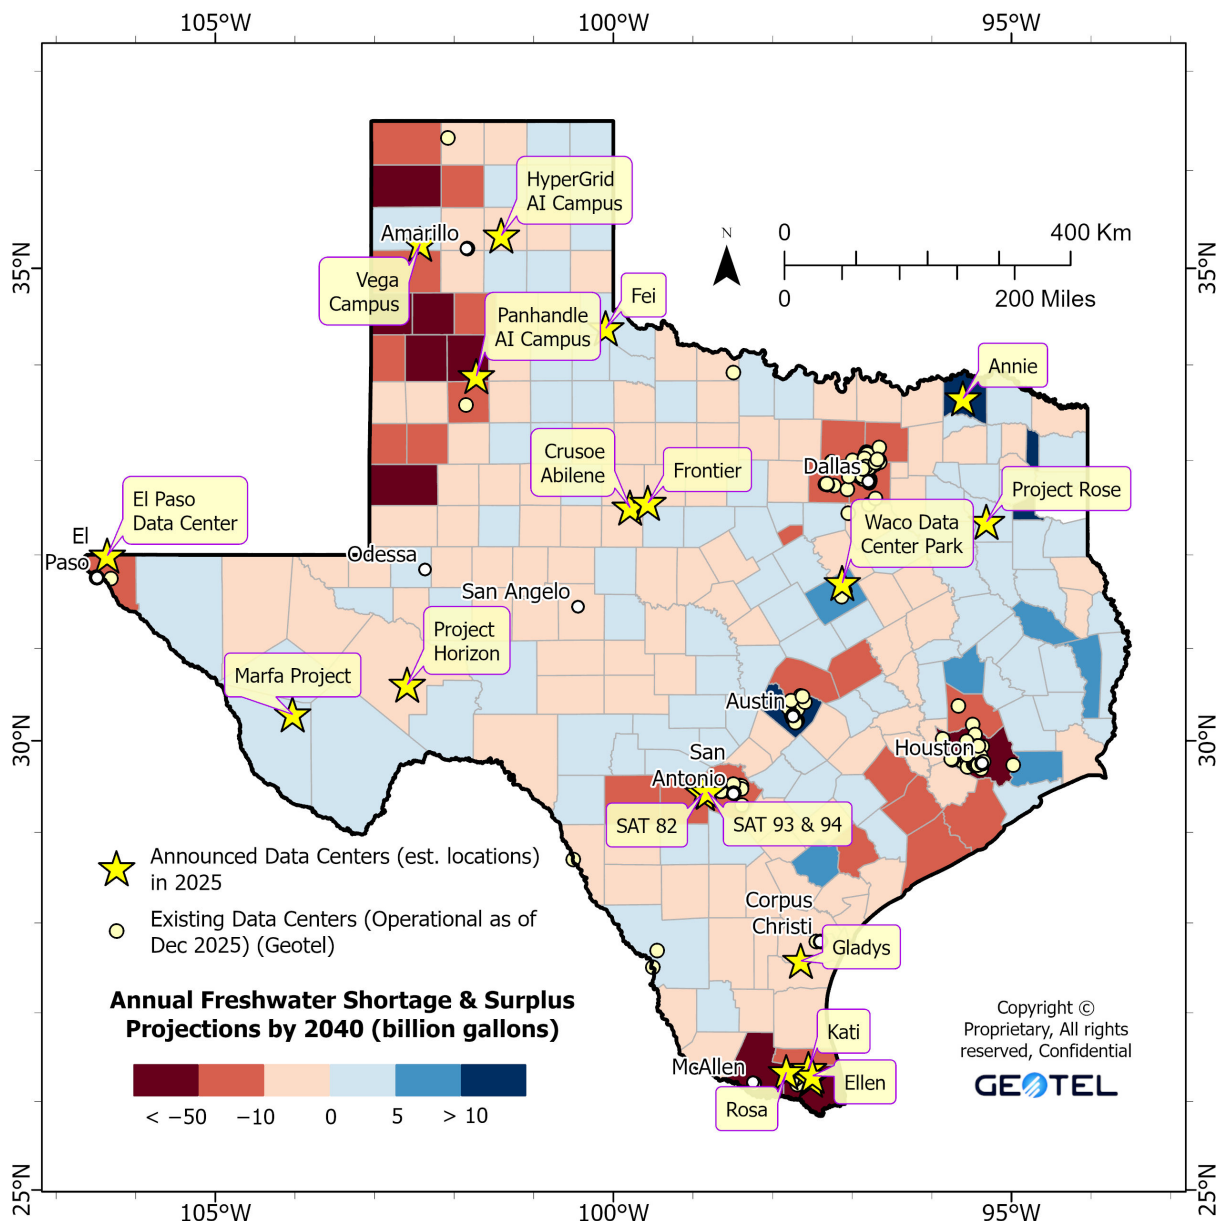


Figure 9: Data Centers and Regional Water Shortage & Surplus Projections by 2040, [34]

COMPASS’s systems-level approach emphasizes scenario resilience, spatial equity, and stakeholder engagement. By modeling interactions across energy, water, land, and infrastructure systems, this framework enables planners to anticipate and mitigate unintended cross-sector consequences of large-load development.

Toward Integrated Planning Frameworks

Transitioning from reactive to proactive planning requires breaking down data and governance silos. Integrated frameworks must combine hydrologic projections, grid capacity models, land-use constraints, and permitting regimes into unified decision-support tools. Such tools should be dynamic, scenario-based, and capable of incorporating uncertainty.

COMPASS is piloting these integrated methods, leveraging geospatial overlays, multi-sector scenario analysis, and structured stakeholder engagement to guide infrastructure investment. The central objective is not merely to assess whether development is technically feasible, but to evaluate whether it is sustainable, strategically aligned, and equitable over the long term.

Future Work

Future research should deepen the technical and institutional foundations required for accurate, forward-looking assessments of data center water use, and evaluation of water sourcing options. One priority is the development of methodologies that translate power draws, compute unit architectures and thermal and reliability considerations to facility cooling and water requirements. As advanced accelerators increasingly dominate data center load profiles, robust modeling will require knowledge of thermal characteristics of compute hardware and compute load demand profiles- such data is usually proprietary; however estimates can be arrived at and are essential for quantifying the full water implications of AI-based computing.

Building on this, future work could also integrate workload-specific thermal and energy models to differentiate the water demands of training, inference, and high-performance computing. Such advances would support development of coupled energy–water–compute system models capable of modeling dynamic interactions between compute hardware, grid characteristics, cooling technologies, and regional hydrologic conditions. Another credible line of inquiry involves comparative technology assessment, evaluating how alternative architectures and advanced cooling systems (e.g., liquid immersion, cold plates) influence water intensity.

In line with the above objectives, our team at the University of Texas at Austin is currently developing a web-based calculator to facilitate water-related planning and decision-making for data centers. The backbone of this calculator will be in-house-developed models and harmonized public datasets to estimate water use and assess water-sourcing options. We envisage that this calculator will be made available to relevant stakeholders (water utilities, industry, local government, community representatives) as required. It should be noted that this will be the first-of-its-kind tool developed for broad use for data center operations.

Future work should also address localized environmental and community impacts, integrating groundwater models, peak-day cooling loads, and land-use dynamics to evaluate how large-scale deployments may affect surrounding agricultural, ecological, and municipal systems. Understanding these localized effects is essential for siting decisions in water-stressed regions.

Finally, additional research is needed on governance and transparency frameworks that facilitate structured data-sharing between semiconductor firms, cloud providers, utilities, and regulators. Developing voluntary or regulatory mechanisms for sharing compute hardware-related data, standardized reporting of embodied water use, and coordinated planning horizons would substantially improve the fidelity of long-term water modeling and support more resilient infrastructure planning.

Contributions

Author	Role	Key Contributions
Dr. Mariam Arzumanyan	Coordination Lead Technical Contributor	Coordination of writing, manuscript drafting and editing, scenario analysis, policy, and regional analysis
Dr. Vaibhav Bahadur	Principal Investigator	Original draft, supervision, validation, manuscript review, overall guidance
Dr. Ning Lin	Principal Investigator	Supervision, manuscript review, validation
Dr. Tahmid Hasan Rupam	Technical Contributor	Energy-water nexus analysis, thermal management requirements
Karey Maynor	Technical Contributor	Indirect water use assessment
Shantanu Katakam	Technical Contributor	Water treatment, quality, and source considerations
Muhammad Usama	Technical Contributor	Cooling technologies, water-energy load calculations
Edna Rodriguez Calzado	GIS Contributor	GIS analysis, Figure 8
Fitsume Wolkeba	Technical Contributor	Water Data and GIS Analysis, Figure 8
Dr. Bridget Scanlon	Technical Contributor	Water shortage and supply analysis, manuscript review
Dr. Justin Thompson	Technical Contributor	Water shortage and supply analysis, manuscript review

Acknowledgments

The authors would like to thank the many colleagues, collaborators, and reviewers whose insights, analyses, and contributions have shaped this white paper. Interdisciplinary collaboration across institutions and sectors made this work possible. We are especially grateful to Dr. Lorena G. Moscardelli, Dr. Mark Shuster, Jay Kipper, Dr. Shuvajit Bhattacharya, Dr. Bridget Scanlon, Dr. Justin Thompson, Roxana Darvari, Dr. Seyyed Hosseini, Linda McCall, Dr. Tingwei “Lucy” Ko, Edna Rodriguez Calzado, David Chapman, Dr. Nur Schuba, Dr. Jani Das, Dr. Yashvi Malhotra, Shipra Mishra, and Mark Blount of the University of Texas at Austin for their support and guidance. We also thank our colleagues at the Texas Advanced Computing Center (TACC), the Energy Institute, and the Cockrell School of Engineering for their technical expertise and critical input.

We would like to thank the Bureau’s Media Production Services team for their support in the design and production of this report. We also gratefully acknowledge the editorial and review

assistance that helped strengthen the clarity and quality of the manuscript.

We extend special thanks to our external reviewers for their thoughtful feedback and contributions, including Marisa Bruno (Hill Country Alliance), Michael Obradovitch (Ecolab), and Anagha Uppal (PowerHouse Texas).

This report would not have been possible without the support of our extended research team and institutional partners.

Contact: compass@beg.utexas.edu

References

- [1] Aterio. *Aterio Report on U.S. Data Center Developments — September 2025*. <https://www.aterio.io/blog/aterio-report-on-us-data-center-developments-september-2025>. Accessed: October 20, 2025. 2025. URL: <https://www.aterio.io/blog/aterio-report-on-us-data-center-developments-september-2025>.
- [2] Baxtel. *Texas Data Centers*. <https://baxtel.com/data-center/texas>. Accessed: October 20, 2025. 2025. URL: <https://baxtel.com/data-center/texas>.
- [3] Mariam Arzumanyan et al. “Geospatial suitability analysis for data center placement: A case study in Texas, USA”. *Sustainable Cities and Society* (2025), p. 106687.
- [4] Ning Lin et al. *Data Center Growth in Texas: Energy, Infrastructure, and Policy Pathways — A White Paper on the Future of Digital Infrastructure and Energy Sustainability in Texas*. Tech. rep. Austin, Texas: Bureau of Economic Geology, The University of Texas at Austin, May 2025. URL: https://www.beg.utexas.edu/files/cee/Data_Center_White_Paper_BEG.pdf.
- [5] Texas Real Estate Research Center. *AI, Land, and Power: Inside Texas’ Data Center Explosion*. Accessed: October 20, 2025. 2025. URL: <https://trerc.tamu.edu/blog/ai-land-and-power-inside-texas-data-center-explosion/>.
- [6] Electric Reliability Council of Texas (ERCOT). *Long-Term Load Forecast Update 2025-2031 and Methodology Changes*. <https://www.ercot.com/files/docs/2025/04/07/8.1-Long-Term-Load-Forecast-Update-2025-2031-and-Methodology-Changes.pdf>. Accessed March 3, 2026. 2025.
- [7] J. Qin and S. Yoon. “The Water Use of Data Center Workloads: A Review and Assessment of Current Practices”. *Resources, Conservation and Recycling* 212 (2025), p. 107872. DOI: [10.1016/j.resconrec.2025.107872](https://doi.org/10.1016/j.resconrec.2025.107872).
- [8] R. Kseibati. “Cooling Innovation and Circularity: Addressing Water Stress in the Age of AI-Driven Data Centers”. MA thesis. Massachusetts Institute of Technology, Center for Real Estate, 2025.
- [9] Steven Carlini. *The 1 MW AI IT rack is coming, and it needs 800 VDC power*. Accessed March 3, 2026. Oct. 2025.
- [10] Brent McDonald. *Data centers evolve to meet AI’s massive power needs*. <https://www.ti.com/lit/ta/ssztdb4/ssztdb4.pdf>. Technical article (SSZTDB4 – November 2025), accessed March 3, 2026. Texas Instruments, 2025.
- [11] Zhihao Zhang, Xuehui Wang, and Yuying Yan. “A review of the state-of-the-art in electronic cooling”. *e-Prime-Advances in Electrical Engineering, Electronics and Energy* 1 (2021), p. 100009.
- [12] J. Friedmann. “Data Center Water Use”. *ICEF AI for Climate Change Mitigation Roadmap (Second Edition)*. 2024.
- [13] Nature Editorial Team. “AI & Robotics Briefing: Data Centres’ Huge ‘Water Footprint’ Becomes Impossible to Ignore”. *Nature News Feature* (Nov. 2023). URL: <https://www.nature.com/articles/d41586-023-03768-y>.
- [14] L. Hoffman and Y. Zhang. “Impact of Generative AI on Water Resources Used to Cool Data Centers”. *International Journal of Progressive Research in Engineering and Management Sciences* 3.11 (2024), p. 36551.
- [15] Electric Reliability Council of Texas (ERCOT). *2024 State of the Grid*. <https://www.ercot.com/files/docs/2025/03/14/ERCOT-2024-State-of-the-Grid.pdf>. Accessed March 3, 2026. Apr. 2025.

- [16] Bridget R Scanlon et al. “Controls on water use for thermoelectric generation: case study Texas, US”. *Environmental science & technology* 47.19 (2013), pp. 11326–11334.
- [17] Margaret Cook. *Thirsty Data and the Lone Star State: The Impact of Data Center Growth on Texas’s Water Supply*. 2025. URL: <https://harcresearch.org/wp-content/uploads/2026/01/Thirsty-Data-Water-Use-and-The-Projected-Data-Center-Boom-in-Texas.pdf>.
- [18] Alejandra Martinez and Jayme Lozano Carver. *Data centers are thirsty for Texas’ water, but state planners don’t know how much they will need*. Accessed March 3, 2026. Sept. 2025.
- [19] Bridget Scanlon and Landon Marston. *Increasing Water Consumption in Data Centers Amplifies Local Water Scarcity*. Accessed March 3, 2026. 2025.
- [20] Adam Zewe. *Explained: Generative AI’s environmental impact*. Accessed March 3, 2026. Jan. 2025.
- [21] Uptime Institute Intelligence. *Uptime Institute Global Data Center Survey 2024*. Tech. rep. Keynote Report 146M; Accessed March 3, 2026. Uptime Institute, July 2024.
- [22] Drew Robb. *Data Center Rack Density: How High Can It Go?* Accessed March 3, 2026. Sept. 2023.
- [23] Caitlin Johnstone. *How Power Density Is Changing in Data Centers and What It Means for Liquid Cooling*. Accessed March 3, 2026. Apr. 2024.
- [24] Navitas Semiconductor. *Nvidia’s Grace Hopper Runs at 700 W, Blackwell Will Be 1 kW: How Is the Power Supply Industry Enabling Data Centers to Run These Advanced AI Processors?* Accessed March 3, 2026. Oct. 2024.
- [25] Miguel Yañez-Barnuevo. *Data Centers and Water Consumption*. Environmental and Energy Study Institute (EESI), accessed March 3, 2026. 2025. URL: <https://www.eesi.org/articles/view/data-centers-and-water-consumption>.
- [26] Bluefield Research. *U.S. Water for Data Centers: Market Trends, Opportunities, and Forecasts, 2025–2030*. <https://www.bluefieldresearch.com/research/u-s-water-for-data-centers-market-trends-opportunities-and-forecasts-2025-2030/>. 2025.
- [27] Alexandra Jonker and Alice Gomstyn. *What Is an AI Data Center?* Accessed March 3, 2026. 2025.
- [28] Juan Pedro Tomás. *AI Data Center vs Traditional Data Center – What Is the Difference?* Accessed March 3, 2026. Mar. 2025.
- [29] TPR Staff. *Texas’ Data Center Dilemma: Growth, Grid Strain, and Water Costs*. Accessed March 3, 2026. 2025.
- [30] Md Abu Bakar Siddik, Arman Shehabi, and Landon Marston. “The environmental footprint of data centers in the United States”. *Environmental Research Letters* 16.6 (2021), p. 064017.
- [31] Liang Chen et al. “Characterizing Data Center Cooling System Water Stress in High Demand Geographies”. *Environmental Science & Technology* 56.15 (2022). Accessed March 3, 2026, pp. 10500–10513. DOI: 10.1021/acs.est.2c01000. URL: <https://par.nsf.gov/servlets/purl/10287008>.
- [32] Amazon. *Amazon Sustainability Report 2024*. <https://sustainability.aboutamazon.com/2024-amazon-sustainability-report.pdf>. 2024.
- [33] AESO (Alberta Electric System Operator). *AESO Data Centre Update*. <https://www.aeso.ca/assets/AESO-Data-Centre-Update-March2025.pdf>. 2025.

- [34] Texas Water Development Board. *2022 State Water Plan: Water for Texas*. <https://www.twdb.texas.gov/waterplanning/swp/2022/docs/SWP22-Water-For-Texas.pdf>. Adopted 7, July 2021; accessed March 3, 2026. July 2021.
- [35] Nuoa Lei and Eric Masanet. "Climate- and technology-specific PUE and WUE estimations for U.S. data centers using a hybrid statistical and thermodynamics-based approach". *Resources, Conservation and Recycling* 182 (2022), p. 106323. ISSN: 0921-3449. DOI: <https://doi.org/10.1016/j.resconrec.2022.106323>. URL: <https://www.sciencedirect.com/science/article/pii/S0921344922001719>.
- [36] A. Shehabi et al. *United States Data Center Energy Usage Report*. Tech. rep. LBNL-1005775, 1372902. Lawrence Berkeley National Laboratory, 2016. DOI: [10.2172/1372902](https://doi.org/10.2172/1372902).
- [37] Arman Shehabi et al. *2024 United States Data Center Energy Usage Report*. 2024. URL: <https://escholarship.org/uc/item/32d6m0d1>.
- [38] ERCOT. *Interval Generation by Fuel Report*. Accessed: Nov. 12, 2025. [Excel Workbook]. 2024. URL: <https://www.ercot.com/gridinfo/generation>.
- [39] Jani Das et al. "Comparative life-cycle assessment of electricity-generation technologies: West Texas case study". *Energies* 17.5 (2024), p. 992.
- [40] U.S. Energy Information Administration. *Electricity - Thermoelectric cooling water data*. Accessed: Oct. 15, 2025. [Online]. 2025. URL: <https://www.eia.gov/electricity/data/water/>.
- [41] Vishnu Sree Shanthanu Katakam and Vaibhav Bahadur. "Reverse osmosis-based water treatment for green hydrogen production". *Desalination* 581 (2024), p. 117588. DOI: <https://doi.org/10.1016/j.desal.2024.117588>.
- [42] Vishnu Sree Shanthanu Katakam and Vaibhav Bahadur. "Heat-mass transfer analogy-based design method for Osmotically Assisted Reverse Osmosis membrane mass exchangers". *Desalination* (2025), p. 119567. DOI: <https://doi.org/10.1016/j.desal.2025.119567>.
- [43] U.S. Department of Energy. *GC Regional Report*. Accessed: 2026-03-01. U.S. Department of Energy, Nov. 2024. URL: https://www.energy.gov/sites/default/files/2024-11/GC%20Regional%20Report_10.29.24.pdf.
- [44] U.S. Department of Energy. *Produced Water from Oil and Gas Development and Critical Minerals Fact Sheet: Oil and Gas Development Challenges to Recovering Critical Minerals from Produced Water — Concentration of Critical Minerals in Produced Water from Shale Plays*. Accessed: 2026-03-01. U.S. Department of Energy, June 2024. URL: https://www.energy.gov/sites/default/files/2024-06/Produced%20Water%20from%20Oil%20and%20Gas%20Development%20and%20Critical%20Minerals%20Fact%20Sheet_6.18.24.pdf.
- [45] Zhuoran Zhang et al. "Comparative techno-economic assessment of osmotically-assisted reverse osmosis and batch-operated vacuum-air-gap membrane distillation for high-salinity water desalination". *Desalination* 532 (2022), p. 115737.
- [46] Texas Water Development Board. *2022 State Water Plan*. <https://2022.texasstatewaterplan.org/>. Accessed February 5, 2026. 2022.
- [47] International Energy Agency (IEA). *AI is set to drive surging electricity demand from data centres while offering the potential to transform how the energy sector works*. <https://www.iea.org/news/ai-is-set-to-drive-surging-electricity-demand-from-data-centres-while-offering-the-potential-to-transform-how-the-energy-sector-works>. [Accessed: 2025-12-03]. 2025.

- [48] McKinsey & Company. *Scaling bigger, faster, cheaper data centers with smarter designs*. <https://www.mckinsey.com/industries/private-capital/our-insights/scaling-bigger-faster-cheaper-data-centers-with-smarter-designs>. Accessed: 2025-12-03, by Erikhans, Kok and Johan, Rauer and Pankaj, Sachdeva and Piotr, Pikul and Dave, Sutton and Rawad, Hasrouni. Aug. 2025.
- [49] U.S. Department of Energy. *Data Centers | Better Buildings Solution Center*. <https://betterbuildingssolutioncenter.energy.gov/sectors/data-centers>. Accessed: 2025-12-03. 2025.

Appendix

The rapid expansion of data centers has introduced significant challenges in water management, both from direct operational requirements and from the indirect demands of electricity generation. To quantify water consumption, we consider the total water use of a data center as the sum of its direct and indirect components:

$$W_{\text{total}} = W_{\text{direct}} + W_{\text{indirect}} \quad (1)$$

Direct water use is primarily associated with cooling systems and other on-site processes, and can be estimated using the water usage effectiveness (WUE) metric, defined as:

$$WUE = \frac{\text{Annual Water Use (gal)}}{\text{Annual IT Energy Use (kWh)}} = \frac{W_{\text{direct}}}{E_{\text{IT}}} \quad (2)$$

From this definition, the annual direct water consumption is calculated as:

$$W_{\text{direct}} = WUE \times E_{\text{IT}} = WUE \times \frac{P_{\text{total}} \times U \times 8760}{PUE} \quad (3)$$

where E_{IT} denotes the annual energy used by the IT load of the facility, P_{total} represents the total installed facility power (kW), U is the utilization factor (0–1), and 8760 is the total number of hours in a year.

Indirect water use is attributable to the water consumed by the electricity generation required to power the data center. It can be estimated based on the grid's water intensity, I_{grid} , using:

$$W_{\text{indirect}} = P_{\text{total}} \times U \times I_{\text{grid}} \quad (4)$$

We obtain the total water footprint of a data center by combining these components, which provides a basis for comparison across different regional grids, capacity scenarios, and efficiency measures. This framework facilitates a sectoral analysis that highlights the relative contributions of direct operational consumption versus grid-dependent water use, and allows for meaningful benchmarking against other industrial and municipal water demands within a region. Moreover, examining variations in WUE, PUE, utilization, and grid water intensity allows researchers to assess the sensitivity of total water consumption to technological and operational improvements, supporting evidence-based planning and sustainable resource management.

Water Use Computation Data Normalization and Texas Share

For each scenario, total water use was normalized to Texas' statewide water demand (in billion gallons per year) to estimate the relative share of data center water consumption:

$$\text{Texas Share (\%)} = \frac{W_{\text{total}}}{W_{\text{Texas}}} \times 100.$$

The model was implemented in Python. All numeric outputs were stored in a structured DataFrame, including:

- Direct, indirect, and total water use (gal)
- Minimum, maximum, and average water use
- Texas share (%) for each scenario

Results were summarized by year, capacity scenario, facility case, and grid case, providing a comprehensive assessment of potential water impacts from future data center growth in Texas.

Table 8: Scenario Definitions for Data Center Water Use Analysis

Scenario Type	Name	Defining Parameter	Reference Notes
Total facility capacity (MW) [47, 48]			
Capacity Scenario	Low	2030: 40,000; 2040: 60,000	Reflects conservative growth assumptions consistent with moderate deployment trends.
Capacity Scenario	Medium	2030: 55,000; 2040: 85,000	Represents the expected market-average expansion of data center capacity based on industry projections.
Capacity Scenario	High	2025: 8,000; 2030: 70,000; 2040: 110,000	Captures a high-growth scenario driven by aggressive AI and cloud infrastructure adoption.
[21, 32, 49]			
Facility Case	Base Case	Utilization 0.85, PUE 1.15, WUE 0.25	Standard efficiency data center reflecting current operational norms.
Facility Case	High Efficiency	Utilization 0.90, PUE 1.10, WUE 0.20	Optimistic efficiency case reflecting state-of-the-art design and cooling technologies.
Facility Case	Low Efficiency	Utilization 0.75, PUE 1.25, WUE 0.35	Less efficient scenario representing older or sub-optimally managed facilities.
Generation mix [40]			
Grid Case	ERCOT Low	Wind 45%, Solar 25%, NG 25%, Coal 3%, Nuclear 2%	Represents a future grid with high renewable penetration, minimizing indirect water use.
Grid Case	ERCOT Moderate	Wind 25%, Solar 25%, NG 40%, Coal 5%, Nuclear 5%	Reflects a balanced grid consistent with projected ERCOT generation forecasts.
Grid Case	ERCOT High	Wind 7%, Solar 22%, NG 46%, Coal 15%, Nuclear 10%	Thermal-heavy grid scenario reflecting historical reliance on fossil and nuclear generation.