TETRA TECH

SUSTAINABLE DATA CENTERS: MODELING WATER AND POWER DEMANDS FOR COOLING

Michael Weller, Ken Nichols
Tetra Tech
8911 N. Capital of Texas Highway, Suite 2310
Austin, TX 78759

Leading with Science®

ABSTRACT

Texas is experiencing accelerated growth in data center capacity driven by the rise of hyperscale computing facilities and the search for available power. Currently, the best available water demand estimates for data centers rely on nation-wide annual averages. This paper presents an analysis using building energy modeling to estimate data center water and power demands on an hourly basis. Multiple data center cooling solutions with varying levels of evaporative cooling were modeled to demonstrate the variability of water consumption based on cooling system selection. The models were run for theoretical sites in Houston and El Paso, which were selected to represent the range of Texas climates from hot and humid, to warm and dry. Significant variations were observed in hourly and seasonal water consumption, indicating the need for more detailed modeling efforts both for more detailed modeling efforts for infrastructure planning and local and time-dependent impact analysis, as well as to highlight sustainable design options available to data centers.

KEYWORDS

Data centers, power usage effectiveness, PUE, water usage effectiveness, WUE, sustainability, hyperscale, evaporative cooling, energy modeling, water modeling, demands.

INTRODUCTION

In 2023, data centers were reported to be responsible for 4.4% of the U.S.'s electricity consumption (Shehabi et al., 2024). The same report estimates that U.S. data center electricity consumption could double or triple by 2028, increasing its share of the country's electricity consumption past 10%. In fact, the total computing capacity of data centers under construction right now is nearly equal to the total data center capacity currently in operation (Batson, 2026). Much of the new data center capacity in the U.S. will be housed in hyperscale data centers that typically support cloud-based computing or artificial intelligence (AI). In search of abundant power and land, hyperscale data centers are increasing capacity in frontier markets such as Texas and the Midwest (Synergy Research Group, 2026). In Texas, hyperscale data centers are expected to more than double the states data center capacity by 2030, potentially causing Texas to surpass Virginia as the largest data center market in the U.S. (Batson, 2026). Increasing energy demand for data centers will increase water demand both at data center sites and at the power generation plants. Texas is a leader in water supply planning, with the Texas State Water Plan updated every five years.

The Draft 2027 Texas State Water Plan lays out the statewide, drought-of-record planning process that evaluates future demand, existing drought-reliable supplies, shortages, and feasible management strategies across the 50-year horizon from 2030 to 2080. Using statewide methodologies and regional water planning group input, the Texas Water Development Board projects Texas population growth from 34.2 million to 52.3 million (+53%), total water demand from 17.4 million acre-feet per year to 18.4 million acre-feet per year (+6%), and existing drought-reliable supplies from 15.5 million acre-feet per year to about 14.0 million acre-feet per year (-10%), largely due to aquifer depletion (Texas Water Development Board, 2026). Under drought-of-record conditions, water user groups face potential shortages of 3.6 million acre-feet per year in 2030 and 5.8 million acre-feet per year in 2080, and even after recommended strategies are implemented, statewide unmet needs remain substantial at 2.4 million acre-feet per year in 2030 and 1.6 million acre-feet per year in 2080. Based on these projections, it will be critical to understand the variables impacting data center water and energy demands so that they can be developed responsibly and sustainably.

DATA CENTER PERFORMANCE

Data center performance metrics are explained in brief in this section. This paper will utilize these metrics for reporting power and water consumption. While water consumption is the primary purpose of this paper, energy performance is included for a wholistic evaluation of sustainability, but also due to the indirect water consumption caused by data centers at the point of power generation.

Energy Consumption

Data center performance metrics are explained in brief in this section. This paper will utilize these metrics for reporting power and water consumption. While water consumption is the primary purpose of this paper, energy performance is included for a wholistic evaluation of sustainability, but also due to the indirect water consumption caused by data centers at the point of power generation.

The Power Usage Effectiveness metric, or PUE, is the ratio of data center's total energy consumption to the energy consumption of just the information technology equipment (ITE) (The Green Grid, 2012). The PUE can be interpreted as the additional amount of power required by the data center to operate the ITE. A theoretical data center with a PUE of 1.0 would mean that all energy to the site was used entirely at the servers and IT equipment. A data center with a PUE of 1.10 means an additional 10% is required to operate the rest of the facility. In 2023, the average data center PUE over a 12-month reporting period in the U.S. was approximately 1.40 (Shehabi et al., 2024).

PUE = Total Data Center Energy (kWh) / ITE Energy Consumption (kWh)

Water Consumption

Water efficiency is often assessed using the Water Usage Effectiveness metric, or WUE. The WUE represents the total volume of water consumed by the data center (including process cooling, potable domestic water, irrigation, etc.) on a per unit of ITE energy basis (The Green Grid, 2011). The units of WUE are Liters per kWh. The energy consumption in the denominator for the WUE formula is the energy consumption of just the IT equipment, the same as the PUE calculation. In 2023, the average data center WUE was approximately 0.38 L/kWh (Shehabi et al., 2024). Note that not all data centers consume water during the normal operation of their cooling systems, so the WUE can vary significantly based on the data center's cooling technology.

WUE = Total Data Center Water Consumption (L) / ITE Energy Consumption (kWh)

Energy and Water Trade-Offs

The dominant consumer of water in a data center is typically the evaporative cooling process, if used as part of the process cooling system. The use of evaporation allows a data center to reduce, or potentially eliminate, the need for compressor-based mechanical cooling (i.e., refrigeration equipment). This means that efforts to reduce water consumption in data center cooling systems may lead to an increase in cooling energy consumption. For societal and state-wide planning purposes, it is important to highlight that an increase in site energy consumption will generally lead to an increase in water consumption at the point of power generation. The amount of water consumed by the electrical grid per unit of energy produced is based on time-of-use, generation technology, and generation cooling technology, among other factors. Estimating this upstream water consumption is outside the scope of this paper, but the results of this paper can be used to calculate these demands based on scalar factors.

MODELING DATA CENTER ENERGY AND WATER DEMAND

Data center energy and water demands can be modeled using building simulation engines that use physics-based models to represent the thermal performance of the data center environment and associated cooling equipment. For this analysis, the APACHE simulation engine in the IES Virtual Environment (IESVE) was utilized. IESVE is capable of modeling the energy and water demand for data center cooling systems over the course of a year at hourly increments so that daily, monthly, and annual demand profiles can be evaluated (Integrated Environmental Solutions, 2026).

The major modeling assumptions used in this study are detailed in the remainder of this section. The three major considerations to evaluate water and power demand for a data center are the external environment (climate), internal environment (ITE design), and cooling plant design.

External Environment (Climate)

The primary weather factors for evaluating evaporative cooling are the outdoor dry-bulb and wet-bulb temperatures. The outdoor dry-bulb temperature is the ambient air temperature measured by a thermometer with a dry sensing bulb (i.e., the "temperature outside”). For engineering purposes, this is the lowest possible temperature at which water can be cooled using ambient air without evaporation. The outdoor wet-bulb temperature is the ambient air temperature measured by a thermometer with a wetted wick covering the sensing bulb. It is an indirect indicator of outdoor air humidity and for engineering purposes represents the lowest possible temperature at which water can be cooled through evaporation. The outdoor air wet-bulb temperature is always lower than the outdoor air dry-bulb temperature, which is why evaporative cooling can be used as supplemental cooling during peak summer conditions.

Houston and El Paso have been selected as the locations for study. These two sites bookend typical Texas climate conditions ranging from hot and humid (Houston) to warm and dry (El Paso). For modeling, hourly weather data was obtained through Typical Meteorological Year (TMY) weather files created from historical weather observations from 2011–2025 at World Meteorological Organization (WMO) weather stations (Lawrie & Crawley, 2026). The TMY data for each site is summarized on psychrometric charts in Figures 1 and 2. A psychrometric chart plots dry-bulb temperature on the horizontal axis and the mass of water vapor in the air (per unit of dry air mass) on the vertical axis. The curved left-hand limit of the psychrometric chart represents air that is fully saturated with water vapor (i.e., air at 100% relative humidity). The region enclosed by the superimposed blue line represents a typical operating environment for the data halls within a data center.

In Figure 1, the TMY data shows many climate hours in Houston close to the left-hand limit of the chart, indicating a high humidity environment. El Paso’s annual outdoor psychrometric conditions in Figure 2 show a drier climate with the majority of hours further away from the maximum saturation line and towards the bottom of the chart. Even though El Paso shows higher peak temperatures with some hours further to the right on the chart, the majority of the hours are to the bottom left of the chart indicating more mild temperature and humidity conditions.

Figure 1 - Psychrometric Chart with Houston TMY Weather Data

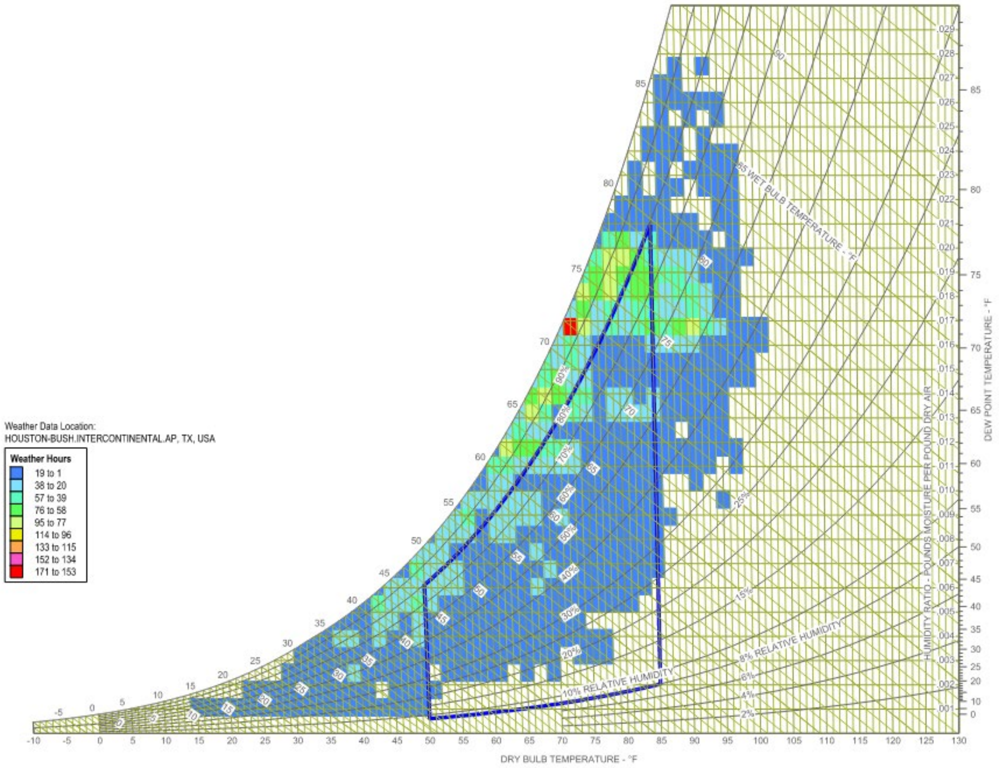

Figure 2 - Psychrometric Chart with El Paso TMY Weather Data

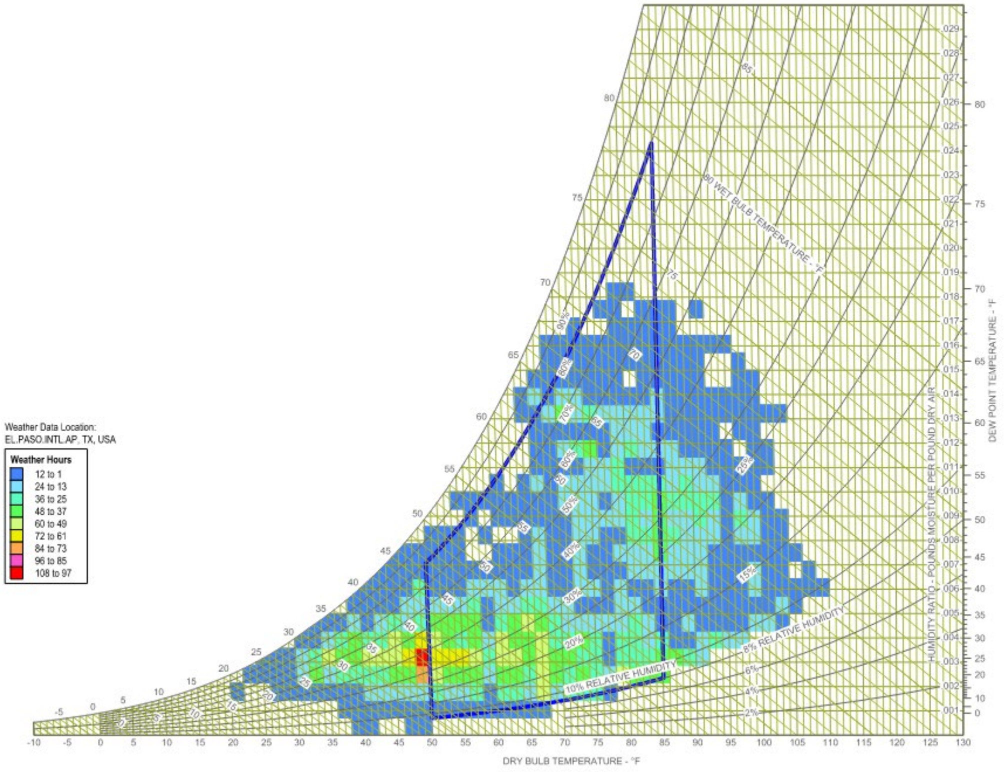

Internal Environment (ITE Design)

Traditionally, data center servers have been cooled using air. The required air conditions entering the servers significantly impact the water and energy performance of a data center by limiting when economization, or free cooling, can be used for cooling. An example of economization includes utilizing 100% outdoor air to cool the data center when ambient temperature and humidity conditions permit. The enclosed region on the psychrometric charts in Figures 1 and 2 represent the indoor environmental conditions used for air-cooled ITE in this analysis. This operating environment used in the mode can be defined using the following ranges:

- Entering server air dry-bulb temperature: 50–85°F
- Entering server air relative humidity: 8–80%
- Air temperature rise through the servers: 20°F

As the power consumption, and heat generation, of servers has increased due to increased chip power and computationally demanding loads (such as AI), many servers require the use of liquid-cooling. Liquid-cooling uses a closed loop of recirculating fluid to cool the server processors before rejecting heat to a secondary cooling loop using a heat exchanger—heat is ultimately rejected to the atmosphere using a variety of cooling technologies. The most common form of liquid-cooling provides the cooling fluid directly to the CPUs or GPUs (termed direct-to-chip cooling) leaving the remaining ITE to be cooled by air. Direct-to-chip cooling is assumed in this study, as other liquid-cooling technologies such as immersion cooling are not currently being deployed at a large scale. The following assumptions are used for liquid-cooled ITE in this analysis:

- Heat removal: 80% by liquid-cooling, 20% by air-cooling
- Entering server liquid temperature: 86°F
- Water temperature rise through servers: 18°F
- Entering server air conditions are the same as listed above for air-cooled ITE

Cooling Plant Design

This paper focuses on the common cooling systems used for hyperscale data centers. Based on reviews of announced hyperscale data center designs and a 2024 study by Lawrence Berkeley National Laboratory (Shehabi et al., 2024), the systems listed in Table 2 were selected for study.

Table 1. Summary of Cooling Systems

System	Water Use	Description
ASE + DEC	Low	Airside economizer (ASE) with direct evaporative cooling (DEC)
WCC + DFC	None	Water-cooled chillers (WCC) with dry fluid coolers (DFC)
WCC + HFC	Medium	Water-cooled chillers (WCC) with hybrid fluid coolers (HFC)
WCC + OCT	High	Water-cooled chillers (WCC) with open cooling towers (OCT)

The airside economizer (ASE) system shown diagrammatically in Figure 3 utilizes outdoor air as the first stage of cooling. During cold outdoor air conditions, an economizer damper modulates to mix outdoor air and recirculated hot-aisle air to maintain server inlet, or cold-aisle, conditions. When the outdoor air temperature is the same as the desired cold-aisle temperature, the system will be providing 100% outdoor air to the data center. As outdoor air temperature increases above the desired cold-aisle temperature, direct evaporative cooling (DEC) in the computer room air handlers will be enabled to cool the air without the use of compressors. Evaporative cooling can either be provided using wetted media or misters spraying water directly into the air path. Supplemental cooling from a compressor-based mechanical cooling system may be required for peak hours.

While the airside economizer system in Figure 3 uses air as the primary cooling means, its use is not limited to air-cooled IT equipment. Recent new construction hyperscale data centers have utilized this design with air-to-water heat exchangers to create cool water for liquid-cooling from the cold-aisle air. Heat from the liquid-cooled ITE is rejected into the hot aisle and removed from the building as normal. This system is also being used to retrofit existing data centers with liquid-cooled ITE.

Figure 3. Example Diagram of Airside Economizer Cooling System

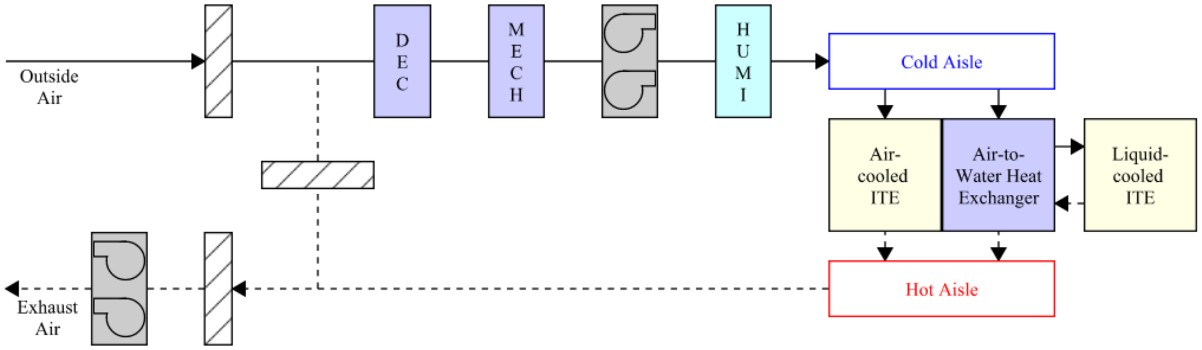

Figure 4 shows an example of a generic closed-loop hydronic cooling system that can provide cooling for both air-cooled and liquid-cooled ITE. In the diagram, the facility water system (FWS) provides cooled water to computer room air handlers (CRAH) for air-cooling as well as coolant distribution units (CDU) for liquid-cooling; it is common to have a single FWS loop serving both air- and liquid-cooled ITE loads. During peak outdoor air temperatures, the chiller removes heat from the FWS loop and rejects it to the condenser water (CWS) loop where it is ultimately rejected to the outdoors through the heat rejection equipment. The heat rejection equipment can either be a dry fluid cooler, a hybrid fluid cooler, or an open cooling tower. When climate conditions permit and the heat rejection equipment can produce cold enough water, the chiller is bypassed to save energy and the FWS loop is cooled by the heat exchanger alone. For this study, the FWS supply temperature setpoint is 75.2°F allowing the hydronic system to cool both the modeled air- and liquid-cooled ITE.

Figure 4. Example Diagram of Waterside Economizer Cooling System

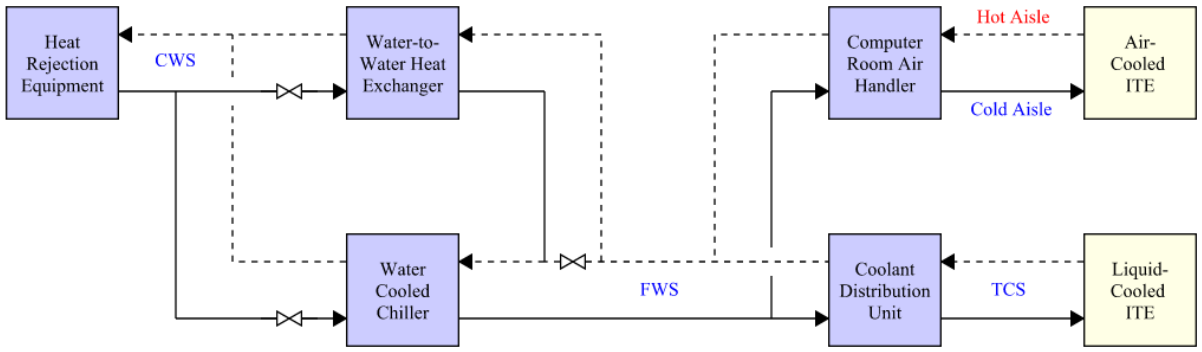

The following options for heat rejection equipment are considered in this study:

- Dry Fluid Coolers (DFC) are large finned tube metal cooling coils that reject heat from the circulating fluid inside the coils to the ambient outdoor air. Variable speed fans blow air over the coils to maintain the desired leaving water temperature. Importantly, the circulating fluid within the dry cooler is not exposed to ambient outdoor air.
- Hybrid Fluid Coolers (HFC) are similar to dry coolers except that they can operate either in dry mode or evaporative mode. Water consumption in evaporative mode can be in the form of spray water applied directly to the surface of the cooling coil, or in the form of adiabatic assist, where the entering air is pre-cooled by evaporation prior to passing over the cooling coil. In both cases, the condenser water is not directly exposed to ambient air as a secondary water stream supplies the water to be evaporated. The use of evaporation increases the efficiency of the heat transfer allowing the hybrid cooler to remove more heat than a dry cooler at warmer ambient temperatures.
- Unlike fluid coolers, which keep the condenser water separated from the ambient air through a cooling coil, an open cooling tower (OCT) exposes the condenser water directly to the ambient air allowing it to undergo evaporation. The water that remains a liquid is cooled by the evaporation process and gets recirculated through the system. The benefit of cooling towers is that this direct form of evaporative cooling is more efficient than the indirect method used in fluid coolers, which allows them to create colder water.

Water consumption from the evaporative cooling process includes water used to directly makeup for the water lost through evaporation, but also water required to make up for blowdown. Blowdown is the periodic discharge of the fluid used for evaporative cooling to prevent excessive buildup of dissolved solids, minerals, or other impurities in the water. Blowdown is represented by the cycles of concentration, which is the ratio of total makeup water to blowdown. For this study, 3.0 cycles of concentration is used to model blowdown for all evaporative cooling equipment.

ANALYSIS AND RESULTS

For this study, the total data center ITE load is assumed to be one (1) megawatt (MW) operating at full capacity for each hour of the year (100% utilization). This was done so that all results presented in this section can be scaled, through post-processing, to any desired data center size or utilization rate. In addition to the ITE load, it is assumed that electrical losses equal to 7% of the ITE load will occur. These losses primary occur at the power conversions at the uninterruptible power supply (UPS) and the various voltage transformers in the electrical distribution system; these losses show up as heat the data center cooling system must remove from the building. Seven percent was selected to represent a typical hyperscale electrical distribution design.

The results for the analysis are presented in the following pairs of figures with the cooling system results for each location shown side-by-side for direct comparison:

- Figures 5 and 6 show the monthly average WUE and PUE
- Figures 7 and 8 show the average annual daily profiles for WUE and PUE
- Figures 9 and 10 show the average WUE and PUE by outdoor dry-bulb temperature

Figure 5 - Monthly WUE Results - Liters per kWh

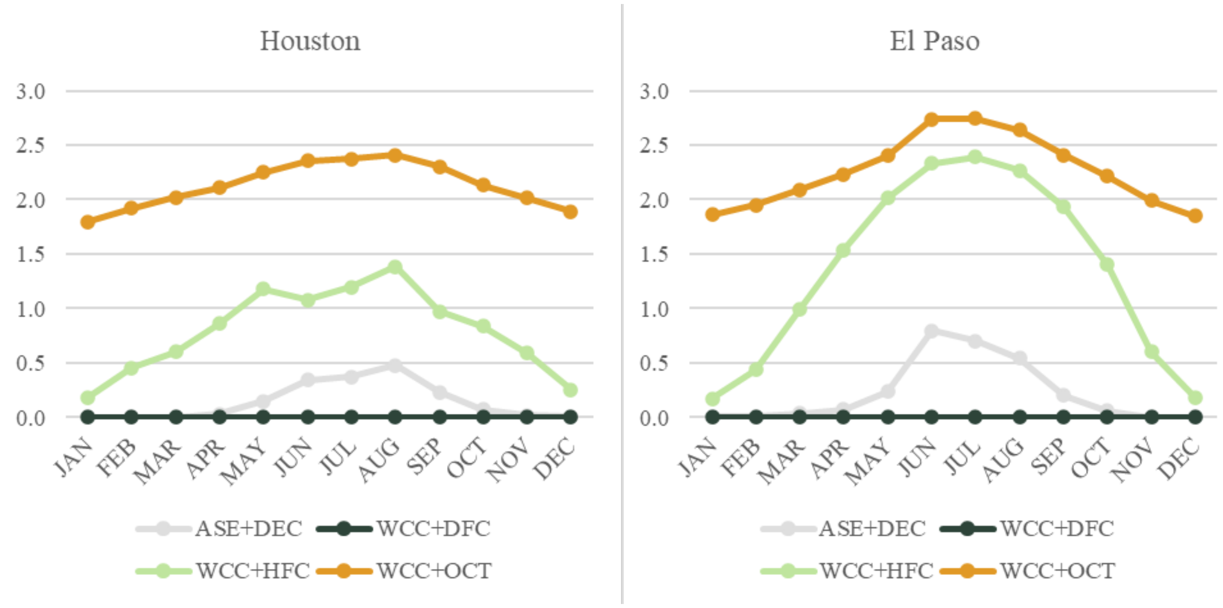

Figure 3 - Monthly PUE Results - kWh per kWh

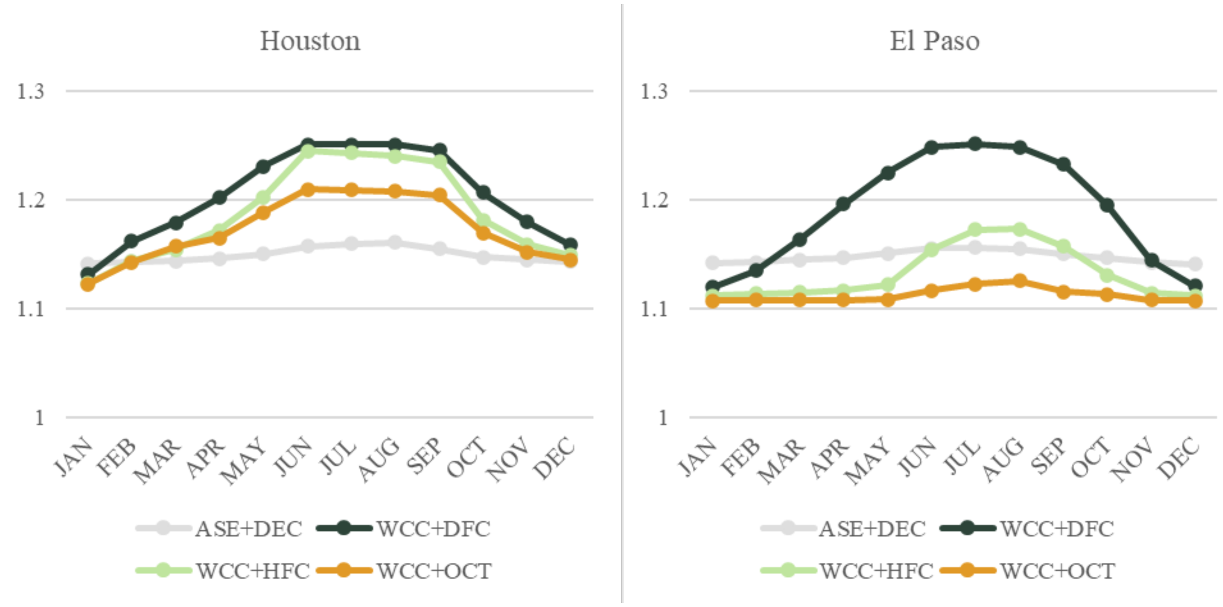

Figure 4 – Average 24-hour Profile WUE Results – Liters per kWh

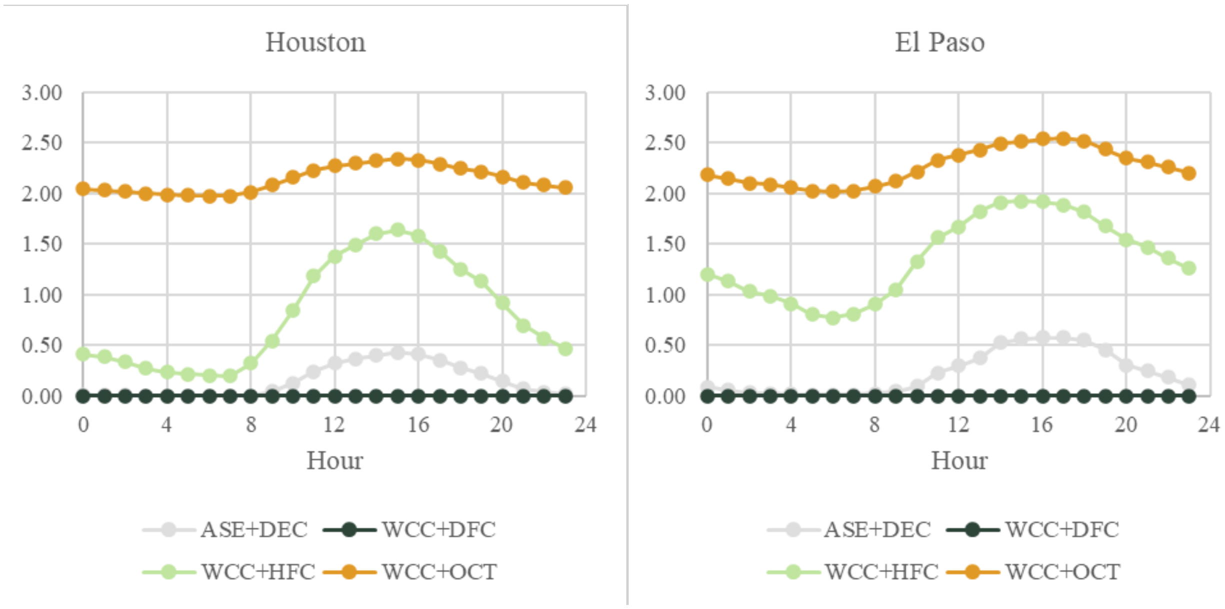

Figure 5 - Average 24-hour Profile PUE Results - kWh per kWh

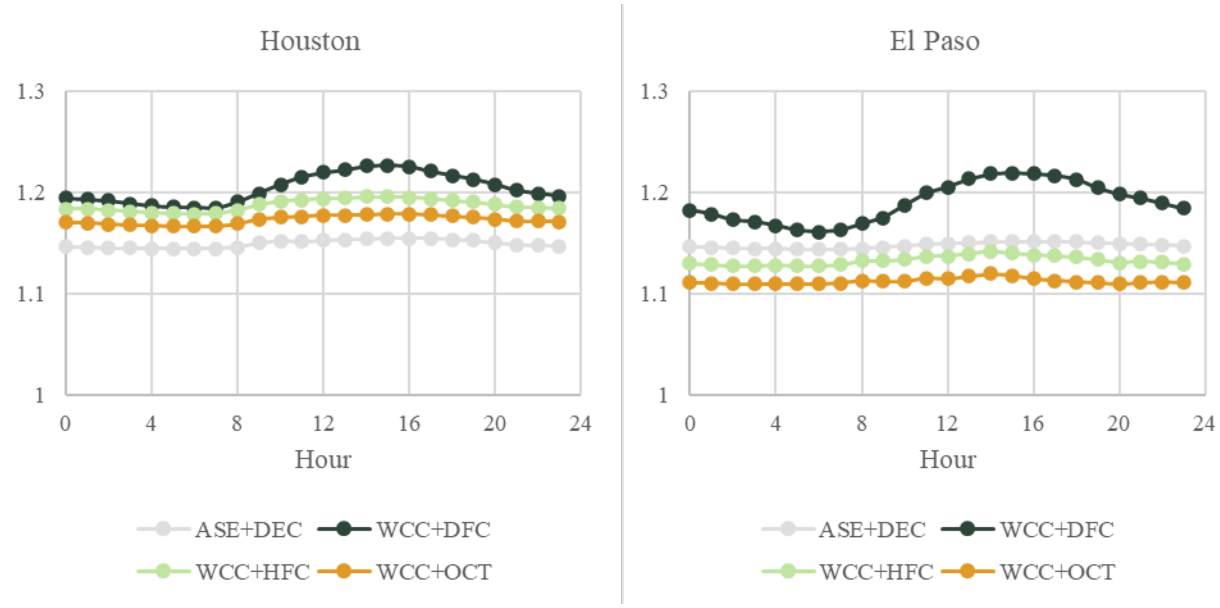

Figure 6 - WUE Results by Outside Air Dry-Bulb Temperature - Liters per kWh

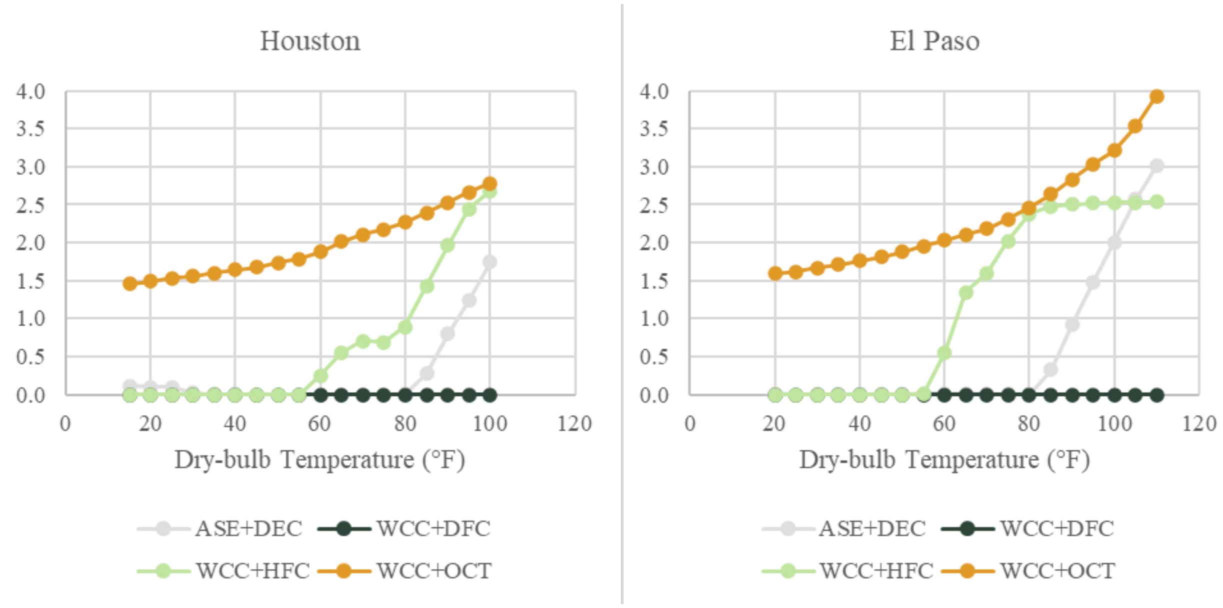

Figure 7 - PUE Results by Outside Air Dry-Bulb Temperature - kWh per kWh

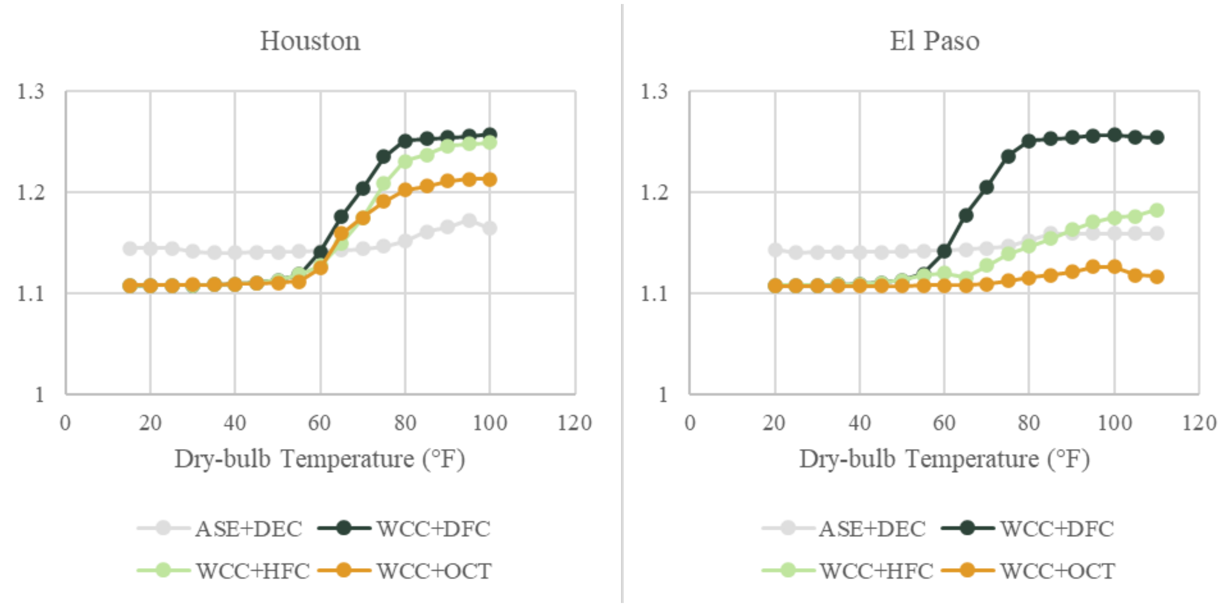

DISCUSSION

For large data centers, a resilient supply strategy will combine development of new supplies with storage and reuse to manage variable seasonal demand. A portfolio approach should be implemented that includes municipal reuse, brackish groundwater, and produced water desalination combined with storage-oriented strategies such as aquifer storage and recovery and off-channel reservoir projects. Data centers and other water users with growing demand should consider large scale storage to smooth peak seasonal demand and protect against drought conditions to improve water system reliability. Strategic siting and coordination with regional planning groups and local stakeholders can help ensure that new infrastructure is developed in a way that is both sustainable and reliable. Texas' biggest long-term pressure is on municipal demand, so siting data centers where they can be served by reclaimed water, brackish groundwater, or other non-traditional supplies can reduce competition with municipal growth while strengthening overall system resilience. If planned early and integrated into regional water planning, new infrastructure and interconnections can be designed to meet data-center needs without undermining existing municipal customers, and in some cases can help justify infrastructure that also improves reliability for growing communities.

REFERENCES

Batson, A. (2026, February 17). North America Data Center Report Year-end 2025. JLL. https://www.jll.com/en-us/insights/market-dynamics/north-america-data-centers

Integrated Environmental Solutions. (2026). IES Virtual Environment (Version 2025.2.0.0). Integrated Environmental Solutions. https://www.iesve.com/software/virtual-environment.

Lawrie, L. K, & Crawley, D, B. (2026). Development of Global Typical Meteorological Years (TMYx). https://climate.onebuilding.org.

Shehabi, A., Smith, S.J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M.A.B., Holecek, B., Koomey, J., Masanet, E., Sartor, D. (2024). 2024 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-2001637

Synergy Research Group. (2026, April 13). U.S. Hyperscale Investment Shifts Decisively Inland. https://www.srgresearch.com/articles/focus-of-us-hyperscale-investment-shifts-dramatically-inland

Texas Water Development Board. (2026). Draft 2027 State Water Plan. https://www.twdb.texas.gov/waterplanning/swp/2027/docs/DraftSWP27-Water-For-Texas.pdf

The Green Grid. (2011). Water Usage Effectiveness (WUE): A Green Grid Data Center Sustainability Metric. The Green Grid.

The Green Grid. (2012). PUE: A Comprehensive Examination of the Metric. The Green Grid.