Water Usage at AI Scale: Insights from a 100 MW Comparative Analysis

White Paper 220

Version 1

Data Center Research & Strategy

by Fabian Janischowsky
Amir Ibraheem, PE
Maria A. Torres Arango

Executive summary

Water consumption has become one of the most pronounced community concerns surrounding the rapid expansion of large-scale AI deployments. Yet the amount of water a data center consumes is *not* determined by its size or workload alone. Cooling technology and design play a significant role in shaping this outcome. Therefore, water use for heat rejection should be viewed as a design choice rather than an inherent requirement. Adiabatic-assist can reduce cooling energy consumption and lower PUE, creating a water-energy tradeoff. However, once the technology cooling system (TCS) supply temperature is optimized, the incremental energy benefit can diminish. We compare four 100 MW data center designs ranging from a traditional air-cooled architecture to optimized liquid-cooled AI architectures, using representative climates and operating strategies to show how cooling technologies influence water efficiency at scale.

RATE THIS PAPER ★★★★★

Key takeaways

1. **Cooling strategy is central to AI infrastructure planning.** The cooling architecture directly impacts a data center's water and energy footprint. It should therefore be evaluated early in the design process based on water and energy availability, cost, operational requirements, regulations and sustainability goals.

2. **Data center water use is determined by how heat is rejected *outdoors*; and can be greatly reduced by using liquid cooling *indoors*.** Outdoor heat rejection systems may use water to improve heat exchange efficiency. Indoor liquid cooling can reduce the proportion of mechanical cooling by lowering chilled-water demand and increase the operating hours in "free-cooling" mode. Together, these effects can significantly reduce water and energy consumption.

3. **Moving from air cooling to liquid cooling consistently improves energy efficiency and can reduce water use in water-constrained locations.** In the 100 MW simulations, moving from traditional air cooling to liquid cooling at a 32°C TCS supply temperature reduced on-site water use by ~8% in Dallas, Texas. It also reduced PUE by ~7% (cooling overhead by ~63%) in both Dallas and Paris, France; demonstrating that the energy-benefits apply across different climates.

4. **Higher liquid cooling temperatures provide substantial additional water savings.** Raising the TCS supply temperature from 32°C to 45°C extends the range of outdoor temperatures in which free cooling can operate. This reduced WUE by ~66% to ~82% depending on location; and PUE by ~0.6% (cooling overhead by ~13%). Alternatively, adjusting equipment count to prevent stranded cooling capacity and reduce CapEx, decreased WUE by ~50% and PUE by ~0.5% (cooling overhead by ~11%) compared to the 32°C baseline.

5. **On-site cooling water use is a design and operating choice, not an unavoidable result of climate.** Our simulations show that implementing liquid cooling at high TCS temperatures maximizes cooling efficiency across locations, to the point where using water to further lower PUE produces negligible improvements. While climate affects cooling performance, it does not determine water use on its own. Heat rejection technology, liquid cooling, TCS temperature, and equipment choices can substantially mitigate water use in both cooler and warmer locations.

Introduction

AI is driving larger and denser data center deployments, which has increased attention on facility resource requirements, including water use. As new facilities are proposed, communities are seeking information about the resource implications of AI infrastructure. Executive Report 16, *How Sustainable Technology, Design, and Operations can Unlock Community Support for Data Centers*, explores this dynamic and how design choices can inform public dialogue.

A common perception is that AI growth drives higher water consumption. This assumption often reflects experience with legacy cooling systems rather than current design practices. Water consumption depends on how a facility is designed and operated. AI factories are engineered to maximize infrastructure efficiency and resource utilization. Accordingly, water use is treated as a critical design variable rather than a fixed consequence of AI growth.

White Paper 211, *10 Ways to Harness the Energy and Water Efficiencies of Direct Liquid Cooling*, outlines how liquid cooling can reduce energy and water use in data centers, along with strategies for achieving those reductions. A starting point for evaluating these opportunities is understanding the different ways water is used within a data center.

How is water used in data centers?

Water use in data centers depends on outdoor heat rejection technology and climatic conditions.[1] Figure 1 contrasts on-site water use for data centers that use cooling towers, air-cooled chillers with adiabatic assist, and waterless dry coolers (no adiabatic).[2] In facilities that use water for heat rejection, this category dominates on-site water consumption. Other uses (grouped for simplicity) include humidification, domestic and sanitary services, fire suppression, and landscaping.

Figure 1

Direct water use comparison for representative data centers with different outdoor heat rejection technologies

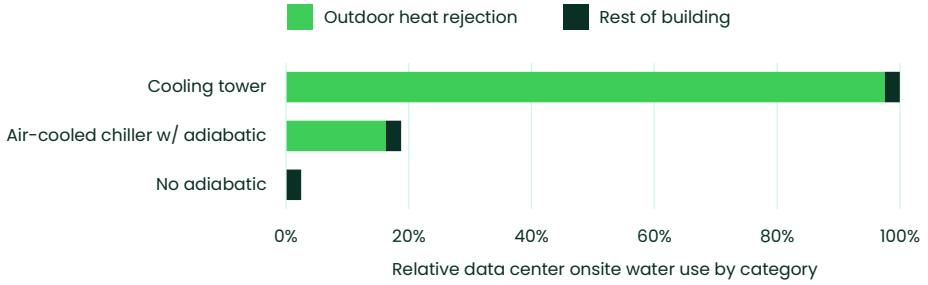

Figure 1 illustrates the variation in on-site water consumption associated with different outdoor heat rejection technologies. Cooling towers can consume 5 to 20 times the amount of water used by dry coolers with adiabatic assist for an equivalent size data center. Water usage effectiveness (WUE), expressed in liters of water consumed per kilowatt-hour of IT energy, provides a standardized basis for comparing cooling related water consumption. Reported WUE values range from approximately 1.7 – 4 L/kWh for data centers with cooling towers; and from approximately 0.2 to 1.5 L/kWh for those using dry coolers with adiabatic assist.[3,4,5] These ranges

[1] Eslami, *et al.*, Optimizing Data Center Cooling Across Climates: Modeling of PUE and WUE, ASHRAE Transactions, 2026

[2] WUE values for this comparison are 0 (zero), 0.5 and 3, for no adiabatic, air-cooled chiller with adiabatic assist and cooling tower, respectively.

[3] 2024 U.S. Data Center Energy Usage Report, Lawrence Berkeley National Laboratory

[4] Quirk *et al.*, Data Centers and the Water-Energy Nexus, ASHRAE Journal, June 2026.

[5] Akinade et al., AI Data Centers and the Water Use Feedback Loop, arXiv:2606.21760, June 2026

indicate that the selection and operation of the outdoor heat rejection system can materially influence a data center's on-site water footprint.

On-site vs. off-site water use

The total water footprint of a data center includes direct and indirect consumption, as shown in Figure 2. *Direct water consumption* refers to water used at the data center facility, primarily for heat rejection. *Indirect water consumption* refers to water used to generate the electricity consumed by the data center. It varies with the local power generation mix. Sourcing electricity from technologies with low energy water intensities, such as wind, solar (photovoltaic), or natural gas,[6] can reduce indirect water consumption.

Figure 2

Components of total data center water footprint

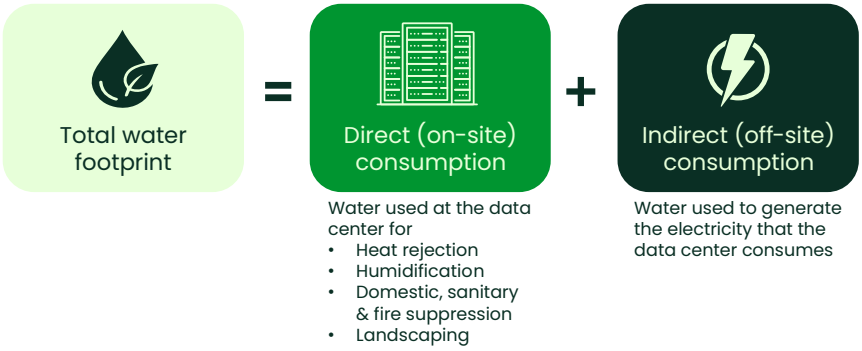

In this paper we examine strategies to reduce direct (on-site) water consumption associated with data center cooling[7]. The analysis uses simulations of representative 100 MW data centers. We present the methodology, simulated scenarios and assumptions. Then, we compare the results, to contrast water use by the choice of outdoor heat rejection and white-space cooling technologies.

A 100 MW comparative analysis

To understand the implications of cooling strategy on water use, we simulated four scenarios, representative of traditional data centers and AI factories. We used Schneider Electric Reference Design 113, *10.2-12.7 MW, Tier III, ANSI, Chilled Water, Liquid-cooled AI clusters (NVIDIA Vera Rubin)*[8], as baseline for the scaled 100 MW AI factory designs. Note, the contribution from the electrical distribution system and other non-cooling systems to power usage effectiveness (PUE) is assumed to be the same across all scenarios investigated.

Methodology

We calculated water and energy consumption for one year of data center operation using proprietary software. We compared the four data center cooling scenarios summarized in Figure 3.

- Scenario 1: Air-cooling – used for traditional data centers.
- Scenario 2: Liquid-cooling, TCS at 32°C – Technology cooling system (TCS) supply temperature is 32°C, representative of early generation AI factories.

[6] Azevedo *et al.*, WP#35 – Water Usage Effectiveness (WUE): A Green Grid Data Center Sustainability Metric, The Green Grid, March 2011.

[7] Other categories including humidification; domestic, sanitary & fire suppression; and landscaping are outside of the scope of our simulations.

[8] Reference Design 113 contains detailed IT layout, power, and cooling design.

- **Scenario 3: Liquid cooling, TCS at 45°C** – Uses the same cooling infrastructure (equipment type and count) as Scenario 2, but raises the TCS supply setpoint to 45°C.
- **Scenario 4: Liquid cooling, TCS at 45°C, reduced equipment count** – Cooling equipment is sized specifically for operation at a 45°C TCS supply temperature, avoiding stranded cooling capacity and reducing CapEx.

Figure 3

Overview of simulated scenarios varying indoor (air vs liquid cooling) and outdoor (no adiabatic vs. adiabatic) heat rejection strategies.



Except for Scenario 3, which uses the same equipment count as Scenario 2 while increasing the TCS supply temperature from 32°C to 45°C, the equipment and setpoints for air and liquid cooling architectures were optimized for minimum energy consumption with N+1 redundancy. Additional details for each scenario are provided in Appendix, Table A1.

Figure 4 shows the representative cooling architectures and how outdoor and indoor technologies are connected.

Figure 4

Representative air and liquid cooling architectures depicting outdoor and indoor cooling technologies used for our traditional data center and AI factory simulations. Water droplets next to outdoor equipment indicate adiabatic assist option.

Note the fanwall provides air cooling for the full white space. Heat sources include the air-cooled racks, the air portion of the liquid-cooled IT racks (4% for our simulations), and other non-IT-related heat sources (e.g. from UPS heat radiation, light bulbs, etc.)

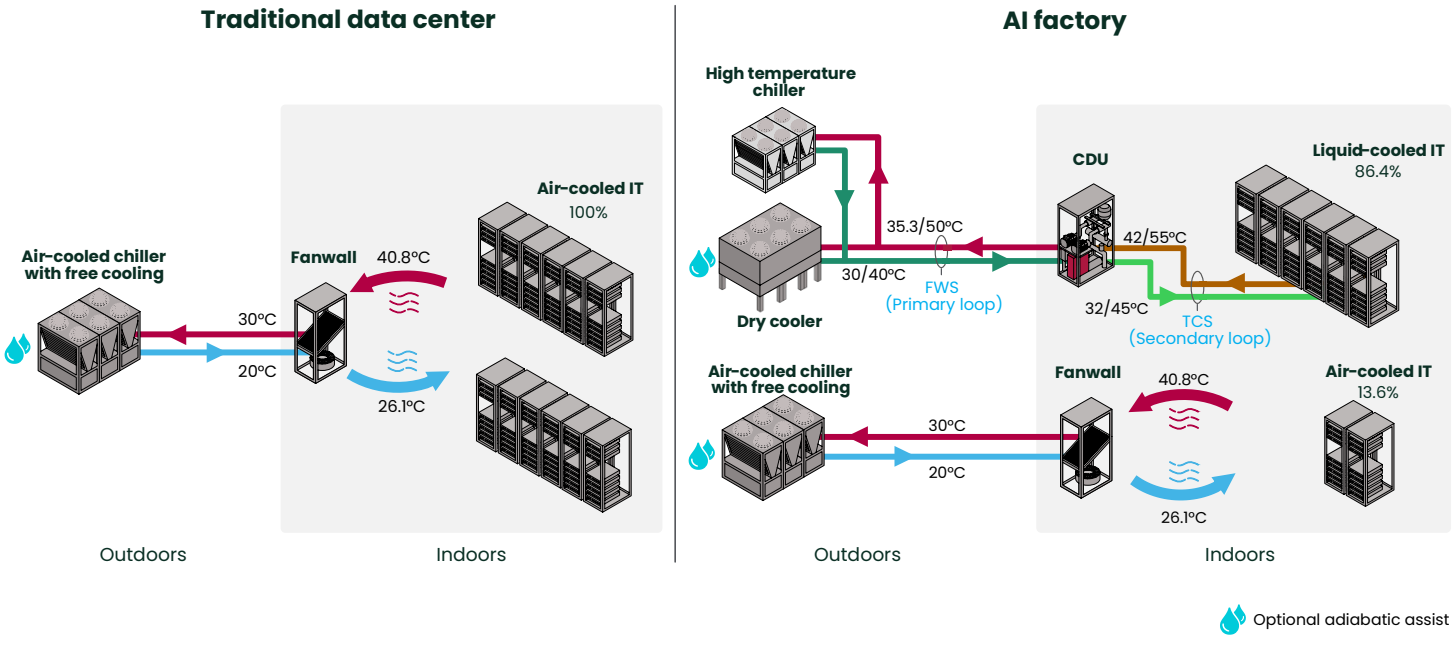

Each scenario was evaluated with and without adiabatic heat rejection to quantify the tradeoff between water consumption and energy efficiency. For the adiabatic cases, the outdoor heat rejection system is equipped with adiabatic assist[9], which uses evaporating water to improve heat rejection performance during periods of high outdoor temperature. Adiabatic assist is automatically triggered when the cooling capacity cannot be met within the same cooling mode. This is, when the system would need to switch between free-cooling to partial free-cooling to full mechanical load. Adiabatic mode helps delay each of these transitions. This is not a fixed setpoint, but changes depending on the environmental conditions and heat rejection operating setpoints.

Our analysis also assessed the impact of location, contrasting the climatic conditions for Paris, France; and Dallas, Texas; based on their annual weather profiles. Temperatures in Paris range from -7°C (90% relative humidity) to 29 °C (42% relative humidity), while temperatures in Dallas range from -9°C (70% relative humidity) to 39°C (27% relative humidity).

Scale implications

As discussed in the introduction, the on-site water consumption for heat rejection is a design choice, driven primarily by the improved heat exchange efficiency of evaporative cooling to reduce energy consumption[10]. At the 100 MW scale, the tradeoffs between water and energy consumption can be significant enough to influence design decisions based on local constraints, operational requirements, and efficiency targets. Figure 5 summarizes on-site water consumption for our *adiabatic* scenarios.

Figure 5

On-site water consumption from cooling IT workloads for the simulated scenarios with varying heat rejection technologies

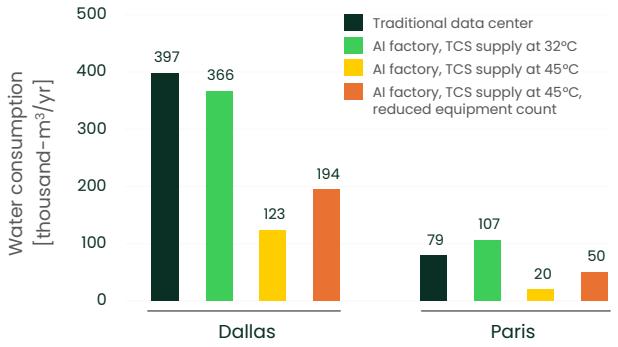

Three key observations emerge from these water consumption results:

- **The highest onsite water consumption occurs in the traditional air-cooled data center scenario in Dallas.** At the 100 MW scale, this annual water consumption is comparable to the water used by ~940 Dallas households[11].

- **Shifting from air cooling to liquid cooling with optimized TCS temperatures (45°C supply) reduces water consumption by at least 50%.** This architecture is characteristic of AI factories, which leverage the higher operating temperatures enabled by liquid cooling. In the context of household water

[9] Adiabatic assist can increase the cooling capacity depending on (1) local environmental conditions: temperature, humidity, altitude, outdoor temperature; and (2) heat rejection operating setpoints: facility water system (FWS) coolant flow rate, type of coolant, supply and return temperatures, equipment load, type of equipment, etc.

[10] When water evaporates (transition from liquid to water vapor), it absorbs ~240 times more heat than air in contact with cooling coils across a 10°C temperature difference can exchange.

[11] This assumes the average Dallas household annual water use is ~422 m³. Dallas Water Utilities, City of Dallas.

consumption above, this represents a positive impact in resource allocation for water constrained locations.

- **Water reduction could be greater by deploying more liquid cooling capacity than is strictly required.** This corresponds to Scenario 3, the AI factory operating with TCS supply at 45°C, relying on the same equipment count as the TCS at 32°C scenario. As data centers become more familiar with liquid cooling and its control schemes, they will be more comfortable raising the TCS temperatures. This scenario provides the versatility to tune the TCS supply setpoint as needed to gain such operations maturity, i.e. to shift to higher TCS temperatures at the risk pace they can manage; or to accommodate additional cooling capacity requirements.

In the next section we discuss the path to reducing water consumption, and how these reductions in water use relate to changes in PUE.

Incremental gains: from air to liquid cooling to optimized liquid cooling

WUE and PUE are interrelated metrics that help assess the water use and operational efficiency of data centers. Our results indicate that **AI factories (using liquid cooling) use less water and energy compared to traditional data centers.** We observe two trends:

1. Moving from traditional (100% air-cooled) data centers to AI factories that implement liquid cooling improves both WUE and PUE.
2. Increasing the TCS IT supply temperature significantly reduces WUE; and modestly improves PUE.

Shifting from air to liquid cooling (Scenario 1 to Scenario 2)

Traditional data centers differ from AI factories in two main ways: (1) they use different IT equipment and (2) they rely on air cooling. It is liquid cooling what enables higher rack density for AI factories. In our simulations, traditional data centers have rack densities of 20 and 40 kW. In contrast, AI factory scenarios relying on liquid cooling have rack densities[12] of up to 188 kW. Supporting these higher-density workloads requires a different cooling architecture.

In practice, transitioning from air to liquid cooling at the 100 MW scale implies the deployment of two separate (dedicated) cooling architectures in the data center: one for liquid-cooled loads (this is the main load), and another for ancillary air-cooled loads as noted in **Figure 4**. The split depends primarily on the IT heat capture ratio and supporting air-cooled systems in the white space. In our AI factory scenarios, the split is 86.6% liquid cooling and 13.4% air cooling. Having two independent cooling loops allows their individual optimization to minimize water and energy consumption. White Paper 133, *Navigating Liquid Cooling Architectures for Data Centers with AI Workloads*, discusses the advantages and requirements for deploying dedicated vs. shared liquid cooling architectures.

Shifting from air to liquid cooling at TCS supply temperature of 32°C, allows the outdoor heat rejection equipment to work with more efficient setpoints, as the temperature requested from these dedicated chilled-water systems is different: 20°C chilled-water supply for air cooled loads and 30°C for the liquid cooled loads[13].

[12] Note, AI factory scenarios also have rack densities of 20 and 40 kW, corresponding to networking air-cooled loads.

[13] This favors a more frequent use of free cooling over mechanical cooling, hence reducing the energy needed. We explain economizer modes in detail in White Paper 132, *Economizer Modes of Data Center Cooling Systems*. The section dedicated to packaged chiller bypass via dry (or evaporative) cooler

Figure 6 shows that the largest benefit is the significant reduction in PUE, by ~5.7% (~62% of cooling overhead) for Paris, to ~7.8% (~64% of cooling overhead) for Dallas.[14] Similarly, shifting from air to liquid cooling at 32°C reduces on-site water for heat rejection by ~7.8% for Dallas.

Figure 6

Comparison of WUE and PUE for traditional air-cooled data center and AI factory with TCS supply temperature at 32°C scenarios

Note, WUE results only account for heat rejection category (rest of building categories are excluded).

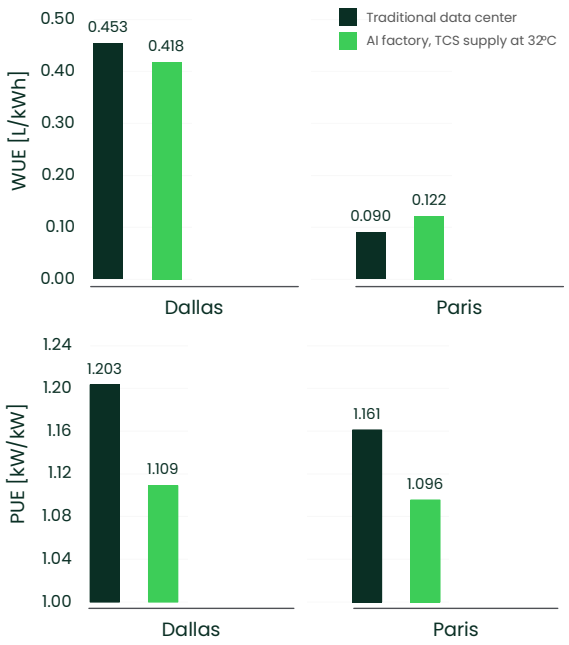

Note there is a slight increase in the water consumption for the Paris scenario with TCS supply temperature at 32°C. This increase results from the combination of (1) the outdoor equipment selected for the liquid cooling loads for the AI factory[15]; and (2) their typical controls scheme triggering adiabatic assist vs high-temperature chillers assistance as temperatures fluctuate during operations. Tuning this control scheme to trigger assistance in a different order or at different temperature set-points results in a different water footprint and should be included in the thermal design and controls commissioning process of the data center.

Later in this section, we examine the tradeoffs of using water to reduce PUE including insights from additional scenarios.

Increasing TCS supply temperature (Scenario 2 to Scenario 3)

Raising the TCS supply temperature from 32°C to 45°C (Scenario 3) significantly reduces WUE by ~66% for Dallas to ~82% for Paris. Concomitantly, PUE is decreased by ~0.9% (reduced cooling overhead by 20%) for Dallas; and by ~0.2% (reduced cooling overhead by ~6%) for Paris, as shown in Figure 7.

While Figure 7 shows that increasing the TCS supply temperature is a powerful lever for reducing water and energy consumption, it is not a trivial choice; but an operational decision that needs to consider multiple requirements. The allowable temperature range is ultimately constrained by the IT equipment manufacturer, and

(page 7), is the mode relevant to this simulation; where the built-in mode automatically switches between mechanical, free cooling, and mixed mode as necessary.

[14] PUE reduction for the adiabatic assisted heat rejection scenarios. A similar PUE reduction of ~6.3% (reduction in cooling overhead by ~57%) is achieved by shifting from air to liquid cooling at 32°C, for no adiabatic heat rejection systems (air-cooled chillers).

[15] Heat rejection equipment (type and count) for AI factory scenarios is different from the equipment for traditional data centers. This implies their energy consumption and controlling schemes are also different.

can affect IT performance, reliability, lifespan, and even warranty. Other considerations include the controls strategy, thermal stability, and available cooling capacity. As data center operators gain experience with liquid cooling and develop confidence in the behavior of their cooling system and controls response, higher TCS temperatures are likely to become more widely adopted.



Figure 7

WUE and PUE summary including results from Scenario 3, TCS supply temperature at 45°C

Note, WUE results only account for heat rejection category (rest of building categories are excluded)

Reducing stranded cooling capacity (Scenario 3 to Scenario 4)

Increasing TCS temperature also reduces the amount of equipment necessary to meet the required cooling capacity. Scenario 3, discussed above, achieves the greatest reductions in water and energy consumption. However, it results in ~33% stranded cooling capacity since the higher TCS temperature enabled more frequent use of free cooling. Scenario 4 adjusts the type and count of equipment to maintain N+1 redundancy, while optimizing for minimum energy consumption. Figure 8 includes this fourth scenario and shows that revising the equipment count can still reduce on-site water and energy compared to the TCS 32°C scenario (Scenario 2).

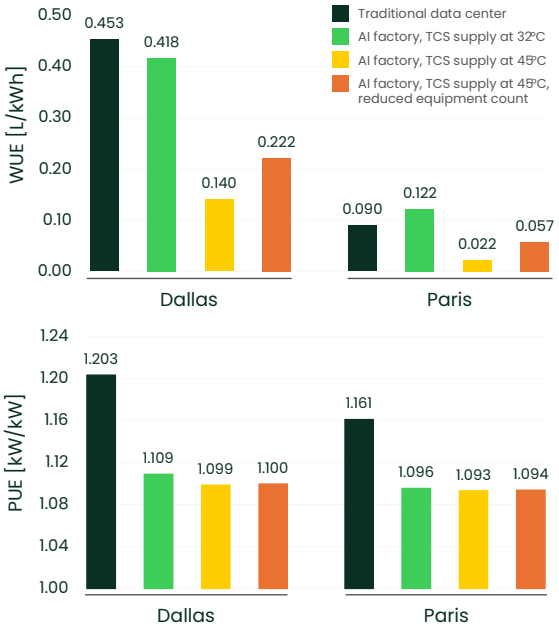

Figure 8

WUE and PUE summary including results from Scenario 4, TCS supply temperature at 45°C, reduced equipment count

Note, WUE results only account for heat rejection category (rest of building categories are excluded)

In Scenario 4, WUE is decreased by ~47% and ~53% for Dallas and Paris, respectively.[16] Similarly, PUE is decreased by ~0.8% (reduced cooling overhead by ~18%) for Dallas, and by ~0.2% (reduced cooling overhead by ~4.2%) for Paris.

AI factories designed to host frontier GPUs running at high TCS temperatures can directly be designed to optimize CapEx. Schneider Electric TradeOff Tool 4, *Data center capital cost calculator*, can assist in the early design stages of the process to evaluate the tradeoffs of liquid cooling architectures and technologies.

TCS supply setpoint & outdoor temperature range for free cooling

Comparing the liquid cooling plant power consumption across Scenarios 2, 3, and 4 provides additional insights into the relationship between TCS supply temperature, cooling capacity, free cooling, and outdoor temperature. **Figure 9** illustrates these relationships by showing liquid cooling plant power consumption as a function of outdoor temperature for the simulated AI scenarios. The key observations include:

- **Raising the TCS temperature from 32°C (Scenario 2) to 45°C (Scenario 3 and Scenario 4) increases the outdoor temperature operation range for free cooling.** As shown in **Figure 9**, the liquid cooling plant power consumption remains low (free-cooling range) for most of the outdoor temperatures relevant to the investigated locations; and increases significantly once chillers are required to supplement the dry coolers at specific outdoor temperatures[17]. **Raising the TCS temperature shifts this onset temperature to the right, allowing free cooling to operate under warmer conditions.**

- **A higher equipment count (Scenario 3 vs. Scenario 4) provides more surface area for heat exchange**[18], slightly shifting the onset of chiller assistance to the right, further extending the range for free cooling.

- **Once a system operates predominantly in free-cooling mode, further increases in the TCS supply temperature increase operational flexibility, and modestly reduce energy consumption.** For locations where outdoor temperatures allow free cooling throughout most of the year, increasing the TCS temperature may still deliver modest energy savings through lower fan power in the outdoor heat rejection equipment. A higher TCS temperature also increases opportunities for heat reuse. This higher quality heat might require less temperature elevation from a potential heat pump or may be directly repurposed. Finally, extending the range for free cooling allows "future proofing" of the data center as the thermal design usually considers maximum temperatures over a 20-year period (ASHRAE weather data). With the ongoing climatic change, extending the range of outdoor temperatures for free cooling via TCS temperature setpoint, provides a "safety margin" towards a hotter climate in the future.

[16] Dallas from 0.42 to 0.22 L/kWh; Paris from 0.12 to 0.06 L/kWh; comparing TCS supply at 32°C and TCS supply at 45°C reduced equipment count scenarios.

[17] This onset temperature can be calculated by subtracting the dry cooler approach temperature from the FWS supply temperature. Onset temperature is 25°C for TCS supply at 32°C; 35°C for TCS at 45°C; and at 37°C for the TCS supply at 45°C scenario with the same equipment count as the 32°C TCS supply scenario.

[18] More surface area for heat exchange is equivalent to reducing the approach temperature for the dry coolers.

Figure 9

Power consumption for the liquid cooling plant under different scenarios showing ambient temperature ranges where dry coolers provide required cooling capacity; and onset for assistance from high-temperature chillers (marked by vertical lines)

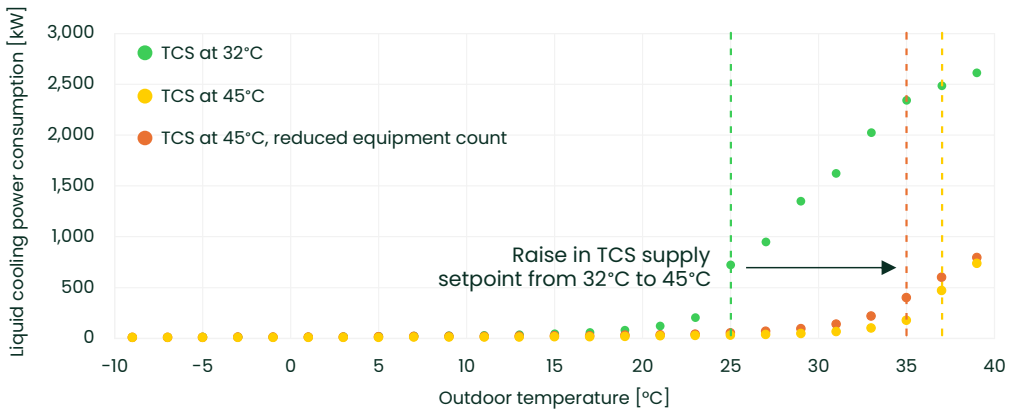

These results agree with earlier observations reported in White Paper 211, discussing ways to optimize liquid cooling design and operation.

Adiabatic vs. no-adiabatic

One of the most notable findings is that high-temperature liquid cooling architectures can operate with little to no cooling-water consumption while maintaining nearly the same PUE. As shown in **Figure 10**, the difference between adiabatic and non-adiabatic operation is negligible for the Paris scenarios and remains very small for the Dallas scenarios operating at higher TCS temperatures.

Figure 10

PUE comparison for adiabatic and no adiabatic scenarios

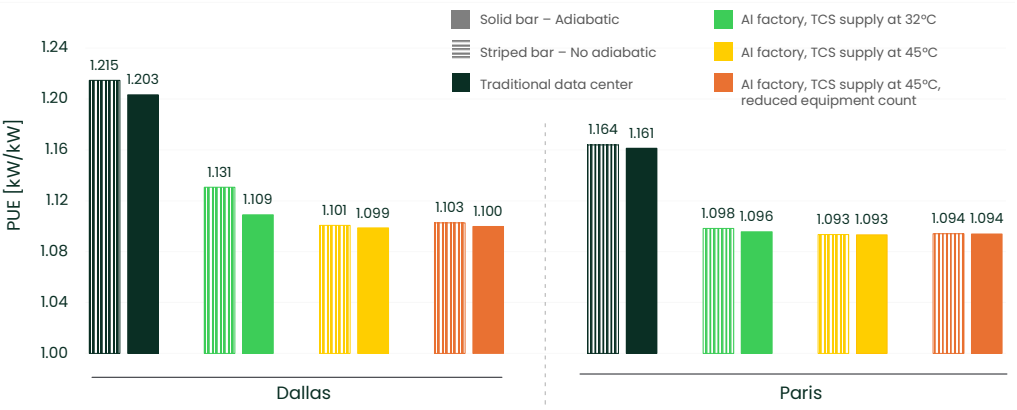

For Paris, implementing adiabatic cooling would not be justifiable from either a technical or an economical perspective, as the gain in PUE would not significantly offset the expense of water. This is evident from a closer examination of PUE and WUE across the Paris scenarios in **Figure 10** and **Figure 8.** These show that moving from a no-adiabatic cooling strategy to one with adiabatic assist only reduces PUE on average by 0.13% (cooling overhead by 2.5%); but increases water use from zero to 0.07 L/kWh.

This observation also holds for the Dallas scenarios operating at high TCS temperatures, where most of the energy savings are realized by extending the free-cooling range rather than through adiabatic cooling.

Across locations, although there is a decrease in PUE for the traditional data center scenarios (air-cooled) when implementing adiabatic cooling (and to some extent in the TCS supply at 32°C Dallas scenario), **the meaningful reduction in PUE comes from adopting liquid cooling**. This again points to the fundamental mechanism in cooling optimization: increasing the temperature difference between the chilled water return and the outdoors; thereby enabling more frequent free-cooling operation.

These results reinforce our earlier observation: **using water through adiabatic assist to improve cooling efficiency is a design choice rather than an operating condition.** Choosing adiabatic cooling or no adiabatic heat rejection should take into account the location, water and energy availability, cost, operational requirements, regulations and sustainability goals.

Next steps

The results presented in this paper show that data center water use is not a fixed consequence of AI growth, but rather the outcome of cooling architecture and operational design choices. Advances in liquid cooling, higher TCS operating temperatures, and optimized heat rejection strategies can significantly reduce both water and energy consumption while enabling the deployment of next-generation AI infrastructure. The following recommendations can help incorporate these considerations into data center planning and operation:

1. **Evaluate water consumption early in the design process**, including opportunities to source water independently so that local communities are not adversely affected by the data center deployment.

2. **Evaluate water and energy performance together.** Perform software simulations of your cooling architecture. Work with your cooling vendor (simulation services / software services) to assess scale and mitigation strategies leveraging cooling technologies.

3. **Balance WUE and PUE objectives when deciding whether to implement adiabatic assist.** The optimal strategy depends on the value placed on incremental PUE improvements relative to water consumption, together with local water and energy costs, resource availability, operational requirements, regulations, and sustainability objectives.

4. **Increase TCS supply temperature to maximize free cooling opportunities.** Evaluate the tradeoffs in compute performance, water use, energy consumption, and capital cost when selecting the IT and cooling infrastructure for your data center. Latest GPUs allow higher liquid cooling temperatures that enable onsite water and energy savings when paired with optimized liquid cooling architectures.

5. **Assess climate-specific opportunities and tradeoffs, considering available resources and regulations.** Location affects both on-site and off-site water use. Heat rejection strategies and liquid cooling can mitigate water use restrictions leveraging dry-cooling approaches. Powering your data center from renewable, low water intensity sources can further reduce its overall water footprint.

6. **Communicate assumptions and performance transparently with stakeholders.** Engage local communities early to identify priorities and actionable mitigation measures. Provide regular progress updates and establish accountability mechanisms to keep water consumption within agreed targets.

About the authors

Fabian Janischowsky is a Simulation Engineer in Schneider Electric's Cooling Innovation team. He analyzes and simulates thermodynamic systems across multiple scales, from individual components to complete system architectures and data center infrastructure. His work focuses on the evaluation and development of next-generation cooling technologies, energy systems and digital infrastructure solutions through advanced modeling and simulation techniques. His expertise includes computational fluid dynamics, heat transfer, thermal management and system-level performance modeling. Fabian holds a Master of Science in Mechanical Engineering from Ostbayerische Technische Hochschule Regensburg.

Amir Ibraheem, PE is a Principal Mechanical Systems Engineer and Cooling Systems Solution Architect in Schneider Electric's Data Center team. He designs and analyzes the thermal architecture of mission-critical facilities, with a focus on scaling high-density AI data centers and developing next-generation Direct-to-Chip (D2C) and liquid cooling solutions. His work bridges advanced modeling and simulation techniques with industry standards like ASHRAE standards to deliver resilient, sustainable, and highly efficient digital infrastructure. Amir's expertise includes physical infrastructure design, advanced thermal management, and strategic performance modeling for AI rack architectures. He is a licensed Professional Engineer with approximately 20 years of experience in consulting engineering and infrastructure design.

Dr. Maria A. Torres Arango is a Research Analyst in Schneider Electric's Data Center Research & Strategy group. She investigates technologies, materials, and infrastructure to guide data center strategies. She contributes to developing analytical TradeOff Tools to evaluate the impact of technology decisions in data centers. Maria's expertise includes materials design and optimization, as well as fundamental studies on materials synthesis processes using X-ray characterization at the National Synchrotron Light Source II, Brookhaven National Laboratory. Maria holds a BS in Aeronautical Engineering from Universidad Pontificia Bolivariana, Colombia; and a MSc in Aerospace Engineering, and a PhD in Materials Science and Engineering from West Virginia University.

Acknowledgements

Special thanks to **Stefano Lena, Florian Hanslik and Wendy Torell** for their expertise and technical review of this white paper.

RATE THIS PAPER ★★★★★

This document is to be considered as an opinion paper presenting general and non-binding information on a particular subject. The analysis, hypothesis and conclusions presented therein are provided as is with all faults and without any representation or warranty of any kind or nature either express, implied or otherwise.

Resources

10 Ways to Harness the Energy and Water Efficiencies of Direct Liquid Cooling
White Paper 211

Economizer Modes of Data Center Cooling Systems
White Paper 132

Navigating Liquid Cooling Architectures for Data Centers with AI Workloads
White Paper 133

How Sustainable Technology, Design, and Operations can Unlock Community Support for Data Centers
Executive Report 16

Browse all white papers and executive reports
Data Center Research & Strategy Papers

Data Center Capital Cost Calculator
TradeOff Tool 4

Browse all TradeOff Tools™
Data Center Research & Strategy Tools

Note: Internet links can become obsolete over time. The referenced links were available at the time this paper was written but may no longer be available now.

Contact us

For feedback and comments about the content of this white paper:

Schneider Electric Data Center Research & Strategy
dcsc@se.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm

© 2026 Schneider Electric. All rights reserved.

Appendix

Table A1

Overview of simulated scenarios

	Scenario 1	Scenario 2	Scenario 3	Scenario 4
Description	Traditional data center with air cooling	AI factory with liquid cooling	AI factory with liquid cooling, increased TCS temperature	AI factory with liquid cooling, increased TCS temperature, reduced equipment count
IT rack density	20 / 40 kW	Server rack density 188 kW (liquid cooled*) Networking rack density 20 / 40 kW (air cooled)	Server rack density 188 kW (liquid cooled*) Networking rack density 20 / 40 kW (air cooled)	Server rack density 188 kW (liquid cooled*) Networking rack density 20 / 40 kW (air cooled)
IT capacity	100 MW air cooled	86.6 MW liquid cooled 13.4 MW air cooled	86.6 MW liquid cooled 13.4 MW air cooled	86.6 MW liquid cooled 13.4 MW air cooled
Operating temperatures	Air cooled loads: Chiller supply / return 20ºC / 30ºC CRAH supply / return 26.1ºC / 40.8ºC	Air cooled loads: Same as traditional Liquid cooled loads: Dry cooler** supply / return 30ºC / 35.3ºC TCS supply / return 32ºC / 42ºC	Air cooled loads: Same as traditional Liquid cooled loads: Dry cooler** supply / return 40°C / 50°C TCS supply / return 45°C / 55°C	Air cooled loads: Same as traditional Liquid cooled loads: Dry cooler** supply / return 40°C / 50°C TCS supply / return 45°C / 55°C
Types & quantity of systems	Air cooled loads: 184 fanwalls 80 air-cooled chillers	Air cooled loads: 32 fanwalls 24 air-cooled chillers Liquid cooled loads: 72 L-to-L CDUs 72 (Dallas) / 72 (Paris) dry coolers 80 (Dallas) / 72 (Paris) high temp. chillers	Air cooled loads: 32 fanwalls 24 air-cooled chillers Liquid cooled loads***: 72 L-to-L CDUs 72 (Dallas) / 72 (Paris) dry coolers 80 (Dallas) / 72 (Paris) high temp. chillers†	Air cooled loads: 32 fanwalls 24 air-cooled chillers Liquid cooled loads‡: 72 L-to-L CDUs 48 dry coolers 24 high temp. chillers (only in Dallas, no adiabatic scenario)†

* Assumes an IT rack heat capture ratio of 96%.
**Chiller assistance is necessary when outdoor temperature is too high for the dry cooler to reject all the heat to the environment.
*** Setting the TCS supply at 45°C while keeping the same equipment type and count as the scenario with TCS supply at 32°C, implies roughly 33% cooling capacity is not utilized (stranded). This is the only scenario where heat rejection has a redundancy greater than N+1.
† When TCS temperature is raised to 45°C, Paris scenarios do not necessitate high-temperature chillers to meet required cooling capacity. Dallas requires high-temperature chillers only in non-adiabatic cases.
‡ Equipment count differs from the other AI factory scenarios 2 and 3, where equipment was kept identical.