When Does Adiabatic Cooling Outperform Dry Cooling? A Twelve-City U.S. Data-Center Siting Assessment at ASHRAE Design Conditions

Vamsi Mokkapati [a],*, Debendra K. Das [b]

[a] *Technical Director, Research and Development, Nimbus Advanced Process Cooling, Anniston, AL, USA*

[b] *Professor of Mechanical Engineering, Emeritus, University of Alaska Fairbanks, Fairbanks, AK, USA*

* *Corresponding author. Email: vmokkapati@nimbus.cool*

ORCID (V. Mokkapati): 0009-0009-3529-9149
ORCID (D. K. Das): 0000-0001-9240-6515

Highlights

- Adiabatic cuts data center cooling electricity 30% (median) across 12 U.S. sites
- Trade ratio ~8–56 gal/kWh on-site water per kWh saved, site-dependent
- Dry cooling not feasible at Phoenix, Las Vegas, Tucson at 0.4% design conditions
- Adiabatic does not reduce operational water; it shifts the burden from grid to site
- Framework pairs design dry-bulb temperature, grid carbon, water stress, tariff ratio

Abstract

U.S. data center operators face growing scrutiny over water and electricity use. Dry cooling has been positioned as the lower-water-impact choice on the basis that it uses no on-site water; this view excludes the upstream water and carbon associated with the additional electricity dry coolers consume. We compared dry and adiabatic cooling for a 10 MW reference facility across twelve U.S. data center hubs selected for adiabatic suitability, spanning hot-arid Southwest, high-desert Mountain West, high-plains, and Mediterranean climates, at three ASHRAE design conditions per location (0.4%, 1%, 2% exceedance); humid climates are addressed in companion work (in preparation). Annual cooling electricity, on-site and upstream water, Scope 2 carbon, and operating cost were computed for the 29 of 36 cases where both technologies are feasible. Adiabatic cooling reduced heat-rejection electricity by a median of 30.5% and maximum 41% under continuous-spray, fan-modulating operation, at a trade ratio ranging from about 8 to 56 gallons (30 to 210 liters) of on-site water per kWh saved, depending strongly on location. This exceeds the upstream water saved per kWh by about an order of magnitude, so adiabatic shifts the water burden from grid to site rather than reducing operational water. At Phoenix, Las Vegas, and Tucson 0.4%, dry cooling cannot reach the 107 °F (41.7 °C) target fluid leaving temperature (FLT) and is unavailable as a primary technology. On an energy-and-water operating-cost basis (excluding demand charges, which would further favor adiabatic), the cost ranking turns on the electricity-to-water tariff ratio more than on climate alone; El Paso, Tucson, Albuquerque, and Reno define the tariff and climate boundary where the adiabatic advantage breaks down. A four-axis framework pairing design conditions, grid carbon, water stress, and the electricity-to-water tariff ratio supports siting choices and bears on current state-level moratorium debates.

Keywords: data center cooling; adiabatic cooling; dry cooling; ASHRAE design conditions; water-energy nexus; Scope 2 carbon; siting framework

Nomenclature

Symbol	Definition
AHRI	Air-Conditioning, Heating, and Refrigeration Institute
ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning Engineers
AZNM	eGRID subregion: WECC Southwest (Arizona–New Mexico)
CAMX	eGRID subregion: WECC California
CO_2e	Carbon dioxide equivalent
DBT	Dry-bulb temperature (°F or °C)
EC	Electronically commutated (fan motor type)
eGRID	EPA Emissions & Generation Resource Integrated Database
EIA	U.S. Energy Information Administration
EPA	U.S. Environmental Protection Agency
FLT	Fluid leaving temperature (°F or °C)
GPM	Gallons per minute
IT	Information technology
LBNL	Lawrence Berkeley National Laboratory
MCWB	Mean coincident wet-bulb temperature (°F or °C)
MWh	Megawatt-hour
NWPP	eGRID subregion: WECC Northwest
PUE	Power usage effectiveness
RMPA	eGRID subregion: WECC Rockies

TMY	Typical Meteorological Year (weather data)
W1–W5	ASHRAE TC 9.9 liquid cooling classes (W3 = 32–46 °C / 90–115 °F)
WBD	Wet-bulb depression (DBT minus MCWB)
WIN	Process water inlet temperature (°F or °C)
WOUT	Process water outlet temperature (°F or °C)
WRI	World Resources Institute
WUE	Water usage effectiveness (L/kWh)

1. Introduction

By 2026, eleven U.S. states are weighing moratoria or restrictions on new data center construction [1]. The trigger varies (water permitting in Arizona, grid capacity in Virginia, fiscal incentives in Georgia), but water consumption sits at the center of most of these debates [3, 13]. The 2024 Lawrence Berkeley National Laboratory report [4] puts U.S. data center electricity at 176 TWh in 2023, with a projected range of 325 to 580 TWh by 2028. Direct cooling water consumption reached 17.4 billion gallons (65.9 billion liters) in 2023 [4]. Indirect water embedded in grid electricity is larger still: AWWA's 2025 technical report [12] estimates this at over 200 billion gallons (757 billion liters) per year. Two-thirds of U.S. data centers built since 2022 are sited in regions of high or extremely high water stress per WRI Aqueduct 4.0 [24].

Two cooling architectures dominate the conversation for facilities in this load class. Dry coolers reject heat through sensible-only heat transfer at the coil surface, consuming no water at the facility boundary. Adiabatic coolers add a wet-mode operating state, in which entering air is pre-cooled toward the wet-bulb temperature through once-through misting; this reduces fan demand at the cost of on-site water consumption. Some industry voices have responded to water-permitting pressure by positioning dry cooling as the lower-water-impact alternative. The argument is simple: dry coolers use no water on site, so they avoid local water stress. The framing relies on a narrow system boundary. Dry coolers typically use more electricity, and that electricity carries its own water and carbon footprint upstream. The upstream cost is rarely reported in vendor literature.

The peer-reviewed literature has not closed this gap. Karimi et al. [15] developed the canonical Phoenix case study showing higher PUE but lower source-WUE for air-cooled relative to evaporatively cooled facilities. Lei and Masanet [5] developed a hybrid statistical and thermodynamic framework, and Lei et al. [16] identified the eight primary determinants of data center water consumption. What is missing is a systematic, location-explicit comparison across the U.S. data center siting envelope at standardized design conditions. We provide one half of that comparison here, covering twelve hot-arid, high-desert, high-plains, and Mediterranean hubs. Companion work (in preparation) covers humid temperate and continental climates, where the technology ranking is different and the engineering comparison is closer.

2. Literature review

2.1. Data center water and energy at scale

Data center water and energy demand has been documented in a sequence of national assessments. The first systematic U.S. estimate appeared in Shehabi et al. (LBNL 2016) [9]; the 2024 update [4], whose headline figures are summarized in Section 1, extended these national estimates. Lei and Masanet [5] developed a hybrid framework for PUE and WUE estimation across U.S. climates. Mytton [10] drew attention to the indirect water footprint of grid electricity. Siddik et al. [11] developed county-level data center water consumption estimates. The picture is one of growing demand on both ends of the water-energy nexus.

More recent work focuses on the spatial mismatch between data center siting and water availability. MSCI's 2025 sustainability research [2] finds 32% of U.S. data centers are sited in regions of high or extremely high water stress per WRI Aqueduct 4.0 [24]. The Environmental Law Institute's 2025 fact sheet [3] documents the resulting state-level policy debates, and the EESI 2025 issue brief [13] tracks the same trend at the federal level. The American Water Works Association's 2025 technical report [12] estimates indirect water associated with U.S. data center electricity at over 200 billion gallons (757 billion liters) per year, an order of magnitude larger than direct use. Jones Lang LaSalle's (JLL) 2026 market report [1] identifies water permitting as a leading siting constraint in 12 of the top 20 U.S. markets, and tracks emerging hubs in West Texas, the Mountain West, and the High Plains where the present analysis is concentrated. Policy responses to the water-stress mismatch are appearing in real markets: S&P Global reports that southern Nevada has restricted evaporative cooling in new data center construction [14].

2.2. Cooling technology comparisons

Karimi et al. [15] led the first peer-reviewed cooling-technology comparison grounded in a real U.S. data center hub. Their Phoenix case study reported higher PUE but lower source-WUE for air-cooled facilities relative to evaporatively cooled ones. The trade-off depended on local grid water intensity. Lei and Masanet [5] extended the framework with a hybrid statistical and thermodynamic model for PUE and WUE estimation across U.S. climates and cooling types. Lei et al. [16] used Sobol sensitivity analysis to identify the eight primary determinants of data center water consumption, with server energy efficiency and grid water intensity ranking above cooling architecture.

Li et al. [17] provide the largest-scale cradle-to-grave analysis to date, extending the system boundary to manufacturing and end-of-life impacts for liquid cooling. The result is that the choice of cooling architecture matters, but its weight in the total environmental footprint varies with location-specific context. Recent industry-side studies [8] have characterised adiabatic pad and misting performance in controlled tests, providing the empirical anchors for the saturation-effectiveness assumptions used in design tools. Cejudo et al. [18] reviewed water reuse systems for cooling tower applications, documenting the trade-off between cycles of concentration and chemical-treatment energy.

2.3. The water-energy nexus and grid water intensity

The water cost of electricity generation has been quantified in successive iterations since the 1980s. Macknick et al. [19] compiled the consumption-factor literature still in routine use: 0.687 gal/kWh (2.60 L/kWh) for coal with cooling tower, 0.198 gal/kWh (0.750 L/kWh) for combined-cycle gas with tower, 0.672 gal/kWh (2.54 L/kWh) for nuclear with tower, approximately 4.5 gal/kWh (17.0 L/kWh) for hydroelectric reservoir evaporation, and below 0.05 gal/kWh (0.19 L/kWh) for solar PV and wind. EPA's eGRID2023 (Revision 2, June 2025) [20] provides current subregion-level fuel mix and emission factors. Weighted by Macknick consumption factors, these yield approximately 0.5 gal/kWh (1.89 L/kWh) as the U.S. average upstream water intensity, computed as a consumption-weighted (not withdrawal-weighted) factor over the eGRID2023 national fuel mix. In the subregions studied here, upstream consumption-weighted intensity ranges from approximately 0.43 gal/kWh (1.63 L/kWh; AZNM, gas-dominant) to 1.93 gal/kWh (7.31 L/kWh; NWPP, hydro-heavy, including reservoir evaporation per Macknick et al.). The U.S. average sits below the subregion minimum studied here because the studied subregions exclude lower-water-intensity Eastern grids dominated by combined-cycle gas and renewables. We note that the upper end of the subregion range (NWPP, hydro-heavy) is sensitive to how hydroelectric reservoir evaporation is allocated; the Macknick consumption factors include hydro reservoir surface evaporation at a typical assignment, but this allocation remains methodologically contested in the life-cycle assessment literature. These values matter for the total-water analysis in Section 4.

2.4. Fan motor technology in cooling systems

Fan power dominates the electricity demand of both dry and adiabatic cooling systems. The transition from AC induction motors to electronically commutated (EC) brushless motors has become a meaningful efficiency lever. Industry analyses report 30 to 70% energy reductions when EC fans replace AC fans, with the largest gains at part-load [6, 21]. EC fans hold near-flat efficiency curves while AC efficiency drops sharply away from the design point. Both technologies in this study use EC motors as standard. The resulting part-load behavior is central to the trade-off we characterize.

2.5. ASHRAE design conditions and liquid cooling classes

ASHRAE Climate Design Data [7] provides the standardized reference for sizing heat rejection equipment, with annual cooling conditions reported at 0.4%, 1%, and 2% exceedance probabilities. ASHRAE Technical Committee 9.9 defines six liquid-cooling classes for data centers (W1 through W5 plus W+), each specifying a permissible range of supply water temperature to the IT equipment. The 107 °F (41.7 °C) target FLT used here matches class W3, which covers the high-density, AI-oriented rack architectures now coming into deployment. Higher FLT targets (W4 at 113 °F (45 °C), W+ at 122 °F (50 °C)) expand the climates where dry cooling remains feasible but require IT equipment qualified for the higher operating temperatures.

2.6. Research gap

Together, the cooling-technology comparison methodology, grid water consumption factors, fan motor improvements, and ASHRAE class system reviewed above provide the analytical building blocks. Karimi et al. [15] cover Phoenix in detail; Lei and Masanet [5] provide a framework but no per-location analysis. The present work delivers the per-location comparison for hot-arid, high-desert, high-plains, and Mediterranean climates.

3. Methodology

3.1. Reference facility

The reference facility is a 10 MW (34.1 MMBtu/h, 2,844 ton) data center running continuously. The cooling load is set by the process water: 5,250 GPM (331 L/s) circulating between WIN = 120 °F (48.9 °C) and WOUT = 107 °F (41.7 °C). The 107 °F target fluid leaving temperature (FLT) matches ASHRAE class W3 liquid cooling and reflects the high-density, AI-oriented rack architectures now coming into deployment. All economic figures are reported as annualized operating cost. The analysis is restricted to the heat-rejection subsystem (dry cooler or adiabatic cooler at the building exterior); chillers, coolant distribution units, IT-side electrical loads, and overall facility PUE are not modeled. Electricity-reduction percentages reported in this paper refer to the heat-rejection subsystem alone, not to total facility electricity.

3.2. Locations and design conditions

We selected twelve locations to span the climates where adiabatic operation captures meaningful wet-bulb depressions, deliberately filtering for adiabatic-suitable territory in service of characterising where this technology delivers its largest gains. The selection includes hot-arid Southwest sites (Phoenix, Las Vegas, Tucson) where dry feasibility is at risk, hot semi-arid border sites (El Paso, Albuquerque), high-desert and high-plains interior sites (Reno, Salt Lake City, Boise, Denver, Colorado Springs, Cheyenne) where wet-bulb depressions at design exceed 30 °F (16.7 °C), and Santa Clara as a Mediterranean reference. The selection prioritises actual U.S. data center hubs and JLL-tracked emerging markets [1] over generic climate representatives. Humid temperate, humid subtropical, and humid continental climates, where dry cooling is more often competitive, are covered in companion work (in preparation). Table 1 lists each location with its eGRID2023 subregion [20], grid CO2 intensity, and water stress classification per WRI Aqueduct 4.0 [24].

Table 1. Reference locations and grid context.

Location	eGRID subregion	Grid CO_2 (kg/kWh)	Water stress (WRI)
Phoenix, AZ	AZNM	0.320	Extremely High
Las Vegas, NV	AZNM	0.320	Extremely High
Tucson, AZ	AZNM	0.320	Extremely High
Albuquerque, NM	AZNM	0.320	High
El Paso, TX	AZNM	0.320	Extremely High
Reno, NV	NWPP	0.288	High
Salt Lake City, UT	NWPP	0.288	High
Boise, ID	NWPP	0.288	Medium-High
Cheyenne, WY	RMPA	0.470	High
Denver, CO	RMPA	0.470	High
Colorado Springs, CO	RMPA	0.470	High
Santa Clara, CA	CAMX	0.195	High

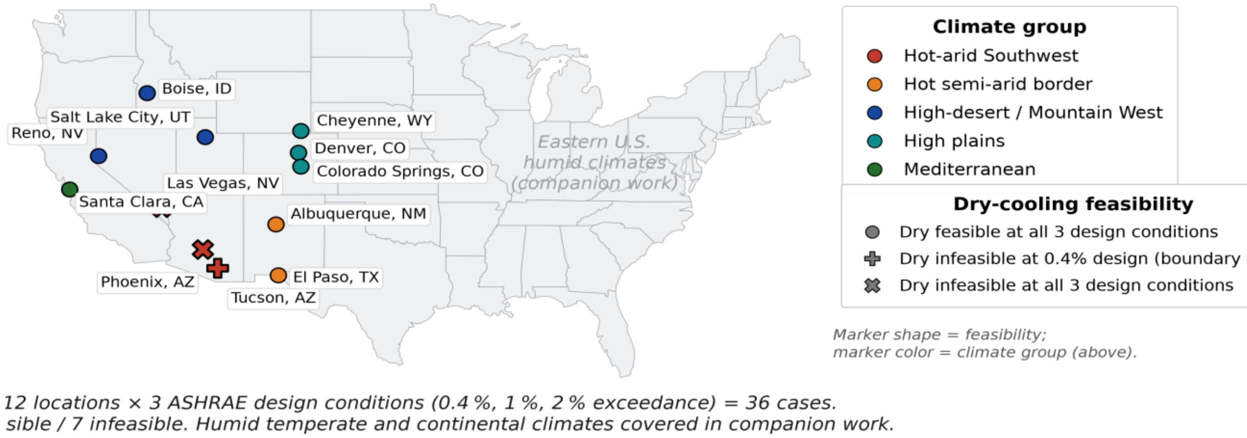

Fig. 1. Study locations across U.S. adiabatic-suitable climate zones. Twelve locations span hot-arid Southwest, hot semi-arid border, high-desert / Mountain West, high-plains, and Mediterranean climates. Phoenix and Las Vegas are infeasible for dry cooling at all three ASHRAE design conditions; Tucson is infeasible at the 0.4 % peak design point only. Humid temperate and continental climates (eastern U.S.) are addressed in companion work (in preparation). Map lines delineate study areas and do not necessarily depict accepted national boundaries.

Each location was evaluated at three ASHRAE Climate Design Data 2025 annual cooling design conditions: 0.4%, 1%, and 2% exceedance. Both the dry-bulb temperature (DBT) and the mean coincident wet-bulb temperature (MCWB) were used. Annual energy and water consumption were accumulated from 2025 TMY hourly weather data binned by DBT.

3.3. Cooling technology models

We computed cooling system sizing and annual energy using in-house engineering tools that combine two components. The dry-cooler component implements a steady-state effectiveness–NTU (ε–NTU) formulation for the air-to-fluid heat exchanger, with effectiveness expressed as a function of the air-side and fluid-side heat capacity rates, the air-side Reynolds number at the fin geometry in use, and the entering-air dry-bulb temperature. Coil heat-transfer coefficients used in the effectiveness calculation are certified to AHRI Standard 410 [22] (Performance Rating of Forced-Circulation Air-Cooling and Air-Heating Coils). We note explicitly that AHRI 410 certifies the rated coil heat-transfer performance only, not the fan-power or annual-integration components of the model. Fan power at design and off-design conditions is computed from manufacturer fan curves (electrical input as a function of volumetric flow at the appropriate air density for site elevation), and is approximately cubic with airflow at moderate part-load while flattening at deep part-load owing to EC motor efficiency characteristics. Annual electricity is accumulated by binning 2025 TMY hourly weather by dry-bulb temperature and applying the fan-modulated power state appropriate to each bin. Specific coil UA values, fan-curve coefficients, and the sizing tool's unit-selection logic are proprietary and not disclosed; the functional dependencies described above are sufficient to interpret the comparative results that follow.

The adiabatic-cooler component adds a wet-mode operating state to the dry-cooler model. In wet mode, entering air is pre-cooled toward the wet-bulb temperature through once-through misting upstream of the coil face. Wet mode activates when the ambient dry-bulb temperature exceeds a Dry Switch Point (DSP) determined from controlled in-house experiments on the spray and coil geometry used in this study; the DSP, the underlying correlation, and the experimental dataset are proprietary to the engineering tool and are not disclosed in this paper. The wet-mode threshold dry-bulb temperature and the water consumption rate during wet operation are computed outputs of the model, not user inputs.

We acknowledge that comprehensive external validation of the wet-mode correlation against an independent peer-reviewed dataset at the W3 target FLT is not currently available, because no published study has reported equivalent measurements at this fluid temperature. The closest peer-reviewed study for the dry-cooled comparison is by Karimi et al. [15], which characterizes an air-cooled colocation facility at Phoenix conditions. Their reported PUE and source-WUE values are directionally consistent with our model output for the Phoenix dry-cooled case, though direct numerical benchmarking requires equating IT load, process temperatures, and facility envelope assumptions across the two studies and is therefore approximate rather than validated. We flag this limitation explicitly: the absence of an external W3-class wet-mode benchmark is part of the gap this paper begins to address, and future work should include side-by-side field measurements against the model.

Three practical points are worth emphasizing. First, dry-cooler power values are higher than vendors typically advertise for nominal-condition operation; advertised values usually correspond to lower process water temperatures and milder ambient conditions, whereas our values reflect the W3 class target FLT under design ambient. Second, the wet-mode threshold sits a few degrees above the lowest ambient temperature at which adiabatic provides benefit, because at lower ambient temperatures the fan demand is already low and the water cost is not justified. Third, the sizing tool selects the smallest unit-and-model configuration that satisfies the FLT feasibility constraint; because equipment is chosen in discrete catalog steps, this can produce non-monotone relationships between design percentile and annual MWh at locations where the selection crosses a unit-count step between design points. Reported water consumption for adiabatic is the total supplied volume during wet-mode hours; the system is once-through with no recirculation, blowdown, or drift, and reported volumes therefore equate supplied with consumed water (i.e., they do not separately account for any non-evaporated overspray that may collect in the drain pan).

Both technologies are modeled with electronically commutated (EC) fan motors. EC motors hold near-flat efficiency curves at part-load, which matters because both systems spend most of their annual hours operating well below the design condition. In wet mode the misting spray runs continuously rather than cycling, and the EC fans modulate down as the pre-cooled air lowers the required airflow; reported adiabatic energy reflects this continuous-spray, fan-modulating operation.

Because the Dry Switch Point (DSP) correlation in this tool is calibrated to properly-maintained spray geometry under controlled conditions, real-world systems may deliver lower performance due to nozzle fouling, partial blockage, drift loss, or off-design face velocities. A formal sensitivity analysis on the correlation is not presented in this paper and is expressed as an acknowledged limitation. Qualitatively, a higher-DSP scenario (later wet-mode engagement) would shift the electricity reduction downward and the trade ratio upward, weakening but not reversing the comparative findings for high-WBD locations. Operators should note that real-world performance depends on maintenance and water quality, and the reported results assume the spray system operating at design effectiveness.

Dry cooling is flagged as infeasible when the 107 °F (41.7 °C) target FLT lies below ambient DBT plus a 2 °F (1.1 °C) minimum approach. Phoenix and Las Vegas exhibit this infeasibility at their 0.4% design conditions and remain infeasible at the 1% and 2% conditions because the tool's density-corrected achievable approach at those conditions cannot deliver the 107 °F (41.7 °C) target FLT. Tucson exhibits infeasibility at the 0.4% design condition (DBT 106.0 °F / 41.1 °C) but becomes feasible at the 1% and 2% conditions (DBT 103.8 °F / 39.9 °C and 101.7 °F / 38.7 °C). The comparison in Section 4 is restricted to the design points where both technologies can reach the target FLT.

3.4. Environmental and economic indicators

Four indicators were computed for each location, design condition, and technology. Annual cooling electricity demand "E" comes from the engineering tool. Scope 2 carbon emissions equal E times the eGRID2023 Revision 2 subregion CO_2 rate [20]. On-site water consumption is the tool output for adiabatic, and zero for dry. Annualized operating cost is electricity (E times commercial tariff) plus water (volume times utility tariff). Commercial electricity tariffs are state-level averages from EIA Electric Power Monthly Table 5.6.A (January 2026 issue) [23]. Water tariffs are from individual utility published rate sheets. Upstream water tied to cooling electricity uses Macknick et al. [19] median consumption factors weighted by the eGRID2023 Revision 2 subregion fuel mix.

3.5. Limitations

The analysis uses annual-average inputs. It does not model hourly dispatch, time-of-use tariffs, demand charges, or peak-load capacity reservation. Demand charges in particular are material: the ~3 MW peak demand difference between dry and adiabatic at a 10 MW facility could add $360,000 to $720,000 per year in capacity charges at typical commercial rates ($10 to $20 per kW per month), which exceeds the annualized operating-cost spreads reported here and would push the cost ranking further toward adiabatic at every feasible site. The reported cost results should therefore be interpreted as energy-and-water operating cost only, not total cost of ownership. eGRID factors are attributional grid averages and do not capture power purchase agreements or projected decarbonization. As a bound on the marginal effect, applying a uniform combined-cycle-gas marginal factor (about 0.40 kg CO_2/kWh) raises the dataset-total Scope 2 savings by roughly 14%, from about 1,800 tCO_2e per year on the attributional basis to about 2,060 tCO_2e per year, and redistributes them: roughly doubling the clean-grid Santa Clara savings while lowering the coal-influenced RMPA-site savings, without changing the qualitative ranking. One operability term is excluded from the adiabatic side: freeze protection at the high-elevation Mountain West and high-plains sites (heat trace or drain-down loads, which are site-specific and beyond the scope of the bin-data energy model). The misting system in this study uses pressurized municipal supply directly through the spray nozzles and does not include a booster pump, so no misting-pump parasitic load applies; control-power auxiliary draw is negligible relative to the fan electrical baseline. Embodied impacts of cooling equipment are excluded. The analysis is restricted to a single FLT target (107 °F (41.7 °C), W3 class). W3 is the coldest ASHRAE class that admits dry cooling under our 2 °F (1.1 °C) minimum approach criterion at the design conditions studied, which means it maximizes both the dry-infeasibility count and the adiabatic electricity-reduction advantage. At looser targets (W4 at 113 °F (45 °C), W+ at 122 °F (50 °C)) the feasibility space for dry cooling expands and several of the seven W3-infeasible cases likely revert to feasible (Phoenix 0.4% at 111.2 °F (44.0 °C) sits below W4, for example). At tighter targets (W2, W1) dry-cooler feasibility narrows further and the comparison would shift toward adiabatic. Future work should repeat the feasibility analysis at W4 and W+ to bound how many of the W3-infeasibility findings carry across the full liquid-cooling class system.

The wet-mode performance correlation used in this study is proprietary to the engineering tool, and neither its functional form nor its empirical coefficients are disclosed in this paper. We acknowledge this constrains independent reproducibility of the absolute values reported here. While absolute values are tool-dependent, the qualitative ranking of locations along the wet-bulb-depression axis is consistent with adiabatic-saturation physics and would be expected to survive correlation re-calibration in direction (the positive WBD-to-reduction relationship), even if absolute magnitudes shift. We do not, however, claim that rankings are immune to correlation bias: the DSP correlation is condition-dependent, so a biased correlation would distort cases non-uniformly along the wet-bulb-depression axis on which the rankings turn. Future work should include side-by-side field measurements at one or more sites to validate the correlation against an external dataset.

A related transparency note: the lead author is employed by Nimbus Advanced Process Cooling, which owns the engineering tool that produced every quantitative result in this paper. No independent re-execution of the underlying calculations was performed within this study; the co-author's contribution was limited to manuscript review and editorial input and did not include re-running the model. Readers should weigh this against the explicit competing-interest declaration. Future work to address this constraint includes field measurement at one or more study sites and independent tool benchmarking.

4. Results

4.1. Cooling electricity

Adiabatic cooling reduced annual cooling electricity at every feasible location and design condition (Fig. 2). Of the 36 combinations of the twelve locations and three design conditions evaluated, 29 are feasible for dry cooling and 7 are not: Phoenix at all three design points, Las Vegas at all three design points, and Tucson at the 0.4% design point exceed the 107 °F (41.7 °C) target FLT plus a 2 °F (1.1 °C) minimum approach, so dry cooling cannot reach the required leaving fluid temperature. The remaining 29 cases form the basis of the comparison. The median reduction was 30.5%, with a maximum of 41% at Colorado Springs' 1% design point and a minimum of -1% at El Paso's 1% design point, where a well-matched dry bank running largely at part load leaves adiabatic with essentially no net electricity benefit. The pattern tracks wet-bulb depression at design. Salt Lake City, Albuquerque, Reno, Denver, Boise, El Paso, Colorado Springs, and Cheyenne all have WBD above 30 °F (16.7 °C) at peak design and achieved per-case reductions spanning 15% to 41% (excluding El Paso 1%). Santa Clara, with peak-design WBD of 26 °F (14.4 °C) and 1% / 2% design WBDs of 23 °F (12.8 °C) and 20 °F (11.1 °C), achieved 8% to 20% depending on the design condition. Tucson sits between these groups: its 1% design point (WBD 38 °F / 21.1 °C) gave a 36% reduction matching the high-WBD cluster, while its 2% design point (WBD 36 °F / 20.0 °C) only reached 12% because the sizing tool used a smaller dry-cooler bank with higher per-unit fan power. These reductions use the calibrated continuous-spray, fan-modulating model. The reported median of 30.5% is a calibrated estimate; its principal uncertainty is the proprietary wet-mode correlation, discussed as a limitation in Section 3.5.

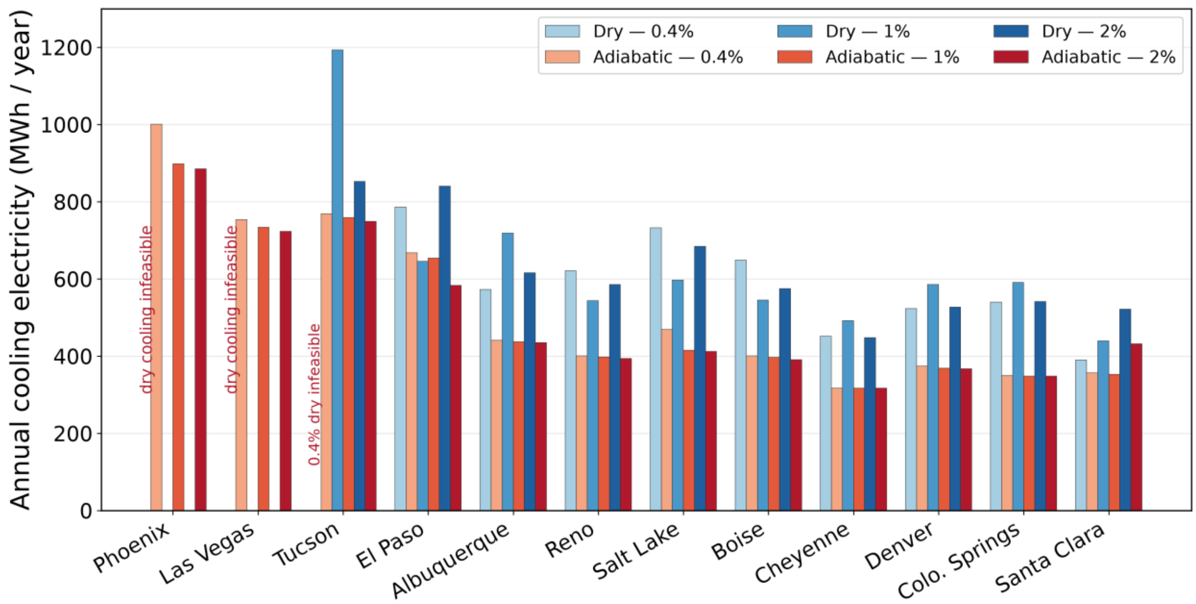

Fig. 2. Annual cooling electricity by location, technology, and ASHRAE design condition. Dry cooling is infeasible at all three Phoenix and Las Vegas design conditions and at Tucson's 0.4 % design condition; only the adiabatic bars appear at those points.

The thermodynamic mechanism is direct (Fig. 3): wet-bulb depression (WBD) at design, defined as the gap between dry-bulb temperature (DBT) and mean coincident wet-bulb temperature (MCWB), sets the maximum air-temperature shift available to adiabatic pre-cooling, with higher WBD permitting a larger shift. The observed relationship between % reduction and WBD across the 29 feasible cases is positive but noisy, because the dry-cooler baseline each case is compared against depends on the installed unit count and on how much of the year that bank runs at part load, where EC-fan efficiency is highest, an engineering effect the continuous WBD axis does not capture. The WBD axis describes the available thermodynamic headroom; realized reductions depend additionally on the size and part-load operation of the dry bank in each case. The cube-law relationship between airflow and fan power amplifies the effect at high WBD. The Mountain West and high-desert sites at 32 to 37 °F (17.8 to 20.6 °C) WBD sit high on the curve. The Mediterranean reference (Santa Clara) at 20 to 26 °F (11.1 to 14.4 °C) WBD sits lower. El Paso, despite the highest peak-design WBD in the dataset (38.5 °F / 21.4 °C), falls below the trend at its 1% condition because its dry bank there runs largely at part load, lowering the dry baseline. That effect is discussed below.

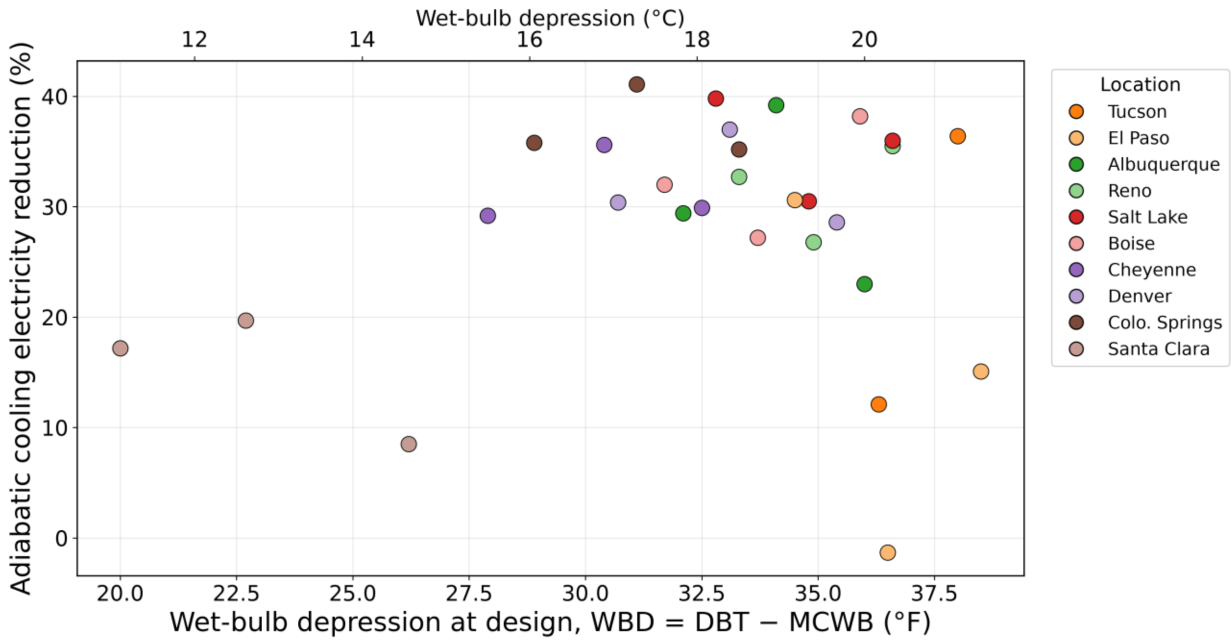

Fig. 3. Adiabatic electricity reduction trends with wet-bulb depression. Each point is one feasible combination of location and design condition; colours identify locations. El Paso 1 % sits below the trend because its dry bank there runs largely at part load, lowering the dry-energy baseline it is compared against (see Section 4.2).

Aggregating across the three design conditions per location (Fig. 4) shows the cross-location pattern compactly. Colorado Springs and Salt Lake City lead with average reductions of 37% and 35%. Boise, Denver, Reno, Cheyenne, and Albuquerque follow as a cluster at 30% to 33%. Tucson averages 24% across its two feasible cases, with the 1% point (36%) on par with the Mountain West cluster but the 2% point far lower. El Paso averages 15%, held down by its near-zero 1% case. Santa Clara averages 15%. Phoenix and Las Vegas do not appear because dry cooling is not feasible at any design condition there, and Tucson's 0.4% point is also outside the feasible set.

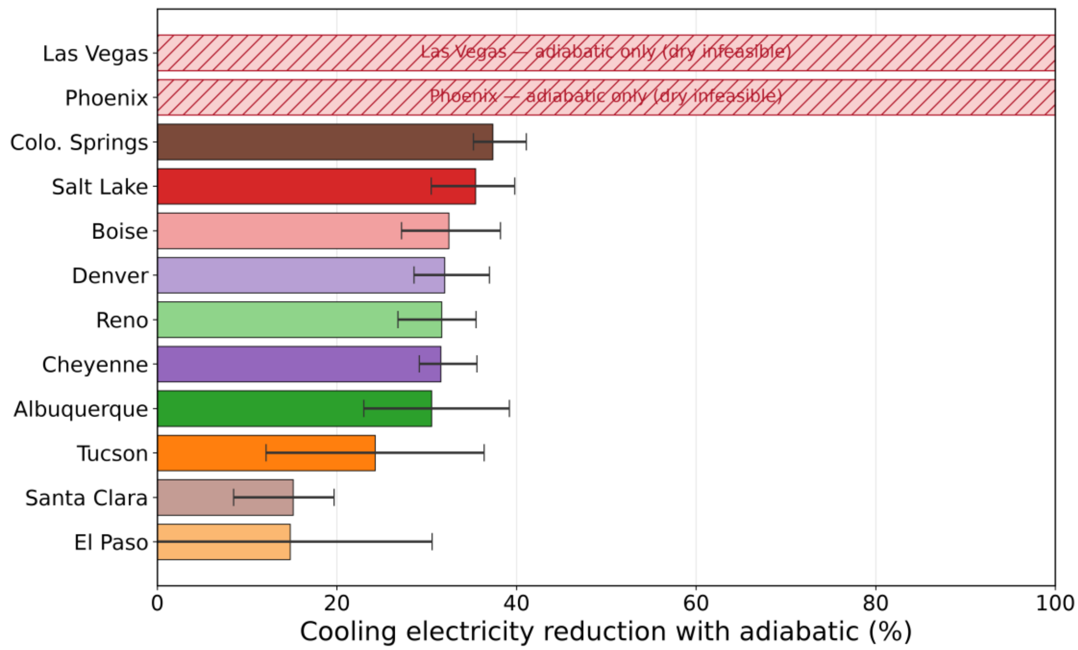

Fig. 4. Adiabatic electricity savings ranked by location. Range bars show the minimum-to-maximum range of electricity reduction across the feasible ASHRAE design conditions (0.4%, 1%, 2%) at each location, with the bar marking the location average; they represent a deterministic range across the three design points evaluated, not a statistical confidence interval. Phoenix and Las Vegas appear with hatched bars and no range bars because dry cooling is infeasible at all three design conditions, so no comparison can be drawn.

El Paso's 1% case warrants further discussion. The dry-cooler sizing tool selected 11 units at the 0.4% design point (5.1 °F / 2.8 °C approach), 9 units at 1% (7.4 °F / 4.1 °C approach), and 8 units at 2% (9.6 °F / 5.3 °C approach), selecting the smallest feasible bank meeting the FLT constraint at each point. The annual energy is approximately linear in unit count within a given fan-curve regime, and the per-unit annual energy changes in steps as the approach crosses fan-curve thresholds. At 0.4% and 1%, per-unit annual MWh remains in a similar band (~72 MWh/unit), so dropping from 11 to 9 units cuts annual dry MWh from 787 to 646. At 2%, the per-unit value steps up to ~105 MWh/unit, and the 8-unit bank produces 841 MWh. The 1% case lands at an optimum for the dry cooler. The adiabatic response under the continuous-spray, fan-modulating model is more modest (668, 654, 583 MWh across the three design percentiles), so at the 1% optimum the adiabatic energy slightly exceeds the dry energy and the reduction is essentially zero (about -1%), the lowest in the El Paso set, rather than a continuation of the WBD-driven trend. More broadly, dry-favorable results at several sites reflect a bank sized for the peak that operates mostly at part load, where EC-fan efficiency keeps annual energy low; this lowers the dry baseline on operating energy but requires a higher installed unit count, and therefore capital cost, which this study does not include (Section 4.4).

4.2. On-site water and the trade ratio

Adiabatic on-site water consumption ranged from approximately 1.26 million gallons (4.77 million liters) per year at Santa Clara's 2% design point to 5.83 million gallons (22.1 million liters) per year at Tucson's 2% design point. Dividing through by the matched electricity savings gives the trade ratio (Fig. 5): from about 8 to 56 gallons (30 to 210 liters) of on-site water per kWh of electricity saved across the feasible cases, with most cases in the 8 to 25 gal/kWh (30 to 95 L/kWh) band. The best case is Boise 0.4% (8.3 gal/kWh / 31.4 L/kWh), where the WBD is largest and the dry-mode baseline is largest; the highest ratios (up to 56 gal/kWh / 213 L/kWh, at Tucson 2%) fall where the dry-mode electricity saving is smallest. El Paso 1% does not yield a finite ratio, because the dry sizing tool's optimum leaves adiabatic with essentially no net electricity savings; this is the same case that produces the near-zero electricity reduction in Section 4.1.

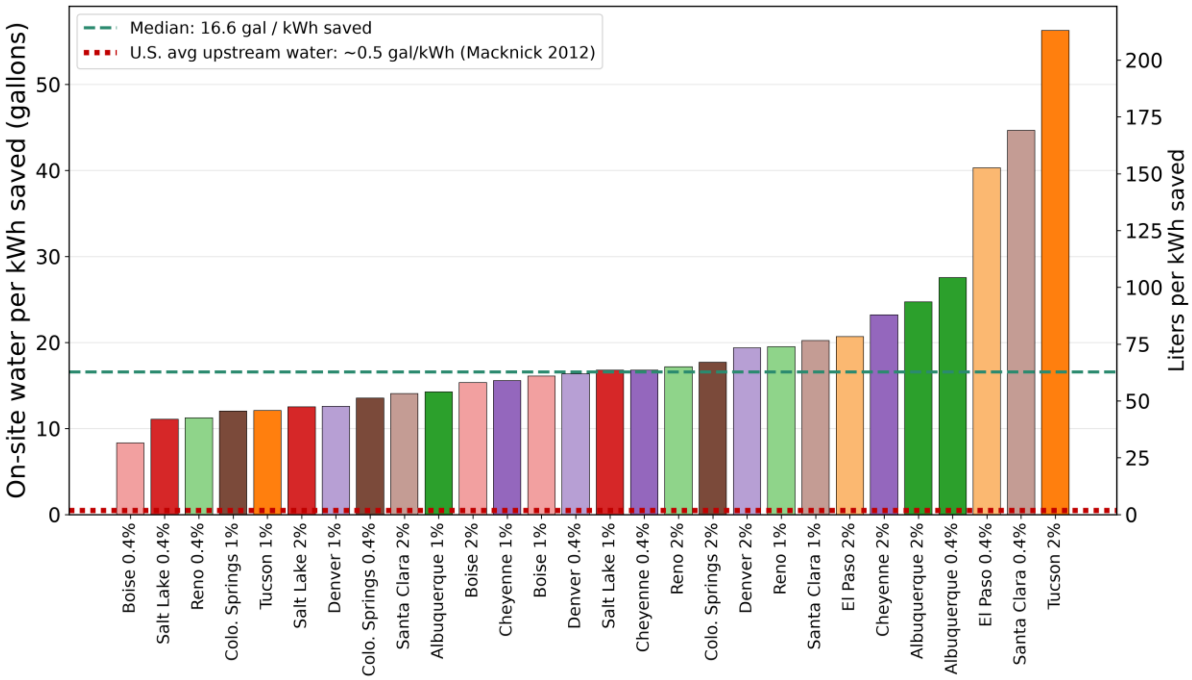

Fig. 5. Water-for-power trade ratio across feasible cases. The dashed green line shows the median (16.6 gal / kWh saved); the dotted red line shows the U.S. average upstream water consumption per kWh from Macknick et al. [19] weighted by eGRID2023 Revision 2 fuel mix.

For comparison, the upstream water consumption associated with one kWh of U.S. grid electricity averages around 0.5 gal/kWh (1.9 L/kWh; Macknick [19] weighted by eGRID2023 mix). The on-site trade ratio is roughly 9 to 39 times the upstream water per kWh, depending on subregion: smallest at the hydro-heavy NWPP sites (Reno, Salt Lake City, Boise), where upstream intensity is highest at 1.93 gal/kWh (7.3 L/kWh), and largest at the gas-dominant AZNM sites at 0.43 gal/kWh (1.6 L/kWh). A naive reader could conclude that adiabatic operation increases operational water use. This conclusion is actually correct on a strict on-site-plus-upstream basis at every site in the dataset, and we explain why it does not undercut the adiabatic case in Section 5.2.

4.3. Scope 2 carbon

Scope 2 carbon savings track the cooling electricity savings weighted by grid carbon intensity (Fig. 6). The largest absolute savings show up where electricity savings meet a carbon-intensive grid. Colorado Springs on RMPA (0.470 kg CO2/kWh) yields the highest location-aggregate savings at 98 tCO2e per year on average (range 89 to 114). Tucson produces the highest single-case carbon savings in the dataset at 139 tCO2e per year at its 1% design point, driven by the largest dry-cooler annual energy in the feasible set. Denver and Cheyenne on RMPA contribute 82 and 69 tCO2e per year respectively. Mountain West sites on NWPP yield 54 to 69 tCO2e per year. El Paso and Albuquerque on AZNM yield 39 and 63 tCO2e per year. Santa Clara on CAMX (0.195 kg CO2/kWh) shows the smallest absolute savings, a location average of 14 tCO2e per year (range 6 to 18), because its grid is the cleanest in the dataset. El Paso's 1% case is the lone negative in the dataset (about −3 tCO2e), because the dry sizing tool's optimum there leaves adiabatic with marginally higher electricity use; it is a sizing artifact, not a climate effect.

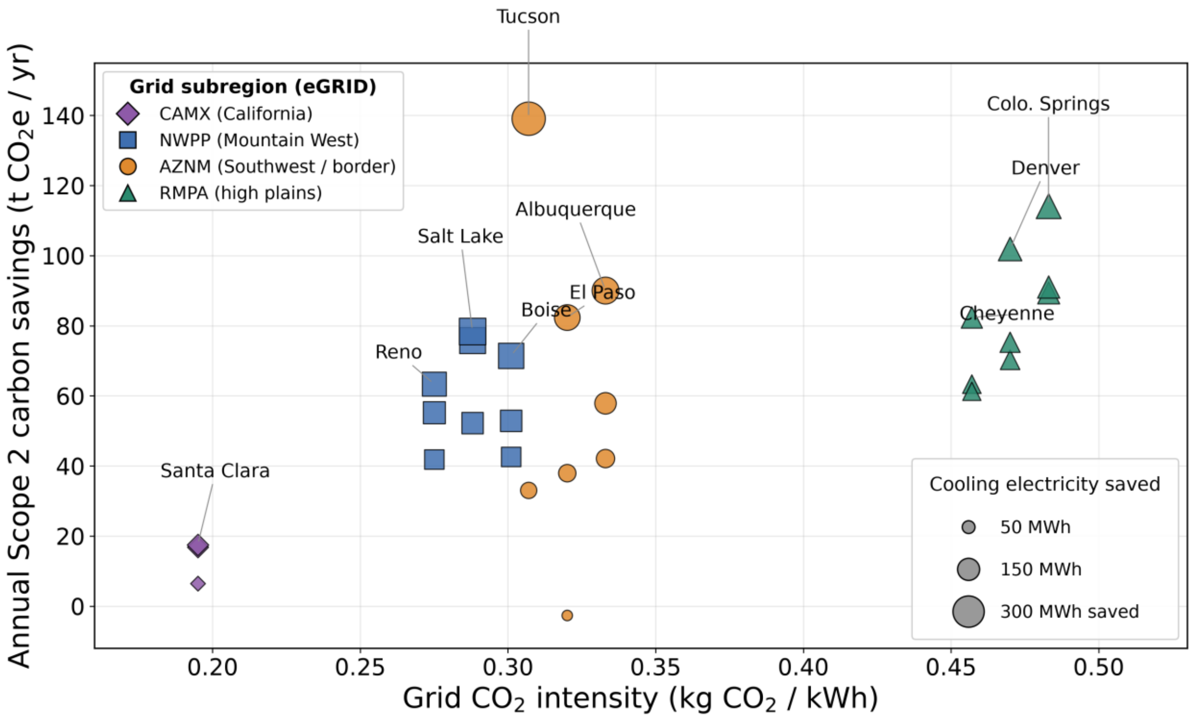

Fig. 6. Annual Scope 2 carbon savings as a function of local grid CO_2 intensity. Point size scales with the magnitude of cooling electricity saved. Carbon savings are largest where high electricity reductions meet a carbon-intensive grid (RMPA at 0.470 kg CO_2/kWh, AZNM at 0.320 kg CO_2/kWh).

4.4. Operating cost

Annualized operating cost ranking is set by the electricity-to-water tariff ratio more than by climate alone. Among the ten feasible locations, six fall on the adiabatic-favored side of that tariff division (Salt Lake City, Boise, Denver, Santa Clara, Cheyenne, and Colorado Springs), with location-aggregate annual savings ranging from about $190 at Denver to $4,700 at Boise. The other four sit on the dry-favored side (El Paso, Tucson, Albuquerque, and Reno). El Paso is the clearest dry-favorable site in the dataset, with dry operation roughly $25,400 per year cheaper on average across the three design points; Tucson and Albuquerque follow at about $11,900 and $11,700 per year cheaper for dry. Reno and Denver sit on the boundary, with Reno favoring dry by only about $480 per year and Denver favoring adiabatic by only about $190 per year; both should be read as close calls rather than firm assignments, consistent with their load-sensitivity, as discussed in Section 4.6. The mechanism at El Paso is tariff economics, not climate: its commercial electricity tariff at $86.90/MWh (Texas state average; the El Paso Electric service area is on the WECC AZNM subregion) is the lowest in the dataset, so the MWh savings adiabatic delivers are worth less in dollars than the water tariff cost of running the misting system.

Capital expenditure is not modeled here. It is the subject of vendor-specific quotation and varies meaningfully with the adiabatic system's water treatment infrastructure and the dry system's coil footprint. The operating-cost ranking and its robustness to tariff variation are characterised in Section 4.6.

4.5. Phoenix, Las Vegas, and Tucson infeasibility

Dry cooling is infeasible at Phoenix and Las Vegas at all three design conditions, and at Tucson's 0.4% design point. Phoenix DBT spans 107.1 to 111.2 °F (41.7 to 44.0 °C) across the three exceedance levels. Las Vegas DBT spans 104.9 to 109.2 °F (40.5 to 42.9 °C). Tucson DBT at 0.4% is 106.0 °F (41.1 °C), with the 1% and 2% points dropping to 103.8 and 101.7 °F (39.9 and 38.7 °C). Six of these seven design points exceed 105 °F (40.6 °C), so with the 2 °F (1.1 °C) minimum approach the 107 °F (41.7 °C) target FLT is thermodynamically unreachable, because a dry cooler cannot drive leaving fluid temperature below entering air temperature. The seventh, Las Vegas at the 2 % condition (104.9 °F / 40.5 °C), is reachable in principle but only with a pathologically oversized bank (roughly 20 units, more than twice any feasible site, with fans near full speed year-round), so it is treated as infeasible as a practical dry-cooling option. Operators at these sites have three options. They can accept higher leaving fluid temperatures, which requires IT equipment qualified for ASHRAE class W4 or W+ operation. They can supplement dry cooling with mechanical refrigeration on hot days, which adds a compressor system with its own cost and complexity. Or they can use evaporative or adiabatic cooling, which is the typical choice. Where local water policy restricts evaporative cooling, this option is itself constrained: southern Nevada's restriction is defined functionally and, read on its face, reaches adiabatic as well as continuously-wet systems [14, 28] (Section 5.4). Tucson is a particularly clean instructive example because dry cooling is feasible at the 1% and 2% conditions but not at the 0.4% peak design point, showing how the operational ambient envelope at a single site can straddle the feasibility boundary.

To confirm the infeasibility classification is not an artifact of the 2 °F (1.1 °C) minimum approach (the tightest approach the engineering tool resolves at standard air density). We re-classified all 36 design points at minimum approaches of 3 °F (1.7 °C) and 5 °F (2.8 °C), spanning the practical field range for dry coolers at design (Table 2). Because 2 °F (1.1 °C) is the tool's floor, these looser approaches are progressively more conservative assumptions; feasibility was re-evaluated by the sizing tool with elevation-dependent density derating. The seven base-case infeasible points remain infeasible at every tested approach. Raising the minimum approach to 5 °F (2.8 °C) adds only Tucson 1% (design DBT 103.8 °F / 39.9 °C), where 3.2 °F (1.8 °C) of headroom can no longer accommodate a 5 °F (2.8 °C) approach; El Paso 0.4% (101.9 °F / 38.8 °C) and all other feasible cases retain feasibility across the range, with the tool confirming an achieved leaving fluid temperature at or below 107 °F (41.7 °C) at site elevation in every case. The infeasibility finding is therefore conservative: 2 °F (1.1 °C) is the most dry-favorable assumption the model permits, and tightening the approach only widens the set of sites where dry cooling cannot reach the W3 target FLT. The full 36-case classification at each approach is provided in the Supplementary Material accompanying this paper.

Table 2. Sensitivity of dry-cooling feasibility to the minimum-approach assumption.

Minimum approach	**Infeasible cases (of 36)**	**Cases**
2 °F (1.1 °C) (base)	7	Phoenix (0.4 / 1 / 2%), Las Vegas (0.4 / 1 / 2%), Tucson 0.4%
3 °F (1.7 °C)	7	Same seven
5 °F (2.8 °C)	8	Above seven + Tucson 1%

4.6. Sensitivity analysis

We tested the robustness of the operating-cost ranking across ±20% variation in electricity and water tariffs. The four corner cases for each site (electricity high or low, water high or low) bound the cost-ratio envelope (Fig. 7). Under the calibrated model the rankings are markedly tariff-sensitive. Only two locations are robust across the full ±20% envelope: El Paso (cost ratio 1.20 to 1.66) and Albuquerque (1.01 to 1.41) are dry-favored at every corner. The remaining eight locations flip within the envelope. Six lean adiabatic at baseline but cross to dry-favored at the low-electricity / high-water corner: Boise (0.83 to 1.03), Salt Lake City (0.84 to 1.08), Colorado Springs (0.86 to 1.15), Cheyenne (0.88 to 1.13), Santa Clara (0.93 to 1.05), and Denver (0.89 to 1.16). The other two, Reno (0.90 to 1.17) and Tucson (0.98 to 1.28), lean dry at baseline but reach an adiabatic corner at high electricity and low water. A complementary check that perturbs the facility load by ±10%, shifting the discrete unit-count selection, flips the winner at three of these boundary cases (Reno 2%, Salt Lake City 2%, and Denver 1%), confirming they are genuine close calls rather than firm assignments.

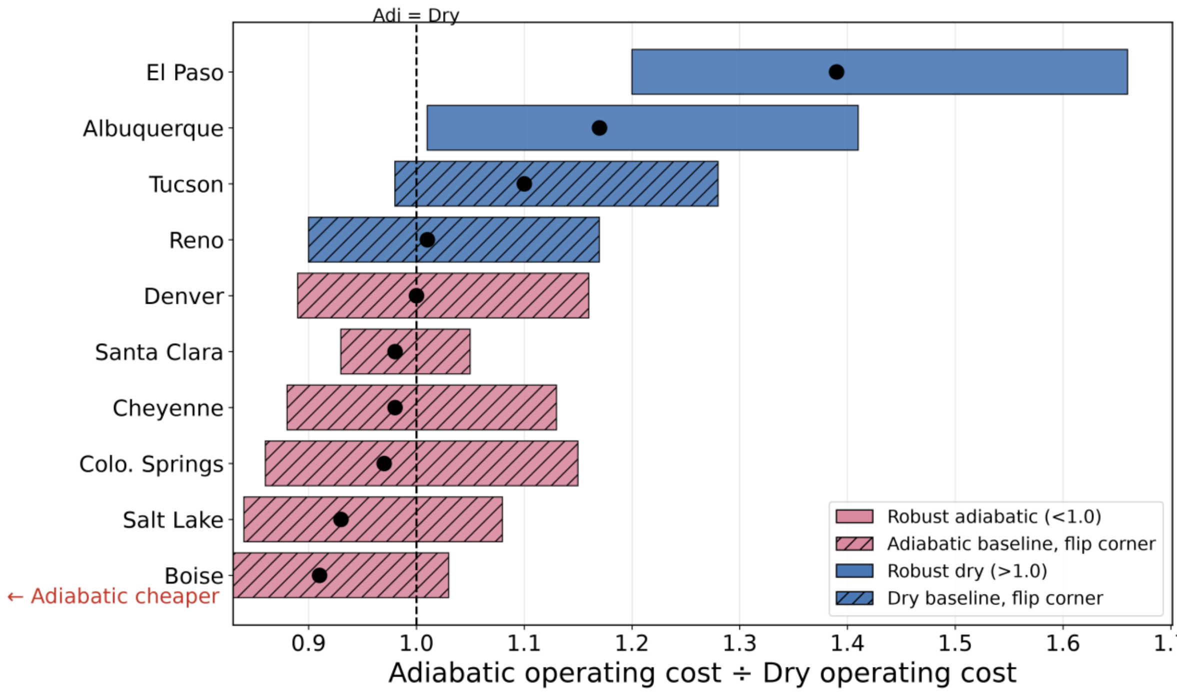

Fig. 7. Cost-ratio sensitivity across ±20 % variation in electricity and water tariffs. Bars span the range across four corner cases at three design conditions; dots indicate the baseline ratio. Solid colour indicates the cost ranking is robust across the full envelope; hatched colour indicates a flip corner exists where the ranking inverts at a combined tariff extreme (low-electricity / high-water for the adiabatic-leaning sites, high-electricity / low-water for the dry-leaning sites (Reno, Tucson).

The practical reading of Fig. 7 is that, under the calibrated model, the cost ranking turns almost entirely on the electricity-to-water tariff ratio rather than on climate. With the smaller electricity savings of continuous-spray operation, only the extreme-tariff sites are robust: El Paso and Albuquerque, where electricity is cheap relative to water, are firmly dry-favored, while the high-tariff Mountain West and high-plains sites lean adiabatic but flip to dry whenever electricity is cheap and water is dear. The water-tariff sensitivity dominates the electricity-tariff sensitivity in absolute dollar terms because adiabatic systems consume water at every wet-mode operating hour, while the electricity savings are bounded by the dry-mode annual energy. This is why nearly every flip sits at the low-electricity / high-water corner, and it is the strongest evidence in the dataset that the tariff ratio, not climate alone, should drive the siting decision.

5. Discussion

5.1. What the trade ratio means

The trade ratio is a useful rule of thumb for operators. It means that swapping dry for adiabatic moves on-site water, from about 8 to 56 gallons (30 to 210 liters) per kWh of cooling electricity, off the operator's electricity bill and onto the operator's water bill. Whether that trade is favorable depends on local water tariff, local water stress, grid carbon intensity, and how much electricity reduction the local climate enables. At the Mountain West and high-desert sites with high WBD, the electricity reduction is large enough that even at conservative carbon and water assumptions the trade is favorable on cost and carbon. At the lower-WBD or low-tariff end of the dataset (Santa Clara, El Paso, Tucson 2%), the trade is less dramatic and in some cases falls on the dry side.

5.2. The total-water question

A reasonable reader might ask whether the adiabatic system, with its on-site water use, ends up consuming more operational water (on-site + upstream) than a dry system once upstream grid water is counted. This comparison is operational only; embodied water in equipment manufacturing is excluded and would run opposite to the operational verdict because dry coolers typically use 2 to 4 times more equipment mass. The answer is yes, at every site in the dataset. Adiabatic uses about 8 to 56 gallons (30 to 210 liters) of on-site water per kWh of electricity saved, while the electricity saved would have used about 0.4 to 1.9 gallons (1.5 to 7.2 liters) of upstream water per kWh depending on subregion. On operational water, the ratio favors dry over adiabatic by a factor of roughly 2.3 to 20.2 across the dataset. We do not claim that adiabatic reduces operational water consumption; an embodied-water analysis at the equipment-mass level would partly offset the operational gap but is beyond the scope of this paper.

This finding does not undercut the adiabatic case; it sharpens it. The choice between technologies is not just a total-water-volume question. It is also a question of where the water comes from. On-site adiabatic water comes from municipal supply at the data center site, often in water-stressed regions. Upstream grid water comes from generating-station cooling, often far from the data center and frequently from sources (large hydroelectric reservoirs, once-through thermal plants on major rivers) where the marginal water cost is structurally lower than municipal supply. Local water stress matters more for siting decisions than national water-mass-balance arithmetic. The honest framing is that adiabatic shifts the water burden rather than reducing it. Whether that shift is acceptable depends on local water permitting, the carbon value of the avoided electricity, and the tariff economics that determine whether the trade remains favorable at the operator's actual bill.

5.3. What is actually better for local communities?

The intuitive reading of the trade ratio is that dry cooling is the community-friendly choice: zero on-site water, no draw from local supply. This view treats the data center fence as the boundary of community impact. It is incomplete in two ways.

First, the additional electricity demanded by dry cooling is not free to the community. It is drawn from a grid the community shares. Recent reporting documents wholesale electricity prices rising by up to 267% over five years in areas of the U.S. with significant data center activity, with the costs flowing to residential ratepayers through capacity charges and grid upgrades [25]. A Union of Concerned Scientists analysis tallied $4.3 billion in 2024 grid-interconnection costs across seven states within PJM Interconnection (the regional transmission organization coordinating wholesale electricity across the Mid-Atlantic and parts of the Midwest) that were passed to ratepayers rather than absorbed by data center operators [26]. A 10 MW data center on dry cooling adds on the order of 3 MW of peak fan demand relative to adiabatic (the approximate magnitude noted in Section 3.5, which sits outside the annual-average cost model). Across a region hosting multiple facilities, that additional load contributes to the utility capacity additions that drive residential bills upward.

Second, the upstream water consumed at the generating station is rarely paid by the data center's host community in any visible way, but it often draws from the same regional watersheds. A thermoelectric plant 200 miles away consuming roughly 0.7 gal/kWh (2.6 L/kWh; mid-range for the subregions studied) from a shared river basin affects the same water table that the community depends on. The "no on-site water" framing obscures this by drawing the system boundary at the facility fence.

What is best for a host community therefore depends on three factors: whether the local water source is scarce or abundant; whether the regional grid draws water from the same watersheds the community uses; and how much of the data center's electricity demand will translate into utility capacity additions that ratepayers absorb. In water-abundant regions served by water-intensive grids, most of the Mountain West outside the Colorado basin, adiabatic cooling is the better community choice by every measure available to this analysis. In extremely water-stressed regions where dry cooling is feasible (El Paso in the present dataset) and where the grid is comparatively clean, the community case for dry cooling strengthens, and the operating-cost result in Section 4.4 happens to align with this. The trade ratio is necessary information for this decision but is not sufficient on its own; the broader community impact requires weighing both ends of the water-energy nexus.

5.4. Where dry cooling is infeasible

Phoenix, Las Vegas, and Tucson's peak design point sit outside the framework of this paper because dry cooling is not feasible at their hottest design conditions at the W3 class target FLT. This is a meaningful point for operators evaluating these markets. The hyperscale buildout in Phoenix, Las Vegas, and the broader Southwest region is real and accelerating. All three sites sit in water-stressed regions, with Phoenix and Tucson on the WRI Aqueduct 4.0 "Extremely High" baseline classification. Las Vegas is a notable exception: the WRI index rates it “Low” because it withdraws little water locally, yet it depends almost entirely on the over-allocated Colorado River, a scarcity expressed in policy rather than in the local-basin metric. Nevada has prohibited Colorado River water for ornamental turf [27], and the Las Vegas Valley has restricted evaporative cooling in new commercial and industrial construction [14, 28]. The logical question becomes: if dry cooling is the path forward for sites where water is constrained, but dry cooling is not feasible at the hottest sites, what is the answer? The answer in these markets is one of three options. Operators can accept higher FLT targets (W4 at 113 °F (45 °C), W+ at 122 °F (50 °C)), which requires IT equipment qualified for higher temperatures. They can hybridize dry cooling with mechanical refrigeration for peak days. Or they can use evaporative or adiabatic systems, accepting on-site water consumption. Tucson is particularly informative because its 0.4% design point is infeasible but its 1% and 2% points are feasible: the operational envelope of a single site can straddle the dry-feasibility boundary, requiring either a hybrid design or W3 operation only under non-peak conditions.

A functional restriction also raises a policy-design question. At Las Vegas the 0.4% design dry-bulb of 109.2 °F (42.9 °C) exceeds the 107 °F (41.7 °C) target FLT, so fully-dry cooling cannot meet the peak; and because open-tower water-cooled chillers are themselves evaporative, the compliant fallback is the air-cooled chiller, which is materially less efficient than either the foreclosed evaporative option or dry and adiabatic cooling (full-load ratings of roughly 1.0 to 1.4 kW/ton, versus 0.5 to 0.65 for water-cooled chillers and 0.05 to 0.35 for dry and adiabatic coolers [29]), with the penalty concentrated at the annual grid peak. A restriction defined by function thus treats a hybrid dry-adiabatic cooler, which runs dry for most of the year and consumes far less water than a continuously-wet tower, identically to an open cooling tower, foreclosing a lower-water, lower-peak-power intermediate. A water-use cap that admits hybrid adiabatic while still excluding continuously-wet systems could capture most of the intended water savings at a smaller energy and carbon penalty.

5.5. A decision framework for siting

Four location attributes drive the right cooling-technology answer: ambient design conditions (which set the electricity savings achievable and dry feasibility), grid carbon intensity (which sets the carbon value of avoided electricity), local water stress (which sets the qualitative cost of on-site water consumption), and the tariff ratio of electricity price to water price (which sets the quantitative cost ranking). The first three are commonly recognized; the fourth is the underappreciated driver. Fig. 8 shows the twelve study locations against the first two attributes; it shows the carbon-intensive Mountain West and high-plains hubs clustering in the adiabatic-preferred region, the cleanest-grid site (Santa Clara) sitting in the low-carbon corner, and the infeasibility sites at opposite ends of the water-stress axis (Phoenix and Tucson at the high extreme, Las Vegas at the low end for the local-basin reasons discussed in Section 5.4). El Paso illustrates the fourth attribute. It sits in the adiabatic-preferred region of Fig. 8 (high WBD, moderate-carbon grid, extremely high water stress), all of which would, on climate alone, point toward adiabatic. But El Paso's realized electricity savings under the calibrated model are modest (a 15% location average, near zero at the 1% sizing optimum), and its commercial electricity tariff ($86.90/MWh, Texas state average) is the cheapest in the dataset, so the MWh savings adiabatic delivers do not earn back the water tariff cost. Both the weaker savings and the low tariff push the ranking to dry cooling.

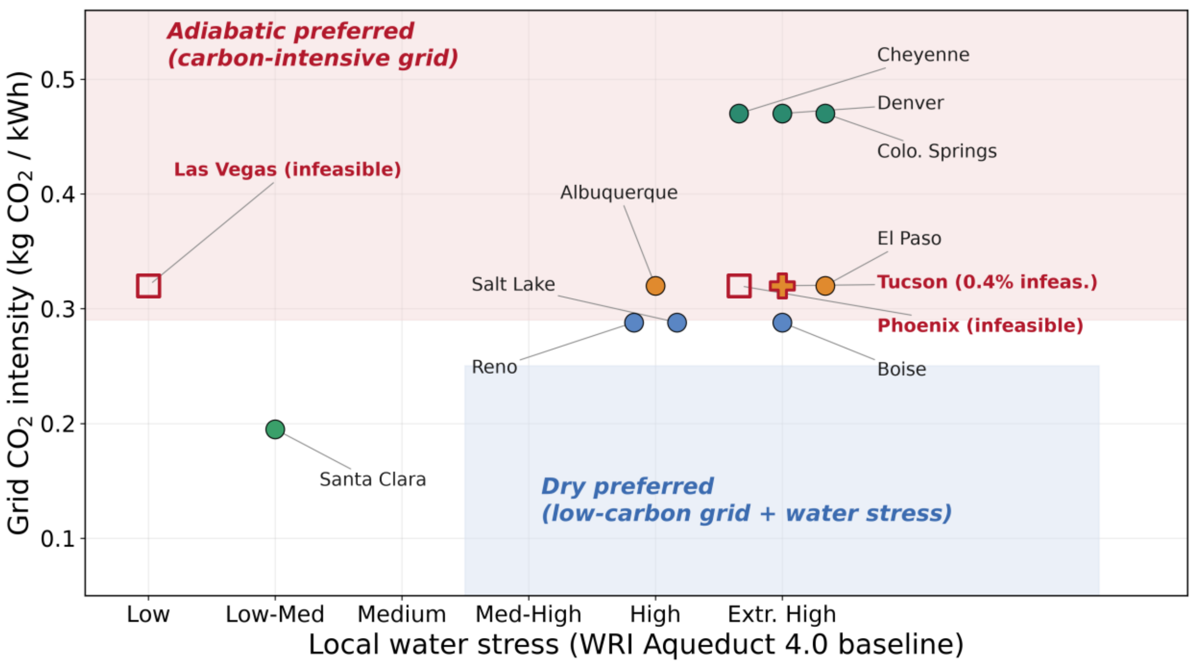

Fig. 8. Cooling technology selection framework. Vertical axis is grid CO_2 intensity from eGRID2023 Revision 2; horizontal axis is local water stress from WRI Aqueduct 4.0. The shaded regions indicate qualitative technology preference; the tariff-ratio dimension (Section 4.6) overlays this framework. Phoenix and Las Vegas (open red squares) and Tucson at 0.4 % design (red plus) sit outside the framework because dry cooling is thermodynamically infeasible at those design conditions. Water stress is the WRI Aqueduct 4.0 baseline indicator (local-basin withdrawal-to-supply); because it reflects local basins, it can understate stress in import-dependent hubs such as Las Vegas, which scores "Low" despite its dependence on the over-allocated Colorado River (Section 5.4).

Four practical recommendations follow. First, at high-WBD sites in the Mountain West and high-plains with carbon-intensive grids, adiabatic is favored on cost at baseline tariffs, but the margin is thin and tariff-sensitive: under the calibrated model only El Paso and Albuquerque are robustly dry-favored, while every other feasible location can flip within a ±20% tariff swing (Section 4.6), so operators should evaluate their own electricity-to-water tariff ratio rather than rely on climate alone. Second, at low-carbon-grid sites (Santa Clara) the trade-off is narrow but adiabatic wins on cost at baseline tariffs. This framing of dry cooling as the lower-water-impact choice ignores upstream water and grid load and therefore does not engage with the full water-energy nexus. Third, at sites where electricity tariffs are low relative to water tariffs (El Paso, Albuquerque) or where the dry cooler's sizing leaves only a small MWh saving (Tucson, Reno), dry cooling is the better choice on cost despite high WBD. Fourth, at Phoenix, Las Vegas, and Tucson's peak design conditions, dry cooling is not an option at the W3 target FLT. Operators in those markets need either a higher FLT target, a hybrid system with mechanical cooling, or, where local water policy permits, an evaporative or adiabatic system.

5.6. Comparison with prior work

Karimi et al. [15] reported a directionally similar trade-off at Phoenix, with higher PUE but lower source water consumption for air-cooled relative to evaporatively cooled facilities. The present analysis extends that single-site finding to a 12-location dataset at standardized design conditions and adds the Mountain West, high-plains, and border hot-arid conditions. Lei and Masanet [5] developed a framework that could generate similar location-explicit estimates but does not provide them; this paper, however, fulfills that need. The Phoenix, Las Vegas, and Tucson 0.4% infeasibility finding is consistent with operational evidence in those markets, where most large facilities use evaporative cooling. Companion work (in preparation) covers humid temperate and continental climates, where the technology ranking is more nuanced and where dry cooling is genuinely competitive.

6. Conclusions

Across twelve U.S. data center hubs evaluated at three ASHRAE design conditions each, adiabatic cooling reduced annual cooling electricity by a median of 30.5% relative to dry cooling across 29 feasible cases. Reductions reached 41% at Colorado Springs' 1% design point. The trade-off ranges from about 8 to 56 gallons (30 to 210 liters) of on-site water per kWh of electricity saved, depending on location. Scope 2 carbon savings range from approximately 6 tCO_2e per year (Santa Clara 0.4%, the cleanest-grid case) to 139 tCO_2e per year (Tucson 1%, the highest dry-cooler-energy case on a moderate-carbon grid). Operational water footprint, counting both on-site and upstream water (embodied water in equipment is excluded), favored dry cooling at every site in the dataset where dry cooling is feasible; adiabatic cooling does not reduce operational water use but shifts the burden from grid to data center site. On an energy-and-water operating cost basis, the ranking is set by the electricity-to-water tariff ratio more than by climate alone; six of ten feasible locations fell on the adiabatic-favored side of this division, and the four dry-favorable sites (El Paso, Tucson, Albuquerque, and Reno) define the edge of the climate-and-tariff envelope where the adiabatic advantage breaks down.

At Phoenix and Las Vegas across all three design conditions, and at Tucson's 0.4% peak design point, dry cooling is infeasible at the 107 °F (41.7 °C) target FLT used in this study. Operators in those markets must choose between higher FLT targets, hybrid systems with mechanical cooling, or, where local water policy permits, evaporative and adiabatic systems; in southern Nevada, where evaporative cooling is restricted in new construction, the practical compliant options are higher-temperature or air-cooled mechanical designs. The result is that dry cooling, despite being positioned as the lower-water-impact choice for water-stressed markets, is not actually available as a primary technology at the hottest U.S. data center hubs.

We close with a siting decision framework combining ambient design conditions, grid carbon intensity, local water stress, and the electricity-to-water tariff ratio. The framework is directly applicable to current state-level data center moratorium debates. These results should be read as a siting guideline for climates where adiabatic cooling is thermodynamically advantaged, not as a general verdict against dry cooling. Future work should extend the analysis to embodied impacts of cooling equipment, hourly marginal grid emissions, projected grid decarbonization pathways, and the tighter FLT targets (ASHRAE W2 and W1) demanded by some emerging cooling architectures. Companion work (in preparation) addresses humid temperate and continental U.S. climates.

CRediT authorship contribution statement

Vamsi Mokkapati: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review and editing, Visualization. Debendra K. Das: Writing – review and editing, Supervision.

Declaration of competing interest

The lead author (V. Mokkapati) is employed in the Research and Development Department of Nimbus Advanced Process Cooling, a manufacturer of both dry and adiabatic cooling systems for industrial and data center applications. The dual product portfolio is relevant to disclose: the company has no commercial preference between the two technologies compared in this paper, and the analysis was conducted independently within the lead author's research and development role. The methodology, data sources, location selection criteria, and analytical conclusions reflect the engineering data and reference framework cited; the findings are not directed or endorsed by the company's sales or marketing functions.

Where the analysis identifies contexts in which dry cooling is the appropriate alternative (El Paso, Tucson 2%, and all infeasibility cases), this is reported alongside the adiabatic-favorable results without selective omission. The co-author (D. K. Das) is Emeritus Professor at the University of Alaska Fairbanks and has no financial or commercial relationship with Nimbus Advanced Process Cooling; his contribution was limited to supervisory guidance, manuscript review, and editorial input. The authors affirm that the technical conclusions stand on their own merits.

Acknowledgments

The underlying engineering tools, performance correlations, and equipment selection methodology used to produce this work are the intellectual property of Nimbus Advanced Process Cooling. The authors hold no intellectual property interest in this work.

The authors thank Nimbus Advanced Process Cooling for use of these internal engineering tools, equipment performance data, and analytical infrastructure that supported this work.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability

The complete results dataset for all 36 combinations of the twelve locations and three design conditions, including raw cooling tool output, environmental factors, tariff data, and the Excel calculation workbook used for the analysis, is provided as Supplementary Material accompanying this paper. All results were generated with the Nimbus native sizing tool, version 4.2 (run 18 June 2026), under the continuous-spray, fan-modulating adiabatic operating mode, and the dataset is frozen at this version. Equipment model codes, the proprietary Dry Switch Point correlation, and wet-mode entering-air temperatures are omitted from the Supplementary Material to protect intellectual property; all performance outputs (annual energy, water, unit counts, and the derived indicators) are reported in full.

Declaration of generative AI in the writing process

During the preparation of this work, the authors used Anthropic’s Claude to assist with language editing and to help draft and refine authors’ own wording; no AI was used to generate or alter data or image content. The authors reviewed, revised, and verified all output, and take full responsibility for the content of the publication.

References

[1] JLL. North America Data Center Report Year-End 2025. Jones Lang LaSalle, Chicago, IL; 2026.

[2] MSCI Inc. When AI meets water scarcity: data centers in a thirsty world. MSCI Sustainability and Climate Research, New York, NY; October 2025.

[3] Environmental Law Institute (ELI). Data centers and water fact sheet. ELI Technical Report, Washington, DC; 2026.

[4] Shehabi A, Smith SJ, Hubbard A, et al. 2024 United States data center energy usage report. Lawrence Berkeley National Laboratory, Berkeley, CA, LBNL-2001637; 2024.

[5] Lei N, Masanet E. Climate- and technology-specific PUE and WUE estimations for U.S. data centers using a hybrid statistical and thermodynamics-based approach. Resour Conserv Recycl 2022;182:106323.

[6] Vertiv. Variable speed drives vs EC fans: cooling system energy efficiency. Vertiv Technical Brief; 2017.

[7] ASHRAE. Handbook of Fundamentals, Chapter 14: Climatic Design Information. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Atlanta, GA; 2025.

[8] Evapco. Adiabatic fluid coolers and refrigerant condensers: impact of pad saturation efficiency on performance. Evapco Technical White Paper; 2018.

[9] Shehabi A, Smith S, Sartor D, et al. United States data center energy usage report. Lawrence Berkeley National Laboratory, Berkeley, CA, LBNL-1005775; 2016.

[10] Mytton D. Data centre water consumption. npj Clean Water 2021;4(1):11.

[11] Siddik MAB, Shehabi A, Marston L. The environmental footprint of data centers in the United States. Environ Res Lett 2021;16(6):064017.

[12] American Water Works Association (AWWA). Cooling the Cloud: Water Utilities in a Data-Driven World. AWWA Technical Report, Denver, CO; 2025.

[13] Environmental and Energy Study Institute (EESI). Data centers and water consumption. EESI Issue Brief, Washington, DC; 2025.

[14] S&P Global. Beneath the surface: water stress in data centers. S&P Global Sustainability Insights, New York, NY; 2025.

[15] Karimi L, Yacuel L, Johnson JD, et al. Water-energy tradeoffs in data centers: a case study in hot-arid climates. Resour Conserv Recycl 2022;181:106194.

[16] Lei N, Lu M, Shehabi A, et al. Determinants of data center water consumption. Resour Conserv Recycl 2025;219:108310.

[17] Li H, et al. Using life cycle assessment to drive innovation for sustainable cool clouds. Nature 2025;641:234-245.

[18] Cejudo CE, Pamintuan BC, Stoughton K, et al. Emerging technologies review: water reuse systems for cooling tower applications. Pacific Northwest National Laboratory, Richland, WA, PNNL-34788; 2023.

[19] Macknick J, Newmark R, Heath G, et al. Operational water consumption and withdrawal factors for electricity generating technologies: a review of existing literature. Environ Res Lett 2012;7(4):045802.

[20] U.S. Environmental Protection Agency (EPA). eGRID2023 Summary Tables, Revision 2. U.S. EPA, Office of Atmospheric Programs, Washington, DC; 2025.

[21] ebm-papst. Embracing EC fan technology in data center cooling. Technical Brief, Data Center Frontier; 2025.

[22] Air-Conditioning, Heating, and Refrigeration Institute (AHRI). Standard 410: Performance Rating of Forced-Circulation Air-Cooling and Air-Heating Coils. Arlington, VA; 2014.

[23] U.S. Energy Information Administration. Electric Power Monthly, Table 5.6.A, Washington, DC; January 2026.

[24] World Resources Institute (WRI). Aqueduct 4.0 Water Risk Atlas. WRI Technical Note, Washington, DC; 2023.

[25] Bloomberg. How AI Data Centers Are Sending Your Power Bill Soaring. Bloomberg News, New York, NY; 2025.

[26] Union of Concerned Scientists. Data Centers Are Already Increasing Your Energy Bills. UCS Analysis, Cambridge, MA; 2025.

[27] Nevada Legislature. Assembly Bill 356 (irrigation of nonfunctional turf with Colorado River water). 81st Session, Carson City, NV; 2021.

[28] Las Vegas Valley Water District. Service Rules, Section 2.0 (Definitions) and Section 3. Effective January 1, 2026.

[29] U.S. Department of Energy, Federal Energy Management Program. Purchasing Energy-Efficient Electric Chillers, Washington, DC; 2023.