# Ranking Low-GWP Refrigerants for Data-Center Chillers Across Climate Zones and a Decarbonizing Grid: A Total Equivalent Warming Impact Analysis

**Authors:** Vamsi Mokkapati [1,2,*] and Sunwoo Kim [1]

[1] Department of Mechanical Engineering, University of Alaska Fairbanks, Fairbanks, AK 99775, USA
[2] Nimbus Advanced Process Cooling, Anniston, AL, USA

*** Correspondence:** vsmokkapati@alaska.edu

**ORCID:** Vamsi Mokkapati, 0009-0009-3529-9149; Sunwoo Kim, 0000-0002-8417-0693

**Date:** 2026-09-23 · **Preprint posted:** 2026-09-23

---

**Preprint: not peer-reviewed.** This manuscript is intended for submission to the *Journal of Building Engineering* (Elsevier) and is posted as a preprint for early dissemination and community feedback. This version has **not** undergone peer review and may differ from later versions; if a peer-reviewed version of record is published, please cite that in preference to this preprint.

**Cite as:** V. Mokkapati, S. Kim, Ranking low-GWP refrigerants for data-center chillers across climate zones and a decarbonizing grid: a total equivalent warming impact analysis, engrXiv preprint (2026).



---

## Highlights

- Eight low-GWP data-center chiller refrigerants ranked by lifetime TEWI.
- Compliant-fluid TEWI margins are within a few percent, too close to decide alone.
- Condenser architecture and control, not the refrigerant, drive lifetime impact.
- A time-varying decarbonizing grid is embedded in the TEWI calculation.
- Cooling-tower water tracks the IT load, not climate, unlike plant electricity.

## Abstract

From 1 January 2027 the American Innovation and Manufacturing Act caps refrigerant global-warming potential (GWP) at 700 in new data-center cooling equipment, forcing a choice among several low-GWP chiller refrigerants. Existing Total Equivalent Warming Impact (TEWI) comparisons cannot settle it: they use a single operating point or a generic comfort-cooling profile, and a static grid emission factor. A data-center chiller instead runs at near-constant load year-round, rejects heat under climate-dependent conditions, and spends 15 to 25 years on a decarbonizing grid. This paper develops a transparent TEWI framework closing all three gaps. Real-fluid coefficients of performance for eight refrigerants, from CoolProp, are coupled to two generic heat-rejection archetypes bracketing practice, a wet-bulb-driven water-cooled condenser and a dry-bulb-driven air-cooled condenser, over annual climate bins for five U.S. stations and three control strategies, with the indirect term summed year by year along regional grid decarbonization trajectories rather than held constant. Two findings follow. The compliant low-GWP fluids differ by only a few percent in lifetime TEWI, within the model's uncertainty, so the fluid is better selected on safety, charge limits, and cost; single-stage transcritical R-744 the one clear exception, 19 to 81 % worse depending on climate and control. What does move lifetime impact is the heat-rejection architecture and its control, which alter plant energy by up to 57 % and leave tower water largely unmoved. Built from open thermophysical, climate, and grid data, the framework lets operators select refrigerant, architecture, and control jointly, not in isolation.

**Keywords:** data center cooling; low-GWP refrigerants; Total Equivalent Warming Impact (TEWI); chillers; grid decarbonization; waterside economizer

## Nomenclature

**Roman and Greek symbols** (alphabetical; Greek follows Roman)

| Symbol | Meaning | Units |
|---|---|---|
| COP | Coefficient of performance (cooling) | - |
| DB | Dry-bulb temperature | °C (°F) |
| $E$ | Annual electricity use | kWh yr$^{-1}$ |
| GWP | Global-warming potential (100-yr basis) | - |
| $L$ | Annual refrigerant leakage fraction | yr$^{-1}$ |
| $m$ | Refrigerant charge | kg |
| MCWB | Mean coincident wet-bulb temperature | °C (°F) |
| $n$ | System (chiller) lifetime | yr |
| $Q$ | Cooling load | kW |
| $t$ | Annual operating hours | h yr$^{-1}$ |
| $T$ | Temperature | °C (°F) |
| $y$ | Year index over the system lifetime | - |

| Symbol | Meaning | Units |
|---|---|---|
| $\alpha_{\mathrm{rec}}$ | End-of-life refrigerant recovery fraction | - |
| $\beta$ | Grid carbon intensity | kg $CO_2$e $kWh^{-1}$ |
| $\Delta T_{\mathrm{rej}}$ | Total approach, driving temperature to condensing temperature | K |
| $\eta_{\mathrm{is}}$ | Compressor isentropic efficiency | - |

**Subscripts**

| Subscript | Meaning |
|---|---|
| ann | Annual value |
| aux | Heat-rejection auxiliaries (fans and pumps) |
| $b$ | Weather (temperature) bin |
| cond | Condenser / condensing |
| drv | Ambient driving temperature (wet-bulb, water-cooled; dry-bulb, air-cooled) |
| min | Minimum permitted value (head-pressure floor) |
| rec | End-of-life recovery |
| rej | Heat rejection |
| $y$ | Year within the system lifetime |

**Abbreviations.** AIM, American Innovation and Manufacturing (Act); ASHRAE, American Society of Heating, Refrigerating and Air-Conditioning Engineers; DB, dry-bulb; HFC, hydrofluorocarbon; HFO, hydrofluoroolefin; IEA, International Energy Agency; LCCP, Life Cycle Climate Performance; MCWB, mean coincident wet-bulb; PUE, power usage effectiveness; TEWI, Total Equivalent Warming Impact; TMY, typical meteorological year; WUE, water usage effectiveness.

## 1. Introduction

Electricity use by data centers is rising steeply. The International Energy Agency projects that global data-center demand will roughly double from about 415 TWh in 2024 to about 945 TWh by 2030, approaching 3 % of world electricity, with artificial-intelligence workloads the fastest-growing component [1]. Cooling is the largest energy draw after the IT load itself

[2,3], so the design and operation of the cooling plant has a first-order effect on both the operating cost and the carbon footprint of the sector, and vapor-compression chilled-water plant remains the workhorse at hyperscale (Section 2.1).

A regulatory constraint now intersects this growth. Under the American Innovation and Manufacturing (AIM) Act, the U.S. Environmental Protection Agency's Technology Transitions Rule caps refrigerant global-warming potential (GWP) at 700 for new data-center cooling equipment from 1 January 2027 [4]. That removes the long-dominant chiller refrigerant R-134a and its higher-GWP relatives from new equipment and forces a choice among lower-GWP fluids, none of which is dominant on every axis: each trades GWP against efficiency, capacity, flammability class, glide, and equipment compatibility (Section 2.3).

Comparing such fluids on climate grounds requires a lifetime, system-level metric rather than GWP alone. Total Equivalent Warming Impact (TEWI) sums the direct emissions from refrigerant leakage and end-of-life loss with the indirect emissions from lifetime electricity use; Life Cycle Climate Performance (LCCP) adds embodied terms [5]. Applied to low-GWP chiller refrigerants, this literature consistently finds the indirect term dominant on today's fossil-heavy grids, so the most efficient fluid rather than the lowest-GWP one tends to minimize total warming impact. That balance is not fixed: as the supplying grid decarbonizes, the indirect term shrinks and the GWP-dependent direct term gains relative weight (Section 2.4).

**The gap.** That shift is where the existing literature stops short. Published TEWI/LCCP refrigerant comparisons for chillers are built on either a single nominal operating point or a generic comfort-cooling/building load profile, and they evaluate against a small set of static power sources (e.g., "coal," "natural gas," "wind") held constant for the equipment lifetime [6,7,8]. None of three things that are decisive for a data-center chiller is represented together: (i) the data-center-specific duty cycle (near-constant high load, ~8 760 h/yr operation, and the elevated chilled-water temperatures now permitted by ASHRAE TC 9.9 thermal guidelines [9]); (ii) the climate-zone dependence of condenser conditions and economizer (free-cooling) hours, which strongly modulate both annual energy use and the share of the year the compressor even runs [10,11]; and (iii) the fact that the grid is decarbonizing over time, so the carbon intensity multiplying the indirect term is not a constant but a declining trajectory across the 15–25 year chiller life [12]. Because the relative weight of the direct and indirect terms shifts as the grid cleans up [7], and because the heat-rejection architecture, whether the condenser is water-cooled or air-cooled, sets both the chiller's condensing temperature and, for evaporative architectures, a climate-dependent water footprint, the lifetime impact of a data-center chiller cannot be read off a GWP table or a first-year COP comparison. It must be computed for the specific intersection of duty cycle, climate, heat-rejection architecture and its control strategy, and grid pathway. To the authors' knowledge, no published TEWI study does this for data-center chillers across climate-resolved heat-rejection architectures and a decarbonizing grid.

**Contribution and novelty.** This paper (1) states that gap; (2) builds a transparent TEWI model for AIM-Act-compliant chiller refrigerants that, to the authors' knowledge, couples for the first time two generic heat-rejection archetypes (a water-cooled and an air-cooled condenser) with a time-varying grid trajectory inside the standard equations; (3) evaluates eight fluids across five real climate stations, two condenser architectures, three control strategies (fixed, floating, and economized), and three eGRID-anchored grid pathways, yielding three results: the compliant refrigerant choice is effectively a tie among a leading low-GWP cluster under both architectures, the climate-dependent levers are the condenser architecture and its control (chiefly the integrated economizer) rather than the refrigerant, and the architecture choice is itself a water-versus-energy trade-off; and (4) distributes the full computational pipeline. The novelty is the integration: bridging the established TEWI methodology to data-center duty, climate-resolved heat-rejection architecture, grid trajectory, and the water-energy trade-off in one open model.

**Scope and inputs.** This work combines a comparative analysis of existing methods, fluids, and published results with the extension of an established methodology (TEWI/LCCP) to a new application domain; it requires no new experiments or primary measurements. All thermophysical inputs come from open property libraries and published COP results [6,7,13]; all grid and climate inputs come from published third-party sources, EPA eGRID2022 [22] and ASHRAE Weather Data Viewer 2025 [23]; the TEWI/LCCP equations are standard [5].

The remainder of the paper synthesizes the relevant literature (Section 2), specifies the framework and its inputs (Section 3), presents the computed results (Section 4), discusses implications and threats to validity (Section 5), and concludes with future work (Section 6).

---

## 2. Background and Related Work

### 2.1 Data-center cooling and its energy weight

Reviews of data-center cooling consistently identify the cooling plant as the dominant non-IT energy consumer and survey the progression from air cooling toward liquid, free, and two-phase strategies [2,3]. Cai and Gou [2] provide a comprehensive review of passive and active strategies, comparing technologies on energy, economic, and environmental indicators, and report representative power-usage-effectiveness (PUE) ranges of roughly 1.4–1.5 for air cooling and 1.1–1.2 for liquid cooling. Nadjahi et al. [3] earlier reviewed thermal-management and innovative cooling strategies, emphasizing the role of free cooling and the difficulty of comparing options on a consistent basis; for the facility stock of that period they report an energy split of 52 % IT equipment, 38 % cooling, and 10 % power distribution and UPS losses. Efficiency has improved since, as the PUE figures above indicate, but on either accounting the cooling plant is the largest energy consumer after the IT load itself, which is what gives refrigerant and heat-rejection choices their leverage. Despite the recent surge of interest in direct-to-chip and immersion cooling, vapor-compression chilled-water plant, often with economizers, remains the most widely deployed heat-rejection system for large facilities, which is why the refrigerant question matters at scale [2].

### 2.2 The regulatory driver: AIM Act GWP limits

The EPA Technology Transitions Rule under the AIM Act establishes sector-specific GWP caps and phase-down dates [4]. For data centers, computer-room air conditioning, and IT-equipment cooling, the GWP cap is 700 with a compliance date of 1 January 2027; comfort-cooling chillers face the same 700 cap from 1 January 2025 [4]. (EPA opened a reconsideration of parts of the Technology Transitions Rule in late 2025; the data-center GWP-700/2027 provision was not among the dates it proposed to change, but the analysis here is conditional on the rule as finalized [4].) The rule therefore makes refrigerant selection an immediate, non-optional engineering decision for new data-center cooling equipment and motivates a rigorous, lifetime-oriented basis for that selection.

### 2.3 Candidate low-GWP refrigerants for chillers

A substantial literature characterizes the leading R-134a replacements. Mateu-Royo et al. [8] performed a theoretical energy-and-carbon analysis of R-1234ze(E) and the non-flammable blend R-515B as low-GWP replacements for R-134a, finding comparable COP to R-134a but lower volumetric capacity, and noting R-515B's A1 (non-flammable) safety advantage. Mansouri, Bisulandu, and Ilinca [6] evaluated R-1234ze(E) against R-134a, R-407C, R-410A, and R-22 in air-cooled vapor-compression chilled-water systems across a range of ambient and evaporating temperatures, reporting that R-1234ze(E) achieved the highest COP and lowest annual TEWI of the fluids studied. These and similar studies establish that the low-GWP candidates are thermodynamically viable, but they evaluate generic chilled-water/AC duty rather than data-center duty, and they do not resolve climate or grid-trajectory effects.

### 2.4 TEWI and LCCP as comparison metrics

TEWI and LCCP are the accepted system-level climate metrics for refrigeration equipment. TEWI sums a direct term (refrigerant charge × GWP × leakage and recovery losses) and an indirect term (lifetime energy use × grid carbon intensity); LCCP additionally includes embodied emissions from refrigerant manufacture and equipment materials [5]. Akinseloyin et al. [7] applied the LCCP framework to an 809 kW water-cooled centrifugal chiller for six R-134a alternatives (R-450A, R-513A, R-516A, R-1234ze(E), R-515A, R-515B), using a thermodynamic chiller model to obtain compressor power and evaluating five static power sources over a 27-year life. They reported that fossil-powered operation made indirect emissions dominate (>90 %), so refrigerant GWP was a minor contributor; but under wind power the GWP-dependent leakage term rose to as much as 74 % of total emissions. Their load profiles were a university campus and a naval-fleet profile (not a data center), and the power sources were fixed in time rather than following a decarbonization trajectory. This is the methodological seam that the present work extends.

### 2.5 Climate-zone and economizer effects

Climate strongly governs chiller energy use and the share of the year that mechanical cooling can be displaced by economization. Studies that resolve climate zones report large differences in achievable free-cooling: water-side economization can deliver substantial energy savings that vary sharply by climate zone and operating condition [10], and air-side economization considerably more in temperate or subtropical climates [11]. A 2024 *Applied Thermal Engineering* study mapped the climate zones favorable to water-side economizers in data-center cooling systems [10], and a 2024 *Buildings* study developed a PUE evaluation tool comparing economizer types across climate zones [11]. These works quantify the climate dependence of energy use but do not connect it to refrigerant-level TEWI/LCCP ranking. Separately, the heat-rejection architecture itself (water-cooled versus air-cooled) sets the condensing temperature and, for evaporative architectures, a climate-dependent water footprint; this is not captured by the single nominal operating point or the generic economizer-fraction inputs of prior chiller TEWI studies, motivating the two-archetype, climate-resolved approach taken here. An evaporatively-assisted *adiabatic fluid cooler* [14] interpolates between the two archetypes; a companion siting assessment by one of the present authors quantifies the resulting water-for-energy trade across twelve U.S. data-center markets [15].

### 2.6 Grid carbon intensity and its trajectory

The indirect emissions term depends on the carbon intensity of the electricity supply, which varies by region and is declining over time as grids decarbonize. The IEA Emissions Factors database provides annual average grid carbon-intensity values by country and the adjustment factors needed for consistent accounting [12]. Combining such present-day intensities with published decarbonization pathways yields the time-varying intensity needed for a lifetime analysis, an input that prior chiller TEWI/LCCP studies held constant [6,7].

### 2.7 Adjacent sustainability pathways

A parallel literature pursues waste-heat recovery from data centers, e.g., feeding district-heating networks [16] or pursuing zero-carbon operation through heat reuse [17]. These pathways reduce net emissions by displacing other heat sources and are complementary to the refrigerant-selection question addressed here; they motivate why minimizing both the energy and the direct-emission footprint of the chiller is worthwhile.

### 2.8 Property data and tools

Transparent TEWI/LCCP work requires consistent thermophysical properties. CoolProp [13] is an open-source library implementing high-accuracy Helmholtz-energy formulations for the relevant fluids and is widely used for reproducible cycle modeling; it (or the equivalent NIST REFPROP formulations) provides all property inputs needed here.

**Synthesis.** The literature establishes (a) that cooling dominates data-center non-IT energy [2,3]; (b) that the AIM Act forces a near-term low-GWP refrigerant choice [4]; (c) that TEWI/LCCP is the right comparison metric and that the direct/indirect balance flips with grid cleanliness [5,7]; and (d) that climate and grid trajectory are large, quantifiable effects [10,11,12]. What is missing is a study that puts these together for data-center chiller duty. Section 3 specifies such a study.

# 3. Proposed Approach and Methodology

## 3.1 Overview

The method is a comparative TEWI analysis that ranks candidate refrigerants for a representative data-center chiller across real climate stations and a decarbonizing grid. Every input is either (a) a thermophysical property or COP value computed from an open library, (b) a generic, public condenser model, (c) climate bin data from published ASHRAE Weather Data Viewer 2025 [23], or (d) the public EPA eGRID grid-intensity dataset [22]. The study bridges the established TEWI methodology to the data-center domain with equipment-, climate-, and time-resolved inputs. The cycle calculations use the open property library CoolProp [13], following published modeling practice [6,7].

## 3.2 System definition and candidate fluids

The reference system is a single vapor-compression chiller sized to a 1 MW cooling block, a round-number reference capacity used to normalize the comparison since TEWI scales with capacity, operating 24/7 at high load with an 18 °C chilled-water supply adopted to represent elevated-temperature data-center chilled-water operation, compatible with the wider IT inlet-temperature envelopes ASHRAE TC 9.9 permits for class A1–A4 equipment [9]. Heat is rejected by one of two generic condenser archetypes that bracket the practical envelope: a water-cooled condenser served by a cooling tower and an air-cooled condenser. Both reject the chiller's condenser duty (about 1.15 MW, i.e., the 1 MW cooling load plus compressor work) and are modelled from generic engineering correlations and published climate data (Section 3.3). The resulting system, its two heat-rejection paths, and the ambient temperature that drives each are shown in Figure 1. Candidate refrigerants, all AIM-Act-compliant except the R-134a baseline used for reference, are listed in Table 1.

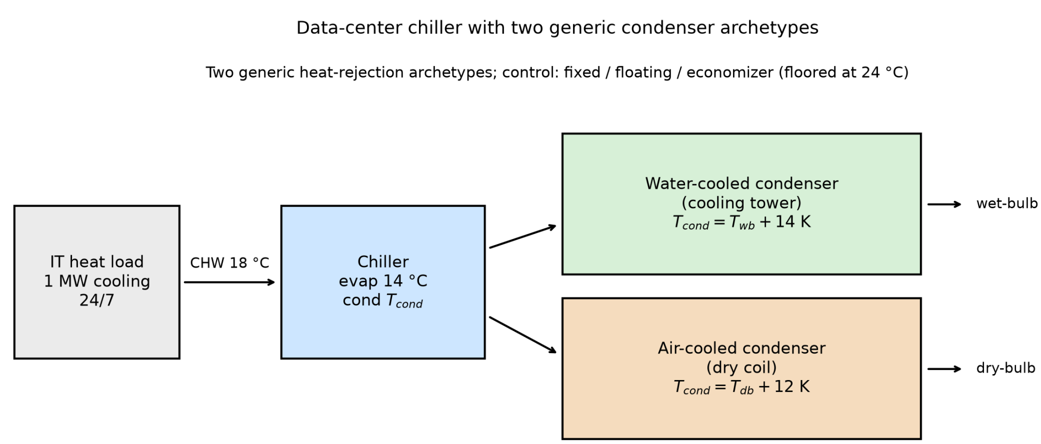


**Figure 1.** The modeled system: a data-center chiller serving a constant 1 MW cooling load, operating 24/7 and rejecting condenser heat through one of two generic heat-rejection archetypes, either a water-cooled (cooling-tower) condenser, wet-bulb-driven, or an air-

cooled condenser, dry-bulb-driven. The condensing-temperature mapping of Section 3.3 is used throughout.

**Table 1. Candidate refrigerants and key attributes (GWP-100 values from the IPCC AR4/AR5/AR6 assessment reports [25–27]; AR5 is the primary basis, with AR4 and AR6 tested as a sensitivity, Section 4.2).**

| Refrigerant | Type | Safety class | GWP-100 (approx.) | Role in study |
|---|---|---|---|---|
| R-134a | HFC | A1 | ~1300–1530 | Baseline (non-compliant; reference) |
| R-513A | HFC/HFO blend | A1 | ~573–674 | Compliant drop-in candidate |
| R-450A | HFC/HFO blend | A1 | ~547–643 | Compliant drop-in candidate |
| R-515B | HFO/HFC blend | A1 | ~293–322 | Compliant, non-flammable |
| R-1234ze(E) | HFO | A2L | ~1–6 | Ultra-low-GWP, mildly flammable |
| R-1234yf | HFO | A2L | <1–4 | Ultra-low-GWP, mildly flammable |
| R-290 (propane) | HC | A3 | ~0.02–3 | Natural, flammable |
| R-744 ($CO_2$) | Natural | A1 | 1 | Natural, transcritical considerations |

*GWP-100 ranges span the IPCC AR4–AR6 bases (basis sensitivity: Sections 4.2, 5.2). AR5 itself lists R-1234ze(E) and R-1234yf only as "<1" without a specific figure and does not tabulate a direct GWP-100 for R-290; this study uses the AR6 values, R-1234yf = 0.5, R-1234ze(E) = 1.4, and R-290 = 0.02, for both the AR5 and AR6 cases for these three fluids. R-1234ze(E) and R-1234yf were not yet in commercial use when AR4 (2007) was published and are not listed in [25]; the AR4-basis values shown for them (≈6, ≈4) are commonly cited industry estimates computed with AR4′s radiative-forcing methodology rather than values drawn directly from [25].*

### 3.3 Governing equations

**Coefficient of performance.** For each fluid and operating condition, the *cycle* COP is obtained from a standard single-stage vapor-compression model (5 K superheat, 2 K subcooling, 0.70 isentropic compressor efficiency, no internal heat exchanger; evaporating temperature 14 °C from the 18 °C chilled-water supply less a 4 K approach), using CoolProp real-fluid properties [13]. The cycle COP is an idealized compressor value; it is converted to a realistic *chiller* COP by a fixed auxiliary derate (0.90, covering motor, mechanical, and auxiliary losses) and a lift-dependent isentropic efficiency that falls from 0.70 at and above a 22 K lift toward a deep-off-design asymptote of 0.35, the design value halved, at 6 K lift. The assumed near-linear decline is an engineering approximation, informed by the fixed-impeller-speed, constant-input-head behavior of centrifugal/screw compressors reported in [18]. The 0.70 is conservative against the 85 % peak vapor-compression efficiency used for centrifugal machines of this capacity class [18], which is quoted before the mechanical and drive losses the 0.90 derate covers separately. The 24 °C head-pressure floor draws on the 65.0 °F minimum entering condenser water at which AHRI 550/590, a water-chilling-package standard, stops crediting condenser relief [19]. That entering-water temperature is connected to the condensing temperature through the condenser approach. For the air-cooled archetype, where that water-cooled mechanism does not apply, the same numerical floor is used as a simplifying, bracketing assumption. The floor never lets the cycle reach 6 K lift, since it sets a minimum realized lift of 10 K where the assumed line gives 0.44, so only 0.44 to 0.70 is exercised. This smoothly caps the low-lift COP (at most ∼11 at the 24 °C head-pressure floor). The design-point value is representative of centrifugal and screw chiller compressors at rated conditions, and the decline at low lift reflects the efficiency such machines lose as they move off-design toward their minimum pressure ratio and surge limit, where capacity control (inlet guide vanes, slide valve, or variable speed) becomes progressively less effective. The two endpoints are an engineering bound rather than a fitted compressor map, chosen so that the resulting chiller COP and its second-law efficiency stay physically plausible across the whole operating range. Both adjustments are refrigerant-independent and bring absolute COP and the control-strategy savings to realistic levels. The sensitivity of the results to both endpoints and to the shape of the curve between them is tested in Section 5.2. The resulting chiller COP corresponds to a second-law (exergetic) efficiency of 55 % at the 26 K design lift, easing to ∼38 % near the 10 K head-pressure floor, both within the realistic range for centrifugal/screw chillers and well short of the Carnot limit. For the near-azeotropic blends, the evaporating and condensing temperatures are taken at the midpoint of the temperature glide, so that a blend sees the same mean temperature difference across each heat exchanger as a pure fluid. This matters only for R-450A, whose 0.56 K glide is far larger than that of R-513A or R-515B (below 0.03 K): taking the dew point at the evaporator and the bubble point at the condenser instead would lower its COP by 2.2 % at 40 °C and 5.4 % at the 24 °C floor, a one-sided penalty that no symmetric uncertainty band can absorb (Section 5.2). Published COP values for the same fluids [6,7,8] are retained as an independent cross-check (Section 3.6).

**Condensing temperature from the condenser.** The chiller's condensing temperature in each weather bin follows the heat-rejection archetype and the ambient driving temperature

$T_{\mathrm{drv},b}$, the wet-bulb for the water-cooled (cooling-tower) condenser and the dry-bulb for the air-cooled condenser:

$$T_{\mathrm{cond},b} = \max\left(\, T_{\mathrm{drv},b} + \Delta T_{\mathrm{rej}}, T_{\mathrm{cond}}^{\mathrm{min}} \,\right), \qquad (1)$$

where $\Delta T_{\mathrm{rej}}$ is the total approach from the driving temperature to the condensing temperature: 14 K for the water-cooled archetype (tower approach + condenser-water loop + condenser approach) and 12 K for the air-cooled archetype. The head-pressure floor $T_{\mathrm{cond}}^{\mathrm{min}}$ is 24 °C (≥10 K lift). Three condenser control strategies are compared: fixed (head pressure held at the 40 °C design minimum, i.e., $T_{\mathrm{cond}} = \max(40^\circ\ \mathrm{C}, T_{\mathrm{drv}} + \Delta T_{\mathrm{rej}})$, which rises in hot weather but is not allowed to float below design); floating (the condensing temperature is allowed to track the ambient down to the 24 °C floor, lowering the lift); and floating with an integrated economizer, which switches the compressor off in any bin whose driving temperature is low enough for the condenser loop to serve the 18 °C chilled-water supply directly through a plate heat exchanger ($T_{\mathrm{drv}} \leq 53^\circ$ F for the water-cooled archetype, $\leq 50^\circ$ F for the air-cooled). Sensitivity to the head-pressure floor is reported in Section 5.2.

**Annual energy.** Annual chiller-plant electricity is the compressor work summed over all bins plus the heat-rejection auxiliaries,

$$E_{\mathrm{ann}} = \underbrace{\sum_b \frac{Q\, t_b}{\mathrm{COP}_b(T_{\mathrm{cond},b})}}_{\text{compressor}} + \underbrace{E_{\mathrm{aux}}}_{\text{tower/condenser fans + pumps}}, \qquad (2)$$

where the sum runs over the site's weather bins $b$ (annual hours $t_b$ summing to 8 760 h), $Q$ is the near-constant data-center cooling load, $\mathrm{COP}_b$ is the cycle COP at that bin's condensing temperature, and $E_{\mathrm{aux}}$ is the refrigerant-independent heat-rejection auxiliary electricity. The auxiliary is modelled generically: a fan term proportional to the heat actually rejected (0.020 kW per kW rejected for the water-cooled tower, 0.035 for the air-cooled condenser; the rejected heat is $Q$ plus compressor work, or $Q$ alone when free-cooling) plus a constant pump term (0.04 kW/kW of cooling for the water-cooled condenser-water and chilled-water pumps, 0.02 kW/kW for the air-cooled chilled-water pump only).

**TEWI.** Following the standard formulation [5],

$$\mathrm{TEWI} = \underbrace{\mathrm{GWP} \cdot m \cdot L_{\mathrm{ann}} \cdot n}_{\text{direct, leakage}} + \underbrace{\mathrm{GWP} \cdot m \cdot (1 - \alpha_{\mathrm{rec}})}_{\text{direct, end-of-life}} + \underbrace{\sum_{y=1}^{n} E_{\mathrm{ann}} \cdot \beta_y}_{\text{indirect, energy}}, \qquad (3)$$

where $m$ is refrigerant charge (kg), $L_{\mathrm{ann}}$ the annual leakage fraction, $n$ the system lifetime (yr), $\alpha_{\mathrm{rec}}$ the end-of-life recovery fraction, and $\beta_y$ the grid carbon intensity (kg $CO_2$e/kWh) in year $y$. The baseline charge $m$ is 300 kg (0.30 kg per kW of cooling) and the baseline lifetime $n$ is 20 years; this charge is held identical across all eight fluids (no per-fluid density/volumetric-capacity adjustment), a simplification immaterial to the near-zero-GWP leaders but bearing on the direct-term share of the mid-GWP blends (R-513A, R-450A, R-515B). $\beta_y$ follows a published decarbonization pathway rather than the present-day constant used in prior chiller

TEWI/LCCP studies [6,7], and the indirect term is summed year by year over the chiller life accordingly. This study reports TEWI; the embodied (manufacturing and disposal) terms that would extend it to LCCP [5] are set to zero (Sections 5.2 and 6). The direct terms carry units of (kg $CO_2$e/kg)·(kg)·(1/yr)·(yr) and the indirect term (kWh/yr)·(kg $CO_2$e/kWh) summed over years, so all three reduce to kg $CO_2$e and are additive.

**Water.** The water-cooled archetype consumes evaporative tower water; the air-cooled archetype consumes none. Makeup water is evaporation, heat rejected divided by the latent heat of vaporization, plus blowdown, evaporation divided by cycles of concentration minus one, the standard EPA cooling-tower balance also underlying data-center WUE estimates [20]. Windage (drift) leaves the tower at the same concentrated dissolved-solids level as blowdown, so it displaces blowdown in the cycles-of-concentration balance rather than adding to the makeup total. Cycles of concentration is set to 6, within the 3 to 12 range reported for a water-cooled chiller without an economizer [20]. Because heat rejected is dominated by the near-constant IT load rather than compressor work, water is reported per kWh of plant electricity (Section 4.4), not IT electricity as in WUE [21].

### 3.4 Climate-zone and grid scenarios

**Climate stations and heat rejection.** Five real locations span the relevant climate range and a major data-center market: Seattle (cool-marine), Phoenix (hot-dry), Miami (hot-humid), Chicago (continental), and Ashburn, VA (mixed-humid, the world's largest data-center hub). For each station the analysis uses ASHRAE annual bin weather data from ASHRAE Weather Data Viewer 2025 [23] for the site's primary airport weather station (WMO 727930, 722780, 722020, 725300, and 724030, respectively): dry-bulb temperature tabulated in 1 °F bins, each bin paired with its mean coincident wet-bulb and its average annual hours of occurrence over the multi-year record, summing to 8 760 h. Bin coverage spans the full annual range observed at each site (for example, -23 °F to 103 °F [−30.6 °C to 39.4 °C] at Chicago and 29 °F to 119 °F [−1.7 °C to 48.3 °C] at Phoenix). The condensing-temperature model of Section 3.3 (Eq. (1)) maps each bin's driving temperature to a condensing temperature, from which the compressor energy, the heat-rejection auxiliaries, and (for the water-cooled archetype) the evaporative water follow. This replaces the single nominal operating point and the prescribed economizer-fraction of earlier chiller TEWI studies with climate-resolved heat rejection under two bracketing architectures.

**Grid trajectories.** Each station's grid is anchored to the present-day $CO_2$e intensity of its EPA eGRID2022 subregion [22]: 0.275 kg $CO_2$e/kWh for Seattle (NWPP), 0.284 for Ashburn (SRVC), 0.354 for Phoenix (AZNM), 0.370 for Miami (FRCC), and 0.456 for Chicago (RFCW). Three forward pathways are applied to each: the present-day intensity held static, and two decarbonization scenarios declining linearly to 50 % and 15 % of it by end of life (≈50 % and ≈85 % reductions, spanning stated-policy to net-zero ambition). The Cartesian product of fluids × stations × control modes × grid pathways defines the comparison matrix.

### 3.5 Outputs and ranking

For each cell, the framework reports annual plant energy (compressor + heat-rejection auxiliaries), direct and indirect TEWI, and the direct-share fraction, then ranks fluids within each station × control-mode × grid scenario. Sensitivity is assessed over leakage rate, recovery fraction, GWP metric (AR4, AR5, and AR6), and chiller lifetime, and a Monte-Carlo propagation of COP uncertainty tests how resolved the refrigerant ranking is (Section 4.2). The deliverables are a decision matrix of the lifetime-optimal refrigerant per condition and a fixed-versus-floating-versus-economized control comparison across the two condenser archetypes.

### 3.6 Validation strategy

Two internal checks are implemented in the pipeline: (i) dimensional consistency of the TEWI terms (Section 3.3); and (ii) limiting-case behaviour (setting the annual leakage to zero and the recovery fraction to unity isolates the indirect term, while holding the grid intensity at zero isolates the direct term). The computed chiller COPs are of the expected magnitude for the elevated chilled-water duty and are consistent with published COP comparisons for these fluids [6,7,8]. Beyond these internal checks, the TEWI accounting itself is benchmarked externally by reproducing a published chiller case (Akinseloyin et al. [7]) through the same `tewi_breakdown()` code with their published per-fluid inputs (GWP, annual energy, and life-cycle parameters); it recovers that study's fluid ranking, indirect-dominant structure, and lifetime magnitude, with the only systematic departure being the embodied term that their LCCP includes and TEWI omits by construction (Section 4.6, Table 5). This benchmark exercises the TEWI summation and direct/indirect bookkeeping. It takes this paper's own COP, condenser-archetype, and climate submodels' annual-energy outputs as a given input; those submodels are supported separately by the internal checks above and the published-COP cross-check.

### 3.7 Code and data availability

The methodology is implemented in an open Python pipeline distributed with this manuscript, in two stages. `compute_tewi.py` builds the refrigerant cycle model on the open-source CoolProp library [13]: the COP of each candidate fluid as a function of condensing temperature, including the surrogate-blend handling described below and an automatic transcritical branch with high-side-pressure optimization for R-744. `compute_tewi_generic.py` is the deterministic driver: for each of the two condenser archetypes it maps every climate bin's driving temperature to a condensing temperature under each control mode (fixed, floating, economizer), evaluates the annual-energy (Eq. (2)) and TEWI (Eq. (3)) equations of Section 3.3 against the time-varying grid trajectory, and emits the plant-energy and TEWI matrices. `sensitivity_generic.py` then runs on those matrices to produce every sensitivity and uncertainty result reported in Sections 4.2, 4.5 and 5.2: the leakage/recovery/GWP-basis/lifetime sweep (420 scenarios: 14 one-at-a-time parameter values (5 leakage + 3 recovery + 3 GWP bases + 3 lifetimes, the others held at baseline) × 30 decision cells), the COP-uncertainty Monte Carlo at four band widths under both independent and common-mode error (Table 6), the isentropic-efficiency and heat-rejection-auxiliary limit reruns, the grid-trajectory shape sweep, the 45-cell economizer-inclusive Monte Carlo, the free-cooling threshold sweep, the blend saturation-convention rerun, and the head-pressure-floor sweep. The only climate input is per-station bin data from published ASHRAE Weather Data Viewer 2025 [23] (dry-bulb, wet-bulb, and annual hours). Every model parameter is an editable constant: the condenser approaches (14 K water-cooled, 12 K air-cooled), the minimum condensing temperature, the free-cooling thresholds, the heat-rejection auxiliary coefficients, charge, leakage, recovery, lifetime, GWP basis, and grid pathways.

`water_generic.py` computes the cooling-tower makeup water of Section 4.4 from the same climate bins. Refrigerant blends require care in an open property library. The build used here does not provide the predefined R-513A, R-450A, or R-515B mixtures, so each blend is represented by its documented surrogate composition: R-513A (R-1234yf/R-134a) uses the library's fitted binary interaction parameters, whereas R-450A and R-515B (R-1234ze(E)-based pairs for which no fitted parameters are available) use an estimated linear (Lorentz–Berthelot) mixing rule. Because the library cannot perform mixture pressure–entropy or near-saturation pressure–temperature flashes, blend cycles are evaluated through its low-level state interface with explicit phase specification and an isentropic compressor-outlet state located by temperature bisection. The resulting blend COPs are physically consistent (R-513A falls between R-1234yf and R-134a, and R-450A and R-515B track R-1234ze(E), of which they are largely composed).

---

## 4. Results

The results are the output of the open `compute_tewi_generic.py` pipeline (Section 3.7): eight refrigerants across five climate stations, two condenser archetypes, three control modes, and three eGRID-anchored grid pathways, using CoolProp real-fluid properties and ASHRAE climate bins [23].

### 4.1 Condenser operating conditions

The condensing temperature follows the ambient driving temperature and the head-pressure control (Section 3.3, Eq. (1); stations in Table 2). The chiller COP follows that condensing temperature (Fig. 2). In fixed mode head pressure is held at the 40 °C design minimum, giving a chiller COP of 6.1 for R-1234ze(E) at every site under the water-cooled archetype; under the air-cooled archetype the condensing temperature rises above 40 °C in the hottest climates, so the fixed-mode COP falls to 6.0 at Miami and 5.5 at Phoenix. In floating mode the condensing temperature tracks the ambient down, lifting the load-weighted seasonal COP; under the water-cooled archetype it rises to 10.2 (Seattle), 9.6 (Chicago), 9.3 (Ashburn), 9.2 (Phoenix), and 7.3 (Miami). The warm-climate values are markedly higher for the water-cooled condenser than for the air-cooled one (e.g., Phoenix 9.2 vs 6.4), the central reason the two architectures diverge in hot climates (Section 4.3); the near-identical curves of the seven subcritical fluids are evident in Figure 2.

**Table 2. Climate stations and free-cooling availability by condenser archetype (0.4 % cooling design DB/MCWB from ASHRAE Weather Data Viewer 2025 [23]; free-cool**

**hours from the 8 760-h annual climate bins). Free cooling is wet-bulb-driven for the water-cooled archetype and dry-bulb-driven for the air-cooled one.**

| Station | Climate | Design DB / MCWB, °F (°C) | Free-cool hours: Water-cooled / Air-cooled |
|---|---|---|---|
| Seattle | Cool-marine | 87 / 66 (30.6 / 18.9) | 5 755 / 4 129 |
| Chicago | Continental | 91 / 74 (32.8 / 23.3) | 5 185 / 4 251 |
| Ashburn | Mixed-humid | 93 / 75 (33.9 / 23.9) | 4 517 / 3 436 |
| Miami | Hot-humid | 92 / 78 (33.3 / 25.6) | 247 / 48 |
| Phoenix | Hot-dry | 111 / 70 (43.9 / 21.1) | 3 812 / 634 |

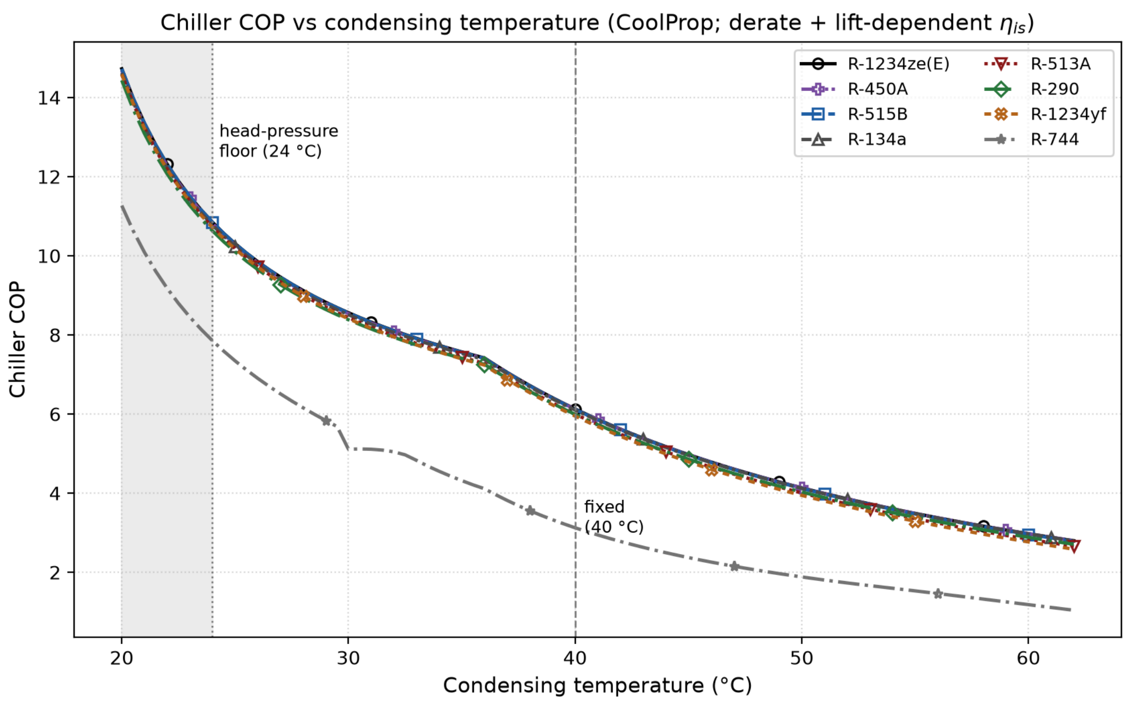


**Figure 2.** Computed chiller COP versus condensing temperature for the eight fluids (CoolProp cycle COP with a 0.90 derate and a lift-dependent isentropic efficiency). The seven subcritical fluids are nearly indistinguishable; transcritical R-744 is far lower and shows the subcritical/ transcritical handoff near 30 °C. The dashed line marks the fixed head-pressure condition (40 °C); the shaded band below the 24 °C head-pressure floor is outside normal operation.

### 4.2 The compliant refrigerant choice is a statistical tie

In the deterministic ranking, R-1234ze(E) gives the lowest lifetime TEWI in all 30 decision cells per condenser archetype (5 stations × 2 control modes × 3 grid pathways), under both the water-cooled and air-cooled architectures, and, for the water-cooled archetype, in all 420 sensitivity scenarios (those cells crossed with the leakage, recovery, GWP-basis, and lifetime sweep) (leakage 1–10 %, recovery 70–95 %, GWP basis AR4/AR5/AR6, lifetime 15–25 yr; see Table 4, Section 4.5). The cells share the same underlying COP curves and differ only in demand- and grid-side multipliers, so the sweep shows the deterministic ordering is insensitive to those assumptions. And the margin is within model uncertainty. The seven subcritical fluids differ by only ~3 % in chiller COP (R-1234ze(E) 6.13 at 40 °C, then R-450A and R-515B 6.12, R-134a 6.10, R-513A 6.00, R-290 5.99, R-1234yf 5.95, with transcritical R-744 far lower at 3.12), so the winner's lead over the runner-up is about one percent or less, smaller than the surrogate-blend and isentropic-efficiency uncertainties the model carries (Section 5.2). A Monte-Carlo propagation of an independent ±5 % COP uncertainty per fluid (2 000 draws over all cells; ±5 % is a plausible scenario band spanning the surrogate-blend and property-library uncertainty discussed above, not a value fitted to measured compressor data, and Table 6 reports the same analysis at ±2/3/8 % for context) makes this explicit (Fig. 3): R-1234ze(E) is the single most-likely optimum but is rank-1 in only 40 % of draws and within 1 % of the cell minimum in 55 %, while R-515B (21 %), R-290 (12 %), R-1234yf (12 %), and R-450A (12 %) form a statistically indistinguishable leading cluster; R-744 is never competitive (0 %). The 40 % assumes the ±5 % COP errors are independent across fluids; under correlated errors the deterministic ordering is preserved and R-1234ze(E) stays nominally first, and under a tighter ±2 % band it leads in 65 % of draws (Table 6, Section 5.2). What is robust across every variant is that the inter-fluid TEWI spread stays only a few percent, far too small to drive the choice. The low-GWP refrigerant choice is therefore effectively a tie among the top cluster, to be settled on safety, charge limits, cost, and availability rather than on lifetime climate impact. The leading cluster is dominated by flammable fluids (R-1234ze(E) and R-1234yf are A2L, R-290 is A3), with R-515B and R-450A its only A1 (non-flammable) members, so the safety and A2L charge-limit overlay flagged in Section 6 is decision-relevant within the cluster.

The literature direct/indirect "crossover" [7] survives only as a reshuffling within this tie: the second-place fluid is the efficient A1 blend R-515B in the warmer, dirtier-grid cases and the ultra-low-GWP HFO R-1234yf in the cooler, cleaner ones, but the spread is a few tenths of a percent. At Ashburn, floating, moderate-decarbonization grid, with the water-cooled condenser, the compliant fluids R-1234ze(E) (6 314 t $CO_2$e), R-515B (6 382 t), R-1234yf (6 383 t), R-290 (6 399 t), R-450A (6 446 t), and R-513A (6 500 t), plus the R-134a baseline (6 631 t), span 5.0 %, with R-744 (8 459 t) 34 % higher (Fig. 4); the air-cooled condenser gives the same ordering, a 4.9 % subcritical span, and a larger 41 % R-744 penalty, the higher condensing temperature widening the transcritical gap. Under the AR6 GWP basis the HFC/HFO blends move closer to the 700 cap (R-513A ≈ 674, R-450A ≈ 643) but remain compliant, and the ranking is unchanged.

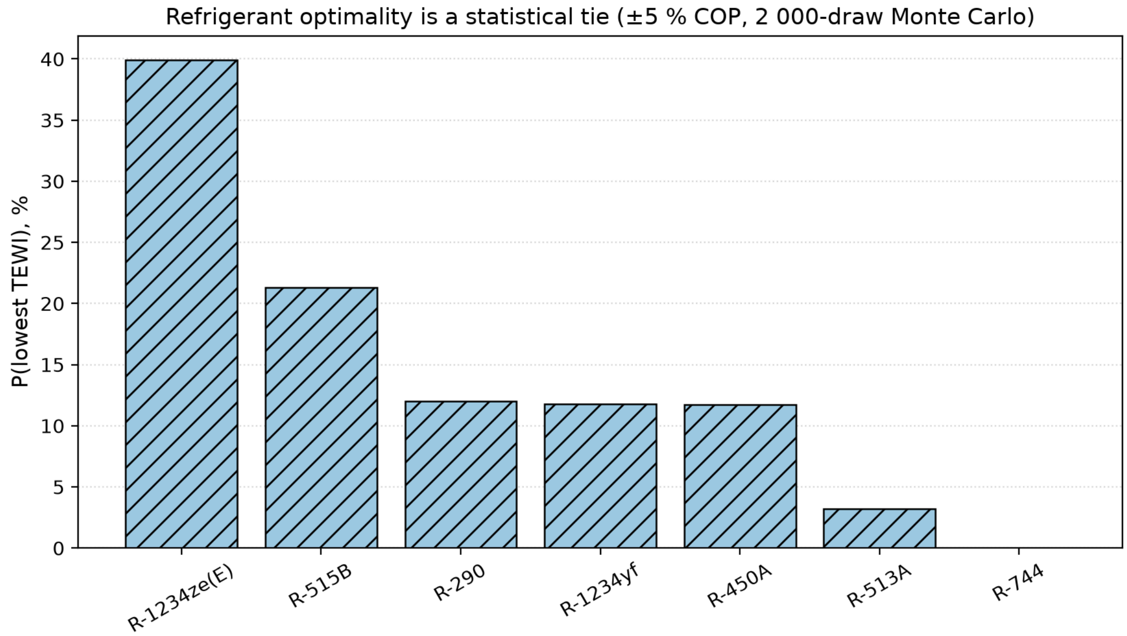


**Figure 3.** Probability that each refrigerant gives the lowest lifetime TEWI under an independent ±5 % COP uncertainty (2 000-draw Monte Carlo over all 30 cells). R-1234ze(E) is the most likely optimum (40 %), but R-515B, R-290, R-1234yf, and R-450A form a statistical tie: the refrigerant choice is not resolved by lifetime climate impact.

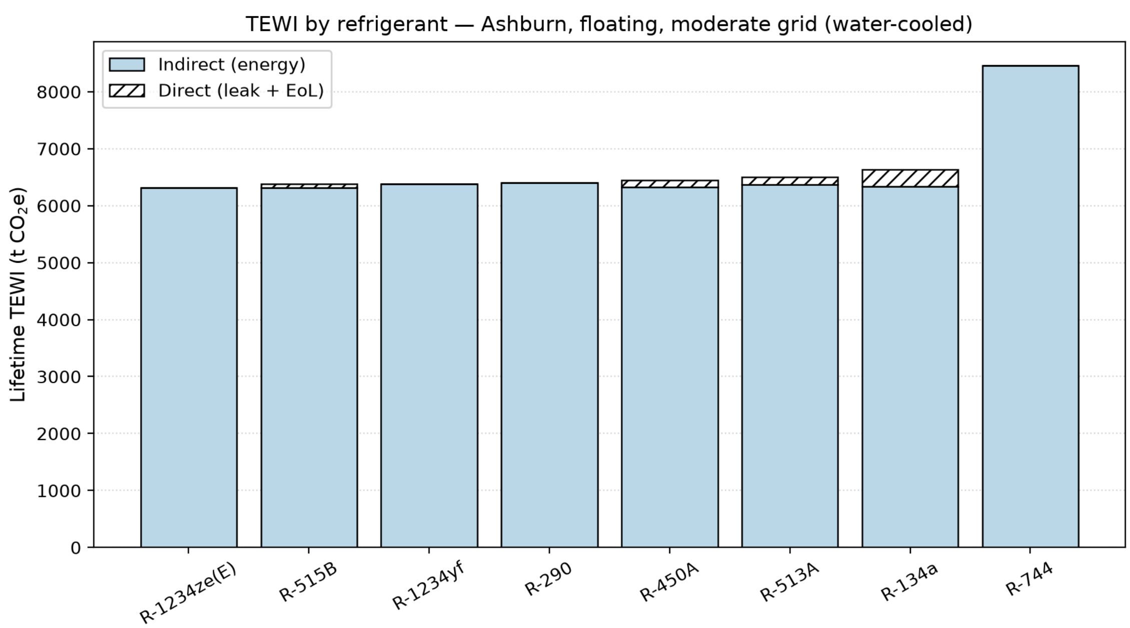


**Figure 4.** Lifetime TEWI by refrigerant (Ashburn, floating, moderate-decarbonization grid, water-cooled condenser), split into indirect (energy) and direct (leak + end-of-life) terms. The

seven subcritical fluids span 5.0 %; the direct term (orange) grows with GWP; transcritical R-744 is 34 % higher.

**4.3 Condenser architecture and control are the climate-dependent levers**

With the refrigerant effectively fixed, the consequential climate-dependent decisions are the condenser architecture and how it is run (Fig. 5, Table 3). Floating the condensing temperature and adding an integrated economizer cut plant energy substantially wherever the ambient driving temperature is low, under either architecture. In the cool and temperate climates both archetypes free-cool heavily: floating saves 23–29 % and the economizer 40–57 % versus fixed head pressure at Seattle, Chicago, and Ashburn. The two architectures diverge in the hot climates, which is where the water-for-energy trade-off appears. In hot-dry Phoenix the water-cooled condenser free-cools 3 812 hours (economizer −43 %); its water use there is no different from any other station (Section 4.4), but the low wet-bulb makes that water the most valuable it is anywhere. The air-cooled condenser there free-cools only 634 hours (economizer −14 %), and once floated runs about 27 % more energy than the water-cooled plant (1.89 vs 1.49 GWh), a gap that widens to ∼60 % with the economizer (1.83 vs 1.13 GWh), yet consumes no water. That gap is the return Phoenix gets on the tower water. In hot-humid Miami neither helps much (≤13 %), because both the wet-bulb and the dry-bulb stay high. The economizer is the larger lever throughout, but floating alone is already worth 9–29 %.

**Table 3. Annual chiller-plant electricity by station, condenser archetype, and control mode (R-1234ze(E), GWh; % versus fixed head pressure in parentheses). Compressor energy is refrigerant-set; the heat-rejection auxiliaries (tower/condenser fans + pumps) are generic and refrigerant-independent (Section 3.3); the economizer shuts the compressor off in the free-cooling hours. Floating percentages are central estimates; Section 5.2 gives their range.**

| Station | Water-cooled: Fixed / Floating / +Econ | Air-cooled: Fixed / Floating / +Econ |
|---|---|---|
| Seattle | 1.98 / 1.40 (−29 %) / 0.86 (−57 %) | 1.96 / 1.40 (−28 %) / 1.01 (−49 %) |
| Chicago | 1.98 / 1.45 (−27 %) / 0.96 (−51 %) | 1.97 / 1.47 (−25 %) / 1.07 (−46 %) |
| Ashburn | 1.98 / 1.48 (−25 %) / 1.06 (−47 %) | 1.97 / 1.51 (−23 %) / 1.18 (−40 %) |
| Phoenix | 1.98 / 1.49 (−25 %) / 1.13 (−43 %) | 2.14 / 1.89 (−12 %) / 1.83 (−14 %) |
| Miami | 1.98 / 1.75 (−12 %) / 1.72 (−13 %) | 2.00 / 1.82 (−9 %) / 1.82 (−9 %) |

The mechanism is the same in both: floating and the economizer act on the compressor by lowering or removing the condensing lift (the heat-rejection auxiliaries are refrigerant-independent and modest), so the savings are reductions in the dominant indirect-TEWI term. The architecture choice is then a water-versus-energy decision (the water-cooled condenser uses less plant energy, the gap concentrated in hot-dry climates, at the cost of cooling-tower water), and an evaporatively-assisted adiabatic cooler occupies the engineered middle ground between the two (the subject of a companion study, in preparation).

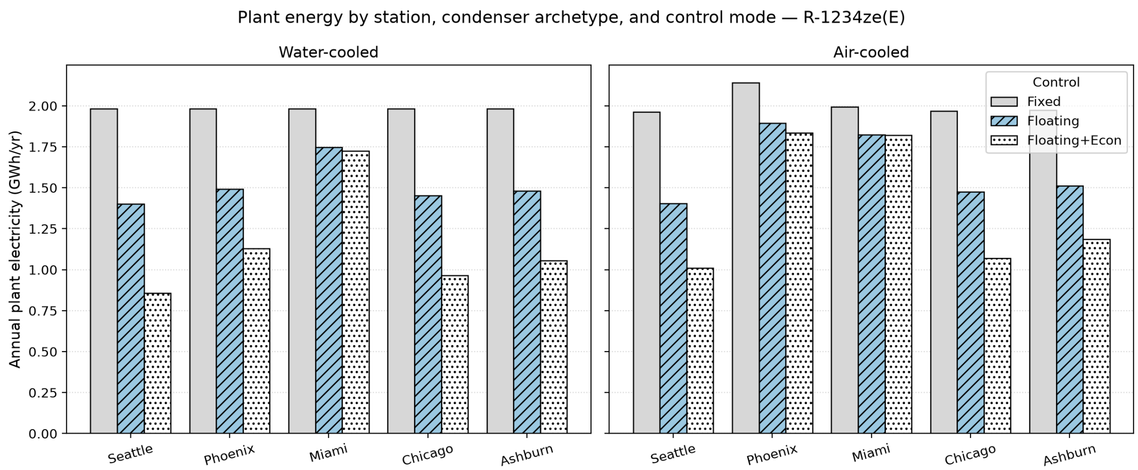


**Figure 5.** Annual chiller-plant electricity by station, condenser archetype, and control mode for R-1234ze(E). Floating and the economizer cut energy wherever the driving temperature is low; the water-cooled and air-cooled archetypes diverge most in hot-dry Phoenix.

### 4.4 The water–energy trade-off between the archetypes

The air-cooled condenser consumes no water; the water-cooled condenser's tower evaporates makeup water year-round, since it always rejects the constant IT load and only intermittently the added compressor heat (Section 3.3). Across all fifteen station-mode cells, annual makeup water for R-1234ze(E) varies by only 12 %, 4.29 to 4.81 million gallons per year (16.2 to 18.2 thousand $m^3$/yr), versus plant electricity, which more than doubles over the same cells, from 0.86 to 1.98 GWh (Table 3): tower duty is dominated by the climate- and control-independent IT load, only modestly by the compressor work that floating and the economizer remove. Water per kWh of plant electricity therefore rises as the plant gets more efficient, from 2.42 gal/kWh fixed to 4.99 gal/kWh with the economizer at Seattle, since electricity falls sharply while water barely does; Miami, with the fewest free-cooling hours, shows the least of either effect. This is the same weak climate sensitivity Lei and Masanet [20] report for cooling-tower systems without an economizer, and for the same reason. The choice here is not a water-versus-energy trade by climate: the two move together on energy and largely independently on water, so architecture and control are best chosen on electricity, with water reported alongside rather than traded against it, particularly where local supply is already stressed [24].

### 4.5 Decision summary

**Table 4. Runner-up fluid and its TEWI margin behind R-1234ze(E), by station and grid pathway (water-cooled archetype, floating control). R-1234ze(E) is the deterministic optimum in every one of the 15 station × grid-pathway cells, in all 30 cells per archetype, and in all 420 sensitivity scenarios, so only the runner-up and the margin are informative. Every margin here is at most 1.09 % (at most 1.18 % water-cooled and 1.32 % air-cooled across all 30 decision cells per archetype), and the runner-up shifts from the efficient A1 blend R-515B toward the ultra-low-GWP HFO R-1234yf as the grid decarbonizes. The air-cooled archetype shows the same shift more weakly, with R-1234yf the runner-up in 3 of its 30 decision cells against 5 of 30 water-cooled.**

| Station | Static grid | Moderate decarbonization | Aggressive decarbonization |
|---|---|---|---|
| Seattle | R-1234yf (+0.86 %) | R-1234yf (+0.86 %) | R-1234yf (+0.86 %) |
| Chicago | R-515B (+0.53 %) | R-515B (+0.69 %) | R-515B (+0.89 %) |
| Ashburn | R-515B (+0.82 %) | R-515B (+1.08 %) | R-1234yf (+1.09 %) |
| Phoenix | R-515B (+0.66 %) | R-515B (+0.87 %) | R-1234yf (+1.09 %) |
| Miami | R-515B (+0.60 %) | R-515B (+0.77 %) | R-515B (+0.97 %) |

The refrigerant decision is thus effectively settled at the *cluster* level (any of R-1234ze(E), R-515B, R-290, R-1234yf, or R-450A is climate-impact-equivalent, to be chosen on safety, charge limits, and cost), while the operationally consequential choices are the condenser architecture and its control strategy (climate-dependent, Section 4.3) and the resulting water budget (Section 4.4).

### 4.6 External validation of the TEWI accounting

Because the headline ranking rests on the TEWI implementation, that accounting was validated against an independent published case rather than only on internal consistency. Akinseloyin et al. [7] report a Life Cycle Climate Performance (LCCP) analysis of an 809 kW (230 RT) water-cooled centrifugal chiller on a campus load profile (27-yr life, 5 %/yr leakage, 85 % end-of-life recovery, ≈26 kg charge, natural-gas-fired grid at 1.004 kg $CO_2$e $kWh^{-1}$). Feeding their published per-fluid inputs (GWP, annual energy, and the parameters above) into our production `tewi_breakdown()` routine (the same code path used throughout this paper) reproduces their results on the three quantities a TEWI model should reproduce (Table 5): the fluid ranking is identical (R-1234ze(E) < R-513A < R-134a); the indirect (energy) term dominates (98–100 % of TEWI here, consistent with their ">90 % of emissions are indirect" on an LCCP basis); and the lifetime magnitude agrees to within 1–9 %, our TEWI lying just below their LCCP. That small shortfall grows monotonically with refrigerant GWP (1.0 % for near-zero-GWP R-1234ze(E), 4.4 % for R-513A, 8.6 % for R-134a), the direction expected of the embodied (refrigerant-manufacturing and equipment-materials) term that LCCP includes and TEWI omits by construction; a shortfall of this size, however, is large to attribute to refrigerant manufacturing alone, and the published inputs do not let us isolate whether it is that term specifically or a different direct-emission convention in the source study. Because the benchmark runs on their published per-fluid energies directly, the identical fluid ranking is a property of those inputs rather than a separate test of this paper's own COP model. This also explains why the headline TEWI differences among the compliant fluids (Section 4.2) are narrower than an LCCP study's would be: the embodied advantage of the lowest-GWP fluid lies outside TEWI scope, a conservative choice that, if anything, understates the case for ultra-low-GWP refrigerants.

**Table 5. External validation against Akinseloyin et al. [7] (809 kW water-cooled centrifugal chiller; their published per-fluid inputs run through this paper's TEWI engine). TEWI excludes embodied terms; LCCP includes them, so TEWI is expected to lie just below LCCP by a GWP-scaling margin.**

| Refrigerant | GWP | Annual energy (kWh) | TEWI, ours ($10^6$ kg) | LCCP, published ($10^6$ kg) | Indirect share (ours) | TEWI vs LCCP |
|---|---|---|---|---|---|---|
| R-134a | 1530 | 131 901 | 3.64 | 3.98 | 98.4 % | −8.6 % |
| R-513A | 674 | 132 690 | 3.63 | 3.79 | 99.3 % | −4.4 % |
| R-1234ze(E) | 1.37 | 130 706 | 3.54 | 3.58 | 100.0 % | −1.0 % |

## 5. Discussion

### 5.1 Implications

For AIM-Act-compliant data-center chillers, the refrigerant decision is effectively settled at the cluster level: R-1234ze(E) is the most-likely lifetime-TEWI optimum, but R-515B, R-290, R-1234yf, and R-450A are climate-impact-equivalent within model uncertainty (Section 4.2). An operator can therefore choose any of the leading low-GWP cluster on safety, charge-limit, cost, and availability grounds without sacrificing lifetime climate performance, a markedly simpler decision than the grid-dependent refrigerant trade-off emphasized in the prior chiller literature [6,7], which here matters only for the narrow second-place ordering (R-515B where the grid stays carbon-intensive, an ultra-low-GWP fluid where it decarbonizes). Transcritical R-744 is the one fluid to avoid in this single-stage architecture, 19 to 81 % worse depending on climate and control mode (Section 4.2).

The decisions that move lifetime impact are therefore operational and mechanical, not chemical. First, the condenser architecture and its control: floating the condensing temperature and an integrated economizer cut plant energy by up to ~57 % wherever the ambient driving temperature is low (under either a water-cooled or an air-cooled condenser), while hot-humid Miami gains least (≤13 %), so the control should be climate-tuned rather than fixed across a fleet. Second, the architecture itself is a water-versus-energy choice: the water-cooled condenser runs cooler (better COP) at the cost of evaporative tower water, whereas the air-cooled condenser uses none at a higher energy cost. That water use is itself nearly climate-invariant (Section 4.4); what varies by climate is where it falls and what it buys: hot-dry markets are typically the most water-stressed, and it is also there that the water-cooled archetype's energy return is largest (0.40 GWh/yr saved at Phoenix in floating mode, versus essentially none at cool-marine Seattle), so the water budget, not the refrigerant GWP, is frequently the binding sustainability constraint. Standards and procurement guidance that rank data-center cooling on refrigerant GWP alone thus optimize the wrong variable; the decision requires the joint refrigerant, control-strategy, and water assessment this model provides. The approach also dovetails with waste-heat-reuse strategies [16,17], which further shrink the indirect term.

### 5.2 Limitations and threats to validity

**COP model and uncertainty.** The chiller COP adjustments described in Section 3.3 are refrigerant-independent, but they are simple stand-ins for a full compressor map, and the magnitude of the lift dependence is the least constrained part of the model. The study was therefore re-run at the two limits that bound it physically: a flat 0.70, representing capacity control that recovers the off-design loss entirely, and efficiency falling in proportion to head, the fixed-speed limit implied by the constant input head of [18], which gives 0.27 at the minimum lift against the 0.44 assumed. R-1234ze(E) remains the deterministic winner in every cell under both, and the subcritical spread stays between 4.1 % and 5.8 %. The assumed curve's fixed-impeller-speed logic follows [18]; for chillers with variable-speed or magnetic-bearing capacity control, which recover more of the off-design loss than a fixed-speed machine, the flat-0.70 limit above is the closer analogue, so the reported floating-mode saving (−29 % at Seattle) is likely conservative for such equipment, nearer its −44 % flat-0.70 counterpart. The control-strategy results divide. Economizer savings hold up wherever free cooling is available, the Seattle figure moving only from −57 % to between −50 % and −61 %, because switching the compressor off does not depend on its off-design efficiency. Floating the condensing temperature alone leaves the compressor running at exactly the low lift where that efficiency is least certain, and is correspondingly sensitive: the Seattle saving spans −6 % to −44 % against −29 % at base, with the other stations behaving the same way. Miami is the exception on both counts, since with few free-cooling hours its economizer result is essentially its floating result and inherits the same range. More consequentially, the seven subcritical fluids differ by only ~3 % in COP, so the deterministic winner leads by about one percent or less (smaller than the combined isentropic-efficiency, surrogate-blend, and property-library uncertainty). The ±5 % COP Monte-Carlo of Section 4.2 shows the refrigerant result to be a statistical tie rather than a determinate ranking, and its assumptions are stress-tested over band width and error correlation (Table 6). R-1234ze(E)‘s rank-1 frequency rises from 40 % at ±5 % to 65 % at ±2 %. Table 6's common-mode column reads 100 % at every band width: a shared multiplicative shift on compressor energy preserves each fluid's compressor-energy ranking exactly, and for these fluids' relative direct-versus-indirect-term magnitudes, that in turn preserves the full TEWI ranking at every band width tested here. This is an upper bound on the independent-draw case; the dominant isentropic-efficiency and derate uncertainties are refrigerant-independent for the same reason (Section 3.3). The ±5 %, independent draw is thus the most tie-favorable case; across all variants the inter-fluid TEWI spread remains a few percent, so the practical conclusion (choose on safety, charge, and cost) is invariant. The independent draw shows that the *nominal* first place is not robustly resolved. The 2 000-draw, fixed-seed estimate carries a binomial sampling error of about ±1 percentage point. The Section 4.2 Monte Carlo is run over the fixed and floating modes only (30 cells); extending it to include the floating-with-economizer mode as well (45 cells) leaves the conclusion intact: R-1234ze(E)'s rank-1 frequency is essentially unchanged (41.9 % water-cooled, 41.5 % air-cooled), and the same four fluids (R-515B, R-290, R-1234yf, R-450A) still form the trailing cluster, though their relative order among themselves shifts slightly, consistent with positions two through four not being resolved.

**Table 6. R-1234ze(E) rank-1 frequency (%) under COP-uncertainty band width and per-fluid error correlation.**

| COP uncertainty | Independent errors | Common-mode (correlated) errors |
|---|---|---|
| ±2 % | 65 | 100 |
| ±3 % | 53 | 100 |
| ±5 % | 40 | 100 |
| ±8 % | 32 | 100 |

**Blend properties.** The HFC/HFO blends are modelled by surrogate compositions because the open CoolProp 7.2.0 build ships no predefined R-513A/R-450A/R-515B mixtures: R-513A uses fitted R-1234yf/R-134a parameters, whereas R-450A and R-515B use an estimated linear mixing rule. The blends' saturation temperatures are taken at the glide midpoint (Section 3.3); rerunning all 60 decision cells with the dew point at the evaporator and the bubble point at the condenser instead lowers R-450A's chiller COP by 2.2 % at 40 °C and 5.4 % at 24 °C, drops it from third place in most cells to sixth or seventh of the eight fluids, level with the R-134a baseline, and changes the winner in none of the 60 cells; R-513A and R-515B move by less than 0.3 %. The resulting COPs are physically consistent but approximate (most so for R-450A and R-515B), and the REFPROP predefined mixtures should be used to firm up the within-cluster ordering; they are not expected to change the cluster membership.

**Heat-rejection model.** The condensing temperature is the ambient driving temperature plus a generic approach (14 K water-cooled, 12 K air-cooled), floored at 24 °C. That floor was swept to 22 °C and 27.5 °C (both archetypes, all stations, all modes): the winning fluid changes in none of the 90 deterministic cells, and the floating-versus-fixed saving shifts by up to 4.8 percentage points (Seattle water-cooled spans −32 % to −25 % against −29 % at base), while the economizer saving, governed by free-cooling availability rather than the floor, barely moves (Section 4.3). The cycles-of-concentration value for the tower water calculation (Section 3.3) is a stated assumption within a cited range; water use scales directly with it, but the finding that it varies far less than plant electricity across modes does not. The heat-rejection auxiliaries (tower/condenser fans and pumps) are generic, refrigerant-independent coefficients (Section

3.3) rather than benchmarked equipment curves, so their robustness was checked by sweep: varying the fan coefficient and the pump power by ±50 %, individually and jointly, changes the winning fluid in none of the 90 deterministic cells and shifts the representative savings only modestly (the Seattle water-cooled economizer saving spans −50 % to −65 % around the −57 % base; hot-climate cells move by up to 7 points). A detailed fan-speed-modulation (affinity-law) and site-specific tower model would refine the absolute auxiliary energy (and is the natural refinement for a companion equipment study, in preparation), but within this ±50 % envelope it changes neither the ranking nor the qualitative mode comparison. The free-cooling threshold (53 °F / 11.7 °C wet-bulb, water-cooled; 50 °F / 10.0 °C dry-bulb, air-cooled) is likewise a stated assumption rather than a fitted value: swept ±3 °F and ±5 °F for each archetype in turn, it changes the winning fluid in none of the 90 deterministic cells, and shifts the economizer-savings estimate by at most 10.5 percentage points at the water-cooled threshold and 7.5 at the air-cooled one, in the expected direction (a higher threshold means more free-cooling hours and larger savings).

**Transcritical R-744.** $CO_2$′s penalty is climate- and control-mode-dependent, 19 % at best (Seattle with the economizer) to 81 % at worst (Phoenix, fixed head pressure, air-cooled), and is specific to the single-stage transcritical cycle modelled here, in which the gas-cooler outlet is held 2 K below the heat-rejection temperature, the same basis as the subcritical fluids' subcooled liquid, and the high-side pressure is optimized for maximum COP over a 7.5–13.0 MPa search range, a realistic design-pressure envelope for commercial $CO_2$ systems, so the design-point COP of 3.12 at the 40 °C reference condition uses the full unconstrained search; the cap applies only above about 53 °C gas-cooler outlet, in the few hottest air-cooled fixed-mode bins, where it understates R-744′s best-case COP by up to roughly 10 %, a small effect on the annual-energy-weighted result given how few hours occur that hot. Ejector, economized, or parallel-compression cycles and a colder gas-cooler approach would narrow the gap: the result describes single-stage transcritical $CO_2$ at these condensing temperatures, not $CO_2$ generally.

**Leakage, recovery, and GWP metric.** Direct emissions scale with the assumed annual leakage (3 %/yr default) and end-of-life recovery (85 %); both are swept (leakage 1–10 %, recovery 70–95 %), and the winner is invariant. Following the standard IIR convention, the end-of-life term applies the recovery fraction to the original charge even though cumulative leakage removes part of it over life; at these rates this slightly overstates the (already small) direct term and is immaterial to the ranking. The 3 % default is generic: hermetic data-center chillers often leak below 1 %/yr, which would make the indirect term even more dominant and the refrigerant choice even more of a tie. GWP-100 is reported on the AR4, AR5, and AR6 bases; the AR6 effect on the HFC/HFO blends is given in Section 4.2, where it does not change the ranking.

**Embodied/LCCP terms.** The study reports TEWI; embodied (manufacturing and disposal) terms are out of scope and set to zero, so LCCP is not computed. The external validation (Section 4.6) makes this gap explicit and bounds it: our TEWI lies 1–9 % below the published LCCP of the same case, a margin that grows monotonically with refrigerant GWP in the direction expected of the embodied refrigerant-manufacturing contribution, though its exact source cannot be isolated from the published inputs alone (Section 4.6). Because embodied emissions gain relative weight as the grid decarbonizes (and would also burden the higher-charge, lower-COP R-744), a full LCCP treatment is a useful extension (Section 6).

**Grid and load.** Grid intensities are anchored to each site's eGRID2022 subregion, but the two decarbonization shapes are uniform linear declines, not region-specific forecasts, and grid-average (not long-run-marginal) factors are used; both are worth revisiting for a 24/7 baseload load (Section 6). For an always-on load, long-run marginal factors typically exceed grid-average in fossil-heavy subregions (e.g., RFCW) and fall below it in hydro-rich ones (NWPP), shifting absolute indirect TEWI by of order ±10–20 %. The trajectory shape was checked directly: rerunning every deterministic cell (both archetypes, all stations, both head-pressure modes, all three pathways; 60 cells) with the linear decline replaced by a front-loaded (square-root) or back-loaded (quadratic) path to the same endpoints shifts each fluid's absolute TEWI by 10 to 23 % (mean −17 % front-loaded, +17 % back-loaded) but changes the winning fluid in none of the 60 cells. This is as expected, since every fluid shares the same $\beta_y$, so a shape change rescales all fluids' indirect terms nearly uniformly. The near-constant full-load, single-1 MW-chiller assumption is reasonable for hyperscale but omits part-load behaviour and multi-chiller staging, and the climate input is ASHRAE bin climatology rather than an hourly TMY, so it captures the annual distribution of conditions but not their sequence (diurnal patterning, thermal inertia).

**Scope.** The study addresses a single vapor-compression chiller with two generic condenser archetypes; it does not model direct-to-chip or immersion architectures, refrigerant cost/availability, or the A2L charge limits pertinent to the recommended cluster. The heat-rejection model is generic, so the energy results are envelope brackets rather than a specific installed design; a real evaporatively-assisted cooler would fall between the air-cooled and water-cooled bounds.

### 5.3 Why the result is non-obvious

AIM-Act compliance can be framed as "pick the lowest GWP under 700" or, following the chiller TEWI literature [6,7], as a grid-dependent trade-off between a higher-COP A1 blend and an ultra-low-GWP fluid. For data-center chillers, neither framing decides the outcome: the compliant low-GWP fluids are so close in lifetime TEWI that the choice among them is a statistical tie (Section 4.2), so impact is governed instead by how heat is rejected and how that is controlled, variables the refrigerant literature does not consider. The counter-intuitive results are that the same control choices (floating, economizer) pay under *both* condenser architectures wherever the driving temperature is low (including hot-dry climates for the water-cooled condenser), yet a high wet-bulb *and* dry-bulb together still block free cooling in hot-humid climates (Miami); and that the water-cooled architecture's energy advantage is bought with exactly the evaporative water that is scarcest where that advantage is largest. Selecting the refrigerant cluster, the condenser architecture, and its control jointly (and per climate) is the actionable contribution.

## 6. Conclusion and Future Work

The imminent AIM Act GWP-700 cap forces a refrigerant decision for new data-center cooling equipment, but for data-center chillers that decision turns out to be the easy part. Coupling CoolProp refrigerant properties with two generic condenser archetypes (water-cooled and air-cooled) across five real climate stations, three control modes, and three eGRID-anchored grid-decarbonization pathways inside the standard TEWI equations, this study finds that the compliant low-GWP refrigerants are separated by only a few percent in lifetime Total Equivalent Warming Impact: R-1234ze(E) leads the deterministic ranking in every cell under both architectures, and in all 420 water-cooled sensitivity scenarios, but only by about one percent or less; under a ±5 % independent-error COP uncertainty it is the most-likely optimum in just 40 % of draws (with R-515B, R-290, R-1234yf, and R-450A statistically indistinguishable behind it). The margin is too small to drive the selection (transcritical R-744 excepted, 19 to 81 % worse depending on climate and control). What depends strongly on climate is operational, not chemical: floating the condensing temperature and an integrated economizer cut plant energy by up to ∼57 % wherever the driving temperature is low (under either architecture), while hot-humid Miami gains least (≤13 %); and the architecture itself is a water-versus-energy choice, the water-cooled condenser trading evaporative water for lower energy most sharply in hot-dry climates. The lifetime-optimal decision is therefore a joint choice of a low-GWP refrigerant (any of the leading cluster, on safety and cost), a condenser architecture, and a climate-tuned control strategy, not a refrigerant-GWP ranking alone.

**Future work.** With two generic condenser archetypes, climate bins, eGRID-anchored grids [22], the integrated economizer, AR6 GWP values, and a COP-uncertainty propagation in the model, the principal remaining refinements need external data. The R-450A/R-515B blend properties should be replaced with the REFPROP predefined mixtures (the deciding margins are within the present surrogate uncertainty); a site-specific evaporatively-assisted (adiabatic) cooler model that interpolates between the air-cooled and water-cooled brackets, with its water-energy trade-off, is the subject of a companion study (in preparation); and the uniform decarbonization shapes should be replaced with region-specific trajectories (e.g., NREL Cambium) and, ideally, long-run marginal rather than grid-average emission factors for a 24/7 baseload load. The TEWI accounting has been benchmarked against an independent published chiller case (Section 4.6, Table 5); extending that benchmark to an end-to-end comparison against an independently modelled chiller (exercising this paper's own COP and condenser submodels) is a further worthwhile check. Longer-term extensions include part-load and compressor-type effects, multi-chiller staging, a water-scarcity-weighted impact metric for hot-dry markets, and a safety/charge-limit overlay for the A2L cluster.

---

## Author Contributions

Conceptualization, V.M.; methodology, V.M.; software, V.M.; validation, V.M. and S.K.; formal analysis, V.M.; investigation, V.M.; writing, V.M., with review and guidance from S.K. All authors have read and agreed to the published version of the manuscript.

## Funding

This research received no external funding.

## Data Availability Statement

The complete computational pipeline is distributed with this manuscript and reproduces every result in this paper, including every sensitivity and uncertainty analysis, on a standard desktop. It comprises seven Python modules: `compute_tewi_generic.py` (the deterministic driver: the two condenser archetypes, the three control modes, and the TEWI evaluation), `sensitivity_generic.py` (every sensitivity and uncertainty result described in Section 3.7), `compute_tewi.py` (the CoolProp refrigerant cycle/COP model), `chiller_model.py` (the chiller-level COP adjustments, the heat-rejection constants, the system parameters, and the grid trajectories), `water_generic.py` (the cooling-tower makeup water of Section 4.4), `make_figures_generic.py` (Figures 1 to 5), and `validate_external.py` (the independent benchmark of Section 4.6), together with the per-station climate bins (`climate_bins/`) and the generated result tables and figures (`results_coolers/`). All inputs are from standard, non-proprietary sources: the open-source CoolProp library for thermophysical properties, published ASHRAE Weather Data Viewer 2025 [23] for climate bins (dry-bulb, wet-bulb, and annual hours per station), and EPA eGRID subregion carbon intensities [22]. Every model parameter (condenser approaches, the head-pressure floor, free-cooling thresholds, the heat-rejection auxiliary-power coefficients, charge, leakage, recovery, lifetime, GWP basis, and grid pathways) is an editable constant.

## Declaration of Competing Interest

Author V.M. is employed as Technical Director at Nimbus Advanced Process Cooling, which manufactures adiabatic and dry fluid coolers, a class of heat-rejection equipment relevant to this application domain; this is disclosed in the interest of full transparency. This study uses no proprietary manufacturer data, products, or methods. The chiller cycle is modelled entirely from open CoolProp thermophysical properties, and heat rejection is represented by two generic, textbook condenser archetypes: a water-cooled (cooling-tower) condenser and an air-cooled condenser, driven by published climate data [23], with all model parameters stated as editable constants (Section 3.3). The refrigerant ranking and the heat-rejection energy results are therefore independent of the corresponding author's employer and fully reproducible by any reader. The company had no role in the analysis or in the decision to publish; the second author (Sunwoo Kim, University of Alaska Fairbanks) declares no competing interests.

## Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the corresponding author used Anthropic's Claude to assist with language editing, to help draft and refine the authors' own wording, and to assist with implementation of the models. The authors reviewed, revised, and verified all outputs, and take full responsibility for the content of the publication.

## References

[1] International Energy Agency, *Energy and AI*, IEA, Paris, France, 2025. https://www.iea.org/reports/energy-and-ai (accessed June 9, 2026).

[2] S. Cai, Z. Gou, Towards energy-efficient data centers: a comprehensive review of passive and active cooling strategies, *Energy Built Environ.* 7 (2026) 206–226. https://doi.org/10.1016/j.enbenv.2024.08.009.

[3] C. Nadjahi, H. Louahlia, S. Lemasson, A review of thermal management and innovative cooling strategies for data center, *Sustain. Comput. Inform. Syst.* 19 (2018) 14–28. https://doi.org/10.1016/j.suscom.2018.05.002.

[4] U.S. Environmental Protection Agency, *Technology Transitions HFC Restrictions by Sector (American Innovation and Manufacturing Act)*, U.S. EPA, Washington, DC, 2023. https://www.epa.gov/climate-hfcs-reduction/technology-transitions-hfc-restrictions-sector (accessed June 9, 2026).

[5] International Institute of Refrigeration, *Guideline for Life Cycle Climate Performance*, version 1.2, IIR Working Party on LCCP, Paris, France, 2016. https://www.cold.org.gr/library/downloads/Docs/Guideline%20for%20life%20cycle%20climate%20performance.pdf (accessed June 9, 2026).

[6] R. Mansouri, B.-J.R.M. Bisulandu, A. Ilinca, Assessing energy performance and environmental impact of low GWP vapor compression chilled water systems, *Energies* 16 (2023) 4751. https://doi.org/10.3390/en16124751.

[7] M.O. Akinseloyin, C. Hall, E.M. Stallbaumer-Cyr, J.A. Morrow, M.M. Derby, Analyzing the carbon dioxide emissions of R134a alternatives in water-cooled centrifugal chillers using the life cycle climate performance framework, *Front. Energy Res.* 11 (2023) 1293993. https://doi.org/10.3389/fenrg.2023.1293993.

[8] C. Mateu-Royo, A. Mota-Babiloni, J. Navarro-Esbrí, Á. Barragán-Cervera, Comparative analysis of HFO-1234ze(E) and R-515B as low GWP alternatives to HFC-134a in moderately high temperature heat pumps, *Int. J. Refrig.* 124 (2021) 197–206. https://doi.org/10.1016/j.ijrefrig.2020.12.023.

[9] ASHRAE Technical Committee 9.9, *Thermal Guidelines for Data Processing Environments*, fifth ed., ASHRAE, Atlanta, GA, 2021.

[10] Y. Jin, X. Bai, X. Xu, R. Mi, Z. Li, Climate zones for the application of water-side economizer in a data center cooling system, *Appl. Therm. Eng.* 250 (2024) 123450. https://doi.org/10.1016/j.applthermaleng.2024.123450.

[11] J.H. Kim, D.U. Shin, H. Kim, Data center energy evaluation tool development and analysis of power usage effectiveness with different economizer types in various climate zones, *Buildings* 14 (2024) 299. https://doi.org/10.3390/buildings14010299.

[12] International Energy Agency, *Emissions Factors 2024: Database Documentation*, IEA, Paris, France, 2024. https://www.iea.org/data-and-statistics/data-product/emissions-factors-2024 (accessed June 9, 2026).

[13] I.H. Bell, J. Wronski, S. Quoilin, V. Lemort, Pure and pseudo-pure fluid thermophysical property evaluation and the open-source thermophysical property library CoolProp, *Ind. Eng. Chem. Res.* 53 (2014) 2498–2508. https://doi.org/10.1021/ie4033999.

[14] M. Lucas, P. Martínez, C.G. Cutillas, P.J. Martínez, J. Ruiz, A.S. Kaiser, B. Zamora, Experimental optimization of the thermal performance of a dry and adiabatic fluid cooler, *Appl. Therm. Eng.* 69 (2014) 1–10. https://doi.org/10.1016/j.applthermaleng.2014.04.033.

[15] V.K.S. Mokkapati, D.K. Das, When does adiabatic cooling outperform dry cooling? A twelve-city U.S. data-center siting assessment at ASHRAE design conditions, engrXiv preprint (2026). https://doi.org/10.31224/7795.

[16] X. Yuan, J. Liu, S. Sun, X. Lin, X. Fan, W. Zhao, R. Kosonen, Data center waste heat for district heating networks: a review, *Renew. Sustain. Energy Rev.* 219 (2025) 115863. https://doi.org/10.1016/j.rser.2025.115863.

[17] L. Zhang, Z. Zhao, B. Chen, M. Zhao, Y. Chen, Zero-carbon development in data centers using waste heat recovery technology: a systematic review, *Sustainability* 17 (2025) 10101. https://doi.org/10.3390/su172210101.

[18] J.J. Brasz, Comparison of part-load efficiency characteristics of screw and centrifugal compressors, in: International Compressor Engineering Conference at Purdue, Purdue University, West Lafayette, IN, 2006, Paper 1827. https://docs.lib.purdue.edu/icec/1827.

[19] ANSI/AHRI Standard 550/590-2023 (I-P), Performance Rating of Water-Chilling and Heat Pump Water-Heating Packages Using the Vapor Compression Cycle, Air-Conditioning, Heating, and Refrigeration Institute, Arlington, VA, 2023.

[20] N. Lei, E. Masanet, Climate- and technology-specific PUE and WUE estimations for U.S. data centers using a hybrid statistical and thermodynamics-based approach, *Resour. Conserv. Recycl.* 182 (2022) 106323. https://doi.org/10.1016/j.resconrec.2022.106323.

[21] The Green Grid, *Water Usage Effectiveness (WUE): A Green Grid Data Center Sustainability Metric*, White Paper #35, The Green Grid, Beaverton, OR, 2011. https://www.thegreengrid.org/system/files/store/WUE_v1.pdf (accessed 2026-09-15).

[22] U.S. Environmental Protection Agency, *Emissions & Generation Resource Integrated Database (eGRID), eGRID2022*, U.S. EPA, Washington, DC, 2024. https://www.epa.gov/egrid (accessed June 17, 2026).

[23] ASHRAE, Weather Data Viewer 2025 (based on ASHRAE Research Project RP-1923, Climatic Design Conditions for the 2025 ASHRAE Handbook—Fundamentals and ASHRAE Standard 169), ASHRAE, Atlanta, GA, 2025.

[24] M.A.B. Siddik, A. Shehabi, L. Marston, The environmental footprint of data centers in the United States, *Environ. Res. Lett.* 16 (2021) 064017. https://doi.org/10.1088/1748-9326/abfba1.

[25] P. Forster, V. Ramaswamy, P. Artaxo, et al., Changes in Atmospheric Constituents and in Radiative Forcing, in: Climate Change 2007: The Physical Science Basis, Cambridge University Press, Cambridge, UK, 2007, Ch. 2.

[26] G. Myhre, D. Shindell, F.-M. Bréon, et al., Anthropogenic and Natural Radiative Forcing, in: Climate Change 2013: The Physical Science Basis, Cambridge University Press, Cambridge, UK, 2013, Ch. 8.

[27] P. Forster, T. Storelvmo, K. Armour, et al., The Earth's Energy Budget, Climate Feedbacks, and Climate Sensitivity, in: Climate Change 2021: The Physical Science Basis, Cambridge University Press, Cambridge, UK, 2021, Ch. 7.