

## | RESEARCH ARTICLE

Article DOI: 10.21474/JNCS01/146

DOI URL: <http://dx.doi.org/10.21474/JNCS01/146>

### ZERO-WATER CLOSED-LOOP DIRECT-TO-CHIP COOLING ARCHITECTURES FOR HIGH-DENSITY AI DATA CENTERS USING ELEVATED COOLANT TEMPERATURES

Anil Mandloi<sup>1</sup> and Prakhi Mandloi<sup>2</sup>

1. Independent Researcher, Phoenix, Arizona, United States.

2. Arizona State University, Phoenix, Arizona, United States.

Corresponding Author: Anil Mandloi

#### | ABSTRACT

High-density AI data centers generate large heat loads while facing growing pressure to limit water use. Traditional evaporative cooling commonly requires 1.5–2.6 million gallons of water per megawatt each year. This paper presents the Multi-Stage Elevated Temperature Closed-Loop Architecture (MET-CLA), which combines direct-to-chip cold plates, dual-stage coolant distribution units, staged dry coolers, and predictive control at coolant temperatures of 40–45 °C. A hybrid steady-state and transient thermal-fluid model, cross-checked against published elevated-temperature benchmarks, indicates that on-site Water Usage Effectiveness can be reduced to approximately 0.02–0.05 L/kWh in temperate and hot-dry climates, while Power Usage Effectiveness remains in the range 1.06–1.12. The architecture is compatible with current high-power GPU platforms and is offered as a practical systems design for lower water intensity.

#### | KEYWORDS

Direct-to-chip cooling; Elevated coolant temperature; AI data center; Water usage effectiveness; Closed-loop architecture; Dry cooler; Power usage effectiveness

#### | ARTICLE INFORMATION

RECEIVED: 08 August 2026

ACCEPTED: 13 September 2026

PUBLISHED: October 2026

#### Introduction:-

The rapid growth of large-language model training and inference has driven rack power densities beyond 100–200 kW, exceeding the practical limits of conventional air cooling [1], [2]. Most existing reviews of liquid cooling in data centers treat indirect liquid cooling in detail, while system-level integration of direct liquid cooling for near-zero on-site water use remains comparatively less developed [2], [3]. Evaporative cooling towers remain energy-efficient in many climates but consume large volumes of water and face increasing regulatory and community opposition [4], [5]. Global average Water Usage Effectiveness (WUE) in data centers sits around 1.8 L/kWh. Efficient sites operate closer to 0.2 L/kWh, and zero operational water consumption is now reported as achievable through closed-loop thermal architectures [5], [6].

Closed-loop liquid cooling, especially direct-to-chip (DTC) cold-plate systems, transfers heat more effectively than air and can support higher coolant temperatures [7], [8]. At these elevated temperatures, ambient air can serve as the primary heat sink through dry coolers [9], [10]. Industry reference designs operating near 45 °C coolant inlet temperature have demonstrated that substantial reductions in on-site water use are possible [4]. Most published designs, however, remain single-loop systems and do not fully integrate multi-stage heat rejection with predictive control [3], [11].

This paper proposes a structured four-stage architecture—the Multi-Stage Elevated Temperature Closed-Loop Architecture (MET-CLA)—that integrates elevated-temperature DTC cooling, dual-loop coolant distribution units (CDUs), staged dry coolers, and a predictive thermal management layer. The primary contribution is the explicit systems-level organization of these elements into a coherent framework aimed at lowering on-site water intensity while remaining compatible with current high-power GPU platforms. The work is presented as a modeling-based systems architecture study rather than an experimental

validation campaign. Fig. 1 contrasts traditional evaporative cooling with the elevated-temperature closed-loop approach that forms the basis of MET-CLA.




**Fig. 1:- Conceptual comparison of traditional evaporative cooling and the elevated-temperature closed-loop direct-to-chip approach.**

**The Proposed Method:-**

**Design Philosophy and System Overview:-**

MET-CLA was developed around three principles: (1) operation at elevated coolant temperatures (40–45 °C) so that dry coolers can reject heat for a large fraction of the year; (2) isolation of the IT coolant loop from the facility loop; and (3) staged heat rejection combined with predictive control to increase free-cooling hours [4], [12], [13]. The architecture consists of four stages, illustrated in Fig. 2.




**Fig. 2:- Schematic of the Multi-Stage Elevated Temperature Closed-Loop Architecture (MET-CLA).**

#### **Four-Stage Architecture:-**

**Stage 1** – Server-level DTC. Cold plates are mounted on GPUs, CPUs, and high-power networking components. The target inlet temperature is 40–45 °C; flow is sized at roughly 1.0–1.5 L/min per kW (commonly about 1.2 L/min/kW at 45 °C). Under design load the coolant typically leaves at 50–55 °C [4], [8], [14].

**Stage 2** – Primary CDU loop. This loop isolates the IT fluid. Heat crosses plate heat exchangers into the secondary loop. Variable-speed pumps limit parasitic power, and fine filtration (approximately 25 µm) with redundant pumps maintains reliability [15], [16].

**Stage 3** – Secondary loop with staged dry coolers. Heat is rejected to ambient air through multiple dry-cooler banks that stage on only when needed. Elevated coolant temperature allows dry coolers to operate effectively for most of the year in temperate and dry climates [4], [17]. Adiabatic assistance is retained only as emergency backup and is not used under normal elevated-temperature operation.

**Stage 4** – Predictive control. A supervisory layer uses real-time sensor data plus short-term weather and workload forecasts to adjust pump speeds, valve positions, cooler staging, and the coolant supply setpoint within the 40–45 °C band [11], [18], [19].

#### **Method:-MATERIALS AND COMPONENT SELECTION:-**

**Coolant.** A single-phase propylene glycol–water mixture (PG-25) was selected for its balance of freeze protection, corrosion inhibition, and thermal performance at elevated temperatures. Propylene glycol was preferred over ethylene glycol because of lower toxicity [9], [10].

**Cold plates.** Copper or copper-tungsten cold plates with microchannel or staggered pin-fin geometries were specified to keep thermal resistance low at higher inlet temperatures. A high-performance thermal interface material is applied between the chip heat spreader and the cold plate [8], [14], [20].

**CDUs and dry coolers.** Primary CDUs include variable-speed pumps, stainless-steel plate heat exchangers, fine filtration, and redundant pumps to maintain isolation and reliability [15]. Staged air-cooled dry coolers provide the main heat-rejection path. Hybrid dry/adiabatic capability is retained only for rare peak conditions [17], [21].

#### **Modeling Approach and Comparative Scenarios:-**

A hybrid steady-state and transient thermal-fluid model was constructed using thermal resistance networks, the effectiveness-NTU method, hydraulic pressure-drop relations, and Typical Meteorological Year (TMY) climate files [13], [22], [23]. The model explicitly accounts for temperature-dependent fluid properties of PG-25, cold-plate thermal resistance, CDU heat-exchanger effectiveness, dry-cooler approach temperatures, and variable-speed fan and pump power. Metrics evaluated were

on-site WUE, facility Power Usage Effectiveness (PUE), annual free-cooling hours, and annual cooling energy consumption. Three comparative scenarios were evaluated under an identical 100 MW IT load (Table 1) to isolate the effect of coolant temperature and heat-rejection strategy. Model outputs were cross-checked against published elevated-temperature liquid-cooling benchmarks [12], [13], [24].

**Table 1:- Comparative scenarios used in the modeling study**

| Scenario | Description             | Coolant Inlet                     | Heat Rejection                  |
|----------|-------------------------|-----------------------------------|---------------------------------|
| A        | Traditional evaporative | Low (chilled water)               | Evaporative cooling towers      |
| B        | Conventional DTC        | $\leq 32\text{ }^{\circ}\text{C}$ | Chillers + partial free cooling |
| C        | Proposed MET-CLA        | 40–45 $^{\circ}\text{C}$          | Staged dry coolers              |

Key modeling assumptions are summarized in Table 2. Performance was examined for temperate, hot-dry, and mixed climates [13], [25].

**Table 2:- Key modeling assumptions**

| Parameter                          | Assumed Value / Approach                                        |
|------------------------------------|-----------------------------------------------------------------|
| IT load                            | 100 MW steady reference                                         |
| Coolant                            | PG-25 (25% propylene glycol)                                    |
| Target inlet temperature (MET-CLA) | 40–45 $^{\circ}\text{C}$                                        |
| Flow rate target                   | $\approx 1.2\text{ L/min per kW at }45\text{ }^{\circ}\text{C}$ |
| Climate data                       | TMY representative profiles                                     |
| TIM and fluid chemistry            | Idealized (high-quality application assumed)                    |

#### Research Design and Data Sources:-

The study follows a mixed-methods design that combines quantitative system modeling with qualitative comparative analysis. The work proceeded in three phases: (1) architectural conceptualization of the four-stage MET-CLA system; (2) quantitative performance modeling of WUE, PUE, cooling energy, and thermal behavior at 40–45  $^{\circ}\text{C}$ ; and (3) comparative evaluation against traditional evaporative cooling and conventional lower-temperature DTC systems [1], [3], [11]. Thermophysical properties of PG-25 were taken from commercial heat-transfer fluid data sheets [9], [10]. Cold-plate and flow-rate parameters were drawn from experimental and industry studies on direct-to-chip cooling [8], [14], [20]. CDU design characteristics were obtained from manufacturer documentation [15]. Climate and operational benchmarks came from published industry analyses [4], [6]. Sustainability and life-cycle impact data were referenced from recent peer-reviewed assessments of data-center cooling technologies [2], [3], [11], [26].

#### Limitations:-

This study is modeling-based and does not include full-scale experimental validation of a complete MET-CLA prototype. Results are therefore indicative performance envelopes rather than measured field data. The validity of the findings rests on established thermal-resistance and effectiveness-NTU formulations, manufacturer fluid properties, published cold-plate parameters, and cross-comparison with industry reference designs operating near 45  $^{\circ}\text{C}$  [4], [8], [13]. Pure dry-cooler operation may still require limited mechanical or adiabatic assistance in hot-humid climates [17], [21]. The life-cycle discussion focuses on the operational phase; manufacturing and end-of-life impacts are not fully quantified [26], [27]. Predictive control performance depends on forecast quality, which was idealized in the present model [18], [19].

#### Results and Discussion:-

Performance of MET-CLA is compared with conventional evaporative cooling and DTC cooling at low coolant temperature ( $\leq 32\text{ }^{\circ}\text{C}$ ). Results are reported for a benchmark 100 MW IT load in temperate, hot-dry, and mixed climates.

#### Water and Power Effectiveness:-

Hyperscale operators have reported large historical reductions in fleet-average WUE through successive cooling-design generations, including closed-loop chip-level designs that consume no evaporative water during normal operation [6], [28]. Table 3 presents the modeled performance indicators for a 100 MW IT load under temperate climate conditions. Under

temperate and hot-dry conditions, MET-CLA produced an on-site WUE in the range 0.02–0.05 L/kWh. This represents a reduction of more than 97% relative to typical evaporative values (1.60–1.85 L/kWh) and is also substantially lower than the conventional DTC range (0.28–0.42 L/kWh). Concurrently, PUE improved to 1.06–1.11. These outcomes are consistent with published elevated-temperature liquid-cooling studies [12], [24], [29].

**Table 3:- Comparative performance metrics (100 MW IT load, temperate climate)**

| Metric                 | Traditional Evaporative | Conventional DTC | MET-CLA (40–45 °C) |
|------------------------|-------------------------|------------------|--------------------|
| On-site WUE (L/kWh)    | 1.60 – 1.85             | 0.28 – 0.42      | 0.02 – 0.05        |
| Annual water (Mgal/MW) | 2.1 – 2.5               | 0.35 – 0.55      | < 0.06             |
| PUE                    | 1.28 – 1.38             | 1.12 – 1.17      | 1.06 – 1.11        |
| Cooling energy savings | Baseline                | 28 – 38%         | 48 – 58%           |




**Fig. 3:- Comparison of water usage effectiveness and power usage effectiveness across the three cooling architectures.**

#### Free-Cooling Hours:-

Elevated coolant temperature substantially increased the number of hours during which ambient air could provide cooling without mechanical refrigeration [12], [17], [25]. Modeled free-cooling hours for MET-CLA reached 7,200–8,100 h/yr in temperate climates, 6,800–7,600 h/yr in hot-dry climates, and 5,900–6,800 h/yr in mixed climates (Table 4). These values are 1.5–3 times higher than those obtained for conventional DTC and more than three times higher than traditional evaporative systems, directly justifying the observed reductions in chiller runtime and cooling energy.

**Table 4:- Estimated annual free-cooling hours by climate type**

| Climate   | Traditional Evaporative | Conventional DTC | MET-CLA       |
|-----------|-------------------------|------------------|---------------|
| Temperate | 1,800 – 2,600           | 4,200 – 5,100    | 7,200 – 8,100 |
| Hot-dry   | 2,100 – 2,900           | 4,800 – 5,600    | 6,800 – 7,600 |
| Mixed     | 1,500 – 2,300           | 3,900 – 4,700    | 5,900 – 6,800 |




**Fig. 4:- Annual free-cooling hours by climate type for the three cooling architectures.**

#### **Cooling Energy and Thermal Behavior:-**

Modeled annual cooling energy for the 100 MW facility is shown in Table 5. MET-CLA reduced cooling energy by approximately 48–58% relative to evaporative systems and by 18–25% relative to conventional DTC. These savings are consistent with the expanded free-cooling hours in Table 4 and with the known sensitivity of cooling energy to facility water temperature (approximately 3–4% savings per 1 °C increase) [12], [29]. Junction temperatures remained within acceptable limits (typically 70–82 °C) at 40–45 °C coolant inlet when cold-plate resistance and flow rates were properly set, consistent with experimental observations at elevated supply temperatures [8], [13], [20].

**Table 5:- Estimated annual cooling energy (MWh) for a 100 MW facility**

| Architecture                     | Temperate       | Hot-dry         | Mixed           |
|----------------------------------|-----------------|-----------------|-----------------|
| Traditional evaporative          | 22,500 – 26,000 | 24,000 – 28,500 | 23,000 – 27,000 |
| Conventional DTC ( $\leq 32$ °C) | 14,800 – 17,200 | 16,200 – 19,000 | 15,500 – 18,200 |
| MET-CLA (40–45 °C)               | 10,800 – 12,600 | 12,100 – 14,300 | 11,500 – 13,500 |




**Fig. 5:- Estimated annual cooling energy consumption for a 100 MW AI data center under different architectures.**

#### Comparison with Industry Reference Designs:-

Table 6 places MET-CLA alongside published characteristics of recent elevated-temperature and zero-water industry designs [4], [6], [28]. The comparison is qualitative and is intended only to clarify where the proposed architecture sits relative to current practice.

**Table 6:- Qualitative comparison with selected industry approaches**

| Feature               | NVIDIA-style 45 °C DTC         | Microsoft Zero-Water Designs | MET-CLA (This Work)            |
|-----------------------|--------------------------------|------------------------------|--------------------------------|
| Coolant inlet target  | ≈ 45 °C                        | Closed-loop, elevated        | 40–45 °C                       |
| Heat rejection        | Dry coolers / limited chillers | Closed-loop + dry coolers    | Explicitly staged dry coolers  |
| Control approach      | Facility-level                 | Facility-level               | Predictive multi-stage staging |
| Architecture emphasis | Reference design               | Zero-water facility design   | Four-stage systems framework   |

#### Discussion:-

The modeling results indicate that organizing elevated-temperature DTC, staged dry coolers, and predictive control into a coherent four-stage architecture can substantially reduce on-site water intensity while preserving acceptable thermal margins [4], [8], [13]. The largest gains appear in temperate and hot-dry climates, where the elevated coolant temperature keeps dry coolers effective for most of the year. The quantitative consistency between the WUE, free-cooling-hour, and cooling-energy results (Tables 3–5) strengthens the internal validity of the modeled performance claims.

Energy reductions arise mainly from lower chiller runtime and more efficient part-load operation of fans and pumps [12], [18], [29]. Refrigeration cycles such as direct expansion and vapor-compression chillers remain common in data centers and are energy intensive when used as the primary heat-rejection path [1], [21]. Direct liquid cooling alternatives, including immersion

---

and spray cooling, have a growing industrial base, but they address a different integration problem than the DTC plus staged dry-cooler path examined here [2], [3], [30].

The increase in free-cooling hours is most pronounced in temperate and hot-dry regions, where ambient dry-bulb temperatures are frequently well below the elevated coolant temperature [17], [25]. Even in a mixed climate, the predictive control layer and the staged dry-cooler design translated into meaningful reductions in mechanical cooling. On implementation, MET-CLA is feasible with existing high-power GPU platforms if cold plates and CDUs are qualified for the higher-temperature regime [8], [14], [20]. Practical deployment still requires attention to fluid chemistry, material compatibility, and long-term reliability of seals and thermal interfaces at 40–45 °C [9], [10].

Because the present results are modeled rather than measured on a full prototype, they should be treated as indicative performance envelopes rather than guaranteed field outcomes. Climate dependence remains a concern in hot-humid environments, where dry coolers are less effective [17], [21]. The operational-phase focus of the present analysis also means that embodied impacts of manufacturing and end-of-life treatment are not fully quantified [26], [27].

### **Challenges and Opportunities:-**

#### **Challenges:-**

The principal challenges remain the high water use of evaporative systems [1], [5], rising AI rack densities [2], [11], residual chiller energy at low coolant temperatures [4], [12], climate limits on free cooling [17], incomplete system-level integration of existing components [3], and reliability risks at elevated temperature [9], [20]. Evaporative towers can use 1.5–2.6 million gallons of water per megawatt of IT load per year, creating sustainability and permitting pressure in water-stressed regions [5], [6]. Rack densities frequently exceed 100–200 kW, rendering air cooling thermally inadequate [1], [2], [12]. Lower coolant temperatures often keep mechanical chillers in service for large portions of the year [12], [29]. In hot or humid regions the temperature difference between ambient air and the coolant becomes insufficient, forcing greater reliance on mechanical cooling or evaporative assistance [17], [21].

#### **Opportunities:-**

Raising the coolant inlet temperature to 40–45 °C enables dry-cooler rejection for most of the year and can reduce on-site WUE to approximately 0.02–0.05 L/kWh in suitable climates [4], [6]. Modeling indicates cooling-energy reductions of roughly 45–60% versus traditional evaporative systems and meaningful savings versus lower-temperature DTC [12], [24]. The combination of DTC, elevated temperatures, and staged dry coolers supports facility PUE values typically in the range 1.06–1.12 [24], [29]. The four-stage structure supports phased deployment in both new and existing facilities [15], [16]. The approach is consistent with elevated-temperature reference designs and hyperscaler initiatives aimed at water-positive or zero-water data centers [4], [6], [28].

### **Recommendations:-**

#### **Data Center Operators:-**

Operators planning new AI facilities or major retrofits should prioritize elevated coolant temperature designs (40–45 °C) combined with closed-loop DTC cooling and dry-cooler-based heat rejection [4], [6]. A staged implementation strategy is advisable: first deploy high-density AI racks with cold plates and primary CDUs, then integrate staged dry coolers and predictive control [15], [16]. Careful fluid management—properly inhibited propylene glycol-water mixtures, continuous chemistry monitoring, and compatible materials—is essential for long-term reliability at elevated temperatures [9], [10].

#### **Technology Developers:-**

Cold-plate and CDU manufacturers should accelerate validation of components specifically optimized for 40–45 °C operation, including lower thermal resistance, acceptable pressure drop, and material compatibility under continuous elevated-temperature service [8], [14], [20]. Dry-cooler manufacturers are encouraged to enhance staging controls and variable-speed fan efficiency, while keeping hybrid adiabatic modes as backup rather than the primary path [17], [21]. Closer collaboration among chip manufacturers, server OEMs, and cooling-system providers is needed so that future accelerators are thermally qualified for elevated coolant inlet temperatures from the design stage [4], [11].

#### **Future Research:-**

Future research should move beyond modeling to instrumented pilot deployments across multiple climate zones [13], [25]. More comprehensive life-cycle assessment studies are needed that include embodied impacts of manufacturing, installation, and end-of-life treatment of cold plates, CDUs, and dry coolers [26], [27]. Research should also explore hybrid architectures that can maintain near-zero water performance in hot-humid climates, possibly through selective and minimal use of adiabatic assistance, and should refine AI-supported predictive control that combines machine learning with physics-based thermal models [11], [18], [19].

---

## Standards and Policy:-

Industry standards organizations are encouraged to update thermal guidelines and water-efficiency metrics to explicitly recognize elevated-temperature closed-loop systems and near-zero WUE designs [1], [28]. Clear classification of coolant temperature regimes (for example, W40 and W45 classes) and standardized testing protocols would support broader adoption and fair performance comparison [4], [16]. Policymakers evaluating data-center projects should consider differentiated water-use requirements that reward facilities adopting verified near-zero-water cooling architectures [5], [6].

## Conclusions:-

This paper has introduced the Multi-Stage Elevated Temperature Closed-Loop Architecture (MET-CLA) as a systems-level approach aimed at near-zero on-site water use in high-density AI data centers. The modeled on-site WUE of 0.02–0.05 L/kWh under temperate and hot-dry conditions corresponds to a reduction of more than 97% relative to traditional evaporative systems (1.60–1.85 L/kWh). The same modeling framework shows PUE values of 1.06–1.11 and cooling-energy reductions of 48–58% compared with the evaporative baseline. These energy and water savings are directly supported by expanded free-cooling hours of 7,200–8,100 h/yr in temperate climates, which reduce mechanical chiller runtime. Junction-temperature results further confirm that thermal margins remain acceptable at 40–45 °C coolant inlet temperatures when cold-plate resistance and flow rates are properly specified [8], [13], [20].

Although the present work is modeling-based rather than experimentally validated at full scale, the quantitative results across multiple climate scenarios provide a coherent and internally consistent justification for the conclusion that elevated-temperature closed-loop architectures with staged dry coolers and predictive control offer a practical pathway toward near-zero on-site water use for high-density AI facilities.

## References:-

- [1] R. Kong, H. Zhang, M. Tang, H. Zou, C. Tian, and T. Ding, “Enhancing data center cooling efficiency and ability: A comprehensive review of direct liquid cooling technologies,” *Energy*, vol. 308, 2024, Art. no. 132846.
- [2] N. Wang, Y. Guo, C. Huang, and S. Shao, “Advances in direct liquid cooling technology and waste heat recovery for data center: A state-of-the-art review,” *Journal of Cleaner Production*, vol. 477, 2024, Art. no. 143872.
- [3] C. Nadjahi, H. Louahlia, and S. Lemasson, “A review of thermal management and innovative cooling strategies for data centres,” *Sustainable Cities and Society*, vol. 37, pp. 517–534, 2018.
- [4] NVIDIA, “Hotter than a hot tub: The 45 °C breakthrough to cool AI’s biggest machines,” NVIDIA Blog, 2026.
- [5] Johnson Controls, “Zero water usage where it really matters: Your data center,” technical brief, 2026.
- [6] Microsoft, “Inside Microsoft’s two-decade push to cut water intensity while scaling for growth,” Microsoft Blog, 2026.
- [7] A. C. Kheirabadi and D. Groulx, “Cooling of server electronics: A design review of existing technology,” *Applied Thermal Engineering*, vol. 105, pp. 622–638, 2016.
- [8] A. Heydari et al., “Experimental evaluation of direct-to-chip cold plate liquid cooling for high-heat-density data centers,” *Applied Thermal Engineering*, vol. 239, 2024, Art. no. 122122.
- [9] Dynalene, Dynalene PG-D Series Technical Data Sheet, Dynalene Inc., 2024.
- [10] The Dow Chemical Company, DOWFROST LC Heat Transfer Fluid Technical Data, 2024.
- [11] D. Al Kez, A. M. Foley, F. W. B. M. Hasan Wong, A. Dolfi, and G. Srinivasan, “AI-driven cooling technologies for high-performance data centres: State-of-the-art review and future directions,” *Sustainable Energy Technologies and Assessments*, vol. 82, 2025, Art. no. 104511.
- [12] Y. Zhang, H. Li, and S. Wang, “The global energy impact of raising the space temperature for high-temperature data centers,” *Cell Reports Physical Science*, vol. 4, no. 9, 2023, Art. no. 101624.
- [13] M. Stahlhut, M. Sekulla, P. Großöhme, A. Auweter, and T. Urbaneck, “Data centers with direct liquid-cooled servers: Experimental analysis and modeling,” *Energy Science & Engineering*, vol. 13, no. 7, pp. 3605–3619, 2025.
- [14] Electronics Cooling, “High-performance copper tungsten cold plates for direct-to-chip single-phase liquid cooling,” *Electronics Cooling*, 2026.
- [15] Trane, Coolant Distribution Unit Brochure, Trane Technologies, 2024.
- [16] ASHRAE, *Liquid Cooling Guidelines for Datacom Equipment Centers*, 2nd ed. Atlanta, GA: ASHRAE, 2014.
- [17] H. Ghadiri Moghaddam et al., “Radiative free cooling for energy and water saving in data centers,” *Energy*, 2024.
- [18] M. K. Patterson, “On energy efficiency of liquid cooled HPC datacenters,” in *Proc. 15th IEEE Intersociety Conf. Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2016, pp. 385–390.
- [19] J. Cho and B. S. Kim, “Evaluation of air management system’s thermal performance for superior cooling efficiency in high-density data centers,” *Energy and Buildings*, vol. 43, no. 9, pp. 2145–2155, 2011.
- [20] Electronics Cooling, “Direct liquid cooling for high-compute servers,” *Electronics Cooling*, 2025.
- [21] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, “A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities,” *Renewable and Sustainable Energy Reviews*, vol. 31, pp. 622–638, 2014.

- 
- [22] J. R. Thome, *Engineering Data Book III*, Wolverine Tube Inc., 2004.
- [23] F. P. Incropera, D. P. DeWitt, T. L. Bergman, and A. S. Lavine, *Fundamentals of Heat and Mass Transfer*, 7th ed. Hoboken, NJ: Wiley, 2011.
- [24] Vertiv, “Quantifying data center PUE when introducing liquid cooling (summary of ASME analysis),” technical report, 2023.
- [25] M. Iyengar et al., “Server liquid cooling with chiller-less data center design to enable substantial energy savings,” in *Proc. IEEE SEMI-THERM*, 2012, pp. 212–217.
- [26] N. Lei and E. Masanet, “Global data center energy and water use: Trends and opportunities,” *Joule*, related industry energy analyses, 2020–2024.
- [27] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, “Recalibrating global data center energy-use estimates,” *Science*, vol. 367, no. 6481, pp. 984–986, 2020.
- [28] The Green Grid, *Water Usage Effectiveness (WUE): A Green Grid Data Center Sustainability Metric*, White Paper, 2011.
- [29] S. Zimmermann, I. Meijer, M. K. Tiwari, S. Paredes, B. Michel, and D. Poulikakos, “Aquasar: A hot water cooled data center with direct energy reuse,” *Energy*, vol. 43, no. 1, pp. 237–245, 2012.
- [30] X. Yuan, X. Zhou, Y. Pan, R. Kosonen, H. Cai, Y. Gao, and Y. Wang, “Phase change cooling in data centers: A review,” *Energy and Buildings*, vol. 236, 2021, Art. no. 110753.