Advanced Cooling Strategies and Waste Heat Recovery for Immersion-Cooled Data Centers in Hot Climates

Chantal Maatouk[a], Richard Khalil[b]

[a] Ecole Supérieure d'Ingénieurs de Beyrouth, Faculty of Engineering and Architecture, Saint Joseph University of Beirut, Beirut, Lebanon, Chantal.maatouk@usj.edu.lb
[b] Ecole Supérieure d'Ingénieurs de Beyrouth, Faculty of Engineering and Architecture, Saint Joseph University of Beirut, Beirut, Lebanon, richard.khalil@net.usj.edu.lb

Abstract:

This study explores advanced cooling solutions for a 2.3 MW immersion-cooled data center operating under extreme Gulf climate conditions, with the objectives of improving energy efficiency, reducing carbon dioxide (CO_2) emissions, and evaluating opportunities for waste heat recovery. The research addresses the challenge of rejecting multi-megawatt heat loads in environments where high ambient and wet-bulb temperatures limit conventional air-based cooling. It also considers the constraint posed by low-grade waste heat (approximately 45–50 °C), which restricts direct reuse without temperature enhancement. A thermodynamic modeling framework was developed to represent the thermal resistance path from chip junction to ambient, applying steady-state energy balances and feasibility criteria. Four heat-rejection strategies were assessed: adiabatic fluid coolers, mechanical-draft cooling towers, air-cooled chillers, and single-effect absorption chillers. Waste heat recovery pathways, including heat pump boosting and district heating, were also evaluated. Results show that a single adiabatic cooler cannot meet peak cooling demand, whereas cooling towers maintain a 2 K approach to wet-bulb temperature and ensure safe device operation. Absorption chillers driven solely by immersion waste heat are infeasible due to insufficient temperature and capacity; however, the addition of a heat pump enables partial operation at reduced efficiency. Evaporative cooling systems reduce electricity consumption and CO_2 emissions by 83–88% compared to air-cooled chillers but increase water use. Economic analysis indicates payback periods of less than one year and levelized cooling costs as low as 10.7 USD per megawatt-hour thermal (MWh_{th}). Overall, immersion cooling combined with optimized evaporative heat rejection offers a robust pathway to low Power Usage Effectiveness (PUE) in hot climates.

Keywords:

Immersion Cooling, Data Center Energy Efficiency, Adiabatic Cooling Systems, Cooling Towers, Waste Heat Recovery, Absorption Chiller, Thermodynamic Simulations, Power Usage Effectiveness (PUE), Environmental Sustainability.

1. Introduction

In recent years, there has been a substantial increase in global data center power consumption. With the growth of digital services, cloud computing, and AI-driven applications, this trend is set to continue. In spite of increasing efficiency, the average Power Usage Effectiveness (PUE) for the industry is still close to 1.5 [1,2]. In this context, modern processors and accelerators have thermal design power consumption increasing rapidly, and high-end graphics processing units have power consumption in excess, pushing rack power densities to tens of kW [3,4]. This is a challenge for sustainability and thermal management, especially in regions where ambient and wet-bulb temperatures are high, making it thermodynamically difficult to reject heat.

Immersion cooling has been recognized as a viable solution to overcome the challenges, providing a way to effectively manage high-density computation. Immersion cooling eliminates air movement at the server level, minimizes thermal resistance between the chip and the coolant, and allows for safe operation [5,6,7]. In addition to efficiency, immersion cooling minimizes water consumption compared to traditional evaporative cooling methods, eliminates hotspots, increases reliability, and allows for the utilization of warm water secondary systems (approximately 45–50 °C) for simplified rejection and utilization of waste heat.

In contrast, traditional air-based cooling architectures face intrinsic performance limits as heat densities increase. Airflow maldistribution, recirculation, and bypass phenomena significantly degrade thermal uniformity

[4, 5], while the volumetric airflow required to cool racks above 20–30 kW becomes increasingly impractical [6, 7, 8]. Even when supplemented by containment strategies, economizers, or high-capacity units, air cooling becomes insufficient under aggressive computational loads or elevated ambient temperatures. Consequently, air-based systems cannot reliably support the thermal requirements of emerging high-density environments and are increasingly replaced by liquid-based approaches [9, 10].

Deploying immersion cooled data centers in hot climates remains highly challenging due to extreme ambient temperatures, such as the Gulf region's 45 - 46 °C design conditions, which significantly reduce cooling feasibility and efficiency [1, 3]. In hot dry environments, adiabatic cooling can reduce water use by up to 70% because water is consumed only during extreme conditions [1], yet its performance still degrades under severe heat. Evaporative cooling towers, capable of producing chilled water near the wet bulb temperature (20–27 °C) on hot days, also lose effectiveness as ambient wet bulb levels rise [2]. Ultimately, climate plays a decisive role: evaporative cooling performs best in hot dry regions, while its efficiency drops in hot humid climates despite high water consumption [3].

While immersion cooling resolves component-level thermal limits, the main challenge shifts to heat rejection, particularly in hot climates such as the Gulf region where ambient temperatures reach 45–46 °C. Under these conditions, dry coolers lose efficiency and even evaporative systems degrade as wet-bulb temperatures rise. In addition, the available waste heat (≈45–50 °C) remains too low to drive single-effect absorption chillers, which require higher temperatures [10]. These constraints emphasize the need to carefully evaluate different heat-rejection strategies in terms of performance, feasibility, and heat-recovery potential.

In this context, the present study examines four alternative cooling configurations for a 2.3 MW single-phase immersion-cooled data center: Adiabatic fluid coolers, which operate primarily in dry mode and employ evaporative pre-cooling under peak ambient conditions; Mechanical-draft cooling towers, which leverage evaporative cooling to approach the ambient wet-bulb temperature and thereby enhance thermal performance; Air-cooled chillers, which provide controlled coolant temperatures independent of outdoor conditions but at the cost of higher electrical consumption; Single-effect absorption chillers, evaluated for their ability to convert immersion waste heat into useful cooling, subject to generator-temperature and feasibility limitations.

The paper begins with a discussion of various immersion cooling technologies, data center thermal management techniques, and heat recovery methods. Subsequently, a unified thermodynamic model of heat transfer from the chip to the ambient environment is presented, followed by the introduction of various performance models of the cooling architectures. The four cooling architectures are compared under peak and yearly conditions using a typical meteorological year data set representative of Saudi Arabia, thereby comparing energy, water, and thermal feasibility of the cooling architectures. The comparative results of all the cooling architectures are obtained from yearly simulations, ensuring that the comparison is realistic.

2. Literature review

2.1. Overview of Immersion Cooling Technologies

Immersion cooling refers to the practice of submerging electronic components directly in a thermally conductive, electrically insulating liquid. This approach bypasses the thermal bottlenecks associated with forced-air systems and allows heat to be transferred directly from components to a dielectric fluid with substantially higher heat capacity than air. Modern implementations include single-phase immersion, where the coolant remains in liquid form, and two-phase immersion, where the fluid undergoes controlled boiling at the component surface, providing high latent-heat-driven heat transfer [5-7].

Recent industrial analyses highlight that immersion cooling has become a leading candidate for high-density and AI-oriented data centers, as it reduces reliance on air-handling equipment and delivers greater thermal uniformity [8]. Market projections show rapid adoption: the global liquid-cooling market grew 52% in 2023 and is expected to exceed $4.8 B by 2028, driven largely by immersion systems that enable power densities exceeding 500 kW per rack [9-10].

2.2. Benefits and Challenges of Single-Phase Immersion Cooling

Single-phase immersion cooling offers clear benefits. Direct submersion in a dielectric liquid enables uniform heat removal and stable thermal conditions, while eliminating server fans and reducing mechanical cooling demand. This allows low PUE values, around 1.03 in optimized systems [11], with relatively simple and stable operation over time [6].

Despite these advantages, several challenges remain. Single-phase fluids have lower heat-transfer performance than two-phase systems, which can limit use in very high-density applications. Long-term concerns such as fluid stability, material compatibility, and leakage risks also persist. [12-14**Error! Reference source not found.**]. In addition, implementing immersion cooling often requires redesigned infrastructure and maintenance practices, making retrofitting existing air-cooled data centers complex and potentially disruptive. [15-16].

2.3. Overview of Dielectric Fluids Used in Immersion Systems

Single-phase immersion cooling relies on dielectric fluids whose properties directly affect performance and reliability. Mineral oils are cost-effective but less efficient, while synthetic hydrocarbons offer better stability and cooling capacity [17]. Ester-based fluids are more environmentally friendly, and fluorinated fluids provide the highest performance, though at higher cost. Key properties such as thermal conductivity, heat capacity, and viscosity govern heat transfer and pumping requirements, making careful fluid selection essential for efficient and reliable operation [16].

2.4. Environmental Assessment Metrics in Data Center Studies

Environmental and economic performance in data centers is typically assessed through a set of standardized metrics, among which Power Usage Effectiveness (PUE) remains the most widely used. Defined as the ratio of total facility energy to IT energy consumption, PUE in conventional air-cooled facilities commonly approaches 1.5, whereas immersion-cooled installations achieve markedly lower values due to the elimination of server fans [11]. Despite its widespread use, PUE provides only a partial view of efficiency: it excludes water consumption, carbon-intensity effects, and workload-level computational efficiency; it can distort comparative assessments in climates where mechanical cooling operates only seasonally; and it does not account for the benefits of waste-heat reuse, even when such reuse reduces overall energy demand beyond the data-center boundary [12-13].

Carbon performance is evaluated through CO_2 emissions and depends on PUE, the local grid's carbon intensity, and water usage. In Gulf regions, water production relies on energy-intensive desalination, which further contributes to emissions. Immersion-cooled facilities powered by low-carbon electricity can therefore achieve significantly lower overall emissions. [15].

3. System architecture and description

The facility considered in this study is a 2.3 MW single phase immersion cooled data center, designed for operation under the extreme thermal conditions of Gulf climates. The heat is transferred from the dielectric fluid to a water loop and then rejected to the ambient. Different heat rejection options are evaluated:

A. Adiabatic cooler,
B. Cooling tower,
C. Air-cooled chiller, and
D. Waste heat recovery system pathways, using absorption chillers

Figure 1 shows a single-phase immersion cooling system where heat from IT equipment is absorbed by a dielectric fluid, transferred via a heat exchanger to a water loop, and then rejected to the environment through different cooling options

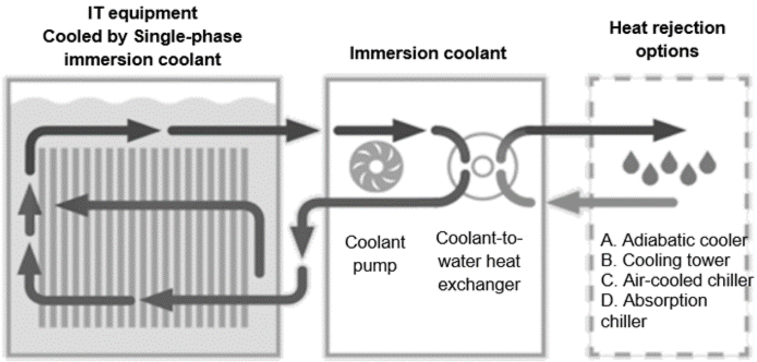

Figure 1. *Single phase liquid immersion cooling*

3.1. Immersion-Cooled Data Center Architecture

The data-center cooling system is organized into 46 single-phase immersion modules, each rated at 50 kW, where servers are fully submerged in a dielectric fluid that removes heat through direct-contact convection after a short conduction path characterized by a 0.01 K/W junction-to-case resistance. Each tank includes an internal plate-type coolant-to-water heat exchanger, through which the dielectric fluid—circulating at 1.25 kg/s under a 35 kPa pressure drop and 70% pump efficiency—transfers thermal energy to the facility's secondary water loop while maintaining coolant temperatures within the 35–50 °C operating envelope. The cooled water is subsequently routed to the external heat-rejection system, ensuring a closed-loop, single-phase cooling process under high-density and high-ambient operating conditions.

3.2. Immersion coolant loop

The hot dielectric coolant exiting the immersion tank is directed through a plate-type coolant-to-water heat exchanger, where heat is transferred to the secondary water loop. The exchanger is modeled using the effectiveness–NTU method and is constrained to a minimum temperature pinch of 5 K to ensure feasible and efficient heat transfer under all operating conditions.

3.3. Dielectric Fluid

Single-phase immersion cooling uses dielectric liquids that provide electrical insulation, thermal stability, and good heat-transfer properties. Common options include synthetic hydrocarbons and natural ester-based fluids. Table 1 presents the properties of some commercial dielectric coolant. The numerical values used in the thermal model fall within these reported ranges and therefore accurately reflect realistic fluids.

Table 1 *Dielectric fluids properties [18]*

Property	CompuZol™ IM2020	NatureCool™ 2000	Thermasolv TM CF2	Used in the study
Specific heat (kJ/kg·K)	2.107	2.3076	1.087	2.0
Thermal conductivity (W/m·K)	0.143	0.1644	0.112	0.15
Density (g/mL @25 °C)	0.81	0.9206	1.815	0.8
Kinematic viscosity @ 40°C (mm^2/s)	9.8	31	1.35	9.8
Flash point (°C)	180–202	325	No flash Point	-
Chemical family	Synthetic hydrocarbon	Natural ester		-

Their high heat capacity and stability at elevated temperatures make these fluids suitable for warm-water, high-density immersion cooling in hot environments. Their low viscosity supports reliable flow and heat transfer. In addition, their non-conductive nature and strong dielectric properties ensure safe operation when directly in contact with electronic components.

3.4. Heat rejection Systems

3.4.1. Adiabatic cooler

An adiabatic fluid cooler combines evaporative pre-cooling with sensible heat exchange. Air passing through wetted pads is cooled toward the wet-bulb temperature before reaching the coil, enabling sub-ambient cooling without refrigeration, though performance remains limited by wet-bulb conditions. As both dry- and wet-bulb temperatures increase, the heat transfer driving force decreases, reducing cooling capacity.

3.4.2. Cooling tower

The cooling tower relies on evaporative heat and mass transfer. It removes heat by bringing warm water into direct contact with air, allowing a small portion to evaporate and cool the rest. The cooled water is recirculated, while saturated air is discharged. Make-up water is needed to replace losses [19].

3.4.3. Air-cooled Chiller

Using a vapor-compression refrigeration cycle, an air-cooled chiller rejects heat to the surrounding air through fans and finned condenser coils. High ambient temperatures in hot climates have a major impact on its performance because they raise condensing pressure and decrease efficiency. Air-cooled chillers typically achieve a coefficient of performance (COP) in the range of 2.5 to 3.5 under hot conditions (35–45 °C ambient), which is lower than water-cooled systems. Despite their decreased efficiency, their primary benefits; simplicity, reduced water consumption, and ease of installation, make them appropriate for areas with limited water resources [20].

3.4.3. Waste heat recovery

Waste heat recovery is key to improving energy efficiency in data centers. In immersion-cooled systems, most of the electrical input is converted into heat, producing a steady thermal output at relatively high temperatures (around 45–50 °C). However, this heat remains low-grade, with limited thermodynamic quality, making direct utilization challenging. [21,23-24].

To make effective use of this available thermal energy, several recovery pathways can be explored. These include thermally driven cooling, temperature upgrading through heat pumps, and direct reuse in heating applications. Each pathway offers distinct advantages, but also faces constraints; primarily linked to temperature levels and fundamental thermodynamic limits.

3.4.3.1. Single-effect absorption chiller

One option is to integrate a single-effect absorption chiller, converting waste heat into useful cooling. At first glance, this approach is appealing, as it could offset part of the electrical cooling demand. In practice, however, it falls short. Lithium bromide ($LiBr–H_2O$) absorption chillers require generator temperatures in the range of 70–95 °C, well above the available 45–50 °C heat source. This gap is not merely technical—it is thermodynamic. Without sufficient temperature, the absorption cycle cannot be sustained, making direct integration infeasible unless the heat is first upgraded [22].

3.4.3.2. Single-effect absorption chiller combined with heat pump

The use of a heat pump combined to an absorption chiller present a solution for heat recovery and upgrade. By lifting the temperature of the waste heat from roughly 50 °C to 80–90 °C, a heat pump effectively unlocks new possibilities [26]. With a typical coefficient of performance (COP) between 3.5 and 5.0, this approach offers an efficient means of upgrading low-grade heat. The resulting higher-temperature stream can then be used either to drive an absorption chiller or to supply heating applications directly. This solution comes at a cost. When combined with an absorption chiller, the system may not outperform conventional vapor-compression solutions, so its benefits must be balanced against higher complexity and cost.

4. Mathematical modeling

The system is modeled as a closed-loop heat removal chain, extending from the IT equipment to the ambient environment. The total heat generated by the servers is assumed constant and fully removed under steady-state conditions.

4.1. Immersion cooling system

In the single-phase immersion cooling system, the thermal behavior is represented using an equivalent resistance model:

$$T_j = T_d + Q_{IT}.R_{th}, \quad (1)$$

R_{th} accounts for chip-to-fluid heat transfer thermal resistance, expressed as :

$$R_{th} = R_{jc} + R_{TIM} + R_{hs} + 1/(h.A), \quad (2)$$

Experimental studies on single-phase immersion cooling systems report coolant flow rates in the range of 2–8 L/min per module, using dielectric fluids with densities of approximately 800–850 kg/m³. The corresponding convective heat transfer coefficient is typically between 200 and 1000 W/m²·K, increasing with flow rate due to enhanced forced convection through heat sink structures [].

The energy balance for the coolant loop is:

$$\dot{Q}_{IT} = \dot{m}_d C_{p,d}\left(T_{d,out} - T_{d,in}\right), \quad (3)$$

Pressure Drop is assumed to be equal to 35 kPa and Pump Power used to circulate the dielectric coolant is calculated based on the following equation:

$$P_{p,d} = \frac{\Delta P.\dot{m}_d}{\eta_{p,d}.\rho_d}, \quad (4)$$

4.2. Coolant-to-water heat exchanger

The heat extracted by the dielectric coolant is transferred to a secondary water loop through a plate heat exchanger:

$$\dot{Q}_{HX} = \dot{Q}_{IT}, \quad (5)$$

$$\dot{Q}_{HX} = \dot{m}_f C_{p,f}\left(T_{f,out} - T_{f,in}\right) = \dot{Q}_{IT}, \quad (6)$$

The exchanger is modeled using the effectiveness–NTU method [25]:

$$C_{min} = \min\,(\dot{m}_d C_{p,d};\ \dot{m}_f C_{p,f}), \quad (7)$$

$$C_r = C_{min}/C_{max}, \quad (8)$$

$$\varepsilon = \frac{1 - \exp[-NTU(1-C_r)]}{1 - C_r \exp[-NTU(1-C_r)]}\,, \quad (9)$$

$$NTU = \frac{UA}{C_{min}}, \quad (10)$$

$$Q_{HX} = \varepsilon\, C_{min}\left(T_{h,in} - T_{c,in}\right) \quad (11)$$

Pressure Drop is assumed to be equal to 60 kPa and Pump Power used to circulate the cooling fluid is calculated based on equation (same as above).

4.3. Heat rejection strategies

4.3.1. Adiabatic fluid cooler

The adiabatic cooler combines evaporative and sensible cooling. The system is modeled in two stages: air cooling across the pad and sensible heat exchange in the coil.

The outdoor air state is first determined from the ambient dry-bulb temperature $T_{db,amb}$, relative humidity expressed in (%) ϕ_{amb}, and atmospheric pressure $Patm$. The humid air properties at the inlet and outlet of the wet pad are calculated based on the psychometric chart relations.

The operating mode is selected according to the ambient conditions. In dry mode, no evaporation occurs and the entering air temperature to the coil is simply, whereas in adiabatic mode the entering air is precooled before the coil according to

$$T^*_{a,in} = T_{db,amb} - \eta_{pad}\left(T_{db,amb} - T_{wb,amb}\right), \qquad (12)$$

where η_{pad} is the adiabatic saturation effectiveness. Typical values of η_{pad} range between 0.6 and 0.8, depending on airflow conditions [27].

The evaporated water rate is obtained from the moisture balance:

$$\dot{m}_{w,evap} = \dot{m}_a\left(\omega^*_{a,in} - \omega_{amb}\right), \qquad (13)$$

and the associated latent heat is

$$\dot{Q}_{lat} = \dot{m}_{w,evap}.h_{fg}, \qquad (14)$$

For an ideal adiabatic saturator, the enthalpy variation across the adiabatic section is almost zero.

$$\dot{Q}_{ad} = \dot{m}_a.\left(h^*_{a,in} - h_{a,amb}\right) \approx 0, \qquad (15)$$

The power consumption of an adiabatic cooler is primarily associated with the operation of circulating pumps and axial or centrifugal fans required to drive air through the heat exchanger. The hydraulic power required for water circulation is calculated based on the volumetric flow rate and the total pressure drop across the hydraulic loop, expressed in equation (4)

The total pressure drop ranges between 30 and 80 kPa, depending on system layout and flow rate, while pump efficiencies are generally in the range of 0.6 to 0.75.

The fan power consumption is calculated from the air volumetric flow rate and the pressure rise required to overcome resistance across the finned heat exchanger and wetted media:

$$P_{fan} = \dot{V}_{air}\,\Delta P_{air}/\eta_{fan} \qquad (16)$$

where $\dot{V}_{air}$ is the air flow rate, ΔP_{air} is the air-side pressure drop, and η_{fan} is the fan efficiency. For adiabatic coolers, the air-side pressure drop typically ranges from **100 to 300 Pa**, increasing when the adiabatic (wet) mode is activated due to the presence of wetted pads. Fan efficiencies are generally between **0.5 and 0.7**.

4.3.2. Cooling tower

The cooling tower is modeled based on a combined heat and mass transfer process in which heat is rejected through both sensible cooling and evaporation. The tower performance is primarily governed by the ambient wet-bulb temperature, which defines the lower limit of achievable water temperature. The cooling tower performance is governed by the Merkel approach [19]:

$$\frac{L}{G}\frac{dT}{h^*(T_w) - h_{air}} = \frac{1}{KaV}dV \qquad (17)$$

where L/G is the liquid-to-air mass flow ratio, $h^*(T_w)$ is the saturated air enthalpy at water temperature, and KaV is the tower mass transfer characteristic. For system-level modeling, a simplified effectiveness approach is more practical.

The cooling tower effectiveness is defined as:

$$\varepsilon_{tower} = \frac{T_{w,in} - T_{w,out}}{T_{w,in} - T_{wb}} \qquad (18)$$

where T_{wb} is the ambient wet-bulb temperature, representing the thermodynamic limit of evaporative cooling [**Error! Reference source not found.**]. In practice, large towers achieve effectiveness values of 0.8–0.9, a lthough lower values are common depending on design and operating conditions [19].

In cooling towers, power consumption is also dominated by pumps and fans. Equations (4) and (16) are used to calculate power consumptions. Typical pressure drops in cooling towers pumps range between 40 and 100 kPa, reflecting the additional resistance introduced by spray systems and elevated hydraulic heads. Pump efficiencies are generally similar to those in adiabatic systems, typically between 0.6 and 0.8. As for the fan, the pressure drop typically lies in the range of 100 to 250 Pa, depending on tower design and air velocity. Fan efficiencies are commonly in the range of 0.5 to 0.7.

4.3.3. Air-cooled chiller

The air-cooled chiller is modeled using a Carnot-based reduced-order approach, where the coefficient of performance (COP) depends on the evaporating and condensing temperature levels. The operating conditions are defined by the chilled-water supply temperature and outdoor dry-bulb temperature, with minimum approach temperatures of 5 K assumed at both the evaporator and condenser.

The ideal performance is given by the Carnot COP, and the real chiller COP is expressed as:

$$COP \ = \eta . COP_{Carnot}, \qquad (19)$$

The compressor power consumption is then calculated from the cooling load:

$$P_{comp} = \dot{Q}_{IT}/COP, \qquad (20)$$

Auxiliary power includes fan and pump consumption, given by equations (4) and (16). This simplified formulation captures the dominant dependence of chiller performance on temperature levels and provides a practical framework for system-level analysis under varying outdoor conditions

4.4. Waste heat recovery pathways

The immersion-cooled data center produces about 2.3 MW of waste heat at 45–50 °C. Despite its magnitude, this heat is low-grade with limited thermodynamic value, meaning only a small portion can be converted into useful cooling or work. To address this, three recovery options are considered: absorption cooling, temperature boosting with heat pumps, and direct heat reuse, each evaluated based on practical feasibility.

4.4.1. Single-effect absorption chiller

A single-effect absorption chiller can convert waste heat into cooling. In this study, it is modeled in Aspen HYSYS using a LiBr–H_2O pair, where hot water from the heat exchanger drives the generator. The cycle then produces cooling through condensation and evaporation, further reducing chilled water temperature. However, its effectiveness depends on the available heat source temperature.

For typical LiBr–H_2O systems, the $COPth$ ranges between 0.65 and 0.75 under nominal conditions. While attractive in principle, this approach is severely constrained by temperature requirements. The absorption cycle requires generator temperatures of at least 70–80 °C, with optimal operation near 90 °C. The available waste heat at 45–50 °C is therefore insufficient to initiate the desorption process.

Consequently, without auxiliary temperature lift, the absorption chiller is not a feasible standalone solution

4.4.2 Single-effect absorption chiller combined with heat pump

To overcome the temperature limitation, a heat pump can be introduced to upgrade the waste heat to a higher, more usable level. The heat pump operates according to the energy balance:

$$\dot{Q}_{hot} = \dot{Q}_{cold} + \dot{W}_{comp}, \qquad (21)$$

where $\dot{Q}_{hot}$is the heat released at the condenser, $\dot{Q}_{cold}$ is the heat absorbed from the source and $\dot{W}_{comp}$ is the compressor work input. The coefficient of performance for the heat pump is:

$$COP_{HP} = \frac{\dot{Q}_{hot}}{\dot{W}_{comp}} \ , \qquad (22)$$

Theoretical performance is bounded by the Carnot limit. For a temperature lift from 45 °C to 85 °C, the ideal COP approaches 9, while practical systems achieve values in the range of 3.5–5.0. This allows the 2.3 MW of low-grade heat to be upgraded to roughly 2.5–2.8 MW of usable thermal energy at higher temperature.

When coupled with an absorption chiller, the effective system performance has an overall cooling COP of approximately 2.5–3.0. However, this remains lower than that of conventional vapor-compression chillers, indicating that the combined system may not offer clear efficiency advantages. In addition, the increased system complexity and capital cost must be carefully considered.

4.4.3. Direct heat reuse

Direct heat reuse offers a simple and efficient way to utilize waste heat, with potential recovery efficiencies of 80–100% in applications such as district heating or hot water. However, its feasibility depends on a nearby and continuous demand for low-grade heat. In practice, it is only relevant when such uses exist close to the data center; otherwise, the low temperature limits its practical value.

5. Results and Discussion

The performance of the four cooling configurations was evaluated under identical operating conditions for a 2.3 MW immersion-cooled data center. The analysis focuses on key performance indicators including PUE, total power consumption, water usage, CO_2 emissions, and operating cost. The modeling framework and assumptions follow the methodology detailed in the system architecture and thermodynamic simulation sections. Table 2 shows the annual results for the four cases.

Table 2. *Summary of annual energy, water, emissions, and OPEX results for the four cooling systems*

	Average annual PUE	Electrical consumption (MWh/yr)	Water usage (m^3/yr)	CO2 emissions (tons/yr)	Energy Cost (USD/yr)	Water Cost (USD/yr)	Total OPEX (USD/yr)
Adiabatic cooler	1.07	1481	17524	900	148141	26285	174427
Cooling tower	1.04	776	31044	509	77604	46566	124170
Air chiller	1.40	8060	0	4435	806042	0	806042
HP +Abs. chiller	1.19	3833	0	2107	383314	0	383314

5.1. PUE Comparison

The annual average PUE clearly distinguishes the performance of the four systems. The cooling tower achieves the lowest value (1.04), followed by the adiabatic cooler (1.07), while the HP–absorption system reaches 1.19. The air-cooled chiller exhibits the highest PUE (1.40), reflecting its reliance on mechanical refrigeration.

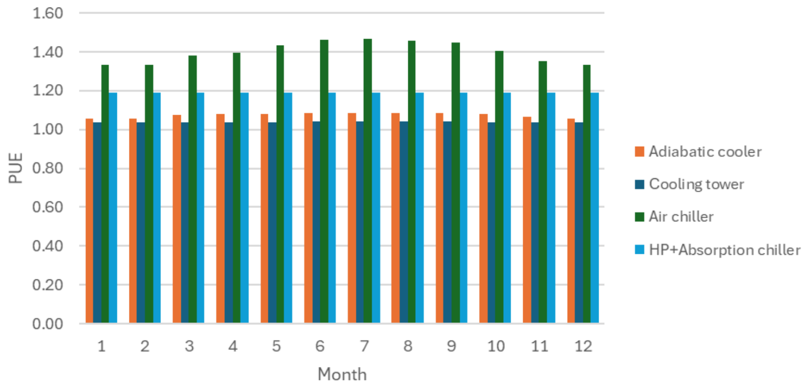

Figure. 2. *Monthly PUE profiles*

This trend is consistent with the thermodynamic behavior of each technology. Evaporative systems operate close to the ambient wet-bulb temperature, enabling efficient heat rejection with low auxiliary demand, whereas the air chiller must overcome a larger temperature lift, resulting in higher compressor power.

Seasonal effects further accentuate these differences: all systems show reduced performance in summer, but evaporative solutions remain relatively stable, while the air chiller is significantly impacted by high dry-bulb temperatures.

5.2. Water Consumption

Water consumption is a critical differentiator between technologies. The cooling tower exhibits the highest water consumption due to continuous evaporative cooling, consistent with its operating principle based on latent heat transfer. This aligns with the discussion in the literature where cooling towers achieve high efficiency at the expense of significant water usage.

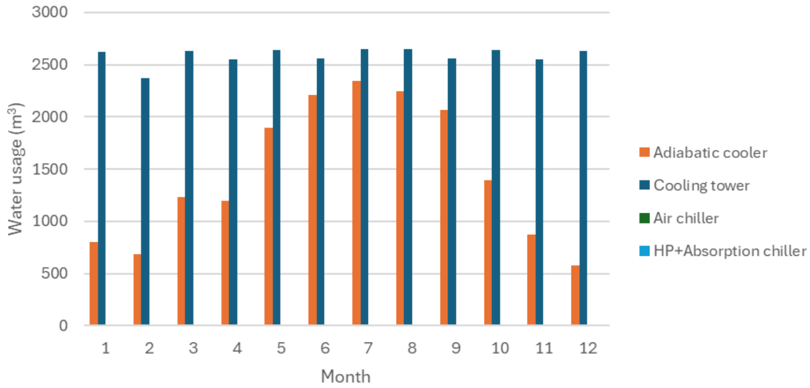

Figure. 3. *Monthly Water Usage profiles*

The adiabatic cooler reduces water usage by operating in dry mode during moderate conditions and activating evaporation only during peak loads. This results in approximately 40–45% reduction in water consumption compared to the cooling tower.

The air chiller and hybrid system eliminate water usage entirely, making them attractive in water-scarce regions such as the Gulf. However, this advantage comes at the cost of significantly higher electrical consumption.

5.3. CO_2 Emissions

CO_2 emissions are directly linked to electricity and water consumption. An electricity emission factor of approximately 0.55 kg CO_2/kWh is adopted for the Saudi grid, while water-related emissions are estimated based on desalination energy intensities (3–5 kWh/m³), resulting in about 2.2 kg CO_2/m³ for supply and 0.44 kg CO_2/m³ for treatment. Under these assumptions, the cooling tower achieves the lowest emissions due to its minimal electrical demand, whereas the adiabatic cooler shows slightly higher values, reflecting its increased auxiliary energy use.

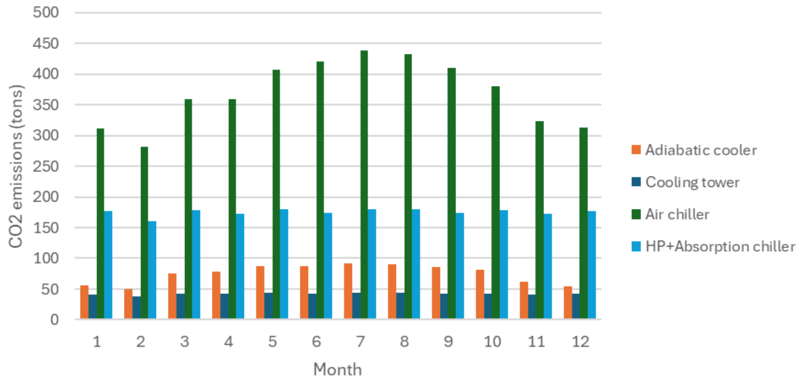

Figure. 4. *Monthly CO_2 emissions profiles*

The air chiller shows the highest emissions, exceeding 4,400 tons/year, more than eight times that of the cooling tower, highlighting the environmental penalty of compressor-based cooling in hot climates. The HP–absorption system reduces emissions relative to the air chiller but remains higher than evaporative solutions, mainly due to the additional electrical demand of the heat pump required to upgrade low-grade waste heat.

5.5. Economic Analysis

The operating expenditure (OPEX) reflects both energy and water costs. The cooling tower achieves the lowest OPEX despite its high-water use, as energy costs dominate. The adiabatic cooler offers a balanced performance with moderate costs. In contrast, the air chiller shows the highest OPEX due to its high electricity demand, making it less suitable for hot climates. The HP–absorption system reduces energy use compared to the air chiller but remains less competitive due to the added complexity and cost of the heat pump and absorption cycle.

6. Conclusion and future work

This study presents a comparative assessment of four cooling strategies for a 2.3 MW immersion-cooled data center under hot climate conditions: adiabatic cooler, cooling tower, air-cooled chiller, and a hybrid heat pump–absorption chiller system. The results confirm that system performance is primarily governed by thermodynamic limits and ambient conditions. Evaporative-based systems consistently outperform mechanical solutions, achieving the lowest PUE and electrical consumption.

Among the evaluated options, the cooling tower delivers the best overall performance, with the lowest PUE, CO_2 emissions, and operating cost, but at the expense of high-water consumption. The adiabatic cooler offers the most balanced solution, reducing water use while maintaining strong energy performance. In contrast, the air-cooled chiller, although water-free and reliable, suffers from high energy consumption, leading to increased emissions and operating costs. The hybrid heat pump–absorption system demonstrates the potential of waste heat recovery, but its performance remains limited.

These results highlight clear trade-offs between energy efficiency, water consumption, environmental impact, and cost. Cooling towers are most suitable in water-abundant regions, while adiabatic systems provide a practical compromise where water resources are moderately constrained. Fully dry or hybrid systems become relevant only in water-scarce environments despite their lower efficiency.

Future work should focus on improving waste heat recovery and exploring advanced cooling solutions. This includes the integration of two-phase immersion cooling for enhanced heat transfer, as well as the use of machine learning techniques to optimize system operation under dynamic conditions. Further research is also needed on low-temperature-driven technologies and hybrid configurations to better utilize low-grade waste heat in data center applications.

Nomenclature

Letter symbols

A	heat transfer area, m²
c_p	specific heat capacity, J/(kg·K)
C	thermal capacity rate, W/K
Cmin	minimum heat capacity rate, W/K
Cmax	maximum heat capacity rate, W/K
h	convective heat transfer coefficient, W/(m²·K)
h_a	moist air enthalpy, kJ/kg(da)
$h^*(T_w)$	Saturated air enthalpy at water temperature, kJ/kg
h_(fg)	latent heat of vaporization, kJ/kg
KaV	Cooling tower mass transfer characteristic
L/G	Liquid-to-air mass flow ratio
$\dot{m}$	mass flow rate, kg/s
NTU	number of transfer units, –
P	power, W
P_{atm}	atmospheric pressure, Pa
Δp	pressure drop, Pa

Greek symbols

η	efficiency
ε	heat exchanger effectiveness
φ	relative humidity
ω	humidity ratio, kg/kg(da)
ρ	Density

Subscripts and superscripts

a	air
abs	absorption chiller
cond	condenser
d	dielectric fluid
evap	evaporator
f	fluid (water of coolant to water heat exchanger)
gen	generator
HP	heat pump
hs	heat sink

$\dot{Q}$ heat transfer rate, W
R thermal resistance, K/W
T temperature, °C
T_j junction temperature, °C
T_wb wet-bulb temperature, °C
T_db dry-bulb temperature, °C
U Overall heat transfer coefficient, W/m²K
$\dot{V}$ volumetric flow rate, m³/s

in inlet
out outlet
rej rejection
TIM thermal interface material
w water

References

[1] International Energy Agency (IEA). Data Centres and Energy – Analysis. Paris, France: IEA; 2020.

[2] ASHRAE. 2016 Handbook—HVAC Systems and Equipment. Atlanta, GA, USA: ASHRAE; 2016.

[3] ASHRAE. 2019 ASHRAE Handbook—HVAC Applications. Atlanta, GA, USA: ASHRAE; 2019.

[4] Data Center Frontier. Hotter Hardware: Rack Densities Test Data Center Cooling Strategies – Available at: <www.datacenterfrontier.com > [accessed 17.3.2026].

[5] Birbarah J., Alzahrani A.R., Ortega A., et al. Experimental investigation of immersion cooling of servers in a two-phase dielectric fluid. Applied Thermal Engineering 2020;167:114736.

[6] Shirazi S., Tian H., He J., et al. Single-phase immersion cooling of data centers: Performance and challenges. Energy Conversion and Management 2020;209:112660.

[7] Rad P., Xu J., Zhang H. Immersion cooling for high performance computing systems: Review and future directions. Renewable and Sustainable Energy Reviews 2019;112:127–145.

[8] Expert Green Prefab Data Centers. Comprehensive Guide to Rack Cooling in Data Centers – Available at: <https://www.greenprefabdatacenters.com/> [accessed 17.3.2026].

[9] MEP Academy. Data Center Cooling Methods Explained: Air Cooling vs Liquid Cooling – Available at: <https://mepacademy.com/data-center-cooling-methods-explained-air-cooling-vs-liquid-cooling/> [accessed 17.3.2026].

[10] Chua H., Toh H., Ng K. A review of thermally activated cooling: Absorption and adsorption cooling. Renewable and Sustainable Energy Reviews 2002;2(3):199–217.

[11] ManufacturingTomorrow. Immersion Cooling and Dielectric Fluids – Available at: <https://www.manufacturingtomorrow.com/> [accessed 2.3.2026].

[12] Vertiv. Immersion Cooling Systems – Available at: <https://www.vertiv.com/> [accessed 2.3.2026].

[13] Inventec Performance Chemicals. Evaluating Immersion Cooling Fluids – Available at: <http://www.inventec.dehon.com> [accessed 2.3.2026].

[14] Data Center Frontier. Meeting Data Center Cooling Demands – Available at: <https://www.datacenterfrontier.com> [accessed 2.3.2026].

[15] Syed Anas N., Banerjee D. Immersion Cooling in Data Centers. In: Proceedings of TFEC; 2025 Mar 9–12; Washington, DC, USA.

[16] Weixue T. Meeting Data Center Cooling Demands. Data Center Frontier; 2025.

[17] Zhai M., Wu X., Wang Y. Performance analysis of silica gel–water adsorption chillers. Energy 2011;36(1):543–551.

[18] Lee S-M., Lee W-J., Cho E.N., et al. Reliability of immersion coolant and PCB materials. IEEE Trans. Components, Packaging and Manufacturing Technology 2025;15(10).

[19] Maatouk C. Study and design of a cooling tower [doctoral dissertation]. Paris: ENMP; 2008.

[20] Tahmasebi M., Nassif N. Energy models for air-cooled chillers. American Journal of Engineering and Applied Sciences 2022;15(10):220–229.

[21] Microsoft. Data center heat reuse – Available at: <https://news.microsoft.com/europe> [accessed 17.3.2026].

[22] Itani N., Dib R., Moubayed N., Maatouk C. LiBr absorption chiller assessment. Engineering Research Express 2025;7:045564.

[23] SEAI. Tallaght District Heating – Available at: <https://www.seai.ie/case-studies/tallaght-district-heating> [accessed 17.3.2026].

[24] Exergi. Stockholm Data Parks – Available at: <https://stockholmdataparks.com> [accessed 17.3.2026].

[25] Incropera F.P., DeWitt D.P., Bergman T.L., Lavine A.S. Fundamentals of Heat and Mass Transfer. 7th ed. Wiley; 2011.

[26] Shrigondekar H., Lin Y-C., Wang C-C. Performance of single-phase immersion cooling. Int. J. Heat Mass Transfer 2023;206:123961.

[27] Jain J.K., Hindoliya D.A. Evaporative cooling pad correlations. J. Inst. Eng. India 2014;95:5–10.

[28] Stull R. Wet-bulb temperature. J. Applied Meteorology and Climatology 2011;50(11):2267–2269.