fluids

MDPI

Review

A Review of Thermal Management in Modern Data Centres: Water Usage Effectiveness and Heat Transfer Coefficients

Andre Cooper and Thi Bang Tuyen Nguyen *

School of Engineering, University of Newcastle, Callaghan, NSW 2308, Australia; andre.cooper@newcastle.edu.au
* Correspondence: thibangtuyen.nguyen@newcastle.edu.au

Abstract

Rapid growth in artificial intelligence, machine learning, and high-performance computing has substantially increased data centre rack power densities, resulting in higher heat generation and more demanding cooling requirements. As water remains widely used in many cooling systems, understanding the relationship between cooling technologies and water consumption is essential for improving cooling efficiency and sustainability. This paper presents a survey of reported water usage effectiveness (WUE) across 83 data centre entries, providing a combined dataset that links WUE with heat-rejection categories. The reported data shows that 23 of these data centres exceed 0.4 L/kWh, which is a sustainability target specified by the Climate Neutral Data Centre Pact for new data centres in water-stressed regions using potable water. Dry facilities employing closed-loop liquid cooling require essentially no water, while evaporative systems typically report water usage effectiveness values up to 2.5 L/kWh. Reported WUE is a facility-level operational metric, set by the proportion of the IT heat load rejected by evaporation, which depends on the heat-rejection topology, ambient wet-bulb conditions, and operating set points. A higher server-side heat transfer coefficient permits a higher coolant supply temperature for a given chip temperature limit, widening the range of ambient conditions under which heat can be rejected without evaporative assistance. Server-side heat transfer is therefore an enabling condition for low WUE rather than a determinant of it. One-dimensional heat transfer models are developed to estimate heat transfer coefficients for different server-level cooling mechanisms widely used for cooling servers within data centres, including air cooling, single-phase immersion cooling, direct liquid cooling, and two-phase immersion cooling. Air cooling, with the lowest heat transfer coefficient, remains widely used in small-scale facilities, whereas direct liquid cooling and two-phase immersion cooling achieve coefficients up to three orders of magnitude higher and are increasingly deployed in high-density installations. These coefficients are used to derive an equivalent evaporative water demand, an upper-bound estimate of the water that would be evaporated in rejecting the heat each mechanism removes; it shares the units of reported WUE but describes thermal capability rather than facility water consumption.

Keywords: data centre cooling; water usage effectiveness (WUE); thermal management; evaporative cooling; convective heat transfer; sustainable cooling

Check for updates

Academic Editors: Ziad Saghir and Florent Duchaine

Received: 17 July 2026
Revised: 4 August 2026
Accepted: 10 August 2026
Published: 14 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

In recent decades, global digitalisation has driven a rapid increase in demand for computing infrastructure. Data centres—specialised facilities that house and operate

information technology (IT) infrastructure for delivering digital applications and services—form the backbone of this digital economy [1]. They broadly fall into two categories: enterprise or colocation facilities, which serve individual organisations or multiple tenants at moderate scale, and hyperscale facilities, operated by cloud and technology providers at massive scale, defined by IBM as housing at least 5000 servers [2]. Alongside growth in total capacity, compute density (the power consumed per rack of servers) has increased substantially, driven by high-performance computing (HPC), machine learning (ML), and artificial intelligence (AI) workloads [3]. This is supported by the fact that the average rack power densities reportedly increased six-fold in a decade, from 2015 to 2024 [4].

This escalation in compute density directly increases thermal loads within server environments. Where heat cannot be removed at a sufficient rate, servers undergo thermal throttling, whereby processing power is deliberately reduced to prevent overheating [5]. Such performance degradation is undesirable, necessitating increasingly effective thermal management strategies. A data centre cooling architecture comprises two subsystems: the server-side cooling mechanism, which transfers heat from IT equipment to a working fluid, and the heat rejection mechanism, which transports that heat to the ambient environment. Historically, low-density facilities relied on air cooling with no water consumption. As power densities increased, evaporative heat rejection became widespread, exploiting the latent heat of evaporation of water to enhance rejection capacity—but inherently consuming water in the process [6].

Contemporary data centres employ a range of *heat-rejection topologies* that differ markedly in their water demand. Note that heat-rejection topology refers to the facility-level means by which heat is transferred to the ambient environment, as distinct from the *server-level cooling mechanism* by which heat is transferred from IT equipment to a working fluid. *Evaporative cooling* systems, characterised by cooling towers and adiabatic coolers consume water continuously during operation [7]. *Free cooling* systems reject heat directly to naturally cool ambient air or water bodies, removing the need for mechanical refrigeration. Where such sources remain sufficiently cool year-round, heat rejection can be achieved without evaporative assistance as well [8]. *Hybrid cooling* systems combine free cooling with evaporative capacity, engaging the latter only when ambient conditions are insufficient for dry rejection, and therefore exhibit water consumption that varies with climate and season [9]. *Closed-loop* (or *dry*) cooling systems use dry coolers to remove heat without evaporative cooling. Liquid-based server cooling is typically employed to remove heat without water consumption [10]. The selection among these topologies is governed by a combination of thermal load, climate, water availability, and regulatory context to directly determine the water intensity of facility operation.

Over the past decade, higher rack densities have driven growing adoption of liquid-based server cooling. For example, water, with approximately 25 times the thermal conductivity and 3500 times the volumetric heat capacity of air, enables significantly greater convective heat transfer coefficients. This reduces the temperature difference required to dissipate a given heat flux, and therefore allows liquid-cooled systems to maintain acceptable chip temperatures at substantially higher coolant temperatures than air cooling permits.

The relationship between server-side heat transfer and facility water consumption is indirect, and is stated here explicitly to avoid overinterpretation. A high chip-to-coolant heat transfer coefficient reduces the temperature difference required to remove a given heat flux, permitting a higher coolant supply temperature for a fixed junction temperature limit. This raises the approach temperature available for sensible heat rejection, increasing the number of annual hours during which dry coolers or free cooling can meet the load without evaporative assistance. Only through this final step does the heat transfer coefficient influence annual water consumption. The water consumed remains determined at facility

level by the heat-rejection topology; the ambient dry-bulb and wet-bulb temperature distribution at the site; cooling tower or dry cooler sizing and approach temperature; the coolant supply temperature and associated control set points; and the reporting boundary adopted for water accounting. A high heat transfer coefficient is therefore necessary but not sufficient for low WUE: it removes a thermal constraint on dry rejection without guaranteeing that dry rejection is adopted, or that ambient conditions permit it year-round. Conversely, a facility in a cool climate may achieve near-zero water consumption with air cooling alone, as several entries in the dataset assembled here demonstrate.

Data centre energy efficiency is commonly benchmarked using Power Usage Effectiveness (PUE), introduced by The Green Grid (data centre efficiency standards body) as the ratio of total facility energy to IT equipment energy [11]. PUE captures only energy consumption and overlooks other environmental impacts, particularly water use. This has motivated the development of complementary metrics. To quantify water consumption, the industry has adopted Water Usage Effectiveness (WUE), also introduced by The Green Grid, in 2011, as the ratio of annual site water consumption to annual IT equipment energy, expressed in litres per kilowatt-hour [12]. Lower WUE values indicate reduced water intensity, with an ideal value of zero corresponding to cooling systems requiring no evaporative water consumption. While WUE reporting has become increasingly common, it remains voluntary, and methodological transparency regarding how reported values are achieved is often limited. As a result, not only is a systematic review of WUE limited, but integrated analysis linking WUE with cooling topologies and data centre scale is also largely lacking.

The emphasis on energy efficiency is reflected in the academic literature. PUE has itself been the subject of sustained critical review: Horner and Azevedo [13] characterise it as an incomplete indicator that captures infrastructure overhead while omitting hardware efficiency, energy productivity, and wider environmental performance, and Yuventi and Mehdizadeh [14] identify ambiguities in its definition and application that limit its value for communicating data centre energy consumption. Reviews of cooling technology follow the same convention. Xu et al. [15] compare air, liquid, and free cooling with respect to thermal management and energy consumption, while Chu and Huang [7] and Zhang et al. [8] review evaporative and free cooling respectively in terms of energy-saving potential; in each case water is addressed qualitatively, if at all. Kong et al. [16] come closest to the present analysis, comparing heat flux ranges, PUE and cost across direct liquid and immersion cooling technologies in order to derive cooling selection guidelines, but the comparison is not extended to water consumption.

Water-focused studies do exist, but operate at a different scale. Mytton [17] and Siddik et al. [18] quantify direct and indirect data centre water consumption at national and watershed level without resolving the cooling architecture responsible for it, while Lei et al. [19] identify cooling system type as one of several determinants of workload-level water use, alongside server efficiency, grid water intensity, and climate zone. Closer to the facility scale, Karimi et al. [20] evaluate PUE and WUE together for two colocation facilities with distinct cooling systems in a hot-arid climate and make recommendations on server-side cooling topology, although the analysis is confined to two sites in a single climate. Taken together, these studies treat water as either absent, aggregated above the level at which cooling decisions are made, or resolved for a small number of facilities, and in none is reported how WUE is connected to the server-level heat transfer behaviour that determines how much heat must ultimately be rejected. Lei et al. [19] reach a similar conclusion, noting that systematic reporting and research on data centre water use and its determinants remain limited.

Even where such variables are analysed together, the underlying thermal behaviour that govern differences between air- and liquid-based cooling systems are often not explic-

itly resolved. It is widely stated that air cooling is insufficient for managing increasing heat loads, while liquid cooling is more effective in reducing water use and improving thermal performance. However, the fundamental link between these observations and heat transfer coefficients remains underexplored.

To the authors' knowledge, this is the first study to combine a facility-scale survey of reported WUE with an explicit treatment of server-level heat transfer coefficients. The study constructs an integrated dataset of WUE and heat-rejection topologies enabling investigation of possible relationships and trend behaviour among these variables. In parallel, simplified heat transfer models with analytical solutions are developed to predict heat flux from electronic chips and the working fluid, which will be used to establish an order-of-magnitude estimate of equivalent evaporative water demand across different server-side cooling mechanisms. By linking fundamental heat transfer principles with industry-reported water metrics, this work aims to provide a clearer physical basis for evaluating water-efficient cooling strategies.

2. Methodology

2.1. Literature Review Methodology

Source identification was conducted using web searches including Google Search, Google Scholar, ScienceDirect, Scopus, IEEE Xplore, and Web of Science. However, WUE is primarily an operational performance metric used in industry practice and is therefore commonly reported in corporate sustainability reporting and technical documentation. Therefore, the majority of reported WUE values originate from industry articles, news reports, sustainability disclosures, and technical white papers rather than peer-reviewed literature. Search terms included "WUE", "water usage effectiveness", "data centre cooling", and "data centre water use", along with related variations.

The search was conducted between November 2025 and June 2026, and the access date for each source is given in the reference list. Because WUE is reported predominantly through corporate sustainability disclosures rather than archival literature, publication dates are not consistently available; where determinable, sources span 2011 to 2025, with the majority published from 2020 onward. Where WUE values were obtained from annual sustainability reports, the value corresponds to the reporting year preceding publication, so the underlying data are typically one year older than the cited source.

Each source identified through this process was subsequently retrieved and reviewed to verify its relevance, accuracy, and suitability for inclusion. Sources were screened to determine whether they reported both a thermal management approach (server-side cooling or heat-rejection method) and an associated WUE value.

Some sources did not report WUE values for individual facilities but instead reported fleet-average WUE values across multiple data centres. In the absence of more detailed facility-level data, these fleet-average values were included in this review. It is acknowledged that fleet-average values may obscure variation between individual facilities and should therefore be interpreted with this limitation in mind.

2.2. Modelling

2.2.1. General Modelling Framework

Typically, in a data centre, the thermal chain begins with electronic components within servers, which generate heat and serve as the primary heat source. This heat is transferred through the thermal interface material and heat sink to a working fluid with a lower temperature than the chip. As the working fluid absorbs heat, it transports the thermal energy away from the chips and transfers it to a secondary cooling loop. The secondary cooling fluid then carries the heat to an evaporative cooling tower (if employed), where

a portion of the water evaporates to reject heat to the atmosphere. This evaporation represents the primary water consumption pathway in the cooling system and introduces the importance of WUE. The process is illustrated in the flow chart in Figure 1.

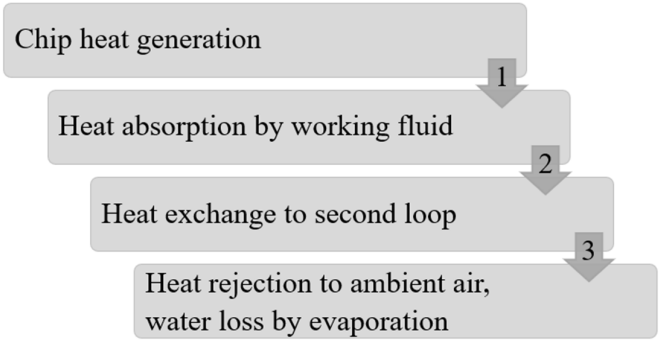


Figure 1. Simplified thermal and water consumption pathway in the cooling system: heat transfer coefficient was estimated for process (1) and water consumption was estimated for process (3). Process (2) is not included in this study.

WUE is defined as the annual water consumed during the heat rejection process relative to the annual IT equipment energy:

$$WUE = \frac{Annual\ Water\ Consumed\ (in\ \text{Litres})}{Total\ IT\ Equipment\ Energy\ (in\ \text{kWh})} \quad (1)$$

As defined here, WUE accounts for water consumed within the site boundary; water associated with off-site electricity generation is excluded, consistent with the Green Grid definition adopted in this study [12].

Under steady-state operation essentially all electrical power delivered to the IT equipment appears as heat, so the quantity of heat requiring rejection is set by the IT load rather than by the server-level cooling mechanism. What the heat transfer capability between a component and the working fluid determines is the temperature at which that heat is delivered to the facility cooling system (heat-rejection topology), and hence the extent to which it can be rejected without evaporation.

The modelling approach adopted here is deliberately reduced order: established analytical correlations are used in place of high-fidelity numerical simulation, trading resolution of the local flow field for the ability to compare cooling mechanisms across a wide parameter space at low computational cost. Explicit physics-based models of this kind have been shown to map system-level performance boundaries at a small fraction of the computational cost of high-fidelity simulation, while retaining the governing physical relationships [21]. The correlations adopted in this study are validated against three-dimensional CFD and experimental data, as summarised in Table 1.

Four representative server-level cooling mechanisms are considered in this analysis, including air cooling, direct liquid cooling (DLC), and single-phase (1-PIC) and two-phase immersion cooling (2-PIC). Among these cooling mechanisms, the first three rely on convective heat transfer, with the heat transfer rate governed by the Reynolds number, thermophysical properties of the working fluid, and the geometrical design of the heat sink. Air cooling remains the most widely deployed server-level cooling method. Fans draw ambient or conditioned air across finned heatsinks mounted directly to processors, dissipating heat by forced convection. Microchannel cold plate DLC circulates a liquid coolant through a metal cold plate in direct contact with the processor surface. The cold plate contains an array of microchannels, which maximise the wetted surface area and promote high convective heat transfer. Heat is absorbed by the coolant and transported

away from the processor. Single-phase immersion cooling is a liquid cooling method in which servers or heated components are fully submerged in a dielectric fluid that removes heat through sensible heating while remaining in the liquid state. Two-phase immersion cooling uses a dielectric fluid that boils at the surface of heated components, absorbing a large amount of heat through its latent heat of evaporation. The vapour rises, condenses on a cooled coil at the top of the tank, and returns as liquid. This mechanism yields substantially higher heat transfer coefficients than either air or single-phase liquid cooling.

2.2.2. Convective Cooling

- *Forced convection*

The following methodology is applied to convective cooling technologies, including air cooling, single-phase liquid cooling, and direct liquid cooling, unless noted otherwise. A schematic diagram of electronic cooling using a fin-based heat sink and a working fluid (either air or water is used in this study) is shown in Figure 2. The heat transfer coefficient (HTC) is a function of Reynolds number, fluid's properties, and geometrical designs. The model assumes steady one-dimensional, incompressible, single-phase flow. The convective HTC for single-phase flow was evaluated using appropriate Nusselt number correlations based on Reynolds number. Temperature dependent properties of the fluids are considered in this study.

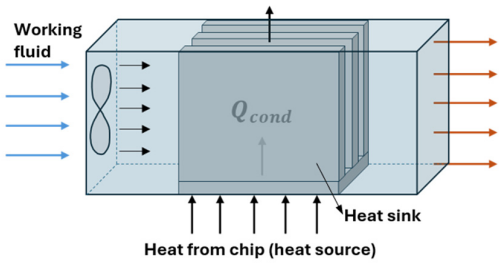


Figure 2. Schematic diagram of electronic cooling using a heat sink and a working fluid.

For the air-cooled fin-based heat sink, operating in the transitional-to-turbulent flow regime ($3000 \leq Re \leq 5 \times 10^6$), the Gnielinski correlation [22] was employed as shown below

$$Nu_{fd} = \frac{(f/8)(Re_D - 1000)Pr}{1 + 12.7(f/8)^{1/2}\left(Pr^{2/3} - 1\right)}, \quad (2)$$

where f is the friction factor corresponding to hydraulically smooth walls and incorporates a correction factor for thermally and hydrodynamically developing flow. Re_D is Reynolds number using the hydraulic diameter and Pr is Prandtl number.

For typical fin-based heat sink, the length of the fins is not usually long enough for the flow to be fully developed; a correction factor to the above equation is used to account for developing flow [22]:

$$Nu = Nu_{fd}\left(1 + \frac{1}{(x/D_h)^{2/3}}\right), \quad (3)$$

where x is the channel length and D_h is hydraulic diameter, which is the ratio between cross-sectional flow area and wetted perimeter. The hydraulic diameter was used to characterise the non-circular cross-sectional geometry for both an air-cooled heat sink and the liquid-cooled microchannel cold plate, with an aspect ratio up to 8 used in the present study.

The same framework, Equations (2) and (3), is used for single-phase immersion cooling. However, for direct liquid cooling, the flow is usually laminar, and so the Nusselt number

was evaluated as a function of channel aspect ratio proposed in [23] (which is in the range between 3.61 and 6.45). This approach captures the geometric dependence of convective heat transfer in rectangular channels.

In both the laminar and turbulent cases, the average convective heat transfer coefficient was then evaluated from the definition of the Nusselt number:

$$h = \frac{Nu * k}{D_h}, \tag{4}$$

where k is the thermal conductivity of the working fluid.

- *Free convection*

For the single-phase immersion cooling case operating under natural convection where the flow is buoyancy-driven: the heated surface warms the adjacent fluid, and the resulting density gradient sets up a convective boundary layer.

In this regime the heat transfer coefficient is a function of the Rayleigh number, the fluid properties, and the geometry of the heated element. The model retains the assumption of steady, single-phase heat transfer, and temperature-dependent fluid properties are considered.

The heat-dissipating elements are treated as isolated vertical cylinders, for which the buoyant boundary layer develops along the element height. The relevant length scale in both the Rayleigh and Nusselt numbers is therefore the vertical height L. The Rayleigh number is defined as:

$$Ra_L = \frac{g\beta\left(T_s - T_f\right)L^3}{\nu^2} * Pr, \tag{5}$$

where g is gravitational acceleration, β the coefficient of thermal expansion, T_s and T_f the surface and bulk-fluid temperatures respectively, ν the kinematic viscosity, and Pr the Prandtl number.

The Nusselt number was evaluated using the Churchill–Chu correlation [24], which provides a single expression valid across the full laminar-to-turbulent range:

$$Nu_L = \left[0.825 + \frac{0.387Ra_L^{1/6}}{\left[1 + \left({}^{0.429}/_{Pr}\right)^{8/27}\right]}\right]^2, \tag{6}$$

Similarly to before, the average convective heat transfer coefficient was then evaluated from the definition of the Nusselt number:

$$h = \frac{Nu_L * k}{\mathrm{L}}, \tag{7}$$

Unlike the forced-convection case, the heat transfer coefficient and the surface-to-fluid temperature difference are not independent: the buoyant driving force $\left(T_s - T_f\right)$ governs h through Equations (5) and (6), while the imposed heat flux requires that h and $\left(T_s - T_f\right)$ together satisfy the surface energy balance:

$$q_{eff} = h\left(T_s - T_f\right), \tag{8}$$

Since h itself depends on $\left(T_s - T_f\right)$, Equations (5)–(8) were solved simultaneously. For a prescribed heat flux q_{eff}, the temperature difference was obtained by fixed-point iteration: an initial $\left(T_s - T_f\right)$ was assumed, h evaluated from Equations (5)–(7), an updated temperature computed from Equation (8), and the process repeated until convergence.

2.2.3. Two-Phase Immersion Cooling

A schematic diagram of electronic cooling using a two-phase immersion approach is shown in Figure 3 where nucleate boiling is expected to dominate the heat transfer process.

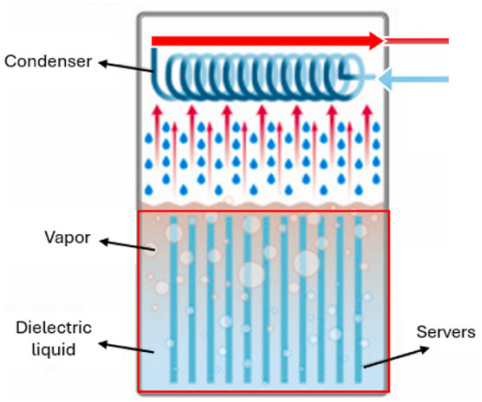


Figure 3. Schematic diagram of electronic cooling using two-phase immersion approach.

The local nucleate boiling heat flux is then evaluated in using the Rohsenow correlation for nucleate boiling [22]:

$$q'' = \mu_l h_{fg} \left[\frac{g(\rho_l - \rho_v)}{\sigma} \right]^{1/2} \left(\frac{c_{p,l} \Delta T_e}{C_{s,f} h_{fg} Pr_l^n} \right)^3, \quad (9)$$

where $\Delta T_e = T_w - T_{sat}$ is excess temperature, $\mu_l$is the liquid dynamic viscosity, h_{fg} is the latent heat of vaporisation, g is gravitational acceleration, ρ_l and ρ_v are the saturated liquid and vapor densities, σ is the surface tension, $c_{p,l}$ is the liquid specific heat and Pr_l is the liquid Prandtl number. The thermophysical properties are all to be evaluated at the saturation temperature T_{sat}.

The parameters $C_{s,f}$ and n are empirical constants that depend on the fluid–surface combination and surface condition. For nucleate boiling of water, the exponent $n = 1$ is typically used [22], while $C_{s,f}$ is treated as a tuning parameter. For boiling between water and copper surfaces, as examined here, values of $0.0068 \leq C_{s,f} \leq 0.0128$ have been reported [22]. In the present work, $C_{s,f} = 0.01$ was adopted for fitting purposes.

An effective nucleate boiling heat transfer coefficient is then obtained by relating the predicted heat flux to the imposed excess temperature:

$$h = \frac{q''}{\Delta T_e'} \quad (10)$$

The Rohsenow correlation does not explicitly account for mass flux effects, as it was originally developed for pool boiling. Its application here therefore assumes that nucleate boiling at the heated surface dominates heat transfer, with convective transport associated with bulk flow considered secondary. As such, it is applied as a first-order estimate of nucleate boiling heat transfer.

A summary is given in Table 1, which outlines the server-side cooling mechanisms considered in this work, their governing equations, and the corresponding facility-level heat-rejection topologies. It is important to note that multiple heat exchangers exist throughout the heat dissipation loop in a data centre. However, this study focuses only on the heat transfer mechanisms directly between the server and the working fluid (referred to as the "cooling mechanism" in Table 1), compared to the "heat-rejection topologies" that represent cooling at the scale of the entire data centre, where WUE is typically evaluated.

Table 1. Server-level cooling mechanism and its relevance in the heat-rejection topology.

Category	Cooling Mechanism	Governing Equations	Validation Source	Compatible Cooling Topologies
Convective cooling	Air cooling	Equations (2)–(4)	Full 3D CFD simulation [25]	Evaporative; Hybrid; Free cooling (climate-dependent)
	1-PIC	Equations (2)–(8)	Forced convection framework inherited from air-cooling validations; qualitative trend support from [26] and natural convection experimental data [27]	Hybrid; Free cooling (liquid) Dry
	DLC	Equations (2)–(4)	Experimental data [28]	Hybrid; Free cooling (liquid); Dry
Two-phase cooling	2-PIC	Equations (9) and (10)	Experimental data [29]	Hybrid; Dry

3. Results

3.1. Review Summary

3.1.1. General Insights into Data Centre Cooling

The choice of heat-rejection topology is not only a technical decision but also an allocation of water, regulatory, and financial risk. Evaporative systems are vulnerable to climate-driven water scarcity. Rising temperatures increase cooling demand and water consumption, intensifying competition with municipal and agricultural users. Regulatory scrutiny is also increasing. For example, California's Water Measurement Regulation requires eligible water users to measure withdrawals and report annual data [30]. Facilities with WUE above 1 L/kWh may face growing pressure to justify consumption or invest in upgrades, while low-water liquid-cooling systems align more closely with emerging sustainability objectives.

The financial implications are substantial. To illustrate, consider a facility with an IT load of 10 MW operating at full capacity, corresponding to an annual IT energy consumption of approximately 87,600 MWh. An evaporative system operating at a WUE of 1.5 L/kWh would consume approximately 131,400 m^3 of water annually, whereas a closed-loop liquid system would consume essentially none. At published Australian charges of roughly AUD 1–3/m^3 [31], this corresponds to an annual difference of approximately AUD 130,000–390,000. This figure is presented as an indicative example rather than a specific case study, and scales linearly with IT load.

Despite these benefits, liquid cooling adoption remains constrained by retrofit costs, downtime, and operational complexity. Existing facilities often require coolant distribution units, leak-detection systems, and specialised maintenance expertise, making implementation challenging for enterprise and colocation operators. The sustainability advantage of liquid cooling also depends on context. Although liquid-cooled systems generally achieve lower PUE through reduced fan energy, additional pumping requirements and the embodied impacts of cooling infrastructure may offset some benefits, particularly in regions with carbon-intensive electricity grids. In contrast, regions with low-carbon electricity, such as Scandinavia, can achieve strong water and carbon performance simultaneously. Consequently, cooling technologies should be evaluated on a region-specific basis.

Hyperscale and HPC operators have adopted liquid cooling more readily due to higher thermal densities, providing a pathway for technology maturation and cost reduction. As systems become more standardised and modular, adoption barriers for smaller facilities are expected to decrease.

Climate and geography strongly influence cooling strategy. Cool climates can support extensive free cooling, enabling WUE values close to zero, although backup cooling is generally required during unusually warm periods. In warmer regions, free cooling is typically seasonal, requiring hybrid operation with evaporative systems and preventing annual WUE from reaching zero. Climate also affects the feasibility of dry cooling. Higher ambient temperatures reduce the temperature difference available for heat rejection, increasing the server-side heat transfer performance required for effective dry operation. This strengthens the case for advanced liquid-cooling technologies such as direct liquid cooling (DLC) and two-phase immersion cooling (2-PIC) in warm climates.

3.1.2. Thermal Management Pathway

An illustration for a multi-scale thermal management framework, from system-level heat rejection to detailed chip-level heat transfer for convective cooling is depicted in Figure 4.

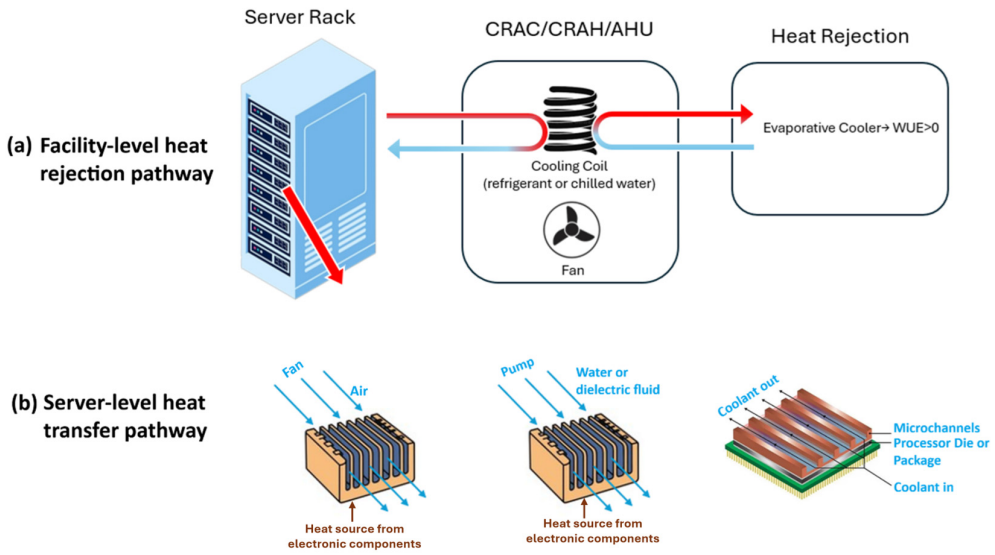


Figure 4. Schematic diagram for multi-scale thermal management in data centres using convective cooling mechanisms: (**a**) facility-level heat-rejection pathway from server to evaporative cooling tower; (**b**) detailed chip-level heat transfer pathway within the server.

For a fin-based heat sink, air flow dissipates heat from servers and gets warmer. Warm air from the server rack flows into units such as a Computer Room Air Conditioner (CRAC), Computer Room Air Handler (CRAH), or Air Handling Unit (AHU), where it is cooled down. The cooled air is sent back to the servers, while the heat is carried away and released outside the building by an evaporative tower. The direct heat dissipation between air flow and the servers is calculated in the following section.

For single-phase immersion cooling technology, liquid is used to remove heat from servers. Dielectric fluid flows through parts attached to the server, taking away heat. The warm liquid then goes to a coolant-to-water heat exchanger, where the heat is passed into water. The heat is finally carried away and released outside the building, keeping the servers cool and working properly.

For direct liquid cooling technology, heat is removed directly from the server using liquid. A small plate with microchannels is attached to the hot parts. Cool liquid flows

through these channels, absorbing heat. The warmed liquid then moves to a heat exchanger, where the heat is passed to water and carried away outside the building.

Figure 5 illustrates the two-phase immersion cooling process. In this system, the servers are placed in dielectric liquid that can boil safely. As the servers heat up, the liquid around them turns into vapor carrying heat away. The vapor then moves to a cooling surface, turns back into liquid and returns to the tank, repeating the cycle to keep the servers cool.

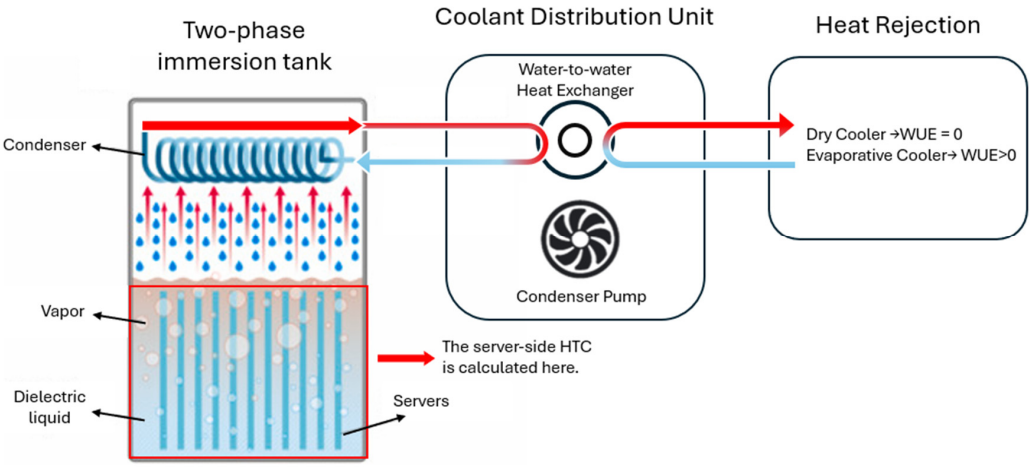


Figure 5. Schematic diagram for two-phase immersion cooling technology in a data centre.

3.2. Water Usage Effectiveness (WUE)

Table 2 summarises 83 data centre (DC) entries included in this study and the corresponding references, including websites and published articles, used for data collection. It should be noted that some entries correspond to data centre fleets comprising multiple facilities. The same data centre names are referenced in Figures 6 and 7.

Table 2. Data centre entries included in this study, their reported WUE values, and the corresponding references.

#	Data Centre (DC)	WUE	Reference
1	LiquidStack Conceptual DC (new)	~0	[32]
2	Microsoft Next-DCs	~0	[33]
3	Various QTS DCs	~0	[34]
4	Ixora's HRM-4 system DCs	~0	[35]
5	Nautilus Data Tech. commercial floating DC	~0	[36]
6	Facebook Luleå DC	~0	[37]
7	Google DC in Hamina, Finland	~0	[38]
8	Various New CyrusOne DCs	~0	[39]
9	Hydro66 DC	~0	[40]
10	LUMI DC	~0	[41]
11	Phoenix (AZ1) DC Campus	~0	[42]
12	Ashburn (VA1) DC Campus	~0	[43]

Table 2. *Cont.*

#	Data Centre (DC)	WUE	Reference
13	Johannesburg I (JNB1) DC Campus	~0	[44]
14	Cyberjaya (KUL1) DC Campus	~0	[45]
15	Melbourne (MEL1) DC Campus	~0	[46]
16	Novva Utah DC	~0	[47]
17	Lefdal Mine DC	~0	[48]
18	BitFury Mega DC	~0	[49]
19	DC4-FR, OpCore	~0	[50]
20	DC3-FR, OpCore	~0	[50]
21	Chindata "X-Cooling" Hyperscale Campus	~0	[51]
22	KDDI small DC (Yokohama)	~0	[52]
23	Microsoft Johor Bahru SEA3 (under construction)	~0	[53]
24	Prime DC Fleet Average	0.002	[54]
25	Allied Control DataTank DC	0.003	[55]
26	DC2-FR, OpCore	0.009	[50]
27	AWS Frankfurt DC	0.01	[56–58]
28	Marsden Park DC	0.01	[59]
29	Microsoft Ireland DC	0.02	[60–62]
30	Microsoft Singapore DC	0.02	[60,61]
31	AWS Melbourne DC	0.02	[56–58]
32	AWS Stockholm DC	0.02	[56–58]
33	AWS Ohio DC	0.03	[56–58]
34	Vint Hill OVHCloud DC	0.04	[63,64]
35	Microsoft Netherlands DC	0.04	[60,61,65]
36	AWS Canada DC	0.04	[56–58]
37	Microsoft Sweden DC	0.05	[60,61]
38	ServerDomes Fleet Average	0.1	[66,67]
39	atNorth Fleet Average	0.1	[68]
40	Microsoft Iowa DC	0.1	[60,61]
41	AWS Ireland DC	0.1	[56–58]
42	AWS Sydney DC	0.12	[56–58]
43	AWS Virginia DC	0.12	[56–58]
44	Gravelines OVHCloud DC	0.13	[63,64]
45	Microsoft Wyoming DC	0.16	[60,61]
46	AWS Oregon DC	0.16	[56–58]
47	Microsoft Virginia DC	0.18	[60,61]
48	Erith OVHCloud DC	0.2	[63,64]

Table 2. *Cont.*

#	Data Centre (DC)	WUE	Reference
49	Qiandao Lake DC	0.2	[69]
50	Luxconnect Fleet Average	0.206	[70,71]
51	AWS Sao Paulo DC	0.23	[56–58]
52	Meta (Facebook) Fleet Average	0.24	[72]
53	Microsoft Texas DC	0.24	[60,61,73]
54	DC5-FR, OpCore	0.25	[50]
55	Limburg OVHCloud DC	0.27	[63,64]
56	Roubaix OVHCloud DC	0.3	[63,64]
57	OVHcloud SYD DC	0.3	[74]
58	Warsaw OVHCloud DC	0.36	[63,64]
59	Green Mountain Fleet Average	0.39	[75]
60	Microsoft Poland DC	0.44	[60,61]
61	AWS California DC	0.51	[56–58]
62	Microsoft Illinois DC	0.52	[60,61]
63	Strasbourg OVHCloud DC	0.57	[63,64]
64	Huawei Cloud Ulanqab DC	0.6	[76]
65	Microsoft Washington DC	0.7	[60,61]
66	NREL HPC DC (new)	0.7	[9]
67	AMS1-NL, Iron Mountain	0.85	[50]
68	Airtunk Fleet Average	0.89	[77,78]
69	EcoDataCenter 1	0.91	[79]
70	AWS Tokyo DC	0.91	[56–58]
71	Hillsboro OVHCloud DC	0.93	[63,64]
72	Equinix Global Fleet Average	0.95	[80,81]
73	Iron Mountain Fleet Average	0.996	[82,83]
74	Beauharnois (BHS8 only) OVHCloud DC	1	[63,64]
75	NREL HPC DC (old)	1.27	[9]
76	Switch DC Fleet Average	1.28	[84,85]
77	MGHPCC DC (Seasonal Low)	1.3	[86]
78	Microsoft Arizona DC	1.52	[60,61]
79	Digital Realty Fleet Average	1.58	[87,88]
80	LiquidStack Conceptual DC (old)	1.59	[32]
81	AWS Asia-Pacific (Singapore) DC	1.68	[56–58]
82	NextDC Fleet Average	2.25	[89,90]
83	MGHPCC DC (Seasonal High)	2.5	[86]

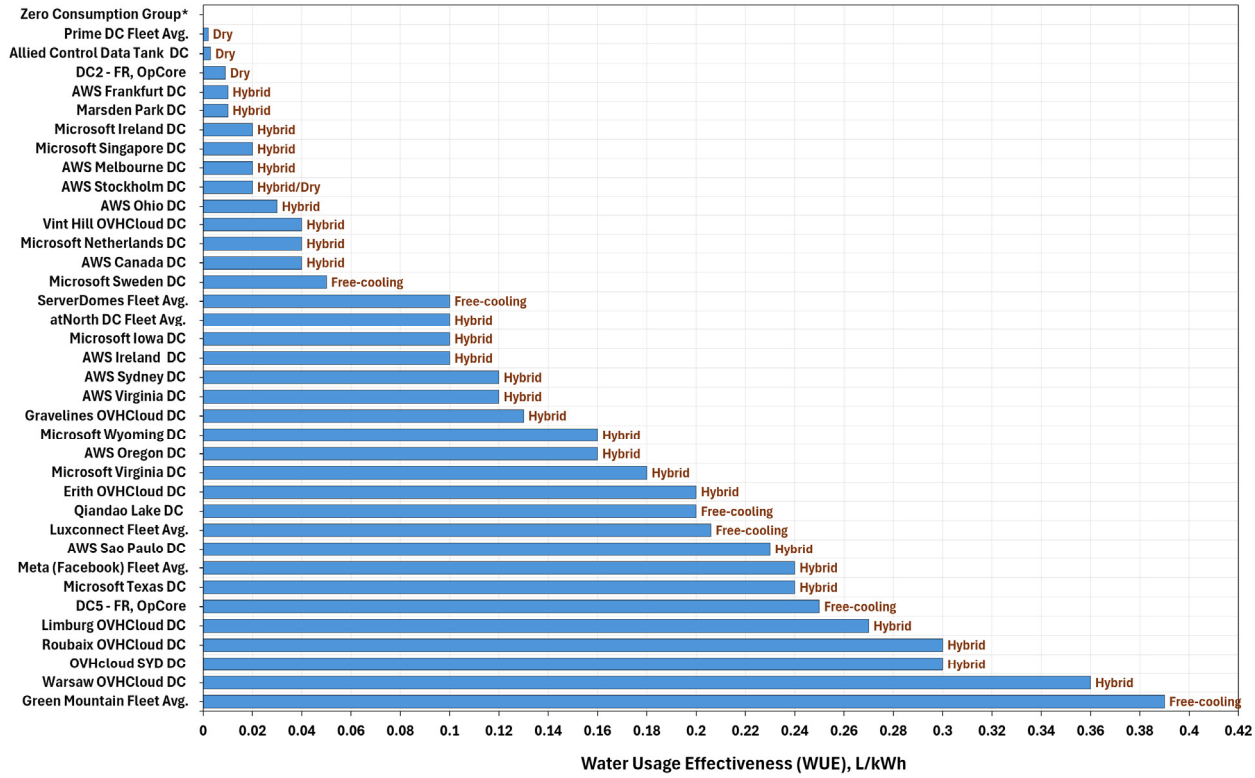


Figure 6. Data centres with WUE up to 0.4 L/kWh. Data labels show heat-rejection topology. The WUE data sources for the data centres shown on the *Y*-axis are listed in Table 2, where the "Zero Consumption Group" comprises DCs 1-23.

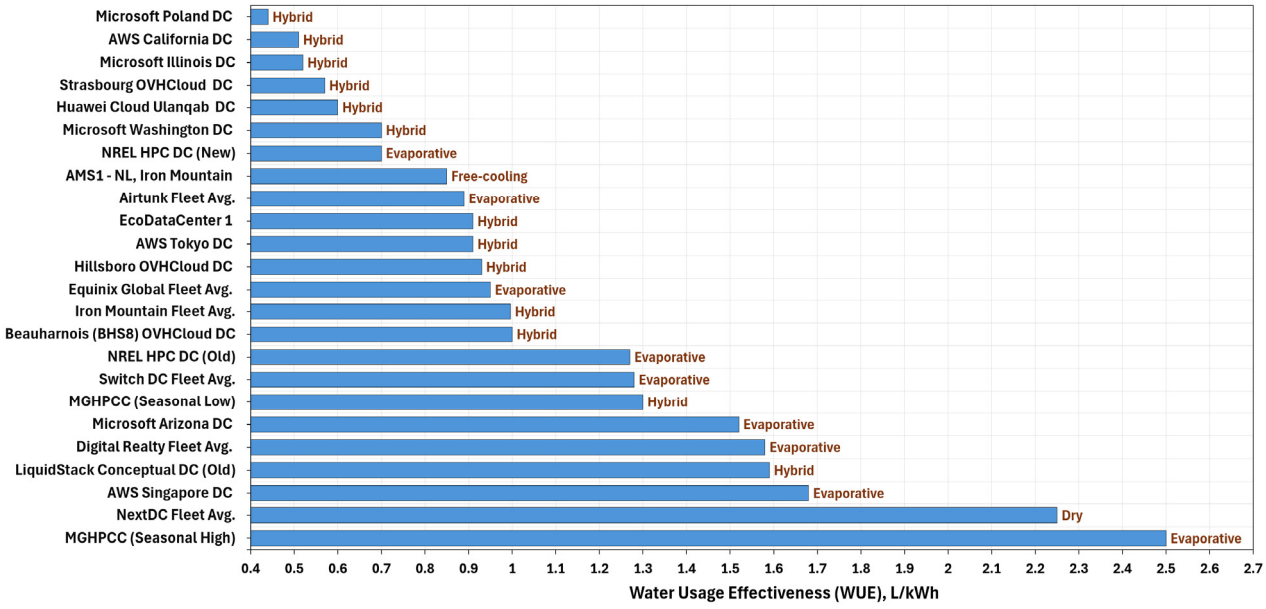


Figure 7. Data centres with WUE above 0.4 L/kWh. Data labels show heat-rejection topology. The WUE data sources for the data centres shown on the *Y*-axis are listed in Table 2.

The dataset, including the WUE and the heat-rejection topologies, is presented in Figure 6 for data centres with WUE up to 0.4 L/kWh and Figure 7 for data centres with WUE above 0.4 L/kWh. The WUE of 0.4 L/kWh (i.e., 0.4 L of water consumed per kilowatt-hour of IT energy) indicates a sustainability target specified by the Climate Neutral Data Centre Pact for new data centres in water-stressed regions using potable water [91]. The

"zero consumption group" mentioned in Figure 6 comprises DCs 1-23 listed in Table 2, with WUE approximately zero.

It is worth noting that the data presented in Figures 6 and 7 are subject to several limitations. Although the definition of WUE has been formally established by the Green Grid, minor differences in interpretation (such as system boundaries or the inclusion of auxiliary water uses) may still occur between organisations. Some WUE values represent averages across multiple data centres or entire facility fleets, which can obscure differences between individual sites as previously discussed. Furthermore, WUE values are typically reported as annual averages, meaning seasonal fluctuations in water consumption are not reported. This is particularly relevant for water resource planning at a regional or utility level, as peak water demand may occur during warmer periods when evaporative cooling systems operate more intensively. WUE represents a relative efficiency metric rather than a measure of absolute water consumption. A large hyperscale data centre with a low WUE may still consume considerably more total water than a smaller facility with a higher WUE, due to differences in IT load. Therefore, WUE indicates the water consumed per unit of computing energy, and total water consumption is determined by both WUE and the IT load.

Reported values also lack a consistent temporal reference, as most sources do not state the year of construction or measurement, and fleet averages aggregate facilities commissioned across different periods. A systematic temporal analysis was therefore not attempted, although the paired entries for the NREL HPC facility and the LiquidStack conceptual design indicate the improvement achievable at a fixed site through a change in thermal management strategy.

3.3. Categorising Heat-Rejection Topology

To further analyse the WUE database developed in this study, four categories of data centres were identified from the results presented in Figures 6 and 7, each exhibiting distinct ranges of WUE. The results shown in Figure 8 were analysed from 10 evaporative DCs, 40 hybrid DCs, 12 free-cooling DCs, and 20 dry DCs. Although the dataset size is very small compared to the number of operational DCs, it serves as an adequate indicator of general performance for each mechanism.

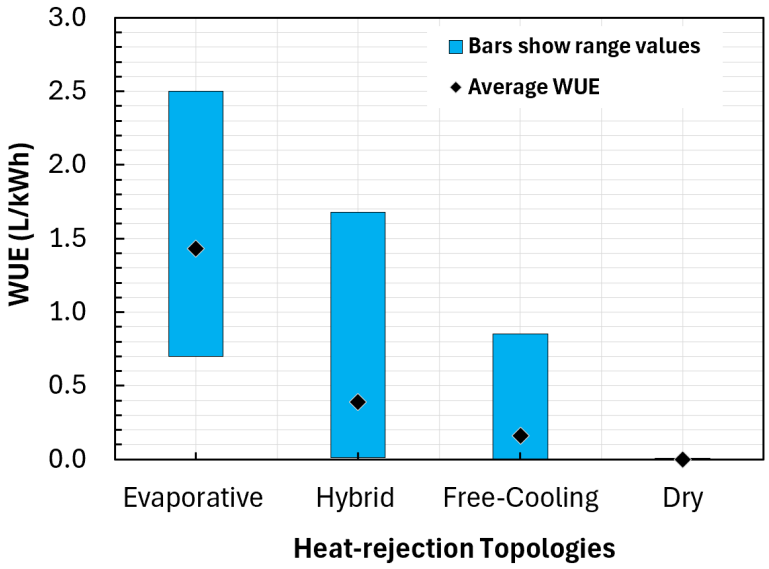


Figure 8. WUE as a function of heat-rejection topology where Hybrid systems combine free cooling with evaporative capacity and Dry or Closed-loop can be used interchangeably. Data markers show average WUE.

Based on the cooling technology characterisation in the literature review, WUE ranges from 0.7 to 2.5 L/kWh for evaporative systems, 0.01 to 2.25 L/kWh for hybrid systems, 0 to 0.85 L/kWh for free cooling, and 0 to 0.009 L/kWh for dry systems. Dry systems consistently achieve the lowest water consumption, with many facilities reporting WUE ≈ 0, while evaporative systems exhibit the highest and most variable values.

3.4. Modelling Results

3.4.1. Heat Transfer Coefficients

The HTC values obtained using Equations (2)–(10) for the four cooling mechanisms described earlier, namely air cooling, single-phase immersion cooling (1-PIC), direct liquid cooling (DLC), and two-phase immersion cooling (2-PIC), are presented in Figure 9.

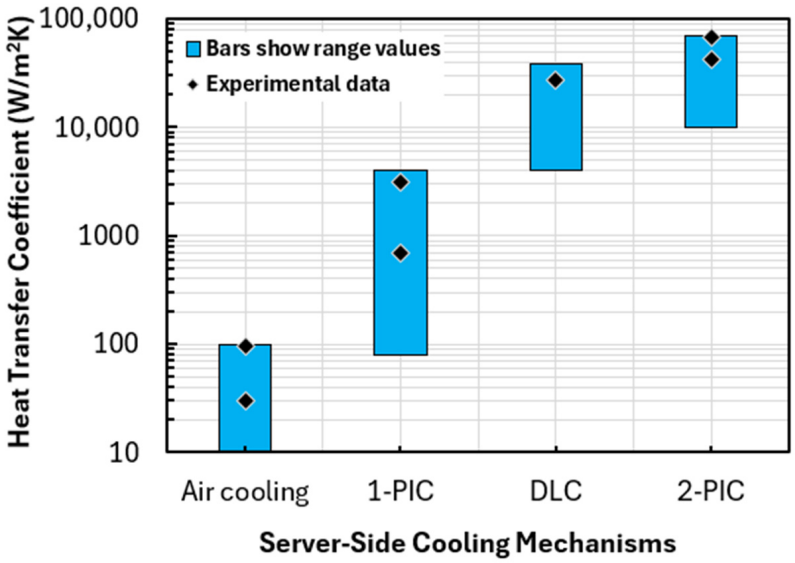


Figure 9. Comparison of heat transfer coefficient (HTC) ranges for four server-side cooling mechanisms. Bars show predicted HTC ranges, and markers show validation values with corresponding experimental studies from [25–29].

For air cooling, the heat transfer coefficient (HTC) was estimated over a Reynolds number range resulting in approximately 10 to 100 W/m^{2}K, which is representative of airflow conditions encountered in fin-based CPU heat sinks with various fin geometries and channel configurations. The lower bound corresponds to natural convection conditions, whereas the upper bound represents advanced heat sink designs, such as cross-cut sinusoidal wavy plate-fin heat sinks [25]. The HTC was determined using the same modelling framework described previously.

The same framework was used to estimate the HTC for the single-phase immersion cooling (1-PIC) mechanism, with a dielectric liquid serving as the working fluid instead of air. In this study, Novec 7100 was selected due to its widespread use in dielectric immersion cooling applications, and its thermophysical properties were obtained from Cheng et al. [26] and Rokonuzzaman et al. [27]. The HTC was found to range from 80 to 4000 W/m^{2}K, corresponding to the lowest and highest flow rates considered, respectively. These two values represent theoretical extremes, where in practice, HTC values are expected to lie well within this range.

For direct liquid cooling (DLC), heat is dissipated through a microchannel cold plate mounted directly onto the electronic chip, which acts as the heat source. In this configuration, HTC predicted by the model depends primarily on the Reynolds number and the microchannel geometry. Nusselt numbers ranging from 3.61 to 6.45 were used for

channels with different cross-sectional geometries and aspect ratios varying from 1 to 8 [23]. A minimum hydraulic diameter of 0.1 mm was selected based on the experimental study of a microchannel cold plate reported in [28], while the flow was restricted to the laminar regime (Re < ~2000). Under these conditions, the predicted HTC ranged from approximately 4000 to 38,000 W/m^{2}K.

Two-phase immersion cooling (2-PIC) relies on nucleate boiling at a heated copper surface, where heat is removed primarily through liquid vaporisation and the associated latent heat transfer. The heat transfer process in this regime is highly complex and governed by multiple interacting phenomena, including bubble nucleation, growth, departure, and replenishment. In the absence of a more widely accepted predictive model for this mechanism, the Rohsenow correlation was adopted as discussed earlier. This model relates the HTC to the wall superheat, defined as the temperature difference between the heated surface and the saturation temperature of the working fluid. Using excess temperatures ranging from 5 to 15 K, the predicted HTC was found to vary from approximately 10,000 to 70,000 W/m^{2}K, respectively, where experimental data was found from [29] between approximately 42,000 and 67,000 W/m^{2}K for water at a saturation temperature of 100 °C.

3.4.2. Heat Dissipation Flux and Equivalent Evaporative Water Demand

The average heat transfer coefficients obtained from Figure 9 were used to estimate heat dissipation from the heat sink as a function of excess temperature, defined as the temperature difference between the chip surface and the working fluid used in Equation (9). For convective cooling methods, including air cooling, single-phase immersion cooling (1-PIC), and direct liquid cooling (DLC), the working fluid temperature is typically the ambient temperature, which varies seasonally. Therefore, excess temperatures ranging from 5 to 70 °C were considered in Figure 10, while values up to 30 °C were used for the two-phase immersion cooling (2-PIC) case. For a fixed working fluid temperature, a higher excess temperature results in higher heat dissipation from the chip. This, in turn, requires higher heat-rejection capacity and water usage effectiveness (WUE) to maintain adequate cooling performance. Consequently, cooling performance generally improves during winter, when lower ambient temperatures (meaning lower working fluid temperatures) increase the excess temperature, and decreases during summer. It should be noted that, although the heat generated by the electrical component equals the heat dissipated to the working fluid under steady-state convective cooling, the heat sink typically has a much larger surface area than the chip. As a result, the heat flux at the heat sink surface is lower than the chip heat generation flux, especially for fin-based heat sinks. Therefore, the heat flux values presented in Figure 10 should be regarded as a lower reference when compared with chip-level heat fluxes. Despite being an order-of-magnitude analysis, Figure 10 provides useful guidance for selecting appropriate cooling technologies. For example, air cooling may be suitable for heat fluxes below 10 kW/m^2, 1-PIC for heat fluxes below 100 kW/m^2, while DLC is preferable for heat fluxes exceeding 100 kW/m^2 and 2-PIC for heat fluxes above 1000 kW/m^2.

Assuming a thermal load of 1 MW, corresponding approximately to 1430 NVIDIAH-100SXM electronic chips, the total chip surface area is approximately 1.164 m^2 with ~ 814 mm^2 each chip. Although the effective heat transfer area provided by a heat sink is typically larger than the chip surface area, the chip area is used here as a reference owing to the absence of a detailed heat sink design. It is further assumed that all heat rejection is achieved through evaporative cooling, with no secondary cooling loop. Under this assumption, the heat transferred directly from the chips to the working fluid is rejected through evaporation. As a result, the calculated water consumption may be regarded as an upper-bound estimate of the water consumption.

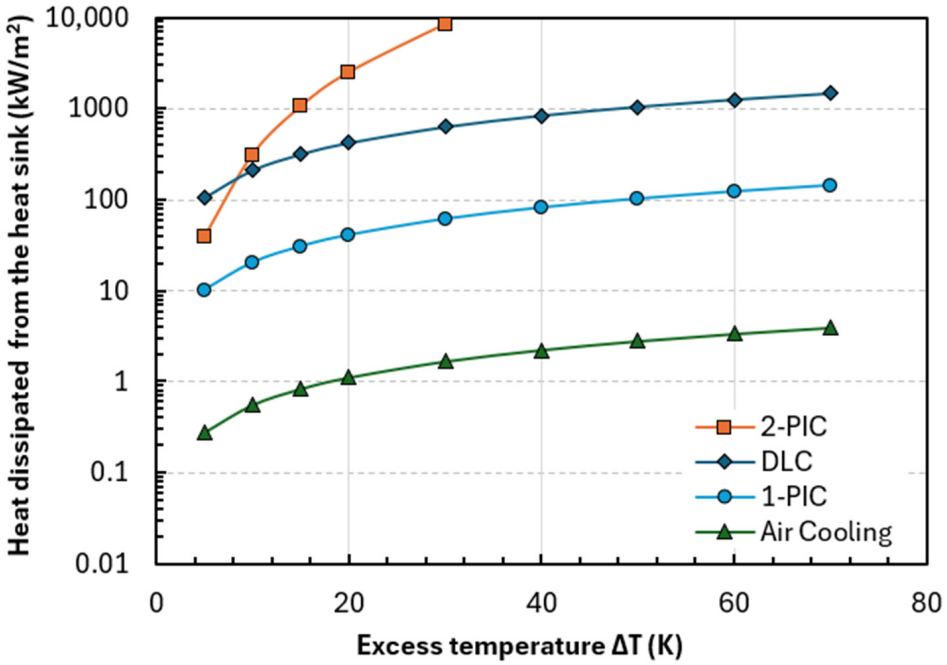


Figure 10. Heat dissipation from heat sink to working fluid for different excess temperatures.

The water usage per megawatt of thermal load was obtained from the following heat balance:

$$Water\ Usage\left[\frac{\mathrm{L}}{\mathrm{kWh}}\right] = \frac{q_{ind} * A_{ind} * 3600}{\rho * L_{evap}} \tag{11}$$

where q_{ind} is indicative heat dissipation flux obtained from Equation (9) for 2-PIC, while for convective heat transfer it is calculated as:

$$q_{ind} = h_{ind} A_{ind} \tag{12}$$

with h_{ind} being the indicative heat transfer coefficient, obtained from the average values presented in Figure 9 for the three convective cooling mechanisms; A_{ind} is the indicative total die area of all 1430 electronic components assumed in this study, and $A_{ind} = 1.164$ m^2 in this study. L_{evap} is latent heat of evaporation for water which is 2260 kJ/kg and ρ is water density which is 998 kg/m^3.

The quantity plotted in Figure 11 is termed the equivalent evaporative water demand, which represents the estimated water usage by evaporation shown in Equation (11) under the above assumptions for different cooling mechanisms. Results are shown for convective cooling with excess temperatures 10, 20, and 60 K, and 10 and 20 K for 2-PIC.

While the analysis is intended as an order-of-magnitude assessment, it provides useful guidance on the expected effectiveness for various cooling technologies. For example, for direct liquid cooling operating with an excess temperature of approximately 20 K, the corresponding equivalent evaporative water demand is approximately 0.4 L/kWh. This value can be used as a first-order estimate of water consumption for data centres with similar thermal management configurations.

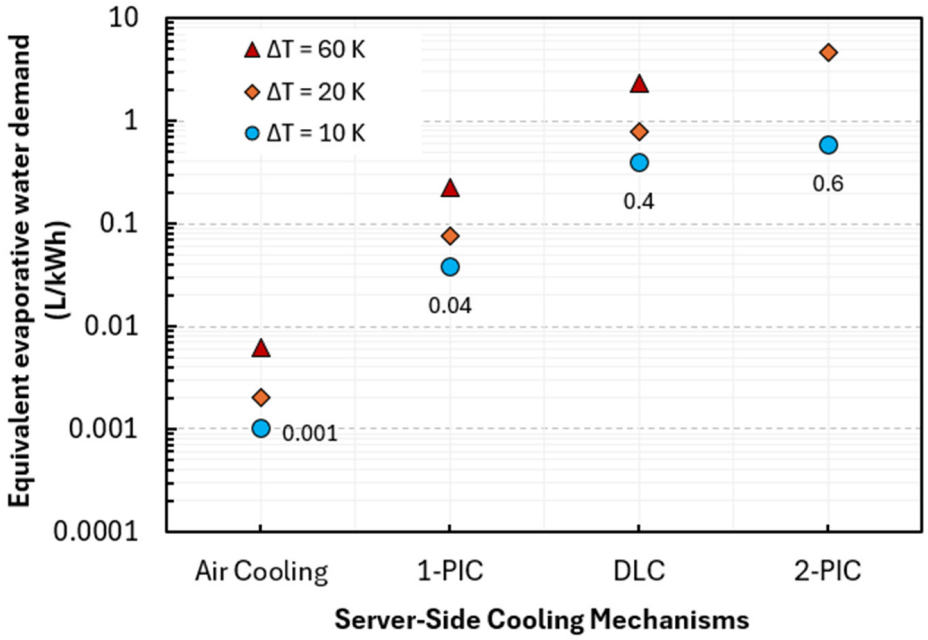


Figure 11. Equivalent evaporative water demand for four server-level cooling mechanisms at selected excess temperatures, evaluated for a 1 MW reference IT load and a total die area of 1.164 m^2. Values are upper bounds, assuming all heat is rejected by direct evaporation with no secondary cooling loop.

4. Conclusions

This study first provides a brief overview of data centre cooling technologies, highlighting that heat-rejection topology strongly influences water consumption, cost, operational risk, and sustainability outcomes. The review further demonstrates that liquid-cooling technologies can substantially reduce WUE, particularly when combined with dry heat rejection, while climate remains a key factor affecting cooling mechanism selection, dry-cooling potential, and achievable WUE.

The literature review shows that WUE is highly dependent on heat-rejection topology, including evaporative, hybrid, free-cooling, and dry cooling. Reported WUE values vary significantly, with a few configurations reaching up to 2.5 L/kWh, while 23 of the 83 data entries exceed 0.4 L/kWh, which is a sustainability target specified by the Climate Neutral Data Centre Pact for new data centres [91].

An order-of-magnitude one-dimensional heat-transfer analysis was conducted for four cooling mechanisms: air cooling, single-phase immersion cooling (1-PIC), direct liquid cooling (DLC), and two-phase immersion cooling (2-PIC). The results indicate that 2-PIC can provide heat transfer coefficients (HTCs) approximately three orders of magnitude higher than those of air cooling. For high-density facilities with substantial heat fluxes from electronic components to the working fluid, cooling technologies with high HTCs are therefore essential for reducing water consumption and enabling dry heat rejection. The estimated HTCs were subsequently used to determine the heat dissipation flux achievable and the equivalent evaporative water demand over a range of excess temperatures. The results provide useful guidance for cooling mechanism selection across applications spanning heat fluxes from a few to several thousand kW/m^2.

Several limitations should be noted. The WUE dataset is drawn predominantly from corporate sustainability disclosures and technical documentation rather than peer-reviewed sources, so reported values are self-declared and not independently verified. Voluntary disclosure introduces a potential selection bias, as operators achieving favourable water performance may be more likely to publish it. A small number of entries correspond to

conceptual designs or facilities not yet in operation and represent design targets rather than measured performance, and the 83 entries represent a small fraction of operational facilities globally, weighted toward Europe, North America, and Australia. The heat transfer analysis is a one-dimensional, steady-state treatment intended to establish relative magnitudes rather than to serve as a design tool. It resolves only the thermal resistance between the electronic component and the working fluid, and does not model the intermediate heat exchangers, coolant distribution units, pumps, or secondary loops present in a real facility. The equivalent evaporative water demand values assume all heat is rejected by direct evaporation with no secondary loop, and are therefore upper bounds rather than expected values.

A further order-of-magnitude analysis was performed to estimate the water consumption associated with each megawatt of thermal load, under evaporative heat rejection assumption. The resulting values provide an upper-bound reference for WUE on a megawatt thermal-load basis. The analysis suggests that low-density facilities employing air cooling may exhibit negligible WUE even during summer conditions. In contrast, high-density facilities using DLC or 2-PIC may consume more than 0.4 L per kWh for each MW thermal load. These values can serve as first-order estimates of water consumption for data centres operating under similar thermal-management configurations.

Author Contributions: Conceptualization, methodology, validation, investigation, resources, data curation, writing-original draft preparation, and writing-review and editing: A.C. and T.B.T.N. Supervision: T.B.T.N. Project administration: A.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: The corresponding author gratefully acknowledges the support of the TRaCE Enterprise Academic Fellowships, which provided teaching relief and enabled the time to undertake this research.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Susnjara, S.; Smalley, I. What Is a Data Center? IBM. Available online: https://www.ibm.com/think/topics/data-centers (accessed on 27 February 2026).
2. Powell, P.; Smalley, I. What Is a Hyperscale Data Center? IBM. Available online: https://www.ibm.com/think/topics/hyperscale-data-center (accessed on 13 March 2026).
3. SuperMicro. What is Liquid Cooling? Available online: https://www.supermicro.com/en/glossary/liquid-cooling (accessed on 27 February 2026).
4. Lee, A.; McCormack, D. Data Centres: Powering the growth of AI and cloud computing. Macquarie Asset Management. Available online: https://www.macquarie.com/assets/macq/mam/insights/2025/data-centers-powering-the-growth-of-ai-and-cloud-computing/data-centres-powering-the-growth-of-ai-and-cloud-computing.pdf (accessed on 27 February 2026).
5. Brooks, D.; Martonosi, M. Dynamic thermal management for high-performance microprocessors. In *Proceedings of the 7th International Symposium on High-Performance Computer Architecture (HPCA-7), 2001, Monterrey, Mexico*; IEEE: Piscataway, NJ, USA, 2001; pp. 171–182. [CrossRef]
6. Jose, L. Data Center Cooling Infrastructure. *WordPress*. Available online: https://dc.mynetworkinsights.com/data-center-cooling-infrastructure/ (accessed on 27 February 2026).
7. Chu, J.; Huang, X. Research status and development trends of evaporative cooling air-conditioning technology in data centers. *Energy Built Environ.* **2023**, *4*, 86–110. [CrossRef]

8. Zhang, H.; Shao, S.; Xu, H.; Zou, H.; Tian, C. Free cooling of data centers: A review. *Renew. Sustain. Energy Rev.* **2014**, *35*, 171–182. [CrossRef]
9. Sickinger, D.; Geet, O.V.; Belmont, S.; Carter, T.; Martinez, D. Thermosyphon Cooler Hybrid System for Water Savings in an Energy-Efficient HPC Data Center. National Renewable Energy Laboratory. 2018. Available online: https://docs.nlr.gov/docs/fy18osti/72196.pdf (accessed on 14 November 2025).
10. Ebrahimi, K.; Jones, G.F.; Fleischer, A.S. A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities. *Renew. Sustain. Energy Rev.* **2014**, *31*, 622–6388. [CrossRef]
11. The Green Grid. PUE™: A Comprehensive Examination of the Metric. 2012. Available online: https://www.thegreengrid.org/en/resources/library-and-tools/20-PUE:-A-Comprehensive-Examination-of-the-Metric (accessed on 13 March 2026).
12. The Green Grid. Water Usage Effectiveness (WUE™): A Green Grid Data Center Sustainability Metric. 2011. Available online: https://www.thegreengrid.org/system/files/store/WUE_v1.pdf (accessed on 17 November 2025).
13. Horner, N.; Azevedo, I. Power usage effectiveness in data centers: Overloaded and underachieving. *Electr. J.* **2016**, *29*, 61–69. [CrossRef]
14. Yuventi, J.; Mehdizadeh, R. A critical analysis of Power Usage Effectiveness and its use in communicating data center energy consumption. *Energy Build.* **2013**, *64*, 90–94. [CrossRef]
15. Xu, S.; Zhang, H.; Wang, Z. Thermal Management and Energy Consumption in Air, Liquid, and Free Cooling Systems for Data Centers: A Review. *Energies* **2023**, *16*, 1279. [CrossRef]
16. Kong, R.; Zhang, H.; Tang, M.; Zou, H.; Tian, C.; Ding, T. Enhancing data center cooling efficiency and ability: A comprehensive review of direct liquid cooling technologies. *Energy Build.* **2024**, *308*, 132846. [CrossRef]
17. Mytton, D. Data centre water consumption. *npj Clean. Water* **2021**, *4*, 11. [CrossRef]
18. Siddik, M.A.B.; Shehabi, A.; Marston, L. The environmental footprint of data centers in the United States. *Environ. Res. Lett.* **2021**, *16*, 064017. [CrossRef]
19. Lei, N.; Lu, J.; Shehabi, A.; Masanet, E. The water use of data center workloads: A review and assessment of key determinants. *Resour. Conserv. Recycl.* **2025**, *219*, 108310. [CrossRef]
20. Karimi, L.; Yacuel, L.; Degraft-Johnson, J.; Ashby, J.; Green, M.; Renner, M.; Bergman, A.; Norwood, R.; Hickenbottom, K.L. Water-energy tradeoffs in data centers: A case study in hot-arid climates. *Resour. Conserv. Recycl.* **2022**, *181*, 106194. [CrossRef]
21. Shi, L.; Zhang, X.; Liu, Y.; Xiao, G.; Ibrahim, M.; Zheng, L.; Wang, T. Optimization of electrified turbocharging: Bridging the energy-drivability trade-off. *Energy* **2026**, *360*, 141522. [CrossRef]
22. Bergman, T.L.; Lavine, A.S. *Fundamentals of Heat and Mass Transfer Eighth Edition*; Wiley: Hoboken, NJ, USA, 2017.
23. Cengel, Y.A. *Heat Transfer A Practical Approach*, 2nd ed.; McGraw-Hill: Columbus, OH, USA, 2002.
24. Churchill, S.W.C.; Humbert, H.S. Correlating equations for laminar and turbulent free convection from a vertical plate. *Int. J. Heat Mass Transf.* **1975**, *18*, 1323–2329. [CrossRef]
25. Nilpueng, K.; Ahn, H.S.; Jerng, D.-W.; Wongwises, S. Heat transfer and flow characteristics of sinusoidal wavy plate fin heat sink with and without crosscut flow control. *Int. J. Heat Mass Transf.* **2019**, *137*, 565–572. [CrossRef]
26. Cheng, C.-C.; Chang, P.-C.; Li, H.-C.; Hsu, F.-I. Design of a single-phase immersion cooling system through experimental and numerical analysis. *Int. J. Heat Mass Transf.* **2020**, *160*, 120203. [CrossRef]
27. Rokonuzzaman, A.S.M.; Erdem, K.; Şahin, B.; Özdemir, M.R. Experimental study of a single-phase immersion cooling system with natural and forced convection. *Int. J. Therm. Sci.* **2025**, *214*, 109868. [CrossRef]
28. Tuckerman, D.B.; Pease, R.F.W. High-Performance Heat Sinking for VLSI. *IEEE Electron Device Lett.* **1981**, *2*, 126–129. [CrossRef]
29. Yang, L.; Hong, F.; Cheng, P. Pool boiling enhancement through a guidance structure mounted above heating surface. *Int. J. Heat Mass Transf.* **2019**, *139*, 751–763. [CrossRef]
30. The California Water Boards. Water Measurement Regulation. Available online: https://www.waterboards.ca.gov/waterrights/water_issues/programs/diversion_use/water_measurement.html (accessed on 20 March 2026).
31. Muray River Council. 2025-2026 Rates and Charges. Available online: https://www.murrayriver.nsw.gov.au/files/assets/public/v/1/documents/ipampr-docs-and-budgets/2025-2026/2025-26-rates-and-charges-explainer-for-website.pdf (accessed on 28 May 2026).
32. LiquidStack. Two-Phase Liquid Immersion Cooling 2022 Case Study. 2022. Available online: https://liquidstack.com/content/uploads/2022/05/2022-Case-Study.pdf (accessed on 14 November 2025).
33. Solomon, S. Sustainable by Design: Next-Generation Datacenters Consume Zero Water for Cooling. *Microsoft*. Available online: https://www.microsoft.com/en-us/microsoft-cloud/blog/2024/12/09/sustainable-by-design-next-generation-datacenters-consume-zero-water-for-cooling/ (accessed on 14 November 2025).
34. QTS. Pioneering Water Efficiency and Sustainability in Data Center Operations. Available online: https://qtsdatacenters.com/resources/pioneering-water-efficiency-and-sustainability-in-data-center-operations/ (accessed on 14 November 2025).
35. Ixora. Unlock Cost Savings and Sustainability with iXora's Advanced Immersion Cooling Solution. Available online: https://www.ixora.eu/ixora-immersion-cooling-roi-analysis/ (accessed on 14 November 2025).

36. Black & Veatch. Commissioning for World's First Zero-Water Consumption Data Center. Available online: https://www.bv.com/projects/commissioning-for-worlds-first-zero-water-consumption-data-center (accessed on 14 November 2025).
37. DPR Construction. Facebook Travels Abroad. Available online: https://www.dpr.com/media/blog/facebook-data-center-goes-online (accessed on 14 November 2025).
38. Reuters. Google Uses Sea Water to Cool Finland Data Center. Available online: https://www.reuters.com/article/business/environment/google-uses-sea-water-to-cool-finland-data-center-idUS2576036493/ (accessed on 14 November 2025).
39. CyrusOne. 2020 Sustainability Report. 2021. Available online: https://www.responsibilityreports.com/HostedData/ResponsibilityReportArchive/c/NASDAQ_CONE_2020.pdf (accessed on 26 November 2025).
40. Data Center Map. Ardent Boden. Available online: https://www.datacentermap.com/sweden/boden/hydro66/ (accessed on 26 November 2025).
41. Robertsén, F. Sustainability in HPC, a Look at LUMI. Available online: https://www.bnl.gov/modsim/events/2024/files/talks/fredrik-robertsen.pdf (accessed on 7 June 2026).
42. Vantage Data Centers. Phoenix (AZ1) Data Center Campus. Available online: https://vantage-dc.com/data-center-locations/north-america/phoenix-arizona/ (accessed on 10 December 2025).
43. Vantage Data Centers. Ashburn (VA1) Data Center Campus. Available online: https://vantage-dc.com/data-center-locations/north-america/ashburn-i-virginia/ (accessed on 10 December 2025).
44. Vantage Data Centers. Johannesburg I (JNB1) Data Center Campus. Available online: https://vantage-dc.com/data-center-locations/emea/johannesburg-i-south-africa/ (accessed on 10 December 2025).
45. Centers, V.D. Cyberjaya (KUL1) Data Center Campus. Available online: https://vantage-dc.com/data-center-locations/apac/cyberjaya-i-malaysia/ (accessed on 10 December 2025).
46. Vantage Data Centers. Melbourne (MEL1) Data Center Campus. Available online: https://vantage-dc.com/data-center-locations/apac/melbourne-australia/ (accessed on 10 December 2025).
47. Novva Data Centers. Water-Free Data Center Cooling. Available online: https://www.novva.com/data-center-services/water-free-cooling/ (accessed on 10 December 2025).
48. Amberside Advisors. Collaboration with innovative Lefdal Mine Data Centers in Norway. AJM Website Design. Available online: https://www.amberside.uk/projects/amberside-advisors'-collaboration-with-innovative-lefdal-mine-datacenter-in-norway (accessed on 10 December 2025).
49. Bitfury. BitFury Launches 40MW Energy Efficient Immersion Cooling Data Center. Available online: https://web.archive.org/web/20200529080624/https://bitfury.com/content/downloads/2015-12-16-tbilisi-data-center-fact-sheet.pdf (accessed on 7 June 2026).
50. Scaleway. Impact Report 2025. Available online: https://www-uploads.scaleway.com/Impact_Report2025_22ee3a8232.pdf (accessed on 7 June 2026).
51. Chindata Group. Chindata and Vertiv Release 'X-Cooling', World's First Zero-WUE Waterless Cooling Technology. Available online: https://www.prnewswire.com/news-releases/chindata-and-vertiv-release-x-cooling-worlds-first-zero-wue-waterless-cooling-technology-301600035.html (accessed on 28 May 2026).
52. GIGABYTE. Japanese Telco Leader KDDI Invents Immersion Cooling Small Data Center with GIGABYTE. Available online: https://www.gigabyte.com/Article/japanese-telco-leader-kddi-invents-immersion-cooling-small-data-center-with-gigabyte (accessed on 28 May 2026).
53. Microsoft. Microsoft to Expand Cloud Region in Johor Bahru, Empowering Southeast Asia's AI Transformation. Available online: https://news.microsoft.com/source/asia/features/microsoft-to-expand-cloud-region-in-johor-bahru-empowering-southeast-asias-ai-transformation/ (accessed on 28 May 2026).
54. Prime Data Centers. 2025 Sustainability Report. Available online: https://primedatacenters.com/wp-content/uploads/2025/08/Prime-Data-Centers_2025-Sustainability-Report_FINAL-optimised.pdf (accessed on 7 June 2026).
55. HPCwire. Allied Control Announces DataTank 1.4MW Container Datacenter. TCI Media Publication. Available online: https://www.hpcwire.com/off-the-wire/allied-control-announces-datatank-1-4mw-container-datacenter/ (accessed on 14 November 2025).
56. AWS. AWS Cloud. Available online: https://sustainability.aboutamazon.com/products-services/aws-cloud (accessed on 26 November 2025).
57. King, C. How is AWS Making Data Centres More Sustainable? *Sustainablity Magazine*. 4 December 2024. Available online: https://sustainabilitymag.com/articles/aws-unveils-next-gen-data-centres-for-ai-computing-demands (accessed on 26 November 2025).
58. About Amazon Team. Local Community Buildings in Ireland to Be Heated by Amazon Data Centre. Amazon.com. Available online: https://www.aboutamazon.eu/news/amazon-web-services/local-community-buildings-in-ireland-to-be-heated-by-amazon-data-centre (accessed on 26 November 2025).
59. Water Services Assoication of Australia. Data Centres and Water in Australia. 2025. Available online: https://majorprojects.planningportal.nsw.gov.au/prweb/PRRestService/mp/01/getContent?AttachRef=SSD-70889211%2120251128T050853.709%20GMT (accessed on 18 December 2025).

60. Microsoft. Modern Datacenter Cooling. Available online: https://datacenters.microsoft.com/wp-content/uploads/2023/05/Azure_Modern-Datacenter-Cooling_Infographic.pdf (accessed on 5 June 2026).
61. Microsoft. Microsoft Datacenters. Available online: https://datacenters.microsoft.com/sustainability/efficiency/ (accessed on 10 December 2025).
62. Microsoft. Ireland Microsoft Commitments. Available online: https://datacenters.microsoft.com/globe/pdfs/sustainability/factsheets/Ireland%20(North%20Europe).pdf (accessed on 7 June 2026).
63. OVHcloud. 2023 Non Financial Performance Statement. Available online: https://corporate.ovhcloud.com/sites/default/files/2023-12/dpef_en_ovhcloud.pdf (accessed on 18 December 2025).
64. Voice, E. OVHcloud Launches Next-Generation of AI-Powered Cooling That Balances Performance and Environmental Responsibility in Its Data Centres. Available online: https://www.ecovoice.com.au/ovhcloud-launches-next-generation-of-ai-powered-cooling-that-balances-performance-and-environmental-responsibility-in-its-data-centres/ (accessed on 18 December 2025).
65. Microsoft. Microsoft Datacenters in Netherlands. Available online: https://local.microsoft.com/wp-content/uploads/2025/10/Microsoft-datacenters-in-Netherlands.pdf (accessed on 7 June 2026).
66. Swinhoe, D. ServerDomes to Use LiquidCool's Immersion Cooling Technology. Data Centre Dynamics. Available online: https://www.datacenterdynamics.com/en/news/serverdomes-to-use-liquidcools-immersion-cooling-technology/ (accessed on 10 December 2025).
67. ServerDome. ServerDomes Energy and Water Savings Report. Available online: https://serverdomes.com/resources/energy-and-water-savings-report/ (accessed on 10 December 2025).
68. atNorth. Sustainability Report 2024. Available online: https://www.atnorth.com/wp-content/uploads/2025/04/Sustainability-Report-2024-web.pdf (accessed on 7 June 2026).
69. Alibaba Group. Alibaba Cloud Launches New Energy-Efficient Qiandao Lake Data Center. Available online: https://www.alibabacloud.com/en/press-room/alicloud-launches-new-energy-efficient-qiandao-lake-data-center (accessed on 14 November 2025).
70. LuxConnect. LuxConnect's Commitment to Water Efficiency: A Leading Example in the Data Center Industry. Available online: https://www.luxconnect.lu/water-efficiency-datacenter-industry/ (accessed on 18 December 2025).
71. LuxConnect. Sustainability—It's Time to Be Green. LuxConnect. Available online: https://www.luxconnect.lu/sustainability/ (accessed on 18 December 2025).
72. Facebook Sustainability. Enhancing Water Efficiency Through Advanced Cooling Technology. Available online: https://sustainability.fb.com/wp-content/uploads/2020/12/FB_Enhancing-Water-Efficiency-Through-Advanced-Cooling-Technology.pdf (accessed on 7 June 2026).
73. Law, M. Microsoft Unveils Zero-Water Cooling for AI Data Centres. DataCentre Magazine. 10 December 2024. Available online: https://datacentremagazine.com/articles/microsoft-unveils-zero-water-cooling-for-ai-data-centres (accessed on 10 December 2025).
74. OVHcloud. OVHcloud Opens New Data Centre in Sydney, Powered by Latest Water Cooling Technology to Sustainably Meet Increasing Demands for High Compute Cloud Solutions. Available online: https://corporate.ovhcloud.com/en/newsroom/new-datacentre-sydney/ (accessed on 10 December 2025).
75. Mountain, G. Green Mountain Sustainability Report 2024. Available online: https://greenmountain.no/wp-content/uploads/Green-Mountain-Sustainability-report-2024.pdf (accessed on 14 November 2025).
76. Huawei Digital Power. Building a Low-Carbon Smart Future | Global Smart Green Data Center Summit 2023 Kicks Off. Available online: https://digitalpower.huawei.com/en/news/data-center-facility/658 (accessed on 28 May 2026).
77. AirTrunk. Airtrunk Enters Malaysia with new 150+ MW Hyperscale Data Centre in Johor Bahru. Available online: https://airtrunk.com/airtrunk-enters-malaysia-with-new-150mw-hyperscale-data-centre-in-johor-bahru/ (accessed on 18 December 2025).
78. AirTrunk. Sustainability Report FY25. Available online: https://airtrunk.com/wp-content/uploads/2025/10/AT_Sustainability-Report-2025-Final.pdf (accessed on 7 June 2026).
79. EcoDataCenter. EcoDataCenter 2023 Sustainability Report. Available online: https://ecodatacenter.tech/hubfs/EcoDataCenter_sustainability_report_2023_20240506_LQ.pdf (accessed on 7 June 2026).
80. Equinx. Equinix 2024 Sustainability Data Summary. 2025 Equinix Inc. Available online: https://www.equinix.com/content/dam/eqxcorp/en_us/documents/resources/data-sheets/ds_sustainability_data_en.pdf (accessed on 26 November 2025).
81. Equinx. Cooling. Available online: https://docs.equinix.com/colocation/about/colo-cooling-intro/ (accessed on 26 November 2025).
82. Better Buildings Better Plants. Iron Mountain Data Centers: Geothermal Cooling System. Better Buildings, U.S. Department of Energy. Available online: https://betterbuildingssolutioncenter.energy.gov/showcase-projects/iron-mountain-data-centers-geothermal-cooling-system (accessed on 18 December 2025).
83. Iron Mountain Data Centers. Iron Mountain Data Centers Sustainability Performance Overview 2024. Available online: https://s204.q4cdn.com/148941814/files/doc_downloads/2024-IMDC-Sustainability-Performance-Overview-FINAL.pdf (accessed on 7 June 2026).
84. switch. AI Factories. Available online: https://www.switch.com/ai-factories/ (accessed on 18 December 2025).

85. switch. ESG 2024 Report. Available online: https://switchdotcom.s3-sites.data.switch.com/media/reports/esg/Switch-ESG-Report-2024.pdf (accessed on 7 June 2026).
86. Sharma, P.; Pegus, P.; Irwin, D.; Shenoy, P.; Goodhue, J.; Culbert, J. *Design and Operational Analysis of a Green Data Center*; IEEE Computer Society: Piscataway, NJ, USA, 2017; Available online: https://homes.luddy.indiana.edu/prateeks/papers/mghpcc-mag.pdf (accessed on 14 November 2025).
87. Digital Realty. High-Density Colocation. Available online: https://www.digitalrealty.com/platform-digital/colocation/high-density-colocation (accessed on 14 November 2025).
88. Zhang, M. How Data Centers Impact the Environment. Dgtl Infra. Available online: https://dgtlinfra.com/how-data-centers-impact-the-environment/ (accessed on 13 March 2026).
89. NEXTDC. Power Play Inside The Data Centres Fuelling Our Digital Economy. Available online: https://www.nextdc.com/hubfs/160617_NEXTDC_Innovation_Final.pdf (accessed on 18 December 2025).
90. NEXTDC. FY25 Environmental, Social and Governance Report. Available online: https://nextdc.com/hubfs/ASX%20Announcements/FY25%20Environmental,%20Social%20and%20Governance%20Report.pdf (accessed on 7 June 2026).
91. Climate Neutral Data Centre Pact. Climate Neutral Data Centre Pact White Paper. 2024. Available online: https://www.climateneutraldatacentre.net/wp-content/uploads/2025/09/White-paper-on-DC-responsible-use-of-water-review-2024.pdf (accessed on 4 August 2026).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.