SUSTAINABLE DATA CENTERS ROADMAP

CHAPTER 2.3

Cooling Technologies

Alexis Abramson

October 2025

Innovation for Cool Earth Forum

2.3 Cooling Technologies

Alexis Abramson

A. Overview

i. History of data center cooling

Cooling infrastructure can be a major challenge for data centers, especially as rising power densities from artificial intelligence (AI), high-performance computing (HPC) and graphics processing unit (GPU)-intensive workloads drive the need for more advanced cooling technologies. The history of data center cooling is closely tied to the evolution of computing itself. As computer systems became more powerful, their thermal output increased, necessitating more sophisticated and efficient cooling methods. The journey from basic ventilation to advanced liquid and AI-optimized cooling reveals a story of continuous innovation driven by performance, reliability and energy efficiency. The American Society for Heating, Refrigeration, and Air Conditioning Engineers (ASHRAE) has been publishing and updating guidelines for data center cooling since 2004 and every 4-6 years thereafter.[1]

Cooling data centers has a lot in common with cooling other buildings, but there are some important differences. Typical residential and commercial buildings gain heat through windows and walls, appliances (e.g., refrigerators), and even people. Circulating cool air through the building via a central air conditioning system (i.e., treating the entire space as a single thermal zone) is the typical way to remove heat. A commonly employed rule of thumb for sizing building cooling systems is 1 ton of cooling for every 400-500 ft^2, which is equivalent to ~8 W/ft^2. In contrast, data centers have no windows, minimal occupancy and highly insulated walls. Their primary heat sources are the computer servers themselves—extremely localized loads that can exceed 100 W/ft^2. (See Chapter 1(A)(ii)(a) of this Roadmap.) This concentrated heat makes broad air circulation inefficient. Instead, cooling systems must be substantially over-provisioned and precisely engineered to prevent local "hotspots" within the equipment. In the mid-20th century, early mainframe computers were housed in large

Figure 2.3-1. *Understanding direct-to-chip cooling in HPC infrastructure: A deep dive into liquid cooling.*

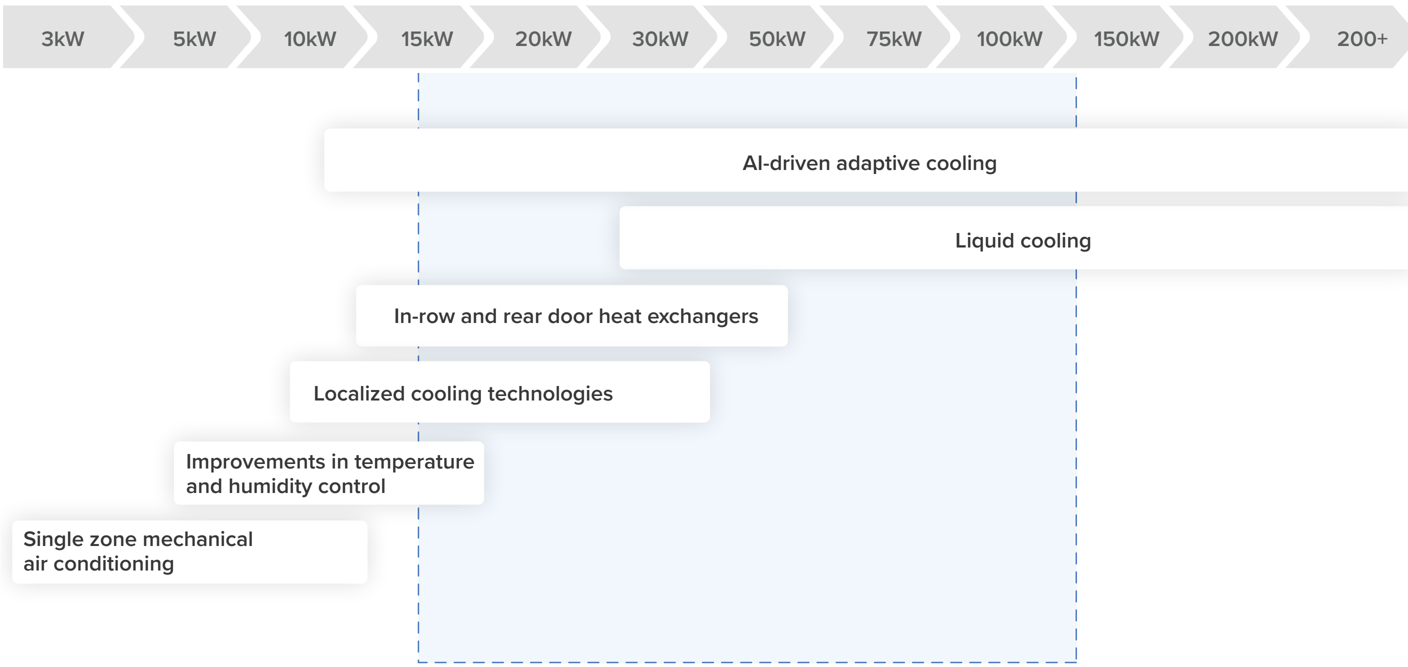

rooms and cooled primarily through ambient air and basic air conditioning systems.[3] Using direct expansion (DX), these systems pulled warm air from the single thermal zone, passed it over a refrigerant coil, and rejected the heat via condensers. Raised floors are used to distribute cool air and ceiling return vents for hot air extraction. As server density increased in the 1990s and 2000s, the inefficiencies of room-level cooling became more apparent. More sophisticated cooling systems, including Computer Room Air Conditioners (CRAC) and later Computer Room Air Handlers (CRAH), allowed for better temperature control and humidity regulation.

In the mid-2000s and into the 2010s, the proliferation of high-density server racks led to adoption of more localized cooling strategies. These servers pack a relatively high compute density, thereby requiring a significantly higher power draw. In high density racks, this power draw can vary from 15-30 kW, while in ultra high-density racks it can reach 100+ kW, leading to a much greater heating load. During this time period, computational fluid dynamics (CFD) modeling became a commonly used tool for facility-level design to simulate airflow and temperature distribution. CFD enables data center designers to optimize cooling strategies, prevent hotspots and improve overall energy efficiency. In-row cooling and rear door heat exchangers (RDHx) became popular for targeting heat removal closer to the source. These heat exchangers can be mounted on the rear door of server racks to remove heat directly before the air can re-enter the data center environment. While they have not captured a significant market share, this technology has emerged and is evolving as an attractive option to address increasingly dense and power-intensive server racks. (See Chapter 2.4.)

With the advent of AI, HPC, and GPU-heavy workloads over the past decade, traditional air cooling approaches have begun to hit their limits. Liquid cooling—through direct-to-chip cold plates, immersion cooling and two-phase systems—has seen an increase in the marketplace. These systems offer superior heat transfer and are often more energy efficient, especially for workloads associated with high-density racks where power densities range from 30-100+ kW per rack.[4]

The selection of specific cooling systems depends not only on data center needs but also on local conditions, enabling them to be more climate responsive and to optimize for performance, costs, efficiency, thermal management, water use and reliability. Globally, the selection of cooling systems can be particularly influenced by geographic climate factors:

- Hot-humid climates (e.g., Singapore): Employ a combination of robust air conditioning, often supported by chilled-water systems. Advanced systems, such as desiccant dehumidification and various forms of liquid cooling, have increasingly been adopted.

- Hot-dry climates (e.g., Phoenix): Rely on evaporative cooling, which is highly efficient and often used in conjunction with traditional air cooling as a backup. Also employed is adiabatic cooling, which uses the evaporative cooling effect to pre-cool intake air. Free cooling (see Figure 2.3-1) may be integrated in the system to take advantage of cooler outdoor temperatures. The availability of water and its potential usage in these systems is an important consideration.[5]

- Temperate climates (e.g., Frankfurt): Deploy hybrid systems that integrate free cooling to leverage moderate outdoor temperatures combined with conventional air cooling mechanical systems optimized for seasonal shifts.

- Cold climates (e.g., Sweden): Adopt free cooling via air-side or water-side economization used most of the year.[6] Ample waste heat generated by servers has also been used for heat recovery applications. (See Chapter 2.4.)

Box 2.3-1

Free cooling aka economization

In cooler climates, free cooling, also known as economization, can eliminate much of the need for energy-intensive, air-cooled mechanical systems. Leading data centers in Scandinavia, Canada and the US Pacific Northwest report significant energy savings using air-side economizers, as well as natural or readily available sources of cold water, to absorb and dissipate the heat.[7] While attractive in theory, some have found free cooling challenging to implement. For example, Microsoft's Project Natick and Subsea Cloud's "Jules Verne" pod have demonstrated how to leverage ocean water for passive cooling, but the project was halted in 2024.[8-11] Meanwhile, in orbit, space data center initiatives, such as the Hewlett Packard Enterprise (HPE) Spaceborne Computer-2 on the International Space Station, utilize the naturally cold environment of space to reject heat; however, challenges remain due to the large surface area required for the radiative heat transfer.[12]

Today, air cooling still dominates in data centers-and is the default choice for many legacy environments because it is widely understood, relatively low cost to install and maintain, and remains effective for facilities with moderate power densities. Even more, new demands of modern data centers are driving an urgent need for more intelligent, adaptable and scalable cooling innovations to sustain future performance and efficiency.

ii. Cooling and its impact

Nearly all electricity consumed by servers is ultimately converted into heat. As server densities rise, so does the cooling load (the amount of heat that must be removed), driving the need for more advanced cooling technologies. Without effective cooling, thermal stress can reduce performance and hardware lifespan, making efficient thermal management an imperative for both reliability and sustainability.
Many efficient cooling technologies, such as evaporative cooling, dramatically reduce electricity use but may also require substantial water, depending on the cooling mechanism.[3] Air conditioning systems that rely on mechanical DX refrigeration use little or no water but demand higher energy input, increasing carbon emissions when powered by fossil fuels. Optimizing for one resource—such as water or energy—can often increase the burden on another.

Power usage effectiveness (PUE) and water usage effectiveness (WUE) provide an indication of the performance of cooling technologies in data centers. PUE is a data center's total energy use divided by the energy its information technology (IT) equipment uses. A lower PUE (closer to 1.0) indicates more efficient cooling and operations. Efficient operators now report PUEs below 1.2, but the global average remains closer to 1.5.[13] WUE represents a data center's water use divided by the energy its IT equipment uses. WUE is typically expressed in liters per kilowatt-hour (L/kWh). Highly efficient data centers can achieve WUE values as low as 0.2 L/kWh or even close to 0 L/kWh, which effectively means the facility uses no water for cooling purposes.[14] Data centers that rely heavily on evaporative cooling are likely to have the highest WUEs, potentially up to 2.5 L/kWh.[15] Target WUEs for a particular data center should be considered in conjunction with PUE, local environmental conditions and water availability. As sustainability becomes a priority, operators and designers must consider both resources to fully appreciate the overall impact.

Cooling systems also contribute to carbon emissions through their use of refrigerants. Some refrigerants, such as hydrofluorocarbons (HFCs), are greenhouse gases with high global warming potential (GWP) and must be managed to avoid leakage or improper disposal. In fact, so-called F-gas leakage-the unintended release of fluorinated gases-can equate to massive amounts of carbon emissions. Various global regulations are already mandating the phase-down and eventual phase-out of high-GWP F-gases in equipment. This is pushing data centers toward lower-GWP alternatives, improved leak-detection protocols and, in some cases, wholesale replacement of legacy systems.[16,17] (See Chapter 3.1-Scope 1 Emissions.)

Lifecycle analyses of data center cooling are now more commonly conducted for quantifying not just operational energy use but also the emissions and resource consumption associated with manufacturing, deploying and disposing of cooling equipment. Such assessments have shown that advanced cooling methods like cold plates and immersion cooling (see below) can reduce greenhouse gas emissions by 15-21% over the entire lifecycle compared to conventional air cooling systems.[18] So while the increasing power demands of modern computing is essentially forcing the industry to adopt more efficient liquid cooling methods, these technologies are inherently better for the climate overall. This makes liquid cooling a compelling and necessary shift for the future of digital infrastructure. A study by Microsoft (2023) confirmed that while liquid cooling systems have higher manufacturing impacts (due to pumps, cold plates and materials), they can offer lower total lifecycle emissions in high-density environments by enabling energy savings and improved PUE.[18] The study also emphasized the need to account for cooling infrastructure replacement cycles, refrigerant GWP and end-of-life impacts when comparing technologies.

Sustainable data center cooling now requires multi-metric optimization, balancing energy, water, carbon and local resource availability, such as water scarcity (if applicable), grid capacity and renewable energy availability. This means moving beyond simply achieving the lowest PUE to performing holistic and complex assessments of energy, water, carbon, and the local environmental and social context of these systems.

B. Current Cooling Technologies and Associated Enhancements

Table 2.3-1. *Data center cooling systems and PUEs.*[6,19-25]

Data Center Cooling System	Typical PUE Range	Notes
Air-Based Cooling		
CRAC units with DX (no enhancements)	1.8–2.5	Small to medium data centers
CRAH units with chilled water (no enhancements)	1.7–2.2	Medium to large data centers
CRAC/CRAH + hot/cold aisle containment	1.5-2.0	May include hot or cold containment or combined
CRAC/CRAH + economizers	1.2-1.6	“Free cooling”; can be air-side or water-side; depends highly on local climate
CRAC/CRAH + modern enhancements	1.04–1.4	May include variable speed drives in compressors; enhanced controls
CRAC/CRAH + adiabatic cooling/pre-cooling	1.2–2.0	Depends on efficiency of primary system and local climate
Direct evaporative cooling	1.1-1.3	Highly dependent on local climate
Liquid-Based Cooling		
Direct-to-chip liquid cooling	1.05-1.2	CRAH still required for heat removal to environment
Single-phase immersion cooling	1.02–1.05	No server fans; no air-based systems
Two-phase immersion cooling	1.01–1.03	Relies on thermosiphon effect rather than pumps

i. Air-based cooling technologies

Air remains the most widely used cooling medium in data centers due to its simplicity, cost-effectiveness and compatibility with existing IT infrastructure. But liquid cooling is gaining traction. In a 2024 survey of almost 1000 data center owners and operators, 95% reported relying predominantly on air-cooled systems, but 94% also reported at least some use of liquid cooling, including 17% who reported substantial use of liquid cooling technologies.[26] Over time, air-based cooling strategies have evolved from room-level temperature control to highly localized, intelligent and climate-responsive systems that improve energy efficiency and thermal performance.

Traditional CRAC units and air handling units (AHUs) are foundational components of air-based cooling. CRACs typically use DX systems to cool air, while CRAHs and AHUs circulate chilled air from centralized systems. These systems supply conditioned air—typically cooled and sometimes dehumidified—either through raised-floor plenums or via overhead ductwork, depending on the facility design. In a raised-floor setup, cool air is pushed into the underfloor space and directed upward through perforated or directional vent tiles, which are strategically placed in front of server racks to deliver air directly to equipment inlets. In overhead systems, ductwork distributes cool air from above, typically aimed down into cold aisles.

To manage heat generated by the equipment, the systems rely on return air plenums, often located in the ceiling, to collect hot exhaust air from the back of the server racks. This hot air is then routed back to the cooling units for reconditioning and recirculation. To improve efficiency and prevent mixing of hot and cold air, many modern facilities incorporate hot aisle containment. This strategy encloses the hot aisle—where exhaust air is expelled—in a physical barrier, such as doors or ceiling panels, which channels the hot air directly into the return system. By maintaining a clear separation between hot and cold air streams, hot aisle containment enhances thermal management and allows for higher cooling efficiency.[27]

In temperate climates, air-side economizers allow filtered outside air to cool the data center when ambient temperatures are favorable, minimizing mechanical cooling. This technique, known as free cooling, can yield substantial energy savings and lower PUE. Microsoft was among the first hyperscalers to report successful deployment of air-side economization in production environments, showing that well-controlled airflow and outdoor air integration could support high-performance operations with PUEs below 1.2.[28] At Alibaba's Qiandao Lake data center (2015), a combination of natural ventilation, free cooling and computational fluid dynamics (CFD) modeling reduced energy consumption significantly, especially during cooler months.[29] Free cooling was achievable for over 90% of the year, leading to cooling cost reductions of approximately 80% and a PUE below 1.3.

Data centers have increasingly adopted evaporative and adiabatic cooling as energy-efficient mechanisms to enhance traditional air-based cooling systems. These systems leverage the natural cooling effect of water evaporation to reduce air temperature. In direct evaporative cooling, outside air is humidified and cooled as it passes through a wetted medium, then it is introduced into the data center. Adiabatic cooling, often used in indirect systems, pre-cools air or a heat exchange surface without introducing moisture directly into the data center environment. Both approaches are particularly effective in dry or temperate climates, where low humidity allows for significant cooling potential. Major operators, such as Amazon, Facebook and Google, have deployed evaporative or adiabatic cooling systems in regions like Iowa, Oregon and Utah, achieving lower PUE and reducing dependence on energy-intensive chillers.[30] These methods are often integrated into multi-mode cooling systems that shift between economization, evaporative cooling and mechanical backup based on real-time environmental conditions.

Case Study 1: Meta's MeeFog—pushing the limits of cooling with evaporative precision

In 2011, Meta (then Facebook) broke new ground with its MeeFog evaporative cooling system, a high-efficiency design that leveraged adiabatic cooling to generate an ultrafine water mist to pre-cool incoming outside air. By using high-pressure atomization to saturate air without over-humidifying it, the system enabled Meta to cool data halls without the traditional chillers used in air-based mechanical systems—even in warmer climates.

The result? A dramatic drop in cooling energy use and a PUE as low as 1.06 at facilities like Prineville, Oregon—years ahead of industry norms.[31] Crucially, the MeeFog system still maintained ASHRAE-recommended temperature and humidity ranges, demonstrating that ultra-efficient evaporative cooling could be both high-performing and reliable. Meta is continuously exploring and implementing upgrades to the MeeFog system. These upgrades are likely to push PUE and WUE even lower, while extending the viability of evaporative cooling to more variable climates.

Recent advances in AI/machine learning (ML) enable dynamic, software-driven thermal management, allowing systems to adapt cooling in real-time based on workload and environmental conditions. In 2022, Meta deployed a reinforcement learning system to optimize the performance of air-handling units located in the penthouse—the mechanical space typically situated on the roof of its data centers that houses key heating, ventilation and air conditioning (HVAC) equipment. This led to a 20% reduction in fan energy usage and improved temperature compliance.[32] In collaboration with Trane Technologies, DeepMind published an article showing actual cooling energy reductions of 9-13% in commercial HVAC systems via reinforcement learning.[33] Earlier, a 2016 DeepMind blog reported up to 40% cooling energy savings and a 15% PUE reduction in Google data centers using ML controls.[34]

ii. Liquid cooling technologies

As data centers face rising rack power densities, traditional air cooling methods are struggling to keep pace. Current hyperscale racks now routinely operate above 100 kW. Looking ahead, Nvidia's roadmap anticipates a stepwise progression to approximately 600 kW by 2027 and possibly beyond megawatt-class densities by 2030.[35] A massive new data center project in Texas, backed by OpenAI and Oracle, will be relying on a liquid cooling system to manage the intense heat from its high-density AI servers. This approach involves a continuous circulation of millions of gallons of chilled water, highlighting how hyperscalers are adopting advanced cooling methods to handle extreme computing loads.[36] Sophisticated liquid cooling technologies will remain essential for managing thermal loads efficiently while improving sustainability metrics like PUE and WUE. For most liquid cooling technologies, WUE values are at or near 0.0 since closed-loop systems can be employed. This is because cooling fluids continuously circulate within a sealed, self-contained circuit without being exposed to the outside environment or losing fluid through evaporation or discharge. Thus, in contrast to water-intensive systems, such as evaporative cooling towers, adiabatic coolers, and open-loop water-cooled chillers (which rely heavily on water consumption to reject heat), liquid cooling technologies reduce or even eliminate water use while maintaining efficient thermal management.[5]

Liquid cooling systems in data centers use a range of fluids, including water-based mixtures for indirect cooling methods (e.g., cold plates and rear-door heat exchangers) and dielectric fluids for direct-contact approaches (e.g., immersion and spray cooling), depending on whether electrical isolation is required. These fluids have a much higher thermal capacity than air, and therefore liquid cooling can be thousands of times more effective at heat removal per unit volume than air-based systems.[37,38] The higher heat transfer efficiency can unlock new levels of compute performance in constrained thermal environments where there might be limitations on heat removal.

Innovations in IT infrastructure and liquid cooling are enabling inlet temperatures to be raised to 30-45 °C (up from 18-24 °C), allowing liquid loops to run at hotter temperatures and reducing mechanical cooling requirements.[39,40] Additionally, these modern IT enhancements are enabling cooling systems to reject heat even in hot climates where passive rejection of low temperature heat would not be thermodynamically feasible due to the lack of sufficient temperature differential. Operating at higher loop temperatures also increases the exergy—or useful work potential—of the waste heat, making it easier and more efficient to transfer heat for secondary uses. This opens the door to a wider range of reuse applications, including district heating, industrial processes or absorption cooling. With the addition of a heat pump, producing steam may even be possible. (See Chapter 2.4.)

The two most commonly employed advanced liquid cooling methods for data centers are direct-to-chip (DTC) (Figure 2.3-3) cooling and immersion cooling. DTC cooling uses cold plates or microchannels mounted directly on processors and GPUs to absorb heat, transferring it through the liquid loop to an external heat exchanger or a dry, air-based cooler.[2] DTC has the advantage of integrating with existing racks but the disadvantage of requiring plumbing for cooling distribution. In 2012, Google was one of the first to deploy cold plate DTC technology to support analytics workloads that pushed power densities beyond 20-30 kW, a level that air cooling alone could not handle. Google laid the groundwork for today's generation of liquid-cooled, high-density racks, providing an early real-world demonstration that chip-level cooling could be scalable, safe and cost-effective in hyperscale environments.

Figure 3.3-4. *Immersion Cooling for Data Center Cooling.*

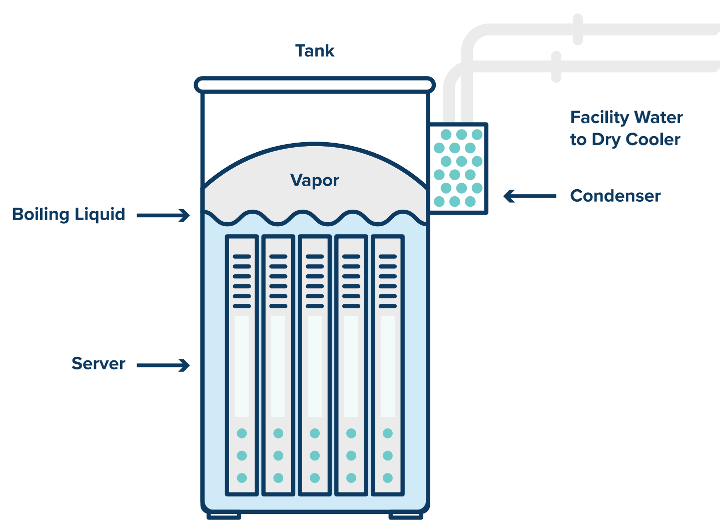

Advances in the DTC sector are accelerating to overcome key technical barriers. For example, ZutaCore's innovative cooling technology, called HyperCool, is a waterless, two-phase, DTC liquid cooling, closed-loop system that uses a non-conductive, dielectric fluid that evaporates directly at the chip level, absorbing heat as it changes phase from liquid to vapor.[41] This approach enables effective cooling of ultra-high-density server racks (up to 250 kW per rack) and supports next-generation chips drawing over 2800 W, while reducing cooling energy use by up to 80%. Another company, Iceotope, offers a solution that combines the compact, targeted efficiency of DTC liquid cooling with the thermal stability and enclosure advantages of tank-based immersion, enabling precise heat removal while protecting IT equipment from environmental contaminants and airflow-related issues.

Case Study 2: DTC with Microsoft and AWS

In 2024, Microsoft began a phased rollout of closed-loop direct-to-chip (DTC) liquid cooling systems at its new AI-focused data centers—making a shift from evaporative cooling in response to the rising heat loads associated with high-density compute.[42] The pilot deployments will be at locations in Phoenix, Arizona and Mount Pleasant, Wisconsin.

Microsoft's goal is to evaluate these systems, considering both energy and water use holistically, to improve overall sustainability while meeting demanding thermal requirements. The deployed system uses a glycol-based coolant in a sealed loop, coupled with dry coolers and adiabatic assist, which pre-cools incoming air using a fine water mist when ambient temperatures allow. While this move resulted in a modest increase in energy consumption and PUE compared to traditional evaporative systems, it reflects a deliberate tradeoff to support long-term water resilience and environmental sustainability of Microsoft's expanding global AI data center fleet.

In 2025, Amazon Web Services (AWS) unveiled one of the most versatile and scalable thermal management strategies to date: a flexible multimodal cooling architecture engineered specifically for next-generation AI workloads. Designed and delivered in just 11 months, the hybrid system uses both air-based cooling and DTC cold plates to extract heat from central processing units (CPUs) and accelerators, circulating a non-evaporative coolant through a sealed-loop system connected to air-based heat rejection units outside the data hall.[43] Unlike immersion systems, this solution maintains compatibility with standard server form factors and supports easy retrofitting across AWS's global infrastructure. It integrates with AWS's building management system for centralized thermal optimization and enables cooling of racks exceeding 80 kW—all while dramatically reducing reliance on water.

AWS reports that their newly launched cooling system can cut cooling energy use by up to 46% during peak cooling periods based on internal comparisons with its legacy system design.[44] While specific PUE values are not disclosed, AWS noted that in 2023, its data centers achieved an average global PUE of 1.15 with the best performing facility reaching 1.04.[45]

In contrast to DTC, immersion cooling (Figure 2.3-4) submerges servers in a non-conductive dielectric fluid that directly absorbs and removes heat, providing uniform cooling to all components. Either the fluid can be circulated through a heat exchanger and returned (single phase immersion), or it can be made to boil at low temperatures on contact with hot components, leading the vapor to rise, condense on a heat exchanger, and return to the bath (two-phase immersion). The technology reduces fan power use to near zero (although pumps are now required to circulate the fluid), supports high rack densities (100-250 kW+), and can extend equipment lifespan by eliminating dust and humidity exposure. And while immersion cooling can deliver superior thermal performance and energy efficiency, its deployment often requires significant modifications to data hall infrastructure, including changes to rack configurations, service workflows and facility layout. As a result, it is typically more practical for use in new builds or purpose-designed expansions. In contrast, DTC liquid cooling is more easily integrated into existing data centers, as it aligns with conventional rack architectures and is therefore better suited for retrofits. Nonetheless, innovation in this space is growing, particularly due to the increasing demands of AI.[46] For example, Green Revolution Cooling (GRC) was one of the first companies to commercialize immersion cooling at scale and has seen deployments across more than 20 countries.[47] Another industry example, Submer, is marketing its single-phase immersion technology using a biodegradable dielectric fluid along with a forced convection heat sink developed in partnership with Intel to cool extremely dense chips, overcoming prior density and heat dissipation barriers.[48]

Many data centers combine liquid and air cooling into hybrid systems, using liquid to manage hotspots (e.g., CPUs, GPUs) and air for bulk heat removal. (See Case Study 2 above.) One additional benefit of the hybrid approach is the extended useful life of air-based cooling infrastructure by enabling it to coexist with targeted liquid systems—minimizing capital disruption. To enhance air-based cooling, rear door heat exchangers (RDHx) are increasingly used in data centers to efficiently cool high-density racks by immediately capturing the hot air that servers expel—known as server exhaust—as it leaves the back of the rack.

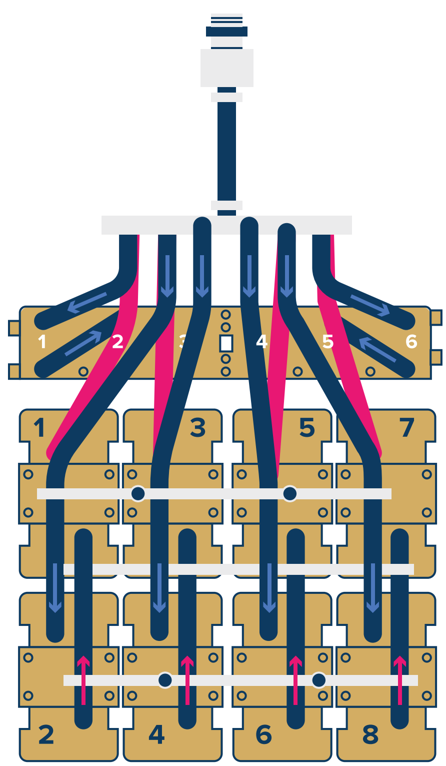

Table 2.3-3. *Direct-to-chip cold plate liquid cooling for high-heat-density data centers.*

This heat is then transferred to a chilled liquid loop, preventing it from entering the data hall and reducing the overall cooling load. RDHx systems enable higher rack densities, reduce reliance on traditional room cooling, and support energy efficiency and sustainability goals. Currently popular for retrofits and new builds alike, RDHx are evolving with smarter controls and higher capacity designs to meet demands of AI and edge computing workloads.

iii. Heat reuse in data centers

As described in more detail in Chapter 2.4, data centers convert nearly all the electricity they consume into heat—typically considered a waste product. But with growing energy and climate concerns, data center operators are increasingly viewing waste heat as a valuable resource. When captured and redirected, this thermal energy can support district heating systems, greenhouse agriculture, industrial processes or building-level heating, offering both carbon reductions and community benefits. One interesting approach is repurposing waste heat to drive adsorption cooling systems, which use heat rather than electricity to regenerate a desiccant material—providing an energy-efficient way to produce chilled water or air for additional cooling loads.[50,51]

In colder regions, the most impactful use of waste heat is often through district heating networks—centralized systems that distribute hot water or steam to buildings for space conditioning and water heating. In 2020, Meta began routing low-grade heat from its Odense data center to the local district heating utility (Fjernvarme Fyn).[52] Through heat pumps, the temperature is raised to meet residential needs. The system provides essentially “free” heat to over 7000 homes (and growing), offsetting fossil fuel use and supporting Denmark’s national climate targets. Following Meta’s lead, Microsoft is launching a heat recovery system at its new Danish data centers to supply thermal energy to the local district heating network. Scheduled for full operation by 2026, the system will deliver surplus heat that is equivalent to thousands of households’ annual demand.[53] The system forms part of Microsoft’s broader strategy to align with Denmark’s clean energy and circular economy goals.[54]

As compute demand grows, especially in temperate and cold regions, heat reuse could become a key pillar of sustainable data center design—transforming facilities from energy consumers into climate-positive infrastructure components.

iv. Other innovations and trends in data center cooling

Beyond air and liquid cooling fundamentals, a wave of emerging technologies and design philosophies are reshaping how data centers manage heat, water and infrastructure resilience. These innovations aim to further reduce energy use, minimize climate impacts and support the deployment of flexible, sustainable data center ecosystems. Technologies being further developed and tested include the following:

- **Thermosyphons** are sealed, gravity-driven heat transfer loops that use phase change (evaporation and condensation) of a working fluid to move heat with no mechanical input. One study suggests these devices have the potential to cut cooling energy use by up to 45% under optimal conditions compared to conventional air-based cooling systems operating under the same heat rejection requirements.[55] They offer near-zero-energy heat rejection and are increasingly integrated into free cooling systems for high-efficiency data centers. Thermosyphon-like technologies also can be found in various cooling solutions available today, but looking ahead, AI-optimized thermosyphons embedded with sensors and adaptive controls could dynamically tune their performance to match varying workloads and environmental conditions.

- **Adsorption cooling** uses materials like silica gel or zeolites that absorb moisture and release heat, then regenerate when heated. These systems can reuse low-grade waste heat to drive their cycles (without requiring significant electrical input)—ideal for pairing with high-efficiency CPUs and heat recovery loops. Recent studies have demonstrated effective integration of adsorption chillers with data center waste heat recovery loops.[50,51]

- **Desiccant systems** remove humidity from intake air before use of mechanical or evaporative cooling, increasing the effectiveness of both strategies and reducing microbial risk—particularly valuable in hot-humid climates. A 2023 simulation study focused on Southeast Asia demonstrated greater than 25% cooling energy savings when using a liquid-desiccant free cooling system compared to a conventional vapor-compression air conditioning system.[56] Desiccant systems could play a role in adaptive, climate-responsive cooling architectures, enabling data centers in hot and humid climates to reduce PUE and increase resilience while maintaining tight control over air quality and corrosion risk.

- **Thermal batteries** store excess or off-peak cooling energy (e.g., in chilled water/ice, reservoirs or solid media) for use during peak demand, improving load balancing and reducing strain on the grid. They are especially useful in regions with variable renewable energy supply. For example, Exowatt's P3 system uses concentrated solar heat stored in solid thermal storage media to provide around-the-clock power and cooling dispatch.[57] When needed, this thermal energy is converted to electricity using a thermophotovoltaic (TPV) generator, which works by first using the stored heat to warm a specialized material until it emits in the infrared to a photovoltaic system, which does the final conversion. A recent National Renewable Energy Laboratory (NREL) technoeconomic analysis of a 5 MW data center demonstrated a thermal energy storage system that reduced the levelized cost of cooling to $5/MWh compared to $15/MWh for conventional chillers and dry coolers.[58]

- **Phase change materials** (PCMs) passively absorb and release large amounts of thermal energy at specific temperatures. Embedded in server racks or walls, PCMs can buffer temperature spikes during demand surges or cooling system failures. For example, Phase Change Solutions offers BioPCM® and Apollo™ Smart Panels, which are widely used in telecom shelters and edge data center applications to stabilize indoor temperatures and reduce cooling energy demand.[59] A 2021 review highlights that integrating PCMs into data center infrastructure enhances thermal management and provides more precise control of temperature, improving overall energy efficiency and reliability.[60]

- **Refrigerants** are under increasing scrutiny due to their GWP. Next-generation systems are shifting toward low-GWP or natural refrigerants (e.g., R-1234yf, carbon dioxide (CO_2), ammonia), in line with global reduction mandates.[16] For example, Vertiv recently launched its Liebert® Evaporative Free Cooling (EFC) unit for data centers, featuring a low-GWP refrigerant (called “R-454B”), which is quickly becoming the new HVAC industry standard to enable significant efficiency gains and carbon footprint reduction.[61] Oak Ridge National Laboratory’s field-scale demonstrations are revealing that replacing conventional high-GWP refrigerants with ultra-low GWP fluids can reduce a data center’s annual CO_2 emissions by up to 40%.[16] In the United States, regulatory frameworks like US Environmental Protection Agency (EPA) Rule 23 requires new residential and light-commercial HVAC systems to use refrigerants with a GWP of 700 or less. In the United Kingdom, under the retained EU F-gas Regulation and domestic Net Zero Carbon Buildings Standard, refrigerants must generally have a GWP of 677 or lower to qualify as low-GWP products.[62]

- **Modular cooling architectures** support rapid deployment, localized upgrades and climate-specific customization. Scaling modular cooling reduces vendor lock-in, accelerates innovation through openness and facilitates global alignment on cooling best practices. Modular units can be factory-built and pre-engineered with liquid or hybrid systems, reducing on-site labor, construction time and permitting challenges. Modular cooling architectures emerged in the late 2000s and early 2010s with early adoption by hyperscalers and vendors offering in-row, rear-door and containerized solutions that enabled faster deployment and greater energy efficiency. However, broad adoption was initially limited by higher upfront costs and integration challenges with legacy infrastructure. A 2024 study found that these systems can be deployed in less than 12 weeks (compared to 18-24 months for traditional builds) with a reduction of on-site labor of over 60%.[63] The Open Compute Project (OCP) is promoting interoperability for modular solutions across Original Equipment Manufacturers (OEMs) and hyperscalers and is enabling faster adoption of high-density, high-efficiency cooling solutions.[64]

- **Advanced heat exchangers** are critical to enhancing liquid-cooled and hybrid systems (see case study inset).

Case Study 3: Liquid heat exchangers at scale—Microsoft pushes thermal boundaries

As the race to cool high-density AI workloads intensifies, Microsoft made major strides in 2024 deploying cutting-edge liquid heat exchanger technologies designed to reduce water use, shrink energy footprints and future-proof data centers against rising thermal loads.

Microsoft's 2024 launch of next-generation liquid-to-air heat exchanger units improves heat rejection by circulating warm coolant from its closed-loop, DTC cooling system through compact ambient-air heat exchangers. This allows servers' heat to be expelled directly to outside air without evaporative cooling, thereby eliminating water use while still efficiently removing thermal energy from the data hall. The system features advanced surface engineering to enhance thermal transfer efficiency.[65] While Microsoft does not provide details, surface engineering may include features such enhanced microchannel designs and/or specialized coatings. This innovation supports Microsoft's broader sustainability targets.

C. Barriers

While innovation in data center cooling is accelerating, widespread adoption of advanced systems faces a range of technical, economic and regulatory challenges. These barriers affect decisions across facility design, operations and long-term investment strategy and must be addressed to scale energy- and water-efficient cooling globally.

Table 2.3-2. Barriers.

Barrier	Emergine solution or mitigation path
Lack of technology maturity leading to inconsistent standards	Open standards (e.g., OCP, ASHRAE Advanced Cooling Solutions); industry-led interoperability efforts
Supply chain constraints, such as long lead times	Localized sourcing; multi-vendor compatibility; simplified system design; modularity
Increased system complexity/reliability requiring monitoring, training, maintenance	AI/ML-powered smart controls; predictive maintenance tools
Capital and operational cost uncertainty	Total cost of ownership (TCO) modeling; case studies; policy-backed green financing tools
Water scarcity	Shift to closed-loop, zero-water systems; adopt advanced dry or air-side cooling
Retrofit constraints, such as space limitations, airflow paths	Use of modular, edge-integrated cooling units; retrofit incentives
Regional regulatory considerations (e.g., building codes, state-level energy codes, water and environmental regulations)	Early engagement with permitting authorities; harmonized environmental and building codes; regional tech pilots
Market/policy levers undervalue energy/water savings	Carbon pricing; water-use transparency mandates; performance-based energy incentives or tax credits

D. Recommendations

Achieving scalable, sustainable cooling across the data center industry will require coordinated engagement from local and national governments, industry, academia and other stakeholders. The recommendations below are organized by stakeholder group and grounded in ongoing industry best practices, emerging research and field-tested case studies.

1. ***Local governments should:***
    a. ***Work with utilities to expand access to non-potable or recycled water sources.***
    b. ***Offer data centers incentive structures****—such as tax abatements or expedited permitting—****tied to clear sustainability performance benchmarks****, such as targets for PUE, WUE and heat recovery ratios.*

2. ***Local governments in regions with naturally favorable climates for cooling should:***
   a. ***Actively promote this advantage to data center operators*** *and develop targeted incentives to attract new facilities, positioning their areas as energy- and cost-efficient locations for sustainable data center development.*
   b. *Accelerate deployment of advanced cooling by* ***streamlining permitting*** *for projects that integrate sustainable thermal management strategies, such as free cooling, heat reuse or closed-loop systems.*
   c. ***Update zoning regulations*** *to enable co-location of data centers with facilities that can use waste heat, such as greenhouses or municipal buildings.*

3. ***Local governments in regions with favorable conditions for thermal integration into district energy systems should work directly with data center operators*** *who understand regional opportunities on such projects.*[52]

4. ***National policymakers should:***
   a. ***Establish the market conditions and regulatory frameworks*** *necessary for the broad adoption of energy- and water-efficient cooling technologies.*
   b. ***Create and enforce minimum energy performance standards for data centers,*** *along with voluntary or mandatory reporting requirements for PUE and WUE; expand programs, such as the US EPA's ENERGY STAR for Data Centers*[66] *or the EU Code of Conduct for Data Centres,*[67] *which can help set performance baselines and identify leaders.*
   c. ***Use government funding to support research and pilot programs*** *for promising but commercially immature technologies, such as two-phase immersion cooling, modular systems, thermal batteries and zero-water liquid cooling.*
   d. ***Provide tax credits, green bonds and procurement incentives*** *to help de-risk early adoption and support widespread deployment of sustainable cooling systems.*

5. ***Universities and research institutions should:***
   a. ***Prioritize studies of novel cooling techniques,*** *including many of the innovations described above, such as desiccant-based systems, thermal batteries, AI-optimized thermal control platforms and climate-specific hybrid systems.*
   b. ***Host experimental testbeds*** *or collaborate with industry to evaluate the performance of emerging solutions in field conditions.*
   c. ***Create open access datasets, simulation tools and digital twins*** *to allow broader communities to model, compare and benchmark advanced cooling approaches; standardizing such tools will improve planning accuracy and reduce design risk.*[68]
   d. ***Create or expand curricula on thermal systems, green data infrastructure and resilient design*** *to train the next generation of engineers and planners.*

6. ***Standards organizations** such as ASHRAE, International Standards Organization (ISO) and OCP should:*
   a. ***Facilitate innovation and interoperability** to evolve their guidance.*
   b. ***Establish uniform testing protocols and certification pathways** to validate performance of new technologies—especially liquid cooling, rear-door heat exchangers and high-efficiency refrigerants.*
   c. ***Push for global alignment on definitions and performance thresholds** to lower costs, reduce vendor lock-in and allow data center operators to deploy advanced cooling with greater confidence across international markets.*

7. ***Cooling equipment manufacturers,** bridging the gap between research and widespread implementation, should:*
   a. ***Invest in research and development (R&D)** focused on compact cold plates, advanced heat exchanger surfaces and system-integrated controls with predictive maintenance and AI optimization capabilities.*
   b. ***Offer comprehensive, modular solutions** that include sensors, telemetry and leak detection to reduce operational complexity.*
   c. ***Prioritize low-GWP and natural refrigerant alternatives,** consistent with the climate goals outlined in the Montreal Protocol Kigali Amendment, particularly ss regulations around refrigerants evolve.*[16]

8. ***Data center developers and operators should:***
   a. ***Integrate design early in the project development process,** especially in siting decisions.*
   b. ***Select locations** that enable the use of free cooling, heat reuse or access to non-potable water and renewable energy.*
   c. ***Install climate-appropriate cooling systems**—such as evaporative cooling in dry regions or air-side economization in temperate zones—in tandem with IT deployment strategies.*
   d. ***Establish facility-level energy and water performance targets and publish sustainability metrics annually.***
   e. ***Evaluate and design for heat reuse opportunities,** either through district heating connections or local use cases like agricultural greenhouses, building heating or industrial preheating.*
   f. ***Set aside dedicated infrastructure for pilot deployments** of emerging cooling systems, allowing testing without disrupting core operations.*

9. ***Utilities should:***
   a. ***Partner with data centers to support load shifting** (to align significant workload periods with cooler times of day).*
   b. ***Integrate waste heat into community heating systems.***

c. ***Offer incentive structures for grid-responsive cooling*** *in which data center cooling systems adjust their operation in response to signals from the electric grid.*

10. ***Environmental organizations should advocate for:***
    a. ***Low-carbon and water use***
    b. ***Transparent reporting***
    c. ***Waste heat reuse***
    d. ***Responsible siting*** *to align with climate and sustainability goals*

11. Investors and financiers should require disclosures *on WUE, PUE and refrigerant use in project finance deals because these metrics directly impact a data center's operational efficiency, climate risk exposure, regulatory compliance and long-term sustainability performance—all of which influence financial returns, reputational risk and alignment with Environmental, Social, and Governance (ESG) commitments.*

12. Insurance providers should reduce risk premiums for data center operators that adopt redundant and fault-tolerant cooling systems*—especially those with active leak detection and real-time monitoring—because these technologies significantly lower the likelihood of costly outages, equipment damage, and water or refrigerant leaks, thereby reducing the insurer's exposure to operational and environmental claims. For most companies, this is a shift from current practice, where premiums often do not fully account for the added risk mitigation these systems provide.*

13. End-use customers*, such as large cloud clients,* ***should shape demand by requiring data centers to meet high-efficiency and low-emissions cooling benchmarks*** *in service agreements, such as maintaining a low PUE (ideally below 1.3) and minimizing greenhouse gas emissions by using low-GWP refrigerants, carbon-free electricity for cooling, and water-efficient or closed-loop systems.*

Realizing the full potential of sustainable data center cooling will require a collective shift—from isolated innovation to coordinated implementation. Each stakeholder group has a distinct and essential role to play, from setting performance standards and funding next-generation technologies to deploying climate-responsive designs and sharing operational data. By aligning incentives, accelerating open collaboration and embedding sustainability into the planning and operation of digital infrastructure, we can ensure that data centers not only meet the demands of the digital age but also contribute to a more resource-efficient and climate-resilient future.

E. References

1. Refrigeration American Society for Heating, and Air Conditioning Engineers (ASHRAE). Data Center Resource Page; ASHRAE, Peachtree Corners, Georgia, https://www.ashrae.org/technical-resources/bookstore/datacom-series (Accessed August 2025).

2. Vertiv. Understanding direct-to-chip cooling in HPC infrastructure: A deep dive into liquid cooling; Vertiv, Thane, Maharashtra, India, https://www.vertiv.com/en-in/about/news-and-insights/articles/educational-articles/understanding-direct-to-chip-cooling-in-hpc-infrastructure-a-deep-dive-into-liquid-cooling/ (Accessed August 2025).

3. Arman Shehabi, Alex Newkirk, Sarah J Smith, Alex Hubbard, Nuoa Lei, Md Abu Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet & Dale Sartor. 2024 United States Data Center Energy Usage Report. Report No. LBNL-2001637, (2024), https://doi.org/10.71468/P1WC7Q.

4. Vertiv. Figure 1 in "Understanding direct-to-chip cooling in HPC infrastructure: A deep dive into liquid cooling"; Vertiv, Thane, Maharashtra, India, https://www.vertiv.com/49ffc2/globalassets/images/on-page-image/800x450/800x450-liquid-cooling-chart_383570_0.jpg (Accessed August 2025).

5. Nuoa Lei, Jun Lu, Arman Shehabi & Eric Masanet. The water use of data center workloads: A review and assessment of key determinants. Resources, Conservation and Recycling 219, 108310 (2025). https://doi.org/10.1016/j.resconrec.2025.108310.

6. Eric Masanet, Arman Shehabi, Nuoa Lei, Sarah Smith & Jonathan Koomey. Recalibrating global data center energy-use estimates. Science 367, 984-986 (2020). https://datacenters.lbl.gov/sites/default/files/Masanet_et_al_Science_2020.full_.pdf.

7. Ali Badiei, Eric Jadowski, Saba Sadati, Arash Beizaee, Jing Li, Leila Khajenoori, Hamid Reza Nasriani, Guiqiang Li & Xin Xiao. The Energy-Saving Potential of Air-Side Economisers in Modular Data Centres: Analysis of Opportunities and Risks in Different Climates. Sustainability 15, 10777 (2023). https://doi.org/10.3390/su151410777.

8. Tamzid Ahmed. World's 5 Bigest Underwater Data Center Projects; Brightlio, Culver City, California, https://brightlio.com/underwater-data-centers/ (2025).

9. Microsoft. Project Natick (Homepage); Redmond, Washington, https://natick.research.microsoft.com/ (Accessed August 2025).

10. Subsea Cloud. Subsea Cloud Homepage; Houston, Texas, https://www.subseacloud.com (Accessed August 2025).

11. You Xiaoying. China Is Putting Data Centers in the Ocean to Keep Them Cool; Scientific American, New York, New York, https://www.scientificamerican.com/article/china-powers-ai-boom-with-undersea-data-centers/ (2025).

12. Hewlett Packard Enterprise. HPE Spaceborne Computer-2 returns to the International Space Station (Press Release); Hewlett Packard Enterprise Spring, Texas, https://www.hpe.com/us/en/newsroom/press-release/2024/01/hpe-spaceborne-computer-2-returns-to-the-international-space-station.html (2024).

13. Douglas Donnellan, Daniel Bizo, Jacqueline Davis, Andy Lawrence, Dr. Owen Rogers, Lenny Simon & Max Smolaks. Uptime Institute Global Data Center Survey 2023: Executive Summary; Uptime Intelligence, New York, New York, https://uptimeinstitute.com/resources/asset/executive-summary-uptime-institute-global-data-center-survey-2023 (2023).

14. Miguel Yañez-Barnuevo. Data Centers and Water Consumption; Environmental and Energy Study Institute (EESI), Washingon, DC, https://www.eesi.org/articles/view/data-centers-and-water-consumption (2025).

15. Andrew Higgins. What Is Water Usage Effectiveness (WUE) in Data Centers?; Equinix, Readwood City, California, https://blog.equinix.com/blog/2024/11/13/what-is-water-usage-effectiveness-wue-in-data-centers/ (2024).

16. United Nations Environment Programme (UNEP) and International Energy Agency (IEA). Cooling Emissions and Policy Synthesis Report; UNEP, Nairobi and IEA, Paris, https://www.unep.org/resources/report/cooling-emissions-and-policy-synthesis-report (2020).

17. European Commission. F-gas legislation; Bruxelles, Belgium, https://climate.ec.europa.eu/eu-action/fluorinated-greenhouse-gases/f-gas-legislation_en (Accessed August 2025).

18. Husam Alissa, Teresa Nick, Ashish Raniwala, Alberto Arribas Herranz, Kali Frost, Ioannis Manousakis, Kari Lio, Brijesh Warrier, Vaidehi Oruganti, T. J. DiCaprio, Kathryn Oseen-Senda, Bharath Ramakrishnan, Naval Gupta et al. Using life cycle assessment to drive innovation for sustainable cool clouds. Nature 641, 331-338 (2025). https://doi.org/10.1038/s41586-025-08832-3.

19. Uptime Institute. Uptime Institute's 2022 Global Data Center Survey Reveals Strong Industry Growth as Operators Brace for Expanding Sustainability Requirements (Press Release); New York, New York, https://uptimeinstitute.com/about-ui/press-releases/2022-global-data-center-survey-reveals-strong-industry-growth (2022).

20. Google. Google data center PUE performance: 2024 PUE Yearly Report; Mountain View, California, https://datacenters.google/efficiency/#2024 (2024).

21. Qian Wei, Jun Lu, Xiaoping Xia, Bin Zhang, Xiang Ying & Leihong Li. Performance and Applicability Analysis of Indirect Evaporative Cooling Units in Data Centers Across Various Humidity Regions. Buildings 14, 3623 (2024). https://doi.org/10.3390/buildings14113623.

22. Fred Rebarber. Quantifying the Impact on PUE and Energy Consumption When Introducing Liquid Cooling Into an Air-cooled Data Center (Blog Post); Vertiv Group Corp., Westerville, Ohio, https://www.vertiv.com/en-us/about/news-and-insights/articles/blog-posts/quantifying-data-center-pue-when-introducing-liquid-cooling/ (2023).

23. Rui Kong, Hainan Zhang, Mingsheng Tang, Huiming Zou, Changqing Tian & Tao Ding. Enhancing data center cooling efficiency and ability: A comprehensive review of direct liquid cooling technologies. Energy 308, 132846 (2024). https://doi.org/10.1016/j.energy.2024.132846.

24. Nuoa Lei & Eric Masanet. Climate- and technology-specific PUE and WUE estimations for U.S. data centers using a hybrid statistical and thermodynamics-based approach. Resources, Conservation and Recycling 182, 106323 (2022). https://doi.org/10.1016/j.resconrec.2022.106323.

25. Arman Shehabi, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Inês Azevedo Nathaniel Horner & William Lintner. United States Data Center Energy Usage Report (LBNL-1005775); Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, California, https://eta-publications.lbl.gov/sites/default/files/lbnl-1005775_v2.pdf.

26. Anthony Sbarra & Laurie Williams. Uptime Institute 2024 Cooling Systems Survey: Direct Liquid Cooling Results; Uptime Institute, New York, New York, https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2024.Cooling.Survey.Report.pdf (2024).

27. Kaifang Shi, Yun Chen, Bailang Yu, Tingbao Xu, Zuoqi Chen, Rui Liu, Linyi Li & Jianping Wu. Modeling spatiotemporal CO_2 (carbon dioxide) emission dynamics in China from DMSP-OLS nighttime stable light data using panel data analysis. Applied Energy 168, 523-533 (2016). https://doi.org/10.1016/j.apenergy.2015.11.055

28. Rich Miller. Microsoft's Chiller-less Data Center; Data Center Knowledge, San Francisco, California, https://www.datacenterknowledge.com/hyperscalers/microsoft-s-chiller-less-data-center (2009).

29. Flann Gao & Luica Mak. Alibaba Cloud launches new energy-efficient Qiandao Lake Data Center; Alibaba Cloud, Hangzhou, China, https://www.alibabacloud.com/en/press-room/alicloud-launches-new-energy-efficient-qiandao-lake-data-center?_p_lc=1 (2015).

30. Nimbus. Adiabatic Cooling Applications: Data Centers; Massillon, Ohio, https://nimbus.cool/resources/articles/adiabatic-cooling-applications-data-centers/ (Accessed August 2025).

31. Open Compute Project. Water Efficiency at Facebook's Prineville Data Center (Blog); Austin, Texas, https://www.opencompute.org/blog/water-efficiency-at-facebooks-prineville-data-center (2013).

32. Chi Zhou, Doris Gao, Lisa Rivalin, Andrew Grier, Gerson Arteaga Ramirez & John Fabian. in Deployable RL: From Research to Practice@ Reinforcement Learning Conference 2024; (Amherst, Massachusetts, 2024; https://openreview.net/pdf?id=3hZL9Vv0Ay).

33. Jerry Luo, Cosmin Paduraru, Octavian Voicu, Yuri Chervonyi, Scott Munns, Jerry Li, Crystal Qian, Praneet Dutta, Jared Quincy Davis, Ningjia Wu, Xingwei Yang, Chu-Ming Chang, Ted Li et al. Controlling Commercial Cooling Systems Using Reinforcement Learning. arXiv:2211.07357 (2022). https://doi.org/10.48550/arXiv.2211.07357.

34. Richard Evans & Jim Gao. DeepMind AI Reduces Google Data Centre Cooling Bill by 40%; Google DeepMind, London, United Kingdom, https://deepmind.google/discover/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40/ (2016).

35. Wayne Williams. Megawatt-class AI server racks may well become the norm before 2030 as Nvidia displays 600kW Kyber rack design; TechRadar, London, United Kingdom, https://www.techradar.com/pro/megawatt-class-ai-server-racks-may-well-become-the-norm-before-2030-as-nvidia-displays-600kw-kyber-rack-design (2025).

36. Christopher Helman & Rashi Shrivastava. Meet The Tiny Startup Building Stargate, OpenAI's $500 Billion Data Center Moonshot; Forbes, Jersey City, New Jersey, https://www.forbes.com/sites/christopherhelman/2025/04/10/meet-the-tiny-startup-building-stargate-openais-500-billion-data-center-moonshot/ (2025).

37. Ali Heydari, Ahmad R. Gharaibeh, Mohammad Tradat, Qusai Soud, Yaman

Manaserh, Vahideh Radmard, Bahareh Eslami, Jeremy Rodriguez & Bahgat Sammakia. Experimental evaluation of direct-to-chip cold plate liquid cooling for high-heat-density data centers. Applied Thermal Engineering 239, 122122 (2024). https://doi.org/10.1016/j.applthermaleng.2023.122122.

38. Nicola Pieretti (Vertiv). Chill factor: Top liquid cooling considerations for high-density environments; Data Center Dynamics, London, United Kingdom, https://www.datacenterdynamics.com/en/opinions/chill-factor-top-liquid-cooling-considerations-for-high-density-environments/ (2025).

39. Vertiv. Liquid cooling options for data centers; Westerville, Ohio, https://www.vertiv.com/en-us/solutions/learn-about/liquid-cooling-options-for-data-centers/ (Accessed August 2025).

40. Mark E Steinke (Committee Chair). Mission Critical Facilities, Data Centers, Technology Spaces and Electronic Equipment: ASHRAE Technical Committee 9.9; American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE), Peachtree Corners, Georgia, https://tpc.ashrae.org/?cmtKey=fd4a4ee6-96a3-4f61-8b85-43418dfa988d (Accessed August 2025).

41. Zutacore. https://zutacore.com/ (Homepage); San Jose, California, https://zutacore.com/ (Accessed August 2025).

42. Steve Solomon. Sustainable by design: Next-generation datacenters consume zero water for cooling; Microsoft, Redmond, Washington, https://www.microsoft.com/en-us/microsoft-cloud/blog/2024/12/09/sustainable-by-design-next-generation-datacenters-consume-zero-water-for-cooling/ (2024).

43. Alex Davies. Next-gen AI demands smarter cooling tech. Here's how AWS delivered it in just 11 months; Amazon, Seattle, Washington, https://www.aboutamazon.com/news/aws/aws-liquid-cooling-data-centers (2025).

44. Mike Wheatley. AWS integrates liquid cooling and simplifies electrical distribution to lower data center power consumption; Silicon Angle Media, Inc., Palo Alto, California, https://siliconangle.com/2024/12/02/aws-integrates-liquid-cooling-simplifies-electrical-distribution-lower-data-center-power-consumption/ (2024).

45. Georgia Butler. AWS global data centers achieved PUE of 1.15 in 2023; Data Center Dynamics, London, United Kingdom, https://www.datacenterdynamics.com/en/news/aws-global-data-centers-achieved-pue-of-115-in-2023/ (2024).

46. Drew Robb. Data Center Cooling Trends for 2025; Upsite Technologies, Albuquerque, New Mexico, https://www.upsite.com/blog/data-center-cooling-trends-for-2025/ (2025).

47. Green Revolution Cooling (GRC). Green Revolution Cooling: Redefining the Efficiency and Sustainability of Data Center Cooling (Homepage); GRC, Austin, Texas, https://www.grcooling.com (Accessed August 2025).

48. Submer.Datacenters that make sense (Homepage); Barcelona, Spain, https://submer.com/ (Accessed August 2025).

49. Pranati Sahoo. Immersion Cooling for Data Center Cooling-Innovation Speaks it All!; Energetica India, Pune, Maharashtra, India, https://www.energetica-india.net/articles/immersion-cooling-for-data-center-cooling—innovation-speaks-it-all (2022).

50. Sumit Kumar Singh, Dibakar Rakshit, K. Ravi Kumar & Anurag Agarwal. Recent advancements and sustainable solutions in adsorption-based cooling systems integrated with renewable energy sources and industrial waste heat: A review. Cleaner Engineering and Technology 23, 100827 (2024). https://doi.org/10.1016/j.clet.2024.100827.

51. Zhaopeng Cui, Shuai Du, Tianhao Zhao, Zhihui Chen & Ruzhu Wang. High-power-density adsorption chiller driven by data center waste heat using encapsulated composite as adsorbent. Energy 311, 133391 (2024). https://doi.org/10.1016/j.energy.2024.133391.

52. Maria Lind Arlaud (Editor in Chief). District energy: The backbone of a flexible, resilient and efficient energy system; State of Green, Copenhagen, Denmark, https://stateofgreen.com/en/publications/district-energy/ (2024).

53. Microsoft. Surplus datacenter heat will be repurposed to heat homes in Denmark (Blog); Redmond, Washington, https://local.microsoft.com/blog/datacenter_heat_repurposed/ (Accessed August 2025).

54. Terje Solsvik. Denmark's European Energy signs long-term power deals with Microsoft; Reuters, London, United Kingdom, https://www.reuters.com/business/energy/denmarks-european-energy-signs-long-term-power-deals-with-microsoft-2024-06-12/ (2024).

55. Tao Ding, Xiaoxuan Chen, Hanwen Cao, Zhiguang He, Jianmin Wang & Zhen Li. Principles of loop thermosyphon and its application in data center cooling systems: A review. Renewable and Sustainable Energy Reviews 150, 111389 (2021). https://doi.org/10.1016/j.rser.2021.111389.

56. Hyo-Lim Park, Soo-Jin Lee, Jae-Hee Lee & Jae-Weon Jeong. A liquid desiccant-assisted free-cooling system for energy-efficient data centers in hot and humid climates. Case Studies in Thermal Engineering 52, 103683 (2023). https://doi.org/10.1016/j.csite.2023.103683.

57. Exowatt. Powering AI with 24-hour Solar (Homepage); Miami, Florida, https://www.exowatt.com/ (Accessed August 2025).

58. National Renewable Energy Laboratory (NREL). High-Performance Computing Data Center; NREL, Applewood, Colorado, https://www.nrel.gov/computational-science/hpc-data-center (Accessed August 2025).

59. Phase Change Solutions (PCS). Data Centers & Telecom; Asheboro, North Carolina, https://phasechange.com/data-center-telecom/ (Accessed August 2025).

60. Xiaolei Yuan, Xuetao Zhou, Yiqun Pan, Risto Kosonen, Hao Cai, Yang Gao & Yu Wang. Phase change cooling in data centers: A review. Energy and Buildings 236, 110764 (2021). https://doi.org/10.1016/j.enbuild.2021.110764.

61. Vertiv. Vertiv Homepage; Thane, Maharashtra, India, https://www.vertiv.com/en-us/ (Accessed August 2025).

62. Vertiv. Evolving Refrigerant Regulations and Applications in Data Center Environments (Vetiv White Paper); Vertiv, Westerville, Ohio, https://www.vertiv.com/48f7e3/contentassets/8eee3437fd734143a36d0395e65e8de7/vertiv-low-gwp-refrigerants-wp-en-na-sl-71228-web.pdf (2025).

63. Jinghan Sun, Zibo Gong, Anup Agarwal, Shadi Noghabi, Ranveer Chandra, Marc Snir & Jian Huang. “Exploring the Efficiency of Renewable Energy-based Modular Data Centers at Scale” in Proceedings of the 2024 ACM Symposium on Cloud Computing, Redmond, Washington. 552-569, https://doi.org/10.1145/3698038.3698544,(2024).

64. Open Compute Project. Scaling Innovation Through Collaboration! (Open Compute Project Homepage); Austin, Texas, https://www.opencompute.org/ (Accessed August 2025).

65. AndyD750 (Azure Infrastructure Blog). Liquid Cooling in Air Cooled Data Centers on Microsoft Azure; Microsoft Community Hub, Redmond, Washington, https://techcommunity.microsoft.com/blog/azureinfrastructureblog/liquid-cooling-in-air-cooled-data-centers-on-microsoft-azure/4268822 (2024).

66. ENERGY STAR. Data Centers (under "Energy Efficient Products"); US Environmental Protection Agency (EPA), Washington, DC, https://www.energystar.gov/products/data_centers (Accessed August 2025).

67. The Joint Research Centre: EU Science Hub. European Code of Conduct for Energy Efficiency in Data Centres; European Commission, Bruxelles, Belgium, https://joint-research-centre.ec.europa.eu/jrc-news-and-updates/eu-code-conduct-data-centres-towards-more-innovative-sustainable-and-secure-data-centre-facilities-2023-09-05_en (Accessed August 2025).

68. Jyotika Athavale, Cullen Bash, Wesley Brewer, Matthias Maiterth, Dejan Milojicic, Harry Petty & Soumyendu Sarkar. Digital Twins for Data Centers. 57, 151-158 (2024). https://doi.ieeecomputersociety.org/10.1109/MC.2024.3436945.).