Technology Review for Carbon Neutrality

Perspective

Accelerating liquid cooling adoption in data centers: necessity, challenges and solutions

David Zhou[1,*], Liguang Du[1], Ran Xu[1], Jiqiang Song[2], Na Chen[1], Shanshan Zheng[2], Yuehong Fan[1]

ABSTRACT

Optimizing cooling strategies of data centers is crucial to the carbon footprint of the world upon the explosive growth of artificial intelligence (AI). This article analyzes the necessity of transitioning from air cooling to liquid cooling in data centers and its challenges, and proposes technological and industry solutions. Liquid cooling technologies are hindered by challenges such as costs, leakage risks, and lack of standards and skilled workforce. The solutions lie in continued improvement in design methodology and technology performance, use of AI-assisted management tools, and advancing standardization and industry ecosystems.

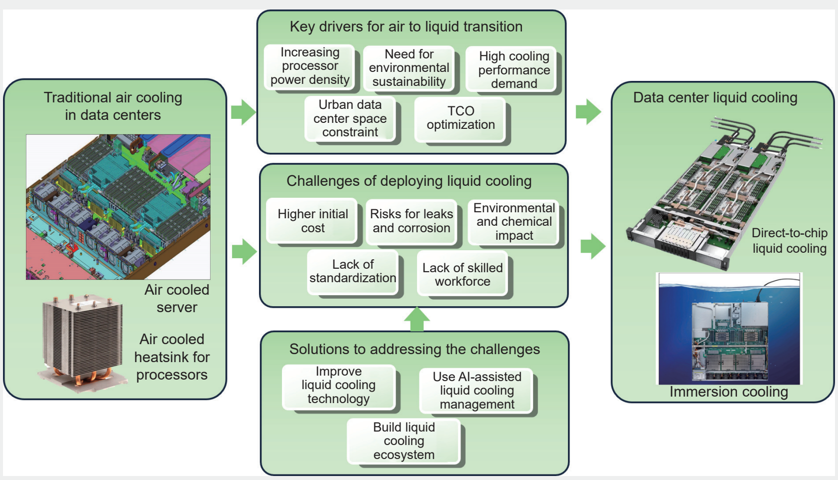

KEYWORDS

Data center liquid cooling; decarbonization; artificial intelligence (AI); green computing

Author affiliations

1 DCAI China, Intel Corporation, Shanghai 200003, China

2 Intel Labs China, Intel Corporation, Beijing 100013, China

*** Correspondence**

David Zhou, david.zhou@intel.com

Received: 12 September 2025

Revised: 22 October 2025

Accepted: 01 November 2025

https://doi.org/10.26599/TRCN.2025.9550014

1. Necessity of transitioning from air cooling to liquid cooling in data centers

The increasing demand for artificial intelligence (AI), big data analytics and cloud services drives the need for more data centers, and consequently, higher energy consumption. The report[1] finds that data centers consumed about 4.4% of total U.S. electricity in 2023 and are expected to consume approximately 6.7% to 12% of total U.S. electricity by 2028. The data center equipment, such as data center servers with graphics processing units (GPU) and AI accelerators, and other IT infrastructures, require plenty of power, which generates a lot of heat. And cooling systems, such as the most often used air conditioning systems for air-cooled data centers, consume a significant portion of the total energy in data centers, with around 40%[2]. This makes cooling one of the most energy intensive aspects of data center operations, and optimizing cooling strategies of data centers is crucial for minimizing facility cooling overhead to maximize computing power budgets and drive a lower power usage effectiveness (PUE) to reduce carbon footprint.

Data center cooling solutions generally fall into two main types: air cooling and liquid cooling, with variations within each. Air cooling methods include traditional air conditioning and free cooling (using naturally cold air or water instead of mechanical refrigeration in a certain season, for example in winter), while liquid cooling can range from liquid-to-air systems to direct-to-chip (i.e., cold plate liquid cooling) and immersion cooling. Hybrid cooling systems combine elements of both air and liquid cooling.

Key drivers are pushing data centers towards liquid cooling (Fig. 1).

Increasing power density and heat generation
Modern IT equipment, especially for AI and high-performance computing (HPC) workloads, generates significantly more heat in compact spaces compared to traditional servers. Liquid cooling excels at removing this concentrated heat efficiently due to the superior heat capacity of liquids compared to air.

Needs for environmental sustainability and energy efficiency
Data centers are under pressure to reduce their energy consumption and carbon footprint. Traditional air-cooling systems consume substantial amounts of electricity, while liquid cooling solutions can significantly reduce power consumption on cooling systems due to its much larger heat capacity and thermal conductivity compared to air cooling.

Demand for higher performance
Efficient cooling with liquid minimizes the risk of thermal throttling and equipment failure, leading to more stable hardware operation and reduced downtime. This increased performance is essential for data centers supporting high-performance computing, AI, and mission-critical applications.

Space optimization and scalability
Liquid cooling systems are more compact than air cooling infrastructure, allowing data centers to maximize rack densities and floor space utilization, potentially achieving cost neutrality over the facility's life cycle[3]. This is crucial for optimizing space in crowded urban areas or edge data centers with limited footprints. Liquid cooling systems can also be more easily scaled to accommodate future growth and changing workload demands.

TCO savings
While the initial investment (i.e., capital expenditure) for liquid cooling could be higher than air cooling, liquid cooling can lead to lower long-term operating expenses (OpEx) compared to air cooling which needs higher operating cost due to high energy consumption of fans and air conditioning units[4]. In the long term, for example, after about 4 years of operation with liquid cooling solution, the TCO with liquid cooling in a data center can decrease by up to 20% compared to air cooling solution[5].

In summary, liquid cooling is an effective cooling solution for the decarbonization in data centers, and it is also a necessary and a must-have solution to address the thermal challenge of high-power GPUs and AI accelerators (for example, when the thermal design power of the GPUs and AI accelerators is greater than 1,200 watts, as the launched Nvidia GB300 products and the beyond) in a data center AI server.

2. Challenges of deploying liquid cooling in data centers

While liquid cooling offers significant long-term energy savings and superior ability to cool higher density data centers, several challenges are hindering its widespread adoption in data centers. These challenges primarily include higher initial cost, risks for leaks and corrosion, environmental and chemical considerations, lack of standardization, and the need for skilled workforce for system operation and maintenance.

Higher initial cost
Liquid cooling systems, especially those involving immersion cooling, often require higher upfront capital investment (than

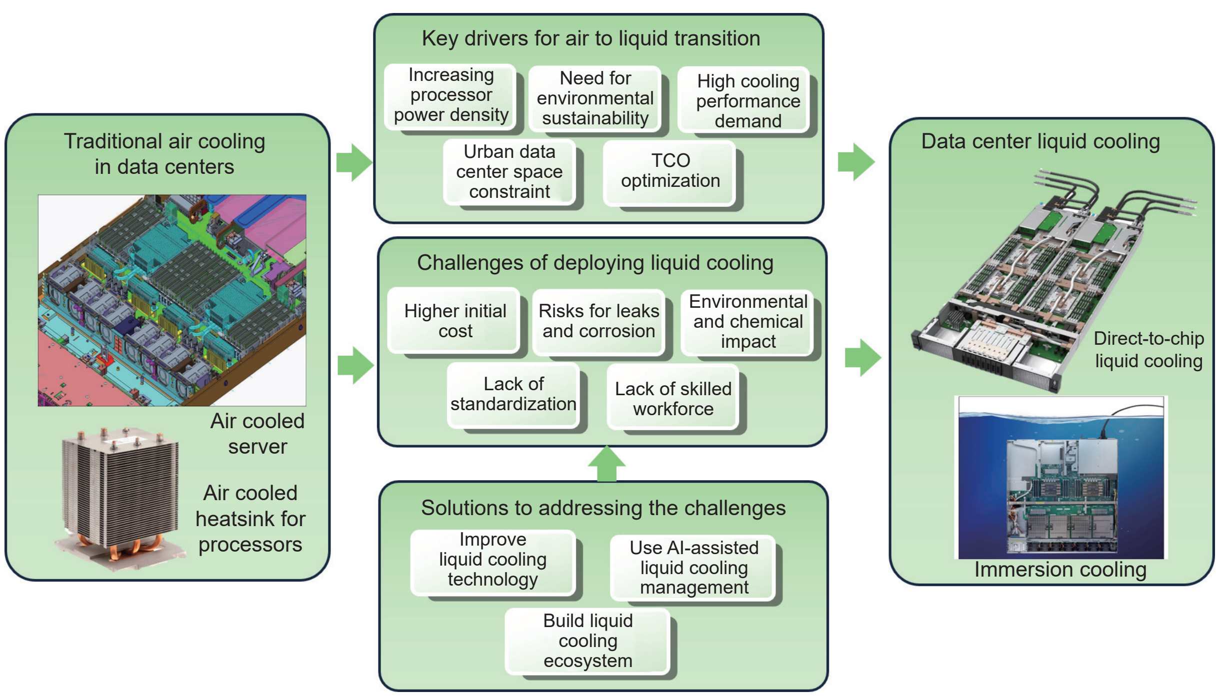

Fig. 1. Data center cooling technology transition.

traditional air-cooling systems) in specialized equipment such as highly customized cold plate for server components, immersion coolant, coolant distribution units (CDUs), plumbing and potentially even modifications to the data center building's structure to handle the weight of the cooling systems. In some cases, data centers may need to integrate liquid cooling with existing air-cooling systems, which can introduce further complexities of the retrofitting and potential inefficiencies if not done properly. For a typical deployment with 80% of the heat in a server rack dissipated by liquid, the initial cost of a liquid cooling solution is about 1.9–2.2 times that with air cooling solution in a data center[6].

Risks for leaks and corrosion

Leaks of liquid coolant in cold plate liquid cooling systems can lead to short circuits, corrosion of server board components and physical damage to IT equipment, resulting in downtime and financial losses. It is crucial to have leak detection and mitigation solutions in place for cold plate liquid cooled servers using water-based coolants (for example, the most often used treated water or propylene glycol liquid fluid)[7].

Liquid coolants need to provide reliable corrosion resistance for metal components and be able to prevent contaminants leached from polymer and elastomeric components, which can cause degradation of the mechanical properties of the wetted materials. It is worth noting that some polymers and elastomeric materials may absorb corrosion inhibitors in the liquid coolant, thereby reducing the concentration of corrosion inhibitors in the liquid coolant used to prevent corrosion, which accelerates the corrosion rate of the wetted materials as not desired.

Liquid's environmental and chemical considerations

When considering the environmental impact of liquid cooling systems in data centers, the selection of liquid coolants is particularly critical. There is growing emphasis on the environmental friendliness of chosen liquids[8]. Key indicators for assessing a liquid's environmental friendliness include its global warming potential (GWP) and ozone depletion potential (ODP). GWP serves as a relative measure of a greenhouse gas's heat-trapping capacity, while ODP indicates the extent of ozone layer depletion. In recent years, there has been a growing emphasis on deploying data center liquid cooling solutions that utilize fluids with low GWP and ODP values. For traditional cooling systems already employing coolants with high GWP and ODP values, it is essential to consider reducing or eliminating the risk of liquid coolant leakage and to establish reliable response protocols for managing leaks once liquid leakage occurs.

There are some chemical requirements for immersion cooling liquids ensuring the chosen liquid is chemically compatible and suitable for use in a data center immersion cooling environment. These chemistry requirements include immersion liquid's color, durability, oxidation potential, odor requirements, chemical purity, safety requirements and flammability[9]. Regarding durability, cooling liquid should meet the liquid property specifications and requirements under operating conditions for 10 years.

Lack of standardization

The lack of standardization in liquid cooling solutions is one of the biggest problems that are hindering liquid cooling adoption in data centers. Most liquid cooling solution vendors are developing diverse solutions with a small volume, which are hard to leverage and scale out across industry. The standardization on liquid cooling parts' interfaces and connections design, test methodology, quality and reliability requirements can get data center ecosystem partners' capabilities well aligned with related standards, contributing to the interchangeability of liquid cooling parts and accordingly initial cost reduction.

Need for skilled workforce for liquid cooling operation and maintenance

Compared to traditional air cooling systems, data center liquid cooling systems are more complex. Their installation, operation, maintenance, and troubleshooting generally require specialized knowledge and skills from operators. Even a small mistake by an unskilled operator can have serious consequences, such as equipment failure, data loss, or even complete system failure.

Data center staff may need to be well trained for new procedures, including leak detection and intervention, liquid coolant quality monitoring, emergency procedures and system maintenance. This can lead to reliance on specialized vendors for support, potentially increasing costs and introducing vendor lock-in.

3. Recommendations for accelerating the adoption of liquid cooling in data centers

To accelerate the adoption of liquid cooling for decarbonization in data centers, several measures can be taken, including improving liquid cooling technology performance, using AI-assisted tools, and building industry ecosystem.

3.1 Advance new liquid cooling design methodologies and technologies

Initial cost reduction through new modular design and assembly methodology

Offering modular designed liquid cooling components for servers, rack and cluster levels and applying a "Lego-style" assembly system can effectively lower the deployment time for data center customers and significantly reduce the initial cost of liquid cooling solutions. This can be done with an integrated liquid cooling design methodology.

Specifically, from the outset, each of the components in a server system — central processing unit (CPU), GPU or AI accelerators, voltage regulator (VR), solid state disk (SSD) or hard drive disk (HDD), peripheral component interconnect express (PCIe) card, open compute project accelerator module (OAM), memories, retime, switch processor, quad small form-factor pluggable (QSFP) and power supply unit (PSU), etc. – can be designed, developed, and optimized for liquid cooling solution in terms of their form factors, power allocation, server and rack system layout, ease of system assembly and disassembly, and ease of maintenance of the liquid cooled systems. Each supplier of the components can provide and ship their liquid cooled components with thermal solutions attached independently. With that, the data center rack and cluster integrators can directly assemble the computing facilities with all the liquid cooled components from different suppliers.

This methodology is not yet ready across the data center ecosystem, but it could be driven by the leading data center customers and leading data center original equipment manufacturers (OEMs) to get this methodology in place by defining the modular design specifications for the liquid cooled components.

Cooling performance improvement

Utilizing additive manufacturing (3D printing) technology to optimize the topology design of the internal flow channels of the single-phase cold plate can significantly improve heat dissipation efficiency and support higher-power CPUs, GPUs, and AI accelerators.

In addition, a study[10] indicated that compared to traditional straight microchannel cold plate designs, cold plates manufactured by the electrochemical additive manufacturing (EAM) exhibited 21% lower thermal resistance and 201% improvement in thermal uniformity compared to the baseline by significantly reducing the temperature gradient and improving temperature uniformity across the chip. This study demonstrates the potential of EAM for fabricating cold plates with complex geometries, marching a significant step forward in thermal management solutions for next-generation high-powered processors.

Direct-to-chip two-phase cold plate cooling is another promising solution with high cooling capacity and efficiency[11]. The two-phase cold plate cooling solution has significant advantages over single-phase cold plate liquid cooling systems, primarily due to the phase change of the coolant fluid (liquid to vapor and back) that enhances heat transfer efficiency. Regarding heat transfer and heat dissipation mechanism, two-phase cold plate cooling has the best possible cooling capability for data center servers, which leads to better energy efficiency (with lower pumping power and potential for heat reuse etc.), flexibility, and lower maintenance needs (fewer corrosion issues and reduced risk of leaks), making it ideal for high-performance computing and data centers. However, there are inherent challenges that must be addressed when developing a two-phase cold plate cooling system[12], including the higher costs associated with the refrigerant and equipment required for high working pressures in a two-phase cold plate cooling system, greater health and safety risks as well as legal-liability risks, which outweigh the perceived cooling performance benefits for industrywide deployments.

Moreover, developing a cost-effective loop heat pipe (with an air-cooled or liquid-cooled condenser) can significantly improve cooling performance of a high-power GPU, CPU and AI accelerators[13].

Leak-proof

A CDU with negative pressure technology can be used in liquid cooling systems to prevent leaks and ensure system reliability[14]. By operating under a vacuum, any potential leak will draw air into the system rather than allowing liquid coolant to leak out, preventing damage to sensitive equipment and minimizing downtime.

3.2 Utilize AI-assisted liquid cooling management

Liquid cooling solution providers must be able to provide not only single liquid-cooled servers or cold plates but also integrated turnkey solutions encompassing the data center infrastructure (such racks, CDUs, piping, etc.), liquid-cooled servers and clusters, and intelligent monitoring and management systems. Based on this intelligent management system, these providers should leverage AI algorithms for dynamic flow adjustment and temperature control to achieve optimal energy efficiency (for example, PUE < 1.1), significantly simplify operations and maintenance, and deliver high density, high reliability, and rapid deployment capabilities for full-stack liquid cooling solutions.

Energy efficiency optimization based on dynamic data analysis

Based on data such as workload predictions, ambient temperature, and real-time power consumption, the AI-driven liquid cooling management system can dynamically adjust pump speed, valve opening, liquid coolant temperature and flow rate, etc. to minimize energy consumption. For instance, reinforcement learning was found to improve chiller operation in Google's AI data centers by minimizing energy consumption based on dynamic data, resulting in energy savings of approximately 9% and 13% respectively at the two live experiment sites[15].

Intelligent operations and fault prediction based on sensor networks

Integrate sensors and software can monitor key parameters of liquid coolant in real time such as flow rate, temperature, pressure, and leaks. Through sensor networks and AI analysis, the AI-driven liquid cooling management system can provide early warning of potential failures such as micro-leaks, pump performance degradation, and pipe blockages, thus improving system reliability and safety while reducing the risk of downtime.

Intelligent prediction based on digital twin technology

Combining digital twin models with AI algorithms can achieve more intelligent decision-making and optimization. Introduce a liquid cooling system operation and maintenance platform based on digital twins[16], which uses a physical engine to simulate massive amounts of sensor data in real time, predicting equipment failures 24 hours in advance and significantly reducing the unplanned downtime rate of liquid cooling systems.

3.3 Build data center liquid cooling ecosystem

The widespread adoption of liquid cooling for data centers requires industry-wide collaboration. A collaborative industry ecosystem can be built that includes server original equipment manufacturers (OEM), chip makers (NVIDIA, Intel, AMD, etc.), liquid coolant suppliers, liquid cooling component suppliers, and data center operators aiming to jointly validate solutions, develop standards, and accelerate the adoption of liquid cooling technology for data centers. For example, promoting standardization of cold plate interfaces, universal quick disconnect (UQD), liquid cooling components' quality and reliability requirements, monitoring protocols, etc. Industry organizations like the Open Compute Project (OCP) and American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) are actively working to develop standards and guidelines for liquid cooling, paving the way for more reliable and consistent performance in the future.

It would also be beneficial for leading industry players to publicly release data center liquid cooling reference solutions for ecosystem partners.

A case in point of such industry ecosystem initiatives is the Data Center Liquid Cooling Innovation & Accelerator Program launched by Intel China in the late of 2023 (Fig. 2). It aims to accelerate the adoption of liquid cooling solutions in data centers by empowering leading liquid cooling independent hardware vendors (iHV) with Intel's liquid cooling technologies and qualification requirement.

Several iHVs' liquid cooling solutions matured for adoption in data centers, including cold plate liquid cooling total solutions, Intel patented G-flow technology[17] for single phase immersion cooling capability enhancement, and UQD interchangeability qualification In 2025, the Intel UQD Interchangeability Alliance was officially established with its first batch of qualified UQD suppliers, which unified the UQD design based on Intel's definition and thus improved the interoperability performance and reliability of the UQD's plugs and sockets. Such initiative significantly improved reliability lowered the cost, and secured the supply chain of UQD, which is the critical part in ensuring easy maintenance of liquid cooled server system in data centers.

4. Conclusion

Liquid cooling solutions in data centers are critical in addressing the challenges of increasing power consumption and carbon emissions from growing GPUs and AI accelerators. While the deployment of data center liquid cooling technologies is hindered by challenges such as costs, leakage risks, and lack of product standardization and skilled workforce, several efforts can be taken to reduce such challenges in terms of liquid cooling technology improvement, AI tools empowerment, and industry collaboration. These efforts will contribute to the wider

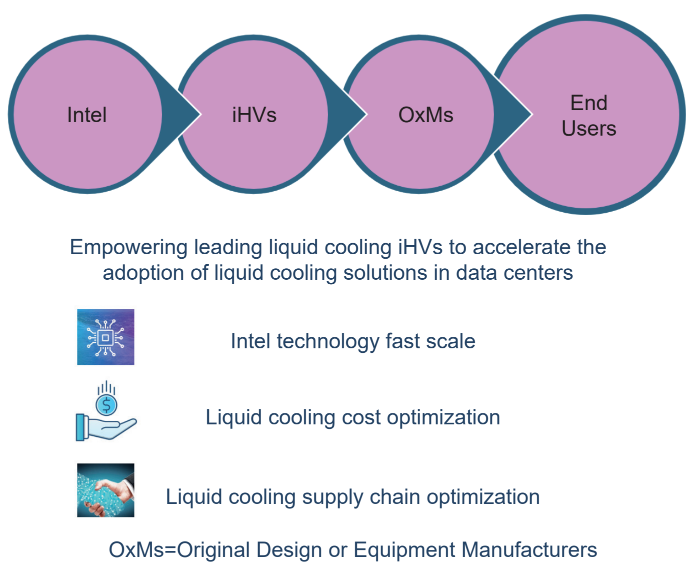

Fig. 2. Intel China Data Center Liquid Cooling Innovation & Accelerator Program as an example of industry ecosystem initiative

adoption of liquid cooling solutions in data centers and eventually higher efficiency and decarbonization of data centers in the coming decade.

Author contribution statement

David Zhou, Liguang Du, Ran Xu, Na Chen, and Yuehong Fan are responsible for Sections 1, 2, and 3. Jiqiang Song and Shanshan Zheng are responsible for Sections 3 and 4. All authors have approved the final manuscript.

Data availability

Not applicable.

Use of AI statement

None.

Declaration of competing interest

All authors are the employees of Intel Corporation.

REFERENCES

1. Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Nuoa, Lei N., Siddik, M. A. B., Holecek, B., Koomey, J., Masanet, E., Sartor, D. (2024). 2024 United States Data Center Energy Usage Report. LBNL-2001637. Lawrence Berkeley National Laboratory, Berkeley, California.
2. Zhang, Y., Shan, K., Li, X., Li, H., Wang, S. (2023). Research and Technologies for next-generation high-temperature data centers–State-of-the-arts and future perspectives. *Renewable and Sustainable Energy Reviews*, 171: 112991.
3. Hassel, William J. (2025) Liquid versus air: life cycle carbon of cooling down AI data centers. Master's Thesis, Harvard University Division of Continuing Education.
4. Zheng, J., Chen, X., Cheng, B., He, Y., Tian, W., Zhong, Y., Yu, J., Zhang, Z., He, Y., Li, Y., et al. (2024). Experimental investigation on single-phase immersion cooling solution for data center application. In Proceedings of the 2024 23rd IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Aurora, CO, USA.: 1–8.
5. Information on https://www.supermicro.org.cn/pressreleases.
6. Schneider Electric report. (2023) Capital cost analysis of liquid-cooled vs. air cooled large data centers.
7. Yaroch, K., Cuthbert, J, Miller, D., Sivapalan, S. T., Barlett, S. (2022). Guidelines for using propylene glycol-based heat transfer fluids in single-phase cold plate-based liquid cooled racks. White Paper on Open Compute Project (OCP).
8. Gullbrand, J., Gore, N., Matteson, J., Langer, E. (2019). ACS Liquid cooling cold plate requirements. White Paper on Open Compute Project.
9. Ahuja, S., Chang, J. Y., Wang, K., Klemes, M., Gullbrand, J., Ahuja, N., Yates, S., Munukutla, A., Mitra, A., Hemmeyer, M. (2022). Base specification for immersion fluids. Open Compute Project,
10. Gupta, G., Castro, D., Madril, J., Ouradnik, T., Winfield, I., Matthews, M., Agonafer, D. (2025). Numerical and experimental investigation of high-powered chips for efficient cooling using optimized electro-chemical additive manufacturing based cold plates. In Proceedings of the 24th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems.
11. Wang, Q., Ozguc, S., Narayanan, A., Bonner, R. W. (2024). Server-level test system for direct-to-chip two-phase cooling of data centers using a low global warming potential fluid. In Proceedings of the 2024 23rd IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Aurora, CO, USA.: 1–6.
12. Najjari, M. R., Turner, C., Mostafavi, K. (2024). Comprehensive analysis of two-phase flow boiling for IT equipment cooling: A system and component-level approach. In Proceedings of the ASME 2024 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems. San Jose, California, USA.
13. Myeong, H. S., Jang, S. P. (2024). Optimization of LHP (loop heat pipe) Geometry for Ultra-high Heat Flux Cooling System. In Proceedings of the 2024 30th International Workshop on Thermal Investigations of ICs and Systems (THERMINIC), Toulouse, France.: 1–4.
14. Chilldyne. (2024) Featured Industry Report: The future of liquid cooling in AI computing and the value of negative pressure technology.
15. Luo, J., Paduraru, C., Voicu, O., Chervonyi, Y., Munns, S., Li, J., Qian, C., Dutta, P., Davis J. Q., Wu, N., et al. (2022). Controlling commercial cooling systems using reinforcement learning. https://doi.org/10.48550/arXiv.2211.07357
16. Brewer, W., Maiterth, M., Kumar, V., Wojda, R., Bouknight, S., Hines, J., Shin, W., Greenwood, S., Grant, D., Williams, W., et al. (2024). A digital twin framework for liquid-cooled supercomputers as demonstrated at exascale. In Proceedings of the SC24: International Conference for High Performance Computing, Networking, Storage and Analysis, Atlanta, GA, USA.: 1–18.
17. Fan Y. H., Wang, C. L., Du, L. G, Zhang, G. Y. et al. (2025) G-flow immersion cooling solution for high-power data center servers, IEEE iTherm2025, US.