

BRIEF 4

Innovative Data-Centre Cooling Technologies in China - Liquid Cooling Solution

KEY MESSAGES

- The increased need to dissipate heat caused by the increased power consumption of IT equipment in data centres calls for energy-efficient cooling solutions. Liquid cooling, with its efficient heat dissipation and high energy-saving characteristics, is becoming greatly preferred in China and is snowballing with successful business cases already on the market.
- Liquid cooling still faces many challenges in the development process. There is an urgent need to promote the development of technology and industry by strengthening industry guidance, standardizing the evaluation system, and improving the industrial ecosystem, among other measures.
- This brief showcases the Alibaba and Tencent data centres' advanced liquid-cooling systems. Of the two leading Chinese internet companies, Alibaba has achieved large-scale deployment of liquid-cooling technology.

DEVELOPMENT OF THE DATA-CENTRE INDUSTRY IN CHINA

In the information age, scientific data has become a strategic resource. Similarly, the scale of data centres has become an indicator of national competitiveness in digital science and technology and a driving force behind the sustainable and high-speed development of information science and technology. With the boom in cloud computing, artificial intelligence, and big data, the overall scale of China's data centres has shown a trend towards rapid growth. As of the end of 2019, the number of data-centre racks in use had reached 3.1 million, up 38% year-on-year. On the market side, according to the China Academy of Information and Communications Technology (CAICT), the size of the Chinese data-centre market has been expanding at an annual growth rate of about 30% since 2012, as shown in Figure 1.

However, the rapid growth in data centres has also created energy-consumption problems. The slower chip-manufacturing process increases power consumption. The TDP (thermal design power) of general CPUs (central processing units) has also increased gradually, from about 100W to 400W. The power of a GPU (graphics processing unit) used for AI (artificial intelligence) training can be as high as 2.6kW. In the future, AI training GPU power is expected to reach 10kW. When dealing with AI deep-learning calculations, GPU is more powerful than CPU. As computing applications become more complex,

GPU will become more popular. As a result, the internal cooling systems of data centres consume enormous amounts of electricity each year, significantly increasing their operating and maintenance costsⁱ. Therefore, how to enhance the heat-dissipation efficiency of data centres and improve their sustainability has become a focus for society generally. The Chinese data-centre industry is actively exploring new cooling solutions, one of the most effective being liquid cooling.

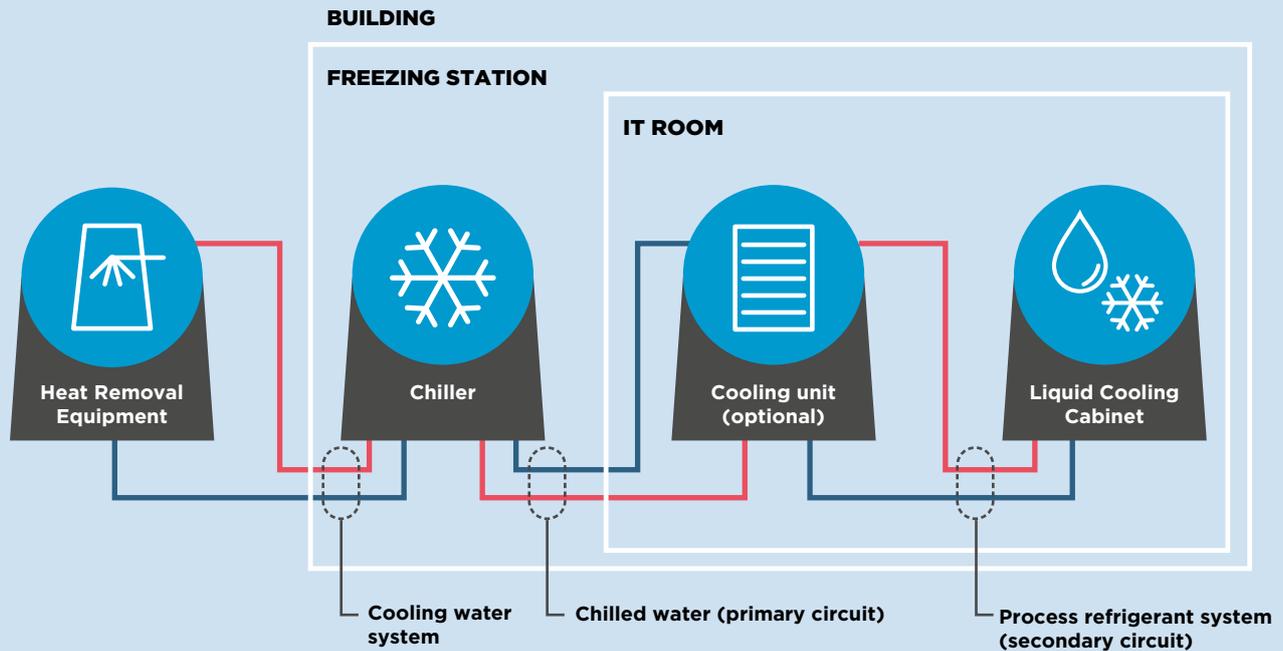
ADVANTAGES OF LIQUID COOLING

Liquid-cooling technology is not a new concept, as it has a long history of application in satellites, rocket engines, and other fields of aerospace industry. In the data-centre field, mainframe computers have been used since the 1960s. With the progress of time and technology, however, the densities of computer chips has greatly increased, computing work has become more and more complex, and the heat load of IT equipment has risen steadily. The ability to provide sufficient air volumes or cooling capacity for data centres is approaching the limit. Traditional air-conditioning systems using air-cooling technology have been unable to meet the cooling requirements of high-density data centres, so the data-centre industry is gradually resorting to liquid-cooling technology to address the issue. The advantages of liquid cooling are becoming more pronounced as it appears more and more frequently in data centresⁱⁱ.

Figure 1. 2012-2019 IDC revenue and growth in China



Figure 2. Liquid cooling general circuit



1. Large heat capacity and efficient heat dissipation

Liquid-cooling technology reduces the temperature of a data centre by using liquid as the medium for heat transfer. This relies on the liquid being in direct contact with the heat source to remove the heat, thus eliminating the need for indirect cooling through the air with an air-cooling solution. Liquid cooling takes most of the heat away through the circulating medium, thus reducing the demand for air from a single server and the overall demand for air circulation in the server room, dramatically reducing the local hot spots caused by the backflow in the server roomⁱⁱⁱ. In addition, the liquid is much more effective than the air in transferring heat, and it can absorb a large amount of heat while keeping the temperature stable, leading to much higher efficiencies in heat dissipation.

2. Reducing energy consumption and expenditure

Data centres consume enormous amounts of electricity in their operations, which is associated with expensive electricity bills. With the shift to liquid-cooling solutions, fans and air conditioners, which consume power continuously, can be completely or partially replaced, resulting in large reductions in electricity consumption and costs.

3. Improving energy efficiency, protecting the environment, and reducing noise

Liquid-cooling systems can reduce a data centre's overall energy consumption and PUE (power usage effec-

tiveness) to a remarkable extent. In addition, since the pumps and other instruments of the liquid-cooling system are quieter than the fans, the noise of liquid-cooling systems is much lower than that of the air-cooling system, enabling the creation of "silent data centres".

INTRODUCTION OF THREE LIQUID-COOLING TECHNOLOGIES

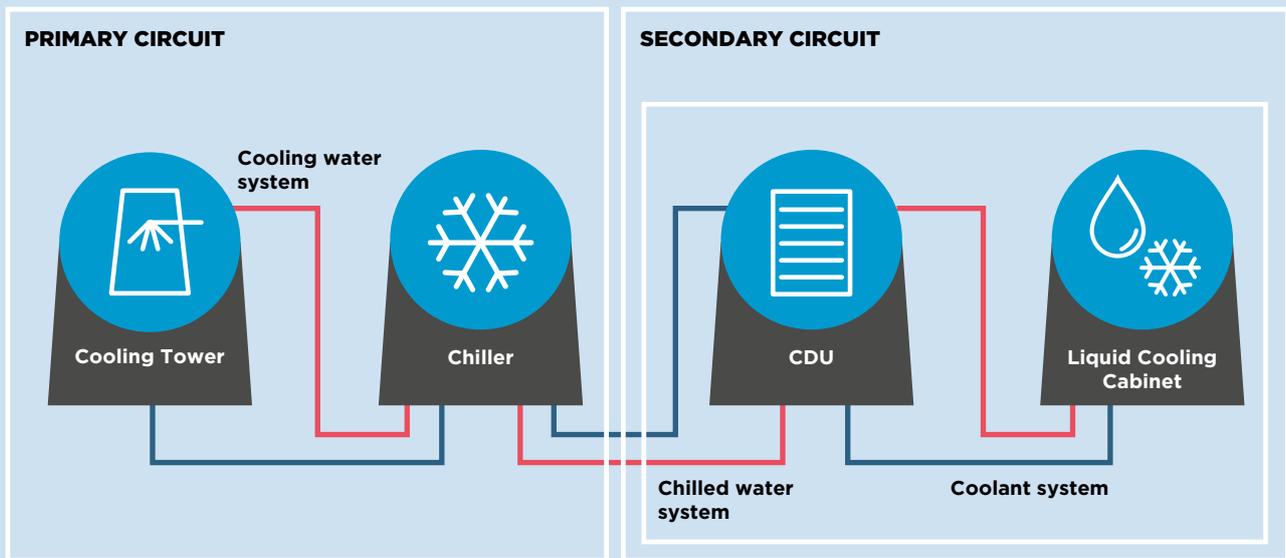
The cooling subsystem of a liquid-cooling system for data communication equipment can be thought of as a liquid circuit in which the cooling liquid exchanges heat with the components. In some cases, the water in the cooling system is provided by the internal CDU (cold distribution unit) of the rack or by an external CDU that serves multiple racks.

The potential liquid-cooling circuit in the data centre and the terminology used are shown in Figure 2. At present, liquid-cooling solutions mainly use one of three technical routes: cold-plate liquid cooling, immersion liquid cooling and spray liquid cooling.

1. Cold-plate liquid cooling

The main deployment method for cold-plate liquid cooling is to configure a water separator on the liquid-cooling cabinet to provide inlet and outlet branch pipes to the liquid-cooling calculation node. The inlet and outlet pipes of the branch pipelines are connected respectively

Figure 3. Cold-plate liquid-cooling system structure



with the calculation node through the joint and the inner cold-plate pipeline of the liquid-cooling calculation node to realize the liquid-cooling cycle in the liquid-cooling calculation node. The liquid from the liquid-cooled computing node converges at the cabinet level, which has one inlet and one outlet connected to the external piping. The connector is linked to an external or built-in CDU in order to circulate the liquid and take heat away from the computing node, as shown in Figure 3. In the cold-plate liquid-cooling system, a large power-consumption component such as the CPU is cooled by the liquid-cooling plate, and other heat-generating devices, such as the hard disk and interface card, are still cooled by the air-cooling system^{iv}.

This type of heat dissipation is higher in density, more energy-efficient, and less noise-generating than air cooling. Since cold-plate liquid-cooling technology does not require the same number of expensive chiller units, it significantly improves the energy efficiency of the data centre while reducing the overall costs^v. In general, air-cooling technology can only support power consumption per cabinet up to 30kW, while cold-plate liquid cooling can reach a total power consumption of 45kW per cabinet with a flow configuration of 60 litres per minute, enabling a higher density data centre deployment^{vi}.

2. Immersion liquid cooling

Immersion liquid cooling is a new type of heat-dissipation technology that has attracted the industry's attention in recent years. Especially at recent global supercomputing conferences, many Chinese and international server companies have demonstrated their immersion liquid-cooling products, dramatically increasing the industry's focus on liquid cooling. The immersion liquid-cooling system is a new, efficient, green and energy-saving data-centre cooling solution.

Immersion liquid cooling has distinctive advantages. Firstly, in an immersion liquid-cooling system, the coolant is in direct and full contact with the heat-generating equipment. As a result, the convection heat resistance is low, and the heat transfer coefficient is high. Secondly, the coolant has high thermal conductivity and a specific heat capacity, with low operating temperature variability. Thirdly, this approach eliminates the need for a fan, thus reducing energy consumption and noise and providing high cooling efficiency. Lastly, the coolant has excellent insulation properties and a high flash point, is non-flammable and non-toxic, and generally has no corrosion. Therefore, this liquid-cooling technology is suitable for large data centres, supercomputing, and industrial and scientific research institutions that have high demands for heat-flow density and green energy savings. It is especially appropriate for data centres located in cold and high-altitude areas with limited space, as well as data

centres that have high requirements for environmental noise and are close to people's offices and living quarters.

3. Spray liquid cooling

The main feature of spray liquid cooling is that the non-corrosive coolant passes directly through the shower plate on the server chassis and is sprayed onto the surface of the heat-generating devices or an extended surface in contact with it. The absorbed heat is then removed and exchanged with the external environment of a large cooling source, thus reaching the targeted temperature of the control system. The spray liquid-cooling system has the characteristics of high device integration, high heat-dissipation efficiency, high energy savings and silent operation. It is one of the most effective means of solving the problem of deploying high-power cabinets in data centres, reducing IT system-cooling costs, and improving energy efficiency^{vii}.

GOOD PRACTICES IN LIQUID-COOLING APPLICATIONS

Immersion liquid cooling—Alibaba Winter Olympics Cloud Data Centre

The Alibaba Group has long focused on energy efficiency improvements of its data centres. Alibaba has deployed a single-phase immersion liquid-cooling server cluster in its 2MW power capacity Winter Olympic Cloud Data

Centre, which can accommodate thousands of servers at scale. As a result, commercialization of this liquid-cooling system has started. On January 6, 2020, Alibaba open-sourced its liquid-cooling data-centre specifications at ODCC (Open Data Center Committee), hoping to develop open standards for large-scale cloud-oriented data-centre infrastructure with partners. Only in this way can ODCC cultivate China's cloud data-centre ecosystem and build a new generation of sustainable and energy-efficient data centres.

Alibaba's development of liquid cooling mainly uses a single-phase immersion liquid-cooling server solution that does not require operators to use traditional air-cooling equipment. Outdoor cooling equipment such as cooling towers or dry coolers can be used all year round, thus significantly reducing energy consumption. Alibaba's immersion liquid-cooled data centres can be deployed in any region of the world. Taking southern China as an example, compared with the PUE1.5 air-cooled data centre, the energy consumption of a liquid-cooled data centre of the same size can be reduced by more than 35%. In other words, when 100,000 servers are running, about 235 million kWh of electricity and 200,000 tons of carbon dioxide emissions can be saved each year.

In terms of unit space, the server no longer uses fans, greatly improving unit density. As the per-unit computing capacity is increased more than tenfold, it can support

Figure 4. Alibaba Winter Olympics Cloud Data Centre (source: Alibaba)



Figure 5. Tencent Liquid-Cooling Data Centre
(source: Tencent)



processors with higher power consumption in the future. In terms of physical space, the immersion liquid-cooling solution can save 75% of the surface occupied by IT equipment. At the same time, the sealed server is no longer affected by vibration, air humidity or airborne dust particles, which improves equipment reliability by 50%. Even if electronic components are submerged in liquid for twenty years, neither the composition of the liquid nor the quality of the electronic equipment changes, significantly prolonging the lifespan of the components and the equipment. Due to the obvious advantages of liquid cooling in terms of efficient heat dissipation, reductions in energy consumption and improved use of space, Alibaba continues to use immersion liquid-cooling in its Hangzhou Renhe Data Centre, which came into operation in 2020.

Cold-plate liquid cooling: Tencent liquid-cooling micro-module data centre

To address the future cooling needs of high-density computing equipment in data centres and to mitigate the impact of increasingly strained power resources on industry development, Tencent engineers have combined liquid cooling and micro-modules to form an innovative data-centre deployment. The application of this liquid-cooling technology can address the cooling needs of high-density computing equipment, further reduce data-centre PUE, and improve the efficiency of power utilization. At the same time, combined with the deployment

of micro-modules, the advantages of productization and modularization can be exploited. This solution can be deployed more flexibly and quickly in batches to cope better with the explosive growth of the data-centre business.

This system uses a cold-plate liquid-cooling technology to take CPU heat away through a cooling water cycle that supports an inlet temperature of 37°C on the primary side. This aspect of the heat dissipation accounts for about 75% of the server's total heat dissipation. Since it supports high-temperature primary water intake, this technology allows natural cooling in most climatic conditions, resulting in significant mechanical and electrical savings, which also reduces the data centre's electricity costs. At the same time, Tencent conducted in-depth customization of the liquid-cooled CDU using loop tubes and isolation valves to make the liquid cooling system more adaptable to micro-modules and the operating requirements in terms of granularity, redundant configuration, and control functions. Tencent has conducted technical research and test verification in cold-plate liquid cooling for many years, and will continue to study the integrated application of liquid cooling and micro-modules.

POLICY DEVELOPMENT AND SUSTAINABILITY ASSESSMENT METHOD

In order to reduce data centres' high levels of energy consumption, the Chinese government has launched policies to regulate such energy consumption^{viii}. Since 2017 the Ministry of Industry and Information Technology (MIIT) has issued a "Notice on Organizing the Application for a National Innovative Model of Industrial Demonstration Bases" every two years. New industries such as data centres, cloud computing and big data were included within the scope of the national model of the industry's demonstration bases. MIIT aimed to select advanced large and super data centres which were models of energy saving, environmental protection, safety, reliability, service capabilities and application levels. Similar high-standard small and medium data centres were also selected. Local governments have introduced policies responding to the call for energy savings and efficiency in order to promote the sustainable development of data centres.

- The "Prohibition and Restriction Catalogue for New Industries in Beijing (2018)", updated by the Beijing city government, prohibits the construction and expansion of data centres in the six districts of the city, though the construction of high-level data centres that meet specific requirements outside the six districts is allowed^{ix}.

- The Shanghai city government has issued its “13th Five Year Plan for Energy Conservation and Climate Change in Shanghai”, which provides that new data centres should be strictly controlled to ensure sustainability. Where construction is deemed necessary, it must ensure sustainable energy savings. On this basis, in 2019 the “Shanghai Internet Data Centre Construction Guidelines” were issued to promote energy efficiency and emissions reductions in various respects, and for the first time to encourage the deployment of liquid-cooling servers.
- The Shenzhen city government has issued a “Notice of an Energy Savings Review of Data Centres” to promote the energy-efficient retrofitting of old data centres and control the increase in the total annual energy consumption of data centres.

Assessments of data-centre sustainability ratings were jointly conducted by industry standard organizations like the Open Data Centre Committee (ODCC) and the Green Grid Committee (TGGC). This has been unanimously approved by large Chinese data-centre operators such as Alibaba, Tencent, China Telecom, China Mobile, China Unicom, GDS and Baidu.

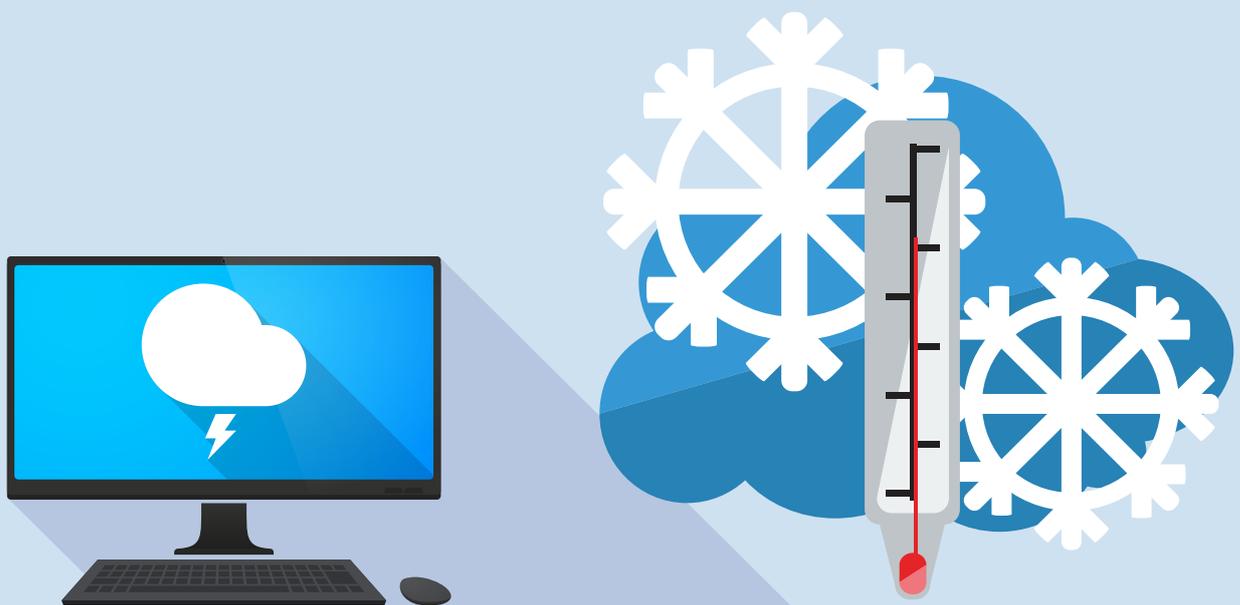
In the eight years since the assessment was introduced, many data centres have participated, but few have been awarded the highest 5A-level green rating. Until recent years, under the active exploration and innovation of Internet companies, 5A-level green data centres began to appear. For example, in 2015, the Baidu Cloud Computing (Yangquan) Centre adopted technologies such as whole

cabinet servers, high-voltage DC + direct power supply, machine-learning control systems, and high-temperature servers to achieve an average annual PUE of 1.23.

In 2016, Alibaba’s Qiandao Lake Data Centre adopted a natural cooling system with lake water, solar panels, high-voltage DC and other technologies to achieve an annual design PUE of 1.28. In 2017, the Tencent Qingpu Trigenation Data Centre adopted technologies such as natural-gas trigeneration, centrifugal frequency conversion chillers, and magnetic levitation chillers to achieve an annual PUE of 1.31. In 2018, the Alibaba Winter Olympics Cloud Data Centre adopted technologies such as non-elevated floor-diffused air supply and fully automated building automation systems to maximize the use of natural cooling sources, achieving an annual PUE of 1.23. In 2019, the first phase of the ByteDance Guantinghu Big Data Industrial Base and the second phase of Tencent’s Bright & China Mobile & GDS Data Centre were awarded the titles of “Operational AAAAA”.

CHALLENGES AND THE WAY FORWARD

The development of liquid-cooling systems still faces many challenges. The first challenge is to adapt the original server rooms. The design of the load-bearing and piping of traditional server rooms derives from the specifications of the traditional cabinets. In the case of liquid cooling, the deployment environment will be very different due to the shift in cooling methods. Deploying a liquid-cooling system in a traditional data centre can pose problems in terms of cost and difficulty of deploy-



ment. For contact liquid cooling such as immersion and spray, compatibility between the liquids and the IT components is an important aspect that needs to be considered. The prolonged immersion of IT components in liquids will have an impact on their functionality and performance, which is a problem we need to be aware and verify. In addition, the safety of the liquid to the human body and the environment is crucial⁸.

Therefore, it is necessary to monitor the electrical and chemical stability of the coolant, the state of the IT equipment, the state of the power equipment, the data centre's energy consumption, and the state of the environment where liquid-cooling systems are concerned. The purpose is to record the status parameters of each link in a timely manner, to ensure the normal operation of all the links in the system, and to initiate backups when an abnormality occurs. The general configured refrigerant-quality monitoring equipment is used to monitor the chemical composition of the additives in the refrigerant, as well as electrical conductivity, the pH value, the corrosion rate and turbidity.

However, the current lack of industry standards for liquid cooling will inhibit the development of the technology. At present, each enterprise's research and development of liquid cooling is in a fragmented state. There is an urgent need for relevant industry-standard organizations to manage the standardization of liquid cooling, establishing clear technical requirements for all aspects of liquids, power supply and distribution, control, security, and operations and maintenance to promote the positive development of liquid-cooling technology.

Undoubtedly, liquid-cooling technology is a huge innovation and constitutes a revolution over traditional data-centre deployments. The transformation of IT equipment and other hardware requires substantial manpower and material investment, and its operation and maintenance also require a period for the accumulation of experience. Therefore, the deployment and use of liquid cooling systems require national guidance and a boost from industry. Currently, relevant Chinese standards organizations and platforms such as the China Communications Standards Association (CCSA) and the Open Data Centre Committee (ODCC) have made great efforts in this area. In addition, it is also necessary to improve the industry value chain of liquid-cooled data centres, promote technical exchanges and cooperation at the national and global level, encourage market development and promote progress with global liquid-cooling technology.

Figure 7. Problems to be solved



CORRESPONDING AUTHORS

Jie Li, China Academy of Information and Communications Technology (CAICT), Cloud Computing & Big Data Research Institute, deputy director

Xiao Wang, Copenhagen Centre on Energy Efficiency, UNEP DTU Partnership

AUTHORS

Liang Guo, CAICT, Cloud Computing & Big Data Research Institute, deputy chief engineer

Lina Xie, CAICT, Cloud Computing & Big Data Research Institute, senior manager of Data Centre Department

Shanyuan Gao, Alibaba Group, Infrastructure Architecture Department, chief architect

Fangyi Mei, Shenzhen Tencent Computer Systems Co., Ltd, IDC Platform Department, senior IDC architect

Hongtao Zhu, Beijing Sankuai Technology Co., Ltd, Infrastructure Technology Unit, director of Server R&D Department

Zhiming Luo, Baidu.inc, SYS, senior system engineer

Thank you for the cooperation and kind support from Cristina Bueti, Reyna Ubeda, Paolo Gemma and the rest of the ITU team.

The Copenhagen Centre on Energy Efficiency functions as the global thematic Energy Efficiency Hub of Sustainable Energy for All (SEforALL), and accordingly works directly to support the SEforALL objective of doubling the global rate of improvement in energy efficiency by 2030.

The Copenhagen Centre fulfills its mission through:

- assisting policy change in countries and cities, with knowledge, insights and technical support
- accelerating action through innovation in project development and finance
- raising the profile of energy efficiency by communicating success stories and supporting outreach.

For more information, please visit

www.energyefficiencycentre.org or contact us at c2e2@dtu.dk.

Regarding our work in Sustainable Data Centres and Smart Energy, please contact Xiao Wang at xwang@dtu.dk

Visit Copenhagen Centre's Knowledge Management System at kms.energyefficiencycentre.org

The Copenhagen Centre on Energy Efficiency is institutionally part of UNEP DTU Partnership (UDP). UDP is a UN Environment Collaborating Centre and a leading international research and advisory institution on energy, climate and sustainable development.

REFERENCES

- i Jie Li, Liang Guo, Bin Lan, Lina Xie. Liquid cooling revolution. The People's Posts and Telecommunications Press. 2019.08:110-116.
- ii Lina Xie, Liang Guo. Discussion on liquid cooling technology and its development. Information Communication Technology and Policy. 2019.02:22-25.
- iii Mustafa A. Kadhim, Nikil Kapur, Jonathan L. Summers, et al. Rack Level Study of Hybrid Liquid/Air Cooled Servers: The Impact of Flow Distribution and Pumping Configuration on Central Processing Units Temperature. 2020.41(19-20):1683-1698.
- iv Andreu Carbó, Eduard Oró, Jaume Salom, et al. Experimental and numerical analysis for potential heat reuse in liquid cooled data centres. 2016.112:135-145.
- v Pin Chen, Souad Harmand, Safouene Ouenzerfi. Immersion cooling effect of dielectric liquid and self-rewetting fluid on smooth and porous surface. 2020.180:32-35.
- vi Ali C. Kheirabadi, Dominic Groulx. Experimental evaluation of a thermal contact liquid cooling system for server electronics. 2018.129:1010-1025.
- vii Schneider Electric Uk. Schneider Electric Releases White Paper on Capital Cost Comparison of Data Centre Liquid Cooling vs Traditional Air Cooling. 2020.
- viii Yue Wang. Data Centre White Paper. Technical Report. CAICT, China Academy of Information and Communications Technology.2018.09:10-15.
- ix Yue Wang, Baohong He, Li Jie. National policies accelerate the green development of data centres. Communications Management and Technology.2015. 06:29-31.
- x Hayk Shoukourian, Torsten Wilde, Herbert Huber, et al. Analysis of the efficiency characteristics of the first High-Temperature Direct Liquid Cooled Petascale supercomputer and its cooling infrastructure. 2017, 107:87-100.