# Key Considerations for Liquid Cooling Architecture for Smart Computing Centres


## Abstract

With the average annual growth rate of AI arithmetic demand exceeding 30%, and the power density of a single cabinet exceeding 50kW, traditional air-cooling is approaching the physical limit. Liquid-cooling technology has become the inevitable choice for smart computing centres by increasing thermal conductivity efficiency by 50 times and reducing PUE to 1.3.

This whitepaper analyses the six key decision-making dimensions of liquid-cooled architecture and provides a technology roadmap for data centre construction.



GOTTOGPOWER
Energy Management Centre


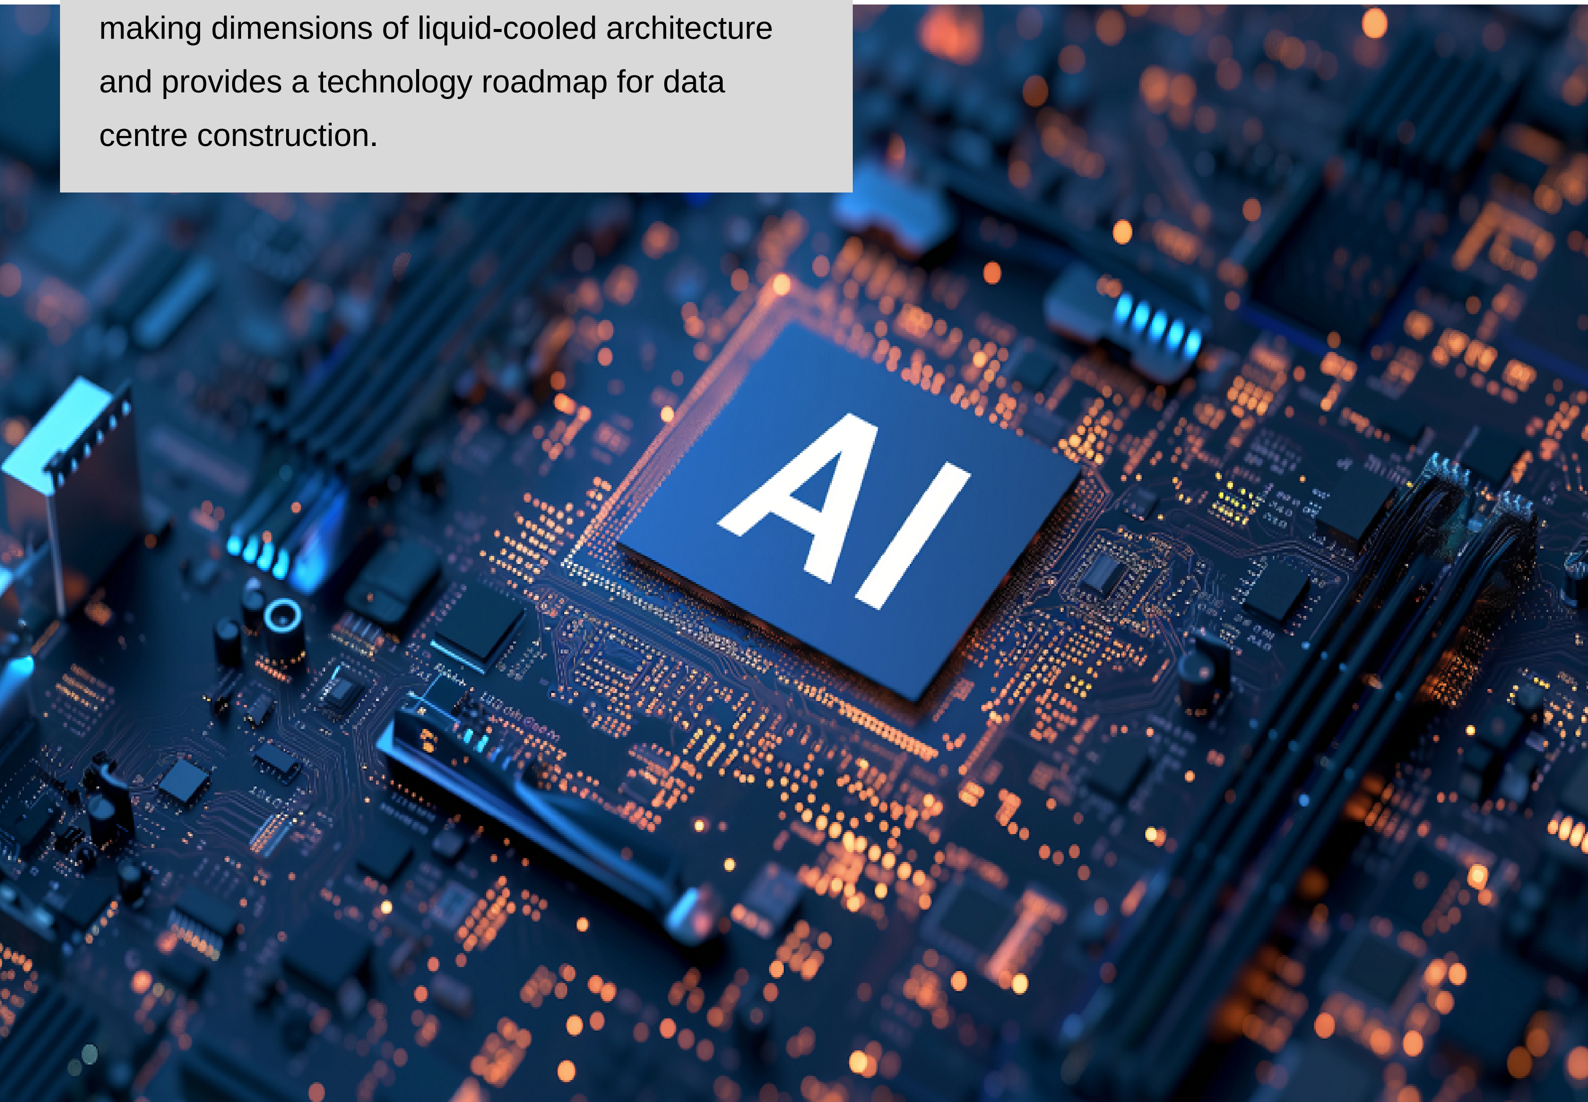

# Introduction

## 1.1 The Rise of Intelligent Computing Centres

Intelligent Computing Centres (ICCs), i.e. core infrastructures focusing on AI training, reasoning, and intelligent computing loads such as High Performance Computing (HPC), are the key engine for the booming digital economy. At a time when artificial intelligence, big data, cloud computing and other technologies are deeply integrated, the ICC bears the heavy responsibility of massive data processing and complex algorithms, providing powerful computing power support for finance, healthcare, transportation, scientific research and many other fields.

## 1.2 Thermal Dilemma of High-Density Computing

With the rapid development of chip technology, the power consumption of CPU, GPU, ASIC and other core computing chips continues to break through the limit, and the power consumption of a single chip has gradually broken through the kilowatt level. The resulting power density of the cabinet continues to rise, and the power density carried by the traditional air-cooled cooling method is far from being able to meet the needs of the Smart Computing Centre. Air as the medium of air cooling, due to its low heat capacity, poor thermal conductivity, in the face of high-power chip cooling is not satisfactory; at the same time, the complex air duct design, limited cooling distance, further constraints on the efficiency of air cooling.

From the perspective of energy consumption, air-cooled systems need to continuously run a large number of high-rotation fans to ensure the cooling effect, which leads to high cooling energy consumption and makes it difficult to reduce the power usage efficiency (PUE) of the data centre. Especially in hot areas, the air-cooled system consumes a large amount of electricity to maintain the temperature of the server room, which significantly increases the operation cost. In addition, the noise pollution generated by the operation of a large number of fans, as well as the space occupied by air-cooled equipment in the server room, also brings many problems to the construction and operation and maintenance of the data centre.

## 1.3 Liquid Cooling: The Way to Break the Dilemma

Under the heat dissipation dilemma of high-density computing, liquid cooling technology, with its unique advantages, has become a key technology path to solve the above problems. By using liquid as the heat transfer medium, liquid cooling technology can take away the heat generated by the chip more efficiently, breaking through the bottleneck of air-cooled heat dissipation, and providing a strong guarantee for the stable operation and sustainable development of the Smart Computing Centre.

## Liquid Cooling Technology

## 2.1 Basic Principle

Liquid cooling technology uses water or specialised coolant as the heat transfer medium to absorb the heat generated during the operation of the equipment by directly or indirectly contacting the heat source (e.g. chips, motherboard components, etc.). After absorbing heat, the coolant is transferred to the external cooling equipment through the circulating pipeline, and after releasing heat in the cooling equipment, the coolant will flow back to the heat source to continue to absorb heat, and so on, to realise the continuous heat dissipation of the equipment.

*Pic 1: Schematic diagram of the working principle of the coolant in liquid cooling technology*

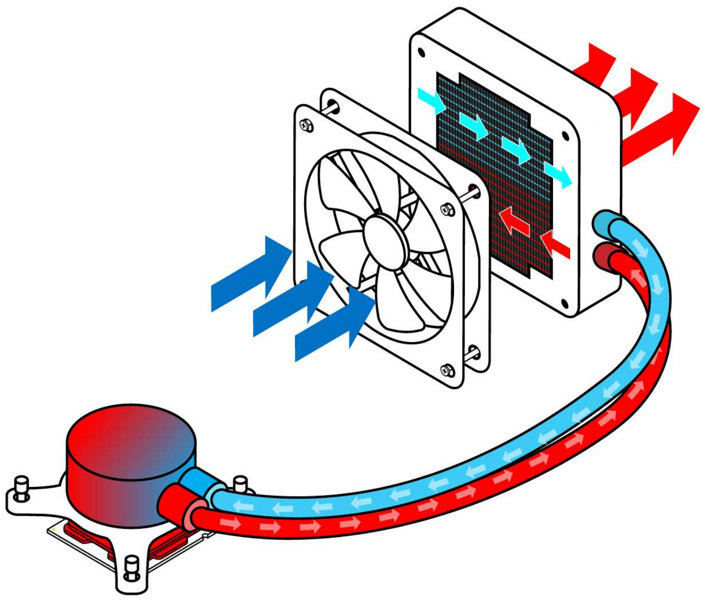

## 2.2 Core Advantages

· Ultra-high thermal efficiency:

The heat capacity of liquid is about 3500 times that of air, and its thermal conductivity is much higher than that of air, which makes liquid cooling technology able to quickly and efficiently cool high power density chips, and effectively deal with the heat dissipation challenges brought about by the surge in chip power consumption.

· **Significant energy saving and consumption reduction:**

Compared with air-cooled systems, liquid-cooled systems reduce the use of a large number of high-rotation fans, which reduces the energy consumption of the cooling system itself. In practice, the PUE of data centres with liquid cooling technology can be reduced to 1.3 or even lower, significantly improving energy efficiency and reducing operating costs.

· **Supports ultra-high density deployment:**

Breaking through the density limitations of air-cooled cooling, liquid-cooled technology can support a single cabinet with a power of 30kW or more, effectively saving space in the data centre and meeting the demand for high-density computing deployments in smart computing centres.

· **Reduced noise levels:**

Reducing or eliminating the use of high-speed fans enables liquid-cooled data centres to operate with significantly lower noise levels, creating a more comfortable working environment for staff.

· **Enhanced computing stability and reliability:**

Liquid cooling system enables more precise temperature control, reduces thermal stress caused by temperature fluctuations in the chip, extends the service life of the equipment, and enhances the stability of system operation.

· **Waste heat recovery potential:**

The heat generated during the liquid cooling process is higher temperature and more concentrated, which facilitates efficient recycling, and can be used for district heating, domestic hot water supply, etc., realising the secondary use of energy.

· **Stronger environmental adaptability:**

Relatively less affected by external ambient temperatures, liquid cooling technology can still ensure the stable operation of data centres in hot areas or scenarios where natural cooling is limited.

# Mainstream Liquid Cooling Architecture

## 3.1 Cold Plate Liquid Cooling

Cold Plate Liquid Cooling uses a metal cold plate (usually made of copper or aluminium) to closely fit on the surface of major heat generating devices such as CPUs and GPUs. The heat generated by the devices is transferred to the cold plate through heat conduction, and the coolant in the internal flow channel of the cold plate takes away the heat.

This solution is relatively small changes to the internal server, mainly replacing the radiator with a cold plate, and the coolant does not directly contact the electronic components, high security, maintenance is more convenient, applicable to the existing air-cooled server renovation and new projects.

However, for non-critical heat generating components such as memory, hard discs, VRMs, etc. in the cabinet, auxiliary air cooling may still be required. Currently, cold plate liquid cooling is the most commercially mature and widely used liquid cooling solution.

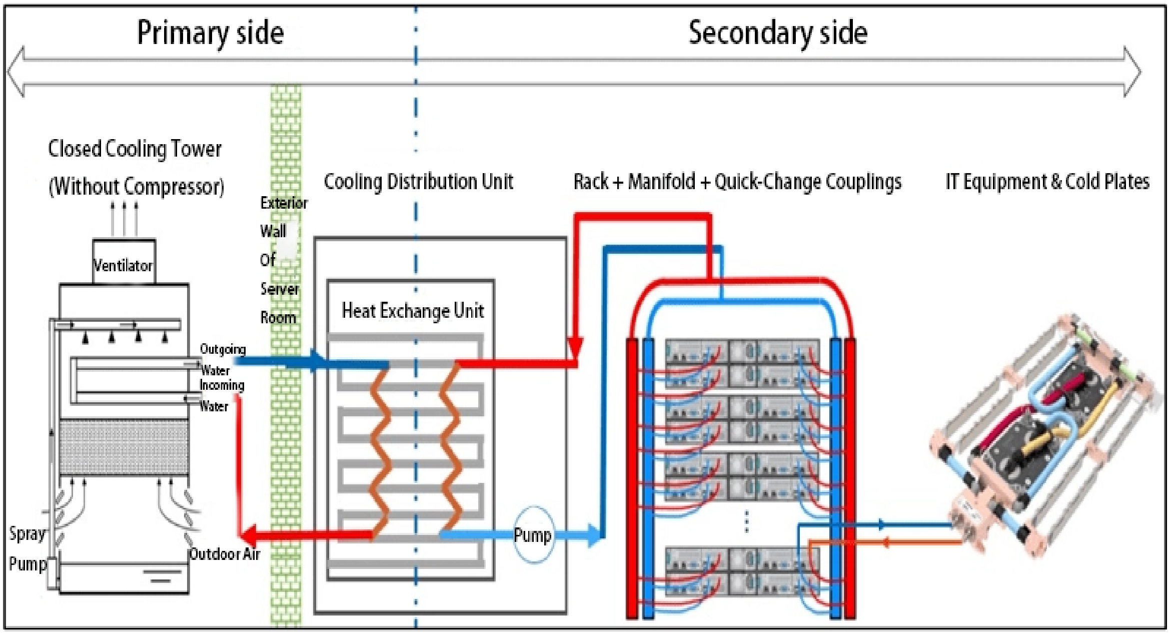


*Pic 2 Schematic diagram of cold plate liquid cooling*

## 3.2 Submerged Liquid Cooling

Immersion liquid cooling completely submerges the entire server or server motherboard in a non-conductive engineered coolant, which can be divided into single-phase immersion and dual-phase immersion.

·Single-phase submerged:

The coolant circulates in the liquid state to absorb heat, after the temperature rises, it is pumped to the external heat exchanger to dissipate the heat, and after cooling, it flows back to continue to absorb heat;

·Dual-phase submerged:

The coolant absorbs heat on the surface of the heating element and then boils and vaporises, using the process of phase change to absorb heat efficiently, and then the vapour rises to the condenser, condenses and liquefies, and then flows back.

*Pic 4 Schematic diagram of immersion liquid cooling*

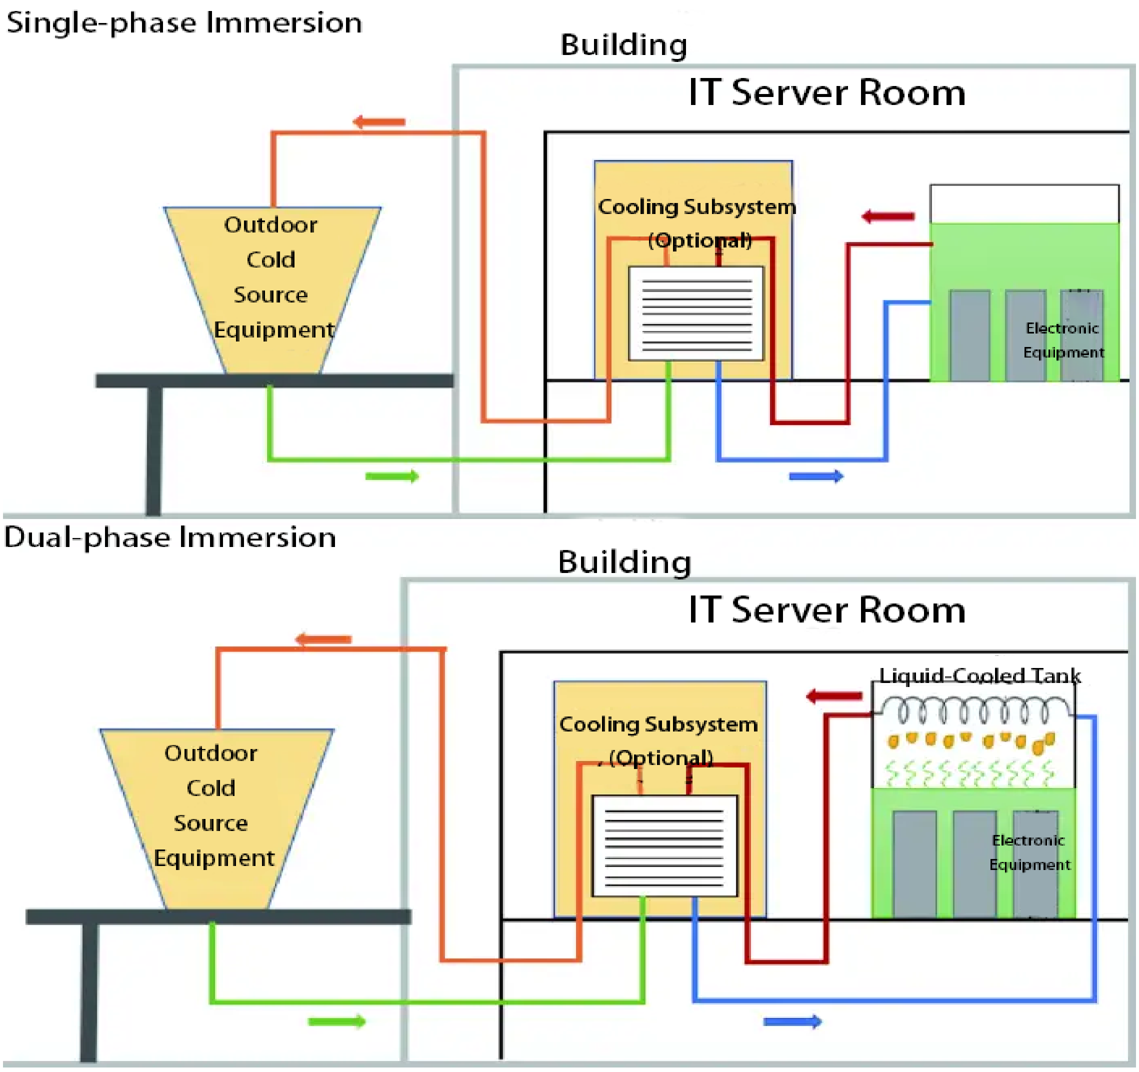

Submerged liquid cooling is extremely efficient, especially the two-phase submerged type, with excellent energy efficiency ratio. It completely eliminates the use of fans to achieve completely silent operation; the equipment can be tightly arranged to achieve ultra-high density deployment; at the same time, the coolant physically isolates the electronic components from the external environment, with dust, moisture and oxidation characteristics, significantly improving equipment reliability.

## Liquid Cooling Architecture Value Realisation

### 4.1 Enabling Higher Computing Power Density

Liquid-cooled architecture can effectively respond to the deployment needs of the next generation of higher-power AI chips and accelerator card clusters, and provide technical support for the Smart Computing Centre to continuously increase the computing power density to meet the growing demand for intelligent computing.

### 4.2 Build a Green and Low-Carbon Data Centre

By significantly reducing PUE and data centre carbon emissions, the liquid-cooled architecture helps achieve the ‘ dual-carbon’ goal, and promotes the development of data centres in the direction of green and low-carbon.

### 4.3 Enhance Business Continuity and Reliability

A more stable operating environment effectively reduces the risk of downtime in the Smart Computing Centre, safeguards business continuity, and provides users with reliable computing power services.

### 4.4 Optimise Total Cost of Ownership

Despite the higher initial investment in liquid-cooled architecture, better Total Cost of Ownership (TCO) can be achieved over the equipment lifecycle through benefits such as energy savings (lower electricity bills) and space savings.

### 4.5 Unlocking the Value of Waste Heat

Waste heat generated by the liquid-cooled architecture can be recycled at scale to create a new revenue stream for the data centre or used to meet its own energy needs, such as district heating, to further reduce operating costs.

## Future Development Trends

### 5.1 Technology Integration and Innovation

Liquid-cooling technologies such as cold plate and submerged liquid cooling will continue to be optimised and upgraded, and new types of high-efficiency and environmentally-friendly coolant will be continuously researched and developed; the exploration of chip-level and micro-channel liquid-cooling technologies will be accelerated to improve heat dissipation efficiency; and the commercialisation of two-phase submerged liquid-cooling will be accelerated with the scope of application gradually expanded.

## 5.2 Standardisation and Openness

Organisations such as the Open Computing Project (OCP) and ODCC will promote the unification of liquid-cooling technology interfaces, cabinets, management specifications, etc., lowering the threshold and cost of deploying liquid-cooling architectures, and promoting collaborative development of the industry.

## 5.3 Intelligence and Automation

AI technology will be deeply integrated into the liquid-cooling system toachieve intelligent tuning and detect potential equipment failures in advance through predictive maintenance; the degree of automated operation and maintenance will be continuously improved to reduce manual intervention and improve operation and maintenance efficiency.

## 5.4 Prefabrication and Modularisation

The degree of prefabrication of liquid-cooled cabinet modules, CDU modules, etc. is further enhanced, and standardised design and production are adopted to achieve rapid deployment and shorten the project construction cycle.

## 5.5 Wider Industry Applications

Liquid cooling technology will gradually expand from large-scale intelligent computing centres and supercomputing centres to enterprise-level high-performance computing, edge computing nodes and other fields, with richer application scenarios.

## 5.6 Scale application of waste heat recovery

The integration of waste heat recovery technology and business model will be more mature, realizing the efficient use of waste heat and commercial operation, and bringing greater economic and environmental benefits to the data centre.

## Conclusion

Liquid-cooled architecture has evolved from a cutting-edge technology to a key pillar supporting the sustainable development of smart computing centres.

In the face of the cooling challenges posed by high-density, high-power computing, liquid-cooling technology provides a practical solution by virtue of its excellent cooling performance, significant energy savings, and ability to support ultra-high-density deployments.

Despite the challenges of cost, technology, operation and maintenance in the implementation process, the popularisation of liquid-cooling technology is accelerating with the continuous maturity of the technology, the gradual unification of the standards, the improvement of the industrial chain, and the increasing prominence of the TCO advantage.

For ICC, embracing liquid-cooling technology is a strategic choice to enhance arithmetic performance, reduce operating costs, practice green and low-carbon development, and win competitive advantages in the future.

In the future, smarter, more efficient and more open liquid-cooling solutions will continue to drive the evolution of ICC to a higher level.

## Contact Us

For questions, feedback and suggestions on the content of this white paper or on the data centre project please contact:

**info@gottogpower.com**

gottogpower.com

0086 18326071160