AIRTRUNK

ACCELERATING A LIQUID-COOLED FUTURE FOR DATA CENTRES

Delivering the energy-saving benefits of liquid cooling to support growing demand for artificial intelligence capacity

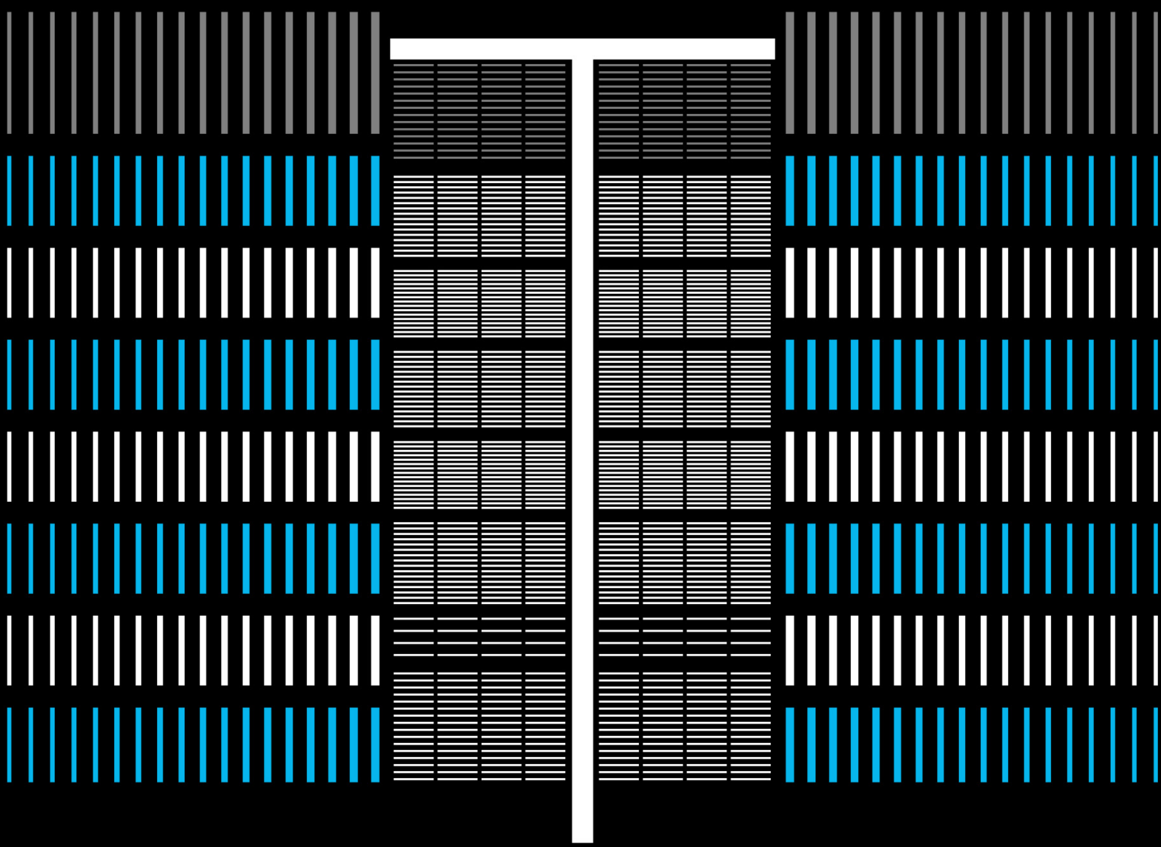

CONTENTS

1. WHY DATA CENTRES NEED LIQUID COOLING NOW

Billions of digital interactions taking place in today's society are powered by data centres. They are the engine rooms of the digital economy. As one of the fastest growing technologies in history, artificial intelligence (AI) is not only reshaping our lives, but also requires a complete rethink of data infrastructure.

The demand for data centre capacity has skyrocketed to support AI. In Asia-Pacific & Japan (APJ) alone, the data centre sector is expected to see significant expansion, with 13.3% annual growth forecast by 2028.[1]

The heat generated in data centres by increasingly powerful computing operations is testing the limits of traditional air-cooling methods. This necessitates a shift towards a new, innovative cooling method - liquid cooling.

How liquid cooling supports sustainable outcomes

Using liquid to cool powerful chips and servers offers a more energy-efficient alternative to traditional air-cooling. Implementing liquid-cooling in data centres offers the potential to deliver sustainable outcomes and environmental benefits for the industry.

AirTrunk predicts that more than 30-60% of all new data centre loads deploying AI hardware will be liquid-cooled. This transition is essential not just for managing the exponential growth of AI workloads, but as an evolution towards making data centres more efficient.

For data centres to meet the demands of AI and cloud computing and to meet new standards in sustainable operations, liquid cooling is a critical innovation for the industry.

Times - and timing - have changed

As recently as last year, public cloud and large technology companies were anticipating a roll out of liquid-cooling technology in 2027 and beyond, starting in Europe and North America before being introduced to the APJ market. In the last few months, this prediction has changed significantly, with liquid-cooling implementation accelerating. The next two years will be particularly critical as the industry moves from the current state of research and development to effective deployment of liquid-cooled data centre capacity, at scale.

What is liquid cooling?

Liquid cooling refers to technology which uses liquids, such as water or special coolants, to absorb and dissipate heat.

Liquid-cooling systems in data centres can be broadly categorised into two types: cold-plate and immersion-cooling systems.

Cold-plate systems rely on a "plate" type heat exchanger, typically made of copper, which is attached directly to major components in a server. Heat generated is removed from the server by liquid circulating through the cold plate.

Immersion cooling involves placing entire servers or specific components directly into non-conductive liquid. The liquid absorbs heat from components and is either passively cooled through natural convection or actively cooled by circulating it through a heat exchanger. A key consideration for immersion cooling is non-conductivity.

[1] Structure Research

Liquid cooling innovation

AirTrunk has been developing deep expertise in liquid-cooling technologies since 2019. With significant investment in resources, research, technology development and solutions, AirTrunk has pioneered innovation in large-scale liquid-cooling systems, able to support tens to hundreds of megawatts.

This investment is paving the way for liquid-cooling ready hyperscale data centres across the APJ region.

In this report, we share AirTrunk's journey in deploying more than 20 megawatts (MW) of liquid-cooled capacity in Malaysia. This project demonstrates our commitment to realising significant energy savings for our customers and ensuring AirTrunk continues to provide industry-leading PUE performance as we expand in tropical APJ climates. We discuss the challenges and possibilities liquid cooling presents, and provide an overview of key criteria for assessing this emerging technology based on our experience.

Liquid cooling development at AirTrunk

2019
Liquid-cooling working group established in AirTrunk to gather industry knowledge of the technology and application.

2020
First customer engagement with design provisions for liquid cooling.

2021
Engagement with multiple industry leaders and manufacturers to understand the nuances of deployment to de-risk future projects, including tours of live deployments to gain an understanding of practical applications.

2022
Significant capital invested in research, development and prototype of CDU and liquid load banks, including comprehensive factory testing.

2022-23
Over 40MW of liquid-cooling technology in-flight. Unique AirTrunk ready cooling design developed.

2023
Further factory testing of secondary distribution systems and rack manifolds with industry partners. Additional liquid- cooled MW in pipeline.

2024
Initial 20+MW liquid-cooling capacity deployed at AirTrunk JHB1.

2. DEMAND FOR AI IS DRIVING DATA CENTRE TRANSFORMATION

AI adoption is driving demand for infrastructure to support rapid growth.

As AI models become more sophisticated, and generative AI moves to multi-modal formats, the processing power required to train and run these models increases dramatically, as the graph below shows.

AI accelerators are the most critical component needed to meet rising demand for the computational capacity AI requires.

Currently, graphic processing units (GPUs) are the most popular type of AI accelerator due to a number of characteristics they have that support efficiency in AI modelling. Future generations of GPUs are expected to be even more sophisticated to keep up with the computing demands of AI applications as they evolve.

Amount of compute used to train notable AI models

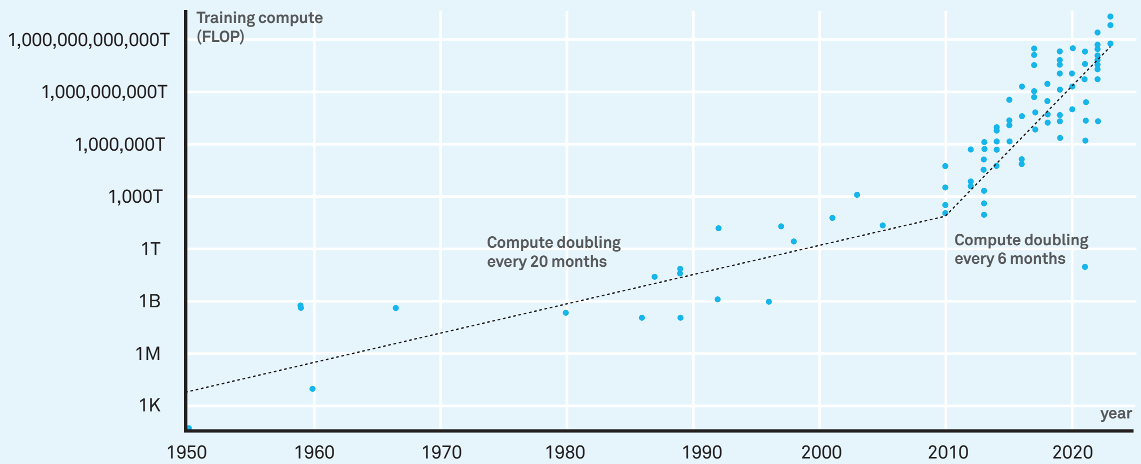

FLOP (floating-point operations) refers to the total number of computer operations used to train an AI system. Computation is estimated based on published results in the AI literature and comes with some uncertainty. Epoch expect most of these estimates to be correct within a factor of 2, and a factor of 5 for recent models for which relevant numbers we not disclosed, such as GPT-4.

Chart: Will Henshall for TIME. Source: Epoch via World in Data

More power in less space increases heat density

Introducing more GPUs into data centres has significantly boosted AI capabilities. This change in data infrastructure also presents a key challenge for data centres – maintaining the air temperature in data halls so that server equipment can operate safely and efficiently.

As AI workloads continue to increase, data centres will need to deploy more capacity. Adding more GPU servers per rack can amplify the computational power available within the same data centre footprint. Not only does this increase overall data centre capacity, but also reduces internal communication latency between each GPU.

Increasing GPU density in this way can create clusters of high-performance computing resources that allow data to move between processors with minimal delay. It also enables data centres to make more efficient use of the resources that support server and AI operations, such as network connections and power supply.

The graph below shows how data halls are able to increase their GPU density to achieve a higher output of computational power within a much smaller gross floor area.

More computing power in less space: data hall GFA reduction vs kW/rack

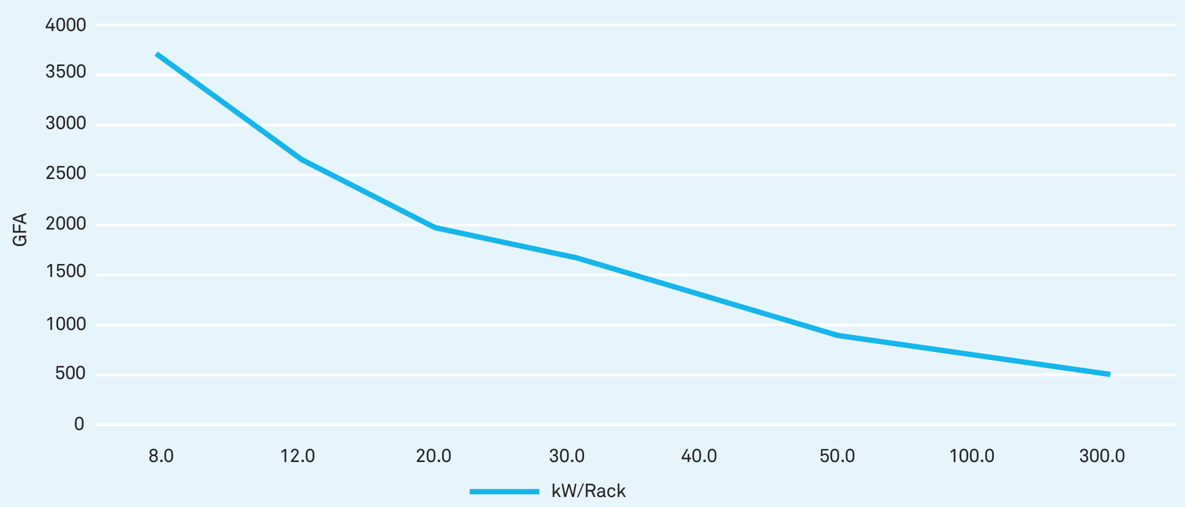

Air-based cooling can't keep up with increasing heat output

Increasing the density of GPUs and computing power output also increases heat output and this is creating issues for the air-based cooling solutions typically used in data centres. Air-based cooling has been found to be effective in maintaining safe and efficient operation of IT hardware at a power density of up to 15 kW/rack. With the GPUs now in use, power density of more than 40 kW/rack is becoming more common in data centres, as shown in the graph above. This change is bringing liquid-cooling solutions to the forefront as the only viable solution for managing high-density servers with capacity beyond 20 kW per rack as shown in the diagram below.

Cooling method applicability tipping point

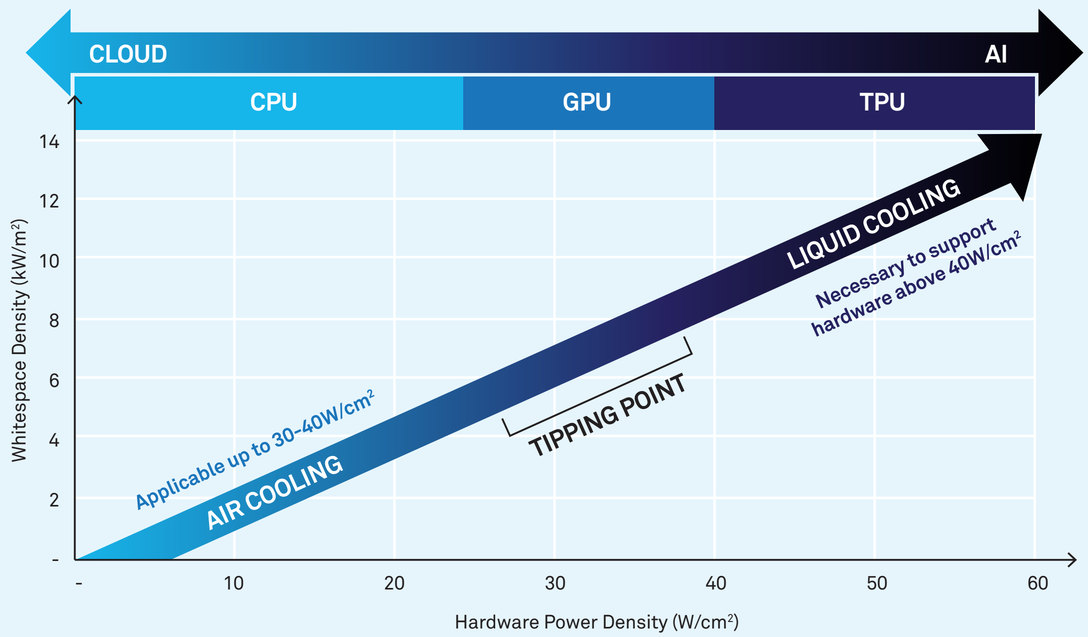

How liquid cooling supports the new era of AI

Liquid-cooling solutions enable data centres to maintain a safe and effective air temperature by dissipating heat coming from the operation of high-density GPUs. When compared to traditional air-cooling systems, liquid-cooling offers multiple benefits.

Superior heat dissipation

Liquids like water have higher thermal conductivity and can dissipate more heat per unit volume than air. This means liquid-cooling systems can be faster and more effective in removing heat from GPU units, maintaining optimal operating temperatures even under intense computational loads.

Energy efficiency

Liquid-cooling systems can operate effectively at higher temperatures. This removes the need for mechanical refrigeration as part of the system, which, in turn, reduces the energy required to run liquid-cooling loads.

Densification

Liquid-cooling systems can be designed to deliver effective cooling in the smaller spaces available in high-density GPU configurations.

Reduced hotspots

High-density GPU configurations can produce uneven heat distribution, with these hotspots negatively impacting performance and lifespan for hardware. Liquid-cooling systems can be designed to target these hotspots directly, ensuring a more even temperature distribution across all components. This precise cooling delivery prevents thermal throttling of GPUs, maintaining consistent performance under higher computing loads.

Noise reduction

With no fans required within the data halls, liquid-cooled systems operate more quietly. This can support compliance with local noise regulations and offer an additional benefit in environments where noise levels affect workplace quality.

3. THE NEXT TWO YEARS WILL BE CRITICAL FOR LIQUID COOLING

Liquid cooling is perhaps the biggest change in data centre technology in more than 30 years. Until now, even taking a water bottle into a data hall was out of the question, demonstrating how this technology shift is changing the way traditional data centres will be designed and operated in the near future.

Exploring new solutions to evolve liquid-cooling technology

With the rapid scaling of data processing demands, many technology providers are continuing to explore the optimal specification for their high-density server rack installations, including the desired power density per rack. This, in turn, will impact their requirements for liquid versus air-cooled capacity and target temperatures for liquid supply.

While server configurations and cooling performance requirements continue to evolve over the next two years, AirTrunk foresees data centre operators developing and testing a wide variety of solutions. As with any new technology, designing effective and safe solutions is challenging as there is no data and limited practical experience available to inform decision making. There is high likelihood of 'false dawns' in fit-for-purpose liquid-cooling systems as these pilot designs move from desk to data hall.

Having already tested and deployed a system to deliver 50MW liquid-cooling capacity, AirTrunk has worked with leading technology providers to understand and meet their configuration and performance needs. This process also addressed design features, and installation and maintenance solutions to minimise risks of damage to equipment or system failures.

In the next section we explore key considerations for ensuring system performance levels and reducing risks.

Liquid-cooling technology solutions

The industry is currently focusing on cold-plate liquid-cooling systems as the solution to adopt liquid cooling moving forward.

Preferred for ease of operation - when compared to immersion systems - these solutions are also compatible with the specifications of current AI hardware.

Plus, immersion system fluids may negatively affect certain chips and components making it a less viable option at present.

Immersion technology may have potential as a future solution as new developments in fluid technology improve compatibility with IT equipment.

4. FROM DESIGN TO DEPLOYMENT: AIRTRUNK CASE STUDY

Case study: a pioneering liquid-cooling solution in Johor Bahru, Malaysia

AirTrunk's customers are some of the world's largest technology companies. As an innovative business, it is evolving the way its data centres are designed, built and operated to better serve their ever-changing needs and requirements. Key to this evolution is ensuring AirTrunk's data centres are equipped to manage the increased computing load and GPU density required for AI, machine learning and high-performance computing applications while supporting greater energy efficiency.

Data centre design for performance and energy efficiency in tropical climates

Adopting liquid-cooling technology supports deployment of AI at scale while targeting better outcomes for energy efficiency. This is particularly important in tropical locations, which are generally subjected to higher Power Usage Effectiveness (PUE) due to increased cooling requirements linked to the climate.

Malaysia's most efficient data centre

AirTrunk's JHB1 campus is the first of several data centres planned for Johor Bahru, Malaysia. Deploying a high-performing, low-risk liquid-cooling solution for this facility was essential to optimise the operating environment for the high-density servers it houses to achieve consistently high levels of performance. JHB1 delivered the initial 50+MW phases in record time, with the first tranche of capacity already commissioned and operational since mid-2024.

Project approach – For JHB1, AirTrunk developed, commissioned, and installed a first of its kind direct-to-chip liquid-cooling technology alongside a traditional indirect evaporative cooling (IEC) system to meet customers' strict performance and risk standards. This required extensive customer consultation as well as prototyping and testing off-site to ensure those standards were met.

This project demonstrates AirTrunk's ability to deploy computing capacity at speed and deliver the performance and reliability its technology customers need to support their rapid growth.

Designing to optimise server performance

The purpose of the cooling systems in data centres is to keep the temperature of all components in the server racks within a range that permits efficient operation of the IT equipment and its ancillary components. If the system exceeds the desired temperature, or the equipment is subject to uneven temperatures, this can impact performance or even damage the hardware installed.

To ensure effective performance of the liquid-cooling system, AirTrunk considered several key design considerations:

CDU design – AirTrunk engaged with its partner Delta Electronics Inc., a global power and thermal management solutions provider, to help develop a bespoke solution to meet AirTrunk's customer-specific requirements. This joint development included several iterations of the design, a range of prototypes, and iterative testing procedures. The main features of the CDUs are the monitoring and continuous protection of water quality within the secondary loop, the control system that ensures continuous supply of the right amount

AirTrunk's environmental commitment

AirTrunk has linked 100% of its multi-billion financing platforms to sustainability commitments. This means the company is accountable for designing, deploying, and operating its data centres in an efficient and sustainable manner.

In 2022 AirTrunk announced its commitment to achieve Net Zero emissions by 2030 for Scope 1 and 2 emissions, and to make further progress on measuring and addressing emissions from embodied carbon.

The investment AirTrunk is making in liquid-cooling technologies for its clients and infrastructure is an important part of our commitment to achieving these targets.

AirTrunk JHB1 initial phases

of liquid, and robust pump and heat exchanger design. Access for maintenance is also a key feature that was developed with the knowledge that the units are located within critical customer areas, the data halls, and not in service corridors or technical areas.

Monitoring liquid quality – Liquid quality is one of the most important factors in system performance as concentration levels of certain components can affect its thermal properties. At JHB1, liquid quality is strictly monitored and maintained in each liquid-cooling subsystem.

Hybrid heat rejection systems – The target temperature of the liquid used is a major factor in the choice of heat rejection technology and the overall energy efficiency of the system. If the system can operate using liquid with a higher supply temperature, this introduces the potential to remove compressors from the cooling system and increase energy efficiency.

Another consideration is whether heat rejection systems for liquid cooling should be shared with non-liquid-cooled loads. In the case of JHB1, AirTrunk took a hybrid approach, with some air-cooled loads serviced by the IEC systems and others serviced by the same facility water used for liquid-cooled loads. This solution takes into account a range of factors including energy efficiency, cost and ease of operation and installation.

Designing to minimise system risk – Developing liquid-cooling systems for JHB1 has enabled AirTrunk to fully explore, evaluate and mitigate the risks of running liquid pipework within a data hall environment. To ensure its customers are comfortable with this cooling technology, AirTrunk designed a solution delivering the cooling benefits of liquid while minimising the risk of system failure or damage to GPUs.

Two key areas of AirTrunk's solution design deliver these outcomes for risk mitigation:

1. **Redundancy –** Designing for redundancy is an important principle to safeguard the continuous delivery of liquid cooling. At JHB1, liquid reaches each server rack via a pipework loop with two CDUs per loop deployed in hot standby mode. With each CDU capable of providing 100% of the flow required for the cooling load, this ensures a stable amount of liquid can be supplied to the servers at any time.
2. **Access for routine and emergency maintenance –** To support safe maintenance procedures and the acceptable blast radius required by the customer, the design features the required number of isolation valves within each pipework loop to disable each component as needed. Another important element is the use of quick-disconnect couplings to connect and disconnect each of the manifolds within the customer rack to avoid any risk of water spillage during upgrades or maintenance procedures.

Success measures

AirTrunk, in collaboration with its customer, developed a new set of concepts applicable to liquid-cooling systems and set the standards for operating systems and maintenance during the operation of the data centre. To deliver service levels to support this outcome, the project established target ranges for system characteristics including flow rate, rate of change, temperature, differential pressure, and water quality parameters.

5. LEVERAGING LIQUID-COOLED SYSTEMS TO ACHIEVE SUSTAINABILITY GOALS

Data centres can enhance their sustainability by using liquid-cooling technology, which is typically more energy-efficient than air cooling.

Liquid-cooled systems can handle higher temperatures, reducing or even eliminating the need for energy-intensive compressors in almost all climates. This higher operational temperature also gives liquid-cooled systems an energy efficiency edge over air-cooled systems, even for those operating at elevated supply air temperatures. Typical comparison numbers for facility water supply temperatures from the heat rejection systems are:

Compressor-less air-cooled

30°C

Compressor-less liquid-cooled

35 to 40°C

This modest 5°C difference has a significant impact, allowing for compressor-less operation throughout the year and reducing the power required by fans from the heat rejection systems, due to the higher allowable approach temperatures.

Unlike elevated supply-air temperature air-cooled systems, liquid-cooled systems do not adversely affect human comfort in the data hall as the heat is contained within the liquid distribution systems. As the technology evolves over the next few years, it is expected that higher temperature liquid systems will be introduced as well as two-phase (2P) fluids, further enhancing energy and water reduction benefits.

Improving PUE

The ideal PUE is 1.0, meaning that all energy is used for computing with none wasted on other processes. Liquid-cooling technologies significantly reduce the energy required for cooling systems, thereby lowering the overall energy consumption, and improving the PUE of a data centre, as summarised in the table below.

Efficient heat rejection

Hybrid dry coolers present as advanced solution for heat rejection in liquid-cooling systems. These systems combine the features of both dry and evaporative cooling processes to optimise thermal management and energy usage.

By primarily using ambient air for the heat rejection process, hybrid dry coolers consume significantly less energy compared to traditional refrigerant-based air conditioning systems. The evaporative cooling component is only used when necessary, minimising use of process water.

Opportunities for heat reuse

For cooling systems where facility water is used to efficiently capture and transport heat away from data centre components, the waste heat can be easily repurposed, rather than simply expelled into the environment.

As liquid-cooling systems use facility water at a higher temperature than air-based systems, the opportunities for reusing the heat from servers become feasible. For example, the excess heat can be used for pre-heating water in industrial processes, or for heating agricultural greenhouses or fish farms. Heat reuse not only reduces the energy consumption of those external processes but also contributes to a reduction in the overall energy footprint for the data centre.

Operating Temperatures

Data centre customers may hesitate to adopt higher fluid temperatures, missing opportunities to reduce cooling energy use and potentially increasing their PUE ratio.

	Typical air-cooled solution	Cold plate liquid cooling	Immersion liquid cooling
Peak PUE	1.40-1.50	1.28-1.32	1.20-1.22
Average annual PUE	1.25-1.35	1.18-1.22	1.14-1.16

Data from calculations and modelling completed by AirTrunk

6. THE FUNDAMENTALS OF LIQUID COOLING SYSTEMS

Properties of air versus water for cooling systems

As the table below shows, water can absorb heat and transport it more efficiently than air can. The typical mass heat capacities of air and water mean water has 4.2 times the mass-based specific heat capacity of air. In addition, water is 784.31 times denser than air and has 3238 times the volumetric-based specific heat capacity of air.

Liquid-cooled systems take advantage of these significant multipliers to remove large amounts of heat.

Heat absorption capacity of air vs water

	Air	Water
Mass heat capacities	1000 (Approx.) Jkg-1K-1	4200 (Approx.) Jkg-1K-1
Density	1.275 kg/m3	1000 kg/m3
Volumetric heat capacity	0.001297 J·K–1·cm–3	4.2 J·K–1·cm–3

Main types of liquid-cooling systems for data centres

Liquid-cooling systems in data centres can be broadly categorised into two types: cold-plate and immersion cooling systems, with each having different setups and components.

Cold-plate systems (Direct to Chip or DCLC)

Cold-plate systems, as the name suggests, rely on a 'plate' type heat exchanger, typically made of copper, which is attached directly to the major components inside a server. Heat generated by GPUs and other components is removed from the server by liquid circulating through the cold plate.

The majority of these systems still require some percentage of heat rejection to be achieved by standard air-cooled systems operating in parallel with the liquid-cooled system. The percentage of liquid cooling varies from 50% to 80% of the total power drawn by the server.

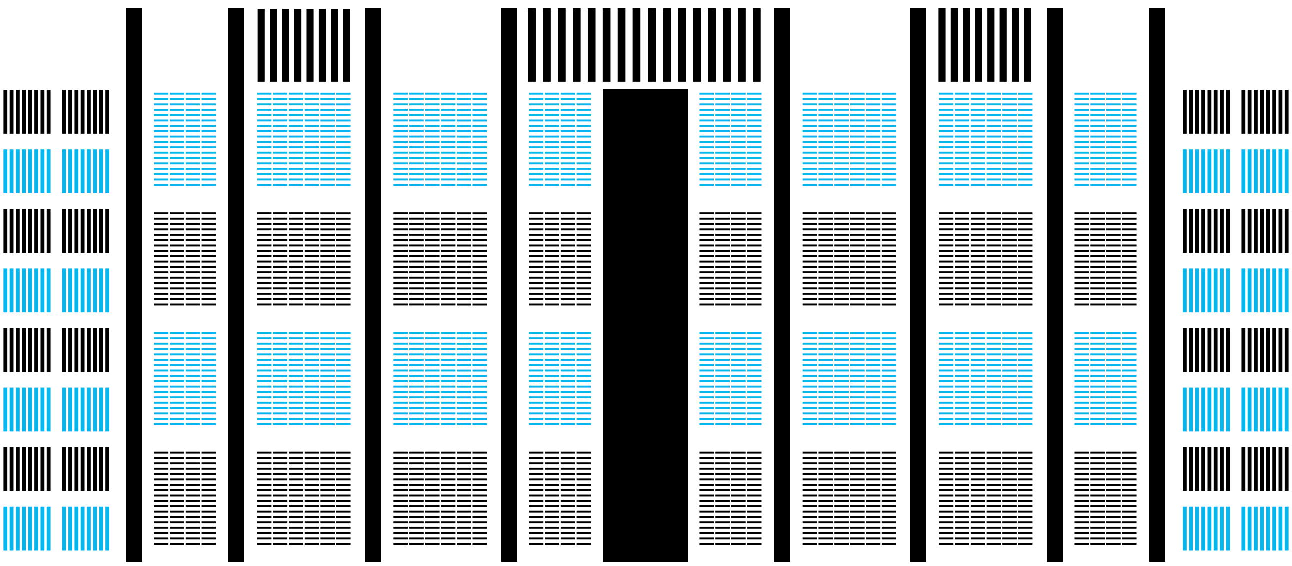

A typical cold plate liquid cooling system consists of four main components:

Cold plate

Liquid-cooling manifold

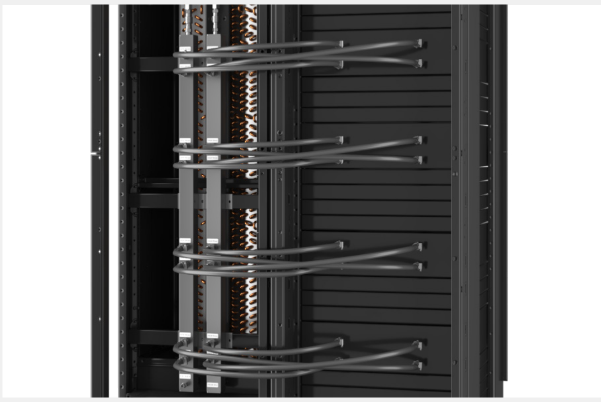

Secondary system piping

CDU

The four components are connected back to a typical heat rejection system consisting of a primary piping system, and heat rejection devices such as cooling towers, hybrid dry coolers or dry coolers.

The CDU distributes liquid of an acceptable quality to the rack manifolds and server cold plates in a controlled manner using a variable speed pump. It isolates the "dirty" primary cooling water from the "clean" secondary system flow using a plate and frame heat exchanger. CDUs are installed in zones related to secondary system size and are configured in either 2N or N+1 configuration.

Liquids used in cold-plate systems

The main consideration in selecting the fluid for cold-plate liquid-cooling systems is to prevent corrosion and the formation of particulates that can affect the flow of the liquid and the heat transfer, particularly around the very sensitive components within the cold plates. Deionised water is a low-cost option that can be produced onsite using reverse osmosis water treatment systems. This fluid will need to be treated with inhibitors to maintain the water chemistry levels required as it can leach ions out of various metals.

A more expensive but stable option are glycol mixtures. OAT-PG-25 is a general-purpose coolant with many applications and recommended by several server manufacturers. It is a pre-mixed organic-acid-technology fluid made up of 75% water and 25% propylene glycol solution, containing a small quantity of inhibitors and is formulated to minimise corrosion of various materials in contact with the solution.

The use of two-phase fluids in cold-plate systems in the near future has potential to increase the capacity of the liquid-cooled servers and increase the density of the liquid-cooled systems.

Immersion cooling

Immersion cooling involves placing entire servers and other components directly into a non-conductive fluid. The fluid absorbs the heat directly from components and is either passively cooled through natural convection, or actively cooled by circulating it through a heat exchanger.

A typical immersion system consists of two main components: the immersion tank and the CDU.

In some cases, the CDU is integrated within the immersion tank which is then connected back to a typical heat rejection system consisting of a primary piping system and heat rejection devices, such as cooling towers or dry coolers. Alternatively, the CDU is provided as a separate component that needs to be interconnected to multiple tanks. In this case the current industry standard is that the interconnection of CDU to tanks is achieved by a single vendor as an "assembly" of immersion tanks and CDU.

Liquid alternatives in immersion cooling

The liquids used to immerse server boards vary widely in composition, availability, and cost and can be single phase or two-phase fluids. A key consideration for immersion cooling is non-conductivity, a property found in mineral oils, fluorocarbon-based liquids, and synthetic liquids and oils. The coolant should also be non-corrosive to all server components, non-hazardous to humans and bio-degradable with zero or very low global warming potential (GWP). The petrochemical industry is currently developing fluids with these characteristics for wider use in the data centre industry.

Immersion tanks

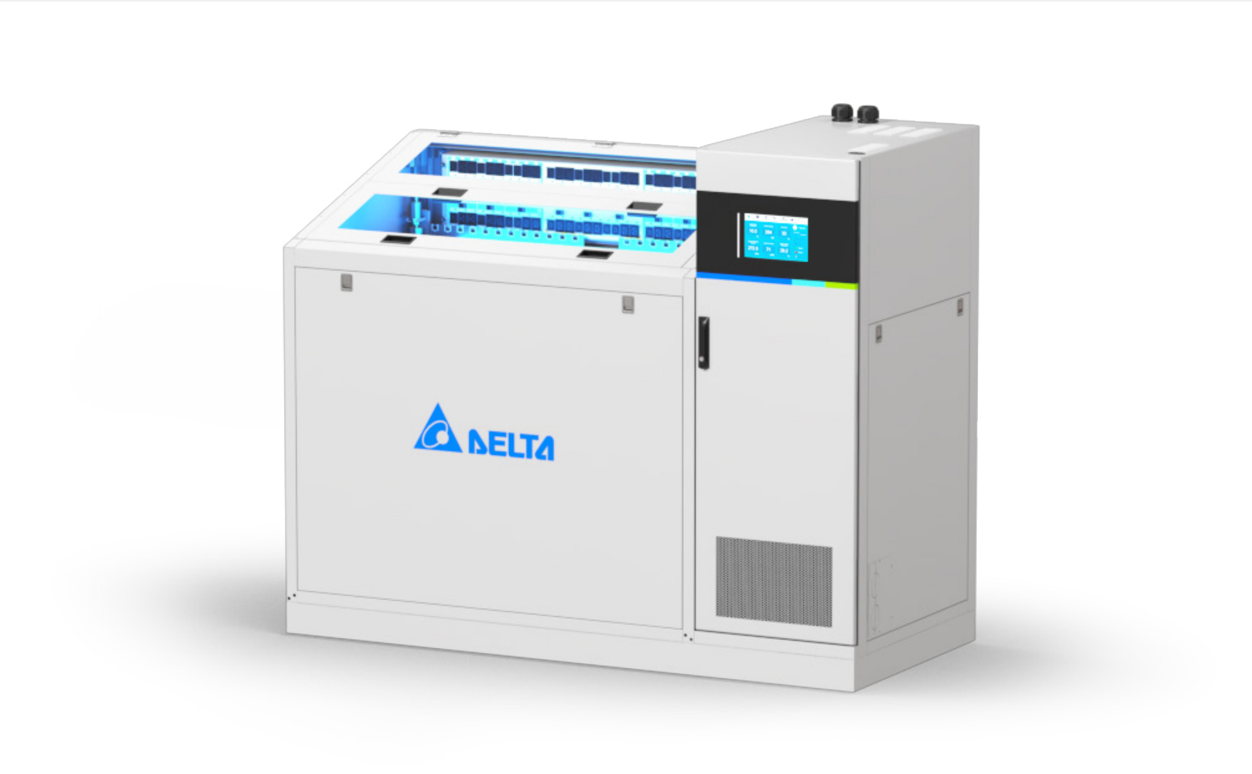

7. SEVEN STEPS TO ACCELERATE PROGRESS IN LIQUID COOLING

While liquid cooling offers a solution to manage the increased heat output from high-density GPU configuration in data halls, effective deployment takes significant investment in research and development. Achieving practical and sustainable outcomes takes a holistic approach, which focuses on realising the potential benefits of implementing liquid-cooling technologies without compromising on fundamental reliability and performance requirements.

Based on this analysis, AirTrunk considers six important factors which will shape future development and implementation of liquid-cooling technology in data centres.

01 Using the highest possible supply liquid temperature

Select the highest feasible supply liquid temperature while still enabling IT components to function at their full performance. The higher the supply temperature, the more opportunities exist for implementing compressor-less heat rejection systems and for reusing heat from the heat-rejection system.

02 Standardise to scale rapidly

Work with industry stakeholders to standardise the main properties of liquid-cooling fluids, operational and quality standards, and the hydronic requirements. Standardising systems, equipment and components in this way supports fast scaling of the technology.

03 End-user collaboration

Collaborate with end users to understand how liquid-cooling technology can facilitate the most efficient use of real estate and technical infrastructure for the sustainable deployment of new AI capacity.

04 Embracing Higher Fluid Temperatures

Despite the benefits of liquid cooling for higher fluid temperatures, many data centre customers remain hesitant. Lowering PUE is vital for improving energy efficiency, and addressing concerns can encourage greater acceptance of higher temperatures. Educating stakeholders and offering support can facilitate this acceptance, enhancing operational efficiency and sustainability.

05 Industry collaboration and regulations

Promote dialogue and collaboration between industry stakeholders, universities, and government agencies, to research and explore liquid-cooling technologies to support sustainable AI deployments. This could include developing policies to set quality and operational standards, adopting minimum energy efficiency thresholds for liquid-cooling facilities, or even mandating minimum adoption levels for liquid-cooling technologies in new data centre developments.

06 Research and development

Continue researching and developing new advancements in technology to further enhance the role of liquid cooling in supporting efficient AI. AirTrunk, as an example, is currently exploring new generation fluids, pre-assembled secondary pipework systems and high-efficiency heat rejection equipment as components for liquid-cooling systems.

07 Invest in new skills

Liquid-cooling technologies require a range of new skills to support development, design, testing, implementation, and maintenance within the data centre industry. Training and upskilling for design, construction and operation teams is required to overcome the challenges in defining, installing, and operating complex hydraulic systems and implementing the stringent water quality requirements for liquid-cooling technology.

GLOSSARY

Air cooling – Using air conditioners and fans to remove heat from servers and equipment in data centres, helping to maintain optimal operating temperatures.

Blast radius – The reliability term used to describe the potential impact on a system in the event of a component or sub-system failure.

Central processing unit (CPU) – The component of a computer system that performs the system's basic operations (such as processing data), that exchanges data with the system's memory or peripherals, and that manages the system's other components.

Cold-plate liquid cooling – These systems use a plate as a heat exchanger, typically made of copper. The plate is attached directly to the major components inside a server. Heat generated by GPUs and other server components is removed from the server by liquid circulating through the cold plate.

Communication latency – Time it takes for data to pass from one point on a network to another.

Coolant Distribution Unit (CDU) – Modular units that exchange heat with the facility water system and circulate the secondary fluid at the desired temperature directly onto the chip or server.

Direct-to-Chip liquid cooling (DCLC) – Another term used for cold-plate liquid cooling.

Facility water – The liquid that circulates in a closed loop between the heat rejection equipment and the terminal units (CDUs, fan-wall units, etc.)

Hot standby mode – A method used to build redundancy into a system to ensure reliability. It involves having one system run simultaneously as a backup for an identical primary system. If the primary system fails, the hot standby system immediately takes over, replacing all the functions of the primary system.

Global warming potential (GWP) – GWP represents how much a given mass of a chemical contributes to global warming, over a given time period, compared to the same mass of CO_2. CO_2's GWP is defined as 1.0. These values are calculated over a 100-year time horizon.

GLOSSARY

Graphic processing units (GPU) – A GPU is a computer chip that renders graphics and images by performing rapid mathematical calculations. Originally, GPUs were used for the rendering of 2D and 3D images, animations and video, but now they have a wider range of uses.

GPUs have thousands of small cores, which are designed to all work in parallel to achieve faster processing. This makes them well suited for AI computing that requires a large number of operations to be done at the same time.

Immersion cooling systems – With immersion cooling systems entire server racks or specific components are placed directly into a non-conductive liquid. The liquid absorbs the heat directly from components and is either passively cooled through natural convection, or actively cooled by circulating it through a heat exchanger.

Indirect Evaporative Cooling (IEC) – IEC lowers ambient air temperature by using water to evaporate. The cooler air is passed through a heat exchanger which is then used to cool the data hall air. By avoiding any mix between ambient and data hall air, this prevents extra moisture being added to the data hall air.

Liquid quality – Liquid quality is an important consideration for liquid-cooling systems as concentration levels of certain components can affect its thermal properties. Issues with liquid quality can also impact the flow of liquid and introduce corrosion risks.

Power Usage Effectiveness (PUE) – A measure of energy efficiency, PUE is a ratio based on how much of the total energy consumed by a data centre goes to computing equipment, rather than cooling and other overhead.

Two-phase (2P) cooling fluid – A fluid that changes phase when exposed to a heat source, vapourising into gas which then flows towards the cooling source, typically a CDU or an external heat rejection unit, where it releases the heat by condensing back into liquid phase.

2P fluids have significantly higher heat transfer per unit of volume compared with single phase fluids. When the cool source is located close to the heat source, 2P systems may not require pumps to move the fluid between both sources.

Water Usage Effectiveness (WUE) – WUE is the metric that measures the litres of water used per energy unit (kWh) of the data centre IT equipment. The lower the WUE figure is for a data centre, the more water efficient it is.

Unless otherwise indicated, material in this publication may be used freely, shared or reprinted, so long as AirTrunk is acknowledged as the source.

About AirTrunk

AirTrunk is a hyperscale data centre specialist creating a platform for cloud, content and large enterprise customers across the APJ region. The company develops and operates data centre campuses with industry leading reliability, technology innovation and energy and water efficiency. AirTrunk's unique capabilities, designs and construction methodologies allow it to provide customers with a scalable and sustainable data centre solution at a significantly lower build and operating cost than the market.

For more information on AirTrunk, visit airtrunk.com