



Don't lose your data centre cool

As compute workloads heat up (pun intended), the pressures placed on data centres are expanding exponentially. Liquid cooling stands to offer a scalable, efficient cooling solution supporting data centre operators in a sustainable manner into the future...

Liquid cooling technology has been around since the 1960s but is now playing an increasingly crucial role for data centre cooling.

"The rapid evolution of processing power and the exponential growth of AI are pushing the boundaries of what air-

cooled systems can handle. Air cooling is increasingly becoming insufficient for effectively dissipating the immense heat generated by modern servers and high-performance computing (HPC) systems," notes Matthew Thompson, Managing Director Europe, Airsys.

Today's data centres are tasked with powering AI and HPC applications, which generate far more heat. "Traditional forms of cooling — particularly air-cooled systems — are simply unable to handle these thermal loads, making liquid cooling a necessity to manage the immense amount of heat generated. Liquid cooling, including both direct-to-chip and immersion approaches, is a powerful solution to meet these demands, offering significantly higher heat rejection capabilities than conventional air cooling," explains Angela Taylor, Chief of

Staff, Head of Strategy, LiquidStack.

Andy Young, CTO, Asperitas, elaborates that "immersion cooling, where IT hardware is fully immersed in a dielectric fluid, offers a scalable, high-performance, and sustainable solution. Unlike traditional air cooling methods, immersion removes the need for traditional air-cooling infrastructure, significantly reducing energy consumption, footprint, and overheads."

According to Chris Carreiro, Chief Technology Officer, Park Place Technologies, as clusters of servers work together and with close proximity stacking, these high-powered systems in a single cabinet make it challenging for facilities to cool the servers.

"AI workloads perform better when you can cluster GPUs together in close proximity, which increases the challenge of capture and removing the heat that is produced," adds Lucas Beran, Director of Product Marketing, Accelcius.

"For example, a rack with 10x 1000W servers could need 500-600 CFM total airflow. The server exhaust feels too hot, and airflow (or CRAC/CRAH units) may be insufficient. This creates hotspots," outlines Carreiro. "Liquid cooling alleviates this problem for dense racks (10-15kW), where air cooling solutions would struggle. The liquid exchanges that heat somewhere outside of the rack and even out of the data centre itself."

"Liquids, due to their significantly higher thermal conductivity compared to air, offer a much more efficient solution for heat removal," adds Thompson. "This makes liquid cooling not just a viable, but a necessary technology for maintaining optimal operating temperatures."

Indeed, while CPUs in conventional servers typically have a TDP of 250-350W, AI accelerators such as Nvidia's GB200 reach 1.2kW per GPU — compared to just 700W for a traditional dual-CPU server. These extreme power densities demand more effective thermal management. As such, Bernie Malouin, Founder and CEO of JetCool Technologies, says that "liquid cooling is no longer a luxury — it's an industry necessity. It enables greater performance, reduces energy consumption with some liquid cooling solutions eliminating the need for ancillary water-intensive infrastructure like evaporative coolers and chillers."

Forward thinking

With AI and high-density chips resulting in unprecedented power demands, alongside grid capacity limitations and competition for power, data centre deployments are being delayed across the UK, sometimes by years, reports Malouin. Indeed, the International Energy Agency (IEA) anticipates data centres to double their energy usage by 2026, forcing the industry to rethink its approach to efficiency.

"Liquid cooling directly addresses these challenges by significantly reducing power consumption, eliminating water consumption, and enabling higher density computing without overloading existing grid capacity," shares Malouin. "Unlike traditional air cooling, which accounts for nearly 40% of a data centre's total energy use, liquid cooling removes heat more efficiently at the source, allowing AI and HPC workloads to run at peak performance with lower power draw."

"AI workloads perform better when you can cluster GPUs together in close proximity, which increases the challenge of capture and removing the heat that is produced," adds Lucas Beran, Director of Product Marketing, Accelcius.

"For example, a rack with 10x 1000W servers could need 500-600 CFM total airflow. The server exhaust feels too hot, and airflow (or CRAC/CRAH units) may be insufficient. This creates hotspots," outlines Carreiro. "Liquid cooling alleviates this problem for dense racks (10-15kW), where air cooling solutions would struggle. The liquid exchanges that heat somewhere outside of the rack and even out of the data centre itself."

"Liquids, due to their significantly higher thermal conductivity compared to air, offer a much more efficient solution for heat removal," adds Thompson. "This makes liquid cooling not just a viable, but a necessary technology for maintaining optimal operating temperatures."

The IT sector is no different to any other industry in the sense of its responsibility to the sustainability and carbon reduction targets set by the UK government.

"Liquid cooling is a necessary shift in the data centre paradigm to work towards these targets whilst also being mindful of the evolving landscape of the IT sector and its ever-deeper integration into our commercial and personal worlds," says Karl Lycett, Product Manager – Climate Control, Rittal. "Liquid cooling is significantly more efficient at transferring heat energy than air and thus, when coupled with the soaring heat loads being seen due to new technologies, must be one of the key tools in our arsenal to ensure data centres have maximum uptime to support our 'always on' lifestyles."

Moreover, as well as the energy-saving benefits, liquid cooling also wins for versatility.

"The underlying infrastructure can be adapted to facilitate future improvements in cooling technology and the expansion

data centres will choose to implement a hybrid approach, which outlines the needs of their installation and will then work with a supplier to choose the right cooling technology for their, in some instances this may be zoning off the environment

"Unlike traditional air cooling, which accounts for nearly 40% of a data centre's total energy use, liquid cooling removes heat more efficiently at the source, allowing AI and HPC workloads to run at peak performance with lower power draw."

of server capacity," explains Sam Evans, Associate Director, Excool. "Essentially, you don't have to swap out the infrastructure as you continue to update your chillers. As rack densities will only increase in the future, we need to find the best cooling solution to reduce temperatures in the quickest and most energy-efficient way — that's what liquid cooling does."

"There's also the question of resilience in a warming climate," highlights Paul Mellon, Operations Director, Stellium Datacenters. "The UK may be mild compared to some parts of the world, but summertime heatwaves and rising average temperatures mean air-cooled data centres often have to crank up the fans to cope. Liquid cooling can be a more stable, predictable solution, regardless of seasonal fluctuations."

Technological approach

With multiple liquid cooling approaches available, there are yet more options for the data centre operator to consider.

"The different methods are suited to different applications and who, how and what the data centre has been designed to do. For example, this is a colocation facility catering to the needs of multiple vendors? A bleeding edge research firm performing computational research? Or a company specific build to support their internal network and more routine operations?" asks Lycett. "In my experience, it seems that many

into a 'normal' load level zone and a more specific HPC environment for more intensive applications."

Taylor agrees that there isn't a single winner in liquid cooling: "the type implemented depends on the data centre's needs, workloads, and infrastructure. That said, as a result of Nvidia making direct-to-chip the standard cooling method for their GB200 NVL72 product, we are seeing a heavier adoption of single-phase direct-to-chip in AI data centres. Direct-to-chip provides a good balance of cooling capacity, ease of deployment, and compatibility with existing data centre



Bernie Malouin, JetCool

We empower IT everywhere

MODULAR • FLEXIBLE • FUTURE-PROOF

Choose Rittal for your IT infrastructure solutions, wherever the location of your requirements.



Explore our IT solutions





www.ittal.co.uk

infrastructure."

"At present, direct-to-chip holds the largest market share, particularly in hyperscale data centres, due to its direct adoption by major server suppliers and the well-established ecosystem supporting DLC deployment. However, it typically requires a larger upfront investment and more extensive infrastructure. This robust ecosystem, combined with strong industry standardisation,

infrastructure, server upgrade expectations, infrastructure adaptability, budget constraints, energy recovery feasibility and long-term operational efficiency," adds Thompson.

"For data centres running the latest and most advanced processors, like GPUs for AI, direct-to-chip (DTC) liquid cooling is leading the pack, with analysts highlighting its dominance as the preferred approach for AI

"By implementing a flexible infrastructure, operators can gradually shift workloads to liquid cooling as technology advances without overcommitting to a single method too soon. This approach ensures compatibility with both current and future IT hardware, prevents unnecessary CAPEX spend, and allows for a phased transition based on actual demand."

accelerates its adoption despite some material and thermal limitations that may arise with next-generation chips," notes Thompson.

Other liquid cooling solutions, such as immersion and liquid spray cooling, are more agile and scalable, allowing the cooling infrastructure to grow alongside the data centre.

"They cater to different customer needs such as existing cooling

workloads," shares Malouin. "While multiple liquid cooling technologies exist, including direct-to-chip, rear-door heat exchangers, and immersion cooling, the best choice depends on a data centre's specific needs, constraints, and long-term strategy."

Decisions, decisions

Implementing liquid cooling in data centres can be highly beneficial, but it requires careful planning and consideration.

"It's important to remember that the data centre industry is a mission critical industry, with downtime being exceptionally costly," highlights Beran. "At the same time, the cost of AI servers has skyrocketed, with costs up to \$3 million per rack of IT. This means that any failures should be considered a risk."

"Implementing the right solution is about more than just cooling — it requires a strategy that prioritises scalability, reliability, and long-term cost savings," notes Malouin. "One of the key considerations when adopting liquid cooling is scalability. Many providers struggle to scale beyond pilot projects due to limited manufacturing, service, and warranty capabilities."

Indeed, Taylor says that operators should look for technologies that meet current and future demands and integrate with

existing or planned facility designs: "they should also prioritise vendors with a strong track record in scalability, and a diverse network of supplier relationships, as the industry is likely to face supply chain constraints in the near term from increased demand."

"Data centre operators should seek scalable solutions that address server adaptability and provide the flexibility needed to grow with demand," agrees Thompson. "Ensuring a strong return on investment (ROI) is crucial, so prioritising a cooling solution with a low Power Usage Effectiveness (PUE) is essential. Additionally, the ability to recover waste heat for hot water and heating applications can transform the data centre into an energy producer, further enhancing its value and sustainability."

Meanwhile, Lycett suggests that it is vitally important to conduct a thorough assessment of the current installation and understand not just what you need now but look ahead as much as is reasonably possible.

"Additionally, it may be important to consider compatibility," adds Lycett. "If you are wishing to implement a new system, how well will this integrate with your existing architecture and data centre design?"

Carreiro warns of vendor lock-in with proprietary solutions: "many liquid cooling providers offer proprietary technologies that only work within their ecosystem. This can lead to limited compatibility with future hardware upgrades; dependence on a single vendor for servicing, spare parts, and expansion; higher long-term costs due to expensive proprietary components; and inflexibility in scaling as the business needs evolve."

"When evaluating liquid cooling solutions, operators must be cautious of systems that do not scale effectively or lack long-term service and warranty support. Many liquid cooling providers do not have the capability to support deployments at scale, leaving operators without reliable service or access to critical maintenance," warns Malouin. "Solutions that still rely heavily on air cooling should also be questioned, as they may fail to deliver the full efficiency and density benefits needed for AI and HPC workloads. Additionally, rigid, one-size-fits-all approaches can be problematic, particularly for legacy data centres that need a phased transition strategy rather than

a complete infrastructure overhaul."

According to Evans, as the liquid cooling industry is still in its early stages, a hybrid approach that combines both air and liquid cooling is the most practical solution, allowing data centre operators to adapt to changing client requirements while maintaining efficiency and cost control.

"By implementing a flexible infrastructure, operators can gradually shift workloads to liquid cooling as technology advances without overcommitting to a single method too soon. This approach ensures compatibility with both current and future IT hardware, prevents unnecessary CAPEX spend, and allows for a phased transition based on actual demand," says Evans. "A well-designed hybrid system integrates controllable air and liquid cooling loops, enabling operators to fine-tune cooling efficiency depending on workload density. By monitoring real-time IT loads and heat distribution, cooling strategies can be dynamically adjusted to optimise performance. Hybrid cooling not only offers resilience but also positions data centres for long-term sustainability. It reduces reliance on energy-intensive air cooling, while taking advantage of liquid cooling's superior heat dissipation as adoption scales. This balanced approach ensures that data centres remain agile, efficient, and ready to support evolving client needs." ■



Chris Carreiro, Park Place Technologies



Karl Lycett, Rittal

NEXT-GEN SOLUTIONS FOR FUTURE-READY DATA CENTRES

RAPIDNET®
ULTRA

Scan for
more info

