



DIRECT-TO-CHIP LIQUID COOLING TECHNOLOGY WHITE PAPER

Single-Phase Direct-to-Chip Liquid Cooling Microconvective vs. Microchannel Liquid-Cooled Cold Plates

This white paper synthesizes insights from industry research to highlight the adoption trajectory and technological advancements in liquid cooling, focusing on how single-phase, direct-to-chip liquid cooling is leading adoption, addressing the challenges posed by accelerated computing and AI-driven workloads.

Table of Contents

- 03 Abstract
- 05 Introduction: Striking the Right Cooling, Power, and Efficiency Balance
- 06 Direct-to-Chip Liquid Cooling Opportunities: Barriers and Promise
- 09 Types of Direct-to-Chip Liquid Cooling
- 10 JetCool: Microconvective Direct-to-Chip Liquid Cooling
- 11 Microconvective Liquid Cooling Advantages
- 12 Final Thoughts

Abstract

This white paper analyzes findings from the latest Dell'Oro Group's Data Center Physical Infrastructure (DCPI) forecast and the Uptime Institute Cooling Systems Survey, exploring the current state of direct liquid cooling (DLC) in data centers. As data center operators face increasing challenges related to cooling performance, energy efficiency, and sustainability, DLC emerges as a promising alternative to air cooling. However, the widespread adoption of DLC is hindered by barriers such as infrastructure, standardization, and cost. This white paper examines the survey findings and highlights the potential of DLC in addressing the cooling requirements of higher density racks and power-hungry chips. Additionally, it introduces JetCool, a leading direct-to-chip liquid cooling manufacturer that outperforms incumbent microchannel liquid cooling solutions with 3X better thermal resistance, enabling efficient cooling of higher power processing.

Disclaimer: This white paper is based on the Uptime Institute Cooling Systems Survey 2023, which collected insights from over 800 respondents, and reflects the independent analysis and opinions of the authors. The information provided in this white paper is for informational purposes only and does not constitute investment or technical advice.



As data center operators face increasing challenges related to cooling performance, energy efficiency, and sustainability, DLC emerges as a promising alternative to air cooling

“

Most data center operators expect air cooling to cede its position as the dominant method in cooling IT hardware within six years.

UPTIME INSTITUTE

Introduction

Striking the Right Power, Cooling and Efficiency Balance

Server power and compute densification drive cooling innovation.

According to the 2023 Uptime Institute Cooling System Survey, direct liquid cooling is predicted to surpass air cooling as the primary method for cooling IT infrastructure within the next six years. As server power continues to escalate and high-density devices integrate into enterprise applications, direct liquid cooling emerges as a promising alternative to cool next-generation devices.

The logo for Uptime Institute, featuring the word "uptime" in a lowercase, sans-serif font with a registered trademark symbol, and the word "INSTITUTE" in a smaller, uppercase, sans-serif font below it, all in white on a dark blue square background.

For more than 25 years, the Uptime Institute has served as the leading authority on Global Digital Infrastructure. Its Tier Standard is universally acknowledged by the IT industry as the global standard for designing, constructing, and operating data centers – the vital foundation of the digital economy.

The introduction of higher power processors in IT architectures, systems, and networks has triggered a swift rise in rack densities and power requirements within data centers. Combined with the emergence of artificial intelligence (AI), advancements in server technology are poised to disrupt the entire data center landscape, prompting enterprises to reassess their existing cooling infrastructure.

Preparing for these density shifts, data center operators are faced with the challenge of striking the right balance between cooling, power, and energy efficiency within compact space constraints. Until recently, conventional air-cooling methods met industry-standard power and spatial requirements, effectively providing efficient cooling solutions. Stagnant power and space parameters facilitated maturation and widespread adoption of air cooling, awarding it as the predominant cooling method in today's data centers.

Nonetheless, escalating server power and compute densification are straining this delicate balance of IT infrastructure, driving data center operators to explore liquid cooling alternatives to cool high-density devices. In response to these evolving needs, direct-to-chip liquid cooling emerges as a viable alternative to traditional air cooling, empowering data centers to support the latest high-density chipsets while upholding performance, energy efficiency and spatial requirements.



Direct-to-Chip Liquid Cooling Opportunities: Barriers and Promise

Market Trends and Drivers

As artificial intelligence becomes increasingly prevalent, the adoption of direct liquid cooling solutions is anticipated to rise to meet the escalating demands. According to Markets and Markets, the data center liquid cooling market is predicted to grow at a 24.4% CAGR from 2023 to 2028, reaching \$7.8 billion. This expansion is primarily attributed to the efficiency of liquid cooling, which can boost server density, decrease energy usage, and lower Power Usage Effectiveness (PUE). As chip density continues to rise, data center operators will require efficient energy consumption to manage thermal challenges effectively.

Market Driver #1: Server Power and Density


Chip manufacturers are racing to enhance processor capacity for AI workloads, leading to an escalation in processor Thermal Design Power (TDP). NVIDIA, a frontrunner in the AI chip market, has released its latest GPUs with a maximum TDP of 700W, underscoring the importance of liquid cooling in enabling compute densification. This trend is supported by projections from the Dell'Oro Group's analysis, which forecasts an 11% CAGR in the Data Center Physical Infrastructure (DCPI) market from 2023 to 2028, surpassing \$46 billion. A significant driver of this growth is the adoption of liquid cooling technologies, spurred by the necessity to accommodate increased rack power densities, particularly with AI and ML workloads expected to range between 60 and 120kW per rack.

Data centers that refresh their servers with air-cooled infrastructure will begin to incur penalties such as reduced peak performance levels, decreased density, and increased energy losses due to server fans and leakage currents. This penalty is already starting to present itself as 68% of respondents reported higher rack densities as the primary driver for direct liquid cooling adoption, according to Uptime Institute. Server power follows as a critical concern. With the introduction of power-hungry CPUs and GPUs, data center operators are recognizing the limitations of air cooling in meeting power and density requirements.

2023 UPTIME INSTITUTE COOLING SYSTEMS SURVEY

FIGURE 1

Which of these do you think are the primary drivers for direct liquid cooling adoption? Choose no more than two. (n=740)



(Responses for "Other" are not included.)

Market Driver #2: Energy Efficiency

Amid the pursuit of sustainable computing, data centers are under increasing scrutiny from governments and legislators to minimize their carbon footprints. According to McKinsey & Company, cooling infrastructure alone constitutes approximately 40% of energy consumption in these facilities, highlighting the urgent need for innovative solutions to promote eco-friendly practices. Effective liquid cooling systems offer a promising avenue to optimize data center efficiency and reduce energy consumption. Notably, 80% of respondents emphasize that energy efficiency plays a pivotal role in determining the sustainability of a cooling

system. By decreasing annual PUE, enterprises can not only cut down on power costs but also align with their sustainability objectives. Suffice it to say, improving data center energy efficiency delivers a multitude of benefits to the entire IT ecosystem - and, more importantly, the planet.

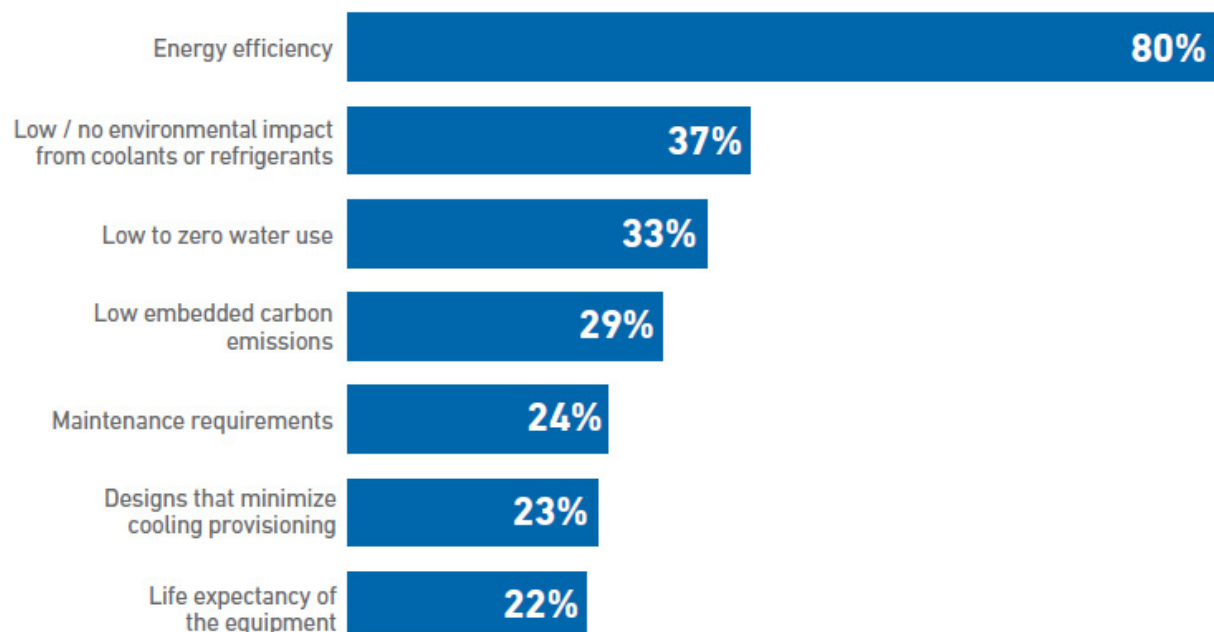
KEY TAKEAWAYS

- Sustainability efforts accelerate liquid cooling adoption
- Compute and rack density exacerbate air cooling methods

2023 UPTIME INSTITUTE COOLING SYSTEMS SURVEY

FIGURE 2

Which of these do you consider to be the most important factors in determining whether a cooling system is sustainable? Choose no more than two. (n=825)



(Responses for "Don't know" are not included.)

UPTIME INSTITUTE COOLING SYSTEMS SURVEY 2023

uptime
INTELLIGENCE

80%

Energy Efficiency

Cooling systems inherently consume a substantial amount of energy. By enhancing energy efficiency, businesses can effectively reduce operational and energy costs as a whole. Consequently, operators prioritize energy efficiency as a critical factor in ensuring the sustainability of data centers.

37%

Standardization

Enterprises seek reliable and standardized liquid cooling solutions that seamlessly integrate into their existing infrastructure, eliminating the need for expensive facility overhauls.

68%

Rack Density

More than half of participants deemed rack density for the primary driver for liquid cooling adoption. Compute densification and the escalation of server power is pushing air cooling to its limit - making way for liquid cooling alternatives.

Barriers to Adoption

Single-phase direct liquid cooling technologies are the most commonly deployed liquid cooling solutions on the market; however, its adoption presents obstacles, including infrastructure compatibility, standardization, and cost concerns. Building on those challenges, previous rack densities and power demands weren't high enough to justify implementation. However, as rack densities rise, these barriers are quickly being overcome as the industry recognizes the technology's ability to meet the thermal management demands of increasingly powerful data centers.

Ease of Retrofitting

For facilities that didn't consider liquid cooling in their initial design, retrofitting or upgrading to any liquid cooling infrastructure can pose a significant financial and labor-intensive challenge. Retrofitting may also result in server downtime during implementation. Consequently, liquid cooling deployments must provide a scalable and simple solution to justify the initial investment and operational disruptions. This approach should minimize downtime, leverage existing infrastructure where possible, and offer modularity to adapt to the evolving needs of the data center.

Lower Capital Expenditures

In addition to addressing retrofitting challenges, it is crucial for direct-to-chip liquid cooling to have affordable initial costs for implementation. Since data centers can have a lifespan of up to 15 years, it is vital to find a solution that can future-proof and grow with the data centers' needs of today and tomorrow.

Ease of Maintenance

After implementing direct liquid cooling, the task of system maintenance arises. Certain liquid cooling systems may necessitate the expertise of cooling specialists for proper upkeep. It is imperative that maintenance procedures remain straightforward and adaptable, enabling any operator to carry out routine tasks as needed.

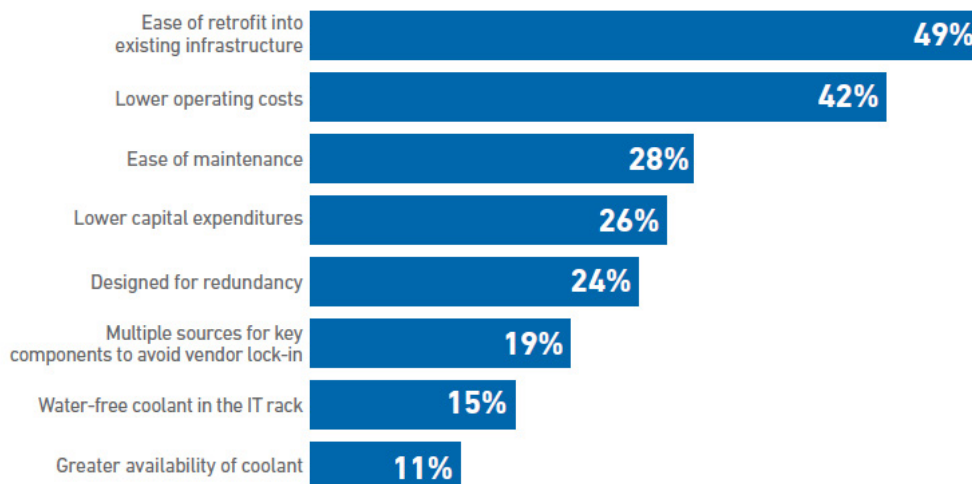
Lower Operating Costs

Transitioning to liquid cooling can entail significant expenses, such as upgrades or new facilities. Liquid cooling solutions must justify adoption by balancing initial expenses with reduced long-term operational costs to ensure viability. Assessing how this upgrade will lead to decreased overall operating expenses, rather than an increase, is crucial when moving from legacy data center infrastructure.

2023 UPTIME INSTITUTE COOLING SYSTEMS SURVEY

FIGURE 3

Which of these factors do you consider most important to determine whether a direct liquid cooling system is viable? Choose no more than two. (n=778)



(Responses for "Don't know" and "Other" are not included.)

Types of Direct-to-Chip Liquid Cooling

Microchannel vs. Microconvective

The direct liquid cooling evolution has largely revolved around two primary technologies –microchannel and microconvective liquid cooling. The main difference between these two cooling methods is how the coolant makes contact with the heated surface or cold plate.


Microchannel Liquid Cooling

Microchannel liquid cooling uses a heat-spreading premise with uniform, parallel cooling (See Graphic 1). This is achieved by utilizing small internal fluid channels that allow for the heat to be removed by a flowing coolant. The use of this heat spreading technique ensures a large cooling surface area, while the small fluid channels facilitate close interaction between the coolant and the heated surface. In microchannel cooling, the coolant indiscriminately cools the entire surface area of the cold plate. In traditional cooling methods, such as air cooling or even microchannel liquid cooling, heat is typically transferred through conduction and convection. However, these methods can be limited in their ability to efficiently handle the high heat loads generated by modern, densely packed electronic components.


Microconvective Liquid Cooling

Microconvective liquid cooling, also known as microjet impingement cooling, uses targeted, perpendicular cooling that pinpoints hot spots on the processor (See Graphic 2). Single-phase jet impingement generates high convective coefficients by way of high-momentum fluid jets impinging perpendicularly to a hot surface. With jet impingement, fluid delivery can be concentrated in proximity to hot spots generated from the semiconductor dies, thereby mitigating the requirement of heat spreading or area enhancements for high-power density heat removal. Further, in certain applications, jet impingement can be implemented with no thermal interface material in a direct-to-die or direct-to-package configuration with no change to the chip assembly procedure. The use of microjets allows for more efficient heat transfer due to increased surface contact and improved convective heat transfer coefficients.

MICROCHANNEL VS. MICROCONVECTIVE



GRAPHIC 1: MICROCHANNEL LIQUID COOLING ARCHITECTURE



GRAPHIC 2: MICROCONVECTIVE LIQUID COOLING ARCHITECTURE

JetCool: Microconvective Direct-to-Chip Liquid Cooling

JetCool is at the forefront of advancing microjet impingement cooling technology, using its patented microconvective cooling approach. Because of this targeted approach, microconvective cooling significantly lowers thermal resistance—three to five times that of the state-of-the-art microchannel cold plate technology—thereby efficiently managing the thermal loads of high-power processors. This introduction explores the key benefits of JetCool's solution, tailored to meet the evolving cooling requirements of data centers equipped with high-density server racks.

37%

Cooler Operation

Lower chip temperature by 37% compared to other competing solutions

30%

Faster Compute

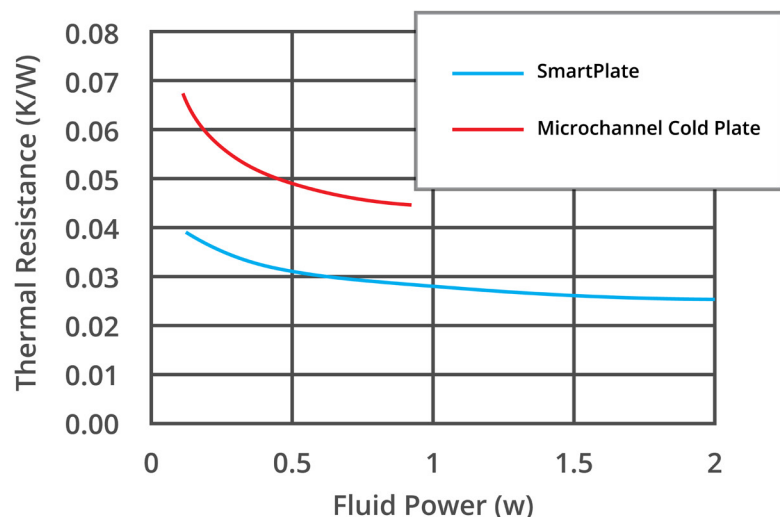
Extract 30% faster processing and support TDPs over 1000W

40%

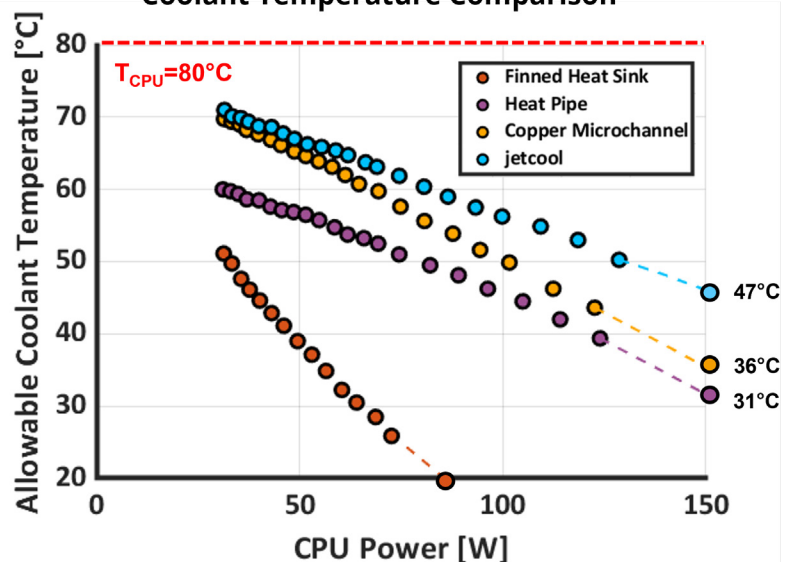
Lower Thermal Resistance

Reduce thermal resistance by 40% compared to leading microchannel cold plates

Microconvective vs. Microchannel Cold Plates Performance on Intel Sapphire Rapids Processors



Coolant Temperature Comparison³




Microconvective Liquid Cooling Advantages



Zero Water Usage & Maximize ESG Efforts

Microconvective cooling technology obviates the need for chillers, cooling towers, and evaporative coolers, eliminating water consumption and an 18% decrease in energy consumption. This innovative solution also helps to prevent the emission of 35 million metric tons of CO₂ annually.



Enhanced Heat Transfer

The direct impingement of high-velocity liquid jets onto the components' surfaces enables efficient and rapid heat transfer. This method enables up to 20% total energy savings in data centers.



Improved Thermal Performance

Microconvective cooling supports over 1,500W TDPs, ensuring electronic components operate within optimal temperature range. This enhances reliability, reduces thermal throttling, and can extend server lifespan.



More Efficient Data Centers

Using inlet coolant temperatures up to 65°C allows data center operators to obviate evaporative coolers, yielding significant power savings and eliminating water consumption.



Final Thoughts

Direct liquid cooling is the market leader for liquid cooling technologies, supporting higher rack densities and server power.

In conclusion, single-phase, direct-to-chip liquid cooling has emerged as the market leader in the liquid cooling sector, driven by its efficiency, reliability, and adaptability. This technology, specifically microconvective liquid cooling, stands out for its ability to directly address the thermal management challenges of modern data centers, providing precise cooling directly to the components that generate the most heat.

By offering significant improvements in energy efficiency, reducing operational costs, and enabling high-density computing configurations without compromising performance, single-phase, direct-to-chip liquid cooling has set a new standard in the industry. Its leadership in the market is a testament to the growing recognition among facilities of the need for innovative cooling solutions that can keep pace with the increasing power demands of contemporary computing environments. This technology not only leads the way in liquid cooling but also paves the path for future advancements in data center design and operation, making it a pivotal solution for the evolving landscape of the digital world.





hello@jetcool.com



jetcool.com



+1 (978) 449-4600