
Supermicro Data Center Liquid Cooling Solution

STEVE MCDOWELL, CHIEF ANALYST
10/7/24

CONTEXT

Supermicro [recently announced](#) a comprehensive, end-to-end liquid cooling solution for data centers. The solution encompasses critical hardware components such as Coolant Distribution Units (CDUs), cold plates, Coolant Distribution Manifolds (CDMs), cooling towers, and integrated management software.

The new solution targets the increasing power demands and cooling challenges posed by high-performance AI and HPC workloads, particularly in large-scale AI factories and CSPs.

TECHNICAL OVERVIEW

Supermicro's new cooling solution enhances cooling efficiency for data centers running dense configurations of high-performance GPUs and CPUs, reducing power consumption and operational complexity.

Here's a look at what the company is delivering.

1. Key Hardware Components:

- **Cold Plates:** Specifically designed for optimal heat dissipation, these cold plates allow liquid to pass through microchannels, maximizing surface area to dissipate up to 1600W of heat generated by next-generation GPUs.
- **Coolant Distribution Manifolds (CDMs):** Configured to enable high-density GPU installations, supporting up to 96 NVIDIA B200 GPUs per rack. Horizontal and vertical CDM designs optimize the use of physical space, ensuring efficient cooling in dense server racks.

- **Coolant Distribution Units (CDUs):** The CDUs support cooling capacities of up to 250kW. Features include hot-swappable pumps and power supplies, allowing continuous operation and minimizing downtime during maintenance.
- **Cooling Towers:** Modular cooling towers are integrated into the system, employing energy-efficient EC fan technology. These towers can be quickly shipped and deployed, shortening the time-to-online (TTO) for new installations or upgrades.

2. Software Integration:

- **SuperCloud Composer:** The management software oversees the entire lifecycle of the liquid cooling system. It monitors the status of components such as CDUs, racks, and cooling towers, optimizing operational costs by adjusting cooling parameters in real-time based on system loads.

3. Power and Space Efficiency:

- Direct Liquid Cooling (DLC) offers significant efficiency improvements over traditional air-cooling systems:
 - Up to 40% reduction in power consumption for cooling infrastructure.
 - There will be an 80% reduction in space, eliminating the need for conventional Computer Room Air Conditioning (CRAC) and Computer Room Air Handler (CRAH) units.
 - Support for warm-water cooling at temperatures up to 113°F (45°C) enables more efficient resource use and the potential for heat reuse in secondary applications like district heating or greenhouse energy.

4. Deployment Scalability:

- The system supports ultra-dense server configurations, such as a 4U server with dual CPUs and 8 NVIDIA HGX GPUs, increasing compute density by 4x per rack.
- Supermicro reports deploying over 100,000 GPUs using its DLC technology for large-scale AI factories and CSP data centers. Servers equipped with DLC can handle power loads approaching 12kW per unit, with individual AI racks generating over 100kW of heat. DLC is crucial in efficiently managing these thermal loads.

Supermicro's liquid cooling solution will impact both the technical and financial aspects of AI infrastructure deployment:

- **Technical Impact:** The high-density configurations and efficient heat dissipation mechanisms will improve overall system performance, allowing for larger and more complex AI model training with reduced hardware footprints. This will help AI-driven enterprises and CSPs scale more efficiently.
- **Financial Impact:** The combined reduction in power consumption, space requirements, and faster deployment timelines will help enterprises lower capital expenditures (CAPEX) and operational expenditures (OPEX).

COMPETITIVE POSITION

Among AI infrastructure providers, only Lenovo's Neptune liquid cooling solution is comparable to what Supermicro announced. There are significant differences, with Supermicro delivering a highly targeted cooling solution for big AI workloads, while Lenovo addresses a broader range of use cases with its Neptune technology.

These solutions will ultimately be judged on their ability to reduce energy usage while increasing data center density. On this front, Lenovo and Supermicro are similar.

Supermicro's liquid cooling solution claims up to 40% reduction in energy consumption for cooling infrastructure and 80% space savings by eliminating the need for CRAC/CRAH units. The system's design helps achieve PUE values of less than 1.1, making it one of the most energy-efficient offerings on the market.

On the other hand, Lenovo Neptune achieves a 30-40% reduction in energy costs associated with cooling. By employing warm water cooling, it eliminates the need for chillers, which significantly reduces energy overhead. Neptune's RDHX units can achieve similar PUE values, especially when deployed in hybrid cooling environments where liquid and air are managed together for different components.

Neptune also focuses on reducing the energy consumed by cooling, but its energy-saving claims align closely with Supermicro's in terms of percentage reduction.

While Dell and HPE both offer liquid cooling solutions, those currently appear to be point products tied to specific AI server offerings and don't broadly address datacenter-wide concerns.

ANALYSIS

Supermicro's liquid cooling solution improves its already strong position in the highly competitive AI infrastructure market. Its new solution will increase its appeal to hyperscalers and large CSPs, potentially drawing business away from traditional data center solution providers that have slower deployment times or less efficient cooling solutions.

The ability to support dense AI training models with fewer physical racks and a smaller data center footprint directly benefits organizations running LLMs or other compute-intensive AI applications. Supermicro's experience with large AI factories (the company is the 'go-to' infrastructure provider for many specialty cloud providers) enhances its credibility as a high-performance data center infrastructure leader.

Supermicro's comprehensive liquid cooling solution is a strong differentiator. The new offering will place Supermicro ahead of competitors in the growing market for high-performance, energy-efficient infrastructure tailored to AI, HPC, and cloud environments.



© Copyright NAND Research.

NAND Research is a registered trademark of NAND Research LLC, All Rights Reserved.

This document may not be reproduced, distributed, or modified, in physical or electronic form, without the express written consent of NAND Research. Questions about licensing or use of this document should be directed to info@nand-research.com.

The information contained within this document was believed by NAND Research to be reliable and is provided for informational purposes only. The content may contain technical inaccuracies, omissions, or typographical errors. This document reflects the opinions of NAND Research, which is subject to change. NAND Research does not warranty or otherwise guarantee the accuracy of the information contained within.

NAND Research is a technology-focused industry analyst firm providing research, customer content, market and competitive intelligence, and custom deliverables to technology vendors, investors, and end-customer IT organizations.

Contact NAND Research via email at info@nand-research.com or visit our website at nand-research.com.