

Immersion Cooling for AI Servers



Overview

In two-phase immersion cooling, a server is dunked into a vat of liquid. The liquid actively boils next to the heat-producing components, cooling them in the process.

Key Takeaways

Immersion cooling submerges AI servers in non-conductive fluids, offering superior heat dissipation, energy efficiency, and high-density deployment compared to air or cold plate cooling.

Overview of Immersion Cooling

Immersion cooling involves submerging servers directly in dielectric, non-conductive fluids that absorb heat from all components, including CPUs, GPUs, and memory modules, without the need for air circulation. This method eliminates thermal resistance from heat spreaders, cold plates, and airflow, allowing AI servers to operate at peak performance for extended periods without throttling. Single-phase immersion keeps the fluid in liquid form throughout operation, while two-phase immersion allows the fluid to boil and condense, offering even higher thermal efficiency but with increased system complexity.

Benefits for AI Workloads

High Thermal Efficiency: Immersion cooling captures nearly 100% of the heat generated by high-TDP AI accelerators, such as NVIDIA Blackwell GB200/GB300 GPUs or TPUs, maintaining stable temperatures under sustained high loads. **Energy Savings:** By removing the need for traditional air-cooling infrastructure like CRAC units and raised floors, immersion cooling reduces overall energy consumption and can achieve PUE values as low as 1.02-1.1.

05.

High-

Density Deployment: Servers can be packed more densely per rack without overheating risks, enabling more compute power per square meter. **Sustainability and Heat Reuse:** The captured heat can be repurposed for district heating or other applications, supporting decarbonization and energy efficiency goals. **Deployment Considerations**
Plug-and-Play Systems: Platforms like GRC's ICEraQ allow rapid deployment of immersion-cooled racks with minimal infrastructure, requiring only power, water service, and a level floor.

Hybrid Environments: Immersion cooling can coexist with air-cooled equipment, enabling gradual upgrades or mixed cooling strategies in existing data centers.

Facility Requirements: While immersion cooling does not require fundamental changes to water systems, data centers must support fluid circulation, heat exchangers, and proper containment.

Comparison with Other Liquid Cooling Methods
Cold Plate (Direct-to-Chip) Cooling: Uses metal plates to transfer heat from components to circulating coolant without direct contact.

It is widely adopted and mature but less efficient than full immersion for extreme TDP workloads. **Immersion Cooling:** Provides complete direct contact heat removal, eliminating the need for heat spreaders and airflow, achieving the highest thermal efficiency and energy savings.

Current Trends

By 2026, immersion cooling is increasingly adopted in AI-focused data centers, particularly for high-power GPUs and AI ASICs, as traditional air cooling cannot handle the thermal load of modern AI servers with TDPs exceeding 1,000W. Leading solutions focus on single-phase immersion for simplicity and reliability, while two-phase systems remain experimental in hyperscale deployments. Immersion cooling is now considered a strategic solution for AI data centers, enabling high-density, energy-efficient, and sustainable operations while supporting the latest high-performance AI hardware.

Article Content

Microsoft Unveils Immersion Cooling to Solve AI Heat Bottleneck | The ...

Microsoft's new system uses a two-phase immersion cooling approach where server components are submerged in a special fluid that boils at a low temperature, carrying heat away from

Data Center Liquid Cooling: The AI Heat Solution

In two-phase immersion cooling, a server is dunked into a vat of liquid. The liquid actively boils next to the heat-producing components, cooling them in the process.

AI Server Liquid Cooling: Cold Plates & Immersion

As AI server TDP surges past 1000W, traditional air cooling fails. Explore cold plate and immersion liquid cooling technologies. Get a Free DFM from Ecotherm.

Data Center Liquid Cooling for AI Workloads (2026) · KAD

An engineering-focused overview of cold plate and immersion liquid cooling in 2026, driven by 1,000W+ AI accelerators like NVIDIA Blackwell and Rubin.

Immersion cooling systems: Advantages and

Immersion cooling absorbs 100% of rack heat — but room cooling remains essential for air quality and humidity management. Rack densities are

Liquid, Immersion, Hybrid: The Cooling Architectures

AI workloads are pushing data center cooling beyond air. Explore why liquid, immersion, and hybrid cooling architectures are emerging as the

Immersion Cooling Data Centers | Hyperscale Builds 2025

Learn why immersion cooling is moving mainstream in hyperscale data centers. Explore benefits, adoption trends, and how AI drives high-density cooling demand.

Optimizing AI Performance with Immersion Cooled Data Centers

The Cisco-GRC solution reduces environmental impact through innovative immersion cooling and optimized infrastructure. By replacing traditional air-cooling with efficient immersion technology,

Immersion Cooling Solution for Sustainable AI Data Centers □NADDOD

As traditional cooling methods incurs high costs in energy water, land use, and carbon footprint, innovative data centers are increasingly adopting immersion cooling—a solution that is more

Cold Plates vs Immersion in AI Data Centers

Cold plates and immersion cooling offer two fundamentally different solutions to these thermal challenges. Engineers deploy cold plates to extend familiar air-assisted liquid systems, while

Rack-Scale Liquid Cooling Solutions | Supermicro

Dive into Liquid Cooling Server solutions, providing efficient thermal management. Ideal for high-performance computing, ensuring optimal operation.

DCX Liquid Cooling Systems

Server Immersion Systems Game changer for the AI driven Datacenter Industry. First Standard Server Immersion Enclosure System dedicated for 19" servers

Direct-to-Chip vs Immersion Cooling

Reviewed by Yoan Guyon, Managing Director at gbc engineers AI and HPC servers have pushed rack power from 5 to 15 kW to 50 kW, 100 kW or more. At those densities, traditional air

Dedicated GPU/AI Servers | 2CRSi

GPU servers can be used in many types of work field, mostly ones who deal with graphic-heavy content, AI learning and inference jobs. Thanks to its stable and fast computing, a GPU server is essentially

Liquid Cooling: The 2026 Mandate for 100kW AI Racks

In 2026, liquid cooling is mandatory for 100kW AI data centers. Learn how key partnerships are solving the power crisis and shaping the future of cooling. Discover more.

LiquidCool Solutions: Immersion-Cooled Rack Servers

Meet soaring compute demands efficiently with LiquidCool Solutions" versatile, immersion cooling for high-performance GPU dense servers. Perfect for data centers, AI, edge computing, and tough

Re-Architecting the AI Server: The Hidden Water Cost of Data Centers ...

As AI data centers adopt liquid cooling, freshwater use is surging—raising environmental justice concerns and straining communities.

Immersion Cooling for Data Centers: What Developers Need to ...

Immersion cooling moves heat from chips to fluid rather than air, enabling power densities that conventional air cooling cannot reach. This guide explains single-phase versus two

Related Video Reference

This video was associated with the source search result. Verify technical details against current product documentation and project requirements.

Contact Us

Continue your research online:

Website: <https://cjccconsultants.co.za>

Technical inquiry: <https://cjccconsultants.co.za/contact.html>

This document is for preliminary research. Verify specifications against current standards, product documentation and project requirements.

