

Are liquid-cooled servers considered AI hardware

Preview

As accelerator thermal design power (TDP) pushes beyond 1,000W in the era of NVIDIA Blackwell and Rubin-class architectures, traditional air cooling has reached its practical and physical limits. Liquid cooling is no longer optional--it is now a baseline requirement for modern AI data. Servers used to train AI models can consume more than 80 kilowatts per rack, and Nvidia's latest GB200 chip, combined with its servers, can require densities of up to 120 kilowatts, according to data from McKinsey. It offers up to 15% better energy efficiency and reduces cooling costs compared to traditional air-cooling systems. The technology also enables higher server density. While liquid cooling is rightly considered an emerging technology, it's not new. Notably, a full-time technician was included in the Cray system purchase for installation, operation and maintenance. Today, many AI servers with accelerators (e., GPUs) used for training LLMs (large language models) and inference workloads, generate enough heat to necessitate liquid cooling. Unlike air, liquid absorbs and transfers heat far more effectively.

AI workloads and HPC require efficient cooling. Discover how immersion and liquid cooling save energy, enhance performance, and future-proof data centers.

Modern AI servers are essentially clustered GPU/accelerator nodes, often with specialized cooling (liquid or advanced air). For example, Dell's new systems

Liquid cooling Liquid cooling comes in various forms, but it's important to understand that liquid cooling is not a single product. It is a system and an ecosystem comprising various components such as

AI demand is causing enterprise hardware revenue to grow, even as PC sales may weaken due to rising memory prices. Data centers are reinventing themselves

But if you're looking to deploy your HPC infrastructure in a facility owned by a data center provider, you should look at what the provider offers in terms of liquid

Future AI infrastructure using the latest accelerators will require this same liquid cooling innovation to address concerns in power efficiency,

The next generation of AI servers pushes the bounds of computational power at the cost of increasing power consumption, requiring the use of liquid cooling. Liquid-cooled servers will need to work

Liquid cooling has become a critical enabler for modern AI data centers as facilities scale to handle

Are liquid-cooled servers considered AI hardware

high-density workloads, such as artificial

Liquid cooling is essential for AI data centers--but deploying it at scale requires advanced design, smart vendors and AI-powered operations to

Google, Microsoft, Meta, Amazon and Alibaba are all investing heavily in liquid-cooled data centers to support AI services. Intel and NVIDIA are

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

Data centers have always relied on air-cooled servers for computing power, but the acceleration of AI deployments is now causing a trend toward

Amazon isn't new to developing its own hardware. The company has created custom chips for general-purpose and AI-specific computing, as well as

Build a future-proof home AI server in 2026 with the latest hardware and software stack. Learn optimal GPU, CPU, and storage choices for AI workloads, plus integration with NVIDIA AI

Liquid cooling is essential for modern AI data centers because it efficiently manages the immense heat from powerful processors. Unlike air, liquid absorbs and

Liquid cooling is essential for AI-driven data centres, efficiently managing the extreme heat generated by high-density AI server racks. It offers

Web: <https://bssoda.pl>

