AMAX GB200 NVL72 AI POD Solution with Liquid-to-Air Cooling
LiquidMax® ALC-B4872

Bringing the New Era of Computing to Every Data Center

AMAX LiquidMax® GB200 NVL72 AI POD with Liquid-to-Air (L2A) Cooling Rack Configuration

NVIDIA GB200 NVL72 AI POD

AMAX's LiquidMax® ALC-B4872 GB200 NVL72 AI POD is pre-configured with 36 Grace Blackwell Superchips, which include 72 NVIDIA Blackwell GPUs and 36 Grace CPUs interconnected by fifth-generation NVLink. NVIDIA Blackwell GPUs feature two reticle-limited dies connected by a 10 terabytes per second (TB/s) chip-to-chip interconnect in a unified single GPU. This platform brings the new era of computing to every data center, delivering unparalleled performance for mainstream large language model (LLM) inference, retrieval-augmented generation (RAG), and data processing. With its scale-out, single-node NVIDIA MGX architecture, AMAX's LiquidMax® ALC-B4872 GB200 NVL72 AI POD enables a wide variety of system designs and networking options to seamlessly integrate into existing data center infrastructure.

In addition, the LiquidMax® ALC-B4872 GB200 NVL72 AI POD includes NVIDIA BlueField®-3 data processing units to enable cloud network acceleration, composable storage, zero-trust security and GPU compute elasticity in hyperscale AI clouds. The LiquidMax® ALC-B4872 GB200 NVL72 AI POD provides up to a 30x performance increase compared to the same number of NVIDIA H100 Tensor Core GPUs for LLM inference workloads, and reduces cost and energy consumption by up to 25x.

Supercharging Next-Generation AI and Accelerated Computing

LLM Inference

30X vs. NVIDIA H100 Tensor Core GPU

LLM Training

4X vs. H100

Energy Efficiency

25X vs. H100

Data Processing

18X vs. CPU

LiquidMax® ALC-B4872 GB200 Liquid-to-Air Solution Key Features

Per Rack:
- 18 x Compute Tray + 9 x Switch Tray
- Liquid cooling

Configuration per Tray:
- 4 x NVIDIA Blackwell GPU + 2 x Grace CPU
- 960GB LPDDR5X + 288GB HBM3e
- 2 x NVIDIA ConnectX®-7 NIC (dual-port 400GbE) +
- 2 x BlueField®-3 DPU (400GbE) NICs

Target Applications:
Datacenter AI Training, LLM Training.

AMAX Liquid-to-Air (L2A) CDU Key Features

	CDU Option 1	CDU Option 2
Racks Managed	Up to 2 NVIDIA GB200 NVL72 racks	Up to 4 NVIDIA GB200 NVL72 racks
Facility Liquid Integration	No	No

AMAX's LiquidMax® ALC-B4872 GB200 NVL72 AI POD with Liquid-to-Air (L2A) Cooling provides an efficient thermal management solution for high-performance data centers. The system circulates liquid coolant to absorb heat from high-performance components within the rack.

This heat is then transferred to the air via a sidecar cooling unit, which expels the hot air into the Computer Room Air Conditioner (CRAC) for final dissipation. This efficient design provides a cost-effective and scalable cooling solution for modern data centers.

Liquid to Air Solution

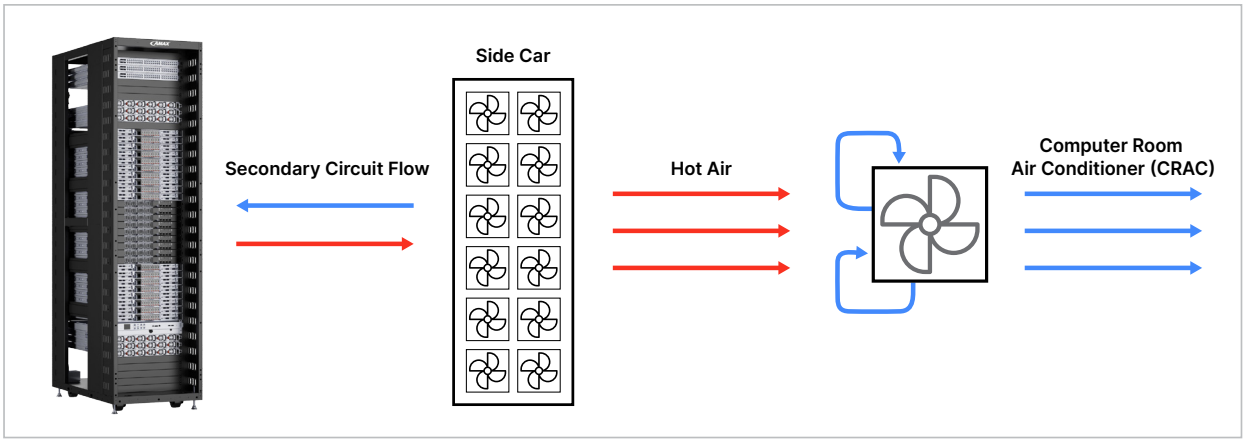

AMAX's L2A CDU is a high-density, airflow-to-power optimized cooling solution, capable of delivering up to 240 kW of cooling capacity within a 2-rack footprint for multiple artificial intelligence (AI), high performance computing (HPC) and high-density enterprise direct liquid-cooled (DLC) servers without requiring facility water.