



## Liquid-Cooled Rack Solution

# AMAX ENGINEERING

With 40 years of engineering expertise, our team specializes in transforming standard IT components into high-performance computing solutions with optimized thermal, electrical, mechanical, and networking design.

## LiquidMax<sup>®</sup> RackScale 64

The LiquidMax<sup>®</sup> RackScale 64 is a high-density, fully liquid-cooled 42U rack solution designed to support advanced AI training and inference. Built to house up to 64x NVIDIA<sup>®</sup> Blackwell GPUs, this system delivers maximum performance and efficient thermal management in a compact footprint.

### Key Features

- Supports up to 64x NVIDIA Blackwell GPUs
- 8x 4U liquid-cooled servers with NVIDIA NVLink™ interconnect
- Integrated CDU and vertical cooling distribution manifold
- Rack-ready system with pre-integrated compute, cooling, and power

### Built for High-Density AI Infrastructure

The RackScale 64 is ideal for data centers and research environments requiring reliable, scalable infrastructure for next-generation AI workloads. The liquid-cooled architecture enables sustained performance during peak compute cycles while reducing energy waste. With 1.4TB of GPU memory per node and total rack capacity of 11.5TB HBM3e, it is engineered for large-scale model development, inference pipelines, and other GPU-intensive applications.



AMAX // DATASHEET

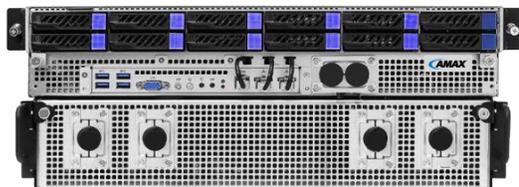
# Specifications

LiquidMax® RackScale 64	
GPU	64x NVIDIA Blackwell GPUs
Cooling	Direct liquid cooling, CDU and manifold included
Architecture	8x 4U compute servers within 44U rack
Power	Integrated power distribution, CDU with dual hot-swap pumps
Memory	11.5TB total HBM3e memory
Networking	NVIDIA Quantum-2 400G InfiniBand or Spectrum-4 400GbE for compute and storage; 10G/1G Out-of-Band management

AceleMax® AXG-428IL	
Form Factor	4U Rackmount
GPU	NVIDIA Blackwell 8-GPU system with NVIDIA NVSwitch™
Topology	Synthetic mode for optimized performance of GPU-to-CPU, GPUDirect RDMA NIC, and GPU Direct Storage
CPU	<ul style="list-style-type: none"><li>• Supports 5th and 4th Gen Intel® Xeon® Scalable Processors</li><li>• Socket: Dual Socket E (LGA 4677)</li></ul>
Chipset	System on Chip
System Memory	<ul style="list-style-type: none"><li>• 16+16 DDR5 DIMM slots (2DPC)</li></ul>
External Drive Bay / Storage	<ul style="list-style-type: none"><li>• Front Side Drive Bay:</li><li>- 8 Hot-swap 2.5" NVMe (PCIe 5.0 x4) drive bays from PCIe Switch</li><li>- 2 Hot-swap 2.5" NVMe (PCIe 5.0 x4) drive bays from CPU</li><li>• Internal Drive Bay: 1 M-key (PCIe 5.0 x4), 1 M-key (PCIe3.0 x4)</li></ul>

## Meeting the Power Demands of Modern AI

Today's AI models require more power and generate more heat than ever before. Traditional air-cooled systems struggle to keep up, leading to inefficiencies and performance risks. The LiquidMax RackScale 64 addresses these challenges with direct liquid cooling that improves thermal management, reduces energy waste, and keeps systems running reliably during sustained high-power operation.



**AceleMax® AXG-428IL**  
(4U Server in LiquidMax RackScale 64)

## Why AMAX

AMAX helps customers deploy dense, liquid-cooled infrastructure built for demanding AI workloads. With 40 years of engineering experience, AMAX offers proven expertise in thermal, mechanical, and electrical design. Each LiquidMax rack is fully assembled, tested, and validated by our team to simplify deployment, improve reliability, and support long-term performance.



**AMAX // DATASHEET**