

Taihang H9236 AI Server



Position

H9236 is a high-performance GPU server based on Intel® Eagle Stream platform. It supports Nvidia H20 modules. It is one of the best choices for AI training and large model inference.

Application

Large model training, large model inference, HPC

Feature

Leading architecture design

- Based on the 4th Gen Intel® Xeon® Scalable Processor, up to 350W TDP
- Support the latest AI training modules
- CPU-Switch supports dual interconnection mode, maximizing the performance of the platform

Flexible System Architecture

- Flexible x16 or x32 PCIe link bandwidth between CPU and Switch
- Support NVidia H20 module
- Full height and half height NIC cards available

Stable product quality

- N+N redundant power supplies
- Stable operation under high ambient temperatures with efficient heat dissipation design
- Extensive experience in developing and debugging multiple modular products on the same platform

Product Specification

Taihang H9236 AI Server

CPU	<ul style="list-style-type: none">• 2x 4th Gen Intel® Xeon® Processor, up to 350W TDP
Memory	<ul style="list-style-type: none">• Up to 32x DDR5 DIMMs, 8 channels per CPU
Storage	<ul style="list-style-type: none">• 10x NVMe SSDs (optional 2x SATA)• 1x onboard M.2 SSD
PCIe Extension	<ul style="list-style-type: none">• 8x FHHL/HHHL (from Switch)• 4x FHHL (from CPU)
Management	<ul style="list-style-type: none">• BMC: AST2600, 1x dedicated RJ45 MGMT port
IO	<ul style="list-style-type: none">• 2x USB, Type-A• 1x PWR BTN with LED• 1x UID BTN with LED• 1x Status LED• 1x Console PORT, Type C• 1x Front VGA
Thermal	<ul style="list-style-type: none">• 15x High-performance hot-swappable fans, N+1 Redundancy
Temperature	<ul style="list-style-type: none">• 5°C - 35°C
PSU	<ul style="list-style-type: none">• 2x 12V PSUs for CPU module, 1+1 Redundancy• 6x 54V PSUs for GPU module, 1+1 Redundancy
Dimension	<ul style="list-style-type: none">• 900mm (D) x 447mm (W) x 352mm (H)