



KR4268V2

4U 8 GPU AI Server with AMD EPYC™ 9004

KR4268V2 with AMD EPYC™ 9004 processors is a powerful 4U AI server featuring the latest PCIe Gen 5.0 architecture. It supports up to 8 FHFL double-width GPUs and 3 additional PCIe GPU cards. Its 10-card pass-through design significantly reduces CPU and GPU communication latency. This system is an ideal solution for advanced AI applications including AI



- 8 FHFL double-width GPU in 4U
- 2x AMD EPYC™ 9004
- Supports additional 2x double-width PCIe GPU, 1x FL single-width PCIe GPU
- 24x 2.5" SSD, up to 16x NVMe U.2
- 11x PCIe 5.0 x16 slots



Enhanced Architecture

- Latest CXL1.1 interconnect technology supports storage-level memory space expansion
- Supports latest 400G NDR infiniband network and 400G OCP 3.0

Flexible Configuration

- Modular design concept supports flexible component-level configuration
- Compatible with the latest AI accelerators in the ecosystem

Advanced Performance

- Supports NVIDIA L40S, H100, and A100 GPUs
- Supports 5nm AMD EPYC™ 9004 processors, up to 192 cores
- Supports 16 NVMe or E3.S hard drives, greatly improving the data access rate
- PCIe topology structure with up to 64GB/s bandwidth channel enables high-throughput data transmission development

Technical Specifications

Model	KR4268-E2-A0-R0-00
Form Factor	4U rack server
GPU	8x FHFL double-width PCIe 5.0 GPU cards; Supports additional 2x FL double-width PCIe GPU; 1x FL single-width PCIe GPU; Supports L40S, H100, A100
CPU	2x AMD EPYC™ 9004 processors, up to 400W
Memory	24x DDR5 4800MT/s RDIMMs
Expansion Slots	11x PCIe 5.0 x16 slots
I/O	Front: 1x USB 3.0, 1x USB 2.0, 1x VGA, 1x Type-C Rear: 2x USB3.0, 1x RJ45, 1x MicroUSB, 1x VGA
OCP	1x OCP 3.0 SFF, Support NCSI
Management	Aspeed 2600
Security	TPM 2.0 (Trusted Platform Module)
Storage	24x 2.5" SSD, up to 16x NVMe U.2; 2x NVMe/SATA M.2
PSU	4x 12V 3000W, 2+2
Dimensions	447mm (W) x 174.5mm (H) x 850mm (D)