



KR6298-E2

6U Extreme AI Server with AMD Instinct™ MI300X 8-GPU



KR6298-E2 is an advanced 6U 8-GPU AI system with AMD Instinct™ MI300X accelerators. With the massive memory capacity and computational power of MI300X, this server is suited for demanding AI workloads and HPC applications, as well as large-scale AI deployments in cloud environments and data centers.

- AMD Instinct™ MI300X 8-GPU in a 6U chassis
- 2x AMD EPYC™ 9004 processors

Extreme Memory and Compute

- 192GB HBM3 memory capacity, 5.3 TB/s peak memory bandwidth, TDP 750W
- Extreme GPU memory and bandwidth for extra-large-scale training models

Supports AI & HPC Workloads

- LLM training, fine-tuning and inference
- Generative AI for text, content, image, video
- Scientific simulations, weather forecasting

Open Ecosystem

- Optimized for GPU-accelerated applications with AMD ROCm software
- Supports mainstream AI frameworks and AI development tools
- Mature open ecosystem enables cost-efficient AI compute solution development

Technical Specifications

Model	KR6298-E2
Form Factor	6U rack server
GPU	MI300X 8-GPU, TDP up to 750W
Processor	2x AMD EPYC™ 9004 processors, max cTDP 400W
Memory	24x DDR5 DIMMs, up to 4800MT/s Supports RDIMM, 3DS RDIMM
Storage	24x 2.5" SSD, up to 16x NVMe U.2
PCIe	8x low-profile x16 + 2x FHHL x16 slots 2x onboard NVMe M.2
OCP	Optional support for 1x OCP 3.0, supports NCSI
Management	DC-SCM BMC management module with Aspeed 2600
TPM	TPM 2.0
PSU	2x 12V 3200W and 6x 54V 2700W, Platinum/Titanium PSU with N+N redundancy
Dimensions	447mm (W) x 263mm (H) x 860mm (D)