



## KR6298V2

# 6U 8-OAM AI Server for Next-Gen AI, LLM

KR6298V2 with 4th Gen Intel® Xeon® Scalable processors is an advanced AI system supporting 8 Habana® Gaudi®2 OAM accelerators. Leveraging the OAM form factor, it delivers scalable, high-speed performance for next-generation AI workloads and LLM.

- 8x fully-connected Habana® Gaudi®2 accelerators
- 2x 4th Gen Intel® Xeon® Scalable
- 32x DDR5, 8TB maximum system memory
- Unprecedented performance to support generative AI, mainstream AI frameworks, development tools and LLM



### Scalable High Performance

- 96GB HBM3, 4.2TB/s converged P2P bandwidth
- Supports AMX/DSA AI engines
- 4000Gbps cross-node communication bandwidth enables pipeline, data parallelism

### Optimized Energy Efficiency

- Refined thermal design, intelligent thermal management, 12V/54V dual power input
- Power dynamic matching for high-power workloads, zero power conversion loss lead to 8% system power reduction

### Open Ecosystem

- Complies with OAM v1.1, UBB v1.5
- Supports mainstream AI frameworks and AI development tools
- Mature open ecosystem enables cost-efficient AI compute solution

## Technical Specifications

<b>Model</b>	<b>KR6288-X2-A0-R0-00</b>
<b>Form Factor</b>	6U rack server
<b>OAM</b>	8x Habana® Gaudi®2 HL225B, 600W
<b>CPU</b>	2x 4th Gen Intel® Xeon® Scalable processors
<b>Chipset</b>	Intel® C740 series
<b>Memory</b>	32x DDR5 DIMMs, up to 4800MT/s
<b>Storage</b>	24x 2.5" SSD, up to 16x NVMe U.2
<b>M.2</b>	2x NVMe/SATA M.2
<b>Expansion Slots</b>	8x low-profile x16 + 2x FHHL x16
<b>I/O</b>	<b>Front:</b> 1x USB 3.0, 1x USB 2.0, 1x VGA <b>Rear:</b> 2x USB 3.0, 1x RJ45, 1x MicroUSB, 1x VGA
<b>OCP</b>	Supports 1x OCP 3.0, supports NCSI
<b>Management</b>	Aspeed 2600
<b>Security</b>	TPM 2.0 (Trusted Platform Module)
<b>PSU</b>	2x 12V 3200W 1+1, 6x 54V 2700W 3+3
<b>Dimensions</b>	482mm (W) x 263mm (H) x 860mm (D)