




# Altos BrainSphere™ R680 F7

Built for AI Training and HPC  
Ultra-High Performance 4U GPU Server

## Powerful Architecture Reinvented for Compute, Storage, and Management



### Dual Processor Architecture

Supports dual 6<sup>th</sup> Gen Intel® Xeon® 6700/6500 Series, up to 350W TDP per CPU



### GPU Expansion

Up to 8 PCIe GPUs, including NVIDIA H200 NVL, RTX PRO™ 6000 Blackwell, L40S



### Large-Capacity Memory

Up to 32 DDR5 DIMM slots, delivering up to 8000 MT/s. Ideal for AI and HPC workloads



### High-Speed Storage

Up to 12x 2.5-inch Gen5 NVMe / SATA / SAS-4 drives. Hot-swappable design



### Enterprise-Grade Power & Cooling

3+1 redundant 3000W titanium PSU, fanless design with high-efficiency cooling



### Remote Management & Security

BMC with IPMI / HTML5 KVM, plus TPM 2.0 & UEFI Secure Boot

## Evolved High-Performance AI GPU Server

### Next-Generation Server Core for AI & HPC

Altos BrainSphere™ R680 F7 is built to handle demanding AI training, HPC, and large-scale data workloads. Powered by dual 6<sup>th</sup> Gen Intel® Xeon® processors, DDR5 memory, and a 4U multi-GPU design, it delivers high-density performance and reliability. Ideal for cloud AI platforms and data centers, the R680 F7 offers the scalability and flexibility needed for real-time analytics, deep learning, and large model training.

### Optimized Thermal, Power, and Manageability

The 4U chassis maintains thermal stability under full load, with streamlined airflow and tool-less fan modules. Redundant 3+1 Titanium 3000W PSUs support efficient, continuous operation in dense deployments. It also features the ASPEED® AST2600 BMC for remote management via IPMI, HTML5 KVM, and SNMP v3, along with built-in security features such as TPM 2.0 and UEFI Secure Boot.

### Extreme Compute Power for AI and HPC

Equipped with Intel® Xeon® 6700/6500 series CPUs, Altos BrainSphere™ R680 F7 delivers strong per-core performance and AI acceleration. It supports up to 8x dual-slot GPUs and is capable of hosting 8x NVIDIA H200 NVL with 4-way NVLink™ bridges—ideal for training large AI models, simulations, and generative workloads. With 32 DDR5 RDIMM/MRDIMM slots supporting up to 8000 MT/s, it ensures high-throughput memory access for latency-sensitive applications like deep learning and analytics.

### Flexible Architecture with Scalable Investment Efficiency

Altos BrainSphere™ R680 F7 offers a more flexible PCIe-based architecture compared to HGX platforms, providing greater configuration scalability to balance performance and cost. This enables phased GPU deployment, supports a broader range of accelerators, and adapts more easily to evolving AI workloads. With the integration of Altos aiWorks solution, the R680 F7 further enhances resource allocation, simplifies AI workflow orchestration, and maximizes return on infrastructure investment.

Contact your Altos Partner now and accelerate your smart computing journey.

[sales@altoscomputing.com](mailto:sales@altoscomputing.com)

# Altos BrainSphere™ R680 F7 Specification

CPU	Dual Intel® Xeon® 6 Processors, TDP up to 350W - Intel® Xeon® 6700-Series Processors - Intel® Xeon® 6500-Series Processors
GPU	Up to 8 PCIe GPUs: NVIDIA H200 NVL (with a 4-way NVLink bridge), RTX PRO™ 6000 Blackwell Server Edition*, L40S
Memory	32 DIMM slots, 8-Channel DDR5 RDIMM/MRDIMM RDIMM Up to 6400 MT/s (1DPC), 5200 MT/s (2DPC) MRDIMM Up to 8000 MT/s**
LAN	2 x 10GbE (Support NCSI function) 1 x 1 GbE Management port
TPM	1 x TPM2.0 module
Drive Bay	Up to 12 x 2.5" (Hot-swappable) Gen5 NVMe/SATA/SAS-4 (SATA and SAS drives require a storage controller card) M.2 support (optional M.2 PCIe card required)
I/O Ports	2 x USB 3.2 Gen1, 1 x VGA, 2 x 10GbE RJ45 ports, 1 x 1 GbE Management port
PCIe Expansion Slots	- 8 x FHFL x16 (Gen5 x16) for GPUs - 4 x FHHL x16 (Gen5 x16) - 3 x LP x16 (Gen5 x16)
Power Supply	3+1 3000W 80 Plus Titanium redundant power supplies
Cooling	Air cool
Operating Properties	Operating: 10°C to 35°C, 8% to 80% (non-condensing) Non-operating: -40°C to 60°C, 20% to 95% (non-condensing)
Security and Manageability	ASPEED® AST2600 BMC TPM 2.0 (FIPS, CC-TCG certification) UEFI Secure Boot SNMP v3 Support
Form factor	4U 448 (W) x 176 (H) x 880 (H) mm 448 (W)x 176(H) x 1000.5(H)mm (includes the dedicated fan module corresponding to the NVIDIA H200 NVL or RTX PRO™ 6000 Blackwell server configuration)
OS	Microsoft® Windows Server, Red Hat Enterprise Linux Server, Ubuntu Server LTS
MC / Safety Compliance	FCC, CE

\* Supports up to 8 x NVIDIA H200 NVL (2x4 NVLink) or 8 x NVIDIA RTX PRO™ 6000; operation requires 25°C ambient temperature.

\*\* MRDIMMs are only supported with Intel® Xeon® 6 Processors with P-cores and in a 1DPC configuration.



## About Altos Computing Inc.

Altos Computing Inc. (abbr. Altos) is established in 2017, is a subsidiary of Acer Group. Leveraging in-house R&D and collaborations with ODM, IHV, and ISV partners, Altos delivers streamlined and cost-effective integrated solutions. Its product portfolio covers AI servers, AI workstations, and AI resource management platforms.

In an era of rapid innovation, Altos provides leading solutions across diverse domains, including High Performance Computing (HPC), Virtual Desktop Infrastructure (VDI), Cloud Infrastructure, and Software-Defined Storage (SDS). Our customers include government agencies, academic institutions, cloud service providers, data center operators, and enterprises.

In a continuing effort to improve the quality of our products, information in this brochure is subject to change without notice.

Images appearing are only representations of some of the configurations available for this model. Availability may vary depending on region. Altos disclaims any liability for errors and omissions in product descriptions. © 2025. All rights reserved.

Altos and the Altos logo are registered trademarks of Altos Computing, Inc. Other trademarks, registered trademarks and/or service marks, indicated or otherwise the properties of their respective owners.