

Lenovo ThinkSystem SR680a V4

Maximum Acceleration for Complex AI



Extreme GPU Acceleration for AI and HPC

The SR680a V4 leverages the 2x Intel® Xeon® 6 processors to provide exceptional bandwidth and low latency for training large language models, generative AI, and HPC simulations. Alongside the NVIDIA HGX B300 NVL8 platform this architecture dramatically reduces time-to-insight for complex workloads. With Intel, the SR680a V4 has performance and efficiency without compromise.

High-Speed Memory and I/O Throughput

With 32 DDR5 DIMM slots operating up to 6400MHz and PCIe Gen5 expansion, the SR680a V4 sustains high GPU utilization for data-intensive pipelines. Eight OSFP links to CX8 switches and multiple high-speed PCIe connections ensure efficient scaling.

Enterprise Reliability and Serviceability

Designed for mission-critical deployments, the SR680a V4 features N+N redundant power supplies, N+1 hot-swap cooling, and hot-swap M.2 boot drives. Its air-cooled design provides thermal headroom for current and next generation GPUs while simplifying service operations. With Intel, the SR680a V4 is built and tested for high compute intensity and scale.

Optimized for Next-Generation AI Workloads

Purpose-built for generative AI, deep learning, and large-scale inference, the SR680a V4 delivers the performance and flexibility needed to meet evolving AI and HPC demands. It is a robust platform for enterprises and cloud providers driving cutting-edge innovation.

AI and HPC Application Examples

- Training and fine-tuning large language models (LLMs)
- Generative AI and deep learning
- High-performance computing (HPC) simulations
- Large-scale inference and data analytics
- Enterprise AI initiatives requiring continuous uptime

Lenovo Makes “AI for All” Possible

Lenovo delivers AI to data deployed inside your data center and at the edge. We have the industry's most comprehensive AI portfolio with over 80 products designed for AI workloads. Lenovo partners with over 50 AI Innovator partners globally to deliver turnkey Hybrid Cloud and Edge AI solutions.

Lenovo

Unique in the industry, Lenovo can deliver AI “from pocket to cloud” with servers, workstations, and AI-enabled devices. Whether you are a public cloud provider or building your AI model on-prem, Lenovo provides systems and solutions that scale to meet your needs. We bring AI to your data where and when you need it the most, in a truly hybrid approach: public, private, or personal.

Specifications

Form Factor	8U rack
Processor	2x Intel® Xeon® 6 processors, up to 350W
Memory	<ul style="list-style-type: none">32x DDR5 DIMM slots, 6400MHzSupports 16 or 32 DIMMs (64GB/96GB/128GB)
GPU	Supports 8x high-performance GPUs: <ul style="list-style-type: none">8x NVIDIA HGX B300 NVL8 GPUs
I/O Expansion	<ul style="list-style-type: none">4x PCIe Gen5 x16 FHHL (connected to CPU)1x OCP slot8x OSFP connectors connected to CX8 switches (1 per SXM module)
Storage	<ul style="list-style-type: none">8x 2.5" hot-swap NVMe SSDs2x M.2 NVMe for boot (RAID via VROC)
Power	<ul style="list-style-type: none">8x 3200W @12V Titanium PSUs (Full N+N redundancy)8x 3800W @54V planned for Wave 2, CFFv5 – Supports multiple input power ranges:<ul style="list-style-type: none">200–220V AC (2900W)221–249V AC (3200W)250–277V AC (3800W)
Cooling	Air-cooled with N+1 hot-swap fan solution
Management	ThinkSystem XClarity Controller 3 (XCC3) with Platform Firmware Resiliency (HW Root of Trust)
OS Support	RHEL, Ubuntu

About Lenovo

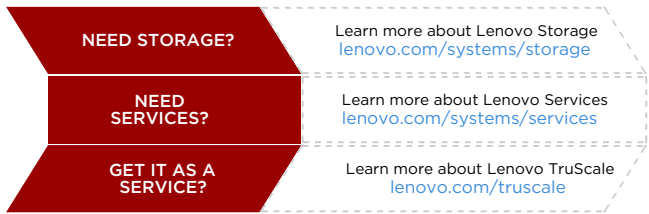
Lenovo (HKSE: 992) (ADR: LNVGY) is a US\$62 billion revenue global technology powerhouse, ranked #171 in the Fortune Global 500, employing 77,000 people around the world, and serving millions of customers every day in 180 markets. Focused on a bold vision to deliver smarter technology for all, Lenovo is expanding into new growth areas of infrastructure, mobile, solutions and services. This transformation is building a more inclusive, trustworthy, and sustainable digital society for everyone, everywhere.

Lenovo Professional Services

Offering a breadth of services, solutions and platforms, Lenovo’s AI Professional Services help businesses of all sizes navigate the AI landscape, find the right solutions, and put AI to work for their organizations quickly, cost-effectively and at scale. Lenovo experts bring AI from concept to reality — from designing AI roadmaps to deploying platforms.

For More Information

To learn more about the Lenovo ThinkSystem SR680a V4, contact your Lenovo representative or Business Partner or visit lenovo.com/thinksystem. For detailed specifications consult the [SR680a V4 Product Guide](#).



© 2025 Lenovo. All rights reserved.

Availability: Offers, prices, specifications and availability may change without notice. Lenovo is not responsible for photographic or typographic errors. **Warranty:** For a copy of applicable warranties, write to: Lenovo Warranty Information, 1009 Think Place, Morrisville, NC, 27560. Lenovo makes no representation or warranty regarding third-party products or services. **Trademarks:** Lenovo, the Lenovo logo, ThinkSystem®, and XClarity® are trademarks or registered trademarks of Lenovo. Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries. Other company, product, or service names may be trademarks or service marks of others. Document number DS0198, published September 2, 2025. For the latest version, go to lenovopress.lenovo.com/ds0198.

