

Lenovo ThinkSystem SR670 V2

From Exascale to Everyscale™ – a modular platform tailored to your enterprise AI requirements


Designed for Your Enterprise

Lenovo ThinkSystem SR670 V2 delivers optimal performance for Artificial Intelligence (AI), High Performance Computing (HPC) and graphical workloads across an array of industries.

Retail, manufacturing, financial services and healthcare industries are leveraging GPUs to extract greater insights and drive innovation utilizing machine learning (ML) and deep learning (DL). Here are a few ways accelerated computing leverages GPUs in different organizations:

- Remote visualization for work-from-home teams
- Ray-traced rendering for photo-realistic graphics
- Powerful video encoding and decoding
- In-silico trials and immunology in Life Sciences
- Natural language processing (NLP) for call centers
- Automatic optical inspection (AOI) for quality control
- Computer vision for retail customer experience

As more workloads leverage the capabilities of accelerators, the demand for GPU's increases. The ThinkSystem SR670 V2 delivers an optimized enterprise-grade solution for deploying accelerated HPC and AI workloads in production, maximizing system performance.



EveryScale Platform means: Versatility

The SR670 V2 features a modular design for ultimate flexibility. With multiple different front compartment options, configurations include:

- Up to eight double-width GPUs with NVLink Bridge
- NVIDIA HGX™ A100 4-GPU with NVLink and Lenovo Neptune™ hybrid liquid cooling performance
- Choice of front or rear high-speed networking
- Choice of local high speed 2.5, 3.5 and; NVME storage



The ThinkSystem SR670 V2 is built on two 3rd Gen Intel® Xeon® Scalable processors and is designed to support the vast NVIDIA Ampere datacenter portfolio:

- NVIDIA HGX™ A100 4-GPU with NVLink
- NVIDIA A100 Tensor Core GPU
- NVIDIA A40 Tensor Core GPU
- NVIDIA A30 Tensor Core GPU

The ThinkSystem SR670 V2 delivers performance optimized for your workload, be it visualization, rendering or computationally intensive HPC and AI.

Most Powerful Compute Platform

The NVIDIA A100 Tensor Core GPU delivers unprecedented acceleration—at every scale—to power the world's highest performing elastic data centers for AI, data analytics, and HPC applications. The A100 can efficiently scale up or be partitioned into seven isolated GPU instances, with Multi-Instance GPU (MIG) providing a unified platform that enables elastic data centers to dynamically adjust to shifting workload demands. A rack of 13 ThinkSystem SR670 V2s can generate up to 2 PFLOPS of compute power.



Cutting Edge Cooling Capability

Traditional air-cooling methods are reaching critical limits. Increases in component power especially on CPU's and GPU's have resulted in higher energy and infrastructure costs, extremely loud systems and heightened carbon footprints.

To combat these challenges and dissipate heat quickly, some models of the SR670 V2 employ Lenovo Neptune™ liquid-to-air hybrid cooling technology. The heat of the NVIDIA HGX™ A100 GPU's is removed through a unique closed loop liquid-to-air heat exchanger that delivers the benefits of liquid cooling such as lower power consumption, quiet operation and higher performance without adding plumbing.

Solutions That Scale

Whether you're just starting with AI or moving into production, your solution must scale with your organization's needs. The ThinkSystem SR670 V2 can be used in a cluster environment using high-speed fabric to scale out as your workload demands increase.

Enabled with Lenovo intelligent Computing Orchestration (LiCO), you can support multiple users and scaling within a single cluster environment. LiCO is a powerful platform that manages cluster resources for HPC and AI applications.

LiCO provides workflows for both AI and HPC, and supports multiple AI frameworks, including TensorFlow, Caffe, Neon, and MXNet, allowing you to leverage a single cluster for diverse workload requirements.

Get a Hands-on, Concierge Experience

Getting started is easy. At the Lenovo AI Innovation Centers, you can test your own Proof of Concept (PoC) on different hardware and software platforms, including the ThinkSystem SR670 V2. Lenovo data scientists and AI solution architects are available to help you along the way. Lenovo can work with you to develop an end-to-end solution for your unique use case with professional services and deep industry partnerships to ensure your success.

Specifications

Form Factor/Height	3U Rack-mount with three modules
Processor	2x 3rd Gen Intel® Xeon® Scalable processors per node
Memory	Up to 4TB using 32x 128GB 3200MHz TruDDR4 3DS RDIMMs per node Intel® Optane™ Persistent Memory 200 Series
Base Module	Up to 4x double-wide, full-height, full-length; FHFL GPUs ; each; PCIe Gen4 x16 Up to 8x 2.5" Hot Swap SAS/SATA/NVMe, or 4x 3.5" Hot Swap SATA (selected configurations)
Dense Module	Up to 8x double-wide, full-height, full-length GPUs each PCIe Gen4 x16 on PCIe switch Up to 6x EDSFF E.1S NVMe SSDs
HGX Module	NVIDIA HGX™ A100 4-GPU with 4x NVLink connected SXM4 GPUs Up to 8x 2.5" Hot Swap NVMe SSDs
RAID Support	SW RAID standard; Intel® Virtual RAID on CPU (VROC), HBA or HW RAID with flash cache options
I/O Expansion	Up to 4x PCIe Gen4 x16 adapters (2 front or 2-4 rear) and 1x PCIe Gen4 x16 OCP 3.0 mezz adapter (rear) depending on the configuration
Power and Cooling	Four N+N redundant hot-swap PSUs (up to 2400W Platinum) Full ASHRAE A2 support with internal fans and Lenovo Neptune™ liquid-to-air hybrid cooling on NVIDIA HGX™ A100
Management	Lenovo XClarity Controller (XCC) and Lenovo Intelligent Computing Orchestration (LiCO)
OS Support	Red Hat Enterprise Linux, SUSE Linux Enterprise Server, Microsoft Windows Server, VMware ESXi Tested on Canonical Ubuntu

Data Center Reliability Leader

At Lenovo, we take a customer-centric approach, which is why ThinkSystem servers consistently rank #1 in reliability*. Also, Lenovo is the leading provider of Supercomputer systems on the TOP500**. The ThinkSystem SR670 V2 provides the latest in performance and reliability in a scalable solution for enterprise and research.

* See [ITIC Reliability Study](#)

** See [Lenovo Celebrates #1 Vendor for TOP500 HPC Systems in 2020](#)

© 2022 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, From Exascale to Everscale, Lenovo Neptune, ThinkSystem, TruDDR4, and XClarity® are trademarks or registered trademarks of Lenovo. Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries. Linux® is the trademark of Linus Torvalds in the U.S. and other countries. Microsoft®, Windows Server®, and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others. Document number DS0123, published April 6, 2021. For the latest version, go to lenovopress.com/ds0123.

For More Information

To learn more about the Lenovo ThinkSystem SR670 V2, contact your Lenovo representative or Business Partner or visit lenovo.com/thinksystem. For detailed specifications consult [SR670 V2 Product Guide](#).

