


ThinkSystem SR675 V3 Sets 20 World Records with New SPECchpc 2021 Benchmark Result Performance Benchmark Result

The Lenovo ThinkSystem SR675 V3 has set twenty new 2-sockets performance world records with the SPECchpc_2021_tny_base, SPECchpc_2021_tny_peak, SPECchpc_2021_sml_base and SPECchpc_2021_sml_peak metric from the SPECchpc 2021 Benchmark.

These new benchmark results, published in new SPEC reports on November 1, 2023, demonstrate that the ThinkSystem SR675 V3 continues Lenovo's leadership with outstanding performance for the server industry.

The SR675 V3, using OpenACC as node-level parallelization model, has achieved the best 1 node, 2 socket, 1 GPU base and peak score as follows:

- **SPECchpc_2021_tny_base = 17.7**
- **SPECchpc_2021_tny_peak = 17.9**



The SR675 V3, using OpenACC as node-level parallelization model, has achieved the best 1 node, 2 socket, 2 GPU base and peak score as follows:

- **SPECchpc_2021_tny_base = 32.3**
- **SPECchpc_2021_tny_peak = 32.6**

The SR675 V3, using OpenACC as node-level parallelization model, has achieved the best 1 node, 2 socket, 3 GPU base and peak score as follows:

- **SPECchpc_2021_tny_base = 45.7**
- **SPECchpc_2021_tny_peak = 46.1**

The SR675 V3, using OpenACC as node-level parallelization model, has achieved the best 1 node, 2 socket, 4 GPU base and peak score as follows:

- **SPECchpc_2021_tny_base = 57.8**
- **SPECchpc_2021_tny_peak = 58.3**

The SR675 V3, using OpenACC as node-level parallelization model, has achieved the best 1 node, 2 socket, 5 GPU base and peak score as follows:

- **SPECchpc_2021_tny_base = 58.9**
- **SPECchpc_2021_tny_peak = 59.3**

The SR675 V3, using OpenACC as node-level parallelization model, has achieved the best 1 node, 2 socket, 6 GPU base and peak score as follows:

- **SPEChpc_2021_tny_base = 68.7**
- **SPEChpc_2021_tny_peak = 69.3**

The SR675 V3, using OpenACC as node-level parallelization model, has achieved the best 1 node, socket, 7 GPU base and peak score as follows:

- **SPEChpc_2021_tny_base = 69.2**
- **SPEChpc_2021_tny_peak = 69.4**

The SR675 V3, using OpenACC as node-level parallelization model, has achieved the best 1 node, 2 socket, 8 GPU base and peak score as follows:

- **SPEChpc_2021_tny_base = 79.9**
- **SPEChpc_2021_tny_peak = 80.4**

The SR675 V3, using OpenACC as node-level parallelization model, has achieved the best 1 node, Base and Peak ACC score as follows:

- **SPEChpc_2021_tny_base = 79.9**
- **SPEChpc_2021_tny_peak = 80.4**

The SR675 V3, using OpenACC as node-level parallelization model, has achieved the best 1 node, 2 socket, 7 GPU base and peak score as follows:

- **SPEChpc_2021_sml_base = 11**
- **SPEChpc_2021_sml_peak = 11.1**

The SPEChpc 2021 Benchmark suite is the industry standard to evaluate hardware-based accelerator devices and the performance of parallel computing workloads. In order to support heterogeneity system architecture, the SPEChpc 2021 benchmark supports multiple programming models:

- Pure MPI
- MPI+OpenMP
- MPI+OpenACC
- MPI+OpenMP with target offload

By properly configuring the CPU based system according to NUMA architecture, the MPI+OpenMP can reduce communication needs and memory consumption in some sub-benchmark of SPEChpc 2021, hence the performance improvement.

The ThinkSystem SR675 V3 server was configured as follows:

- 2x AMD EPYC 9654 processors (96 cores, 2.40 GHz)
- 768 GB memory (24x 32GB RDIMMs at 4800MHz)
- 8x NVIDIA H100 PCIe 80GB GPUs
- Red Hat Enterprise Linux 8.6

The results are current as of November 1, 2023.

To view details of the results, see the following SPEC web pages:

SPEChpc using the OpenACC model 1 node, 2 socket, 1 GPU base and peak:
<https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00248.html>

SPEChpc using the OpenACC model 1 node, 2 socket, 2 GPU base and peak:
<https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00247.html>

SPEChpc using the OpenACC model 1 node, 2 socket, 3 GPU base and peak:
<https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00246.html>

SPEChpc using the OpenACC model 1 node, 2 socket, 4 GPU base and peak:
<https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00245.html>

SPEChpc using the OpenACC model 1 node, 2 socket, 5 GPU base and peak:
<https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00244.html>

SPEChpc using the OpenACC model 1 node, 2 socket, 6 GPU base and peak:
<https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00243.html>

SPEChpc using the OpenACC model 1 node, 2 socket, 7 GPU base and peak:
<https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00242.html>

SPEChpc using the OpenACC model 1 node, 2 socket, 8 GPU base and peak and 1-node, base and peak ACC:
<https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00241.html>

SPEChpc using the OpenACC model 1 node, 2 socket, 7 GPU base and peak Small:
<https://spec.org/hpc2021/results/res2023q4/hpc2021-20230905-00256.html>

To view all SPEChpc 2021 results, go to
<https://spec.org/hpc2021/results/>

About the ThinkSystem SR675 V3

The Lenovo ThinkSystem SR675 V3 is a powerful and versatile server designed to meet the demands of modern data center environments. The server delivers optimal performance for Artificial Intelligence (AI), High Performance Computing (HPC) and graphical workloads across an array of industries. The SR675 V3 is a versatile GPU-rich 3U rack server that supports eight double-wide GPUs including the new NVIDIA H100 and L40S Tensor Core GPUs, or the NVIDIA HGX H100 4-GPU offering with NVLink and Lenovo Neptune hybrid liquid-to-air cooling.

The server is based on the new AMD EPYC 9004 Series processors (formerly codenamed "Genoa", "Genoa-X" and "Bergamo") with up to 160 PCIe lanes and up to 3TB of the latest DDR5 memory. Multiple drive options using SAS/SATA and NVMe with hot-swap capabilities and XClarity system management software enable changes to be made quickly with ease. The versatile design doesn't stop at storage; the SR675 V3 includes support for multiple options for GPU and PCIe to satisfy graphics, speed, and budget requirements.

About SPEChpc 2021

High Performance Computing (HPC) systems are getting built with an increased level of heterogeneity. The numerous types of accelerators bring in tremendous extra computing power, while at the same time introduce big challenges in performance evaluation and characterization. More complications are added to the problem when multiple parallel and accelerator programming models have been developed with each only supporting a subset of the computing devices.

The SPEChpc 2021 Benchmark Suite address these challenges by providing a set of application benchmark suites using a comprehensive measure of real-world performance for the state-of-the-art HPC systems. They offer well-selected science and engineering codes that are representative of HPC workloads and are portable across CPU and accelerators, along with certain fair comparative performance metrics.

SPEChpc 2021 focuses on compute intensive parallel performance across one or more nodes, which means these benchmarks emphasize the performance of the following components:

- Processors & GPUs - The CPU chips and optionally, an acceleration device such as a GPU
- Memory - The memory hierarchy, including caches and main memory
- Interconnects - The communication between nodes of a cluster
- Compilers - C, C++, and Fortran compilers, including optimizers
- MPI - The MPI implementation.

The SPEChpc 2021 benchmark suite is broken out into four workloads, Tiny, Small, Medium and Large:

- SPEChpc2021 Tiny workloads use up to 60 GB of memory and are intended for use on a single node using between 1 and 256 ranks.
- SPEChpc2021 Small workloads use up to 480 GB of memory and are intended for use on one or more nodes using between 64 and 1024 ranks
- SPEChpc2021 Medium workloads use up to 4 TB of memory and are intended for use on a mid-size cluster using between 256 and 4096 ranks
- SPEChpc2021 Large workloads use up to 14.5 TB of memory and are intended for use on a larger clusters using between 2048 and 32,768 ranks

Learn more

To learn more about solutions for HPC applications, contact your Lenovo Sales Representative or visit <https://www.lenovo.com/us/en/servers-storage/solutions/hpc/>

To find out more about SPEC, visit <https://www.spec.org>

To learn more about the Lenovo ThinkSystem SR675 V3 server, visit the SR675 V3 product web page: <https://www.lenovo.com/us/en/p/servers-storage/servers/racks/thinksystem-sr675-v3/7d9r1000na>

Related product families

Product families related to this document are the following:

- [2-Socket Rack Servers](#)
- [SPEChpc Benchmark Results](#)
- [ThinkSystem SR675 V3 Server](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2024. All rights reserved.

This document, LP1851, was created or updated on November 21, 2023.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1851>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1851>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

Lenovo Neptune®

ThinkSystem®

XClarity®

The following terms are trademarks of other companies:

AMD and AMD EPYC™ are trademarks of Advanced Micro Devices, Inc.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

SPEC® and SPECchpc™ are trademarks of the Standard Performance Evaluation Corporation (SPEC).

Other company, product, or service names may be trademarks or service marks of others.