Lenovo ThinkSystem Servers Continue to Lead the Industry in Performance and Customer Value

Positioning Information

For many customers, the performance of their servers is paramount, and as a result Lenovo strives to design its systems to maximize performance. The proof points to show the success we have achieved doing this are the #1 benchmarks we have achieved. We produce systems that cater across all workloads that our customers require, achieving #1 benchmarks across a broad range of applications types.

As of April 1, 2024 we currently hold a total of **555 world records** for performance. This number consists of:

- **184** third-generation Intel Xeon Scalable processors
- **45** second-generation Intel Xeon Scalable processors
- **48** fourth-generation Intel Xeon Scalable processors
- **3** fifth-generation Intel Xeon Scalable processors
- **53** third-generation AMD EPYC processors
- **2** using second-generation AMD EPYC processors
- **39** using first-generation Intel Xeon Scalable processors
- **181** fourth-generation AMD EPYC processors

These results were achieved on a wide variety of workloads that include Applications, Data Management, Application Development and Testing, Infrastructure, and Engineering/Technical computing. Lenovo achieved these world-records across multiple ThinkSystem platforms ranging from single-socket to 8-socket, spanning Intel and AMD architectures. This demonstrates Lenovo's ability to design systems that create value for customers across architectures, workloads, and industries.

While specific ThinkSystem models were used for these benchmark tests, much of the same technology and attention to detail is prevalent throughout the ThinkSystem product line, providing exceptional performance and reliability across the product line.

This article summarizes the outstanding performance of the world record ThinkSystem servers.

Applications current world records

Lenovo ThinkSystem SR950, SR860 V2, SR850 V2, SR650 V3, SR635 V3, SR645 V3, SR655 V3, SR665 V3 and SR860 V3 servers achieved **31** current world records in benchmarks designed to test system performance across a range of application workload types. Application types include Collaboration, Enterprise Resource Management, Supply Chain, Customer Resource Management, and other typical data center applications.

- SAP Sales and Distribution (SAP SD 2T) Benchmark
- SPEC CPU 2017 Benchmark

SAP Sales and Distribution (SAP SD 2T) Benchmark

The **1** ThinkSystem world record result on the SAP Sales and Distribution Benchmark is:

Performance

ThinkSystem SR655 (**1** current world record)

- 1-Socket

About this benchmark: SAP Sales and Distribution (SD) Standard Application Benchmarks test the hardware and database performance of SAP applications and components. SAP Application Performance Standard (SAPS) is a hardware-independent unit of measurement that describes the performance of a system in creating the order, creating a delivery note for the order, displaying the order, changing the delivery, posting a goods issue, listing orders, and creating an invoice.

Why it matters: If you are running SAP Business Suite applications, a leadership benchmark score means this system is the highest-performing server in the industry for processing your business transaction workflows in an SAP environment.

For more information see the Performance Benchmark Reports section.

SPEC CPU 2017 Benchmark

The **30** ThinkSystem world record results on the SPEC CPU 2017 benchmark include:

SPEC CPU 2017 Performance

ThinkSystem SR950:

- 6-Socket — **1** current world record
- 3-Socket — **4** current world records

ThinkSystem SR950 V3:

- 8-Socket — **1** current world record

ThinkSystem SR860 V2:

- 4-Socket – **1** current world record

ThinkSystem SR860 V3:

- 4-Socket – **3** current world record

ThinkSystem SR655 V3:

- 1-Socket — **8** current world records

ThinkSystem SR665 V3:

- 2-Socket — **12** current world records

About this benchmark: SPEC CPU 2017 contains SPEC's next-generation, industry-standardized, CPU intensive suites for measuring and comparing compute intensive-performance, stressing a system's processor, memory subsystem and compiler. CPU 2017 has 43 sub-benchmarks, organized into four suites. SPEC designed these suites to provide a comparative measure of compute-intensive performance across the widest practical range of hardware using workloads developed from real user applications.

Why it matters: If you are running compute-intensive workloads, a world-record benchmark score means this server is the highest performing server in its class in terms of how fast a server completes a CPU-intensive task (speed) and/or how much a server can accomplish in a certain amount of time (throughput or rate measurement).

For more information see the Performance Benchmark Reports section.

Data Management current world records

Lenovo ThinkSystem SR860 V2, SR950, SR665, SR655, SR655 V3, SR650 and SR650 V3 servers delivered **79** current world records on benchmarks designed to test the performance of servers performing Structured Data Management analysis, Structured Data Analytics, and Unstructured Data Analytics.

- SAP HANA (BWoH) Benchmark
- STAC-M3 Benchmarks
- TPC-E Benchmark

SAP HANA (BWoH) Benchmark

The **10** ThinkSystem world record results on the HANA BWoH benchmark in a single-node setup include:

Performance

ThinkSystem SR950:

- 4-Socket (10.4B Records v3) — **1** current world record

ThinkAgile HX7821:

- 4 Socket – **1** current world record

ThinkAgile VX 4U Certified Node:

- 8 Socket – **2** current world record

ThinkSystem SR860 V2

- 4 Socket (1.3B Records v3) – **1** current world record
- 4 Socket (5.2B Records v3) – **1** current world record
- 4 Socket (7.8B Records v3) – **1** current world record

ThinkSystem SR650 V2

- 2 Socket (1.3B Records v3) – **1** current world record
- 2 Socket (5.2B Records v3) – **1** current world record

ThinkAgile VX650 V3 Certified Node:

- 2 Socket (7.8B Records Virtualized v3)– **1** current world record

About this benchmark: The SAP BW Edition for SAP HANA (BW/HANA) Standard Application Benchmark is the latest addition to the list of SAP BW benchmarks and goes well beyond the scope and features used in the BW-AML benchmark. The benchmark is designed to fully utilize the new capabilities of SAP HANA to process the benchmark workload. The benchmark consists of three phases: data load, query throughput, and query runtime.

* The current SAP BW edition for SAP HANA benchmark is version 3. Existing v1 and v2 results are still valid; however, SAP will no longer accept new benchmark results for versions v1 and v2.

Why it matters: If you are running SAP HANA, a #1 benchmark score means this server is the highest performing server in its class for processing SAP HANA memory-intensive database and analytics workloads.

For more information see the Performance Benchmark Reports section.

STAC-M3 Benchmarks

The **56** ThinkSystem world record results on the STAC-M3 benchmark include:

STAC-M3

ThinkSystem SR950 (**16** current world records):

- 4-Socket (16 Shasta suite)

ThinkSystem SR860 V2 (**26** current world records)

- 4-Socket (3 records Single Node Antuco suite)
- 4-Socket 2 node (Antuco 8 records, 3 yr Kanaga 15 records)

ThinkSystem SR650 (**1** current world records):

- 2-Socket (1 Antuco suite)

ThinkSystem SR650 V2 (**4** current world records)

- 2-Socket (4 Antuco suite)

ThinkSystem SR650 V3 (**9** current world records)

- 2-Socket (9 Single-Node Antuco suite)

About this benchmark: The STAC-M3 Benchmark suite is the industry standard for testing solutions that enable high-speed analytics on time series data. The STAC-M3 benchmarks measures challenging areas such as time-series analytics, risk simulations, and processing of very-high-speed data. The key metric is query response time. In particular, STAC benchmarks test high-speed analytics on time-series data — tick-by-tick market data. The benchmark is used by large global banks, brokerage houses, exchanges, hedge funds, proprietary trading shops, and other market participants.

Why it matters: If you are running high-speed financial services or securities workloads, a #1 benchmark score means this server is the highest performing server in its class for processing high-speed analytics and financial transactions.

For more information see the Performance Benchmark Reports section.

TPC-E Benchmark

The **13** ThinkSystem world record results on the TPC-E benchmark include:

TPC-E Performance and Price/Performance

ThinkSystem SR665 (**3** current world record):

- 2-Socket

ThinkSystem SR655 (**2** current world record):

- 1-Socket

ThinkSystem SR655 V3 (**3** current world record):

- 1-Socket

ThinkSystem SR650 V3 (**5** current world record):

- 2-Socket

About this benchmark: The TPC-E benchmark is designed to enable users to more objectively measure and compare the performance and price of various OLTP systems. The TPC-E benchmark uses a database to model a brokerage firm with customers who generate transactions related to trades, account inquiries, and market research. Although the underlying business model of TPC-E is that of a brokerage firm, the database schema, data population, transactions, and implementation rules have been designed to be broadly representative of modern OLTP systems in general.

Why it matters: If you are running On-Line Transaction Processing (OLTP) or decision-support workloads and databases, a #1 benchmark score means this server is the highest-performing server in its class for data-intensive OLTP transactions and data-intensive queries.

For more information see the Performance Benchmark Reports section.

Application Development & Testing current world records

Lenovo ThinkSystem SR950, SR860 V2, SR665 V3, SR655 V3, SN550 V2 and SR860 V3 servers hold **60** current world records in benchmarks designed to test server performance in a Java development and testing environment.

- SPECjbb2015

SPECjbb2015

The **60** ThinkSystem world record results on the SPECjbb2015 benchmark include:

Performance

ThinkSystem SR950:

- 8-Socket — **2** current world records (Windows)
- 6-Socket — **10** current world records (8 Linux / 2 Windows)
- 3-Socket — **12** current world records (Linux)

ThinkSystem SR860 V3

- 4-Socket – **14** current world record (10 Linux / 4 Windows)

ThinkSystem SR665 V3:

- 2-Socket — **6** current world record (Windows)
- 2-Socket — **4** current world record (Linux)

ThinkSystem SR655 V3:

- 1-Socket — **6** current world record (Windows)

ThinkSystem SN550 V2:

- 2-Socket (14 nodes) — **6** current world record (2 Linux / 4 Windows)

About this benchmark: The SPECjbb 2015 benchmark has been developed from the ground up to measure performance based on the latest Java application features. It is relevant to all audiences interested in Java server performance, including JVM vendors, hardware developers, Java application developers, researchers and members of the academic community.

Why it matters: If you are interested Java server performance, a world-record benchmark score means this server is the highest performing server in its class for Java application response time and throughput needs.

For more information see the Performance Benchmark Reports section.

Infrastructure current world records

Lenovo ThinkSystem SR950, SR860 V2, SR850, SR665, SR665 V3, SR655, SR655 V3, SR645, SR645 V3, SR635, SN850, SN550, SN550 V2 servers earned **36** current world records on benchmarks designed to test the performance of servers utilized in an IT/Web infrastructure/AI role. Workload types include Networking, Systems Management, Virtual Desktop, File & Print, Media Streaming, and Web Serving.

- SPEC Power
- SPEC VIRT_SC 2013
- SPECvirt Datacenter 2021

SPEC Power

The **22** ThinkSystem world record results on the SPECpower_ssj2008 benchmark include:

Performance

ThinkSystem SR950 (**3** current world records):

- 8-Socket (2 Linux / 1 Windows)

ThinkSystem SR860 V3 (**2** current world record)

- 4-Socket (1 Linux/ 1 Windows)

ThinkSystem SR850 V3 (**2** current world record):

- 4-Socket (1 Linux/ 1 Windows)

ThinkSystem SN850 (**3** current world records):

- 4-Socket/7-node (4S/7N, 1 Linux / 2 Windows)

ThinkSystem SN550 (**2** current world record):

- 14-node (14N, Windows)

ThinkSystem SN550 V2 (**1** current world record):

- 14-node (14N, Linux)

ThinkSystem SR635 V3 (**2** current world record)

- 1-Socket (1 Linux / 1 Windows)

ThinkSystem SR645 V3 (**4** current world record)

- 2-Socket (1 Linux / 3 Windows)

ThinkSystem SR655 V3 (**3** current world record)

- 1-Socket (2 Linux / 1 Windows)

About this benchmark: The SPECpower_ssj 2008 benchmark suite measures the power and performance characteristics of server-class computer equipment. It is used to compare power and performance among different servers and serves as a tool set for use in improving server efficiency.

Why it matters: The IT industry, computer manufacturers, and governments are increasingly concerned with the energy use of servers. This benchmark provides a means to measure power (at the AC input) in conjunction with a performance metric. This helps IT managers consider power characteristics along with other selection criteria to increase the efficiency of data centers. For those concerned about energy savings, a world-record benchmark score means the server provides the best performance in its class relative to power consumption.

For more information see the Performance Benchmark Reports section.

SPEC VIRT_SC 2013

The **11** ThinkSystem world record results on the SPEC virt_sc 2013 benchmark include:

Performance

ThinkSystem SR950 (**3** current world records):

- 8-Socket
- 4-Socket

Performance-per-watt

ThinkSystem SR950 (**3** current world records):

- 8-Socket
- 4-Socket

ThinkSystem SR650 (**1** current world record):

- 2-Socket

Server Performance-per-watt

ThinkSystem SR950 (**3** current world records):

- 8-Socket
- 4-Socket

ThinkSystem SR650 (**1** current world record):

- 2-Socket

About this benchmark: The SPEC virt_sc 2013 benchmark measures the end-to-end performance of all system components including the hardware, virtualization platform, and the virtualized guest operating system and application software. SPEC virt_sc 2013 is the second-generation SPEC VIRT benchmark for evaluating the virtualization performance of datacenter server consolidation, including enterprise-class workloads.

Why it matters: If you virtualize multiple workloads, a world-record benchmark score means this system is the highest performing server in its class for memory-intensive virtualized environments.

For more information see the Performance Benchmark Reports section.

SPECvirt Datacenter 2021

3 ThinkSystem world record results on the SPECvirt Datacenter 2021 benchmark is:

ThinkSystem SR665 (**1** current world record):

- 2-socket, 4 node

ThinkSystem SR665 V3 (**2** current world record):

- 2-socket, 4-node

About this benchmark: The SPECvirt® Datacenter 2021 benchmark is the next generation of virtualization benchmarking for measuring performance of a scaled-out datacenter. The SPECvirt Datacenter 2021 benchmark is a multi-host benchmark using simulated and real-life workloads to measure the overall efficiency of virtualization solutions and their management environments.

The SPECvirt Datacenter 2021 benchmark differs from the SPEC VIRT_SC® 2013 benchmark in that SPEC VIRT_SC benchmark measures single host performance and provides interesting host-level information. However, most of today's datacenters use clusters for reliability, availability, serviceability, and security. Adding virtualization to a clustered solution enhances server optimization, flexibility, and application availability while reducing costs through server and datacenter consolidation.

Why it matters: If you virtualize multiple workloads, a world-record benchmark score means this system is the highest performing server in its class for memory-intensive virtualized environments.

For more information see the Performance Benchmark Reports section.

Engineering/Technical current world records

Lenovo ThinkSystem SD650-N V2, SD650 V2, SR950, SR860 V2, SR670, SR670 V2, SR665, SR665 V3, SR655, SR655 V3, SR650, SR650 V2, SR860 V3 and SE450 servers earned **349** current world records in benchmarks designed to test a system's performance running various Engineering/Technical/HPC/ AI workloads.

- SPEC OMP 2012
- SPECmpiM 2007
- SPEC ACCEL
- MLPerf
- SPEChpc 2021

SPEC OMP 2012

The **11** ThinkSystem world record results on the SPEC OMP 2012 benchmark include:

Performance

ThinkSystem SR950 (**1** current world record):

- 3-Socket

ThinkSystem SR665 V3 (**6** current world record):

- 2-Socket

ThinkSystem SR655 V3 (**2** current world record):

- 1-Socket

ThinkSystem SR860 V3 (**2** current world record):

- 4-Socket

About this benchmark: The SPECompG 2012 benchmark is designed for measuring performance using applications based on the OpenMP 3.1 standard for shared-memory parallel processing. The benchmark includes 14 scientific and engineering application codes, covering everything from computational fluid dynamics (CFD) to molecular modeling to image manipulation.

Why it matters: If you are running scientific or engineering applications, a leadership benchmark score means this system is the highest performing server in its class for these workloads.

For more information see the Performance Benchmark Reports section.

SPECmpiM 2007

The **67** ThinkSystem world record results on the SPECmpiM 2007 benchmark include:

Performance

Medium **38** current world records

ThinkSystem SR950 (**3** current world records):

- 8-Socket (Medium Metric 1-node)
- 6-Socket (Medium Metric 1-node)
- 3-Socket (Medium Metric 1-node)

ThinkSystem SR860 V2 (**12** current world records)

- 4-Socket (Medium Metric 1-node)
- 4-Socket (Medium Metric1-node)
- 4-Socket (Medium Metric 2-node)
- 4-Socket (Medium Metric 2-node)
- 4-Socket (Medium Metric 3-node)
- 4-Socket (Medium Metric 3-node)
- 4-Socket (Medium Metric 3-node)
- 4-Socket (Medium Metric 3-node)
- 4-Socket (Medium Metric 4-node)
- 4-Socket (Medium Metric 4-node)
- 4-Socket (Medium Metric 4-node)
- 4-Socket (Medium Metric 4-node)

ThinkSystem SR665 (**11** current world records):

- 2-Socket (Medium Metric 3-node base and peak)
- 2-Socket (Medium Metric 3-node base and peak)
- 2-Socket (Medium Metric 4-node peak)
- 2-Socket (Medium Metric 5-node base and peak)
- 2-Socket (Medium Metric 5-node base and peak)
- 2-Socket (Medium Metric 5-node base and peak)
- 2-Socket (Medium Metric 5-node base and peak)
- 2-Socket (Medium Metric 6-node base and base)
- 2-Socket (Medium Metric 6-node base and base)
- 2-Socket (Medium Metric 6-node peak and peak)
- 2-Socket (Medium Metric 6-node peak and peak)

ThinkSystem SR665 V3 (**8** current world records):

- 2-Socket (Medium Metric 1-node base and peak) (4 records)
- 2-Socket (Medium Metric 2-node base and peak) (4 records)

ThinkSystem SR655 V3 (**4** current world records):

- 1-Socket (Medium Metric 1-node base and peak) (2 records)
- 1-Socket (Medium Metric 2-node base and peak) (2 records)

Large **29** current world records

ThinkSystem SR860 V2 (**16** current world records)

- 4-Socket (Large Metric 1-node base and peak) (4 records)
- 4-Socket (Large Metric 2-node base and peak) (4 records)
- 4-Socket (Large Metric 3-node base and peak) (4 records)
- 4-Socket (Large Metric 4-node base and peak) (4 records)

ThinkSystem SR665 (**9** current world records):

- 2-Socket (Large Metric 6-node base and base) (2 records)
- 2-Socket (Large Metric 5-node base and base) (2 records)
- 2-Socket (Large Metric 4-node base)
- 2-Socket (Large Metric 3-node base)
- 2-Socket (Large Metric 2-node base)
- 2-Socket (Large Metric 1-node base and peak) (2 records)

ThinkSystem SR655 (**4** current world records):

- 1-Socket (2 records Large Metric 2-node base and peak) (2 records)
- 1-Socket (2 records Large Metric 1-node base and peak) (2 records)

About this benchmark: The SPECmpiM 2007 benchmark suite is used to evaluate MPI-parallel, floating point, compute-intensive performance across a wide range of cluster and SMP hardware. This suite gives users the most objective and representative benchmark suite for measuring and comparing high-performance computer systems.

Why it matters: If you are running compute-intensive technical workloads, a leadership benchmark score means this server is the highest performing server in its class, taking into account the CPUs, MPI library, communication interconnect, memory architecture, compilers, and file system performance.

For more information see the Performance Benchmark Reports section.

SPEC ACCEL

The **20** ThinkSystem world record results on the SPEC ACCEL benchmark include:

Open MP

Think System SR860 V2 (**2** current world records)

- SR860 V2 4S

ThinkSystem SR655 (**2** current world record):

- 1-Socket/1-Node (1S/1N)

Open CL

Think System SR860 V2 (**2** current world records)

- SR860 V2 4S

ThinkSystem SR665 (**4** current world records):

- 2-Socket/1-Node (2S/1N)

ThinkSystem SR655 (**2** current world record):

- 1-Socket/1-Node (1S/1N)

Open ACC

Think System SR860 V2 (**2** current world records)

- SR860 V2 4S

ThinkSystem SR665 (**4** current world records):

- 2-Socket/1-Node (2S/1N)

ThinkSystem SR655 (**4** current world record):

- 1-Socket/1-Node (1S/1N)

About this benchmark: The SPEC ACCEL benchmark suite tests performance with computationally intensive parallel applications running under the OpenCL, OpenACC, and OpenMP 4 target offloading APIs. The suite exercises the performance of the accelerator, host CPU, memory transfer between host and accelerator, support libraries and drivers, and compilers.

Why it matters: If you are running servers with accelerators (GPUs, coprocessors), a #1 benchmark score means a solution incorporating this server (equipped with a specific accelerator and supporting software) is the highest performing solution in its class.

For more information see the Performance Benchmark Reports section.

MLPerf

Lenovo ThinkSystem SR650 V2, SR670 V2, SD650 V2, SD650-N V2 and SE450 servers delivered **137** current world records on MLPerf Inference v2.0, v2.1, v3.0 benchmarks designed to test Artificial Intelligence inference performance, that is, to measure how fast systems can process inputs and produce results using a trained model.

MLPerf is a consortium of industry-leading Artificial Intelligence organizations who share the goal of creating fair and meaningful benchmarks for measuring performance of Machine Learning hardware, software and services. Due to it's extensive support, MLPerf is fast becoming the machine learning benchmark of choice for the industry.

The world record results on the MLPerf benchmark include:

MLPerf v0.2 Inference Closed Datacenter Division: https://mlcommons.org/en/inference-datacenter-20/

SR670 V2 2S with 8x80GB A100 PCIe (#2.0-066) Fence by Accelerator type and number (**5** records)

- Image Classification, 99.0% Accuracy, Server Scenario
- Object Detection (large), 99.0% Accuracy, Server Scenario
- Speech-to-Text, 99.0% Accuracy, Server Scenario
- Speech-to-Text, 99.0% Accuracy, Offline Scenario
- Natural Language Processing, 99.0% Accuracy, Server Scenario

SR670 V2 2S with 4x80GB A100 SXM (#2.0-067) Fence by Accelerator type and number (**5** records)

- Image Classification, 99.0% Accuracy, Offline Scenario
- Medical Imaging, 99.0% Accuracy, Offline Scenario
- Medical Imaging, 99.9% Accuracy, Offline Scenario

- Speech-to-Text, 99.0% Accuracy, Server Scenario
- Speech-to-Text, 99.0% Accuracy, Offline Scenario

SR650 V2 2S with 2xA16 GPU (2 physical GPUs, but 8 logical ones) (#2.0-068) Fence by Accelerator type and number (**10** records)

- Image Classification, 99.0% Accuracy, Offline Scenario
- Image Classification, 99.0% Accuracy, Server Scenario
- Object Detection (large), 99.0% Accuracy, Offline Scenario
- Object Detection (large), 99.0% Accuracy, Server Scenario
- Medical Imaging, 99.0% Accuracy, Offline Scenario
- Medical Imaging, 99.9% Accuracy, Offline Scenario
- Speech-to-Text, 99.0% Accuracy, Server Scenario
- Speech-to-Text, 99.0% Accuracy, Offline Scenario
- Natural Language Processing, 99.0% Accuracy, Offline Scenario
- Natural Language Processing, 99.0% Accuracy, Server Scenario

SD650 V2 with 2xIntel(R) Xeon(R) Platinum 8380 CPU (no accelerator) (#2.0-070) Fence by No accelerator and 2xCPU (**5** records)

- Image Classification, 99.0% Accuracy, Offline Scenario
- Image Classification, 99.0% Accuracy, Server Scenario
- Object Detection (large), 99.0% Accuracy, Offline Scenario
- Object Detection (large), 99.0% Accuracy, Server Scenario
- Natural Language Processing, 99.0% Accuracy, Offline Scenario

MLPerf v2.0 Inference Closed Edge Division: https://mlcommons.org/en/inference-edge-20/

SE450 2S with 2x80GB A100 PCIe (#2.0-071) Fence by Accelerator type and number (**5** records)

- Image Classification, 99.0% Accuracy, Single Stream;
- Object Detection (small), 99.0% accuracy, Single Stream
- Object Detection (small), 99.0% accuracy, Offline
- Object Detection (large), 99.0% accuracy, Single Stream
- Speech-to-text, 99.0% accuracy, Single Stream

SE450 2S with 2xA30 (#2.0-072) Fence by Accelerator type and number (**17** records)

- Image Classification, 99.0% Accuracy, Single Stream;
- Image Classification, 99.0% Accuracy, Multi Stream;
- Image Classification, 99.0% Accuracy, Offline Stream;
- Object Detection (small), 99.0% accuracy, Single Stream
- Object Detection (small), 99.0% accuracy, Multi Stream
- Object Detection (small), 99.0% accuracy, Offline
- Object Detection (large), 99.0% accuracy, Single Stream
- Object Detection (large), 99.0% accuracy, Multi Stream
- Object Detection (large), 99.0% accuracy, Offline

- Medical Imaging, 99.0% accuracy, Single Stream
- Medical Imaging, 99.0% accuracy, Offline Stream
- Medical Imaging, 99.9% accuracy, Single Stream
- Medical Imaging, 99.9% accuracy, Offline Stream
- Speech-to-text, 99.0% accuracy, Single Stream
- Speech-to-text, 99.0% accuracy, Offline
- Natural Language Processing, 99.0% accuracy, Single Stream
- Natural Language Processing, 99.0% accuracy, Offline

MLPerf v2.0 Training: https://mlcommons.org/en/training-normal-20/

SR670 V2 2S with 4x80GB A100 SXM (#2.0-2076) Fence by Accelerator type and number (**1** records)

- Image Segmentation

SD650 V2-N with 4x80GB A100 SXM (#2.0-2077) Fence by Accelerator type and number (**1** records)

- Object detection, heavy-weight

MLPerf v2.1 Inference Closed Datacenter Division: https://mlcommons.org/en/inference-datacenter-21/

SR670v2 2S with 8x80GB A100 PCIe (#2.1-0079) Fence by Accelerator type and number (**2** records)

- Image Classification, 99% Accuracy, Offline
- Speech-to-text, 99% accuracy, Offline

SE450 2S with 2x80GB A100 PCIe (#2.1-0080) Fence by Accelerator type and number (**2** records)

- Image Classification, 99% Accuracy, Server Scenario
- Image Classification, 99% Accuracy, Offline

MLPerf v2.1 Training: https://mlcommons.org/en/training-normal-21/

SR670v2 2S with 8x80GB A100 PCIe (#2.1-2058) Fence by Accelerator type and number (**2** records)

- Image Classification, 99.0% Accuracy, Server Scenario
- Object detection, heavy-weight

SR670v2 2S with 4x40GB A100 SXM4 (#2.1-2057) Fence by Accelerator type and number (**3** records)

- Object detection, heavy-weight
- Natural language processing (NLP)
- Recommendation

MLPerf v3.0 Training Closed Datacenter Division: https://mlcommons.org/benchmarks/training

SR670v2 2S with 8x80GB A100 PCIe (#3.0-2019) Fence by Accelerator type and number (**1** record)

- Image Classification, 99% Accuracy, Offline

MLPerf v3.0 Inference Closed Datacenter Division: https://mlcommons.org/en/inference-datacenter-30/

SR670v2 2S with 8x80GB A100 PCIe (#3.0-0062) Fence by Accelerator type and number (**7** records)

- Image Classification, 99.0% Accuracy, Server Scenario
- Medical imaging, 99.00% Accuracy, Offline
- Medical imaging, 99.90% Accuracy, Offline
- Speech-to-text, 99.00% Accuracy, Server Scenario
- Speech-to-text, 99.00% Accuracy, Offline
- Natural Language Processing, 99.00% Accuracy, Server scenario
- Natural Language Processing, 99.90% Accuracy, Server scenario

MLPerf v3.0 Inference Closed Edge Division: https://mlcommons.org/en/inference-edge-30/
SE450 with 4x QAIC100 Lite (#3.0-0064) Fence by Accelerator type and number (**8** records)

- Image classification, 99% Accuracy, Single stream scenario
- Image classification, 99% Accuracy, Multi stream scenario
- Image classification, 99% Accuracy, Offline
- Object detection, 99% Accuracy, Single stream scenario
- Object detection, 99% Accuracy, Multi stream scenario
- Object detection, 99% Accuracy, Offline
- Natural Language Processing, 99% Accuracy, Single stream scenario
- Natural Language Processing, 99% Accuracy, Offline

SE350 with 1x QAIC100 Pro (#3.0-0063) Fence by Accelerator type and number (**8** records)

- Processing, 99% Accuracy, OfflineImage classification, 99% Accuracy, Single stream scenario
- Image classification, 99% Accuracy, Multi stream scenario
- Image classification, 99% Accuracy, Offline
- Object detection, 99% Accuracy, Single stream scenario
- Object detection, 99% Accuracy, Multi stream scenario
- Object detection, 99% Accuracy, Offline
- Natural Language Processing, 99% Accuracy, Single stream scenario
- Natural Language, 99.9% Accuracy, Single stream scenario

MLPerf v3.1 Training Closed Datacenter Division: https://mlcommons.org/benchmarks/training

SR675 V3 Server with 4x96GB SXM5 H100 (#3.1-2043) Fence by Accelerator type and number (**6** records

- Image Classification (ResNet50)
- Medical Image Segmentation (3D-UNet)
- Object Detection, heavy weight (Mask R-CNN)
- Object Detection, light weight (SSD)
- Speech Recognition (RNN-T)
- Natural Language Processing (BERT-large)

MLPerf v3.1 Inference Closed Datacenter Division: MLPerf v3.1 Training Closed Datacenter Division: https://mlcommons.org/benchmarks/inference-datacenter

SR675v3 2S with 8x80GB H100 PCIe (#3.1-0102) Fence by Accelerator type and number (**3** records)

- Natural Language Processing, 99% Accuracy, Offline (Bert)
- Object detection, 99% Accuracy, Offline (RetinaNet)
- Speech-to-text, 99.0% Accuracy, Server Scenario (RNN-T)

SR665v1 2S with 5x QAIC100 Pro (#3.1-0103) Fence by Accelerator type and number (**8** records)

- Natural Language Processing, 99% Accuracy, Offline (Bert)
- Natural Language Processing, 99% Accuracy, Server scenario (Bert)
- Natural Language Processing, 99.9% Accuracy, Offline (Bert)
- Natural Language Processing, 99.9% Accuracy, Server scenario (Bert)
- Image Classification, 99.0% Accuracy, Offline (ResNet)
- Image Classification, 99.0% Accuracy, Server scenario (ResNet)
- Object detection, 99% Accuracy, Offline (RetinaNet)
- Object detection, 99% Accuracy, Server scenario (RetinaNet)

MLPerf v3.1 Inference Closed Edge Division: https://mlcommons.org/benchmarks/inference-edge

SE450 with 2x L40-48GB-PCIe (#3.1-0104) Fence by Accelerator type and number (**12** records)

- Medical imaging, 99.0% Accuracy, Offline (3D-UNet)
- Medical imaging, 99.0% Accuracy, Single stream scenario (3D-UNet)
- Medical imaging, 99.9% Accuracy, Offline (3D-UNet)
- Medical imaging, 99.9% Accuracy, Single stream scenario (3D-UNet)
- Natural Language Processing, 99% Accuracy, Offline (Bert)
- Natural Language Processing, 99% Accuracy, Single stream scenario (Bert)
- Image Classification, 99.0% Accuracy, Offline (ResNet)
- Image Classification, 99.0% Accuracy, Multi stream scenario (ResNet)
- Image Classification, 99.0% Accuracy, Single stream scenario (ResNet)
- Object detection, 99% Accuracy, Offline (RetinaNet)
- Object detection, 99% Accuracy, Multi stream (RetinaNet)
- Object detection, 99% Accuracy, Single stream (RetinaNet)

SE450 with 4x QAIC100 Standard (#3.1-0105) Fence by Accelerator type and number (**3** records)

- Natural Language Processing, 99.0% Accuracy, Single stream scenario (Bert)
- Image Classification, 99.0% Accuracy, Offline (ResNet)
- Object detection, 99% Accuracy, Offline (RetinaNet)

MLPerf v4.0 Inference Closed Datacenter Division: https://mlcommons.org/benchmarks/inference-datacenter/
SE455 with 4x48GB NVidia L40 PCIe (#4.0-0058) Fence by Accelerator type and number (**10** records)

- Medical Imaging, 99.0% Accuracy, Offline Scenario (3d-unet)
- Medical Imaging, 99.9% Accuracy, Offline Scenario (3d-unet)
- Natural language processing, 99.9% Accuracy, Offline Scenario (Bert)

- Natural language processing, 99.9% Accuracy, Server Scenario (Bert)
- Image classification, 99% Accuracy, Offline Scenario (ResNet50-v1.5)
- Image classification, 99% Accuracy, Server Scenario (ResNet50-v1.5)
- Object Detection, 99% Accuracy, Offline Scenarion (RetinaNet)
- Object Detection, 99% Accuracy, Server Scenarion (RetinaNet)
- Speech-to-Text, 99% Accuracy, Offline Scenario (RNNT)
- Speech-to-Text, 99% Accuracy, Server Scenario (RNNT)

SR670v2 2S with 8x QualCOMM QAIC –100 Ultra (#4.0-0059) Fence by Accelerator type and number (**2** records)

- Natural language processing, 99.0% Accuracy, Offline Scenario (Bert)
- Natural language processing, 99.9% Accuracy, Server Scenario (Bert)

MLPerf v4.0 Inference Closed Edge Division: https://mlcommons.org/en/inference-edge/

SE450 with 2x NVidia L40 (#4.0-0057) Fence by Accelerator type and number (**11** records)

- Medical Imaging, 99.0% Accuracy, Offline Scenario (3d-unet)
- Medical Imaging, 99.0% Accuracy, Single Stream Scenario (3d-unet)
- Medical Imaging, 99.9% Accuracy, Offline Scenario (3d-unet)
- Medical Imaging, 99.9% Accuracy, Single Stream Scenario (3d-unet)
- Natural language processing, 99.0% Accuracy, Offline Scenario (Bert)
- Natural language processing, 99.0% Accuracy, Singel Stream Scenario (Bert)
- Image classification, 99% Accuracy, Single Stream Scenario (ResNet50-v1.5)
- Image classification, 99% Accuracy, Multi Stream Scenario (ResNet50-v1.5)
- Image classification, 99% Accuracy, Offline Scenario (ResNet50-v1.5)
- Speech-to-Text, 99% Accuracy, Single Stream Scenario (RNNT)
- Speech-to-Text, 99% Accuracy, Offline Scenario (RNNT)

About this benchmark: The MLPerf inference benchmark measures how fast a system can perform machine learning (ML) inference using a trained model. The MLPerf inference benchmark is intended for a wide range of systems from mobile devices to servers. MLPerf Inference is a benchmark suite for measuring how fast systems can process inputs and produce results using a trained model.

Why it matters: If you are running machine learning (ML) or artificial intelligence (AI) workloads using trained models, a #1 benchmark score means this server is the highest performing server in its class for how fast systems can train models to a target quality metric.

For more information see the Performance Benchmark Reports section.

SPEChpc 2021

The **114** ThinkSystem world record results on the SPEChpc 2021 benchmark include:

Tiny Suite (62 records)

- ThinkSystem SR655 V3 1S 1-node 1GPU ACC base and peak (**2** current world records)
- ThinkSystem SR655 V3 1S 1-node 2GPU ACC base and peak (**2** current world records)
- ThinkSystem SR675 V3 2S 1-node 1GPU ACC base and peak (**2** current world records)
- ThinkSystem SR675 V3 2S 1-node 2GPU ACC base and peak (**2** current world records)

- ThinkSystem SR675 V3 2S 1-node 3GPU ACC base and peak (**2** current world records)
- ThinkSystem SR675 V3 2S 1-node 4GPU ACC base and peak (**2** current world records)
- ThinkSystem SR675 V3 2S 1-node 5GPU ACC base and peak (**2** current world records)
- ThinkSystem SR675 V3 2S 1-node 6GPU ACC base and peak (**2** current world records)
- ThinkSystem SR675 V3 2S 1-node 7GPU ACC base and peak (**2** current world records)
- ThinkSystem SR675 V3 2S 1-node 8GPU ACC base and peak + Best 1 node ACC base and peak (**2** current world records)
- ThinkSystem SR655 V3 1S 1-node Peak MPI (**1** current world records)
- ThinkSystem SR655 V3 1-node 1S OMP base and peak (**2** current world records)
- ThinkSystem SR655 V3 2-node 1S OMP base and peak (**2** current world records)
- ThinkSystem SR665 V3 1-node 2S MPI base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 1-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 2-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 3-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 4-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 5-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 6-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR860 V3 1-node 4S MPI base and peak + Best 1 node MPI base and peak (**4** current world records)
- ThinkSystem SR860 V3 1-node 4S OMP base and peak + Best 1 node MPI base and peak (**4** currer world records)
- ThinkSystem SR950 8S 1-node 8 CPU OMP base (**1** current world records)

Small Suite (40 records)

- ThinkSystem SR675 V3 2S 1-node 7GPU ACC base and peak (**2** current world records)
- ThinkSystem SR655 V3 1S 1-node MPI peak (**1** current world records)
- ThinkSystem SR655 V3 1S 1-node OMP base and peak (**2** current world records)
- ThinkSystem SR655 V3 1S 2-node OMP base and peak (**2** current world records)
- ThinkSystem SR665 V3 2S 1-node MPI base/peak (**2** current world records)
- ThinkSystem SR665 V3 2S 1-node OMP base/peak (**2** current world records)
- ThinkSystem SR665 V3 2S 2-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 2S 3-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 2S 4-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 2S 5-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 2S 6-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR860 V3 4S 1-node MPI base and peak + Best 1 node MPI base and peak (**4** current world records)
- ThinkSystem SR860 V3 4S 1-node OMP base and peak + Best 1 node MPI base and peak (**4** currer world records)
- ThinkSystem SR950 8S 1-node 8 CPU OMP Base (**1** current world records)

Medium Suite (12 records)

- ThinkSystem SR665 V3 4-node 2S OMP base/peak and base/peak (**4** current world records)

- ThinkSystem SR665 V3 5-node 2S OMP base/peak and base/peak (**4** current world records)
- ThinkSystem SR665 V3 6-node 2S OMP base/peak and base/peak (**4** current world records)

About this benchmark: The SPEChpc 2021 benchmark was designed to provide a comprehensive measure of real-world performance for High Performance Computing (HPC) systems. Offering science and engineering codes that are representative of HPC workloads and are portable across CPU and accelerators, the benchmark includes four suites, Tiny, Small, Medium, and Large, enabling fair vendor-neutral comparisons of the performance of different HPC systems, ranging from a single node to hundreds of nodes with support multiple programming models, including MPI, MPI+OpenACC, MPI+OpenMP, and MPI+OpenMP with target offload.

Why it matters: 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 Suites 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.

For more information see the Performance Benchmark Reports section.

Conclusion

The portfolio of Lenovo ThinkSystem servers continues its dominance of data center performance, with **555** world record benchmarks (as of April 1, 2024). This outstanding performance was achieved across multiple configurations and a variety of workloads and industry benchmarks.

To learn more about ThinkSystem servers, go to

Performance Benchmark Reports

Each ThinkSystem benchmark has a Lenovo Performance Benchmark Report. These Performance Benchmark Reports detail the specific benchmark, benchmark result, and hardware/software configuration used for that benchmark result. The reports also provide a link to the specific benchmark results page (e.g., SPEC, TPC, SAP, etc.). View all Performance Benchmark Reports.

The following Lenovo ThinkSystem benchmark world records are current as of April 1, 2024.

- SAP SD 2T Reports
- SAP BWoH Reports
- TPC-E Reports
- SPECvirt Datacenter 2021 Reports
- SPECvirt_sc2013 Reports
- SPEC CPU 2017 Reports
- SPECpower_ssj2008 Reports
- STAC-M3 Reports
- SPEC ACCEL Reports
- SPEC OMP2012 Reports
- SPEC MPI2007 Reports
- SPECjbb2015 (Linux) Reports
- SPECjbb2015 (Windows) Reports
- MLPerf Reports
- SPEChpc2021 Reports

SAP SD 2T Reports

- FSC https://www.sap.com/dmc/benchmark/2023/Cert23072.pdf
- SR655 1S - https://www.sap.com/dmc/benchmark/2021/Cert21070.pdf (1 record)

SAP BWoH Reports

- SR950 4S 10.4B Records v3 - https://www.sap.com/dmc/benchmark/2019/Cert19014.pdf (1 record, 3 KPIs)
- SR860 V2 4S 1.3B Records v3 - https://www.sap.com/dmc/benchmark/2020/Cert20036.pdf (1 record 3 KPIs)
- SR860 V2 4S 5.2B Records v3 - https://www.sap.com/dmc/benchmark/2021/Cert21017.pdf (1 record 2 KPIs)
- SR860 V2 4S 7.8B Records v3 - https://www.sap.com/dmc/benchmark/2021/Cert21049.pdf (1 record 2 KPIs)
- ThinkAgile HX7821 4S 24.7B Records v3 - https://www.sap.com/dmc/benchmark/2020/Cert20026.pdf (1 record, 3KPIs)
- ThinkAgile VX 4U Certified Node 8S 41.6B Records Virtualized v3 - https://www.sap.com/dmc/benchmark/2022/Cert22026.pdf (1 Overall record, 3 KPIs)
- ThinkAgile VX 4U Certified Node 8S 41.6B Records Bare-Metal v3 - https://www.sap.com/dmc/benchmark/2022/Cert22025.pdf (1 Overall record, 3 KPIs)
- SR650 V2 2S 1.3B Records v3 - https://www.sap.com/dmc/benchmark/2021/Cert21035.pdf (1 Overal record, 3 KPIs)
- SR650 V2 2S 5.2B Records v3 - https://www.sap.com/dmc/benchmark/2021/Cert21053.pdf (1 Overal record, 3 KPIs)
- ThinkAgile VX650 V3 Certified Node 2S 7.8B Records Virtualized v3 - https://www.sap.com/dmc/benchmark/2023/Cert23030.pdf (1 record, KPI)

TPC-E Reports

- SR665 2S - http://www.tpc.org/4090 (3 record - 16-core performance and 2P & Overall price/performance)
- SR655 1S - http://tpc.org/4089 (2 records - 64-core performance and 64-core price/performance)
- SR650 V3 2S EMR - http://tpc.org/4095 (3 WRs- 2P & Overall performance and 128-core price/performance)
- SR655 V3 1S Bergamo - http://tpc.org/4094 (1 WRs- 1P performance)
- SR655 V3 1S Genoa - http://tpc.org/4093 (2 records- 1P & Overall performance and 1P price/performance)
- SR650 V3 2S SPR http://tpc.org/4091 (2 records- 2P & performance and 120-core price/performance)

SPECvirt Datacenter 2021 Reports

- SR665 2S 4-node - https://www.spec.org/virt_datacenter2021/results/res2021q3/virt_datacenter2021-20210809-00002-perf.html
- SR665 v3 4-node - https://www.spec.org/virt_datacenter2021/results/res2023q4/virt_datacenter2021-20231006-00010-perf.html
(2 records - best 2 CPU and best overall score)

SPECvirt_sc2013 Reports

Performance

- SR950 4S - http://www.spec.org/virt_sc2013/results/res2019q2/virt_sc2013-20190312-00118-perf.html (1 record-Best 4 CPU)
- SR950 8S - http://www.spec.org/virt_sc2013/results/res2019q2/virt_sc2013-20190611-00119-perf.html (2 records-Best 8 CPU and Best overall)

Performance Per Watt

- SR650 2S - http://www.spec.org/virt_sc2013/results/res2017q4/virt_sc2013-20171018-00102-ppw.html (1 record-Best 2 CPU)
- SR950 4S - http://spec.org/virt_sc2013/results/res2019q2/virt_sc2013-20190312-00118-ppw.html (2 records-Best 4 CPU and Best overall)
- SR950 8S - http://www.spec.org/virt_sc2013/results/res2019q2/virt_sc2013-20190611-00119-ppw.html (1 record- Best 8 CPU)

Server Performance Per Watt

- SR650 2S - http://www.spec.org/virt_sc2013/results/res2017q4/virt_sc2013-20171018-00102-ppws.html (1 record- Best 2 CPU)
- SR950 4S - http://www.spec.org/virt_sc2013/results/res2019q2/virt_sc2013-20190312-00118-ppws.html (2 records-Best 4 CPU and Best overall)
- SR950 8S - http://www.spec.org/virt_sc2013/results/res2019q2/virt_sc2013-20190611-00119-ppws.html (1 record- Best 8 CPU)

SPEC CPU 2017 Reports

- SR950 6S SPEC_speed_int_base2017 - http://spec.org/cpu2017/results/res2019q2/cpu2017-20190401-11612.html
- SR950 3S SPEC_speed_int_base2017 - http://spec.org/cpu2017/results/res2019q2/cpu2017-20190401-11610.html
- SR950 3S SPEC_speed_fp_base2017 - https://www.spec.org/cpu2017/results/res2019q2/cpu2017-20190319-11383.html
- SR950 3S SPEC_rate_int_base2017 - https://www.spec.org/cpu2017/results/res2019q2/cpu2017-20190319-11385.html
- SR950 3S SPEC_rate_fp_base2017 - https://www.spec.org/cpu2017/results/res2019q2/cpu2017-20190319-11381.html
- SR860 V2 4S SPECspeed_int_base_energy2017 - https://spec.org/cpu2017/results/res2020q4/cpu2017-20201026-24300.html
- SR860 V3 4S SPECspeed_fp_base_energy2017 -https://spec.org/cpu2017/results/res2023q3/cpu2017-20230731-38300.html
- SR860 V3 4S SPECrate_int_base_energy2017 - https://spec.org/cpu2017/results/res2023q3/cpu2017-20230731-38301.html
- SR860 V3 4S SPECrate_fp_base_energy2017 - https://spec.org/cpu2017/results/res2023q3/cpu2017-20230731-38302.html
- SR950 V3 8S SPECspeed_int_base2017 - https://spec.org/cpu2017/results/res2023q4/cpu2017-20231204-40187.html
- SR655 V3 1S SPECcpu rate_int_base_energy_2017 - https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36762.html
- SR655 V3 1S SPECcpu rate_int_peak_energy_2017 - https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36762.html
- SR655 V3 1S SPECcpu rate_fp_base_energy_2017 -

https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36764.html

- SR655 V3 1S SPECcpu rate_fp_peak_energy2017 - https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36764.html
- SR655 V3 1S SPECcpu speed_fp_base_energy - https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36761.html (1→ 2 records - best 1 CPU and best overall score)
- SR655 V3 1S SPECcpu speed_fp_peak_energy - https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36761.html (1→ 2 records - best 1 CPU and best overall score)
- SR665 V3 2S SPECcpu speed_int_base_energy_2017 - https://spec.org/cpu2017/results/res2023q1/cpu2017-20230116-33599.html
- SR665 V3 2S SPECcpu speed_int_peak_energy_2017 - https://spec.org/cpu2017/results/res2023q1/cpu2017-20230116-33599.html
- SR665 V3 2S SPEC_speed_fp_base_energy_2017 - https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36752.html
- SR665 V3 2S SPEC_speed_fp_peak_energy_2017 - https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36748.html
- SR665 V3 2S SPEC_rate_int_base_energy_2017 - https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36749.html (1→ 2 records - best 2 CPU and best overall score)
- SR665 V3 2S SPEC_rate_int_peak_energy_2017- https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36749.html (1→ 2 records - best 2 CPU and best overall score)
- SR665 V3 2S SPEC_rate_fp_base_energy_2017 - https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36877.html (1→ 2 records - best 2 CPU and best overall score)
- SR665 V3 2S SPEC_rate_fp_peak_energy_2017 - https://spec.org/cpu2017/results/res2023q2/cpu2017-20230522-36877.html (1→ 2 records - best 2 CPU and best overall score)

SPECpower_ssj2008 Reports

- SN850 4S 7-node https://www.spec.org/power_ssj2008/results/res2018q3/power_ssj2008-20180828-00852.html (Windows) 2 WRs (best 7node, best 7node 28cpu linux)
- SN850 4S 7-node https://www.spec.org/power_ssj2008/results/res2019q3/power_ssj2008-20190731-00999.html (Linux) 1 WR (best 4cpu 2u Linux score)
- SN550 14 nodes - https://www.spec.org/power_ssj2008/results/res2017q4/power_ssj2008-20171011-00794.html (Windows) 2WRs (best 14node score,best,14node 28cpu windows score)
- SR950 8S - https://www.spec.org/power_ssj2008/results/res2018q3/power_ssj2008-20180828-00853.html (Windows) 1WR (best 8cpu 4u windows score)
- SR950 8S - https://www.spec.org/power_ssj2008/results/res2019q2/power_ssj2008-20190312-00928.html (Linux) 2WRs (best 8cpu 4u linux score, best 8cpu 4u score)
- SN550 V2 14 nodes - https://www.spec.org/power_ssj2008/results/res2022q2/power_ssj2008-20220426-01172.html (Linux) 1WRs (best 14node score,best 14node 28cpu windows score)
- SR645V3 2S(WINDOWS) https://spec.org/power_ssj2008/results/res2023q3/power_ssj2008-20230523-01266.html 3 WRs (best 2cpu 1u windows score, best 2cpu 1u score, best 1u score)
- SR645V3 2S(LINUX) https://spec.org/power_ssj2008/results/res2023q3/power_ssj2008-20230523-01265.html 1 WR (best 2 cpu 1u linux score)
- SR655V3 1S(WINDOWS) https://spec.org/power_ssj2008/results/res2023q3/power_ssj2008-

20230523-01269.html 1 WR (best 1cpu 2u windows score)

- SR655V3 1S(LINUX) https://spec.org/power_ssj2008/results/res2023q3/power_ssj2008-20230524-01270.html 2 WRs (best 1cpu 2u Linux score, best 1cpu 2u score)
- SR635V3 1S(WINDOWS) https://spec.org/power_ssj2008/results/res2023q3/power_ssj2008-20230523-01267.html 1 WR (best 1cpu 2u windows score)
- SR635V3 1S(LINUX) https://spec.org/power_ssj2008/results/res2023q3/power_ssj2008-20230523-01268.html 1WR (best 1cpu 2u linux score)
- SR850V3 4S(WINDOWS) https://spec.org/power_ssj2008/results/res2024q1/power_ssj2008-20240208-01366.html 1 WR best 4cpu 2u windows score
- SR850V3 4S(LINUX) https://spec.org/power_ssj2008/results/res2024q1/power_ssj2008-20240208-01367.html 1 WR best 4cpu 2u linux score
- SR860V3 4S(WINDOWS) https://spec.org/power_ssj2008/results/res2024q2/power_ssj2008-20240304-01379.html 1 WR best 4cpu 4u windows score
- SR860V3 4S(LINUX) https://spec.org/power_ssj2008/results/res2024q2/power_ssj2008-20240304-01378.html 1 WR best 4cpu 4u linux score

STAC-M3 Reports

- SR950 4S - https://stacresearch.com/KDB171024 combined with https://stacresearch.com/KDB170629 (**16** records)
- SR860 V2 4S - https://stacresearch.com/KDB201109 (Single-Node Antuco **3** records)
- SR860 V2 4S 2-node - https://stacresearch.com/KDB210428 (Antuco 8 records, 3 yr Kanaga 15 records)
- SR650 2S - https://stacresearch.com/KDB190320b (1 records)
- SR650 V2 - https://stacresearch.com/KDB210317 (4 records)
- SR650 V3 2S - https://stacresearch.com/KDB230707 (Single-Node 2S Antuco, 9 records, Sapphire Rapids)

SPEC ACCEL Reports

- SR665 2S OpenACC - https://www.spec.org/accel/results/res2021q1/accel-20210223-00152.html (Base)
- SR665 2S OpenACC - https://www.spec.org/accel/results/res2021q1/accel-20210223-00152.html (Peak)
- SR665 2S OpenCL - https://www.spec.org/accel/results/res2021q1/accel-20210223-00153.html (Base)
- SR665 2S OpenCL - https://www.spec.org/accel/results/res2021q1/accel-20210223-00153.html (Peak)
- SR655 1S OMP – https://www.spec.org/accel/results/res2021q2/accel-20210512-00158.html (base)
- SR655 1S OMP – https://www.spec.org/accel/results/res2021q2/accel-20210512-00158.html (peak)
- SR655 1S OpenCL - https://www.spec.org/accel/results/res2021q2/accel-20210512-00157.html (base)
- SR655 1S OpenCL - https://www.spec.org/accel/results/res2021q2/accel-20210512-00157.html (peak)
- SR655 1S OpenACC - https://www.spec.org/accel/results/res2021q2/accel-20210512-00156.html (base)
- SR655 1S OpenACC - https://www.spec.org/accel/results/res2021q2/accel-20210512-00156.html

(peak)

- SR860 V2 4S OpenMP -- https://www.spec.org/accel/results/res2020q4/accel-20200917-00149.html (base)
- SR860 V2 4S OpenMP -- https://www.spec.org/accel/results/res2020q4/accel-20200917-00149.html (peak)
- SR860 V2 4S OpenCL -- https://www.spec.org/accel/results/res2020q4/accel-20200917-00148.html (base)
- SR860 V2 4S OpenCL -- https://www.spec.org/accel/results/res2020q4/accel-20200917-00148.html (peak)
- SR860 V2 4S OpenACC -- https://www.spec.org/accel/results/res2020q4/accel-20200917-00147.html (base)
- SR860 V2 4S OpenACC -- https://www.spec.org/accel/results/res2020q4/accel-20200917-00147.html (peak)

SPEC OMP2012 Reports

- SR950 3S - http://spec.org/omp2012/results/res2019q2/omp2012-20190312-00163.html
- SR860 V3 4S - https://www.spec.org/omp2012/results/res2023q3/omp2012-20230719-00214.html (base)
- SR860 V3 4S - https://www.spec.org/omp2012/results/res2023q3/omp2012-20230719-00214.html (peak)
- SR665 V3 1S - https://www.spec.org/omp2012/results/res2023q3/omp2012-20230719-00215.html (energy_base)
- SR665 V3 1S - https://www.spec.org/omp2012/results/res2023q3/omp2012-20230719-00215.html (energy_peak)
- SR665 V3 2S - https://www.spec.org/omp2012/results/res2023q2/omp2012-20230517-00213.html (base)
- SR665 V3 2S - https://www.spec.org/omp2012/results/res2023q2/omp2012-20230517-00213.html (peak)
- SR665 V3 2S - https://www.spec.org/omp2012/results/res2023q3/omp2012-20230719-00216.html (energy_base)
- SR665 V3 2S - https://www.spec.org/omp2012/results/res2023q3/omp2012-20230719-00216.html (energy_peak)
- SR655 V3 1S - https://www.spec.org/omp2012/results/res2023q2/omp2012-20230516-00212.html (base)
- SR655 V3 1S - https://www.spec.org/omp2012/results/res2023q2/omp2012-20230516-00212.html (peak)

SPEC MPI2007 Reports

- SR950 (Medium Metric) 8S - https://www.spec.org/mpi2007/results/res2019q2/mpi2007-20190312-00619.html
- SR950 (Medium Metric) 6S - https://www.spec.org/mpi2007/results/res2019q2/mpi2007-20190312-00618.html
- SR950 (Medium Metric) 3S - https://www.spec.org/mpi2007/results/res2019q2/mpi2007-20190312-00616.html
- SR860 V2 4S (Medium Metric 1-Node) - https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20200917-00662.html (base)
- SR860 V2 4S (Medium Metric 1-Node) -- https://www.spec.org/mpi2007/results/res2020q4/mpi2007-

20200917-00662.html (peak)

- SR860 V2 4S (Medium Metric 2-Node) – https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00670.html (base)
- SR860 V2 4S (Medium Metric 2-Node) – https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00670.html (peak)
- SR860 V2 4S (Medium Metric 3-Node) – https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00665.html (base)
- SR860 V2 4S (Medium Metric 3-Node) – https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00665.html (peak)
- SR860 V2 4S (Medium Metric 4-Node) – https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00669.html (base)
- SR860 V2 4S (Medium Metric 4-Node) -- https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00669.html (peak)
- SR860 V2 4S (Large Metric 1-Node) -- https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20200917-00657.html (base)
- SR860 V2 4S (Large Metric 1-Node) -- https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20200917-00657.html (peak)
- SR860 V2 4S (Large Metric 2-Node) – https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00667.html (base)
- SR860 V2 4S (Large Metric 2-Node) – https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00667.html (peak)
- SR860 V2 4S (Large Metric 3-Node) – https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00668.html (base)
- SR860 V2 4S (Large Metric 3-Node) – https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00668.html (peak)
- SR860 V2 4S (Large Metric 4-Node) – https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00671.html (base)
- SR860 V2 4S (Large Metric 4-Node) -- https://www.spec.org/mpi2007/results/res2020q4/mpi2007-20201020-00671.html (peak)
- SR665 2S (Large Metric 6-Node) - https://www.spec.org/mpi2007/results/res2020q2/mpi2007-20200416-00641.html
- SR665 2S (Large Metric 5-Node) - https://www.spec.org/mpi2007/results/res2020q2/mpi2007-20200416-00646.html
- SR665 2S (Large Metric 4-Node) - https://www.spec.org/mpi2007/results/res2020q2/mpi2007-20200416-00645.html
- SR665 2S (Large Metric 3-Node) - https://www.spec.org/mpi2007/results/res2020q2/mpi2007-20200416-00644.html
- SR665 2S (Large Metric 2-Node) - https://www.spec.org/mpi2007/results/res2020q2/mpi2007-20200416-00643.html
- SR665 2S (Large Metric 1-node) - https://www.spec.org/mpi2007/results/res2021q1/mpi2007-20210223-00672.html (base)
- SR665 2S (Large Metric 1-node) - https://www.spec.org/mpi2007/results/res2021q1/mpi2007-20210223-00672.html (peak)
- SR655 1S (Large Metric 1-node) - https://www.spec.org/mpi2007/results/res2021q2/mpi2007-20210511-00683.html (base)
- SR655 1S (Large Metric 1-node) - https://www.spec.org/mpi2007/results/res2021q2/mpi2007-20210511-00683.html (peak)

- SR655 1S (Large Metric 2-node) - https://www.spec.org/mpi2007/results/res2021q2/mpi2007-20210511-00680.html (base)
- SR655 1S (Large Metric 2-node) - https://www.spec.org/mpi2007/results/res2021q2/mpi2007-20210511-00680.html (peak)
- SR665 2S (Medium Metric 3-node) – https://www.spec.org/mpi2007/results/res2021q1/mpi2007-20210223-00676.html (base)
- SR665 2S (Medium Metric 3-node) – https://www.spec.org/mpi2007/results/res2021q1/mpi2007-20210223-00676.html (peak)
- SR665 2S (Medium Metric 4-node) – https://www.spec.org/mpi2007/results/res2021q1/mpi2007-20210223-00675.html (peak)
- SR665 2S (Medium Metric 5-node) – https://www.spec.org/mpi2007/results/res2021q1/mpi2007-20210223-00674.html (base)
- SR665 2S (Medium Metric 5-node) – https://www.spec.org/mpi2007/results/res2021q1/mpi2007-20210223-00674.html (peak)
- SR665 2S (Medium Metric 6-node) – https://www.spec.org/mpi2007/results/res2020q2/mpi2007-20200416-00650.html (base and base)
- SR665 2S (Medium Metric 6-node) – https://www.spec.org/mpi2007/results/res2021q1/mpi2007-20210223-00678.html (peak and peak)
- SR655 V3 1S (Medium Metric 1 node) - https://www.spec.org/mpi2007/results/res2023q1/mpi2007-20230208-00690.html (base and peak)
- SR655 V3 1S (Medium Metric 2 node) - https://www.spec.org/mpi2007/results/res2021q2/mpi2007-20210511-00681.html (base and peak)
- SR665 V3 2S (Medium Metric 1 node) - https://www.spec.org/mpi2007/results/res2023q1/mpi2007-20230208-00689.html (base and peak)
- SR665 V3 2S (Medium Metric 2 node) - https://www.spec.org/mpi2007/results/res2023q1/mpi2007-20230208-00688.html (base and peak)

SPECjbb2015 (Linux) Reports

- SR860 V3 4S Distributed Critical-JOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230906-01171.html (2 records – best 4 CPU Linux and best 4 CPU)
- SR860 v3 4S Distributed Critical-jOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230906-01170.html (2 records – best 4 CPU Linux and best 4 CPU)
- SR860 V3 4S Composite Max-jOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230906-01168.html (2 records – best 4 CPU Linux and best 4 CPU)
- SR860 v3 4S Compostie Critical-jOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230906-01168.html (2 records – best 4 CPU Linux and best 4 CPU)
- SR860 V3 4S MultiJVM Critical-JOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230906-01169.html (2 records – best 4 CPU Linux and best 4 CPU)
- SR950 3S Distributed Max-JOPS - http://spec.org/jbb2015/results/res2018q3/jbb2015-20180829-00315.html (2 records – best 3 CPU Linux and best 3 CPU)
- SR950 3S Distributed Critical - http://spec.org/jbb2015/results/res2019q2/jbb2015-20190314-00401.html (2 records – best 3 CPU Linux and best 3 CPU)
- SR950 3S MultiJVM Max-JOPS - http://spec.org/jbb2015/results/res2019q2/jbb2015-20190314-00415.html (2 records – best 3 CPU Linux and best 3 CPU)
- SR950 3S MultiJVM Critical - http://spec.org/jbb2015/results/res2019q2/jbb2015-20190314-00429.html (2 records – best 3 CPU Linux and best 3 CPU)
- SR950 3S Composite Max-JOPS - http://spec.org/jbb2015/results/res2019q2/jbb2015-20190314-

00442.html (2 records – best 3 CPU Linux and best 3 CPU)

- SR950 3S Composite Critical - http://spec.org/jbb2015/results/res2018q3/jbb2015-20180829-00318.html (2 records – best 3 CPU Linux and best 3 CPU)
- SR950 6S Distributed Max-JOPS - http://spec.org/jbb2015/results/res2018q3/jbb2015-20180829-00311.html (2 records – best 6 CPU Linux and best 6 CPU)
- SR950 6S Distributed Critical - http://spec.org/jbb2015/results/res2018q3/jbb2015-20180829-00312.html (2 records – best 6 CPU Linux and best 6 CPU)
- SR950 6S MultiJVM Max-JOPS - http://spec.org/jbb2015/results/res2019q2/jbb2015-20190314-00421.html (2 records – best 6 CPU Linux and best 6 CPU)
- SR950 6S MultiJVM Critical - http://spec.org/jbb2015/results/res2019q2/jbb2015-20190314-00405.html (2 records – best 6 CPU Linux and best 6 CPU)
- SN550 V2 28S Distributed Max-JOPS - https://www.spec.org/jbb2015/results/res2022q2/jbb2015-20220504-00805.html
- SN550 V2 28S Distributed Critical JOPS - https://www.spec.org/jbb2015/results/res2022q2/jbb2015-20220504-00804.html
- SR665 V3 2S Distributed critical-jOPS: https://www.spec.org/jbb2015/results/res2023q4/jbb2015-20231115-01185.html (2 records – best 2 CPU Linux and best 2 CPU)
- SR665 V3 2S MultiJVM critical-jOPS: https://spec.org/jbb2015/results/res2024q1/jbb2015-20240124-01220.html (2 records – best 2 CPU Linux and best 2 CPU)

SPECjbb2015 (Windows) Reports

- SR655 v3 1S MultiJVM MaxjOPS: https://www.spec.org/jbb2015/results/res2023q1/jbb2015-20230120-01014.html
- SR655 v3 1S MultiJVM Critical-jOPS: https://www.spec.org/jbb2015/results/res2023q1/jbb2015-20230121-01015.html
- SR655 v3 1S Composite Max-jOPS: https://www.spec.org/jbb2015/results/res2023q1/jbb2015-20230120-01011.html
- SR655 v3 1S Composite Critical-jOPS: https://www.spec.org/jbb2015/results/res2023q1/jbb2015-20230120-01010.html
- SR655 v3 1S Distributed MaxjOPS: https://www.spec.org/jbb2015/results/res2023q1/jbb2015-20230120-01013.html
- SR655 v3 1S Distributed Critical-jOPS: https://www.spec.org/jbb2015/results/res2023q1/jbb2015-20230120-01012.html
- SR665 V3 2S MultiJVM MaxjOPS: https://www.spec.org/jbb2015/results/res2022q4/jbb2015-20221019-00895.html
- SR665 V3 2S MultiJVM Critical-jOPS: https://www.spec.org/jbb2015/results/res2022q4/jbb2015-20221019-00894.html
- SR665 V3 2S Distributed Max-JOPS: https://www.spec.org/jbb2015/results/res2022q4/jbb2015-20221019-00892.html
- SR665 V3 2S Distributed Critical JOPS: https://www.spec.org/jbb2015/results/res2022q4/jbb2015-20221019-00890.html
- SR665 V3 2S Composite Max-jOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230906-01172.html
- SR665 V3 2S Composite Max-jOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230906-01172.html
- SR860 V3 4S MultiJVM Max-JOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230726-01146.html

- SR860 V3 4S MultiJVM Critical-JOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230726-01145.html (2 records - best 4 CPU Windows and best 4 CPU)
- SR860 V3 4S Distributed Max-JOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230726-01147.html -(2 records – best 4 CPU Windows and best 4 CPU)
- SR860 V3 4S Distributed Critical-JOPS: https://www.spec.org/jbb2015/results/res2023q3/jbb2015-20230726-01144.html - (2 records – best 4 CPU Windows and best 4 CPU)
- SR950 6S MultiJVM Max-JOPS - http://spec.org/jbb2015/results/res2019q2/jbb2015-20190313-00379.html
- SR950 6S MultiJVM Critical - http://spec.org/jbb2015/results/res2019q2/jbb2015-20190313-00379.html
- SR950 8S MultiJVM Max-JOPS - http://spec.org/jbb2015/results/res2019q2/jbb2015-20190314-00440.html
- SR950 8S MultiJVM Critical-JOPS -http://spec.org/jbb2015/results/res2019q2/jbb2015-20190313-00380.html
- SN550 V2 28S Distributed Max-JOPS - https://www.spec.org/jbb2015/results/res2022q2/jbb2015-20220504-00807.html (2 records – best 14node 1CPU Windows and best 14node 1CPU)
- SN550 V2 28S Distributed Critical JOPS - https://www.spec.org/jbb2015/results/res2022q2/jbb2015-20220504-00806.html (2 records – best 14node 1CPU Windows and best 14node 1CPU)

MLPerf Reports

- SR670 V2 2 processors, 8x80GB A100 PCIe (#2.0-066) Fence by Accelerator type and number https://mlcommons.org/en/inference-datacenter-20/ (5 records)
- SR670 V2 2 processors, 4x80GB A100 SXM (#2.0-067) Fence by Accelerator type and number https://mlcommons.org/en/inference-datacenter-20/ (5 records)
- SR650 V2 2 processors, 2xA16 GPU (2 physical GPUs, but 8 logical ones) (#2.0-068) Fence by Accelerator type and number https://mlcommons.org/en/inference-datacenter-20/ (10 records)
- SD650 V2 with 2xIntel(R) Xeon(R) Platinum 8380 CPU (no accelerator) (#2.0-070) Fence by Accelerator type and number https://mlcommons.org/en/inference-datacenter-20/ (5 records)
- SE450 2 processors, 2x80GB A100 PCIe (#2.0-071) Fence by Accelerator type and number https://mlcommons.org/en/inference-edge-20/ (5 records)
- SE450 2 processors, 2xA30 (#2.0-072) Fence by Accelerator type and number https://mlcommons.org/en/inference-edge-20/ (17 records)
- SR670 V2 2 processors, 4x80GB A100 SXM (#2.0-2076) Fence by Accelerator type and number https://mlcommons.org/en/training-normal-20/ (1 records)
- SD650 V2-N 4x80GB A100 SXM (#2.0-2077) Fence by Accelerator type and number https://mlcommons.org/en/training-normal-20/ (1 records)
- SR670v2 2S with 8x80GB A100 PCIe (#2.1-0079) Fence by Accelerator type and number https://mlcommons.org/en/inference-datacenter-21/ (2 records)
- SE450 2S with 2x80GB A100 PCIe (#2.1-0080) Fence by Accelerator type and number https://mlcommons.org/en/inference-datacenter-21/ (2 records)
- SR670v2 2S with 8x80GB A100 PCIe (#2.1-2058) Fence by Accelerator type and number https://mlcommons.org/en/training-normal-21/ (2 records)
- SR670v2 2S with 4x40GB A100 SXM4 (#2.1-2057) Fence by Accelerator type and number https://mlcommons.org/en/training-normal-21/ (3 records)
- SR670v2 2S with 8x80GB A100 PCIe (#3.0-0062) Fence by Accelerator type and number https://mlcommons.org/en/inference-datacenter-30/ (7 records)
- SE450 with 4x QAIC100 Lite (#3.0-0064) Fence by Accelerator type and number

https://mlcommons.org/en/inference-edge-30/ (8 records)

- SE350 with 1x QAIC100 Pro (#3.0-0063) Fence by Accelerator type and number https://mlcommons.org/en/inference-edge-30/ (8 records)
- SR670v2 2S with 8x80GB A100 PCIe (#3.0-2019) Fence by Accelerator type and number https://mlcommons.org/benchmarks/training (1 record)
- SR675 V3 Server with 4x96GB SXM5 H100 (#3.1-2043) Fence by Accelerator type and number https://mlcommons.org/benchmarks/training (6 records)
- SR675v3 2S with 8x80GB H100 PCIe (#3.1-0102) Fence by Accelerator type and number https://mlcommons.org/benchmarks/inference-datacenter (3 records)
- SR665v1 2S with 5x QAIC100 Pro (#3.1-0103) Fence by Accelerator type and number https://mlcommons.org/benchmarks/inference-datacenter (8 records)
- SE450 with 2x L40-48GB-PCIe (#3.1-0104) Fence by Accelerator type and number https://mlcommons.org/benchmarks/inference-edge (12 records)
- SE450 with 4x QAIC100 Standard (#3.1-0105) Fence by Accelerator type and number https://mlcommons.org/benchmarks/inference-edge (3records)
- SE455 with 4x48GB NVidia L40 PCIe (#4.0-0058) Fence by Accelerator type and number https://mlcommons.org/benchmarks/inference-datacenter/ (10 records)
- SR670v2 2S with 8x QualCOMM QAIC –100 Ultra (#4.0-0059) Fence by Accelerator type and number number https://mlcommons.org/benchmarks/inference-datacenter/ (2 records)
- SE450 with 2x NVidia L40 (#4.0-0057) Fence by Accelerator type and number https://mlcommons.org/en/inference-edge/ (11 records)

SPEChpc2021 Reports

Tiny Suite

- 2 WRs - SR655 V3 https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00175.html (1S 1-node 1GPU ACC base and peak)
- 2 WRs - SR655 V3 https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00174.html (1S 1-node 2GPU ACC base and peak)
- 2 WRs – SR675 V3 https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00248.html (2S 1-node 1GPU ACC base and peak)
- 2 WRs – SR675 V3 https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00247.html (2S 1-node 2GPU ACC base and peak)
- 2 WRs – SR675 V3 https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00246.html (2S 1-node 3GPU ACC base and peak)
- 2 WRs – SR675 V3 https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00245.html (2S 1-node 4GPU ACC base and peak)
- 2 WRs – SR675 V3 https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00244.html (2S 1-node 5GPU ACC base and peak)
- 2 WRs – SR675 V3 https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00243.html (2S 1-node 6GPU ACC base and peak)
- 2 WRs – SR675 V3 https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00242.html (2S 1-node 7GPU ACC base and peak)
- 4 WRs – SR675 V3 https://spec.org/hpc2021/results/res2023q4/hpc2021-20230822-00241.html (2S 1-node 8GPU ACC base and peak + Best 1 node ACC base and peak)

- 2 WRs - SD650-N V2 https://spec.org/hpc2021/results/res2021q4/hpc2021-20210908-00022.html (2-node 2S 4GPU ACC base and base)
- 1 WRs - SR655 V3 https://spec.org/hpc2021/results/res2023q1/hpc2021-20230207-00189.html (1S 1-node Peak MPI)
- 4 WRs - SR665 V3 https://spec.org/hpc2021/results/res2022q4/hpc2021-20221016-00134.html (1-node 2S MPI base/peak and base/peak)
- 4 WRs – SR860 V3 https://spec.org/hpc2021/results/res2023q3/hpc2021-20230823-00252.html (1-node 4S MPI base and peak + Best 1 node MPI base and peak)
- 2 WRs - SR655 V3 https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00172.html (1-node 1S OMP base and peak)
- 2 WRs - SR655 V3 https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00173.html (2-node 1S OMP base and peak)
- 4 WRs - SR665 V3 https://spec.org/hpc2021/results/res2022q4/hpc2021-20221016-00135.html (1-node 1S OMP base/peak and base/peak)
- 4 WRs - SR665 V3 https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00184.html (2-node 2S OMP base/peak and base/peak)
- 4 WRs - SR665 V3 https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00181.html (3-node 2S OMP base/peak and base/peak)
- 4 WRs - SR665 V3 https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00185.html (4-node 2S OMP base/peak and base/peak)
- 4 WRs - SR665 V3 https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00182.html (5-node 2S OMP base/peak and base/peak)
- 4 WRs - SR665 V3 https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00183.html (6-node 2S OMP base/peak and base/peak)
- 4 WRs – SR860 V3 https://spec.org/hpc2021/results/res2023q3/hpc2021-20230823-00251.html (1-node 4S OMP base and peak + Best 1 node MPI base and peak)
- 1 WRs - SR950 https://spec.org/hpc2021/results/res2021q4/hpc2021-20210908-00025.html (8S 1-node OMP base)

Small Suite

- SR675 V3 2S 1-node 7 GPU ACC base and peak https://spec.org/hpc2021/results/res2023q4/hpc2021-20230905-00256.html
- SR655 V3 1S 1-node MPI peak - https://spec.org/hpc2021/results/res2023q1/hpc2021-20230207-00190.html
- SR665 V3 2S 1-node MPI base/peak and base/peak – https://spec.org/hpc2021/results/res2022q4/hpc2021-20221016-00133.html
- SR655 V3 1S 1-node OMP base and peak - https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00170.html
- SR655 V3 1S 2-node OMP base and peak - https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00171.html
- SR665 V3 2S 1-node 2S OMP base/peak and base/peak – https://spec.org/hpc2021/results/res2022q4/hpc2021-20221010-00132.html
- SR665 V3 2S 2-node 2S OMP base/peak and base/peak – https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00177.html
- SR665 V3 2S 3-node 2S OMP base/peak and

base/peak – https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00176.html

- SR665 V3 2S 4-node 2S OMP base/peak and base/peak – https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00180.html
- SR665 V3 2S 5-node 2S OMP base/peak and base/peak – https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00179.html
- SR665 V3 2S 6-node 2S OMP base/peak and base/peak – https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00178.html
- SR950 8S 1-node 8 CPU – https://spec.org/hpc2021/results/res2021q4/hpc2021-20210908-00024.html
- SR860 V3 4S 1-node MPI base/peak + Best 1node MPI base and peak https://spec.org/hpc2021/results/res2023q3/hpc2021-20230823-00249.html
- SR860 V3 4S 1-node OMP base/peak + Best 1node MPI base and peak https://spec.org/hpc2021/results/res2023q3/hpc2021-20230823-00250.html

Medium Suite

- SR665 V3 4-node 2S OMP base/peak and base/peak - https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00186.html
- SR665 V3 5-node 2S OMP base/peak and base/peak - https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00187.html
- SR665 V3 6-node 2S OMP base/peak and base/peak - https://spec.org/hpc2021/results/res2023q1/hpc2021-20230206-00188.html

References

- Standard Performance Evaluation Corporation

Related product families

Product families related to this document are the following:

- ThinkEdge SE450 Edge Server
- ThinkSystem SD530 Server
- ThinkSystem SD650-N V2 server
- ThinkSystem SE350 Edge Server
- ThinkSystem SN550 Server
- ThinkSystem SN550 V2 Server
- ThinkSystem SN850 Server
- ThinkSystem SR250 Server
- ThinkSystem SR635 V3 Server
- ThinkSystem SR645 Server
- ThinkSystem SR645 V3 Server
- ThinkSystem SR650 Server
- ThinkSystem SR650 V2 Server
- ThinkSystem SR650 V3 Server
- ThinkSystem SR655 Server
- ThinkSystem SR655 V3 Server
- ThinkSystem SR665 Server
- ThinkSystem SR665 V3 Server
- ThinkSystem SR670 Server
- ThinkSystem SR670 V2 Server
- ThinkSystem SR850 Server
- ThinkSystem SR850 V2 Server
- ThinkSystem SR860 Server
- ThinkSystem SR860 V2 Server
- ThinkSystem SR860 V3 Server
- ThinkSystem SR950 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, LP1145, was created or updated on April 11, 2024.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  https://lenovopress.lenovo.com/LP1145
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at https://lenovopress.lenovo.com/LP1145.

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®
ThinkAgile®
ThinkSystem®

The following terms are trademarks of other companies:

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.

Windows® is a trademark of Microsoft Corporation in the United States, other countries, or both.

TPC and TPC-E are trademarks of Transaction Processing Performance Council.

SPEC®, SPEC ACCEL®, SPEC CPU®, SPEC MPI®, SPEC OMP®, SPEC VIRT_SC®, SPEC VIRT®, SPEChpc™, SPECjbb®, and SPECpower_ssj® are trademarks of the Standard Performance Evaluation Corporation (SPEC).

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