



Unleashing AI Excellence: Lenovo ThinkSystem Dominates MLPerf v5.1 with Exceptional Inference Performance

Article

Artificial intelligence is transforming how enterprises innovate, compete, and grow—but scaling AI requires more than deploying a model. It demands enterprise-ready, sustainable, and outcome-driven infrastructure. Lenovo continues to lead the way, delivering Smarter AI for All through Lenovo Hybrid AI Advantage™, the industry’s most comprehensive hybrid AI platform.

The latest MLPerf® Inference v5.1 results from MLCommons® underscore Lenovo’s leadership in AI infrastructure. Lenovo ThinkSystem servers—powered by Intel® Xeon® 6 processors and NVIDIA GPUs—achieved exceptional performance across multiple inference scenarios, including generative AI, recommendation engines, and real-time applications. These systems set benchmarks for low latency, high throughput, and scalability, earning Top 3 rankings and 1st place in fifteen categories.

Whether you’re a retailer optimizing ad ranking, a healthcare provider accelerating diagnostics, or a tech innovator deploying complex language models, Lenovo ThinkSystem servers deliver faster time to value and real-world business outcomes. With Lenovo and NVIDIA, organizations can build, train, and scale AI confidently—unlocking innovation and driving the future of smarter AI.

Lenovo’s achievements

Along with of twenty-seven participating organizations including: AMD, Azure, Broadcom, Cisco, Dell, Google, Intel and HPE, Lenovo results are impressive:

Across numerous MLPerf Inference tests, these results highlight the capabilities of Lenovo ThinkSystem servers in various AI inference scenarios. Whether you're developing cutting-edge AI models or processing large datasets, Lenovo's server configurations provide the scalability and performance you need to drive innovation forward.

The following table provides a breakdown of our results for Datacenter and Edge.

Table 1. Lenovo’s achievements

System	Total Categories Submitted	First Place Finishes	Second Place Finishes	Third Place Finishes
ThinkEdge SE100	12	12	0	0
ThinkSystem SR680a V3	5	2	0	1
ThinkSystem SR780a V3 (Intel)	9	1	2	2

Our benchmarks display Lenovo’s consistent improvement to achieve best-in-class results for our customers. These key results show how powerful our systems are in those specific categories:

- **ThinkSystem SR680a V3 (8x B200-SXM-180GB, TensorRT)**
 - llama2-70b-99 server (7/10)
 - llama2-70b-99 offline (1/10)
 - llama2-70b-99.9 server (3/11)
 - llama2-70b-99.9 offline (1/11)
 - whisper (7/14)
- **ThinkSystem SR780a V3 (8x B200-SXM-180GB, TensorRT)**
 - 1-8b server (4/9)
 - 1-8b offline (1/9)
 - stable-diffusion-xl server (4/8)
 - stable-diffusion-xl offline (6/9)
 - 1-405b server (2/13)
 - 1-405b offline (9/13)
 - 1-405b interactive (2/10)
 - rgat (4/5)
 - whisper (5/14)
- **ThinkEdge SE100 (1x RTX-2000E-Ada-16GB)**
 - resnet50 singlestream (1/1 - no competitors)
 - resnet50 multistream (1/1 - no competitors)
 - resnet50 offline (1/1 - no competitors)
 - retinanet singlestream (1/1 - no competitors)
 - retinanet multistream (1/1 - no competitors)
 - retinanet offline (1/1 - no competitors)
 - 3d-unet singlestream (1/1 - no competitors)
 - 3d-unet offline (1/1 - no competitors)
 - 3d-unet-99.9 singlestream (1/1 - no competitors)
 - 3d-unet-99.9 offline (1/1 - no competitors)
 - stable-diffusion-xl singlestream (1/1 - no competitors)
 - stable-diffusion-xl offline (1/1 - no competitors)

Conclusion

The insights from the latest MLPerf benchmarks are critical for stakeholders in the generative AI and machine learning ecosystem, from system architects to application developers. They provide a quantitative foundation for hardware selection and optimization, crucial for deploying scalable and efficient AI/ML systems. Future developments in hardware and software are expected to further influence these benchmarks, continuing the cycle of innovation and evaluation in the field of machine learning.

Professionals in the field are encouraged to consider these results in their future hardware procurement and system design strategies. For further discussion or consultation on leveraging these insights in specific use cases, engage with our expert team at aidiscover@lenovo.com.

For more information

For more information, see the following resources:

- Explore Lenovo AI solutions:
<https://www.lenovo.com/us/en/servers-storage/solutions/ai/>
- Engage the Lenovo AI Center of Excellence:
<https://lenovo-ai-discover.atlassian.net/servicedesk/customer/portal/3>

Explore Lenovo MLPerf Results:

- [Lenovo ThinkSystem Servers Continue to Lead the Industry in Performance and Customer Value](#)
- [Lenovo Strengthens AI Leadership with Standout MLPerf Training 5.0 Results](#)

- [Lenovo Leads the Pack in MLPerf Inference: Datacenter 5.0 Benchmarks with Industry-Best AI Performance](#)

MLCommons®, the open engineering consortium and leading force behind MLPerf, has now released new results for MLPerf benchmark suites:

- Benchmark results: <https://mlcommons.org/benchmarks/inference-datacenter/>
- Latest news about MLCommons: <https://mlcommons.org/news-blog>

Author

David Ellison is the Chief Data Scientist for Lenovo ISG. Through Lenovo's US and European AI Discover Centers, he leads a team that uses cutting-edge AI techniques to deliver solutions for external customers while internally supporting the overall AI strategy for the Worldwide Infrastructure Solutions Group. Before joining Lenovo, he ran an international scientific analysis and equipment company and worked as a Data Scientist for the US Postal Service. Previous to that, he received a PhD in Biomedical Engineering from Johns Hopkins University. He has numerous publications in top tier journals including two in the Proceedings of the National Academy of the Sciences.

Related product families

Product families related to this document are the following:

- [Artificial Intelligence](#)
- [MLPerf Benchmark](#)

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 2025. All rights reserved.

This document, LP2330, was created or updated on November 11, 2025.

Send us your comments in one of the following ways:

Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2330>

Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2330>.

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 Hybrid AI Advantage

ThinkEdge®

ThinkSystem®

The following terms are trademarks of other companies:

AMD is a trademark of Advanced Micro Devices, Inc.

Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Azure® is a trademark 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.