



Lenovo Al Supports MLPerf, Helping Customers Make Informed Decisions Faster Article

Lenovo is committed to supporting MLPerf

As a leader in Artificial Intelligence, Lenovo is committed to supporting MLPerf, helping customers make better-informed decisions for their AI workloads.

MLPerf is a consortium of AI leaders from academia, research labs, and industry. Its mission is to "build fair and useful benchmarks" that provide unbiased evaluations of training and inference performance for hardware, software, and services—all conducted under prescribed conditions.

Lenovo understands the complexity of deploying AI solutions that solve real business challenges. For example, AI is dramatically transforming retail. Some key AI trends in retail and real-life applications where retailers are using AI, including

- Up to 4X increase annually for the next three years in autonomous shopping, such as autonomous checkout and NanoStores, a wholly automated convenience store
- Over \$100B in annual shrinkage using AI for loss prevention, including ticket switching, misscanning, and security
- Up to 50% of top retailers use AI for store analytics to disclose demographic analysis, customer engagement, price matching, and stock-outs.

Choosing the best infrastructure for your workload should not be a barrier to your AI Initiatives. Fundamentally, Lenovo is making continual progress with supporting more models and benchmarks in the MLPerf portfolio. This momentum includes extending and enhancing our AI infrastructure portfolio so customers can make decisions with better, faster insights. Our goal through MLPerf Inference 2.0 - Training is to bring clarity to infrastructure decisions so customers can focus on the success of their AI deployment overall.

MLPerf 2.0 highlights

Some recent highlights supporting the MLPerf 2.0 efforts include:

- Expanded Lenovo infrastructure systems testing (seven inference and three training)
- Support for NVIDIA A100 Tensor Core GPUs on MLPerf Inference 2.0 AI Training Benchmark
- Engineering enhancements with Lenovo Neptune, including liquid-cooled CPUs, GPUs, DIMMs, and storage
- Optimized Lenovo Server for AI
 - Lenovo ThinkSystem SR670 V2 with 8x 80 GB A100
 - Lenovo ThinkSystem SR670 V2 with 4x SXM A100 (80GB/500W)
 - Lenovo ThinkSystem SD650-N V2 with 4x SXM A100 (80GB/500W)
- Expanded benchmark submission, now including MiniGo score for Reinforcement Learning
 - Improved Performance: Lenovo achieved the best scores for 4xSXM servers for MaskRCNN (Object Detection) and 3D-Unet (medical imaging)



Figure 1. Lenovo ThinkSystem SR670 V2 configured to support eight NVIDIA A100 GPUs

MLPerf results

Lenovo achieved the following improvements in comparison with previous rounds:

- 3% in Resnet score (SXM Server)
- 9% in MaskRcnn score (SXM Server)
- 18% for PCIe
- 42% in RNN-T (Recurrent Neural Network Transducer)

Lenovo demonstrated AI performance across various infrastructure configurations, including NVIDIA A16 GPUs and NVIDIA Triton[™] Inference Server, running on Lenovo ThinkSystem platforms with liquid-cooled CPUs. We also showcased the efficiency and performance of our air-cooled systems like the SR670 V2, providing both PCIe and HGX deployment options in a standard data center platform that enterprises of all sizes can quickly deploy.

These engineering feats reinforce the performance benefits of Lenovo's ThinkSystem Servers with Neptune[™] that delivers Exascale grade performance in a standard, dense data center footprint.

Lenovo and NVIDIA

Lenovo and NVIDIA collaborate extensively in Artificial Intelligence. We partner through Lenovo Al Innovation Centers, where we work to ensure the success of our mutual customers with their Al initiatives. Here, customers get access to Lenovo and NVIDIA Al experts to consult on projects, the proper infrastructure to run a proof of concept, and proof of ROI before deployment.

We also collaborate with the Lenovo AI Innovators program, where we have multiple AI software companies whose codes are optimized to support Lenovo servers with NVIDIA GPUs.

For more information

For more information, see the following resources:

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

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

- Benchmark results: https://mlcommons.org/en/training-normal-20/
- Latest news about MLCommons: https://mlcommons.org/en/news/mlperf-training-2q2022/.

Related product families

Product families related to this document are the following:

- Artificial Intelligence
- MLPerf Benchmark
- ThinkSystem SD650-N V2 server
- ThinkSystem SR670 V2 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 2025. All rights reserved.

This document, LP1616, was created or updated on June 29, 2022.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at: https://lenovopress.lenovo.com/LP1616
- Send your comments in an e-mail to: comments@lenovopress.com

This document is available online at https://lenovopress.lenovo.com/LP1616.

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® Neptune® ThinkSystem®

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