

# Lenovo Strengthens AI Leadership with Standout MLPerf Training 5.0 Results

## Article

MLPerf is the industry's gold standard for measuring AI performance, developed by a collaboration of leading tech companies and academic institutions. It benchmarks real-world AI workloads like large language models, computer vision, and more. With the release of MLPerf Training 5.0, Lenovo has once again proven its strength in AI infrastructure delivering standout results.

Lenovo continues to prove its leadership in AI infrastructure with a strong showing in the latest MLPerf Training 5.0 benchmarks. Lenovo ThinkSystem SR680a V3 and SR780a V3 systems delivered impressive performance across a range of AI training workloads—including large language models, generative AI, and computer vision. **This is the first Lenovo submission with a multi-node cluster: 4 nodes SR680a V3 with 8x H200-SXM-141GB in each.**



Figure 1. Lenovo ThinkSystem SR680a V3

## Highlights from MLPerf Training 5.0

Lenovo proudly participated in the MLPerf Training 5.0 benchmark suite, achieving outstanding results with the state-of-the-art Lenovo ThinkSystem SR680a V3 and SR780a V3 systems, with the latest NVIDIA H200 and B200 GPUs.

### ThinkSystem SR680a V3 with 32x H200-SXM-141GB (4 Nodes, 8 GPUs in each)

Here are the key outcomes with the ThinkSystem SR680a V3:

- LLMs - LaMA2-70B-LoRA: Training Time: 6.94 minutes
- Computer Vision - RetinaNet: Training Time: 9.64 minutes
- Generative AI - Stable Diffusion (Training): Training Time: 79 minutes

The Lenovo ThinkSystem SR680a V3 with 32x H200-SXM-141GB (4 Nodes, 8 GPUs in each) is an NVIDIA Certified system built for large-scale AI training, the SR680a V3 showed its muscle across some of the most demanding benchmarks earning **1st place in LLMs**.

These results confirm the SR680a V3's capability to handle advanced AI workloads at scale—ideal for organizations training cutting-edge models that require massive compute.

### Lenovo ThinkSystem SR780a V3 with 8x B200-SXM-180GB (Single Node)

Performance results for the powerful Lenovo ThinkSystem SR780a V3 have been officially published and can be viewed on the MLCommons website, showcasing its capabilities in the industry-standard MLPerf Training benchmarks:

<https://mlcommons.org/benchmarks/training/>

## Why Lenovo?

Lenovo's strong benchmark performance is no accident—it's the result of thoughtful design and a clear focus on what AI users need most.

Here are a few key reasons why Lenovo systems continue to stand out:

#### 1. Scalability That Fits Your Needs

Whether you're training massive LLMs or building models on a single node, Lenovo systems offer the flexibility to grow with your AI projects.

#### 2. Versatility Across Workloads

From computer vision and generative AI to fine-tuning large models with LoRA, Lenovo delivers consistent, dependable performance.

#### 3. Powered by the Best

Equipped with the latest NVIDIA H200 and B200 GPUs, ThinkSystem platforms are built to tackle today's toughest AI challenges with speed and efficiency.

## A Future-Focused Approach to AI

The MLPerf Training 5.0 benchmarks reinforces the rapid evolution of AI hardware and software, highlighting the potential for significant improvements in performance and efficiency. Our contributions to these benchmarks reflect our commitment to innovation and pushing the boundaries achieving business outcomes with AI.

As we continue to innovate, the insights gained from MLPerf Training 5.0 will guide us in developing cutting-edge AI solutions that address the most pressing challenges across industries. The transformative power of AI is evolving fast and is becoming increasingly evident. Whether its supporting cutting-edge research, helping businesses deploy AI at scale, or building tools for the next generation of generative models, Lenovo is all-in on powering real progress.

We are excited to be at the forefront of this revolution.

## Conclusion

Lenovo's latest results in the MLPerf Training 5.0 benchmarks highlight the strength and versatility of its AI systems. With the SR680a V3 leading in large-scale training and the SR780a V3 offering a compact and powerful solution, Lenovo is ready to support organizations at every stage of their AI journey. Backed by powerful hardware, smart design, and proven performance, Lenovo continues to be a reliable partner for teams building and scaling the future of AI.

Professionals in the field are encouraged to consider these results in their future hardware procurement and system design strategies.

## Author

**Traci Parker** is the Worldwide Solutions Marketing Manager for Enterprise IT and AI at Lenovo. She specializes in hybrid cloud, infrastructure modernization and AI solutions. She has more than 15 years of experience as a Marketing Manager and Product Marketing Manager across high-tech, fin-tech and healthcare industries.

## Related product families

Product families related to this document are the following:

- [MLPerf Benchmark](#)
- [ThinkSystem SR680a V3 Server](#)
- [ThinkSystem SR780a V3 Server](#)

## Notices

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

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

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

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

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

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

© Copyright Lenovo 2025. All rights reserved.

This document, LP2277, was created or updated on August 15, 2025.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:  
<https://lenovopress.lenovo.com/LP2277>
- Send your comments in an e-mail to:  
[comments@lenovopress.com](mailto:comments@lenovopress.com)

This document is available online at <https://lenovopress.lenovo.com/LP2277>.

## 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®

ThinkSystem®

The following terms are trademarks of other companies:

Approach® is a trademark of IBM in the United States, other countries, or both.

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