

From CPU to GPU: building the right datacenter for your AI workload needs



Are you ready to seize the opportunity of AI?

If you want to run demanding AI inferencing workloads you need a reliable, high-performance platform that can meet a range of evolving needs while balancing power and efficiency.

And you need it fast.

AI won't wait

96% of CIOs anticipate increased investment in AI.¹

79% of leaders say their organizations need to adopt AI to stay competitive.²



Datcenters share capabilities

Whether in-house or contracted public/private, organizations rely on datacenters to power AI across their organization.

- Public models
- Public services (e.g. Chat-GPT, Gemini)
- Privately built/trained models
- Licensed models (enterprise, industry)

How will you power ever-growing AI?

CPUs - The latest CPUs are being designed for AI. They can be more efficient than GPUs for small- to mid-sized AI model development, testing, and training.



GPUs - GPUs can be faster and reduce overall total cost of ownership (TCO) when developing and training larger models.



Get the best of both worlds

Why compromise? Embrace a flexible platform that offers all the benefits of CPUs with the ability to efficiently scale up to GPU power as needed.

AMD EPYC™ processors and AMD Instinct™ accelerators address the spectrum of AI workloads

Deploy with the confidence of open standards and leadership security features.

AMD Instinct Accelerators (GPU) and AMD EPYC processors (CPU) are designed to excel across different AI workloads:

AMD EPYC processors (CPU) are ideal for...

- Mixed workload inference deployments
- Classic machine learning
- Small to medium models
- Batch, offline & small-scale real-time inference

AMD Instinct accelerators (GPUs) are ideal for...

- AI training
- Dedicated AI deployments
- Medium to large Gen AI models
- Large-scale, real-time inference

ThinkSystem SR685a V3

- Powerful 8U server
- Features two AMD EPYC™ processors
- Includes up to 8 AMD Instinct™ accelerators

ThinkSystem SR675 V3

- Versatile GPU-rich 3U rack server
- Features one or two AMD EPYC™ processors
- Includes Lenovo Neptune® hybrid liquid-to-air cooling³

The right CPU host is essential to maximize the value of large GPU Investments

Perfect partners

AMD EPYC CPUs complement the GPUs with:

- High frequency and performance for massive volumes of data pre-processing
- High memory bandwidth, high performance I/O

CPU-based AI use cases

- Recommendation systems
- ML image recognition and analysis
- Generative AI for enterprise-class LLM/SLM systems (e.g. HR Chatbot, document summarization)
- Mixed, AI-enabled Apps (e.g. Database, collaboration, design enhanced with AI)



Lenovo and AMD offer performance and efficiency leadership for both CPUs & GPUs.

Together, we will work with you to help identify the best path to solve your AI datacenter challenges.

[Learn More](#)



¹ Microsoft and LinkedIn, "2024 Work Trend Index Annual Report," May 2024
² Microsoft and LinkedIn, "2024 Work Trend Index Annual Report," May 2024
³ NVIDIA HGX H200 4-GPU offering with NVLink and Lenovo Neptune hybrid liquid-to-air cooling

© 2025 Lenovo. All rights reserved. Lenovo and the Lenovo logo are trademarks of Lenovo. AMD, the AMD Arrow logo, AMD EPYC, AMD Instinct, and combinations thereof are trademarks of Advanced Micro Devices, Inc. All other trademarks are the property of their respective owners.