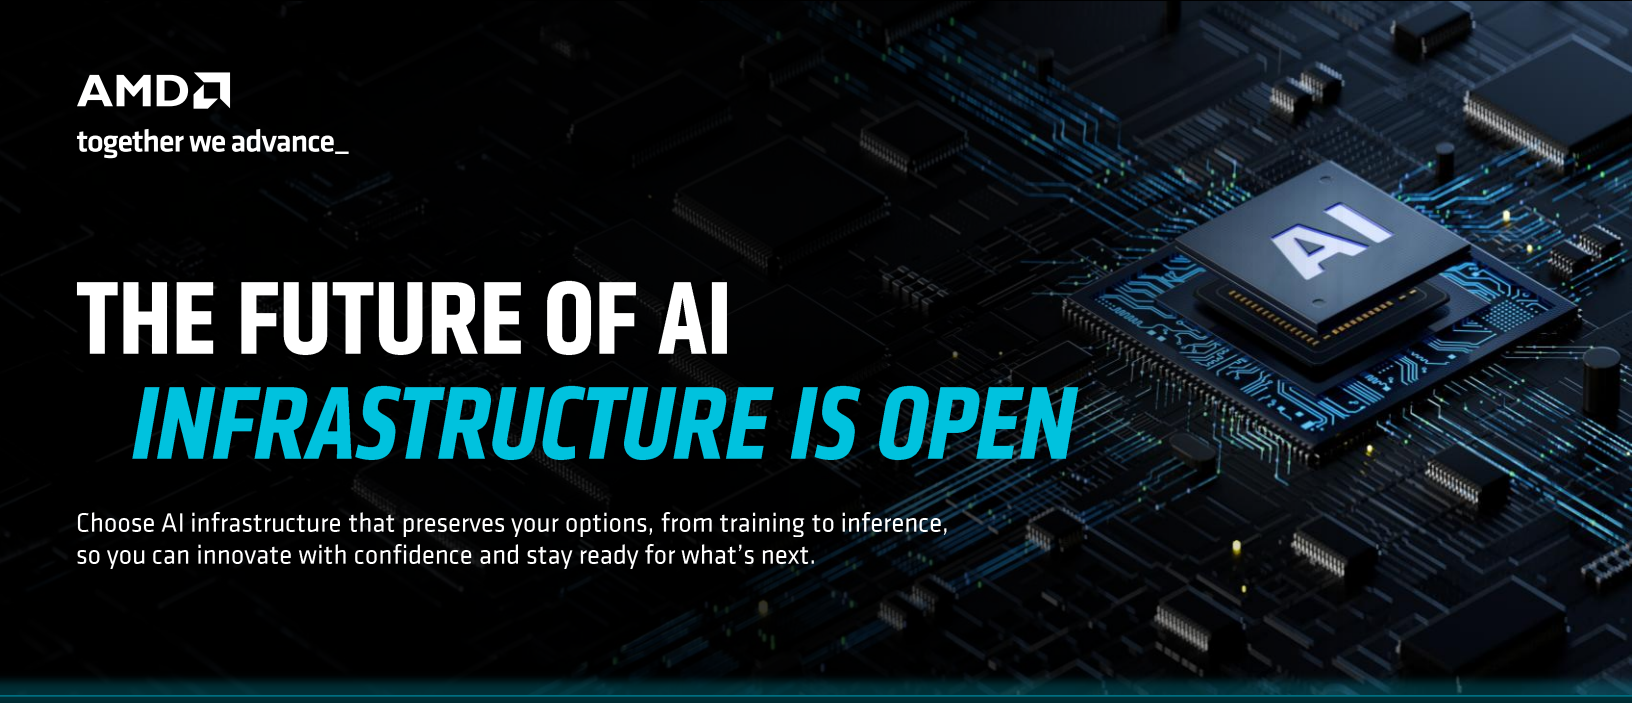
AMD
together we advance_



THE FUTURE OF AI
INFRASTRUCTURE IS OPEN

Choose AI infrastructure that preserves your options, from training to inference, so you can innovate with confidence and stay ready for what's next.

BUILD YOUR DATA CENTER INFRASTRUCTURE WITH THE FUTURE OF AI IN MIND

AI workloads, including agentic AI, are complex and place significant demands on the infrastructure that powers them. As those demands grow, the infrastructure decisions you make now can shape your long-term ability to adopt and scale AI. The pressure to deliver AI value quickly can often make the most familiar path feel like the natural choice. However, it's important to consider how to preserve flexibility as your needs change.

IDC projects that next year, for the first time, annual worldwide investments in infrastructure for AI inferencing will exceed those for AI training.[1]

79% of organizations require AI to be integrated into their general-purpose workloads, rather than treating AI inferencing as a stand-alone capability.[1]

THE MYTHS AND REALITIES SHAPING AI INFRASTRUCTURE DECISIONS

The strongest infrastructure decisions come from a clear view of the facts. Often, the greatest value lies in the willingness to question what's familiar. When it comes to proprietary and open ecosystems, separating myths from realities can help you choose a path that preserves your flexibility and allows you to adapt as business needs evolve.

MYTH	REALITY
Time to deployment is the most important factor.	It's worth weighing both deployment speed and long-term innovation. Proprietary and open ecosystems can get you started, but the difference reveals itself over time in how much flexibility you retain as your needs change. Open, turnkey AI solutions can help enable fast deployments while preserving your options, so you don't need to choose between moving quickly today and innovating tomorrow.
Open ecosystems are difficult to deploy and support.	Open ecosystems are designed to work with the leading AI frameworks and tools your organization already uses, so adoption doesn't require starting from scratch. When they are backed by mature OEM, software, and partner ecosystems, they can offer enterprise-level support when scaling.
Once you train on a proprietary platform, you're locked in for everything downstream.	Training and inference are two separate infrastructure decisions, and they don't need to run on the same platform. Because open ecosystems support industry-standard frameworks and tools built for portability, you have the flexibility to run inference where it makes the most sense for your business.
Proprietary ecosystems are the most innovative.	Innovation tends to accelerate with the number of people contributing to it. Open AI ecosystems draw on contributions from a wide range of companies, developers, and research communities, creating a steady stream of ideas and improvements that benefit everyone building on them.

THE ADVANTAGES OF AN OPEN ECOSYSTEM WITH AMD

Seamless portability and compatibility
AMD supports leading AI frameworks and inference environments, including PyTorch, Hugging Face, ONNX Runtime, vLLM, SGLang, and llama.cpp. Through HIP, developers can port code that runs on both AMD and other hardware, helping to reduce migration friction.

Tap into a broad network of trusted partners
AMD works with hyperscalers like Microsoft Azure and Oracle Cloud, and OEMs including HPE, Lenovo, and Supermicro, to build and optimize open AI infrastructure.

Scale with open networking and interconnect standards, backed by AMD leadership in the Ultra Accelerator Link (UALink) Consortium.

Preserve the freedom to innovate and adapt as AI technologies, workloads, and business priorities evolve.

UNISON'S YUCA SUPERCOMPUTER SHOWS WHAT'S POSSIBLE WITH AN OPEN ECOSYSTEM

As UNISON's research needs grew, its team adopted AMD GPUs alongside AMD ROCm™ software, which supported a smooth transition for researchers already working in CUDA—no rebuilding required.

RESULTS



Cut research time from years to hours



Streamlined CUDA migration with AMD ROCm software



Reliable 24/7 operations[2]

> We trained our users so that with ROCm it was transparent. They could now use AMD GPUs through the set of tools that ROCm provides and continue running their CUDA code just as they did on NVIDIA. The ROCm training made the transition for users from NVIDIA technology to AMD GPUs much easier for our researchers.
>
> Maria Del Carmen Heras Sanchez, Director, High-Performance Computing Area, UNISON

BUILD FOR TODAY. BE READY FOR TOMORROW.

Discover how an open ecosystem can help you build for innovation today, while staying ready for what's next.

LEARN MORE

[1]IDC, White Paper, sponsored by AMD, The Future of AI Inferencing Lies in Open, Flexible Infrastructure Stacks, #US54516826, July2026.
[2]GD-18u: All performance and cost savings claims are provided by University of Sonora and have not been independently verified by AMD. Performance and cost benefits are impacted by a variety of variables. Results herein are specific to University of Sonora and may not be typical. GD-181 The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18u.

© 2026 Advanced Micro Devices, Inc.