

The Dell Technologies logo is positioned in the top left corner. It features the word "DELL" in a white, bold, sans-serif font, with a stylized "E" that has three diagonal lines extending from its top. To the right of "DELL", the word "Technologies" is written in a smaller, white, sans-serif font. The background of the entire slide is a close-up, high-angle photograph of a green printed circuit board (PCB) with various electronic components, including integrated circuits and capacitors, under a blue and purple light. The components are slightly out of focus, creating a sense of depth and technical complexity.

DELL Technologies

The AMD logo and tagline are located in the top right corner. The AMD logo consists of the letters "AMD" in a white, bold, sans-serif font, followed by a square icon containing a stylized white "A" shape. Below the logo, the tagline "together we advance_" is written in a smaller, white, sans-serif font. The background of the slide is a close-up, high-angle photograph of a green printed circuit board (PCB) with various electronic components, including integrated circuits and capacitors, under a blue and purple light. The components are slightly out of focus, creating a sense of depth and technical complexity.

AMD
together we advance_

Drive Powerful AI Insights with an Open Ecosystem

A solid purple horizontal bar is located below the main heading.

Dell PowerEdge XE9680 Server with AMD Instinct™
MI300X Accelerator and ROCm™ 6 open software platform

With the growing popularity of ChatGPT and other generative AI (GenAI) applications, artificial intelligence (AI) is rapidly becoming pervasive. And more organizations want in on the action. A recent IDC study shows that:



77% of organizations

are either exploring AI use cases or already significantly investing in AI initiatives.¹

Organizations looking to capitalize on accelerated insights and other AI benefits need to act soon, if you haven't already. The time is right, especially with an open ecosystem heavily engaged in the advancement of AI, and a new solution that is ready to drive AI innovation — and your business — forward.

¹ IDC, Future Enterprise Resiliency and Spending Survey, Wave 6, July 2023.



The road forward is open

With GenAI and other advanced applications on the rise, chip technology is becoming more powerful and sophisticated. But proprietary systems, at the foundational level, have created barriers for developers on the front line.

This is about to change. Dell Technologies and AMD are working together to accelerate your capabilities with infrastructure solutions built on open-source technologies. Running the AMD Instinct™ MI300X accelerator and the AMD Radeon Open Compute (ROCm™) open software platform on the Dell PowerEdge XE9680 Server provides you with a powerful ecosystem option — an open path forward — to spark opportunity and unleash your potential.

Gain performance and flexibility with Dell Technologies and AMD

Dell Technologies and AMD are helping businesses across the globe accelerate outcomes with a system that will set developers up for success.



The Dell PowerEdge XE9680 Server is a data-processing powerhouse designed especially for AI tasks. It supports eight accelerators, ideal for machine learning (ML) / deep learning (DL) training and inferencing workloads, especially for those training large language models (LLMs).

The AMD Instinct MI300X accelerator, powered by 3rd Gen AMD CDNA™ (Compute DNA) architecture, delivers the greatest high-bandwidth memory capacity in the GPU market, providing 1.5 terabytes of coherent GPU accelerator memory per server for training the largest AI models. It also enables you to increase inferencing capacity using fewer servers.

Unlock acceleration with ROCm 6

ROCm, in turn, unleashes the computational power of AMD Instinct. An open, flexible and scalable software platform for accelerated computing, ROCm is equipped with drivers, APIs and other GPU tools that support code optimization for AMD Instinct accelerators and simplify portability for existing GPU accelerated code.

It also radically simplifies your experience. ROCm is optimized at the device level with a freely shared code base that can evolve as systems and their parts evolve. This open flexibility also enables effortless porting and almost zero coding.

An ecosystem without borders

Built on open standards and interoperability, the AMD Instinct MI300X accelerator and ROCm are natively supported by a variety of AI frameworks and libraries. This provides greater choice and flexibility for those wanting to harness the power of AMD and still maintain compatibility with existing software systems.




AMD clears the road, providing the freedom to advance your projects at speed.

ROCm advantages

- Unlocks GPU power to accelerate computational tasks
- Optimized for ML/DL/GenAI workloads at scale
- Delivers out-of-the-box support for leading models and frameworks (including PyTorch[®], TensorFlow[™] and others)
- Enables rapid adoption of new innovations (including native support in JAX, Open Neural Network Xchange (ONNX), OpenAI Triton, etc.)
- Drives innovation, differentiation and collaboration with connection to the OSS development community
- Provides a flexible foundation for open portability

Learn more about [What's new in ROCm](#).



Easily Port and Optimize Applications for AMD Instinct with ROCm

- Enables low-friction software porting to AMD. HIP allows developers to create portable applications for the AMD Instinct MI300X from single-source code.
- Offers all features with commonly used libraries and operators
- Quickly ports and optimizes most applications for performance

Growing Upstream Support

ROCm provides upstream support for today's most popular frameworks, including but not limited to:

- PyTorch
- TensorFlow
- Open Neural Network Xchange (ONNX)
- JAX



Save valuable time with open portability

ROCm delivers peace of mind to developers who might be wary of a new ecosystem and the impact it can have on existing code. A key advantage of the ROCm platform is that you can port applications with minimal rework to your existing code. Porting is simple, and you save a tremendous amount of time.

Port existing applications to the AMD Instinct MI300X accelerator in four different ways:

1 Leverage existing AI frameworks that have upstream support for ROCm.

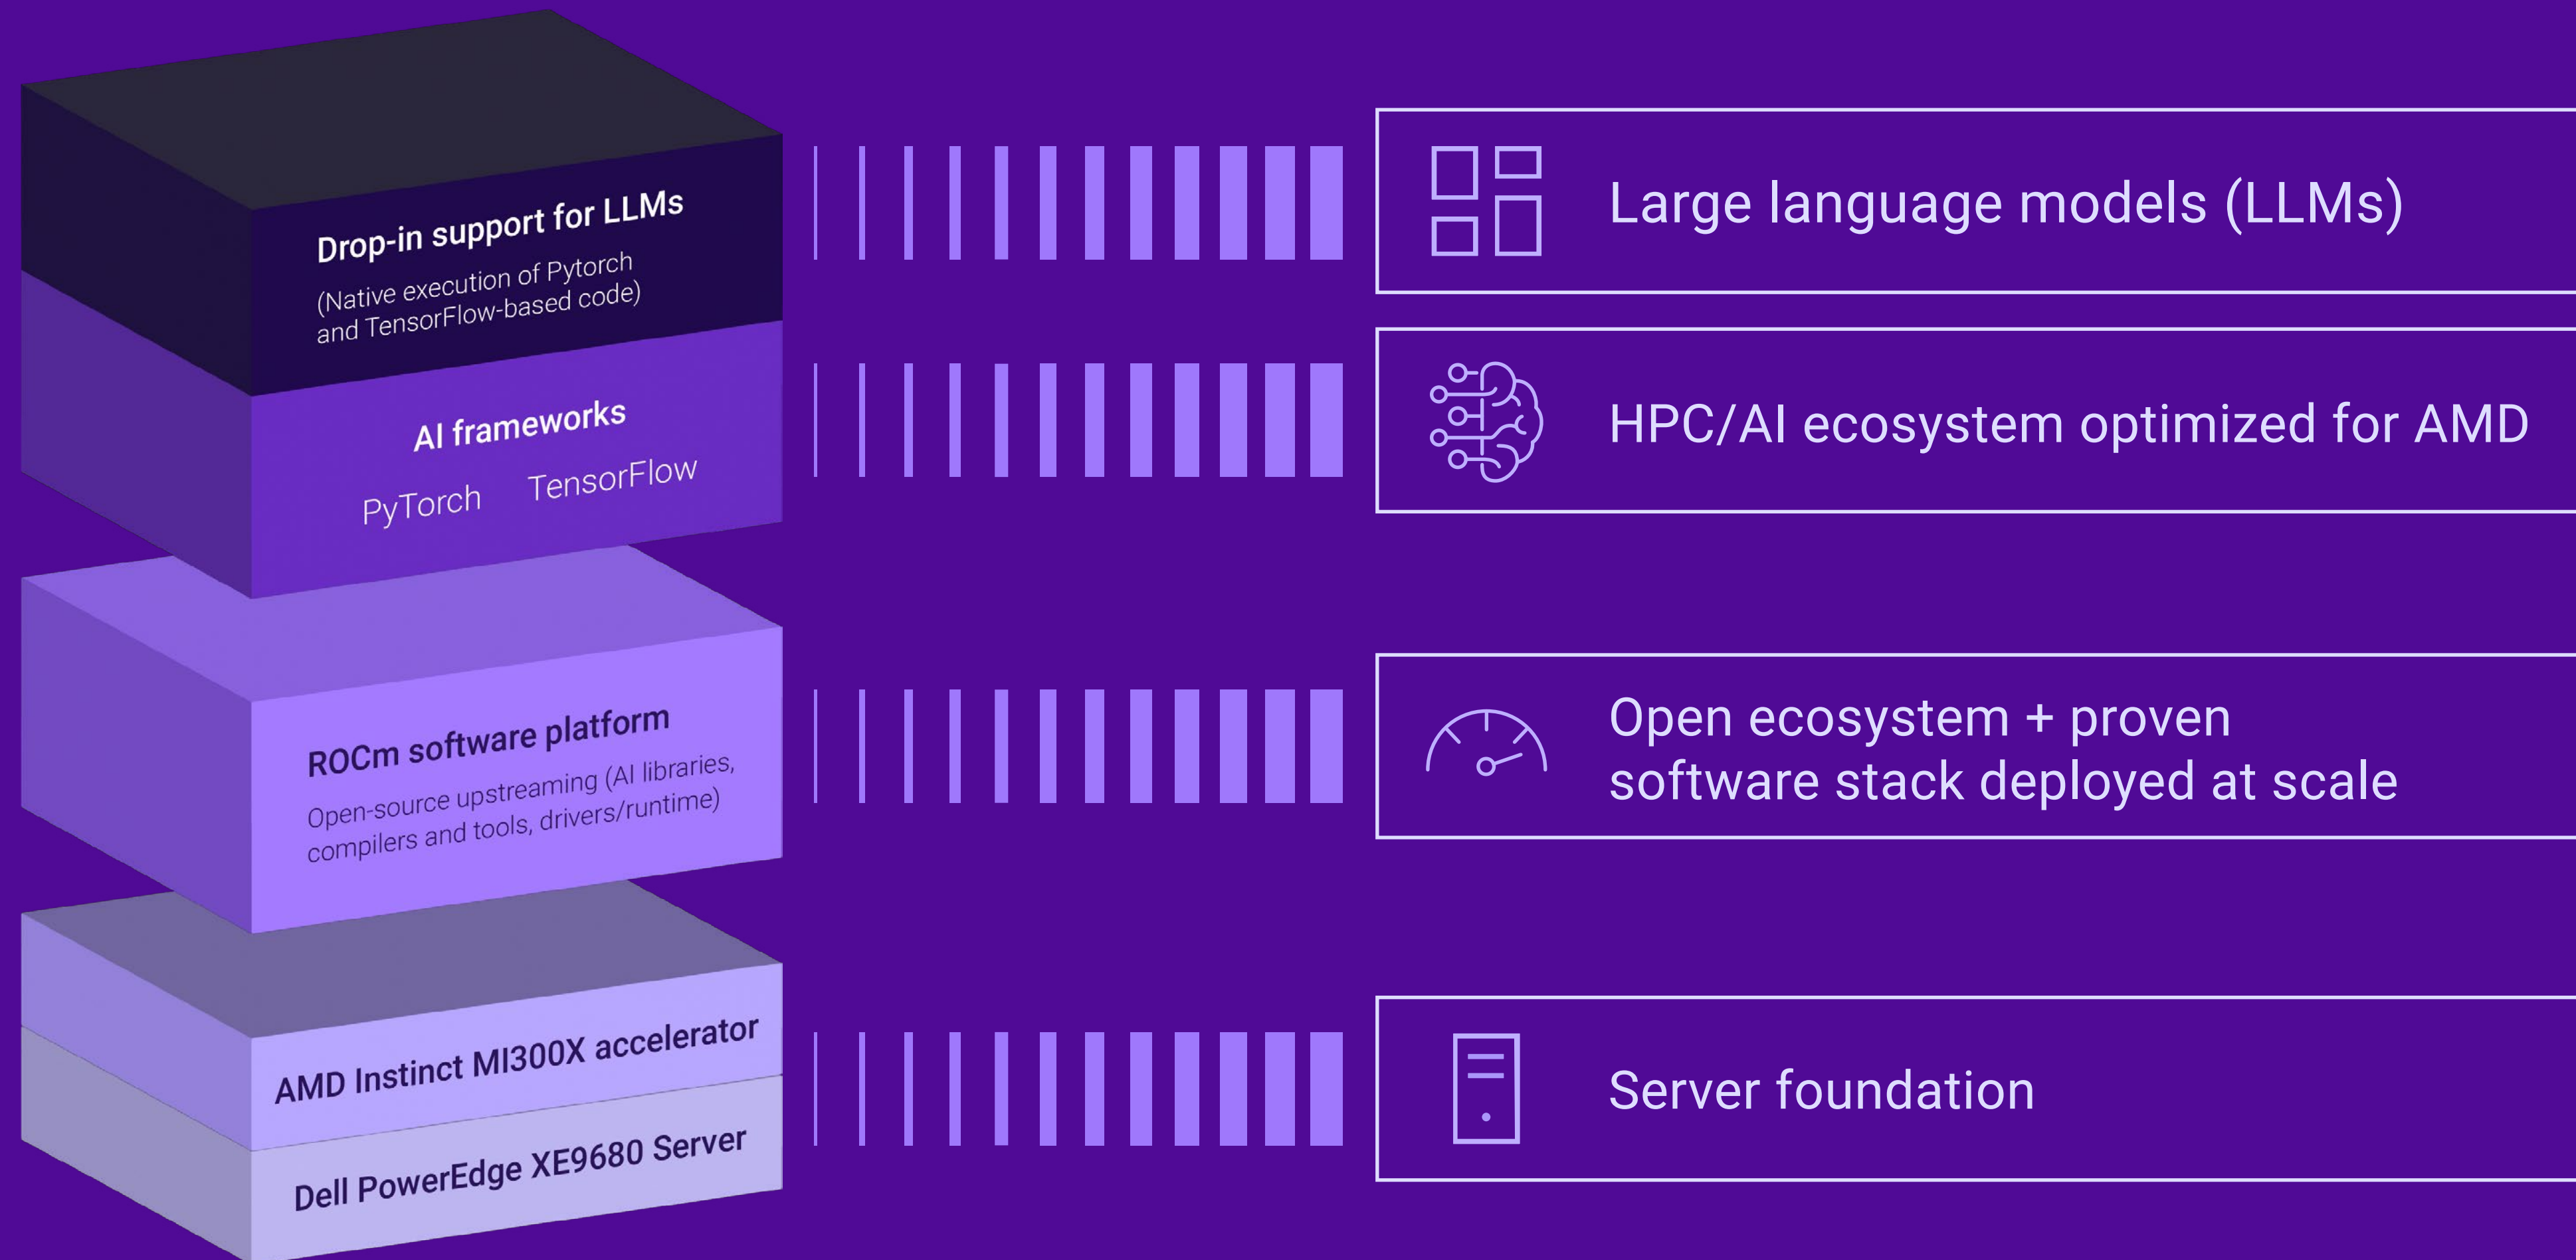
2 Use drop-in ML kernels that are optimized for the AMD Instinct MI300X. Or take advantage of the AMD HIP conversion tool, which enables you to optimize CUDA applications on ROCm-based GPUs.

3 Use mirrored versions of existing CUDA library software repositories that are optimized and maintained for ROCm.

4 Access AMD prepackaged AI applications and containers.

Get started today

Accelerate existing AI and new GenAI projects with an open framework from Dell Technologies and AMD. Power your workloads on a proven Dell PowerEdge foundation that supports the latest AMD accelerators, including the AMD Instinct MI300X accelerator. [Visit the ROCm Developer Hub.](#)




The AMD Instinct MI300X accelerator is ready to deploy


Delivering breakthrough performance, the AMD Instinct MI300X accelerator delivers the speed, scalability and flexible, open ecosystem that's needed today, especially as businesses deal with resource, complexity, latency and architecture challenges. The AMD Instinct MI300X accelerator provides a distinct advantage, in terms of:




Performance: Enable new AI workflows with an unprecedented 192GB of stacked HBM3 memory per GPU operating at 5.3TB per second.



Density: Reduce the need for data movement with state-of-the-art die-stacking, chiplet technologies packing more compute and memory capacity per server.



Open: Ease the way to results with an ecosystem designed for accessibility and adaptability.



Unleash your AI advantage with the Dell PowerEdge XE9680 Server

You can accelerate your path to faster, smarter outcomes with a server platform designed specifically for AI. Enhance data-driven business, deliver visual outcomes and expedite business-wide AI operations with a solution that accelerates insights, delivers trusted AI and simplifies operations.



Trusted AI

Reduce risks and save time as you accelerate your AI lifecycle with a trustworthy, high-quality solutions infrastructure. The PowerEdge XE9680 offers built-in protection from the silicon root of trust all the way to asset retirement.



Simplified operations

Effectively control and manage your AI infrastructure and performance workloads anywhere with AI infrastructure automation. Dell Technologies provides intelligent systems that work together and independently to accelerate tasks and simplify your experience.



Accelerated insights

Drive improved performance across the AI lifecycle and for delivery of AI operations at the speed of business. Dell Technologies delivers innovative compute to help you harness AI with maximum performance.

Drive Innovation Forward

Dell Technologies and AMD are working hard to drive innovation forward for our valued customers. Together, we bring silicon diversity to the market in a viable solution that empowers developers. Our joint solution enables you to:



- **Get to market faster** with more options and fewer roadblocks.
- **Drive faster insights** with an open ecosystem that doesn't hold you back.
- **Choose what's best for you.** Rely on the Dell supply chain advantage to consistently bring innovative technology to market when it's most needed.

Move full speed ahead with Dell Technologies and AMD

AI is here now, but its influence is only going to expand and accelerate moving forward. For businesses looking to take advantage of its capabilities, the possibilities seem endless.

That's why you don't want your potential restricted by technological roadblocks that can sometimes get in your way. Instead, keep your options open with a powerful solution stack dedicated to AI workloads.

The AMD Instinct MI300X accelerator and ROCm software on the Dell PowerEdge XE9680 Server is the answer. Speed the path to new applications, get to market faster and start driving outcomes that matter to your business and the world at large.

Learn more at Dell.com/GPU

Copyright © 2023 Dell Inc. or its subsidiaries. All Rights Reserved. Dell and other trademarks are trademarks of Dell Inc. or its subsidiaries. AMD, Instinct™, ROCm™, and combinations thereof are trademarks of Advanced Micro Devices, Inc. PyTorch® is a trademark or registered trademark of PyTorch or PyTorch's licensors. TensorFlow™ and any related marks are trademarks of Google Inc. CUDA® is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries. Other trademarks may be the property of their respective owners. Published in the USA 12/23 eBook.

Dell Technologies believes the information in this document is accurate as of its publication date. The information is subject to change without notice.