

BUILT FOR YOUR BUSINESS OUTCOMES

# SCALE ENTERPRISE AI WITH FLEXIBLE, FULL-STACK SOLUTIONS

## WHY ENTERPRISES STRUGGLE TO SCALE AI

- Aging Infrastructure Limits AI Potential**  
Legacy systems can drain power, take up more space than what's needed, and lack the performance modern AI workloads demand.
- Data Readiness Remains Uneven**  
Many enterprises are lacking unified, usable data, making it difficult to know where they stand on the AI adoption curve.
- Closed Architectures Slow Innovation**  
Rigid systems restrict integration, blocking the plug-and-play flexibility needed to scale AI across environments.

## OVERCOMING MODERNIZATION BARRIERS WITH AMD

**78%** of ITDMs cite enabling AI and ML use cases as their top priority.<sup>1</sup>

- Quickly Scale AI**
  - Trusted by many of the world's largest organizations to enable AI workloads.
  - Deliver transformative performance gains and enable faster innovation, with Day One model support.
  - Empower end-to-end AI orchestration across the business to help accelerate scale and adoption.

**70%** of ITDMs cite scaling infrastructure to meet changing demand as their top challenge.<sup>2</sup>

- Increase Compute Capabilities**  
Whatever your compute needs, AMD has a solution.  
AMD provides CPUs, GPUs, DPUs, and adaptive accelerators to power AI from cloud to edge to endpoint.

**69%** of BDMs cite reducing costs as their top priority.<sup>3</sup>

- Effectively Manage TCO**  
Consolidate rack space without compromising performance by replacing legacy servers with solutions powered by AMD EPYC™ processors.  
Experience up to 67% lower TCO over 3 years by modernizing your data center with AMD servers.<sup>4</sup>

**68%** of ITDMs cite increasing flexibility as their top IT challenge.<sup>5</sup>

- Deploy Easily**
  - Experience open, flexible platforms built for choice and portability.
  - Equip teams with seamless interoperability across environments for faster integration.

## THE COMPETITION STRUGGLES TO POWER AI WORKLOADS

- Unleash the Power of Performance**  
AMD EPYC™ server CPUs deliver unmatched compute density and energy efficiency, enabling enterprises to train, infer, and scale faster while maintaining performance and reliability. AMD Ryzen™ AI PRO processors bring AI acceleration to the endpoint, helping organizations extend AI performance to every device in the enterprise.
- Accelerate AI Outcomes**  
AMD Instinct™ GPUs are built for next-generation AI training and inference, powering the world's largest models with breakthrough scalability and energy efficiency. With AMD Instinct™ GPUs, teams can train larger models, complete jobs faster, and deploy AI at a more sustainable pace than with conventional accelerator solutions.

- End-to-End Optimization for Modern AI**  
AMD ROCm™ software is an open software stack that provides developers with freedom, flexibility, and cross-platform optimization that proprietary GPU ecosystems can't offer. Developers get deep hardware-software optimization, and full observability across the AI pipeline, and seamless portability across cloud, on-prem, and edge systems.

- Flexible Foundation for Enterprise AI**  
AMD delivers a unified, open, and highly scalable compute foundation spanning CPUs, GPUs, adaptive SoCs, and DPUs, that enables organizations to deploy, optimize, and scale workloads consistently from edge to cloud. This unified approach helps remove the architectural silos that limit innovation, offering the flexibility and performance enterprises need to accelerate modernization, adopt AI confidently, and future-proof their infrastructure.

## YOUR TRUSTED PARTNER TO DRIVE BUSINESS

- Broad Portfolio of Leadership Compute Engines**  
Advance your AI initiatives with the broadest end-to-end AI compute portfolio, empowering scalable, high-performance AI infrastructure.
- Open Ecosystem for Rapid Co-Innovation**  
Accelerate innovation through openness and a customer-first approach to collaboration.
- Flexible, Full-Stack Solutions**  
Solve for your specific needs with optimized, integrated platforms from silicon to systems.

**“** We have found that AMD has the most powerful processor on the market, and that helps us build systems that can increase the throughput for each server while reducing the hardware investment costs. <sup>6</sup>

**Professor Minh Hoai Nguyen, Principal Research Scientist & Head of Smart Edge, VinAI**

**➔ READ MORE**

## DRIVE BETTER ENTERPRISE AI EXPERIENCES WITH LEADERSHIP PERFORMANCE AND POWER EFFICIENCY.

**➔ LEARN MORE**

<sup>1</sup>5-Point Likert Scale; Watchtower AI-powered research platform powered by Intercept, Global, n=122, 676,300.  
<sup>2</sup>5-Point Likert Scale; Watchtower AI-powered research platform powered by Intercept, Global, n=75,170.  
<sup>3</sup>5-Point Likert Scale; Watchtower AI-powered research platform powered by Intercept, Global, n=3104,240.  
<sup>4</sup>SP9xTTCO-003 This scenario contains many assumptions and estimates and, while based on AMD internal research and best approximations, should be considered an example for information purposes only, and not used as a basis for decision making over actual testing. The AMD Server & Greenhouse Gas Emissions TCO (total cost of ownership) Estimator Tool - version 1.12, compares the selected AMD EPYC™ and Intel® Xeon™ CPU based server solutions required to deliver a TOTAL PERFORMANCE of 391000000 units of SPECrate™2017\_int\_base performance as of October 10, 2024. This estimation compares a legacy 2P Intel Xeon 28 core Platinum 8280 based server with a score of 391 versus 2P EPYC 9965 (192C) powered server with an score of 3030 (https://spec.org/cpu2017/results/res2024q3/cpu2017-20240923-44833.pdf) along with a comparison upgrade to a 2P Intel Xeon Platinum 8592 (64C) based server with a score of 1130 (https://spec.org/cpu2017/results/res2024q3/cpu2017-20240701-43948.pdf). Actual SPECrate™2017\_int\_base score for 2P EPYC 9965 will vary based on OEM publications.  
<sup>5</sup>5-Point Likert Scale; Watchtower AI-powered research platform powered by Intercept, Global, n=653,830.  
<sup>6</sup>AMD, LUMI Case Study, 2023.