

Cautionary Statement

This presentation contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) such as the features, functionality, performance, availability, timing and expected benefits of AMD's current products, future products and product roadmaps, which are made pursuant to the Safe Harbor provisions of the Private Securities Litigation Reform Act of 1995.

Forward-looking statements are commonly identified by words such as "would," "may," "expects," "believes," "plans," "intends," "projects" and other terms with similar meaning. Investors are cautioned that the forward-looking statements in this presentation are based on current beliefs, assumptions and expectations, speak only as of the date of this presentation and involve risks and uncertainties that could cause actual results to differ materially from current expectations. Such statements are subject to certain known and unknown risks and uncertainties, many of which are difficult to predict and generally beyond AMD's control, that could cause actual results and other future events to differ materially from those expressed in, or implied or projected by, the forward-looking information and statements. Investors are urged to review in detail the risks and uncertainties in AMD's Securities and Exchange Commission filings, including but not limited to AMD's most recent reports on Forms 10-K and 10-Q.

AMD does not assume, and hereby disclaims, any obligation to update forward-looking statements made in this presentation, except as may be required by law.

AMD @ CES



together we advance_



Solving the World's Most Important Challenges

Accelerating AI from Silicon to Solutions



Powers the Daily Lives of Billions



Cloud



Healthcare



Industrial



Transportation



Connectivity



PCs



Entertainment



Science

AI is Everywhere



Cloud

Healthcare

Industrial

Transportation


Connectivity

PCs

Entertainment

Science

AI is for Everyone




Professionals



Creators



Students



Scientists



Gamers




Developers

Within 5 Years

5 Billion

People Using AI Daily



From Zetta to
Yotta Scale Compute

10,000x

Increase In AI Compute



10+
YottaFLOPS

Unlocking AI's Potential Requires **Massive Compute Everywhere**



Data Center




Personal Computing



Edge / Physical AI



Cloud



Personal
Computing



Healthcare



Physical AI



Space



Science &
Education

Cloud is the Fastest Path From Idea to Impact

Develop



Integrate



Deploy



Scale



Cloud is the Epicenter for AI Development



Chatbots &
Copilots

Personal
Assistants

Generative
Media

AI Agents &
Automation



Smart
Search

Coding
Assistants


Together with our Partners

AMD is Enabling the AI Era


AI Innovators




AMD Partners




AI Model Innovation and Adoption are Accelerating Compute Demand



AI Model Innovation and Adoption are Accelerating Compute Demand



Source: Based on AMD internal analysis and data as of Dec. 2025.




Robi Rahman and David Owen (2024), "The training compute of notable AI models has been doubling roughly every six months". Published online at epoch.ai. Retrieved from: <https://epoch.ai/data-insights/compute-trend-post-2010>.

The Blueprint for Yotta Scale Infrastructure



Leadership
Compute Capability



Open Rack
Architecture



**AMD
PENSANDO**


Ultra Ethernet
Consortium



ESUN



Open Fabric
to Scale AI




Rack Scale Efficiency
& Serviceability



Turnkey
Solutions

AMD “Helios” AI Rack





AMD Instinct™ MI455X GPUs

AMD EPYC™ "Venice" CPU

AMD “Helios”

The Open Rack Platform Towards Yotta Scale AI

2.9 Exaflops

AI Compute

31 TB

HBM4 Memory

43 TB/s

Scale Out Bandwidth

2nm / 3nm

Advanced Process

4,600


“Zen 6” CPU Cores

18,000

GPU Compute Units





AI COMPUTE PERFORMANCE



AI COMPUTE PERFORMANCE

A Massive Leap in AI Performance with **MI455X**



AMD Data Center Portfolio


Hyperscale AI



72 GPU Rack Scale

“Venice” + MI455X


Enterprise AI



8 GPU Drop-In Ready

“Venice” + MI440X

Sovereign AI & HPC



Hybrid Compute

“Venice-X” + MI430X

AMD
EPYC

“Venice”



AMD
INSTINCT

MI400 Series



AMD
PENSANDO

“Vulcano”

Open Ecosystem Drives AI Innovation

Open Hardware



Open Software



Open Ecosystem

Scale Up
Network

Scale Out
Network

Infrastructure

Software
Stack

Software
Ecosystem



ESUN

Ultra Ethernet
Consortium



AMD
ROCm

PyTorch
 Hugging Face

LLM
 SGL

AMD ROCm™ Platform

Adopted by the Most Popular Developer Communities

AI Framework



Inference Engines



Model Hub



Local LLM



Image & Video Gen



RL & Finetuning




Open Physical AI





Extending the Leadership Roadmap



AMD Instinct™
MI300A/X


2023



AMD Instinct™
MI325X


2024

AMD Instinct™
MI350 Series



2025

AMD Instinct™
MI400 Series



2026

AMD Instinct™
MI500 Series



2027

Next-Gen MI500 Series


Next Big Leap in AI Performance

AMD CDNA™ 6 | Advanced 2nm Process | HBM4E

>1000x

AI Performance in 4 Years

AI COMPUTE PERFORMANCE



AI Experiences are Evolving from Cloud to AI PCs



Private and Secure

Personal to Your Usage

Proactively Works for You

Performance You Can Depend on

AI PC Experiences Everywhere



Revolutionary
Content Creation



Assisted
Coding



Immersive
Collaboration




Enterprise
Productivity



Personal
Assistance


AMD is at the Forefront of AI PC Innovation

First
x86 NPU




Ryzen™ 7000
Series

First
**Copilot+
x86 Processor**



Ryzen™ AI 300
Series

First
**200B Model
Capable x86 Processor**



Ryzen™ AI Max
Series

INTRODUCING

AMD Ryzen™ AI 400 Series

Powering Next Generation AI PCs

12 / 24

“Zen 5” CPU Cores / Threads

60 AI TOPS

AMD XDNA™ 2 NPU

16

AMD RDNA™ 3.5 GPU Cores

8,533 MT/s

Memory Speed


Up To **1.7x**

Faster
Content Creation

Up To **1.3x**

Faster
Multitasking

VS INTEL CORE ULTRA 9 288V



AMD Ryzen™ AI 400

Available Starting Q1 2026



acer

ASUS

DELL
Technologies

hp

Lenovo

 Beelink®

COLORFUL

GIGABYTE™

LG

 **mechrevo**

msi

NEC

AMD Ryzen™ AI Max


The PC Reimagined for Creators, Gamers, and AI Developers



Content Creation



Gaming



AI Development

AMD Ryzen™ AI Max

The PC Reimagined for Creators, Gamers, and AI Developers

16 / 32

“Zen 5” CPU Cores / Threads

50 AI TOPS

AMD XDNA™ 2 NPU

40

AMD RDNA™ 3.5 GPU Compute
Units

128 GB

Unified Memory



AMD Ryzen™ AI Max+

Windows and Linux Support
Works with 1000s of Applications

vs Apple Macbook Pro M5

Up To
1.4x
Faster
AI Performance
tokens per second



Up To
1.8x
Faster
Content Creation

Asus Flow Z13

vs Nvidia DGX Spark

Up To
1.7x
Tokens/sec per Dollar
LM Studio: GPT-OSS 120B



Framework Desktop

Up To
1.8x
Tokens/sec per Dollar
LM Studio: GLM 4.5 Air

AI Development & Creation Unleashed

Infinite Possibilities with AMD Ryzen™ AI Max



ASUS

Lenovo™

acer

framework

xiaomi

abee

ASRock

Beelink®

CORSAIR

Emdoor^{DIGI}

GEEKOM

GMKtec

华红集团
HUARONG GROUP

MINISFORUM

nimo


六联智能
SIXUNITED

BOSGAME

INTRODUCING

AMD Ryzen™ AI Halo Processors

AI Developer Platform Available Q2 26



Go from Idea
to Workflow in Minutes

Linux and Windows Support


Optimized Applications

Preloaded Models

INTRODUCING

AMD Ryzen™ AI Halo Processors

AI Developer Platform Available Q2 26



Full AMD ROCm™ Platform Support

Dev-Ready Apps Optimized:
LM Studio, ComfyUI, VS Code, and more

Optimized Models Included:
GPT-OSS, FLUX.2, SDXL, and more

Day-0 Support for Leading AI Models

AMD Powers Over
1 Billion
Gamers & Creators



Advancing Healthcare



Genomic &
Computational Biology




Precision Medicine &
Drug Discovery



Diagnostics &
Imaging




Medical Devices &
Robotics



Patient
Care

Advancing Physical AI

High Performance x86 CPUs | Market-Leading Adaptive SoCs



Aerospace

Healthcare

Industrial

Robotics

Transportation

Smart Cities

Over 20 Years of Robotics Innovation

Industrial



Healthcare



Physical AI

Demands Real-Time Intelligence and Rapid Action



Perception

Sensor Fusion Understands
the World in Real-Time



Decision



AI Compute that Picks the
Right Answer Fast



Action

Deterministic Control that
Moves the System Instantly

The Blueprint for Physical AI Architecture



Leadership Compute
Architectures



Vision & Sensor
Processors



AMD
ROCm

 **ROS 2**™

yocto ·  MoveIt

PyTorch  OpenXen Project

N A V 2  OPEN
NAVIGATION

Open-Source
Ecosystem



Advancing Space Technology

Partnering to Further Breakthroughs and Discovery

BLUE ORIGIN

•esa



Jet
Propulsion
Laboratory

NASA



Communities Receive Earlier,
More Accurate Warnings

Development of Cleaner Energy
and Sustainable Materials

Faster Innovation
Across Industries

Lifesaving Vaccines
Reach Patients Sooner

AMD is the Compute Foundation for Modern Science



LUMI



2.7
Million Cores

Evaluating and Forecasting Geohazards and Natural Disasters

380
Petaflops

Predictive Modeling Under Varying Weather Conditions

Time to Generate Climate Forecasts Reduced by 85%



3.1
Million Cores

478
Petaflops

Developing Higher Performance Batteries

Optimizing Industrial Facilities

Improving Accuracy of Dynamic Studies for CO2 Storage





9
Million Cores

1.4
Exaflops

ORBIT-2 Delivers Hyper-Resolution Earth System Predictions Using Scalable AI

Produces Hyper-Detailed Global Weather Predicting Enabling Actionable Insights

Delivers Accuracy in the Range of ~0.98–0.99 When Compared to Observational Data





11
Million Cores

Pre-Emptively Optimized Antibody to Neutralize Broad Diversity of SARS-CoV-2 Variants

1.8
Exaflops

Anticipates Viral Evolution to Reduce Need for Constant Redevelopment

Improves Biodefense Preparedness and Antibody Therapy Resilience



Genesis Mission



First AI Factory for Science
Multi-Tenant AI Cloud Platform

A New Age of AI Accelerated
Innovation and Discovery



Next-Generation Exascale Supercomputer
Accelerating the Speed of Scientific Discovery

Advancing Education





Pledge to Advance
U.S. AI Education

\$150M
Investments



Engaging with Institutions
to Build Next-Gen AI Skills

800+
Institutions Worldwide



Empowering Everyone
to Shape the Future of AI

150,000+

Students in 2026 Through
AMD AI Education Program

AMD x HACK CLUB

Empowering AI Learning





Winning Team Armtender


Emme McDonald

Ruzanna Gaboyan

Afia Ava



AI Everywhere for Everyone



Unmatched Technology Portfolio



Open Ecosystem

Deep Co-Innovation

DISCLAIMER & ATTRIBUTIONS

DISCLAIMER: 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.

© 2025 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD CDNA, AMD RDNA, Radeon, Ryzen, Versal, XDNA and combinations thereof are trademarks of Advanced Micro Devices, Inc. Microsoft is a registered trademark of Microsoft Corporation in the US and/or other countries. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners. Certain AMD technologies may require third-party enablement or activation. Supported features may vary by operating system. Please confirm with the system manufacturer for specific features. No technology or product can be completely secure.

ENDNOTES

- GD-164: Day-0 driver compatibility and feature availability depend on system manufacturer and/or packaged driver version. For the most up-to-date drivers, visit AMD.com.
- GD-243: Trillions of Operations per Second (TOPS) for an AMD Ryzen processor is the maximum number of operations per second that can be executed in an optimal scenario and may not be typical. TOPS may vary based on several factors, including the specific system configuration, AI model, and software version.
- GD-250: Based on AMD chip shipment figures, AMD has powered over 1 billion gaming devices with processors and graphics for desktop PCs, notebook PCs, and gaming consoles from 2008 to 2025.
- MI350-047B: Based on engineering projections by AMD Performance Labs in September 2025, to estimate the peak theoretical precision performance of seventy-two (72) AMD Instinct™ MI455X GPUs "Helios" AI Rack using MXFP4 dense Matrix datatype vs. an 8xGPU AMD Instinct MI355X platform using the MXFP4 dense Matrix datatype. Results subject to change when products are released in market.
- MI500-001: Based on engineering projections by AMD Performance Labs in December 2025, to estimate the peak theoretical precision performance of AMD Instinct™ MI500 Series GPU powered AI Rack vs. an AMD Instinct MI300X platform. Results subject to change when products are released in market.
- SH0-42: Testing as of Oct/Nov 2025 by AMD CMPL. The systems tested were configured as follows: Ryzen AI Max+ 395, 55W, Asus Flow Z13, best performance, 8x4GB @8000MHz, AMD Radeon 8060S graphics, VBS=ON, SAM/ReBar=ON INT vs. an Apple M5 10 Core CPU, Apple Macbook Pro 2025 14in, 24GB RAM, Balanced, 10 core GPU, VBS=not supported SAM/Rebar=not supported Int, power=DC. The applications tested include: 7zip, Handbrake, Blender CPU, VRay 6.0, Corona, and LM Studio. System manufacturers may vary configurations yielding different results. Results may vary.

ENDNOTES

- SHOP-26: Testing as of November 2025 by AMD. All tests conducted in LM Studio 0.3.30 (Build 2). Vulkan llama.cpp v 1.57.1 used with Ubuntu 24.0.4.3 and therock-gfx1151-7.9rc1 for AMD Ryzen™ AI Max+ 128GB. CUDA llama.cpp v1.57.1 used with pre-installed DGX OS (based on Ubuntu 24.04.3) Driver Version 580.95.05 and CUDA Toolkit Version 13.0 Update 2 for NVIDIA DGX Spark. Flash Attention = ON in all cases. Token/s: sustained performance average of multiple runs with specimen prompt “How long would it take for a ball dropped from 10 meter height to hit the ground?”. Models tested: OpenAI GPT-OSS 120B, OpenAI GPT-OSS 20B, GLM 4.5 Air and DeepSeek R1 Distill Llama 70b. Tokens per second per dollar measured using market pricing of \$2500 for the HP Z2 G1a and \$4000 for the NVIDIA DGX Spark as of November 2025. AMD Ryzen™ AI Max+ 395 PRO on an HP Z2 Mini G1a with 128GB memory and NVIDIA DGX Spark with 128GB memory. Performance may vary.
- SHOP-27: Testing as of November 2025 by AMD. All tests conducted in LM Studio 0.3.30 (Build 2). Vulkan llama.cpp v 1.57.1 used with Ubuntu 24.0.4.3 and therock-gfx1151-7.9rc1 for AMD Ryzen™ AI Max+ 128GB. Flash Attention = ON in all cases. MMLU and GPQA scores as-reported from research papers and github repos. Cloud-quality statement from OpenAI “The gpt-oss-120b model achieves near-parity with OpenAI o4-mini on core reasoning benchmarks.” AMD Ryzen™ AI Max+ 395 PRO on an HP Z2 Mini G1a with 128GB memory. 200 billion parameters require 128GB of unified memory. The AMD Ryzen™ AI Max+ was the first x86 processor to launch with 128GB of unified memory. Performance may vary.
- GPT-2: Testing done as of November 2025 by AMD to measure multitasking performance by running Procyon Office Suite with Microsoft Teams on “Balanced” power mode. Configuration for AMD Ryzen™ AI 9 HX 470 processor: AMD reference board, Radeon™ 890M integrated graphics, Graphics Driver 25.20.24-251007a, 32GB LPDDR5x-8533 memory, Windows 11 Home. Configuration for Intel Core Ultra 9 288V processor: HP OmniBook Ultra Flip, Arc 140V integrated graphics, Graphics driver 32.0.101.8136, 32GB LPDDR5x-8533 memory, Windows 11 Home. System manufacturers may vary configurations, yielding different results.
- GPT-10: Testing done as of December 2025 by AMD to measure content creation performance in the following applications and benchmarks on “Best Performance” power mode: Blender (CPU Classroom), Cinebench nT, Handbrake, PugetBench for Photoshop, PugetBench for DaVinci Resolve Studio, 7zip. Configuration for AMD Ryzen™ AI 9 HX 470 processor (28W): ASUS Zenbook S16, Radeon™ 890M integrated graphics, Graphics Driver 25.20.32-251114n, 32GB LPDDR5x-8533 memory, Windows 11 Pro. Configuration for Intel Core Ultra 9 288V processor (30W): HP OmniBook Ultra Flip, Arc 140V integrated graphics, Graphics driver 32.0.101.7026, 32GB LPDDR5x-8533 memory, Windows 11 Home. System manufacturers may vary configurations, yielding different results.