

AMD FINANCIAL ANALYST DAY 2025

Mark Papermaster
EVP & CTO




together we advance_

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 products, including "Helios" rack-scale platform; leadership CPU core roadmap; data center GPU architecture roadmap; AMD's ability to accelerate GPU performance gains; high-speed SerDes interconnect; chiplet and packaging leadership; AMD Infinity Fabric™ roadmap; Gaming GPU architecture roadmap; NPU inference IP Roadmap or client and embedded; AMD's opportunities in AI; and AMD's focus areas, 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.


Industry's Broadest Solution Portfolio



Industry's Broadest Solution Portfolio



We Called the Play



Financial Analyst Day • 2015

We Delivered

Industry Leadership “Zen” x86 CPU


- Consistent Leadership Performance
- Chiplet Enabled Core Density Leadership
- High-Performance & Energy Optimized

AMD Instinct™ GPU with Leadership Inference Performance

- 2.5D Integrated HBM & 3D Chiplets Drove Memory & Compute Density Leadership
- TCO Advantages for Key AI Workloads

Breakthrough SoC Approach


- Scalable Infinity Fabric Enabled Modular Architecture
- First Chiplet Design Methodology



We Are Accelerating **Optimizing AI Across Portfolio**

- **Compute Leadership Roadmap from Laptops to Rack-Scale**
- **Driving a Truly Open Software & Hardware AI Ecosystem**
- **SerDes, Optical & Networking Investments**
- **Confidential Computing for Trusted AI Processing**

Leadership CPU Core Roadmap




*All roadmaps subject to change.

AMD Instinct™: Data Center GPU Architecture Roadmap



*All roadmaps subject to change.

Accelerated GPU Performance Gains



- 2x Per Two Year Performance Trend Pre-2025 (16b Math)
- 2x Per One Year Performance Trend Commences in 2025
- AI-Optimized Low Precision Math (FP8 & FP4) Provide Even Higher Gains

High-Speed SerDes Interconnect

40 Gb/s

First Server PCIe® 5.0
10s of Millions in Production

112 Gb/s

First 112G Ethernet FPGA
Leading Edge PCIe 6.0
PAM4

224 Gb/s

Leading Edge PCIe 7.0
PAM4
2026 “Helios” Rack
Interconnect

Next Gen

Rack Transition Point
Copper & Optical Solutions

2022

2026

AMD Infinity Fabric Roadmap

3rd Generation

Up to 8-Way
GPU Connectivity

Coherent Memory for
Exascale Computing

PCIe® 5.0

4th Generation

Heterogeneous Chiplet

CXL™ 2.0 Memory Expansion

Support for Neural Engine

Introducing 5th Generation

Scale Up & Scale Out Leadership

Unified System Level Coherency

Standard Based Open Ecosystem
(UALink™, CXL 3.1 & UCle™)

PCIe 6.0


Next-Gen

Continued Scale

2022

2026

Chiplet & Packaging Leadership



Solution Example

“Helios” Rack


Link at Rack-Scale – Scale Up & Scale Out

- AMD Infinity Fabric for CPU to GPU link
- UALink™ Over Ethernet for Scale-Up
- UALink for GPU to AI Network Interface Card (NIC)
- AI NIC for Scale Out with Ultra Ethernet (UEC)

Hardware-Software-Application Co-Optimization

- AMD ROCm™ Inference & Training Scale with Compute & Communication
- Day-0 Optimization of Most Popular Models & Frameworks





AMD Instinct™ GPUs



- CDNA 5
- 5th Gen Infinity Fabric at 112 Gb/s
- 3.5D Packaging
- 224 Gb/s SerDes

EPYC™ CPUs



- “Zen” 6
- 5th Gen Infinity Fabric
- 2.5D Packaging
- 64 Gb/s SerDes


AI NICs




- 800G Ethernet
- 224 Gb/s SerDes

Trusted Computing

- Top Priority for IT Administrators: Protecting AI Models, Data & Weights
- Solution Enables End-to-End Encryption
- Confidential Computing Including Memory, Storage & I/O Devices




Gaming GPU Architecture Roadmap



2021

2025

NPU Inference IP Roadmap for Client & Embedded





AI Beyond the Data Center

One Software Stack Enables AI from Cloud to Edge

AI Models & Algorithms



Libraries

Compilers & Tools

Runtime



AMD Instinct™



EPYC™



Radeon™ AI



Ryzen™ AI



Embedded

Train Once • Infer Everywhere

Investments for the Future



Data Center AI Compute Leadership

Rack-Scale GPU &
Open Ecosystem



Agentic AI

Leadership CPUs to
Accelerate Agents



Edge AI Explosion

Integrated CPU,
GPU & NPU



Physical AI

Leadership
Embedded FPGA,
CPU & GPU



Quantum Computing

Hybrid Quantum with
AMD Compute Cluster

Innovation for the AI Era

**Product
Leadership**

**Broadest
IP Portfolio**

**Open
Ecosystems**

**Relentless
Execution**

Disclaimer & Attribution

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 Instinct, EPYC, Pensando, Radeon, ROCm, Ryzen, Versal, Xilinx, and combinations thereof are trademarks of Advanced Micro Devices, Inc. CXL is a registered trademark of Compute Express Link Consortium, Inc. OpenAI is a trademark of OpenAI, Inc. PCIe® is a registered trademark of PCI-SIG Corporation. UCIE is a trademark of Universal Chiplet Interconnect Express, Inc. Ultra Accelerator Link and UALink are trademarks of the UALink Consortium. 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.