



# High Performance & Adaptive Computing

Solving the world's most important challenges

Journée AuDACES 2026

- Patrick Masse / Field Support Mgr  
• [patrick.masse@amd.com](mailto:patrick.masse@amd.com)
- Anthony Clary / BDE Public Sector  
• [anthony.clary@amd.com](mailto:anthony.clary@amd.com)

• 18 juin 2026



# Strategic Pillars

## Vision

High performance and adaptive computing is transforming our lives

## Strategy

1

Extend leadership in foundational technology and IP

2

Provide competitive software, tools and open-source enablement

3

Leverage core tech across platforms to achieve economies of scale

4

Build effective ecosystems with disciplined partnerships


5

Build scalable, cross-portfolio GTM engine


## Advancing end-to-end AI infrastructure



Data Center



End User Devices



Embedded

# A propos d'AMD



**57**

Years

Founded May 1, 1969  
HQ Santa Clara, CA

**31,000+**

employees worldwide

~14,000 + in North America

**#167**

Fortune 500

Among 8 companies with >30%  
annualized 10-year  
shareholder return

**~\$34.6 B**

annual revenue FY 2025

5X annual revenue growth  
over 10 years

**18,800**

patent assets worldwide<sup>1</sup>

Spans hardware, software, AI,  
and related technologies

**\$6.5B**

annual R&D spend

25% of revenue invested  
in R&D in 2024

<sup>1</sup>AMD 10 K annual report filing, February 5, 2025. Includes patents and pending patent applications



# Ressources AMD secteur Public



**Philippe Grégoire**  
Commercial ESR / Défense



**Patrick Masse**  
Spécialiste Serveurs  
Grenoble



**Benjamin Pajot**  
Spécialiste HPC (Centre  
d'Excellence)  
Grenoble



**Cyril Laurie**  
Directeur Avant Vente  
IDF



**Mathieu Gontier**  
Spécialiste HPC (Centre  
d'Excellence)  
Aix en Provence



**Anthony CLARY**  
Commercial Secteur Public FR



**Arnaud Touré**  
Serveur / Client / GPU  
IDF

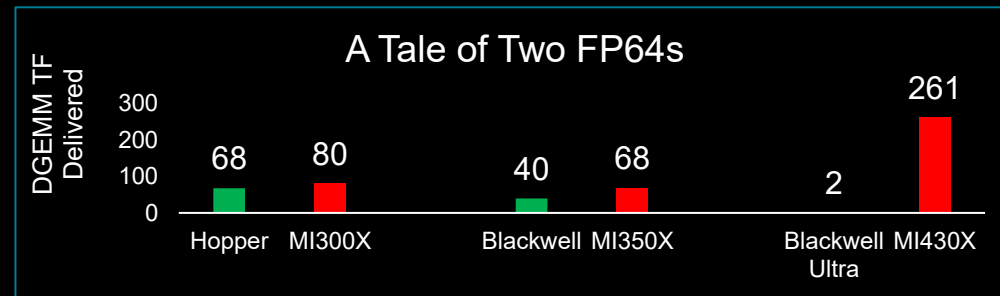


**Apostolos Tsetoglou**  
Avant Vente Clients / Serveurs  
Rouen



# AMD SOVEREIGN AI & HPC PROGRAM LEADERSHIP

- **Lux** DOE|ORNL 1st U.S. AI factory powered by AMD Instinct™ MI355X GPUs
- **Discovery** Next-Gen DOE|ORNL flagship system designed to drive breakthroughs in science, energy, and national security. Powered by AMD Instinct™ MI430X GPUs + AMD EPYC™ “Venice” CPUs + HPE Slingshot
- **Alice Recoque** Expected to be the largest European Sovereign AI system to-date. Powered by AMD Instinct™ MI430X GPUs + AMD EPYC™ “Venice” CPUs.
- **El Capitan** World’s most powerful supercomputer designed to excel on the most challenging HPC & AI workloads. Powered by AMD Instinct MI300A APUs.
- **Frontier** World’s 1<sup>st</sup> system to cross the Exascale barrier with 1.1 exaflops of performance



The Journey to Exascale [Brochure](#)  
Precision at Scale [White Paper](#)

[AMD.com/HPC](https://www.amd.com/HPC)



2022 | MI250X  
*World’s First Exascale*



2024 | MI300A  
*World’s Most Powerful*



2026 | MI350X  
*1<sup>st</sup> U.S. AI Factory*



2027 | MI430X  
*Next-Gen European Sovereign AI & HPC system*



2028 | MI430X  
*Next Gen U.S. Sovereign AI & HPC system*

Source: <https://top500.org/lists/>, November 2025

# Education Supérieure et Recherche

## IA Clusters in France :

- **PostGenAI@PARIS (Sorbonne université) – 35 M€**
- **PR[AI]RIE – PSAI (Université Paris Sciences et Lettres) – 75 M€**
- **MIAI Cluster (Université Grenoble Alpes) – 70 M€**
- **Hi! PARIS Cluster 2030 (Institut Polytechnique de Paris) – 70 M€**
- **ENACT (Université de Lorraine) – 30 M€**
- **DATAIA-Cluster (Université Paris Saclay) – 20 M€**
- **ANITI IA Cluster (Université de Toulouse) – 20 M€**
- **3IA Côte d'Azur 2030 (Université Côte d'Azur) – 20 M€**
- **SequoIA (Université de Rennes) – 20 M€**



## AMD University Program

Hub for educators, researchers and students to access AMD resources, programs and news

# AMD et le gouvernement français

Signature d'une lettre d'intention entre le gouvernement français et AMD le 16 avril 2026

Le Gouvernement Français et AMD renforcent leur collaboration pour accélérer l'innovation en intelligence artificielle, la recherche et le développement d'un écosystème ouvert en France.

Objectif :

- Mettre à disposition de l'enseignement supérieur, de la recherche et de l'espace des moyens de calcul GPU AMD (ressources matérielles et accès à distance aux ressources AMD AAC (AMD Accelerator Cloud))
- Accès au futur supercalculateur Alice Recoque.
- Mise en place du programme AMD University Program
- Accompagnement de la communauté (événement Hackathons, POC, démonstrateurs, etc..)
- Education : AMD Developer Cloud, AI Academy



<https://www.amd.com/en/newsroom/press-releases/2026-4-16-amd-and-the-french-government-announce-plans-to-advance.html>

# AMD University Program - AUP



## Students

**Affordable access  
to AMD products**

Build industry relevant skills and  
knowledge with AMD products



## Educators

**University resources  
and support**

Deliver best-in-class teaching with  
AMD educational resources and  
support



## Researchers

**Leading-edge  
AMD technologies**

Advance state-of-the-art  
research

# Accélérer l'apprentissage IA avec AUP

📅 30. Juin, 10h00

⚙️ ⚙️ **Durée : 90-minutes (présentation 60 minutes + lab 30 minutes)**

🔗 🖥️ **Zoom : [https://amd.zoom.us/webinar/register/WN\\_cKORYwJGSWCywu7jiFI59w](https://amd.zoom.us/webinar/register/WN_cKORYwJGSWCywu7jiFI59w)**

## De programmes et modules de laboratoire en IA prêts à l'emploi :

- Le **cloud d'apprentissage AMD** (open source + plateforme matérielle).
- Options de déploiement (systèmes locaux ou clusters distants).
- Parcours pratiques pour intégrer **les flux de travail modernes en IA** (LLM, IA multimodale, systèmes agents, calcul GPU) dans vos cours.
- Nous montrerons également comment les universités / écoles peuvent **rapidement permettre un enseignement pratique de l'IA sans construire une infrastructure à partir de zéro.**

## Le Programme Université AMD (AUP) propose :

- Un matériel pédagogique et tutoriels préparés pour vous.
- Une pile logicielle entièrement ouverte en open source (y compris la plateforme basée sur JupyterHub)
- Des environnements de calcul évolutifs (local ou à distance)
- Des opportunités de **support matériel / co-investissement** pour des partenaires sélectionnés
- Et, pour les ateliers pratiques associés, un accès à distance au AUP Learning Cloud sera fourni.

## Objectif de cette session :

Selon **format sur invitation uniquement**, cette session vise à explorer :

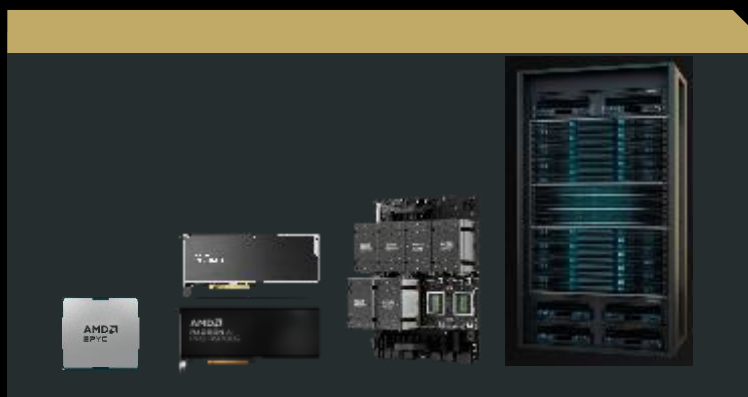
- Utilisation potentielle dans votre enseignement et votre recherche
- Opportunités de collaboration conjointe avec AMD
- Les prochaines étapes vers un engagement plus profond au sein de l'écosystème de la PU

Lien zoon



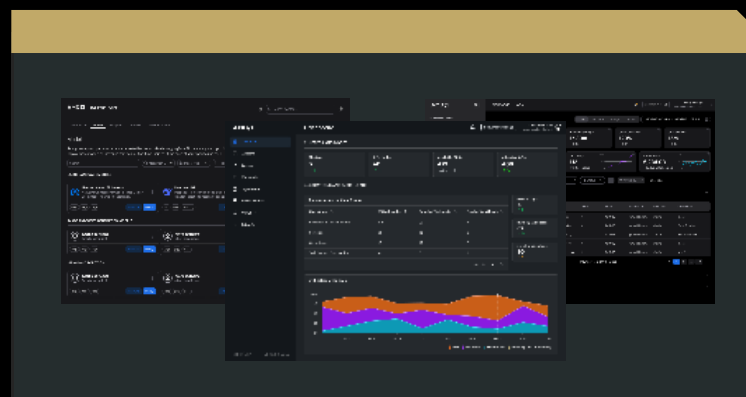
# AMD Enterprise AI Solutions

Seamless Enablement of High Value Enterprise Applications



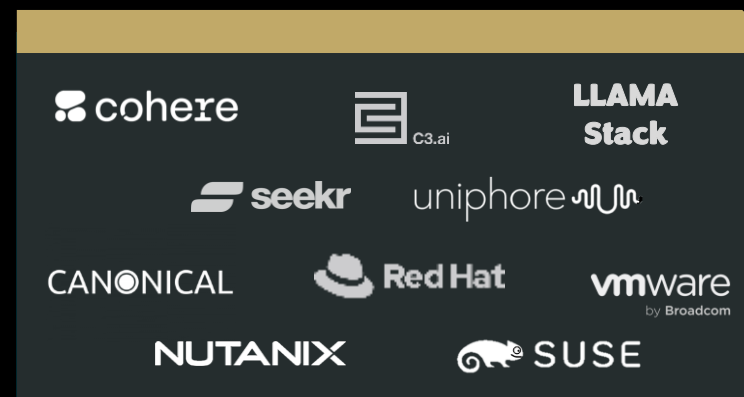
## High Performance Platforms

Delivering performance and value; On-prem & Cloud



## Enterprise AI Software

Complete solutions for deploying & customizing industry-standard AI models




## Strategic Partnerships

Enable high-impact use cases optimized for AMD platforms

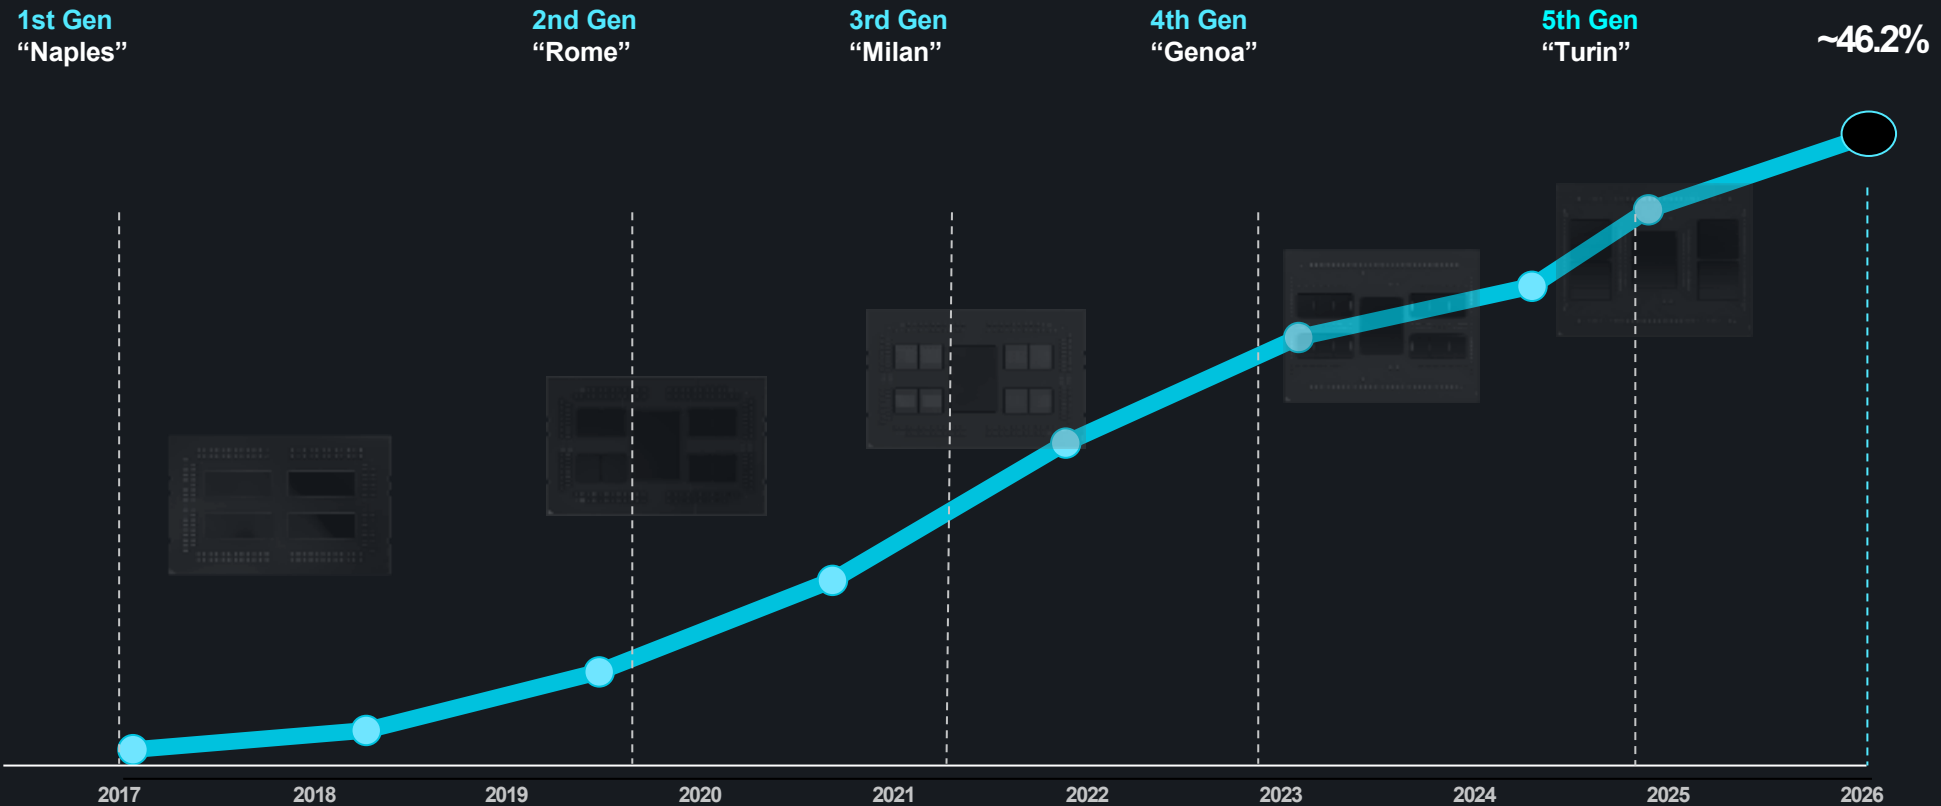


# AMD EPYC™ Processors



# AMD EPYC™ Trusted to Power Over 46% of the World's Servers

## >18x Server CPU Market Share Growth

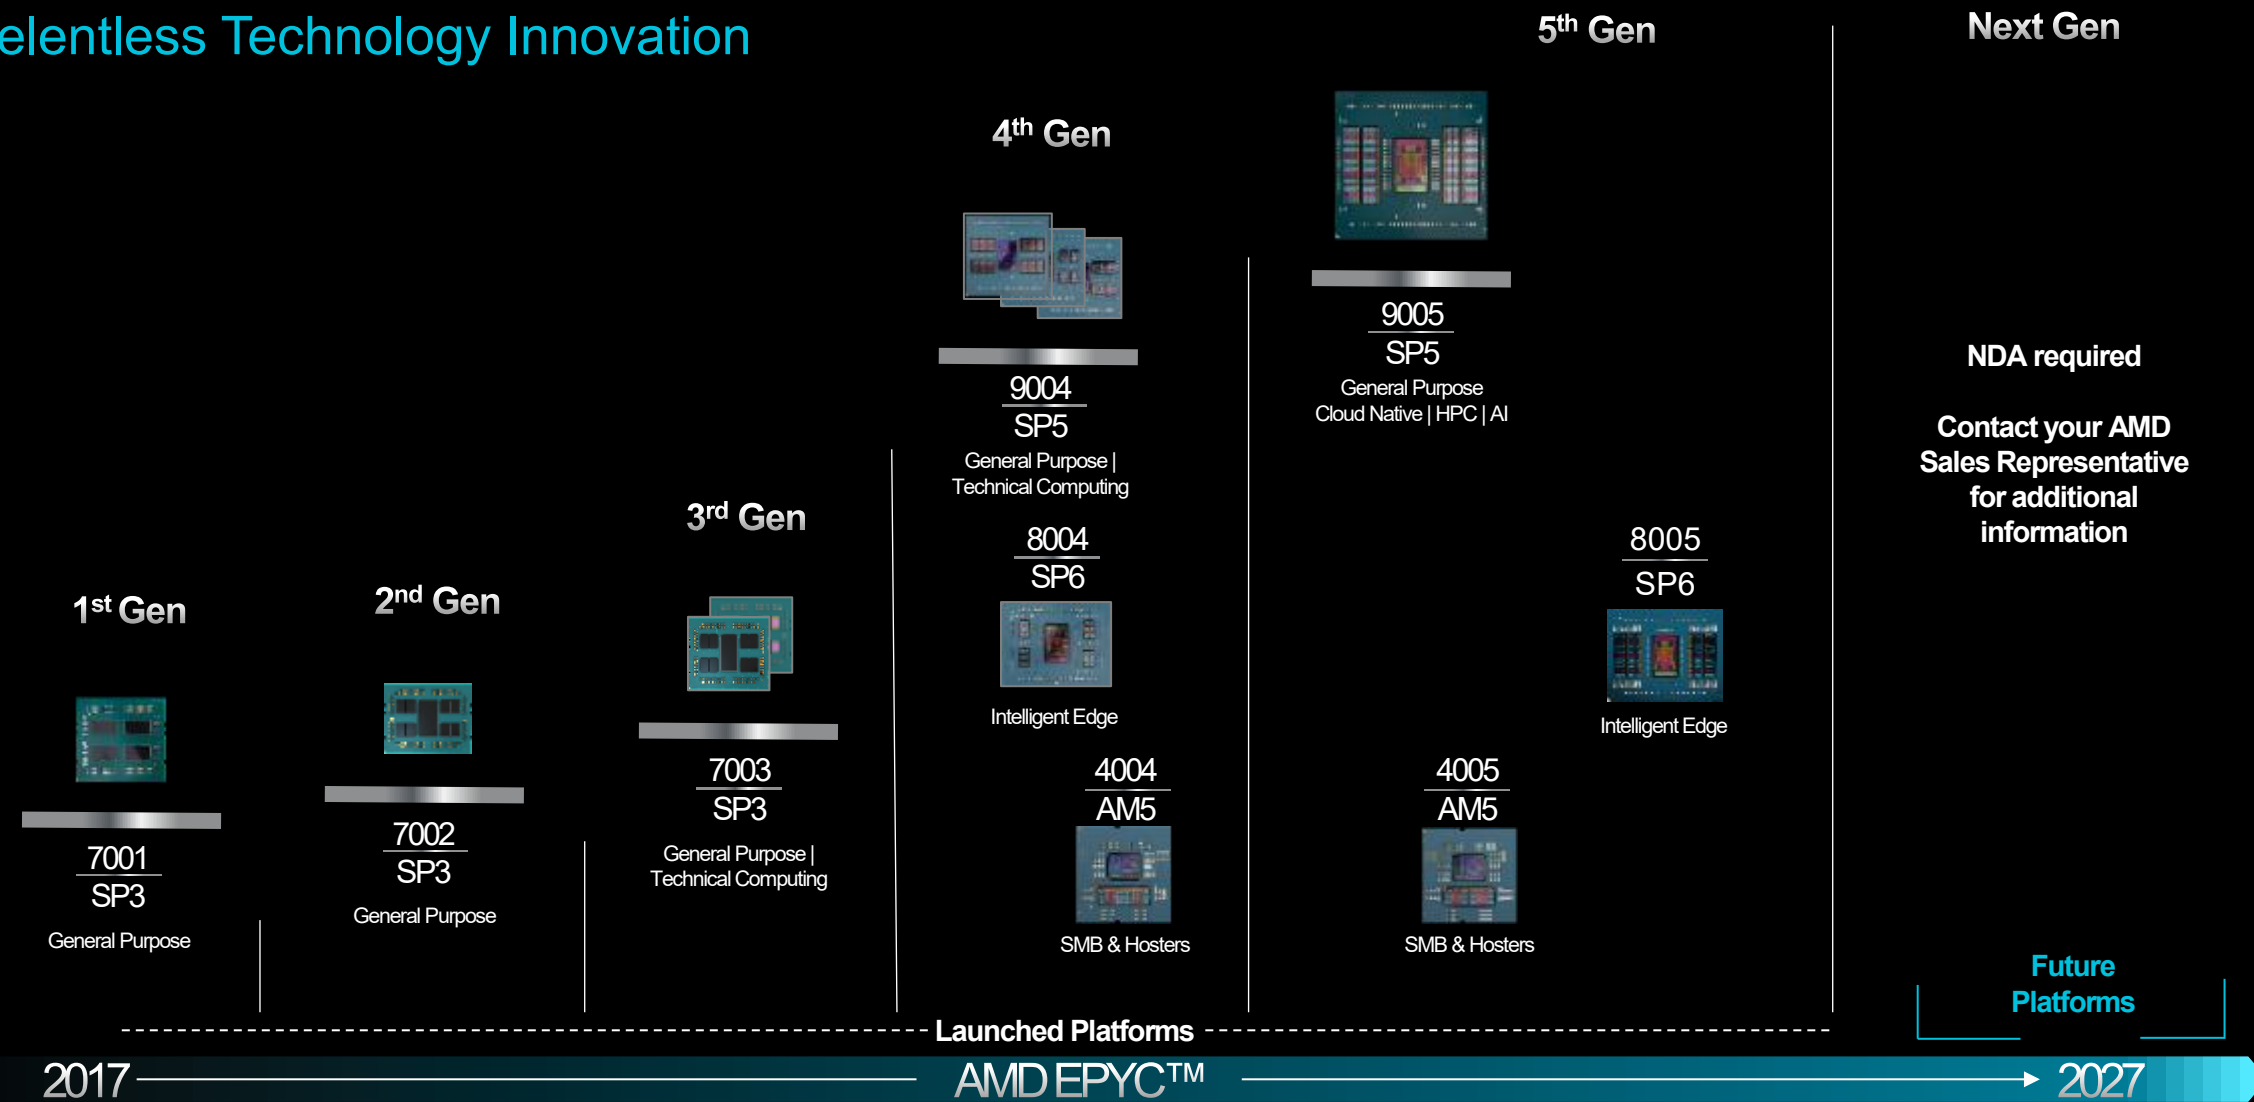


Source: Mercury Research Sell-in Revenue Shipment Estimates, Q4'25

EPYC-055D - Mercury Research Sell-In Revenue Shipment Estimates, Q1 2026. Revenue share of 46.2%, unit share of 27.4%.

# AMD EPYC™ Server CPU Roadmap


## Relentless Technology Innovation



# 5<sup>th</sup> Gen AMD EPYC™ Processors

Formerly codenamed “Turin”

## World’s best CPU for cloud, enterprise & AI



**TSMC 3/4nm**

Up to **192 cores**  
Up to **384 threads**

Up to **5GHz**

**AVX512**  
full 512b data path

**17%**  
Enterprise IPC Uplift

**37%**  
HPC/AI IPC Uplift

**SP5 Platform**  
Compatible with “Genoa”

# AMD EPYC™ CPUs for AI

## AI Inference



Generative AI



Recommendation Systems



Natural Language Processing



Classical Machine Learning

## Enabling AI



Retrieval Augmented Generation [RAG]



Pre & Post - Processing



Data Cleansing & Synthesis



Speculative Decoding

## Agentic AI



Tools (General Purpose Applications)



Opportunistic Agents [SLM]



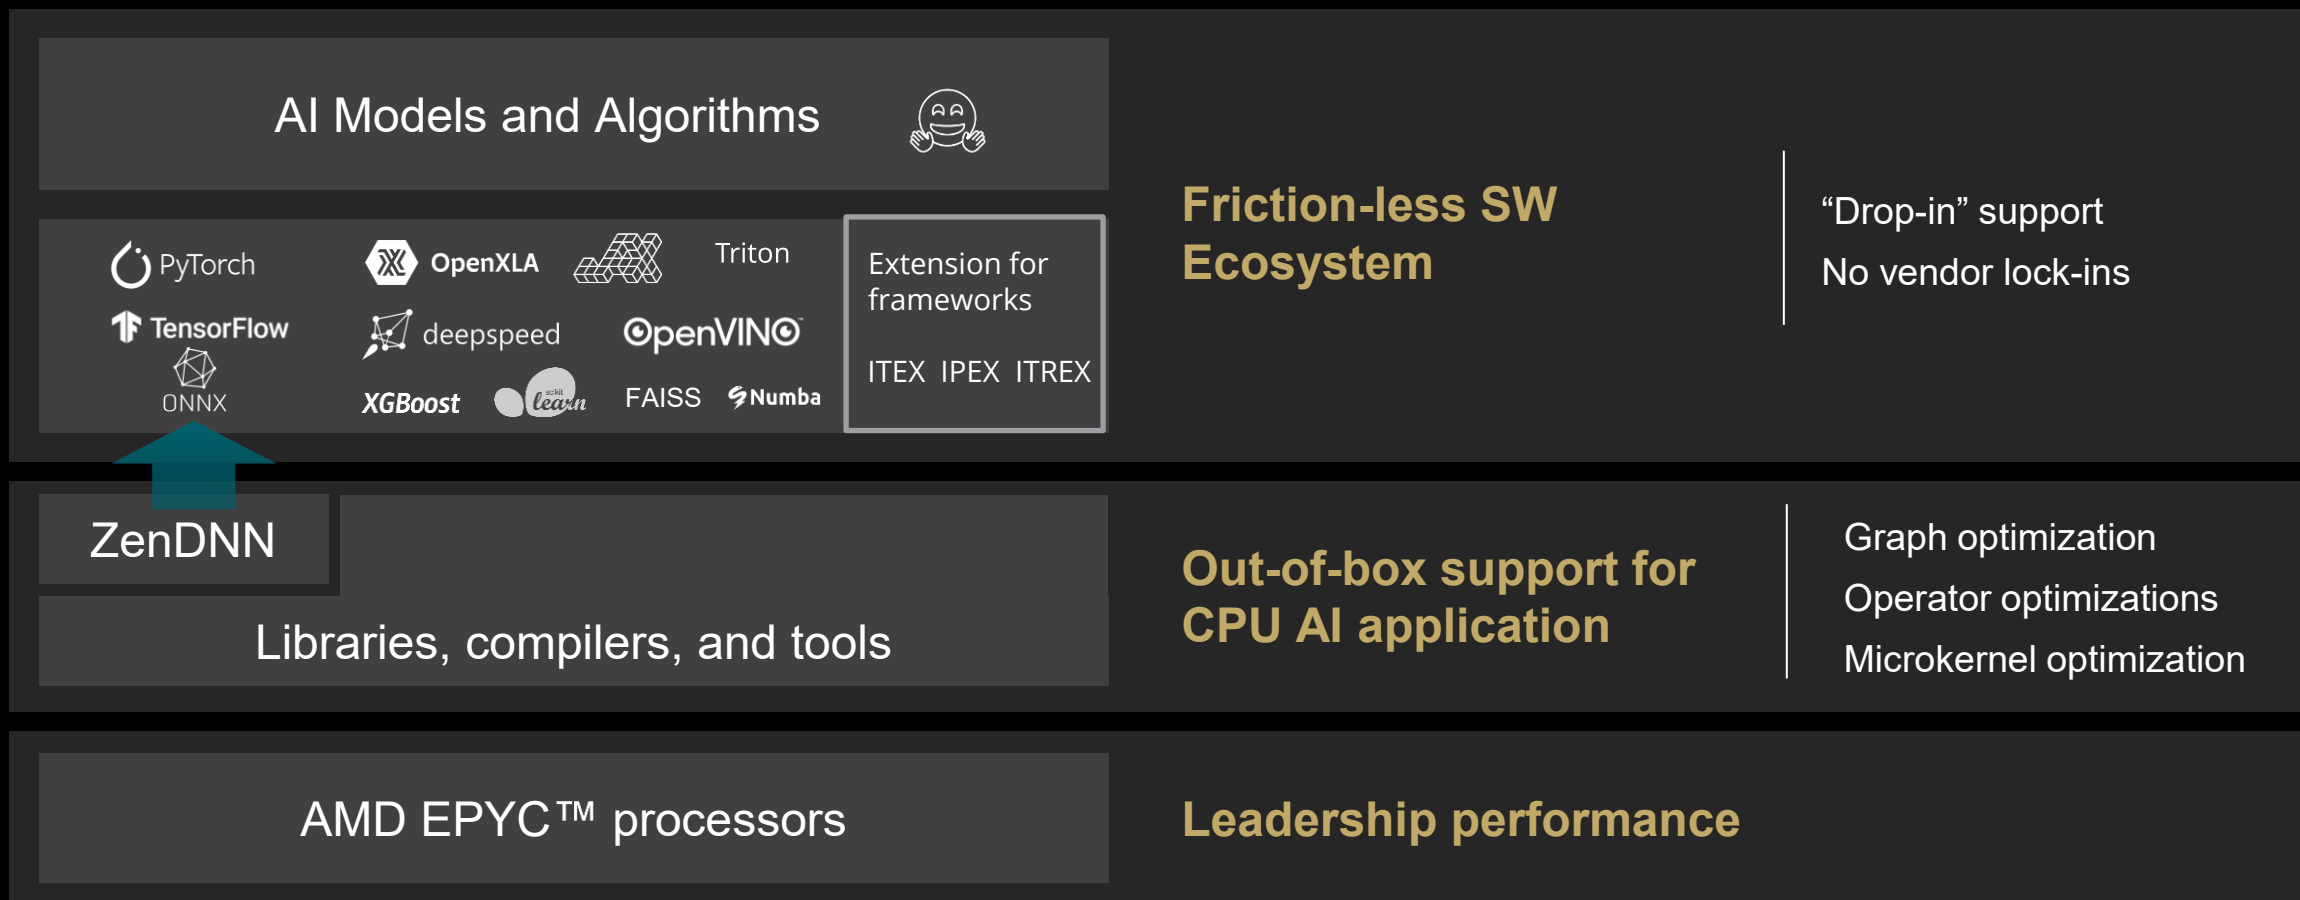
Orchestration, Observability, Monitoring



Storage, Vector Database, Memory

# AI on AMD EPYC™ Processors


If your code runs on x86, it runs on AMD EPYC™ processors



Use or mention of third-party marks, logos, products, services, or solutions herein is for informational purposes only and no endorsement by AMD is intended or implied. GD-83c



# AMD INSTINCT™ ACCELERATORS



# AI Leadership GPUs



## Performance

MXFP4 / MXFP6 / FP8 /  
FP16 / BF16 / FP32  
throughput

## Memory Capacity

Large models require  
massive memory

## Scalability

Infinity Fabric™  
supports Multi-GPU  
scaling efficiency

## Open Software Ecosystem

Broad framework  
support  
Scalable deployment  
Comprehensive  
developer tools

## Cost

Leadership  
performance per watt  
Low cost per token

# Right Size Your AI

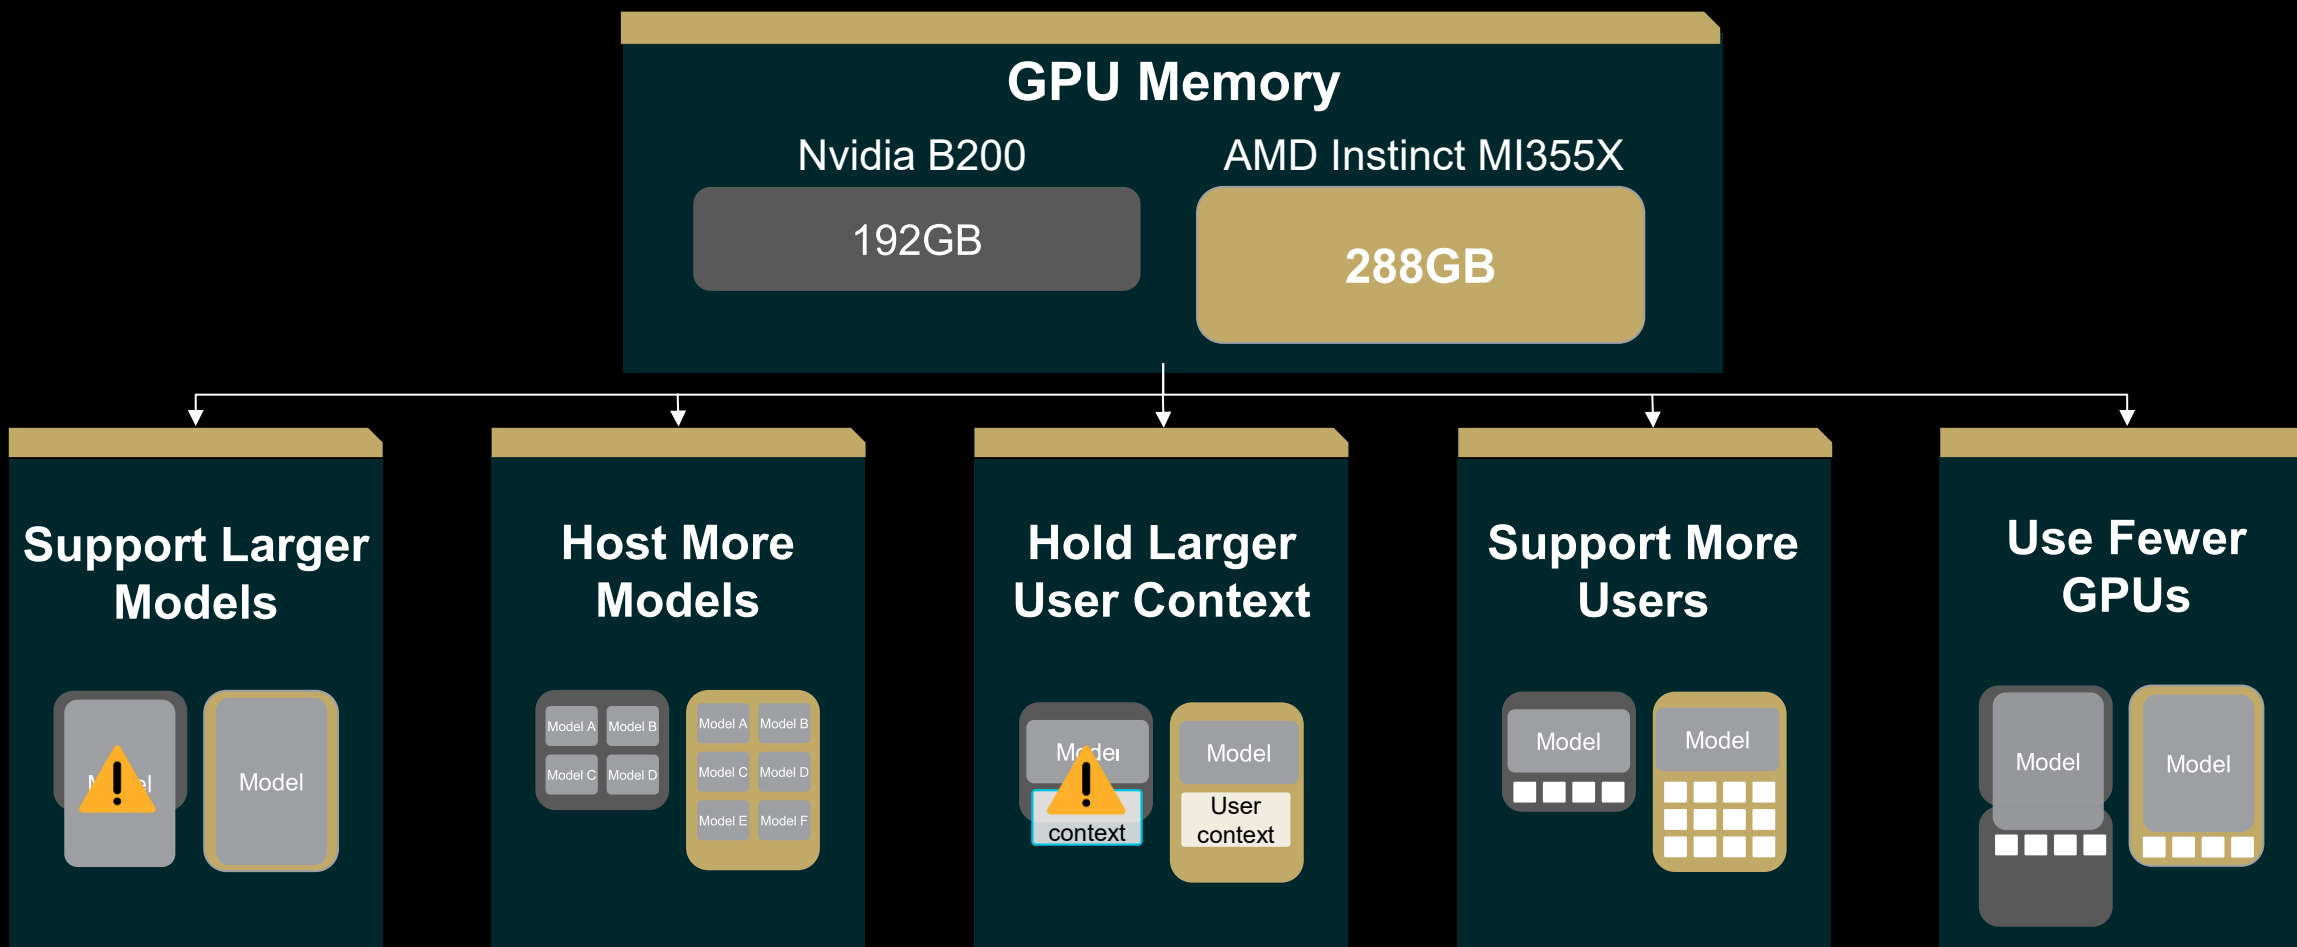
## GPU Options Aligned to Your Business Needs



GPU	AMD Radeon™ AI PRO R9700 (PCIe)	AMD Instinct™ MI350P (PCIe®)	AMD Instinct™ MI355X x8 (UBB)
<b>Form Factor</b>	Workstation/Server	Server	AI-Optimized Server
<b>Memory</b>	32GB	144GB	2,300GB
<b>Concurrent Users</b>	100s	1,000s	10,000s
<b>Model Size</b>	20B+	175B	1T+
<b>Rack Power</b>	Air   10-20kW	Air   15-30kW	Air or Liquid   30-60kW
<b>Use Case</b>	<ul style="list-style-type: none"> <li>• Edge AI, retail</li> <li>• Image generation</li> <li>• Traditional AI (CNN, DLRM)</li> <li>• Small language model inference</li> </ul>	<ul style="list-style-type: none"> <li>• Mainstream LLM inference</li> <li>• Content generation</li> <li>• RAG-enabled chat</li> </ul>	<ul style="list-style-type: none"> <li>• Large model inference</li> <li>• Model fine tuning</li> <li>• Small model training</li> <li>• Code generation</li> </ul>

# AMD Instinct™ GPU Memory Capacity

Brings Real World Benefits



Graphics for illustrative purposes only

# AMD Instinct™ MI355X GPU



Instinct™ MI355X	vs. Nvidia B200	Why it matters
MEMORY CAPACITY	1.6x	Supports larger models
FP6 Tensor/MXFP6 Matrix	Up to 2.2x	Better enterprise AI support
PEAK FP64/PEAK 32	Up to 2.1x	Better enterprise HPC support

# AMD Instinct™ MI350P PCIe® Card

## Deploy and Scale in Existing Data Center Infrastructure

### Built for Existing Enterprise Infrastructure

- 10.5" FHFL Dual-Slot Design
- 600W Passive Air-cooled
- 450W Power Cap Configurable
- PCIe Gen 5 Host Interface
- Enterprise AI Software
- Inference Microservices Support



### Leadership GPU Architecture

- AMD CDNA™ 4 Architecture
- 128 Compute Units
- Support for MXFP8, MXFP6, MXFP4<sup>1</sup>
- 144GB HBM3E Capacity
- 4.0 TB/s Peak Bandwidth
- Up to 4 Physical Partitions at 36GB Each
- Dedicated Video and JPEG HW Decode

### Open Enterprise Ready Software

- GPU Independence with Cross Platform Interoperability
- License-Free, Open Source, Standards-based Software
- Broad Ecosystem to support Custom Software Stacks

<sup>1</sup>AMD Scaled data formats conform to OCP standard MX data types.





**ROCm SW**



# AMD

## Leadership Software Stacks



**Industry's Premier  
Open-Source AI Stack**

**Native Support in  
Largest Open-Source  
Projects and Day-0  
Models**

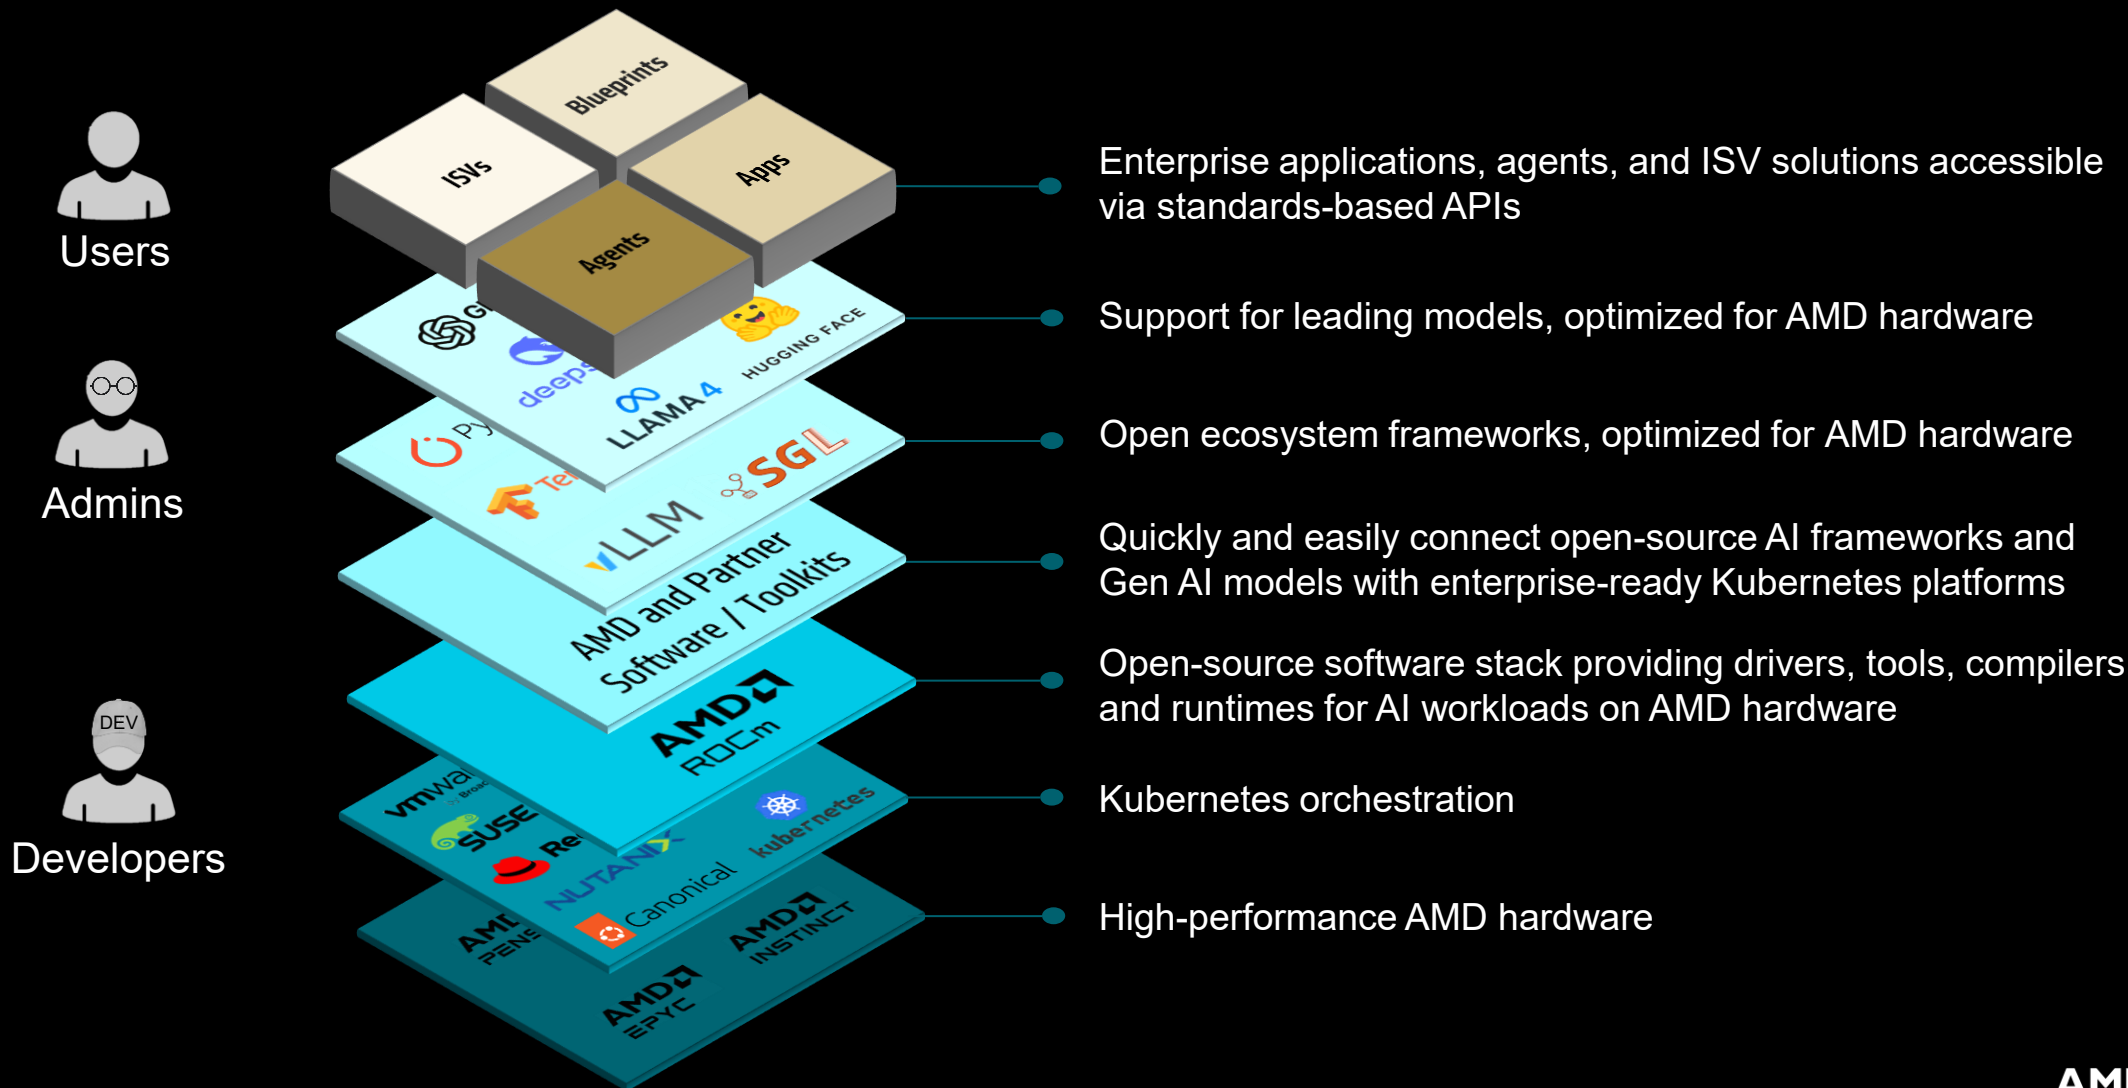
**Enterprise-Optimized  
Solutions**

**Empowering an  
Expanded Set of  
Developers**

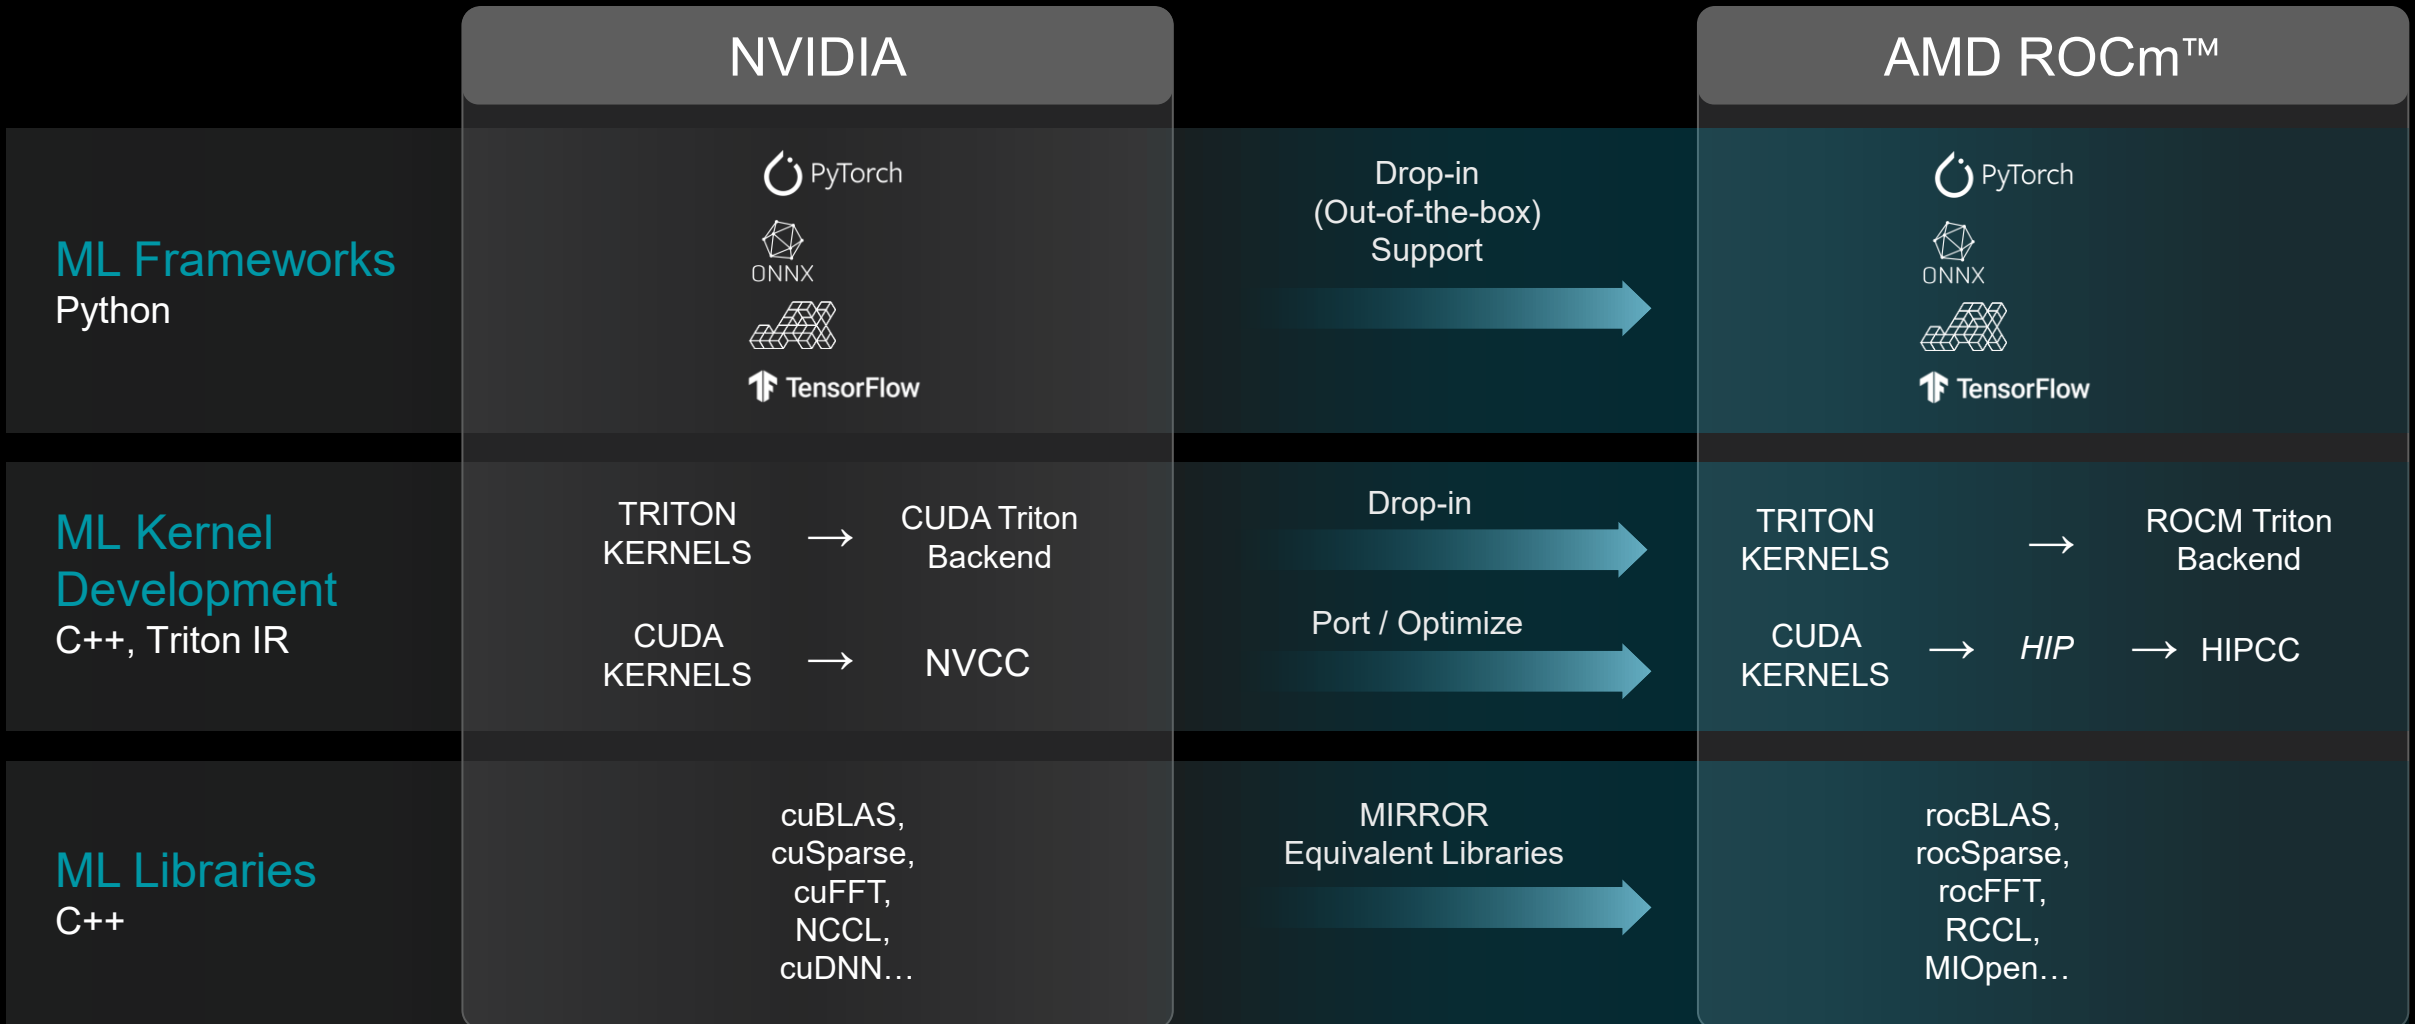
**Unified Software  
Ecosystem**

# Powering Enterprise AI Across the Full Stack

## Open, Optimized, Enterprise Ready

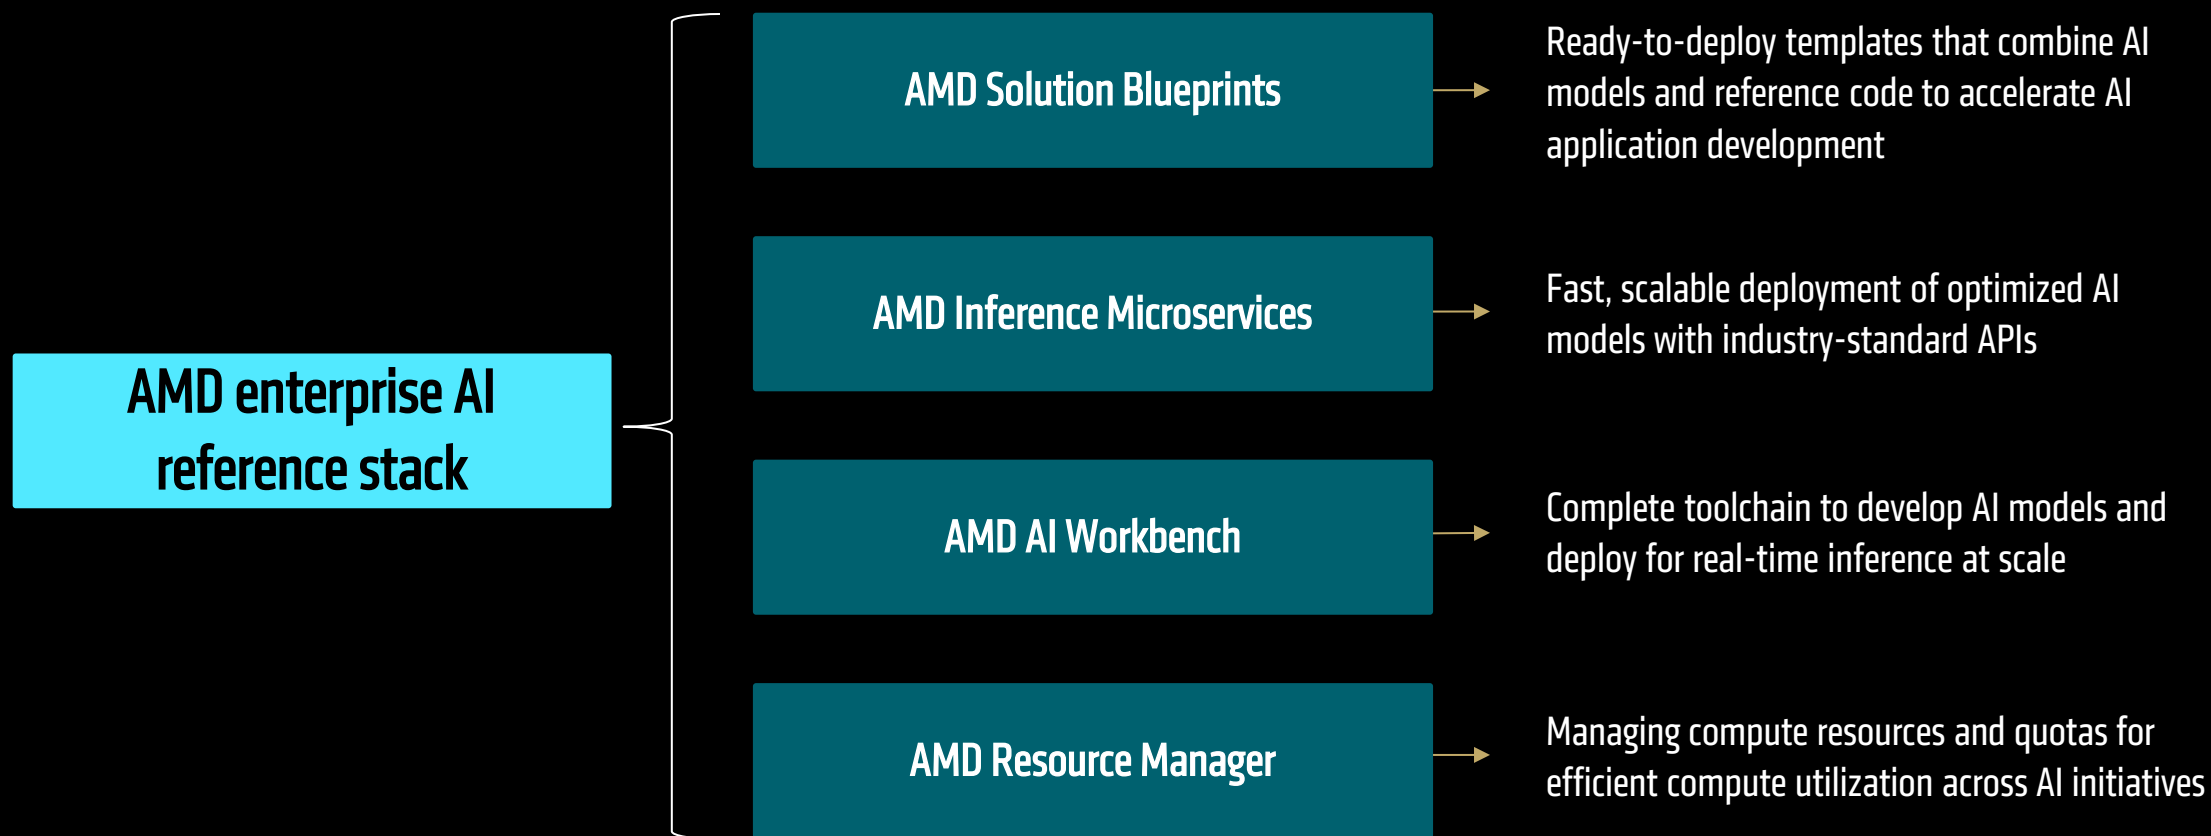


# Transitioning AI Workloads to AMD GPUs

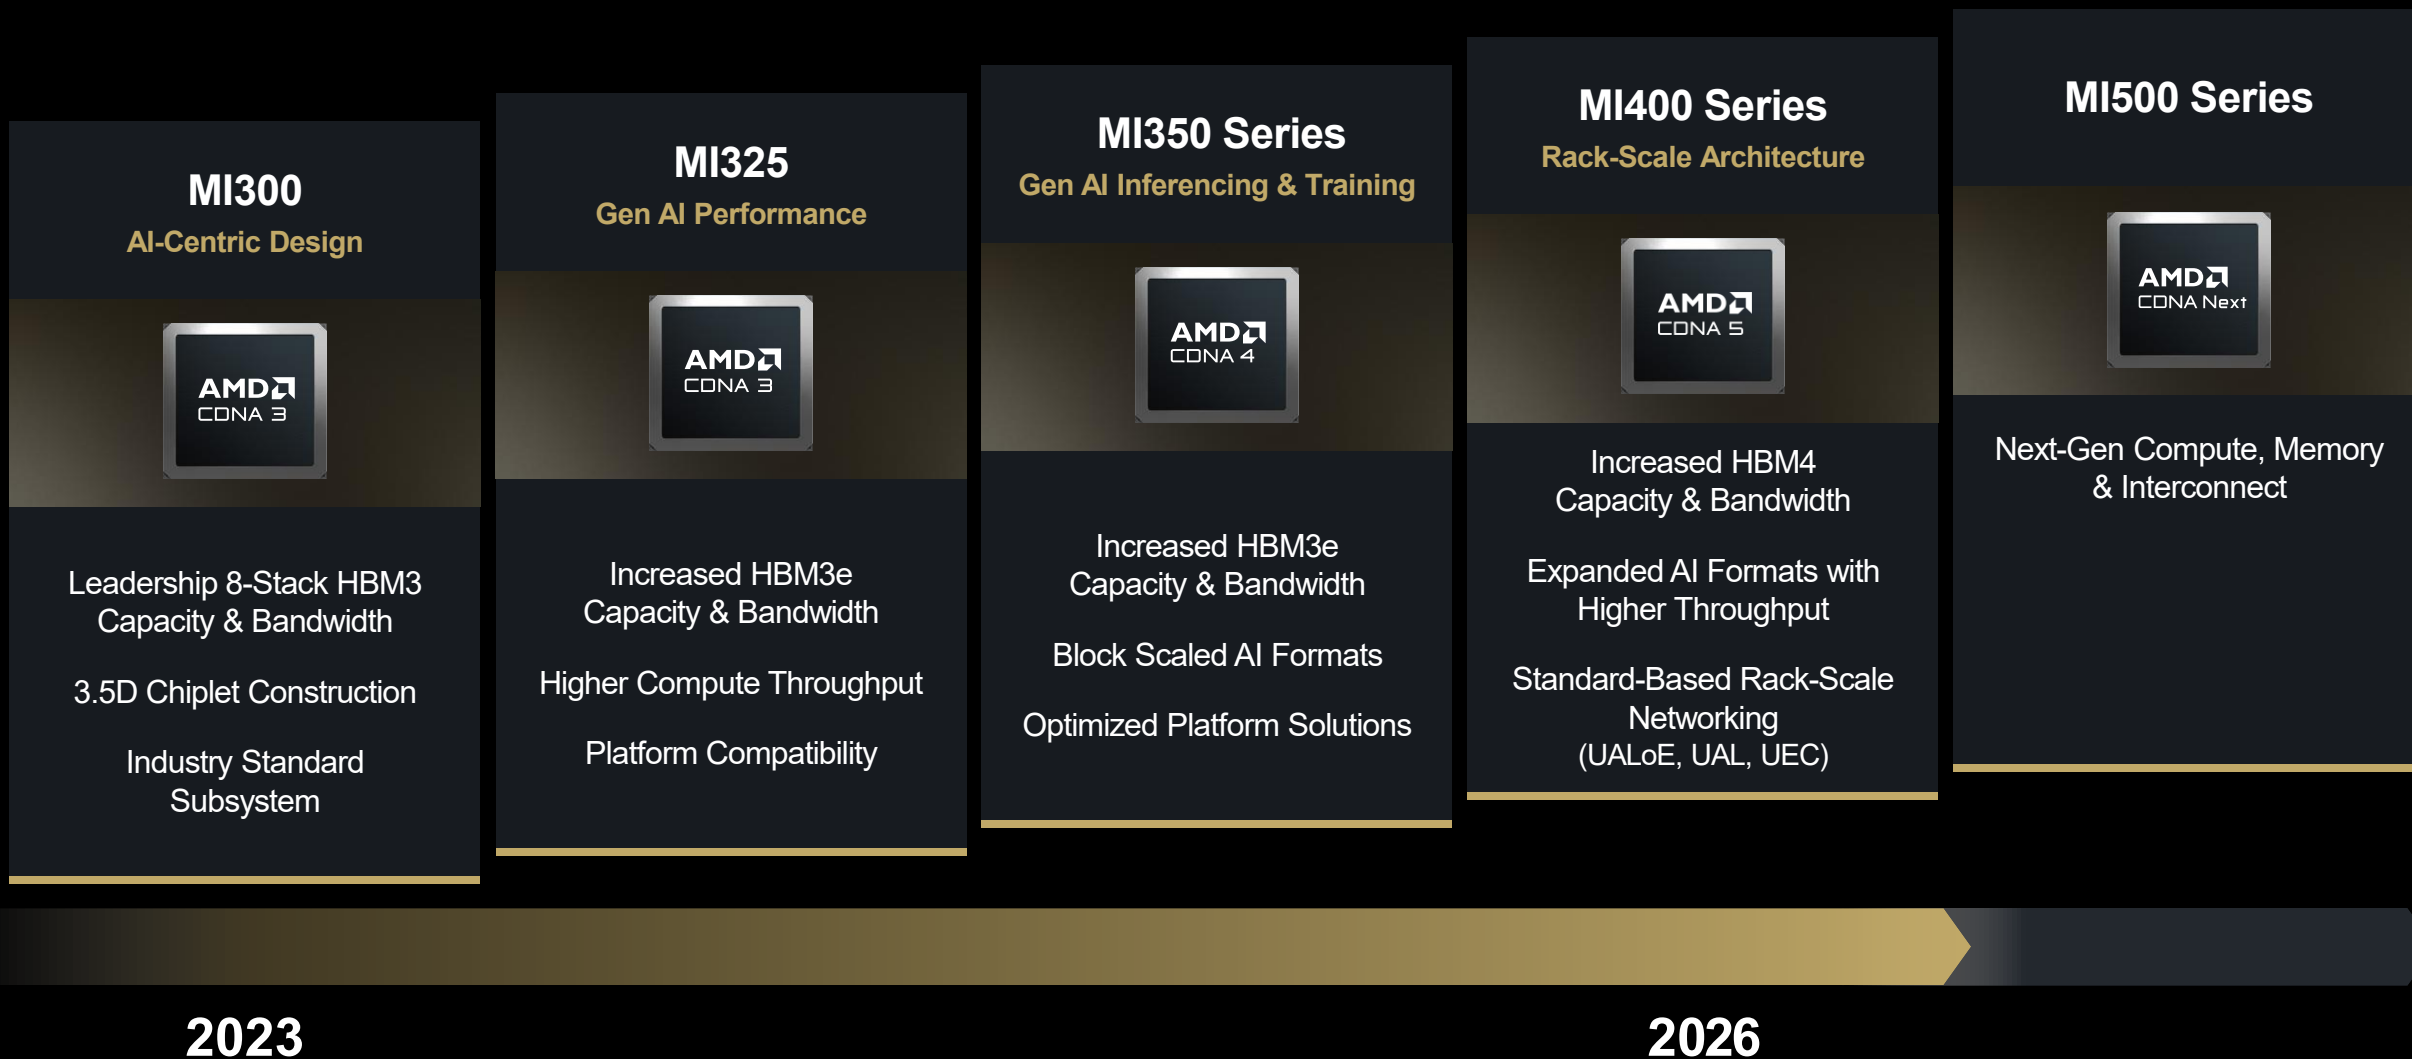


# AMD enterprise AI reference stack

Removes underlying complexity, making AI deployments simple



# AMD Instinct™ : Data Center GPU Architecture Roadmap




2023

2026

# LEARN MORE

- [AMD AI Solutions](#)
- [AMD HPC Solutions](#)
- [AMD Instinct™ GPUs](#)
- [AMD ROCm Software](#)
- [AMD ROCm™ Software Developer Hub](#)
- [AMD ROCm Blogs](#)
- [AMD Enterprise AI Suite](#)
- [AMD University Program](#)



**AMD** 