AMD

FINANCIAL ANALYST DAY 2022

together we advance_

Adaptive and Embedded Leadership

Victor Peng
President, Adaptive and Embedded Computing Group

Cautionary Statement

This presentation contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) including, but not limited to, the timing, availability, features, functionality and expected benefits of AMD's adaptive computing products; AMD's momentum and TAM; AMD's silicon roadmap; and AMD's new revenue opportunities and path forward, 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.

XILINX MOMENTUM

Data Center

5G Telecommunications

Automotive

Platform Transformation

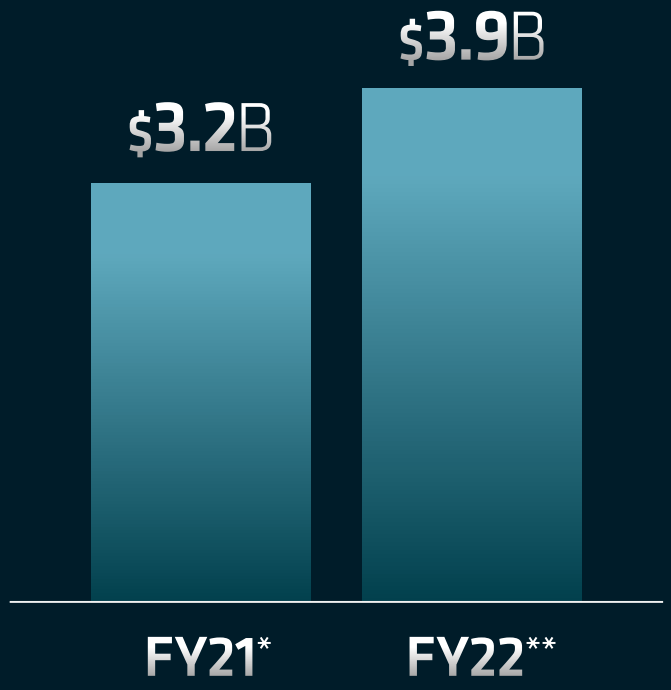

CY25 TAM

Building on Momentum as Part of AMD

*Xilinx fiscal year ended April 3, 2021
**Xilinx trailing 4 quarters ended March 26, 2022

AECG MISSION

Exceeding customer expectations with
high performance, adaptive, and intelligent solutions
for the data center, edge, and endpoints

AECG STRATEGY

- Delivering leadership adaptive compute products and technology
- Increase data center growth in networking, compute, and AI acceleration with adaptive and customizable solutions
- Drive growth in embedded markets by enabling customer differentiation with a broad product portfolio and AI technology

SIGNIFICANT TAM EXPANSION

Data Center*
$13B

Embedded
$33B

Communications
$32B

Automotive
$27B

Over 3X Xilinx CY25 TAM

*Included in Data Center TAM
Based on AMD Internal Data

ADAPTIVE COMPUTING
DATA CENTER LEADERSHIP

Compute Acceleration

Adaptive SmartNIC

Computational Storage

General Purpose

Deployed Across 10 of the Largest Hyperscalers

Alibaba Cloud · aws · IBM · Microsoft Azure · Tencent Cloud

ADAPTIVE COMPUTING
COMMUNICATIONS LEADERSHIP

Metro/Transport and Converged Access

5G Core

Central Unit

Distributed Unit

Radio Unit

Deployed at 6 of the Top 7 5G Wireless Equipment Manufacturers

cisco FUJITSU \Orchestrating a brighter world NEC NOKIA SAMSUNG

ADAPTIVE COMPUTING

AUTOMOTIVE LEADERSHIP

Autonomous Control

ADAS/AD Domain Controller

Surround View

Automated Parking ECU

External Sensors

Lidar

Radar

Forward-Looking Camera

Internal Safety and Comfort

Infotainment

Occupant Monitoring

Designed-in at the Top 10 Manufacturers

BYD · Continental · DAIMLER · MAGNA · SUBARU · TESLA · ZF

ADAPTIVE COMPUTING
EMBEDDED LEADERSHIP

Healthcare and Sciences

5M Units Deployed in Medical Equipment in Last 12 Months*

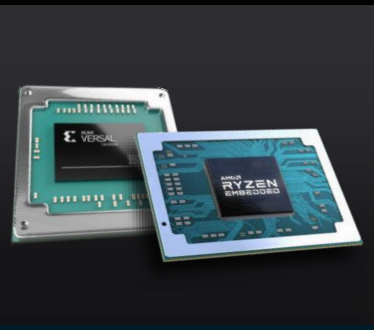

Industrial

#1 Logic IC Vendor in the Industrial Market**

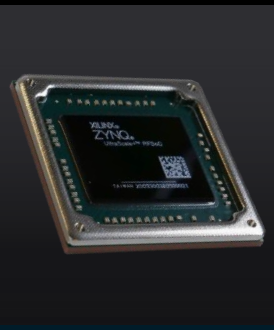

Aerospace and Defense

30+ Years of Innovation, Longevity of Supply and Reliability

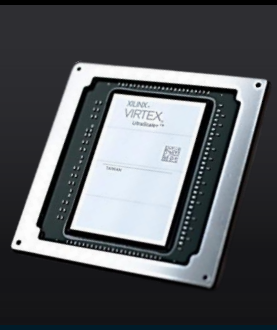

Test, Measurement and Emulation

#1 Logic IC Vendor for Test and Measurement**

Edge Server

Performance Leadership and Advanced Reliability

6000+ Unique Customers Across Diverse Embedded Markets

ADVANTECH · GE · INTUITIVE SURGICAL · OLYMPUS · SIEMENS Gamesa RENEWABLE ENERGY · Tektronix

*Based on AMD Internal Data

**According to Omdia Industrial Semiconductor Market Tracker Database 4Q21, Xilinx holds the #1 spot for Logic IC vendor on the top ten semiconductor suppliers list for the industrial market, which captures 2019-2020 revenue and market share.

SILICON ROADMAP

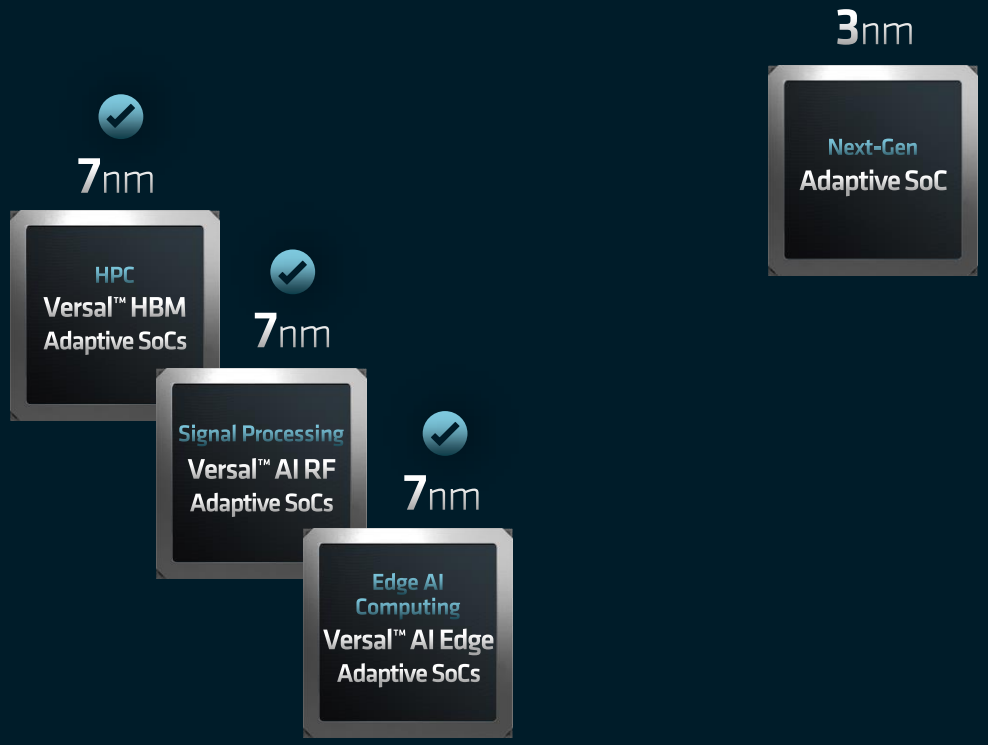

2021 → 2025

Adaptive Silicon

Data Center and Edge

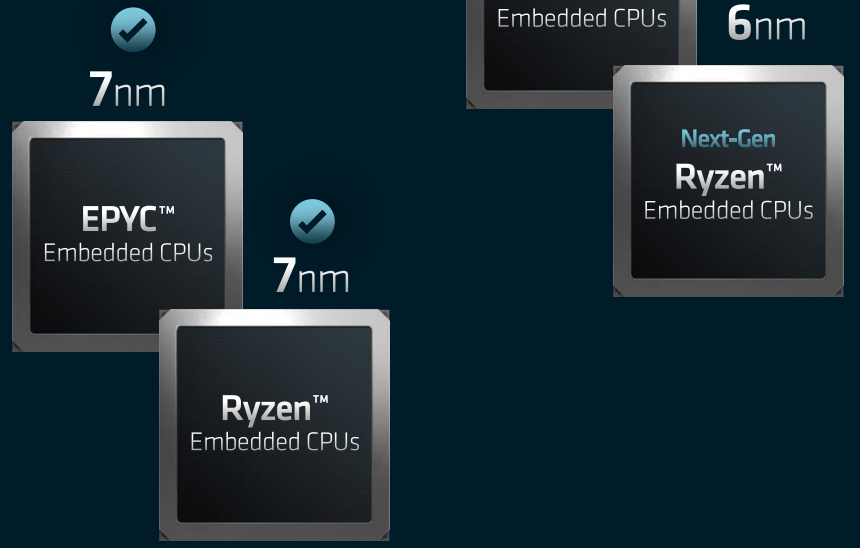

2021 → 2024

Embedded x86 Processors

Networking, Storage & Automotive

All roadmaps are subject to change.

Automotive

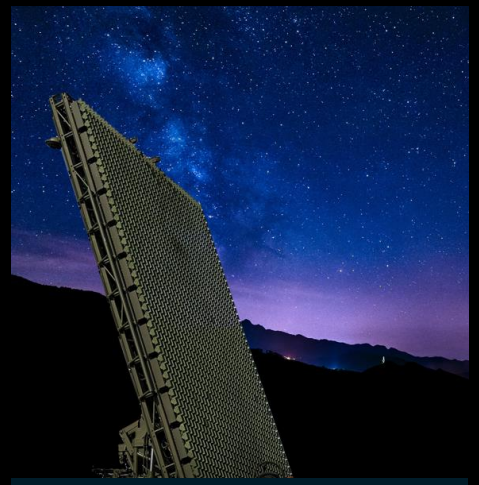

Aerospace and Defense

Data Center

Communications

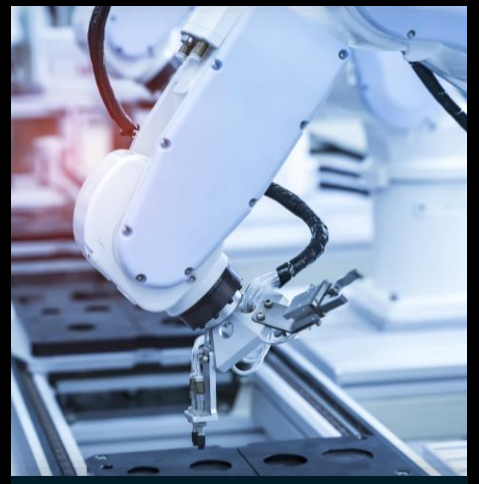

Industrial and Vision

Healthcare and Sciences

AECG GROWTH ACCELERATION

- Leadership adaptive computing platforms
- Sustained double digit growth
- Diverse high-margin markets

AMD + XILINX OPPORTUNITIES

AI Inference and Training

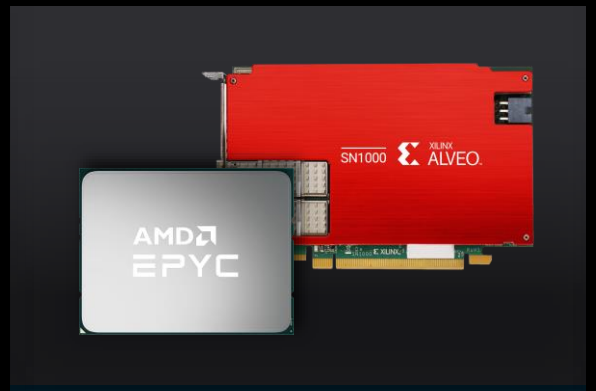

Data Center and Communications

Automotive

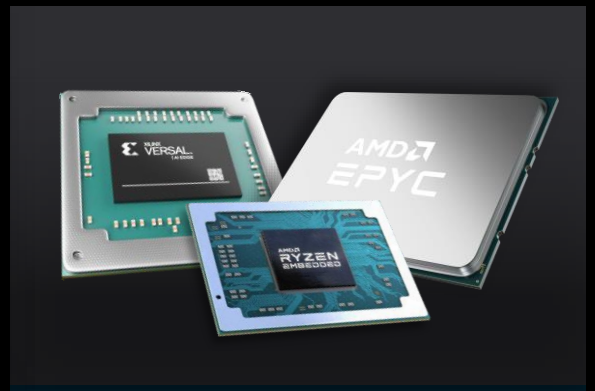

Embedded

Leverage Expanded Customer Base, Leadership Products, and AI Platforms


PERVASIVE AI

PERVASIVE AI

Commercial & Enterprise

EPYC™, Ryzen™ PRO CPUs

- Server
- PC
- Workstation

Cloud Data Center

EPYC™ CPUs , AMD Instinct™ GPUs

- AI Training
- AI Inference

Digital Home

Ryzen™, Radeon™ CPUs

- PCs and Consoles
- Metaverse

PERVASIVE AI

Commercial & Enterprise
EPYC™, Ryzen™ PRO CPUs

- Server
- PC
- Workstation
- Edge

Intelligent Factory
Zynq® Adaptive SoCs

- Machine Vision
- Predictive Maintenance
- AI Robotics

Smart Retail
Zynq® Adaptive SoCs

- Cashier-less payment
- Customer Reidentification
- Inventory Management

Cloud Data Center
EPYC™ CPUs , AMD Instinct™ GPUs, Alveo™ Accelerators, Kintex® FPGAs

- AI Training
- AI Inference
- Security
- Automatic Speech Recognition

Transportation
Zynq® Adaptive SoCs

- ADAS
- Autonomous Vehicles

Smart City
Zynq® Adaptive SoCs

- Security Endpoints
- Traffic Optimization
- Smart Grid

Digital Home
Ryzen™, Radeon™ GPUs, Zynq® Adaptive SoCs

- PCs and Consoles
- Metaverse
- Smart Home

Healthcare & Life Sciences
Zynq® Adaptive SoCs

- Imaging and Diagnosis
- Surgical Robotics

Communications
Versal® Adaptive SoCs

- 5G Wireless
- Networking Security

PERVASIVE AI

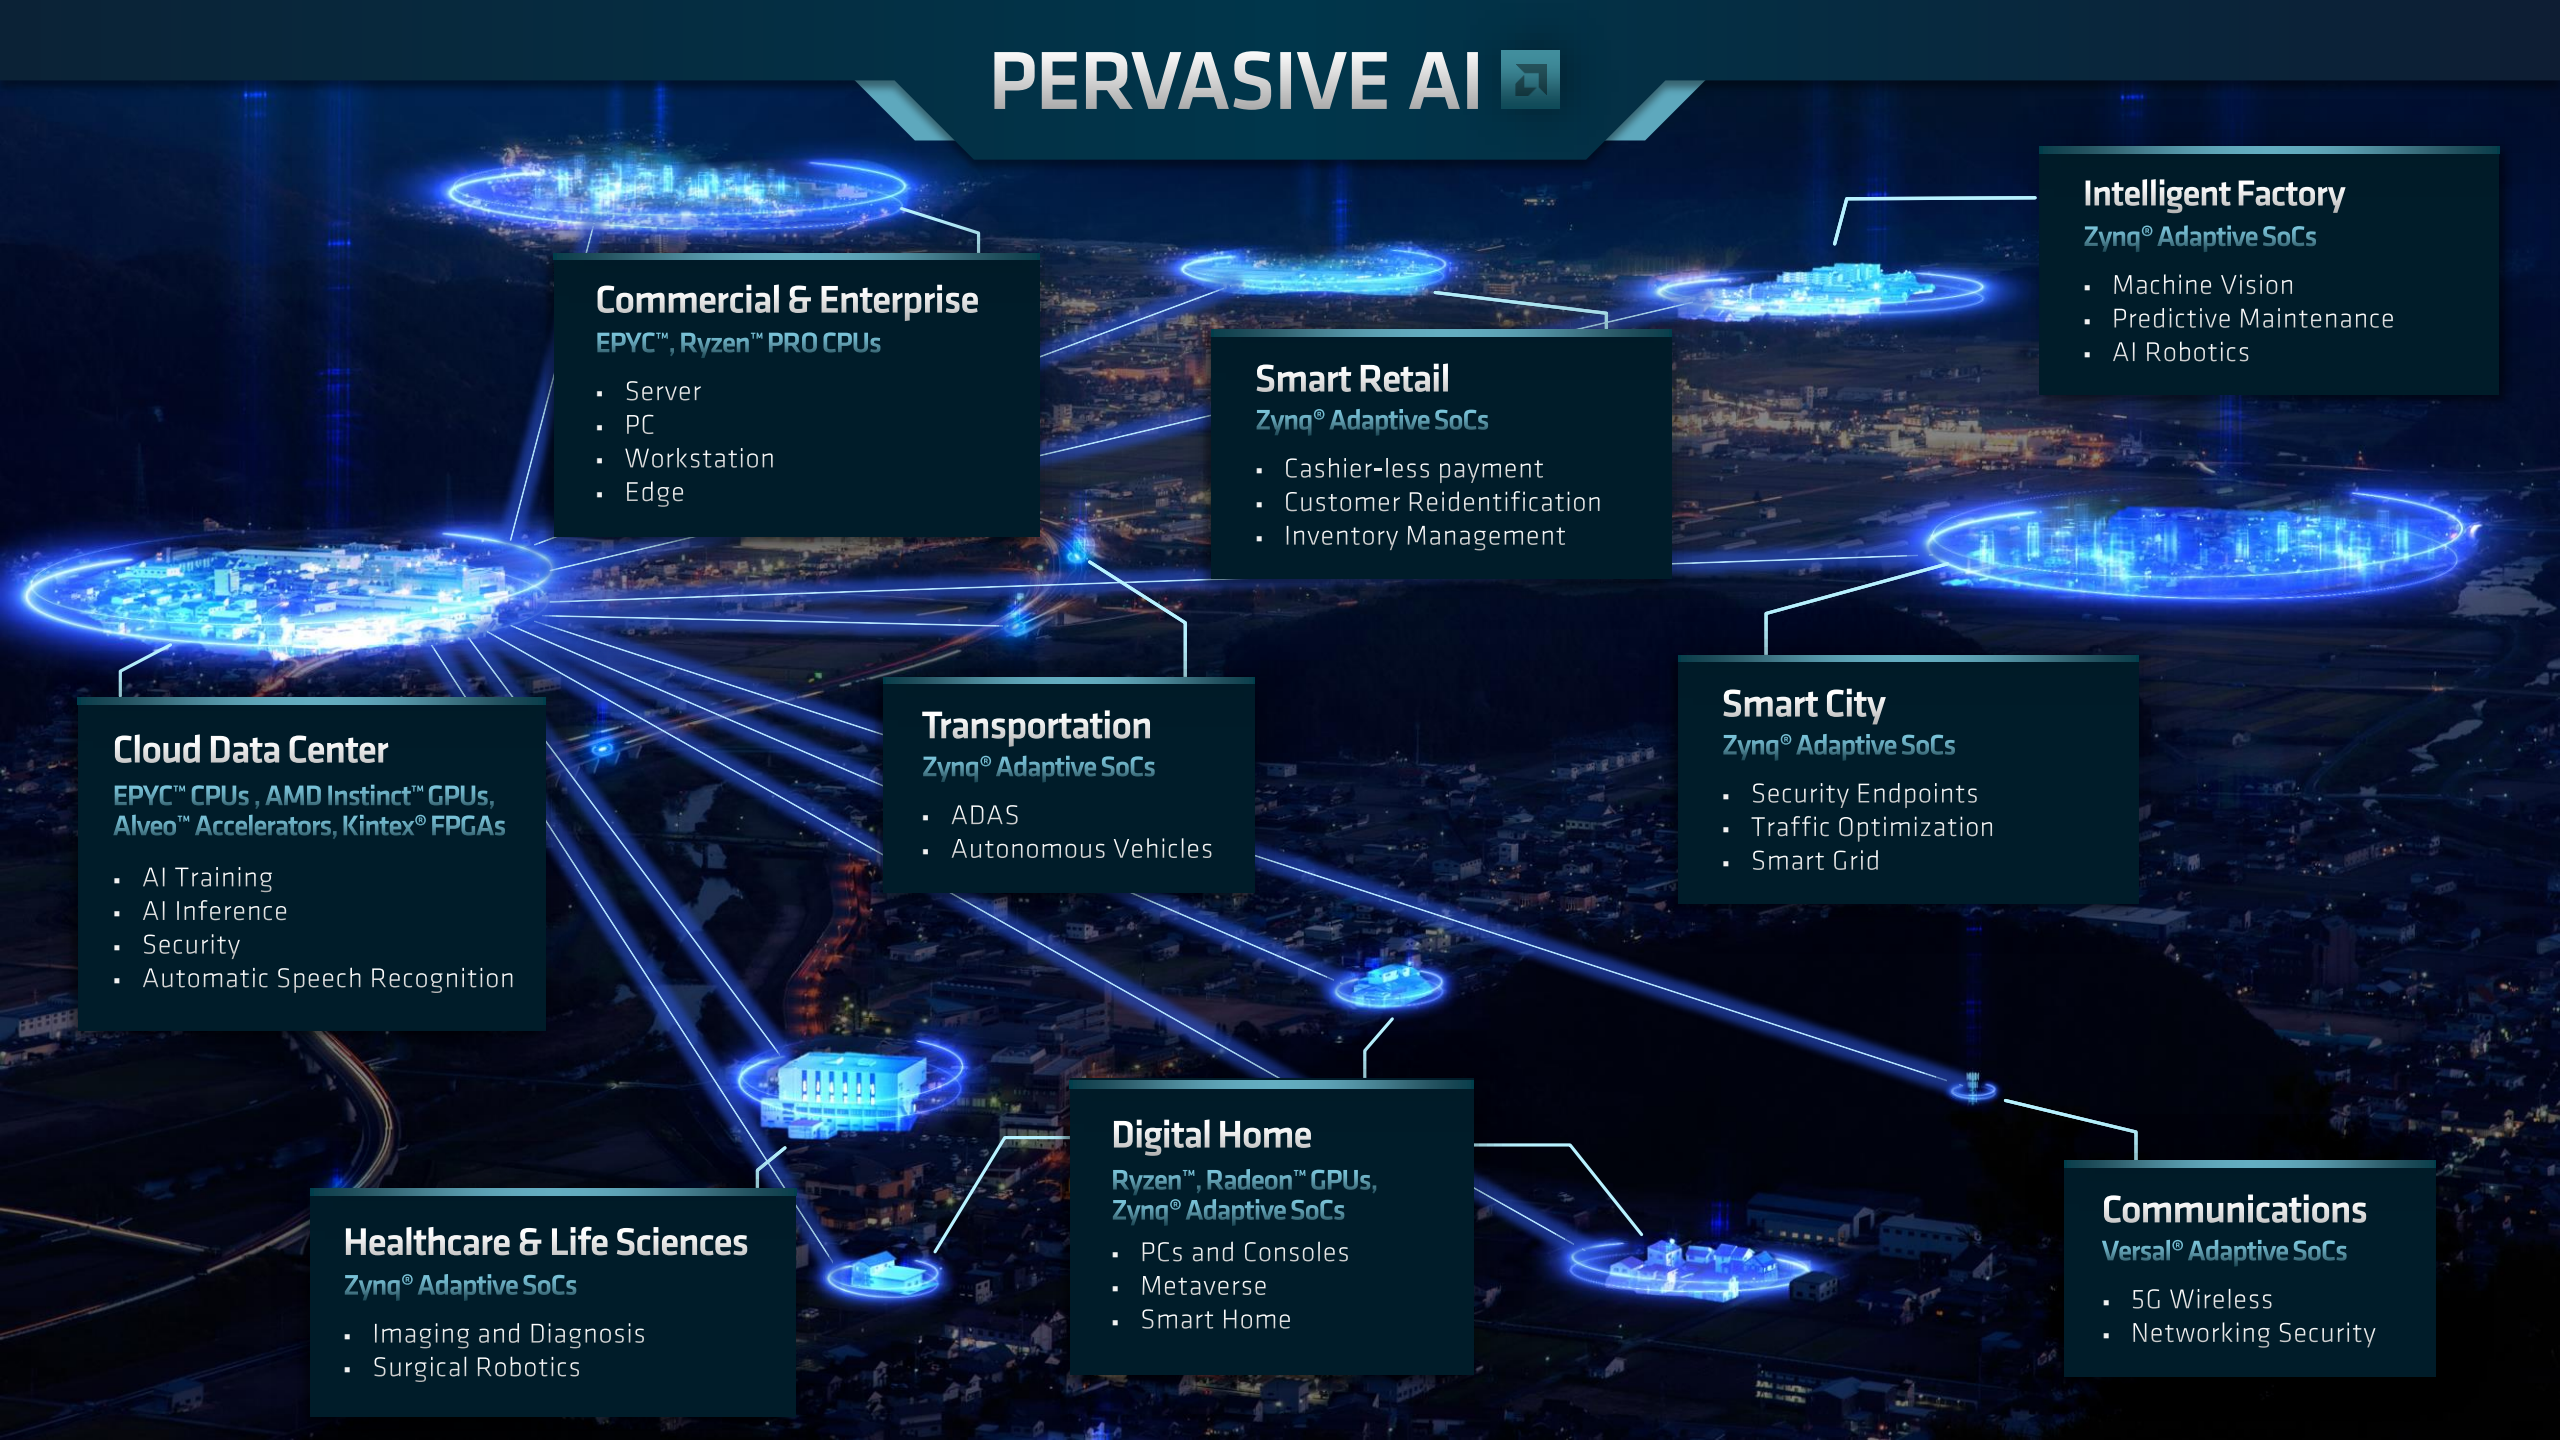

Commercial & Enterprise
EPYC™, Ryzen™ PRO CPUs

- Server
- PC
- Workstation
- Edge

Intelligent Factory
Zynq® Adaptive SoCs

- Machine Vision
- Predictive Maintenance
- AI Robotics

Smart Retail
Zynq® Adaptive SoCs

- Cashier-less payment
- Customer Reidentification
- Inventory Management

Cloud Data Center
EPYC™ CPUs , AMD Instinct™ GPUs, Alveo™ Accelerators, Kintex® FPGAs

- AI Training
- AI Inference
- Security
- Automatic Speech Recognition

Transportation
Zynq® Adaptive SoCs

- ADAS
- Autonomous Vehicles

Smart City
Zynq® Adaptive SoCs

- Security Endpoints
- Traffic Optimization
- Smart Grid

Digital Home
Ryzen™, Radeon™ GPUs, Zynq® Adaptive SoCs

- PCs and Consoles
- Metaverse
- Smart Home

Healthcare & Life Sciences
Zynq® Adaptive SoCs

- Imaging and Diagnosis
- Surgical Robotics

Communications
Versal® Adaptive SoCs

- 5G Wireless
- Networking Security

PERVASIVE AI STRATEGY

- Leadership AI IP with AMD CDNA™ and scalable AIE architecture
- Broaden AMD AI product portfolio across cloud, edge, and endpoint applications
- Unified AI stack to empower developers across the AMD portfolio

AMD XDNA: ADAPTIVE ARCHITECTURE IP

- Dataflow architecture optimal for AI and signal processing applications
- Highly-scalable array of engines with local memory and data movers
- Leverages deep expertise of compiling algorithms to FPGAs and adaptive SoCs

AI Engine
(AIE)

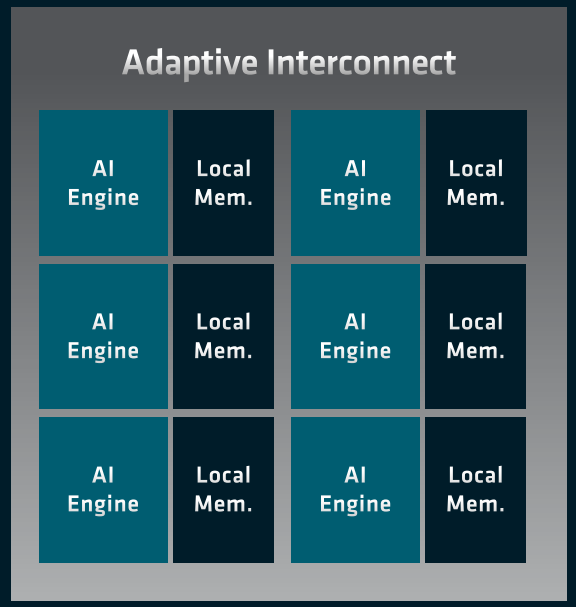

High-Performance and Energy Efficiency for AI and Signal Processing

FPGA Fabric
(FPGA)

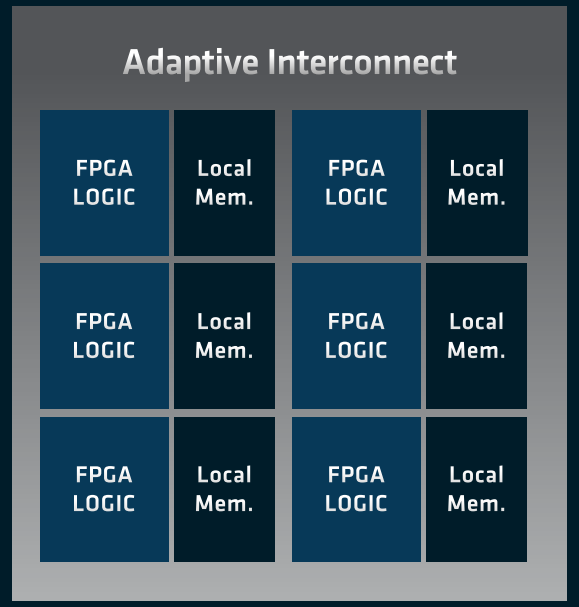

Leading FPGA for broad set of applications and AI

Source: The McClean Report, May 2022

AMD AI ENGINE: ADAPTIVE DATAFLOW PROCESSOR

Deep Neural Network (DNN)

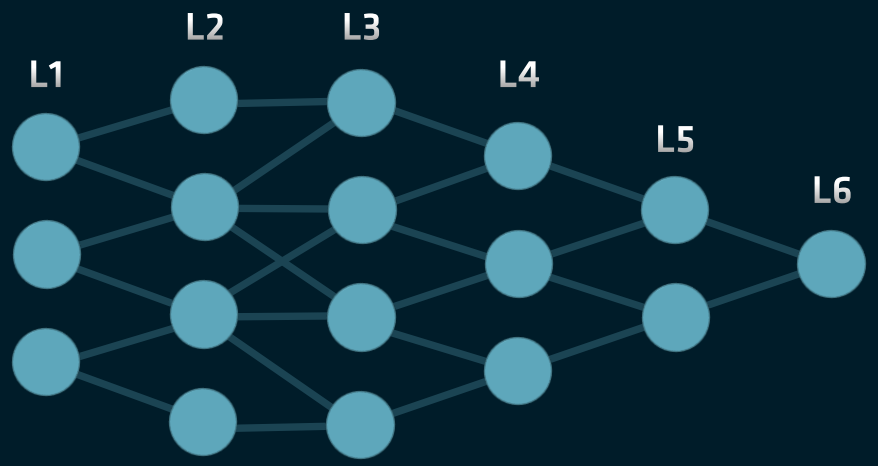

Data “flows” from layer to layer, connections between layers are often “sparse”

DNN Runs Optimally on AIE

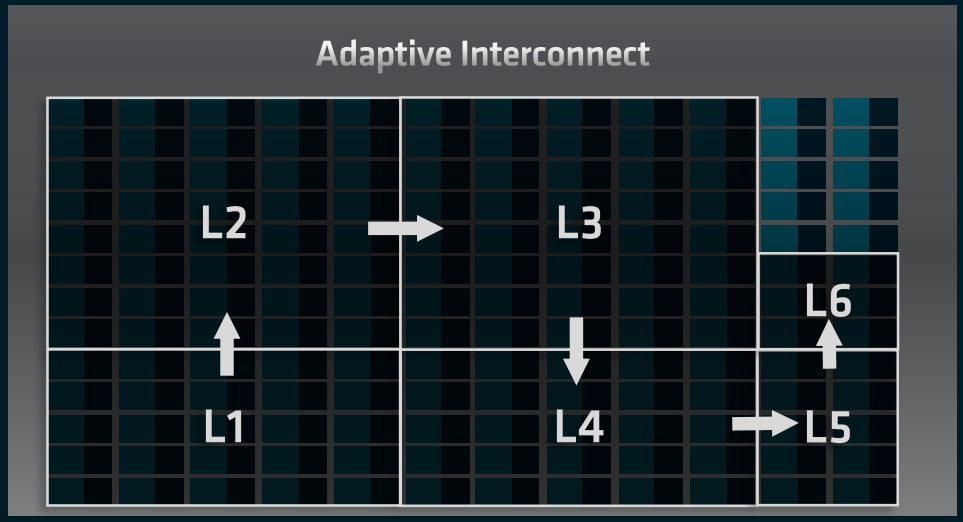

Dataflow architecture, sparsity, efficient datatypes deliver high performance and low power

High-Performance, Energy Efficient, and Customizable for AI Workloads

AI APPLICATION COVERAGE

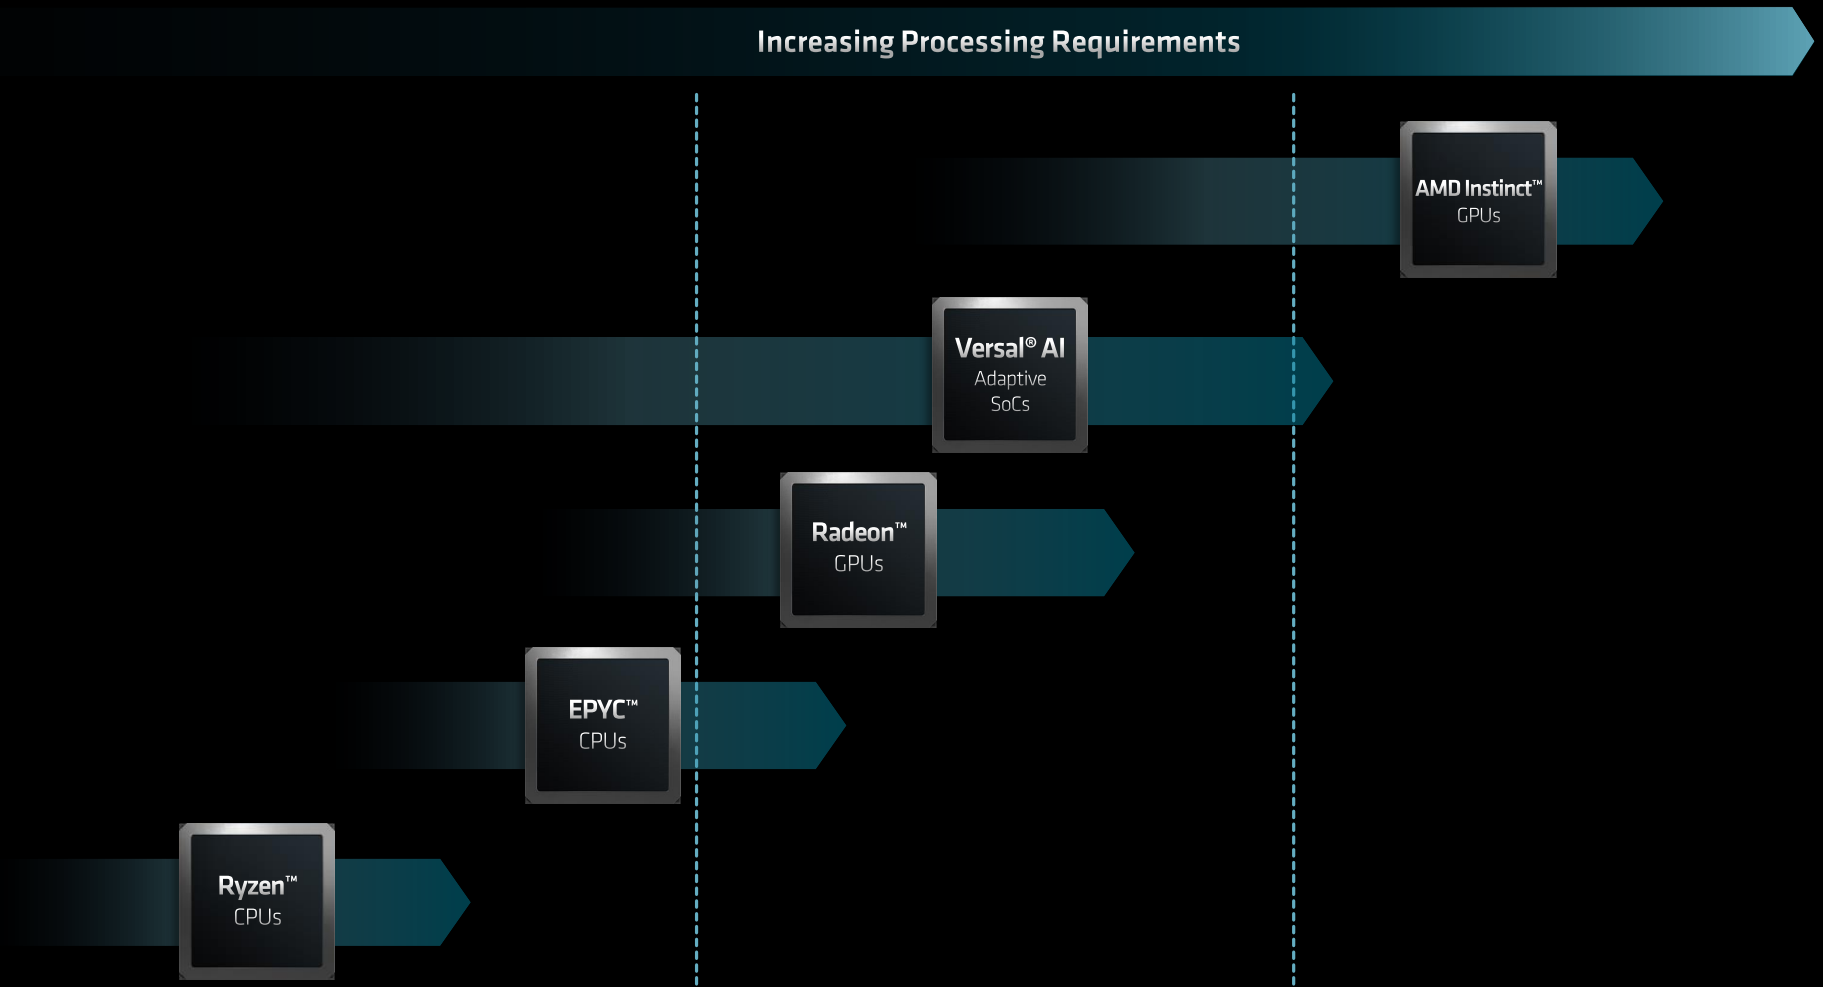

Training	**Small-medium models (EPYC™ CPUs only)**	**Medium-large models**	**Large-very large models**
Inference	**Mid-size models, mid-range performance and memory BW, high memory capacity** (e.g., face recognition, recommendation systems)	**Larger models, high performance and more memory BW** (e.g., video analytics)	**Very large models, efficient data types, and high memory BW** (e.g., NLP, large recommendation systems)

AI APPLICATION COVERAGE

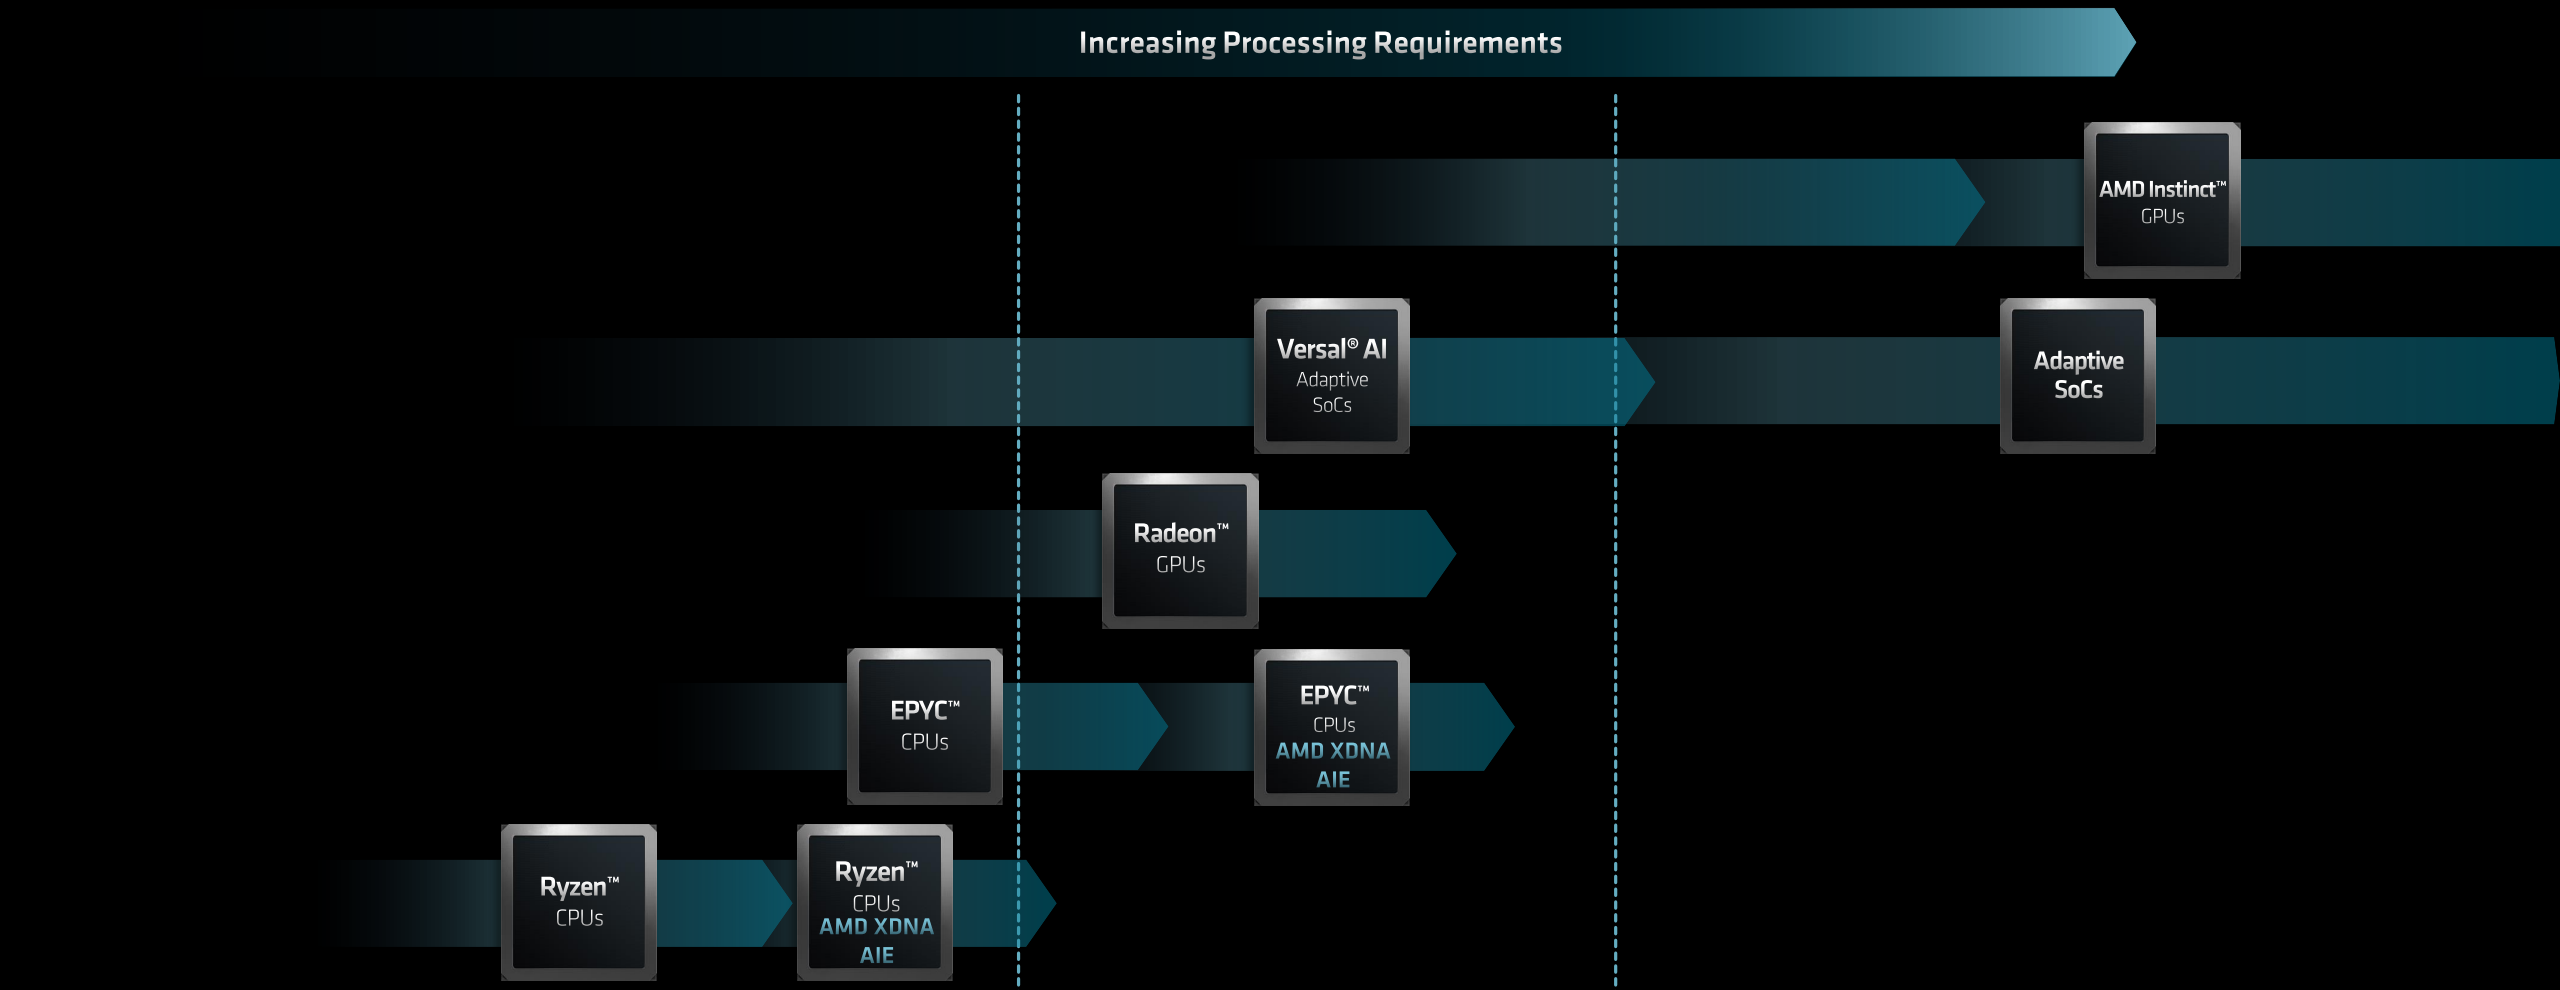

Training	**Small-medium models (EPYC™ CPUs only)**	**Medium-large models**	**Large-very large models**
Inference	**Mid-size models, mid-range performance and memory BW, high memory capacity** (e.g., face recognition, recommendation systems)	**Larger models, high performance and more memory BW** (e.g., video analytics)	**Very large models, efficient data types, and high memory BW** (e.g., NLP, large recommendation systems)

AMD AI SOFTWARE TODAY

CPU Stack

- Optimized Inference Models
- WinML** ONNX Runtime | Pytorch | TensorFlow
- ML Graph Compiler
- ZenDNN, AOCL Optimized Library
- ZEN Studio (AOCC)
- Windows Runtime | Linux Runtime

Ryzen™ CPUs | EPYC™ CPUs

ROCm™ Platform

- Optimized Inference Models
- ONNX Runtime | Pytorch | TensorFlow
- MIGraphX: AI Development Tools
- MIGraphX: ML Graph Compiler
- MIOpen, ROCBLAS
- ROCm™ HIP Compiler and tools
- Runtime

AMD Instinct™ GPUs | Radeon™ GPUs

Vitis™ AI Platform

- Optimized Inference Models
- ONNX Runtime | Pytorch | TensorFlow
- Vitis AI Development & Deployment Tools
- Vitis AI ML Graph Compiler
- Vitis ML Libraries
- Vitis™ SW Platform (AIE Compiler, HLS and tools)
- Runtime
- CNN Overlay | Transformer Overlay | ...

Versal® Adaptive SoCs | Zynq® Adaptive SoCs

*Ryzen uses RadeonML and WinML
**WinML for Ryzen Only

AI Development Tools: Quantizer, Sparsity/Pruning Tool
AI Deployment Tools: Inference Server

AMD UNIFIED AI STACK 1.0

Unified Inference Frontend (UIF) for AI Developers

Optimized Inference Models

WinML* | ONNX Runtime | Pytorch | TensorFlow

Unified AI Development & Deployment Tools (Quantization, Pruning, Inference Server)

ML Graph Compiler	MIGraphX: ML Graph Compiler	Vitis AI ML Graph Compiler
ZenDNN, AOCL Optimized Library	MIOpen, ROCBLAS	Vitis ML Libraries
ZEN Studio (AOCC)	ROCm™ HIP Compiler and tools	Vitis™ SW Platform (AIE Compiler, HLS and tools)
Windows Runtime / Linux Runtime	Runtime	Runtime
		CNN Overlay / Transformer Overlay / ...

Ryzen™ CPUs | EPYC™ CPUs | AMD Instinct™ GPUs | Radeon™ GPUs | Ryzen™ CPUs AMD XDNA AIE | Versal® Adaptive SoCs | Zynq® Adaptive SoCs

*WinML for Ryzen Only

All roadmaps are subject to change.

AMD UNIFIED AI STACK 2.0

Seamless Workload Partitioning with UIF, Graph Compiler and Library APIs

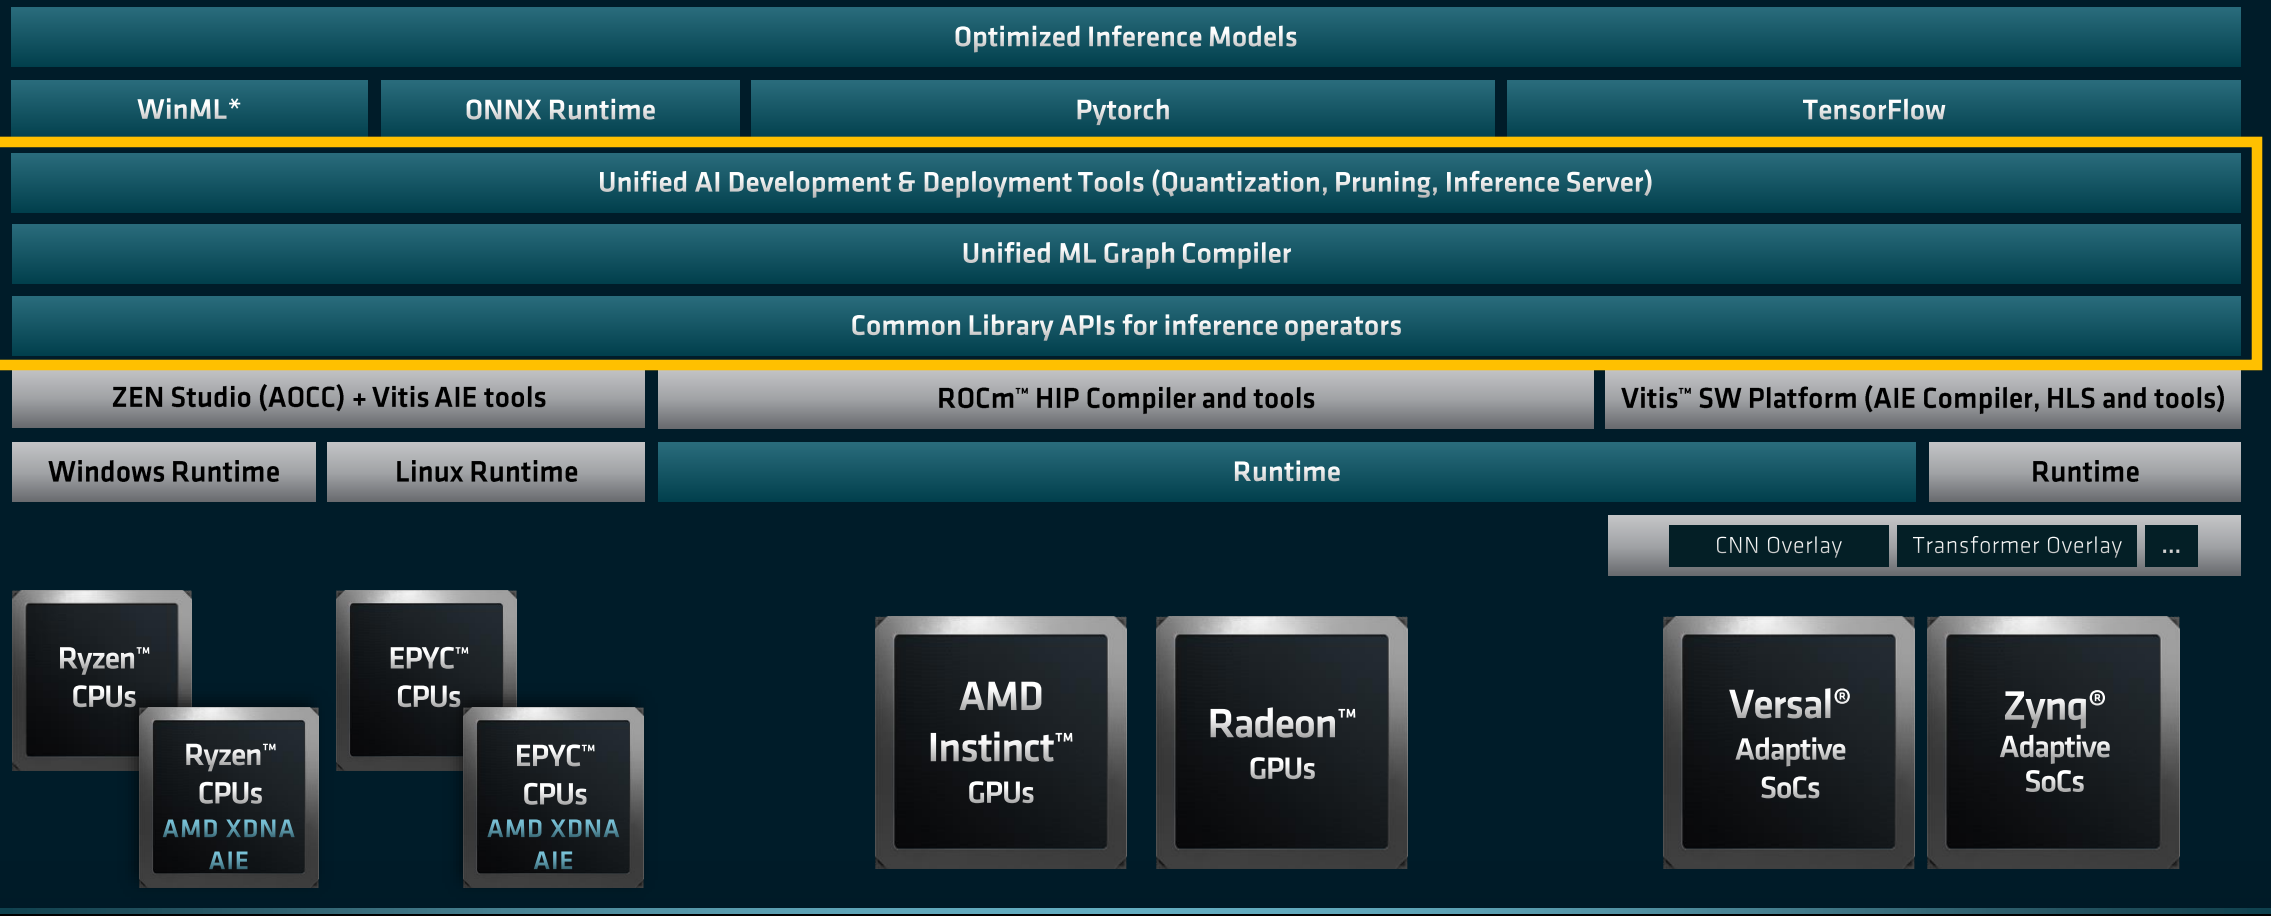

*WinML for Ryzen Only

All roadmaps are subject to change.

AMD + XILINX OPPORTUNITIES

AI Inference and Training

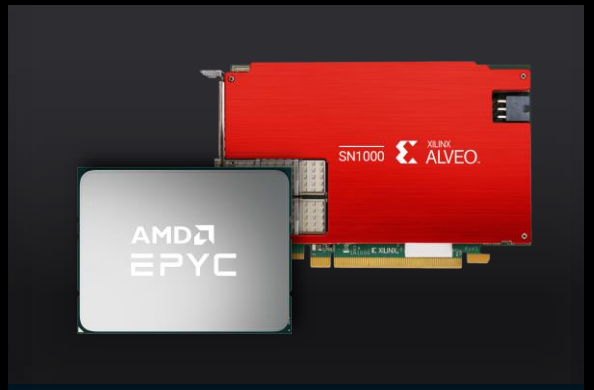

Data Center and Communications

Automotive

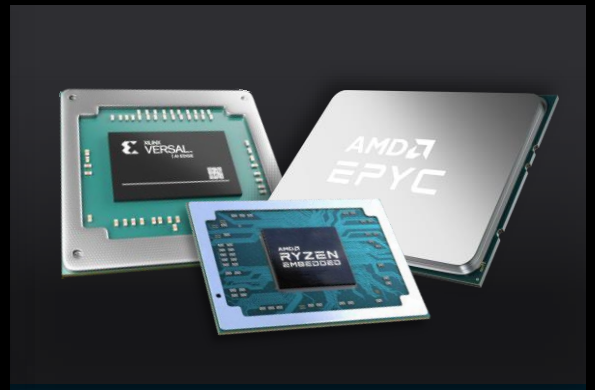

Embedded

>$10B IDENTIFIED REVENUE OPPORTUNITIES

Based on AMD Internal Data

ACCELERATING AMD GROWTH

TAM Expansion

Large and diversified markets

Leadership Products

Broad portfolio of platforms

Pervasive AI

Drive cloud, edge and endpoints to scale