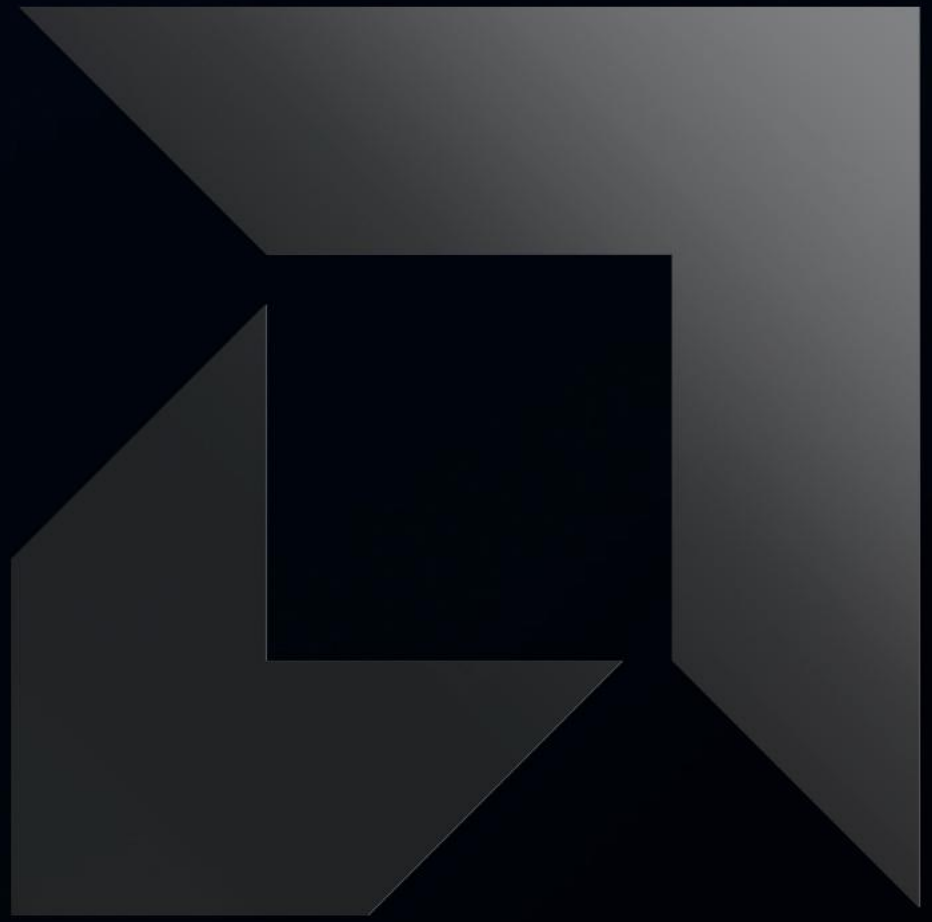

Silicon Photonics at Scale:

Bridging Advanced Packaging and Optical Connectivity for the AI Era

Robert Fox
October 2nd, 2026

AMD
together we advance_

Public

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 and technology as well as technology trends, innovation and 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.

Public

The Challenge: AI is Driving Massive Compute Demand

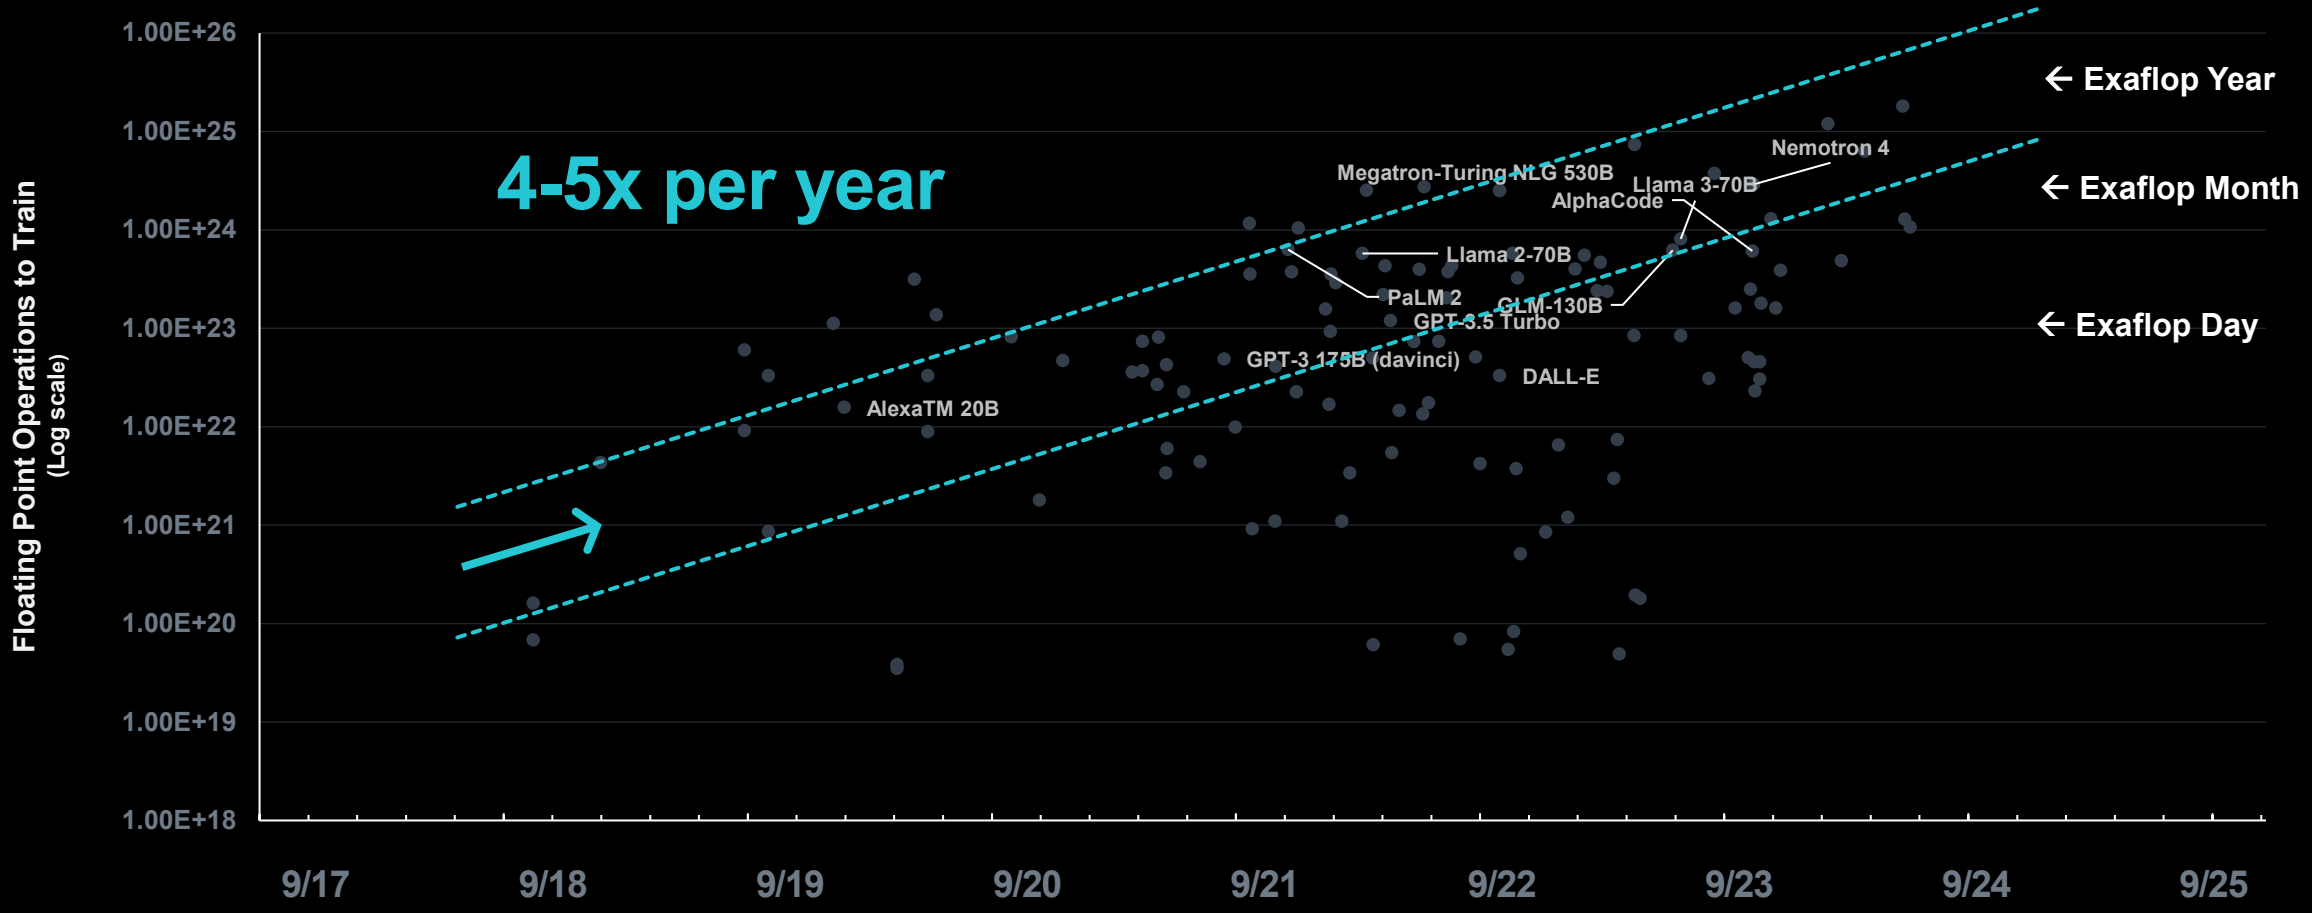

Jaime Sevilla and Edu Roldán, "Training Compute of Frontier AI Models Grows by 4-5x per Year," Epoch AI, May 28, 2024. [Online].
Available: https://epochai.org/blog/training-compute-of-frontier-ai-models-grows-by-4-5x-per-year
Exaflop labeling inspired by discussions with Argonne National Laboratory.

AMD

Public

AI Model Innovation and Adoption are Accelerating Compute Demand

Training Flops Increasing 4-5x Every Year

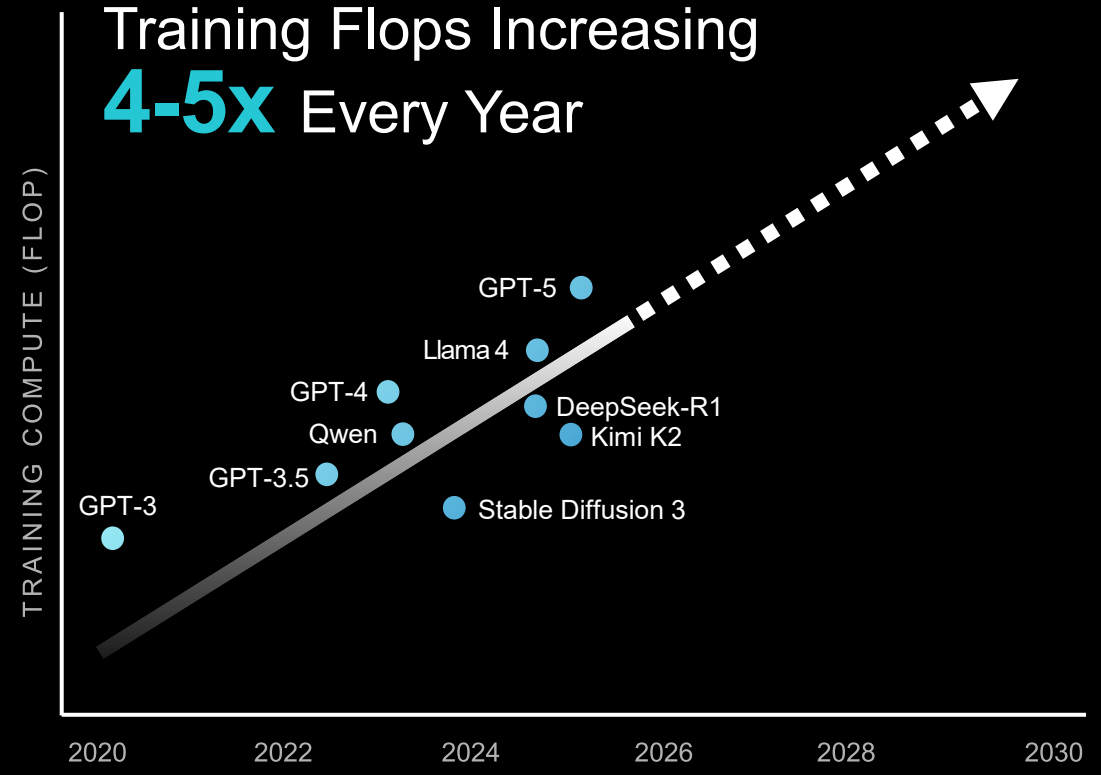

Source: Based on AMD internal analysis and data as of Dec. 2025.

Inference Tokens Increased 100x in the Last 2 Years

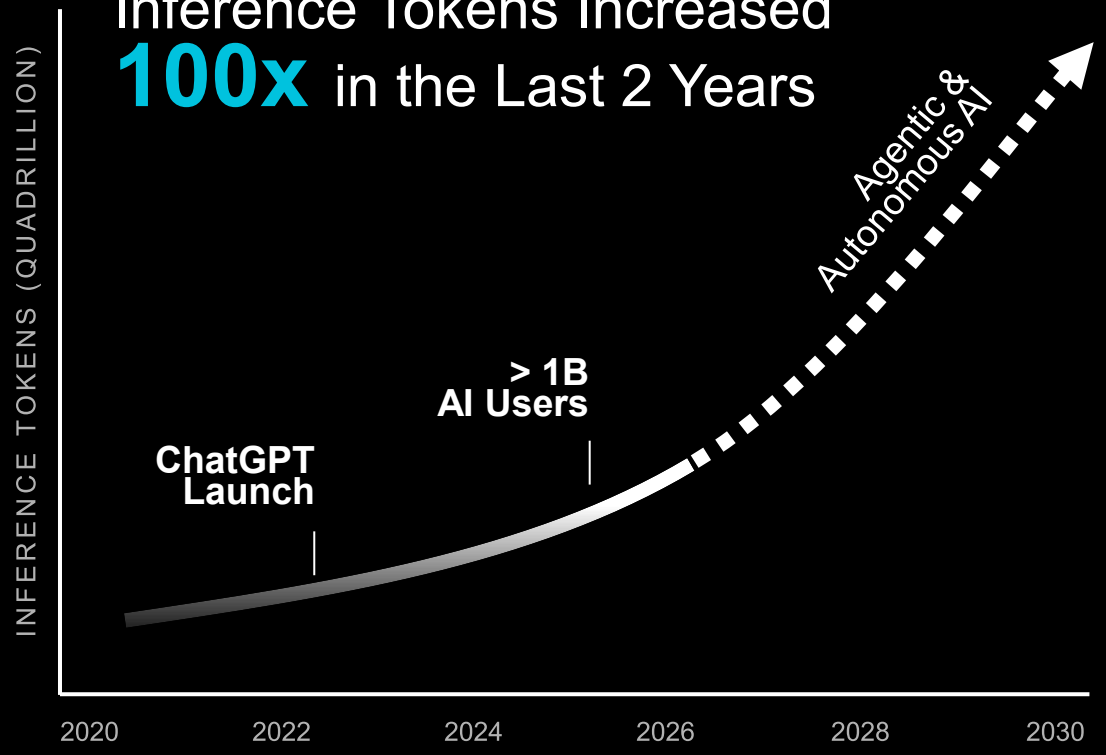

Robi Rahman and David Owen (2024), "The training compute of notable AI models has been doubling roughly every six months". Published online at epoch.ai. Retrieved from: 'https://epoch.ai/data-insights/compute-trend-post-2010'.

[AMD-CES2026: https://www.amd.com/en/blogs/2026/from-ces-2026-to-yottaflops-why-the-amd-keynote.html]

AMD
together we advance_

Public

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

4,600
“Zen 6” CPU Cores

18,000
GPU Compute Units

260 TB/s
Scale Up Bandwidth



OPEN Compute Project®

ULTRA ACCELERATOR LINK™

Ultra Ethernet Consortium

[AMD-CES2026: https://www.amd.com/en/blogs/2026/from-ces-2026-to-yottaflops-why-the-amd-keynote.html]

Public

Energy-to-Train Frontier AI Models

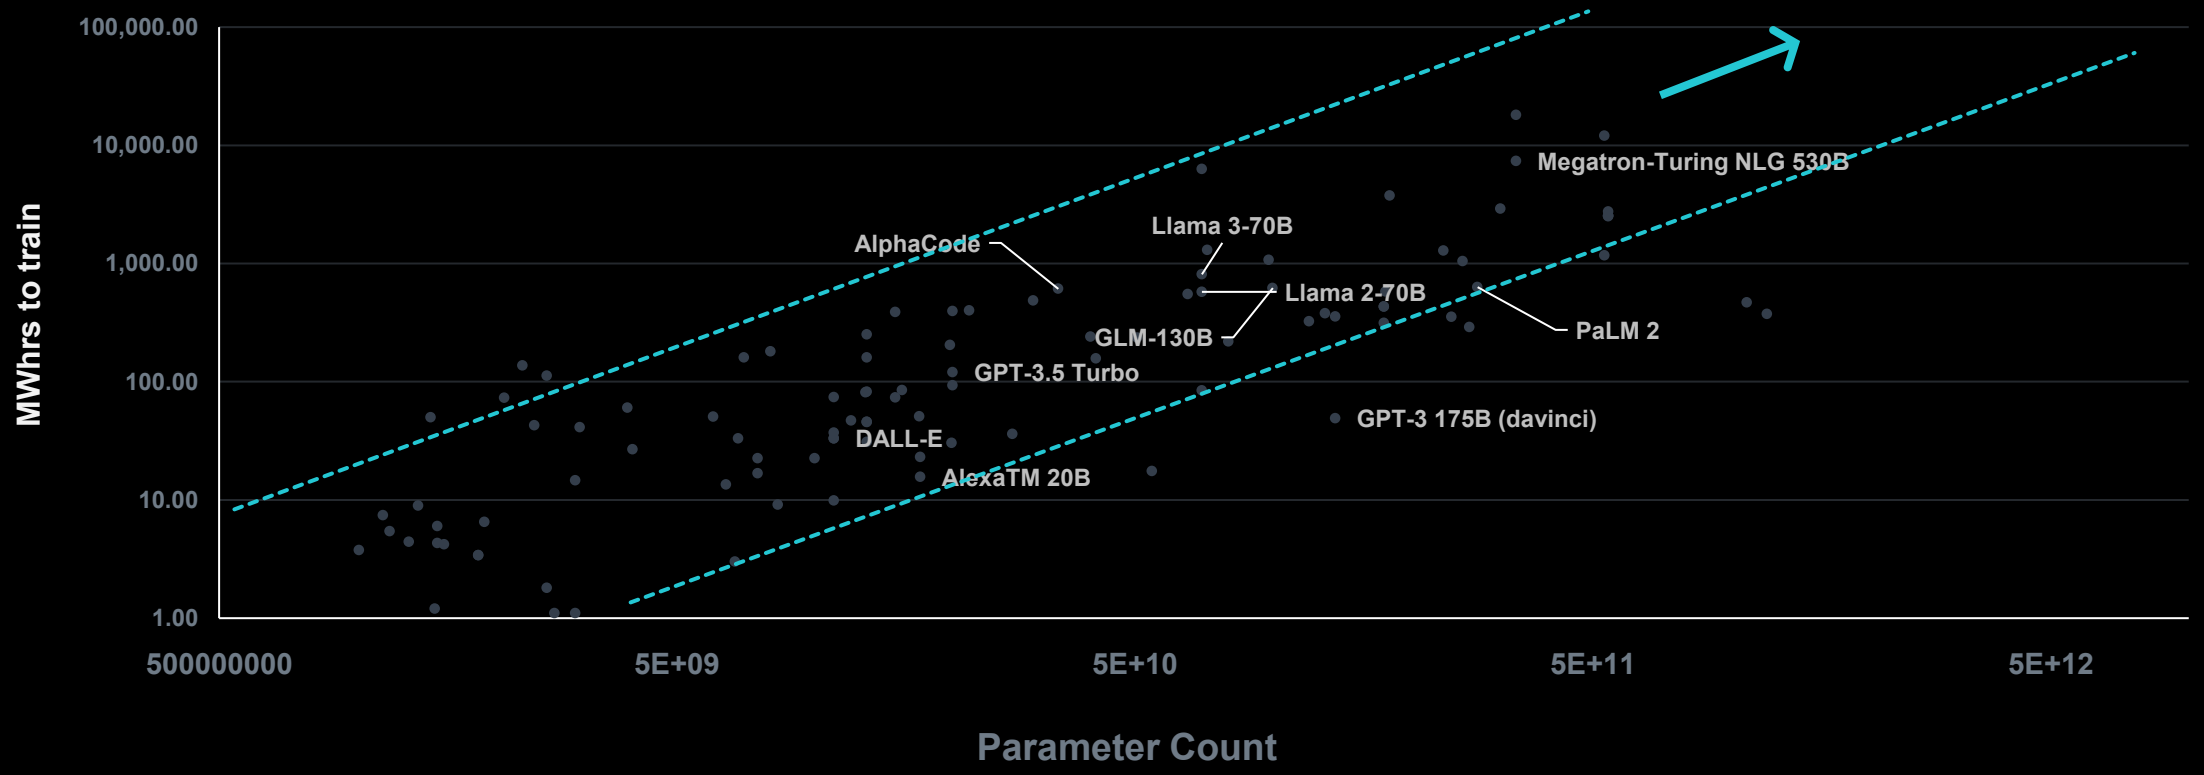

Exponential growth in model sizes driving massive increase in energy required for training

Upper bound on training requirements is yet to be determined

Based on published parameter counts of leading training models; and AMD internal calculations approximating energy to train

AMD

Public

The Need for Disaggregation

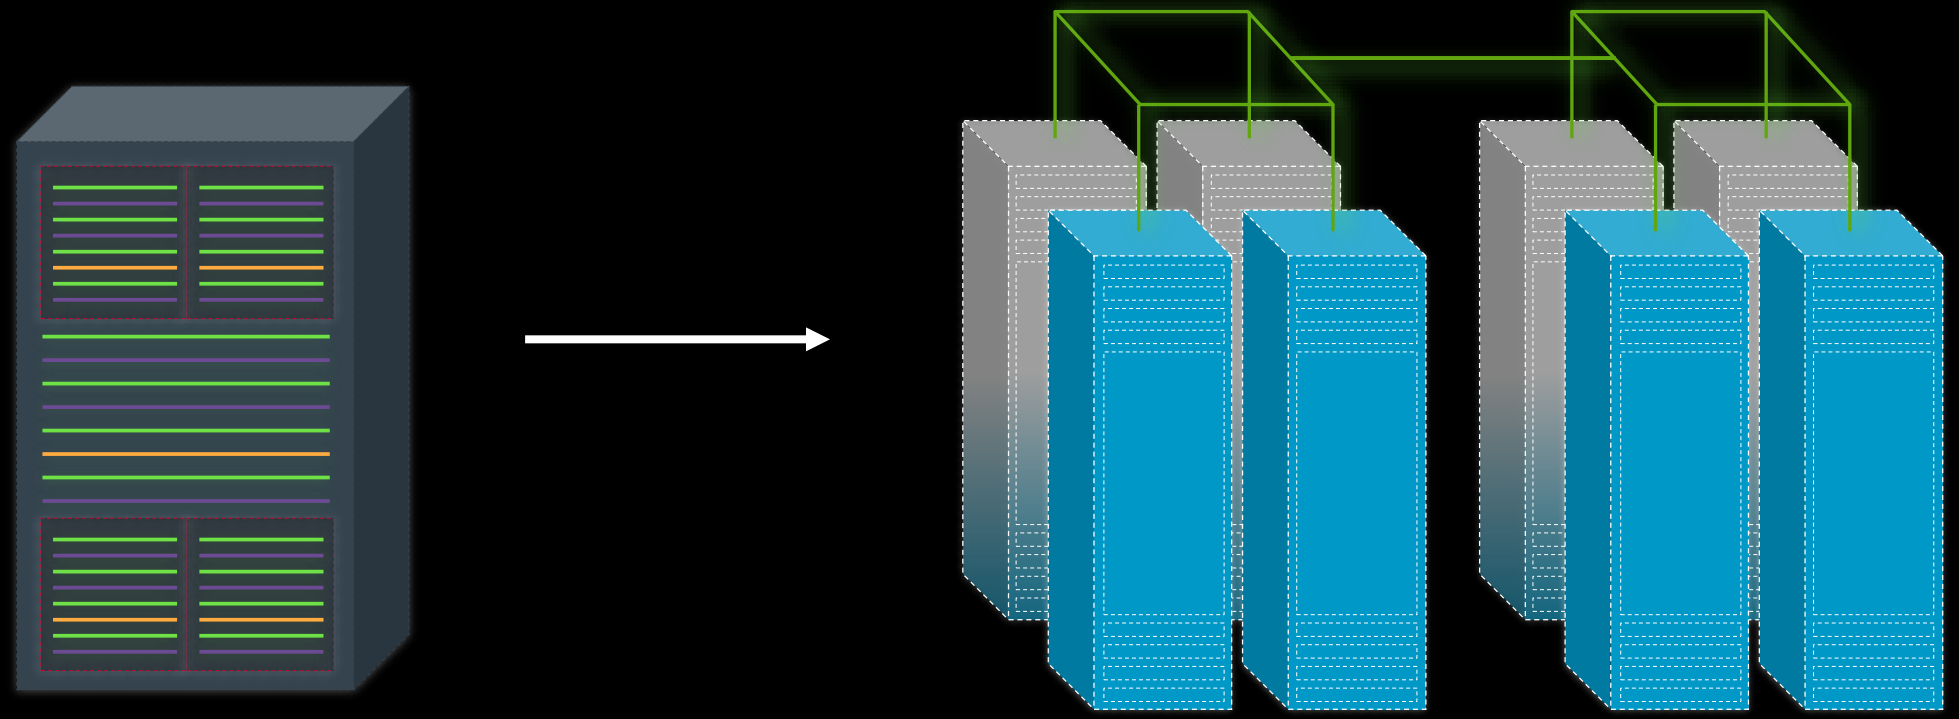

Electrical/Copper
Limits on single-node compute

Optical Scale-up
Enables multiple node/in-row compute

AMD

Public

Scaling to Meet the AI Challenge Requires Innovation on all Fronts

Unprecedented pace of compute growth driven by AI

Reducing energy of compute, memory, and network is key to meeting demand

- Tight integration of large 3D chiplet-based accelerators with SW-driven architecture
- Tightly coupled high bandwidth memory with low energy/bit
- Optical interconnects to support network bandwidth and reach

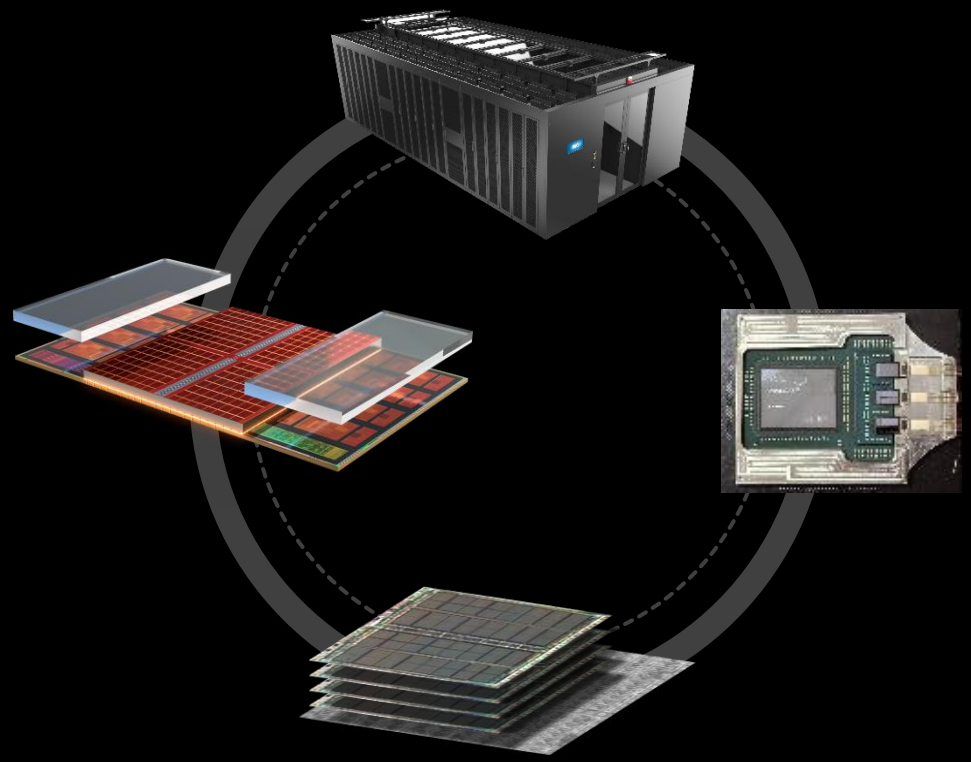

Holistic architecture and multi-disciplinary innovations are the key to growth

AMD

Public

Chiplet and Packaging Leadership

Multichip

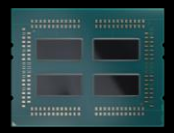

MI100
2.5D + HBM2

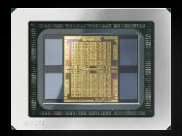

MI250
2.5D + HBM2e

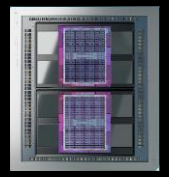

MI350
3.5D/COWOS-S + HBM3e

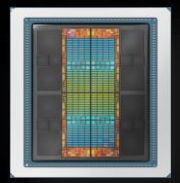

MI400
3.5D/COWOS-L + HBM4



30x
Compute & Memory

3-6x
Beyond Moore's Law

2017 — **2026**

AMD

Public

From DAC to CPO



Extra Reach with Lower Power

Direct Attach Copper "DAC"

Active Copper Cable "ACC"

Active Electrical Cable "AEC"

Active Optical Cable "AOC"

On Board Optics "OBO"

Co-Packaged Optics "CPO"

Co-Packaged Optics Direct Drive "CPO-DD"

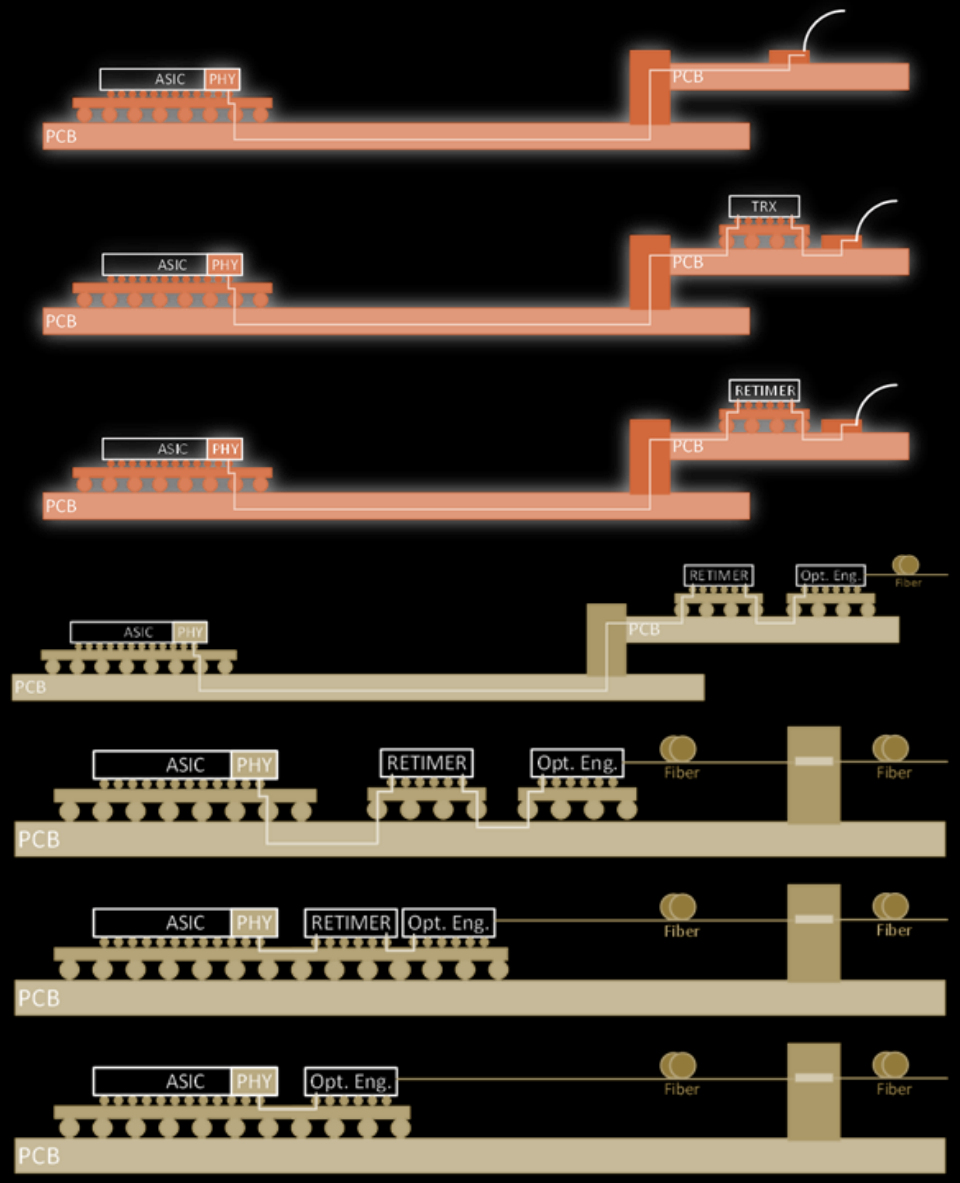

[S. Ardalan, BCICTS 2025]

AMD
together we advance_

Public

Example of Progress in Integrated Photonics

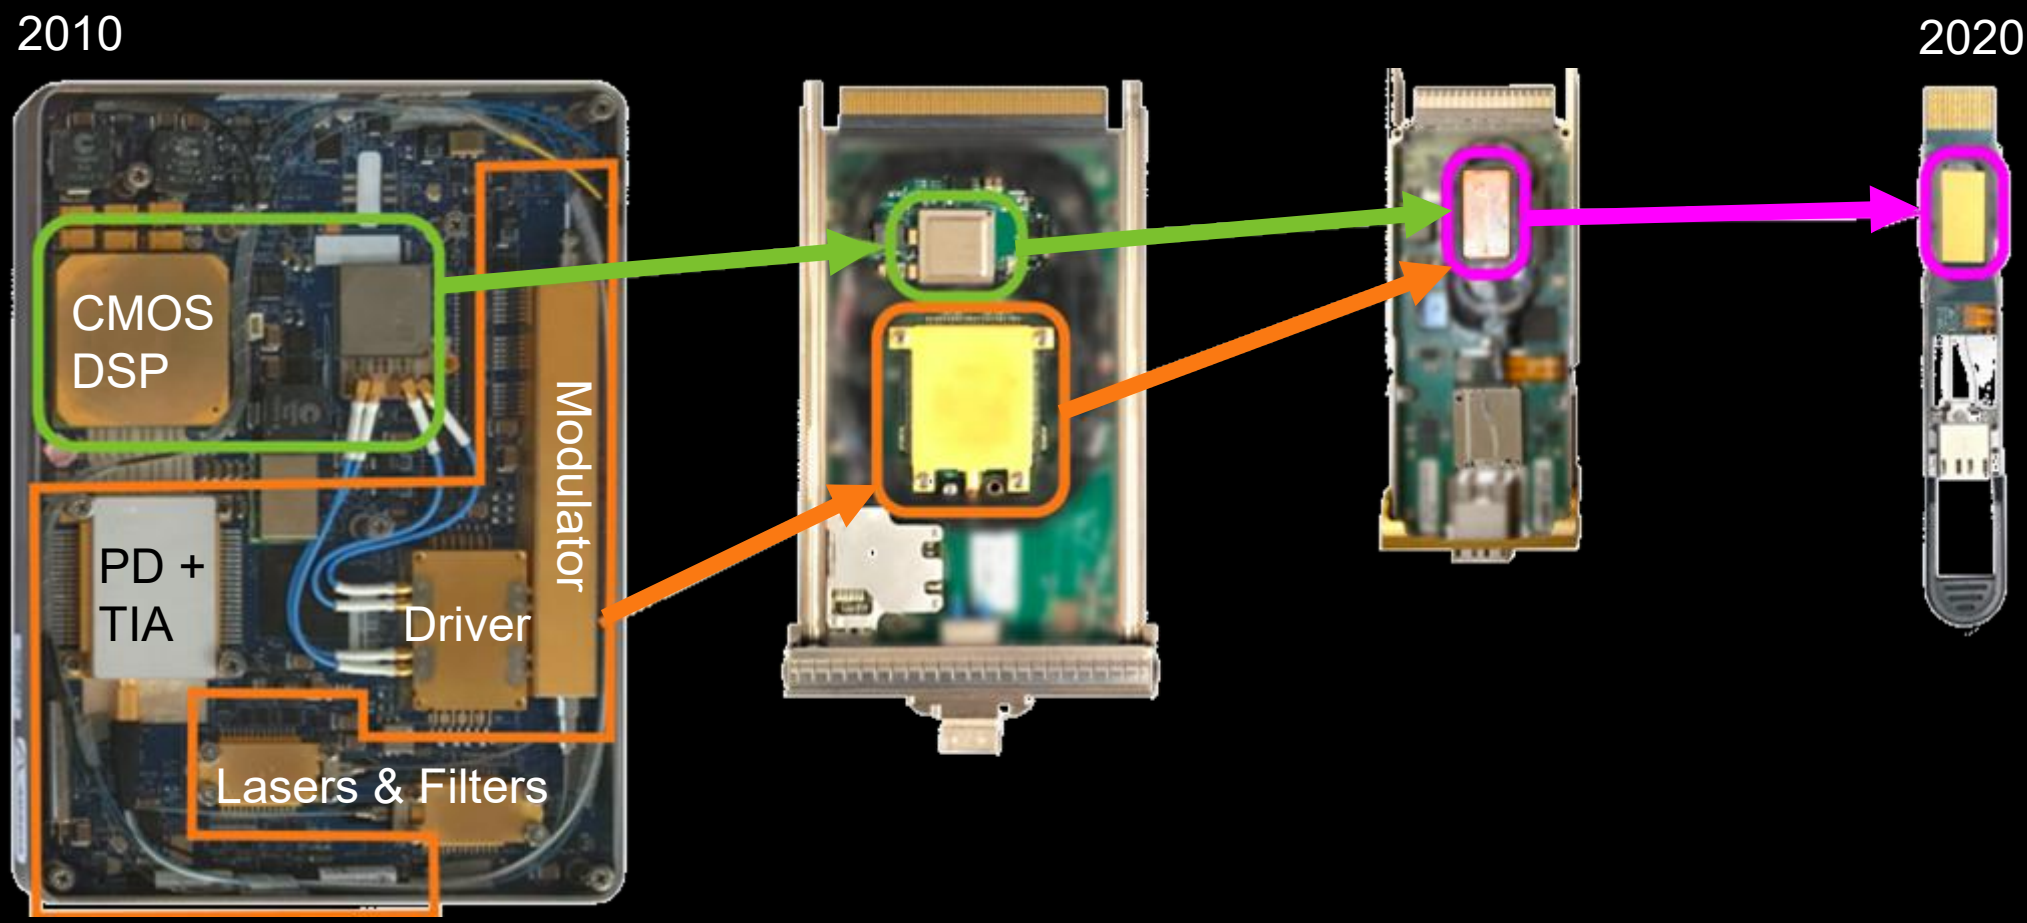

https://acacia-inc.com/acacia-resources/coherent-for-service-provider-edge-and-access-network-applications/

AMD
together we advance_

Public

Co-Packaged Optics

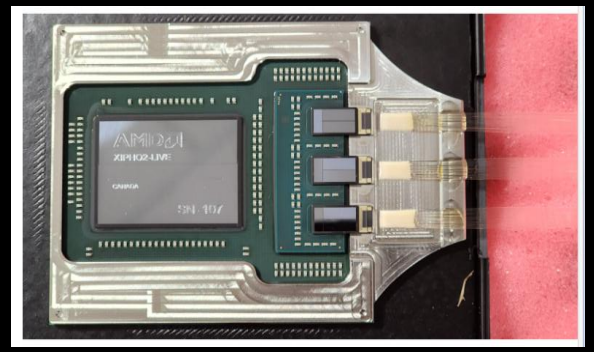



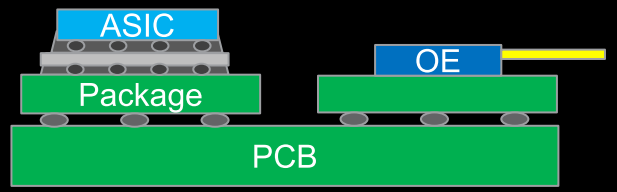

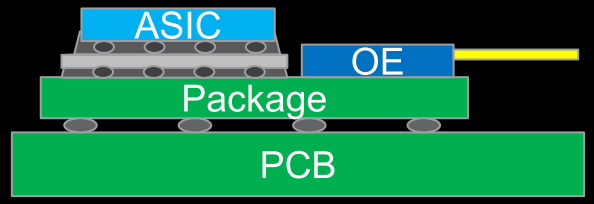

Pluggable Module	On-board Optics (OBO)	Co-Packaged Optics (CPO)
>15pJ/bit	~10-15pJ/bit	~5pJ/bit

- CPO provides ~3X improvement in energy efficiency
- CPO also enables higher bandwidth and compact design

AMD

Public

Key Metrics for CPO System-Level Optimization

- **All-in energy efficiency** (pJ/bit)
  - ASIC TX PHY
  - Optical TX PHY (or driver if direct-drive)
  - Optical RX PHY (or TIA if direct-drive)
  - Control systems (microring tuning, etc.)
  - ASIC RX PHY
  - Laser output power * wall plug efficiency

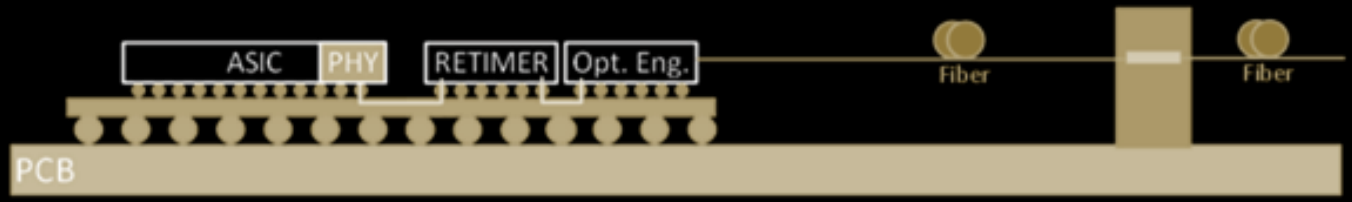

- **Link Budget** (dB), **Reach** (meters), number & Type of **Connectors** in optical path (#)
- Aggregate **Bandwidth**, Bandwidth per fiber (Gbps)
- **End-to-end latency** (ns)
- **Bandwidth/shoreline** (Gbps/mm) & bandwidth/area (Gbps/mm^2)
  - ASIC PHY
  - Electrical Fanout
  - Optical Chiplet
  - Fiber coupling + Keep-out zones
- **Ecosystem**

AMD
together we advance_

Public

OCI-MSA

Founding Members

Industry alignment around an optical scale-up.

Est. March 2026

Future AI infrastructure built on open standards.



Optical Scale-Up

Architecture that delivers the efficiency required to scale AI infrastructure.

Move Scale-up from In-Rack to Multi-Rack & Multi-Row

For more information and to request a copy of the OCI GEN1 specification please see www.oci-msa.org

Public

OCI-MSA Implementation

OPTICAL IMPLEMENTATION FOR SCALE-UP

Optical Scale-up only limited by RADIX



OCI-MSA
Est. March 2026

Scalable lowest PWR

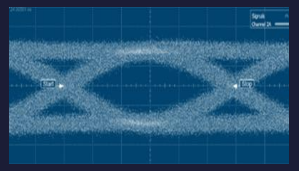

NRZ

Standard Lasers

WDM

Reduce Cost w/o RADIX impact

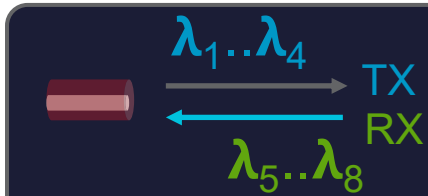

Bi-Di one fiber

- Lower Power
- Lower Cost
- Lower Latency
- High Bandwidth per Fiber

Move Scale-up from In-Rack to Multi-Rack & Multi-Row

For more information and to request a copy of the OCI GEN1 specification please see *www.oci-msa.org*

Public

Future System-in-Package Architecture

- Advanced packaging enables maximally efficient integration of compute elements and memory
- System level communication accomplished with low-power, high-bandwidth optical

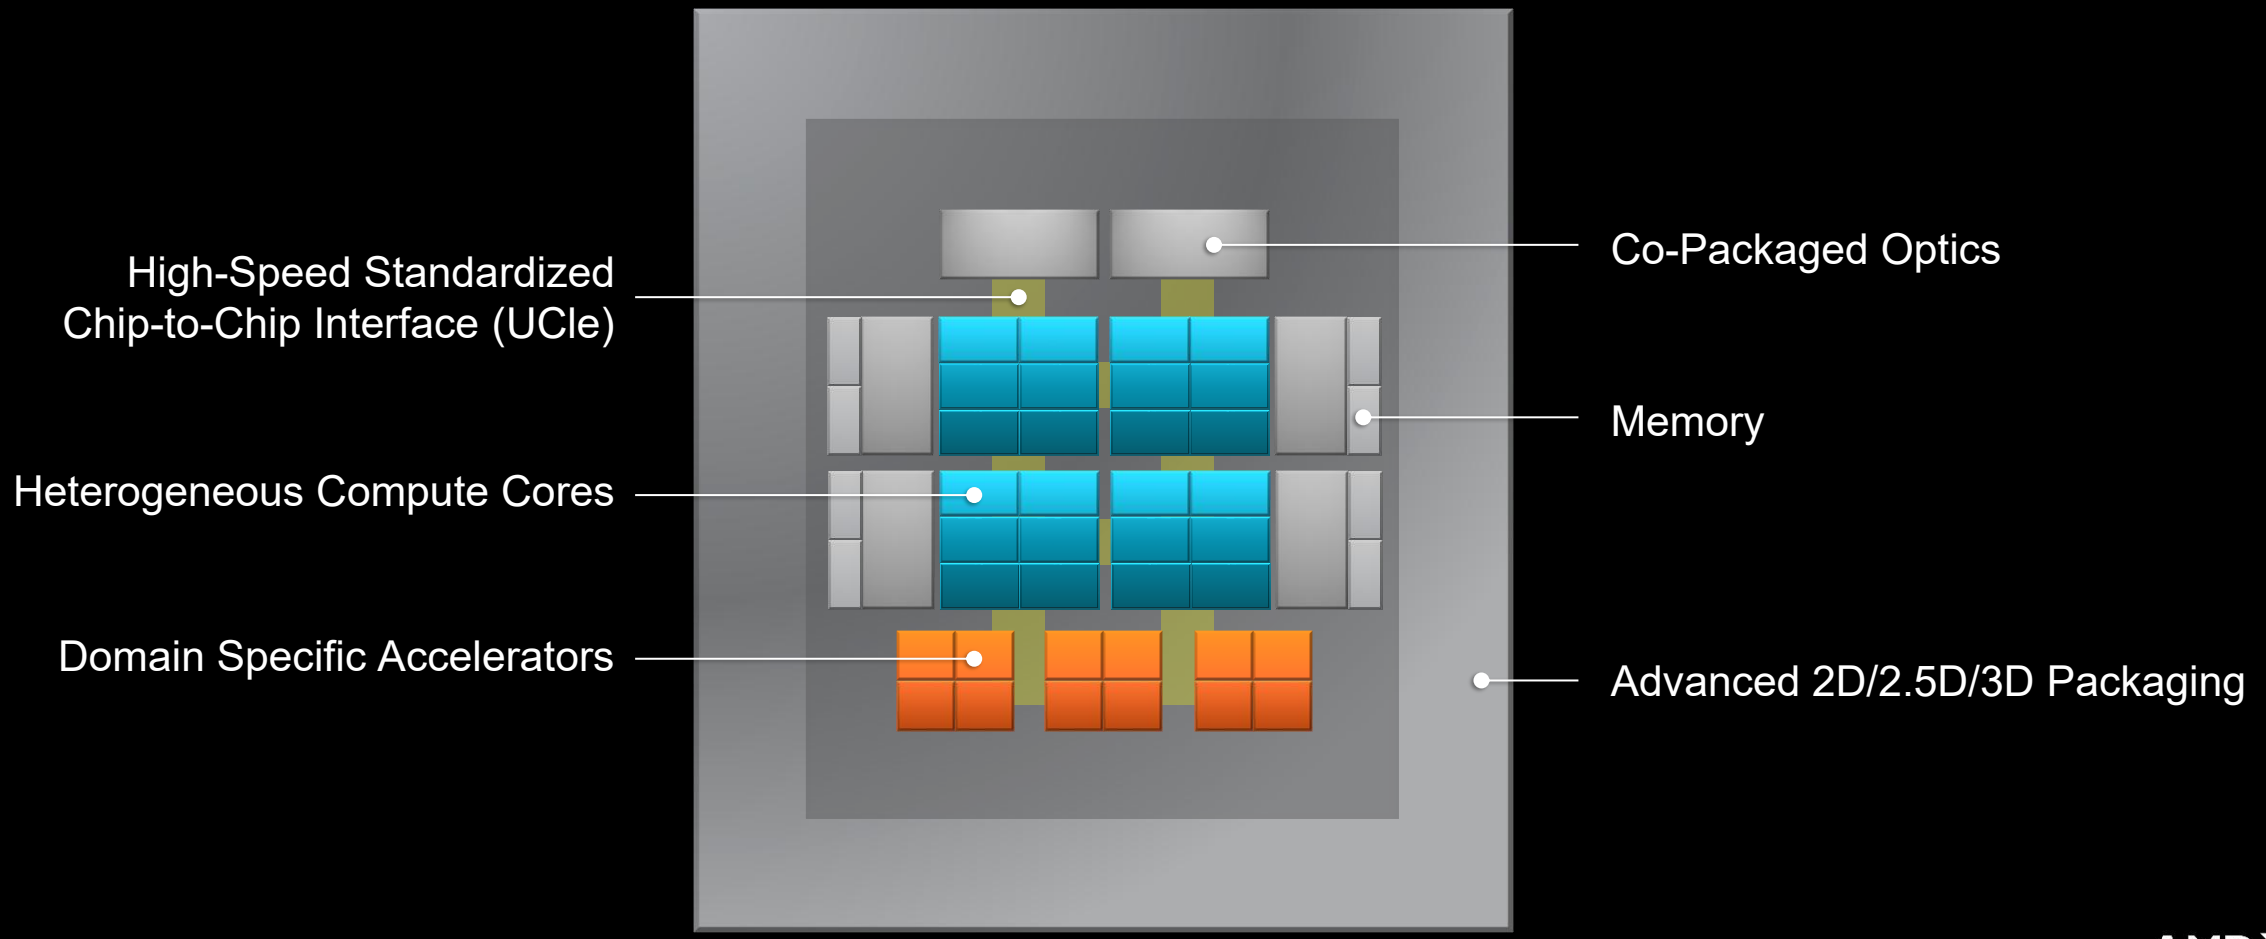

AMD

Public

Multi-Physics Modeling

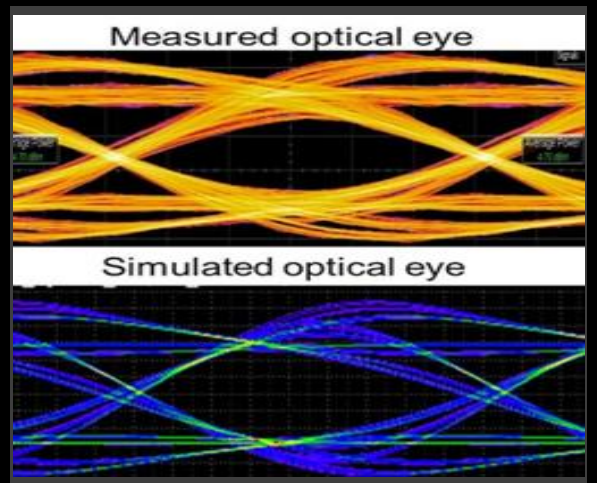

Optical

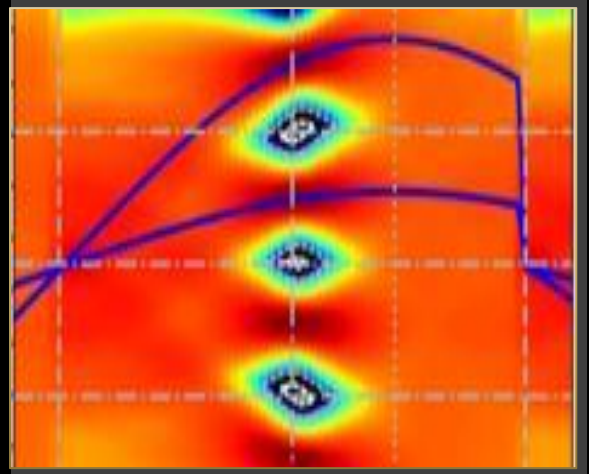

Electrical

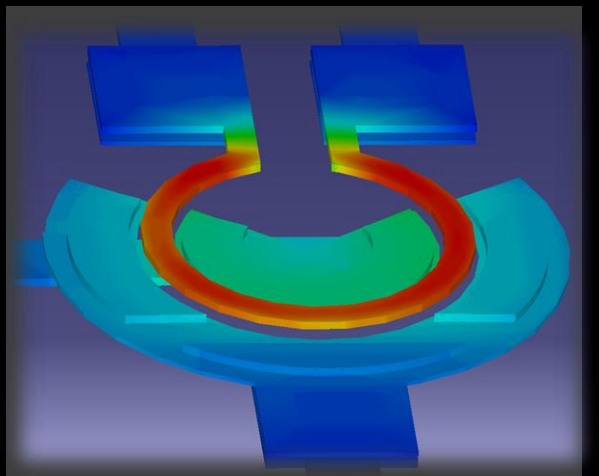

Thermal

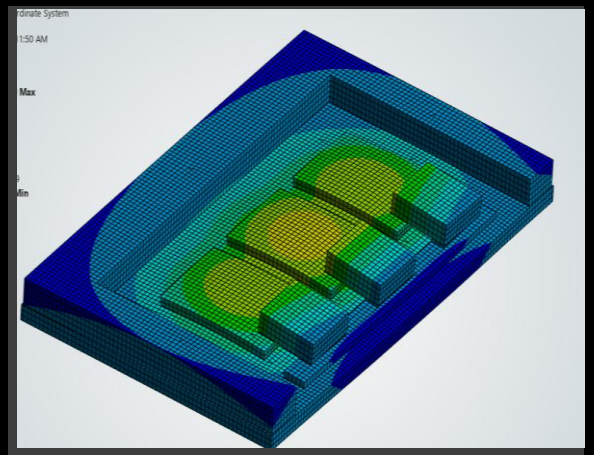

Mechanical

Optics design requires co-optimization across multiple disciplines

Public

Multi-Scale Models

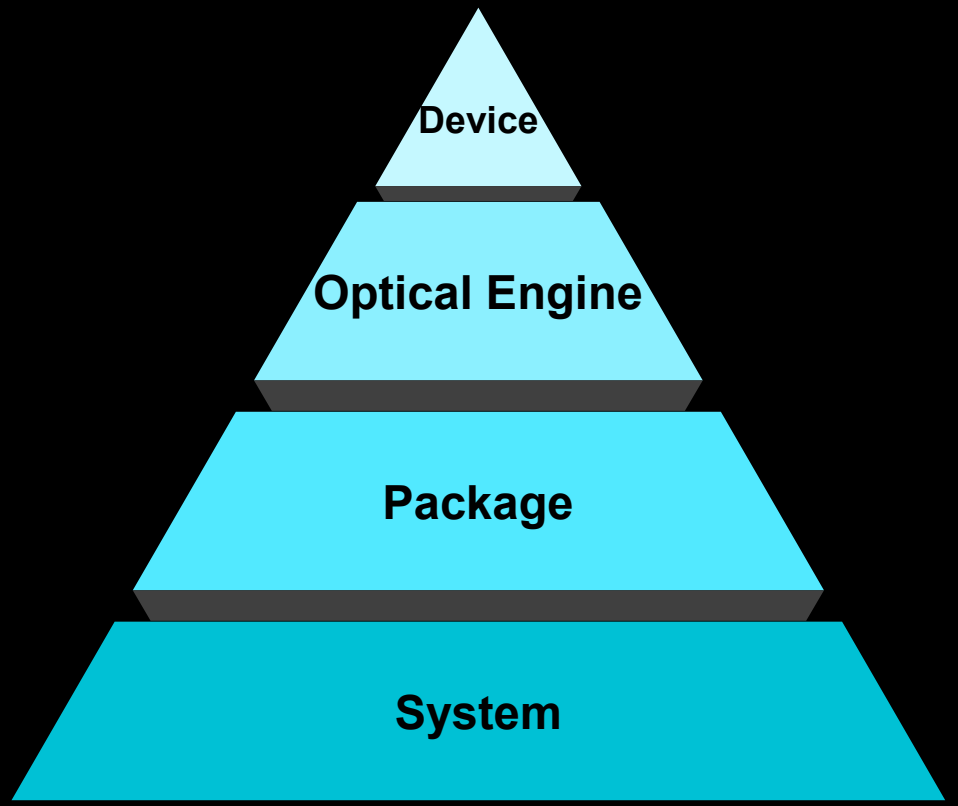

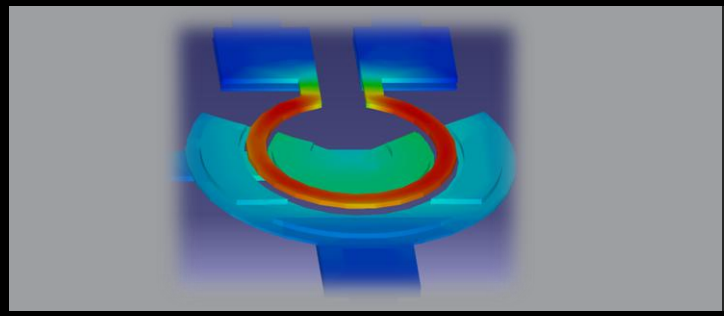

Device

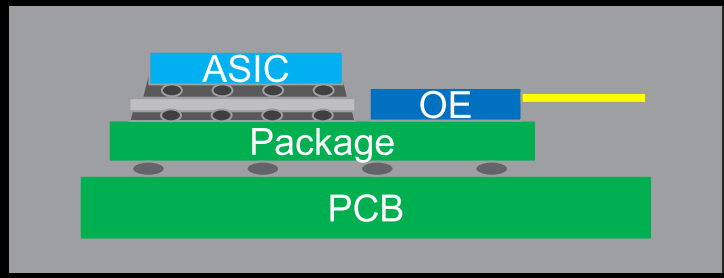

Package

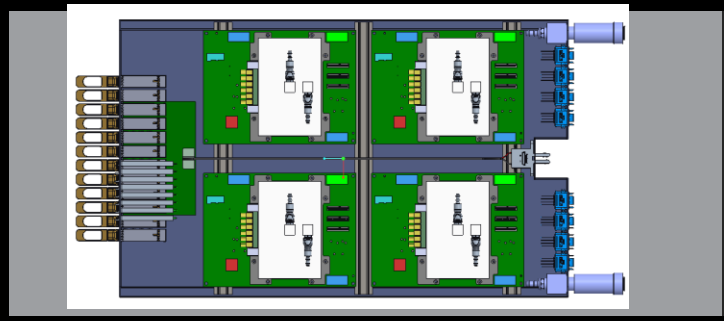

System

Optical design involves multi-scale models that comprehends chiplet, package, and system interactions

AMD
together we advance_

Public

Manufacturing Silicon Photonics at Scale

- **Reliable optical AI systems require multiple, simultaneous Integration solutions**
  - **Electrical:** ASIC and EIC interfaces
  - **Optical:** PIC, FAU, couplers, fibers
  - **Thermal: c**ompute density and optical stability
  - **Mechanical: a**ssembly, test, and rework

- **Optics integration and packaging are coupled problems, not independent workstreams**
  - **Alignment: l**ow-loss coupling at scale
  - **Connectors: d**etachable optical I/O
  - **Thermals: s**tability near compute
  - **Test: a**utomated electro-optical test
  - **Reliability: l**ifetime of the full stack

The Optical Engine is an interacting system
A solution in one area can change the requirements in another

AMD
together we advance_

Public

Silicon Photonics at Scale: Bridging Packaging and Optics

Laser power and wavelength → **Fiber / FAU** alignment and handling → **Coupler** loss and contamination → **PIC / EIC** device and interconnect → **Package** stress and cycling → **Connector** mating durability

a CPO Module with ASIC and Fiber Optic Connectors
b ELSFP Laser Modules
c Passive Fiber Optic Connectors
d Faceplate
e Fiber Optic Cables

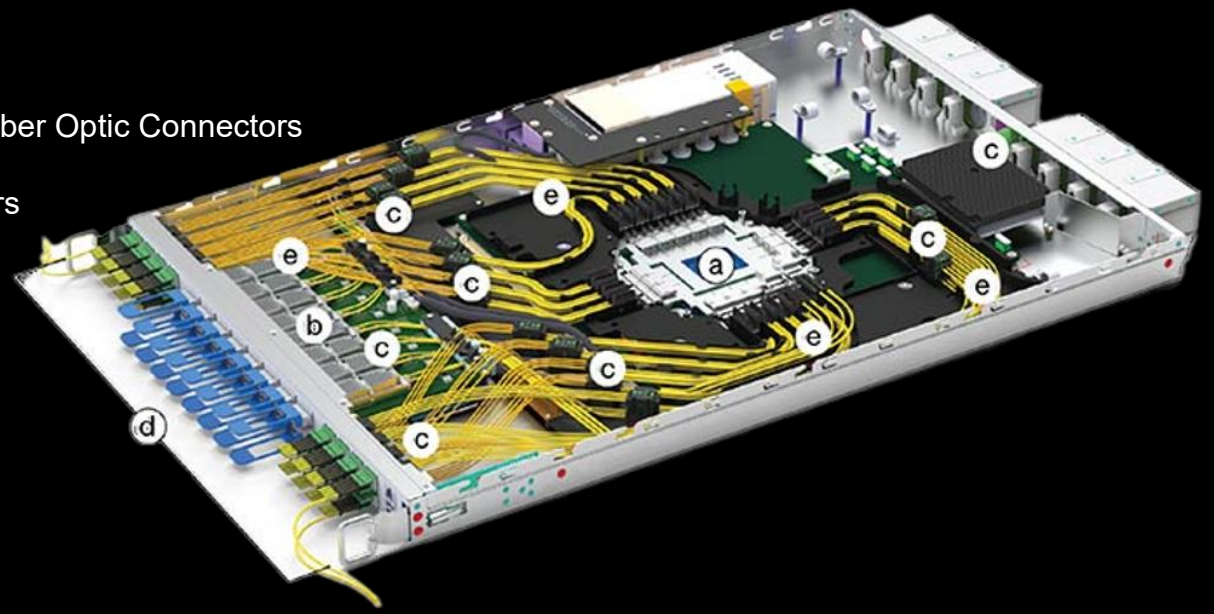

Optical performance depends on every interface in the fully-assembled system

https://www.photonics.com/Articles/The-New-Optical-Interface-Novel-Connector/a71803

AMD
together we advance_

Public

Silicon Photonics at Scale: Bridging Packaging and Optics

- AI has created an Interconnect inflection point
- Silicon Photonics is an Integration platform
- The Optical Engine is becoming part of the Package
- Alignment, thermals, test, and reliability are coupled
- Manufacturing scale is the next innovation frontier

The challenge is no longer proving that Silicon Photonics works
The challenge is scaling it into high-yielding, reliable AI systems

AMD
together we advance_

Acknowledgements Deepak Kulkarni, Matthew Streshinsky



Public

End notes

MI350-059: Based on calculations by AMD Performance Labs as of July 2025, using the total memory capacity of a single AMD Instinct MI350X / MI355X 8x GPU Platform (2304 GB HBM3) vs a single Nvidia B200 8x GPU Platform (1440GB). Calculations performed with FP4 precision datatype at (.5) bytes per parameter, to determine the minimum number of 8x GPU platform(s) required to run a Generic Inference model at 4.19T parameters (AMD) and 2.62T parameters (NVIDIA). Stated results include a calculated 10% overhead. Server manufacturers may vary configurations, yielding different results. Results may vary based on GPU memory configuration, LLM size, and any potential variances in GPU memory access or the server operating environment. MI350-059

MI350-060: Based on calculations by AMD Performance Labs as of July 2025, using the total memory capacity of a single AMD Instinct MI350X / MI355X 8x GPU platform (2304GB HBM3) vs a single Nvidia B200 8x GPU Platform (1440GB). Calculations performed with FP4 precision datatype at (.5) bytes per parameter, optimizer states at 16 bytes, memory for weights and gradients at 135B (AMD) and 84B (NVIDIA), memory for activations at 2160GB (AMD) and 1334GB (NVIDIA), and total memory required to run the training model at 2295GB (AMD) and 1428GB (NVIDIA), to determine the minimum number of 8x GPU platform(s) required to run a Generic Training model at 135B parameters (AMD) and 84B parameters (NVIDIA). Calculations include a 10% overhead. Server manufacturers may vary configurations, yielding different results. Results may vary based on GPU memory configuration, LLM size, and any potential variances in GPU memory access or the server operating environment. MI350-060

AMD

Public

Copyright and disclaimer

©2026 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow logo, AMD Helios™, AMD Instinct™and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate releases, for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS." AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION


AMD