

# AMD Multi-die Architecture for Exa-Class HPC and AI Systems

Forum Teratec 2025  
May 22<sup>nd</sup>, 2025

Jose Noudohouenou

**AMD**   
together we advance\_

# AMD PLATFORM FOR ACCELERATED COMPUTING

## INVESTING IN CORE TECHNOLOGY FOR LEADERSHIP IN HPC & AI

### COMPUTE

WORKLOAD-OPTIMIZED  
COMPUTE ARCHITECTURE  
w/ WIDE RANGE OF DATA  
FORMAT SUPPORT


AMD  
**CDNA**

FP64    FP32  
FP16    BF16  
FP8     INT8

And more...


### MEMORY

HIGHEST CAPACITY  
MEMORY AND BANDWIDTH  
AVAILABLE IN THE  
INDUSTRY



### NETWORKING

ADVANCING NETWORK  
BANDWIDTH WITH  
SUPPORT FOR INDUSTRY  
STANDARD & CUSTOM  
TECHNOLOGIES



### SOFTWARE

FRICITIONLESS SW  
ECOSYSTEM W/DROP-IN  
SUPPORT FOR LEADING  
PROGRAMMING MODELS &  
AI FRAMEWORKS


OpenMP

 kokkos

 RAJA

 PyTorch

 TensorFlow

 deepspeed

 Triton


 OpenXLA

# Powering the Top 2 Supercomputers in the world

2 Generations of Instinct™ - MI300A & MI250X leading Top500



5 of Top10

The GREEN 500 

15 of Top25

#1 EL CAPITAN

#2 FRONTIER

#5 eni

#8 LUMI

#10 Lawrence Livermore National Laboratory

GENCI  
INES  
69 GF/W

Lawrence Livermore National Laboratory  
63 GF/W

FRONTIER  
63 GF/W

Top 10

4.2 <sub>EF</sub>   61%	1.2 <sub>EF</sub>   18%
AMD Instinct™	Nvidia


AMD Instinct™ MI300A 69 <sub>GF/W</sub>	AMD Instinct™ MI250X 63 <sub>GF/W</sub>
--	--

# AMD Chiplet Ecosystem



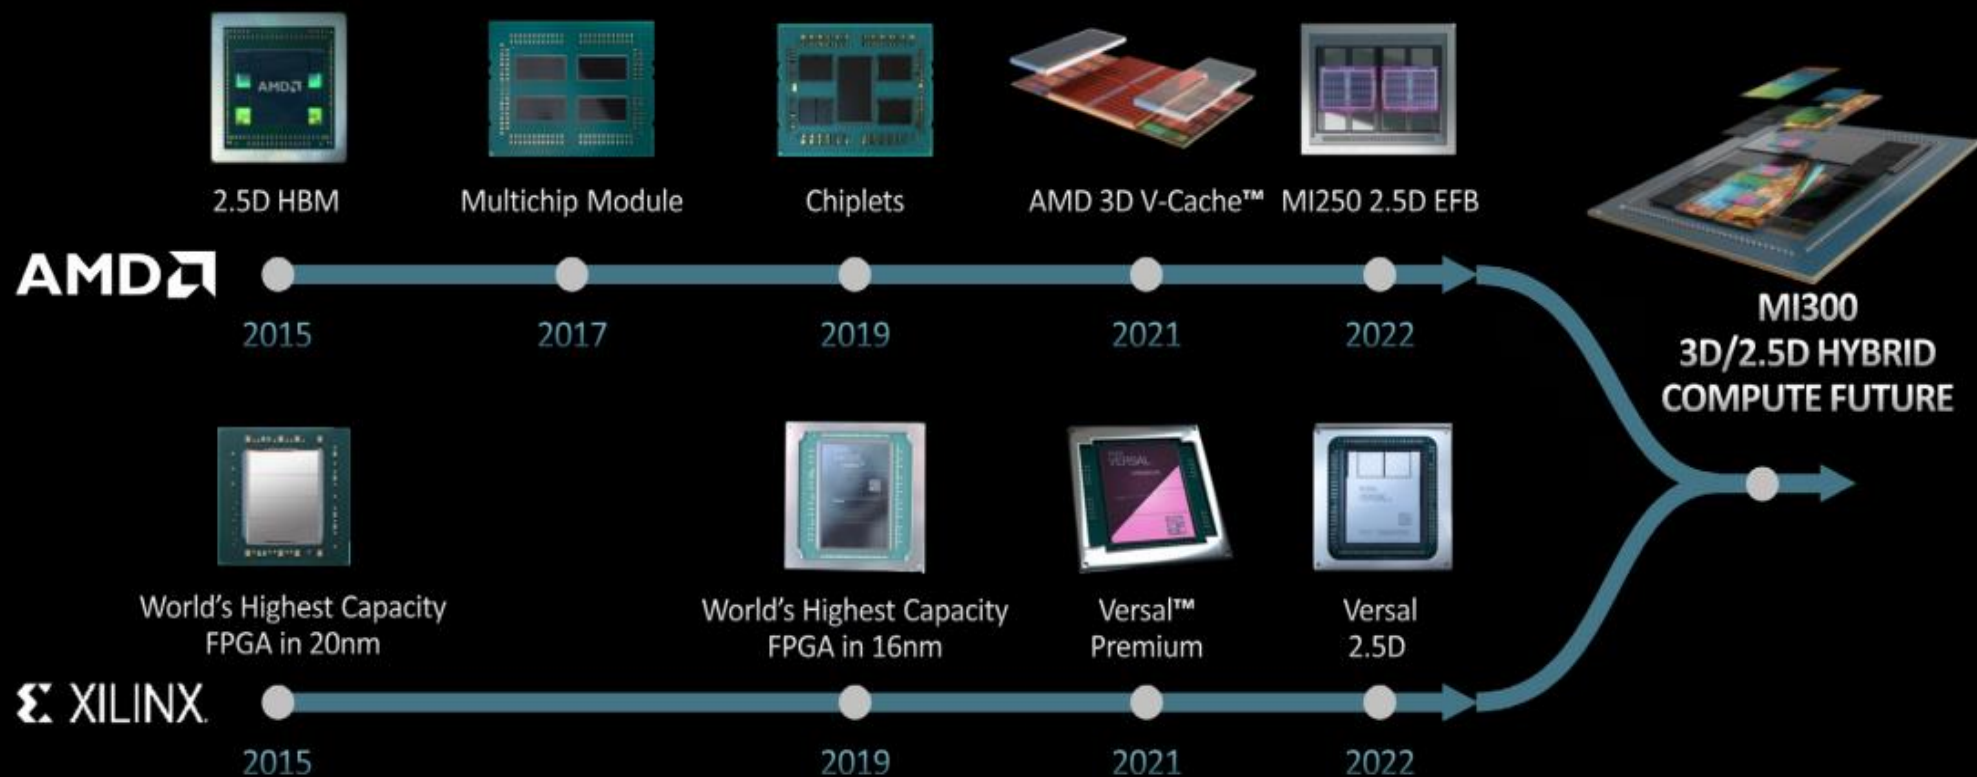
# Chiplet

- Break up the monolithic silicon into smaller functions
- Modular integrated circuits that are combined (co-packaged) to create larger, more complex semiconductor devices
- Provide designers and manufacturers with flexible and scalable design capabilities to meet the design requirements of modern SoCs



MI300 Series Modular Chiplet Package  
HPC and AI Accelerator for Exa-Class Systems

# AMD Chiplet Innovations



- More than 10 years of innovation in chiplet architecture.
- In 2019, AMD's 2.5D chiplet technology was introduced with the AMD Ryzen and AMD EPYC processors.
- In 2023, AMD released the Instinct™ MI300X AI accelerator that incorporates the latest 2.5D and 3D technology.

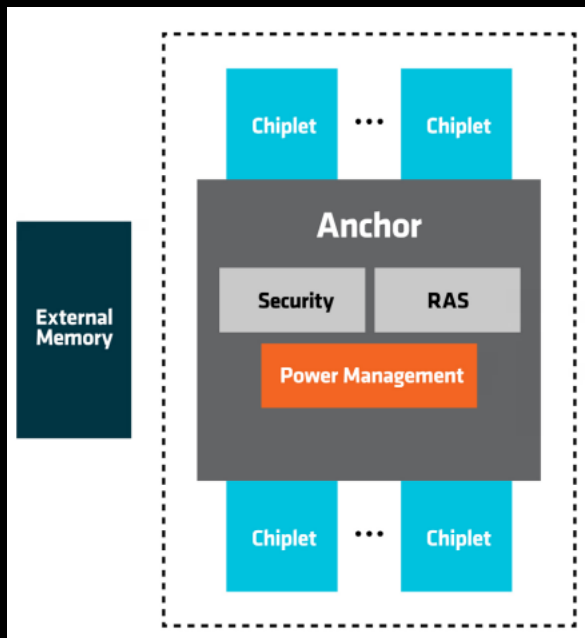
# Chiplet Capabilities

A robust chiplet ecosystem needs to support a diverse range of chiplet types, each offering unique functionalities within an SoC:

- I/O capable chiplets
- Accelerator chiplets
  - Custom arithmetic engines
  - Compression/encoding engines
  - Networking engines
- Compute cores
  - CPU cores
  - GPU cores

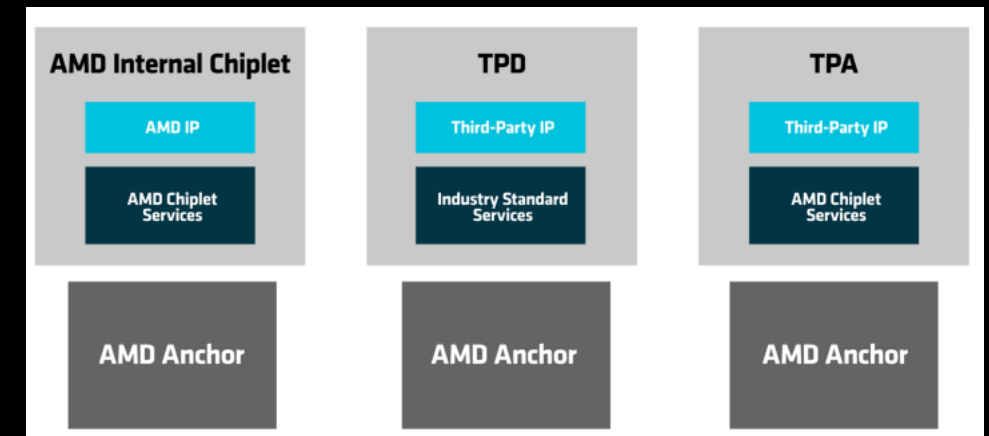
# Building Chiplet-based SoCs

- An SoC (System-on-Chip) built with chiplets relies on a central chiplet called an anchor. This anchor orchestrates the essential system level functions:
  - Power Management
  - Security
  - Reliability, Availability, and Serviceability (RAS)
  - Interconnectivity



SoC Integration Using Chiplets

- Chiplet integration models by AMD
  - Internal Chiplets
  - Third-Party Die (TPD)
  - Third-Party Adapted die (TPA)



AMD, TPD, and TPA Chiplet Integration Diagrams



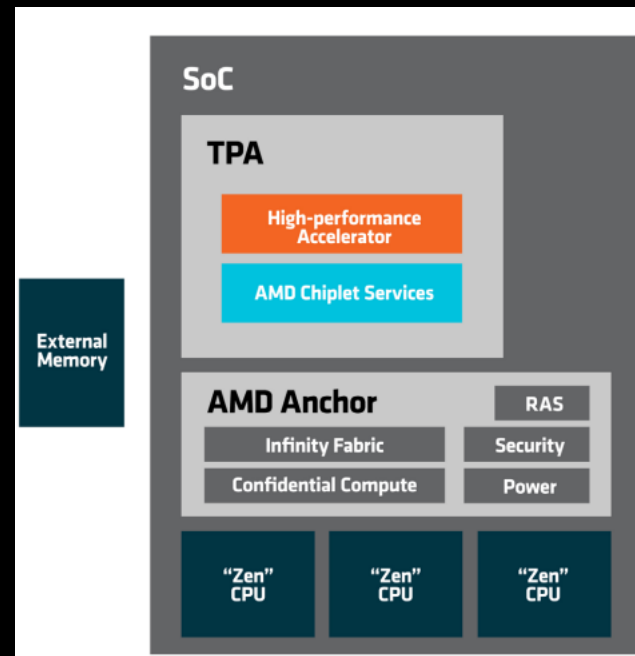
# Services for Chiplet Integration

- Data Path communication: communication services
  - Memory access: Efficiency access to internal, external and virtual memories.
  - Caching and Coherency: Support to ensure all processors and accelerators in the SoC see the same view of memory
- Power management:
  - Simple static power states
  - Dynamic power management
  - Chiplet-level power management
- Reliability, Availability, and Serviceability (RAS):
  - Error reporting
  - Isolation and recovery
  - Monitoring and preventive mitigations
- Thermal management:
  - Chiplet level thermal management
  - SoC thermal failure management
- Security:
  - Centralized security controls
  - Chiplet-level security

# Chiplet Integration: Example

Example: integrating a cutting-edge accelerator from an external vendor into an AMD “Zen”-based SoC

- AMD’s TPA integration model is used:
  - Interfaces with the AMD SoC through AMD’s Chiplet Communications
  - Providing premium access to AMD’s datacenter-class Anchor infrastructure



Enhancing “Zen”-Based SoCs  
with Third-Party Acceleration

# Building Exa-Class HPC and AI Systems



# Some AMD Instinct™ Accelerators




Launch Year	Accelerator	# Computing Units	Memory (HBM)	Power(W)
2024	MI325X	304	256	1000
2023	MI300X	304	192	750
2023	MI300A (APU)	228	128	760 (CPU+GPU)
2021	MI250X	220	128	560
2021	MI250	208	128	560

Thermal: Liquid Cooled & Air Cooled

- Over the years:
- More compute units
  - More memory (and BW)
  - More power-hungry GPUs
  - New datatypes support at HW level  
(well suited for AI applications/benchmarks)

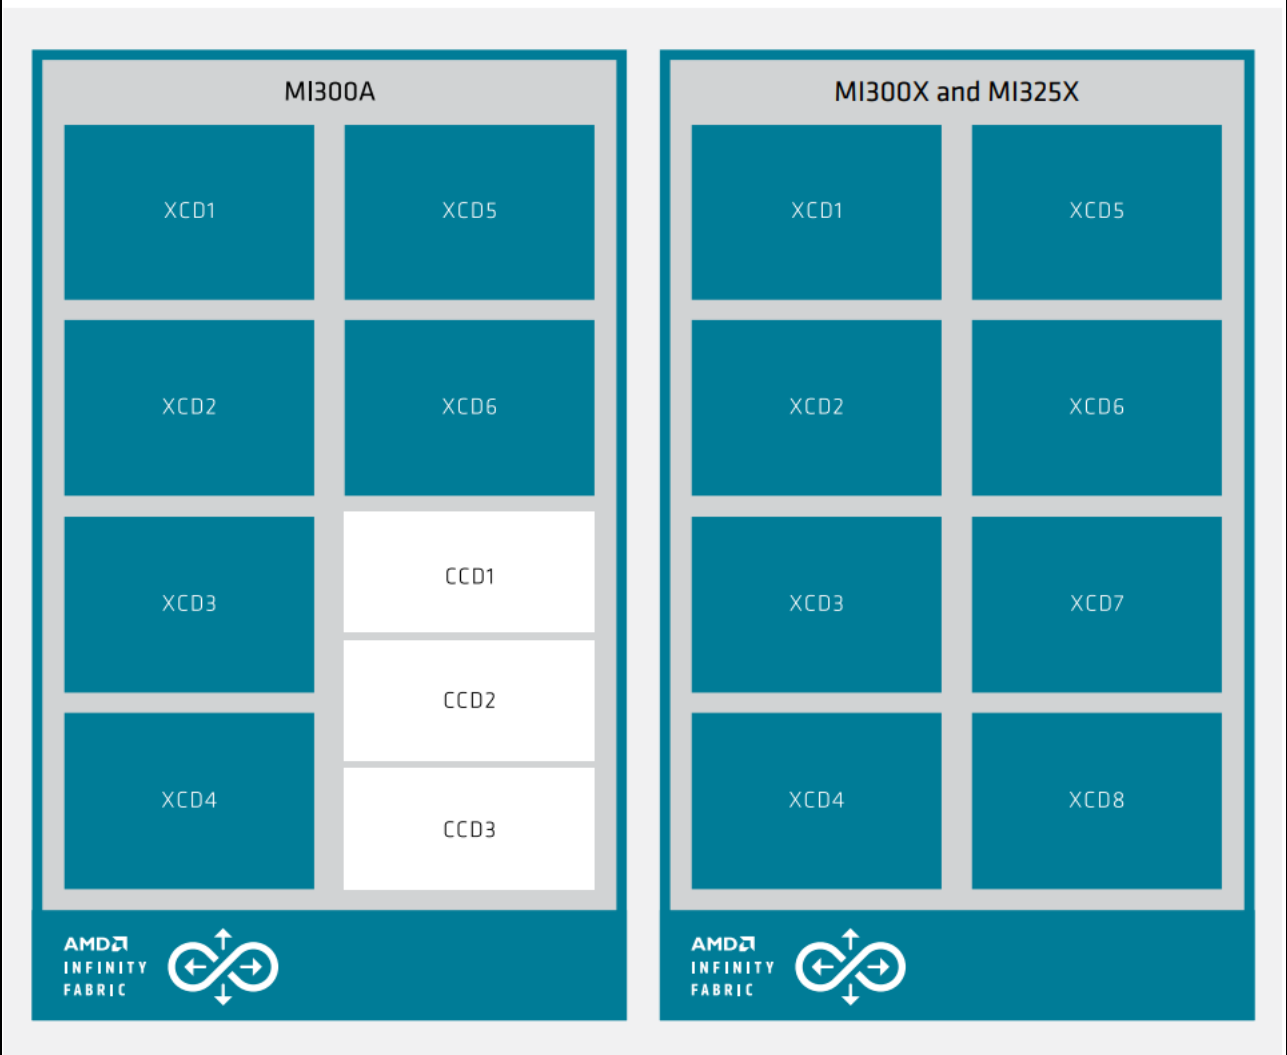
[https://en.wikipedia.org/wiki/AMD\\_Instinct](https://en.wikipedia.org/wiki/AMD_Instinct)

# Building AMD Instinct™ Accelerators: Example of MI300 GPU Series




Advanced 3D Package and chiplet-based construction of the AMD Instinct™ MI300 Series processors

# Building AMD Instinct™ Accelerators: Example of MI300 GPU Series

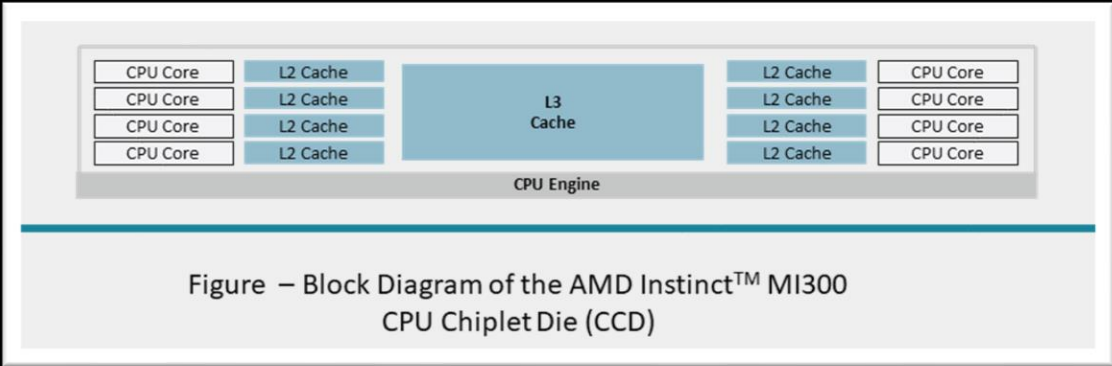


Block diagram of the AMD Instinct™ MI300A APU, MI300X and MI325X discrete GPUs

# MI300: LEADERSHIP HPC & AI SOLUTIONS




Accelerated Compute Die (XCD) Cache	AMD Instinct™ MI300A	AMD Instinct™ MI300X
L1 Cache / CU	32 KB	32 KB
# Active CU / XCD	38	38
L2 Cache Shared Between CUs	4 MB	4 MB
CDNA™ 3 Accelerated Compute Dies (XCD)	6	8
LLC Cache Shared Across XCDs	256 MB (also shared with CCDs on MI300A)	256MB



CPU Chiplet Die (CCD) Cache	AMD Instinct™ MI300A
L2 Cache / CPU Core	1MB
# Active CPU Cores / CCD	8
L3 Cache Shared Between GPU Cores	32 MB
INSTINCT™ CPU Chiplet Dies (CCD)	3
LLC Cache Shared Across CCDs	256 MB (also shared with XCDs on MI300A)


# AMD CDNA™ 3 ARCHITECTURE MEMORY ARCHITECTURE DIAGRAM



XCD: Accelerated Complex Die  
IOD: I/O dies  
HBM: High-Bandwidth Memory  
LLC: Last Level Cache

**LEGEND**

- Green line is L2 to XCD  
51.6 TB/s (Aggregate)
- Yellow line is LLC to XCD  
17.2 TB/s (Aggregate)
- Red line is LLC to XCD Bandwidth  
5.3 TB/s (Aggregate)



XCD: Accelerated Complex Die  
CCD: CPU Complex Die



# MI300: LEADERSHIP HPC & AI SOLUTIONS

## MI300A: APU

## MI300X: Discrete GPU

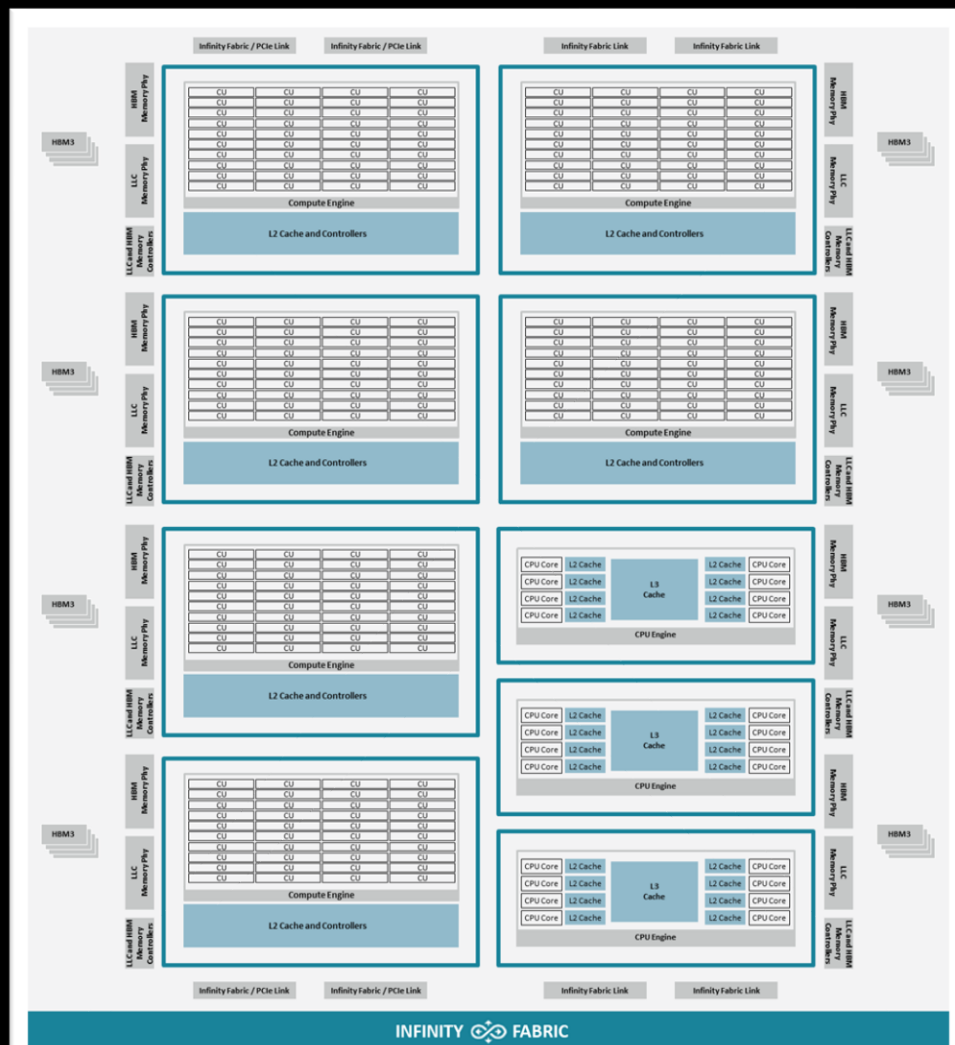


Figure – Block diagram of the AMD Instinct™ MI300 multi-chip module (AMD Instinct™ MI300A) in the SH5 socket form factor accelerator which comprise of three CPU Chiplet Dies (CCD) and six Accelerated Compute Dies (XCD) as illustrated.

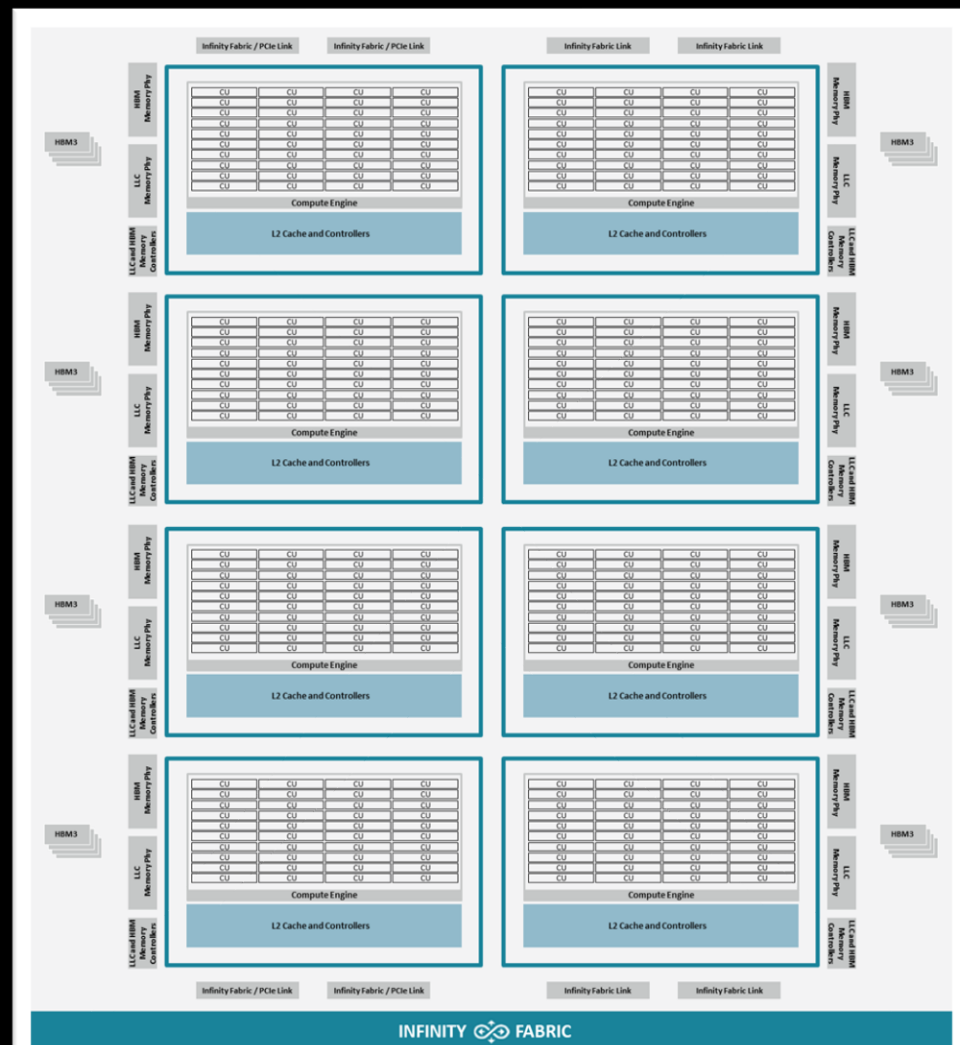


Figure – Block diagram of the AMD Instinct™ MI300 multi-chip module (AMD Instinct™ MI300X) in the OAM form factor accelerator which comprise eight Accelerated Compute Dies (XCD) as illustrated.

# Advantages of Adopting Multi-die Architecture



# Building Multi-die Architecture for Exa-Class HPC and AI Systems: Advantages

## Benefits of using Chiplet over monolithic silicon

- Cost-effective: smaller chiplets have higher manufacturing yields, reduce waste, and lower costs
- Scalability and flexibility: chiplets are modular and can be mixed and matched for specific needs, allowing for easier upgrades and customization.
- Faster innovation: separate chiplet development allows for parallel development and faster innovation cycles
- Performance: chiplets provide scalability and flexibility that allows for specialized designs

AMD is committed to capitalizing on the advantages of chiplets


# ENABLING MULTIPLE WORKLOADS FOR OPTIMAL GPU UTILIZATION

## MI300 PARTITIONING


### MI300A - APU

- Maximize GPU utilization with 3 partitions
- NPS Modes\* (NPS1)

Single partition




Three partitions




### MI300X - OAM

- Maximize GPU utilization with up to 8 partitions
- NPS modes\* (NPS1, NPS4)


Single partition




Two partitions



Four partitions



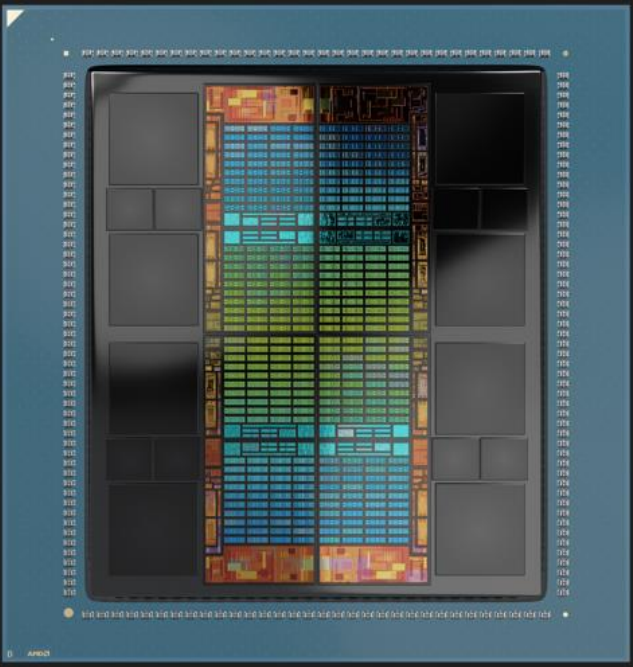
Eight partitions



\* → Memory partitioning can only be changed via a re-boot

# AMD Instinct™ MI300X GPU Partitioning


Support for up to 8 SR-IOV Virtual Functions per GPU

GPU Partition Options	1 (SPX)	2 (DPX)	4 (QPX)	8 (CPX)
	<div>GPU Instance 192GB</div>	<div>GPU Instance 96GB</div> <div>GPU Instance 96GB</div>	<div>GPU Instance 48GB</div> <div>GPU Instance 48GB</div> <div>GPU Instance 48GB</div> <div>GPU Instance 48GB</div>	<div>GPU Instance 24GB</div> <div>GPU Instance 24GB</div> <div>GPU Instance 24GB</div> <div>GPU Instance 24GB</div> <div>GPU Instance 24GB</div> <div>GPU Instance 24GB</div> <div>GPU Instance 24GB</div> <div>GPU Instance 24GB</div>

Partitioning mode selected applies to all MI300X GPUs on UBB8 Node



# AMD INSTINCT™ MI300X UBB8 GPU VIRTUALIZATION



SINGLE 8GPU 1.5TB VM INSTANCE PER NODE  
Large AI Training

EIGHT 1GPU 192GB VM INSTANCES PER NODE  
Large AI Inference

KVM HYPERVISOR SUPPORT  
Ubuntu Host, Ubuntu Guest OS

SR-IOV VIRTUAL FUNCTIONS

INFINITY FABRIC™ INTERCONNECT SUPPORT

# Multi-die Architecture Design Benefits (1)

- GPU partitioning
  - Compute units partitioning
  - GPU memory partitioning
- Benefits
  - Ease GPU partitioning even at HW (bare-metal) level
  - Optimizing resource allocation (assigning workloads/tasks to specific partitions)
  - Allowing more efficient resource utilization and reducing the need for whole, more power-hungry GPUs
  - Reducing power consumption



## Multi-die Architecture Design Benefits (2)

- Give AMD opportunities to address the widest set of customer needs
- Ease to design and manufacture specialized machines and new datatype supports (at HW level) for example:
  - Machines dedicated to HPC
    - Maintaining FP64 at HW level
  - Machines specialized for AI
    - Reducing FP64 at HW level and integrating more and new datatype units like FP4, FP6, FP8, etc. needed by AI applications/benchmarks/workloads
    - Ease GPU memory (HBM) size increasing since large AI applications require more memory



Available 2025

# AMD Instinct™ MI350 Series

## Continued Gen AI Leadership

**AMD**  
CDNA 4


3nm  
Process Node

Up to  
288GB  
HBM3E

**NEW**  
FP4 / FP6  
Datatype Support

# Summary (1)

AMD uses multi-die architectures for designing Exa-Class HPC and AI Systems



- Benefits of considering that approaches includes:
  - Cost-effective
  - Scalability and flexibility
  - Faster innovation
  - Performance

## Summary (2)

AMD continues cultivating a strong chiplet ecosystem that maximizes the benefits of heterogeneous integration

- Proposed solutions include:
  - Bringing in external proficiency while ensuring compatibility, performance, and reliability within our SoCs (AMD intermediate third-party die (iTPD))  
Example: Integrating I/O subsystem chiplets from external vendors with AMD anchor chiplet
  - Facilitating the smooth incorporation of third-party capabilities (third-party adapted die (TPA))  
Example: enables customers to enhance both functionality and performance by integrating advanced accelerators from third-party vendors

AMD actively supports the UCle standard and are committed to its development to achieve seamless communication and compatibility among various providers and chip types

Collaboration readiness: partnering with AMD for custom chiplet development

**Questions?**

# DISCLAIMERS AND ATTRIBUTIONS

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's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18

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 CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

© 2025 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow logo, Radeon™, Instinct™, EPYC, Infinity Fabric, ROCm™, 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.

