

AMD INSTINCT™ MI300X & MI300A ACCELERATOR

AMD INSTINCT MI300X

AMD Instinct MI300X GPU は、次世代の AMD CDNA 3 アーキテクチャに基づいており、最も要求の厳しい AI および HPC アプリケーションでの優れた効率とパフォーマンスを実現します。304 個の高スループットコンピューティングユニットや新しいデータタイプのサポート、写真とビデオのデコードなどの AI 固有の機能、及び GPU アクセラレータでは前例のない 192GB の HBM3 メモリが搭載されているのが特徴です。

AMD INSTINCT MI300A

AMD Instinct MI300A は、AI や HPCでの処理を強化するように設計されており x86プロセッサと AMD Instinct アクセラレータ、HBM3 メモリを 1つのパッケージに統合した APU(アクセラレータド プロセッシング ユニット)です。CPU と GPU がメモリアドレス空間を共有し内部インターコネクットの帯域を圧迫せず高性能化を実現しており、電力効率の高いチップレットテクノロジーによる高度なパッケージングが特徴です。

AMD Instinct MI300 Series accelerators specifications and features

	MI300X GPU	MI300A APU
ARCHITECTURE	AMD CDNA 3	AMD CDNA 3
ACCELERATED COMPLEX DIES (XCD)	8	6
COMPUTE UNITS	304	228
STREAM PROCESSORS	19,456	14,592
MATRIX CORES	1,216	912
MAX ENGINE CLOCK (PEAK)	2,100 MHz	2,100 MHz
AMD“ZEN 4”CPU CHIPLETS (CCD)	NA	3
TOTAL“ZEN 4”X86 CORES	NA	24
UNIFIED MEMORY ACROSS CPU & GPU	NA	Yes

AMD Instinct MI300 Series vs NVIDIA H100 SXM5


	MI300X GPU	H100 SXM5	MI300A APU
PERFORMANCE (PEAK THEORETICAL)			
FP64	81.7 TF	33.5 TF	61.3 TF
FP32	163.4 TF	66.9 TF	122.6 TF
TF32 (SPARSITY)	1,307.4 TF	989.6 TF	980.6 TF
FP16 (SPARSITY)	2,614.9 TF	1,978.9 TF	1,961.2 TF
BF16 (SPARSITY)	2,614.9 TF	1,978.9 TF	1,961.2 TF
FP8 (SPARSITY)	5,229.8 TF	3,957.8 TF	3,922.3 TF
MEMORY			
MEMORY CAPACITY	192 GB HBM3	80 GB HBM3	128GB HBM3
MEMORY BANDWIDTH (PEAK)	up to 5.3 TB/sec	up to 3.35 TB/sec	up to 5.3 TB/sec
TDP			
MAX POWER	750W	700W	550W 760W

AMD Instinct MI300 Series Server - Coming Soon

8GPU SYSTEM

Powered by


AMD Instinct MI300X



QUAD APU SYSTEM

Powered by

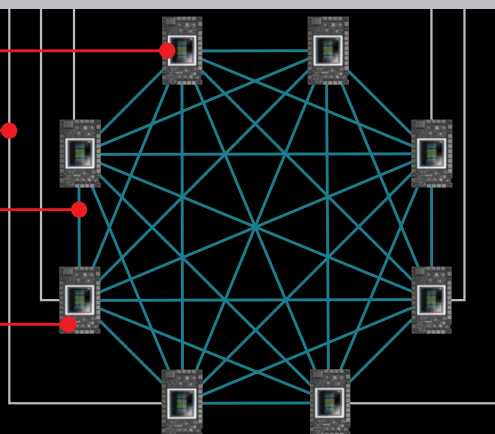
AMD Instinct MI300A



Model number	HPCT RS8E52-MI300X	HPCT RS4AI4-MI300A
GPU Support	<ul style="list-style-type: none"> • AMD Instinct MI300X Platform with 8 MI300 x OAM GPUs 	<ul style="list-style-type: none"> • 4x AMD Instinct MI300A APUs in SH5 sockets
CPU / APU Support	<ul style="list-style-type: none"> • 2x AMD EPYC 9004 /9005 Series processors 	
Memory Slots & Capacity	<ul style="list-style-type: none"> • 12-channel DDR5 memory support • 24 DIMM slots for up to 6 TB ECC DDR5-4800 RDIMM 	<ul style="list-style-type: none"> • 128 GB HBM3 memory (5.3 TB/s) per APU, 512 GB total per system • Last-level 256 MB AMD Infinity Cache shared between CPU and GPU Cores
Form Factor	<ul style="list-style-type: none"> • 8U rackmount 	<ul style="list-style-type: none"> • 4U rackmount
Expansion Slots	<ul style="list-style-type: none"> • 8 PCIe 5.0 x16 low-profile slots connected to GPU via PCIe switch • 2 PCIe 5.0 x16 full-height full-length slots • Optional 2 PCIe 5.0 x16 slots via expansion kit 	<ul style="list-style-type: none"> • 8 PCIe 5.0 x16 (or x8) slots or 16 Optional PCIe 5.0 x8 slots • 2 OCP PCIe 5.0 x16 AIOM slots
Storage	<ul style="list-style-type: none"> • 12 PCIe 5.0 x4 NVMe U.2 drives • Optional 4 PCIe 5.0 x4 NVMe U.2 drives • 2x M.2 NVMe boot drive • 2 hot-swap 2.5" SATA drives 	<ul style="list-style-type: none"> • 8 x2.5" U.2 NVMe hot-swap drives or 24x 2.5" SAS/SATA drives (optional) • 2 M.2 NVMe boot drives
Power Supplies	<ul style="list-style-type: none"> • 6x or 8x 3000W N+N redundant Titanium-Level power supplies 	<ul style="list-style-type: none"> • 2+2 redundant 1600W hot-swappable Titanium-Level power supplies


Universal Baseboard Connectors

- MI300X OAM Accelerators
- PCIe Gen 5 x16 I/O connectivity
- Fully-meshed 128GB/s bidirectional Infinity Fabric Connectivity
- 192GB HBM3 Memory / OAM (1.5TB Total)



AMD Infinity Platform Architecture

4APUs connected with AMD Infinity Fabric links



AMD Infinity Fabric

PCIe Gen 5 x16

A2ZEON 日本総代理店 ANSYS Discovery Live 代理店