



AMD INSTINCT™ MI300X ACCELERATOR

LEADERSHIP PERFORMANCE FOR AI & HPC

AMD 
together we advance_

Breakthrough Discrete GPU for AI and HPC

305 AMD CDNA™ GPU Compute Units, 1,216 Matrix Cores, and 19,456 Stream Cores


192GB HBM3

With 5.3 TB/s peak theoretical
memory bandwidth²

Up to

163.4 TFLOPS

peak theoretical single precision
(FP32) HPC performance¹



Up to

81.7 TFLOPS

peak theoretical double-precision
(FP64) Performance¹

Up to

5.2 PFLOPS

peak theoretical 8-bit
precision (FP8) with sparsity
for AI Performance³

AMD CDNA™ 3 Architecture

AMD CDNA™ 3 is the dedicated compute architecture underlying AMD Instinct™ MI300 Series accelerators. It features advanced packaging with chiplet technologies—designed to reduce data movement overhead and enhance power efficiency.

4th Gen Infinity Architecture

Next-gen AMD Infinity Architecture, along with AMD Infinity Fabric™ technology, enables high-throughput unification of AMD GPU chiplet technologies with stacked HBM3 memory in single devices and across multi-device platforms.

AMD Instinct™ MI300X Platform

The AMD Instinct™ MI300X is offered in an industry-standard 8 GPU UBB 2.0 platform with a high-bandwidth AMD Infinity Fabric™ and PCIe® Gen 5 enabling superior performance, reliable and fast system deployments in existing infrastructures.

AMD ROCm™ Ecosystem

With the ROCm open software platform, HPC and AI communities can now gain access to an array of different open compute programming languages, frameworks, compilers, libraries and tools.