




Home > Products > Instinct > MI200 Series

AMD Instinct™ MI200 Series Accelerators

Powering AI and HPC at Scale




ON THIS PAGE ▾

Feedback



Driving Exceptional HPC and AI Performance

AMD Instinct™ MI200 Series accelerators enable exceptional performance for the data center, at any scale—from single-server solutions up to the world's **largest**, Exascale-class supercomputers.¹



Under the Hood

AMD Instinct accelerators are built on AMD CDNA™ architecture, which offers Matrix Core Technologies and support for a broad range of precision capabilities—from the highly efficient INT8 and FP8 to the most demanding FP64 for HPC.

[Learn More >](#)

Meet the Series

Explore AMD Instinct MI250X, MI250, and MI210 accelerators.

AMD Instinct MI250X Accelerators

AMD Instinct MI250X accelerators power some of the world's top supercomputers.

[View Specs >](#)

220 CUs

220 GPU Compute Units

128 GB

128 GB HBM2e Memory

3.2 TB/s



3.2 TB/s Peak Memory Bandwidth

400 GB/s



400 GB/s Peak Aggregate Infinity Fabric

Specs Comparisons

Up to 383 TFLOPS peak theoretical half-precision (FP16) performance with up to 1.6X more memory capacity and bandwidth than competitive GPUs for the most demanding AI workloads^{2,3,4}



Delivers up to 4X advantage vs. competitive GPUs, offering up to 47.9 TFLOPs FP64 and up to 95.7 TFLOPs FP64 Matrix peak theoretical performance.²



*The TF32 data format is not IEEE compliant and not included in this comparison.

AMD Instinct MI250 Accelerators

AMD Instinct MI250 accelerators drive exceptional HPC and AI performance for enterprise, research, and academia-related use cases.

[View Specs >](#)

208 CUs

208 GPU Compute Units

128 GB

128 GB HBM2e Memory

3.2 TB/s

3.2 TB/s Peak Memory Bandwidth

100 GB/s

100 GB/s Peak Infinity Fabric™ Link Bandwidth

Performance Benchmarks

[View All](#)

AMD Instinct MI210 Accelerators

AMD Instinct MI210 accelerators power enterprise, research, and academic HPC and AI workloads for single-server solutions and more.

[View Specs >](#)

104 CUs

104 GPU Compute Units

64 GB

64 GB HBM2e Memory

1.6 TB/s

1.6 TB/s Peak Memory Bandwidth

100 GB/s

100 GB/s Peak Infinity Fabric™ Link Bandwidth

Performance Benchmarks

[View All](#)

AMD ROCm™ Software


AMD ROCm™ software includes a broad set of programming models, tools, compilers, libraries, and runtimes for optimizing AI models and HPC workloads targeting AMD Instinct accelerators.

[About ROCm](#)

[ROCm Developer Hub](#)

Case Studies

[View All](#) >



DriveNets Advances Open Backend Networking for AI with AMD

DriveNets worked with AMD to deliver an open, full stack backend networking solution for AI advanced by AMD Instinct™ MI350 Series GPUs.

July 26, 2026

WWT Helps Customers Make the Right Decision Faster with AMD

WWT customers keep pace with rapid technology change and make the right decision faster, with AMD EPYC™ Server CPUs and AMD Instinct™ GPUs.

July 26, 2026

AT&T Achieves Outstanding Efficiency for AI Training with AMD

AT&T trains telecom AI on AMD Instinct™ MI355X GPUs, running a full model on a single GPU to power inference for millions of customers.

July 02, 2026

Find Solutions

Find a partner offering AMD Instinct accelerator-based solutions.

AMD Instinct Solutions

Contact Sales >

Resources

Blogs

Read the latest blogs on AMD Instinct accelerators.

[Read Blogs >](#)

Case Studies

Read the latest case studies on how customers are leveraging AMD Instinct accelerators.

[Read Now >](#)

Documentation

Find solution briefs, white papers, programmer references, and more documentation for AMD Instinct accelerators.

[Find Docs >](#)

AMD Instinct Education & Research Initiative (AIER) 2.0

AIER 2.0 offers scientists, researchers, and academics easier access to AMD Instinct accelerators, AMD ROCm software, and dedicated support.

[Explore AIER 2.0 >](#)

Stay Informed

Sign up to receive the latest data center news and server content.

[Sign Up](#)



