

March 22, 2022



AMD Instinct™ Expands Ecosystem and Delivers Exascale-Class Technology for HPC and AI Applications

– Powered by AMD CDNA™ 2 architecture and AMD ROCm™5, new AMD Instinct MI210 GPUs accelerating insights and discovery for mainstream users –

SANTA CLARA, Calif., March 22, 2022 (GLOBE NEWSWIRE) -- [AMD](#) (NASDAQ: AMD) today announced the availability of the [AMD Instinct™ ecosystem](#) with expanded system support from partners including ASUS, Dell Technologies, Gigabyte, HPE, Lenovo and Supermicro, the new [AMD Instinct™ MI210 accelerator](#) and the robust capabilities of [ROCm™ 5 software](#). Altogether, the AMD Instinct and ROCm ecosystem is offering exascale-class technology to a broad base of HPC and AI customers, addressing the growing demand for compute-accelerated data center workloads and reducing the time to insights and discovery.

“With twice the platforms available compared to our previous generation accelerators, growing customer adoption across HPC and AI applications, and new support from commercial ISVs in key workloads, we’re continuing to drive adoption of the AMD Instinct MI200 accelerators and ROCm 5 software ecosystem,” said Brad McCredie, corporate vice president, Data Center GPU and Accelerated Processing, AMD. “Now with the availability of the AMD Instinct MI210 accelerator to the MI200 family, our customers can choose the accelerator that works best for their workloads, whether they need leading-edge accelerated processing for large scale HPC and AI workloads, or if they want access to exascale-class technology in a commercial format.”

“The Lumi supercomputer powered by AMD EPYC processors and AMD Instinct MI200 accelerators will provide a generational leap in performance for large-scale simulations and modeling as well as AI and deep learning workloads to solve some of the biggest questions in research,” said Pekka Manninen, Director of the LUMI Leadership and Computing Facility, CSC. “We’ve utilized AMD Instinct MI210 accelerators to get hands on experience with the Instinct MI200 family, preparing our scientists to tackle the many challenging and complex projects they will run once Lumi is fully deployed.”

Powering The Future of HPC and AI

The AMD Instinct MI200 series accelerators are designed to power discoveries in exascale systems, enabling researchers, scientists and engineers to tackle our most pressing challenges, from climate change to vaccine research. The [AMD Instinct MI210 accelerators](#) specifically enable exascale-class technologies for customers who need fantastic HPC and AI performance in a PCIe® format. Powered by the AMD CDNA™ 2 architecture, AMD Instinct MI210 accelerators extend AMD performance leadership in double precision (FP64) compute on PCIe form factor cards¹. They also deliver a robust solution for accelerated deep learning training offering a broad range of mixed-precision capabilities based on the AMD Matrix Core Technology.

Driving the ROCm Adoption

An open software platform that allows researchers, scientists and engineers to tap the power of AMD Instinct accelerators to drive scientific discoveries, the AMD ROCm platform is built on the foundation of numerous applications and libraries powering top HPC and AI applications.

With ROCm 5, AMD extends its software platform by adding new hardware support for the AMD Instinct MI200 series accelerators and the AMD Radeon™ PRO W6800 professional graphics card, plus Red Hat® Enterprise Linux® 8.5 support, increasing accessibility of ROCm for developers and enabling outstanding performance across key workloads.

Additionally, through the [AMD Infinity Hub](#), the central location for open-source applications that are ported and optimized on AMD GPUs, end-users can easily find, download and install containerized HPC apps and ML frameworks. AMD Infinity Hub application containers are designed to reduce the traditionally difficult issue of obtaining and installing software releases while allowing users to learn based on shared experiences and problem-solving opportunities.

Expanding Partner and Customer Ecosystem

As more purpose-built applications are optimized to work with ROCm and AMD Instinct accelerators, AMD continues to grow its software ecosystem with the addition of commercial ISVs, including Ansys®, Cascade Technologies, and TempoQuest. These ISVs provide applications for accelerated workloads including Computational Fluid Dynamics (CFD), weather forecasting, Computer Aided Engineering (CAE) and more. These updates are on top of existing application support in ROCm which includes HPC, AI and Machine Learning applications, AMBER, Chroma, CP2K, GRID, GROMACs, LAAMPS, MILC, Mini-HAAC, NAMD, NAMD 3.0, ONNX-RT, OpenMM, PyTorch, RELION, SPECfem3d Cartesian, SPECfem3d Globe, and TensorFlow.

AMD is also enabling partners like ASUS, Dell Technologies, Gigabyte, HPE, Lenovo, Supermicro, and System Integrators including Colfax, Exxact, KOI Computers, Nor-Tech, Penguin and Symmetric to offer differentiated solutions to address next generation computing challenges. Supercomputing customers are already taking advantage of the benefits offered via these new customer wins including the Frontier installation at [Oak Ridge National Laboratory](#), [KTH/Dardel](#), [CSC/LUMI](#) and [Cines/Adastra](#).

Enabling Access for Customers and Partners

The [AMD Accelerator Cloud](#) offers customers an environment to remotely access and evaluate AMD Instinct accelerators and AMD ROCm software. Whether it's porting legacy code, benchmarking an application or testing multi-GPU or multi-node scaling, the AMD Accelerator Cloud gives prospective customers and partners quick and easy access to the latest GPUs and software. The AMD Accelerator Cloud is also used to power various events such as hackathons and ROCm training sessions offered to both existing and prospective customers, allowing developers to hone their skills and learn how to get the most out of AMD Instinct accelerators.

MI200 Series Specifications

Model	Compute Units	Stream Processors	FP64 FP32 Vector (Peak)	FP64 FP32 Matrix (Peak)	FP16 bf16 (Peak)	INT8 (Peak)	HBM2e ECC Memory	Memory Bandwidth	Form Factor
AMD Instinct MI210	104	6,656	Up to 22.6 TF	Up to 45.3 TF	Up to 181.0 TF	Up to 181.0 TOPS	64GB	Up to 1.6 TB/sec	PCIe®
AMD Instinct MI250	208	13,312	Up to 45.3 TF	Up to 90.5 TF	Up to 362.1 TF	Up to 362.1 TOPS	128GB	3.2 TB/sec	OCP Accelerator Module (OAM)
AMD Instinct MI250x	220	14,080	Up to 47.9 TF	Up to 95.7 TF	Up To 383.0 TF	Up to 383.0 TOPS	128GB	3.2 TB/sec	OCP Accelerator Module (OAM)

Supporting Resources

- Learn more about [AMD Instinct™ MI200 series accelerators](#)
- Learn more about [AMD Instinct™ MI210 accelerator](#)
- Learn more about [AMD CDNA™ 2 architecture](#)
- Learn more about [AMD ROCm™ open software platform](#)
- Learn more about [AMD Infinity Hub](#)
- Learn more about [AMD Instinct ecosystem partners](#)
- Follow AMD on [Twitter](#)

About AMD

For more than 50 years AMD has driven innovation in high-performance computing, graphics and visualization technologies. Billions of people, leading Fortune 500 businesses and cutting-edge scientific research institutions around the world rely on AMD technology daily to improve how they live, work and play. AMD employees are focused on building leadership high-performance and adaptive products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit the AMD (NASDAQ: AMD) [website](#), [blog](#), [LinkedIn](#) and [Twitter](#) pages.

©2022 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD CDNA, AMD Instinct, Radeon, ROCm and combinations thereof are trademarks of Advanced Micro Devices, Inc. PCIe is a registered trademark of PCI-SIG Corporation. Red Hat, Red Hat Enterprise Linux, and the Red Hat logo, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the U.S. and other countries. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

¹ MI200-41 - Calculations conducted by AMD Performance Labs as of Jan 14, 2022, for the AMD Instinct™ MI210 (64GB HBM2e PCIe® card) accelerator at 1,700 MHz peak boost engine clock resulted in 45.3 TFLOPS peak theoretical double precision (FP64 Matrix), 22.6 TFLOPS peak theoretical double precision (FP64), and 181.0 TFLOPS peak theoretical Bfloat16 format precision (BF16), floating-point performance.

Calculations conducted by AMD Performance Labs as of Sep 18, 2020 for the AMD Instinct™ MI100 (32GB HBM2 PCIe® card) accelerator at 1,502 MHz peak boost engine clock resulted in 11.54 TFLOPS peak theoretical double precision (FP64), and 184.6 TFLOPS peak theoretical half precision (FP16), floating-point performance.

Published results on the NVidia Ampere A100 (80GB) GPU accelerator, boost engine clock of 1410 MHz, resulted in 19.5 TFLOPS peak double precision tensor cores (FP64 Tensor Core), 9.7 TFLOPS peak double precision (FP64) and 39 TFLOPS peak Bfloat16 format precision (BF16), theoretical floating-point performance. The TF32 data format is not IEEE compliant and not included in this comparison.

<https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/nvidia-ampere-architecture-whitepaper.pdf>, page 15, Table 1.

Contacts:

Aaron Grabein

AMD Communications

+1 512-602-8950

Aaron.Grabein@amd.com

Laura Graves

AMD Investor Relations

+1 408-749-5467

Laura.Graves@amd.com



Source: Advanced Micro Devices, Inc.