

November 8, 2021



AMD Unveils Workload-Tailored Innovations and Products at The Accelerated Data Center Premiere

— AMD launches AMD Instinct™ MI200 series accelerators, previews 3rd Gen AMD EPYC™ processors with AMD 3D V-Cache, and provides new details on expanded set of next-generation EPYC™ processors powered by “Zen 4” and “Zen 4c” CPU cores —

— Meta adopts EPYC™ CPUs for its data center —

SANTA CLARA, Calif., Nov. 08, 2021 (GLOBE NEWSWIRE) -- AMD (NASDAQ: AMD) held the virtual [Accelerated Data Center Premiere](#), launching the new AMD Instinct™ MI200 series accelerators, the world’s fastest accelerator for high performance computing (HPC) and artificial intelligence (AI) workloadsⁱ, and provided a preview of the innovative 3rd Gen AMD EPYC™ processors with AMD 3D V-Cache. AMD also revealed new information about its next generation “Zen 4” processor core and announced the new “Zen 4c” processor core, both of which will power future AMD server processors and are designed to extend the company’s leadership products for the data center.

“We are in a high-performance computing megacycle that is driving demand for more compute to power the services and devices that impact every aspect of our daily lives,” said Dr. Lisa Su, president and CEO, AMD. “We are building significant momentum in the data center with our leadership product portfolio, including Meta’s adoption of AMD EPYC and the buildout of Frontier, the first U.S. exascale supercomputer which will be powered by EPYC and AMD Instinct processors. In addition, today we announced a breadth of new products that build on that momentum in next-generation EPYC processors with new innovations in design, leadership, 3D packaging technology, and 5 nm high-performance manufacturing to further extend our leadership for cloud, enterprise and HPC customers.”

Meta Adopts EPYC CPUs [03:09 – 05:29]

AMD announced Meta is the latest major hyperscale cloud company that has adopted AMD EPYC CPUs. AMD and Meta worked together to define an open, cloud-scale, single-socket server designed for performance and power efficiency, based on the 3rd Gen EPYC processor. Further details will be discussed at the Open Compute Global Summit later this week.

Advanced Packaging Driving Data Center Performance [05:35 – 18:00]

AMD previewed the use of innovative 3D chiplet packaging technology in the data center with the first server CPU using high-performance 3D die stacking. The 3rd Gen AMD EPYC processors with AMD 3D V-Cache, codenamed “Milan-X,” represents an innovative step forward in CPU design and packaging, and will offer a 50% average performance uplift across targeted technical computing workloadsⁱⁱ.

- 3rd Gen EPYC with AMD 3D V-Cache will offer the same capabilities and features as the 3rd Gen EPYC processors and they will be drop-in compatible with a BIOS upgrade, delivering easy adoption and performance enhancements.
- Microsoft Azure HPC virtual machines featuring 3rd Gen EPYC with AMD 3D V-Cache are available today in Private Preview, with broad rollout in the coming weeks. More information on performance and availability is available [here](#).
- 3rd Gen EPYC CPUs with AMD 3D V-Cache will launch in Q1 2022. Partners including Cisco, Dell Technologies, Lenovo, HPE and Supermicro are planning to offer server solutions with these processors.

Delivering Exascale Class Performance for Accelerated Computing [18:02 – 31:50]

AMD [launched the AMD Instinct MI200 series accelerators](#). Based on the AMD CDNA™ 2 architecture, the MI200 series accelerators are the most advanced accelerators in the worldⁱⁱⁱ and compared to the competition provide up to 4.9x the peak performance for HPC workloads^{iv} and 1.2X the peak flops of mixed precision performance for leadership AI training, helping to fuel the convergence of HPC and AI.

- Utilized in the Frontier supercomputer at [Oak Ridge National Laboratory](#), the HPC and AI performance capabilities in AMD Instinct MI200 series accelerators will be key in enabling researchers and scientists to accelerate their time to science and discovery.

“Zen 4” Powered Data Center, Designed for Leadership Performance [31:52 – 36:22]

AMD provided new details on the expanded next generation AMD EPYC processors codenamed “Genoa” and “Bergamo.”

- “Genoa” is expected to be the world’s highest performance processor for general purpose computing. It will have up to 96 high-performance “Zen 4” cores produced on optimized 5nm technology, and will support the next generation of memory and I/O technologies with DDR5 and PCIe® 5. “Genoa” will also include support for CXL, enabling significant memory expansion capabilities for data center applications. “Genoa” is on track for production and launch in 2022.
- “Bergamo” is a high-core count CPU, tailor made for cloud native applications, featuring 128 high performance “Zen 4c” cores. AMD optimized the new “Zen 4c” core for cloud-native computing, tuning the core design for density and increased power efficiency to enable higher core count processors with breakthrough performance per-socket. “Bergamo” comes with all the same software and security features and is socket compatible with “Genoa.” “Bergamo” is on track to ship in the first half of 2023.

You can watch the [full video here](#) and learn more about all of the products discussed [during the event here](#).

Supporting Resources

- Learn more about the [AMD Instinct™ accelerators](#)
- Learn more about [AMD EPYC™ processors](#)
- Learn more about the [Oak Ridge National Laboratory’s Frontier supercomputer](#)
- Become a fan of AMD on [Facebook](#)
- Follow AMD on [Twitter](#)
- Connect with AMD On [LinkedIn](#)

About AMD

For more than 50 years AMD has driven innovation in high-performance computing, graphics and visualization technologies — the building blocks for gaming, immersive platforms and the datacenter. Hundreds of millions of consumers, leading Fortune 500 businesses and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work and play. AMD employees around the world are focused on building great 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](#), [Facebook](#), [LinkedIn](#) and [Twitter](#) pages.

AMD, the AMD Arrow logo, EPYC, AMD Instinct, Infinity Fabric and combinations thereof are trademarks of Advanced Micro Devices, Inc. PCIe is a registered trademark of PCI-SIG Corporation. Other names are for informational purposes only and may be trademarks of their respective owners.

This press release contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) such as the features, functionality, performance, availability, timing and expected benefits of AMD products including the 3rd Gen AMD EPYC™ processors with AMD 3D V-Cache and AMD Instinct™ MI200 series accelerators as well as the “Zen 4” data center roadmap, which are made pursuant to the Safe Harbor provisions of the Private Securities Litigation Reform Act of 1995. Forward-looking statements are commonly identified by words such as “would,” “may,” “expects,” “believes,” “plans,” “intends,” “projects” and other terms with similar meaning. Investors are cautioned that the forward-looking statements in this press release are based on current beliefs, assumptions and expectations, speak only as of the date of this press release and involve risks and uncertainties that could cause actual results to differ materially from current expectations. Such statements are subject to certain known and unknown risks and uncertainties, many of which are difficult to predict and generally beyond AMD’s control, that could cause actual results and other future events to differ materially from those expressed in, or implied or projected by, the forward-looking information and statements. Material factors that could cause actual results to differ materially from current expectations include, without limitation, the following: Intel Corporation’s dominance of the microprocessor market and its aggressive business practices; global economic uncertainty; loss of a significant customer; impact of the COVID-19 pandemic on AMD’s business, financial condition and results of operations; competitive markets in which AMD’s products are sold; market conditions of the industries in which AMD products are sold; cyclical nature of the semiconductor industry; quarterly and seasonal sales patterns; AMD’s ability to adequately protect its technology or other intellectual property; unfavorable currency exchange rate fluctuations; ability of third party manufacturers to manufacture AMD’s products on a timely basis in sufficient quantities and using competitive technologies; availability of essential equipment, materials, substrates or manufacturing processes; ability to achieve expected manufacturing yields for AMD’s products; AMD’s ability to introduce products on a timely basis with expected features and performance levels; AMD’s ability to generate revenue from its semi-custom SoC products; potential security vulnerabilities; potential security incidents including IT outages, data loss, data breaches and cyber-attacks; uncertainties involving the ordering and shipment of AMD’s products; AMD’s reliance on third-party intellectual property to design and introduce new products in a timely manner; AMD’s reliance on third-party companies for design, manufacture and supply of motherboards, software and other computer platform components; AMD’s reliance on Microsoft and other software vendors’ support to design and develop software to run

on AMD's products; AMD's reliance on third-party distributors and add-in-board partners; impact of modification or interruption of AMD's internal business processes and information systems; compatibility of AMD's products with some or all industry-standard software and hardware; costs related to defective products; efficiency of AMD's supply chain; AMD's ability to rely on third party supply-chain logistics functions; AMD's ability to effectively control sales of its products on the gray market; impact of government actions and regulations such as export administration regulations, tariffs and trade protection measures; AMD's ability to realize its deferred tax assets; potential tax liabilities; current and future claims and litigation; impact of environmental laws, conflict minerals-related provisions and other laws or regulations; impact of acquisitions, joint ventures and/or investments on AMD's business, including the announced acquisition of Xilinx, and ability to integrate acquired businesses; AMD's ability to complete the Xilinx merger; impact of the announcement and pendency of the Xilinx merger on AMD's business; impact of any impairment of the combined company's assets on the combined company's financial position and results of operation; restrictions imposed by agreements governing AMD's notes and the revolving credit facility; AMD's indebtedness; AMD's ability to generate sufficient cash to meet its working capital requirements or generate sufficient revenue and operating cash flow to make all of its planned R&D or strategic investments; political, legal, economic risks and natural disasters; future impairments of goodwill and technology license purchases; AMD's ability to attract and retain qualified personnel; AMD's stock price volatility; and worldwide political conditions. Investors are urged to review in detail the risks and uncertainties in AMD's Securities and Exchange Commission filings, including but not limited to AMD's most recent reports on Forms 10-K and 10-Q.

ⁱ MI200-01: World's fastest data center GPU is the AMD Instinct™ MI250X. Calculations conducted by AMD Performance Labs as of Sep 15, 2021, for the AMD Instinct™ MI250X (128GB HBM2e OAM module) accelerator at 1,700 MHz peak boost engine clock resulted in 95.7 TFLOPS peak theoretical double precision (FP64 Matrix), 47.9 TFLOPS peak theoretical double precision (FP64), 95.7 TFLOPS peak theoretical single precision matrix (FP32 Matrix), 47.9 TFLOPS peak theoretical single precision (FP32), 383.0 TFLOPS peak theoretical half precision (FP16), and 383.0 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), 46.1 TFLOPS peak theoretical single precision matrix (FP32), 23.1 TFLOPS peak theoretical single precision (FP32), 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), 19.5 TFLOPS peak single precision (FP32), 78 TFLOPS peak half precision (FP16), 312 TFLOPS peak half precision (FP16 Tensor Flow), 39 TFLOPS peak Bfloat 16 (BF16), 312 TFLOPS peak Bfloat16 format precision (BF16 Tensor Flow), 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.

ⁱⁱ MLNX-021R: AMD internal testing as of 09/27/2021 on 2x 64C 3rd Gen EPYC with AMD 3D V-Cache (Milan-X) compared to 2x 64C AMD 3rd Gen EPYC 7763 CPUs using cumulative average of each of the following benchmark's maximum test result score: ANSYS® Fluent® 2021.1, ANSYS® CFX® 2021.R2, and Altair Radioss 2021. Results may

vary.

iii MI200-31: As of October 20th, 2021, the AMD Instinct™ MI200 series accelerators are the “Most advanced server accelerators (GPUs) for data center,” defined as the only server accelerators to use the advanced 6nm manufacturing technology on a server. AMD on 6nm for AMD Instinct MI200 series server accelerators. Nvidia on 7nm for Nvidia Ampere A100 GPU. <https://developer.nvidia.com/blog/nvidia-ampere-architecture-in-depth/>

iv MI200-02: Calculations conducted by AMD Performance Labs as of Sep 15, 2021, for the AMD Instinct™ MI250X accelerator (128GB HBM2e OAM module) at 1,700 MHz peak boost engine clock resulted in 95.7 TFLOPS peak double precision matrix (FP64 Matrix) theoretical, floating-point performance. Published results on the NVidia Ampere A100 (80GB) GPU accelerator resulted in 19.5 TFLOPS peak double precision (FP64 Tensor Core) theoretical, floating-point performance. Results found at: <https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/nvidia-ampere-architecture-whitepaper.pdf>, page 15, Table 1.

Contact:

Aaron Grabein
AMD Communications
(512) 602-8950
aaron.grabein@amd.com

Laura Graves
AMD Investor Relations
(408) 749-5467
laura.graves@amd.com



Source: Advanced Micro Devices, Inc.