November 16, 2020

AMD EPYC™ Processors and New AMD Instinct™ MI100 Accelerator Redefine Performance for HPC and Scientific Research

—Support for next generation AMD EPYC processors, codenamed 'Milan,' in new HBv3 Virtual Machines announced by Microsoft—

—AMD Instinct™ MI100 accelerator is first accelerator to use new AMD CDNA architecture dedicated to high-performance computing (HPC) workloads—

SANTA CLARA, Calif., Nov. 16, 2020 (GLOBE NEWSWIRE) -- During this year's SC20 virtual tradeshow, AMD (NASDAQ: AMD) is showcasing its leadership in the high performance computing (HPC) industry. It launched the new AMD Instinct™ MI100 accelerator with ROCm™ 4.0 open ecosystem support and showcased a growing list of AMD EPYC™ CPU and AMD Instinct accelerator based deployments, and highlighted its collaboration with Microsoft Azure for HPC in the cloud. AMD also remains on track to begin volume shipments of the 3rd Gen EPYC processors with "Zen 3" core to select HPC and cloud customers this quarter in advance of the expected public launch in Q1 2021, aligned with OEM availability.

The new AMD Instinct™ MI100 accelerator, is the world's fastest HPC GPU accelerator for scientific workloads and the first to surpass the 10 teraflops (FP64) performance barrier [1]. Built on the new AMD CDNA architecture, the AMD Instinct MI100 GPU enables a new class of accelerated systems for HPC and AI when paired with 2nd Gen AMD EPYC processors. Supported by new accelerated compute platforms from Dell, HPE, Gigabyte and Supermicro, the MI100, combined with AMD EPYC CPUs and ROCm 4.0 software, is designed to propel new discoveries ahead of the exascale era.

"No two customers are the same in HPC, and AMD is providing a path to today's most advanced technologies and capabilities that are critical to support their HPC work, from small clusters on premise, to virtual machines in the cloud, all the way to exascale supercomputers," said Forrest Norrod, senior vice president and general manager, Data Center and Embedded Solutions Business Group, AMD. "Combining AMD EPYC processors and Instinct accelerators with critical application software and development tools enables AMD to deliver leadership performance for HPC workloads."

AMD and Microsoft Azure Power HPC In the Cloud

Azure is using 2nd Gen AMD EPYC processors to power its HBv2 virtual machines (VMs) for HPC workloads. These VMs offer up to 2x the performance of first-generation HB-series virtual machines[2], can support up to 80,000 cores for MPI jobs[3], and take advantage of 2nd

Gen AMD EPYC processors' up to 45% more memory bandwidth than comparable x86 alternatives[4].

HBv2 VMs are used by numerous customers including The University of Illinois at Urbana-Champaign's Beckman Institute for Advanced Science & Technology which used 86,400 cores to model a plant virus that previously required a leadership class supercomputer and the U.S. Navy which rapidly deploys and scales enhanced weather and ocean pattern predictions on demand. HBv2 powered by 2nd Gen AMD EPYC processors also provides the bulk of the CPU compute power for the OpenAI environment Microsoft announced earlier this year.

AMD EPYC processors have also helped HBv2 reach new cloud HPC milestones, such as a new record for Cloud MPI scaling results with NAMD, Top 20 results on the Graph500, and the first 1 terabyte/sec cloud HPC parallel filesystem. Across these and other application benchmarks, HBv2 is delivering 12x higher scaling than found elsewhere on the public cloud.

Adding on to its existing HBv2 HPC virtual machine powered by 2nd Gen AMD EPYC processors, Azure announced it will utilize next generation AMD EPYC processors, codenamed 'Milan', for future HB-series VM products for HPC.

You can see more about the AMD and Azure collaboration in this video with Jason Zander of Azure and Lisa Su of AMD.

AMD Continues to Be the Choice for HPC

AMD EPYC processors and Instinct accelerators have the performance and capabilities to support numerous HPC workloads across a variety of implementations. From small clusters at research centers, to commercial HPC, to off premise and in the cloud, to exascale computing, AMD continues to provide performance and choice for HPC solutions.

Hewlett Packard Enterprise (HPE), CSC Finland and EuroHPC recently introduced a new pre-exascale system, LUMI. Based on the HPE Cray EX supercomputer architecture, LUMI will use next generation AMD EPYC CPUs and Instinct accelerators and is expected to provide a peak performance of 552 petaflops when it comes online in 2021, making it one of the fastest supercomputers in the world.

Beyond LUMI, AMD powered HPC systems continue to grow in volume. Since SC19, there have been more than 15 supercomputing systems announced using AMD EPYC CPUs, Instinct GPUs, or both. A highlight of the systems includes

- Chicoma – Los Alamos National Laboratory – this system is based on the HPE Cray EX supercomputer architecture and uses 2nd Gen AMD EPYC CPUs, combined with 300 terabytes of system memory for COVID-19 research,
- Corona – Lawrence Livermore National Laboratory – this system was recently upgraded with funding from the Coronavirus Aid, Relief and Economic Security (CARES) Act, adding nearly 1,000 AMD Instinct MI50 accelerators, pushing peak performance to more than 11 petaFLOPS,
- Mammoth – Lawrence Livermore National Laboratory – the "big memory" cluster uses 2nd Gen AMD EPYC Processors to perform genomics analysis and graph analytics required by scientists working on COVID-19.

- Northern Data – a distributed computing system in Europe that is using AMD EPYC CPUs and Instinct accelerators for large scale HPC applications such as rendering, artificial intelligence and deep learning,
- Pawsey Supercomputing Centre – Using the HPE Cray EX supercomputer architecture and future AMD EPYC CPUs and AMD Instinct accelerators, the supercomputer at Pawsey will be Australia's most powerful supercomputer.

In addition, AMD is also powering the following supercomputers: Anvil and Bell – Purdue University, Big Red 200 – Indiana University, Bridges 2 – Pittsburgh Supercomputing Center, CERN, European Centre for Medium-Range Weather Forecasts, Expanse – San Diego Supercomputer Center, Goethe University Frankfurt, IT4Innovations National Supercomputing Center, Jetstream 2 – Indiana University, Mahti – CSC, Mangi – University of Minnesota, National Oceanic and Atmospheric Administration, Red Raider – Texas Tech University, TinkerCliffs – Virginia Tech.

"With the Expanse supercomputer, our goal is to give scientists and researchers cloud-like access to a high-performance machine that can handle everything from astrophysics to zoology," said Michael Norman, Director of the San Diego Supercomputer Center. "The 2nd Gen AMD EPYC processors have helped us achieve fantastic performance with Expanse, enabling our researchers to do more science than before. We also have a great collaboration with AMD and have worked together to create a forum for AMD HPC customers to share experiences, information and more, to better benefit HPC research."

Paving the Path to Exascale Computing

To help researchers start on the path to exascale, AMD has provided Oak Ridge National Labs access to the new AMD Instinct MI100 accelerator, which delivers a giant leap in compute and interconnect performance. The Instinct MI100 accelerator enables a new class of accelerated systems and delivers true heterogeneous compute capabilities from AMD for HPC and AI. Designed to complement the 2nd Gen AMD EPYC processors, and built on the AMD Infinity Architecture, the AMD Instinct MI100 delivers true heterogeneous compute capabilities from AMD for HPC and AI.

"Frontier, powered by AMD, represents a huge increase in computational power compared to today's systems. It's going to allow scientists to answer questions that we didn't have the answer to before," said Bronson Messer, director of science, Oak Ridge Leadership Computing Facility. "The ability to run molecular simulations that aren't just a few million atoms, but a few billion atoms, provides a more realistic representation of the science, and to be able to do that as a matter of course and over and over again will lead to a significant amount of important discoveries."

AMD continues to provide the performance, capabilities and scale needed to power current and future HPC workloads, no matter if they are helping students at a research center, improving aerodynamic efficiency for an auto manufacturer, or providing valuable insights for critical medical breakthroughs. Read more about the AMD presence at SC20 and its HPC capabilities here.

Supporting Resources

- Read quotes from AMD HPC customers about their experience with AMD products
- Learn more about AMD Instinct™ Accelerators

- Learn more about AMD CDNA Architecture
- Learn more about AMD HPC Solutions here
- Learn more about the AMD 2nd Gen EPYC™ Processor
- Follow AMD on Twitter

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 data center. 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 and Twitter pages.

CAUTIONARY STATEMENT

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 next generation AMD EPYC™ Processors and AMD Instinct™ MI100 Accelerator and the expected timing and benefits of AMD partnerships, 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; the ability of third party manufacturers to manufacture AMD's products on a timely basis in sufficient quantities and using competitive technologies; expected manufacturing yields for AMD’s products; the availability of essential equipment, materials or manufacturing processes; AMD's ability to introduce products on a timely basis with features and performance levels that provide value to its customers; global economic uncertainty; the loss of a significant customer; AMD's ability to generate revenue from its semi-custom SoC products; the impact of the COVID-19 pandemic on AMD’s business, financial condition and results of operations; political, legal, economic risks and natural disasters; the impact of government actions and regulations such as export administration regulations, tariffs and trade protection measures; the impact of acquisitions, joint ventures and/or investments on AMD's business, including the announced acquisition of Xilinx, and the failure to integrate acquired businesses; AMD’s ability to complete the Xilinx merger; the impact of the announcement and pendency of the Xilinx merger on AMD’s business; potential security vulnerabilities; potential IT outages, data loss, data breaches and cyber-attacks; uncertainties involving the ordering and shipment of AMD’s products; quarterly and seasonal sales patterns; the restrictions imposed by agreements governing AMD’s notes and the

revolving credit facility; the competitive markets in which AMD's products are sold; market conditions of the industries in which AMD products are sold; 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 the design, manufacture and supply of motherboards, software and other computer platform components; AMD's reliance on Microsoft Corporation 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; the potential dilutive effect if the 2.125% Convertible Senior Notes due 2026 are converted; future impairments of goodwill and technology license purchases; AMD's ability to attract and retain qualified personnel; AMD's ability to generate sufficient revenue and operating cash flow or obtain external financing for research and development or other strategic investments; AMD's indebtedness; AMD's ability to generate sufficient cash to service its debt obligations or meet its working capital requirements; AMD's ability to repurchase its outstanding debt in the event of a change of control; the cyclical nature of the semiconductor industry; the 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; the efficiency of AMD's supply chain; AMD's ability to rely on third party supply-chain logistics functions; AMD's stock price volatility; worldwide political conditions; unfavorable currency exchange rate fluctuations; AMD's ability to effectively control the sales of its products on the gray market; AMD's ability to adequately protect its technology or other intellectual property; current and future claims and litigation; potential tax liabilities; and the impact of environmental laws, conflict minerals-related provisions and other laws or regulations. 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 Quarterly Report on Form 10-Q for the quarter ended September 26, 2020.

---

[1] 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 double precision (FP64), 46.1 TFLOPS peak single precision matrix (FP32), 23.1 TFLOPS peak single precision (FP32), 184.6 TFLOPS peak half precision (FP16) peak theoretical, floating-point performance. Published results on the NVidia Ampere A100 (40GB) GPU accelerator resulted in 9.7 TFLOPS peak double precision (FP64). 19.5 TFLOPS peak single precision (FP32), 78 TFLOPS peak half precision (FP16) theoretical, floating-point performance. Server manufacturers may vary configuration offerings yielding different results. MI100-03

[2] Source: https://azure.microsoft.com/en-us/blog/introducing-the-new-hbv2-azure-virtual-machines-for-high-performance-computing/

[3] Source: https://azure.microsoft.com/en-us/blog/azure-hbv2-virtual-machines-eclipse-80000-cores-for-mpi-hpc/

[4] AMD EPYC™ 7002 Series processors have 45% more memory bandwidth than Intel Scalable processors in the same class.

Contacts:
Gary Silcott
AMD Communications
+1 512-602-0889
Gary.Silcott@amd.com

Laura Graves
AMD Investor Relations
+1 408-749-5467
Laura.Graves@amd.com

Source: Advanced Micro Devices, Inc.