

November 18, 2019



AMD Delivers Best-in-Class Performance from Supercomputers to HPC in the Cloud at SC19

— San Diego Supercomputer Center, Swiss ETH, AWS and others leverage record breaking performance of 2nd Gen AMD EPYC™ processors —

— AMD CTO Mark Papermaster to detail “The Last Mile to Exascale” —

— 2nd Gen AMD EPYC processors enter latest Top 500 list with Atos and GENCI —

DENVER, Nov. 18, 2019 (GLOBE NEWSWIRE) -- At SC19, the premier annual event for supercomputing, [AMD](#) (NASDAQ: AMD) is extending its performance lead in high-performance computing (HPC) with a range of new customer wins in top research systems worldwide, new platforms supporting AMD EPYC™ processors and Radeon Instinct™ accelerators and the newly announced ROCm 3.0, bringing support for new compilers and HPC applications.

“AMD enters SC19 as the processor provider for the upcoming Frontier supercomputer, expected to be the highest performance supercomputer in the world when it is delivered in 2021,” said Mark Papermaster, executive vice president and chief technology officer, AMD. “Show attendees can spend time with the same foundational AMD technologies that will go into that exascale-class system this week. From high performance AMD EPYC CPUs and Radeon Instinct GPUs working together across high-speed interconnects, to our open software ecosystem, all delivered by the biggest names in supercomputing, it’s all here at SC19.”

At SC19, Papermaster will be on a panel along with CTOs from Cray, CERN and Xilinx to [discuss](#) the unique innovations in exascale systems and beyond, including new technical breakthroughs across hardware, software and programming tools.

New Amazon EC2 Compute-Optimized Instances to Use 2nd Gen AMD EPYC

Amazon Web Services (AWS) is expanding its use of AMD EPYC processors, announcing the [upcoming launch of two new Amazon Elastic Compute Cloud \(Amazon EC2\) compute-optimized instances](#). These instances, C5a and C5ad, will be powered by custom 2nd Gen AMD EPYC processors, running at frequencies up to 3.3Ghz. C5a and C5ad will come in 8 virtualized sizes with up to 96 vCPUs which will provide additional choices to help customers optimize both cost and performance for a variety of compute intensive workloads, including batch processing, distributed analytics, and web applications. Both instances will be available in bare metal variants which will enable customers to run applications that benefit from direct access to the processor and memory resources of the underlying server. These instances will have 192 logical processors on 96 physical cores, which will be twice the largest instance size offered in the EC2 compute-optimized instance family. The C5a and C5ad bare metal instance will be able to utilize 100 Gbps network bandwidth and will be

compatible with Elastic Fabric Adapter, enabling customers to scale up High Performance Computing and other large compute intensive workloads. The new instances will be available soon across multiple AWS regions.

AMD Powers Supercomputers to New Levels and EPYC Joins TOP500

HPC organizations are continuing to adopt the 2nd Gen AMD EPYC processor and Radeon Instinct accelerators for more powerful and efficient supercomputing systems. The 2nd Gen EPYC processors provide twice the manufacturing application performanceⁱ and up to 60% faster Life Sciences simulations than competing solutionsⁱⁱⁱⁱⁱ, while the Radeon Instinct GPU accelerator provides up to 6.6 peak theoretical TFLOPS Double Precision performance for HPC workloads. Both products support PCIe[®] 4.0 enabling high-bandwidth interconnect for faster compute across heterogenous systems.

The latest customers deploying AMD processors and accelerators include:

- [Atos](#), a global leader in digital transformation, is supplying two BullSequana XH2000 supercomputers based on 2nd Gen EPYC to Météo-France for operational weather forecasting and research in atmospheric, ocean and climate science.
- Atos and the French national high-performance computing organization, GENCI, announced that the latest extension of its Joliot-Curie supercomputer, located and administered by the CEA teams at its TGCC (Very Large Computing Centre), and based on Atos' BullSequana XH2000 solution and 2nd Gen EPYC is now operational.
- The extension of the Joliot-Curie supercomputer joined the 54th edition of the TOP500, making it the first AMD EPYC 7H12 64-core 280W^{iv} equipped supercomputer in the global ranking.
- Cray, a Hewlett Packard Enterprise company, recently announced two new supercomputers equipped with the 2nd Gen AMD EPYC processors and based on the Shasta™ supercomputer architecture, [ARCHER2](#) and [Vulcan](#).
- Dell Technologies is supplying the [San Diego Supercomputer Center](#) and its Expanse supercomputer with Dell EMC PowerEdge servers and 2nd Gen EPYC Processors.
- ETH Zurich is using AMD EPYC 7742 processors in its Euler VI system.
- NEC is supplying the German weather forecasting service, Deutscher Wetterdienst, a system using 2nd Gen EPYC processors in conjunction with NEC SX-Aurora TSUBASA vector engines.

“We chose the AMD EPYC 7742 because the processor not only provided broad support for most common software applications our researchers use, but it provides the raw performance, impressive memory and I/O bandwidth and most importantly the price/performance that is critical to meeting the demands of these researchers,” said Christian Bolliger, scientific IT services, ETH Zurich. “With the Euler VI system using 2nd Gen AMD EPYC processors, our users now have access to a system that will provide them the capabilities needed to advance their research.”

Bringing Supercomputing to the Cloud with AMD EPYC

The HPC industry is evolving to support new workloads, greater demands for performance and most importantly, easier access for long term or temporary use. This is being done through the Cloud, which for the first time can provide users with the same levels of performance, at lower use costs, compared to what they can deploy on-premise.

Microsoft Azure [announced](#) previously unobtainable levels of performance for computational fluid dynamics (CFD) using an Azure HB cloud instance running on an 1st Gen AMD EPYC-processor based system.

Now, Azure is pushing the boundaries of HPC in the cloud even further with the preview of [Azure HBv2 virtual machines](#) for high-performance computing. Based on the AMD EPYC 7742 processor, these virtual machines provide customers with access to supercomputer performance, supporting 200Gbps HDR InfiniBand, and up to 80,000 cores for a single job, in the ease and simplicity through the Cloud.

AMD Introduces ROCm 3.0

Community support for the pre-exascale software ecosystem continues to grow. This ecosystem is built on ROCm, the foundational open source components for GPU compute provided by AMD. The ROCm development cycle features monthly releases offering developers a regular cadence of continuous improvements and updates to compilers, libraries, profilers, debuggers and system management tools. Major development milestones featured at SC19 include:

- Introduction of ROCm 3.0 with new innovations to support HIP-clang – a compiler built upon LLVM, improved CUDA conversion capability with hipify-clang, library optimizations for both HPC and ML.
- ROCm upstream integration into leading [TensorFlow](#) and PyTorch machine learning frameworks for applications like reinforcement learning, autonomous driving, and image and video detection.
- Expanded acceleration support for HPC programming models and applications like OpenMP programming, LAMMPS, and NAMD.
- New support for system and workload deployment tools like Kubernetes, Singularity, SLURM, TAU and others.

A Growing Hardware Ecosystem

As HPC systems are asked to support more demanding workloads, the need for heterogeneous compute powered by CPUs and accelerators is critical for modern HPC systems. AMD partners are creating platforms that support this demand, as well as traditional CPU only computing:

- [GIGABYTE announced](#) four new G-Series GPU servers that support 2nd Gen AMD EPYC processors, the G292-Z22, G292-Z42, G482-Z50 and G482-Z51. The G482-Z51 can support up to eight PCIe 4.0 GPU cards, giving customers great “AMD + AMD” options for a range of accelerated computing workloads.
- [The world-record shattering](#) HPE ProLiant DL325 Gen10 and DL385 Gen10 servers are joined by [new Gen10 Plus models](#) significantly enhancing performance and efficiency for vital workloads like virtualization, HPC and Big Data with up to 64 processor cores, 3200 MT/s memory for 9 percent faster memory performance, 2X greater I/O bandwidth with PCIe 4.0 support and 2.4X greater storage capacity versus previous generations.
- With the new Penguin Altus[®] XE4218GT, supporting up to 8 GPUs, Penguin customers have an “AMD + AMD” solution that uses the PCIe 4.0 support in the 2nd Gen EPYC and Radeon Instinct MI50 to power machine learning, big data analytics, and similar workloads.
- Tyan also announced new platforms in its HPC-focused Transport HX product line and database-focused Transport SX product line powered by 2nd Gen AMD EPYC

processors.

With the introduction of PCIe 4.0 support in the 2nd Gen AMD EPYC processors and Radeon Instinct GPU accelerators, AMD has led the enablement of the PCIe 4.0 ecosystem by working closely with key industry partners. The 2nd Gen EPYC supports PCIe 4.0 enabled devices such as high speed ethernet and InfiniBand® interconnects, NICs and switches, compute accelerators (GPUs and FPGAs) and storage (NVME SSD) devices. Some of the partners offering PCIe 4.0 products for 2nd Gen EPYC include:

- Broadcom Thor NIC for 200 GB ethernet.
- Mellanox ConnectX-6 NIC showing ~400 GB/s InfiniBand performance.
- Samsung Gen4 PM1733 NVME SSD – Showcasing 2x of IOPS over the Samsung Gen3 SSD.
- Xilinx Alveo U50, U280 FPGAs.

Additional Resources

- [AMD EPYC](#) on AMD.com
- Learn more about AMD Radeon Instinct™ [MI50](#) accelerators
- Learn more about the ROCm 3.0 open software platform [here](#)
- Learn about how AMD EPYC performs with HPC software [here](#)
- Follow AMD datacenter developments on Twitter [@AMDServer](#)
- Follow AMD Radeon Instinct™ on Twitter [@RadeonInstinct](#)

About AMD

For 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](#), and [Facebook](#) and [Twitter](#) pages.

Cautionary Statement

This press release contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) including the features, functionality, availability, timing, deployment and expectations of AMD future products and technologies as well as future collaborations and the expected benefits of those collaborations, 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," "intends," "believes," "expects," "may," "will," "should," "seeks," "intends," "plans," "pro forma," "estimates," "anticipates," or the negative of these words and phrases, other variations of these words and phrases or comparable terminology. Investors are cautioned that the forward-looking statements in this document are based on current beliefs, assumptions and expectations, speak only as of the date of this document 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 may limit AMD's ability to compete effectively; AMD relies on third parties to manufacture its products, and if they are unable to do so on a timely basis in sufficient quantities and using competitive technologies, AMD's business could be materially adversely affected; failure to achieve expected manufacturing yields for AMD's products could negatively impact its financial results; AMD has a wafer supply agreement with GF with obligations to purchase all of its microprocessor and APU product requirements, and a certain portion of its GPU product requirements, from GLOBALFOUNDRIES Inc. (GF) with limited exceptions. If GF is not able to satisfy AMD's manufacturing requirements, its business could be adversely impacted; the success of AMD's business is dependent upon its ability to introduce products on a timely basis with features and performance levels that provide value to its customers while supporting and coinciding with significant industry transitions; if AMD cannot generate sufficient revenue and operating cash flow or obtain external financing, it may face a cash shortfall and be unable to make all of its planned investments in research and development or other strategic investments; the loss of a significant customer may have a material adverse effect on AMD; AMD's receipt of revenue from its semi-custom SoC products is dependent upon its technology being designed into third-party products and the success of those products; global economic uncertainty may adversely impact AMD's business and operating results; AMD's operations are subject to political, legal and economic risks and natural disasters which could have a material adverse effect on AMD; government actions and regulations such as export administration regulations, tariffs and trade protection measures, may limit AMD's ability to export its products to certain customers; AMD products may be subject to security vulnerabilities that could have a material adverse effect on AMD; IT outages, data loss, data breaches and cyber-attacks could compromise AMD's intellectual property or other sensitive information, be costly to remediate and cause significant damage to its business and reputation; AMD's operating results are subject to quarterly and seasonal sales patterns; AMD may not be able to generate sufficient cash to service its debt obligations or meet its working capital requirements; AMD has a large amount of indebtedness which could adversely affect its financial position and prevent it from implementing its strategy or fulfilling its contractual obligations; the agreements governing AMD's notes and the Secured Revolving Line of Credit impose restrictions on AMD that may adversely affect its ability to operate its business; the markets in which AMD's products are sold are highly competitive; the conversion of the 2.125% Convertible Senior Notes due 2026 may dilute the ownership interest of its existing stockholders, or may otherwise depress the price of its common stock; uncertainties involving the ordering and shipment of AMD's products could materially adversely affect it; the demand for AMD's products depends in part on the market conditions in the industries into which they are sold. Fluctuations in demand for AMD's products or a market decline in any of these industries could have a material adverse effect on its results of operations; AMD's ability to design and introduce new products in a timely manner is dependent upon third-party intellectual property; AMD depends on third-party companies for the design, manufacture and supply of motherboards, software and other computer platform components to support its business; if AMD loses Microsoft Corporation's support for its products or other software vendors do not design and develop software to run on AMD's products, its ability to sell its products could be materially adversely affected; and AMD's reliance on third-party distributors and AIB partners subjects it to certain risks. 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 28, 2019.

Links to third party sites are provided for convenience and unless explicitly stated, AMD is not responsible for the contents of such linked sites and no endorsement is implied.

AMD, the AMD Arrow logo, EPYC, Radeon, Radeon Instinct, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.

ⁱ Based on AMD internal testing of ANSYS FLUENT 19.1, Im6000_16m benchmark, as of July 17, 2019 of a 2P EPYC 7742 powered reference server versus a 2P Intel Xeon Platinum 8280 powered server. Results may vary. ROM-42

ⁱⁱ AMD Internal testing as of 30July2019 of a 2P AMD EPYC 7742 powered reference platform versus a 2P Intel Platinum 8280 powered production server, on GROMACS version 2019.3 benchmark. Results may vary. ROM-113

ⁱⁱⁱ "Best-in-class based on industry-standard pin-based (LGA) X86 processors. EPYC-08"

^{iv} EPYC 7H12 processor boost frequencies may be achieved only with a cooling solution that meets group 'Z' requirements. Achievable boost frequencies may vary depending on the effectiveness of the actual cooling solution. ROM-282

PR Contact:
Gary Silcott
+1 (512) 602-0889
Gary.Silcott@amd.com

Investor Contact:
Laura Graves
laura.graves@amd.com
+1 (408) 749-5467



Source: Advanced Micro Devices