

November 16, 2021



AMD Processors Accelerating Performance of Top Supercomputers Worldwide

Growing preference for EPYC™ processors resulted in the number of AMD-powered supercomputers growing 3.5x year-over-year

SANTA CLARA, Calif., Nov. 16, 2021 (GLOBE NEWSWIRE) -- During this year's Supercomputing Conference 2021 (SC21), [AMD](#) (NASDAQ: AMD) is showcasing its expanded presence and growing preference in the high performance computing (HPC) industry with the exceptional innovation and adoption of AMD data center processors and accelerators. Customers across the industry continue to expand their use of AMD EPYC™ processors and AMD Instinct™ accelerators to power cutting-edge research needed to address some of the world's biggest challenges in climate, life sciences, medicine, and more.

Growing preference for AMD is showcased in the latest [Top500 list](#). AMD now powers 73 supercomputers, compared to 21 in the November 2020 list, a more than 3x year-over-year increase. Additionally, AMD powers four out of the top ten most powerful supercomputers in the world, as well as the most powerful supercomputer in EMEA. Finally, AMD EPYC 7003 series processors, which launched eight months ago, are utilized by 17 of the 75 AMD powered supercomputers in the list, demonstrating the rapid adoption of the latest generation of EPYC processors.

"The demands of supercomputing users have increased exponentially as the world seeks to accelerate research, reducing the time to discovery of valuable information," said Forrest Norrod, senior vice president and general manager, Data Center and Embedded Solutions Business Group, AMD. "With AMD EPYC CPUs and Instinct accelerators, we continue to evolve our product offering to push the boundaries of data center technologies enabling faster research, better outcomes and more impact on the world."

AMD has also been recognized in the annual HPCwire Readers' and Editors' Choice Awards at SC21. The company won ten awards including Best Sustainability Innovation in HPC, Best HPC Server Product and the Outstanding Leadership in HPC award presented to President and CEO Dr. Lisa Su.

Expanding Customer Base

AMD is engaged broadly across the HPC industry to deliver the performance and efficiency of AMD EPYC and AMD Instinct products, along with the ROCm™ open ecosystem, to speed research. Through high-profile installations like the ongoing deployment of Oak Ridge National Laboratory's "Frontier" supercomputer, AMD is bringing the compute technologies and performance needed to support developments in current and future research across the world. Highlights of "Frontier" and other new HPC systems in the industry include:

- [“Adastra,”](#) an HPE supercomputer that will have two partitions powered by AMD CPUs and accelerators, was recently announced by GENCI, the French national agency for HPC, and CINES, the National Computing Center for Higher Education. The first partition is based on next-generation AMD EPYC processors code named “Genoa” and the second partition is based on 3rd Gen AMD EPYC processors and AMD Instinct MI250x accelerators.
- [Argonne National Laboratory’s “Polaris” testbed supercomputer](#), powered by AMD EPYC 7003 series processors, enabling scientists and developers to tackle a range of artificial intelligence (AI), engineering and scientific projects.
- A new supercomputer built by HPE using AMD EPYC CPUs to advance weather forecasting and climate research for the [National Center of Meteorology](#) in the United Arab Emirates. HPE also updated [Eni’s supercomputer](#) to accelerate the discovery of energy sources using AMD EPYC processors.
- [Oak Ridge National Laboratory’s “Frontier” exascale computer](#) – which is powered by optimized 3rd Gen AMD EPYC processors and AMD Instinct MI250x accelerators.
- The [Texas Advanced Computing Center](#) at The University of Texas at Austin launched Lonestar6, a Dell Technologies supercomputer powered by AMD EPYC 7003 series processors.
- [University of Vermont’s Advanced Computing Core](#), powered by AMD EPYC processors and AMD Instinct accelerators, driving research into COVID-19 and solutions to future potential threats to global health.
- [Washington University’s advanced clustering technologies](#), powered by AMD EPYC processors, studying COVID-19 and home to the Folding@home project.

A Year of Breakthrough Products and Research

This year AMD [launched](#) its AMD EPYC 7003 series processor, the world’s highest-performing server processor.¹ Since then, there has been overwhelming adoption from partners across the industry who are driving discoveries in biomedicine, predicting natural disasters, clean energy, semiconductors, microelectronics and more.

Expanding on the features of the EPYC 7003 series processor, AMD recently previewed the [3rd Gen EPYC processor with AMD 3D V-cache](#). By utilizing innovative packaging technology, which layers the L3 cache in EPYC 7003 series processors, AMD 3D V-Cache technology offers enhanced performance for the technical computing workloads prevalent in HPC. Microsoft Azure HPC virtual machines featuring 3rd Gen EPYC with AMD 3D V-Cache are [currently available](#) in Private Preview and will be available globally soon.

AMD also unveiled [the world’s fastest HPC and AI accelerator](#)², AMD Instinct MI250X. Designed with the AMD CDNA™ 2 architecture, the [AMD Instinct MI200 series accelerators](#) deliver up to 4.9x the peak FP64 performance versus competitive data center accelerators, which is critical for HPC applications like weather modeling². The AMD Instinct MI200 series accelerators are also the first to have over 100GB high-bandwidth memory capacity, delivering up to 3.2 terabytes per second, the industry’s best aggregate bandwidth³.

Supporting Resources

- Find more AMD HPC & AI information and customer testimonials on the [AMD HPC and AI Solutions Hub](#).
- Learn more about [AMD EPYC Processors](#) and [AMD Instinct Accelerators](#)
 - Read more about [AMD Exascale Computing Technologies](#) and [AMD HPC Solutions](#)
 - 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](#) and [Twitter](#) pages.

AMD, the AMD Arrow logo, AMD CDNA, EPYC, AMD 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.

¹ MLN-016: Results as of 01/28/2021 using SPECrate@2017_int_base. The AMD EPYC 7763 a measured estimated score of 798 is higher than the current highest 2P server with an AMD EPYC 7H12 and a score of 717, <https://spec.org/cpu2017/results/res2020q2/cpu2017-20200525-22554.pdf>. OEM published score(s) for 3rd Gen EPYC may vary. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.

² 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.

³ MI200-07: Calculations conducted by AMD Performance Labs as of Sep 21, 2021, for the AMD Instinct™ MI250X and MI250 (128GB HBM2e) OAM accelerators designed with AMD CDNA™ 2 6nm FinFet process technology at 1,600 MHz peak memory clock resulted in 128GB HBM2e memory capacity and 3.2768 TFLOPS peak theoretical memory bandwidth performance. MI250/MI250X memory bus interface is 4,096 bits times 2 die and memory data rate is 3.20 Gbps for total memory bandwidth of 3.2768 TB/s ((3.20 Gbps*(4,096 bits*2))/8). The highest published results on the NVidia Ampere A100 (80GB) SXM GPU accelerator resulted in 80GB HBM2e memory capacity and 2.039 TB/s GPU memory bandwidth performance. <https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/a100/pdf/nvidia-a100-datasheet-us-nvidia-1758950-r4-web.pdf>

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.