November 18, 2024

AMD Accelerates Exascale Computing to New Heights Powering the Fastest Supercomputer Ever, El Capitan

— El Capitan, powered by the AMD Instinct MI300A APU, becomes the second AMD supercomputer to surpass the Exascale barrier, placing #1 on the Top500 list —

— AMD continues setting the standard for HPC, powering 50 percent of the top ten fastest and 40 percent of the ten most energy efficient supercomputers in the world—

— IBM and AMD announce collaboration to deploy AMD Instinct MI300X accelerators as a service on IBM Cloud —

SANTA CLARA, Calif., Nov. 18, 2024 (GLOBE NEWSWIRE) -- Today, AMD (NASDAQ: AMD) showcased its ongoing high performance computing (HPC) leadership at Supercomputing 2024 by powering the world's fastest supercomputer for the sixth straight Top500 list.

The El Capitan supercomputer, housed at Lawrence Livermore National Laboratory (LLNL), powered by AMD Instinct™ MI300A APUs and built by Hewlett Packard Enterprise (HPE), is now the fastest supercomputer in the world with a High-Performance Linpack (HPL) score of 1.742 exaflops based on the latest Top500 list. Both El Capitan and the Frontier system at Oak Ridge National Lab claimed numbers 18 and 22, respectively, on the Green500 list, showcasing the impressive capabilities of the AMD EPYC processors and AMD Instinct GPUs to drive leadership performance and energy efficiency for HPC workloads.

"We are thrilled to see El Capitan become the second AMD powered supercomputer to break the exaflop barrier and become the fastest supercomputer in the world. Showcasing the incredible performance and efficiency of the AMD Instinct MI300 APUs, this groundbreaking machine is a testament to the dedicated work between AMD, LLNL and HPE," said Forrest Norrod, executive vice president and general manager, AMD. "At AMD, we are driving the future of computing with leadership performance and capabilities that will continue to define the convergence of HPC and AI for years to come."

"El Capitan is crucial to the National Nuclear Security Administration's core mission and significantly bolsters our ability to perform large ensembles of high-fidelity 3D simulations that address the intricate scientific challenges facing the mission," said Rob Neely, director of LLNL's Advanced Simulation and Computing program.

Bronis R. de Supinski, LLNL's chief technology officer for Livermore Computing adds, "Leveraging the AMD Instinct MI300A APUs, we've built a system that was once unimaginable, pushing the absolute boundaries of computational performance while maintaining exceptional energy efficiency. With AI becoming increasingly prevalent in our

field, El Capitan allows us to integrate AI with our traditional simulation and modeling workloads, opening new avenues for discovery across various scientific disciplines."

AMD – Powering HPC and AI
AMD compute engines continue to power the most important supercomputers and deliver exceptional performance for technical computing for enterprises and national laboratories around the world.

The latest AMD EPYC™ 9005 Series processors are the best server CPUs for enterprise, AI and cloud[i], providing up to 37 percent better generational IPC performance for HPC and AI workloads[ii]. These processors also provide up to 3.9X faster time to insights compared to the competition for science and HPC applications that solve the world's most challenging problems[iii].

AMD Instinct accelerators provide leadership performance for the data center, at any scale—from AI solutions to Exascale-class supercomputers. The AMD Instinct MI300X and MI325X accelerators enable leadership AI performance and memory capabilities, while the AMD Instinct MI300A APU puts CPU and GPU cores, and stacked memory together into a single package, enabling new levels of efficiency and performance for HPC and AI workloads.

Additionally, AMD EPYC Processors and AMD Instinct accelerators are being used to power many new supercomputing and AI projects and deployments including:

- Eni, the Italian energy company, recently announced the HPC 6 supercomputer powered by AMD EPYC CPUs and AMD Instinct GPUs. HPC 6 is one of the world's most powerful supercomputers dedicated to industrial applications and is now the fifth fastest supercomputer in the world.
- The University of Paderborn is in the process of taking delivery and installing a new supercomputer powered by the latest 5th Gen AMD EPYC CPUs.
- Sigma2 AS will replace two of three nationally owned supercomputers in Norway with an HPE Cray Supercomputing EX system powered by 5th Gen AMD EPYC CPUs. When fully installed, this supercomputer is expected to be the fastest system in Norway.
- IBM and AMD have announced a collaboration to deploy AMD Instinct MI300X accelerators as a service on IBM Cloud. This offering, expected to be available in the first half of 2025, aims to enhance performance and power efficiency for Gen AI models such as high-performance computing applications for enterprise clients. The collaboration will also enable support for AMD Instinct MI300X accelerators within IBM's watsonx AI and data platform, as well as Red Hat® Enterprise Linux® AI inferencing support.
- AMD Instinct MI300A APUs will also power a next-generation supercomputer system for Japan's National Institutes for Quantum Science and Technology (QST). The system, built by NEC Corporation, will use 280 AMD Instinct MI300A APUs to drive AI and scientific research for the National Institutes for Quantum Science and Technology, and the National Institute for Fusion Science.

Leading the Exascale Era
As the only company that powers multiple exascale supercomputers, AMD is continuing to drive performance and energy efficiency leadership for HPC installations around the world.

El Capitan, the world's most powerful supercomputer and the first exascale-class machine for the National Nuclear Security Administration (NNSA) stands as the premiere computing resource for the NNSA Tri-Labs — LLNL, Los Alamos and Sandia National Laboratories. It will be used to advance scientific discovery and national security, providing the vast computational power necessary to ensure the safety, security and reliability of the nation's nuclear deterrent without testing. This state-of-the-art system marks a monumental leap forward in HPC, enabling unprecedented modeling and simulation capabilities essential for NNSA's Stockpile Stewardship Program that certifies the aging nuclear stockpile, and other critical nuclear security missions, such as nonproliferation and counterterrorism.

LLNL and the other NNSA Tri-Labs are also using El Capitan and its companion system, Tuolumne, to drive AI and machine learning-assisted data analysis, further propelling LLNL's AI-driven goals of creating scientific models that are fast, accurate and capable of quantifying uncertainty in their predictions. El Capitan will apply AI to high energy density problems such as inertial confinement fusion research, while Tuolumne will be used for unclassified open science applications including climate modeling, biosecurity/drug discovery, and earthquake modeling.

Beyond El Capitan, AMD and HPE also power the first exascale supercomputer, Frontier. Housed at Oak Ridge National Lab and powered by AMD EPYC CPUs and AMD Instinct GPUs, Frontier is the second fastest computer in the world with 1.35 exaflops of performance. Frontier continues to enable researchers to tackle complex scientific problems, from climate modeling and biomedical research to training large language models, further demonstrating its significant contribution to advancing scientific discovery and fueling breakthroughs in AI.

These world-leading systems provide immense computational power that significantly contribute to a wide range of research, including materials science, climate modeling and AI model development. By empowering researchers across diverse fields and fueling the development of AI models, El Capitan and Frontier are shaping the future of science and technology and enabling solutions to address critical issues across the globe, underscoring the commitment at AMD to provide the high-performance computing resources necessary to power the next generation of scientific discovery and innovation.

Visit the AMD booth #2731 at Supercomputing 2024 to learn more about AMD solutions for HPC and speak with AMD experts.

Supporting Resources

- Learn more about AMD EPYC Processors
- Learn more about AMD Instinct Accelerators
- Follow AMD on X
- Connect with AMD on LinkedIn

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 X pages.

[i] EPYC-029C: Comparison based on thread density, performance, features, process technology and built-in security features of currently shipping servers as of 10/10/2024. EPYC 9005 series CPUs offer the highest thread density [EPYC-025B], leads the industry with 500+ performance world records [EPYC-023F] with performance world record enterprise leadership Java® ops/sec performance [EPYCWR-20241010-260], top HPC leadership with floating-point throughput performance [EPYCWR-2024-1010-381], AI end-to-end performance with TPCx-AI performance [EPYCWR-2024-1010-525] and highest energy efficiency scores [EPYCWR-20241010-326]. The 5th Gen EPYC series also has 50% more DDR5 memory channels [EPYC-033C] with 70% more memory bandwidth [EPYC-032C] and supports 70% more PCIe® Gen5 lanes for I/O throughput [EPYC-035C], has up to 5x the L3 cache/core [EPYC-043C] for faster data access, uses advanced 3-4nm technology, and offers Secure Memory Encryption + Secure Encrypted Virtualization (SEV) + SEV Encrypted State + SEV-Secure Nested Paging security features. See the AMD EPYC Architecture White Paper (https://library.amd.com/l/3f4587d147382e2/) for more information.

[ii] 9xx5-001: Based on AMD internal testing as of 9/10/2024, geomean performance improvement (IPC) at fixed-frequency.
- 5th Gen EPYC CPU Enterprise and Cloud Server Workloads generational IPC Uplift of 1.170x (geomean) using a select set of 36 workloads and is the geomean of estimated scores for total and all subsets of SPECrate®2017_int_base (geomean), estimated scores for total and all subsets of SPECrate®2017_fp_base (geomean), scores for Server Side Java multi instance max ops/sec, representative Cloud Server workloads (geomean), and representative Enterprise server workloads (geomean).
“Genoa” Config (all NPS1): EPYC 9654 BIOS TQZ1005D 12c12t (1c1t/CCD in 12+1), FF 3GHz, 12x DDR5-4800 (2Rx4 64GB), 32Gbps xGMI;
“Turin” config (all NPS1): EPYC 9V45 BIOS RVOT1000F 12c12t (1c1t/CCD in 12+1), FF 3GHz, 12x DDR5-6000 (2Rx4 64GB), 32Gbps xGMI
Utilizing Performance Determinism and the Performance governor on Ubuntu® 22.04 w/ 6.8.0-40-generic kernel OS for all workloads.
- 5th Gen EPYC generational ML/HPC Server Workloads IPC Uplift of 1.369x (geomean) using a select set of 24 workloads and is the geomean of representative ML Server Workloads (geomean), and representative HPC Server Workloads (geomean).
“Genoa” Config (all NPS1) “Genoa” config: EPYC 9654 BIOS TQZ1005D 12c12t (1c1t/CCD in 12+1), FF 3GHz, 12x DDR5-4800 (2Rx4 64GB), 32Gbps xGMI;
“Turin” config (all NPS1): EPYC 9V45 BIOS RVOT1000F 12c12t (1c1t/CCD in 12+1), FF 3GHz, 12x DDR5-6000 (2Rx4 64GB), 32Gbps xGMI
Utilizing Performance Determinism and the Performance governor on Ubuntu 22.04 w/ 6.8.0-40-generic kernel OS for all workloads except LAMMPS, HPCG, NAMD, OpenFOAM, Gromacs which utilize 24.04 w/ 6.8.0-40-generic kernel.
SPEC® and SPECrate® are registered trademarks for Standard Performance Evaluation Corporation. Learn more at spec.org.
[iii] 9xx5-022: Source: https://www.amd.com/content/dam/amd/en/documents/epyc-technical-docs/performance-briefs/amd-epyc-9005-pb-gromacs.pdf

Media Contacts:
Aaron Grabein
AMD Communications
+1 512-602-8950
aaron.grabein@amd.com

Mitch Haws
AMD Investor Relations
+1 512-944-0790
mitch.haws@amd.com

Source: Advanced Micro Devices, Inc.