March 4, 2020

Next-Generation AMD EPYC™ CPUs and Radeon™ Instinct GPUs Enable El Capitan Supercomputer at Lawrence Livermore National Laboratory to Break 2 Exaflops Barrier

— Collaboration with HPE targets leadership HPC and AI processing performance starting in early 2023 —

— El Capitan is set to be the second all AMD CPU and GPU exascale system in the U.S., and is expected to be more powerful than today's 200 fastest supercomputers combined —

SANTA CLARA, Calif., March 04, 2020 (GLOBE NEWSWIRE) -- AMD (NASDAQ: AMD) today joined Lawrence Livermore National Laboratory (LLNL) and HPE, in announcing that El Capitan, the upcoming exascale class supercomputer at LLNL, will be powered by next generation AMD EPYC™ CPUs, AMD Radeon™ Instinct GPUs and open source AMD ROCm heterogeneous computing software.

With delivery expected in early 2023, the El Capitan system is expected to be the world's fastest supercomputer with more than 2 exaflops of double precision performance. This record setting performance will support National Nuclear Security Administration requirements for its primary mission of ensuring the safety, security and reliability of the nation's nuclear stockpile.

The AMD based nodes will be optimized to accelerate artificial intelligence (AI) and machine learning (ML) workloads to potentially enable the expanded use of AI and ML into the research, computational techniques and analysis that benefits NNSA missions.

"El Capitan will drive unprecedented advancements in HPC and AI, powered by the next generation AMD EPYC CPUs and Radeon Instinct GPUs," said Forrest Norrod, senior vice president and general manager, Datacenter and Embedded Systems Group, AMD. "Building on our strong foundation in high-performance computing and adding transformative coherency capabilities, AMD is enabling the NNSA Tri-Lab community—LLNL and the Los Alamos and Sandia national laboratories—to achieve their mission critical objectives and contribute new AI advancements to the industry. We are extremely proud to continue our exascale work with HPE and NNSA and look forward to the delivery of the most powerful supercomputer in the world expected in early 2023."

AMD Technology Powering El Capitan
AMD is applying its experience from the high-performance computing industry and advancing its existing CPU and GPU designs for El Capitan, enabling the system to be architected for maximum performance and ease of use.

AMD technology within El Capitan includes:

- Next generation AMD EPYC processors, codenamed “Genoa” featuring the “Zen 4” processor core. These processors will support next generation memory and I/O sub systems for AI and HPC workloads,
- Next generation Radeon Instinct GPUs based on a new compute-optimized architecture for workloads including HPC and AI. These GPUs will use the next-generation high bandwidth memory and are designed for optimum deep learning performance,
- The 3rd Gen AMD Infinity Architecture, which will provide a high-bandwidth, low latency connection between the four Radeon Instinct GPUs and one AMD EPYC CPU included in each node of El Capitan. As well, the 3rd Gen AMD Infinity Architecture includes unified memory across the CPU and GPU, easing programmer access to accelerated computing,
- An enhanced version of the open source ROCm heterogenous programming environment, being developed to tap into the combined performance of AMD CPUs and GPUs, unlocking maximum performance.

“This unprecedented computing capability, powered by advanced CPU and GPU technology from AMD, will sustain America’s position on the global stage in high performance computing and provide an observable example of the commitment of the country to maintaining an unparalleled nuclear deterrent,” said LLNL Lab Director Bill Goldstein. “Today’s news provides a prime example of how government and industry can work together for the benefit of the entire nation.”

“We are pleased to partner with AMD to bring together HPE’s Cray Shasta architecture with latest AMD EPYC CPUs and Radeon Instinct GPUs and provide leadership technology in support of the mission critical HPC and AI workloads for Lawrence Livermore National Laboratory,” said Peter Ungaro, senior vice president and general manager, HPC and Mission Critical Solutions (MCS) at HPE. “In our three U.S. exascale wins this last year, which includes the U.S. Department of Energy’s Frontier at Oak Ridge National Laboratory, the HPE systems and technologies along with AMD processors and GPUs have become a unique solution for high-performance and exascale-class computing and we look forward to continuing our collaboration to further bolster innovation.”

Additional Resources

- AMD Exascale Computing Technologies
- HPE Announcement
- Lawrence Livermore Labs Announcement
- Follow AMD datacenter developments on Twitter @AMDServer

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 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.

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 an EPYC™ CPU and Radeon™ Instinct GPU powered exascale class supercomputer with open source AMD ROCm heterogeneous computing software, 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 presentation are based on current beliefs, assumptions and expectations, speak only as of the date of this presentation 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; AMD's 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; AMD's ability to generate sufficient revenue and operating cash flow or obtain external financing for research and development or other strategic investments; the loss of a significant customer; AMD's ability to generate revenue from its semi-custom SoC products; global economic uncertainty; political, legal and economic risks and natural disasters; government actions and regulations such as export administration regulations, tariffs and trade protection measures may limit our ability export our products to certain customers; potential security vulnerabilities; potential IT outages, data loss, data breaches and cyber-attacks; the ability of GLOBALFOUNDRIES Inc. to satisfy AMD's manufacturing requirements; 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 secured credit facility; the competitive markets in which AMD's products are sold; the potential dilutive effect if the 2.125% Convertible Senior Notes due 2026 are converted; the 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; and AMD's reliance on third-party distributors and add-in-board partners. 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 Annual Report on Form 10-K for the year ended December 28, 2019.

AMD, the AMD Arrow logo, EPYC, 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.

Contact:
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