

December 12, 2016



AMD introduces Radeon Instinct: Accelerating Machine Intelligence

AMD speeds deep learning inference and training with high-performance Radeon Instinct accelerators and MIOpen open-source GPU-accelerated library

SUNNYVALE, CA -- (Marketwired) -- 12/12/16 -- [AMD](#) (NASDAQ: AMD) today unveiled its strategy to accelerate the machine intelligence era in server computing through a new suite of hardware and open-source software offerings designed to dramatically increase performance, efficiency, and ease of implementation of deep learning workloads. New Radeon™ Instinct accelerators will offer organizations powerful GPU-based solutions for deep learning inference and training. Along with the new hardware offerings, AMD announced MIOpen, a free, open-source library for GPU accelerators intended to enable high-performance machine intelligence implementations, and new, optimized deep learning frameworks on AMD's ROCm software to build the foundation of the next evolution of machine intelligence workloads.

Inexpensive high-capacity storage, an abundance of sensor driven data, and the exponential growth of user-generated content are driving exabytes of data globally. Recent advances in machine intelligence algorithms mapped to high-performance GPUs are enabling orders of magnitude acceleration of the processing and understanding of that data, producing insights in near real time. Radeon Instinct is a blueprint for an open software ecosystem for machine intelligence, helping to speed inference insights and algorithm training.

"Radeon Instinct is set to dramatically advance the pace of machine intelligence through an approach built on high-performance GPU accelerators, and free, open-source software in MIOpen and ROCm," said AMD President and CEO, Dr. Lisa Su. "With the combination of our high-performance compute and graphics capabilities and the strength of our multi-generational roadmap, we are the only company with the GPU and x86 silicon expertise to address the broad needs of the datacenter and help advance the proliferation of machine intelligence."

At the AMD Technology Summit held last week, customers and partners from 1026 Labs, Inventec, SuperMicro, University of Toronto's CHIME radio telescope project and Xilinx praised the launch of Radeon Instinct, discussed how they're making use of AMD's machine intelligence and deep learning technologies today, and how they can benefit from Radeon Instinct.

Radeon Instinct accelerators feature passive cooling, AMD MultiGPU (MxGPU) hardware virtualization technology conforming with the SR-IOV (Single Root I/O Virtualization) industry standard, and 64-bit PCIe addressing with Large Base Address Register (BAR) support for multi-GPU peer-to-peer support.

Radeon Instinct accelerators are designed to address a wide-range of machine intelligence applications:

- ***The Radeon Instinct MI6 accelerator based on the acclaimed Polaris GPU architecture will be a passively cooled inference accelerator optimized for jobs/second/Joule with 5.7 TFLOPS of peak FP16 performance at 150W board power and 16GB of GPU memory***
- ***The Radeon Instinct MI8 accelerator, harnessing the high-performance, energy-efficient "Fiji" Nano GPU, will be a small form factor HPC and inference accelerator with 8.2 TFLOPS of peak FP16 performance at less than 175W board power and 4GB of High-Bandwidth Memory (HBM)***
- ***The Radeon Instinct MI25 accelerator will use AMD's next-generation high-performance Vega GPU architecture and is designed for deep learning training, optimized for time-to-solution***

A variety of open source solutions are fueling Radeon Instinct hardware:

- ***MIOpen GPU-accelerated library: To help solve high-performance machine intelligence implementations, the free, open-source MIOpen GPU-accelerated library is planned to be available in Q1 2017 to provide GPU-tuned implementations for standard routines such as convolution, pooling, activation functions, normalization and tensor format***
- ***ROCm deep learning frameworks: The ROCm platform is also now optimized for acceleration of popular deep learning frameworks, including Caffe, Torch 7, and Tensorflow*, allowing programmers to focus on training neural networks rather than low-level performance tuning through ROCm's rich integrations. ROCm is intended to serve as the foundation of the next evolution of machine intelligence problem sets, with domain-specific compilers for linear algebra and tensors and an open compiler and language runtime***

AMD is also investing in developing interconnect technologies that go beyond today's PCIe Gen3 standards to further performance for tomorrow's machine intelligence applications. AMD is collaborating on a number of open high-performance I/O standards that support broad ecosystem server CPU architectures including X86, OpenPOWER, and ARM AArch64. AMD is a founding member of CCIX, Gen-Z and OpenCAPI working towards a future 25 Gbit/s phi-enabled accelerator and rack-level interconnects for Radeon Instinct.

Radeon Instinct products are expected to ship in 1H 2017. For more information, visit Radeon.com/Instinct.

Supporting Resources

- ***Learn more about Radeon Instinct at Radeon.com/Instinct***
- ***Watch the new Radeon Instinct teaser video [here](#)***
- ***Read how Alibaba Cloud is using AMD technology [here](#)***
- ***Read how Google Cloud Platform is using AMD technology [here](#)***
- ***Read how University of Toronto's CHIME radio astronomy project is using AMD technology [here](#)***
- ***Learn more about [ROCm](#)***
- ***Become a fan of [AMD](#) on Facebook***
- ***Follow Radeon Instinct on Twitter [@RadeonInstinct](#)***

About AMD

For more than 45 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, Radeon 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.

**** Tensorflow support is expected to be available January 2017.***

Cautionary Statement

This press release contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) including the features, functionality, expectations, timing and availability of AMD Radeon Instinct products, 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 document 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 may limit AMD's ability to compete effectively; AMD is party to a wafer supply agreement with GLOBALFOUNDRIES Inc. (GF) with obligations to manufacture products at GF with certain exceptions. If GF is not able to satisfy AMD's manufacturing requirements, its business could be adversely impacted; 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; 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; the markets in which AMD's products are sold are highly competitive; AMD may not be able to generate sufficient cash to service its debt obligations or meet its working capital requirements; AMD has a substantial

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 (Secured Revolving Line of Credit) impose restrictions on AMD that may adversely affect its ability to operate its business; 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 Add-in-Board 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 24, 2016.

Contact Information

Chris Hook
AMD Communications
512-578-9727
chris.hook@amd.com

Laura Graves
AMD Investor Relations
408-749-5467
laura.graves@amd.com

Source: Advanced Micro Devices