

November 15, 2016



AMD Radeon Technology Will Be Available on Google Cloud Platform in 2017

Google Compute Engine will expand its cloud computing offerings to support AMD FirePro S9300 x2

SALT LAKE CITY, UT -- (Marketwired) -- 11/15/16 -- At SC16, [AMD](#) (NASDAQ: AMD) announced that Radeon™ GPU technology will be available to Google Cloud Platform users worldwide. Starting in 2017, Google will use AMD's fastest available single-precision dual GPU compute accelerators, Radeon-based AMD FirePro™ S9300 x2 Server GPUs, to help accelerate [Google Compute Engine](#) and [Google Cloud Machine Learning](#) services.(1) AMD FirePro S9300 x2 GPUs can handle highly parallel calculations, including complex medical and financial simulations, seismic and subsurface exploration, machine learning, video rendering and transcoding, and scientific analysis. Google Cloud Platform will make the AMD GPU resources available for all their users around the world.

"Graphics processors represent the best combination of performance and programmability for existing and emerging big data applications," said Raja Koduri, senior vice president and chief architect, Radeon Technologies Group, AMD. "The adoption of AMD GPU technology in Google Cloud Platform is a validation of the progress AMD has made in GPU hardware and our Radeon Open Compute Platform, which is the only fully open source hyperscale GPU compute platform in the world today. We expect that our momentum in GPU computing will continue to accelerate with future hardware and software releases and advances in the ecosystem of middleware and libraries."

As part of AMD's continuing investments in GPU computing, the company revealed yesterday a new release of Radeon Open Compute Platform (ROCm) featuring software support for new GPU hardware, new math libraries, and a rich foundation of modern programming languages, designed to speed development of high-performance, energy-efficient heterogeneous computing systems. The news bolsters ROCm's position as the most versatile open source platform for GPU computing, targeting high-performance computing across a range of advanced applications from academic and scientific research to commercial deployments. Visit AMD at SC16, booth 1431 to see demonstrations of a wide range of recent GPU computing advances, including:

At SC16 in booth 1431, AMD will showcase:

- ROCm Technology Cluster running Machine Learning Code on Supermicro® servers
- Porting the CUDA application Caffe via HIP Porting Tool
- Ray-tracing and VR visualization for HPC with AMD FirePro™ S9300 X2 & Radeon™ R9 Nano GPUs
- OpenMP 4.5 Interoperability targeting multiple GPUs & platforms
- IBM Power8 server with AMD FirePro™ S9170 Server GPU running ROCm
- Penguin Computing Tundra™ Extreme ARMv8 ThunderX based server with Radeon™ RX 460 running ROCm

- In-situ rendering with Headless OpenGL/EGL Interop OpenCL on ROCm

Supporting Resources

- [AMD Developer Central](#)
- Learn more about [Heterogeneous System Architecture](#)
- Learn more about [GPUOpen](#)
- Learn more about [ROCm](#)
- Become a fan of [AMD](#) on Facebook
- Follow AMD on Twitter [@AMD](#)
- Follow Radeon™ Pro on [Twitter](#)
- Follow Radeon™ graphics on [Twitter](#)

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, FirePro, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. ARM is a registered trademark of ARM Limited in the UK and other countries. OpenCL is a trademark of Apple Inc. used by permission by Khronos Group, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.

Cautionary Statement

This press release contains forward-looking statements concerning Advanced Micro Devices, Inc. ("AMD") including Google's use of AMD FirePro™ S9300 x2 Server GPUs to help accelerate Google's Compute Engine and Google Cloud Machine Learning services; the expectation that Google Cloud Platform will make the AMD GPU resources available for all their users around the world; the expected benefits from Google's use of AMD's GPU resources; and the expectation that the momentum in AMD GPU computing will continue to accelerate with future hardware and software releases and advances in the ecosystem of middleware and libraries, 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 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; We are 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 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 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 24, 2016.

¹ AMD internal testing as of March 2016. System configurations may vary, yielding different results. The AMD FirePro S9300 x2 delivers 13.9 TFLOPS of peak single-precision compute performance. As of March 18th, 2016, NVIDIA's fastest Tesla board is the Tesla K80, which delivers up to 8.73 TFLOPS of single-precision compute performance. For AMD, calculation based on a formula using a combination of shader units and clock speeds. For Nvidia, source of data as of March 18, 2016: www.nvidia.com/object/tesla-k80.htm FP-189

Contact Information

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

Alina Ostrovsky
AMD Investor Relations

408-749-6688

alina.ostrovsky@amd.com

Source: Advanced Micro Devices