

November 14, 2016



AMD Releases New Version of ROCm, the Most Versatile Open Source Platform for GPU Computing

Support for new GPU hardware, math libraries and modern programming languages further unlock the power of GPU computing

SALT LAKE CITY, UT -- (Marketwired) -- 11/14/16 -- At SC16, [AMD](#) (NASDAQ: AMD) today announced a new release of Radeon Open Compute Platform (ROCm) featuring software support of new Radeon 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. AMD also announced planned support of OpenCL™ and for a wide range of CPUs in upcoming releases of ROCm, including support for AMD's upcoming "Zen"-based CPUs, Cavium ThunderX CPUs, and IBM Power 8 CPUs. The advances further cement ROCm's position as the most versatile open source platform for GPU computing.

"Radeon Open Compute is a platform for a new era of GPU problem-solving, designed to harness the power of open source software to unlock new solutions for HPC and hyperscale computing," said Raja Koduri, senior vice president and chief architect, Radeon Technologies Group, AMD. "Today's release of ROCm gives developers consummate flexibility in where and how they use GPU compute. We're excited to show off some of these incredible applications at SC16."

"An open source approach to heterogeneous computing can help bring the benefits of high-performance computing directly to developers, providing the flexibility for them to use available compute resources and extract the best possible performance from their available hardware," said Denise Dumas, vice president, Platform Engineering, Red Hat. "As a leader in helping to nurture and maintain open source communities, we are pleased to see that AMD has elected to take a community-focused approach to ROCm."

"The combination of AMD's GPU support of ThunderX® along with the breadth of features of the ROCm software release is a valuable addition to Cavium's overall solution for the HPC and Hyperscale server market," said Larry Wikelius, Vice President of the Software Ecosystem and Solutions Group at Cavium, Inc. "Today's announcement by AMD is an excellent complement to the innovative features of ThunderX® that include 48 ARMv8 cores, integrated I/O and dual socket configuration support that our customers and partners are deploying today in these markets"

The new release of ROCm introduces a wide range of updates, including:

- *Expanded GPU support* - ROCm now supports all Polaris architecture-based graphics products, including the Radeon™ RX 460, 470 and 480 graphics cards, and the Radeon™ Pro WX 7100, 5100 and 4100 GPUs, growing the rich portfolio of devices

supported by ROCm devices. The Polaris architecture is specifically designed to benefit low-level programming, helping developers to extract the most from the hardware.

- *ROCm Virtualization of the GPU hardware via OS Containers and Linux[®]'s Kernel Virtual Machine (KVM)* - ROCm now supports Docker containerization, allowing end-users to simplify the deployment of an application in ROCm-enabled Linux server environments. ROCm also supports GPU Hardware Virtualization via KVM pass-through to allow the benefits of hardware-accelerated GPU computing in virtualized solutions.
- *Heterogeneous Compute Compiler (HCC)* - HCC is a single source ISO C++ 11/14 compiler for both CPU and GPU, with support for the C++17 "Parallel Standard Template Library". It is built on a rich compiler infrastructure including LLVM-based GCN ISA code generation with assembler and disassembler support.
- *Heterogeneous-Compute Interface for Portability (HIP)* - HIP enables developers to port CUDA applications to ROCm using HIPIFY which automates the conversion to the HIP kernel language and runtime API, creating portable applications that can run on virtually any GPU using either NVIDIA's CUDA Compiler or HCC.
- *New Math Acceleration Libraries* - ROCm introduces support for new advanced math acceleration libraries with support for BLAS, FFT and N-dimensional tensor contractions.

In addition to today's release of ROCm, AMD also announced that it is working to expand the ROCm ecosystem through increased server CPU support, and planned support of OpenCL. Upcoming releases of ROCm are expected to support AMD "Zen"-based x86 CPUs, ARM AArch64 CPU architecture starting with Cavium ThunderX processors, as well as IBM Power 8 CPUs and servers to meet customers' growing hyperscale and HPC needs. OpenCL is being built on a ROCm runtime and compiler foundation which give much greater low-level control of the GPU via its direct-to-metal access.

The new version of Radeon Open Compute is available now at <https://github.com/RadeonOpenCompute/ROCm>.

Demonstrations at SC16

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

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