Dr. Timour Paltashev, Advanced Micro Devices, Santa Clara, California

AMD ROCm: Open Source Platform for High Performance Computing and Machine Intelligence

Abstract

ROCm is the first open-source Exascale-class platform for accelerated computing that's also programming-language independent. It offers the UNIX® philosophy of choice, minimalism and modular software development to GPU computing. Users are free to choose or even develop tools and a language run time for their applications.

ROCm is built for scale, it supports multi-GPU computing and has a rich system run time with the critical features that large-scale application, compiler and language-run-time development requires.

AMD's collaboration with and contributions to the open-source community are a driving force behind ROCm platform innovations. This industry-differentiating approach to accelerated compute and heterogeneous workload development gives our users unprecedented flexibility, choice, and platform autonomy.

Since the ROCm ecosystem is comprised of open technologies: frameworks (Tensorflow / PyTorch), libraries (MIOpen / Blas / RCCL), programming model (HIP), inter-connect (OCD) and up streamed Linux® Kernel support – the platform is continually optimized for performance and extensibility. Tools, guidance, and insights are shared freely across the ROCm GitHub community and forums.