

Why this new AMD FirePro Cluster is important for OpenCL



FirePro S9150 cluster

Then it hit the doormat:

"AMD is proud to collaborate with ASUS, the Frankfurt Institute for Advanced Studies, (FIAS) and GSI to support such important physics and computer science research," said David Cummings, senior director and general manager, professional graphics, AMD. "This installation reaffirms AMD's leading role in HPC with the implementation of the **AMD FirePro S9150 server GPUs in this three petaFLOPS supercomputer cluster**. AMD and ASUS are enabling OpenCL applications for critical science research usage for this cluster. We're committed to building our HPC leadership position in the industry as a foremost provider of computing applications, tools and technologies."

You read more [here](#) and the official news [here](#).

Why is this important?

It could be that there is more flops for the same price, as AMD hardware is cheaper? Nice, but secondary.

That it runs OpenCL? We like that, but from a broader perspective this is not the most important.

It is important because it creates more diversity in the world of HPC. Currently there are a few XeonPhi clusters and only one big AMD FirePro S10000 cluster. The rest is Nvidia Tesla or CPU only. **With more AMD clusters the HPC market is democratised.** That means that more software will be written in vendor-neutral software like OpenCL (with high-level software/libraries on top), and prices of HPC accelerators will not be kept high.

How to further democratisise the HPC world?

We started with porting Gromacs to OpenCL, and we will continue to port large projects to OpenCL. This software will simply run on XeonPhi, Tesla and FirePro with just little porting time, reducing costs in many ways. **We can not do it alone, but together we can.** Start by telling us which software needs to be ported from OpenMP to OpenCL or OpenMP 4, or from CUDA to OpenCL. And if you are porting open source software to OpenCL, drop us a line for free advice and help with testing the software.

And the best you can do to break the monopoly of CUDA, is to simply buy AMD or Intel hardware. The price difference is enough to buy lots of extra FLOPS and to pay for a complete porting project to OpenCL of a large application.