

Question: do we work with CUDA?



Answer: Yes, actually a lot!

The company was built on OpenCL and we are still work with the language a lot - from embedded GPUs and FPGAs to high-end GPUs. Like OpenCL unjustly isn't associated with clusters full of professional GPUs, we were not associated with CUDA. I can tell many of our customers have found us to build high performance software in CUDA.

Breaking with the past is not easy due to associations that seem to stick. With the name change from StreamComputing to Stream HPC some years ago, we wanted to enforce that break with being "the OpenCL company". For some time we were much more pragmatic in solving the problems of our customers, which resulted in making software in MPI and CUDA - sometimes an unexpected direction as the customer initially chose OpenCL.

We also started hiring people who only knew CUDA (but expect them to learn OpenCL), as the right algorithm and the right processor is more important. Internships with CUDA, large CUDA-projects, seeking better relations with Nvidia and such - all have been going on for years. And we like it as much as we like OpenCL - both have unique advantages.

So if you have questions about CUDA, don't be afraid that you hurt us - we're happy to help you get fast software.