

---

## OpenMP Target Device Offloading for the SX-Aurora TSUBASA Vector Engine

Tim Cramer<sup>1</sup>, Manoel Römmer<sup>1</sup>, Boris Kosmynin<sup>1</sup>, Erich Focht<sup>2</sup>, Matthias Mueller<sup>1</sup>

<sup>1</sup>IT Center, RWTH Aachen University, Germany

<sup>2</sup>NEC Cooperation, Stuttgart, Germany

{cramer, roemmer, kosmynin, mueller}@itc.rwth-aachen.de

erich.focht@emea.nec.com

Driven by the heterogeneity trend in modern supercomputers, OpenMP provides support for heterogeneous systems since 2013. Having a single programming model for all kinds of accelerator-based systems decreases the burden of code porting to different device types. The acceptance of this heterogeneous paradigm requires the availability of corresponding OpenMP compiler and runtime environments supporting different target device architectures. The LLVM/Clang infrastructure is designated to extend the offloading features for any new target platform. However, this supposes a compatible compiler backend for the target architecture. In order to overcome this limitation we present a source-to-source code transformation technique which outlines the OpenMP code regions for the target device. By combining this technique with a corresponding communication layer, we enable OpenMP target offloading to the NEC SX-Aurora TSUBASA vector engine, which represents the new generation of vector computing.

**Keywords:** HPC, OpenMP, Offloading, Vector Computing, SIMD.