



Contribution ID: 69

Type: **Talk**

TensorFlow as HPC Infrastructure?

Tuesday, 4 December 2018 13:30 (30 minutes)

TensorFlow is the system driving Google's ML efforts. Many components make up this system, including a sophisticated user-friendly development environment, highly optimized language features and compilers, ultra-high performance custom chips called Tensor Processing Units (TPU), and scalable deployment on the world's devices. TPU pods may well eclipse traditional performance boundaries of the top HPC systems at a much lower cost. We will review TensorFlow's SW and HW environment, which begs the question of how usable it might be for HPC.

Presenter Biography

Bio

Peter Braam is a scientist and entrepreneur focused on large-scale computing. After obtaining a PhD in mathematics under Michael Atiyah at Oxford on Gauge Theories and Magnetic Monopoles, he was an academic at several universities including Oxford, CMU and Cambridge. Peter created the Lustre parallel file system, which has become a key product for large-scale HPC. Peter founded and sold several successful startup companies, and held executive roles at public companies subsequently. From 2013, Peter has been contributing to architecture for data processing in the SKA telescope and is doing research and design for solutions in the area of data-intensive computing.

Primary author: Dr BRAAM, Peter

Presenter: Dr BRAAM, Peter

Session Classification: HPC Technologies

Track Classification: HPC Technology