



Contribution ID: 171

Type: **Talk**

The Role of Storage for efficient Machine Learning

Monday, 30 November 2020 15:45 (30 minutes)

Nowadays, many organizations are trying to find ways to converge classic HPC and AI. There are generally good reasons to do this because of significant similarities between HPC and AI workloads and workload scaling. However, for AI workloads to perform well on clusters, it is also important to be aware of the differences for AI workloads (especially DeepLearning) compared to classic HPC workloads. One of the most important differences are the requirements for storage systems for AI compared to classic HPC. This talk will provide an overview of the special storage system challenges that come with AI workloads, how to characterize and simulate them and especially how to overcome them to ensure that the GPUs can run efficiently instead of stalling on storage access.

Student?

Supervisor name

Supervisor email

Primary author: BREUNER, Sven (Excelero)

Presenter: BREUNER, Sven (Excelero)

Session Classification: HPC

Track Classification: HPC Technology