



Contribution ID: 117

Type: **Talk**

The GekkoFS Burst Buffer File System

Wednesday, 2 December 2020 12:45 (30 minutes)

Many scientific fields increasingly use high-performance computing (HPC) to process and analyze massive amounts of experimental data while storage systems in today's HPC environments have to cope with new access patterns. These patterns include many metadata operations, small I/O requests, or randomized file I/O, while general-purpose parallel file systems have been optimized for sequential shared access to large files. Burst buffer file systems create a separate file system that applications can use to store temporary data. They aggregate node-local storage available within the compute nodes or use dedicated SSD clusters and offer a peak bandwidth higher than that of the backend parallel file system without interfering with it. We present GekkoFS, a temporary, highly-scalable file system which has been specifically optimized for the aforementioned use cases. GekkoFS provides relaxed POSIX semantics which only offers features which are actually required by most (not all) applications. GekkoFS is, therefore, able to provide scalable I/O performance and reaches millions of metadata operations already for a small number of nodes, significantly outperforming the capabilities of common parallel file systems.

Student?

No

Supervisor name

Supervisor email

Primary author: Prof. BRINKMANN, André (Johannes Gutenberg University Mainz)

Presenter: Prof. BRINKMANN, André (Johannes Gutenberg University Mainz)

Session Classification: HPC

Track Classification: Storage and IO