Centre for High Performance Computing 2022 National Conference



Contribution ID: 17

Type: Talk

The role of storage target allocation in applications' I/O performance with BeeGFS - a medium-sized cluster experience

Friday, 2 December 2022 14:00 (30 minutes)

Parallel file systems are at the core of HPC I/O infrastructures. Those systems minimize the I/O time of applications by separating files into fixed-size chunks and distributing them across multiple storage targets. Therefore, the I/O performance experienced with a PFS is directly linked to the capacity to retrieve these chunks in parallel. In this work, we conduct an in-depth evaluation of the impact of the stripe count (the number of targets used for striping) on the write performance of BeeGFS, one of the most popular parallel file systems today. We consider different network configurations and show the fundamental role played by this parameter, in addition to the number of compute nodes, processes and storage targets.

Through a rigorous experimental evaluation, we directly contradict conclusions from related work. Notably, we show that sharing I/O targets does not lead to performance degradation and that applications should use as many storage targets as possible. Our recommendations have the potential to significantly improve the overall write performance of BeeGFS deployments and also provide valuable information for future work on storage target allocation and stripe count tuning.

Primary authors: Dr BOITO, Francieli (University of Bordeaux, Inria); Dr PALLEZ, Guillaume (Inria); Dr TEYLO, Luan (Inria)

Presenters: Dr BOITO, Francieli (University of Bordeaux, Inria); Dr TEYLO, Luan (Inria)

Session Classification: HPC

Track Classification: Storage and IO