

Data Technologies E3

Wednesday, 6 September 2017 16:15 (1 hour)

The first part of hands-on exercises aims to improve understanding of basic parameters in IO systems:

- network and media latency
- access patterns
- OS caching
- bottlenecks and optimization strategies for local and remote data access.

Few essential Linux tools will be introduced to monitor and measure IO performance avoiding bias introduced by OS caching. Students will experience and measure the impact of latency and access patterns on IO performance.

The second part covers the concept of parallelism and redundancy in storage system. We will apply the technology of Cloud storage systems to store and retrieve files in our local desktop cluster using a distributed hash table to locate files or file fragments and a REST interface to do GET, PUT or DELETE operations on these.

The exercises conclude with the implementation and performance tuning of a RAID verification algorithm.

Summary

Presenter: PETERS, Andreas (CERN)

Track Classification: Data Technologies