

DAOS as HPC Storage: Exploring Interfaces

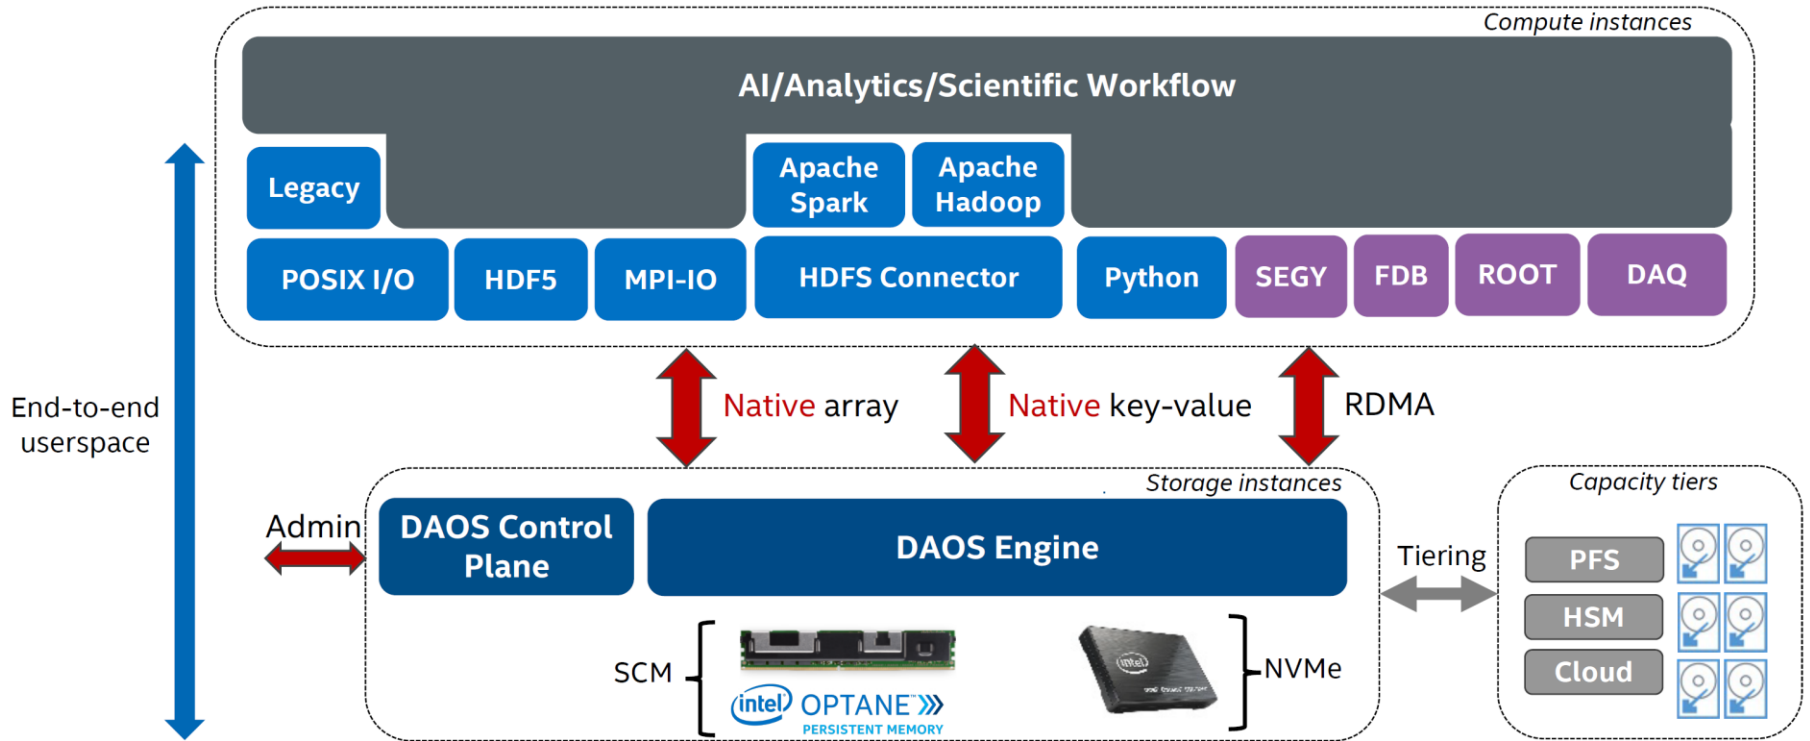
Adrian Jackson

a.jackson@epcc.ed.ac.uk

Nicolau Manubens

ECMWF/EPCC PhD Student

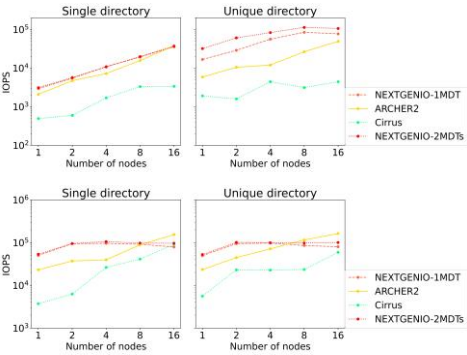

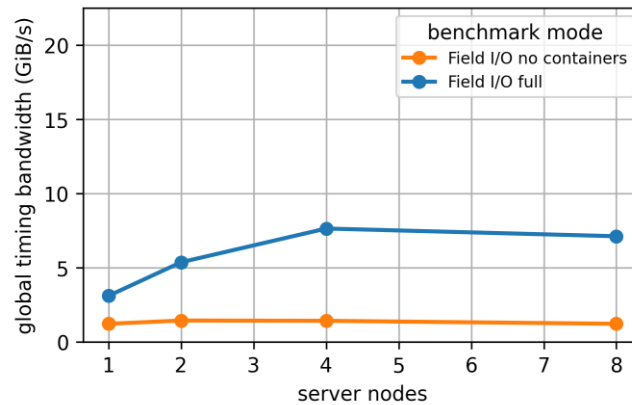
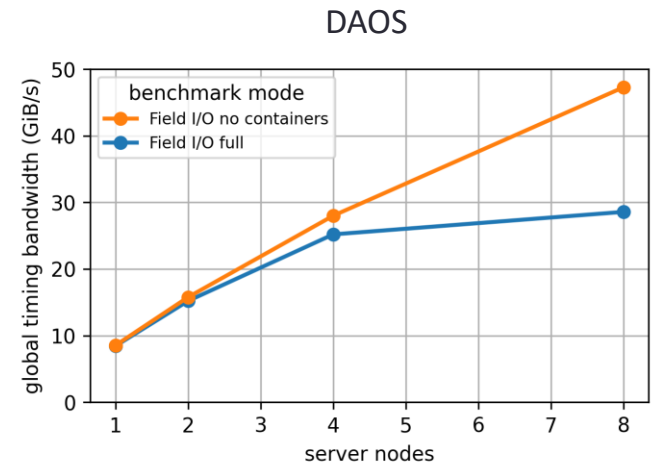
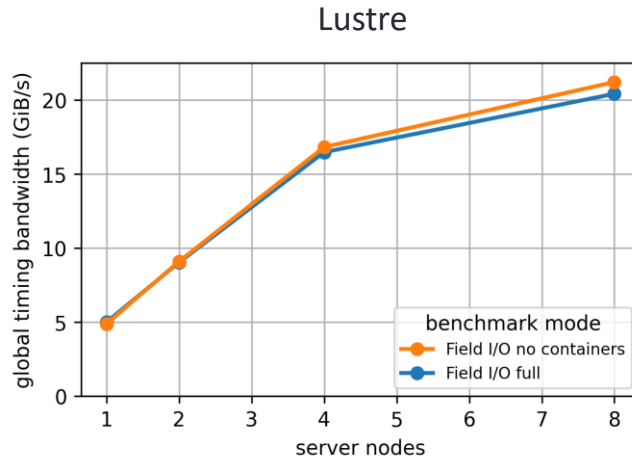
DAOS access models



- Path to solution with no persistent memory
 - Metadata in RAM
 - Write Ahead Log (WAL) on NVMe

1MB separate read/write phases

Read




Write



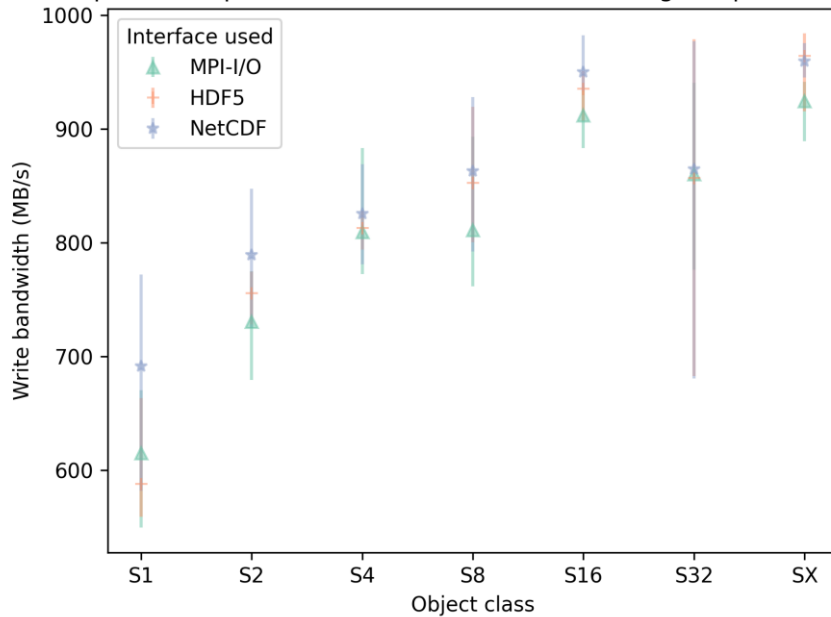
benchio

- Application-like I/O benchmark
- MPI distribution
- Different I/O interfaces
 - MPI-I/O
 - NetCDF
 - HDF5
 - ADIOS2
 - DAOS
 - POSIX

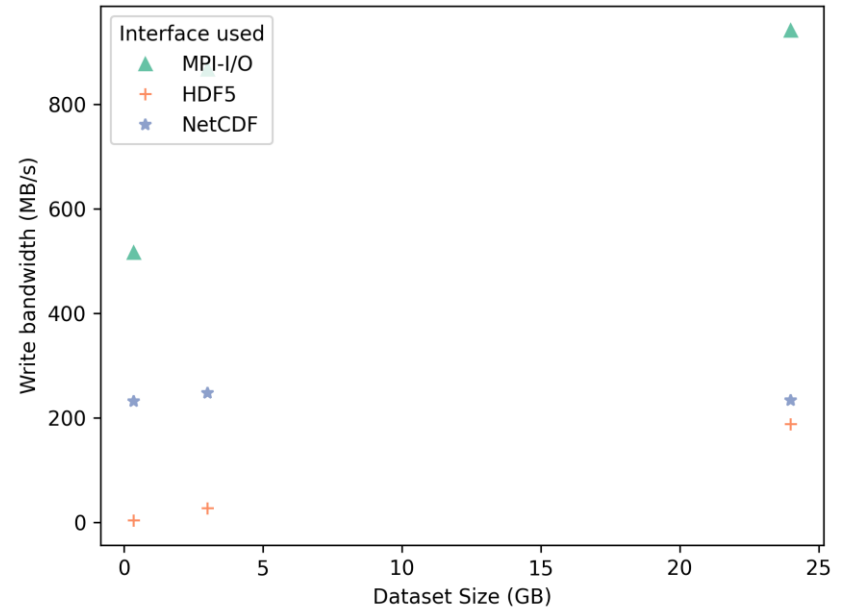


Interfaces/Configurations dFuse

benchio Performance on DAOS using 4 client nodes (48 processes per node) and 2 server nodes (2 engines per node)




benchio Performance on DAOS using 4 client nodes (48 processes per node) and 2 server nodes (2 engines per node) with a default container



Configurations/Performance DAOS API

benchio Performance on DAOS using 4 client nodes (48 processes per node) and 8 server nodes (2 engines per node)



benchio Performance on DAOS using 16 client nodes (48 processes per node) and 8 server nodes (2 engines per node)

