Coupled Storage System for Efficient Management of Self-Describing Data Formats (CoSEMoS)

ISC High Performance 2020

Michael Kuhn
michael.kuhn@informatik.uni-hamburg.de

2020-06-23

Scientific Computing
Department of Informatics
Universität Hamburg
https://wr.informatik.uni-hamburg.de

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

ISC High Performance

SHAPING TOMORROW

ISC 2020 DIGITAL

JUNE 22-25

#ISC20

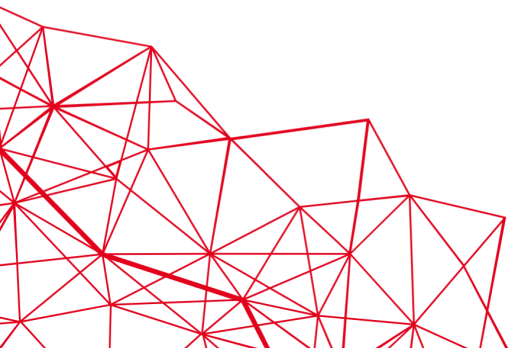

- Performance is determined by computation and I/O
  - Numerical applications often write output periodically
- Vast amounts of data written to parallel distributed file systems
  - Sizes in the range of hundreds of PB and throughputs of multiple TB/s
- Self-describing data formats (SDDFs) are widely used to exchange data
  - Structural information is encoded in the files themselves
  - Files can be accessed and interpreted without prior knowledge

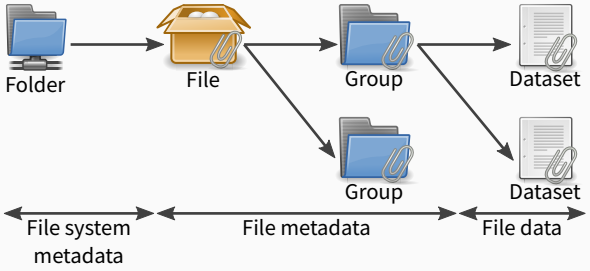

Problems

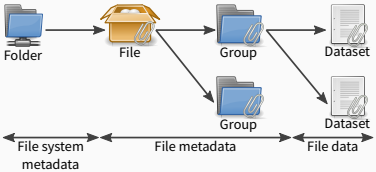

1. Weak treatment of different types of metadata
   - Two different types of metadata
     - File system metadata is stored on the metadata servers
     - File metadata is stored within files on the data servers
   - Strict separation of metadata leads to inefficient file access
     - Data servers are commonly optimized for streaming I/O
   - Reading data requires access to file system metadata, file metadata, and file data
     - Operations have to be performed sequentially
2. Static I/O semantics
3. Inefficient data placement

Approach

1. Global metadata management

- Storage system and self-describing data formats will be closely coupled
  - File system metadata and file metadata handled by the metadata servers
  - Optimize access to both types of metadata using database technologies
- Storage system can handle different kinds of metadata in an optimal way
  - Access to file metadata gives storage system further opportunities
- Novel data management approaches via a data analysis interface
  - Currently no unified way to link and query file metadata across multiple files
  - For instance, “average temperature over the last 12 months for all experiments”

2. Adaptable I/O semantics
3. Intelligent storage selection

Architecture

- Unmodified SDDF APIs
  - Backward compatible
  - Optional extensions
- Work in progress
  - HDF5 plugin and ADIOS2 engine
  - Database client/backend
- Data analysis interface
  - Direct access to backends

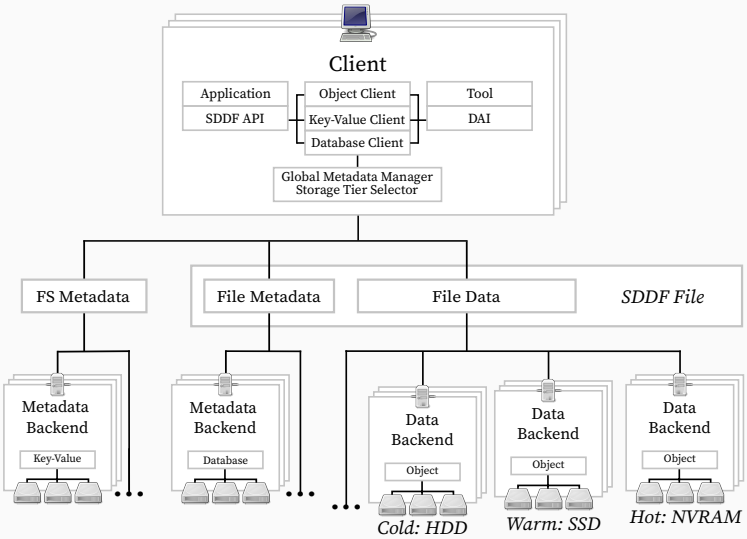

Questions and suggestions are always welcome!

Contact me at michael.kuhn@informatik.uni-hamburg.de

More information is available at https://cosemos.de/

SPONSORS

Follow us on Twitter at #ISC20 !

ISC 2020 DIGITAL MEDIA SPONSORS

ISC High Performance
SHAPING TOMORROW

ADMIN Network & Security
ADMIN HPC
ENTERPRISE AI NEWS & INSIGHTS FOR THE AI JOURNEY
HPC wire
inside HPC
LINUX MAGAZINE
Primeur Magazine
SCIENCE NODE™
SCIENTIFIC COMPUTING WORLD
THE NEXT PLATFORM
The Register®

Follow us on Twitter at #ISC20 !