EXIOS – Unified Exascale I/O System

Julian M. Kunkel, Michael Kuhn, Thomas Ludwig,
André Brinkmann,
Norbert Ritter,
Jens Dittrich,
Christian Schwede, Luis Kornblüh, Malcolm Muggeridge

julian.martin.kunkel@informatik.uni-hamburg.de

University of Hamburg,
University of Mainz,
Saarland University,
DKRZ, MPI-M, Xyratex

2012-02-28

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

informatik die zukunft

Are you satisfied with the current data access paradigm?

We are not.

Let’s talk about a vision for a future storage system.

Goals

Replace dumb block storage with an intelligent system

- Lightweight, modular & portable I/O architecture
- Semantical and domain specific data access
- Utilization of heterogeneous I/O systems

EXIOS Architecture

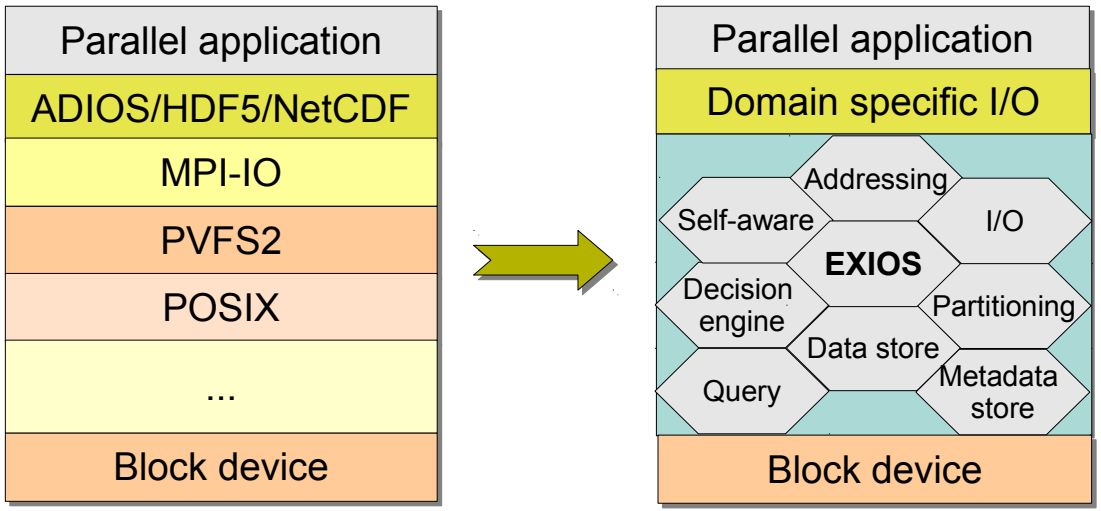

Flaws of the current I/O stack that are solved with EXIOS

- Reimplementation of features – just provide modules
  - Redundant optimizations
- Loss of semantical information through layers
- Portability issues – imagine EXIOS as MPI / POSIX

Semantical Data Access

Features of EXIOS

- Manages semi-structured data
- EXIOS understands content of “files” & metadata
- Abolish file system hierarchy
  - Access what you need (e.g. like with a search engine)
  - A mapping to the traditional namespace can be done
- Database alike “semantical” access
  - Complex data access patterns possible (compare to SQL)
  - Simplifies post-processing / access across traditional files
- Application specific access semantics (e.g. concurrency)
- High-level I/O can be implemented as thin layer above EXIOS

Utilization of Heterogeneous I/O Systems

Features of EXIOS

- Migrate, replicate and organize data across storage back ends:
  - Node-local discs, gold storage, tapes, databases, NoSQL, ...
- Self-aware: know the characteristics of the storage back ends
- Dynamic “on-disk format”
  - Do not worry about the location inside “byte streams” any more!
  - Exploit semantical knowledge/observed patterns automatically
- More self* properties

- This vision is very ambitious, we'll have a plan towards the goal
- We explicitly support climate science as a first use case
- Any help is welcome:
  - Whole development will be made available to the public
  - Open community for requirements, design, ...
  - Implementations for modules can be provided by anyone
- How would you like to access data in the future?
  - Let me/us hear your opinion...