

EXIOS

Unified Exascale I/O System

Topics: Data management and Exploration, Programming

Prof. Thomas Ludwig¹, Julian M. Kunkel, Michael Kuhn,
Prof. André Brinkmann²,
Prof. Norbert Ritter¹,
Prof. Jens Dittrich³
Xyratex, Max Planck Institute for Meteorology, DKRZ

¹ Universität Hamburg,

² Johannes Gutenberg-Universität Mainz,

³ Universität des Saarlandes

2012-07-09



Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

informatik
die zukunft

Aims of the Project

Mission

Developing a novel I/O access paradigm to utilize *exascale storage*¹

Goals

- Provide semantical & domain-specific access
- Develop a unified I/O architecture
- Prototype an intelligent storage system
- Utilize heterogeneous I/O landscapes

¹Massive parallelism and highest data volume

Semantical & Domain-Specific Access

Providing a convenient interface

- Close to application semantics
 - Natural data exploration
 - Direct support for complex inter-“file” I/O access
 - Support for different views based on scientific metadata
- ⇒ Capabilities beyond POSIX and hierarchical namespace

Enhanced storage system capabilities

- Native support for scientific data
- Integration of scientific metadata

Unified I/O Architecture

Unification into a modular architecture

- Fuse existing storage technologies
- Merge intermediate I/O layers

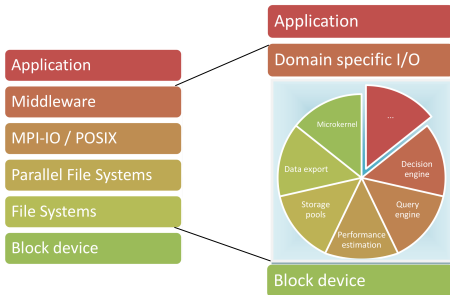


Figure: From layered architecture towards a modular architecture

Prototype an Intelligent Storage System

Intelligent storage

- Smart data placement
 - Replication / transformation / migration based on access patterns
- Performance-aware
 - Embedded performance estimation
 - Use the combination of storage technologies fitting best
- Query optimization

Development approach for the prototype

- Implementation: Leveraging existing solutions
- Climate science as the initial application driver
- Legacy support and ported applications / post-processing tools

Utilization Heterogeneous I/O Landscapes

EXIOS enables superior performance by

- Providing a better abstraction to scientific data
- Integrating intelligence into the storage
- Exploiting available information
- Combining expertise into a unified architecture

Community Effort

The acceptance of a novel access paradigm depends on the community!

- Integral collaborative approach & dissemination!
 - Initial collaboration: **EIOW/EOFS** and **seven SPPEXA** projects

