EXA4MIND

# Supercomputing, Databases and Data Ecosystems: The Best of Three Worlds for Efficient Data-Driven Workflows

Stephan Hachinger

Funded by
the European Union

Visit us at exa4mind.eu

# EXA4MIND Consortium / Project: Who we are.

# EXA4MIND consortium
# “Extreme Data Analytics for Mining Data Spaces”

- 10 partners, 6 countries, 40 months
- EU contribution: € 4.9M
- Overall budget: € 6M
- Coordinator

  Jan Martinovič (IT4I)
- Science & Co-Design Coordinator

  Stephan Hachinger (LRZ)

# What and why?

# A fast gap analysis...

Big Data Analytics –
Hot Topic Since Years

... yields great results with AI and database technology; yields more with more computing power.

# A fast gap analysis...

Big Data Analytics –
Hot Topic Since Years

... yields great results with AI and database technology; yields more with more computing power.

Excellent Supercomputing Facilities with HPC, Cloud, GPU and Quantum Systems

... do simulations and complex workflows but not so much work on Databases & Extreme Data Analytics.

# A fast gap analysis...

Big Data Analytics –
Hot Topic Since Years

... yields great results with AI and database technology; yields more with more computing power.

Excellent Supercomputing Facilities with HPC, Cloud, GPU and Quantum Systems

... do simulations and complex workflows but not so much work on Databases & Extreme Data Analytics.

European Data Spaces, EOSC and EUDAT

... can do great research data management but rarely address general Big/Extreme Data.

# A fast gap analysis...

Big Data Analytics –
Hot Topic Since Years

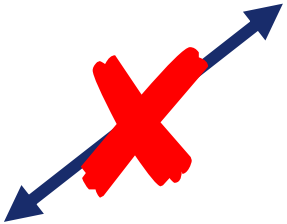

... yields great results with AI and database technology; yields more with more computing power.

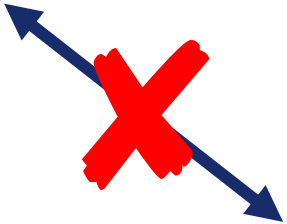

Excellent Supercomputing Facilities with HPC, Cloud, GPU and Quantum Systems

European Data Spaces, EOSC and EUDAT

... do simulations and complex workflows but not so much work on Databases & Extreme Data Analytics.

... can do great research data management but rarely address general Big/Extreme Data.

# … motivating EXA4MIND

Supercomputing


Computational Workflows

LEXIS2 Platform

EU Data Ecosystems

EUDAT

# … motivating EXA4MIND





DBMS & Object Stores
Advanced Queries
Analytics




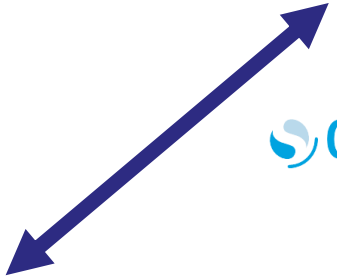

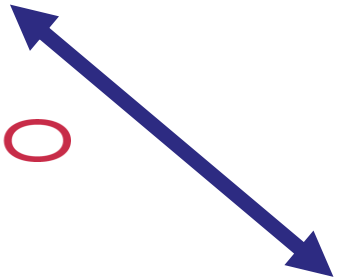

Supercomputing


Computational Workflows


EU Data Ecosystems

# … motivating EXA4MIND

Milvus

iRODS

DBMS & Object Stores
Advanced Queries
Analytics

OpenSearch

MariaDB

MINIO

Supercomputing


Computational Workflows

LEXIS2 Platform

EU Data Ecosystems

Supporting eosc

EUDAT

DATA SPACES SUPPORT CENTRE

# What do we provide?

Tested concepts & boilerplates to set up your *Extreme Data /* computing workflows with the best data backends and data-sharing possibilities.

EXA4MIND

## Dataset Connectivity and FAIR support

### Dataset Index and FAIR Portal (Portals with PID Support)

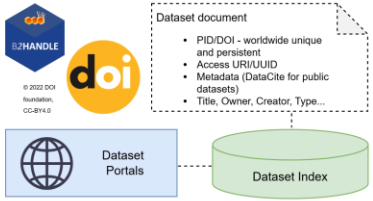


### ID-to-Dataset (and Version) Resolution

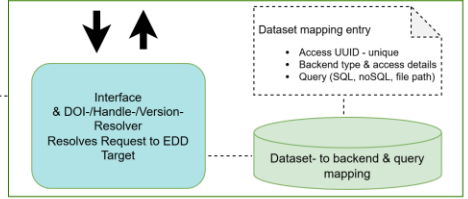


### Dataspace Connectors and Adaptors

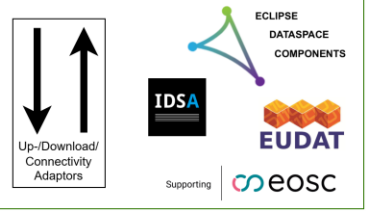


## Toolboxes

### Validation Toolbox


### Preprocessing Toolbox


### Analytics & AI Toolbox


## AQIS

### Workflow Catalogue

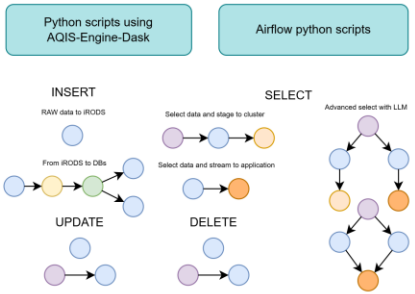


### AQIS Engine (with Interfacing)


### Database Catalogue


### Inference

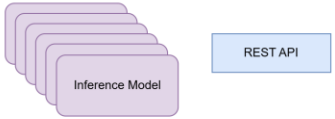


### Cache Support

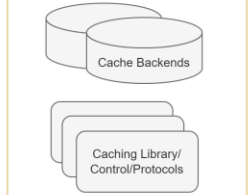


### Data System Adaptors

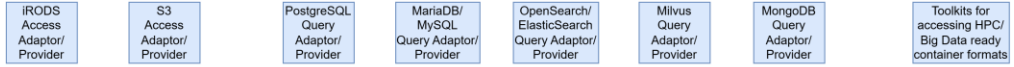


## Database / Object Store / Storage Layer - Data Systems Recipes

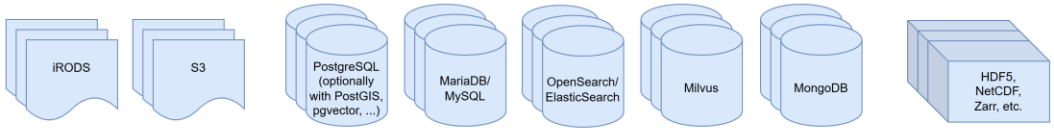


## Compute

### HPC / Cloud / AI Systems Access

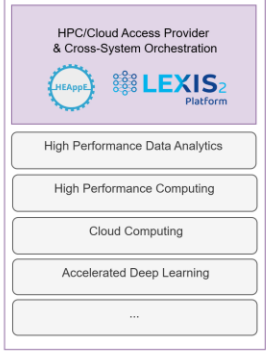


### Cross-System Data Movement

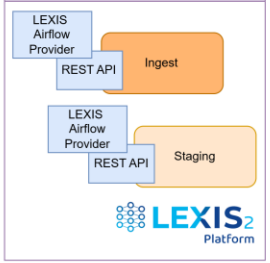

# A „flexibly-deployable platform“: selecting modular parts & instantiating them to implement …

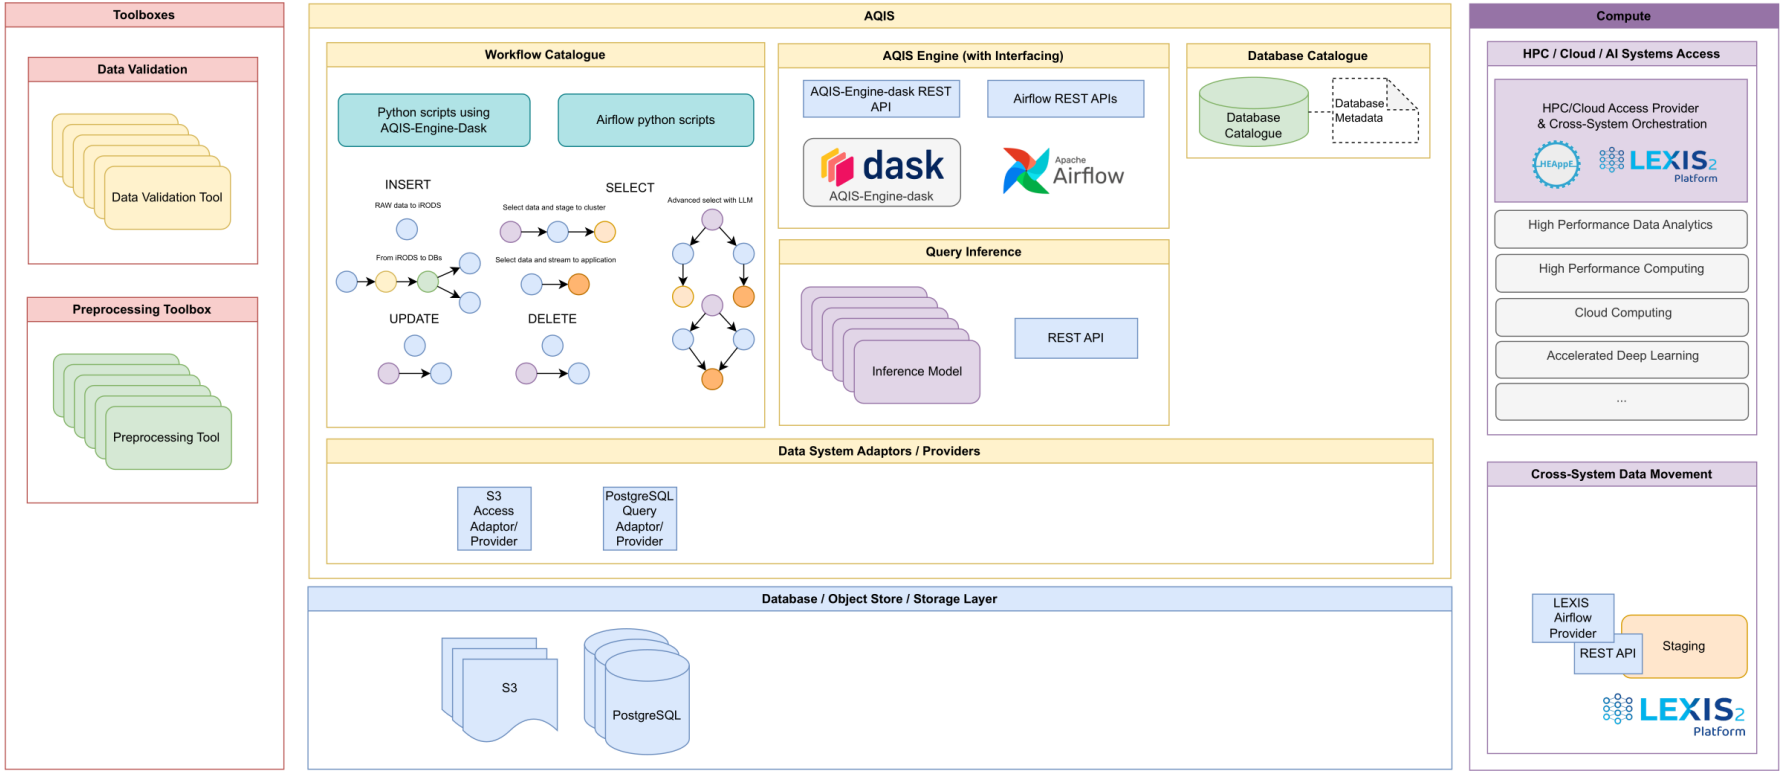

# … as an example …

… a website where engineers can explore simulation data & trigger more HPC simulations.

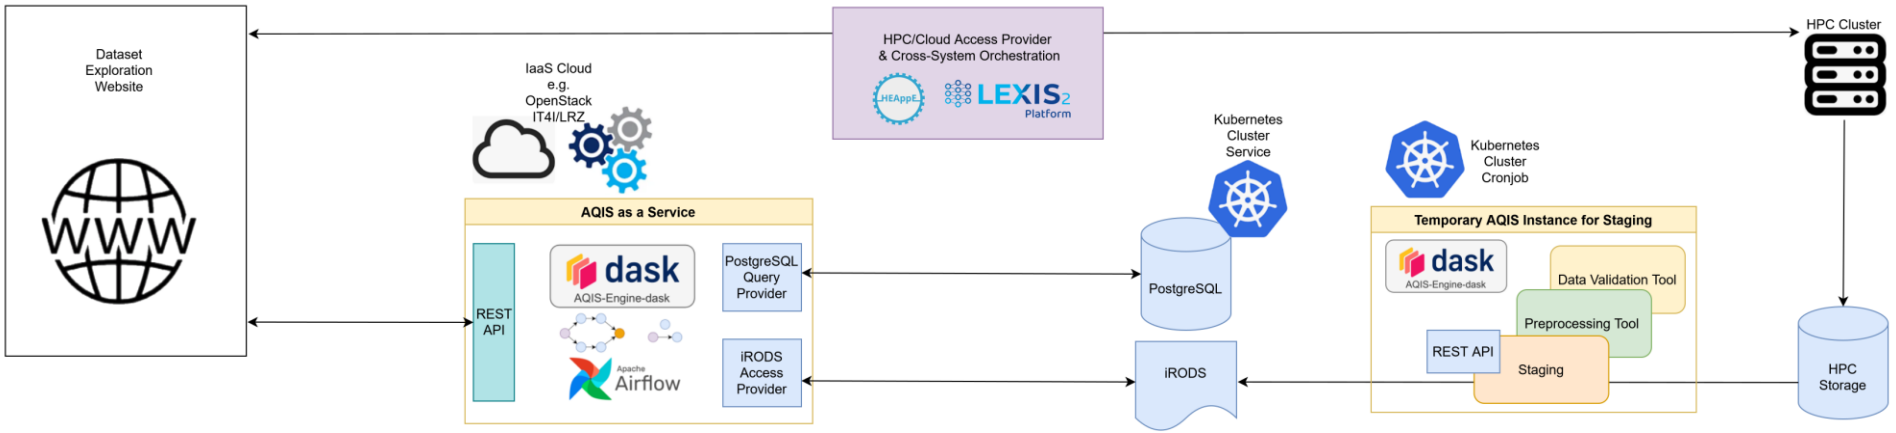

# Application cases within the project

## SCIENTIFIC – Molecular simulations
WP4

ADAMS4SIMS, an innovative Automated Data Mining System for Systematic Improvement of Molecular Simulations accuracy and predictability will combine services and on-/offline data sources.

Data: 0.1-1 PB range

## INDUSTRY – Automotive Car Sensor Recordings
WP5

Advanced Driver Assistance Systems (ADAS) need Extreme Data annotation, search and analytics. Methods, AI models and framework are co-developed in EXA4MIND.

Data: 1-100 PB range

## SME – Smart farming / viticulture
WP6

Easy and quick ingestion of large amounts of agriculture focused data, as well as efficient and space-saving re-distribution of computed/ingested (meta-)data.

Data: 0.01-1 PB range

## SME – Health
WP6

Demonstration of viability of Extreme Data Database (EDD) for Extreme Data processing use cases with health and societal datasets, based on direct user interaction.

Data: 0.01-0.1 PB range

# Platform releases

opencode.it4i.eu/exa4mind

docs.exa4mind.eu

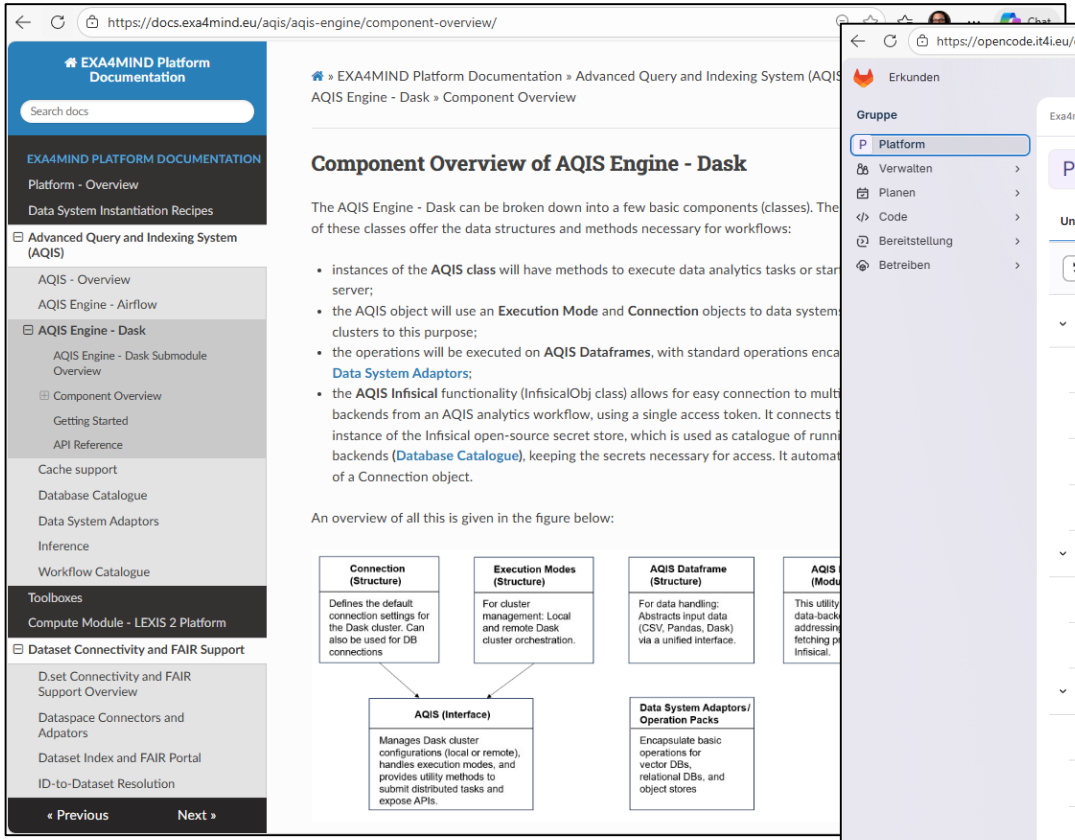


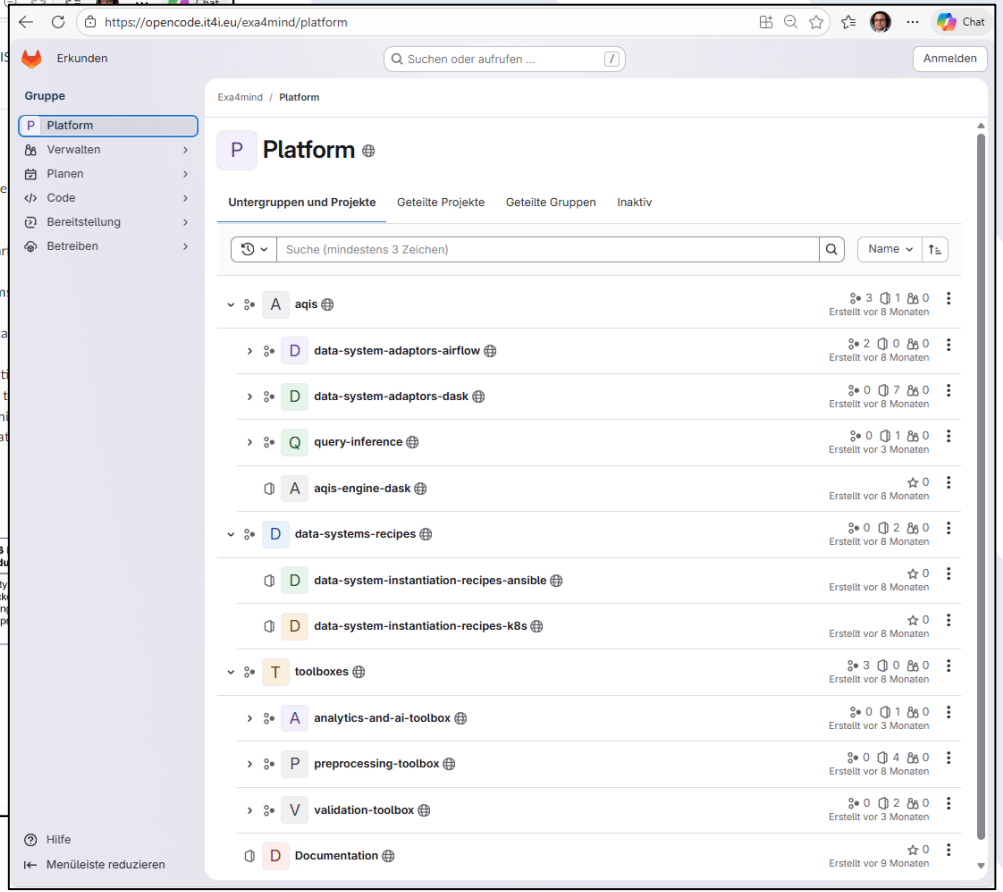

Are you serious – do you really wanna use such stuff on HPC?

# Yes, we are serious about understanding a wealth of data backends ☺

## Why diversify data backends at supercomputing centres?

- More and more customers want it (used to Cloud/Business world maybe)
- Databases can have advantages, e.g.:
  - Data indexing for fast access
  - Vector databases for similarity search
- Object stores can have advantages:
  - Data sharing functionalities
  - Codes can stay the same as in Cloud


EXA4MIND

Tapping Data Spaces in HPC:
Enabling, Exploring and Benchmarking Database & Object Store Usage

M. Hayek[a], S. Hachinger[a], H. Gurbanov[a], V. Pauw[a], M. Golasowski[b], J. Martinovič[b], P. Karagoz[c], I.H. Toroslu[c]

[a]Leibniz Supercomputing Centre (LRZ), Bavarian Academy of Sciences & Humanities, Garching near Munich, Germany; [b]IT4Innovations National Supercomputing Center (IT4I), VŠB – Technical University of Ostrava, Czech Republic; [c]Department of Computer Engineering, Middle East Technical University (METU), Ankara, Turkey

OUR RESEARCH

UNDERSTANDING THE LIMITATIONS OF OBS

ENVIRONMENTS USED FOR BENCHMARKING

MAIN TAKEAWAY FROM BENCHMARKS

IS DBMS USAGE ON A HPC SYSTEM POSSIBLE?

OVERCOMING DATA ANALYTICS OBSTACLES WITH EXA4MIND

OBS USAGE ON A HPC SYSTEM

REFERENCES & FURTHER INFORMATION

ACKNOWLEDGEMENT

Poster: Hayek et al, SC Asia ‘26

# Benchmarking of DBMS and OBS on HPC clusters

| Target system / Direction | *PostgreSQL DBMS* on LRZ Compute Cloud | *PostgreSQL DBMS* on CoolMUC-4 | CoolMUC-4 Scratch *cp* (for comparison) |
|---|---|---|---|
| Rate from system to CoolMUC-4 Scratch | 54.6 MB/s | 792 MB/s | 2610 MB/s |
| Rate from CoolMUC-4 Scratch to system | 28,3 MB/s | 24,0 MB/s | 2610 MB/s |

**Table 1.** Data rates obtained communicating with an LRZ-Cloud or internal PostgreSQL server from the LRZ CoolMUC-4 HPC system. Colours indicate low (red) to high (green) rates. Rates for a file copy (cp) from the Scratch file system to itself are shown for comparison (right col.). A "SELECT *" type query is used for reading, and a "COPY" for writing.

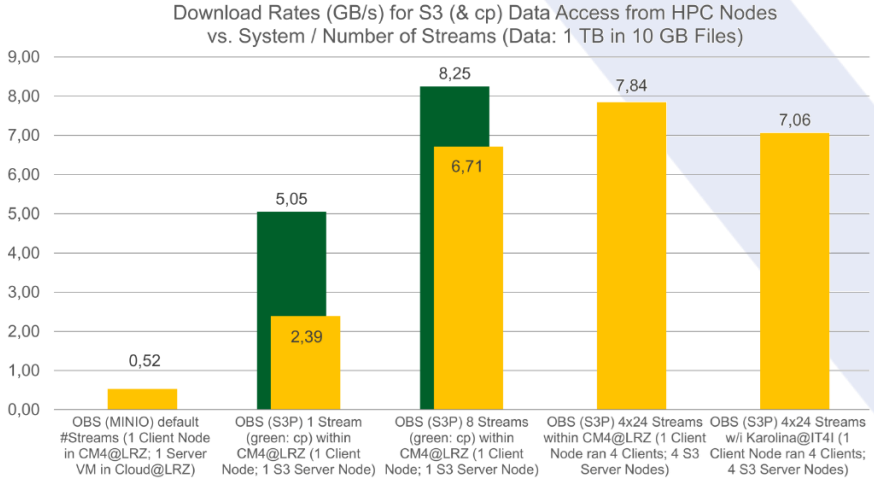


**Figure 2.** Rates (yellow bars) retrieving data from S3Proxy OBS (MINIO where noted) on nodes of LRZ's CoolMUC-4 and IT4I's Karolina (rightmost bar). We use s5cmd as a client and a thread number as indicated (number of streams); rates obtained with one or eight cp processes/ streams (green bars) copying the same data from SCRATCH to /dev/null are given to roughly indicate max. read bandwidth. In case of 4x24 streams, four server nodes are used and four client instances (24 threads each) on one node, to test whether client-node network bandwidth can be utilized.

# Object stores for warm data enable federation & sharing

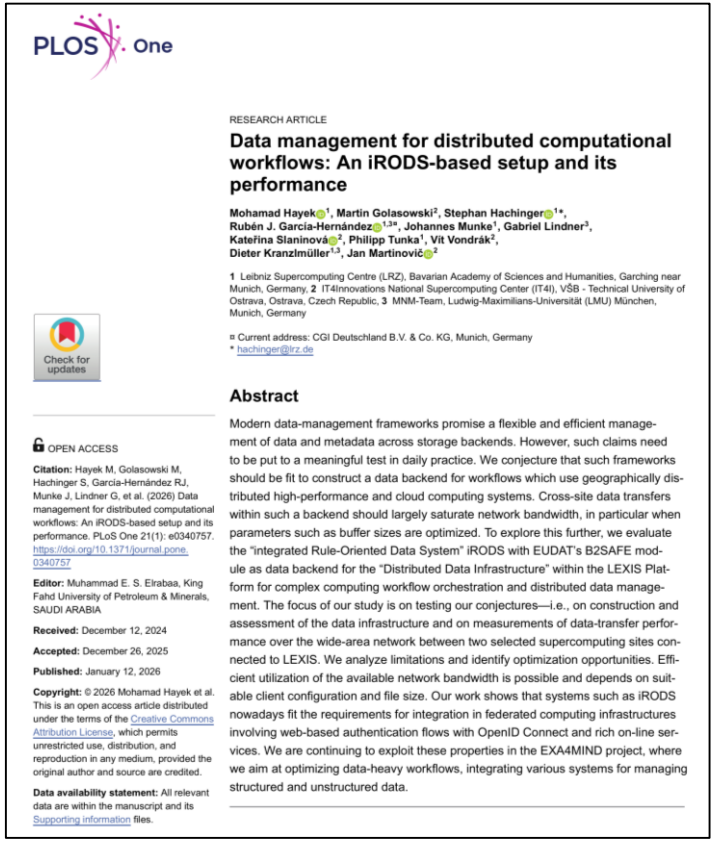

PLOS One

RESEARCH ARTICLE

**Data management for distributed computational workflows: An iRODS-based setup and its performance**

**Mohamad Hayek[1], Martin Golasowski[2], Stephan Hachinger[1]*, Rubén J. García-Hernández[1,3¤], Johannes Munke[1], Gabriel Lindner[3], Kateřina Slaninová[2], Philipp Tunka[1], Vít Vondrák[2], Dieter Kranzlmüller[1,3], Jan Martinovič[2]**

**1** Leibniz Supercomputing Centre (LRZ), Bavarian Academy of Sciences and Humanities, Garching near Munich, Germany, **2** IT4Innovations National Supercomputing Center (IT4I), VŠB - Technical University of Ostrava, Ostrava, Czech Republic, **3** MNM-Team, Ludwig-Maximilians-Universität (LMU) München, Munich, Germany

¤ Current address: CGI Deutschland B.V. & Co. KG, Munich, Germany
* hachinger@lrz.de

Check for updates

OPEN ACCESS

**Citation:** Hayek M, Golasowski M, Hachinger S, García-Hernández RJ, Munke J, Lindner G, et al. (2026) Data management for distributed computational workflows: An iRODS-based setup and its performance. PLoS One 21(1): e0340757. https://doi.org/10.1371/journal.pone.0340757

**Editor:** Muhammad E. S. Elrabaa, King Fahd University of Petroleum & Minerals, SAUDI ARABIA

**Received:** December 12, 2024

**Accepted:** December 26, 2025

**Published:** January 12, 2026



**Data availability statement:** All relevant data are within the manuscript and its Supporting information files.

## Abstract

Modern data-management frameworks promise a flexible and efficient management of data and metadata across storage backends. However, such claims need to be put to a meaningful test in daily practice. We conjecture that such frameworks should be fit to construct a data backend for workflows which use geographically distributed high-performance and cloud computing systems. Cross-site data transfers within such a backend should largely saturate network bandwidth, in particular when parameters such as buffer sizes are optimized. To explore this further, we evaluate the "integrated Rule-Oriented Data System" iRODS with EUDAT's B2SAFE module as data backend for the "Distributed Data Infrastructure" within the LEXIS Platform for complex computing workflow orchestration and distributed data management. The focus of our study is on testing our conjectures—i.e., on construction and assessment of the data infrastructure and on measurements of data-transfer performance over the wide-area network between two selected supercomputing sites connected to LEXIS. We analyze limitations and identify optimization opportunities. Efficient utilization of the available network bandwidth is possible and depends on suitable client configuration and file size. Our work shows that systems such as iRODS nowadays fit the requirements for integration in federated computing infrastructures involving web-based authentication flows with OpenID Connect and rich on-line services. We are continuing to exploit these properties in the EXA4MIND project, where we aim at optimizing data-heavy workflows, integrating various systems for managing structured and unstructured data.

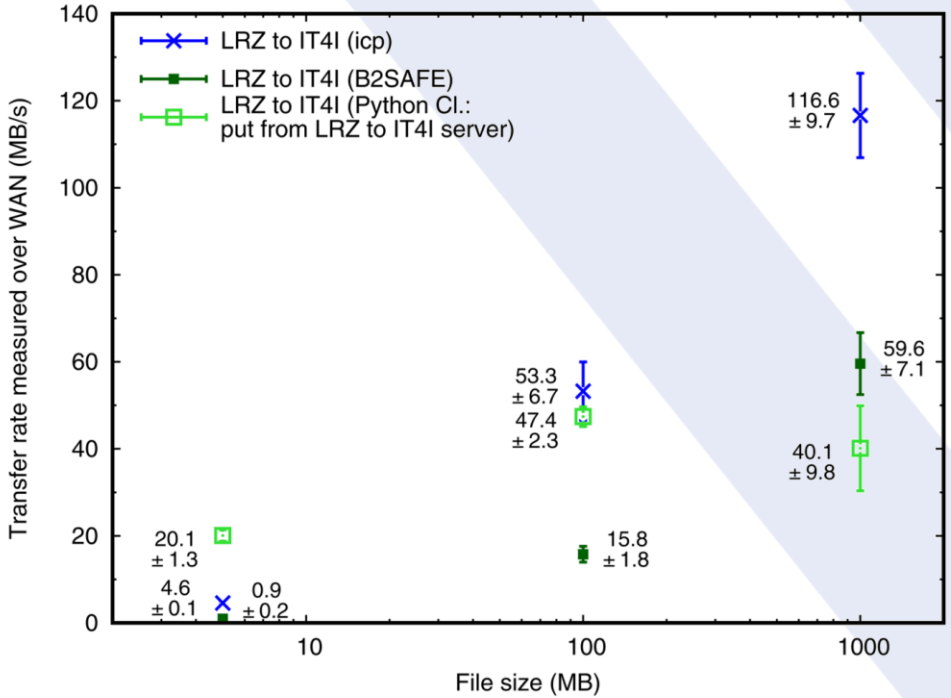


cf. iperf3 in tested env: 140-180 MB/s

For going HPC, EXA4MIND has staging tools

Some speed tests, is that all you show?

# ADAMS/IDA4SIMS – a molecular dynamics platform for comparing simulations to experiments and improving them

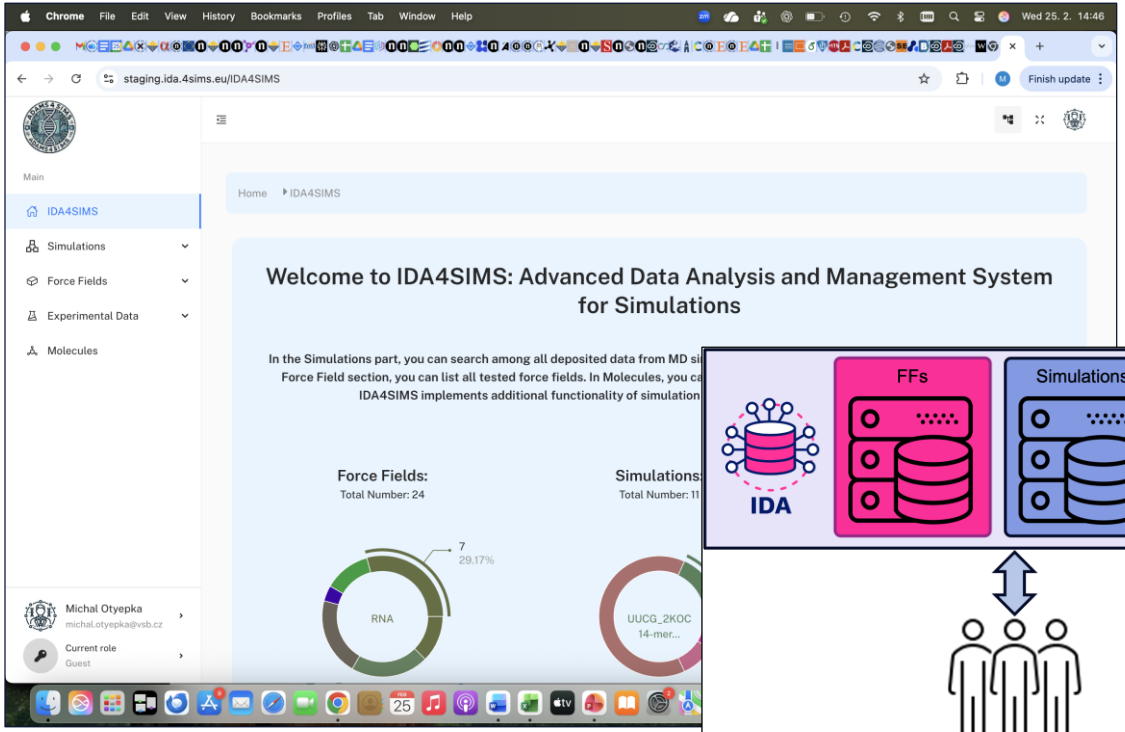


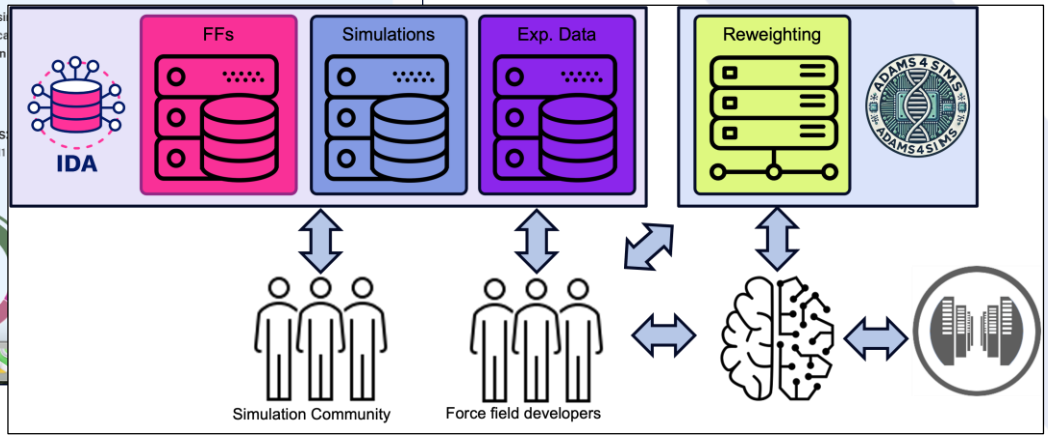

# Self-driving cars

## Workflows on HPC for image segmentation, annotation, …

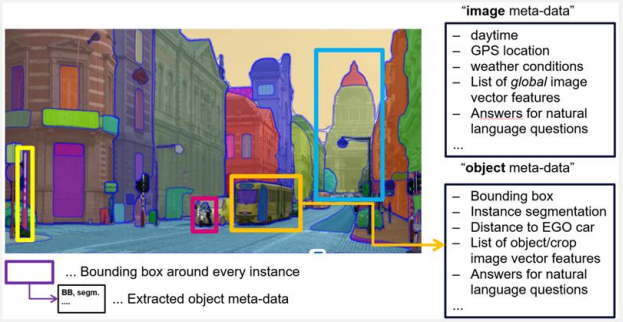


- Extract automatically information about scene
- Including object locations and classes and properties via AI

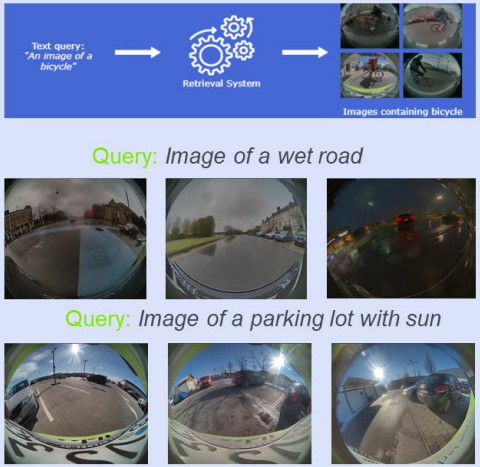


- Use natural language querying on top of extracted metadata
- And retrieve results from petabytes of data in realtime

## (Natural language) query capabilities for system validation

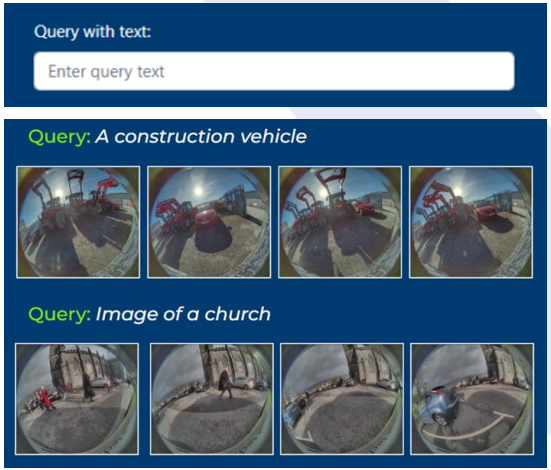


## EXA4MIND also helps with data transfer from test vehicles and company systems to HPC cluster

# Smart viticulture

Aquaview: Soil moisture prediction from satellite and weather data

EXA4MIND: satellite-data cache (iRODS + STAC) with Cloud-Optimised GeoTIFFs.

More: Friday @EnvComp Workshop.

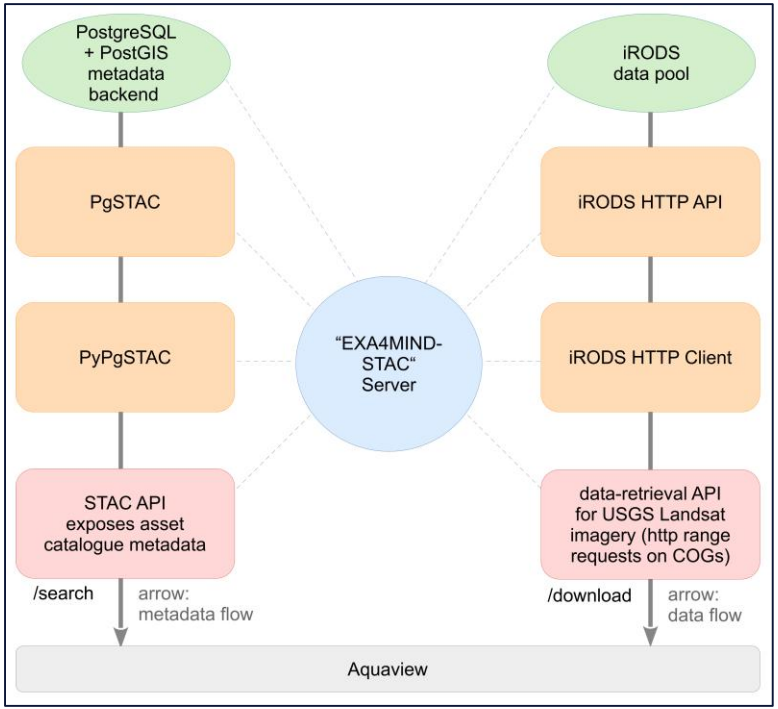


| Metric | Existing (PROD) | iRODS | iRODS Improvement |
|---|---|---|---|
| Total processing time | 4m 33.80s | 40.50s | 6.76× faster |
| Total data transferred | 20.40 GB | 0.26 MB | 79.522× less |
| Mean data downloaded per parcel | 77.50 MB | 0.90 KB | 100.00% reduction |
| Mean time per parcel | 16.6s | 2.4s | 6.80× faster |
| Std dev time per parcel | 36.80s | 10.80s | — |
| Weighted time per hectare | 1.02 s/ha | 0.15 s/ha | 6.64× faster |
| Mean STAC query time per parcel | 0.36s | 0.16s | 2.20× faster |

# Thanks for your attention – please test our modules!

Visit us at exa4mind.eu

EXA4MIND on YouTube

Project Site

E4M Kickstarter Initiative

Open Source Repository

Funded by the European Union

Keep in touch! @exa4mind /exa4mind