Karlsruhe Institute of Technology

Software and Services for Structured Data
Large Scale Data Management LSDF

Rainer Stotzka,
D. Tonne, T. Jejkal, V. Hartmann, H. Pašić, T. Schmidt, F. Rindone, S. Ochsenreither, M. Weber
W. Mexner, M. Hagelstein, P. Vagovic, T. dos Santos Rolo, T. Farago
J. van Wezel, M. Hardt, A. Garcia, R. Kupsch, S. Bourov, A. Streit

Institute for Data Processing and Electronics | Institute for Synchrotron Radiation | Steinbuch Centre for Computing

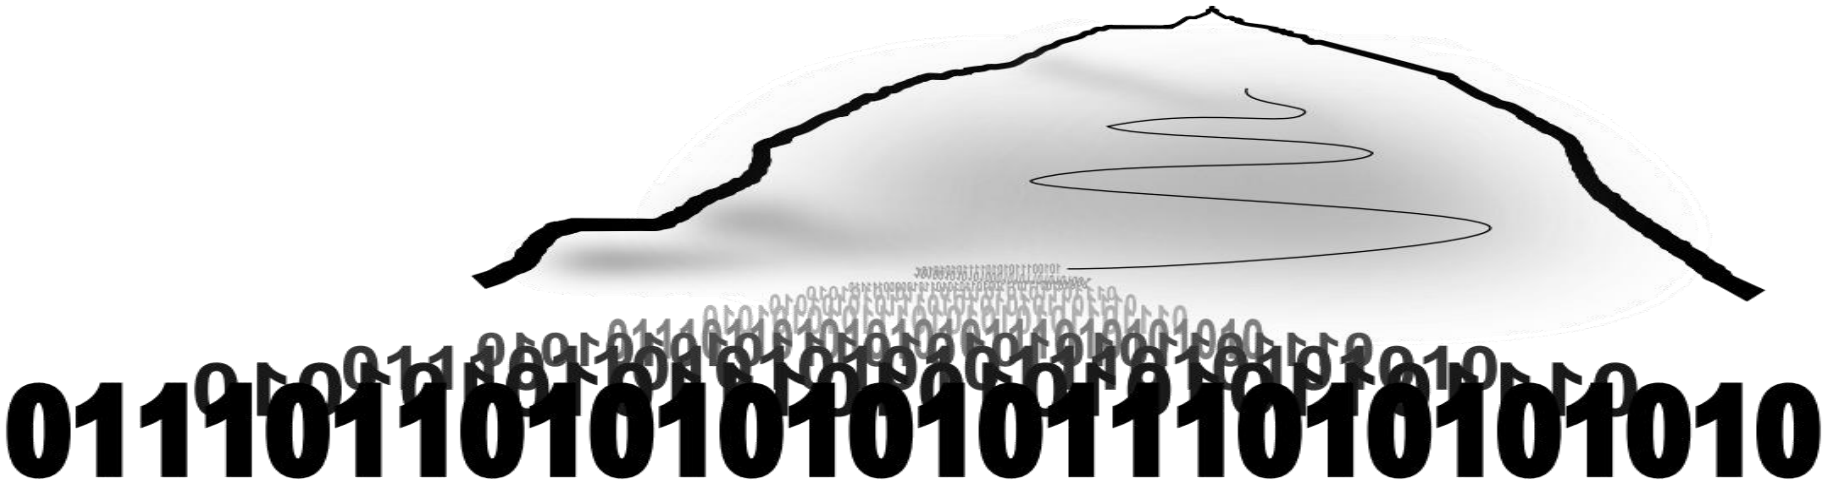

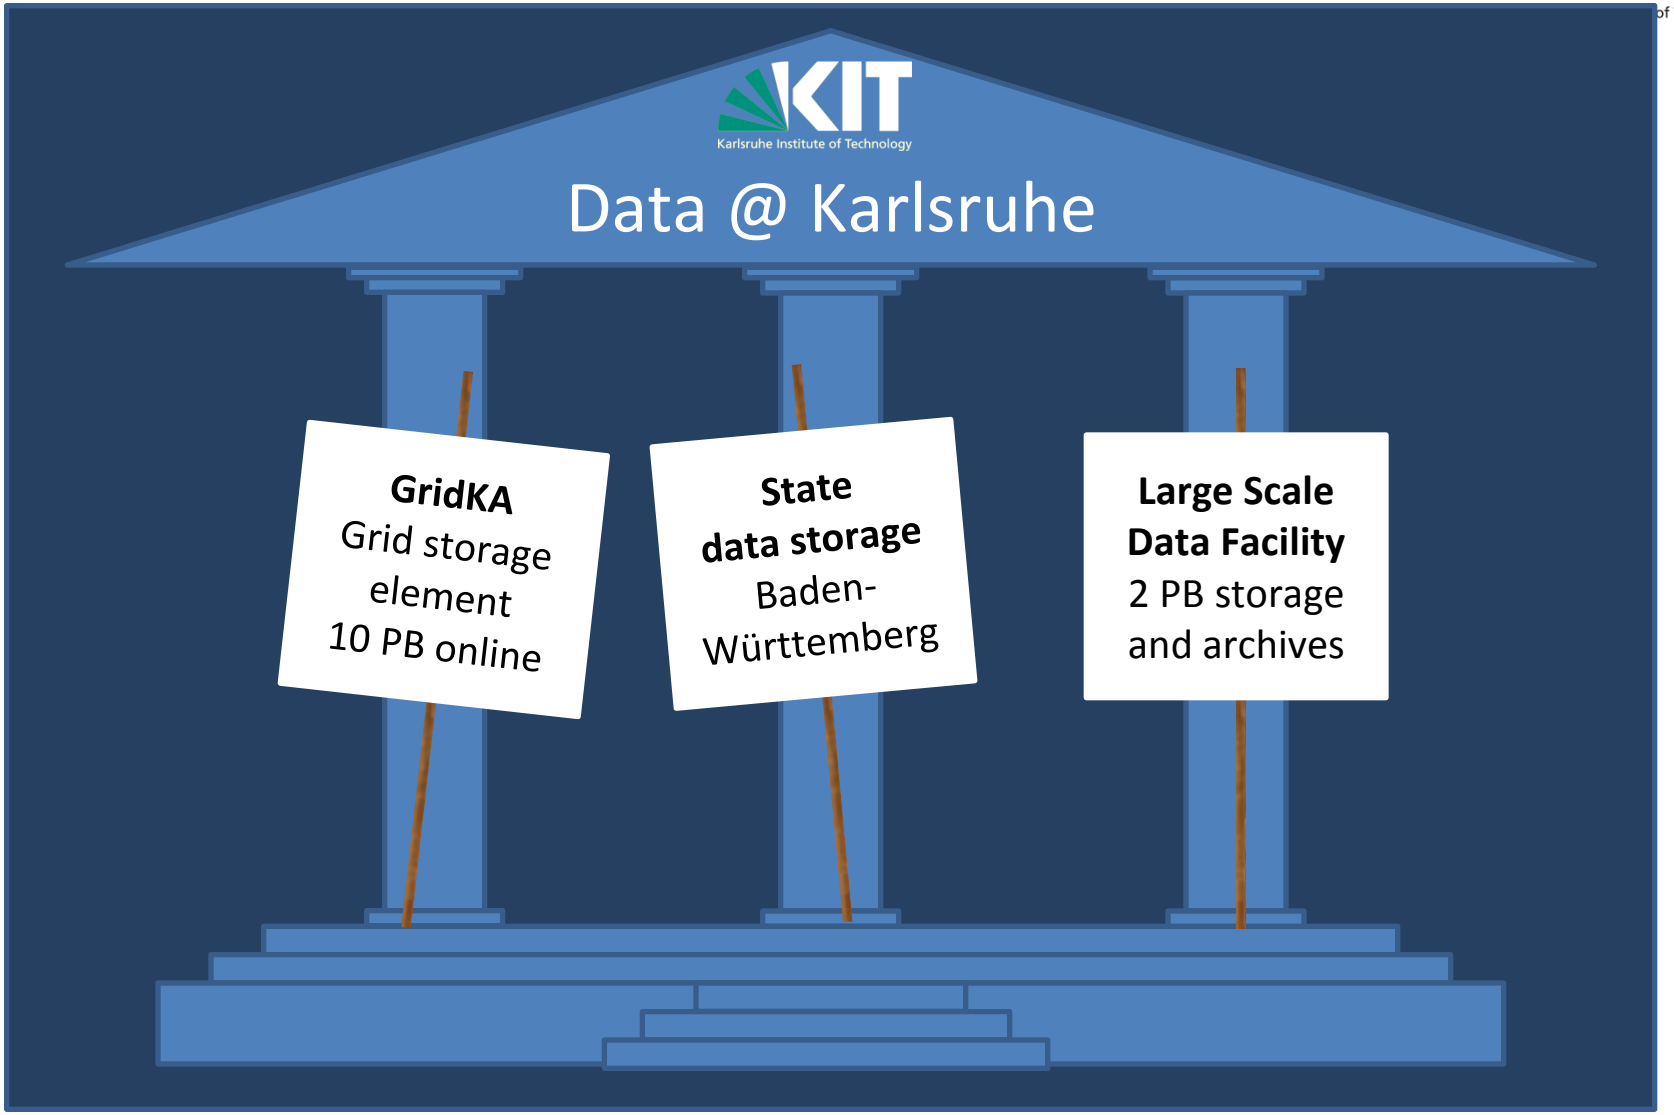
Karlsruhe Institute of Technology
Data @ Karlsruhe
GridKA
Grid storage element
10 PB online
State data storage
Baden-Württemberg
Large Scale Data Facility
2 PB storage and archives

LSDF objectives

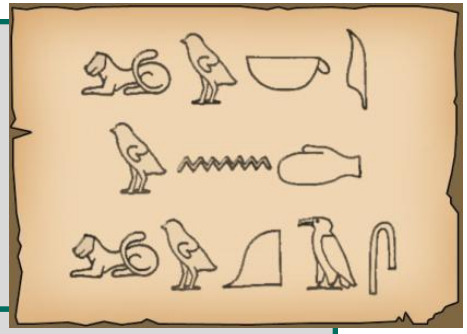

Storage

- Dedicated for science data
- ExaByte scale data
- To archive data, long term sustainability (10 yrs. – ?)

Interactivity

- To enable scientists to gain better scientific results by providing
  - Data intensive analysis
  - Added value services for data intensive processing
- To provide high performance access, high throughput
- “Barrier free” access (easy-to-use)
- Sustainability and interoperability

“Guidelines”:

- PARADE White Paper (2009): Strategy for a European Data Infrastructure
- ESFRI Data Management Task Force (2009): e-IRG Report on Data Management
- OAIS (2002): Reference Model for an Open Archival Information System
- High Level Experts Group (2010): Riding the Wave – European Commission Report on Scientific Data
- HLEG-SD (2010?): Note on Data Services infrastructure
- Microsoft Research (2009): The Fourth Paradigm: Data-Intensive Scientific Discovery

Focus: ANKA Image BLs

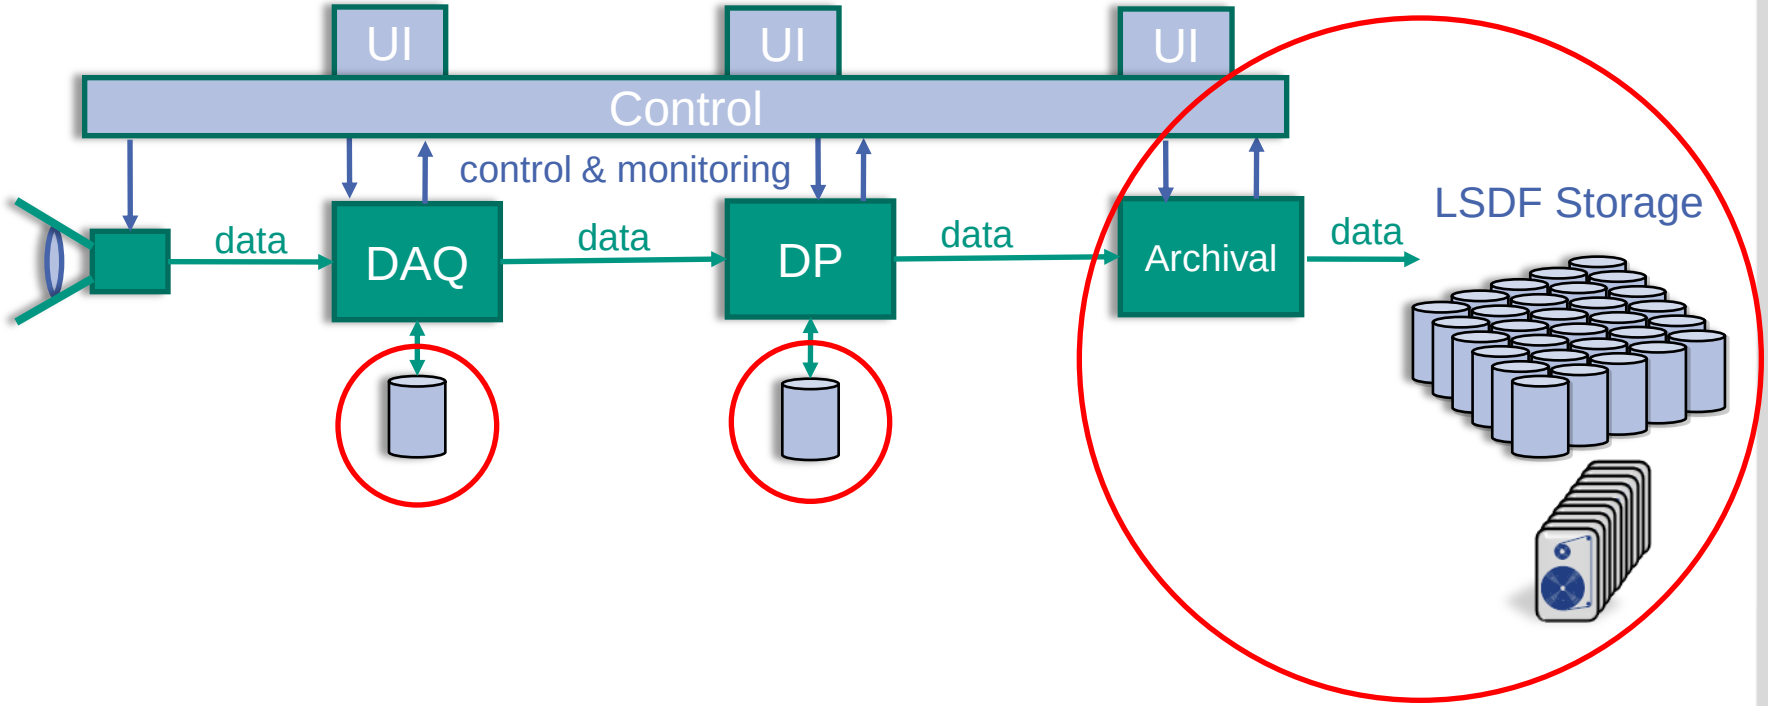

Unstructured vs. Structured Data

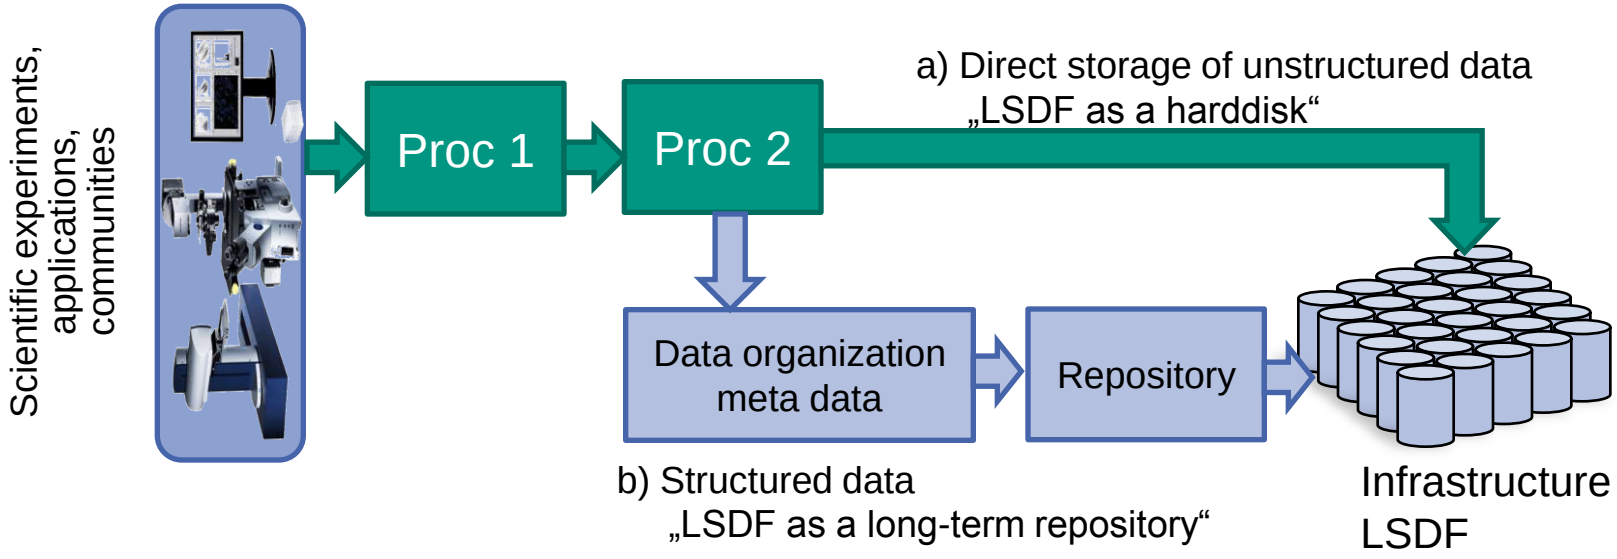

a) Unstructured data

- Harddisk
- “Content“ and “Structure” unknown to provider
- User is responsible for data life cycle and curation
- Backup, seldom bit preservation
- Local access

b) Structured data

- Data is archived with known meta data and data organization
- Clear separation between storage organization, data organization and project meta data
- Provider is responsible for life cycle, preservation and curation
- World-wide access

Separation

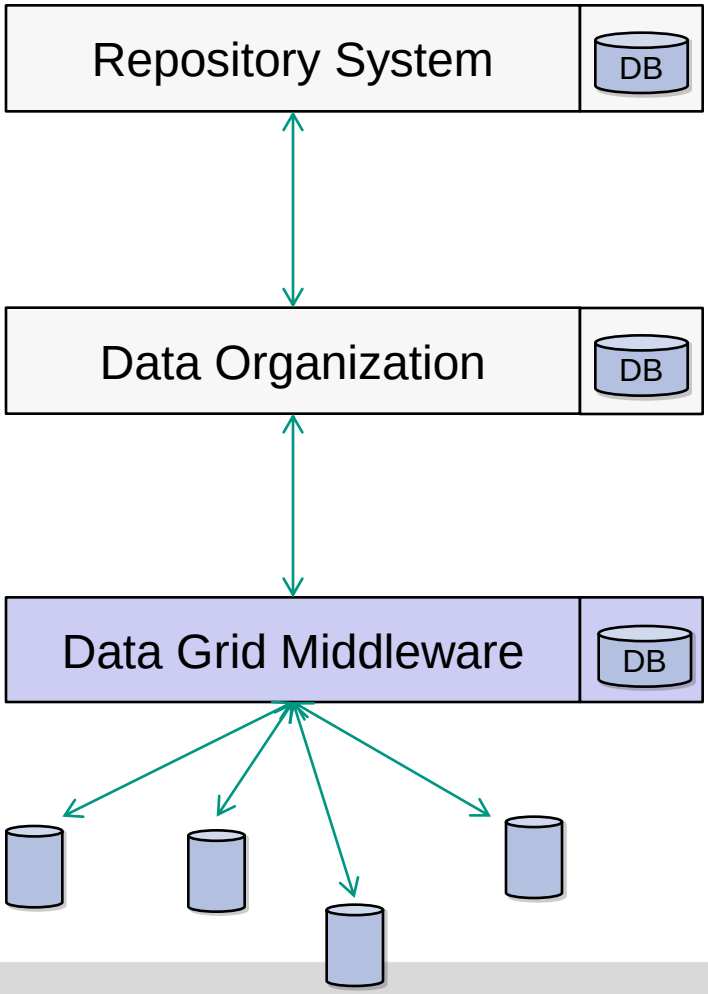

Scientific meta data
(KIT DataBrowser, Fedora repository, dSpace,…)

Content structure and
file meta data

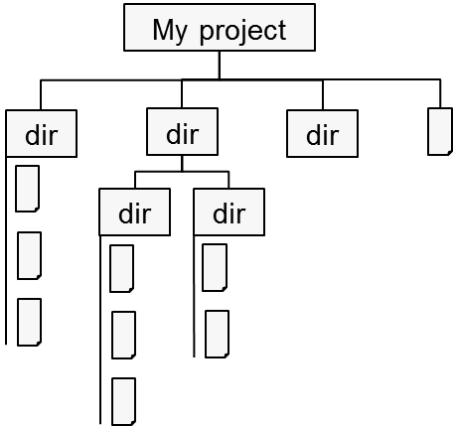

Virtualization of storage resources
(dCache, iRODS, Merritt, Active Circle…)

Storage resources,
heterogeneous, distributed

DataBrowser for structured data archives

Hierarchical Catalog System for Structured Data (Repository system)

- Sustainable
- Easily extensible
- Independent of data formats
- Enhanced performance: distribution of access
- Safety by redundancy

- Easy-to-use?

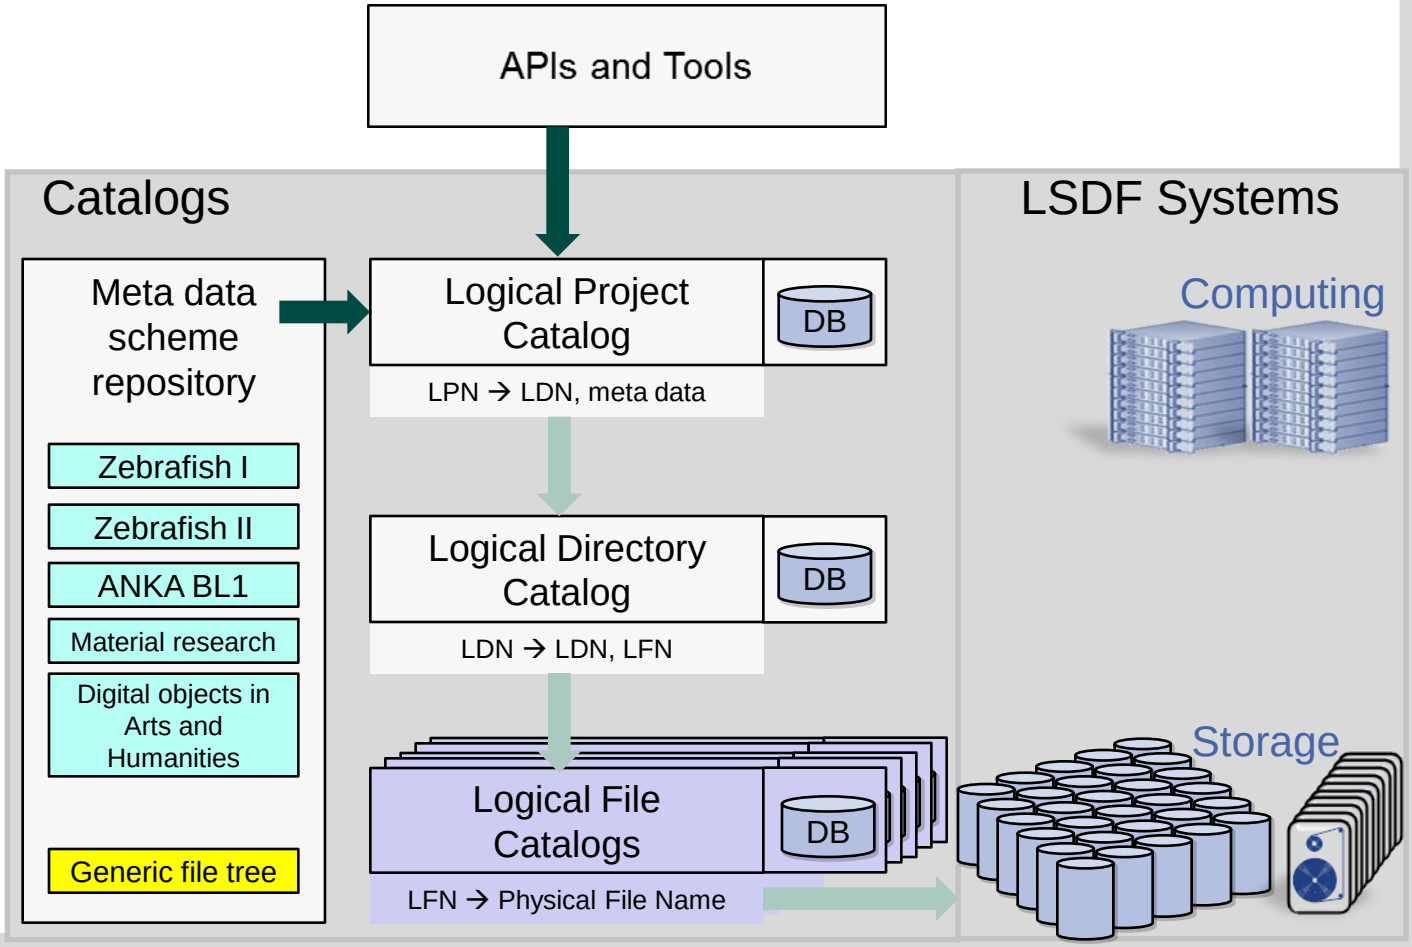

KIT DataBrowser

Repository system for large scale data
- High-throughput for large data sets

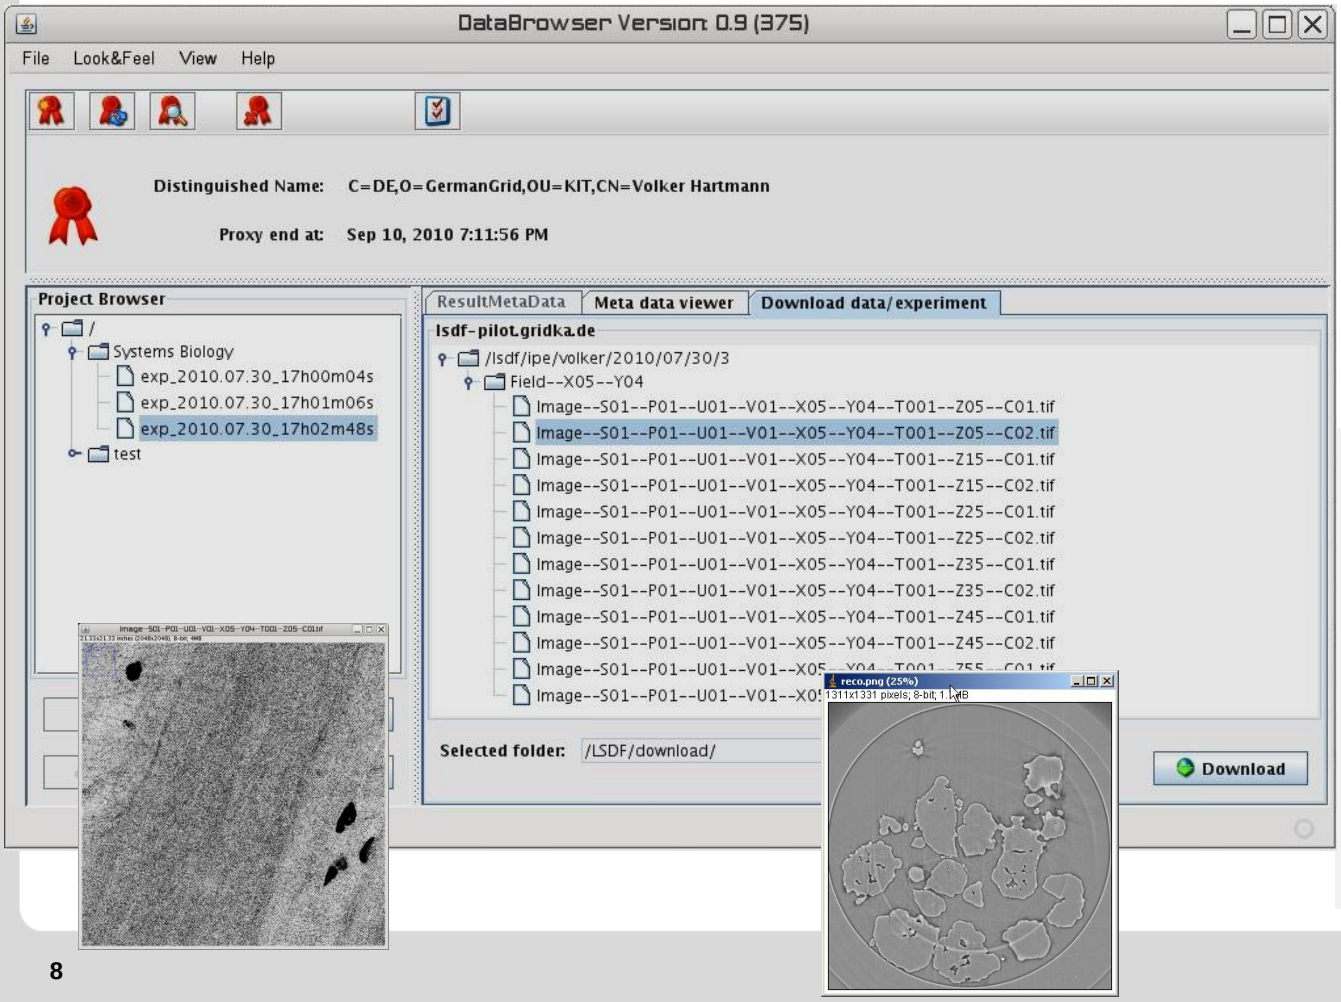

- Easy-to-use
- Extensible
- World-wide accessible
- Stable

Functions:
- Data management
- Queries in meta data cataloges
- Up-/Download
- Control of data analysis + vis. workflows

Data Transfer

Abstract Data Access Layer Application Programming Interface (ADALAPI)

- Java class library
- Seamless application access to LSDF
- Independent of transfer protocol and location
- Protocols and filesystems
  - local files,
  - gsiftp
  - sftp
  - http(s)
  - hdfs
  - iRODS
  - REST (Fedora)
- Authentifikation: X.509 certificates, user/passwd
- Performance up to 210 MB/s, 10 GE, gsiftp

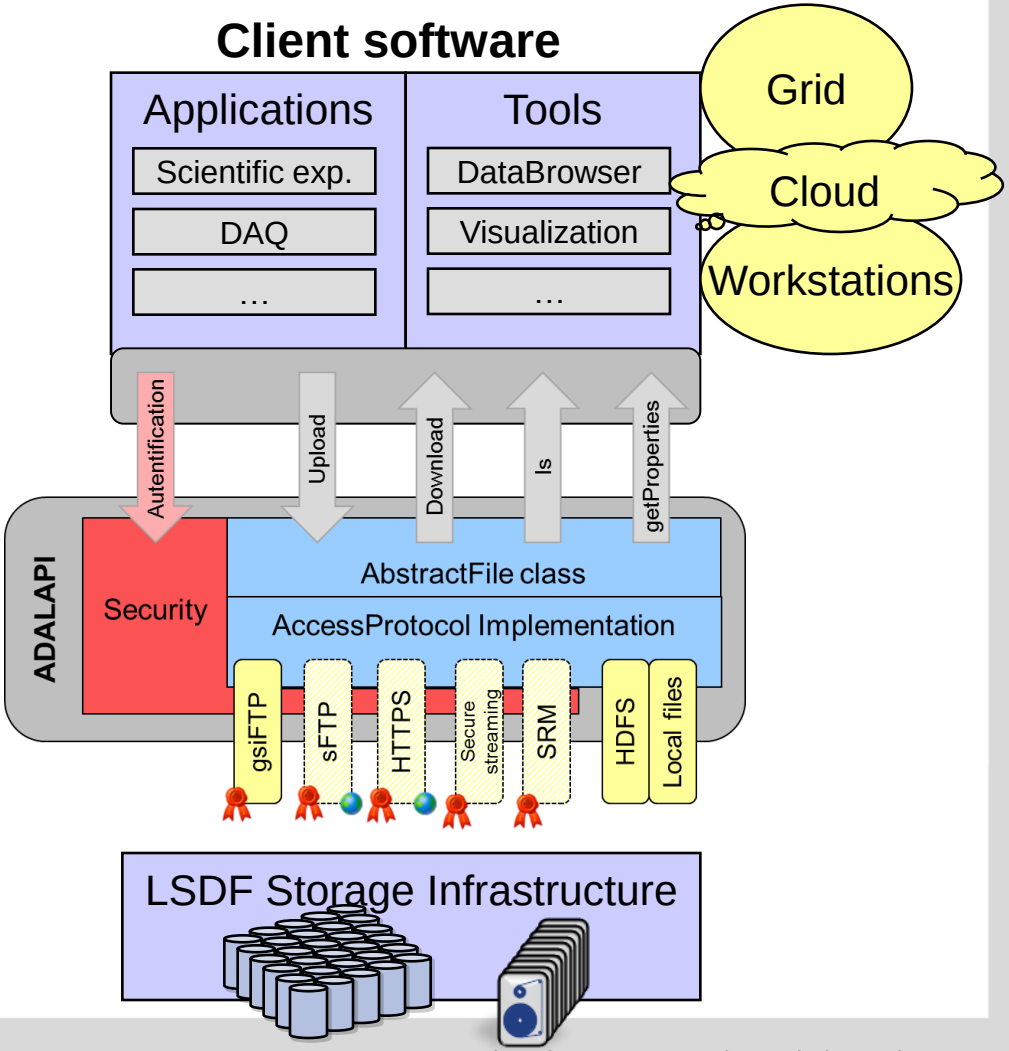

Generic Workflow Execution Framework

DataBrowser

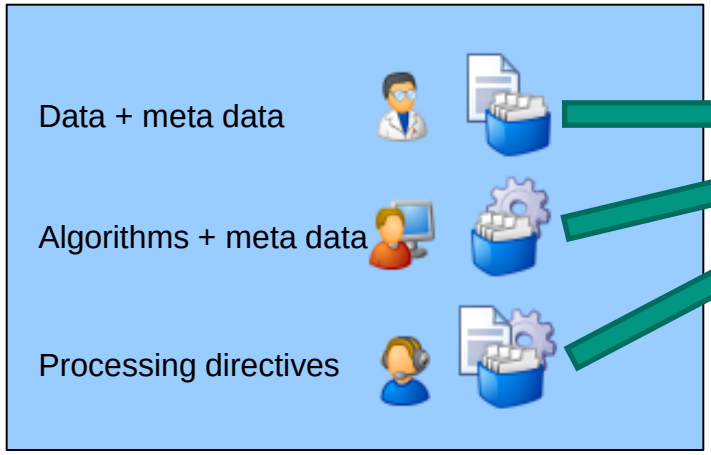

LSDF + hadoop

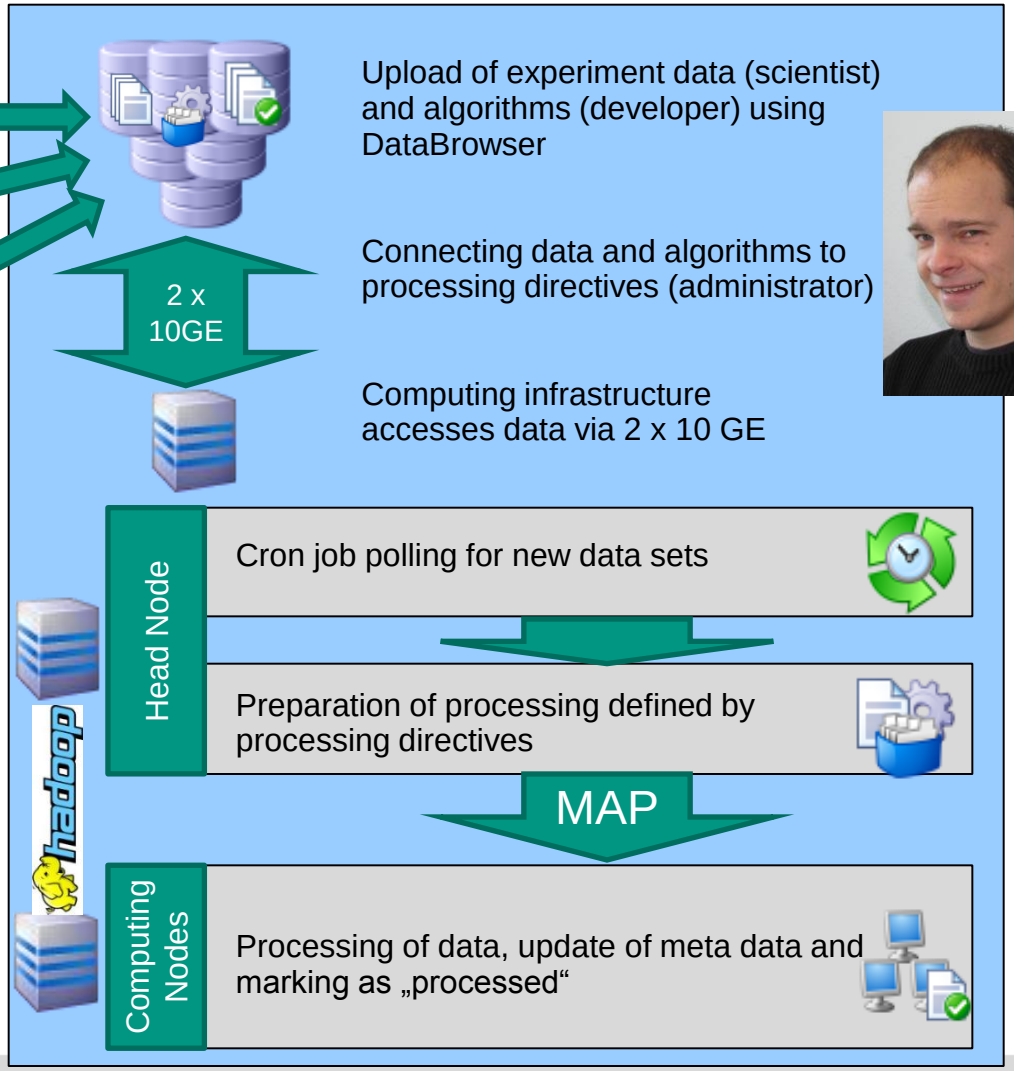

- Environment für arbitrary meta data driven workflows
  - Development
  - Deployment
  - Ingest
- Automatic processing based on “processing directives”
- Monitoring and error processing
- Algorithm database

DataBrowser und Lif2Raw Workflow

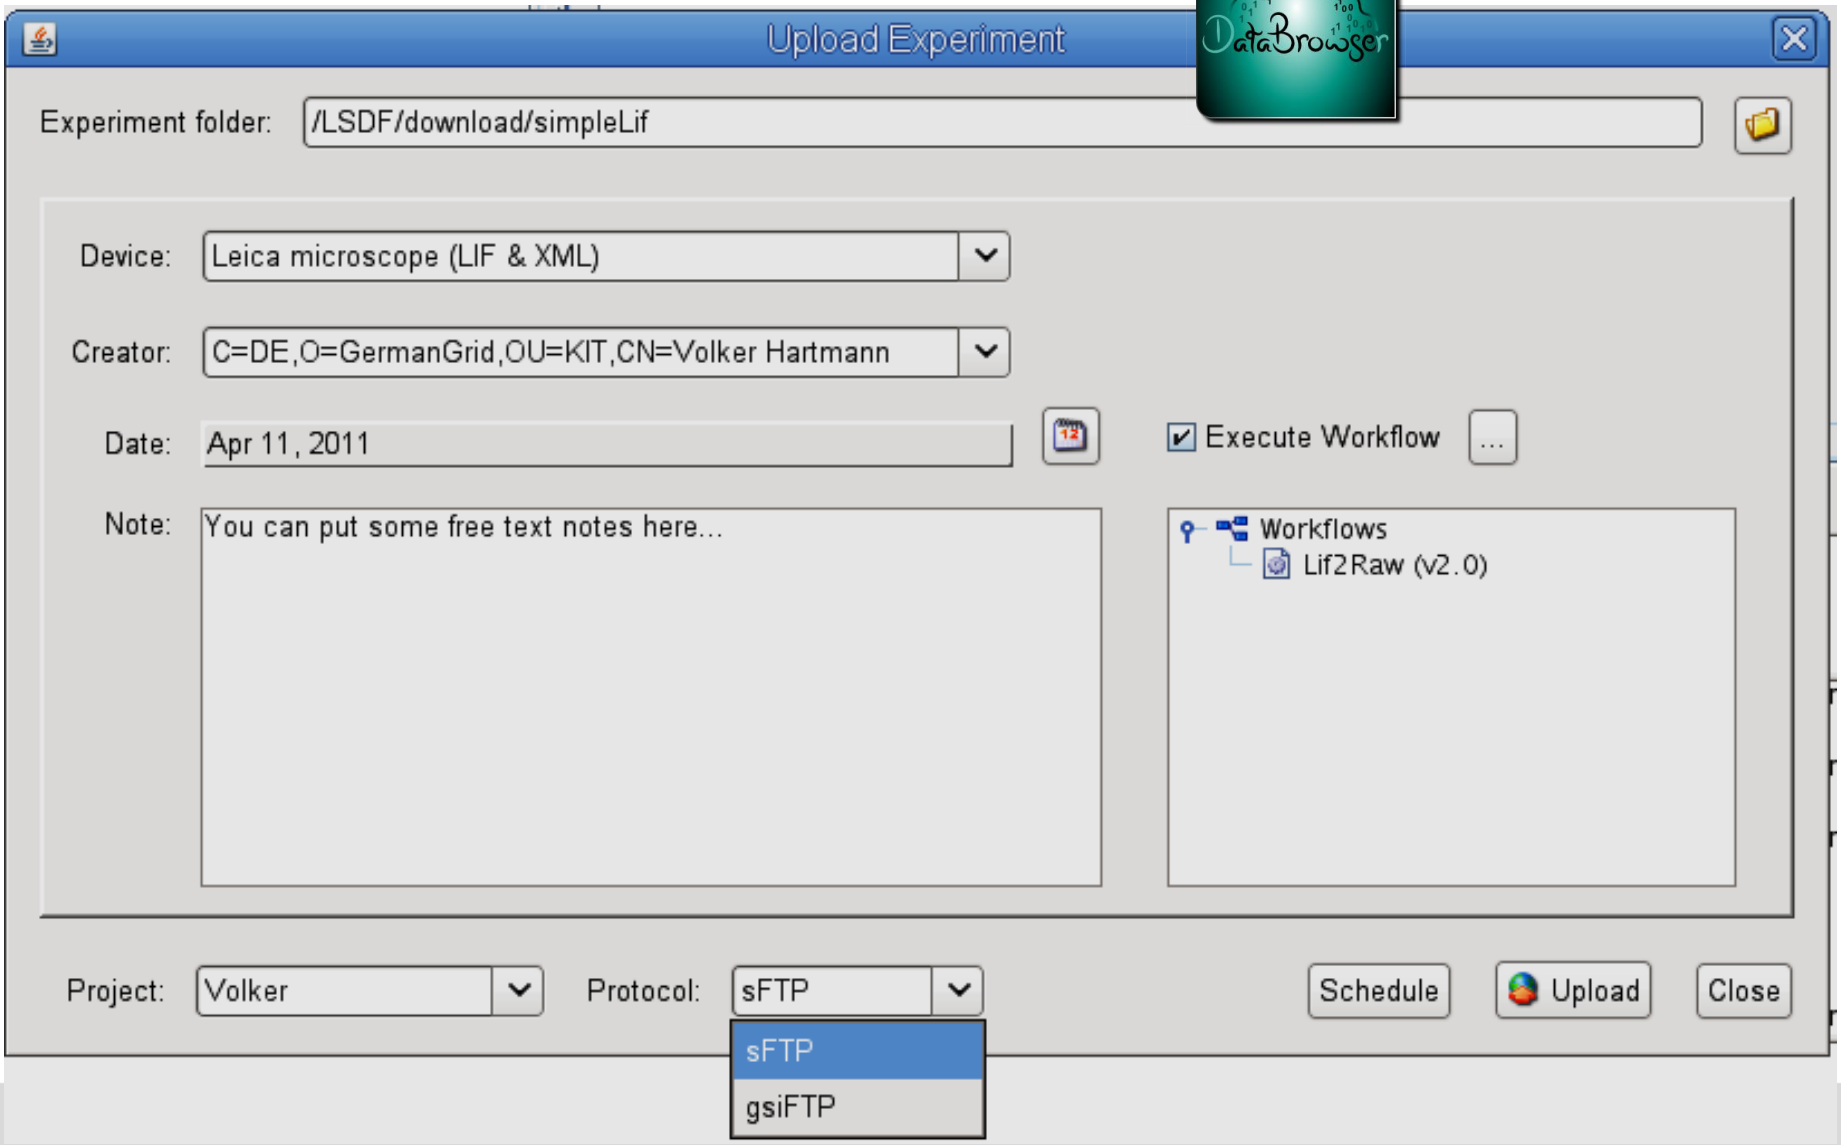

DataBrowser und Lif2Raw Workflow Monitoring

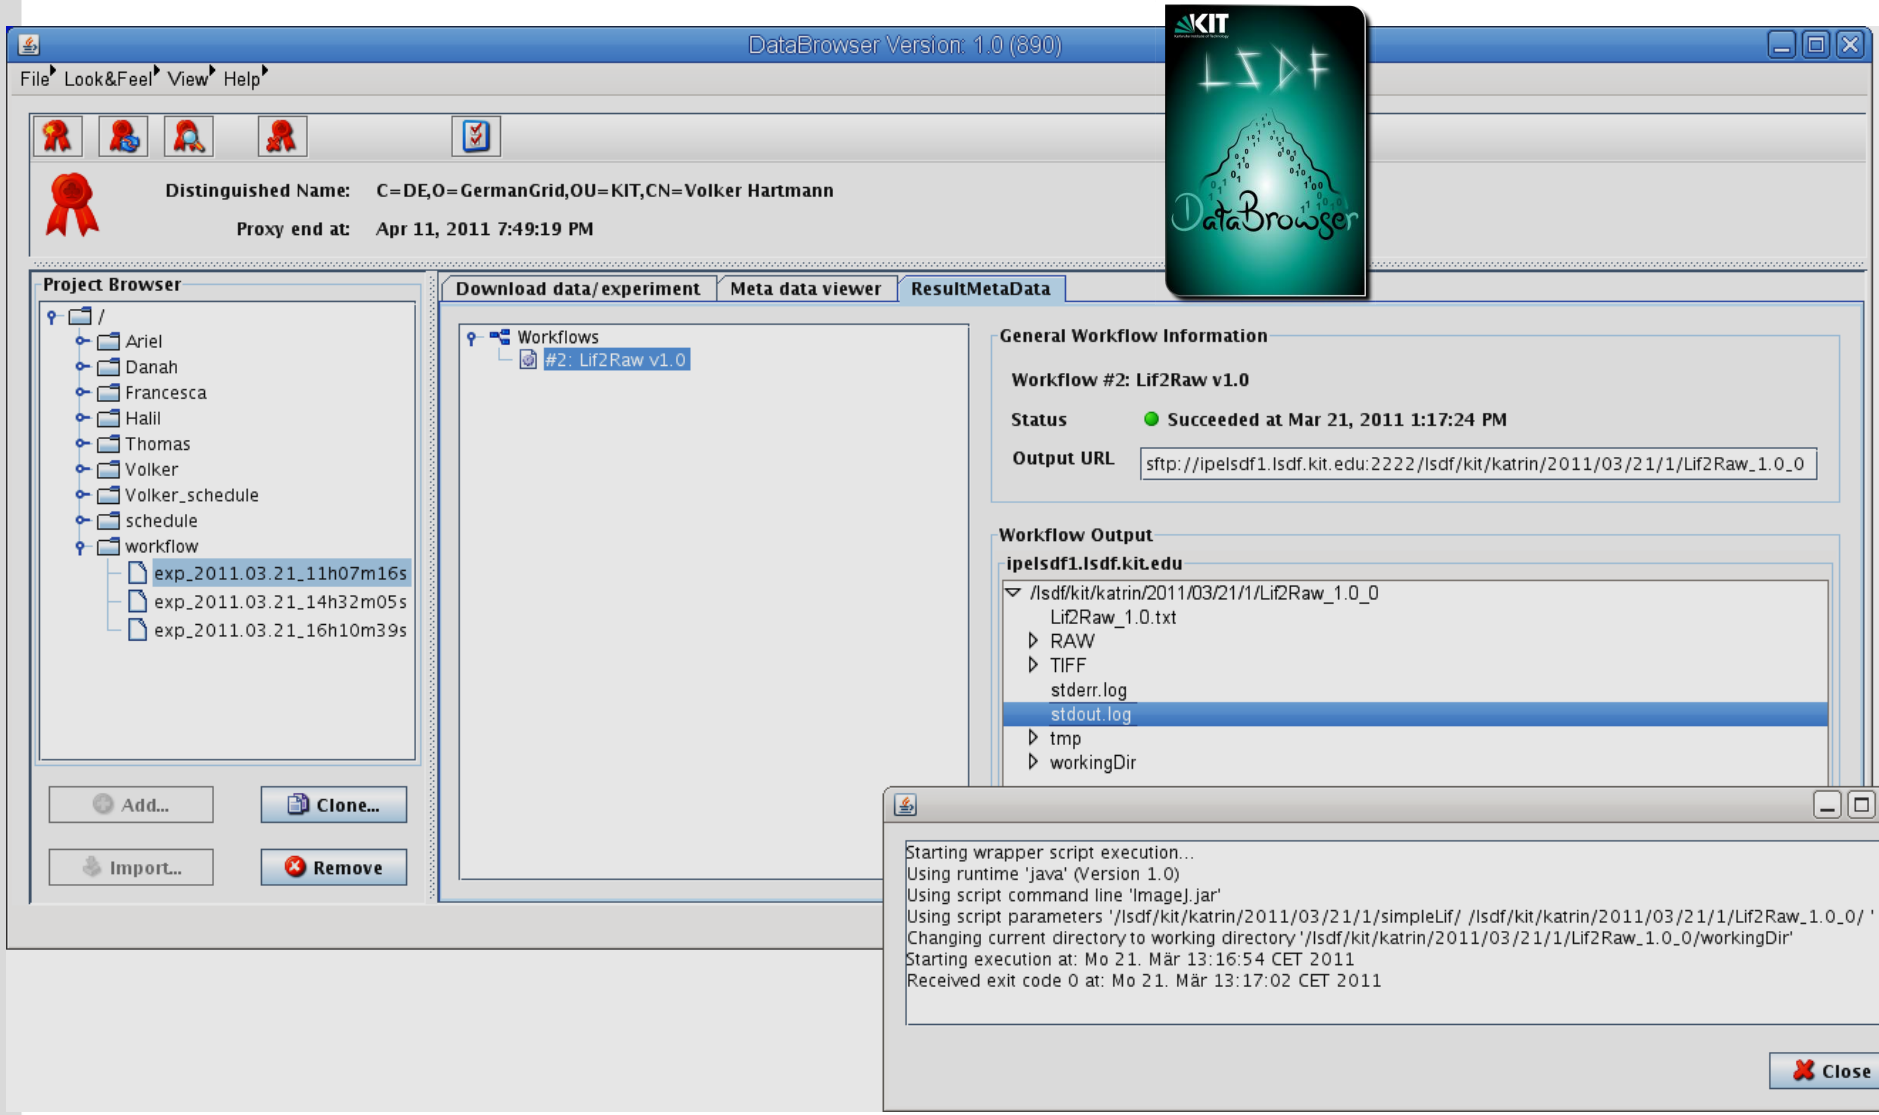

Data Management for ANKA

IPE DataBrowser

IPE MetaData

IPE Workflow

PhD Student Halil Pašić

- Design of the data life cycle management for image BLs
- Transparent connection to LSDF
- LSDF archiving (0,5 PB/a)

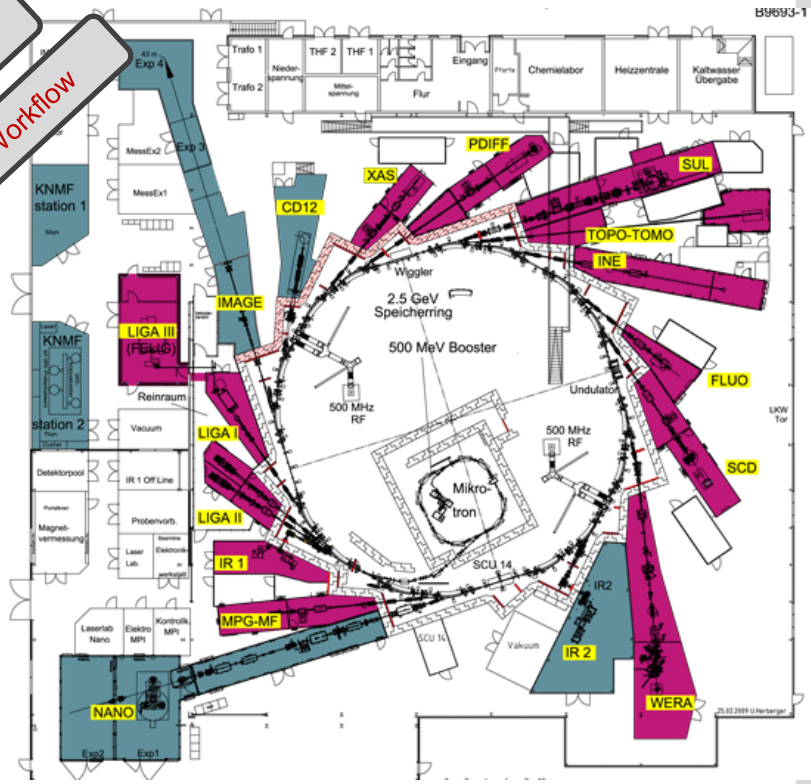

- Definition and Implementation of data formats and APIs for BL data access
  → in cooperation with Eugen Wintersberger (DESY) + Tomas Farago (ANKA)

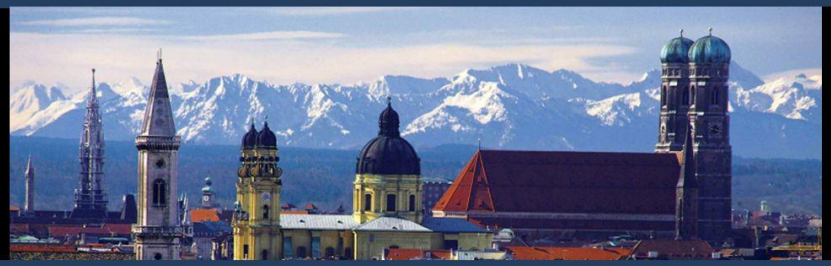

PDP 2012

20th Euromicro International Conference on Parallel, Distributed and Network-based Processing

15-17 February 2012 in Garching, Germany

www.pdp2012.org

Topics

- Parallel Computing
- Distributed and Network-based Computing
- Data Intensive Computing
- Models and Tools
- Systems and Architectures
- Advanced Algorithms and Applications

Program Chair

Rainer Stotzka, KIT, Karlsruhe, Germany
Michael Schiffers, LMU, Munich, Germany
Yiannis Cotronis, UA, Athens, Greece

Special Sessions

- Cloud Computing for Compute and Data Intensive Applications
- On-chip Parallel and Network-based Systems
- Energy-aware Systems
- GPU Computing and Hybrid Computing
- Security in Networked and Distributed Systems
- Next Generation of Web Computing
- Modeling, Simulation, and Optimization of Peer-to-peer environments
- Grid, Parallel and Distributed Bio-informatics Applications
- Cross-Fertilisation of Parallel and Metaheuristics Techniques
- Parallel and distributed storage systems

Important Dates

Papers due:	July 25, 2011
Acceptance notification:	October 10, 2011
Camera-ready versions:	November 11, 2011