LSDMA All Hands Meeting 2013 WP1 and WP2

Report

Patrick Fuhrmann, DESY

Preliminary remarks on WP1 and WP2

- dCache members of the WP1 and WP2 team at DESY (including the WP1/2 lead) started work for LSDMA May this year, as this was the end of the European Middleware Initiative funding period.
- We
  - started focusing on WP1 (fast start)
  - and setting the scene for WP2 (slow start).
- Many thanks to Marcus, helping us with a quick-start and for easing DSIT management structures especially helping WP1/2 with the 18 month plan.

Result from the Monday DLCL – DSIT meeting

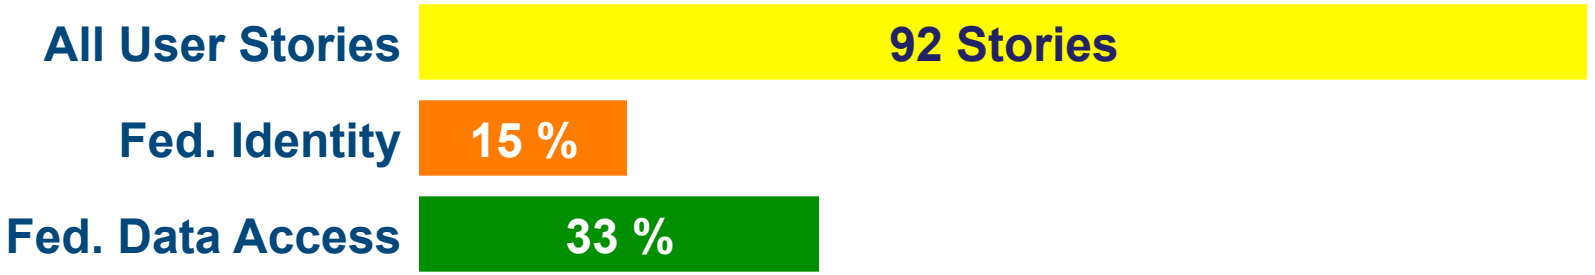

- WP1 (Federated Identity)
  - secure data access and storage plus data privacy and data sharing.
- WP2 (Federated Data Access)
  - Reliable, fast, automatic data transfers
  - Remote access to storage systems for read and write
    - from automated systems
    - from analysis facilities
  - Local cache system management for reading and writing
  - Database access and data - metadata relations

How did we compose the 18 month plan

Initial WP activity list

Interest of DSIT partner

WP extraction from User Stories

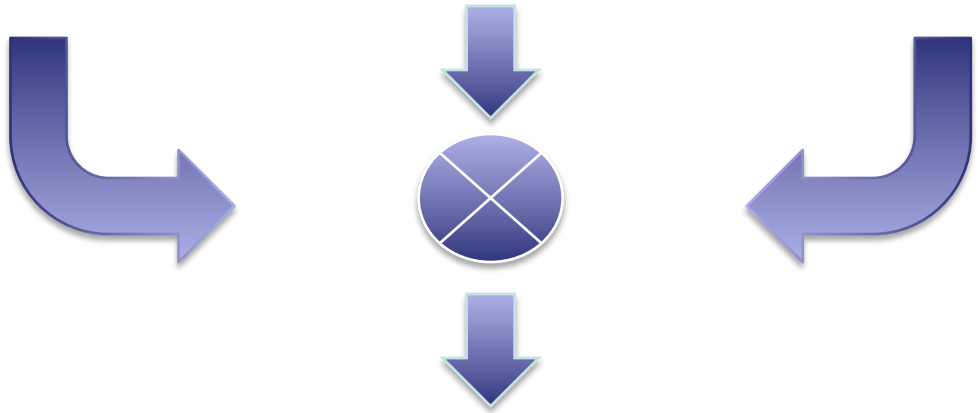

Work Package Activities

Feasibility has only be checked briefly yet.

Moving on with the WP 1/2 reports

WP1 Federated Identity

WP1 administrative activities

- Very active team
  - KIT : Arsen, Marcus
  - Jülich: Bernd, Shiraz
  - DESY: Paul, Patrick
- Organized Events:
    - AAI Symposium at DESY
    - 3 Technical and 2 regular phone meetings
    - F2F at DESY just recently
- Paper submitted, under the lead of Arsen

Overview WP1

- My Goal for WP1:
  - All German scientists should be associated with an IdP a.s.a.p.
  - Everyone associated with an IdP should be able to use German E-science resources. (Initially focusing on Storage Resources)
  - Authorization should still be controlled by the Virtual Organization or the Resource Providers.
- The provided technical solution should have (technically) the capacity of being integrated into European and worldwide infrastructures.
- We are well aware that legal issues are actually the main problems on our way to a German or even European solution but for now, we are focusing on the technical issues.
- Paul will elaborate on WP1 details later.

WP2 Federated Data Access

LSDMA

We might choose a new title

Something with Big Data

WP2 administrative activities

- Large team
  - KIT: Thomas
  - DESY: Tigran, Paul, Christian, Patrick
  - Jülich: Bernd, Shahbaz
  - THW: Hermann, Tom, Jana, Leonie
  - GSI: Dennis, Thomas
- Organized Events:
  - Workshop with Oracle and NetAPP developers (at DESY) on NFS 4.2 and FedFS.
  - Workshop with developers from SLAC, CERN, Midwest Tier II and DESY (at DESY) on contributing to the xRootd federation (for LHC computing only.)
  - Phone meeting with Globus Online Management on WP2 and WP1 on the possible future use of GO transfer service for German Scientists. Considering to work with IGE (Helmut Heller) to go for a proposal in the H2020 Framework.
- Accepted paper at CHEP 2013
- Not a single phone meeting ☺. Local site interests are just too divers.

18 Month plan tasks for WP2

- Data transfer tools
  - Globus Online
  - FTS and FTS3
- Data Federations
  - Centrally managed file system federations: FedFS
  - Dynamic Federation based on http/WebDAV.
- New Technologies: Cloud Protocols and Object Stores
  - CDMI and Object Storage Interfaces for dCache and UNICORE
  - Evaluation of CDMI for KIT big data a applications
- Data Acquisition support (Performance and Data Formats)
  - High performance Data Sources (high frequency, large CCD's)
  - Support of HDF5 and Nexus
- Database access from UNICORE
- File Systems evaluations (Lustre ZFS)

http://wiki.gsi.de/cgi-bin/viewauth/LSDMA/WP2FederatedDataAccess-2013+2014

Moving on with collaborative efforts

LSDMA

- Wide area transfers
- Data Federations
  - NFS 4
  - Dynamic HTTP/WebDAV federation
- Cloud protocols
  - WebDAV
  - CDMI
    - Data
    - Metadata

Wide Area Data Transfers

Two ‘ready to use’ options

- US Service by Globus
- Will never become a product
- Will never run in Europe
- dCache has frequent contact with GO developers.

File Transfer Service (FTS3)

- CERN software product (Former EMI)
- FTS long history in transfering LHC data, FTS3 in preproduction phase
- There will be an FTS 3 instance at KIT anyway in the near future

Data Federations

Federating File Systems with NFS 4

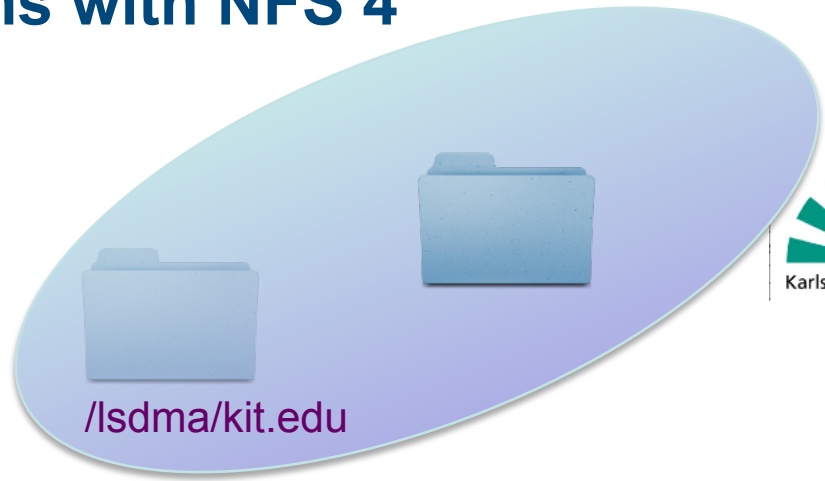

- FedFS allows for centralized administration
- Authentication yet unclear

Dynamic Federation

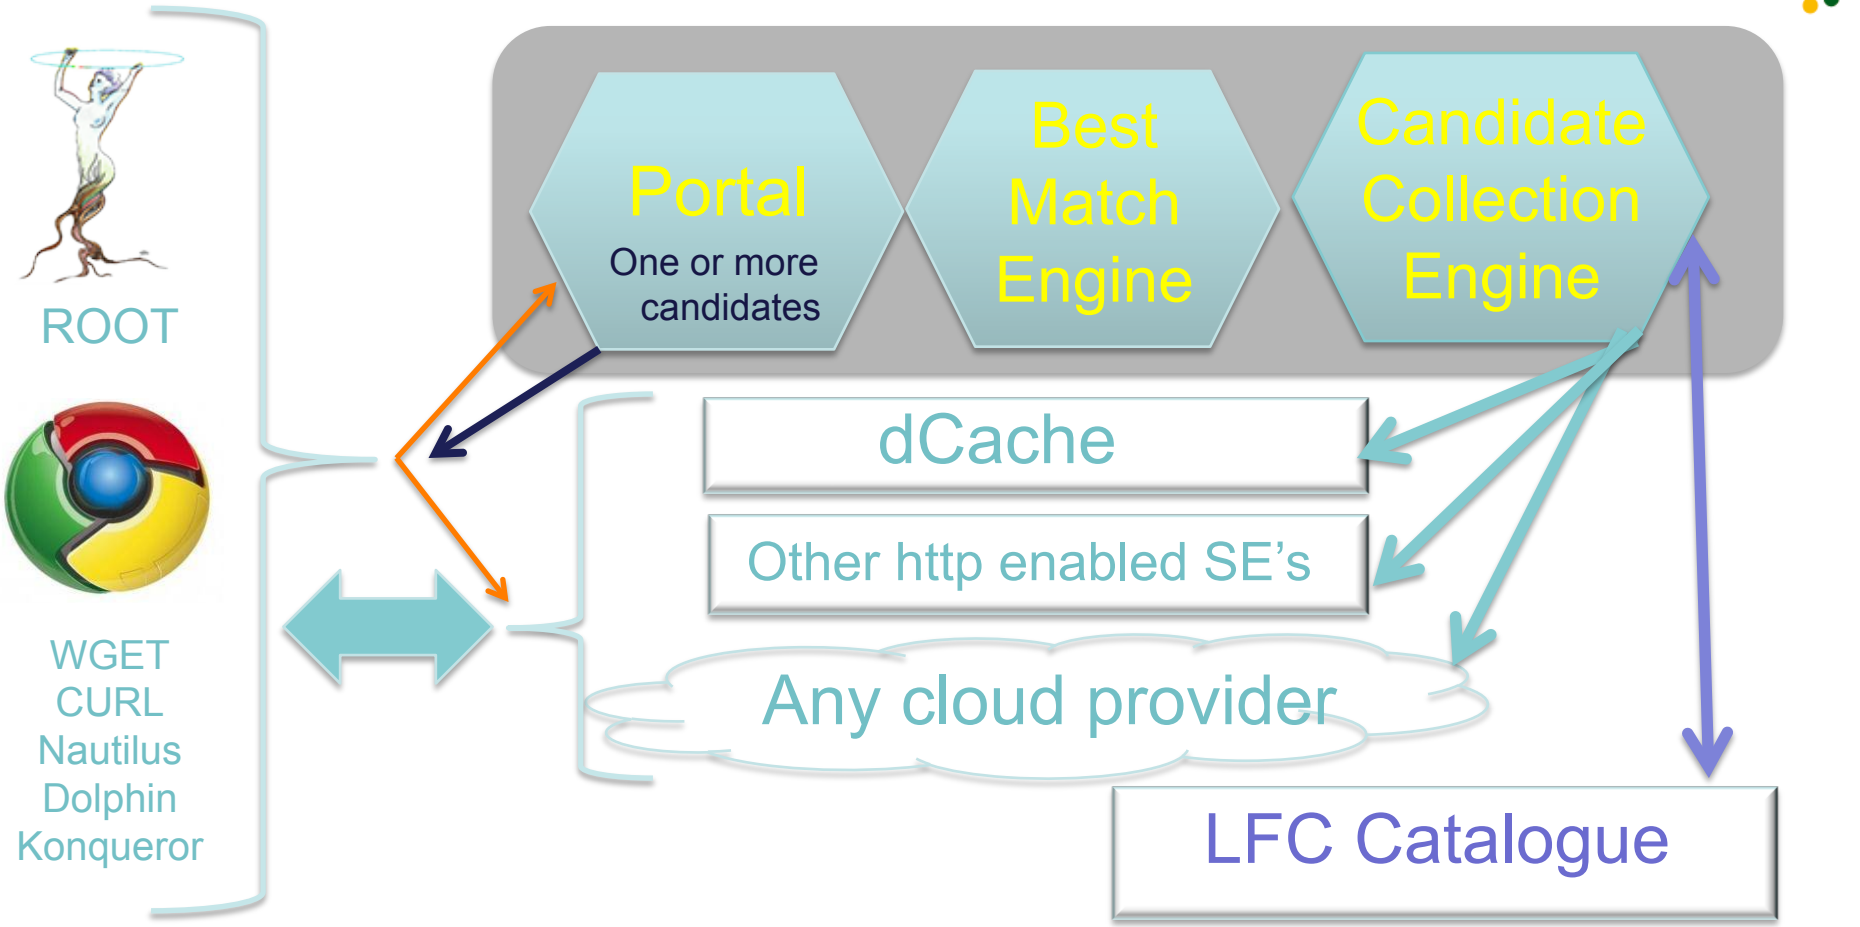

Cloud storage in LSDMA

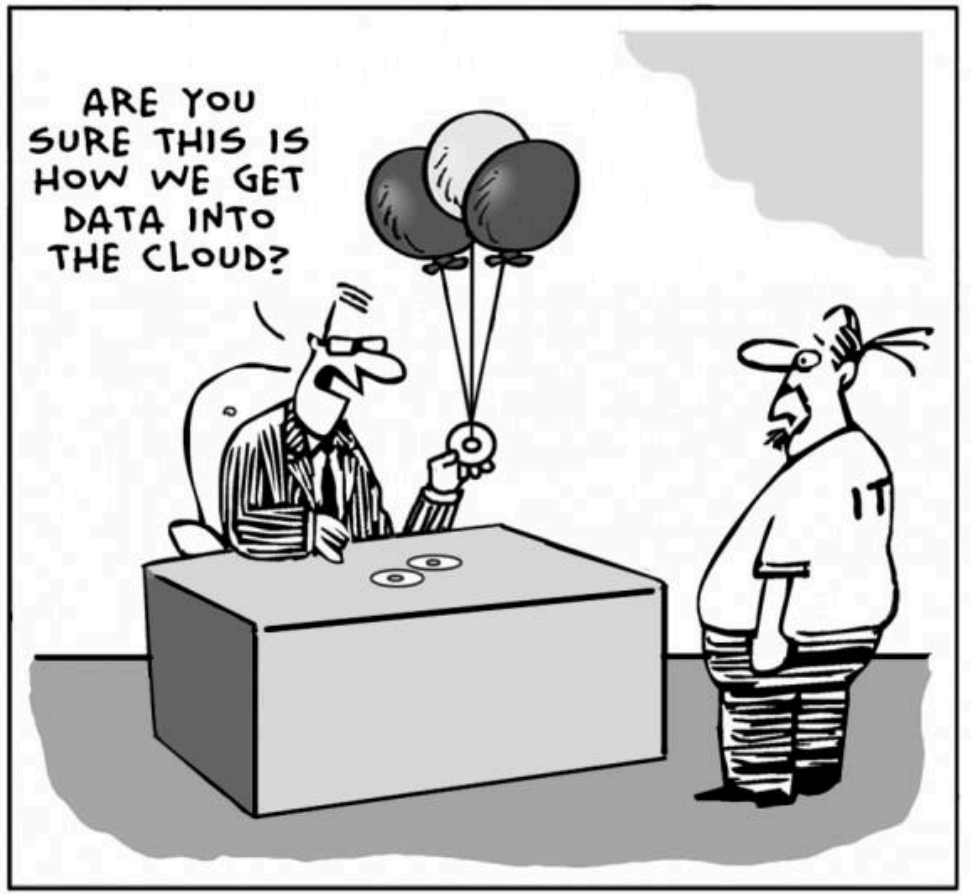

LSDMA collaboration on cloud storage

LSDMA

Step I: dCache and the HTW Berlin (Hermann, Tom, Leonie, Jana)

- We are offering a real cloud like service for students using WebDAV. "Unlimited storage space for free".
- In return we get their input on
  - Requirements for a 'young cloud'
  - Which mobile clients work and don't work with our Cloud Service.
  - Student contribute to fixing Open Source mobile WebDAV clients.
- This allows us to adjust to e.g. Fabians (Energy) requirements (reported this morning) "It's not about performance and file size, but about sharing and privacy"

In Progress

LSDMA

There is a better Cloud protocol then WebDAV

What is CDMI

LSDMA

DataDirect NETWORKS

One initiative to prevent Vendor Lock-in, is SNIA's Cloud Data Management Interface (CDMI). This is a set of pre-defined RESTful HTTP operations "for assessing the capabilities of the cloud storage system, allocating and accessing containers and objects, managing users and groups, implementing access control, attaching metadata, billing, moving data between cloud systems, exporting data, etc."

LSDMA collaboration on cloud storage

Step II: Implementing the CDMI data server into dCache

In Progress

Step III: UNICORE implements the CDMI client

Step IV: UNICORE implements the CDMI server

Planned

STEP V: CDMI and metadata

LSDMA

- CDMI allows arbitrary meta data to be assigned to data.
- CDMI allows to run meta data queries against its data repository.
- dCache supports meta data to be associated to data.
- dCache intends to support the meta data feature of CDMI.
- This can either be used as a backend for ICAT or as a replacement,
- The task is important for the DESY Photon Science community.
- Following yesterdays DLCL – DSIT meeting, this is a nice solution for other communities as well.

Planned

Many more ‘single partner’ activities

LSDMA

- UNICORE : Database access for UNICORE
- GSI: Evaluation of Lustre performance
- dCache
  - Support of HDF5 and other containers in dCache (read/mod/write)
  - Doing data management with large amounts of small files by (internally) creating containers to improved tertiary storage device performance. Completely transparent for the user.

Summary

- WP 1 and 2 are getting up to speed.
- Monday meeting was extremely helpful to adjust requirements.
- WP1 is technically feasible but will face legal fights.
- WP1 will require European and Worldwide collaboration.
- WP2 software and service components are already available to a large extend.
- Not solved yet
  - DLCL are already picking technologies
    - E.g. iRODs for KIT <-> DKRZ
    - E.g. Mongo DB access using mechanisms not talking to WP1
    - E.g. Key Technologies: Distributed analysis and storage
    - E.g. DARIAH-DE: “Building Storage Federation” (KIT, Juelich, Goettingen, ..)
  - How to get customers to consider alternative services and software components.
- The communication between the data lab. projects and the WPes is not yet sufficiently professional.
  - DLCL choose their solutions by themselves anyway
  - DSIT WPs don’t have a chance to present alternative solutions (maybe not wanted)

The END

Further reading:

http://wiki.gsi.de/cgi-bin/viewauth/LSDMA/WP1FederatedIdentityManagement-2013+2014

http://wiki.gsi.de/cgi-bin/viewauth/LSDMA/WP2FederatedDataAccess-2013+2014