

DSIT WP1 WP2

Federated AAI and Federated Storage Report for the Autumn 2014 'All Hands Meeting'



Content



- WP1

- GSI: GSI Web Services accessible via IdP credentials
- GSI: Plan to integrate with UNITY (setting up a unity server at GSI)
- DESY: dCache, Globus Transfer (online). SAML Translation
- KIT: LDAP Facade
- KIT: Non Web Login using federated identity systems.
- Jülich: UNITY
- DESY and KIT part of INDIGO-DataCloud project management with AAI and Storage (WP4)

- WP2

- Jülich: "Database access in UNICORE" (Master Thesis completed)
- GSI: Lustre evaluation
- KIT: Something with data transfers by Jos (request pending)
- DESY: CDMI in dCache (Master Thesis Completed, Jana)
- DESY: Small Files for tape access (In production for 'friendly users')
- DESY: OwnCloud integration with dCache (In production for IT people)
- DESY: Fast ingest for dCache (good progress)
- DESY: Dynamic Federation: Production system (adopted by LHCb and Intesity Frontier)
- Bits and Pieces
 - FedFS: no further interest
 - We would be able to offer FTS3 (but no interest yet). Globus Transfer is easier.

Publications



- Combining the X.509 and the SAML Federated Identity Management Systems, 2014
 - Recent Trends in Computer Networks and Distributed Systems Security
 - M Hardt, A Hayrapetyan, A Memon, P Millar,
- Evaluating the Performance and Scalability of the Ceph Distributed Storage System
 - IEEE Big Data 2014
 - D Gudu, M Hardt, A Streit

Master Thesis Work

- Implementierung eines Cloud Interfaces in dCache
 - HTW Berlin (Jana Weschenfelder)
- Database Access in UNICORE
 - Forschungszentrum Jülich

Presentations



- 5.5.2014 Patrick: Berlin Cloud Event on dCache/OwnCloud
- 16.6.2014 LSDMA Demo in Berlin (all of us)
 - Please check indico for details
- 2.9.2014 Paul: GridKA School presentation on AAI
- 22.9.2014 Paul: Terena Storage Task Force, Uppsala
- 24.9.2014 Patrick: EGI Storage Session
- 13.10.2014 Paul: Atlas Workshop
- 14.10.2014 Patrick: KM3Net @ Paris

WP1 Federated Identity (Management)



- GSI
 - Setting up an IdP at GSI which will be connected to the DFN AAI federation
 - Getting the LSDMA Wiki at GSI to work with the German IdPs
 - Support GSI customers in working towards accepting SAML for their services
- DESY
 - Getting the DESY IdP to work with the DFN Federation
 - Using Globus Transfer against storage services with DFN SLCS service
- KIT
 - Credential translation for Web SSO and ECP.
 - Getting login service (non web services) to work
- Jülich
 - UNITY support

- Dennis is preparing GSI enabling Web Based usage of SAML.
- The plan is to establish a test-bed to allow the distributed communities that GSI supports to understand Fed Identity and to see if integrating it makes sense for them.
- A testbed IdP and to support GSI communities in adding SAML support to their services.
 - Technically things are prepared
 - Still waiting for the documents to be signed.
- Plans: Setting up a UNITY server at GSI for
 - Group Management
 - Running UNITY as IdP including for homeless users
- Once GSI joins DFN production AAI, Dennis will install the AAI plug-in on the main wiki. With this, anyone from DFN-AAI can use the Wiki by authenticating with their local IdP.

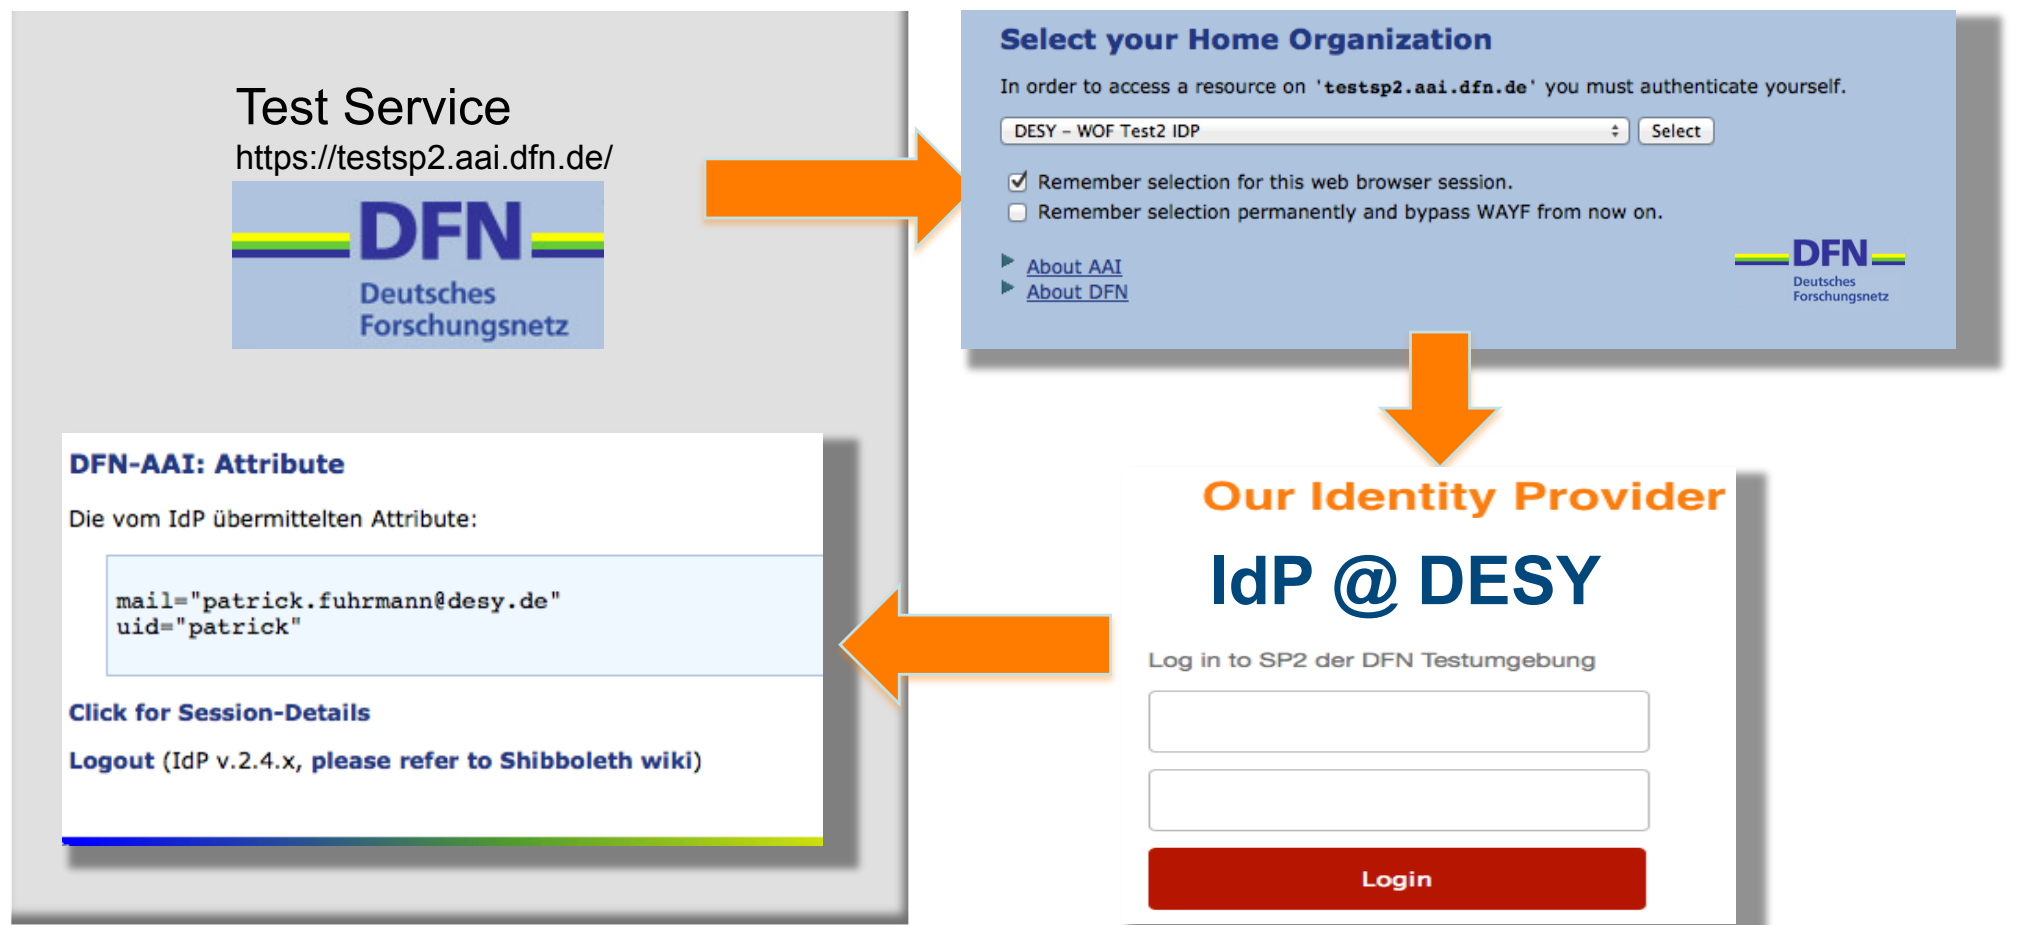
BTW: In case you plan to do the same



What you need :

- DFN Documentation
 - <https://www.aai.dfn.de/dokumentation/identity-provider/konfiguration/>
- Shibboleth documentation (wiki.shibboleth.net)
- DFN Test SP
 - <https://testsp2.aai.dfn.de/>
 - Provides log files for debugging
- Ask DSIT for help

The DFN Test SP in action

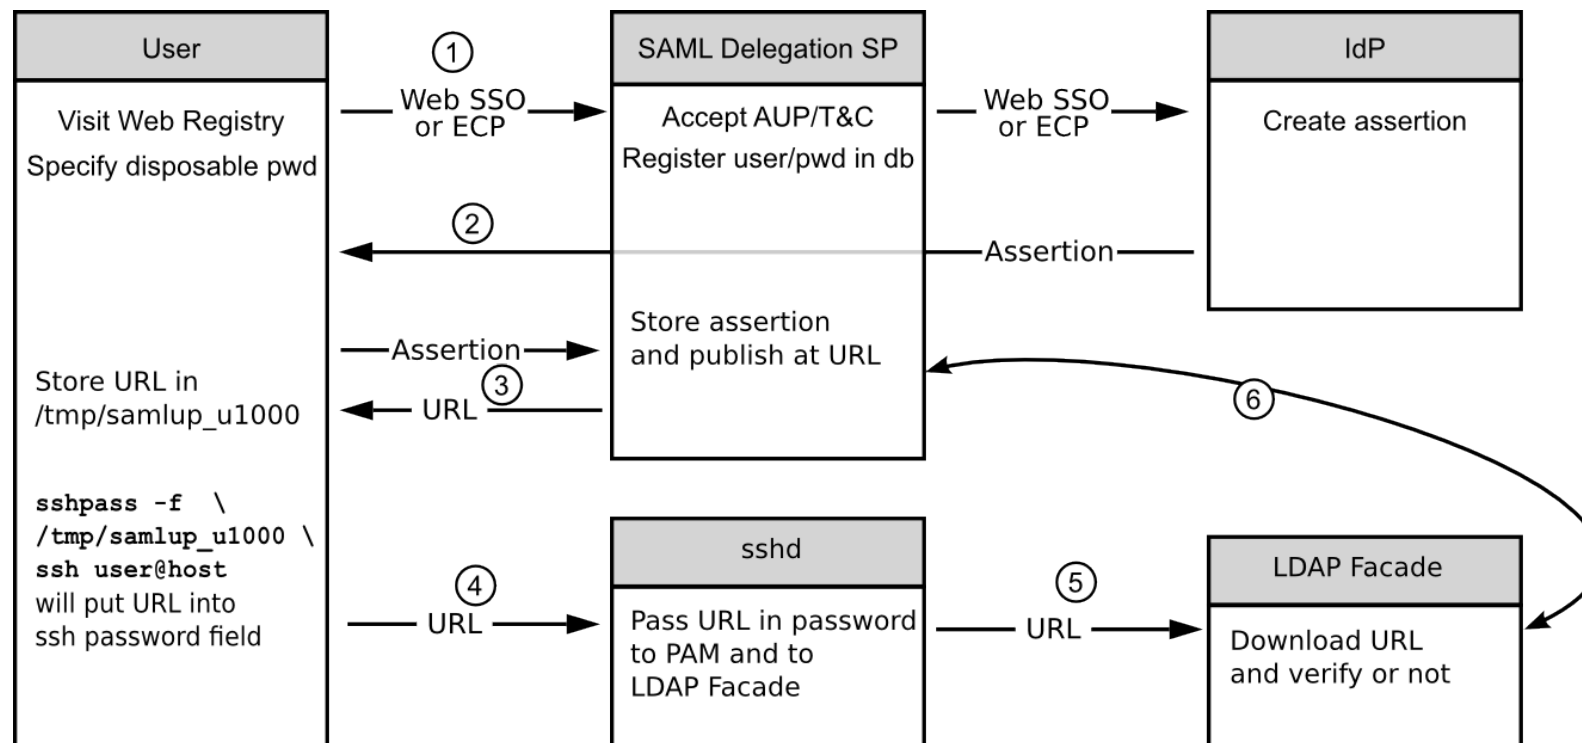


Bureaucratic part



- Sign an agreement with the DFN AAI federation.
- You agree to fulfill a certain number of requirements concerning your IdP policies. (Basic and enhanced)
 - Identify of an entity (user) within the IdP has to be verified by a well specified procedure.
 - Guarantee decommissioning
 - etc

- **High Level Objective:** Ssh login with SAML credentials



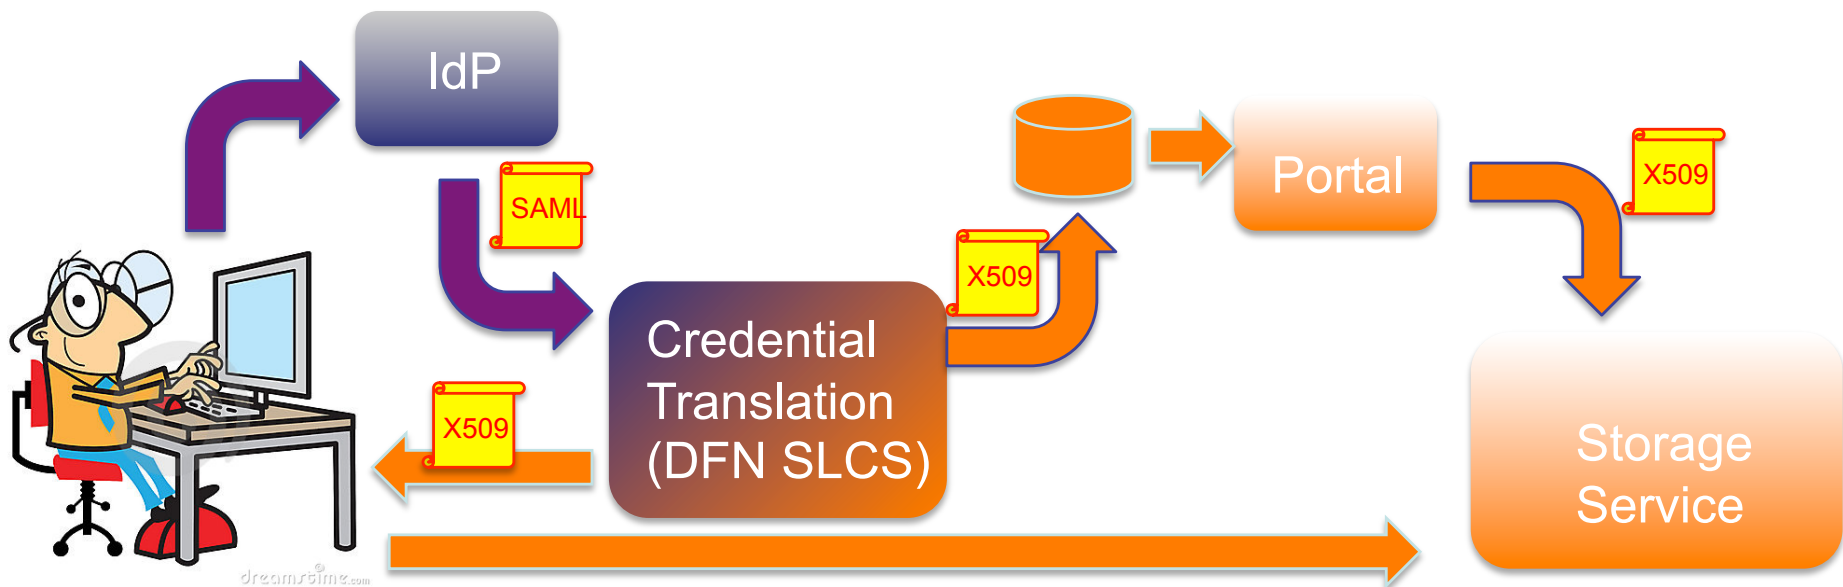
KIT and DESY:



- **High Level Objective:** Using DFN credential translation service to allow users to utilize X509 based resources through your local IdP.
Two use cases
 - Arsen: X509 credential is used directly (user to service)
 - Paul: X509 credential is given to a portal. The portal acts on behalf of the user.
- **Requires:** Arsen tries to get the DFN Short Lived Credential Service (SLCS) to work for him.
 - Turned out the DFN doesn't support ECP (Extended Client or Proxy)
 - Arsen in heavy contact with DFN to get this fixed.
 - Seems he has to build a system first before the upgrade their Shibboleth instance. (mimicking their setup to provide ECP)
- Please check our Arsens and Pauls Demos from the meeting in Berlin.
 - <https://indico.desy.de/conferenceDisplay.py?confId=10101>

Credential translation for user and portals

- User starts with SAML federated Identity (IdP)
- Service needs an X.509 Certificate
- Service is either directly accessed by user or via a portal



and whatever has to do with storage and data access.

- Database access from UNICORE (Master Thesis)
- Cloud protocol evaluation and implementation (Master Thesis)
- Dynamic Federation
- Scientific Storage Cloud
 - Fast Ingest
 - Sync-n-share interface
 - Media Aware Storage (e.g. Tape access for small files)
 - CDMI cloud protocol
 - New Customer (Jade)

Jülich: Database access in UNICORE



- At Jülich, a master thesis on "Database Access in UNICORE" was completed.
- The service supports SQL databases (e.g. MySQL) and Cassandra as an example for a NoSQL database. The security details such as database usernames and passwords are configured by the administrator, so the user has a seamless experience.
- The code is currently under review and will be included in the next major UNICORE release (early 2015). The service will be available both via a SOAP/WSRF interface and a RESTful interface.

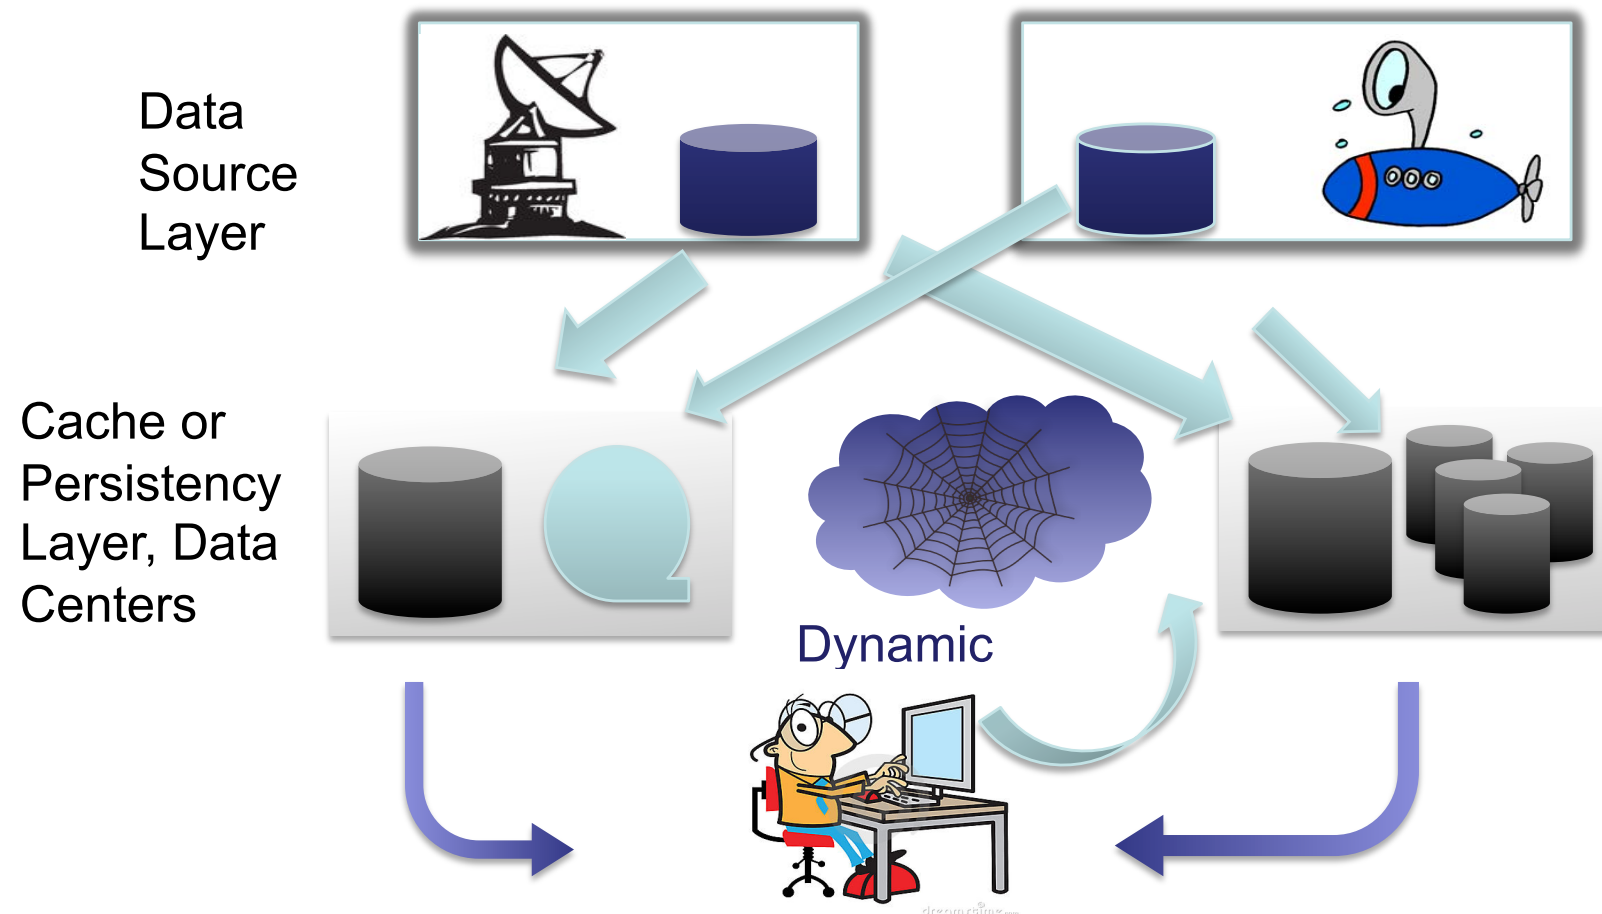
DESY/HTW Berlin: CDMI Server in dCache



- DESY/HTW Berlin: Jana finished her Master Thesis on Cloud Storage Protocols.
- Result: CDMI (w/o metadata) integrated into the dCache code repository.
- We became member of the Storage Network Industry Association (SNIA) to be able to feed back our code into the main stream SNIA repository.
- Code needs to be revised and will become available in one of the next feature releases.
- Next Step:
 - Add meta data operations.
 - Bernd is now considering to integrate the CDMI client into UNICORE so that we have a sparing partner.

The dynamic storage federations

Reminder: DynFed use case (Berlin DEMO)



Dynamic Federation



- Technology “Production ready”
- Was presented at the Berlin LSDMA Demo Meeting
- 9. Sep :
 - Quote by Fabrizio Furano : “LHCb starts today setting up WebDAV everywhere”
 - This is a basis for using DyncFed for LHCb.
- Next Candidate : Intensity Frontier at FERMIlab
- No customers found in LSDMA yet

dCache : The Scientific Storage Cloud

- Improving fast ingest of data into dCache
- Adding sync-n-share as regular door to dCache (Own Cloud)
- Improving media-awareness of dCache concerning tape (small files)
- New Customer for dCache

dCache: The Scientific Storage Cloud



High Speed
Data Ingest



Fast Analysis
NFS 4.1/pNFS

Resilient Mgr
Small Files



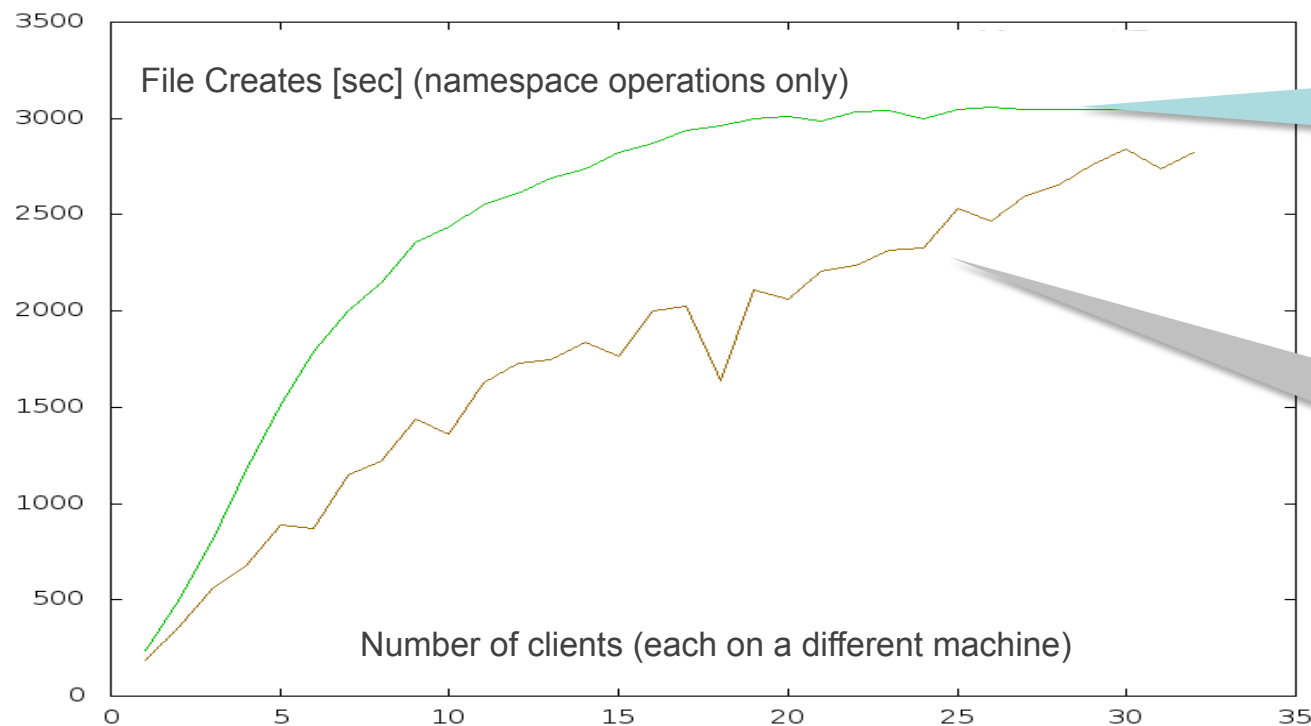
Sync & Share
by WebDAV

Wide Area Transfers
(Globus Online, FTS)
by GridFTP



Fast data ingest into dCache

- Name space operations for ingesting data into dCache could be significantly improved.
- We now seem to be limited by Postgres Performance




Server
8 Cores

Server
16 Cores
Using Hyper-
threading

Stolen from Tigran

OwnCloud dCache integration into infrastructure



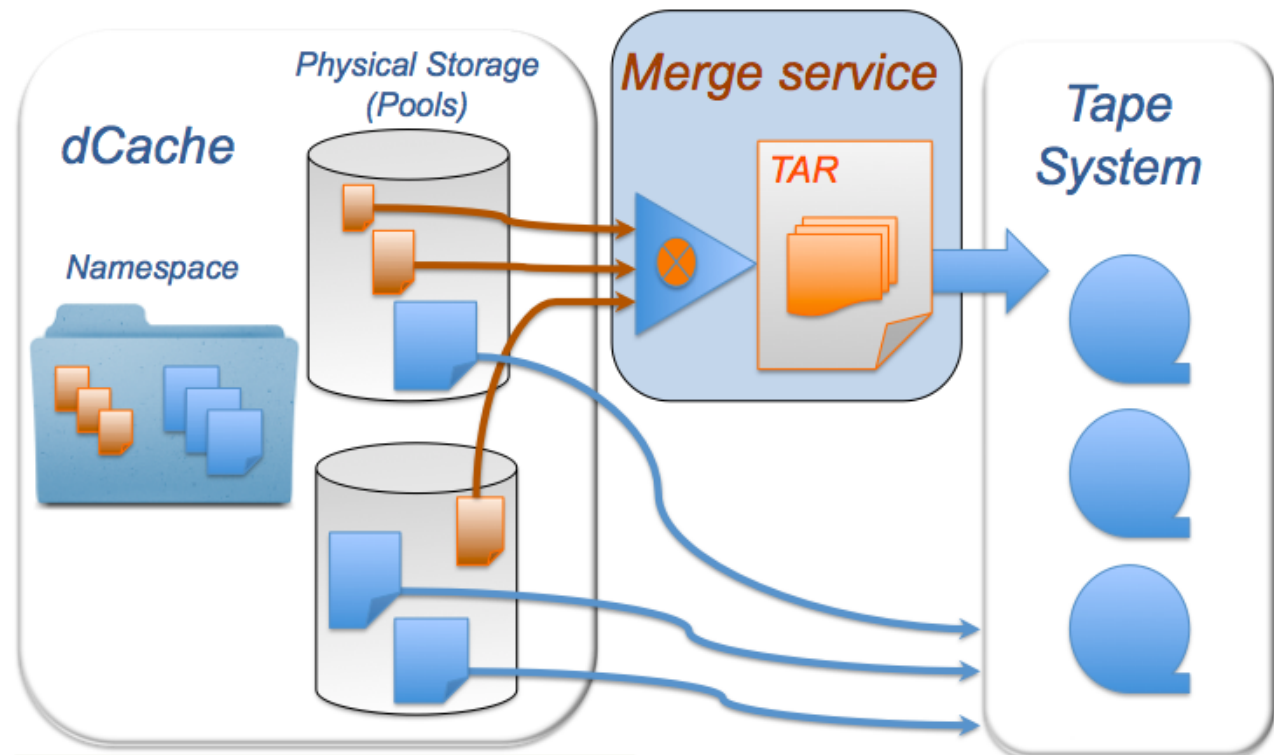
dCache scientific cloud (dCache-Own Cloud)



- Own Cloud software running on top of dCache
- All sync-n-share advantages of OwnCloud
 - Syncing with all kind of mobile devices and OS
 - Easy sharing of files and folders with 'identified users' or public
 - In use by a variety of sites.
- All dCache 'managed storage' advantages of dCache.
 - Different quality of services (scratch/regular/high resiliency/tape)
 - Extends access to more protocols (nfs/WebDAV/GridFTP)
- Moved into production for DESY IT people.
- Photons science (CFEL) customers following soon.

Small File support dCache media-awareness

- dCache bundles small files when data is scheduled to be migrated to tape system.
- Restore to disk happens transparently.
- Namespace only shows original files. User not bothered with file collections.
- Already second improved version available for photon science customers.
- In production for 'friendly' customers.



New customer for the scientific storage cloud



- Jade project within “Supercomputing and Modeling for Human Brain, SMHB”, Jülich.
- Storage collaboration with Aachen
- Using dCache features for
 - File hopping between Jülich and Aachen
 - Using GridFTP for wide area
 - pNFS (NFS4.1) for direct access
 - WebDAV
 - Tape access for achieving
 - Multi Tier access (SSD)
- Waiting for
 - Own Cloud integration becoming part of dCache distribution
 - SAML or SLCS integration



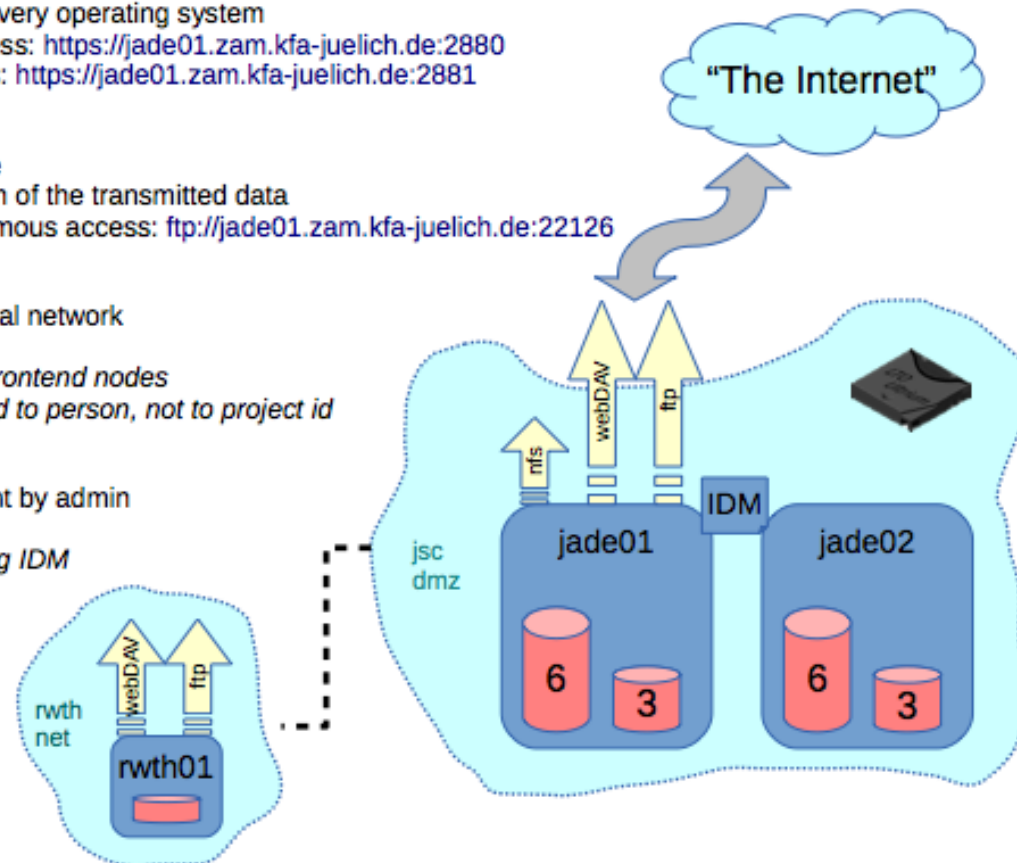
Slides by Bastian Tweddell

The SMHB Jade Project

Slides by Bastian Tweddell



- Supported access protocols
 - WebDAVs
 - World-wide accessible
 - Supported by nearly every operating system
 - Authenticated rw-access: <https://jade01.zam.kfa-juelich.de:2880>
 - Anonymous ro-access: <https://jade01.zam.kfa-juelich.de:2881>
 - ftp
 - World-wide accessible
 - Caution: No encryption of the transmitted data
 - ftp anonymous:anonymous access: <ftp://jade01.zam.kfa-juelich.de:22126>
 - pNFS
 - Only within JSC internal network
 - Outlook:
 - Availability on SC frontend nodes
 - Authorization bound to person, not to project id
- Identity Management (IDM)
 - Manual user management by admin
 - Outlook:
 - Integration into existing IDM
 - HBP
 - JSC



Bits and Pieces



- The **FedFS** evaluation has been stopped due to a complete lack of interest.
- We can make the the **FTS3** file transfer service available if needed. It is software which has to be run by the corresponding community or associated site or lab. It is not a services like Globus Transfer.
- dCache signed a MoU with an industry partner in Switzerland to handle License issues and compensation for dCache working
- DESY and KIT in H2020 Proposal 'INDIGO Data Cloud' on AAI and storage.
- Planning various CHEP'15 abstracts.

END