

# Welcome ... and stuff



7<sup>th</sup> International dCache Workshop

Patrick Fuhrmann



# Preview

dCache.org 

- News
- More news
- And maybe even more news

# Premiere



- Not at a WLCG location but a University.  
(Many thanks to Hermann)
- I wasn't involved in the logistics: All done by Christian, Tanja and Melanie. Which doesn't really make me less nervous.



## General News

## Bye Bye, for now



- Tanja Baranova left the dCache team. She got a new job, here in Berlin.
- Doris Ressmann from KIT is no longer leading the German Support team . She is still at KIT but now leading a developers team. CB took over.

**And Good Luck**

# First Asian Pacific Workshop

dCache.org 

collocated with the

## International Symposium on Grids & Clouds 2013

*Collaborative Simulation, Modelling, Data Analytics*



17 - 22 March 2013, Academia Sinica, Taipei, Taiwan

### The first Asian Pacific dCache Workshop

17 March 2013 - Taipei

#### Main Topics

- dCache Installation & Configuration
- NFS4.1/pNFS
- HTTP/WebDAV
- Security
- Hardware Life Cycle
- Tertiary Storage Access
- dCache Features
- Master Classes



With participants from

- Australia
- Taiwan
- Thailand
- Japan
- India
- UK and
- Germany



- Now seriously “Open Source”, ( $\leq$  AGPL)
  - Code in GITHUP
  - Starting with developers tutorials (yesterday)
    - 28 participants
    - 0 suicides (as far as we know)
  - Reminder:
    - Anyone can join
    - However, code still has to pass a review process

- Annual dCache developers workshop starting Thursday
  - Plans for the coming year.
  - So please take the opportunity and talk to the developers if you have a particular wish.
  - Release process is currently reviewed
    - In terms of release intervals
    - In terms of information flow (EGI, 1.9.12 incident)







## Collaborations, funding and things



EMI came to an end April 30<sup>th</sup>



Consequences for dCache



# MeDIA

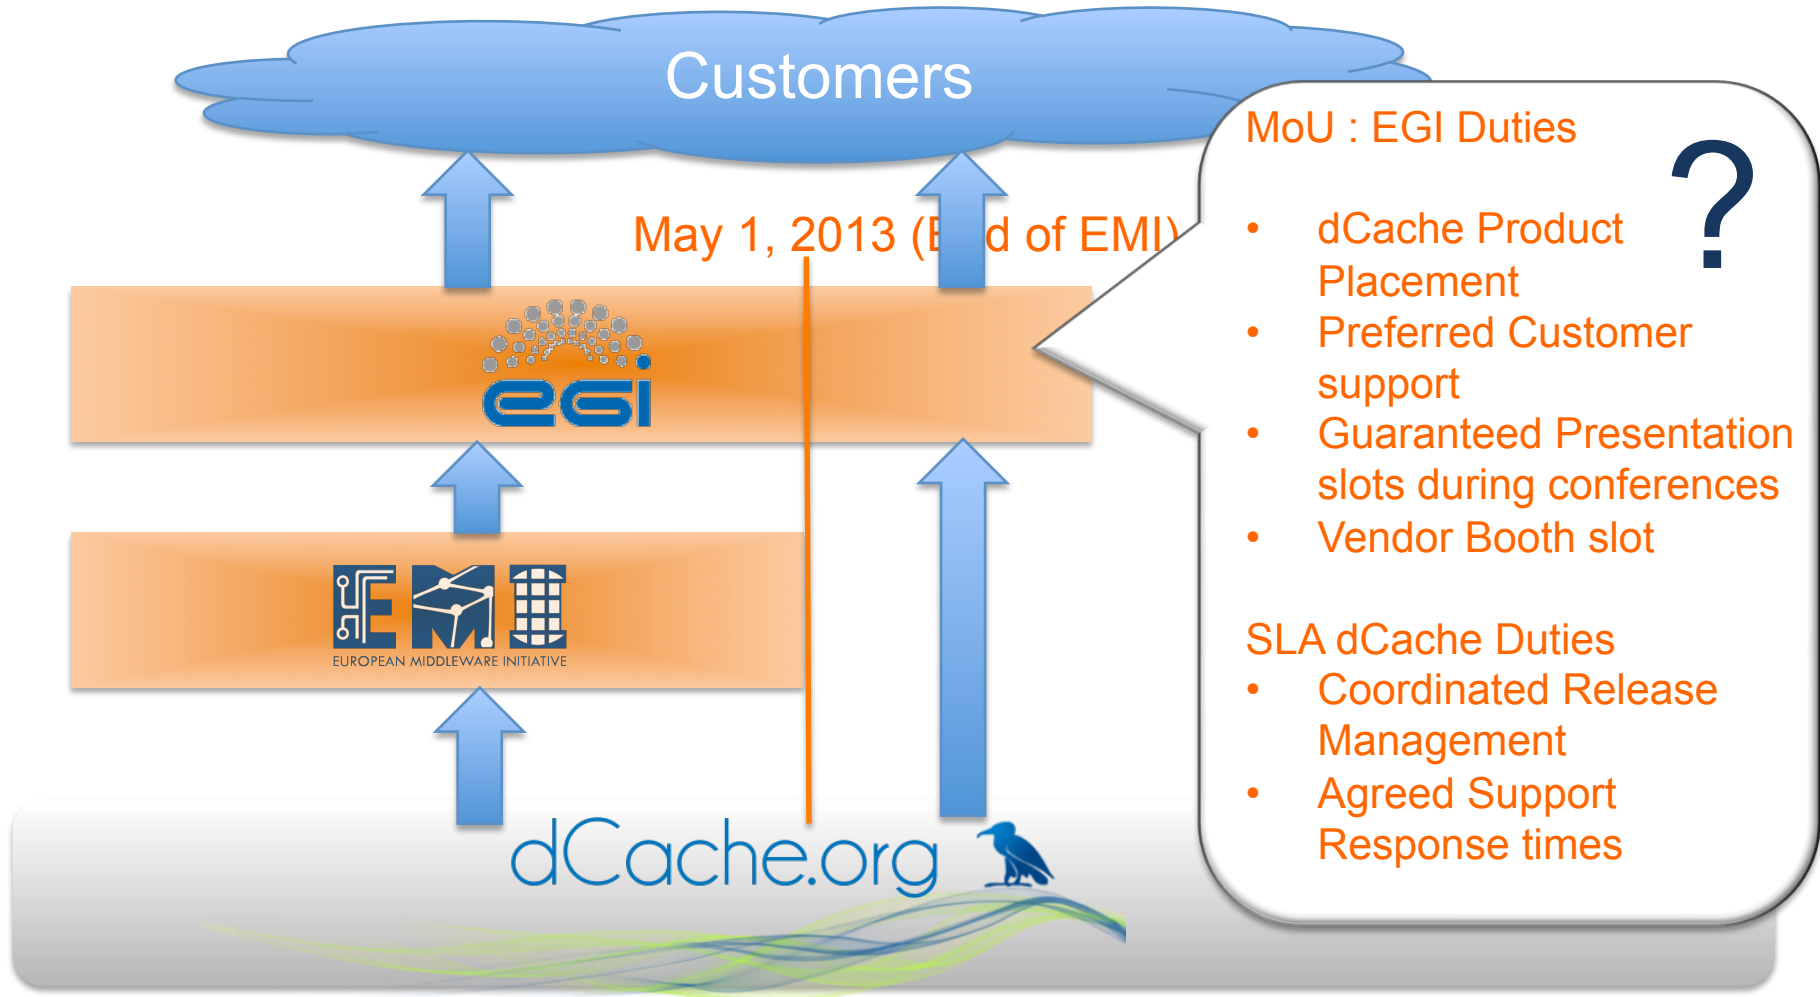
## Middleware Development and Innovation Alliance

- Foundation meeting: 22 April 2013 in Rome
- 37 Participants
- Most middle-ware groups agreed to contribute
- dCache will give it a try
- Main goal is to give European middleware a common contact.
  - The hope is that such a group would simplify requesting more funding in the future.




## In terms of deployment

# In terms of deployment



## Question to the audience

- Anyone has a particular strong opinion on dCache via EGI or not ?





## In terms of funding

# In terms of funding

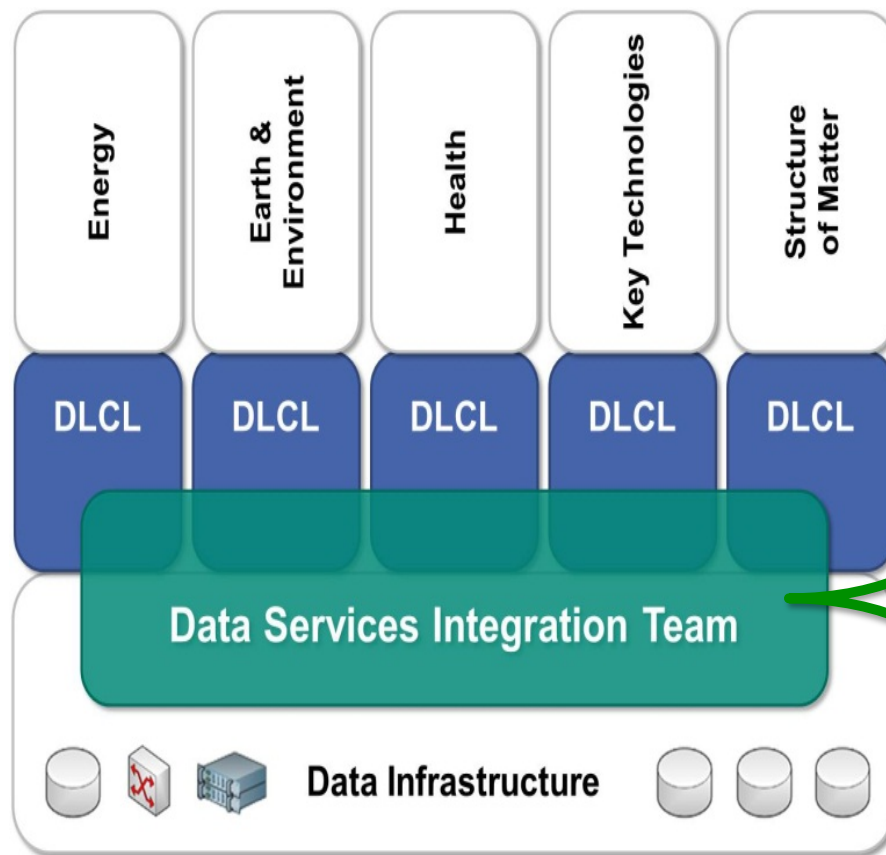


- Dropped EMI funding is completely compensated by LSDMA
- Large Scale Data Management and Analysis
- Technical Lead is Christopher Jung, the former lead of the dCache German Support Group.

LSDMA extends the data services for research of the Helmholtz Association of research centers in Germany with community specific Data Life Cycle Labs (DLCS). ..... The DLCLs are complemented with a Data Service Integration Team (DSIT). It provides generic technologies and infrastructures for multi-community use, based on research and development in the areas of data management, data access and security, storage technologies and data preservation.



# LSDMA and dCache



## Work packages

- WP1: Federated Identity
- WP2: Federated Storage
- WP3:
- WP4: Archives
- WP5:
- WP6:

Collaboration with

dCache.org 



HTW Berlin

- In two areas
  - Implementing Cloud protocols (Jana)
  - dCache/DESY and HTW Berlin will offer an “Online Data Service”
    - To HTW Students first and hopefully followed by
    - The DFN (Deutsches Forschungs Netz)
- Pilot: HTW Students only. (guinea pigs)

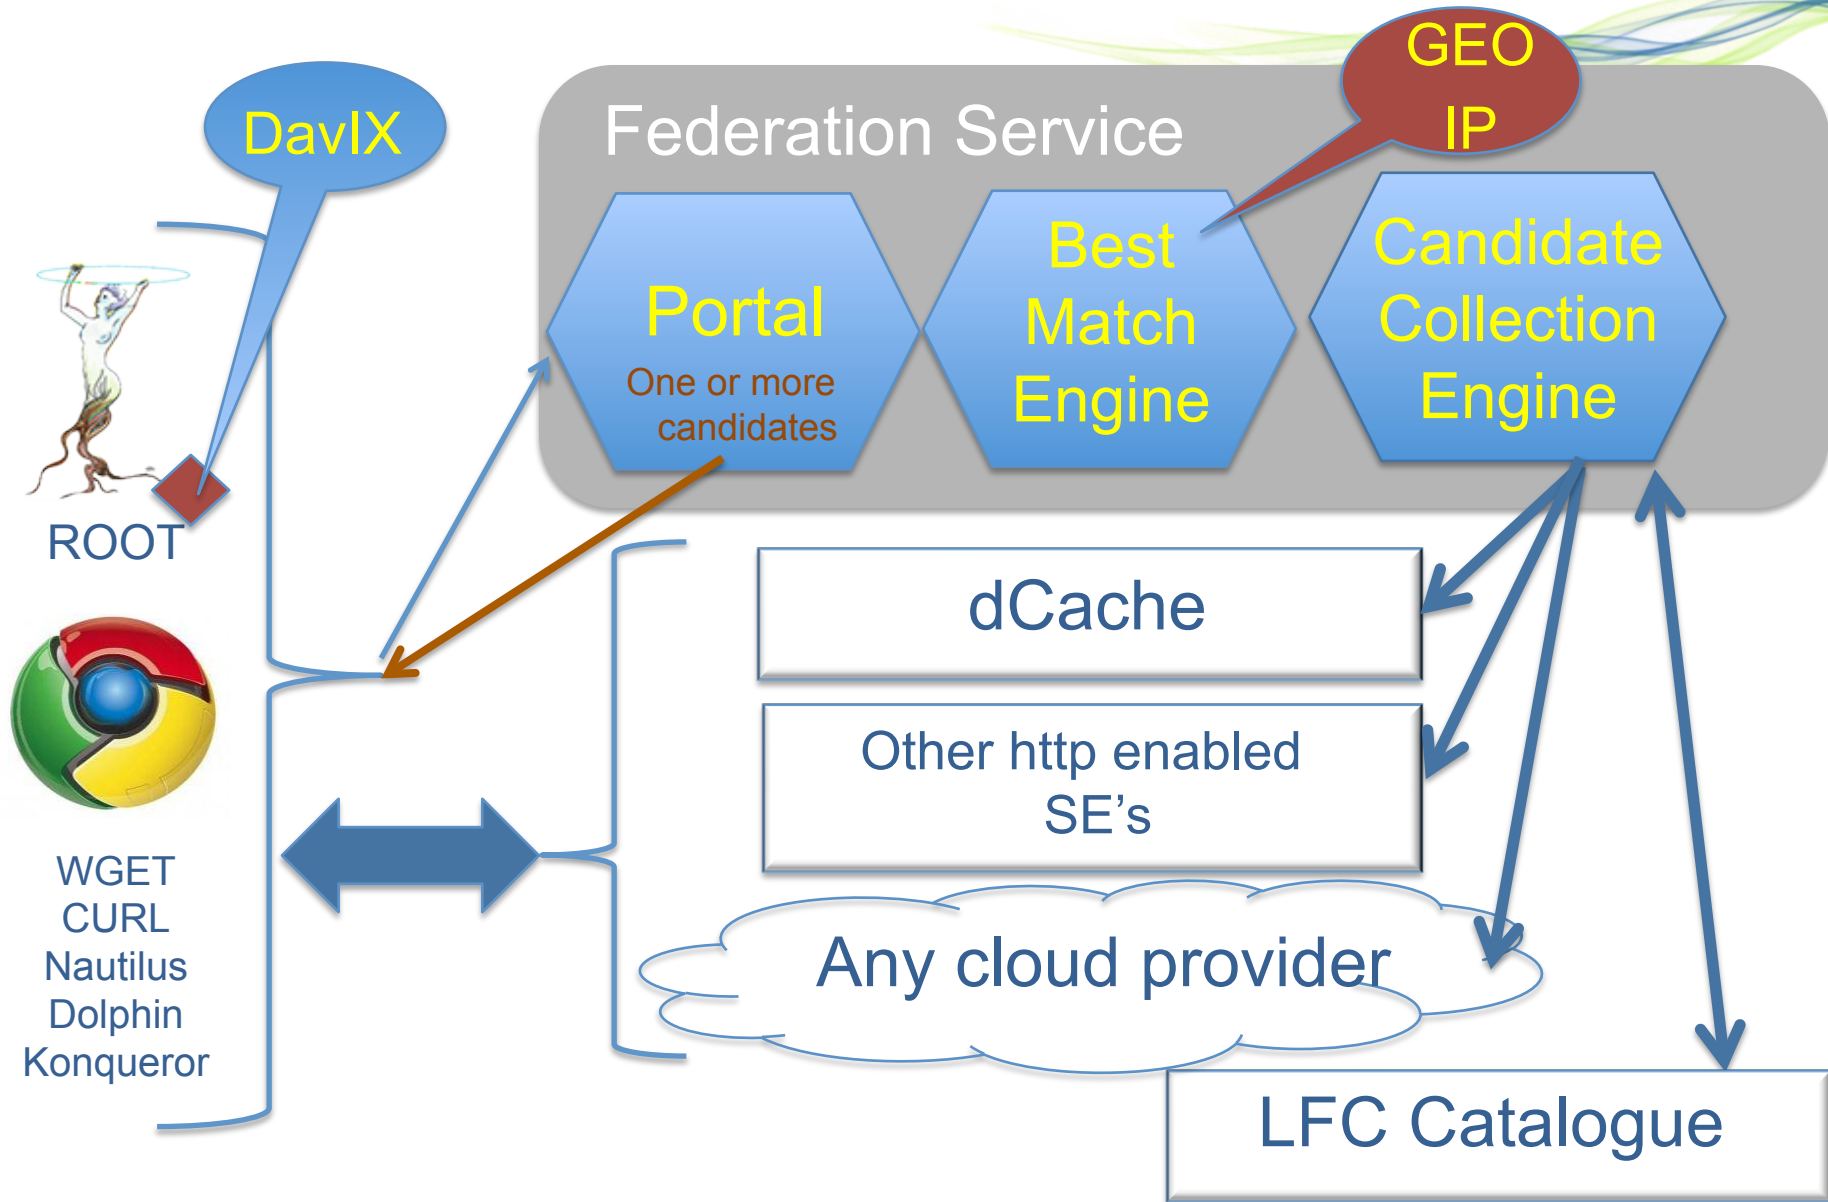
# More Collaborations: dCache and CERN/DM

dCache.org



- CERN and DESY are developing and running a federation service
  - Software updated by CERN DM
- Tigran contributes by integrating the DavIX library into ROOT
  - Essentially an optimized http/WebDAV driver for ROOT
- Status:
  - ROOT integration still in progress
  - Negotiations with WLCG experiments
  - Attracting non WLCG communities (LSDMA)

# Reminder



# Collaborations: Globus Online

(dCache-leaks: taken from e-mail exchange  
between Paul and Steve Tuecke)

- UDP/UDT with NAT traversal
- Mass storage optimizations concerning “Tape backend”
- Firewall friendly improvements
- Data Channel Security Context
- Adding MLSC
- Prototype of a single port GridFTP
- HTTP data channel
- Some more



Stolen from Paul



THE FRENCH NATIONAL INSTITUTE FOR AGRICULTURAL RESEARCH (Pierre Leonard)



**Intensity Frontier**

Muon and Neutrino experiments at FERMILAB






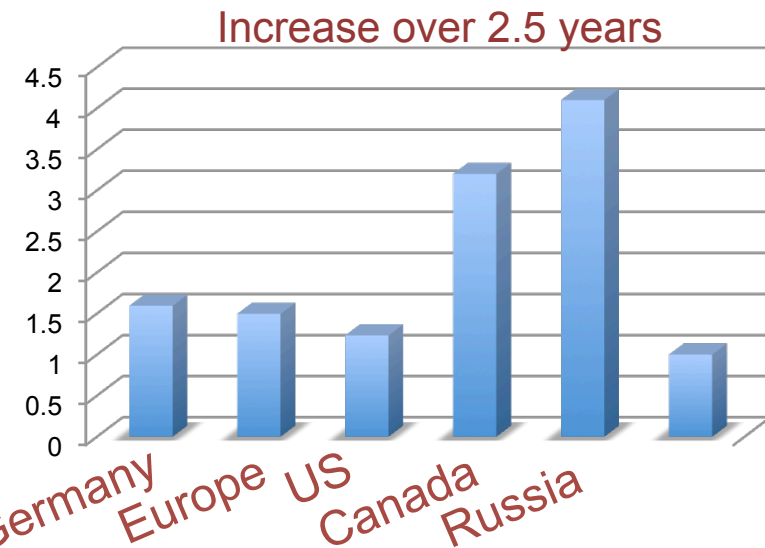
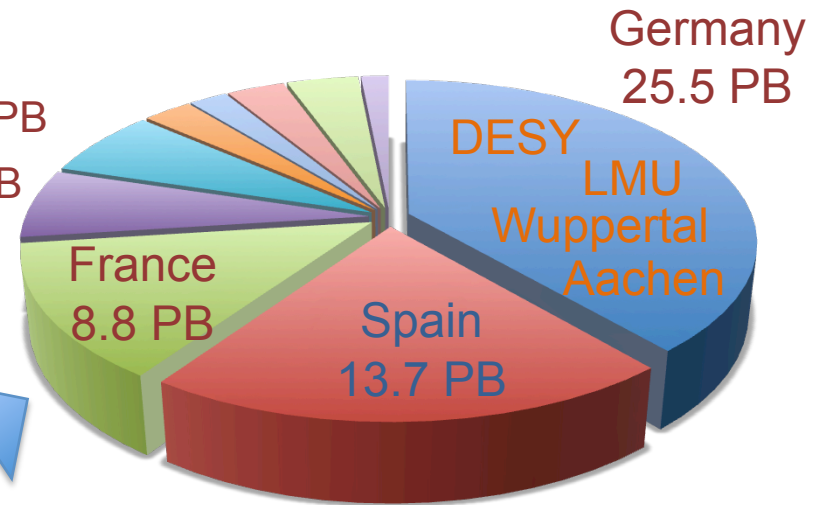
# About deployments

# Deployment (according to BDII)

dCache.org



Nordics 4.5 PB  
Netherlands 4.5 PB






## More deployments news



- 2 new Tier I's in Russia.
  - Dubna (CMS)
  - Kurchatov (Alice, Atlas and LHCb)
  - dCache plus Enstore
- FERMIlab CMS Tier I will be split into a Tape and one Disk-only dCache following CMS recommendations. Other CMS Tier I's will use only a single dCache. More on that tomorrow.

# New candidate for the smallest dCache deployment

dCache.org



# Moving on: The agenda of our “workshop”



# On the agenda

## ★ Kind external contributions ★

- Ron on LDAP and gPlazma
- Fabio on his CMS Tier III

## ★ A lot of theory: a new episode of our masterclasses ★

- Gerd and Paul on
  - Messages and
  - Queues

## ★ Some theory you might need. ★

- Tigran on NFS and ACL's
- Paul on gPlazma II
- Gerd on databases in dCache

# On the agenda

## ★ Very 'hands-on' and hopefully useful ★

- Al, Christian, Dmitry, Paul and Karsten on
  - Alarms, Billing Plots, Log formats and customization
  - Branding your dCache
  - Tuning your SRM
  - Joining the xrootd federation
  - Bits and pieces you might not yet know

## ★ Fish bowl ★

- Feedback from you to us

# On the agenda

## ★ Social event

- Dinner
- Coffee breaks
- Plenty of time to talk



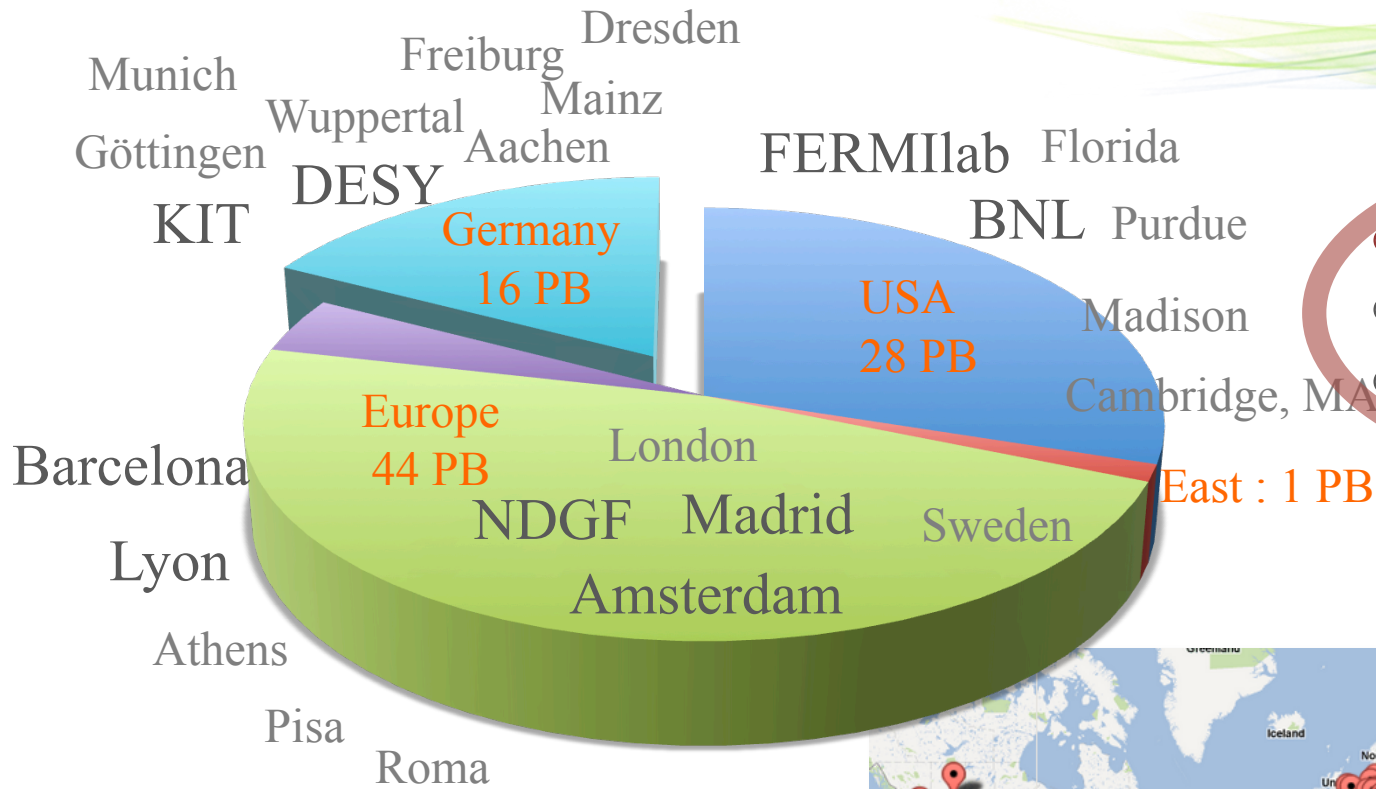


# Enjoy

further reading

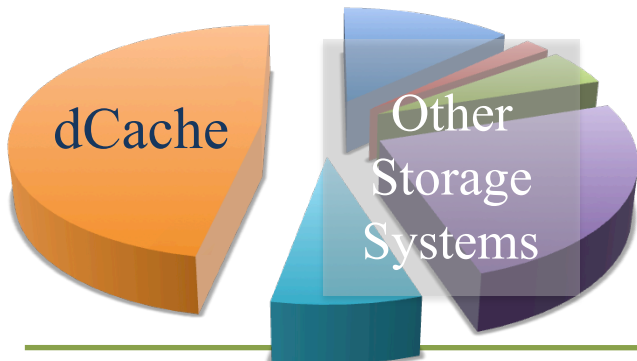
[www.dCache.org](http://www.dCache.org)

# WLCG Deployments



- 100 PB in total
- 7 - 9 Tier I's
- 60 others

Stolen from Tigran







## The dCache labs



# Questions

- Configuration dCache with 2 different interfaces
- Bigger pools, as disks get larger. More pools make more trouble.
- Pools on parallel file systems
- Integrating xrootd auth in dCache w/o plugins
- Discussion of release concept (too fast)

# Content



- CMS Disk / Tape separation
- dCache supporting federated IdM
- Multi Tier Storage
- Small file support to optimize tape
- Object Store
- Scientific Storage Cloud





- CMS is planning to strictly split disk and tape storage elements at their Tier I's.
  - With the available network bandwidth in the OPN, data is rather taken from one other Disk-Tier-1, than from Tape.
- Can be solved in various ways in dCache
  - Have two dCache's: FERMIlab, CMS Tier I
  - Having the namespace twice, pretending to be two different dCaches towards Phedex. Internally the data repository is (almost) only exist once.

# CMS Disk / Tape separation (cont)

dCache.org



- Easy case:
  - Get files from Tape to Disk: BringOnline
- Challenging cases:
  - Migrate disk-only files to Tape
  - Accept files from remote to local disk-dCache, which are already in dCache on tape.
  - Deleting the tape copy only
- Plan
  - PIC is setting up a test instance.
  - We can begin right away with two completely independent copies in one dCache.
  - We can work on the optimization gradually.

# Federated Identities



- General issue:
  - Use credentials from site-A to access data at site-B.
- Plenty of possible combinations
  - SAML or X509 including conversion (e.g. STS)
  - Web-based (including ECP Profile)
  - Generic (no portals involved)
  - And all possible combinations
- We will agree on an example setup
  - Relying party: dCache for sure.
  - Likely SAML support
  - Details need to be negotiated in LSDMA WP1

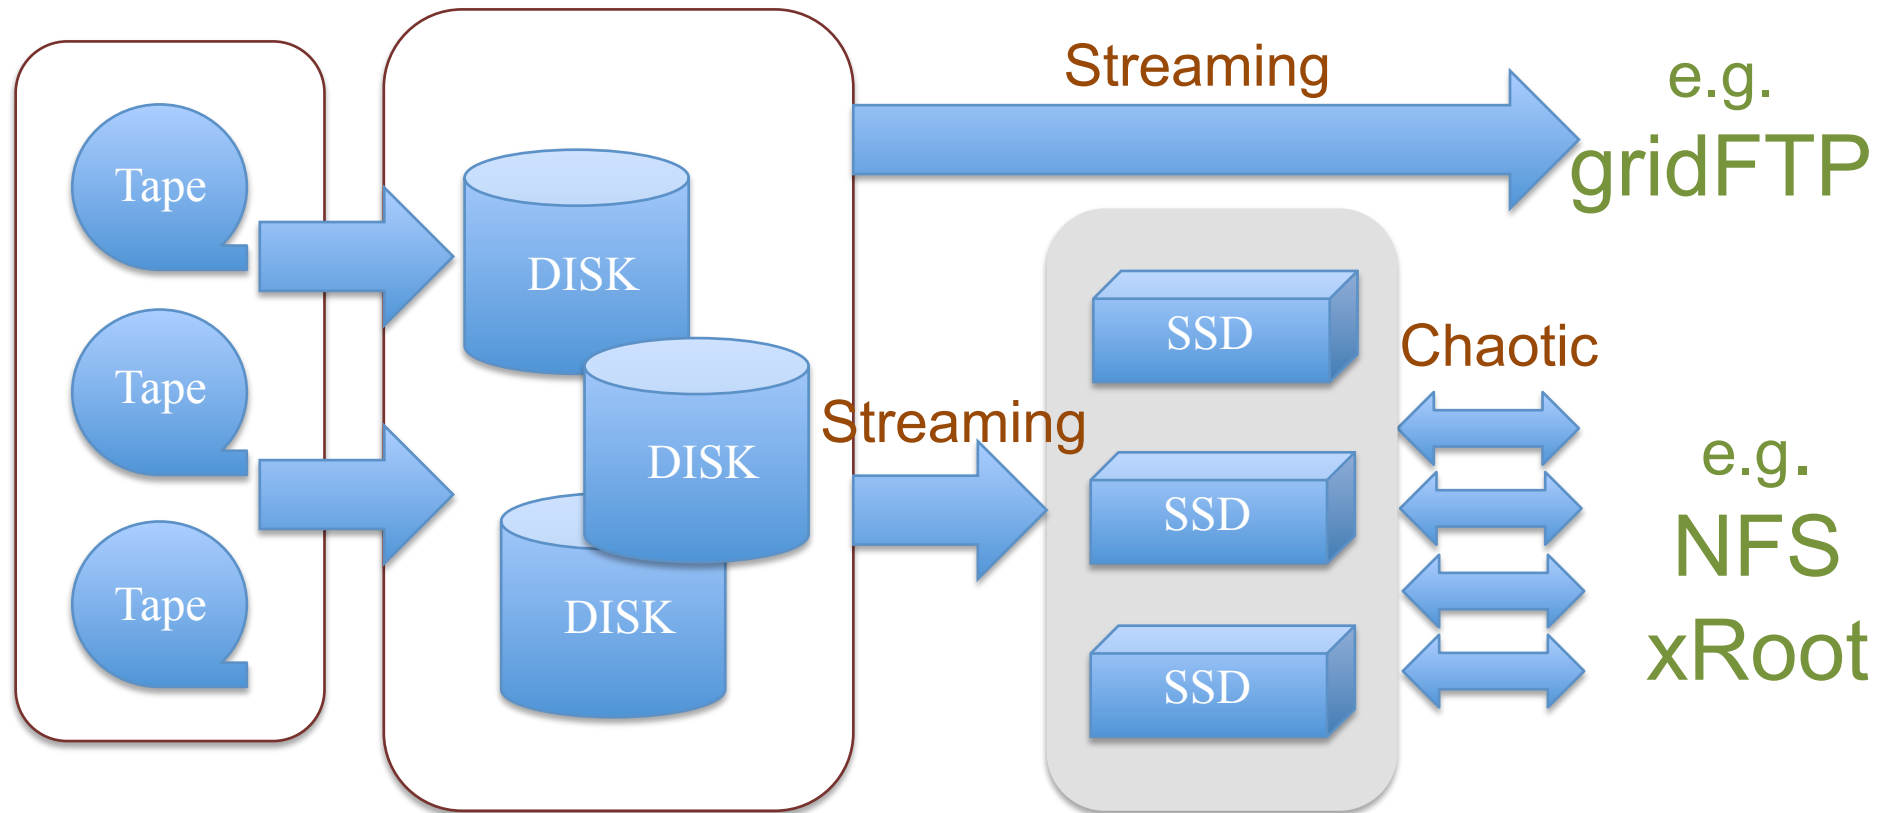
# Example

dCache.org






# Multi Tier Storage




# Multi Tier Storage

- Can already be configured, but



- Tigran : Better would be



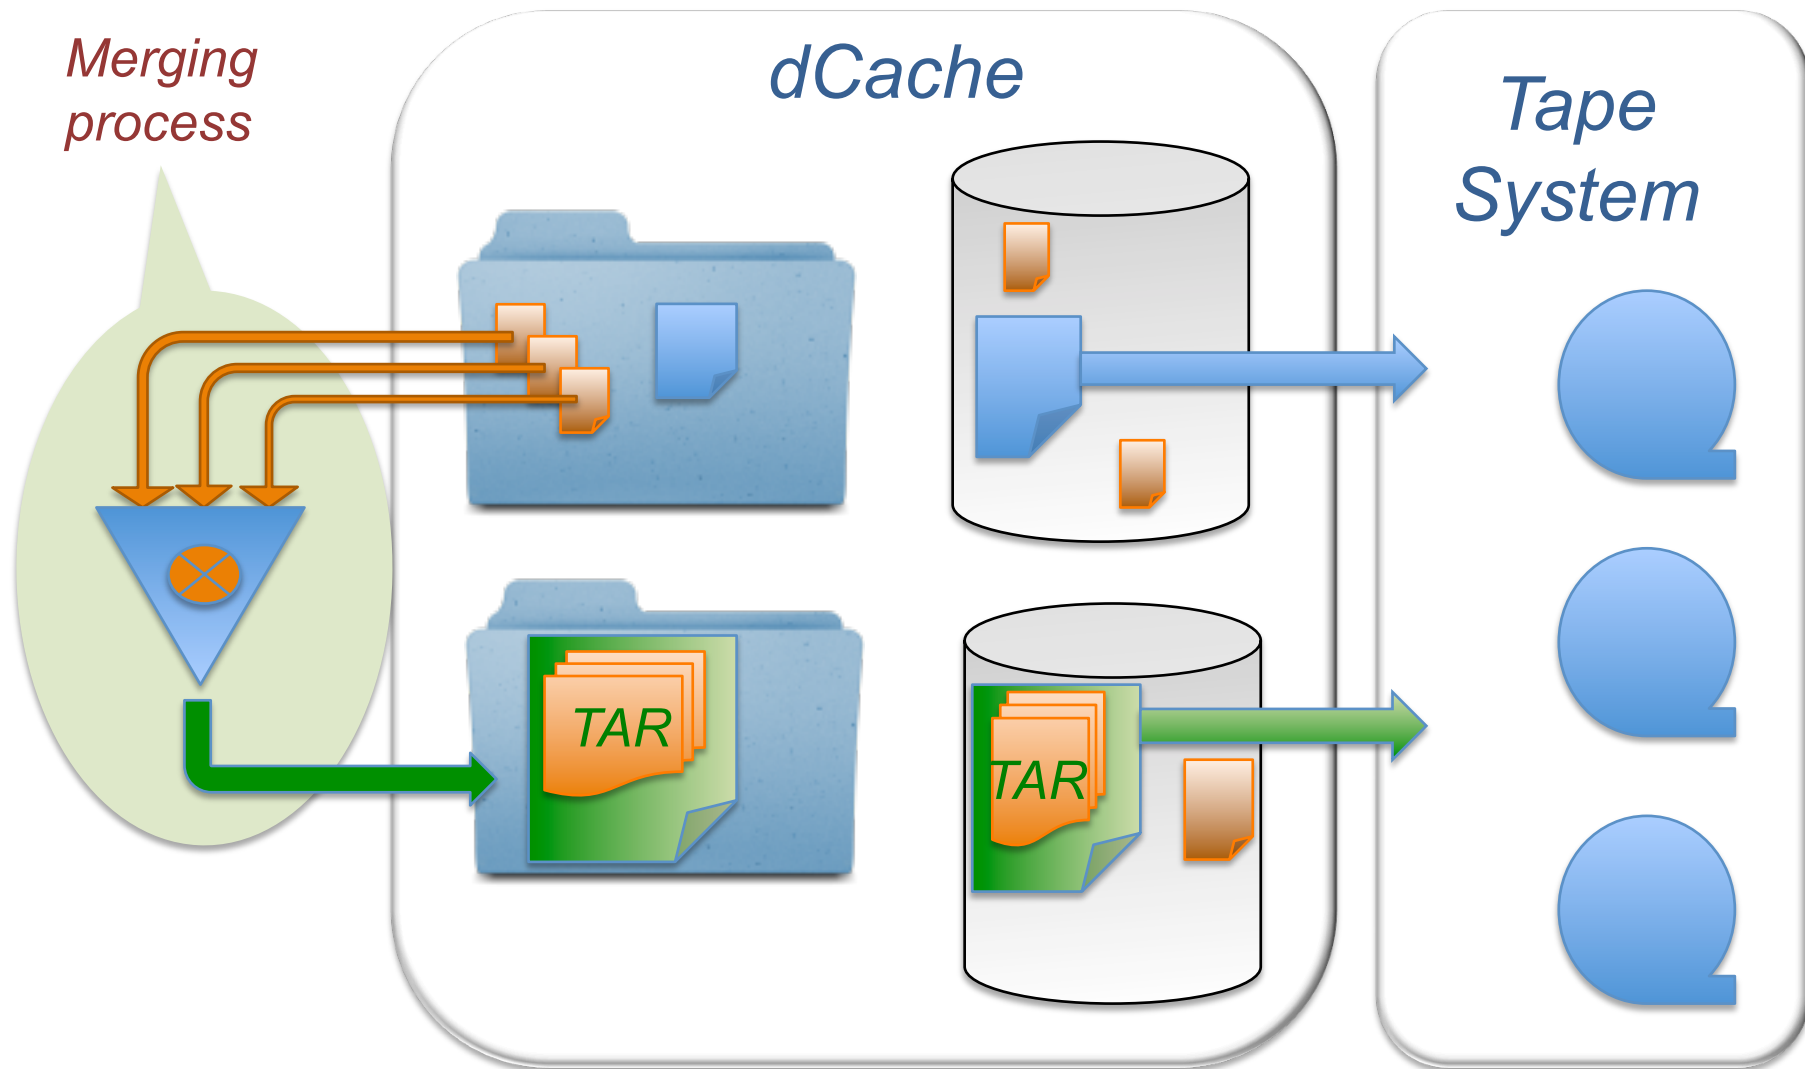
- Will be done, if we find resources

# Small File Support

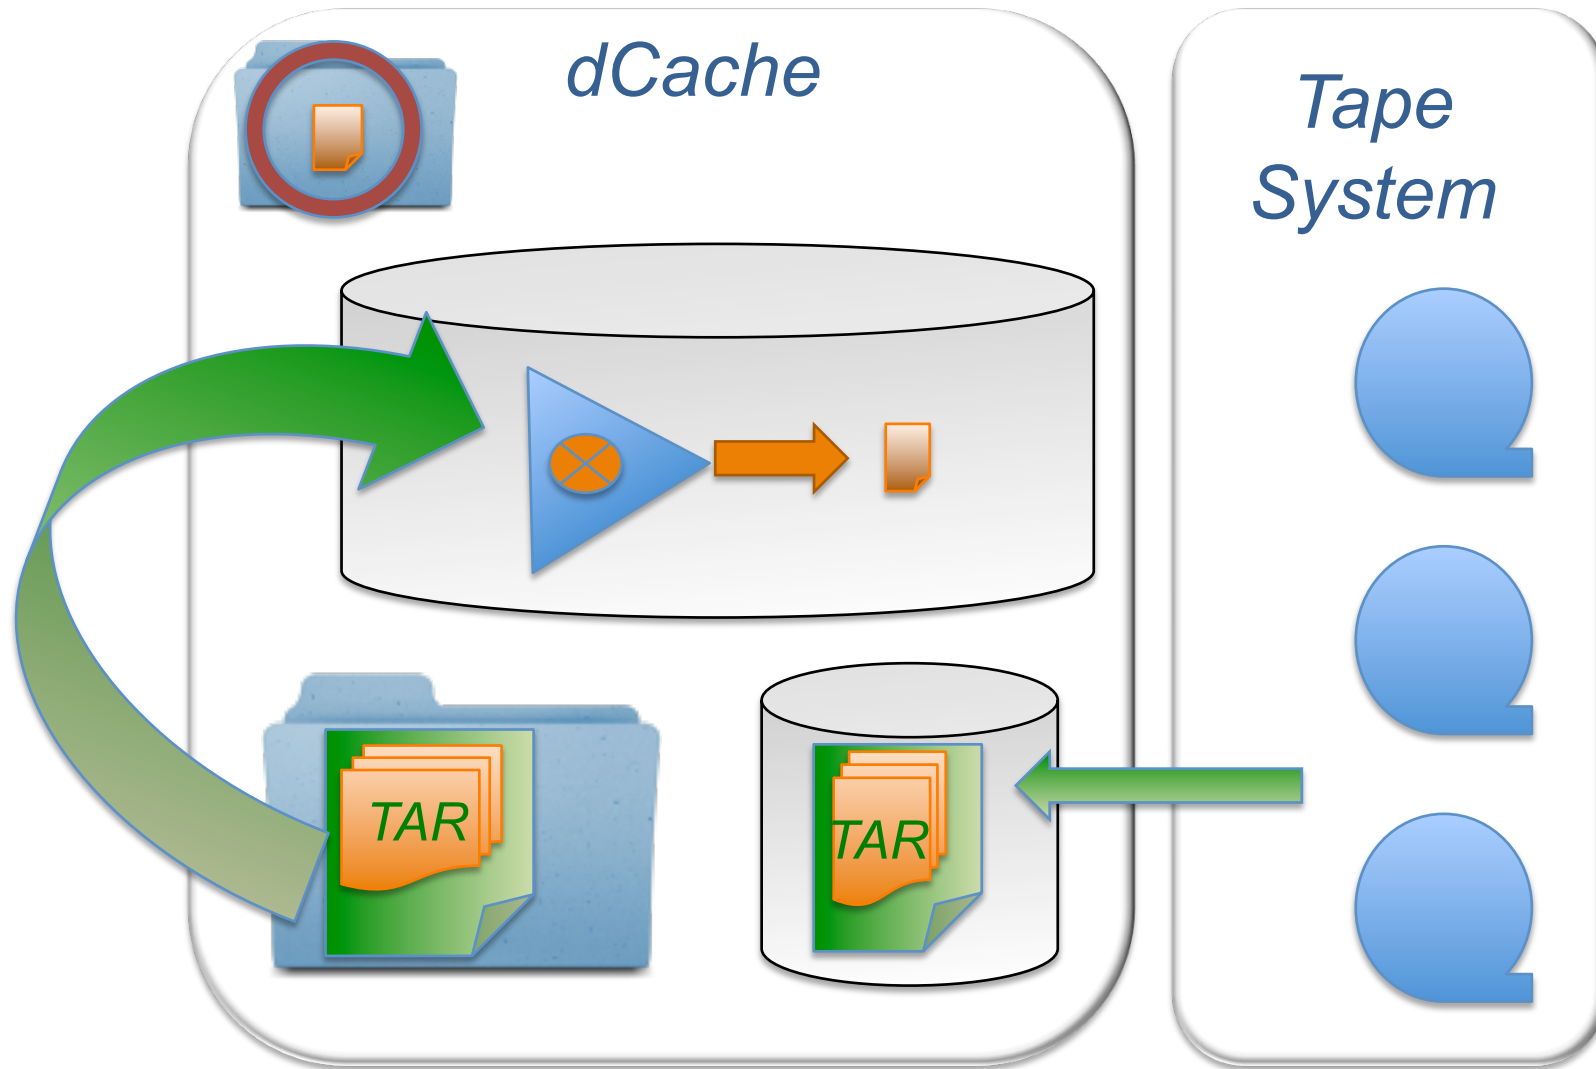


- Transparent for the user:
  - We ‘tar’ or ‘cpio’ files before they are flushed to tape.
  - We extract the correct file from the archive if needed.
- Options:
  - Only the requested file is extracted, or
  - when the first file of a container is requested, dCache could extract all files of the container.
- As the container file is still on disk for awhile after the first file has been extracted (depending on space availability), subsequent requests for small files will be handled w/o further tape access.
- We could even pin recalled containers for some time.
- On top service. Runs on already supported dCache versions.

# Merging small files



# Extracting small file(s)



# Scientific Storage Cloud

