dCache.org

Evolution, by tackling new challenges.

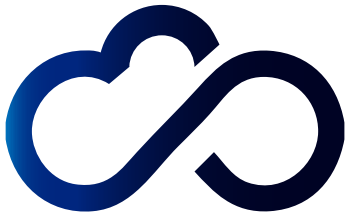

Patrick Fuhrmann

On behave of the project team

dCache strategy

- After 10 years of storage technology support, we feel the responsibility for sites using dCache.
- Therefore our main concern is the efficiency and evolution of those sites in terms of new hardware and software technologies.
- And not to forget: 3 of those sites are actually the authors of that technology.
- We try to achieve this in various ways :
- We try to acquire National and European funding and we partner with projects and sites to provide a sustained support infrastructure. (dCache.org)
- We focus on activities, allowing sites to use our technology for all their customers, not only a particular community. “Alessandra Forti presentation at the WLCG WS”
- One crucial prerequisite is to provide industry standard interfaces and protocols to your storage.
  - Collaborating with CERN DM on various topics in that direction.
  - Great success with http, even in WLCG (See presentations by Oliver Keeble and Johannes Elmsheuser)
- Evaluating new trends in hardware and software, which we might integrate in dCache.
- Exploring new communities to broaden the spectrum of our services.

Who are we ?

The dCache.org collaboration

dCache.org

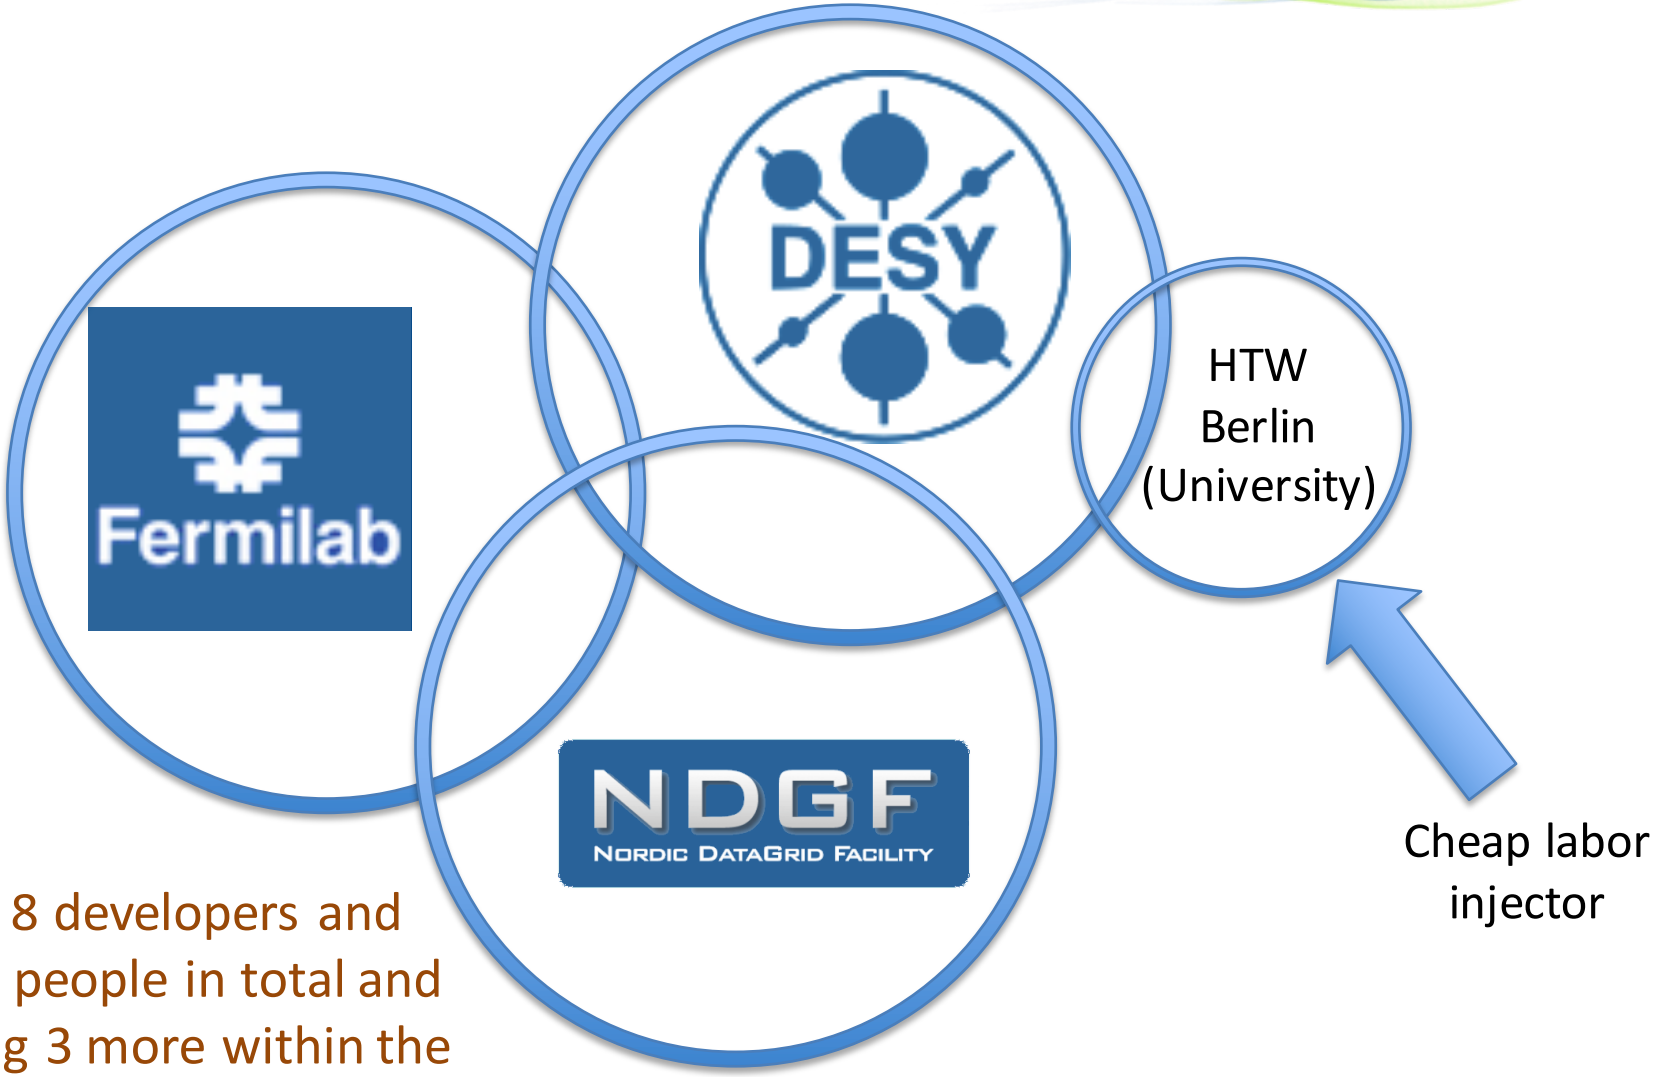

About 8 developers and support people in total and expecting 3 more within the next 3 months.

Funding and high level objectives

Funding and Objectives

dCache.org

EUROPEAN MIDDLEWARE INITIATIVE

LSDMA

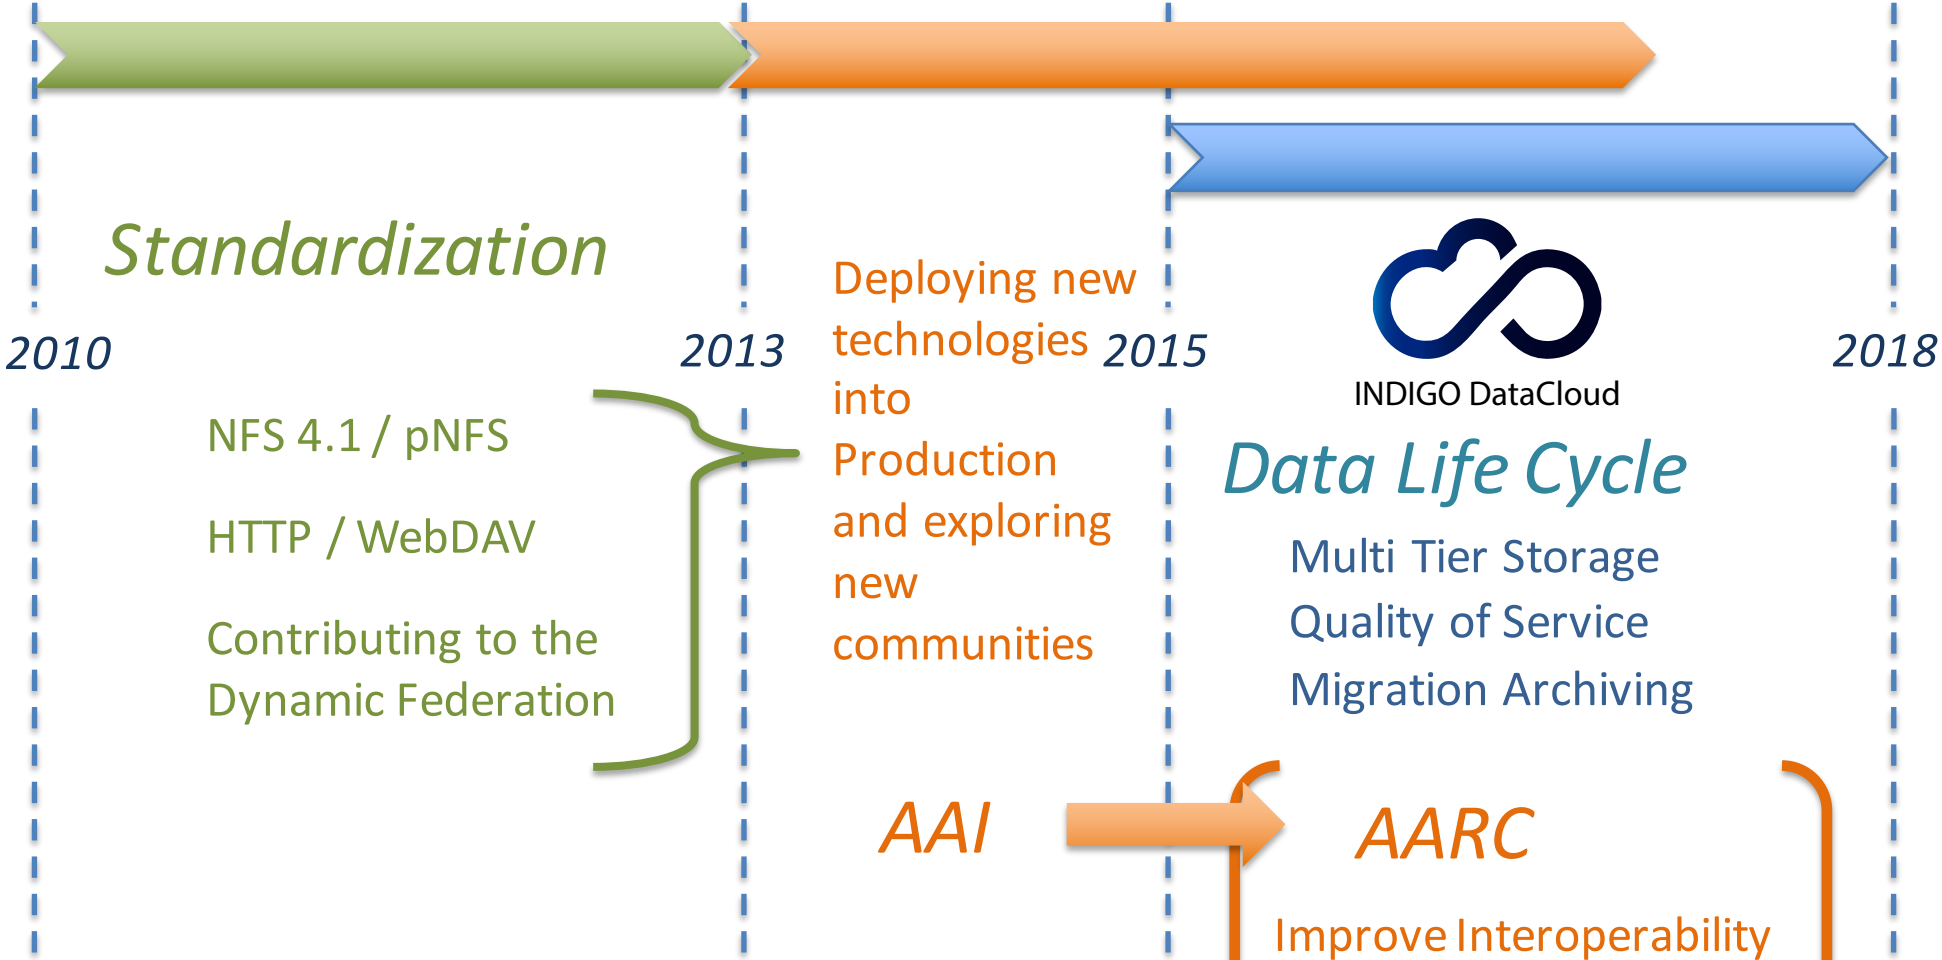

INDIGO Data Cloud Cheat Sheet

dCache.org

- 11 ++ Million Euros
- 30 months duration
- 26 partners
- *The project aims for an Open Source Data and Computing platform targeted at scientific communities, deployable on multiple hardware, and provisioned over private and public e-infrastructures.*
- About 800.000 Euro for dCache.
- ~ 2 more FTEs
- Major objectives for dCache is :
- *"Data LifeCycle Support"* and
- *"Software Defined Storage"*

INDIGO DataCloud

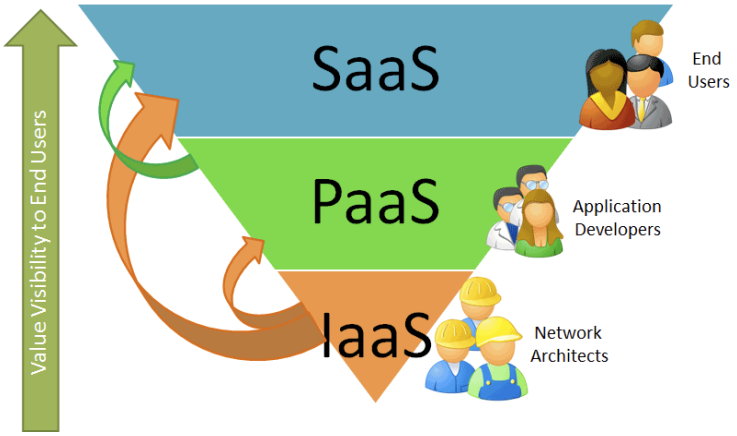

More interesting Challenges

Exploring new communities

dCache.org

- Intensity Frontier (IF) at Fermilab.
  - Quote “Craig Group” (plenary talk)

Nice

- dCache
  - Highly distributed storage with central name space
  - Much lower cost (~$100/TB), ~4PB shared by IF experiments
  - Read / Write interfaces, but does not look like usual file systems
  - Accessible from off-site
  - A cache (optionally front-end to tape system) -- old files are flushed

Hm, actually it does …

HOW ?

NFS 4.1 / pNFS

dCache.org

- Reminder
  - Parallel NFS
  - Clients are directly receiving data from distributed storage nodes.
  - Industry standard, pNFS client in the Linux Kernel.
- Already in use for smaller groups at DESY.
- Slowly migrating CMS Grid worker nodes at DESY to NFS4.1/pNFS data access.
  - Encouraging results (next slide)
- Time consuming, as bugs or misunderstandings are still found in the Linux driver implementation.
  - Disadvantage of standards ☺

Job Efficiency (NFS – dCap)

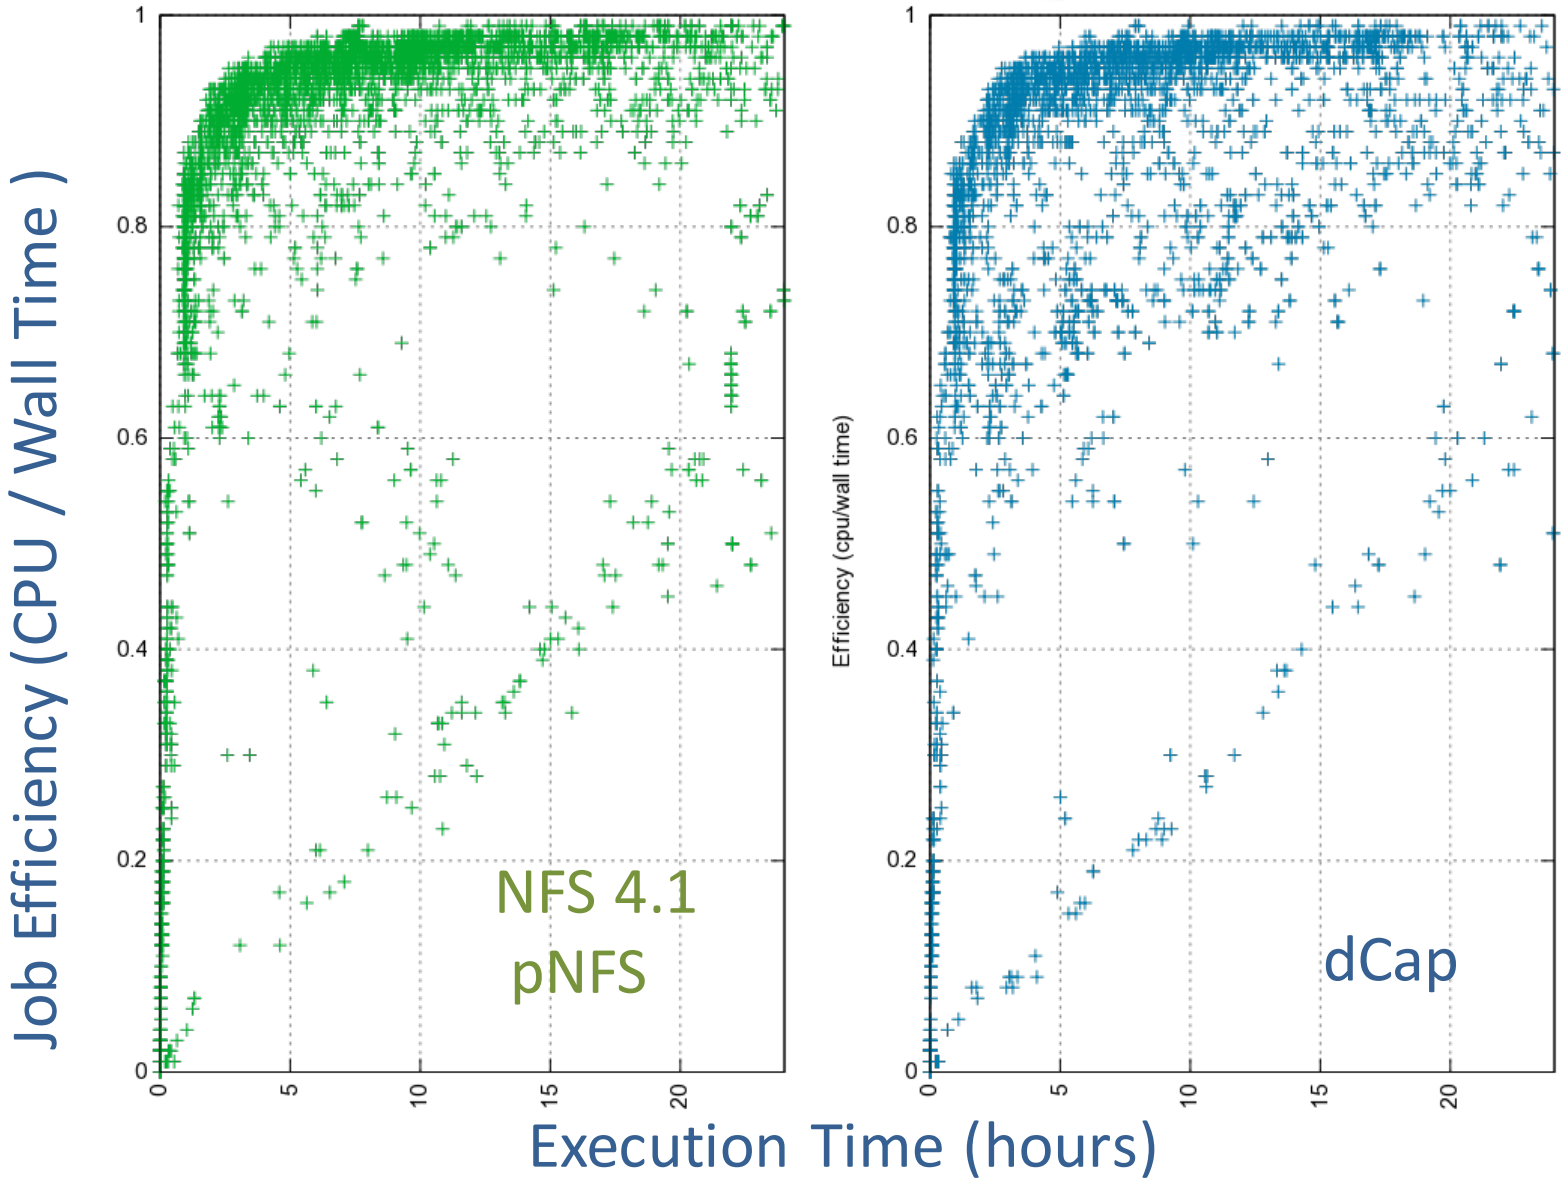

Exploring more …

- German support for the Human Brain Project (SMHB)
  - Jülich – Aachen Research Alliance
  - Distributed dCache between Aachen and Jülich
  - dCache's ability to select pools close to the client or to move data closer to the client made it a perfect match for their requirements.

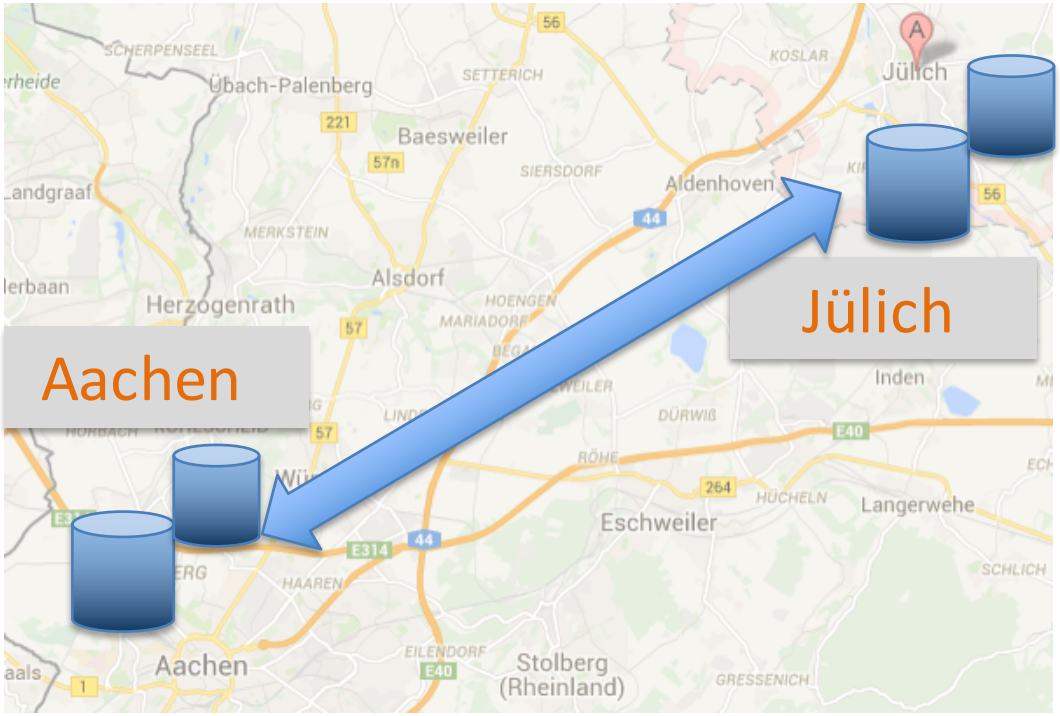

- Two cities, one system.
  - Similar to NDGF (4 Countries one system)
  - Second copy automatically generated at the other location.
  - Or second location just used as a cache.

Projects in HPC

dCache.org

UNICORE

HPC jobs on supercomputer

VAVID

MoSGrid

HPC jobs get access to dCache storage.

Requirement : CDMI

dCache.org

- ISO/IEC Standard
- Important features for the HPC use cases:
  - File selection based on meta – data
    - (not file name based)
  - Supporting remote 'data lifecycle'
    - Bring to / release from fast storage
    - Allow tape migration
    - …
- Required by EGI Fed Cloud
- Supported by INDIGO Data Cloud
- See presentation on CDMI by Paul Millar

Scientific Data Cloud

First Implementation of the Idea : DESY CLOUD

Scientific Data Cloud

dCache.org

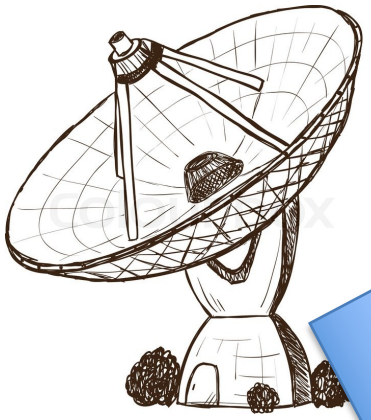

High Speed
Data Ingest

Fast Analysis
NFS 4.1/pNFS

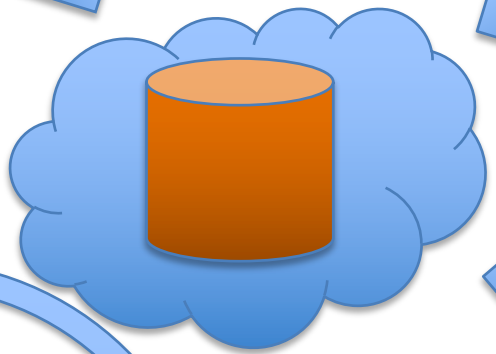

Wide Area Transfers
(Globus Online, FTS)
by GridFTP

Sync'ing and Sharing with OwnCloud

See Paul's presentation

Small file migration to tape

dCache.org

dCache

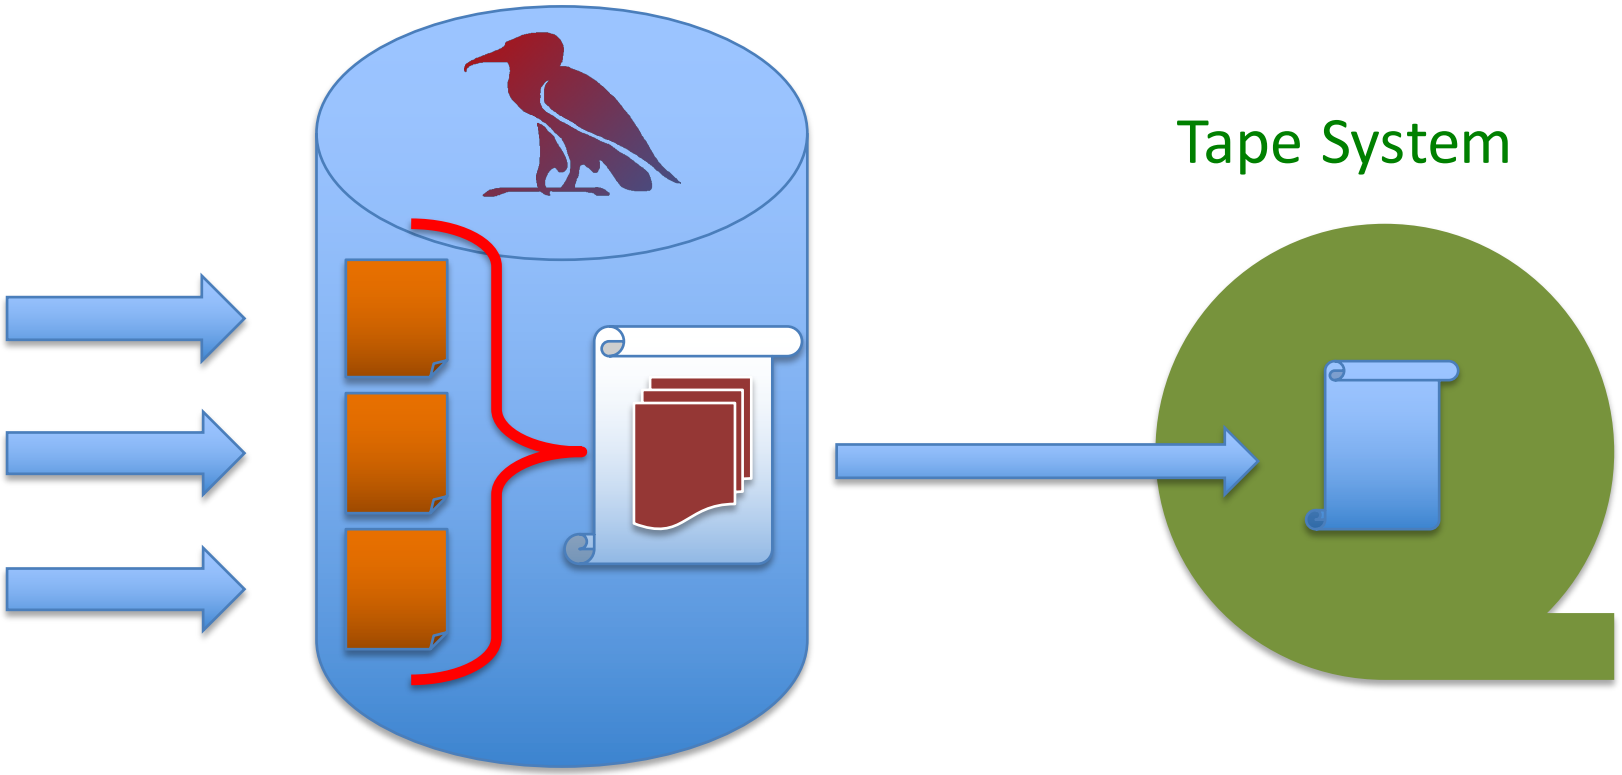

- Currently used by
  - DESY light sources
  - DPHEP
  - NEXT : DESY CLOUD
- See also Poster and Presentation by Karsten Schwank

Responding to new technologies

Response to ceph

dCache.org

- CEPH complements dCache perfectly.
  - Simplifies operating dCache disks.
  - dCache accesses data as object-store anyway already.
- dCache is evaluating a ‘two step approach’.
  - Each pools sees it own object space in CEPH
  - All pools have access to the entire space, which is a slight change of dCache pool semantics.
- Would merge CEPH and dCache advantages
  - Multi Tier (Tape, Disk, SSD)
  - Multi protocol support for a common namespace.
    - All protocols see the same namespace
  - All the dCache AAI features
    - Support for X509, Kerberos, username/password

Summary

dCache.org

- “On Top” funding secured again for 3 more years.
- Storage services based on standards extended our user base towards HPC and ‘long tail of science’ communities and helps sites to reduce software stack costs.
- Wider user base broadens our feature set.
- Continue to investigate new hardware and software technologies and will make them available to our customers.

Don’t forget

dCache.org

The END

further reading

www.dCache.org