

dCache: status and plans

Paul Millar

ATLAS GridKa Cloud T1/T2 F2F
DESY, Hamburg, Germany



dCache server releases

... along with the series support durations.



8th International dCache workshop



dCache German Support

- Group of volunteer dCache admins

Answer questions on mailing list.

Share and publish knowledge on site operations.

Organise and help run dCache tutorials:

GridKa school (KIT, Karlsruhe);

ISGC (ASGC, Taipei);

dCache workshops (various locations).

- Would like to see role of this group grow

see German dCache sites to be exemplary

Communication

- **EGI:**
Patrick in TCB; Paul leads FedCloud AAI; Christian part of “UMD Release Team”
 - **WLCG:** Member of various storage-related groups
 - **CERN Data Management:** direct communication
dCache organised data-management session at EGI Amsterdam meeting.
 - **SLAC:** xrootd collaboration with direct f2f meetings and conferences.
 - **Standards:** OGF, SNIA
 - **Industry:** direct communication with NFS client developers
-

Sustainability: funding

- dCache.org partners:
 - Commitment from DESY, Fermilab and NEIC
 - LSDMA:
 - project continues until end of 2016
 - Horizon 2020:
 - dCache.org participating in two proposals:
Zephyr and **IndigoDataCloud**
-

Sustainability: other communities

... only a selection

- **DESY**: “dCache storage cloud”
providing sync-and-share capabilities
- **Fermilab**: Intensity Frontier (NFS)
- **JADE**: Jülich-Aachen Data Exchange

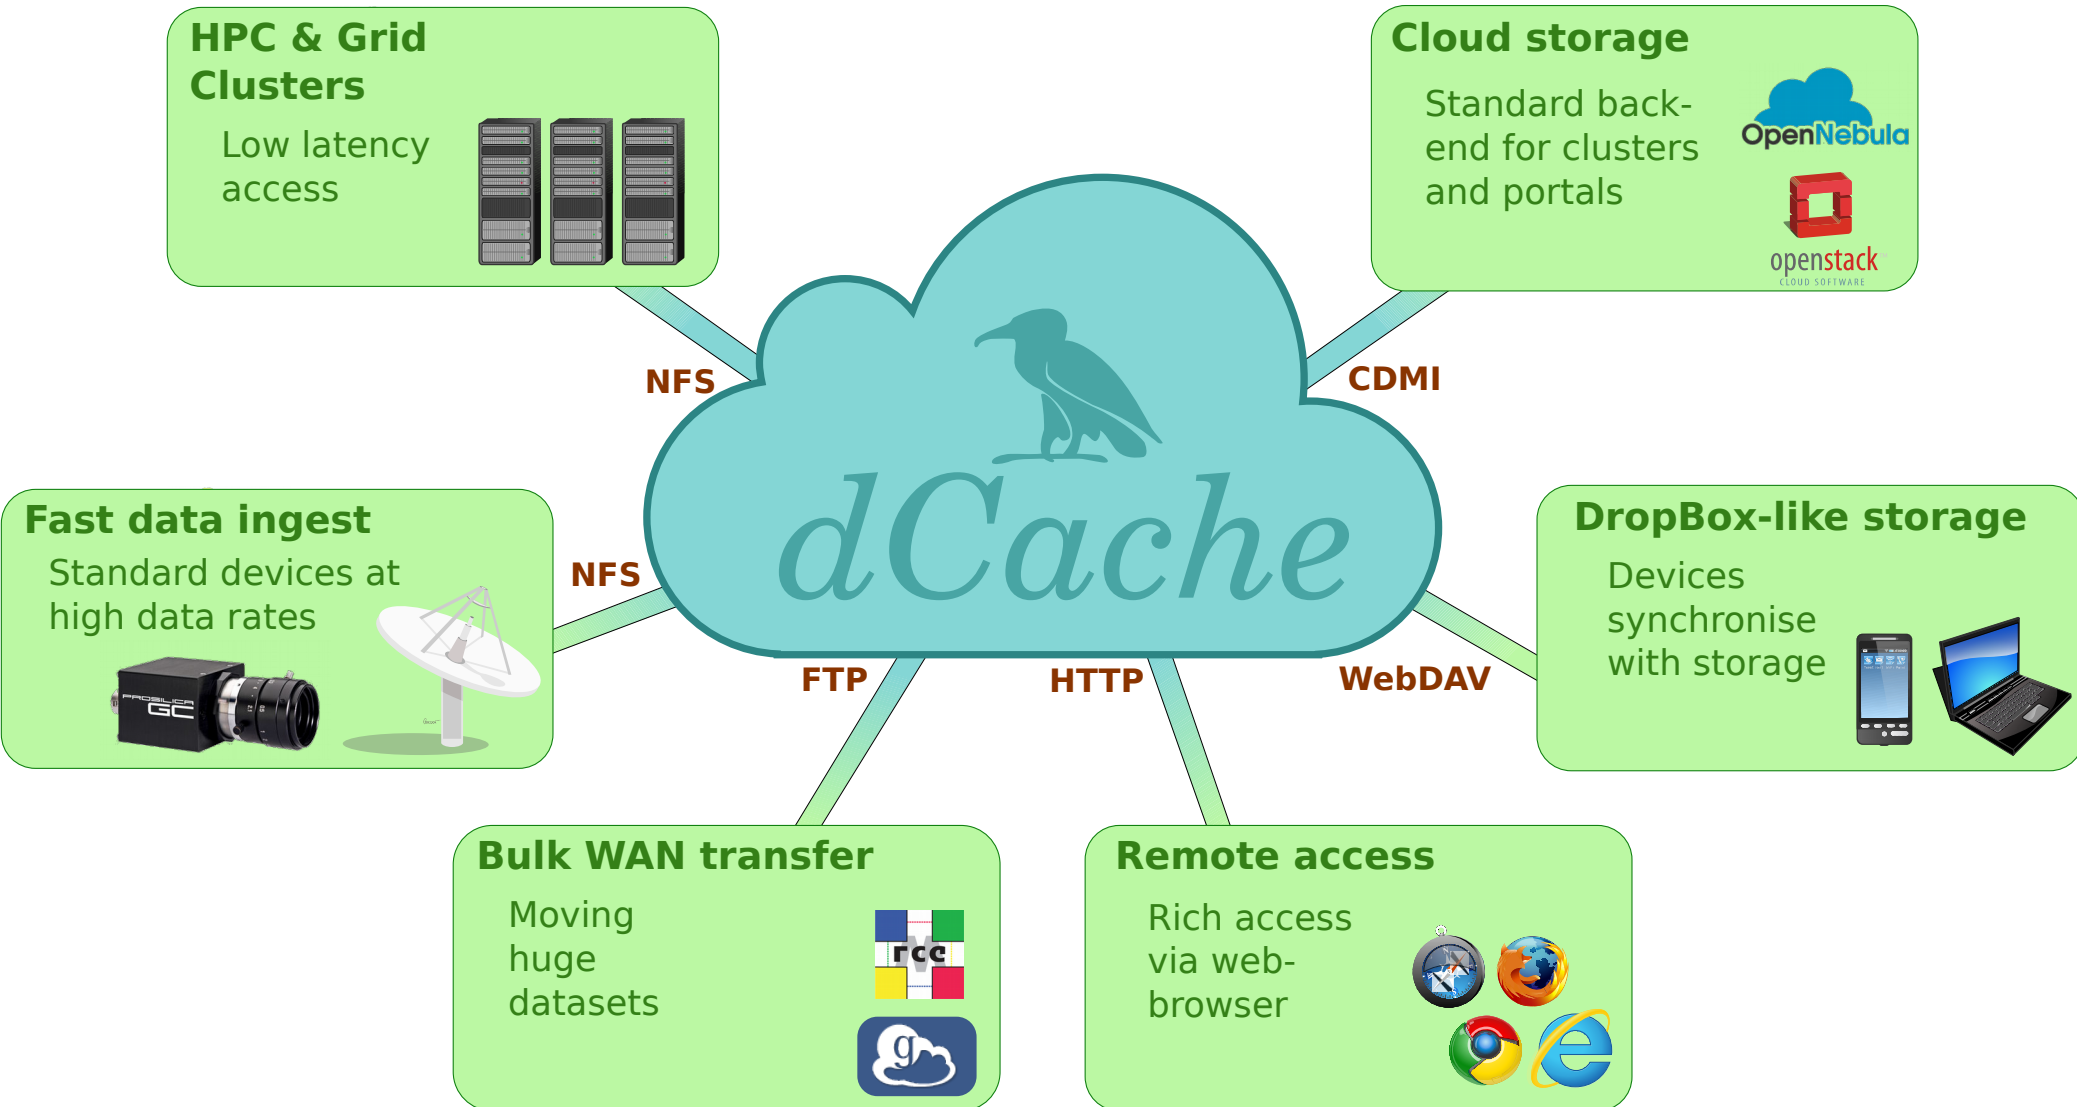


Storage for *Supercomputing and Modelling for the Human Brain* (SMHB), candidate for the *Human Brain Project* (HBP)

- **Commercial:**

DESY directorate signed contracts with an industry partner in Switzerland for dCache support.

dCache the scientific cloud



Activities

- **HTTP ecosystem:**

- Collaboration with CERN; HTTP federation – harmonising xrootd and WebDAV federations.
- With 2.10, dCache supports WebDAV 3rd-party transfers.

- **NFS:**

- Belle I & II support completely based on NFS
 - Photon science: in production for many years
 - CMS (@DESY): start 2nd phase, intent is that all reading is via NFS on a “CHEP time-scale”.
 - ATLAS (@DESY): mid-term roll out similar reading via NFS support.
-

Activities: HTTP Federation

- Project in collaboration with CERN
 - All SEs in federation provide WebDAV access.
 - Central server provides an **aggregate view**
 - Assume that if files exist in multiple servers, they are identical replicas
 - Client sees all available files
 - When reading data, the client is **redirected** to “best” replica.
-

Activities: HTTP Federation

- ATLAS has two prototypes:

federation.desy.de: small number of endpoints

Canadian-Australian fed: sites in CA and one in Melbourne

- Federates path as exported by storage system; e.g.,

```
http://federation.desy.de/fed/atlasdisks/atlasdatadisk/rucio/mc12_8TeV/00/00/A0D.01226672._003195.pool.root.1
```

- Next step: investigate providing a FAX-like view; e.g.,

```
http://federation.desy.de/fed/atlasdisks/rucio/mc12_8TeV/A0D.01226672._003195.pool.root.1
```


Activities

- AAI
 - Mid-term activity at CERN to get rid of X.509 for end-users
 - dCache.org already started investigations (ahead of demand) → work supported by LSDMA
 - As CERN joins SWITCH, so DESY is joining DFN-AAI
 - xrootd 3rd-party copy:
 - We're evaluating the protocol docs and demand
-

dCache with ownCloud

- Use ownCloud on top of dCache, via NFS
 - Files in dCache **owned by the user** (*not* ownCloud process)
 - Users can write data into dCache
 - Immediately** visible through ownCloud.
 - Users can write data into ownCloud (sync client)
 - Immediately** visible through dCache
 - Limitations:
 - If user shares data with you, you can only read that through ownCloud.
 - If you set ACL in dCache, not reflected in ownCloud
 - Service is **live**: currently limited to DESY-IT (as a beta test).
-

How to store small files on tape



CDMI: managing cloud storage

- **Network protocol** for Cloud storage
 - initially by SNIA, now an ISO standard
 - with many, many features
 - Limited vendor uptake:
 - Catch-22: demand and availability
 - Some **IAAS** systems use CDMI internally,
 - the EGI FedCloud has CDMI as a common requirement
 - **Preliminary support for dCache** from student project,
 - Not available now, but plan to integrate (after code review)
 - What is the demand?
-

Software Defined Storage & QoS

- dCache can already provide **differentiated QoS** (Quality of Service):
 - Different files can have different replication factors, multi-tier (SSD, HDD, tape) usage, utilise different hardware
 - Currently these QoS attributes are most configured by the **dCache admin**.
 - We are investigating SDS to allow:
 - Modification of QoS after data is written,
 - Allow users finer grain control of QoS choices.
-

Summary

- Future of dCache is secure
 - Sites should plan their upgrade to 2.10
 - Continue expanding dCache admin community
 - Continue to provide innovative features, based on what will best help dCache users
-

Thanks for listening ... any questions?
