

## dCache Scientific Cloud

Paul Millar

Workshop on Cloud Services for File Synchronisation and Sharing, CERN



# Need a sync-n-share service at DESY

- Hard **requirements**:

- Easy to use,

- Store everything at DESY,

- Integrate with existing infrastructure.

- Anticipated **future usage**:

- change data between syncing and non-syncing storage,

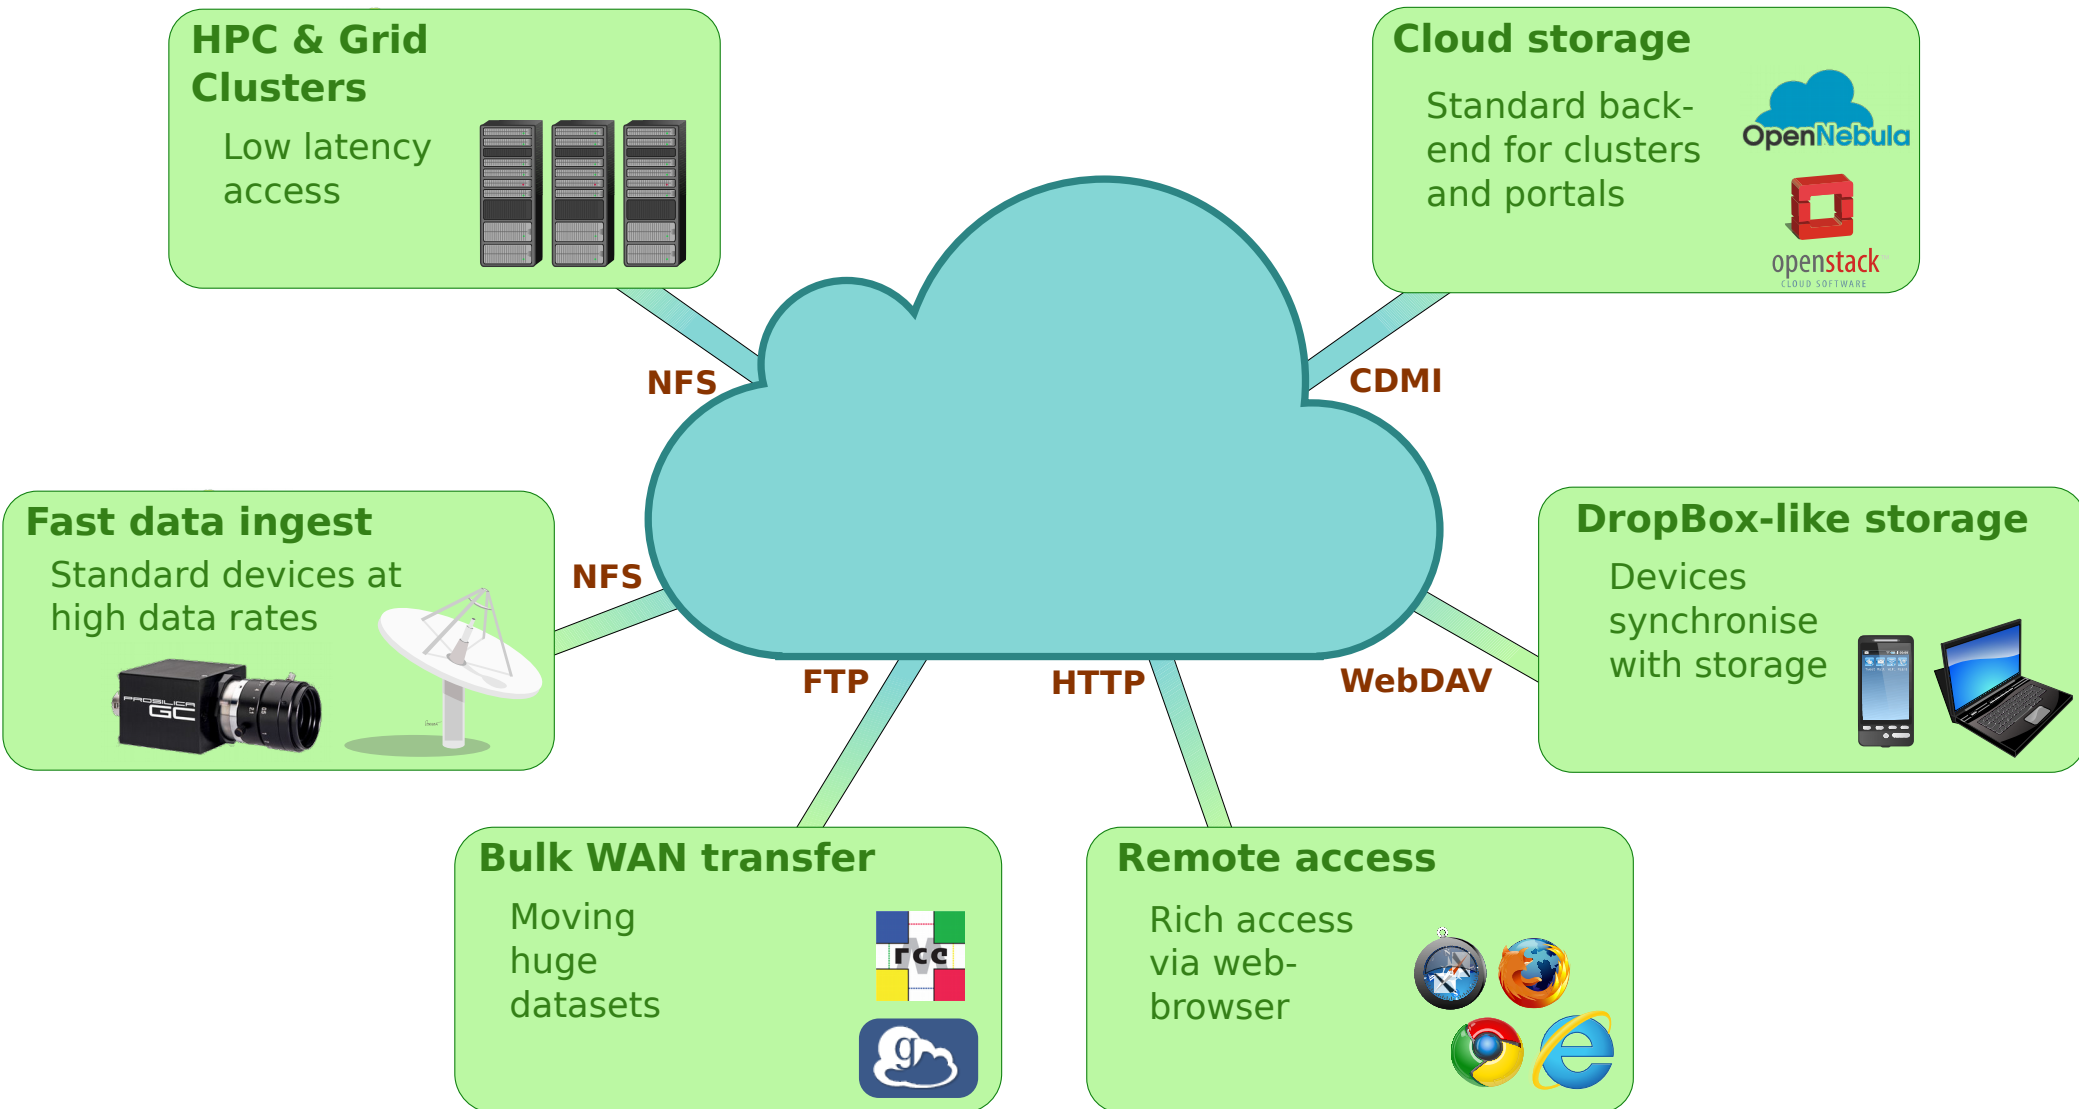
- share data without syncing,

- 3<sup>rd</sup> party transfers between sites,

- immediate access to sync space from compute facilities.

---

# A scientific cloud vision




# How we solved it at DESY

- Looked around, chose two open-source projects:
    - **dCache: powerful managed storage** system
      - Proven integration with scientific data life-cycle.
      - Multiple production instances with tens-of-Petabyte capacity.
      - “Hot” data can be stored on SSDs, “cold” on cheaper HDDs.
      - Provides tape integration; automatic and transparent data migration.
      - ... but no sync and share facilities.
    - **ownCloud: popular front-end**
      - Our collaborators adopting ownCloud makes it more attractive,
      - ... but assumes storage is managed.
  - Combining these two gives DESY the best of both worlds:
    - dCache is mounted on ownCloud server with **NFS v4.1/pNFS**.
    - Integrated with DESY Kerberos, LDAP and Registry.
-



# What is dCache?


## LHC data stored on each storage system




- dCache (96 PB)
- DPM (34 PB)
- EOS (0 PB)
- StoRM (20 PB)
- CASTOR (14 PB)
- BeStMan (7.6 PB)
- Globus FTP (6.1 PB)
- ARC (0.01 PB)
- xrootd (22 PB)

Source: BDII (2014-11-14)







5 FTEs



2 FTEs



1.5 FTEs



dCache


Core team


**Student mentor programme**

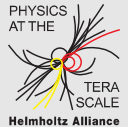



Hochschule für Technik und Wirtschaft Berlin  
3 students


**Collaborations**














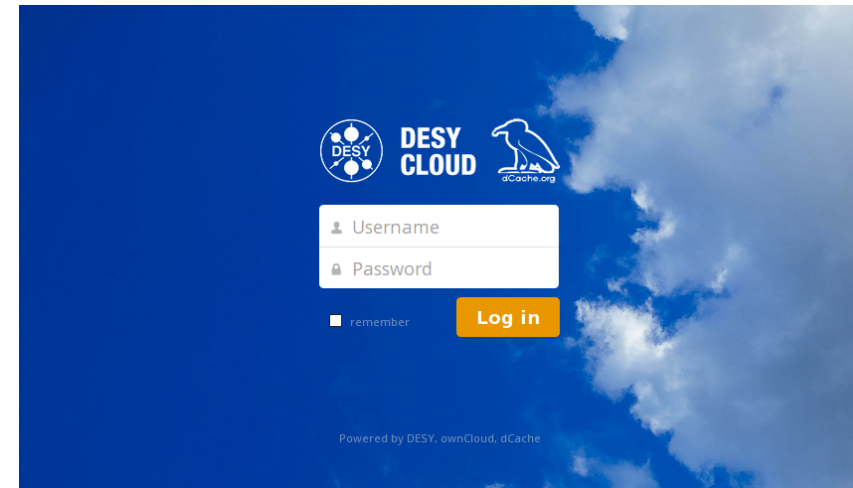






# The DESY Cloud service



- Status: **production**, but for IT only;  
in two weeks “by invitation”  
(selected power-users);  
1<sup>st</sup> January general availability.
- Required minor **patches** to ownCloud & dCache.
- Changes pushed into dCache:
  - Running **unpatched dCache** release
  - Have a blueprint for **any site** to deploy ownCloud+dCache.
- Changes pushed **upstream to ownCloud**:
  - Not all were accepted for v7, so running patched version.
  - Will upgrade to v8 when released.





# Development and future work

- Files in dCache have **user-ownership**, not ownCloud:
    - Plan to expose files directly from dCache: NFS mount, 3<sup>rd</sup> party transfers, direct access from any grid worker-node, ...
    - Couldn't fix ownCloud: work-around within dCache
  - **Consistency** between ACLs and shares:
    - dCache ACLs to honour ownCloud shares and vice versa
  - **Redirection** support for sync-client:
    - ownCloud server proxying data is bottleneck; want syncing to be more efficient by taking data from where its stored.
-

# NFS v4.1/PNFS vs ownCloud (currently)



# ownCloud: currently vs with redirect



## Experience: problems with ownCloud

- If underlying FS disappears, **all sync-clients delete all data.**
  - If underlying FS returns **EIO** on read, sync-client creates 0-length file: **impossible to recover.**
  - Bulk delete through web interface is **unreliable** (under investigation).
  - Admin interface **awkward** with  $O(5k)$  users.
-

## Not just ownCloud ...

- dCache team hosted a **two-day workshop** with project- and technical-lead of DCORE
    - Provides cloud storage with features beyond ownCloud
    - Some “big name” customers
  - Initial “lite integration” by December  
(includes redirection support)
  - Then providing “tight integration” with shared namespace
-

Thanks for listening ... any questions?





# Backup slides



# Thinking about sync-and-share

- Like other systems, small fraction of data is “hot”
    - SSDs provide better performance, but can't afford only SSDs; nice to have system that places hot data on SSDs, cold data on HDD.
  - Amazon had a smart idea: allow people to choose how much to pay
    - Let users choose between Normal and Glacial QoS; e.g., disable sync for Glacial-like storage but allow access via web interface
-



**DESY  
CLOUD**



 Username


 Password

remember

**Log in**




Powered by DESY, ownCloud, dCache

# WLCG dCache instances (only WLCG sites shown)









# DESY now a photon science lab



# Over 10 years “Big Data” experience

Era	Disk cache	Grid Storage	Generic Storage	Cloud Storage
Additional Communities				
Additional Authentication	<p>Trusted host</p>	<p>X.509, Kerberos</p>	<p>Username+PW</p>	<p>SAML, OpenID, OAuth, Token, ...</p>