





dCache - recent developments

Paul Millar

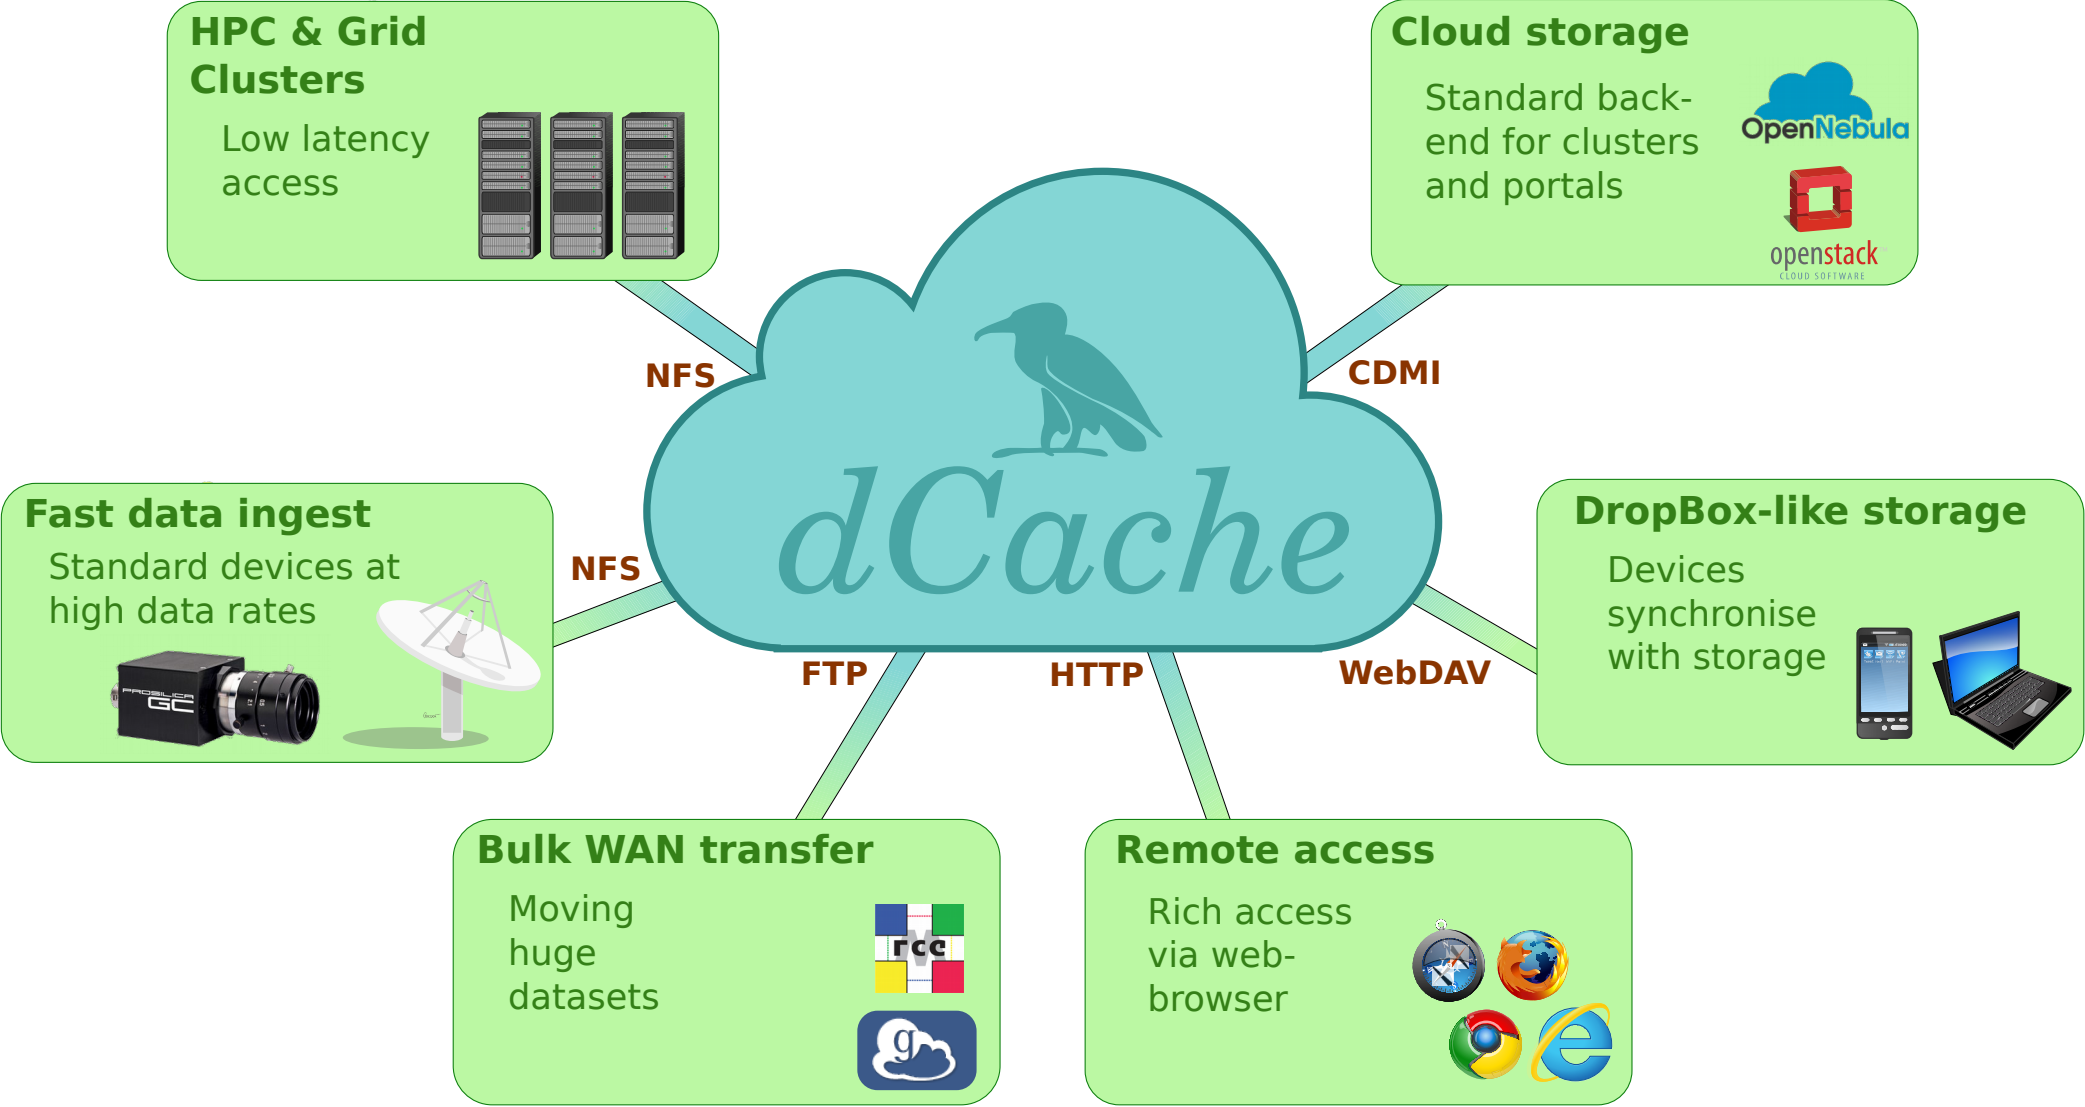
8th Annual Helmholtz Alliance Workshop on "Physics at the Terascale"
DESY, Hamburg, Germany



dCache evolution for Big Data and Cloud

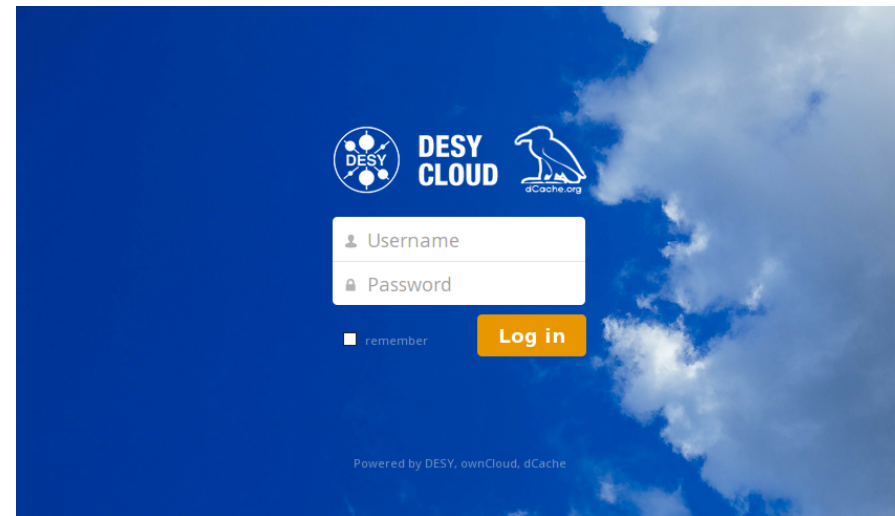
Era	Disk cache	Grid Storage	Generic Storage	Cloud Storage
Additional Communities				
Additional Authentication	Trusted host	X.509, Kerberos	Username+PW	SAML, OpenID, OAuth, Token, ...

dCache the scientific cloud



The DESY sync-and-share service


- Looked around, chose two open-source projects:
 - **dCache**: powerful managed storage
 - Proven integration with scientific data life-cycle
 - ... but currently 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.
 - In the future, users can access files either directly through dCache or via ownCloud.



The audience!



Small files and tape-storage



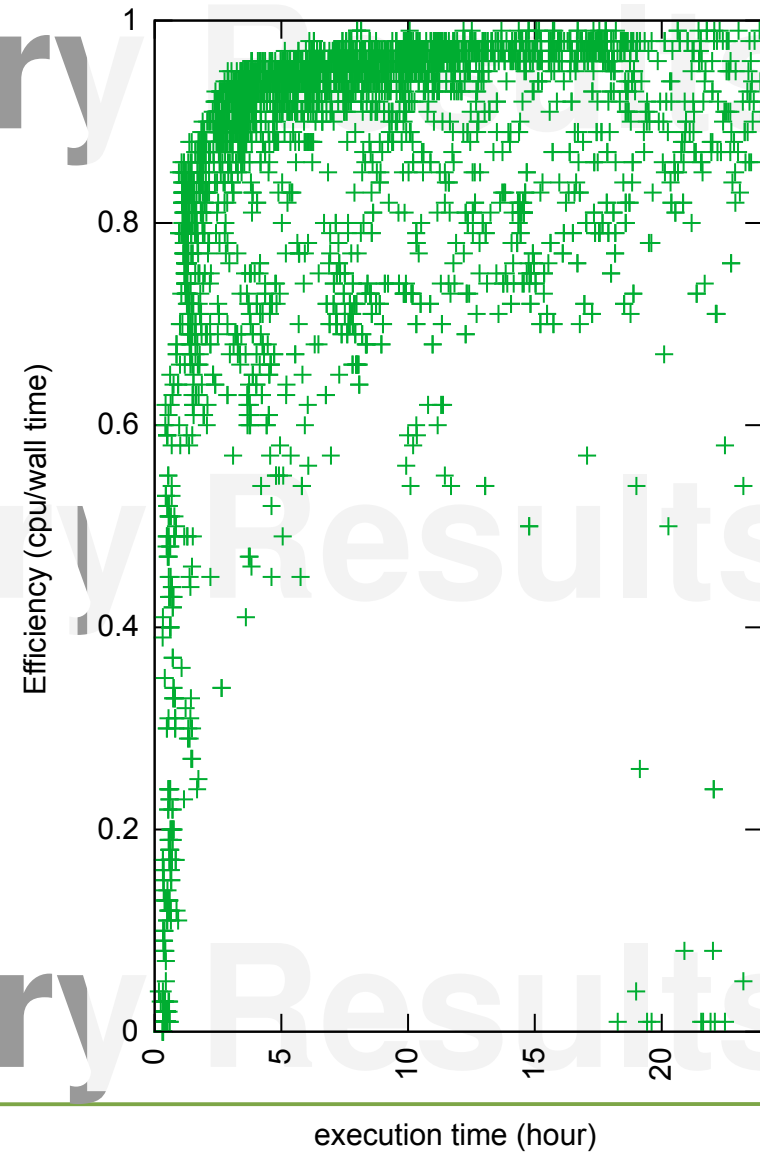
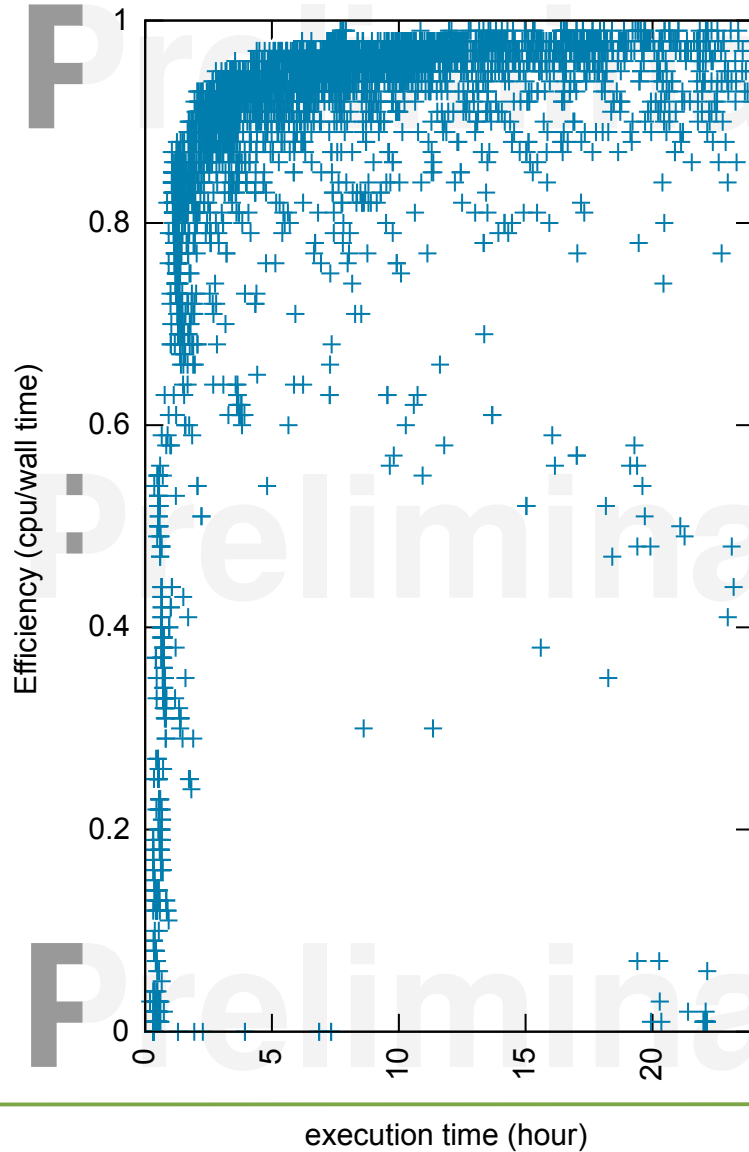
The future is NFS...

- Adopting **industry standard** allows us to support non-grid users:
 - The only way to support BELLE and Photon Science communities.
 - At DESY, existing users **moving to NFS**:
 - HERA (DP-HEP) now only available through NFS,
 - NFS mounted CMS storage on BIRD/NAF,
 - Starting migrating CMS Tier-2 jobs to use NFS for reading (currently ~50% of jobs).
 - Fermilab **Intensity Frontier** uses NFS-mounted dCache.
-

CMS job efficiency by access protocol

DCAP (6664 recs.)

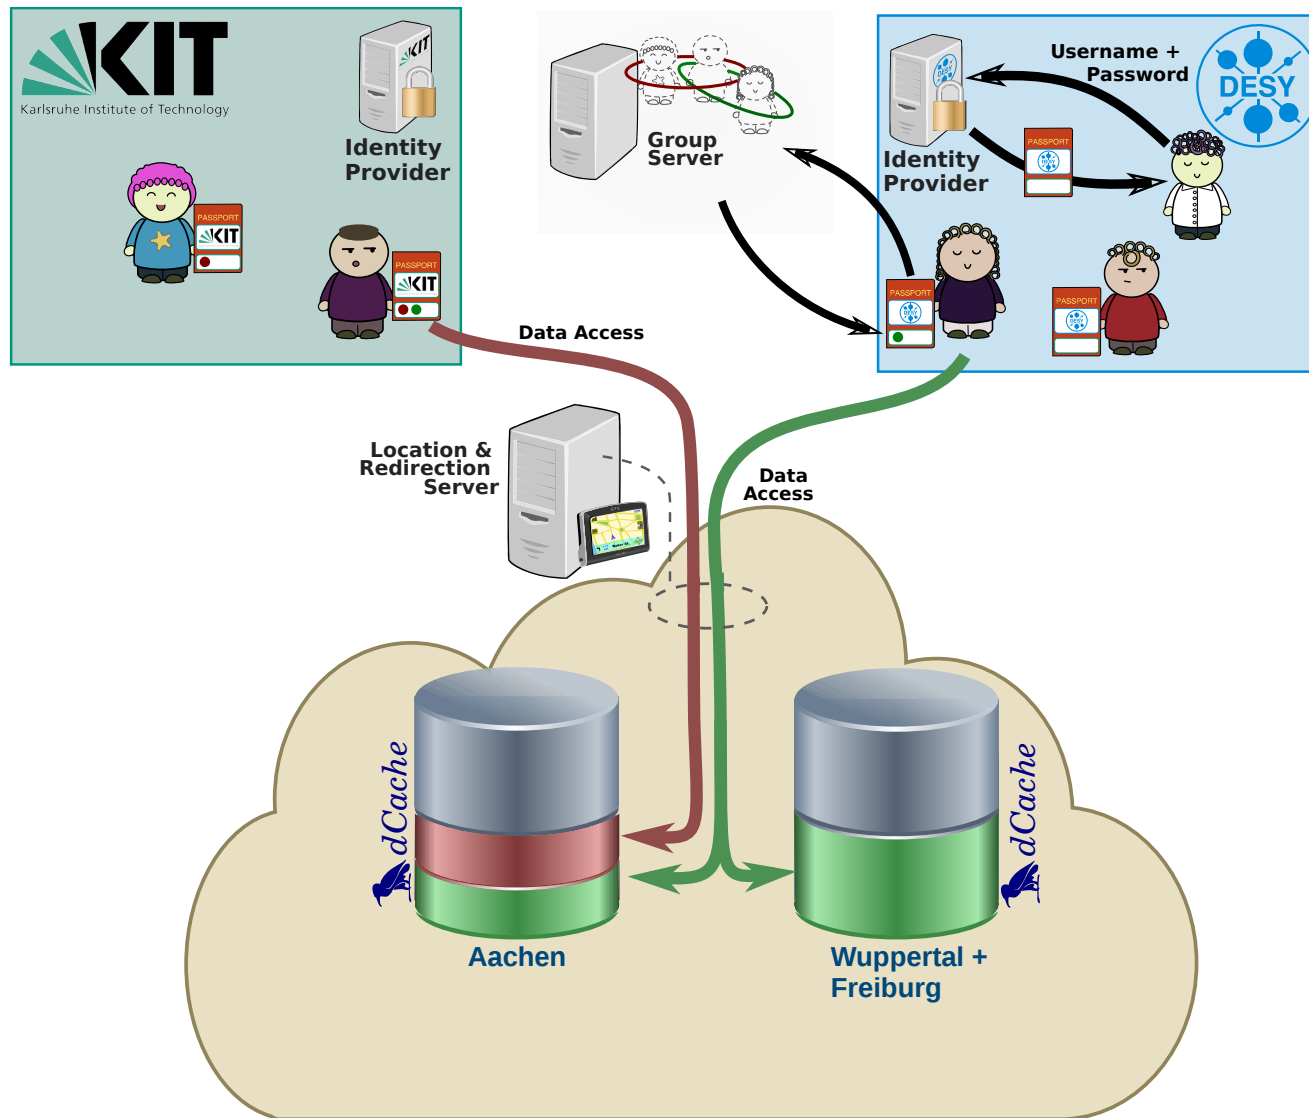
NFSv4.1 (4884 recs.)



Fast data ingest

- Current x-ray cameras write **single-frame files**:
 - Writing small(-ish) files at 20 Hz – 30 Hz, but rates will likely rise to 100 Hz or higher.
 - dCache was designed to accept larger files written less often:
 - Places less stress on the file-system.
 - Working on dCache to **improve write speed**.
 - Work mostly focused on NFS, but most improvements benefit all protocols.
-

Moving away from X.509



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.
-

Working with industry

- Working with **HDD manufactures:**

Hard disks with built-in ARM™ processor and Ethernet adaptor.

- Working with commercial **cloud provider:**

Integrating dCache into their software to provide large-scale storage.

- Working with **appliance provider:**

They want to have boxes with dCache support; or boxes with dCache pre-installed.

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: non-grid 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 platform for the *Human Brain Project* (HBP).

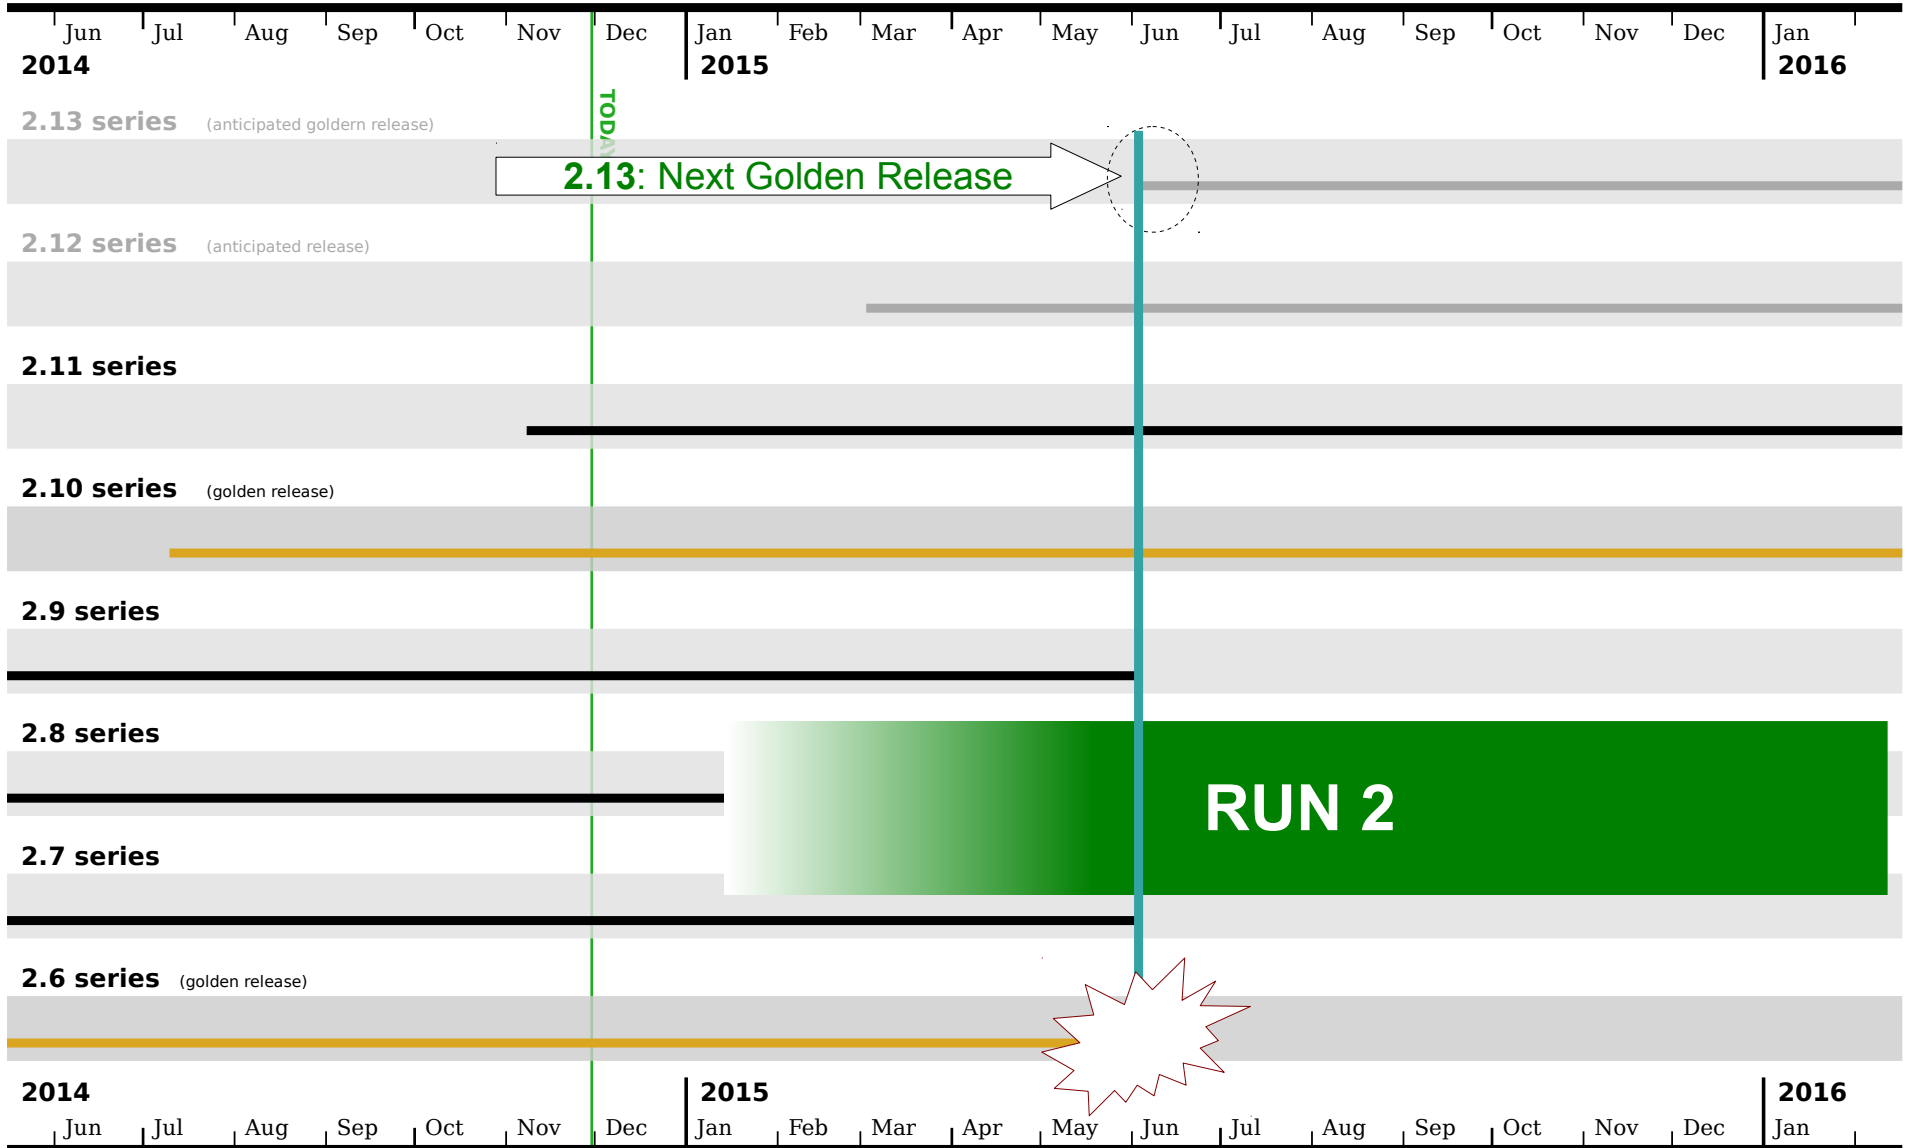
- **Commercial**:
(see earlier slide)

8th International dCache workshop



dCache server releases

... along with the series support durations.



Summary

- Future of dCache is secure,
 - Considerable activity, aimed at bringing real benefits to end-users,
 - Make sure your dCache site is upgrading to 2.10!
-

Communication

- **EGI:**
Patrick in TCB; Paul leads FedCloud AAI, participating in EGI AAI pilot project; 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
-

Thanks for listening ... any questions?

Backup slides



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).
-

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?
-

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.
-

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
