



WLCG Workshop 2017 Report

Data Management Session

July 2017 GDB

Oliver Keeble, Andrea Manzi

12/07/2017



Themes

- Data Steering Group
 - Overview of medium term priority topics
- Storage Resource Reporting
- Object Stores
- SKA Data Challenges
- Infrastructure architectures
 - Caches, Regional Federations, Multi-sites storages and Disk-less sites

Data Steering Group

- 7 strategic areas identified
 - Infrastructure diversity
 - Common global strategy
 - Quality of Service
 - Infrastructure Architecture
 - Operations
 - Analysis
 - AAI, Federated Identity
- <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGDataSteeringGroup>

Data Steering Group: Priorities

- Infrastructure diversity
 - In particular, object stores (Atlas, CMS, xrootd)
- Archival storage implementations (disk/tape) and their optimal exploitation (Atlas, Alice, Castor/EOS)
- Caching storage (Atlas, CMS)
- post-X509 AAI, Federated identity (longer term)

Storage Resource Reporting

- Further iteration on proposal
 - Intended to enable storage accounting and experiment operations in the post-SRM era
- R1: provide used/free for spaces via a non-SRM protocol
- R2: provide a structured data file
- ~~• R3: provide subdir resource reporting~~
- R4: provide storage dumps in a rationalised fashion (common to all experiments)


Object stores

- Distinguishing features of object stores
- Ceph
- Use in commercial cloud
- Replication, fault tolerance, erasure codes
- RAL experience (*Echo*)
 - Xrootd and gridftp plugins for Ceph
 - Issues with Xrootd direct I/O and requests for metadata



SKA Data Challenges

SKA: A Leading Big Data Challenge
for 2020 decade



Infrastructure Architectures [1/2]

Caches in OSG

- Caches are (supposedly) easier than traditional storage (data loss vs eviction)
- Experience with Xrootd Cache proxy (Stash Cache)
- Integration with CVMFS - metadata from CVMFS (extended with X509 authz), data from Xrootd cache

Regional federations

- Canadian project involving federation of cloud/dav endpoints using Dynafed
- Info on integration into Rucio and deployment of the infrastructure

Infrastructure Architectures [2/2]


Multisite storage

- Distributed dCache as used in NDGF, midwest T2 and the Russian demonstrator
- Overview of benefits and challenges of this approach

Diskless sites

- Report on consolidation of a group of 4 Romanian T2s who moved all storage to one of the sites. (each site using DPM)
- Intended to reduce costs
- Two sites concentrated on event-generation and MC and moved their storage to their partner

Questions?



References:

<https://indico.cern.ch/event/609911/sessions/238310/#20170620>