

Rucio session summary

ATLAS T* Jamboree

Martin Barisits, CERN

Rucio Session

- 3 presentations
 - Rucio Introduction, News & Status – Vincent
 - Rucio & DMLite – Mario
 - Rucio storage migration plan – Cedric

- Questions/comments welcome
 - ph-adp-ddm-lab@cern.ch

Rucio

- Rucio – next version of ATLAS DDM system
- Target deployment: early 2014 during LS1
 - Preliminary changes in DQ2 to facilitate the final migration
- Concepts see conceptual model document
- Collecting use cases for cross-checking and performance evaluation
 - Site use cases?
- Storage: Life without SRM?
 - WebDAV, xroot supported by 85% of sites (soon ~99%, e.g. storm)
 - SRM Space tokens? Can we get rid of them?
 - Functionality (like ACLs) are covered by scopes in Rucio
 - Accounting: Space usage published every 30 min in a local file
 - Publication of dumps (file index) for accounting & consistency

Rucio & DMLite

- Historically there are two ways to access DDM: Clients & Python API
- Rucio plugin for DMLite offers additional ways to access Rucio data
 - It maps Rucio namespace to a POSIX-compliant file system
 - Redirects replica access with appropriate protocols and paths
- Access:
 - WebDAV frontend
 - Mount like a regular file system (Linux, OSX)
 - HTTP client (curl, webbrowser, ...)
- Plugin communicates with Rucio server, files are accessed per WebDAV
- Next steps:
 - Functional testing done, now scaling tests
 - Possibility to create and retrieve data

Rucio storage migration plan 1/2

- ❑ New naming convention for files on storage
- ❑ Deterministic function → No need to rely on external catalogs (LFC)
 - ❑ End of use LCG file catalog in 2014
 - ❑ `f('scope:filename')` → `'basepath/scope/L1/L2/.../LN/filename'`
 - ❑ L1 ... LN: are the n^{th} 2-hexadecimal-digits from the hash value of the LFN (not the file content)
 - ❑ Assuming 32k files per directory, N defines a maximum number of files per scope
- ❑ Comparison between N=2 and N=3 on current namespace
 - ❑ N=2 → max ~2B files per scope; average of 2800 files per directory
 - ❑ N=3 → max~500B files per scope; average of 12 files per directory
- ❑ **Proposal: N=2; Even with order of magnitude more files (10 year scale) it is fine**

Rucio storage migration plan 2/2

- Migration strategy: Renaming of files → Works transparently
- Sites need to provide protocols which allow renaming (e.g. WebDAV or xroot)
- Renaming infrastructure in place
 - Tested on a few sites
 - Large scale tests coming
- Procedure for tape needs to be discussed
- **Schedule:**
 - January 2013: New files in DQ2 will be stored with the new naming convention (But still registered to LFC)
 - Before end of April: Migration of all old files to new convention on a couple of sites
 - May to end of year: Bulk migration