

Distributed OCDB using CVMFS

Raffaele Grosso

June 26, 2014



OCDBs on CVMFS are based on two things:

- 1 a synchronized copy of the AliEn OCDBs in CVMFS;

```

Ideal/
MC/ Residual/
Full/
/cvmfs/alice-ocdb.cern.ch/calibration/
data/ 2009/OCDB/
...
2013/OCDB/
```

- 2 a cvmfs package containing the list of OCDB files at a given time (**full OCDB snapshot**), used by the CDB framework when querying for an object.

```

Ideal.list.gz
MC/ Residual.list.gz
Full.list.gz
2009.list.gz
...
2013.list.gz
/cvmfs/alice.cern.ch/x86_64.../Packages/OCDB/v5-xx-Rev-yy/ data/
bin/ getUriFromYear.sh
getOCDBFilesPerRun.sh
```

$\underbrace{\hspace{15em}}_{\$OCDB_PATH}$



Implementation:

- 1 Synchronization between CVMFS OCDBs and AliEn OCDBs is based on the framework having OCDB uploads followed by publishing to CVMFS OCDBs (code in master: 675228755233), via a script to be available on authorized machines.
- 2 In the cvmfs OCDB case the list of valid files is uploaded from the full OCDB snapshot.



- When the user (production) uploads an AliRoot package, it can choose which OCDB package to upload with it.
- Hiding all OCDB objects produced after a certain moment in time is one of the main advantages which CVMFS OCDBs offer.
- The main difference w.r.t. snapshot files currently used by productions is a single text file global snapshot, as opposed to per-run snapshots.
- It also is a critical point. It requires clear answers to the following questions:
 - 1 who is the main **user**?
 - 2 **when** should the full OCDB snapshot be produced?
 - 3 how flexible should the schema be w.r.t. **retrofitting** objects?

Possible answers to these questions:

- 1 Main user: **MC productions**
- 2 OCDB objects are produced and uploaded:
 - automatically at the end of run by the shuttle
 - later, manually, as result of CPasses
 - still later, manually, to fix specific issues or make objects consistent with changes in the code

The OCDB snapshot (OCDB objects list and CVMFS package containing it) should be bound to a (MC) production. It should be produced **shortly before the production** is launched.

- 3 The working flow is quite **open to retrofitting** of CDB objects
 - AliRoot package not bound to specific OCDB snapshot
 - possibility to manually edit the OCDB objects list before packaging it

