

Distributed OCDB using CVMFS

Raffaele Grosso

March 20, 2014



- A synchronized¹ copy of the AliEn OCDBs is maintained on CVMFS:

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

- It can offer two major advantages:
 - 1 distributed OCDBs (fast retrieval of CDB objects)
 - 2 full OCDB snapshots (CDB object versions fixed for all runs and calibration types)

¹Full synchronization is not yet implemented. Plan: after a preliminar synchronization every AliEn OCDB update will be done via a proxy, which will also signal the new CDB objects to CVMFS.



- CDB files on CVMFS are accessed as in a local filesystem. To run the productions' macros (`sim.C/rec.C`) transparently, the CDB framework silently converts "`alien://`" storages to "`local://`" ones and changes the URLs appropriately.
- This mode is triggered by the `$OCDB_PATH` environment variable which is set to the base path of the AliRoot CVMFS package, once entered in the package.
- Using the usual AliEn URLs for OCDBs, user and productions on CVMFS will access CVMFS distributed OCDB files transparently, provided an AliRoot package is entered, with the corresponding OCDB package dependency:

```
/cvmfs/alice.cern.ch/bin/alienv enter
```

```
VO_ALICE@AliRoot::v5-05-76-AN,VO_ALICE@OCDB::v5-05-76-AN
```



OCDBs on CVMFS

```
rgrosso@Lxplus0150:~$ /cvmfs/alice.cern.ch/bin/alienv enter VO_ALICE@AliRoot::v5-05-76-AN,VO_ALICE@OCDB::v5-05-76-AN
[AliRoot/v5-05-76-AN OCDB/v5-05-76-AN] ~ -> aliroot
*****
*                                     *
*           W E L C O M E  t o  R O O T           *
*                                     *
*   Version   5.34/08           31 May 2013   *
*                                     *
*   You are welcome to visit our Web site   *
*           http://root.cern.ch             *
*                                     *
*****

ROOT 5.34/08 (v5-34-08@v5-34-08, Feb 17 2014, 18:00:58 on linuxx8664gcc)

CINT/ROOT C/C++ Interpreter version 5.18.00, July 2, 2010
Type ? for help. Commands must be C++ statements.
Enclose multiple statements between { }.

WELCOME to ALICE

root [0] gSystem->Getenv("OCDB_PATH")
(const char* 0x7fff94e939f6)/cvmfs/alice.cern.ch/x86_64-2.6-gnu-4.1.2/Packages/OCDB/v5-05-76-AN"
root [1] AliCDBManager *cdb = AliCDBManager::Instance()
I-AliCDBManager::Init: AliEn classes enabled in Root. AliCDBGrid factory registered.
root [2] cdb->SetDefaultStorage("alien://folder=/alice/data/2012/OCDB")
root [3] cdb->SetRun(170000)
I-AliCDBLocal::QueryCDB: Querying files valid for run 170000 and path "*" into CDB storage "local:///cvmfs/alice.cern.ch/calibration/data/2012/OCDB"
I-AliCDBLocal::QueryCDB: 249 valid files found!
```



ALICE

CVMFS AliRoot packages contain also the **lists of all the OCDB files** present at the moment when the package was created and two executables:

- 1 to **extract the year from the run number** and
- 2 to **extract the list of files valid for the given run** from those lists.

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

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

- In this way the version of all CDB objects for **all runs** is fixed, once the AliRoot tag has been chosen. AliRoot packages on CVMFS come each with a “**full OCDB snapshot**”, providing “schema evolution” in addition to the advantages already provided by the “per-run OCDB snapshot” file.
- The full set of CDB objects for a reconstruction job (~ 200 CDB objects) is retrieved in $\sim \frac{1}{2}$ minute.



- The list OCDBs' files included in the CVMFS AliRoot package is automatically build when tagging the package.
- The changes in the CDB framework were mainly in two directions:
 - changes in the creation of the storage from `AliCDBManager::GetStorage` to convert `alien://` to `local://` URLs using `$OCDB_PATH` and the executables to create the list of valid files;
 - changes in the local storage (`AliCDBLocal`, some in `AliCDBManager`) in order to let it use the list of valid file ids for the given run, as it was already the case for the AliEn storage. inconsistency apparent.

