

Rucio-DIRAC Integration at Belle II

3rd Rucio community Workshop

C. Serfon (BNL)





Introduction

- Belle II uses Dirac for its workload + Data Management
- For reason explained in talk yesterday Belle II plans to move to Rucio
- In order to be able to use Rucio, one needs a RucioFileCatalogClient plugin in Dirac that would replace the LCGFileCatalogClient currently used by Belle II
- In the following slides, I'll focus on the prototype of a RucioFileCatalogClient (RFC) developed at BNL for Belle II



RucioFileCatalogClient overview

- Belle uses the LFC (LCG File Catalog) as replica catalog and not the Dirac File Catalog (DFC) :
 - i.e. only the file replicas catalog functionalities are used, not the metadata ones that are handled by an external service (AMGA)
 - For this reason, the prototype developed doesn't include those functionalities that are only found in the DFC like the query by metadata
- The plugin preserves the hierarchical behaviour of the LFC :
 - Files in LFC → Files in Rucio
 - Directories containing files in LFC → Datasets in Rucio
 - Directories containing directories in LFC → Containers in Rucio
- The plugin has been made as experiment agnostic as possible, even though the scope determination is experiment dependant



Complications

- The concept of scope is unknown in the LFC, therefore it needs :
 - Either be part of the LFN (e.g. /grid/belle/MC)
 - Or there must be any deterministic function that can be used to extract it from the LFN

→ This part has to be implemented independently for each collaboration

- Mixture of files and directories in the same directory cannot be handled, e.g. :

```
lfc-ls -l /grid/belle/MC/fab/
drwxrwxr-x 1 292      147          0 Oct 13 2015 merge
-rw-rw-r-- 1 292      147          45240230 Jan 10 2017 output_15634.root
drwxrwxr-x 2 353      147          0 Jun 23 2017 prerelease-00-09-00b
drwxrwxr-x 1 353      147          0 Dec 12 2017 prerelease-01-00-00b
```

(Not a problem for Belle II since these are test files that can be cleaned)



Complications

- The plugin uses the Rucio client API that therefore needs to be installed on the Dirac server :
 - Possibility to use the Rucio REST API would need further development (no plan to be provided by BNL)
- Some methods cannot have the same behaviour as the LFC (more details in the next slides)



Read methods

- Methods can easily be mapped to Rucio methods
- Methods implemented :
 - getReplicas
 - listDirectory
 - getFileMetadata (only returns LFC metadata : checksums, GUID, filesize)
 - exists
 - getFileSize
 - isDirectory
 - isFile
- Behaviour is the same as in the LCGFileCatalogClient plugin



Write methods

- Methods implemented :
 - addFile :
 - This method reconstructs a hierarchy similar to the LFC (e.g. when one create a new file `/belle/ddm/functional_tests/releaseXXXX/test.dst`, it will create the dataset `/belle/ddm/functional_tests/releaseXXXX` if it doesn't exist as well as all the parents)
 - The replica of the file is added on the specified SE
 - A rule is set on the dataset as soon as it's created (e.g. rule is created on `/belle/ddm/functional_tests/releaseXXXX` with RSE expression ANY and grouping NONE)
 - addReplica :
 - Add the replica + create a rule on the file
- Discussion needed to know if the rules must have finite lifetime or not



Delete methods

- Delete methods are the most complicated ones because of the different concept in the LFC/Dirac and Rucio, e.g. :
 - removeReplica : In Rucio, not possible if the file belongs to a dataset and there is a rule on this dataset. In the RFC wouldn't do anything
 - removeFile : In LFC only removes the file if it has no replicas, otherwise trigger exception. For Rucio, no check if other replicas exist : Remove file just correspond to detach a file from a dataset



Next steps

- The prototype is in the test phase at BNL.
 - Already able to submit jobs that upload the files to Rucio using `dirac-dms-add-file`
 - Need to be validated with Belle II production system
- Once extensively validated, will submit a PR to Dirac repository
- If anyone else is interested in testing, just let me know (already got some contact with ESCAPE)



Test job

```
echo $!
cp RucioFileCatalogClient.py $DIRAC/DIRAC/Resources/Catalog/
mkdir etc
cp rucio.cfg etc
cp ca.crt etc
export PYTHONPATH=/cvmfs/atlas.cern.ch/repo/ATLASLocalRootBase/x86_64/rucio-clients/1.20.8/lib/python2.7/site-packages:$PYTHONPATH
export RUCIO_HOME=$PWD
export RUCIO_ACCOUNT=root
export RUCIO_AUTH_TYPE=x509_proxy
/bin/dd if=/dev/urandom of=test.txt.$! bs=1000000 count=1
dirac-dms-add-file -ddd /belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.$! ./test.txt.$! BNL-TEST1-SE
```

```
dirac@bldiracvm08 ~]$ rucio list-rules test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9
ID ACCOUNT SCOPE:NAME STATE[OK/REPL/STU]
-----
ba0945ab3aae48a09a79bcb23ec1d4c9 root test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9 OK[10/0/0]
```

```
dirac@bldiracvm08 ~]$ rucio list-files test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9
```

SCOPE:NAME	GUID	ADLER32	FILESIZE	EVENTS
test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.1	(None)	ad:9d5fa98f	1.000 MB	
test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.11	(None)	ad:1feed17f	1.000 MB	
test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.13	(None)	ad:cf42010a	1.000 MB	
test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.15	(None)	ad:b1c5c3fb	1.000 MB	
test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.17	(None)	ad:8341c40e	1.000 MB	
test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.19	(None)	ad:cc938327	1.000 MB	
test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.3	(None)	ad:a5ef1e93	1.000 MB	
test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.5	(None)	ad:172866e5	1.000 MB	
test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.7	(None)	ad:372467e5	1.000 MB	
test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9/test.txt.9	(None)	ad:6d83f0ea	1.000 MB	

```
Total files : 10
Total size : 10.000 MB
dirac@bldiracvm08 ~]$ rucio list-rules test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9
ID ACCOUNT SCOPE:NAME STATE[OK/REPL/STU]
CK] RSE_EXPRESSION COPIES EXPIRES (UTC) CREATED (UTC)
-----
ba0945ab3aae48a09a79bcb23ec1d4c9 root test:/belle/ddm/functional_tests/releaseXXXX/DB000001/SCOPE1/prod00001/e1/4S/r00000/test/0000001/udst/test_cedic9 OK[10/0/0]
ANY 1 2020-03-03 13:50:05
```



Conclusion

- RFC prototype available and being tested
- Prototype is as experiment agnostic as possible
- More details/results will be shown during the next Dirac workshop