

Report on DIRAC usage for SuperB use cases

Giacinto Donvito
INFN-BARI

Outline

- Testbed description
- Simple test jobs
 - Features exploited
- Analysis use case
 - Test results
 - What is still missing
- Data Management test
 - Catalogue and other info
 - Put/Get/delete features
 - Replica
 - High level features
 - Interactive usage
 - Metadata
 - Ancestors
- Future works
- People involved

Testbed description

- 1 “production” Dirac server installed @CNAF with DIRAC v5r13p2 (stable release)
- 1 “test” Dirac server installed @INFN-BARI with DIRAC v6rxx (still under development)
- Production Grid sites configured on DIRAC server
 - We configured a sub-set of sites officially supporting SuperB activities
 - There are sites from gLite and OSG grid
 - At the moment we only tested job submission on gLite sites
 - Both Computing Element and Storage Element of each site should be configured
- Thanks to interaction with DIRAC developers we find the right server configuration needed to exploit all the required features

Simple job test

- The job is described through a JDL file
- Main features tested:
 - InputSandbox (using Dirac File Catalogue)
 - OutputSandbox (using Dirac File Catalogue)
 - InputData (using Dirac File Catalogue)
 - OutputData (using Dirac File Catalogue)
 - Parametric Jobs (with and without Parametric InputData)
- All tests are executed mainly using CLI interface for submission, and both CLI and Web Interface for monitoring the status

Simple job test: details

- The Job submitted to DIRAC, is executed by mean of a wrapper script that takes care of:
 - Preparing the environment before job execution
 - Staging the output after job execution
 - Checking and reporting the exit code of each step
- The executable of the job is registered in the Dirac File Catalogue and copied by the user on a SuperB Storage Element before job submission
 - As the Input Sandbox too
- The Output Sandbox and the Output Data of the job is automatically registered in the catalogue by the wrapper
 - It is possible to choose in the JDL the SE where to store the output:
 - You can use more than one SE: if the first fail the second will be used, etc

Simple job test: details (2)

- Parametric Jobs:
 - A single DIRAC JDL can describe several (10-100-1000...) jobs
 - The end user submit only one JDL and all the jobs are automatically created by the system and should be monitored separately
 - Parameters could be:
 - Numbers (e.g. from 1 to 100)
 - List of strings (e.g. "aa ab ac ad ...")
 - The Input and the output files could also be parametric in order to deal with different Input/Output files per each job
- InputData:
 - It is possible to send the job on the computing elements close to the SE where the files are stored, using one or few files registered on the Dirac File Catalogue
 - This is valid for each sub-jobs in a parametric job if the files are parametric
- The wrapper of the job check the status of each step for each sub-job
 - If a sub-job fail, only that job will be automatically resubmitted

Analysis use case

- Executable used:
 - runPhrReduce (Thanks to Elisa Manoni)
- Input Files:
 - 1260 files
 - ~190GB in total
- Sandbox
 - Both input and output sandbox are registered on the Dirac File catalogue
- Output file:
 - 3 files for each job (1 root + 2 txt)
- Grid Site used:
 - CNAF, INFN-Pisa
- Sub-jobs
 - 126 jobs
 - Each job will process 10 input files

Analysis use case (2)

- It is quite inefficient to deal with very small input files:
 - So we packed 10 input file together and stored in the catalogue (~1.6GB each file)
 - Each job processed only one “packed file”
 - The name of each input file contains a string that is “parametric” (e.g. “aa ab ac ...”)
- Each file is replicate using DIRAC File Catalogue in both sites
 - Each job could be scheduled on both sites
- The output files are automatically stored in a directory in the catalogue that is specific for each job
- The executable is transferred as InputSandbox while we exploited the SuperB software installation already available in each site

Analysis use case – Results

- The analysis is carried out using a script that takes care of:
 - Untarring input file
 - Setting up the needed environment variables
 - Running the real executable with the needed parameters
 - Tar the output files
 - No need to interact with grid service!
- 126 job correctly executed in less than 30 min
 - From submitting to having the results back in the home directory
- The jobs were correctly distributed between CNAF (~60%) and INFN-PISA (~40%)
- No jobs failures from the user point of view

Analysis use case – Still missing

- To automatize the “input package builder”
 - Packing the small input files
 - Uploading to the grid and registering in the catalogue
 - Automatically writing the JDL
 - Automatically writing the user script
 - Retrieving all the output and merging together
-
- Typically those operation could be covered by GANGA or similar tool
 - We need to explore this option before writing new code!

Catalogue

- All the test has been carried on with DIRAC File Catalogue
 - This will be interesting because we will have both File Catalogue features and Metadata feature
- Using the Storage Element already configured in DIRAC instance
 - Configuring the SE is not at all a trivial task!!
 - We succeeded to use both OSG and gLite SE
- The typical grid file catalogue features are available
 - copy-and-register, replicate, delete, get, etc...

Catalogue – 2

- The DIRAC File Catalogue has a good number of added features:
 - Checksum
 - CreationDate
 - ModificationDate
 - Owner
 - OwnerGroup
 - Size
- DIRAC provides both “batch” and “interactive” command line
 - Web interface still to be tested
- The system is able to provide also an index of the Storage Element configured
 - They can be used with user-friendly names

Catalogue – 3

- It is possible to have both catalogue and storage metadata:
`dirac-dms-lfn-metadata /superbvo.org/user/g/donvito/test3`
{'Failed': {}},
'Successful': {'/superbvo.org/user/g/donvito/test3': {'Checksum': 'ao683fo4',
'ChecksumType': 'Adler32',
'CreationDate': datetime.datetime(2011, 6, 17, 9, 24, 48),
'FileID': 14L,
'GID': 1,
'GUID': '8531C86D-BD55-A983-3F09-124523587E1F',
'Mode': 775,
'ModificationDate': datetime.datetime(2011, 6, 17, 9, 24, 48),
'Owner': 'gdonvito',
'OwnerGroup': 'user',
'Size': 967L,
'Status': 1,
'UID': 2}}}}

Put/Get/delete features

- PUT:

- `dirac-dms-add-file /superbvo.org/user/g/donvito/dir_test/test2 /lustre/donvito/dirac/bashrc BARI-INFN`

- The storage element name should be one of those configured in DIRAC

```
{'Failed': {},
```

```
'Successful': {'/superbvo.org/user/g/donvito/dir_test/test2': {'put': 7.4923808574676514,  
                                                                'register': 0.52035999298095703}}}
```

```
dirac-dms-lfn-replicas /superbvo.org/user/g/donvito/dir_test/test2
```

```
{'Failed': {},
```

```
'Successful': {'/superbvo.org/user/g/donvito/dir_test/test2': {'BARI-INFN': 'srm://storm-  
se-01.ba.infn.it:8444/srm/manager?SFN=/superbvo.org/superbvo.org/user/g/donvito/dir_test/  
test2'}}}
```


Put/Get/delete features

- GET:
 - `dirac-dms-get-file /superbvo.org/user/g/donvito/dir_test/test2`
 - The Storage Element is automatically chosen
- DELETE:
 - `dirac-dms-remove-lfn`
 - Remove LFN and *all* associated replicas from Storage Elements and File Catalogs.
 - `dirac-dms-remove-lfn-replica /superbvo.org/user/g/donvito/dir_test/test2 FERRARA-INFN`
 - Remove replica of LFN from specified Storage Element and File catalogs.

REPLICA features

- Replica:

- Quite an easy task for the end user:

```
-bash-3.2$ dirac-dms-lfn-replicas /superbvo.org/user/g/donvito/dir_test/test2
```

```
{'Failed': {},
```

```
'Successful': {'/superbvo.org/user/g/donvito/dir_test/test2': {'BARI-INFN': 'srm://storm-  
se-01.ba.infn.it:8444/srm/manager?SFN=/superbvo.org/superbvo.org/user/g/donvito/dir_test/  
test2'}}}
```

```
-bash-3.2$
```

```
-bash-3.2$ dirac-dms-replicate-lfn /superbvo.org/user/g/donvito/dir_test/test2 FERRARA-INFN
```

```
{'Failed': {},
```

```
'Successful': {'/superbvo.org/user/g/donvito/dir_test/test2': {'register': 0.5209200382232666,
```

```
-bash-3.2$ 'replicate': 15.390527009963989}}}
```

```
-bash-3.2$ dirac-dms-lfn-replicas /superbvo.org/user/g/donvito/dir_test/test2
```

```
{'Failed': {},
```

```
'Successful': {'/superbvo.org/user/g/donvito/dir_test/test2': {'BARI-INFN': 'srm://storm-  
se-01.ba.infn.it:8444/srm/manager?SFN=/superbvo.org/superbvo.org/user/g/donvito/dir_test/test2',
```

```
'FERRARA-INFN': 'srm://grid2.fe.infn.it:8444/srm/manager?  
SFN=///superbvo.org/user/g/donvito/dir_test/test2'}}}
```


High level features: interactive usage

- The interactive usage is much more intuitive (and fast) when it is needed a long “data-management session”

```
-bash-3.2$ dirac-dms-filecatalog-cli
```

```
Starting DIRAC FileCatalog client
```

```
File Catalog Client $Revision: 1.17 $Date:
```

```
FC:/> help
```

```
Documented commands (type help <topic>):
```

```
=====
```

```
add      chgrp    exit  guid  meta   replicas  rmreplica  
ancestor chmod     find  id    mkdir   replicate size  
ancestorset chown    get   lcd   pwd     rm        unregister  
cd        descendent group  ls    register rmdir    user
```

```
FC:/> cd /superbvo.org/user/g/donvito/dir_test/
```

```
FC:/superbvo.org/user/g/donvito/dir_test>ls -l
```

```
-r-----rwx o gdonvito user      967 2011-06-22 17:10:45 test2
```

```
-r-----rwx o gdonvito user      967 2011-06-21 09:39:34 test3
```


High level features: Metadata

- The DIRAC File Catalogue give also metadata features embedded on the same tool:

```
FC:/> meta set /superbvo.org/user/g/donvito/test3 Creator  
GiacintoDonvito
```

```
FC:/> meta set /superbvo.org/user/g/donvito/test3 Group GSuperB
```

```
FC:/> meta get /superbvo.org/user/g/donvito/test3  
Group : GSuperB  
Creator : GiacintoDonvito
```


High level features: Metadata – 2

- The DIRAC File Catalogue also provide metadata “search” features:

```
FC:/superbvo.org/user/g/donvito>meta set dir_test Owner string
```

```
/superbvo.org/user/g/donvito/dir_test Owner string
```

```
FC:/superbvo.org/user/g/donvito>meta get dir_test
```

```
!Owner : string
```

```
Meta1 : test_dir
```

```
!NewMetaInt : 0
```

```
FC:/superbvo.org/user/g/donvito>
```

```
FC:/superbvo.org/user/g/donvito>meta set dir_test NewMetaInt 3
```

```
/superbvo.org/user/g/donvito/dir_test NewMetaInt 3
```

```
FC:/superbvo.org/user/g/donvito>meta get dir_test
```

```
!Owner : string
```

```
Meta1 : test_dir
```

```
!NewMetaInt : 3
```

```
FC:/superbvo.org/user/g/donvito>meta set dir_test Owner Giacinto
```

```
/superbvo.org/user/g/donvito/dir_test Owner Giacinto
```


High level features: Metadata – 3

```
FC:/superbvo.org/user/g/donvito>find NewMetaInt>2  
Owner=Giacinto
```

```
Query: {'Owner': 'Giacinto', 'NewMetaInt': {'>': 2}}  
/superbvo.org/user/g/donvito/dir_test/test2  
/superbvo.org/user/g/donvito/dir_test/test3
```

```
FC:/superbvo.org/user/g/donvito>find NewMetaInt>2  
Owner=DGiacinto
```

```
Query: {'Owner': 'DGiacinto', 'NewMetaInt': {'>': 2}}
```

```
FC:/superbvo.org/user/g/donvito>find NewMetaInt>0  
Owner=DGiacinto
```

```
Query: {'Owner': 'DGiacinto', 'NewMetaInt': {'>': 0}}  
/superbvo.org/user/g/donvito/test_2_dir/test1  
/superbvo.org/user/g/donvito/test_2_dir/test2
```


High level features: Metadata – 4

- How to do it not interactively

```
-bash-3.2$ cat test_cli
```

```
find NewMetaInt>o Owner=DGiacinto  
exit
```

```
-bash-3.2$ dirac-dms-filecatalog-cli < test_cli
```

Starting DIRAC FileCatalog client

File Catalog Client \$Revision: 1.17 \$Date:

FC:/> Query: {'Owner': 'DGiacinto', 'NewMetaInt': {'>': o}}

/superbvo.org/user/g/donvito/test_2_dir/test1

/superbvo.org/user/g/donvito/test_2_dir/test2

High level features: Ancestors

- The DIRAC File Catalogue give also “File Provenance” features embedded on the same tool:

```
FC:/superbvo.org/user/g/donvito/test_2_dir>ancestorset test1 /superbvo.org/user/g/donvito/  
dir_test/test3
```

Added 1 ancestors to file /superbvo.org/user/g/donvito/test_2_dir/test1

```
FC:/superbvo.org/user/g/donvito/test_2_dir>ancestorset test2 test1
```

Added 1 ancestors to file /superbvo.org/user/g/donvito/test_2_dir/test2

```
FC:/superbvo.org/user/g/donvito/test_2_dir>ancestor test2 2
```

```
/superbvo.org/user/g/donvito/test_2_dir/test2
```

```
1    /superbvo.org/user/g/donvito/test_2_dir/test1
```

```
2      /superbvo.org/user/g/donvito/dir_test/test3
```

```
FC:/superbvo.org/user/g/donvito/test_2_dir>descendent /superbvo.org/user/g/donvito/dir_test/  
test3 2
```

```
/superbvo.org/user/g/donvito/dir_test/test3
```

```
1    /superbvo.org/user/g/donvito/test_2_dir/test1
```

```
2      /superbvo.org/user/g/donvito/test_2_dir/test2
```


Job Submission: feedback

- It is needed to learn the JDL language
- Powerful all-in-one system for managing DM and Job Submission over a grid environment
- Stage-in and stage-out features
 - With failover capabilities
- Efficient monitoring mechanism for each step of the job
 - The exit status of: stage input, stage output, application execution is checked
 - If something fails, the job is re-submitted until the maximum number of resubmission is reached
- Peek job capabilities
 - A “tail” of the output of the job every X minutes
- Data-driven and load balancing job brokering
- Well integrated with gLite grid infrastructure
- Documentation is the main weak point

DMS: feedback

- Interesting features integrated in the same tool
 - Grid File Catalogue
 - Replica Cli
 - Metadata & ancestor
- It is flexible enough to cover different use cases
- Documentation is the main weak point
- Wildcard (*) not usable for bulk operations

DIRAC Test planned

- Configuring and Testing all the SuperB sites
 - Testing OSG job submission
- Porting the Monte Carlo Production use case into DIRAC
- FTS Usage
- Scalability tests
 - Both DMS and job submission
- Adding a detailed monitoring feature
 - Building a python script that exploit the DIRAC monitoring class
- Performance evaluation
- DB replica for HA purposes
- Try to build a simple automatic procedure to replicate a large dataset of files and check performance and reliability

People involved

- People working on DIRAC test
 - Giacinto Donvito – INFN-Bari
 - Bruno Santeramo – INFN-Bari
 - Armando Fella – INFN-Pisa
- Thanks to:
 - Matteo Manzali – INFN-Ferrara