



A replication feature for the *AliEn submit* command

Alice Offline Week

Steffen Schreiner (CERN, PH/AIP)



Motivation

Currently, whenever I run a job in AliEN on the grid

I will receive the output file/archive on **one** Storage Element

Why can't I specify, to have a certain number of **copies / replicas** immediately?

Why can't I tell AliEn to use several certain Storage Elements ?

Let's just tell AliEn within the jobs JDL file, what we want...



Let AliEn do the replication right away

```
[pcepalice10.cern.ch:3307] /pcepalice10/user/a/ali/ > submit jobTest.jdl
Jun 24 03:04:12 info Submitting job '/pcepalice10/user/a/ali/bin/echo.bash --round20
02-02--run 00076--event'...
Jun 24 03:04:12 info There is no time to live (TTL) defined in the jdl... putting th
e default '6 hours'
Jun 24 03:04:12 info There is no price defined for this job in the jdl. Putting the
default '1.0'
Jun 24 03:04:12 info Input Box: {}
Jun 24 03:04:12 info Command submitted (job 491)!!
Job ID is 491 - 0
[pcepalice10.cern.ch:3307] /pcepalice10/user/a/ali/ > top
JobId Status Command name Submitthost
 491 WAITING /pcepalice10/user/a/ali/bin/echo.bash ali@pcepalice10
.cern.ch
[pcepalice10.cern.ch:3307] /pcepalice10/user/a/ali/ > top
JobId Status Command name Submitthost
 491 SAVING /pcepalice10/user/a/ali/bin/echo.bash ali@pcepalice10
.cern.ch
[pcepalice10.cern.ch:3307] /pcepalice10/user/a/ali/ >
```

```
[pcepalice10.cern.ch:3307] /proc/ali/491/job-output/ > whereis .alien_archive.491.PCEPA
LICE10::CERN::TESTSE.1245805519.zip
Jun 24 03:05:52 info The file 491/job-output/.alien_archive.491.PCEPALICE10::CERN::T
ESTSE.1245805519.zip is in
SE => pcepalice10::CERN::TESTSE fn =>file://pcepalice10.cern.ch/alie
ntmp/SE_DATA/14/18935/1638ba2e-605b-11e0-a144-001e0b24002f
SE => pcepalice10::CERN::TESTSE2 fn =>file://pcepalice10.cern.ch/tmp
/pcepalice10.org/SE2/DATA/14/18935/1638ba2e-605b-11e0-a144-001e0b24002f
[pcepalice10.cern.ch:3307] /proc/ali/491/job-output/ >
```



How to choose Storage Elements ?

If not specified otherwise, the job output should go to 'close' Storage Elements

'Close' might mean not to copy the output from France to India

Yet, geographical distance is not enough

Network bandwidth and response time don't follow geographical distance

Infrastructure and performance changes over time



How can a JobAgent know about 'close' SEs ?

Maybe the site you're running has available SEs

An **SE_Election agent** should generate a list based on

- Availability
- IP based decisioning - network organizational distance
- Response time / round trip time
- Country / organization



How can a JobAgent know about 'close' SEs ?

Maybe the site you're running has available SEs

JobAgent

An **SE_Election agent** should generate a list based on

- Availability
- IP based decisioning - network organizational distance
- Response time / round trip time
- Country / organization

MonALISA



Extending the AliEn *submit*

The AliEn JobAgent considers a JDL specification, how many copies are demanded and which Storage Elements to use

If the site on which the JobAgent is running has Storage Elements defined, it will use these first

Thereafter the JobAgent will request additional SEs from the new functionality in MonALISA



Extending the AliEn *submit*

Run your job, and the output goes to ...

- a JDL specified number of k copies while AliEn (and MonALISA) decides which SEs to use
- a JDL specified number of k copies on n specified SEs while $k \leq n$ and AliEn will use any SEs in the list
- a JDL specified number of k copies on n specified SEs while $k > n$ and AliEn will choose the remaining SEs and ask MonALISA if necessary



k replicas of the job output

AliEn should take care about the finding of k proper (close) SEs

Without any further action a user will have k copies of the job output

```
myJob.jdl
```

```
...
```

```
outputArchive = "... , copies=k "
```

```
....
```

```
outputFiles = "... , copies=k "
```

```
...
```



k replicas of the job output on certain SEs

A list of n SEs is specified in the JDL

The job output should be copied on k SEs, while using the specified list

```
myJob.jdl
...
outputArchive = "... @pcepalice::CERN::testSE,
                  pcepalice::CERN::testSE2,
                  copies=2 "
...
```



Release

- The AliEn JobAgent module will know how to process JDL user specifications up from the AliEn release, coming within the next weeks
- It will know how to handle information coming from MonALISA and account for it in order to store job results
- The SE_Election agent in MonALISA will start with an easy approach and is improved step by step without making further changes in AliEn necessary



Outlook (nothing planed so far)

The behavior could be extended to the commands

add Upload files directly on several SEs

mirror Mirror on several SEs in one step

Both features could use the SE_Election feature

An inverted 'close' criteria could deliver proper SEs for security backups on wide distance or across organizations