

Table of Contents

Manage Completed jobs.....	1
----------------------------	---

Manage Completed jobs

Completed jobs are a class of jobs which leave the input files in Processed status, that is they are not accounted as Failures (aka they don't add files in MaxReset), but is a symptom of potential issues in data management.

When the job finishes to process the input files, Dirac generates a so called Request to move the output files to the right place, register them to the FileCatalog and to the Bookkeeping. If, for some reasons, the Request fails, a new one is issued as a backup solution and the files are moved to a so called FAILOVER. It can happen that these requests get stuck and the production accounts the input files as Processed but with pending operations to be completed.

The simplest way to check if a job has problematic output files is to issue this command:

```
dirac-bookkeeping-job-input-output --Output <JobID> | dirac-bookkeeping-get-file-descendants
```

for jobs in Reco, Turbo or Stripping productions, which have descendants. Another useful command to issue is:

```
dirac-bookkeeping-job-input-output --Output <list of JobID> | grep -v log | grep -v root | dirac
```

This command takes the output files of jobs and tell if they are somewhere else than their right final destination. In principle one could replicate them by hand with:

```
dirac-dms-replicate-to-run-destination <lfn> --SE <Storage>
```

where <Storage> can be one among Tier1-DST, Tier1-BUFFER, Tier1-RDST, etc.

This manual replication must be issued carefully, because the system could be just lagging behind and a removal of a file could end up in a failing request.

-- MarcoCorvo - 2016-10-25

This topic: LHCb > CompletedJobs

Topic revision: r2 - 2016-10-25 - MarcoCorvo



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback