

# Table of Contents

<b>Closing procedure for productions.....</b>	<b>1</b>
Manage files in.....	1
Manage files in Assigned status.....	2
Manage files in Unused status.....	2
Closing a production.....	3
Particularly odd situation.....	3

# Closing procedure for productions

These steps should be done at the end of a production, that is when there are no more running jobs. There could be still some Completed jobs which btw must be treated separately.

- Check file status on Transformation Monitor tab in the web portal of Dirac
- Look for file statuses like MaxReset, Unused or Assigned (Removed or Problematic are set elsewhere)
- Try to figure out if some of these files are recoverable

These files can be checked using the command

```
dirac-transformation-debug <prodID> --Status <MaxReset|Assigned|Unused> --Info <files|jobs|flush>
```

## Manage files in

Many are the reasons why a file can be in MaxReset, namely:

- file is corrupted
- applications finish with errors
- applications are unable to download input

If a file is corrupted (e.g. when a code 16 is returned by DaVinci) there's not much we can do. Just keep the file in MaxReset to avoid picking it up in case of further reprocessing.

If the application finishes with errors (e.g. DaVinci exits with code 134) the problem could be only a glitch or a mismatch in , e.g., a conddb tag. In this case a reset can be tried with:

```
dirac-transformation-reset-files <prodID> --LFNs=<list of lfns>
```

In case of multiple input files (Merging or Turbo productions principally but also Reconstruction ones), it can happen that only one file is corrupted or problematic. The only thing to do is to isolate the affected file and reset the others. `dirac-transformation-debug` should tell which one is the file to isolate:

```
dirac-transformation-debug <prodID> --Status MaxReset --Info jobs
```

and you can expect a result like:

```
[LHCbDirac v8r3p10] ~ $ dirac-transformation-debug 53934 --Status MaxReset --Info jobs
Transformation 53934 (Idle) of type DataStripping (plugin ByRunWithFlush, GroupSize: 1) in Real
BKQuery: {'StartRun': 183890L, 'ConfigName': 'LHCb', 'EndRun': 184372L, 'EventType': 90000000L, '
3 files found with status ['MaxReset']

1 LFNs: ['/lhcb/LHCb/Collision16/RDST/00053882/0004/00053882_00044998_1.rdst'] : Status of corre
142091039 142110901 142125753 142143210 142163591 142176729
6 jobs terminated with status: Failed; Requests done; DaVinci Exited With Status 134

1 LFNs: ['/lhcb/LHCb/Collision16/RDST/00053882/0000/00053882_00000637_1.rdst'] : Status of corre
141358742 141410782 141446879 141478686 141530009 141567063
6 jobs terminated with status: Failed; Job stalled: pilot not running; DaVinci v41r2 step 1
1 jobs stalled with last line: (LCG.IN2P3.fr) EventSelector SUCCESS Reading Event record
1 jobs stalled with last line: (LCG.IN2P3.fr) EventSelector SUCCESS Reading Event record
1 jobs stalled with last line: (LCG.IN2P3.fr) EventSelector SUCCESS Reading Event record
1 jobs stalled with last line: (LCG.IN2P3.fr) EventSelector SUCCESS Reading Event record
1 jobs stalled with last line: (LCG.IN2P3.fr) EventSelector SUCCESS Reading Event record
```

## ClosingProcedure < LHCb < TWiki

```
1 jobs stalled with last line: (LCG.IN2P3.fr) EventSelector      SUCCESS Reading Event record
1 LFNS: ['/lhcb/LHCb/Collision16/RDST/00053882/0004/00053882_00041497_1.rdst'] : Status of corre
141924608 141937899 141963800 142001696 142022765 142069366
6 jobs terminated with status: Failed; Requests done; DaVinci Exited With Status 134

Summary of failures due to: Application Exited with non-zero status
ERROR ==> /lhcb/LHCb/Collision16/RDST/00053882/0004/00053882_00044998_1.rdst was Partial (last ev
ERROR ==> /lhcb/LHCb/Collision16/RDST/00053882/0004/00053882_00041497_1.rdst was Partial (last ev
```

If the situation is **not** pathological, e.g. the files keep staying in a non final status without evidence of a situation, try to check the following:

```
dirac-transformation-debug <prodID> --Status MaxReset --Info files | dirac-production-check-descendants
```

This command line provides a summary of files which are the descendants of the ones in MaxReset status, if any, and gives also some suggestions on how to cure the situation. `dirac-production-check-descendants` also produces a list of affected files named `CheckDescendantsResult_<ProdID>.txt`. This file can be used when fixing problems, e.g.:

```
grep InFailover CheckDescendantsResults_53197.txt | dirac-dms-replicate-to-run-destination --SE
```

This command selects those files which are in Failover, instead of their final destination, and replicates them to the right SE.

## Manage files in Assigned status

A file can remain stuck in Assigned status due to internal glitches of the system. Like for MaxReset files, the command to issue is:

```
dirac-transformation-debug <prodID> --Status Assigned --Info files | dirac-production-check-descendants
```

Again the output of this command gives you some hints on how to fix problems.

## Manage files in Unused status

A file can also remain Unused basically for two reasons:

- the files doesn't have an active replica
- the run to which the file belongs has not yet finished the online processing (basically HLT2)

To verify in which case the file falls one can try:

```
dirac-transformation-debug <prodID> --Status Unused --Info flush
```

which says if the run is still in the HLT2 processing phase or:

```
dirac-transformation-debug <prodID> --Status Unused --Info files | dirac-dms-lfn-replicas
```

which says if the file has a replica.

Another useful chain of command is:

```
dirac-transformation-debug <Prod2> --Status Unused --Info files | dirac-bookkeeping-get-file-ances
```

where Prod2 is, for example, a Merging production and Prod1 is a Stripping one. In this case it's possible to understand whether the files are still Unused because the ancestors file where in some particular status, e.g. MaxReset.

## Closing a production

If there are no more files in a bad shape that could be fixed, there is a final check on a given production that must be performed:

```
dirac-production-check-descendants <prodID>
```

For Reco, Turbo or Turcal production, the command can complete in few hours, but for Stripping it can take much longer. To solve this problem these are the steps to perform when closing a production with final Merging steps (Stripping or Turbo), given that <prod1> is the Stripping/Turbo and <prod2> the Merging:

```
dirac-bookkeeping-get-files Prod <prod1> Visibility No | dirac-production-check-descendants <pr
```

This would verify if the non-merged files still in the BK have been merged or not. This is not enough as there could be files in the FC but not in the BK which are not picked up by the previous command. To solve this there's another useful command (will be released soon, not yet available with the current LHCbDirac release v8r3p7):

```
dirac-loop TurboFileTypes.txt 'dirac-dms-list-directory /lhcb/LHCb/Collision16/@arg@/000<prod1> |
```

this currently works for Turbo and should be adapted for Stripping.

This command can be followed by a final

```
dirac-dms-check-fc2bkk Last
```

in case some files are not in the BK.

If everything's ok, the final steps are to go to the Dirac web portal and mark the production "Complete", then set it "Done" in the Production Requests tab. This commands will leave the production in a "suspended" status for a week, just to let the experts have the time to resume it or to fix other issues.

## Particularly odd situation

It can happen that a job gets killed or finishes in a very bad shape while performing the last operations on the output files, that is while moving them to their final destination. In this case the job is marked as Failed, but still some files could have been replicated. In principle the system should take care of issuing a Removal request for the output files that made it through, but this doesn't happen always.

Best thing to do in this case is to remove the output files and reset the input Unused.

First check files descendants:

```
[localhost] ~ $ dirac-bookkeeping-job-input-output --Output 142393519 | dirac-bookkeeping-get-fil
Got 8 LFNs
Getting descendants for 8 files (depth 1) : completed in 0.1 seconds
NotProcessed :
  /lhcb/LHCb/Collision16/CHARMCHARGED.MDST/00053884/0000/00053884_00003714_1.charmcharged.mdst
  /lhcb/LHCb/Collision16/CHARMKSHH.MDST/00053884/0000/00053884_00003714_1.charmkshh.mdst
  /lhcb/LHCb/Collision16/CHARMMULTIBODY.MDST/00053884/0000/00053884_00003714_1.charmmultibody.m
  /lhcb/LHCb/Collision16/CHARMSPECPARKED.MDST/00053884/0000/00053884_00003714_1.charmspecparked
  /lhcb/LHCb/Collision16/CHARMSPECPRESCALED.MDST/00053884/0000/00053884_00003714_1.charmspecpre
```

## ClosingProcedure < LHCb < TWiki

```
/lhcb/LHCb/Collision16/CHARMTWOBODY.MDST/00053884/0000/00053884_00003714_1.charmtwobody.mdst  
/lhcb/LHCb/Collision16/LEPTONS.MDST/00053884/0000/00053884_00003714_1.leptons.mdst  
/lhcb/LHCb/Collision16/LOG/00053884/0000/00003714/DaVinci_00053884_00003714_1.log
```

then check the status of the replicas:

```
[localhost] ~ $ dirac-dms-replica-stats --Last  
Got 8 LFNs  
Getting replicas for 8 LFNs : completed in 0.2 seconds  
6 files found without a replica  
2 files found with replicas  
  
Replica statistics:  
0 archives: 2 files  
0 replicas: 6 files  
1 replicas: 2 files  
  
SE statistics:  
  CNAF-BUFFER: 2 files  
  
Sites statistics:  
  LCG.CNAF.it: 2 files
```

Two out of eight files, in this example, were copied even if the system didn't set the replica flag correctly. Remove these two files:

```
dirac-dms-remove-files --LFNs=<list of lfns>
```

and reset the input Unused

```
[localhost] ~ $ dirac-transformation-reset-files <ProdID> --LFNs=<list of lfns>
```

-- MarcoCorvo - 2016-10-07

---

This topic: LHCb > ClosingProcedure

Topic revision: r6 - 2016-10-28 - MarcoCorvo



Copyright &© 2008-2022 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback