

Table of Contents

2018 dataflow transformations.....	1
Replication of RAW.....	1
FULL Stream.....	1
LOWMULT, NOBIAS.....	1
TURCAL.....	1
TURBO.....	1
2017 dataflow transformations.....	2
stripping28.....	2
RAWReplication.....	2
FULL (90000000).....	3
.Reco17.....	3
Turbo (94000000).....	4
TURCAL (95100000).....	4
NOBIAS (96000000), LOWMULT(90600000).....	5
2016 dataflow transformations.....	6
RAW files replication.....	6
RDST replication for FULL stream (90000000).....	6
Stripped DST replication from FULL stream (90000000).....	6
TURBO.MDST replication for TURBO stream (94000000).....	6
TURCAL reco output replication (FULLTURBO.DST) (95100000).....	6
Removing files from disk BUFFER.....	6
Removal of RAW from BUFFER.....	6
Removal of RDST from BUFFER for FULL stream.....	7

2018 dataflow transformations

The dataflow is exactly the same as for 2017. The transformation/production numbers are kept in a google spreadsheet

https://docs.google.com/spreadsheets/d/1egYsMWP5rDAzOoEWOAkzWkGV7jjEA0Iyw-WxRpWS9Pg/edit?usp=share_link

Everything is discussed in Jira Tasks (linked in the spreadsheet)

Replication of RAW

RAWReplication (1 tape + 1 disk) for FULL, LOWMULT, NOBIAS and TURCAL ReplicateToRunDest (1 disk) for TURBO

FULL Stream

RDST output replicated to local site (ReplicateToLocalSE RDST)

The RAW will be removed from disk after the Stripping (RemoveReplicasWithAncestor)

LOWMULT, NOBIAS

The output DST can be replicated (LHCbDSTBroadcast) and the input RAW files removed

TURCAL

The RAW goes through Reco and Turbo. The output can be replicated to local RDST (ReplicateToLocalSE)
The RAW file is removed from buffer after Reco/Turbo

TURBO

The output is replicated with LHCbDSTBroadcast, and the input RAW removed

2017 dataflow transformations

Jira task for 2017 data replication: <https://its.cern.ch/jira/browse/LHCBPO-116>

List of standing transformations: <https://its.cern.ch/jira/browse/LHCBDIRAC-650>

stripping28

Staging of full and RDST:

Some files are already on disk (because not removed):

183890:184372 (Beam6500GeV -VeloClosed-MagUp) : 57323 files

184592:184666 (Beam6500GeV -VeloClosed-MagUp) : 11919 files

184683:185611 (Beam6500GeV -VeloClosed-MagDown) : 123647 files

/LHCb/Collision16//RealData/Reco16//RDST

Plugin: ReplicateWithAncestors

DestSE = Tier1-Buffer

The EndRun should be the beginning of what is already staged (183163). The StartRun will move (automatically maybe with option TargetFilesAtDestination)

Transformation number: 59245 (182354:183163)

Removal when processed:

Transformations 59588 (to follow 59557),59589 (to follow 59559), 59590 (to follow 59561)

```
for 59557: dirac-dms-add-transformation --BK /LHCb/Collision16/Beam6500GeV-VeloClosed-MagUp/RealData/Reco16//RDST
```

```
for 59559: dirac-dms-add-transformation --BK /LHCb/Collision16/Beam6500GeV-VeloClosed-MagDown/RealData/Reco16//RDST
```

```
for 59561: dirac-dms-add-transformation --BK /LHCb/Collision16/Beam6500GeV-VeloClosed-MagDown/RealData/Reco16//RDST
```

Follow startRun-Endrun evolution

Output Replication:

Transformation 59591

```
dirac-dms-add-transformation --BK /LHCb/Collision16//RealData/Reco16/Stripping28//ALL.DST,ALL.MDST
```

RAWReplication

1 tape copy (+ CERN-RAW), 1 on disk 2 transformations: one per polarity for the FULL, TURCAL, NOBIAS and LOWMULT streams

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagUp/RealData/90000000,95100000,96000000,90600000/RAW

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagDown/RealData/90000000,95100000,96000000,90600000/RAW

Plugin: RAWReplication

For the Turbo , one per polarity, only one disk copy (+ CERN-RAW)

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagUp/RealData/94000000/RAW

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagDown/RealData/94000000/RAW

Plugin: ReplicateToRunDestination Destination: Tier1-Buffer

The run destination is set by RAWReplication

Replication with the 2016 data shares until run 194248: 63149 (FULL, TURCAL,NOBIAS,CEP) 63150 (TURBO)

FULL (90000000)

Reco17

No staging needed, since the data is fresh

Replicate the output:

/LHCb/Collision17//RealData/Reco17//RDST

Plugin: ReplicateToLocalSE

DestSE = Tier1-RDST DQFlags=OK,UNCHECKED

```
dirac-dms-add-transformation --Plugin ReplicateToLocalSE --DestinationSEs Tier1-RDST --BK /LHCb/
```

Replication of output

/LHCb/Collision17//RealData/Reco17/Stripping29//ALL.DST,ALL.MDST

LHCbDSTBroadcast

(there is no CALIBRATION.DST anymore)

```
dirac-dms-add-transformation --Plugin=LHCbDSTBroadcast --BK=/LHCb/Collision17//RealData/Reco17/S
```

Removal of the input

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagUp/RealData/Reco17//RDST

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagDown/RealData/Reco17//RDST

Plugin: RemoveReplicasWithAncestors

FromSE =Tier1-Buffer

RAWReplication

ProcessingPass =Stripping29

DQFlag=OK,UNCHECKED

Follow startRun-Endrun evolution

dirac-dms-add-transformation --Plugin=RemoveReplicasWithAncestors --BK=/LHCb/Collision17/Beam6500

Turbo (94000000)

Replicate the output:

/LHCb/Collision17//RealData/Turbo/94000000/TURBO.MDST

Plugin: LHCbDSTBroadcast DQFlags=OK,UNCHECKED

Removal when processed:

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagUp/RealData/94000000/RAW

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagDown/RealData/94000000/RAW

Plugin: RemoveReplicasWenProcessed

ProcessingPass = Turbo

FromSE = Tier1=Buffer

DQFlag=OK,UNCHECKED

Follow startRun-Endrun evolution

TURCAL (95100000)

output replication (FULLTURBO.DST)

/LHCb/Collision17//RealData/Reco17/Turbo04/95100000/FULLTURBO.DST

ReplicateToLocalSE, destination Tier1-RDST

* Removal when processed*

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagUp/RealData/95100000/RAW

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagDown/RealData/95100000/RAW

Plugin: RemoveReplicasWenProcessed

ProcessingPass = Reco17/Turbo04

FromSE = Tier1=Buffer

DQFlag=OK,UNCHECKED

Follow startRun-Endrun evolution

NOBIAS (96000000), LOWMULT(90600000)

Replicate the output:

/LHCb/Collision17//RealData/Reco17/96000000,90600000/ALL.MDST

Plugin: LHCbDSTBroadcast DQFlags=OK,UNCHECKED

```
dirac-dms-add-transformation --Plugin LHCbDSTBroadcast --BK /LHCb/Collision17//RealData/Reco17/9
```

Removal when processed:

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagUp/RealData/96000000,90600000/RAW

/LHCb/Collision17/Beam6500GeV-VeloClosed-MagDown/RealData/96000000,90600000/RAW

Plugin: RemoveReplicasWhenProcessed

ProcessingPass = Reco17

FromSE = Tier1=Buffer

Follow startRun-Endrun evolution

Because there is one production per Stream, we need one transformation per stream, otherwise it can't find the production

```
dirac-dms-add-transformation --Plugin RemoveReplicasWhenProcessed --FromSEs Tier1-BUFFER --BK /LH
```

```
dirac-dms-add-transformation --Plugin RemoveReplicasWhenProcessed --FromSEs Tier1-BUFFER --BK /LH
```

2016 dataflow transformations

This give a list of DMS transformations required for the 2016 data flow. Several streams have to be handled, whose event types are listed here:

- 90000000 : FULL stream
- 94000000 : TURBO stream
- 95100000 : TURCAL stream

RAW files replication

Two transformations exist: 51431 for MagUp and 51432 for MagDown

```
dirac-dms-add-transformation --plugin RAWReplication --BK /LHCb/Collision16//RealData/90000000,94
```

This replicates the RAW files to an external (to CERN) Tier1-RAW SE according to shares defined in the CS path /Operations/Default/Shares/RAW (arbitrary units). RAW files are also replicated to one of the Tier1-BUFFER SEs chose between CERN-BUFFER and that at the same site chosen in the Tier1-RAW.

This transformation sets the destination for a given run. All files from a run are replicated to the same SEs, and subsequently all files created during the processing of the data will be uploaded to that same site.

RDST replication for FULL stream (90000000)

Transformation number 51648

```
dirac-dms-add-transformation --Plugin ReplicateToLocalSE --DestinationSEs Tier1-RDST --BK /LHCb/C
```

Stripped DST replication from FULL stream (90000000)

Transformation number 52357, with 2 disk replicas only for the time being

```
dirac-dms-add-transformation --Plugin LHCbDSTBroadcast --BK /LHCb/Collision16//RealData/Reco16/St
```

TURBO.MDST replication for TURBO stream (94000000)

No transformation yet, to be updated if more than one file type or more than one processing pass

```
dirac-dms-add-transformation --Plugin LHCbDSTBroadcast --BK /LHCb/Collision16//RealData/Turbo02a/
```

TURCAL reco output replication (FULLTURBO.DST) (95100000)

Uploaded to Tier1-DST, no replication and no archive

Removing files from disk BUFFER

Removal of RAW from BUFFER

One transformation per magnet polarity and also change transformation when the run range changes

FULL stream: transformation number 52288 (MagUp), 52290 (MagDown). End run set for both of them as Stripping steps will change when fixed

```
dirac-dms-add-transformation --Plugin RemoveReplicasWhenProcessed --FromSEs Tier1-BUFFER --BK /LH
dirac-dms-add-transformation --Plugin RemoveReplicasWhenProcessed --FromSEs Tier1-BUFFER --BK /LH
```

For changing the end run:

```
dirac-production-set-runs <prod> --End <runNumber>
```

TURCAL stream: ttransformation number 52353 (MagUp), Archived for MagDown, to be relaunched when we get new data

```
dirac-dms-add-transformation --Plugin RemoveReplicasWhenProcessed --FromSEs Tier1-BUFFER --BK /LH
dirac-dms-add-transformation --Plugin RemoveReplicasWhenProcessed --FromSEs Tier1-BUFFER --BK /LH
```

TURBO stream: transformations not existing

```
dirac-dms-add-transformation --Plugin RemoveReplicasWhenProcessed --FromSEs Tier1-BUFFER --BK /LH
dirac-dms-add-transformation --Plugin RemoveReplicasWhenProcessed --FromSEs Tier1-BUFFER --BK /LH
```

Removal of RDST from BUFFER for FULL stream

Transformation number 52287 (MagDown), 52289 (MagUp). End run set, new transformations should be launched with new stripping

```
dirac-dms-add-transformation --Plugin RemoveReplicasWhenProcessed --FromSEs Tier1-BUFFER --BK /LH
dirac-dms-add-transformation --Plugin RemoveReplicasWhenProcessed --FromSEs Tier1-BUFFER --BK /LH
```

-- PhilippeCharpentier - 2016-07-07</verbatim>

This topic: LHCb > DataFlow2015

Topic revision: r23 - 2018-06-06 - ChristopheHaen



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