

Reprocessing Christmas 2017

2017 Run 2 reprocessing

<https://twiki.cern.ch/twiki/bin/view/Atlas/Winter2017Run2Reprocessing>

David South (DESY)
Kostas Nikolopoulos (Birmingham)

ADC Meeting
December 20th 2016



Current release 21 configuration

> Latest tests are up and running using release 21.0.13, detailed on new twiki <https://twiki.cern.ch/twiki/bin/view/Atlas/Winter2017Run2Reprocessing>

- Use this twiki page to get the full history or release 21 tests and bug fixes: <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/Release21DataReprocessingValidation>
- As usual we test the full reprocessing workflow using two runs, one from 2015 and one from 2016
**AOD HIST DRAW_ZMUMU DRAW_EGZ DRAW_TAUMUH DRAW_EMU DESDM_SGLEL DESDM_SLTTMU DESDM_MCP
DESDM_CALJET DESDM_PHOJET DESDM_EGAMMA DAOD_IDTIDE DRAW_RPVLL DESDM_EXOTHIP**
- And then: **DESDM_ZMUMU DAOD_ZMUMU DESDM_EGZ DAOD_EGZ DAOD_TAUMUH DESDM_EMU DAOD_EMU**
- As well as testing: **Reco_tf, AODMerge_tf, ESDMerge_tf, HISTMerge_tf**

- Job options for 2015 and 2016:

```
preExec: "all:DQMonFlags.enableLumiAccess=False;"
```

```
postExec: "e2d:from AthenaCommon.AppMgr import ServiceMgr; import  
MuonRPC_Cabling.MuonRPC_CablingConfig; ServiceMgr.MuonRPC_CablingSvc.RPCMapfromCool=False;  
ServiceMgr.MuonRPC_CablingSvc.CorrFileName="LVL1confAtlasRUN2_ver016.corr";  
ServiceMgr.MuonRPC_CablingSvc.ConfFileName="LVL1confAtlasRUN2_ver016.data";"
```

```
postExec: "r2e:from AthenaCommon.AppMgr import ServiceMgr as svcMgr;  
svcMgr.AthenaPoolCnvSvc.MaxFileSizes=["15000000000"];"
```

and for 2015 only:

```
preExec: "r2e:from LArConditionsCommon.LArCondFlags import larCondFlags;  
larCondFlags.OFCShapeFolder.set_value_and_lock("4samples1phase");"
```

- Also running now with **--athenaMPMergeTargetSize 'ESD':0.0**, which prevents the **tmp.ESD** from merging in the m-core job

Remaining release 21 issues

- > Release 21 is characterised by issues with the enormous **tmp.ESD**, as we cannot run with **RAWtoALL** due to the high memory requirements (up to 4GB/core)
 - Hence now turning off this merge in our m-core jobs and increasing the **domMaxFileSize** to 15GB
 - But even then we are still battling with this sometimes (recent ALFA prod needed split inputs)
 - In addition, **ESDMerge_tf** was also leaking 1MB/event, now fixed <https://its.cern.ch/jira/browse/ATLASRECTS-3711>
 - Nevertheless, this is still happening and jobs are crashing because of large ESDs..
- > One outstanding issue seen in release 21.0.13 (so far..):
 - Transient crash in HLT monitoring, <https://its.cern.ch/jira/browse/ATR-15479>, some progress here, but this is a big hit in job completion efficiency
 - Already seen in 21.0.9 but still not fixed and attempts tail is long, with some jobs completing on 14th try
 - Must be fixed before full production
- > Also, not related to release, but issue of all jobs failing at SLAC – needs ticket?
 - e.g. Job <http://bigpanda.cern.ch/job?pandaid=3132894832> and log: http://aipanda057.cern.ch/media/filebrowser/dd9fb46c-73b1-4e8d-a835-7b99085fc55f/data16_13TeV/tarball_PandaJob_3132894832_SLACXRD_MP8/athena_stderr.txt
 - Error: `/cvmfs/atlas.cern.ch/repo/ATLASLocalRootBase/swConfig/./atlantis/menu.sh: No such file or directory`

Christmas 2016 plan

Request ID	Task table	Submitted by	Description	Runs	Reco tag Output formats	Merge tag(s) and step(s)	Status	Notes
10321	10321	dsouth	Slice test for data15_13TeV run 284285	1	r8965 All formats described above	p2938 AOD p2937 DESD (x2) p2938 DAOD (x2) p2939 HIST (x3)	running	
10322	10322	dsouth	Slice test for data16_13TeV run 304008	1	r8966 All formats described above	p2938 AOD p2937 DESD (x2) p2938 DAOD (x2) p2939 HIST (x3)	running	
10324	10324	dsouth	2015 express stream runs	10	r8965 AOD.HIST	p2939 HIST (x2)	prepared	
10325	10325	dsouth	2016 express stream runs	15	r8966 AOD.HIST	p2939 HIST (x2)	prepared	
10327	10327	dsouth	Pixel studies using Z to mu mu	3	r8966 DAOD_ZMUMU	-	prepared	
10323	10323	dsouth	DRAW production of five 2016 runs used to check release 21 alignment	5 and 5	r8876 DAOD_EGZ, DESDM_EGZ DAOD_ZMUMU, DESDM_ZMUMU	--	prepared	

- All 21.0.13 requests prepared, two usual slice tests up and running
- Expect slice test to be done on timescale of about 7-10 days (60k jobs, average attempt number about 1.5, each one running on average 3 hours, say ~ 270 CPU hours)
- An additional run 306310 maybe added, another ~ 30k jobs - will see how others progress
- Express stream small, AOD and HIST only, certainly should run within current resources
- Others even smaller, and staging from tape working fine