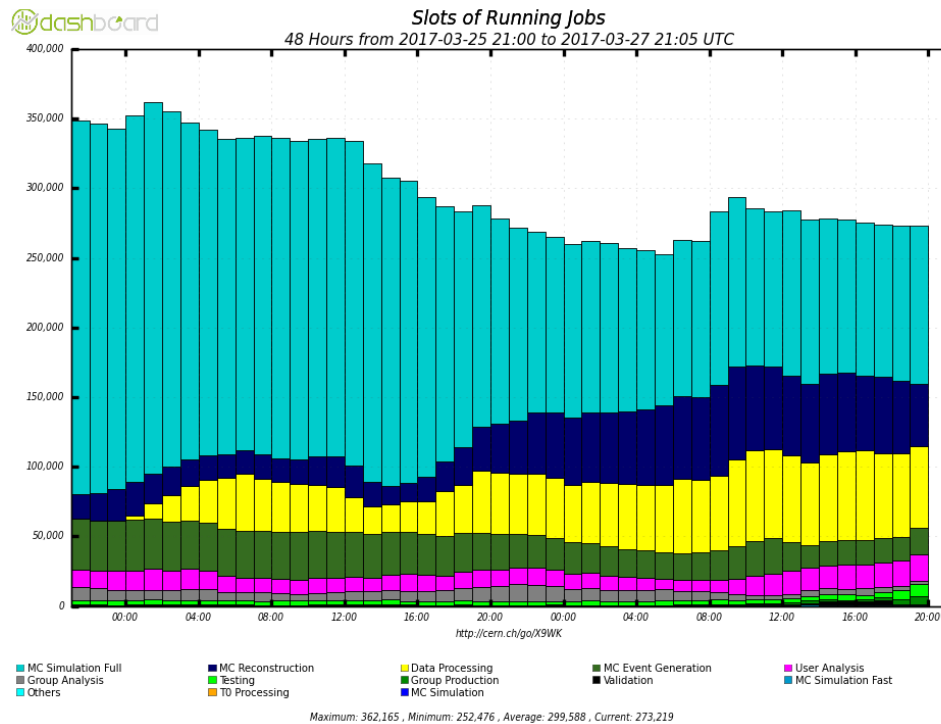


# Reprocessing Status

## 2017 Run 2 bulk reprocessing

[https://twiki.cern.ch/twiki/bin/view/Atlas/Winter2017Run2Reprocessing#Bulk\\_processing](https://twiki.cern.ch/twiki/bin/view/Atlas/Winter2017Run2Reprocessing#Bulk_processing)



David South (DESY)

ADC Weekly  
March 28th 2017



# Release 21.0.19 production

> Total of 18 physics\_Main runs (17 runs from 2016, run 284285 from 2015)

- Full spread of formats also produced in reprocessing workflow; here are the HIST files:

`data15_13TeV.00284285.physics_Main.merge.HIST.r9214_p3069_p3069_p3069`

`data16_13TeV.*.physics_Main.merge.HIST.r9214_p3069_p3069_p3069`

> Express stream (10 runs 2015; 17 runs 2016)

- `data16_13TeV.*.express_express.recon.AOD.r9214/`

`data16_13TeV.*.express_express.merge.HIST.r9214_p3069_p3069/`

> DRAW production (5 runs from 2016)

- `data16_13TeV.*.physics_Main.recon.DAOD_EGZ*r9214/`

`data16_13TeV.*.physics_Main.recon.DESDM_EGZ*r9214/`

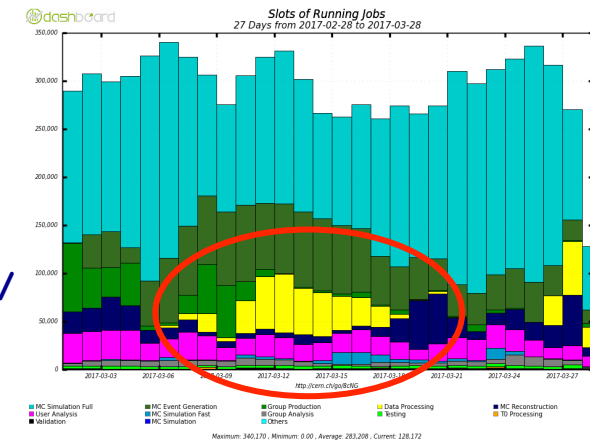
`data16_13TeV.*.physics_Main.recon.DAOD_ZMUMU*r9214/`

`data16_13TeV.*.physics_Main.recon.DESDM_ZMUMU*r9214/`

> CosmicCalo production (6 runs from 2016)

- `data16_13TeV.*.physics_CosmicCalo.recon.AOD.r9214/`

`data16_13TeV.*.physics_CosmicCalo.merge.HIST.r9214_p3069_p3069/`



# Problems and issues seen in release 21.0.19

- > 1) Bug in Tile reconstruction software, wrong correction applied for 50ns data
  - Fixed by **TileRecUtils-00-09-80-01**, scheduled for 21.0.19.1 and then 21.0.20
- > 2) Calo DQ monitoring fix, conditions only
  - **DetStatusDEFECTS-RUN2-BLK-UPD2-01**  
**DetStatusDEFECTLOGIC-RUN2-BLK-UPD4-01**
  - Will need new global conditions tag **CONDBR2-BLKPA-2017-05**
  - But this one will not require running again for the already submitted `physics_Main` runs
- > 3) Amendment needed to monster AOD-reduction preExec used in reco jobs  
<https://its.cern.ch/jira/browse/ATLASRECTS-3820>
  - Additional containers should be removed, to allow jet/met derivation to run correctly
  - Another configuration issue discovered in flavour tagging
  - Observe the carnage by clicking on “preExec” here:  
[https://ami.in2p3.fr/app?subapp=amiTags\\_compare&userdata=r9214,r9264](https://ami.in2p3.fr/app?subapp=amiTags_compare&userdata=r9214,r9264)
- > 4) Concern about track linking in CombinedMuonTrackParticle  
<https://its.cern.ch/jira/browse/ATLASRECTS-3988>
  - Conclusion was to proceed as is and can be AODFixed at some point

# Release 21 configuration for reprocessing

> Full reprocessing workflow produces:

**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 DRAW\_TOPSLMU**

▪ Then: **DESDM\_ZMUMU DAOD\_ZMUMU DESDM\_EGZ DAOD\_EGZ DAOD\_TAUMUH DESDM\_EMU DAOD\_EMU**

> Release: **21.0.20**

Compiler: **gcc62**

Geometry tag: **ATLAS-R2-2016-01-00-01**

Conditions tag: **CONDBR2-BLKPA-2017-05**

> 2016 runs: **179** runs of **physics\_Main**

- Periods A-L except H (low mu) and J (ALFA), grouped together in blocks of between 27 and 36 runs
- Periods H (4 runs) and J (7 runs): **physics\_MinBias** stream to be considered

> 2015 runs: **112**; 106 **physics\_Main** and finally **6 physics\_MinBias** (period B)

> Other streams to consider as done in 2015 reprocessing:

- **express\_express, physics\_CosmicCalo, physics\_ZeroBias, physics\_L1Calo, debugrec\_hlt**

# Release 21 configuration for reprocessing

## > Now using same job options for 2015 and 2016:

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

```
preExec: "from InDetRecExample.InDetJobProperties import
InDetFlags; InDetFlags.useDynamicAlignFolders.set_value_and_lock(True);"
```

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

```
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"];"
```

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

## > Also running with full (huge) AOD reduction job options, including additional statements to run producing HIST, as detailed here:

<https://its.cern.ch/jira/browse/ATLASRECTS-3820>

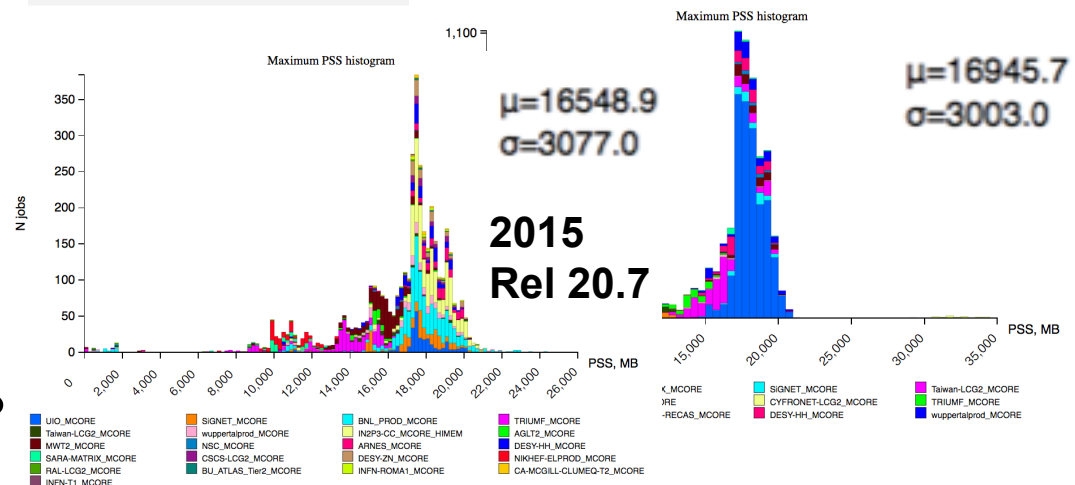
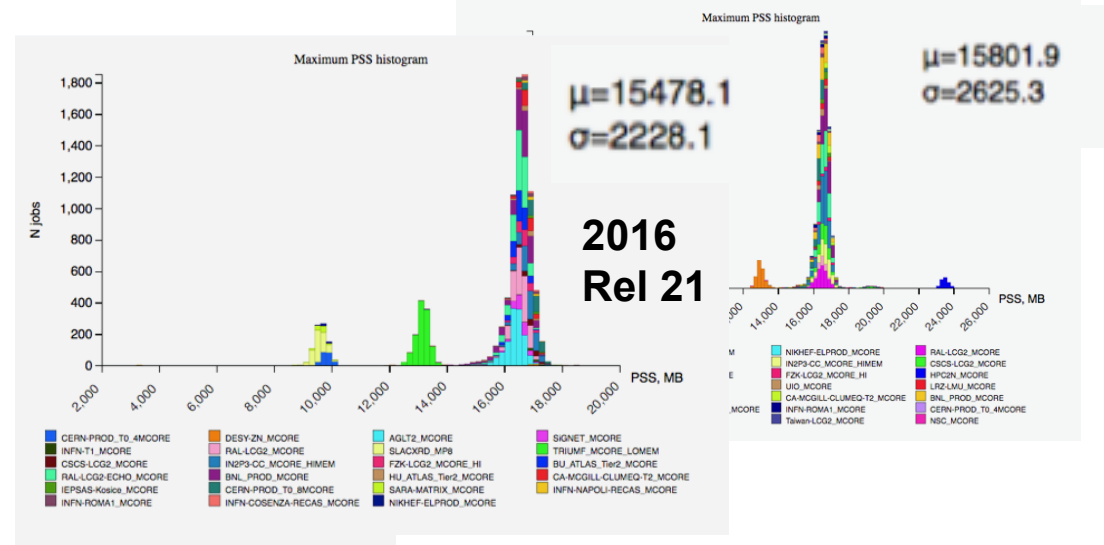
## > All reco tasks now run on m-core (including DRAW); all merges run on s-core

# DRAW reco configuration

- New DRAW reco job configuration running over multiple (up to 100) inputs, configuration will be used for remainder of bulk reprocessing
  - Up to factor of 100 reduction in no. of DRAW reco jobs, and most run within an hour or so
  - **Also now running on m-core**
- Problem found in that setting NE=3000 was not passed to JEDI from the `tgtNumEventsPerJob=3000` in `project_mode`
  - Always taking 100 inputs, DRAW\_ZMUMU jobs ending up with too large DESD\_ZMUMUs
  - Fixed by changing entry in “Number of Files” box to “0” from “-1” (wrong default)
  - Changed by hand for running jobs, and two more runs submitted to validate it
  - Works, will be changed for all future submissions

# Current status and schedule

- Usual slice test run submitted on Saturday, following last tests of preExec
- Final ten runs of 2016 sent on Sunday, progressing well. Additional two runs from period L sent yesterday to validate the DRAW reco fix
- Peak slots yesterday about 60k, currently about 40k
- ADC report memory usage too high, here some total plots approaching 16GB
- Last data: these are the highest mu, toughest runs in 2016 (we imagine)
- 2015 reprocessing in release 20.7 was pretty much the same, in fact a little worse
- We only used ~ 50k last time?



## Next

- > Suggestion on event forking does not work: it breaks monitoring (no HIST)
  - Also not really validated in release 21 and “classic” mode (rather than RAWtoALL)
  
- > Logical next step would be add rest of period L, then I+K, then F+G..
  - However: memory issue maybe prevent this
  
- > We could come at 2016 from both sides, i.e. add in period A now, assuming that earlier periods are lighter, lower mu etc
  - Then we would (hopefully) have a mix of memory needs, lowering the median..
  - There is also a demand to have 2015 low mu rather soon, which should be no problem
  
  
- > There is also the idea to force some of the tasks to 2GB/core (increasing the amount of swapping going on?)
  
- > Also: How is the frontier / squid situation?



# Extras