Reprocessing Status

2017 Run 2 reprocessing

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

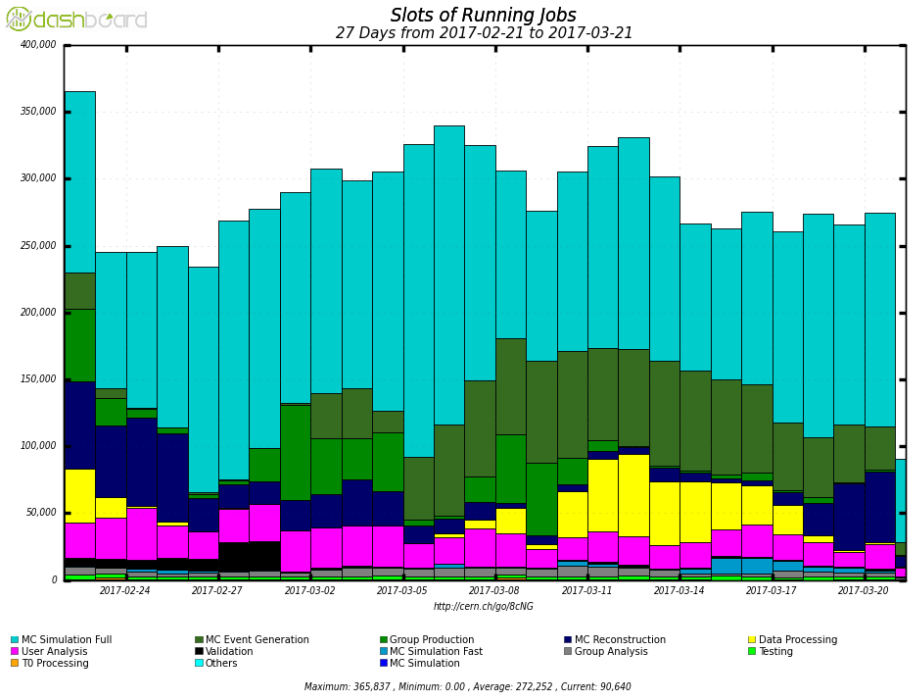

David South (DESY)

ADC Weekly
March 21st 2017

HELMHOLTZ ASSOCIATION

DESY

What's available so far in release 21.0.19

- Total of 18 `physics_Main` runs submitted (17 2016, run 284285 from 2015)
  - Majority done inside a week, one reco task remains (last 22 jobs) and tails of some merges
  - Full spread of formats also produced in reprocessing workflow; HIST files:
    data15_13TeV.00284285.physics_Main.merge.HIST.r9214_p3069_p3069_p3069
    data16_13TeV.*.physics_Main.merge.HIST.r9214_p3069_p3069_p3069 (16/17 available)

- 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/

There's already (at least) 3 known release 21.0.19 issues

- 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
  - Additional containers should be removed, to allow jet/met derivation to run correctly:
    Fix tested available for this part

  Two new issues found related to the AOD-reduction preExec:

  - Another configuration issue discovered in flavour tagging this morning, as well as a (probably non-critical) muon reco issue
  - This means either resubmitting the `physics_Main` runs already done, *OR* running some `AODtoAOD` job that strips this information from the AOD – to be clarified
- Some issues seen in DQ and PhysVal (jet calibration, primarily) are believed to be understood or not showstoppers

Performance on the grid

- A significant increase was seen on squid hits, with failover to RAL resulting in a cascade effect and increasing load on Frontier server
  - Majority of cases coming from s-core reco jobs, which are the DRAW ones
  - Also not ruled out the AOD and ESD merge jobs, which for some reason need conditions like RPC and TILE, and where the AMI tags now have a `postInclude` "`all:RecJobTransforms/UseFrontier.py`"
- There's a line of investigation into why such large conditions payloads are being loaded at in the reco jobs: Suspicion was that some DCS data is no longer cached – *any update on this?*
- New DRAW reco job configuration running over multiple (up to 100) inputs per job has been validated, configuration will be used for remainder of bulk reprocessing
  - Factor of 100 reduction in number of DRAW reco jobs, and most run within an hour or so
  - Also now running on m-core

Performance on the grid: A few more things

- I have seem quite a number of staging issues, but thanks to Rod and Ivan these have been fixed quite quickly
  - It is however often still the case that a job will try only the one site when other replicas are available and the go to exhausted
- Transfer issues again at GLASGOW and RAL, but this seems to have been resolved now
- A couple of cases of bad RAW inputs
  - One physics_Main, where they have 0 events https://its.cern.ch/jira/browse/ATLASJT-356
  - Another in CosmicCalo, revisting an old issue of LB number after the end of the run https://its.cern.ch/jira/browse/ATLASRECTS-3843
- There's some memory issue at INFN (again, Rod knows)
  - HIST merges will fail 15 times on INFN_T1 due to lost heartbeat, not reassigned, fixed only by changing ramcount to 3000 (default is 2000)
- Also had one first step HIST file lost on eos, which meant redoing all merges again for that run to be consistent
  - (Although this allowed the new DRAW configuration to be tested)

Current status and schedule

- I now have next to nothing running on the grid, most 21.0.19 tasks completed
- Before submitting further `physics_Main`
  - New conditions tag needed for additional monitoring
  - New r-tag with updated AOD reduction preExec
  - Some things are still not 100% resolved
- I sincerely hope this will come in the next days
  - If we still want a new test for these cached DCS conditions updates, we can do that today with the existing configuration
  - More delays we will need more slots as originally planned (but MC is getting more now..)
- New release will be needed for 50ns 2015 data, but this is far away in the schedule at this point

Current status and schedule

- I now have next to nothing running on the grid, most 21.0.19 tasks completed
- Before submitting further `physics_Main`
  - New conditions tag needed for additional monitoring
  - New r-tag with updated AOD reduction preExec
  - Some things are still not 100% resolved
- I sincerely hope this will come in the next days
  - If we still want a new test for these cached DCS conditions updates, we can do that today with the existing configuration
  - More delays we will need more slots as originally planned (but MC is getting more now..)
- New release will be needed for 50ns 2015 data, ~~but this is far away in the schedule at this point:~~ Scrap that, turns out the minBias data is needed for CP recommendations for 2017 running, so please can you stage:

```
data15_13TeV:data15_13TeV.00267358.physics_MinBias.merge.RAW
data15_13TeV:data15_13TeV.00267359.physics_MinBias.merge.RAW
data15_13TeV:data15_13TeV.00267360.physics_MinBias.merge.RAW
data15_13TeV:data15_13TeV.00267367.physics_MinBias.merge.RAW
data15_13TeV:data15_13TeV.00267385.physics_MinBias.merge.RAW
data15_13TeV:data15_13TeV.00267599.physics_MinBias.merge.RAW
```

Extras

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: **21.0.19**
  Compiler: **gcc62**
  Geometry tag: **ATLAS-R2-2016-01-00-01**
  Conditions tag: **CONDBR2-BLKPA-2017-04**
- 2016 runs: **182** in current list, validating today, all **physics_Main**
  - First 17 runs across year, then second half of the year, followed by the first half
- 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
  - Removed the `'ToolSvc.InDetSCTRodDecoder.+ERROR.+Unknown.+offlineId.+for.+OnlineId'` ignorePatterns field from the r-tag
- The main reco task runs on m-core; everything else on s-core
- 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