

# Reprocessing / ProdSys2.

Some issues arising from recent reprocessing campaigns

Lily Asquith (Uni Sussex)  
David South (DESY)

ProdSys2 Meeting  
May 6th, 2015



# Recent lost files and automatic recovery procedure

File loss iteration	Run number	Number of events according to RAW	Original lost file(s) (all four outputs are subsequently declared lost)	Replacement Files
1	259003	7,525,203	HIST.05326657._000059.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326657._004793.pool.root.1
			NTUP_MUONCALIB.05326657._000400.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326657._004794.pool.root.1
	259165	9,794,826	NTUP_MUONCALIB.05326679._004176.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326679._006035.pool.root.1
			NTUP_MUONCALIB.05326679._004613.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326679._006036.pool.root.1
	259424	6,709,377	NTUP_MUONCALIB.05326699._001820.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326699._003893.pool.root.1
	259555	7,243,353	HIST.05326703._002469.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326703._003616.pool.root.1
	259756	8,372,076	NTUP_MUONCALIB.05326725._000883.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326725._003920.pool.root.1
NTUP_MUONCALIB.05326725._002219.pool.root.1			AOD/ESD/HIST/NTUP_MUONCALIB.05326725._003922.pool.root.1	
2	259003	7,525,203	AOD.05326657._002614.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326657._004795.pool.root.1
			HIST.05326657._003160.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326657._004796.pool.root.1
			NTUP_MUONCALIB.05326657._003374.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326657._004797.pool.root.1
			HIST.05326657._004710.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326657._004798.pool.root.1
	259424	6,709,377	HIST.05326699._002186.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326699._003896.pool.root.1
259555	7,243,353	AOD.05326703._003547.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326703._003618.pool.root.1	
3	259424	6,709,377	NTUP_MUONCALIB.05326699._000447.pool.root.1	AOD/ESD/HIST/NTUP_MUONCALIB.05326699._003897.pool.root.1

- Fifteen lost reco output files, corresponding to lost inputs to merge jobs
  - Three separate incidents, only a few days after file creation
- Mixture of AOD, HIST, NTUP\_MUONCALIB files, details on our twiki [https://twiki.cern.ch/twiki/bin/view/Atlas/ReprocessingCampaigns2015 - Problems\\_with\\_this\\_production](https://twiki.cern.ch/twiki/bin/view/Atlas/ReprocessingCampaigns2015 - Problems_with_this_production) and in the JIRA ticket: <https://its.cern.ch/jira/browse/ADCSUPPORT-4196>
- Recovery of files by Tadashi; but this currently only treats reco part
- Task goes from **done** to **running** then back to **done**: Reactivates merges!
- Merged ds then get broken due to picking up the replacement files

# Recent lost files and automatic recovery procedure (2)

Task Name	Request	Step	Task ID	Priority	Total Jobs	Done Jobs	Events	Status	Submit time	Timestamp	Provenance	JIRA
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331_p2331	2402	Reco	5385414	900	1	1	6709377	done	May 04 10:21	May 05 06:11	AP	199812
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331_p2331	2402	Reco	5385413	900	2	2	6709377	done	May 04 10:20	May 05 03:11	AP	199811
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331	2402	Reco	5385412	900	65	64	6709377	done	May 04 10:20	May 04 21:31	AP	199810
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2334	2402	Reco	5384872	990	635	634	6709377	done	May 04 05:50	May 04 20:31	AP	199712
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2335	2402	Reco	5384871	900	317	317	6709377	done	May 04 05:50	May 04 19:31	AP	199711
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331_p2331	2402	Reco	5382984	900	1	1	6711590	obsolete	May 03 15:31	May 04 10:11	AP	199701
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331_p2331	2402	Reco	5382983	900	2	2	6711590	obsolete	May 03 15:31	May 04 10:11	AP	199700
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331	2402	Reco	5382979	900	65	65	6711590	obsolete	May 03 15:31	May 04 10:11	AP	199699
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2334	2402	Reco	5382972	900	646	634	6700863	obsolete	May 03 15:30	May 04 05:43	AP	199695
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2335	2402	Reco	5382775	900	318	318	6711590	obsolete	May 03 14:20	May 04 05:42	AP	199689
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331_p2331	2402	Reco	5371894	900	1	1	6602841	obsolete	Apr 30 15:41	May 03 15:15	AP	199092
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331_p2331	2402	Reco	5371893	900	2	2	6602841	obsolete	Apr 30 15:41	May 03 15:15	AP	199091
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331	2402	Reco	5371892	900	82	63	6602841	obsolete	Apr 30 15:41	May 03 15:15	AP	199090
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2334	2402	Reco	5371880	900	715	635	6711540	obsolete	Apr 30 15:26	May 03 15:20	AP	199080
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2335	2402	Reco	5371802	900	334	317	6709377	obsolete	Apr 30 15:11	May 03 14:03	AP	199027
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2334	2402	Reco	5345003	900	655	634	6709377	obsolete	Apr 23 19:16	Apr 30 15:11	AP	193215
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2337	2402	Reco	5342446	900	42	0	0	aborted	Apr 23 12:12	Apr 23 16:11	AP	192663
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2335	2402	Reco	5329484	900	318	318	6711571	obsolete	Apr 21 13:45	Apr 30 14:53	AP	190511
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2334	2402	Reco	5329461	900	61	0	0	aborted	Apr 21 13:44	Apr 22 15:31	AP	190488
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331_p2331	2402	Reco	5329420	900	2	2	6711571	obsolete	Apr 21 13:42	Apr 30 15:24	AP	190457
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331_p2331	2402	Reco	5329418	900	3	3	6711571	obsolete	Apr 21 13:41	Apr 30 15:24	AP	190456
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2331	2402	Reco	5326783	900	65	65	6711571	obsolete	Apr 20 18:56	Apr 30 15:24	AP	190360
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2332	2402	Reco	5326754	900	716	635	6711571	obsolete	Apr 20 18:55	Apr 30 18:49	AP	190337
data15_cos.00259424.physics_CosmicMuons.merge.r6622_p2329	2402	Reco	5326702	900	0	0	0	aborted	Apr 20 18:52	Apr 21 15:11	AP	190306
data15_cos.00259424.physics_CosmicMuons.recon.r6622	2402	Reco	5326699	900	3897	3171	6715947	done	Apr 20 18:52	May 04 03:11	AP	190305

Showing 1 to 25 of 25 entries

- Current solution is to **obsolete** the original merge tasks and resubmit them
  - In the above example the merges had to be submitted four times
- We only notice the lost inputs when we run the merges, i.e. when we use them – maybe there are more lost reco outputs? (e.g. we don't merge ESDs)
- We only realised this week that the recovery procedure is affecting the merged datasets – maybe there are other merged ds also affected?

## Recent lost files and automatic recovery procedure (3)

- > A possible solution would be that the child merge files created using the lost files would also be removed and remade containing the new files instead.
- > Example:
  - The file `NTUP_MUONCALIB.05326725._000883.pool.root.1` is lost
  - All `AOD/ESD/HIST/NTUP_MUONCALIB.05326725._000883.pool.root.1` are declared lost and replaced by new files `AOD/ESD/HIST/NTUP_MUONCALIB.05326725._003920.pool.root.1`
  - ✓ Up to this point is what is currently done
  - The next step would be to declare **ANY** merge file that contains a `05326725._000883` file lost and replace it by a new merge file containing the new `05326725._003920` file
- > Another problem here: Some lost file declarations in Rucio (specifically the formats that were not actually lost) are not propagated to AMI
  - Some progress understanding this, interplay between `declare-bad` and `detach-dids`
  - Details and example in <https://its.cern.ch/jira/browse/PRODSYS-451>

# Containers of containers

- > Problem using containers of containers in **SampleHandler/EventLoopGrid**, JIRA ticket on this: <https://its.cern.ch/jira/browse/ATLASG-58>
- > From our side, AMI interface is identical and we create Physics Containers exactly as we did before: change is in the background
- > Recent containers made, then fixed to contain datasets rather than containers  
`data12_8TeV.periodAllYear.physics_Muons.PhysCont.DAOD_ZMUMU.repro17_v02`  
`data12_8TeV.periodB.physics_Muons.PhysCont.AOD.repro17_v02`  
(for run 204442, the unmerged AODs are added to the container due to the b-tagging issue)
- > At the same time these containers were also remade with datasets  
`data12_8TeV.periodAllYear.physics_Egamma.PhysCont.DAOD_ZEE.repro17_v01`  
`data12_8TeV.periodB.physics_Egamma.PhysCont.AOD.repro17_v01`  
`data12_8TeV.periodB.physics_JetTauEtmiss.PhysCont.AOD.repro16_v05`
- > Current default is to continue to fill with containers and then ask ADC to run script to replace the containers with the datasets (tids)
- > Long term plan?

# Finished status of data reprocessing tasks

- > Data tasks should, in principle, never go to **finished**, as we want 100% success rate. They should at least not go to **finished** automatically
- > Example:
  - Reco task goes to **finished** not **done**, because last job failed after 15 attempts (e.g. not enough memory)
  - Merges start, and HIST merge goes through three steps (factor 50 each time) producing 97, 2, 1 files
  - After adding more retries, the reco then gets to **done**, so the merges pick up this extra file
  - For ESD, AOD, NTUP we “don't care”, because it is just added
  - But for HIST (and TAG) we want one file, so this makes 98, 3, 2 files in each step
- > Second example:
  - First HIST merge step has gone to **finished**, rather than **done**, due to a missing input.
  - The second and third steps then follow and get to **done**.
  - Missing file is then recovered: first step goes to **done**.
  - Second and third steps then re-run, again resulting in 2 files in the final merged dataset
- > Current solution is to **obsolete** the final (third) merge step and resubmit to get 1 file again
- > Only affected datasets are when the final merge step should contain 1 file

# Finished status of data reprocessing tasks

- > The issue can be solved by
  - 1) Automatically fixing such final step merges, some “auto-re-merge” procedure
  - 2) Stopping data reprocessing tasks going to **finished** in the first place, which is preferable, as we retain more control: we know if something is not 100% successful
    - Note that neither solution fixes the duplicate files in merged datasets caused by the current recovery procedure
- > Option 2) could be implemented as a new state, **tocheck** (or similar), which would need manual intervention to set it to **finished** – there are jobs we know cannot complete, so would still have some **finished** tasks
- > It could also have a 10 day timeout, so that if no action is taken the task does go to **finished** to avoid jobs hanging in the system forever
- > The option of not starting merges before reco is **done** is not desirable, as if we have 6000 jobs and we’re merging in blocks of 5, this should start as soon as possible