



# Data processing

Offline review  
Feb 2, 2011

# Productions, tools and results

- Three basic types of processing
  - RAW
  - MC
  - Trains/AODs
- I will go through these with as few details as possible
- I will not discuss the AliRoot evolution
  - This deserves a long separate presentation

# MC productions

- MC productions (on behalf of PWGs)
  - Request in AliRoot Savannah (as tasks)
  - Macros and JDLs (supplied by the requestor) in MonALISA
  - The priority and validity of the request are discussed by PB
    - Exceptions – for short and urgent productions this step is usually short-circuited, e.g. first physics, overnight, weekend (the Friday evening syndrome)

# MC production statistics

- MC productions (on behalf of PWGs)
  - All normal requests and exceptions have been handled
  - We have run 161 different MC production cycles throughout 2010 (effectively one each 2 days), quite unevenly spread
  - Total of 969Mio events

# MC production in 2010

- Recent PB decisions
  - Reduction of MB cycle to 10% of RAW statistics
  - Compensated by PWG specific signals/configuration requests
  - The effective reduction in processing time/resources in general will be negligible, if any
- Likely, the number of MC production requests will grow this year...

# MC production – what follows

- QA PWG1 train
- AODs filtering train
  - Luckily all of the above without tenders (yet)
  - ESD and AOD are still consistent
- The process is simpler than for RAW data, but still not fully automatic – AOD trains always have ‘last minute changes’ and usually are delayed with respect to the ESD processing
  - The usual culprit is excessive memory consumption

# RAW data productions

- Two types – Physics and Special
  - Pass 1 (with or without Pass0) – automatic, quasi online, always latest AliRoot revision
  - Pass2+ and special – AliRoot savannah as tasks
  - The tasks usually follow the status of calibration and code updates necessary for the pass to start
  - A complicated, long and ever changing procedure...

# RAW data productions (2)

- Special passes
  - Detector calibration runs
    - Usually no problems with these
  - Specific 'quality check' physics runs or detector calibration passes (calorimetry)
    - With updated AliRoot version and/or OCDB
    - Either precursors of or post-Pass(x)
    - Tend to be repetitive and follow each other quickly i.e. 'Try again after a bug fix or OCDB update'
    - Not necessarily small statistics



# RAW data passes and statistics

- Total of 52 RAW data passes in 2010
- 2,197,284,148 events (as of last evening)
  - Does not include the Pass0 processing, but includes the detector calibration data

# RAW data reco specifics

- More complex than MC, 3 basic reasons
  - Detector code is better ‘adapted’ to MC processing
  - More events per job – memory leaks and reconstruction code exceptions are appearing here
  - The run conditions change during the period, the code needs to be adapted (more on this on the next slide)

# RAW data reco specifics (2)

- Pass 1 processing
  - Run always with the latest available AliRoot revision
  - This catches running reco problems, as seen during the quasi-online reconstruction
  - Fixes follow weekly
- Ideally (and this is also in the CM) Pass2 should follow shortly after Pass1, effectively the latest code, only calibration update
  - Same AliRoot revision, should be OK...

# RAW data reco specifics (3)

- What happens in reality
  - Long period of calibration after the period is over
  - Numerous changes in the code, usually relevant to the next running period
    - And often specific to it
  - OCDB updates – same as above
- As a consequence for Pass2 the conditions have changed considerably
  - Case in point LHC10e(2)

# RAW data – what happens next

- Analysis trains
  - PWG1 QA train – running after each pass
  - Tag merging
  - AOD filtering (several trains), delayed with respect to production up to a month
    - These are with tenders...
    - Running with AN tag
    - And are repeated several times

# Trains statistics

- 280 Analysis trains in 2010
  - about 1/2 are merging trains (reduces fragmentation)
  - Each train has a unique configuration – detectors included/excluded depending on readiness
  - Always on the bleeding edge of AliRoot AN tag...
- None of the trains are calibration...with the exception of Pass0

# Summary of production in 2010

- 161 MC cycles
- 52 RAW cycles
- 280 analysis trains
- Almost each of the above has unique configuration
  - We've managed reduce the macros complexity only for Pass1(+) reco
- A small demo of the cascading of tasks

# General risk factors

- Configuration
  - Error in the processing macros
- OCDB
  - Wrong or out-of-synch update
- Code
  - Insufficient testing prior to submission
- All of the above – never enough time
- New hooks and pre/post-processing are being added continuously



## General risk factors (2)

- Validation and tests - the results are not verified quickly enough
  - Given the vast amount of output, this is not surprising
- The 'validation pause', introduced long ago is simply impractical
  - The production must finish for the next step (or cycle) to begin

# Beware of 'good ideas'

- Accidental deletion of 900 GeV data
  - Someone in 2007 had the 'good idea' to introduce recycle tape pools
  - Basically, automatic data deletion in a storage system was entrusted to a script
  - It was just a question of time before something important went there... and it did

# Summary

- The ALICE production system was able to cope with all processing requests in 2010
  - Which were ready to start
- We have a relatively flexible approach and tools to carry out the production
  - This includes the code/OCDB update practices
  - But this seem to just increase the complexity of the tasks

## Summary (2)

- The general trend is to push processing toward later tasks
  - Online calibration → Pass0
  - Reco QA → PWG1 train
  - AOD filtering → AOD filtering (with tender)
  - User analysis tasks → ESDs and AODs with tender
- Consequences – see Federico's and Yves' slides