Distributed Data Analysis in the ATLAS Experiment: Challenges and Solutions

Johannes Elmsheuser[1],
Daniel van der Ster[2]
for the ATLAS collaboration
[1]Ludwig-Maximilians-Universität München, Germany, [2]CERN, Switzerland

1. Introduction

The ATLAS experiment at the LHC at CERN is recording and simulating several 10's of PetaBytes of data per year. To analyse these data the ATLAS experiment has developed and operates a mature and stable distributed analysis (DA) service on the Worldwide LHC Computing Grid.
The service is actively used: more than 1400 users have submitted jobs
in the year 2011 and a total of more 1 million jobs run every week. Users are provided with a suite of tools to submit Athena, ROOT or generic jobs to the grid, and the PanDA workload management system is responsible for their execution. The reliability of the DA service is high but steadily improving; grid sites are continually validated against a set of standard tests, and a dedicated team of expert shifters provides user support and communicates user problems to the sites.

2.ATLAS Analysis

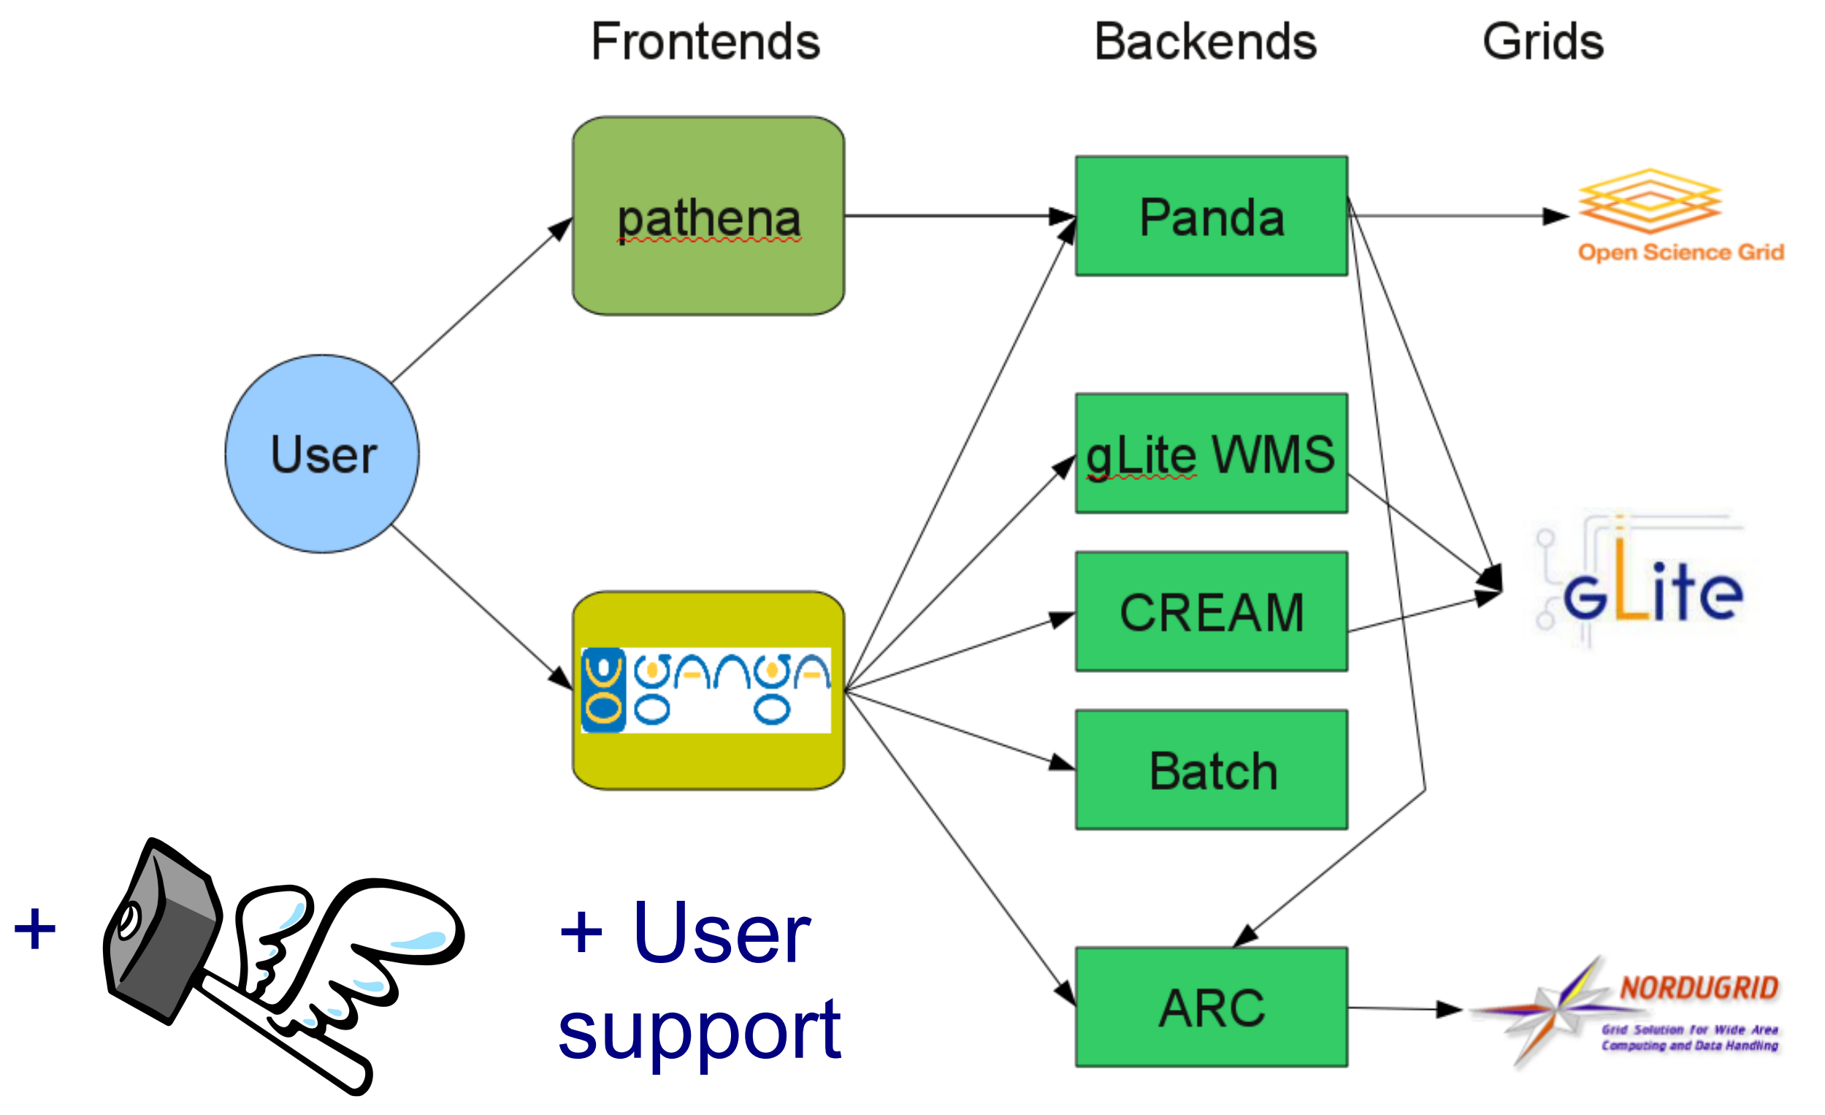

ATLAS analysis in a nutshell:

- Detector and MC data distributed worldwide by the data management system DQ2 to sites in hierarchical tier structure.
- Experiment software Athena is installed at all sites
- User analysis code and tasks: ,,Job to Data".
- Distributed analysis tools: Ganga and the Panda clients
- Output of user jobs are stored on site scratch disks

3. Analysis usage

April 2011-April 2012:

- More than 1500 ATLAS grid users
- 300-400k finished jobs per day
- 20-40k running jobs
- Job duration is rather short: only 10-60 min

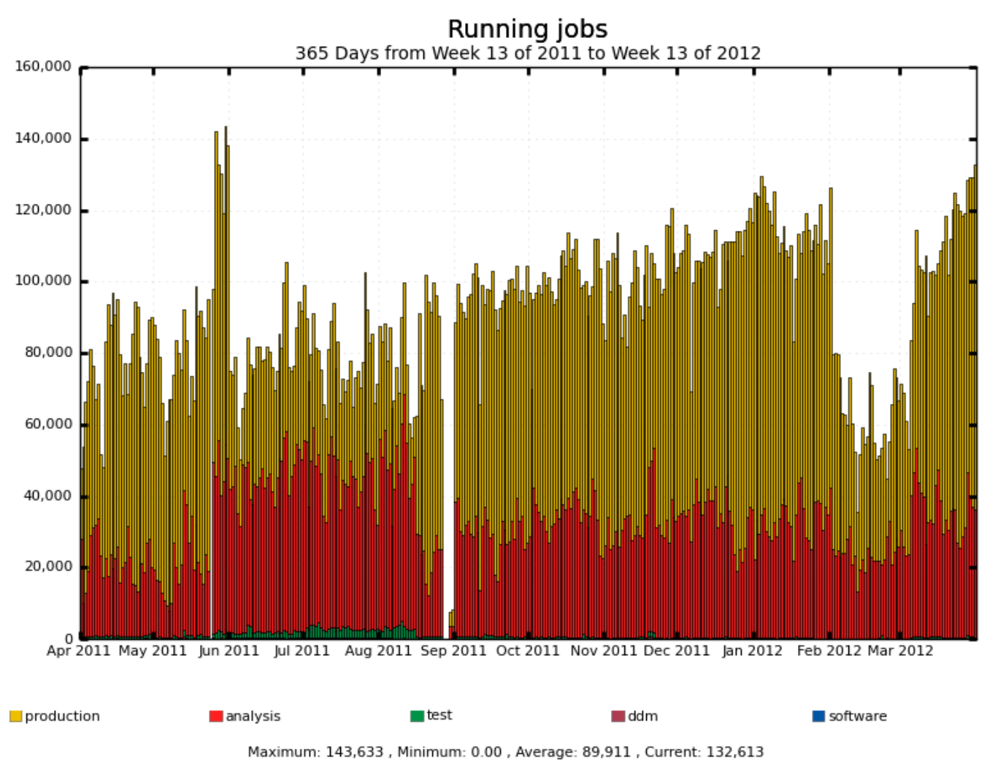

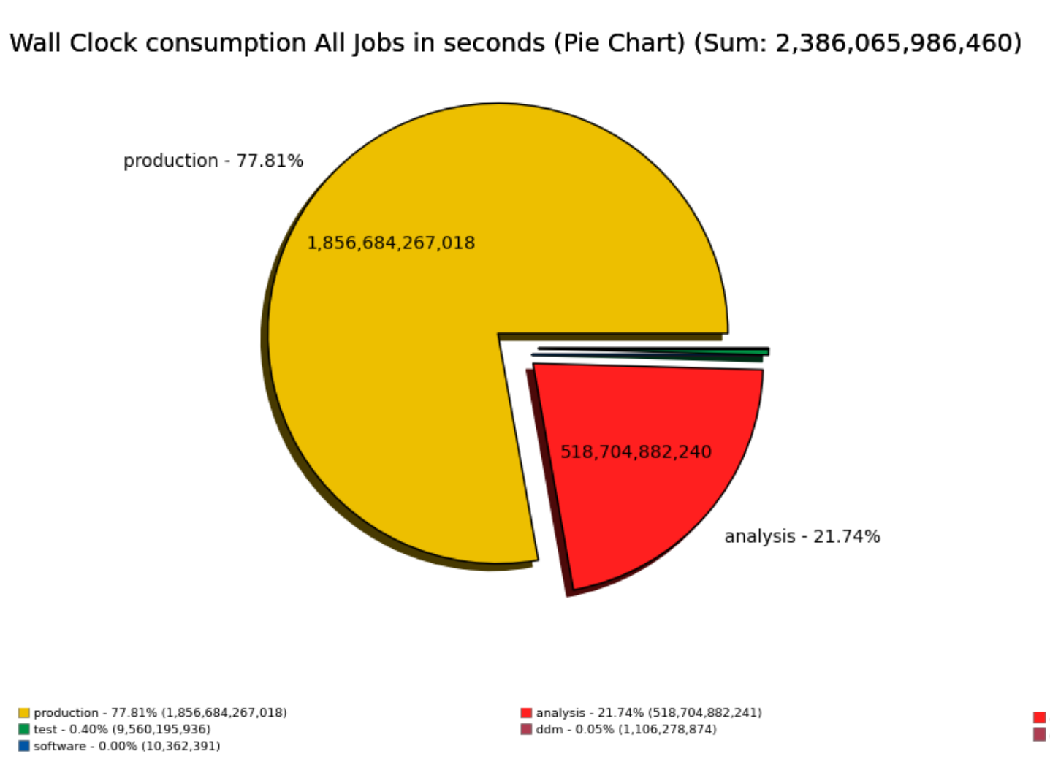

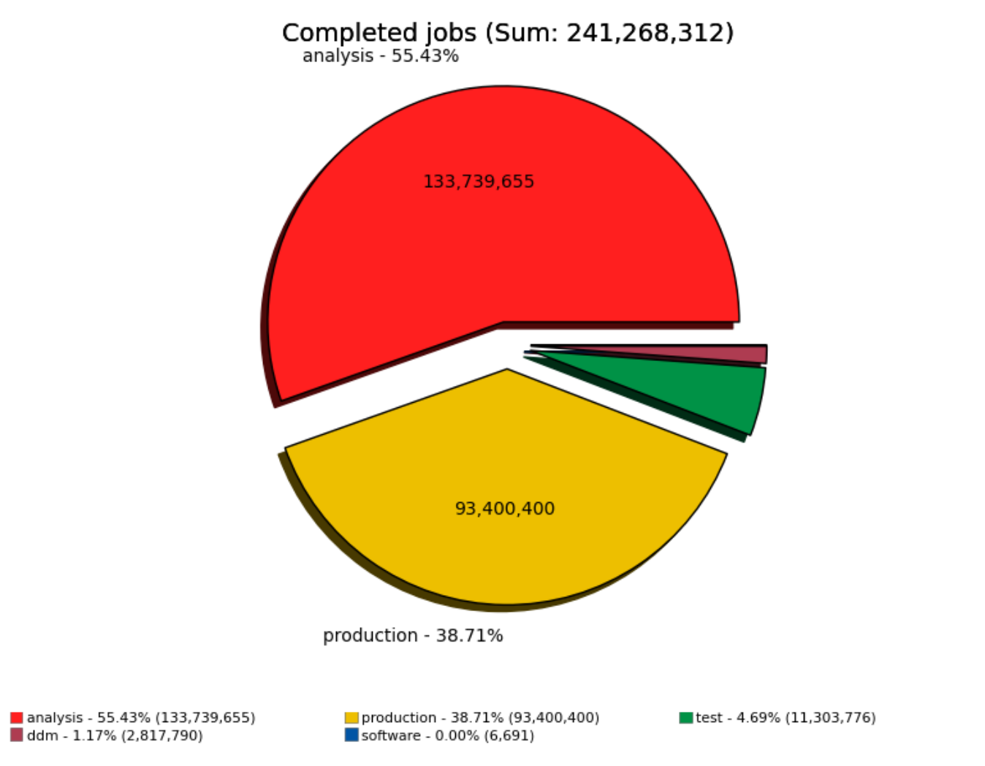

4.Analysis efficiency

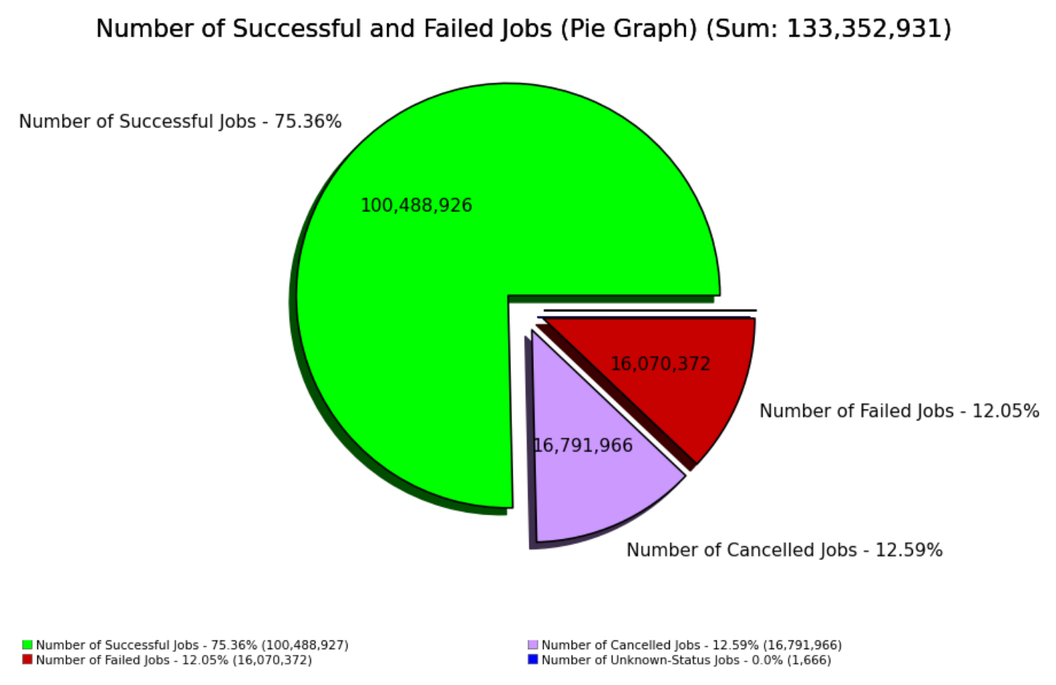

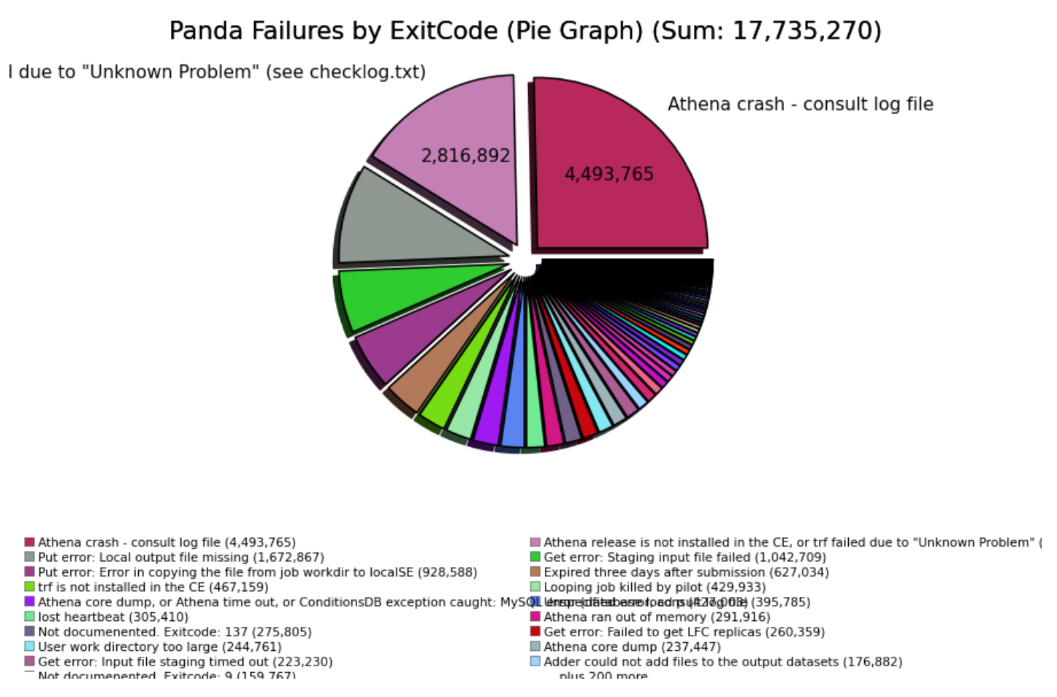

April 2011-April 2012
Job errors and re-tries:

- about 12% job failures
- All errors categorized into re-triable and non-re-triable errors.
- In this first version of the categorization 20% of failures re-triable, i.e. jobs are automatically rescheduled and 66% succeed after 2nd attempt.

Errors categories:

- temporary site problems like SE input or output problems
- Conditions data access problem

HammerCloud site exclusion:

- Sites are tested with automatic test jobs
- Automatic site black- and white-listing
- Started in June 2011 and improved overall job efficiency

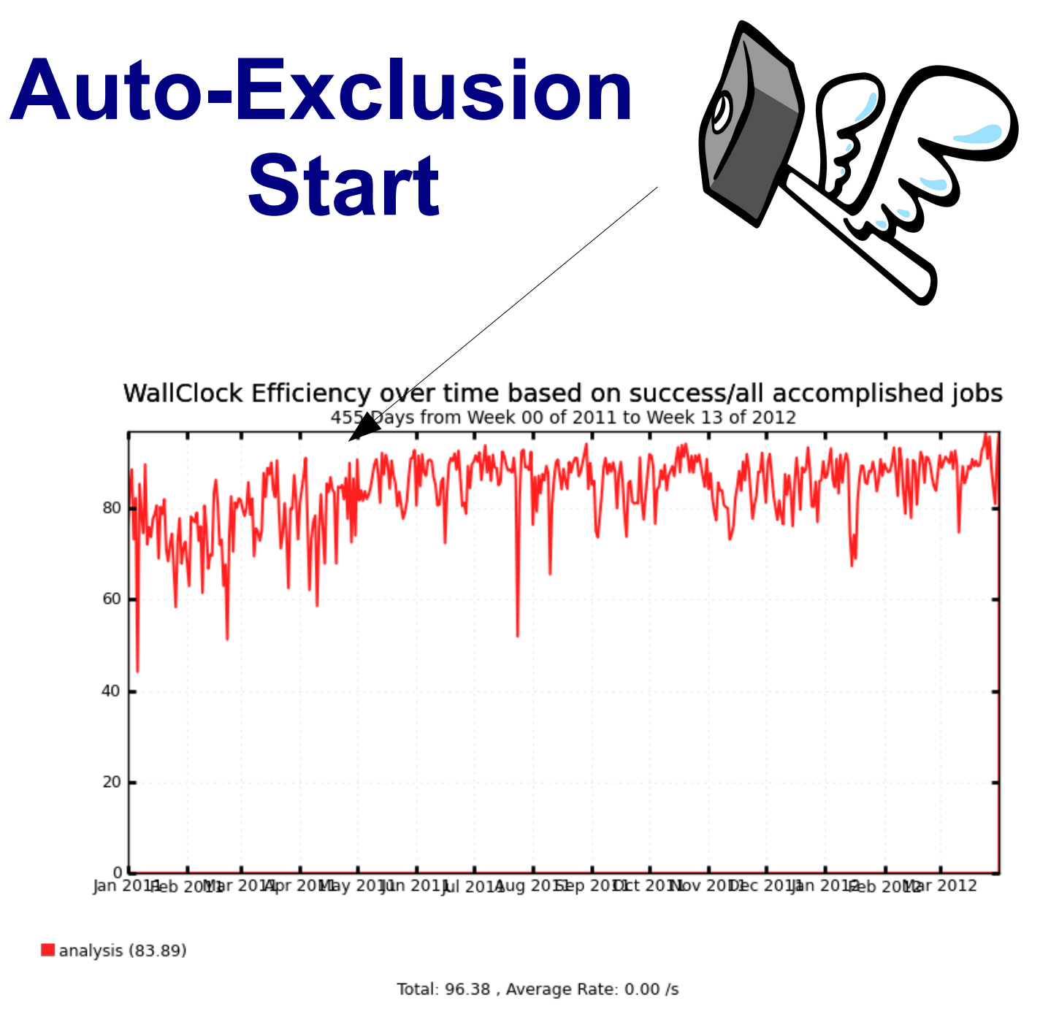

Analysis formats:

- Athena based analysis (AOD/ESD/RAW) is approx. 1/3
- ROOT based analysis: NTUP in various flavours and user ntuples is approx. 2/3

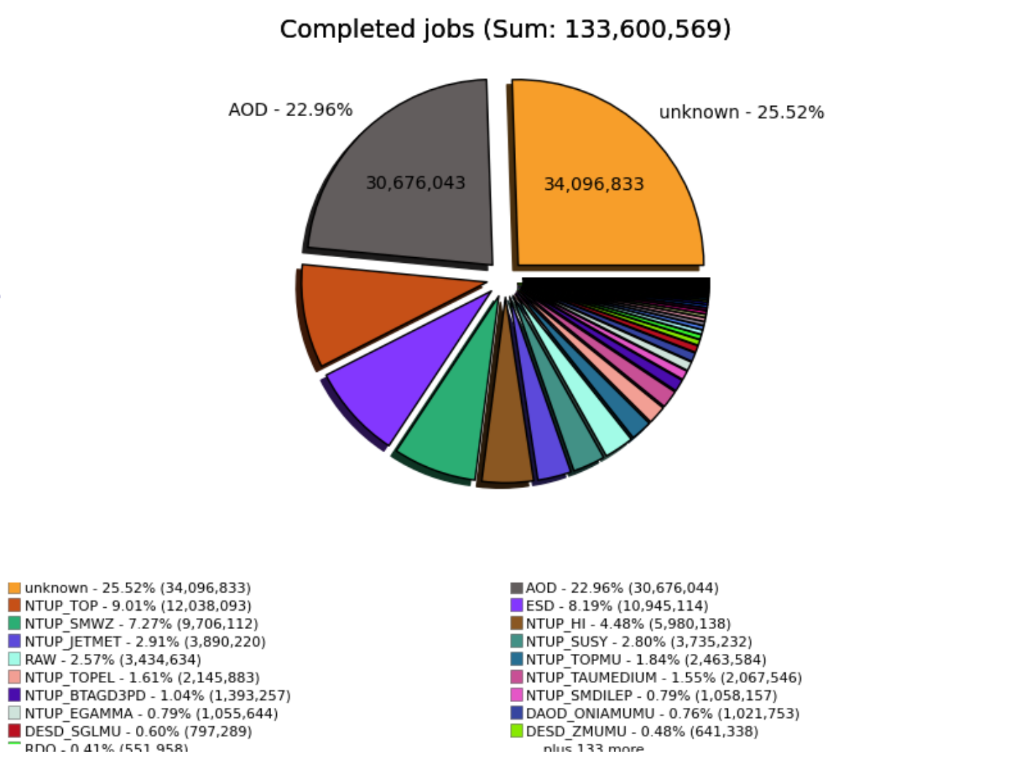

Job rescheduling PD2P:

- If jobs are waiting too long at a site, datasets are replicated and rebrokered to another more empty site

5. Future plans

As shown in the previous section the overall job efficiency has improved with the introduction of the HammerCloud system. Continuous monitoring of sites and the central infrastructure with new test cases will help to increase the efficiency even more. Furthermore increasing the number of re-triable errors will also lead to higher efficiencies. For this the error categorization and identification in the various tools needs to be improved.

Slow job submission speed is sometimes reported by users. This can be improved by moving parts of the job splitting and input data detection algorithms from the client tools to the Panda server.

More and more analysis is done using ROOT with ROOT n-tuple input. Improved job error detection and input file handling for this work-flow is planned and in the works by dedicated client tools or job wrapper on the grid.

In conclusion, the distributed analysis system of the ATLAS experiment is processing 300000-400000 jobs per day and is working well. Steady improvements in the tools have been and will be made to improve the user experience and job efficiency in this the very complex world wide system.

Poster presented at the 19th International Conference on Computing in High Energy and Nuclear Physics (CHEP), New York, United States of America, May 21-25, 2012.