

Atlas Report 27.01.2015

- Rucio:
 - Lost file recovery is now working
 - Loss at SARA of 500k files last week was good stress test
 - Reminder: please create JIRA tickets for lost files for the next few weeks
 - Developers are working on monitoring and an improved permissions model
 - For now the “Recovery” activity can be monitored on the DDM dashboard
- Production :
 - lack of running jobs due to small issues in PANDA and JEDI
 - some issue with misconfigured jobs accessing ATLAS web instead of CVMFS
- Main outstanding issues in production (from GDP ops)
 - PRODSYS-351: The dominant error rate in the last days is due to the submission of large simulation samples where the secondary minbias pileup or cavern background inputs could not be reused. This happens only to big tasks, each job uses few secondary inputs but when they are all used, the next jobs should start reusing the inputs.
 - **Distribution of secondary inputs:** The tasks cannot be assigned to the clouds because the primary and the secondary input datasets are not located in the same cloud. This need manual intervention to create rules ad-hoc.
 - **Data file loss recoveries:** Atlas can now recover missing outputs resubmitting the corresponding jobs from finished or completed tasks
 - **Destination full:** some group tasks are defined with a specific destination endpoint. Is this endpoint gets full then all tasks fail including all retries which impacts the production.
 - **Tasks cannot be splitted between clouds:** For large productions the input datasets from the container could be in several clouds. In Prodsys1 the request created several tasks which were assigned to several clouds. Now this is not always possible, concentrating the execution large task to one cloud causing troubles to move the inputs to one cloud and spreading jobs with MCP.

Atlas Report 27.01.2015

- From WLCG 22.01.2015
 - Prodsys-2 has been fully validated, took several weeks to fully understand the comparison of the physics distributions from Prodsys-1 and Prodsys-2 datasets
 - Rucio is fairly stable, although monitoring is still lacking some information, as the data-loss and data-recovery information
 - For the last two weeks, the production and analysis fully use the grid resources, although production has some hiccups occasionally (lack of tasks, APF failure). Most of the production runs multicore. Analysis is using 50% of the resources.
 - data loss at SARA: 0.5M files were lost due to raid failure. The recovery procedure in Rucio is working well and fast, but the relevant information on the files/datasets removed from the catalogs needs to be obtained from the Rucio log files for now. The report on the physics projects affected is being prepared.
 - Data recovery in Prodsys-2/JEDI will be tested on the affected tasks in the following few days, and the plan for automatic recovery will be defined after.
 - multicore queues deployment on sites is being followed in jira ADCSUPPORT-4117
 - the data lifetime policy has been applied on both T1 and T2 sites, the order of 3PB of data has been secundarized
 - FTS issues: staging on castor did not work for all the files, callback to Rucio were missing, cancellation of requests was not working properly. All fixed in the latest release being deployed this week.
 - MC15 simulation still not ready, schedule not clear yet. The MC14 tasks are not enough to fill the grid, so we need to wait for the big campaign before the production will use all the resources.



Status of ASAP

600 Hours from 2015-01-01 00:00 to 2015-01-26 00:00



