

GDP Introduction & Communications with Shifters

Andreu Pacheco Pagés,

v07



ATLAS Facilities and Shifters Jamboree – 5th December 2014

What is GDP?

- GDP was the acronym for Grid Data Processing
- Now the name has changed to DP (Distributed Production) to move away from the idea that all production in Atlas is done only on the Grid.
- Also DP combines with the DA (Distributed Analysis) activity.
- I personally prefer to keep the name informally and to keep the GDP for historical reasons on this talk.

What are the primary goals for GDP?

- To monitor and to keep the production activity running optimally in the system.
- To check for issues affecting a portion or whole of the production.
- To ensure that critical requests of the production managers (MC, group & derivation, data and HLT reprocessing) are followed and issues addressed to the experts.
- To check the balance between resources available and requests submitted.

What are the primary activities for GDP?

- Checking the status of the production system and task errors from the monitoring.
- Digging into tasks, jobs and log files to understand where the problem is.
- Increasing manually job limits and retry submission limits, moving jobs between sites, priorities.
- Killing jobs not being finished for too long before timeouts to speed up task completion or blocking system.
- Moving data between sites to avoid deadlocks.

What are the primary activities for GDP? (2)

- Assisting the ADC Ops coordination on daily issues (9am daily meeting)
- Opening tickets to developers for any issue that needs improvement and/or automation.
- Keeping a direct communication channel open with production managers.
- Validating new resources being added to the system. New MCORE, high memory queues.

Who participates to the Distributed Production activity?

- Distributed Production is a collaborative work between many people.
 - Experts: Rod, Andrej, Sasha, Alastair, Kaushik, Tadashi, Paul, Peter, Ales...
 - MC Production, Reprocessing and Group Production who handles the requests and submits the tasks.
 - Central operations people who monitor, handles the changes and track the ongoing production issues.
 - Shifters, Site and cloud support teams who report, diagnose and fix the detected problems at the sites.

What are the issues of Distributed Production?

- Why there are not enough load on the grid?
- Which critical tasks need manual intervention, like more than 4GB of memory or reprocessing jobs looping indefinitely due to a bug.
- Do we have saturation of the capacity of transferring the outputs from the grid to the Tier1s?
- Are there inconsistencies in the configuration of the resources in AGIS due to changes in the sites, new resources or new requirements?

Tracking of production tasks

- Production Managers (MC, Reprocessing, Group) want some tracking of their tasks. This means that if tasks fail for operational issues, they want them to be fixed without their intervention.
- One of the discussions is about “who tracks the tasks that are being submitted to the production system”.
- In principle the physics group who requested the sample should track the progress in addition to who submitted.
- Production managers in collaboration with GDP will provide a priority list of tasks to shifters that are important to be completed. This list should be limited.

Tracking of production tasks (2)

- I am not completely sure that the distributed production activity can track efficiently in real time issues from running jobs in the system. The production system was running 132 kJobs simultaneously as daily average.
- Usually GDP track issues from tasks reported explicitly by production managers, operations, shifters or site managers.
- Problematic tasks can be usually caught by requesting the summary of failed job errors in the last 12 hours:
[http://bigpanda.cern.ch/errors/?
jobtype=production&sortby=count](http://bigpanda.cern.ch/errors/?jobtype=production&sortby=count)

GDP Tools

- The main tool for GDP is BigPanda as the shifters use and in general all tools available from <http://adc-monitoring.cern.ch/>
- The second tool for GDP is the DEFT interface where the production requests are submitted. We use: <https://prodtask-dev.cern.ch/>
- GDP is still using direct SQL database access for some activities, Panda scripts and the API interface to JEDI. In the understanding that after some time the views or actions could be implemented in BigPanda and DEFT interfaces.

Communication with shifters

- Shifters play a central role in the detection of issues related to distributed production, but their primary goal is to look at the evolution of crashes per site.
- There is a day-to-day procedure which is kept updated to detect site, storage, network or task execution issues.
- GDP and Shifters share the Jira tracker <https://its.cern.ch/jira/browse/ADCSUPPORT> and there is a special GDP account to assign tickets and can use ELOG.
- Communication with GDP support is available by the mailing list: atlas-adc-gdp@cern.ch

When to contact GDP?

- When detecting that a task have jobs failing for temporary reasons (site, infrastructure) and there is a reason to increase the maximum number of job retries of the failed jobs.
- When a task looks stuck. Active for more than 3 months for example or just not having any finished job for a week. All this indicates a pathological problem somewhere.
- When there is a possible solution for an identified task issue. Shifter can detect that jobs fail by CPU or memory issues and GDP could increase the limits, move the task, move data, increase priority, etc... - **Failing jobs block the system** -

What are the challenges in the next months?

- Getting experienced with the operation of ProdSys and Rucio to regain the production level reached with Prodsys1 and DQ2..
- The new data reduction framework with trains: 52 derivations.
- Early identification of issues related to configuration conflicts, task definitions and resource availability.
- Integration of Analysis and Production.
- Efficient tracking of the production tasks.
- ₁₃ Efficient tracking of the production activity.

Backup slides