

ATLAS SUMMARY

Johannes Elmsheuser
for the GangaAtlas team

Ludwig-Maximilians-Universität München


8 Feb 2012
Ganga Developer Days



ALL ATLAS JOBS, OCT 2011-JAN 2012


Completed jobs

122 Days from Week 39 of 2011 to Week 05 of 2012

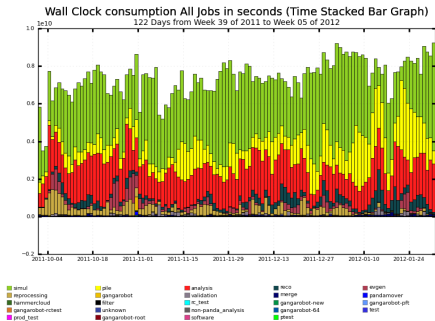


Running jobs


122 Days from Week 39 of 2011 to Week 05 of 2012



ALL ATLAS JOBS WALLCLOCK, OCT 2011-JAN 2012




Wall Clock consumption Good Jobs in seconds (Pie Chart) (Sum: 829,608,901,214)



DA JOBS EFFICIENCY, OCT 2011-JAN 2012

Number of Successful and Failed Jobs (Time Stacked Bar Graph)

121 Days from Week 39 of 2011 to Week 05 of 2012




■ Number of Successful Jobs ■ Number of Cancelled Jobs ■ Number of Failed Jobs ■ Number of Unknown-Status Jobs

Maximum: 942,405 , Minimum: 113,764 , Average: 440,793 , Current: 244,321

Number of Successful and Failed Jobs (Pie Graph) (Sum: 53,776,763)

Number of Successful Jobs - 73.73%



Number of Failed Jobs


Number of Cancelled Jobs - 14.08%

■ Number of Successful jobs - 73.73% (39,647,535)
■ Number of Failed Jobs - 12.2% (6,559,095)

■ Number of Cancelled Jobs - 14.08% (7,570,133)
■ Number of Unknown-Status jobs - 0.0% (0.00)

DA JOBS SITE AND FAILURE CODES: OCT 2011-JAN 2012


Completed jobs (Sum: 53,776,763)



BNL-ATLAS - 8.57% (4,606,863)
 AGLT2 - 3.89% (2,090,425)
 UKI-LT2-QMUL - 3.14% (1,689,063)
 RAL-LCG2 - 2.81% (1,510,176)
 FZK-LCG2 - 2.42% (1,302,108)
 DESY-HH - 2.31% (1,240,908)
 IN2P3-CC - 2.24% (1,205,267)
 DESY-ZN - 2.09% (1,121,994)
 UKI-SCOTGRID-GLASGOW - 2.00% (1,074,955)
 INFN-NAFNI-LBTI-AS - 1.87% (1,003,700)

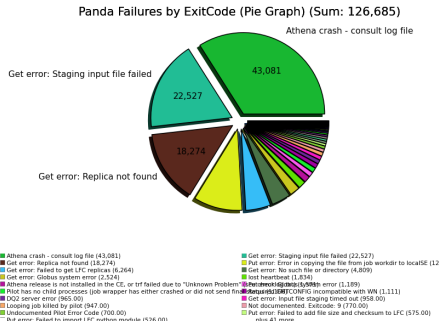
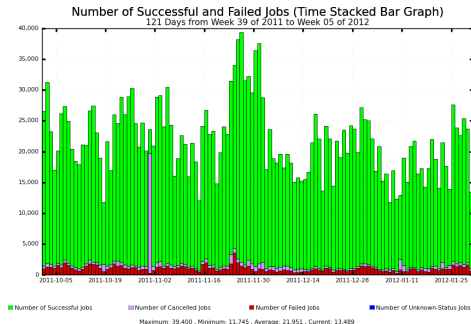
MWTF2-UC - 4.85% (2,606,930)
 CERN-PROD - 3.60% (1,933,306)
 IN2P3-CC-T2 - 3.11% (1,670,363)
 WIT2 - 2.48% (1,336,135)
 TOKYO-LCG2 - 2.33% (1,250,358)
 INFN-T1 - 2.25% (1,210,387)
 NIKHEF-ELPROD - 2.20% (1,182,131)
 UNI-FREIBURG - 2.07% (1,111,690)
 SARA-MATRIX - 1.98% (1,063,001)
 Other 80 sites

Panda Failures by ExitCode (Pie Graph) (Sum: 7,172,318)
 failed due to "Unknown Problem" (see checklog.txt)




Athena crash - consult log file (1,430,395)
 Put error: Local output file missing (1,426,566)
 Put error: Error in copying the file from job workdir to localSE (426,635)
 Trf is not installed in the CE (252,398)
 Looping job killed by pilot (170,699)
 Get error: input file staging timed out (149,936)
 Athena ran out of memory (118,659)
 Not documented. Exitcode: 9 (93,276)
 Adder could not add files to the output datasets (66,445)
 Athena release is not installed in the CE, or trf failed due to "Unknown Problem" (436,216)
 Get error: Staging input file failed (348,052)
 Athena core dump, or Athena time out, or ConditionsDB exception (365,790)
 Get error: Failed to get LFC replicas (165,790)
 Lost heartbeat (142,999)
 Athena core dump (113,659)
 Not documented. Exitcode: 137 (91,908)
 Not documented. Exitcode: 139 (66,283)

GANGAROBOT AFT: OCT 2011-JAN 2012



GANGA AND ATLAS DA



- 60-80 GangaAtlas users per week
- Strong competitor: command line tools: prun/pathena/pbook for Panda submission
- In total: ≈ 700 DA users per week
- Also: non negligible analysis share off-grid

BIGGEST GANGA(ATLAS) CUSTOMER: HAMMERCLOUD

<http://hammercloud.cern.ch/>

<https://twiki.cern.ch/twiki/bin/viewauth/IT/HammerCloud>

- ATLAS sites are tested with analysis functional test jobs (AFT) and production functional test jobs (PFT): 3-6 jobs per hour per site
- Job driver: Ganga sessions based on GangaRobot
- Job results are stored in MySQL DB and processed and displayed using Django
- Queues are black-listed/white-listed depending on the job efficiency
- ATLAS cloud squad notified by email about failures, results pushed to Site Status Board and Efficiencies used in ADC (ATLAS Distributed Computing)
- Integral part of ADC distributed analysis
- Stress test on demand to test new sites configurations
- Similar setup for CMS and LHCb

FURTHER AUTOMATIC SERVICES

- Ganga is the backbone of the TAG skimming service
- TAGs are small event summary files which can be used for fast pre-selection
- TAG skimming service is used e.g. for RAW data skimming
- Positive feedback about MultiTasks package

RECENT GANGAATLAS/GANGAPANDA UPDATES

Major updates in 5.6.x and 5.7.x:

- Integration of prepared state: user code tarball
- Production System Transformation support in GangaPanda for HammerCloud PFT
- Support for Panda rebrokerage added
- Support for RootCore (Build System for ATLAS ROOT analysis)
- Minor Tasks/MultiTasks fixes
- Minor GangaPanda and scripts/athena fixes
- 64bit job support
- GangaPanda merger step support
- Various TAG support updates

General Problem(s):

- Still perception that there are 2 systems in ATLAS: Ganga vs. Panda, which is wrong: Ganga is a configuration/monitoring job tool, Panda is a workload management tool
- Ganga usage (IPython) seems to be too complicated ?
- But command line submission is pretty similar to prun/pathena

Code change suggestions: (in no particular order)

- User code tarball storage location:
 - Athena.prepare() creates tarball and excludes large files/root files
 - Prepared state copies tarfile to shared dir
 - AthenaPandaRTHandler adds missing files from inputsandbox config
 - Only last file is uploaded, but 2 'incomplete' tarball are stored
- Job submission speed optimization
- ATLAS DDM is undergoing major change in a couple of month
DQ2 → Rucio
 - Compatibility API has been promised, unclear what code changes will be necessary
 - DQ2Dataset, DQ2OutputDataset, DQ2JobSplitter changes
 - All shared worker node code in Local/Batch/LCG/Cream needs to be adapted

Code change suggestions: (in no particular order)

- Common WN code for LCG/Batch/Local with Panda Pilot (runAthena, runGen, etc...)
- Improve ROOT/ROOTCore support on GangaPanda
- Can Tasks be made even easier ? Command line support ?
- Proper CVMFS cond DB support in Local/Batch/LCG/Cream
- Command line submission scripts/athena improvements
- Concentrate on GangaPanda usability for ROOT, Athena and easy Local/Batch workflows
- AthenaMC, AtlasLocalDataset, TAG, EventPicking, AMIDataset ?
- GangaCore I/O changes should offer some new features
- Write stable Tests !