



Data Reprocessing: Lessons learned

-

Pavel Nevski



Reprocessing Highlights

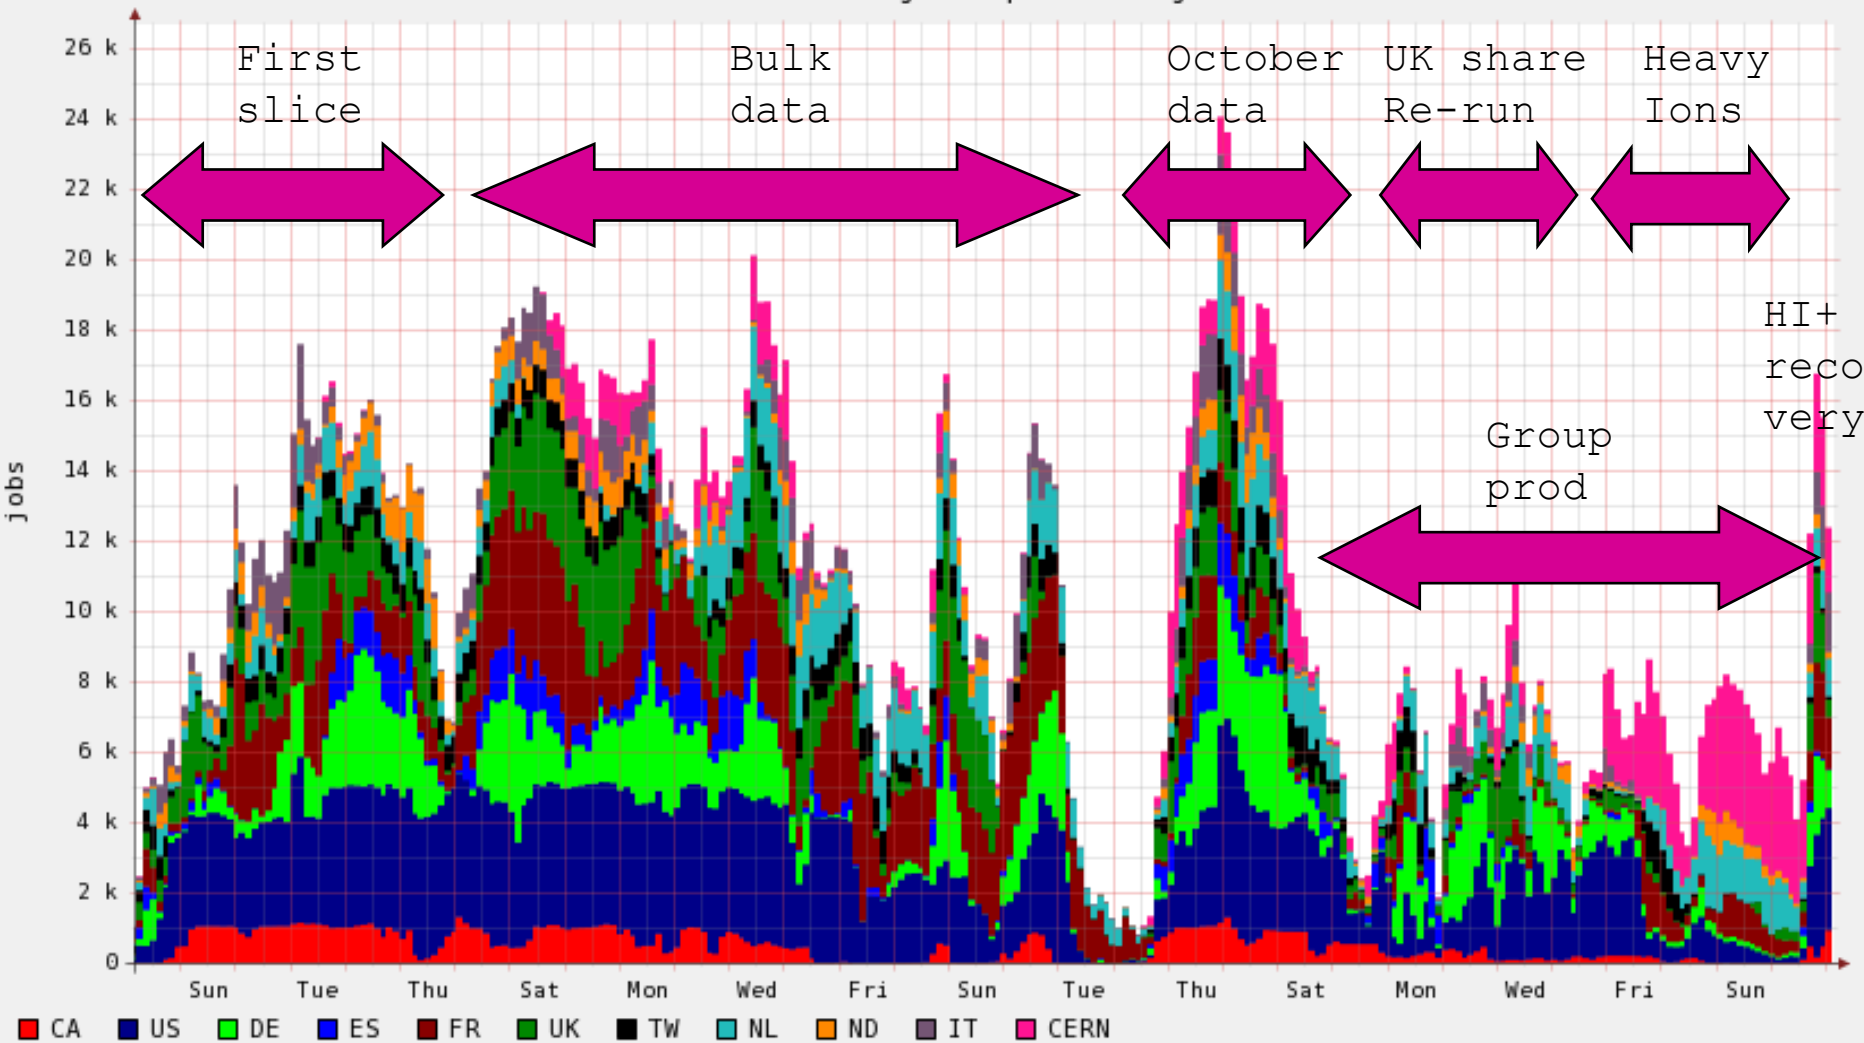
- Goal: to reconstruct with an improved software and calibration all data collected by the ATLAS detector since the LHC start-up in 2009
 - “Full reprocessing”: specially prepared Software release
 - software release not tested by Tier 0
- Several different component were included
 - Express stream processing, delivering results for beam spot determination and Data Quality assessment
 - Reprocessing of Derived RAW data produced at Tier0
 - First pass Heavy Ion data processing since last week
- Unprecedented scale - 1.2 M jobs altogether
 - High quality - 46 jobs aborted in main campaign, all recovered
- Running simultaneously with full MC reprocessing and other activity
- Including MC, up to 10k active tasks simultaneously, 30K tasks total

Main Reprocessing Cycles



World Wide - running - reprocessing - month

PROTOCOL / TOBI OETIKER



Range from Sat Oct 30 02:24:00 2010 UTC to Tue Nov 30 03:12:00 2010 UTC
 Generated by TRIUMF-LCG2 (times in UTC)

Progress in campaign organization

- Main task submission done by Jonas S. using bulk request option
 - Better campaign planning and more flexibility
- All merging steps are generated using production scripts, tested in several past reprocessing campaigns.
- Conditions DB release preparation is improved
 - Incremental DB release is possible
- Automate job abortion limited to know patterns (TP converter...)
- A procedure to move tasks from a busy T1 to other Tiers, developed and tested for reprocessing,
- Reprocessing status is exposed on-line on the web pages.
- Simultaneous Ntuple production for several physics working groups
 - Helps to exercise same reprocessing machinery in advance
- BCT is now running some reprocessing task
 - useful for early problem discovery

Sites Validation

- Site validation started as soon as a candidate release was available
- Site validation was run before and after express stream processing
- Several sources of differences have been found
 - Differences due to software instability (Muons, Inner tracker)
- Remains an issue, at least in a long term (AMD vs Intel, but not only)
- The difference excludes an automation of comparison procedure
 - Some human analysis is need, manual operation
- Site validation allows to get more uniform batch parameter setting
- Although not perfect, but very useful
- Also use to test Conditions DB release consistency



Some observations

- Site readiness improved due to a series of site validation tests.
 - More uniform memory requirements were implemented
- No more "Missing input" problem leading to data losses
- Tape pre-staging test was run on the eve of the reprocessing
 - Few sites fully used this opportunity to tune the system
- However, tape staging is still a problem sometimes
 - File status does not reach the DDM in some cases
 - Manual intervention is needed
- Number of attempts per successful job is increasing again:

2008	2009	spring 2010	summer 2010	Fall 2010
1.8	1.1	1.34	1.4	1.2

 - most of this is probably due to the Panda server crash



Software issues

Reprocessing software need a special approach,
several problems are in particular boring

- TRF error handling and reporting
 - Still have to dig in log files, time consuming
- Attempt to always produce a stack trace is a pitfall
 - Stack trace often breaks after i.e. a memory problem and the job goes into an infinite loop, no stack trace anyway
 - As the reason is unclear we are repeating such job many times
- Looping Reconstruction Software
 - 23 jobs were aborted on the first two cycles
 - No permanently looping jobs in the October cycle



Dataflow issue

- In this reprocessing campaign we were running a 3 step dataflow:
 - RAW→ESD,Hist/Ntuples
 - Several ESD→AOD,DESD,Ntuples
 - Merging/super-merging→ AOD,TAGs,DESD
- However, now the ESD produced in the first step are 1-4 GB:
 - ESD merging gives a factor less then 2
 - Running one more step introduces just an extra delay
- We should switch to the Tier0 dataflow now on
 - Load on LFC will not increase much



Conditions DB access crisis

- Bulk of the conditions data are supplied as a DB release
 - SQLite replicas for most of the database-resident data
 - DB release is structured (SQLite file per run) to minimize data copied by a job to the worker node
 - DB release build is automated and includes some validation
 - Proved to be a very successful and robust solution
- An algorithm (Magnetic Field service) was known to accessed the Oracle directly, bypassing interval of validity service
- While reprocessing October data we discovered that some new algorithms appeared to use Oracle directly for the luminosity information
 - The connection remained open after the data were read
- Debug stream contains many luminosity blocks per file, thus saturating server capacity



Site issues

- A key element for a timely reprocessing is the prompt diagnostics and a fast reaction to any site issue.
- We made a good progress in the past in improving communications between central team and the sites.
- Once a problem is found sites are working aggressively to correct the issue.
- However, the length of the reprocessing campaign still strongly depends on sites performance.
- Sometime a lengthy downtime (even scheduled !) creates an extra problem where it can be avoided. Please try to really minimize downtimes during reprocessing campaigns



Conclusions

- Reprocessing infrastructure developed over the past years was exercised for the physics data reprocessing on a new scale
- Identified weaknesses (error reporting, infinite jobs, conditions data access, staging blocking) should be addressed more carefully
- Overall the reprocessing campaign was very successful and delivered results on time