


Grid Data Processing

Andrej Filipčič - Jožef Stefan Institute
Rod Walker - LMU


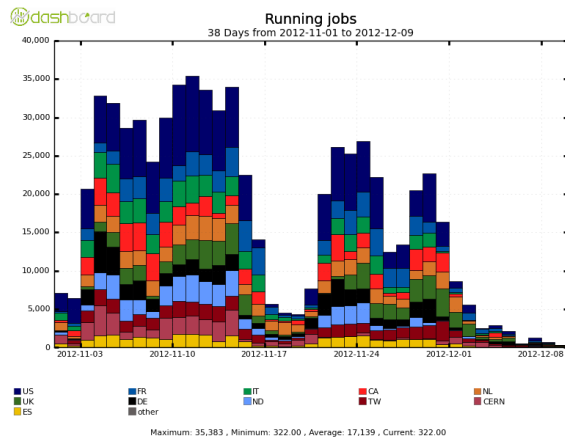
Reprocessing report

- Started on 2012.11.2
- Not quite finished yet, but most on 2012.12.1
- Inputs: mostly from disk, some from tape
- Remaining jobs:
 - ➔ Athena failures
 - ➔ Site failures (lost files)

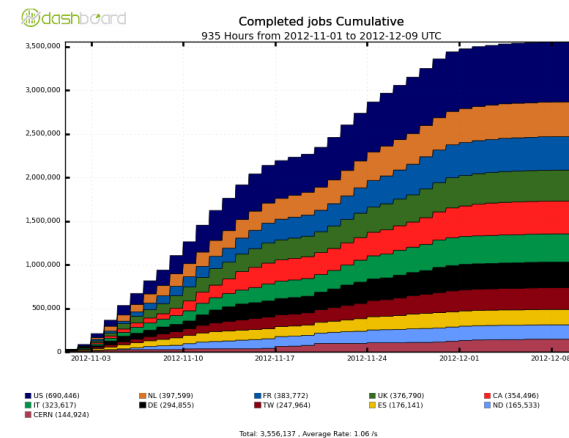
Jobs per processing cloud



Reprocessing statistics



Processing cloud



Destination cloud

- 30k simultaneous jobs
- Second part lower due to lower no of running jobs on T1s
- Multi-cloud processing extensively used – many sites ran foreign jobs
- Few T2Ds added to speed up reprocessing – not all to let MC pile-up and group production progress
- Job types:
 - ➔ 2.7M Reconstruction
 - ➔ 0.8M merging....

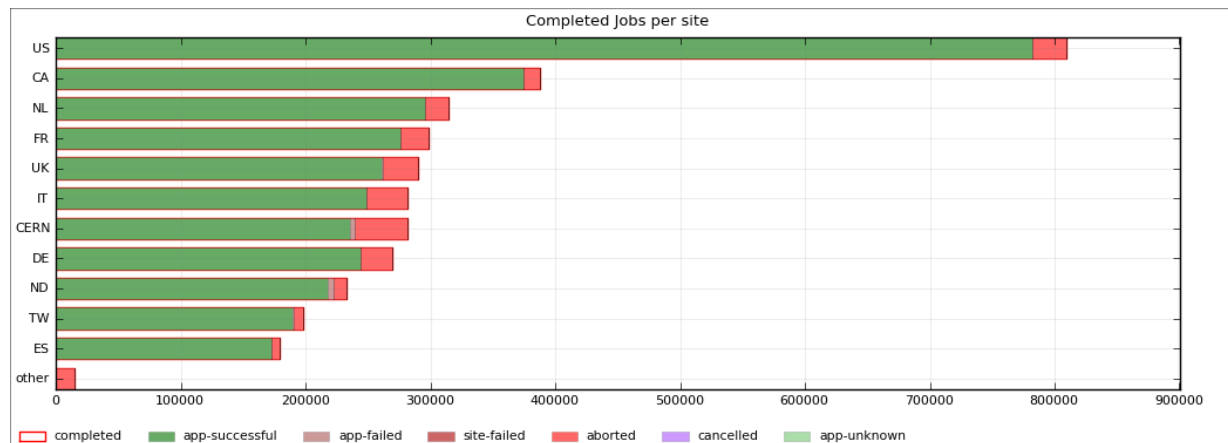
Problems

- Tier-1 downtimes


- ➔ Power outage at RAL
- ➔ Tape performance issues at FZK

- Full DATADISK at several sites

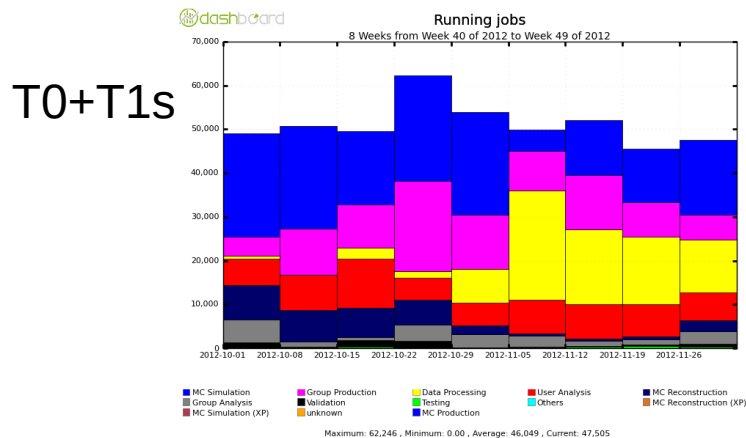
- ➔ Need to balance the number of activated reprocessing tasks
- ➔ Finish with period and cleanup temporary datasets before the next period was activated



Reprocessing and resources



- Reprocessing mostly took MC simulation slots at T1s
- More Group production done T2s in December
- MC pile-up was reduced on T1s but also on T2s



PRODDISK issues

- PRODDISK used for
 - On T1: Multi-cloud processing
 - On T2/T3: for all production jobs
- Fat jobs (many inputs, large outputs):
 - 35-50TB on large sites → significant fraction of disk space
 - “permanent” space reservation → smaller DATADISK
 - Wildly varying usage pattern: from 10 to 100%

PRODDISK removal

- Goal: use DATADISK for temporary datasets
- Inputs datasets: migrated already on 5 T1s
 - ➔ Needs a change in panda server only
 - ➔ 90% of PRODDISK was used by inputs
 - ➔ PRODDISK space reservation lowered on 2 T1s and given to DATADISK
 - ➔ Continue with the rest of the T1s after reprocessing is complete
- Output datasets:
 - ➔ A change in the pilot code is needed
 - ➔ Testing in January
- After all T1s are migrated, continue with the rest of the sites and remove PRODDISK completely.

Downtime handling

- Two mechanisms:
 - ➔ Offline: a site does not receive new jobs
 - ➔ Brokeroff: a site (cloud) does not receive new tasks
- Problem with recovery:
 - ➔ Takes a long time (several hours) to run and complete the functional tests
 - ➔ When brokeroff, new tasks are not assigned, after downtime it takes a long time for new tasks to fully activate
- Changes:
 - ➔ Brokeroff/offline handling (see Jarka's slides) → Short downtimes should affect the production/analysis for much shorter period (check)

AGIS/Schedconfig

- AGIS will become the ultimate information system for ATLAS
 - a unique source for all the site settings, capabilities
 - All the services will migrate to AGIS as primary source of information/configuration
- Schedconfig:
 - Up to now, the site/queue settings stored in svn
 - AGIS and Schedconfig in sync
 - Schedconfig will use AGIS as the source of information
- svn access will be disabled → web api/interface for site settings/configuration (will be announced when ready)

Site parameters

● Memory: requirement still 2GB/core (physical)

- ➔Memory: 2GB, but sometimes more is used, up to 4GB vmem. Even the newest hardware rarely capable of more than 2GB/core.
- ➔Sites/queues are requested to publish maxmemory settings
- ➔If the sites can provide more memory for special tasks, this should be published
- ➔New mechanisms to limit physical memory (slurm/cgroups), potentially useful later, or ulimit through pilot (see Alessandras's slides)

● Walltime/cputime:

- ➔2 days, not enough for many tasks (long evgens, long pile or group production jobs)
- ➔Prodsys: scout job (10) walltime is used to set the walltime of the bulk
- ➔Long jobs only assigned to sites with maximum walltime long enough
- ➔VO card: cputime 3120 min, walltime 5760
- ➔ATLAS jobs typically use up to 2 days of cputime, scaling in SPECs to as low as 1 day on some sites does not make sense

● Hospital queues: queues with 96h limit, used for jobs which do not fit default queue walltime

● Recommendation:

- ➔Set longer cputime/walltime for production jobs (2.2 days at least, 4 days if possible at least in T1s)
- ➔Keep it short for analysis jobs (2 days?)

● Local disk size:

- ➔Some jobs use much more than foreseen 20GB
- ➔Maxwdir (in addition to maxinputsize): maximum space available to a job (default site setting: maxinputsize + 2GB)
 - Some jobs need much more space for outputs than assumed default of 2GB

Conclusions

- Reprocessing ran smoothly in spite of some site outages
- The tape copies were mostly not used, the availability of those is unclear
- The ATLAS VO card needs a revision, mostly for job time limits
- Activities to reduce the spacetoken related limits on-going
- DATADISK is in crisis. Data placement strategy needs a revision.