ATLAS operation for Q4 2011

GDB

Cédric Serfon

Ludwig-Maximilians Universität, München

December 14th 2011

Contents

DDM activity during last 3 months

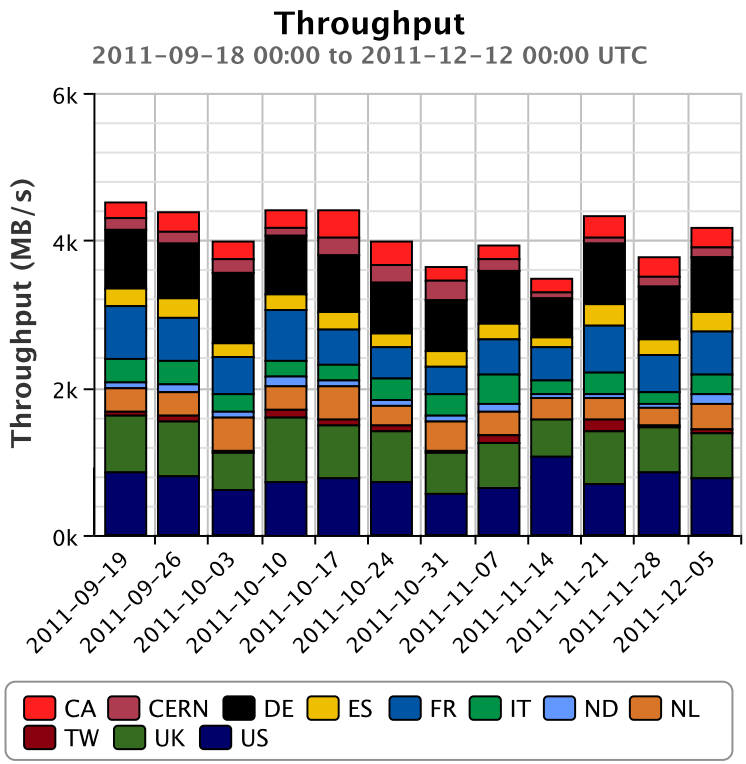

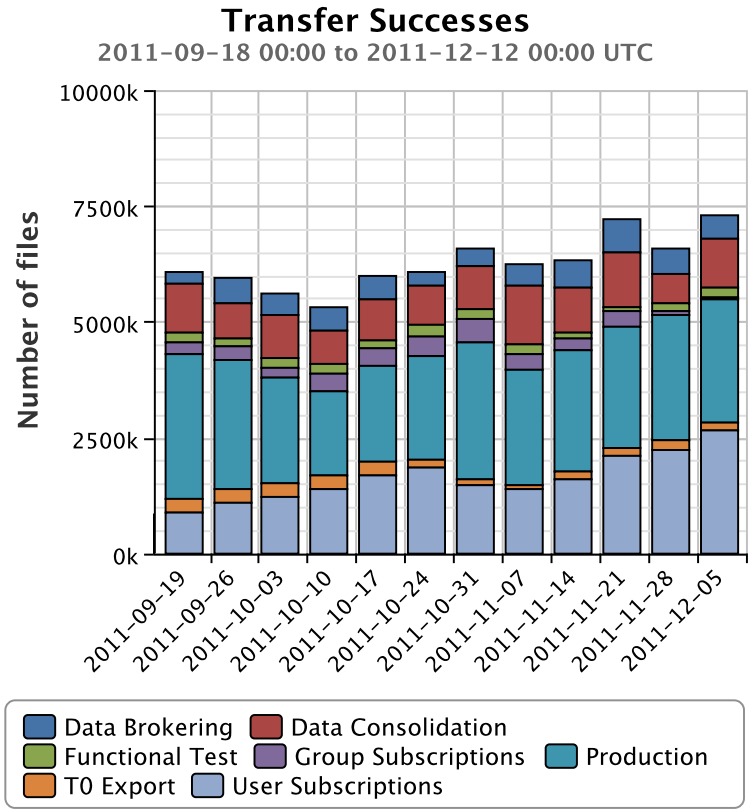

- Throughput constantly around 4 GB/s.
- Transfer rate (in term of number of files) ∼1M files/day, largely dominated by production and users subscriptions.

Tier0 export

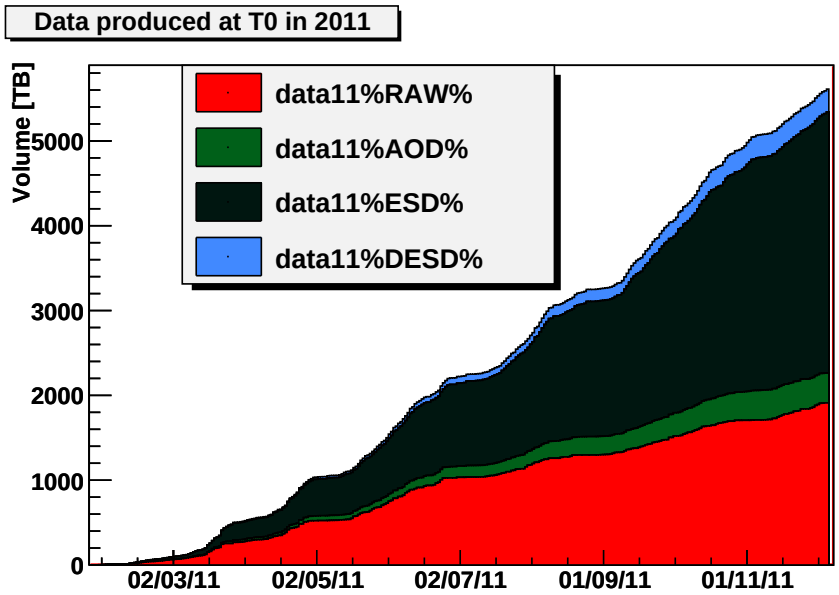

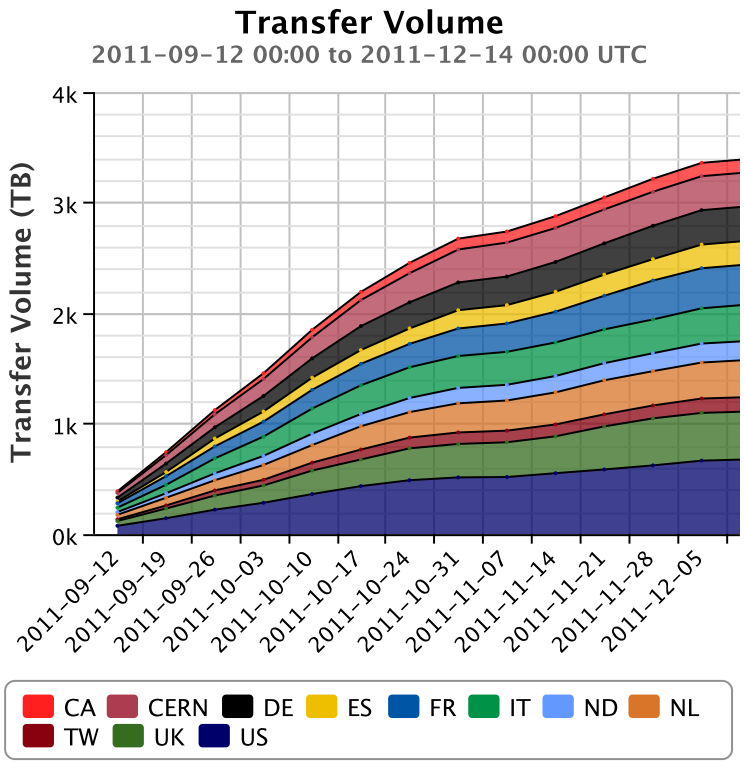

- About 3.5 PB exported from T0 to T1s in the last 3 months.
- Change of slope end of October in right plot due to end of p-p data taking.

Production during last 3 months

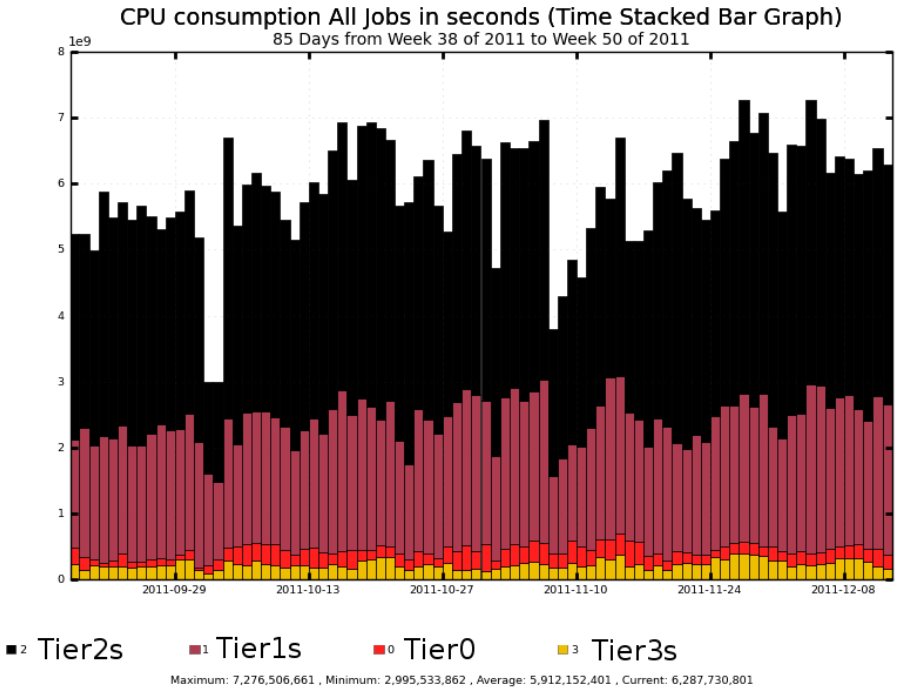

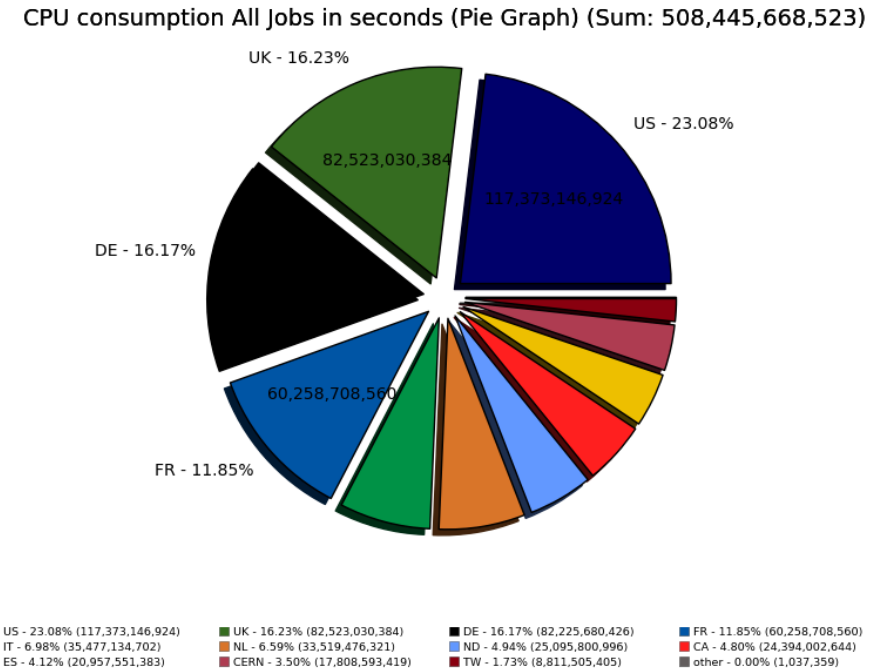

- 60% of the CPU comsumption on T2s.
- More than 50% of the CPU consumption comes from 3 clouds (US, UK, DE).

Production during last 3 months

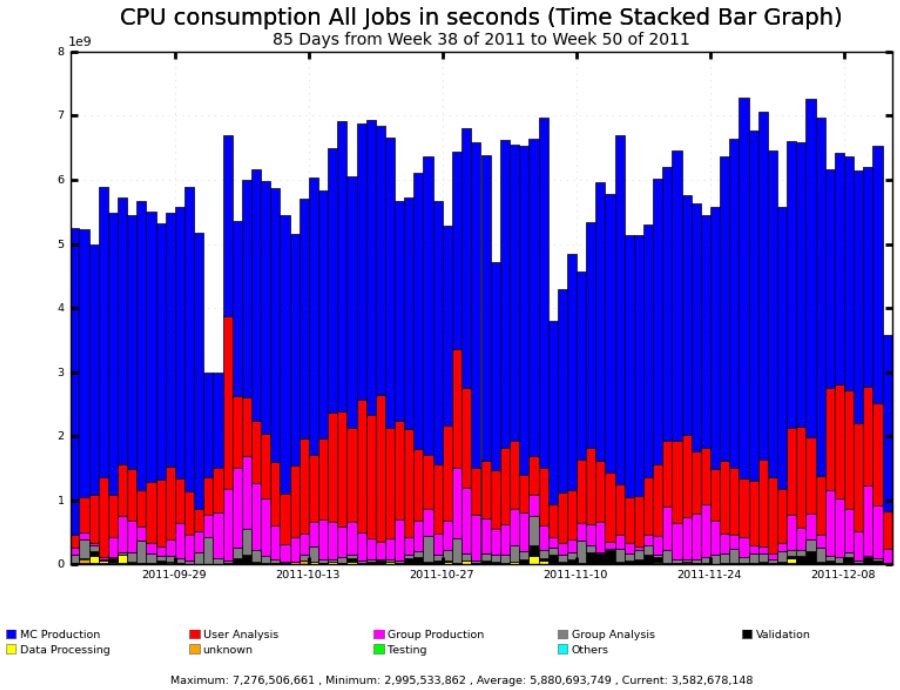

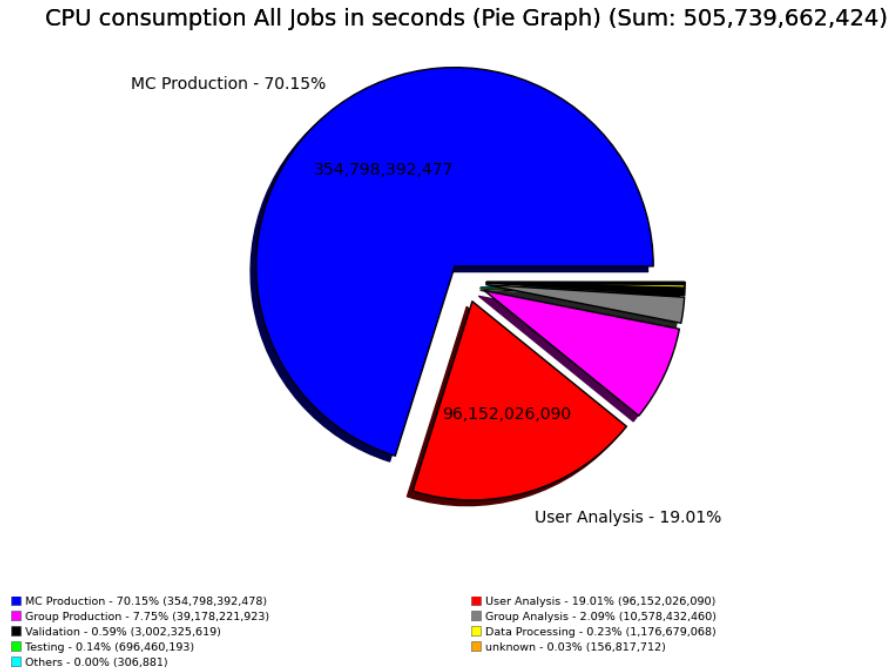

- MC production is still the biggest activity with 70%.
- Central production for groups starts to take a significant part of the resources (now ∼8%).

Reprocessing

pp reprocessing

- Part II (for data from Jul. 1-Aug. 24) of release 17 reprocessing ended on end of September.
- After the end of the reprocessing, large group production and export of data to T2s.

HI reprocessing

- Reprocessing of Heavy Ions runs from last year beginning of October
  - Input : Volume reprocessed $\sim$0.5 PB of ESD.
  - Output : $\sim$32 TB of NTUP_HI.

DB issues

- Multiple problems with databases (ATLR and ADCR).
- High load on ATLR (offline conditions DB) linked to T0 reco + HLT repro. Situation relaxed after 2 new nodes were added.
- High load on ADCR (ATLAS Distributed Computing DB) due to very high activity coming from DQ2 :
  - A lot of work was done by DQ2 devs and DBAs to optimize queries and clean-up tables. Situation looks better.
  - In a near future, hope that new hardware + 1 additional node will relax the situation even more.

Alarm tickets

- 18 Alarm tickets submitted during the last 3 months (excluding tests)
- Split by support unit :
  - 9 tickets for CERN : 77069, 77049, 76770 LFC issues. 77065, 76039 LSF issues. 74838, 74448 Castor issues. 76519 AFS problem. 75234 DB issue.
  - 3 tickets for CNAF : 76663 premature end of downtime. 75601 LFC issue. 74429 GPFS problem.
  - 3 tickets RAL : 75823, 75597, 74686 Castor issue.
  - 1 ticket for SARA : 76628 dCache issue.
  - 1 ticket for IN2P3-CC : 75609 SRM issue.
  - 1 ticket for ASGC : 74758 LFC issue.

Network congestion

- Since 6 months, we have started to break the clouds boundaries by setting up T2Ds (i.e. T2s qualified to get their data directy from T1s). Right now more than 50% of all ATLAS T2s are T2Ds.
- Start to see some problems :
  - 1Gb/s bandwidth of TRIUMF research network fully saturated. Will be increased.
  - Problem with transfers from US T2s to KIT (firewall on the general 10Gb/s IP uplink overloaded).
  - Degraded connectivity between IN2P3 and BEIJING/TOKYO and other foreign T2s under investigation.
- Other problems observed : degradation of transfers CERN→AGLT2 due to change of the path from CERN to AGLT, or PIC→T2 due to switch to LHCOne.

Other issues

- One problem with synchronization VOMRS→VOMS, fixed now.
- Different problems linked to CAs :
  - Change of UK CA ("CN=UK e-Science CA 2B" vs "CN=UK e-Science CA") generating a lot of problems.
  - Some users with cert from Israeli and Chinese CA's had problem to get files from some T2s using Bestman2 SRM (related to an "email" entry in the CA file itself). Fixed.

LFC consolidation

- Goal is to consolidate all LFCs at CERN (see previous ATLAS reports in GDB to know the reasons).
- Lots of tests realized to validate the procedures during last summer. End of September, the merging script developed by F. Furano validated by ATLAS → Ready to start the consolidation.
- Many people involved during the consolidation of one LFC : ATLAS central operation people, DBAs, cloud squads, T1s people, LFC devs, IT-PES (thanks to all of them).
- 4 LFCs merged in October/November :
  - W41 : SARA (NLT1) (∼16M replicas).
  - W44 : Taiwan (ASGC) (∼5M replicas).
  - W46 : GridKa (KIT) (∼34M replicas).
  - W48 : CNAF (∼16M replicas).

LFC consolidation

- Now ∼105M replicas registered in the central LFC.
- Some slowness starts to be observed (look-up, deletion, insertion). Need to evaluate if the consolidation can continue.
- Plan to go to central registration in Panda + IP authentication.
- If no scaling problem foreseen, continue the merges (2/month) after Oracle 11g migration. Still 14 LFCs to merge.

Castor to EOS migration

- Migration done in 2 steps for grid space :
  - First migration during summer of data hosted on ATLASDATADISK. Switch of catalogs on September 13th.
  - Second migration in October of data hosted on ATLASGROUPDISK. Switch of catalogs on November 28th.
- In total, ∼8M files migrated.
- In the meantime, migration of non-grid data (castort3).
- Some problems of SRM overload when transfers. Problem circumvented by using SRMless (gsiftp transfers) in FTS.
- In the near future, EOS should benefit of the new FTS version (2.2.8 under test in ATLAS) that supports checksum validation with gsiftp.

Data acceptance status

- Since September, ATLAS review every month the 'availability' in analysis of all T2s to decide where to send data.
- Status definition :
  - T2Ds with >90% availability labeled alpha.
  - T2s with >90% availability labeled bravo.
  - T2s with availability between 80% and 90% labeled charly.
  - T2s with <80% availability labeled delta.
- Number of datasets sent to the site depends on the status (e.g. alpha get 2.53% per site, bravo 1.27%...)
- This is in addition to Panda Dynamic Data Placement (PD2P).

LHCOne

- The ATLAS approach is to evaluate LHCONE in a managed way and gain expertise in liaison with Network providers.
- Agreement on a list of sites that we want to use for initial evaluation of LHCOne :
  - T1s : ASGC, BNL, CERN, PIC, SARA, TRIUMF.
  - T2s : AGLT2, (DESY-HH), LAL, (LRZ-LMU), Napoli, Prague, Tokyo, Toronto
- Those sites must have perfSONAR-PS deployed + Need to setup dedicated FTS channels.
- Measurements will be done before and after establishing connectivity to LHCOne (some sites are already connected to LHCOne).

CVMFS

- Almost 1/2 of ATLAS sites have deployed CVMFS.

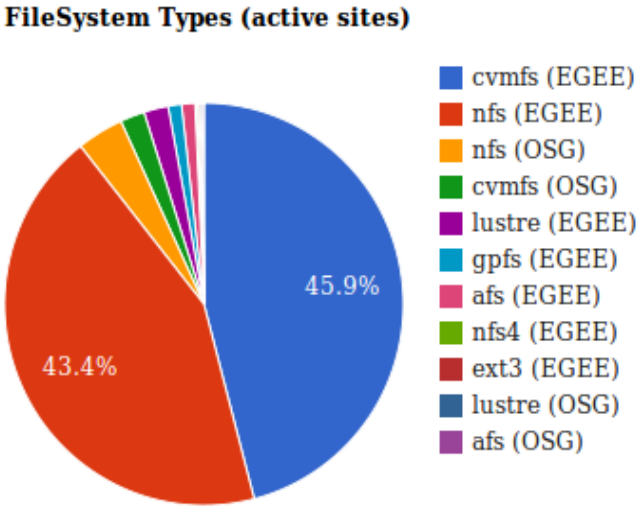

- Most of the remaining sites plan to deploy before Xmas.
- Few should deploy after Xmas and very few have no plans.
- Some sites waiting for development (missing features).

Conclusions

- Good performance of ATLAS computing that leads to nice physics results :-).
- But a high number of ALARM in last quarter ($\sim$1/week !).
- Some heavy operations successfully conducted (e.g. Castor to EOS migration, LFC consolidation). Help from many non-ATLAS people. Thanks a lot to them !
- Quite a few of network issues observed. LHCOne (currently being evaluated in ATLAS) should help in fixing these problems.
- Future challenges : 11g migration, continue LFC consolidation.