# ATLAS Quarterly Report and Plans

Dario Barberis

CERN & Genoa University/INFN

# Outline

- Organisation update
- Tier-0 data-taking activities
- Data export
- Prestaging tests
- Database access problems
- Pledged vs installed capacity
- Plans

# Organisation news

- The ATLAS Collaboration Board met last Friday and took the following decisions (among others):
  - Dario Barberis was re-appointed as Computing Coordinator from March 2009 until February 2010
  - David Quarrie was re-appointed as Software Project Leader from March 2009 until February 2010
  - Kors Bos was elected Deputy Computing Coordinator from March 2009 until February 2010
    - becoming Computing Coordinator from March 2010 until February 2011
- Hans von der Schmitt's term in office as Database Coordinator ended on 30 Sept after 2.5 years. He is replaced by:
  - Giovanna Lehmann Miotto (CERN) for online databases
  - Elizabeth Gallas (Oxford) for offline databases

# Tier-0 and data-taking activities

- We are taking continuously cosmic ray data since several months and until 3rd November
  - With only short breaks for detector work (and LHC data!)
- The Tier-0 is coping well with nominal data rates and processing tasks
  - A few Castor glitches are usually sorted out with the Castor team within a very reasonable time
- In November hardware detector commissioning work will restart
  - But cosmic data-taking will carry on with partial read-out

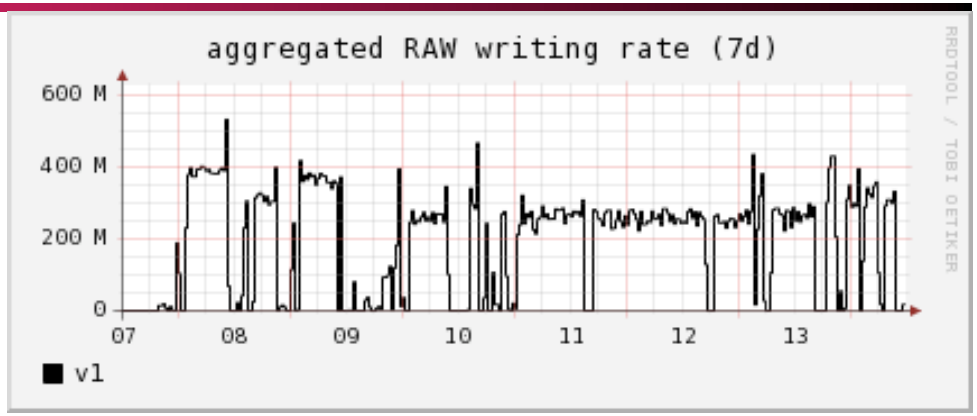


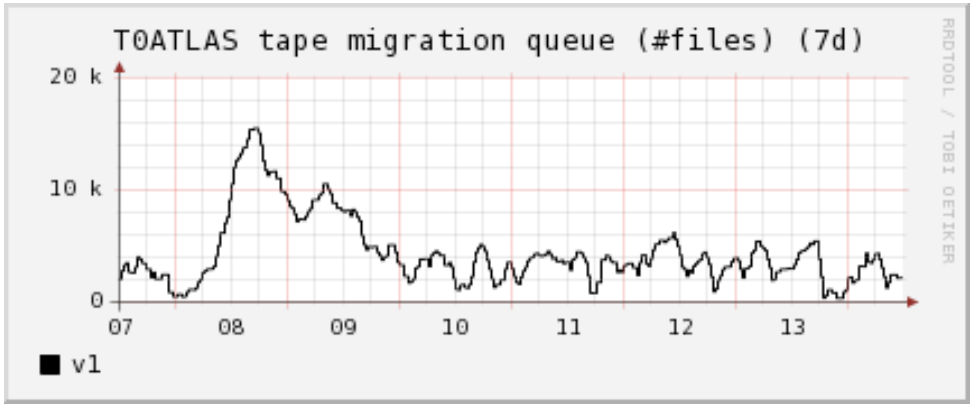


# Data export

- We export all raw and processed data from Tier-0 to Tier-1s and Tier-2s according to the computing model
  - The system can sustain the peak rate of 1.2 GB/s for an indefinite time
- Data distribution patterns are periodically revised as data types (triggers) and processing needs change

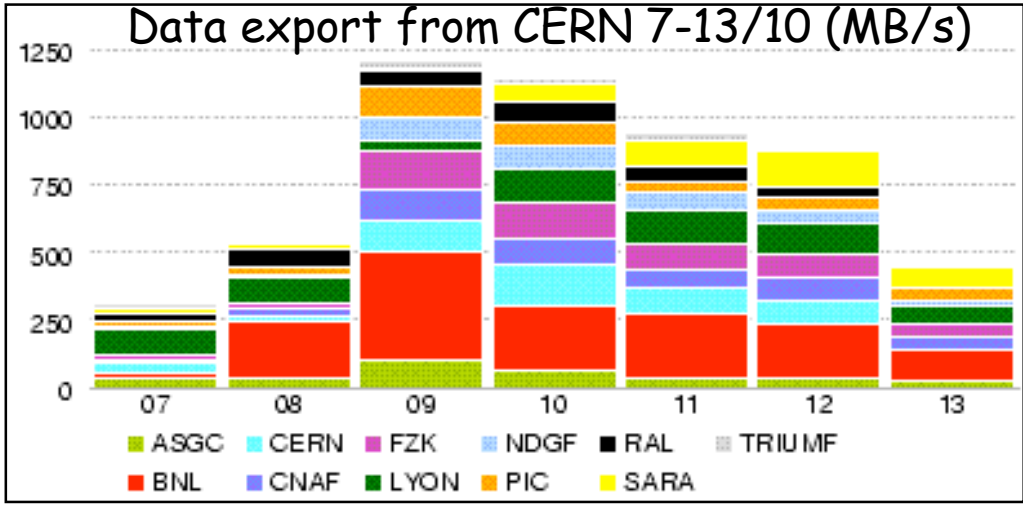


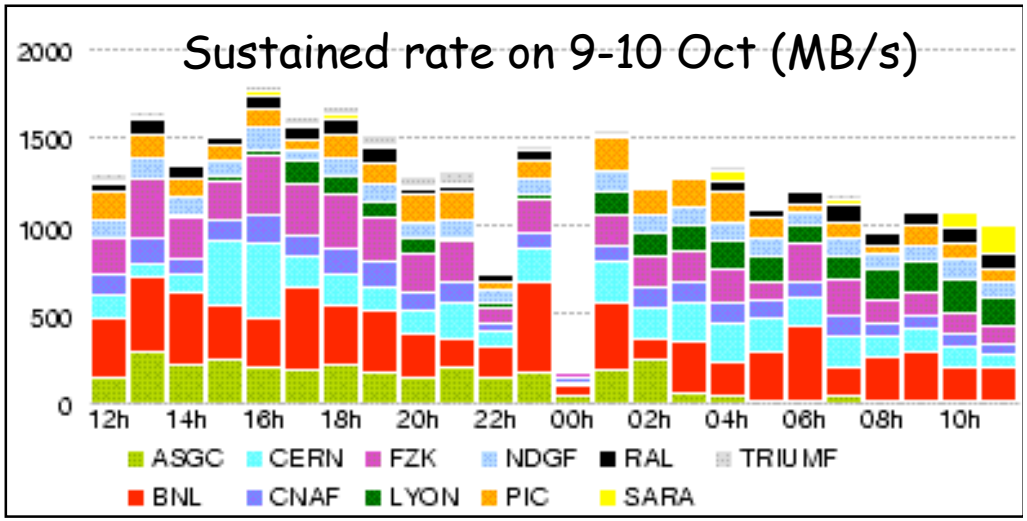

# Prestaging tests

- We started during the summer prestaging tests at all Tier-1s
  - Recalling whole datasets at a time (up to 10 TB)
- Performance varies a lot as tape back-ends are different at each site
  - After a few tries, most sites are mostly OK
- Outstanding (different) problems at PIC, FZK and SARA
- This exercise also showed that the number of available tape drives varies a lot from site to site
  - There is no point in having 1000s of processing cores if they cannot be fed at the correct rate with data
  - Example:
    - Our reprocessing tasks consume 1.6 MB of raw data every ~7 real seconds
    - One needs a total read rate from tape of 400-500 MB/s to keep 1000 cores busy
      - Including x2 contingency

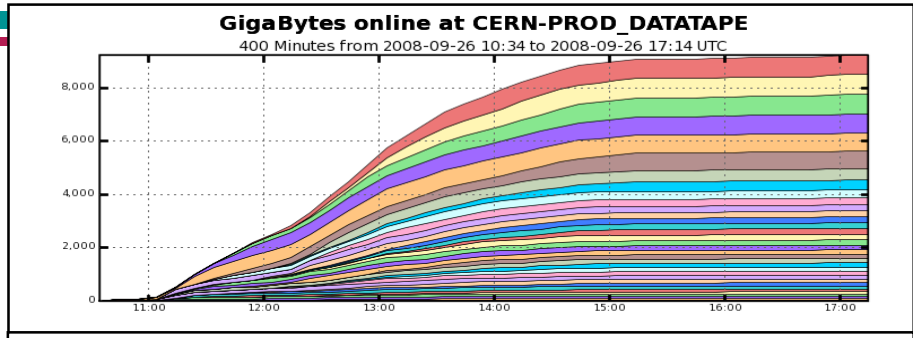


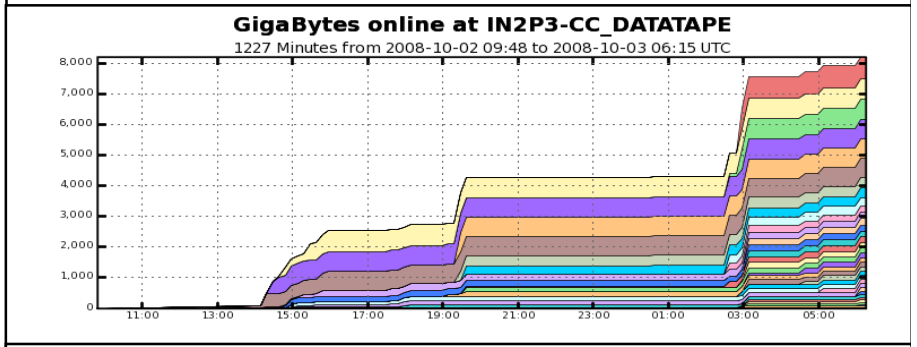


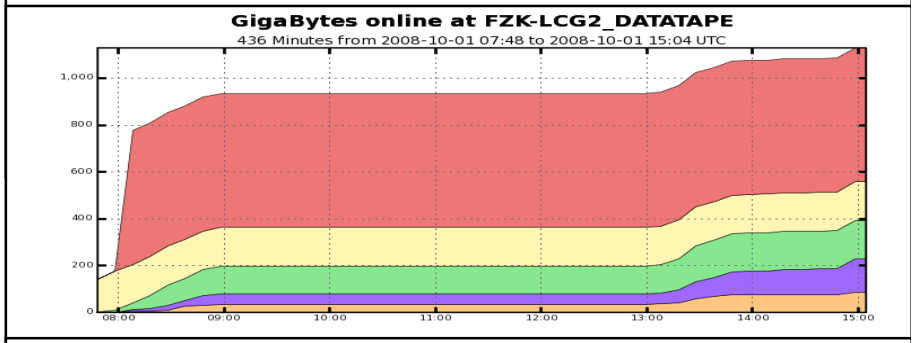


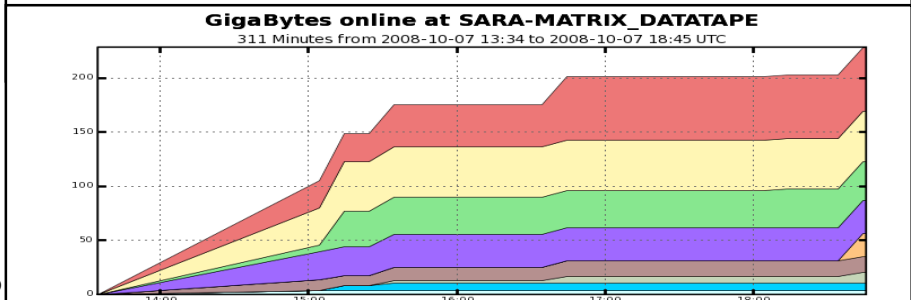

# Database access problems

- Early tests of database scalability did not indicate there would be any problem with reprocessing at Tier-1s
  - More recent tests instead showed a serious limitation when more than a few 10s (up to 100) jobs start simultaneously, as they all access conditions data from Oracle databases
- Two factors differed between these tests:
  - Oracle Streams are now used to move data from CERN to Tier-1s
  - DCS (Detector Control System) data are now accessed by reconstruction tasks
- Actions in progress:
  - Task force to analyse data access patterns from the Oracle server side with ATLAS and CERN DBAs
  - Activity to instrument Athena to log database access and data volumes
  - Action on detector code developers to revise and optimise their database access patterns
  - (Last but not least) Exploration of the SQLite technology for reprocessing tasks
    - Dump all data for a given run to an SQLite file and use it locally for all jobs
    - Reduces the database access by a factor of several 100s (the number of files in a run)

# Disk at CERN and Tier-1s (2008)

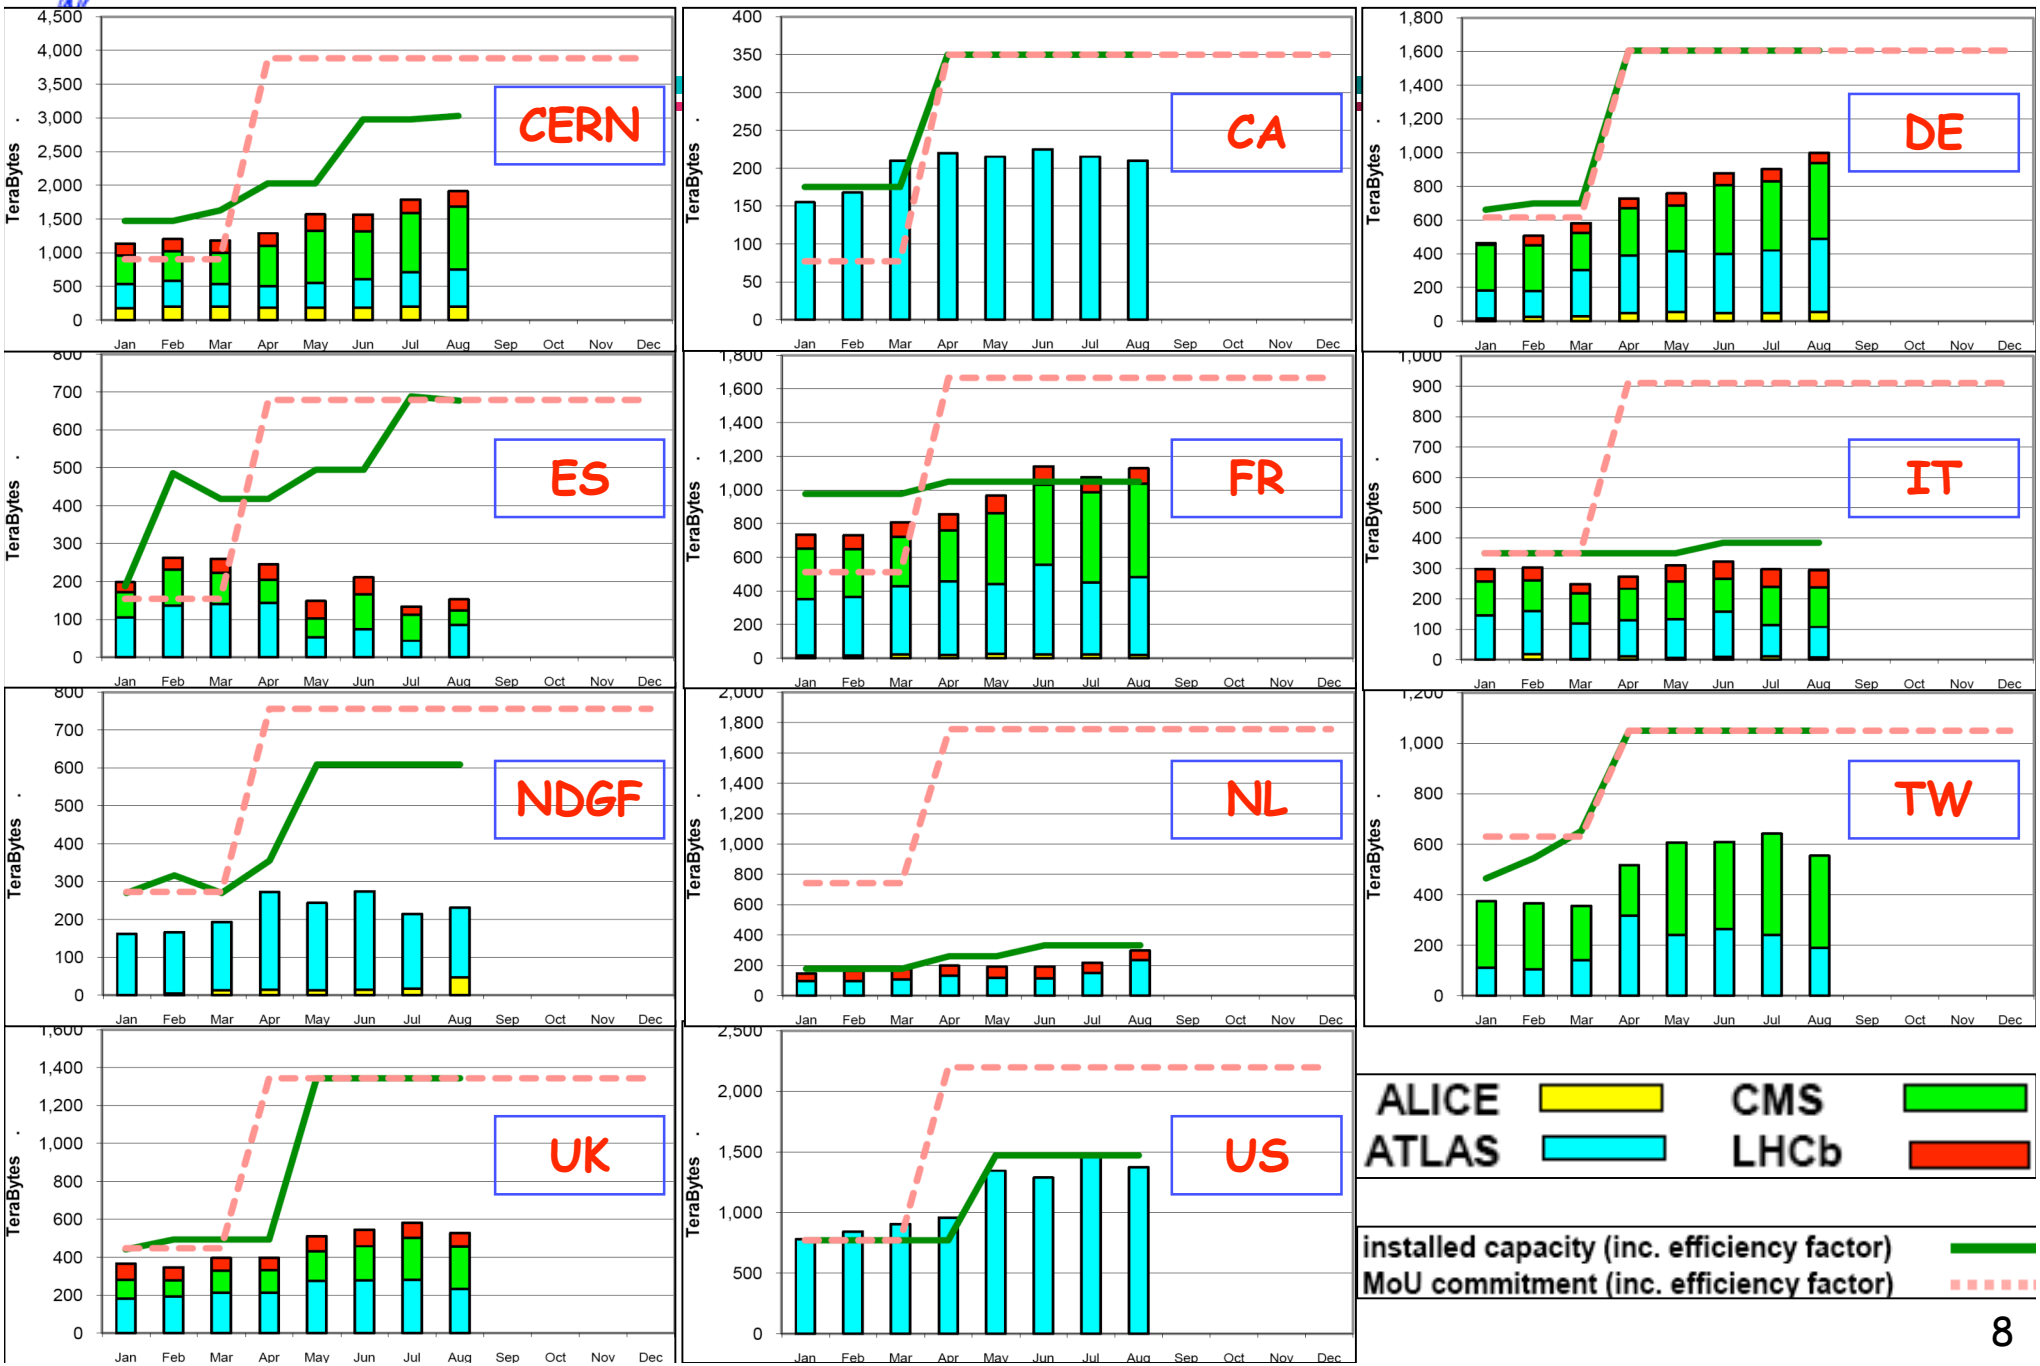

# Plans

- Software releases:
  - 14.X.Y releases
    - Bug fixes only for HLT/Tier-0 and Grid operations
  - Release 15.0.0
    - February 2009. Include feedback from 2008 cosmic running.
    - Base release for 2009 operations.
- Cosmic runs:
  - Complete detector:
    - Continuing till early November 2008
    - Restarting late March 2009
  - Partial read-out:
    - All the time
- Collision data:
  - Ready to go from April 2009 for what concerns ATLAS Software & Computing