ALICE report for February phase of CCRC'08

Latchezar Betev (ALICE)

March 4, 2008

Data management – focus and status

- Quasi-online RAW data replication T0->T1
  - OK, ongoing
- Replication of specific data (ESDs, calibration and special runs) to T2s
  - Ongoing, done on request by the detector groups
- Deployment of xrootd-enabled storage across all ALICE T1s and as many T2s as possible
  - All T1s have deployed or are in the process of deploying
  - 6 T2s have deployed considerable storage capacities (in production), more are joining

Data management – T1s

- Three virtual SEs per T1 site with different storage directories
  - T1D0 for RAW
  - T0D1 for ESDs
  - T1D0R for complementary data (keeping 60MB/sec constant rate out of T0)
- Implemented, OK with respect to the access patterns
- Accessible through gridftp and xrootd

Data management – storage type

- dCache – stable, new development for support of advanced xrootd features is ongoing
  - CCIN2P3, GridKA, NDGF, NL-T1
- CASTOR2 – real success story
  - Stable at CERN
    - Fixes transparently and quickly put in production
  - CNAF installation (in one week!) working
    - Many thanks to CNAF and CERN experts!
  - RAL - ongoing

Data management - files

- MSS–friendly large RAW filesizes
  - Exercise started with ‘old style’ 1GB RAW data chunks
  - Since 2 days, DAQ has switched to 10GB files for tests
  - The RAW data is not yet ‘standard’, still non-negligible proportion of short runs with small files (calibration) – expected to improve rapidly
- ESD ‘consolidation’ is in the works, but may become a non-issue with large RAW data size
- Still no good solution for user files
  - However RAW and production data should start to dominate the average filesizes in MSS

Data management - replication

- Gradual inclusion of all T1s and number of T2 sites
  - T1s – copy of RAW data, T2s – copy of ESDs

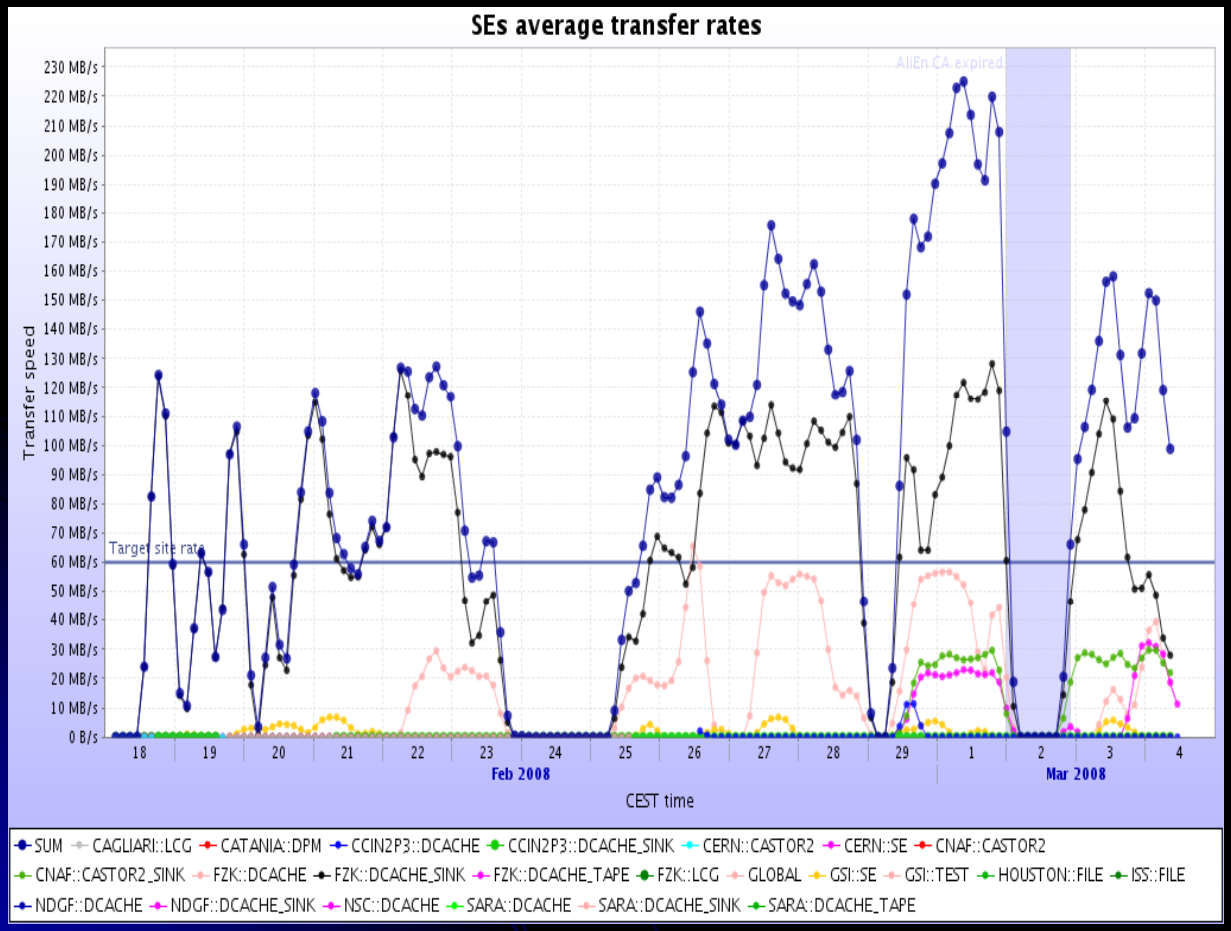

The target rate of 60MB/sec was achieved, average over 2 weeks is 80MB/sec
Uneven transfer profile – proxy corruption problem + ALICE software updates

Data management - replication (2)

- Simultaneous replication together with ATLAS/ CMS LHCb not a problem
- Short interruptions of transfers (up to one day) is manageable – the data is still ‘hot, on disk’ at the source
- The FTS is only ‘VO-expert level’ tool

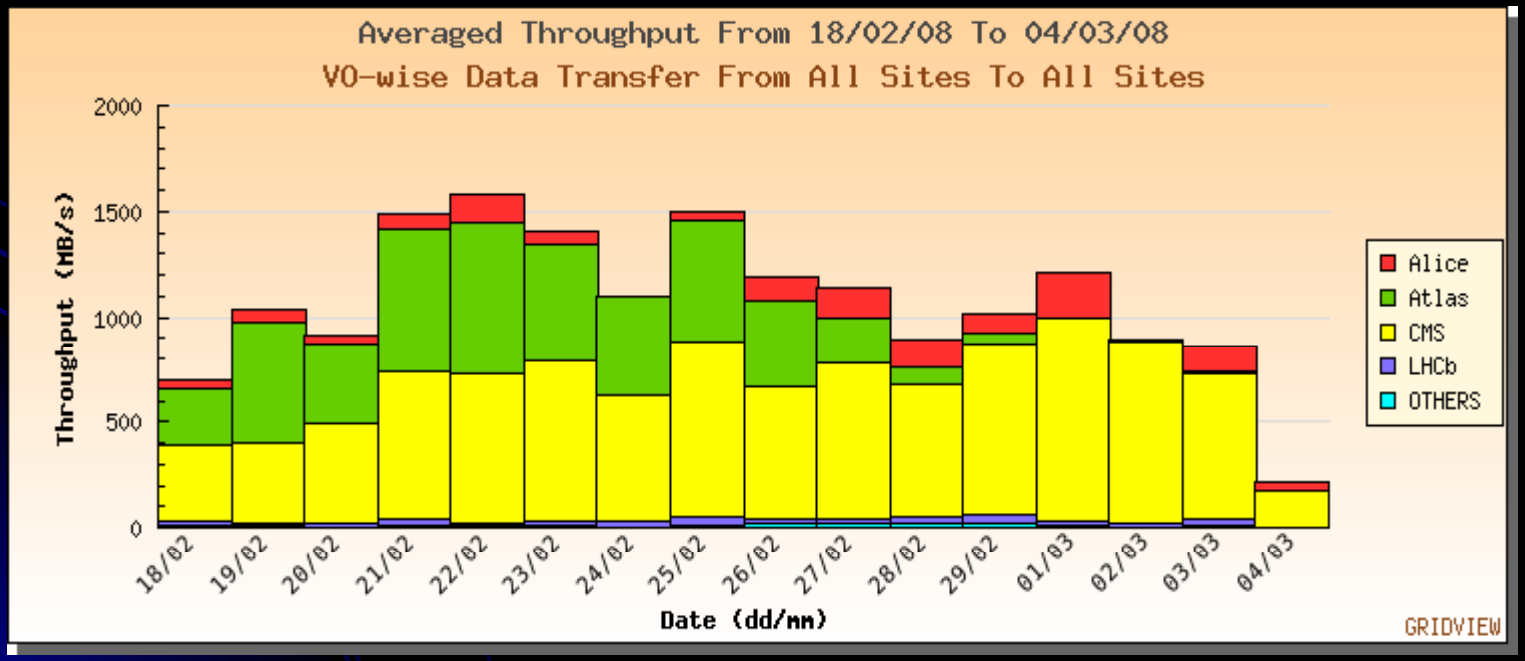

RAW data registration

- Commissioning exercise RAW data volume almost exactly the expected 60TB (surprisingly)
  - We may go to about 70TB until the end of the exercise, however the capacity is available both at T0 andT1s

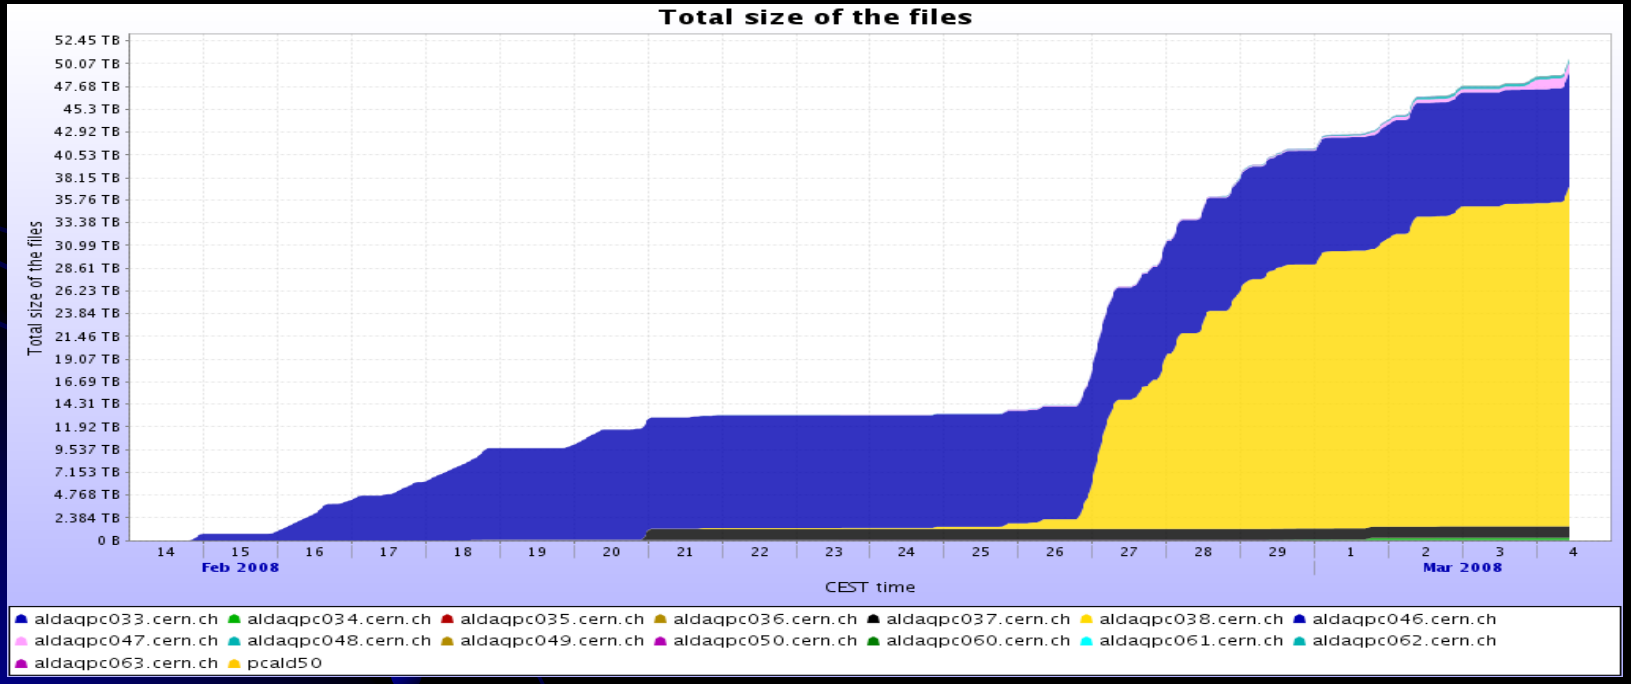

RAW data access

- Top priority given to detector experts for Grid and CAF processing of calibration/special detector data
- The standard ESD production is delayed due to the data type and specific conditions
  - Not yet in quasi-online production regime

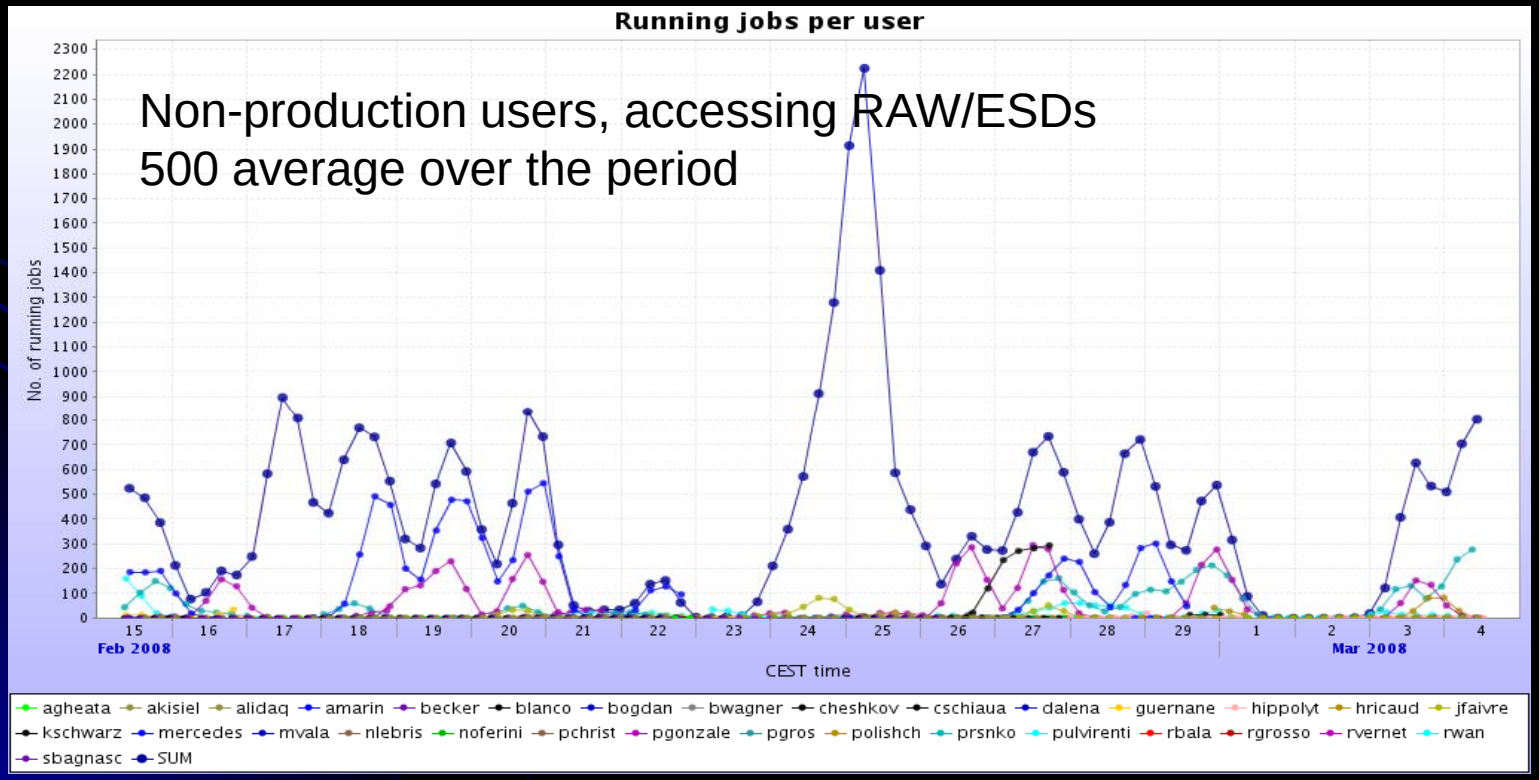

Storage setting up

- Somewhat delayed on several sites
- A good lesson for us to start site preparations much more in advance
  - However some of the code was not yet ‘in production’
- The novelty of xrootd-enabled storage was the predominant factor
  - dCache sites were ready first
- The issues were not related to available capacity, but configuration

Problems and response

- Extremely good reaction times – most of the issues are solved within hours of reporting
  - Having problems marked as ‘urgent’ helps ☺
  - Experts support at T1s and T2s continues to be excellent
  - ALICE dedicated support at the T1s is still an issue
- Direct communication with site experts is still the fastest way to go
  - And probably the only way for VO-specific issues
- The standard tools: eLog, GGUS were also used extensively

Summary and outlook

- For ALICE, the CCRC exercise has fulfilled its purpose
  - Focus on data management
  - Brings all experiments together
  - Controlled tests, organization
- ALICE is continuing this week and possibly next (depending on detector plans) with RAW data from the experiment
  - No additional storage/processing resources needed yet, however we reserve the right to ask
- Data transfers and storage will not stop after the commissioning exercise – DC mode from now on