CCRC'08 – May phase and beyond

Latchezar Betev

ALICE Offline week, April 11, 2008

CCRC'08 short plan

- Set of common exercises planned since Dec'2007
- Involving all LHC experiments
- Two major periods – February 2008, May+ 2008
- General goals
  - Test the LCG data management utilities (FTS, SRM) and storage
  - Test of data transfers (bandwidth, CASTOR2@CERN, storage@T1)
  - Each individual experiment tests its computing model (from RAW to ESD)
  - Experiments can put additional tasks in

ALICE in February phase - RAW data management

- Registered 82TB from 14/02 to 09/03 (26 days)
  - This is 70% of our p+p data rate
- Data replication – see Costin's presentation
- The CERN (T0) and T1 storage is validated for RAW and replication
  - We had sufficient MSS capacities, all ALICE T1s (but RAL) participated in the exercise
  - Still to validate the T1 storage with massive data production
- In general – no major concerns regarding RAW registration and replication

Data volumes and replication

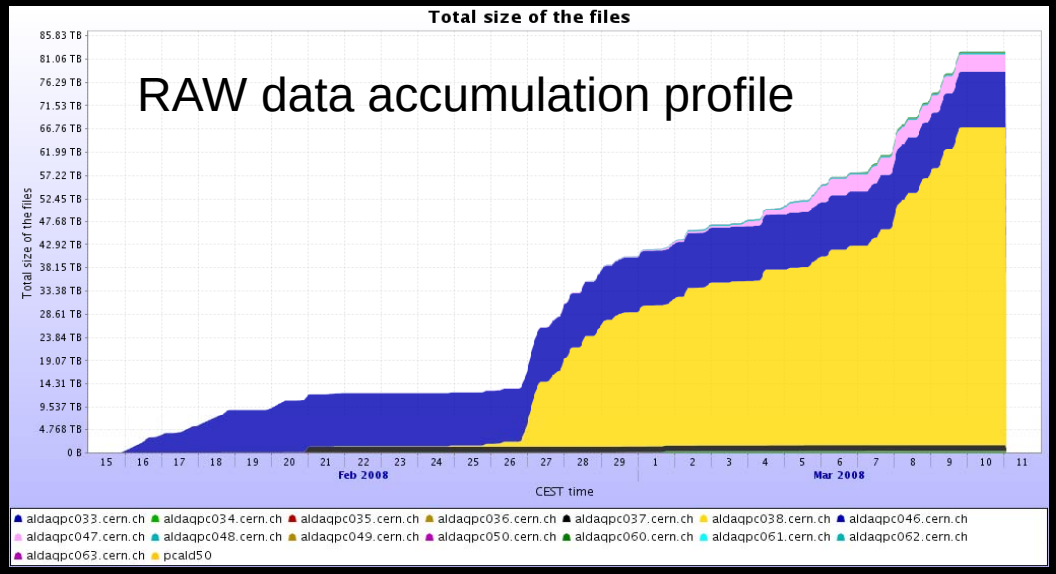

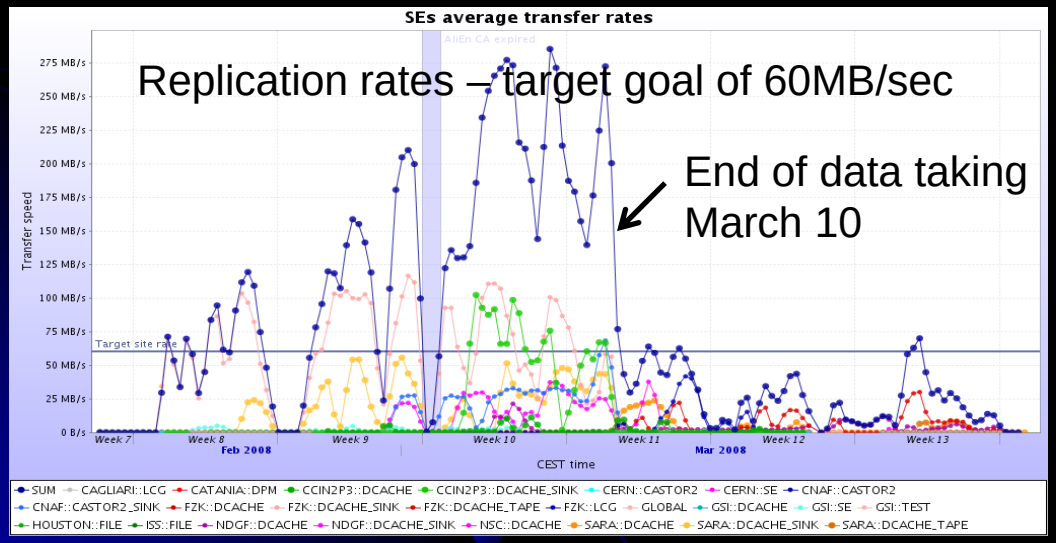

Total of 82TB in 90K files
0.9GB/file

To improve MSS performance switch to larger (10GB) RAW data chunks

5/6 T1 centres participating
Sufficient storage capacity
Good overall stability of storage and replication tools
Good T2 centres involvement

Conditions data gathering - Shuttle

- Fully operational during the February/March commissioning exercise
- Extensively reported on in the morning session
- ALICE is in very good shape!

CCRC'08 May phase

- Global - whatever was not tested in February should be included now
  - This is mostly concerning data management
- For ALICE
  - Same as in February/March - data registration, replication T0->T1, shuttle
- Focus on quasi-online data reconstruction T0/T1
  - The strict AliRoot release policy will help a lot here

Storage requirement

- Additional resources needed for May exercise (80% p+p scenario)
- Disk will store ESDs from RAW
  - Assuming ESD+other files 20% of RAW

Tier 1 site	Disk space (TB)	Tape space (TB)
CCIN2P3 (15%)	3	14
CNAF (15%)	3	14
GridKA (45%)	9	44
NDGF (15%)	3	14
RAL (5%)	1	5
T1-NL (5%)	1	5
TOTAL	20	96

Storage deployment

- Request for additional storage will be sent to the T1s today
  - GridKA has already put the storage in place (thanks Kilian)
- Should not be a problem, as the resources must be in place by now (WLCG compliance), question of configuration
- Only one new SE – RAL
- Keeping the same directory structure
  - T1D0 for RAW
  - T0D1 for ESDs
  - T1D0R for complementary data (keeping 60MB/sec constant rate out of T0
    - The use of this will be kept to a minimum

Storage deployment – T2s

- Progressing rather well – we have 11 SEs at T2s in production, total of 150TB
  - We are also filling it rapidly…
- MC production primary copy is now stored at T2
- Replication of specific RAW runs and ESDs ongoing

Replication

- February phase – the relative replication rate of the T1 sites was not enforced
- For May, we will follow on the replication rate of all sites
  - According to the table on slide 7
  - Will validate the storage performance wrt expectations (writing and reading – if reco works as expected)
  - Will validate the FTS ‘VO-shares’ mechanism (policy changed in April)

RAW data production

- Quasi-online reconstruction is a ***must***
  - Systematically this was never done (for good/bad reasons)
- Output of the reconstruction needs to be carefully re-evaluated (we will not have sufficient storage)
  - Currently we write *.root

Grid updates

- Ongoing – VO-box migration to SL(C)4 / gLite 3.1
- AliEn and AliRoot build – finally on SLC4
- Re-deployment (AliEn v.2-15) on all sites
- Same for all AliRoot/ROOT/GEANT3 versions needed
- We must accomplish all of the above ***before*** the end of the month

Summary

- May phase of CCRC'08 for ALICE will be very similar to the February phase
  - Special emphasis on quasi-online reconstruction
  - And storage management validation
- From Grid point of view, it is a common exercise for all 4 LHC experiments (and all 4 will participate)
  - There will be some contention for resources, but nothing critical is expected
- Depending on the machine schedule CCRC may develop in a full-fledged data taking/processing