



ALICE CCRC'08 post-mortem

Latchezar Betev

Friday the 12+1-th June 2008

General tasks list

- Registration of data in CASTOR2 (T0) and on the GRID
- Replication T0->T1
- Conditions data gathering and publication on the GRID
- Quasi-online reconstruction
 - Pass 1 at T0
 - Pass 2 at T1s
 - Replication of ESDs to CAF/T2s
- Quality control
- MC production and user analysis at CAF/T2s

May period

- ALICE detector activities
 - Finalization of detector installation
 - Tuning of detector parameters
 - Calibration and alignment data
- ALICE Offline upgrades
 - New VO-box installation (see 'Critical services' presentation)
 - New AliEn version (see 'Critical services' presentation)
 - Tuning of reconstruction software
 - Exercise of 'fast lane' calibration/alignment procedure
 - Other...

Data volumes

- Despite the massive upgrades of ALICE Grid services, steady registration or RAW from experiment + replication
 - Many thanks to CERN IT/FIO for the efficient CASTOR2 upgrade (coinciding with ALICE DAQ and AliEn central services upgrade)

2543 runs		113644 files	84.25 TB
-----------	--	--------------	----------

- Almost exactly as expected (96TB)
 - Written to CASTOR2 and replicated
- File sizes still low, also expected
 - Short calibration runs

File sizes

- 10GB/chunk RAW data size tests are very promising
 - Recent run

```
/alice/data/2008/LHC08b/000037956/raw/ > ls -al
-rwxr-xr-x  alidaq  alidaq           530440 Jun 12 20:34 08000037956000.0.tag.root
-rwxr-xr-x  alidaq  alidaq      9776226240 Jun 12 20:55 08000037956000.10.root
-rwxr-xr-x  alidaq  alidaq      9769266079 Jun 12 20:55 08000037956000.20.root
-rwxr-xr-x  alidaq  alidaq      9770767560 Jun 12 20:55 08000037956000.30.root
-rwxr-xr-x  alidaq  alidaq      9769218809 Jun 12 20:55 08000037956000.40.root
-rwxr-xr-x  alidaq  alidaq      9768620353 Jun 12 20:55 08000037956000.50.root
-rwxr-xr-x  alidaq  alidaq      9777456300 Jun 12 20:37 08000037956000.60.root
-rwxr-xr-x  alidaq  alidaq      3391443105 Jun 12 20:35 08000037956000.70.root
```

- No degradation of writing (rfio) and reading (xootd) observed (not expected to see any)
- Data reconstruction, streaming from file servers, is ongoing

File sizes (2)


- Large RAW chunk size - simple way to address many pending issues
 - MSS requirements for large files and sparse access to tapes
 - Fewer entries in the Grid file catalogue
 - Efficient (regarding MSS) file sizes of secondary and tertiary files
 - 10GB RAW -> ~1GB ESDs -> ~100MB AOD (to be consolidated)
- Testing will continue also next week
 - Put as default in DAQ afterwards

Replication

- Only RAW data replication running (no padding)
- Replicated immediately after data registration in CASTOR2@CERN
 - Data on disk, no staging from tape
- Replication manager functions
 - Number of transfer scheduled according to the ALICE T1 storage capacity share
 - Issue with the logic – if there are no pending transfers, always the site with largest share (in ALICE – GridKa) gets the transfer – to be fixed

Replication (2)

- 60MB/sec rate (when data present) as per ALICE share
- Total replicated 63TB (out of 84TB)
 - Part of the data deliberately excluded



Conditions data

- Comprehensive test of conditions framework readiness
 - Collection of data from experiment
 - DAQ/DCS/HLT databases
 - Detector calibration software running in DAQ framework
 - External conditions sources – LHC machine parameters, magnet status, etc...
 - Publication of data through Shuttle framework in OCDB

Conditions data (2)

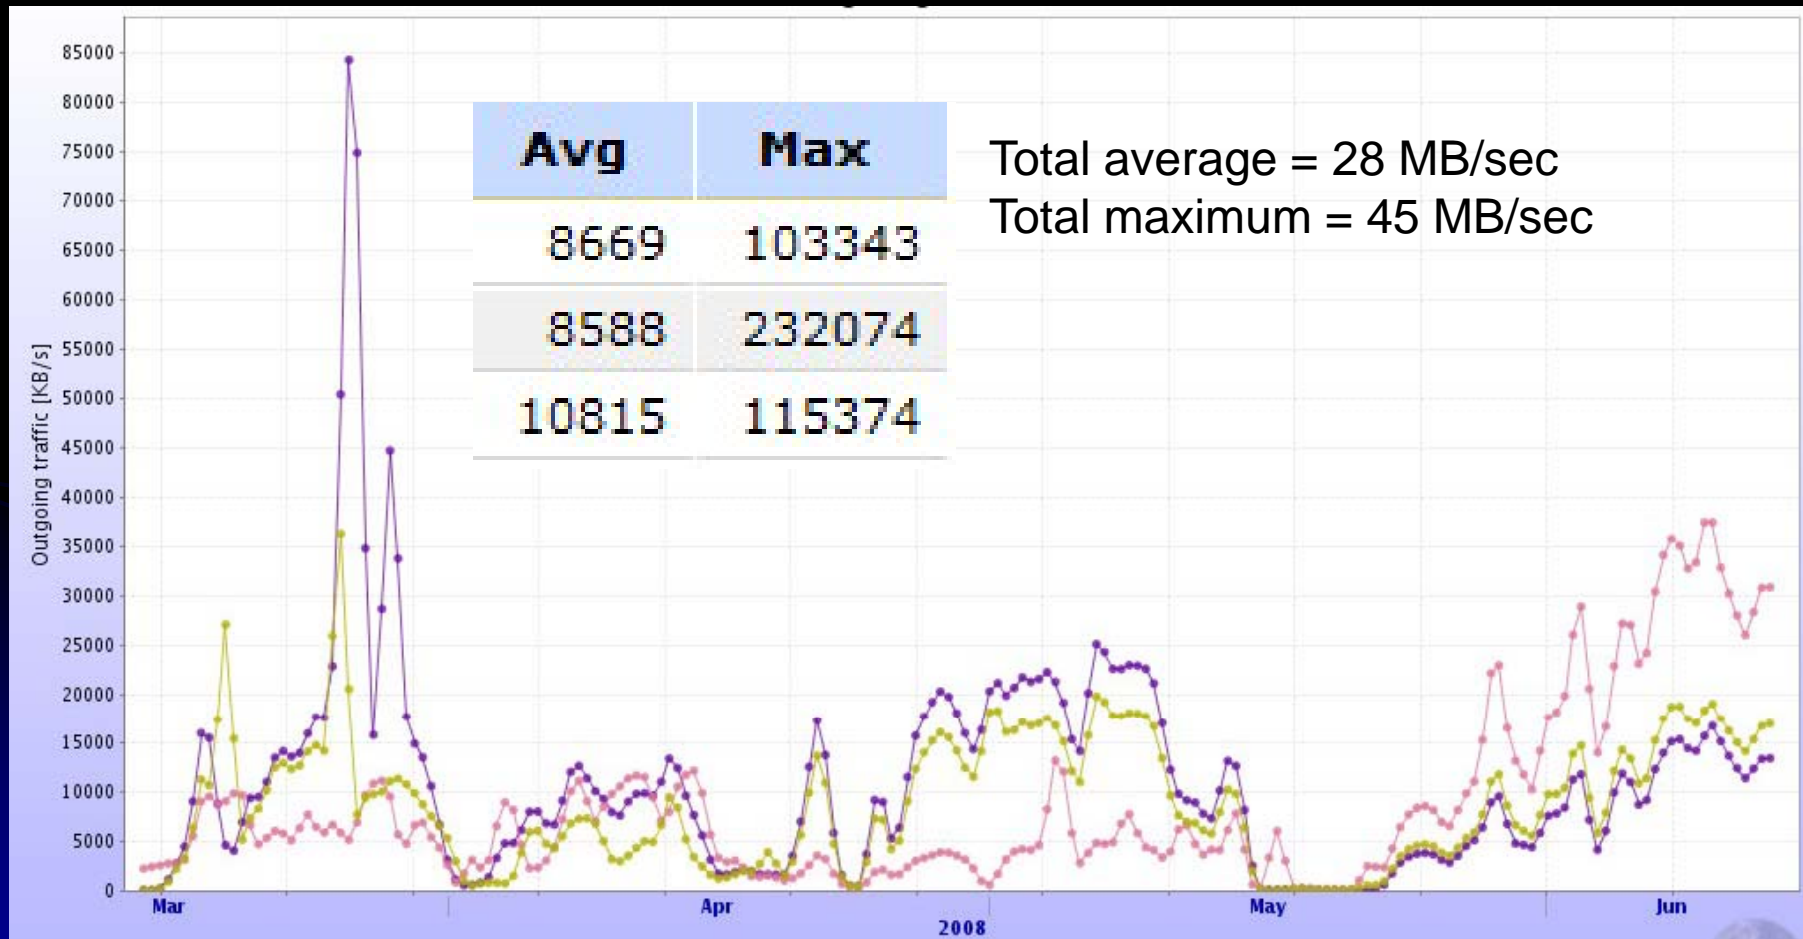
- Reminder – OCDB is
 - ROOT objects, published in a special section of the AliEn file catalogue
 - Associated metadata for fast search and selection of condition parameters
 - Stored on Grid storage and replicated to T1s
 - Accessed by all reconstruction and some analysis tasks

Conditions data (3)

- The OCDB structure has been in operation since mid-2007
- In the past 3 months it was stress-tested with special Grid jobs with full OCDB access
- Tests include failover in case of inaccessibility of storage replicas
 - Up to a limiting case with only one replica available

Conditions data (4)

- Operation of CERN instance



Data reconstruction - emphasis

- Fast MC production for first physics with various LHC startup scenarios
 - And analysis for first publication
- Fast analysis of detector calibration data
 - Essentially immediately after data taking (same day/night)
 - Crucial for feedback to detector experts

MC for first physics

PDC 08/LHC08b6	First physics pp, Phojet, No field, 10 TeV	Completed
PDC 08/LHC08b5	First physics pp, Pythia6, No field, 900GeV	Completed
PDC 08/LHC08b4	First physics pp, Pythia6, No field, 10TeV	Completed
PDC 00/LHC00b3	First physics pp, Phojet, 5kG, 10TeV	Completed
PDC 08/LHC08b2	First physics pp, Pythia6, 5kG, 900GeV	Completed
PDC 08/LHC08b1	First physics pp, Pythia6, 5kG, 10TeV	Completed

- Number of events per cycle – ‘first few days of LHC running’
- High priority production – 10 days for all
- Fast re-processing of data also tested

Data reconstruction - emphasis

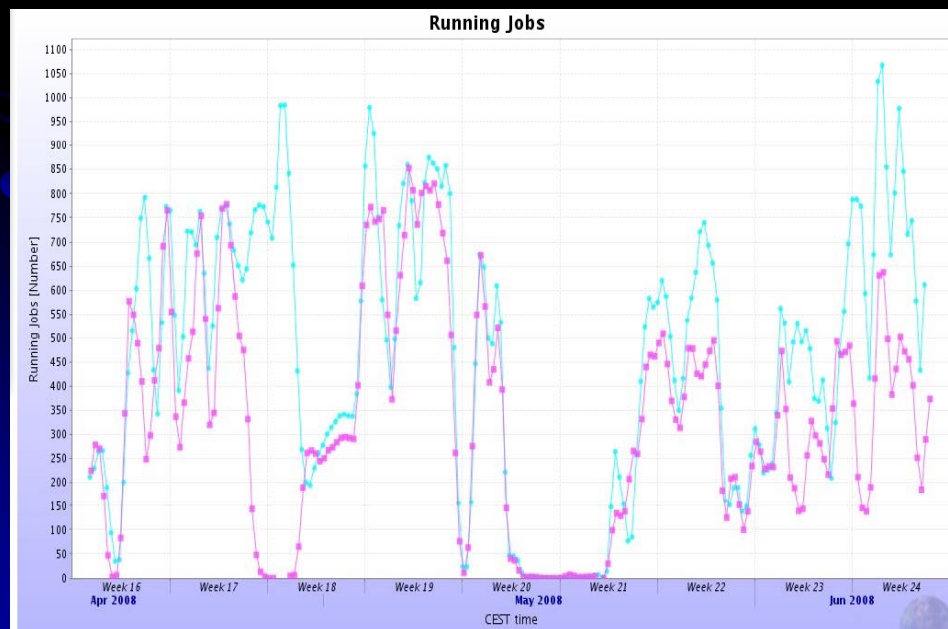
- Fast MC production for first physics with various LHC startup scenarios
 - And analysis for first publication
- Fast analysis of detector calibration data
 - Essentially immediately after data taking (same day/night)
 - Crucial for feedback to detector experts

Analysis

- Critical data (RAW/ESDs) immediately replicated to CAF@CERN
 - In addition to specific T2s with large detector communities
- Fast lane analysis is generally performed on CAF
- In addition, the user Grid analysis was ongoing

Reco/analysis duration/efficiency

- Difficult to predict – the job duration is determined by the quality of the RAW data
 - Amount of noise in the detector – big events, but little content
 - Trigger selectivity
 - Special event types



Average efficiency for reconstruction jobs (last 2 months)

Jobs efficiency (cpu time / wall time)				
Farm	Last value	Min	Avg	Max
CERN-L	67.98	0	71.42	101.3
CERN_gLite	63.75	0	69.39	100
Total	65.86		70.4	

Critical services and monitoring

ALICE Wiki: <http://twiki.cern.ch/twiki//bin/view/ALICE/CCRC08>

- See 'Critical services' report by Patricia Mendez, Pablo Saiz

- Many thanks to the site experts for the excellent and quick support

Summary

- The May phase of CCRC'08 was centered on detector activities and upgrades of the ALICE Grid installations
- Collection, replication and processing of calibration and alignment data from the experiment
- Testing and tuning of Grid services ongoing – generally smooth operation
- June and beyond period – full experiment data taking and reconstruction at T0/T1s
 - Including ongoing MC production and analysis at T2s
- **We are (reasonably) ready for LHC data**