



ATLAS

Report and Plans

Dario Barberis

CERN & Genoa University/INFN





Outline

- Tier-0 data-taking activities
- Data export
- Prestaging tests
- Database access issues
- Resources for 2009
- Functional tests
- Plans



Tier-0 and data-taking activities


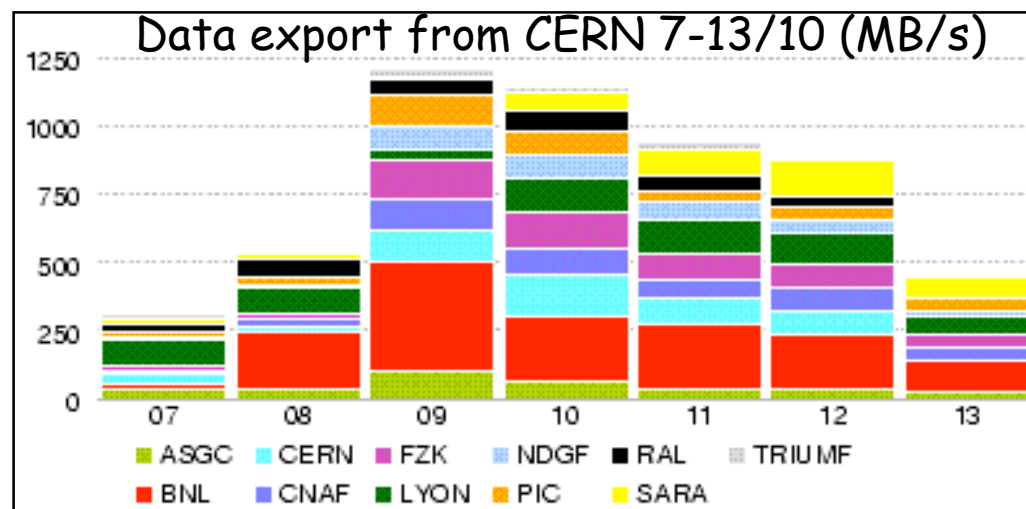
- We took continuously cosmic ray data for several months and until 3rd November
 - With only short breaks for detector work (and LHC data!)
- The Tier-0 coped well with nominal data rates and processing tasks
 - A few Castor glitches were usually sorted out with the Castor team within a very reasonable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





Data export




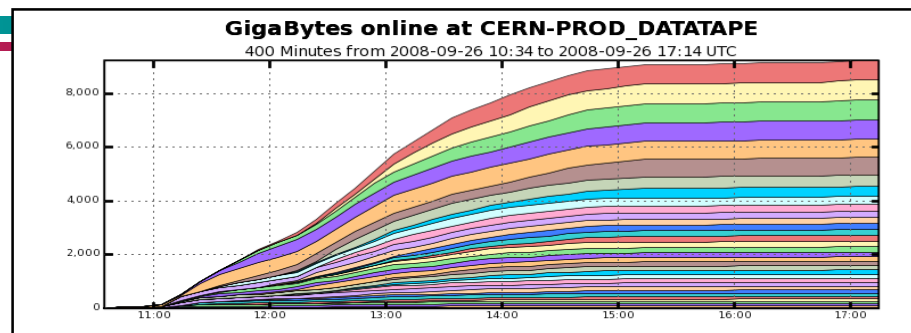
- We export all raw and processed data from Tier-0 to Tier-1s and Tier-2s according to the computing model
 - The system can sustain the peak rate of 1.2 GB/s for an indefinite time
- Data distribution patterns are periodically revised as data types (triggers) and processing needs change





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 some dCache sites
- 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
 - 10 tape drives, including x2 contingency





Database access issues

- 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 undertaken:
 - 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
 - Throttling of the job submission rate to avoid excessive concurrent database access
 - As a result, we are now able to reprocess 2008 cosmics data with no known database limitation
- Exploration of the FroNTier technology for reprocessing tasks at Tier-1s and analysis tasks at Tier-2s
 - Cache all data for a given run and use it locally for all jobs
 - Reduces the database access by a factor of several 100s (the number of files in a run)



Resources for 2009

- The C-RRB Scrutiny Group broadly approved our computing capacity requests for 2009

		Tier-0	CAF	Tier-1s	Tier-2s	Total
CPU/MSI2k	Scrutiny	7.6	5.8	29.4	30.3	73.1
	Request	7.6	5.8	28.4	27.0	68.8
Disk/PB	Scrutiny	0.7	3.3	21.4	14.6	40.0
	Request	0.7	3.3	20.9	13.3	38.2
MS/PB	Scrutiny	8.6	1.5	15.1	0	25.1
	Request	8.6	1.1	15.8	0	25.5

- Their report nevertheless includes a few factual mistakes that we would like NOT to propagate to the LHCC and elsewhere
 - We did NOT change our computing model in 2008. The increased request for resources at CERN (mainly for the CAF) comes from the experience built with FDR exercises and cosmic ray data-taking and the time needed for the prompt calibration/alignment loop
 - There is NO proliferation of event formats. The recently defined DPDs (Derived Physics Data) were already foreseen in the Computing TDR as skims of events in the same global format as ESDs and AODs.




Functional Tests


- DDM Functional Test
 - we get a VERY BIG one coming up
 - next 2 weeks we accumulate a lot of data in the T1's (on DATADISK)
 - in week 3 we then launch T1-T1 transfers all at once
 - this will test DDM at an order of magnitude higher rates
 - otherwise we shall always run the DDM FT
 - to keep testing all sites
 - and this is badly needed, unfortunately!
- MC Production Functional Test
 - running at low rate now: 200 jobs/cloud/week
 - will increase the rate to become more significant
 - will be brokered indiscriminately so all T2 sites get jobs
 - we should always run them, at lower priority than real MC
 - need to check if priorities work
- Distributed Analysis Functional Test
 - will be expanded to all clouds (now in IT, DE, US...)
 - but needs people watching in (almost) all Tier-2's
 - need to identify real user analysis use-cases (we just have a few different ones now)
 - need to be tested with Panda and WMS backend in EU clouds
 - need to check if priorities work



DDM FT: examples



Examples of Tier-1-Tier-1
and Tier-1-Tier-2
Functional Test transfers
on 17 Nov 08 (1 hour)





Plans for 2009

- November-December 2008
 - preparation for re-processing
 - scale-up MC and Analysis FT
 - Super throughput challenge
 - Panda installation at CERN
- January-March 2009
 - re-processing
 - further scale-up MC and Analysis FT
 - Panda running from CERN now
- April
 - data clean-up
 - preparations for new real data
- May
 - Cosmics and detector commissioning - ready for collisions