

ATLAS Tier0 Operations



ATLAS Software & Computing Workshop
Mon, 04 April 2011

Tier0 status and plans

Getting ready for 400Hz data taking

- reconstruction time at the Tier0
- storage capacity
- resources to run jobs
- RAW compression

Improvements in monitoring

Tier0 status report

Concerns on software side: pileup, events size

- many s/w testing campaigns (“Big Chain Tests”, BCTs), frequent processing of the same data with several different configuration tags
 - Run 167607 (last year, $\mu \sim 2-3$)
 - Run 178044, 178109 (recent, $\mu \sim 8$)
- Reco wall time now well below 15s for all streams
- with 3000 job slots, allows on average (400Hz DAQ rate) @ (50% LHC efficiency)
- see [presentation](#) by Thomas Kittelmann @ ATLAS Weekly meeting

Tier0 status report

Resource extension in progress

- staged extension of the main CASTOR production pool t0atlas
 - to cope with I/O, more disk buffer comes “for free” (currently 8GB/s concurrent I/O)
 - 74 servers (~700GB)
 - (currently) 106 servers (~1PB)
 - (eventually) 116 servers (1.2PB)
- taken from the atldata pool

Tier0 next steps

Current LSF status:

- atlast0 cluster
- ~3400 cores: 3000 for reco slots, 300 for CAF/TMS, rest for RAW, AOD, DESD, NTUP, HIST merging, beamspotproc and other special tasks

LSF reconfiguration: in case of high load, Tier0 will “spill over” to the ATLAS public share

- priority over other jobs
- “transparent” reduction of the share of users’ and Grid jobs
- no pre-empting
- currently up to 50% of the public share foreseen to be shared

timescale: 1-2 weeks

Tier0 next steps

RAW compression

- achievable volume reduction: 50%
- event-level compression example code provided by Wainer Vandelli
 - preserving event and file structure
 - using tdaq online libraries
- combined RAW compression and merging to be tested on Tier0

timescale: 1-2 weeks

- final deployment has also to wait for offline s/w development
 - modified tdaq libraries have to be ported to offline release to read back compressed events

Tier0 monitoring

conTZole 1.0

- User oriented, user configurable
- Rock solid reliably designed
- Built using the modern Web 2.0 technology
 - AJAX (jQuery, xmlhttp), Python
- Most recent improvements
 - Global Tier-0 processing statistics
 - API to access Tier-0 data from 3rd party applications

Tier0 monitoring

