

The logo consists of the letters 'E' and 'S' in a bold, white, sans-serif font, positioned on a black background.

Experiment Support

CERN IT
Department

LHCb Experiment Operations Report

Stefan Roiser

IT-ES-VOS


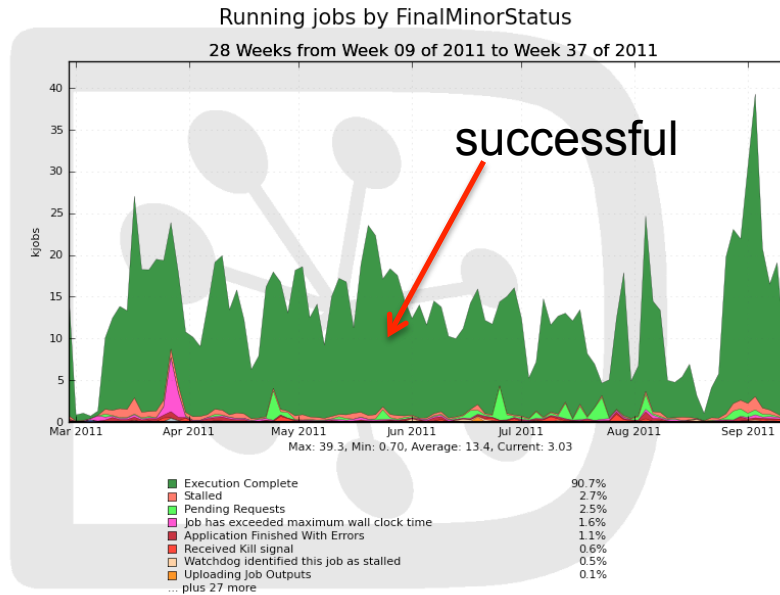
on behalf of the LHCb Grid Operations team

- News and major changes
- LHCb grid operations in last half year
- Reprocessing campaign


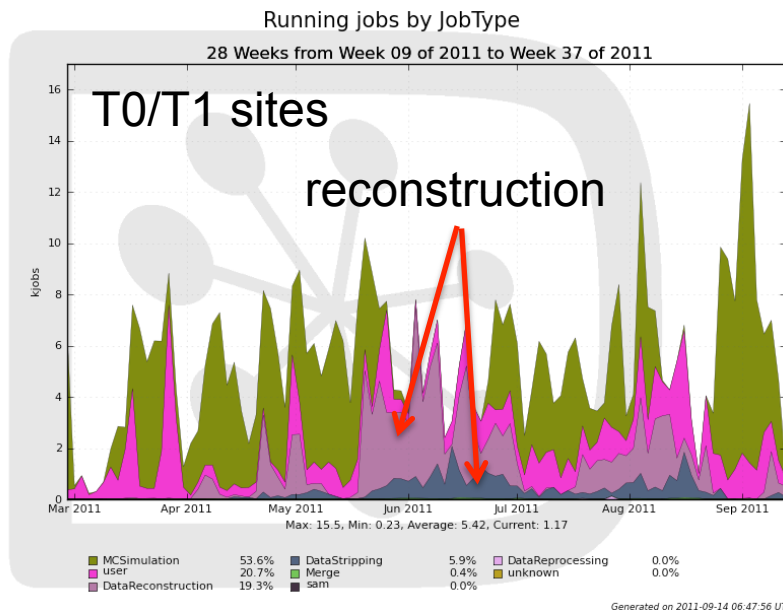
- No more T1 sites installed Cvmfs since last report
 - Cnaf has been testing on T2 instance
 - In2p3, Gridka & Sara are missing
- What about deployment status in WLCG?
 - The full gain for LHCb can be only achieved when all sites have migrated

- LHCb went down from 9 space tokens to 3
 - i.e. User, Disk and Tape
 - Simplification and more flexibility for data management
 - Increase the # of spindles / ST for throughput
- Implemented through “draining” of old STs
 - Done on all T1 sites
 - Especially for dcache sites the re-allocation of space to new ST is not straight forward
 - LHCb needs to trigger the site after cleaning

- Running jobs at all sites since March




- Running jobs T0/T1 and T2 sites
 - “holes” at T0/T1 files with MonteCarlo jobs
 - First quarter we had 50 % MC on T0/T1 and T2s
 - LHCb is coping well with “live” reconstruction
 - For reprocessing more resources needed -> see later



- Reconstruction activities after technical stop continue very smooth
 - E.g. Full Reconstruction jobs


Running jobs by FinalMinorStatus

144 Hours from 2011-09-08 00:00 to 2011-09-14 00:00 UTC



Running jobs by Site

144 Hours from 2011-09-08 00:00 to 2011-09-14 00:00 UTC



- Some problems with disk space before TS
 - Fixed by re-allocation from “old” STs
- In addition copies of DST files have been reduced from 7 to 4
- Still LHCb is falling short on disk space. Compared to requested space
 - T0: 400 TB short
 - T1s: 1,5 PB short

- Job runtime environment setup at Cern
 - For 48 core nodes overload of Disk I/O
- 3D streaming failing at IN2P3
 - Mitigated by using other T1 instances in Dirac
- Memory problems for stripping applications
 - Increased limit at T1 sites needed <- Many thanks

- File staging problems at Gridka
 - Chain of several problems
 - Disk pools in front of tape were increased for better performance
- FTS transfer problems
 - It would be beneficial if FTS could handle also the staging of files

- Starting by the end of the month
 - Whole 2011 dataset to be reprocessed
 - i.e. ~ 400 Tbytes of raw data
- Several changes in processing model
 - Using “ROOT” format for stripping jobs
 - Decrease of memory footprint to “normal” levels
 - (we still need special queues b/c of current data)
- Tentatively run until Xmas shutdown

- Dynamically attach/detach T2 sites to T1 storage whenever needed
 - Attach to help on processing raw
 - Detach in case of not needed anymore / problems
 - T2 sites can be anywhere, but once attached to T1.x not possible to re-attach to T1.y
- Tested successfully with 20 T2 sites

- New processing model will result in increased storage usage
 - Mitigation of problems implemented by “throttling mechanism” of staging
- In addition MC11 activities to be launched around fall
 - Monte Carlo will be done mainly on T2 sites
 - Because of new model less T2 sites available

- Very calm and smooth operation last half year in terms of data processing, but ...
 - ... it's always calm in the eye of the storm



We are here

- With reprocessing starting end of this month a much higher activity likely until end of year