LHCb Operations

Philippe Charpentier
CERN
On behalf of the LHCb Operations team

Activities since December

- Start of large MC productions for 2011 analysis
- Steady user activity (2700 jobs on average in parallel)
- Additional stripping of 2011 data started early March

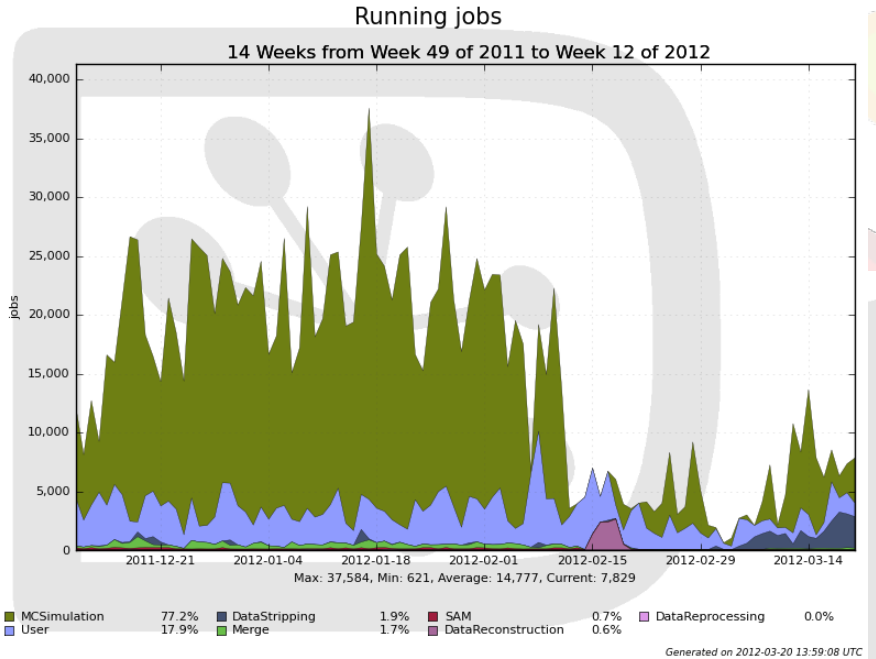

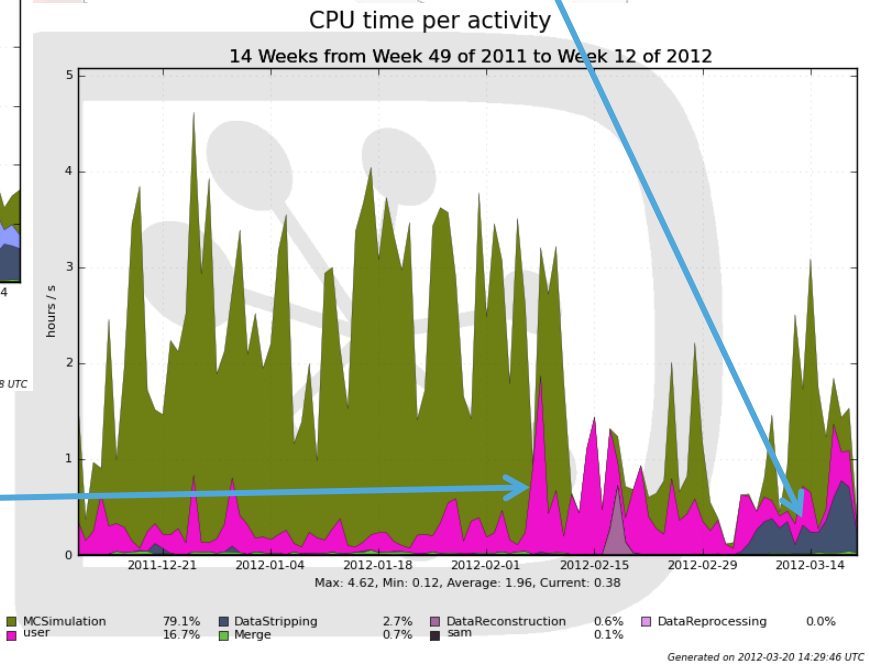

Moriond preparation

Site usage

- 92 sites used
  - Incl. non Grid sites
- 11.6% at CERN
- Many non-pledged sites
- Only 5.8% CPU wasted from failing jobs

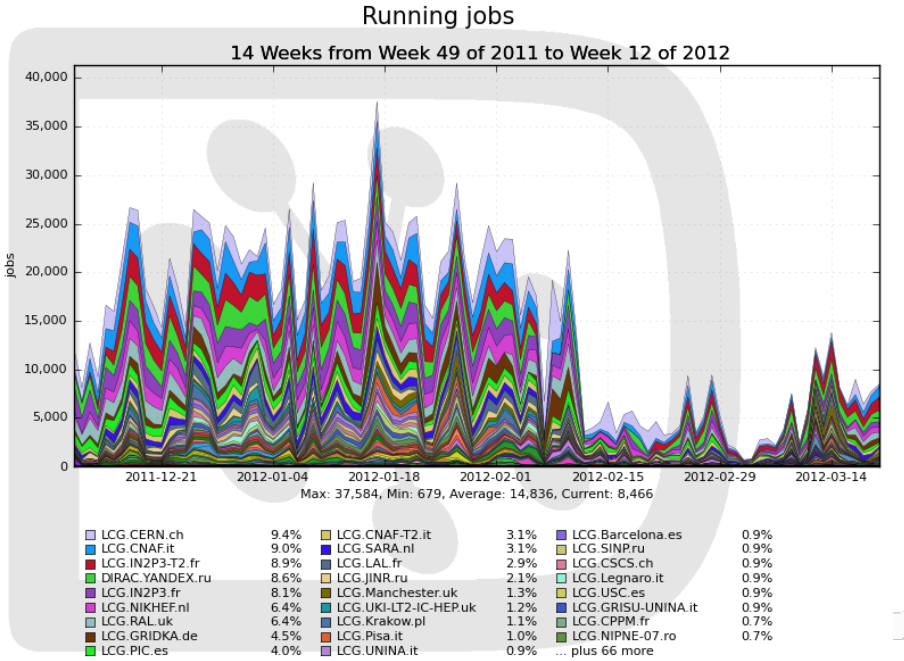

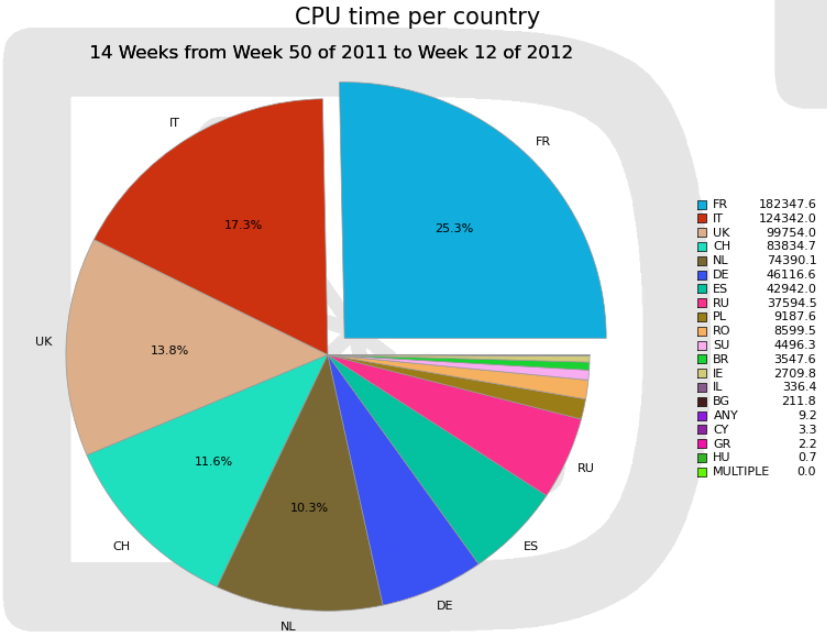

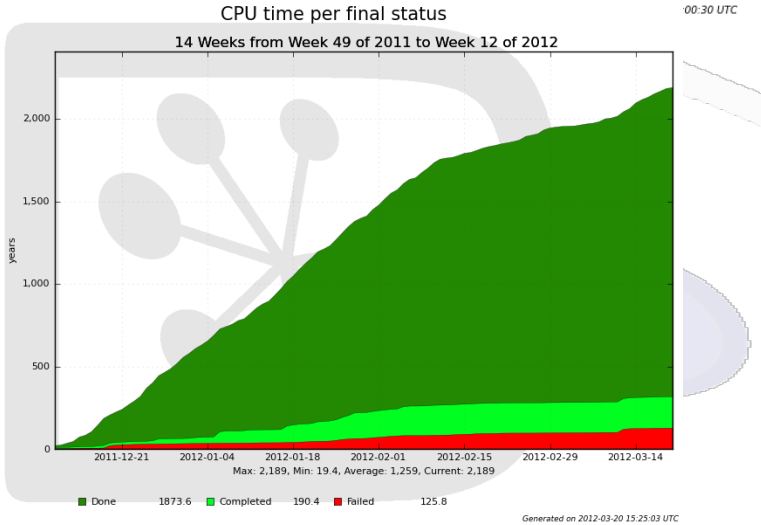

Currently: additional stripping

- Heavy recall from tape (SDST+RAW files from Winter reprocessing): 800 MB/s aggregated
  - Using an internal throttling based on cache size
- 60% done in 2 weeks
  - 10 more days expected + small tails as usual

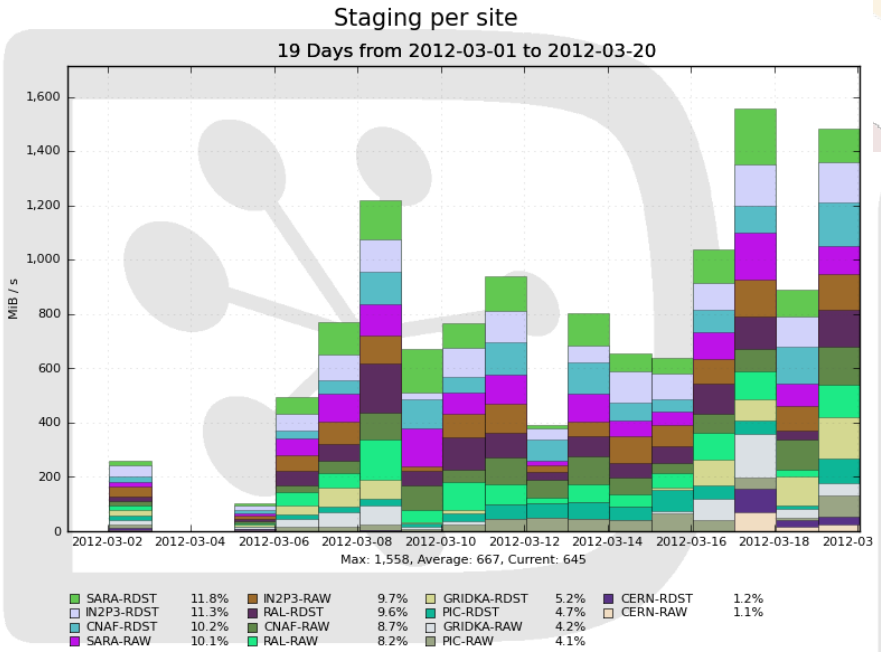

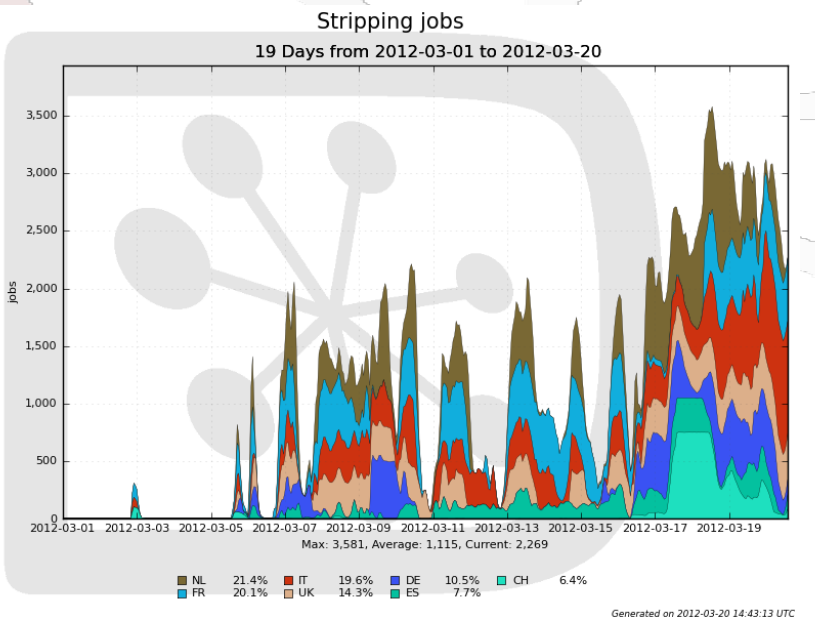

Storage usage

- 3.4 TB used on LHCb-Disk space tokens
- 1.4 TB free
  - Thanks to those sites that have already installed part of 2012 pledges!
  - We were critical at some sites but we got very positive attitude from sites

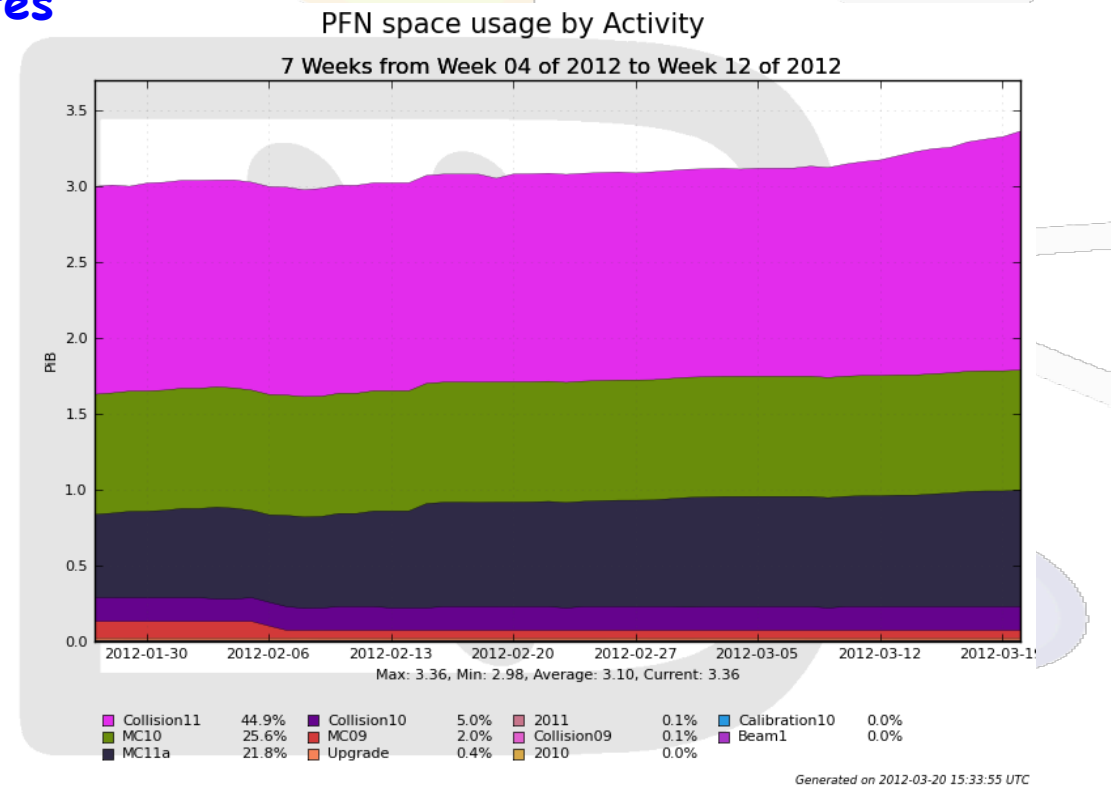

LHCb Operations

Issues and open points

- Critical on disk space: we had to fiddle with datasets and ask sites for advanced installation.
- Site outages: normal and expected, but sometimes really long, even when scheduled
- CVMFS improved dramatically the software distribution
  - Still waiting for KIT (imminent for weeks)
  - Some hiccoughs though from time to time with squids
- Tape recall and cache management:
  - We need to know the actual size of the cache read buffers
    - Not available with SRM for dCache
  - Better share disk servers with T0D1 storage classes
    - Much better throughput
  - On our side: implement pinning / unpinning by jobs rather than current arbitrary pinning

LHCb Operations

Plans for 2012

- Continuous, increasing (?) user analysis
- Somewhat steady load of MC production
  - 4 TeV beam MC (2012 conditions)
    - Comes as an addition to existing 3.5 TeV MC
  - Substantial MC for upgrade simulation
  - Running at all sites with low priority
    - real data processing and user jobs have preemption (Tier1 and Tier2s)
- Real data processing @ CERN + Tier1s
  - Essentially "quasi-online" reconstruction and stripping
  - No partial re-processing foreseen
    - Except maybe first few days/weeks
  - End-of-year re-processing/stripping starting (as in 2011) in September/October
    - In parallel with data taking and first pass reco/stripping
    - Using Tier2s linked to Tier1s (data download)
  - Followed by 2011 data re-processing/stripping

LHCb Operations

2012 data

- Extended physics program
  - Approved and supported by LHCC (Charm physics)
  - Additional 1.5 kHz of HLT output rate: 4.5 kHz total
- Additional load on CPU for reconstruction/stripping
  - Not a concern for quasi-online
  - Extend the time for re-processing
- Additional load on storage
  - Reduce number of replicas (already down to 4)
    - More aggressive reduction of older datasets (to 2 replicas)
  - Reduce number of archives for MC (from 2 to 1)
    - Tape shortage
  - Some data selections will not be included until 2013
    - "Data freezing" (i.e. not available for physics but recorded)
      - But events will be reconstructed though (too much overlap with hot physics for online streaming)
    - Open the selections during 2013
      - Additional disk resources needed

2012 real data workflow

- Additional production processing for Charm physics
  - Swimming (moving event's primary vertex and recomputing trigger and selection result)
  - Highly time consuming, but run on very few events
  - Output negligible in size (MDST format)
- Femto-stripping negligible in terms of CPU and storage, not in number of jobs
  - Calibration purpose, already there in 2011

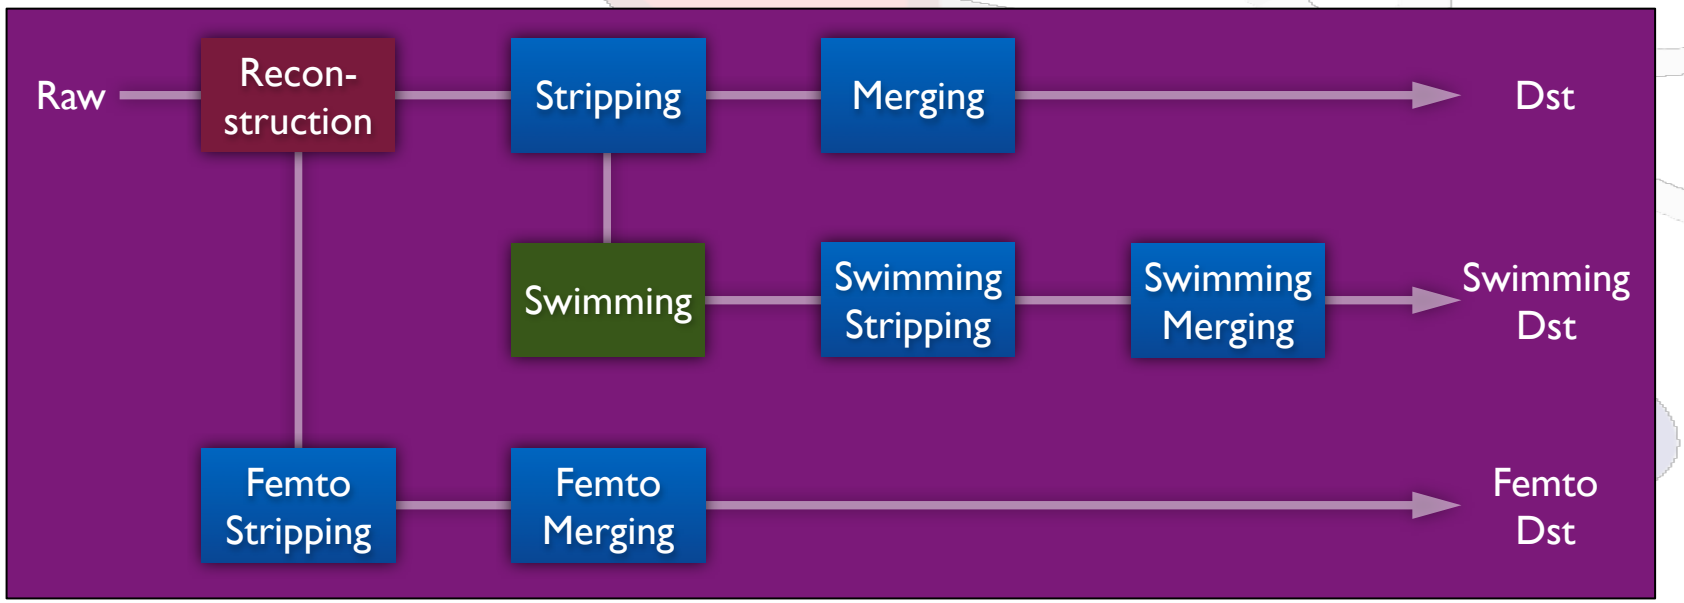

Infrastructure

- LHCb doesn't rely on any of the WN middleware installed
  - Uses LCG-AA distribution
    - Currently still gLite, testing EMI in progress
- 4 GB virtual memory now OK
  - Limit per process, not process group (5GB in this case)!
    - If possible only kill the offending process, not the job!
  - Applications are all fixed now (pb with ROOT buffers)
- Moving towards direct CREAM submission
  - On all Tier1s now, several major Tier2s
  - Not yet at CERN, fix required in Dirac
  - WMS used mostly for pilots running MC outside Tier1s
    - …but can use CREAM submission as well
    - No need for any additional installation at the site
- CVMFS welcome everywhere
  - Better for us, better for sites!
- SL6 being certified, but not foreseen in production during 2012 data taking
  - Requires MW python bindings for 2.7

LHCb Operations

Conclusions

- Quiet smooth running during last quarter
- Usual hiccoughs here and there
  - Site configuration issues
    - Queues, storage, CVMFS…
  - Application failures (very small indeed)
- Concerns with disk installation
  - Not 100% needed in April
  - We cannot wait for November for an increase of disk space
    - Progressive increase from now to October OK though
    - If not possible for some sites, we ask for installation of next years' pledge (2013 in November 2012)