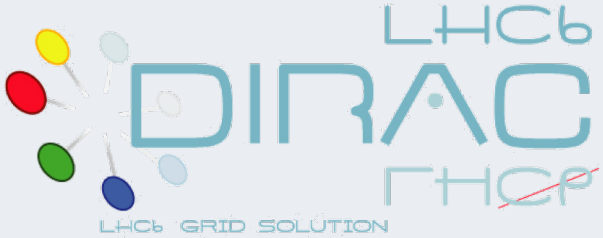

LHCb and DIRAC

Philippe Charpentier
CERN

What's new since DUW in Warsaw?

- In fact not so much, unfortunately!
  - Mostly due to manpower limitation
  - … and pressure from operations

- Outline
  - LHCbDirac operations in the last year
  - What still needs to be commissioned
  - What still needs to be developed

- Apologies
  - Many slides are exactly the same as last year!
    - This indicates LHCb is struggling with manpower

LHCb activities in the last year

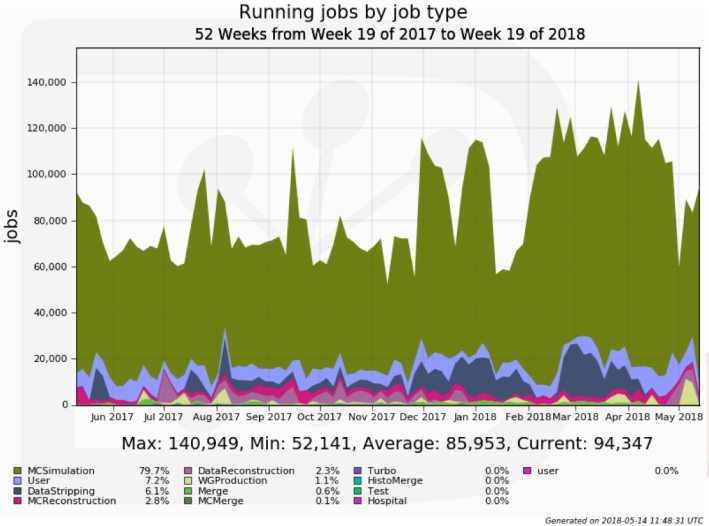

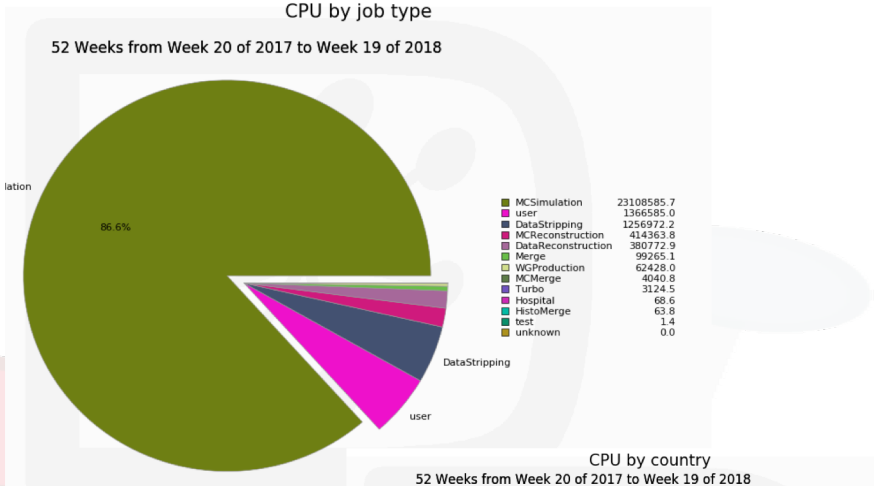

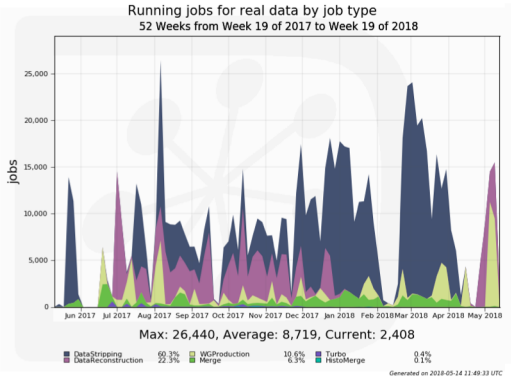

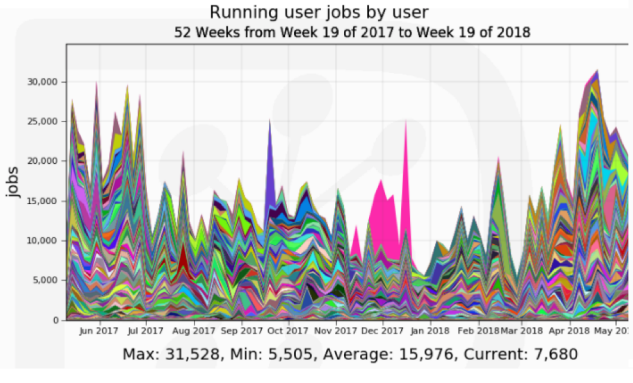

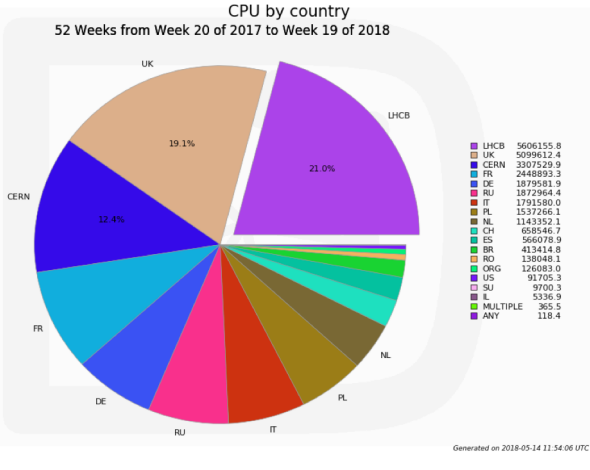

Handling the CNAF outage

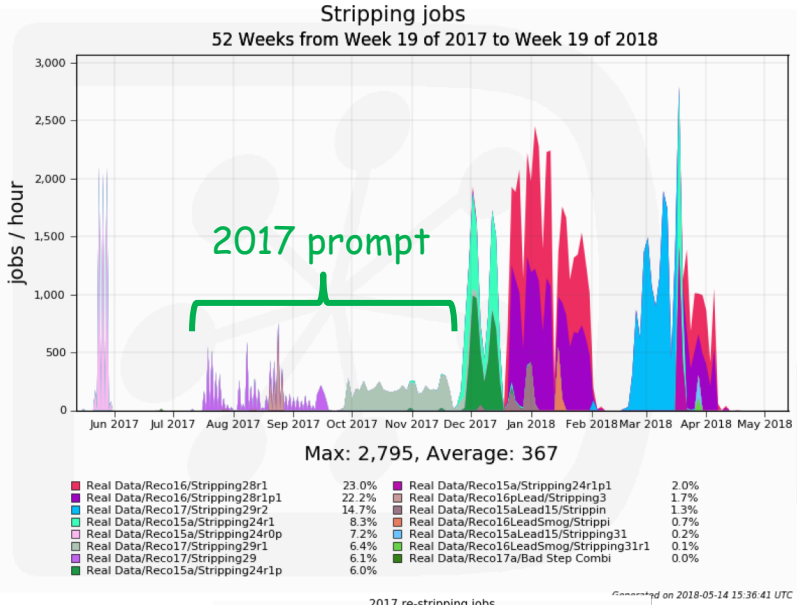

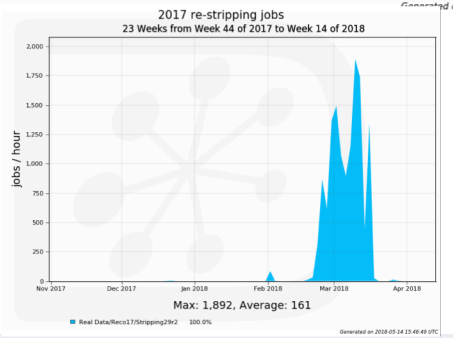

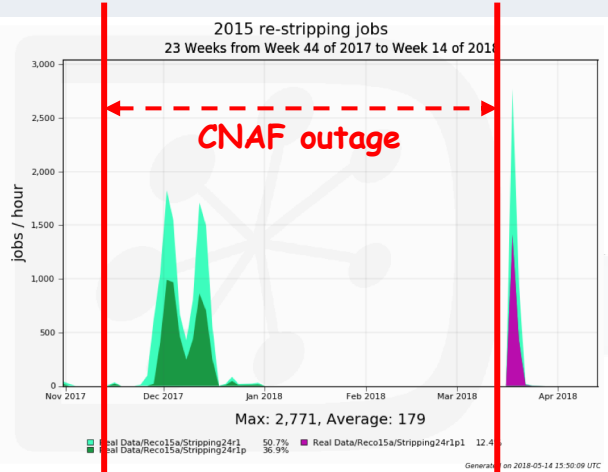

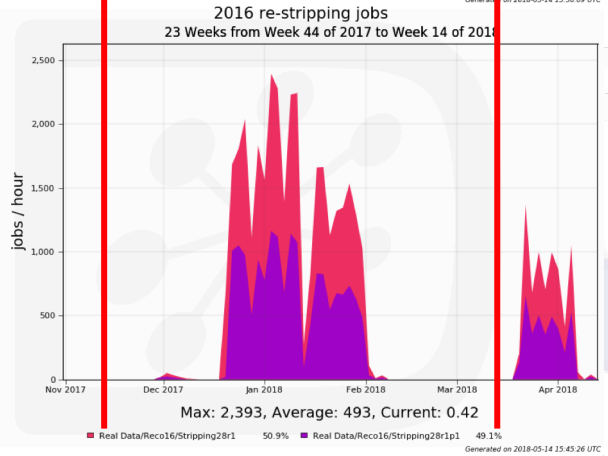

Usage of the LHCb HLT farm

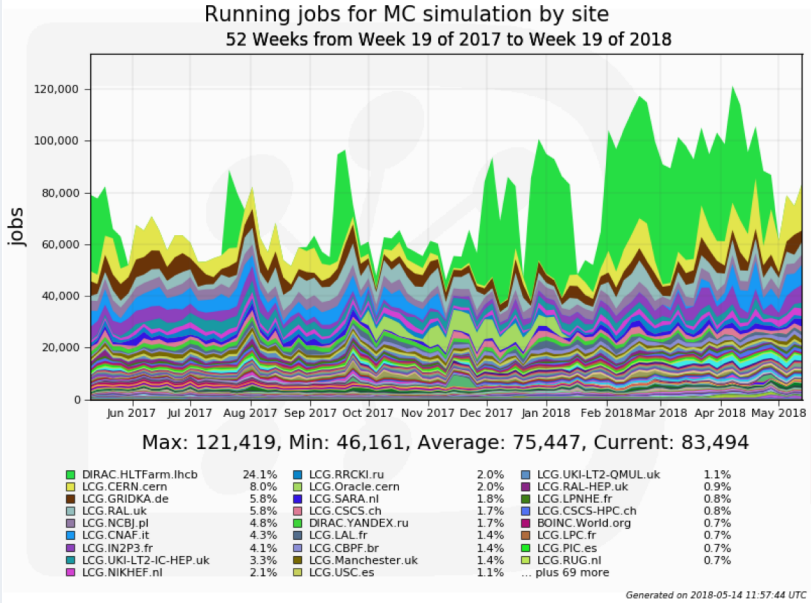

- Used during technical stops
- Used even during data taking as background task

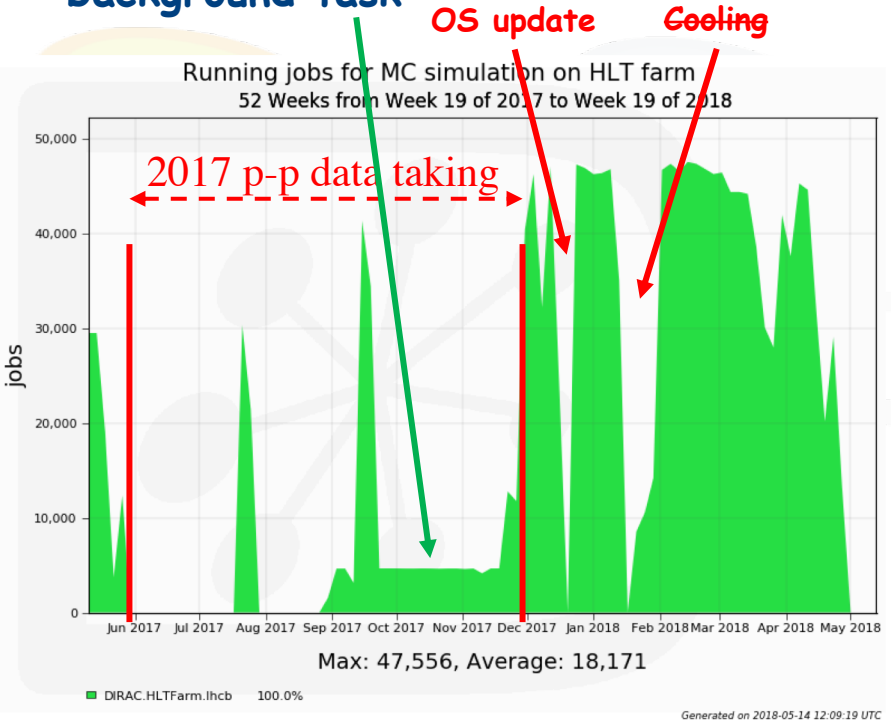

- HLT farm uses Pilot 3.0
- Jobs are stopped gracefully when needed

Activities: comments

- Always dominated by MC simulation
  - Heavily using the HLT farm
    - Even during data taking (background of MC tasks in parallel with HLT1/HLT2)
    - Top CPU provided integrated over the year
    - Graceful stop using signalling to Gaudi (stop after current event)
    - Heavy load on CERN EOS storage as the HLT farm has no external access
    - Heavy load at CERN for MCReconstruction jobs (done on the CERN batch)
- Steady user load (average ~16000 jobs, was 5000 last year!)
  - 482 unique users have used LHCbDirac in the year (was 470 last year)!
- Data reconstruction and stripping campaigns
  - Took place when processing (application and calibration) was ready
    - Stripping started quite late in 2017 (mid-July)
  - Huge full Run2 re-stripping campaign during the YETS
    - Due to a bug in the code of the previous stripping passes
    - Main load was on data management for pre-staging (see earlier talk on DM in LHCb)

Possible improvements

From 2017 talk

- Less deterministic site assignment
  - Currently sites are assigned at job creation (mesh processing)
    - Impossible to react to site downtime or overload once jobs are waiting
  - What about late site binding, e.g. at matching time
    - Assign jobs to sites hosting the data
    - At matching time, check whether there are jobs for the originating site
      - If yes, get it
      - If not, use the "mesh processing" information to check other task queues
    - This would allow much more flexibility and reactivity
- Task queue agnostic pilot submission
  - Use feedback from pilots for throttling pilot submission
    - Rather than use task queue information
    - Anyway pilots don't match jobs that they have been submitted "for"
    - "VAC-like" model for site director
    - Being worked on by Andrew
  - Would allow a better pilot dissemination
    - And faster pilot submission

Code exists
Not commissioned yet

What is left from 2016~~7~~?

From 2017 talk

- Bulk submission is not yet there
  - This is a MUST: user jobs, MC productions
    - Currently ganga playing tricks with input sandbox uploaded to User storage
      - Pb if SE is overloaded or in downtime
    - Improve submission time at the client level
      - Mostly for user jobs
- Pilot filling mode is far from optimal
  - No maintenance of TimeLeft utility
    - MJF usage still very limited (also at site level)
- Multiprocessor jobs still not in use
  - Not a must as LHCb jobs are not too much memory-hungry, but still…
- Split jobs and tasks statuses
  - This is VO-dependent, but should be implemented
  - Task final status is not necessarily that of the job
    - A job may be failed but the task successful and vice-vers (more rare but still exists)
- "Completed" job status should be changed !!!

Ready To be commissioned

Filling mode disabled

Would help on HLT farm

What are we expecting to come soon?

From 2017 talk

- Pilot logging is more and more eagerly expected
  - Large fraction of jobs without any pilot logfiles (not using a CE)
    - And even on CE, the lifetime of logfiles is short (a week)
- New FTS system still not commissioned
  - Hopefully will be more reliable
  - Was already developed 2 years ago, just after Ferrara WS…
- Extend multi-protocol usage
  - Still very limited now
  - Careful deployment as there may be site-related issues
- Usage of priorities for jobs should be revisited
  - Better documentation first: how does it work is far from clear
  - In LHCb we have hard time to get top priority jobs running and use MCSimulation for filling up sites
  - Need to better control user jobs and shares

v6r20

xroot

???

Conclusions

- LHCb is running successfully a lot of workflows through LHCbDirac
- More and more new platforms are popping up
  - Usage of Pilots 3.0 (see Andrew's talk) should be generalized
- Some improvements / simplifications may help scaling
  - Late site binding
  - Vac-like pilot submission
  - Bulk job submission
- Long-standing developments should be included ASAP
  - Pilot logging
  - New FTS system
- All is a matter of lack of manpower
  - A lot of development for DIRAC and LHCbDIRAC done in LHCb
  - Most developers also participate in certification and operations
  - … and we need to keep an efficiently running system!