

LHCb Computing



Stefan Roiser

LHCC WLCG Referees

21 Feb '17

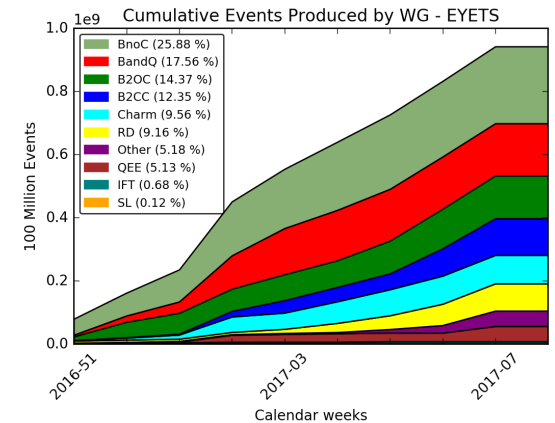




Operations – Job metrics



- Continued running at very high CPU/wall time efficiency
- Job success rate at 93 % slightly lower than usual
 - Main reasons: Heavy ion reconstruction running out of time → run on VMs@CERN
 - + ~ 3% user jobs, fix CVMFS libraries with afs dependencies, SRM overload

HLT usage for offline during EYETS



- EYETS program for MC production progressing very well
 - ~ 950 Mio events produced since Xmas (to compare with 2.5 Billion events / year)
- New tool available to monitor simulation resources by physics working group
 - Plots on work, disk usage, events produced, event types, production types

2017 pledges & 2018 requests

- Change of strategy for data stripping campaigns
 - Skip “incremental stripping” campaigns
 - 2015 data re-stripping during ‘16/’17 EYETS becomes incremental stripping
 - 2016 re-stripping of ‘16/’17 EYETS replaces ‘16 synchronous stripping on disk



- Many thanks to funding agencies for additional efforts with 2017 pledges !!!!
 - Still short on pledged disk but ...
 - change of strategy for stripping campaigns will allow to accommodate the shortage of pledged disk space

- 2017 resources re-submitted to CRSG with 20 % increase mainly for disk
- 2018/19 resources inline with constant budget increase of resource requests for cpu/disk/tape on top of 2017



20 Sep '16

LHCC Referees


- 2018 requests assumed with “flat budget increase”
 - CERN provisionally announced no resource increase in 2018

Run 2 Optimizations – Software


- Application improvements via backports of Run3 upgrade:
 - RICH reconstruction improvements by 20-30 %
 - Re-write of tracking Kalman filter fit improves by factor 2
- Software stack built with SSE 4.2
 - Upgrade concentrates on vectorization → more backports possible
 - ~ 4% of grid WNs older than 8 years → adapt grid submission
- Fast simulation
 - LHCC Poster →
 - Customized use of sub-detectors
 - Only simulate excl. final state of interest

Fast simulation options in LHCb from ReDecay to fully parametrised


Dominik Müller¹ and Benedetto Siddi²
¹The University of Manchester, ²Università degli studi di Ferrara/CERN on behalf of the LHCb collaboration




MANCHESTER
1824
The University of Manchester



LHCb
LHCb



INFN
Istituto Nazionale di Fisica Nucleare
Sezione di Ferrara

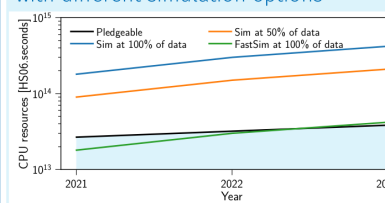


CERN

Introduction

- ▼ For LHCb upgrade will need to simulate more complex events and in a bigger amount
 → estimated resources will not be enough (see plot).
- ▼ Simulation time completely **dominated by detector transport** ($O(95\%–99\%)$!)
- ▼ Increased CPU resources and new Geant versions will not suffice
 → dedicated fast simulation options necessary
- ▼ Existing fast simulation options in LHCb:
 - ▼ Customize used subdetectors to the analysts needs ($O(40\%–90\%)$ reduction.)
 - ▼ Only simulate the exclusive final state of interest ($O(95\%–99\%)$ reduction.)
- ▼ Currently in development: shower libraries for the calorimeters, ReDecay and Delphes.
 - ▼ To be implemented in the official simulation of LHCb, Gauss [M. Clemencic et al. 2011 J. Phys. Conf. Ser. 331 032023]

Estimated CPU resources needed by LH with different simulation options



ReDecay


Motivation

- ▼ Many LHCb analyses involve exclusive decays e.g. $D^0 \rightarrow K^- \pi^+$.

Parametric simulation using Delphes

Motivation

- ▼ Feasibility studies for physics analysis.




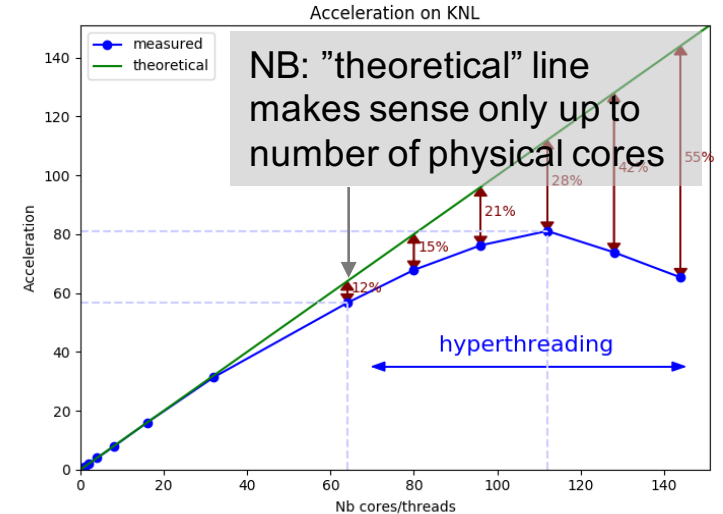
21 Feb '17

Run 2 Optimizations – Infrastructure & Computing Model

- 1/3 of Turbo data parked on tape for later analysis
- Optimize number of streams for Turbo data for faster analysis I/O access
- Reduction of data disk replicas and simulation data
- 2016 stripping will replace previous version on disk
 - Analysis work needs to adapt accordingly
- Introduction of “working group productions” ramping up
 - Production of further analysis formats handled centrally by production team

Run 3 upgrade – Software

- First results of multi-threaded task-parallel “mini reconstruction” available:
 - Very good scaling with number of physical processor cores
 - Big improvements in memory usage
- Next steps:
 - Performance improvements via code vectorization and event model re-organisation
 - Integrate more algorithms
 - Demonstrator of technical changes by March '17



Run 3 upgrade – Training

- The vast majority of the LHCb code base needs to be touched for the upgrade version of the framework
- → involvement of sub-system developers is mandatory
- Approached by training and hands-on hackathons

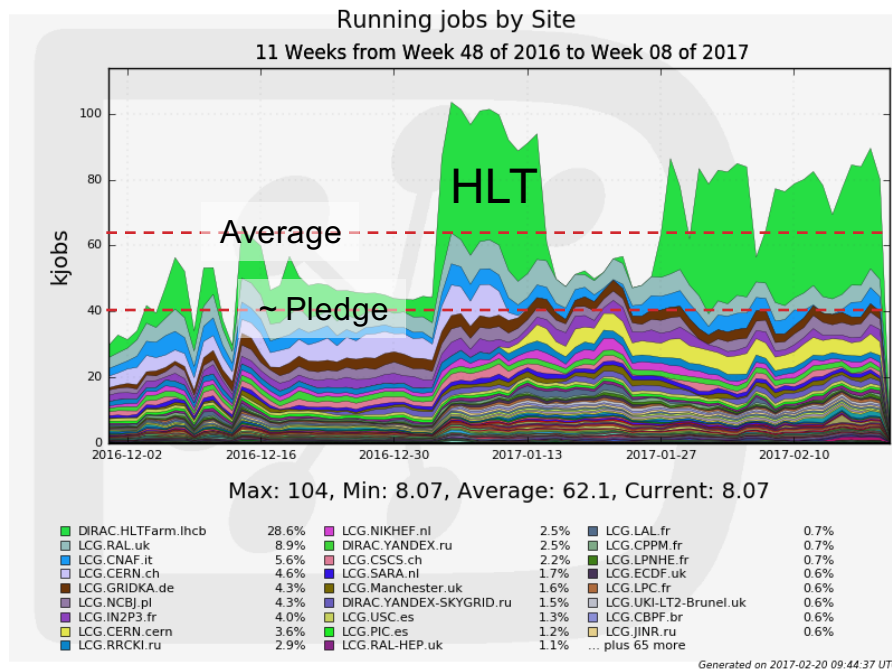


Summary

- Very good usage of resources in 2016 + EYETS
 - HLT farm is a big boost to the CPU resources
- Several Run 2 resource optimizations and mitigation measures are in place, more to come
 - Parking of Turbo data on tape for later analysis
 - Software optimization – code vectorization, fast simulation
 - Replacement instead of parallel usage of stripping versions on disk
- First encouraging results of Run 3 upgrade activities
 - Multi-threaded, task-parallel framework functional and performant


Executive Summary

LHCb – Operations & Run 3 Upgrade




- Very good usage of HLT farm for simulation workflows since Xmas break
 - When available ~ doubles LHCb CPU resources

- First version of multi-threaded and task-parallel framework for Run 3 functional and performant
 - Very good scaling with physical processor cores and big improvements in memory usage



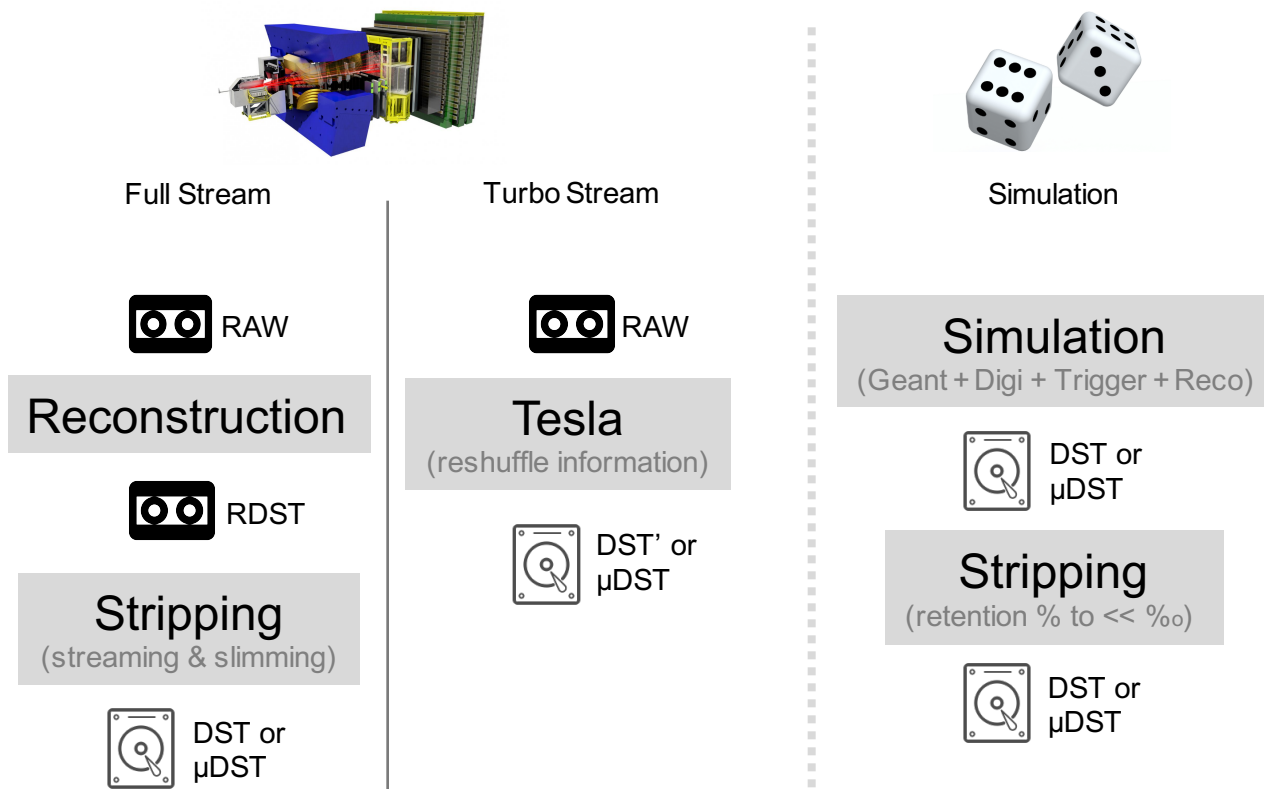
LHCb – Run 2 Resource Optimization

- Run 3 upgrade software backported to Run 2 applications
 - Parts of reconstruction and tracking with up to 30 % improvements
- Software stack build with SSE 4.2 vectorization enabled
 - Focus for upgrade software stack → more backports possible
- First fast simulation options available in production
 - Eg. Customized use of sub-detectors
- Changes in data processing model
 - Reduction of disk replicas, skipping of “stripping” campaigns
- Parking of 1/3 of Turbo data on tape for later analysis



Backup

Main LHCb processing workflows



File Type	Event Size (kB)
RAW	50
RDST	40
Data DST	65
Data μ DST	10
MC DST	200
MC μ DST	60