

LHCb Software & Computing

Concezio Bozzi

WLCG/LHCC Referees


4 June 2019



Operations

Running jobs by JobType


13 Weeks from Week 07 of 2019 to Week 21 of 2019



89 % of running jobs are MCSimulation @ > 98 % CPU/Wall time efficiency

Max: 124, Min: 59.4, Average: 103, Current: 89.1


MCSimulation	78.8%	MCFastSimulation	10.5%	user	6.0%	MCFastSimulation	0.7%	Hospital	0.2%	MCMerge	0.1%	WGProduction	0.1%	DataReconstruction	0.0%	HistoMerge	0.0%	unknown	0.0%
--------------	-------	------------------	-------	------	------	------------------	------	----------	------	---------	------	--------------	------	--------------------	------	------------	------	---------	------



Generated on 2019-05-29 12:00:32 UTC

Running jobs by Site

13 Weeks from Week 07 of 2019 to Week 21 of 2019




Max: 124, Min: 59.4, Average: 103, Current: 99.6

DIRAC.HLTFarm.lhcb	38.7%	LCG.NIKHEF.nl	1.4%	LCG.JINR.ru	0.6%
LCG.CERN.cern	12.0%	LCG.LAL.fr	1.3%	LCG.Beijing.cn	0.6%
LCG.RAL.uk	7.7%	LCG.PIC.es	1.1%	VAC.Cambridge.uk	0.6%
LCG.CNAF.it	7.5%	LCG.UKI-LT2-IC-HEP.uk	1.1%	LCG.Manchester.uk	0.6%
LCG.GRIDKA.de	3.4%	LCG.DESYHH.de	1.0%	LCG.SARA.nl	0.6%
LCG.NCBJ.pl	2.9%	LCG.UKI-LT2-QMUL.uk	0.9%	LCG.CPPM.fr	0.5%
LCG.IN2P3.fr	2.4%	LCG.CBPF.br	0.9%	LCG.LPNHE.fr	0.5%
LCG.CSCS.ch	1.6%	LCG.BEER.cern	0.8%	LCG.GLASGOW.uk	0.5%
LCG.RRCKI.ru	1.6%	LCG.USC.es	0.7%	...	plus 59 more

Generated on 2019-05-31 15:21:44 UTC

CPU days used by FinalMajorStatus


13 Weeks from Week 08 of 2019 to Week 21 of 2019



Generated on 2019-05-31 15:47:21 UTC

CPU days used by FinalMajorStatus

13 Weeks from Week 08 of 2019 to Week 21 of 2019





Generated on 2019-05-31 15:48:29 UTC


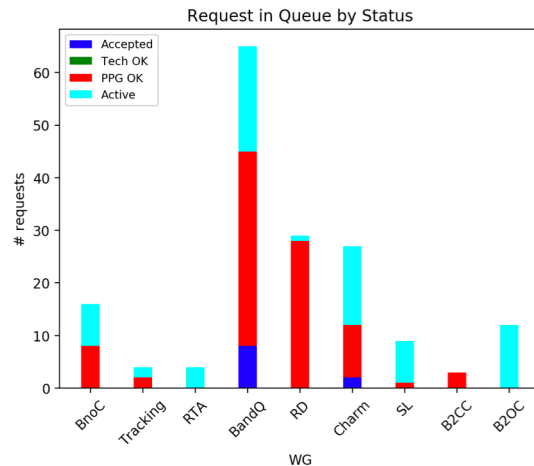

Data productions in 2019

Legacy stripping campaigns for all Run1 and Run2 data under way.


- Some over-timing observed and fixed, incremental campaigns for Run1 and 2018 data under validation
- Then, full re-stripping of 2015 and 2016 data due to ECAL re-calibration
- Then, incremental re-stripping of 2017 data
- Plan to finish processing of data by end 2019
 - ~8 weeks for each Run 2 data taking year and Run1 combined
 - running time is dominated by staging from tape



MC productions




- 2017 and 2018 production at full steam
- New Sim10 simulation framework under validation
- Large sample produced with fast simulation, now fast/full ~50:50
- Clear dependence of produced number of events on the amount of fast simulation
- Most productions are filtered before being written to disk




Opportunistic Computing

- Resources not pledging to WLCG continue to be successfully used
- Significant usage of BEER at CERN and CERN cloud
- Equivalent to another T1
- [HLT farm not shown]



HPC centers

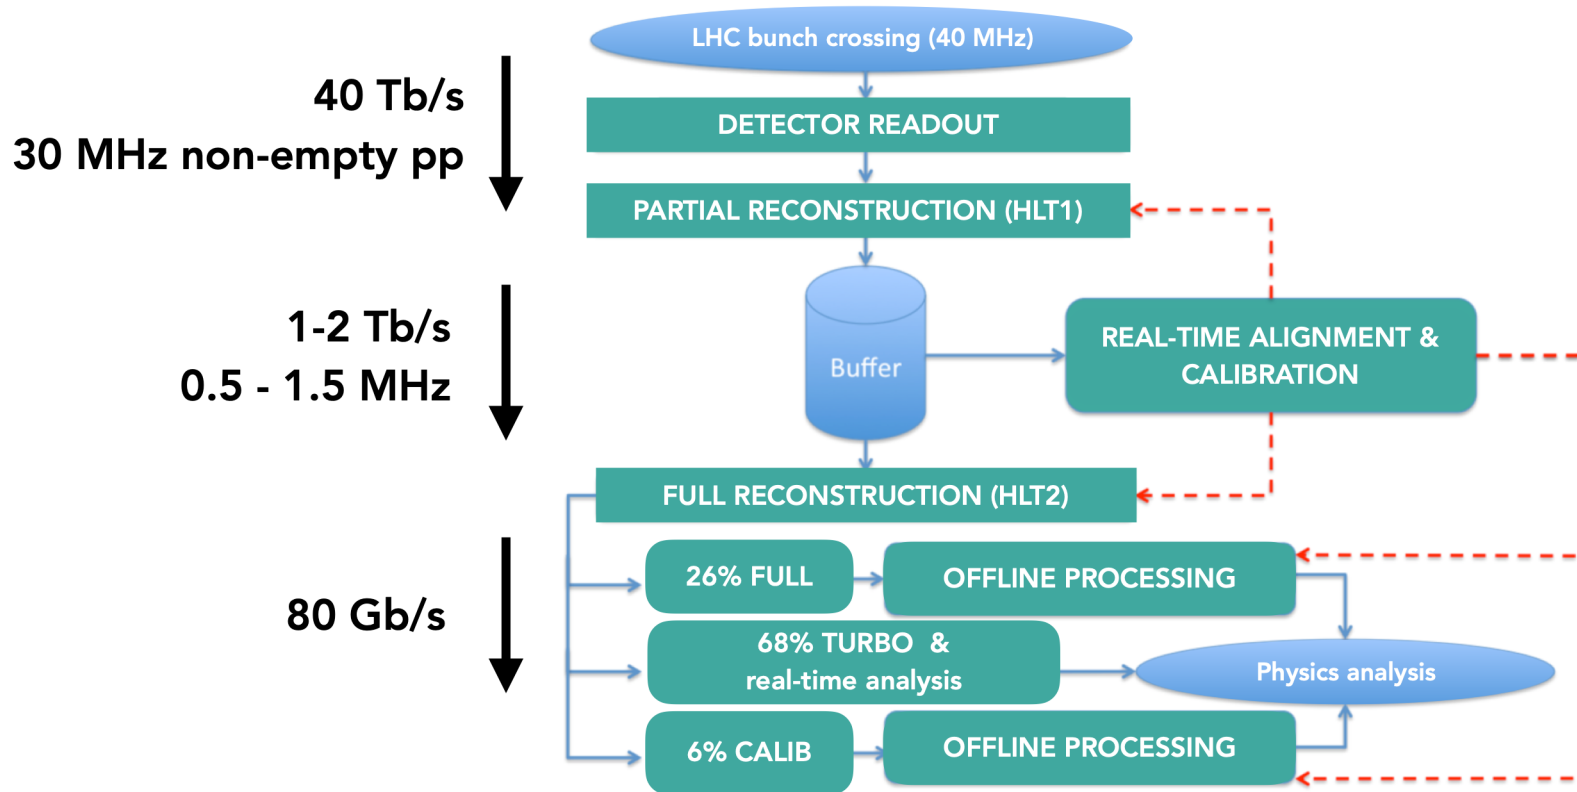
- Currently two sites
 - CSCS in Switzerland
 - pledging to WLCG
 - SantosDumont in Brazil
- LHCb successfully uses HPC centers when
 - Worker nodes have external connectivity
 - CVMFS is available
- Working on relaxing these two requirements
- Usage of non-x86 architectures is and will remain problematic
- [see also Simone's talk]



Run 3 Upgrade

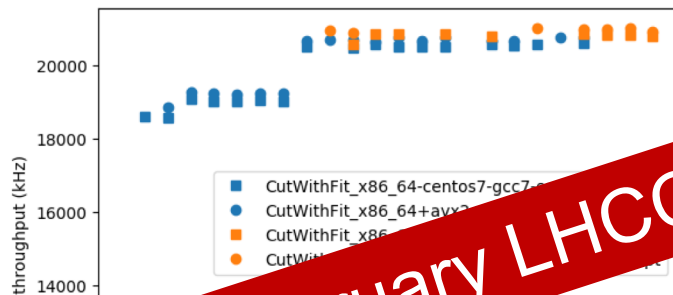
(see talk by V.V. Gligorov at [in-depth review](#) for a complete picture)

Reminder of Run3 data processing

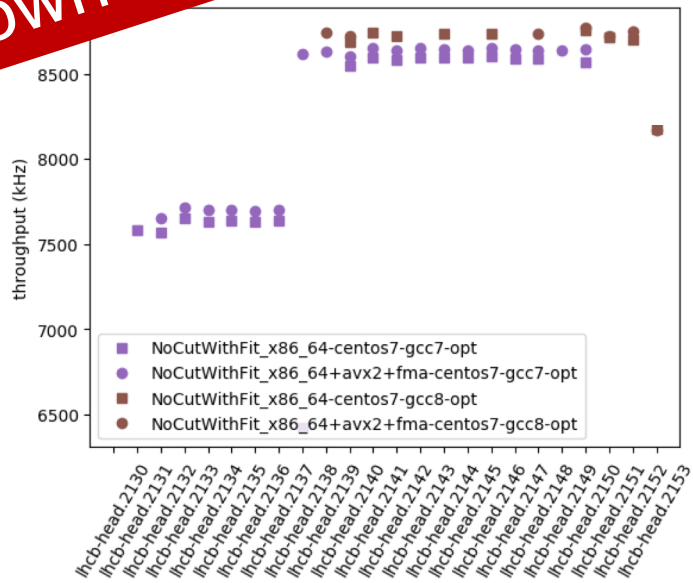


Upgrade Software Status


- New real-time analysis (RTA) project in charge of the HLT and reconstruction/selection
- Strict collaboration with computing project, which takes care of software engineering aspects
 - New scheduler, framework improvements, better memory management, new interfaces to conditions data and detector description (DD4HEP)
- Steady improvements in the past months for HLT1
- More improvements in the pipeline
 - Refactoring of algorithms
 - Working on data structures to enable vectorization
- Moving focus on HLT2 application



Slide shown at February LHCC

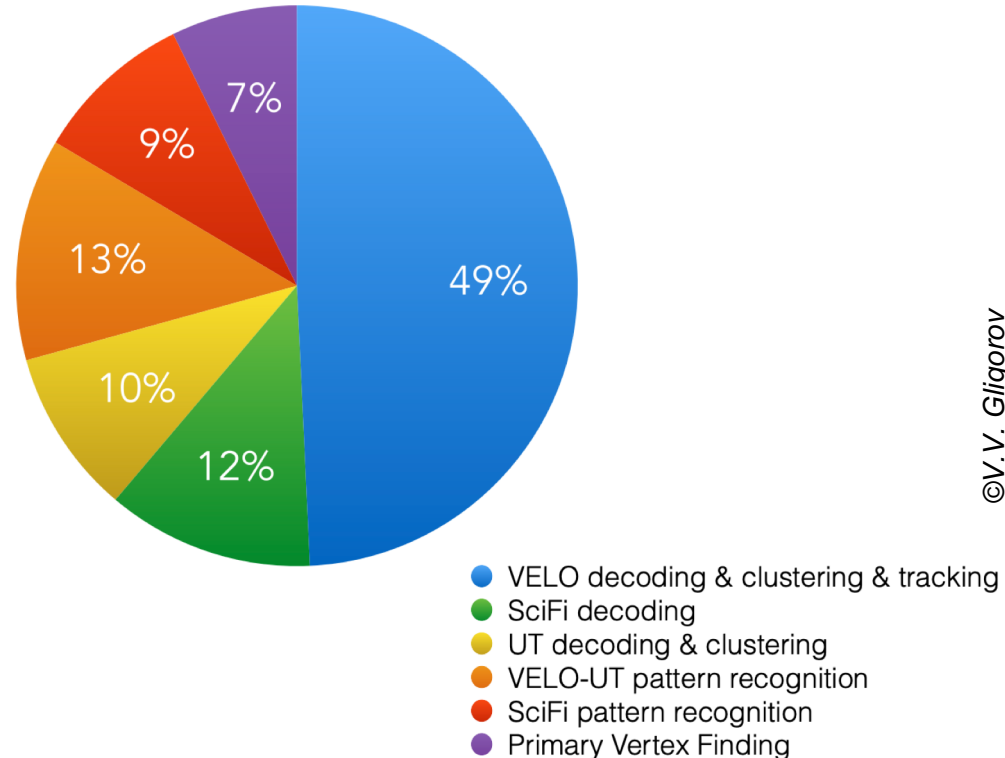


Status of partial reconstruction (full physics programme)



Status of HLT1 reconstruction (full physics programme)

- «Psychological» 30MHz threshold passed
- Some bits still missing, but global picture should not change too much
 - Muon reconstruction
 - Kalman fit (VELO tracks only)
 - Trigger selection logic
- Further optimizations being worked on
 - e.g. VELO clustering on FPGA in back-end electronics



Summary

- Computing operations continue at high usage and efficiency
 - on average 2x above WLCG pledged resources
 - MC production requests dominates CPU work
 - Working on expanding the pool of “opportunistic” resources (e.g. HPCs...)
- Run 3 upgrade work progressing well
 - Modern software engineering techniques paid off
 - usage of vector registers
 - rearrangement of data structures
 - refactoring/rewriting of algorithms
 - Factors gained in processing speed with almost no loss in physics performance