

LHCb Software & Computing

Concezio Bozzi

WLCG/LHCC Referees

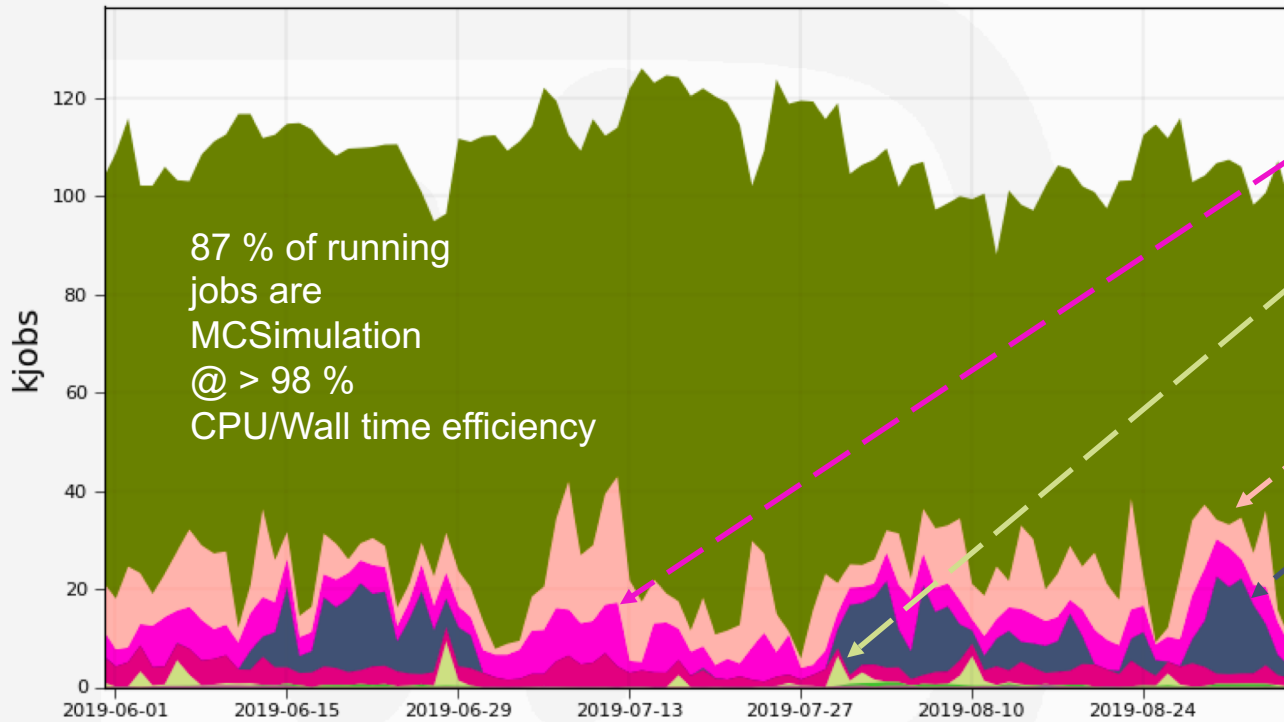
10 September 2019



Operations


Running jobs in all sites

13 Weeks from Week 21 of 2019 to Week 35 of 2019



Max: 126, Min: 88.1, Average: 109, Current: 98.0


MCSimulation	78.1%	MCRreconstruction	2.9%	MCMerge	0.0%	unknown	0.0%
MCFastSimulation	8.5%	WGProduction	0.5%	MCRreprocessing	0.0%		
user	5.2%	Merge	0.2%	DataReconstruction	0.0%		
DataStripping	4.6%	Hospital	0.1%	test	0.0%		



Generated on 2019-09-06 12:09:33 UTC

Running jobs in all sites


13 Weeks from Week 21 of 2019 to Week 35 of 2019



Max: 126, Min: 88.1, Average: 109, Current: 97.5


DIRAC.HLTfarm.lhcb	40.5%	LCG.NIKHEF.nl	1.6%	LCG.Beijing.cn	0.8%
LCG.CERN.cern	8.7%	LCG.RRCKI.ru	1.6%	LCG.Manchester.uk	0.7%
LCG.RAL.uk	7.4%	LCG.CBPF.br	1.1%	LCG.JINR.ru	0.6%
LCG.CNAF.it	6.3%	LCG.UKI-LT2-IC-HEP.uk	1.0%	LCG.USC.es	0.6%
LCG.IN2P3.fr	3.9%	CLOUD.YANDEX.ru	1.0%	LCG.SARA.nl	0.6%
LCG.GRIDKA.de	3.6%	LCG.DESYHH.de	0.9%	VAC.Cambridge.uk	0.6%
LCG.NCBJ.pl	2.7%	LCG.PIC.es	0.9%	LCG.CPPM.fr	0.6%
LCG.CSCS.ch	1.7%	LCG.RHEA.cern	0.9%	VAC.Manchester.uk	0.5%
LCG.LAL.fr	1.6%	LCG.UKI-LT2-QMUL.uk	0.8%	... plus 69 more	

CPU days used in all sites
13 Weeks from Week 21 of 2019 to Week 35 of 2019



Generated on 2019-09-06 12:05:12 UTC

CPU days used in all sites
13 Weeks from Week 21 of 2019 to Week 35 of 2019




Generated on 2019-09-06 12:03:27 UTC


Opportunistic resources and HPC centers

- Using two HPC sites that are complying with LHCb requirements* * Network connectivity and CVMFS
- Core-hours allocated (with ATLAS and CMS) on the Marconi KNL cluster in CINECA (Bologna)
 - Validating Gauss Multi-Processing to run on KNL nodes
 - Setting up DIRAC
- Sizeable contributions from opportunistic resources

Jobs in HPC centers




Jobs on opportunistic resources





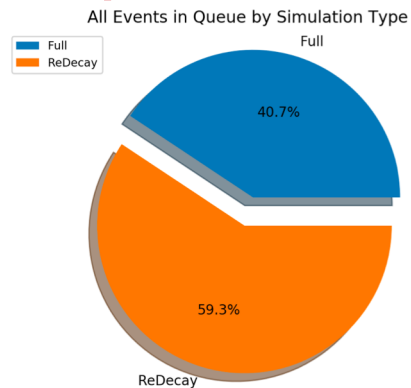
Data productions in 2019

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

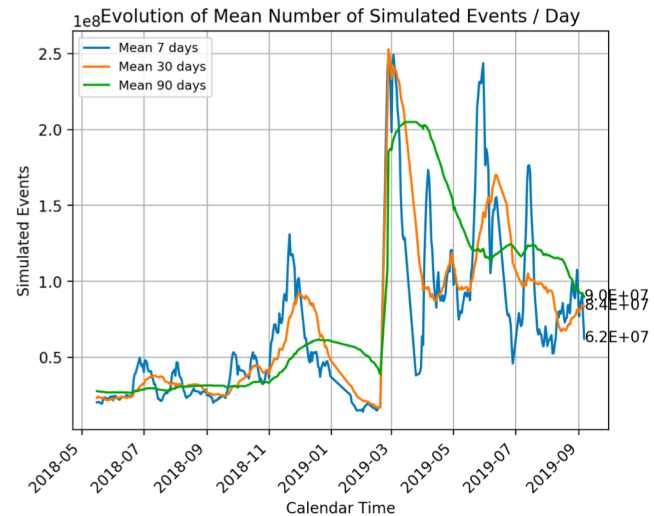



- Currently validating re-stripping of 2015 and 2016 data
- Then, re-stripping of 2017 data
- Plan to finish processing of data by end 2019 (tight)

MC productions



- 2017 and 2018 production at full steam.
- Some production with Run1 conditions still going on
- Fast:full simulation ~ 40:60
- Most productions are filtered before being written to disk



Containers

- Usage of containers (Singularity) recently implemented in DIRAC
- Essential for long-term preservation of production workflows
- First application: reconstruction of Run1 MonteCarlo
 - Running SLC5 executable on CentOS7
 - Currently implemented at CERN. Tier1 and Tier2 sites are being asked to setup and deploy containerised workflows


Towards Run3

Disclaimer: software trigger is covered by the Real-Time Analysis project and not discussed here. See <https://indico.cern.ch/event/776231/>

Simulation

- Preparing «Sim10»: significantly faster (~40%) than current «Sim09», due to Geant4 v10 and RICH physics optimisation
 - Currently under validation by the physics working groups
- Significant progress achieved on «Gaussino» simulation framework
 - complies with and profits from functional/multi-threaded Gaudi
 - interfaced with multi-threaded Geant4
 - Included in nightly builds and PR tests
 - tracker-only simulation can be run on minimum-bias events

Sim10 validation jobs



Gaussino nightly build

lhcb-gaussino - build: 568 (2019-09-06) [Compare with previous build](#) [Compare with other slot](#) [Browse files](#)

Test build of Gaussino and future Gauss available on: [cvmsf](#)

Project	Version	x86_64-centos7-gcc8-opt	
		Completed at 00:44:45	Platform info
LCG	96		
Gaudi	master	build	tests
Geant4	v104r2-mt	build (980)	tests (0)
LHCb	modern_sim_components	build	tests (3)
Gaussino	lhcb_master	build (9)	tests (1)
Gauss	Futurev2	build (2)	tests (2)
DBASE	None	build	

Core SW, offline analysis

- Core software: working on
 - Detector description: validation of upgrade geometry, migration to the DD4HEP toolkit
 - conditions access: migration to a gitLab-based conditions DB completed, now algorithms need to be refactored for efficient conditions access
 - Integration of Gaudi within the online environment: efficient reading of multi-packet events and interfacing with online control system
- Offline Analysis Task Force (OATF) established
 - System design until the end of 2019, implementation afterwards
 - Major task: overhaul of user analysis workflow
 - Centralised production instead of «chaotic» user submissions
 - Building on the experience gained with «working group productions» in the past few years

Offline computing requests for 2021

- Preliminary requests have been sent to the C-RSG
- Same model as in LHCb Upgrade Computing Model TDR
 - Minor adjustments following latest prescriptions on instantaneous (1×10^{33}) and integrated (3fb^{-1} baseline, 7fb^{-1} contingency) luminosities
 - Contingency used for tape requests only
 - Large increase

CPU Power (kHS06)	2020	2021
Tier 0	98	112
Tier 1	328	367
Tier 2	185	205
Total WLCG	611	684
HLT farm	10	50
Yandex	10	50
Total non-WLCG	20	100
Grand total	631	784

Disk (PB)	2020	2021
Tier0	17.2	20.7
Tier1	33.2	41.4
Tier2	7.2	8.0
Total	57.6	70.1

Tape (PB)	2020	2021 (baseline)	2021 (contingency)
Tier0	36.1	56	85
Tier1	55.5	96	147
Total	91.6	152	232

Summary

- Smooth and efficient usage of computing resources
 - on average 2x above WLCG pledged resources
 - MC production requests dominates CPU work
 - Sizeable contribution from opportunistic resources
- Trying to add more flexibility
 - Containerised workflows
 - HPC centers
- Preparing for Run3
 - Speeding up simulation
 - Progress in detector description, access to conditions, integration with online
 - Offline Analysis Task Force set up
 - Preliminary 2021 computing requests
- Run4: no major changes in computing model
- Run5 will be challenging (Upgrade-II)