



LHConCRAY: ATLAS report

Gianfranco Sciacca

AEC - Laboratory for High Energy Physics, University of Bern, Switzerland

Friday 27 January 2017



u^b

^b
UNIVERSITÄT
BERN

AEC
ALBERT EINSTEIN CENTER
FOR FUNDAMENTAL PHYSICS

LABORATORIUM FÜR HOCHENERGIEPHYSIK
LHEP
UNIVERSITÄT BERN



15 1. Introduction

16 The challenges posed by the High Luminosity LHC (beyond 2020) to the LHC computing are
17 substantial and cannot be met by means of a purely evolutionary approach. Estimates show that
18 already in the first year of the HL-LHC (2027), compared to 2016, we can expect requirements
19 on disk and CPU to be up to one order of magnitude higher than what we can expect from
20 technology improvement, considering a flat level of funding. Efforts are ongoing worldwide to
21 find solutions that will help to meet the demand. It is clear that a single simple solution cannot
22 be expected, and optimisation must occur in several realms: from computing models to software
23 performance in the experiments; from infrastructure evolution and optimisation to efficient use
24 of manpower.

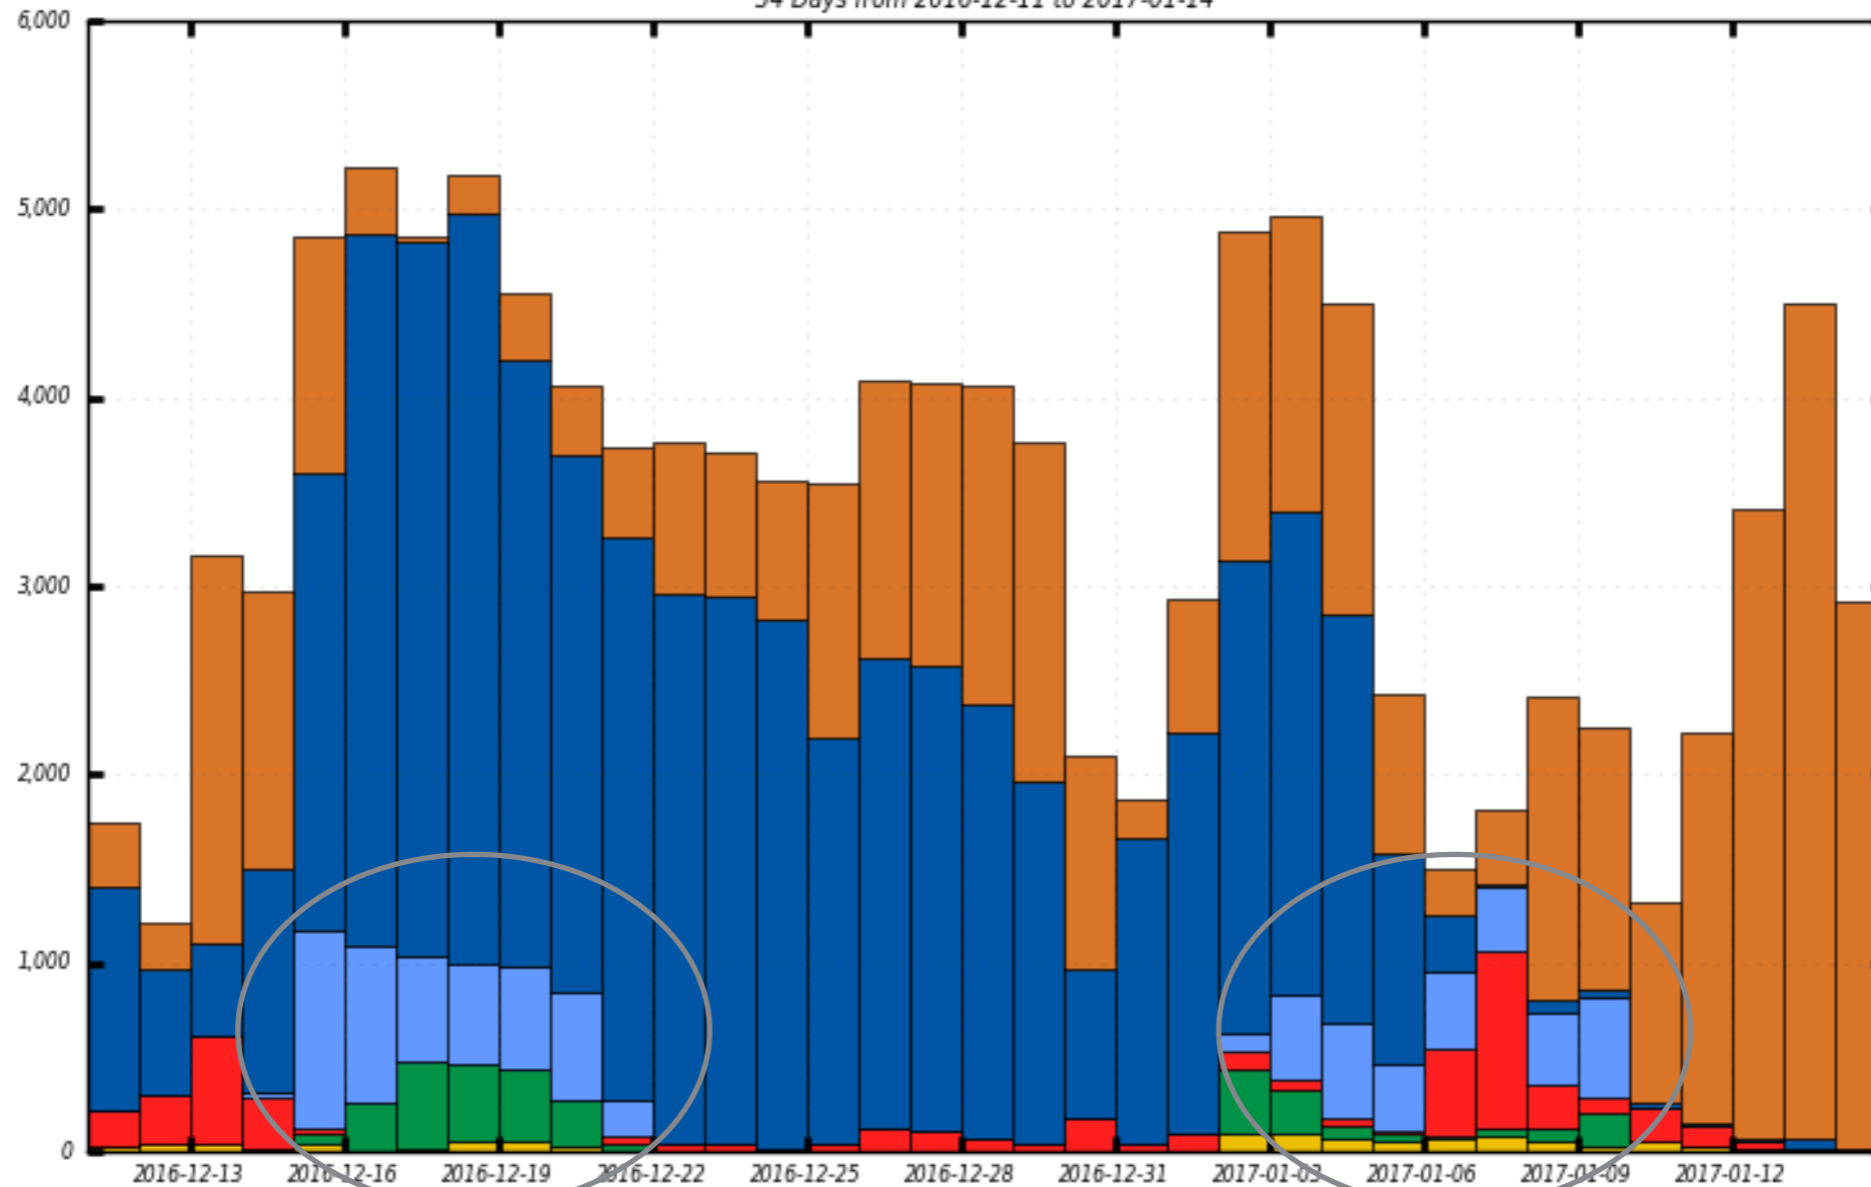
29 A consolidation of the LHC
30 computing to fewer but bigger sites should result in significant operational cost savings. We
31 investigate the use of academic cloud resources [2] and HPC centres for ATLAS computing for
32 potential cost savings, or obtain more computing for the same money. General purpose HPC
33 centres seem a good alternative to dedicated clusters: typically featuring high-end hardware,
34 efficient operational models, and might offer profiting from economy of scales when procuring
hardware. We report on the efforts in integrating high-end Cray systems in WLCG.



LHConCRAY: ATLAS slots running



Slots of Running Jobs
34 Days from 2016-12-11 to 2017-01-14



Legend: CSCS-LCG2_MCORE (orange), CSCS-LCG2 (dark blue), CSCS-LCG2-HPC_MCORE (light blue), ANALY_CSCS (red), CSCS-LCG2-HPC (green), ANALY_CSCS-HPC (yellow)

Maximum: 5,224 , Minimum: 1,207 , Average: 3,376 , Current: 2,919




LHConCRAY: good/fail WC efficiency



WallClock Efficiency based on success/all accomplished jobs

34 Days from 2016-12-11 to 2017-01-14






LHConCRAY: CPU/WC efficiency



Efficiency Good Jobs
34 Days from 2016-12-11 to 2017-01-14






LHConCRAY: CPU/WC efficiency



Average Efficiency Good Jobs





LHConCRAY: single core stats



➔ Single core:

jobstatus (6)

activated (9)

cancelled (319)

closed (327)

failed (3962)

finished (13757)

throttled (252)

Overall error summary

Category:code	Attempt list	Nerrors	Sample error description
ddm:100	jobs	2	Setupper._setupSource() could not get VUID of prodDBlock
ddm:200	jobs	1085	expected output EVNT.10246081._071932.pool.root.1 is missing in pilot XML
exe:-1	jobs	2071	Failed in data staging: Failed checking source replica srm://storage01.lcg.cscs.ch:8443/srm/managerv2?SFN=/pnfs/lcg.cscs.ch/atlas/atlasdatadisk/rucio/panda/1e/10/DAOD_EXOT3.10321519._000121.pool.root.1.panda.um_001
exe:15	jobs	887	No events to process: 999000 (skipEvents) >= 0 (inputEvents of
exe:38	jobs	1	Event count check for DAOD_STDM4 to DAOD_STDM4_MRG failed: found 10209 events, expected 14577; Logfile error in log.STDM4Merge: "SysError in <TFile::ReadBuffer>: DAOD_STDM4.10347298._000089.pool.root.1.panda.um_006 (Inpu (truncated)
exe:65	jobs	278	Non-zero return code from generate (64); Logfile error in log.generate: "Sherpa_i FATAL UNKNOWN Exception is caught"
exe:68	jobs	2	Fatal error in athena logfile: "G4 exception at line 9669 (see jobReport for further details)"
exe:9	jobs	1	Input file argument globbed to NO input files - probably the file(s) are missing
exe:91	jobs	452	LRMS error: (91) Job failed
jobdispatcher:100	jobs	2	lost heartbeat : 2016-12-15 19:51:55
pilot:1008	jobs	2671	Job failed for unknown reason
pilot:1112	jobs	1	Exception caught: [Errno 5] Input/output error, Traceback (most recent call last): File "/scratch/snx2000tds/wlwg/arc_sessiondir/RFwKdMOf2ipnHy6lqpQ1NrpABFKDmABFKDmvUiKdMABFKDm1KY6Um/pilot3/Monitor.py", line 1507, in monitor_job self.__che
pilot:1187	jobs	3	Payload metadata is not available
pilot:1213	jobs	43	LRMS error: (257) Job timeout
transformation:1	jobs	69	Unspecified error, consult log file
transformation:251	jobs	2	



LHConCRAY: 8core stats



➔ Single core:

jobstatus (6)

activated (58)

cancelled (149)

closed (2112)

failed (930)

finished (5265)

merging (1)

Overall error summary

Category:code	Attempt list	Nerrors	Sample error description
ddm:200	jobs	157	expected output AOD.10096183._000112.pool.root.1 is missing in pilot XML
exe:-1	jobs	385	LRMS error: (271) Job was cancelled
exe:11	jobs	45	Missing AthenaMP outputs file athenaMP-outputs-RD0toRD0Trigger-r2t (probably athena crashed)
exe:65	jobs	44	Non-zero return code from AODtoDAOD (65); Logfile error in log.AODtoDAOD: "Segmentation fault: Event counter: 6228; Run: 284500; Evt: 5147833; Current algorithm: AODFi Function: unknown"
exe:66	jobs	2	File data16_13TeV.00300908.physics_Main.merge.AOD.f708_m1606._lb0506._0003.1 did not pass corruption test
exe:68	jobs	3	Fatal error in athena logfile: "G4 exception at line 17472 (see jobReport for further details)"
exe:91	jobs	124	LRMS error: (91) Job failed
pilot:1008	jobs	734	Job failed for unknown reason
pilot:1213	jobs	1	LRMS error: (257) Job timeout
taskbuffer:113	jobs	2	merge job failed
transformation:1	jobs	99	Unspecified error, consult log file



They should be repeated on the optimised system and in parallel performance tests for I/O intensive workloads should be carried out, including stress tests with a production-like mix of experiment workloads at the scale of several thousand cores.

Given the number of complexities we had to deal with, this study has not been in the position yet to show that a HPC is necessarily a cost effective solution for the future of LHC computing.