


LHCb Computing


LHCC status report





Operations September 2014 to November 2014


- Running jobs by activity
 - Montecarlo simulation continues as main activity
 - ☆ Productions for 2015 trigger optimisation started mid-October
 - User jobs
 - Data Stripping
 - ☆ Production tests for Stripping21





Operations September 2014 to November 2014

- CPU usage by site - business as usual
 - ~50% CERN + Tier1 (including RRCKI)
 - ~12% HLT + 5% Yandex
 - ~33% Tier2






Ramp up of storage use at RRCKI

- Start filling with new data from 2015 simulation

PFN space usage by StorageElement

10 Weeks from Week 35 of 2014 to Week 45 of 2014






Ramp up of Tier2-D usage

- Place new data also at T2-D
 - Enable user jobs

User jobs data usage at T2-Ds



Generated on 2014-11-01 21:59:38 UTC

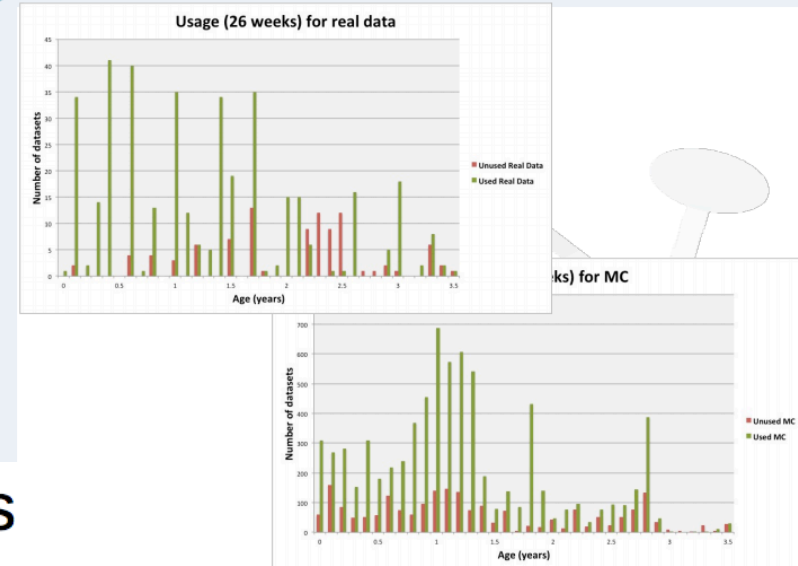


Studies of data popularity data

- Many interesting features in data popularity data
 - Starting to analyse them to evolve replication and cleanup strategies
 - Studying use of MVA techniques to predict future use



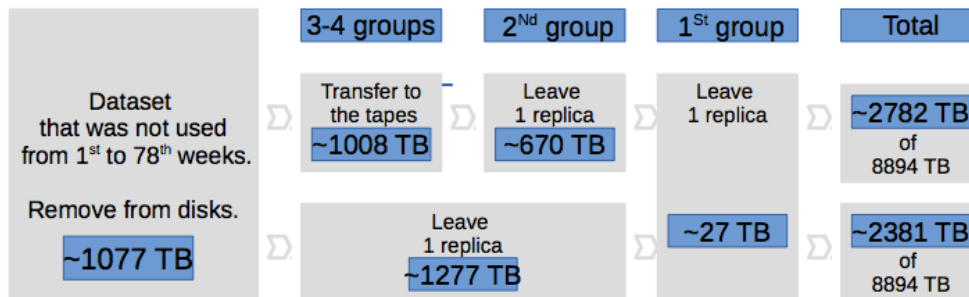

Age vs usage




4th Computing workshop, Philippe.Charpentier@cern.ch

16

Storage saving strategies examples



Age of unused datasets



4th Computing workshop, Philippe.Charpentier@cern.ch

15



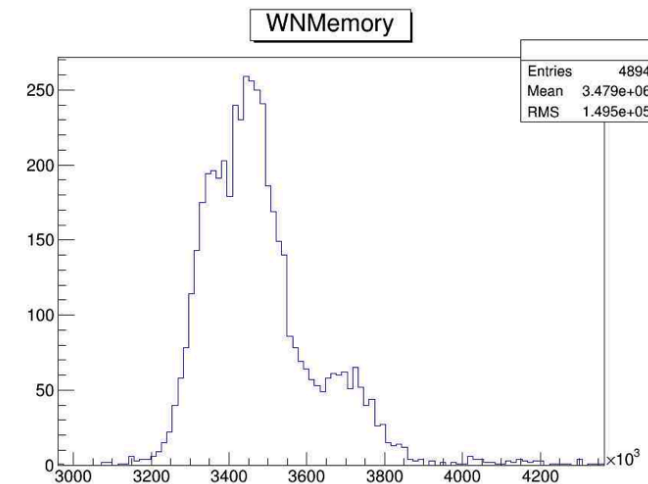
Data processing plans, next quarter

- Full restripping of 2011, 2012 data (Stripping21)
 - Applying latest calibrations
 - Slow and painful commissioning
 - ☆ Intended to be legacy dataset
 - ☆ Large scale deployment of MDST, many minor issues to be addressed
 - ☆ Intended also as commissioning of initial stripping for 2015 data
- Simulation of 2015
 - Continue campaign of 2015 simulation for
 - ☆ Tuning of HLT bandwidth division
 - ☆ Systematics studies for “early measurements”
- Reprocessing of 2010 data postponed
 - Tiny addition to 2011+2012 statistics
 - ☆ Systematics limited
 - ☆ On hold for now, pending valid use case



Stripping 21 status

- Legacy dataset, intended to replace all previous ones
 - ~1500 stripping lines (c.f. 900 in Stripping 20)
- Big success in MDST migration
 - S21 selects 30-35% of 5kHz HLT rate
 - Uses 20% of RAW bandwidth, despite adding Reco information
 - ☆ MDST.DST safety net adds a further 10%
 - However, calculation of variables for MDST takes time
 - ☆ Code that was previously run in analysis, so overall gain, but slows down production
 - ☆ ~1s/event, twice slower than S20
- Memory (VMEM) still a worry, but acceptable
 - N.B. using ROOT5
 - ☆ ROOT6 adds ~700MB
- Schedule:
 - Ready to launch later this week
 - ~6 weeks of processing





Migration to ROOT 6

- LHCb has made major contribution to ROOT 6 validation
 - Decision to expose users to it from early on
 - ☆ Using ROOT 6.00.00 as default since June
 - Not without problems, especially in Grid deployment
 - ☆ Xrootd incompatibility
 - ☆ SLC6.6 incompatibility
 - ☆ Issues reading new data with ROOT 5
 - All issues promptly addressed and resolved
 - ☆ Thanks to Root team, middleware providers, grid sites, USERS!
- VMEM footprint of dictionaries remains a major concern
 - Some improvements available in next release
 - Work programme on LHCb side to decouple “interactivity” from production use
 - ☆ Related to use of python for configuration
- Baseline for 2015 remains ROOT 6
 - But dependent on results of benchmarking of HLT in production configuration on HLT farm
 - ☆ Benchmarking to take place in coming weeks



Online calibration : all strategies defined

JOB	How often	Const. update	Expected update	Where	Input data	Job type	Time
RICH calibration	Per run	Same run	Per run	Hlt2, offline	All events in one run	Online analysis task (online Brunel)	O(1min) from when histograms available
OT t0 calibration	Per fill or day	Next run	Under study (after TS in run 1)	Hlt1, hlt2, offline	100k – 1M events	Online analysis task (online Brunel)	O(1min) from when histograms available
VELO alignment	Per fill	Next run	Per fill (not each fill)	Hlt1, hlt2, offline	50k halo tracks + 50k coll. ev.	Online alignment framework. Kalman alignment	20min on 8CPU → O(1min) in online farm
TT, IT, OT alignment	Per fill	Next fill. Next run after TS or Mag. swaps	each 2-4 week	Hlt1, hlt2, offline	50k D ⁰ + overlap tracks	Online alignment framework (kalman alignment)	1h on 8 CPU → O(1min) in online farm
RICH alignment	Per fill	None	None		50k selected ev.	Online alignment framework (RICH alignment)	Time requirement under study
MUON alignment	Per day/ week	None	None		20k J/ψ	Online alignment framework (kalman alignment)	1h on 8 CPU → O(1min) in online farm



HLT Software and calibration commissioning

Status in the pit

- Previous Hlt commissioning week (6-10 October).
 - Successfully ran Hlt1 over the L0 triggered data on disk at the pit
 - And Hlt2 over the Hlt1 output
 - OT cosmons also run.
- HLT now automatically reads all run-by-run calibrations.
- Plans for upcoming TED runs (21-24 November)
 - Use to commission the online alignment and calibration

Mika Vesterinen



Tesla timescale

- New extension of the selection reports has already allowed much more information to be included.
- This means that further implementation can be quick enough to be included in the consultation and allows use case to be brought forward and extended.

End of the summer:

Have a version of Tesla that can be processed by DaVinci and use common analysis tools.

Now:

Consultation with selection report requirements and implementation of needed variables in extended reports.

End of 2014:

Use case with analysts comparing offline and Tesla products.


April 2015:

J/psi and charm production first use in anger as a validation.




Tier0 + Tier1 pledged resources


T0+T1 CPU 2015



T0+T1 Disk 2015



T0+T1 Tape 2015




2015 T0+T1	CPU HS06	Disk Tbytes	Tape Tbytes
CERN	36000	5500	11200
France	23000	1880	4360
Germany	19600	2340	3960
Italy	23600	2720	6870
Netherlands	15661	1570	2773
Russian Federation	14200	1260	1480
Spain	7670	761	1541
UK	35400	3510	7110
Total	175131	19541	39294
Requested	154000	17200	34900
Difference	13.7%	13.6%	12.6%



Tier2 pledged resources

T2 CPU 2015



T2 disk 2015



2015 Tier2	CPU HS06	Disk Tbytes
France	12323	404
Germany	3400	4
Italy	7875	0
Latin America	3183	300
Poland	3500	0
Romania	4900	323
Russian Federation	1539	80
Spain	3000	1
Switzerland	7000	250
UK	13861	602
Total	60581	1964
Requested	66000	1900
Difference	-8.2%	3.4%



Education



The CMS (Compact Muon Solenoid) experiment is one of two large general-purpose detectors built on the Large Hadron Collider (LHC). Its goal is to investigate a wide range of physics such as the characteristics of the Higgs

Explore CMS >

For education purposes, the complex primary data need to be processed into a format (examples below) that is good for simple applications. Get in touch if you wish to build your own applications similar to those shown here

- Cern Open Data portal to access and analyse public data from all LHC experiments
- LHCb joined with the data samples already made public for the International Masterclass program
- Virtual machine image to download the data and analyse it.
- Other LHCb samples and applications will be added in the future, according to the LHCb open data policy.



properties of the Higgs-like particle, extra dimensions of space, unification of fundamental forces, and evidence for

Explore ATLAS >



The LHCb (Large Hadron Collider beauty) experiment aims to record the decay of particles containing b and anti-b quarks, known as B mesons. The detector is designed to gather information about the identity

Explore LHCb >



opendata.cern.ch

ABOUT SEARCH EDUCATION RESEARCH

Search

LHCb > LHCb Derived Datasets

LHCb Derived Datasets

This collection contains LHCb simplified data events to be used for the masterclass exercises.

LHCb event file for real measurement

This file contains about 60k events pre-selected using loose criteria to contain D0 meson candidates decaying into a Kaon and a Pion.

Collection: LHCb-Derived-Datasets DOI: 10.7483/OPENDATA.LHCb.E7EJJUWR

LHCb masterclass event files for the event display

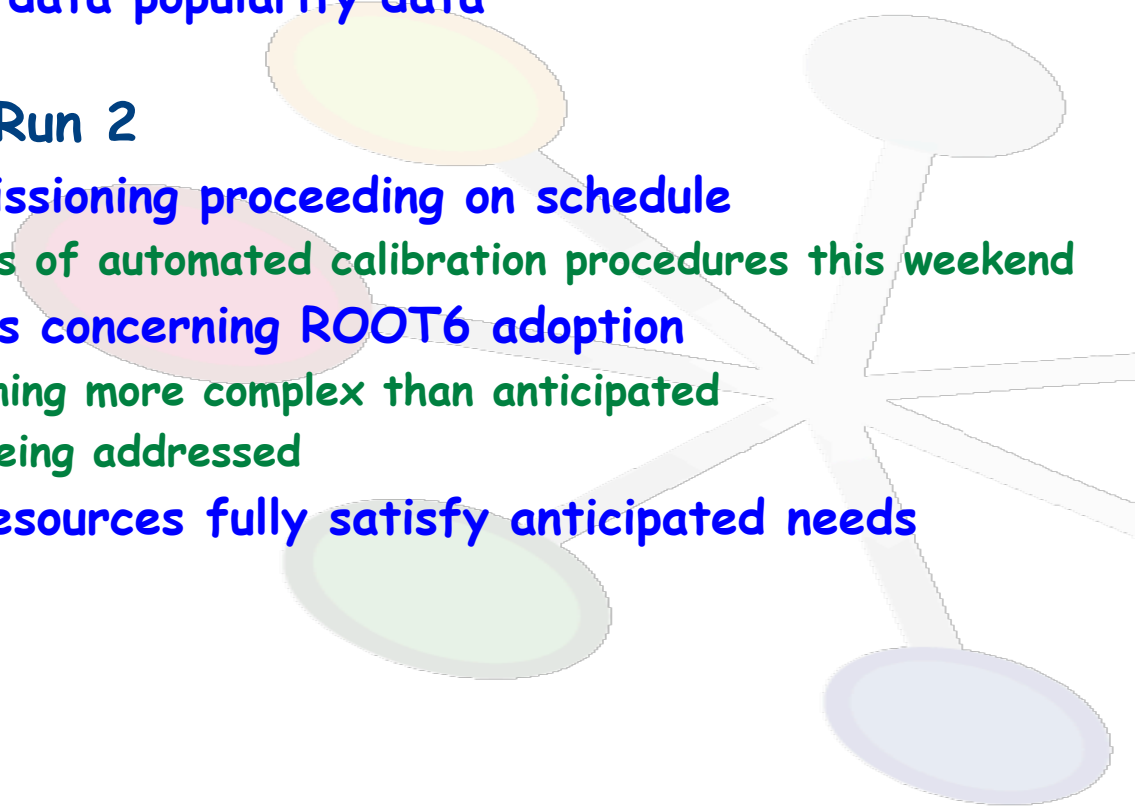
Each file contains 30 events recorded with LHCb detector. The events have been pre-selected using loose criteria to contain D0 meson candidates decaying into a Kaon and a Pion.

Collection: LHCb-Derived-Datasets DOI: 10.7483/OPENDATA.LHCb.6BDP.9K6B




- **Operations**
 - **Business as usual, no major issues**
 - **Stripping 21 much delayed but now on the starting blocks**
 - **Learning to use data popularity data**

- **Preparations for Run 2**
 - **Software commissioning proceeding on schedule**
 - ☆ **Full scale tests of automated calibration procedures this weekend**
 - **No showstoppers concerning ROOT6 adoption**
 - ☆ **But commissioning more complex than anticipated**
 - ☆ **VMEM issue being addressed**
 - **2015 pledged resources fully satisfy anticipated needs**



LHCb Computing

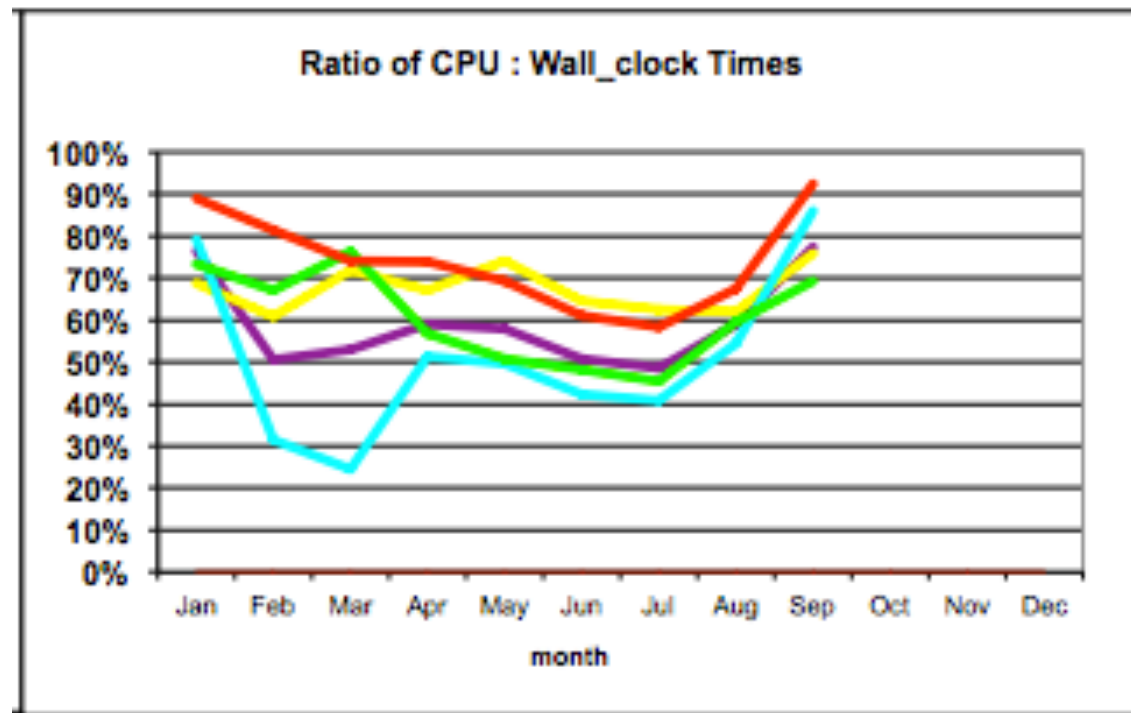
Backup









Job "efficiency" at CERN

- Plot shown not measuring efficiency, also idle resources permanently allocated to VO
 - VO Boxes, Openstack
- OK since September, now excludes VO boxes



ALICE  CMS 
ATLAS  LHCb  site average - cpu:wall_clock ratio 