LHCb Computing

2015Q1 status report

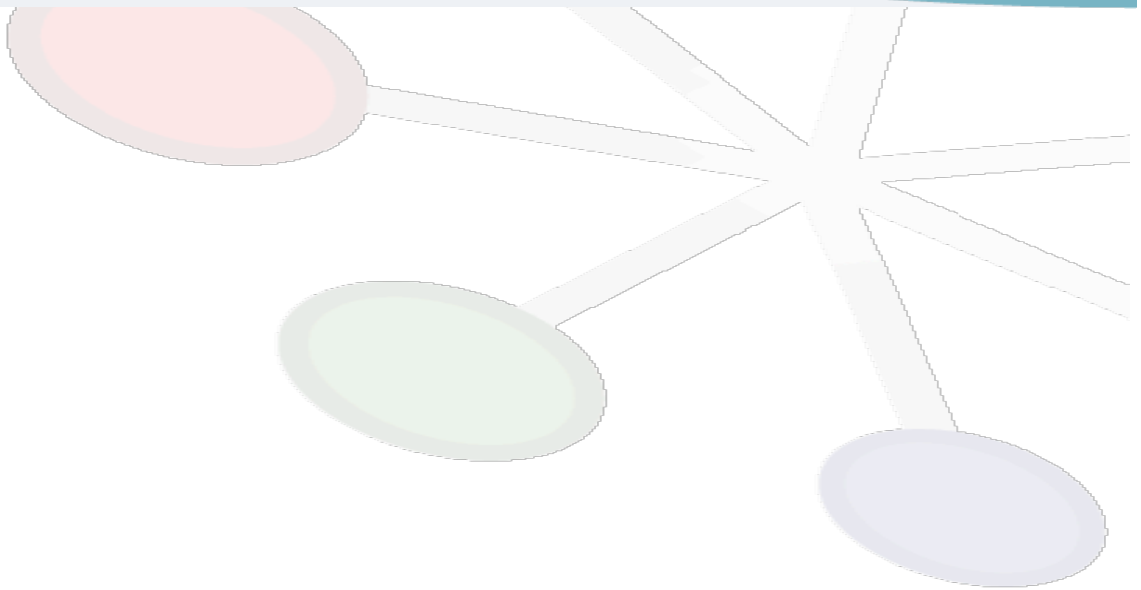

Computing activities in last quarter

- Completion of "legacy" stripping of Run1 data
  - "Stripping 21"
- Ongoing simulation for run 1 analysis
  - "Sim08"
- Tuning of simulation for 2015 analysis
  - "Sim09", also for future Run1 analyses
- Preparation of HLT and prompt reconstruction for Run2
  - Split HLT
  - Online calibration
- Consolidation of DIRAC framework

Stripping 21 productions (finally!) closed last week

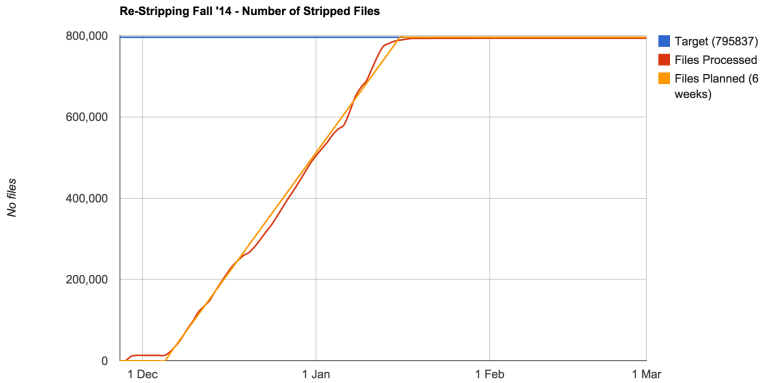

Status of S21

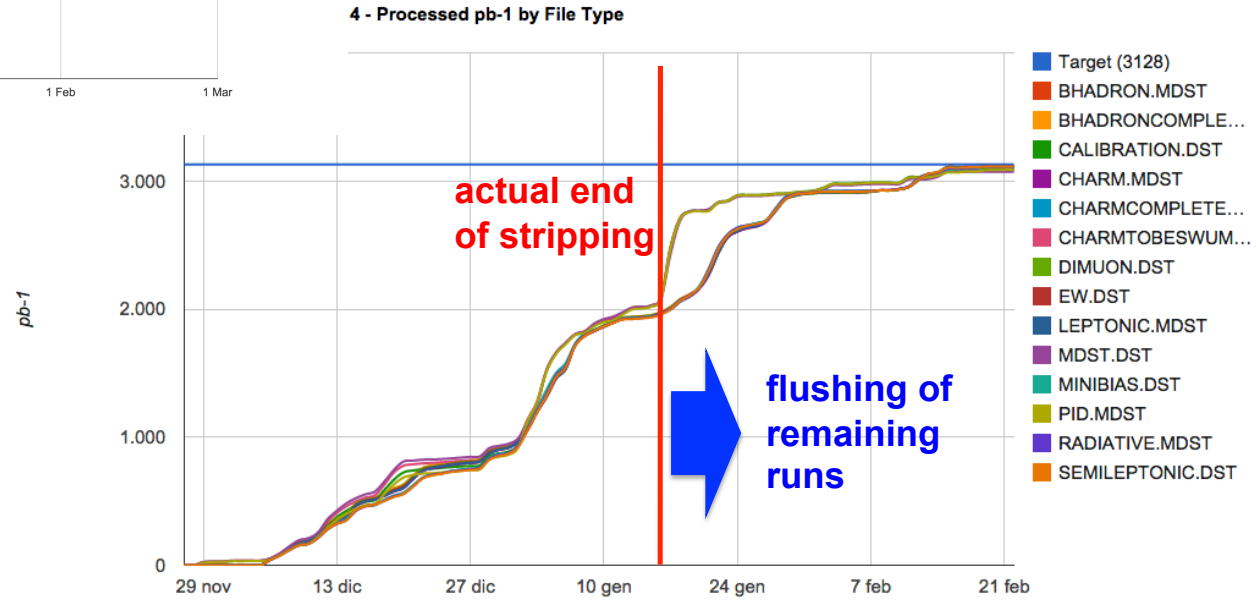

- Stripping of 2011+2012 data took ~6 weeks as in the original plans
- Commissioning of S21 took 3 months!
  - some actions will be taken to speed up this phase in the future (and make it safer)

Data loss: redundancy versus effort

Lost data in S21

c.f. total luminosity:
2012: 2fb^{-1} 2011: 1fb^{-1}

- A major storage failure at one of the Tier1s caused the loss of a small fraction of data
  - lost files are scattered in the whole data taking period
  - recovering them would mean to re-run 1/3 of the stripping
  - a lot of human effort would be required in order to perform a complete recover of lost data

Stream	2012	2011
Bhadron	0.342 pb^{-1}	0.160 pb^{-1}
BhadronCE	0.147 pb^{-1}	0.400 pb^{-1}
Calibration	24.395 pb^{-1}	0.550 pb^{-1}
Charm	0.729 pb^{-1}	0.069 pb^{-1}
CharmCE	0.406 pb^{-1}	0.762 pb^{-1}
CharmTBS	27.308 pb^{-1}	0.405 pb^{-1}
Dimuon	0.254 nb^{-1}	0.070 pb^{-1}
EW	0	0.154 pb^{-1}
Leptonic	0.431 pb^{-1}	0.156 pb^{-1}
MSDT.DST	29.479 pb^{-1}	0.044 pb^{-1}
MiniBias	1.327 pb^{-1}	1.069 pb^{-1}
PID	0.534 pb^{-1}	0.341 pb^{-1}
Radiative	0.534 pb^{-1}	0.190 pb^{-1}
SemiLeptonic	1.217 pb^{-1}	0.091 pb^{-1}

Decided not to recover these data

These were intermediate temporary files during processing, single disk copy
Reinforces conviction that minimum of two copies are needed for important data
Regeneration is an expensive option (when it's possible)

Simulation for Run 1 analysis: constant rate since > 1 year

Ongoing MC productions (I)

- Continue to provide current level of support for Sim08{a,…}, THE productions for physics analysis with 2011, 2012 and 2013 data until Sim09 supersedes it
  - Extend support for legacy productions for reprocessed 2010 data
  - Keep adding new compatible features
  - Support will drop to data preservation level once Sim09 starts

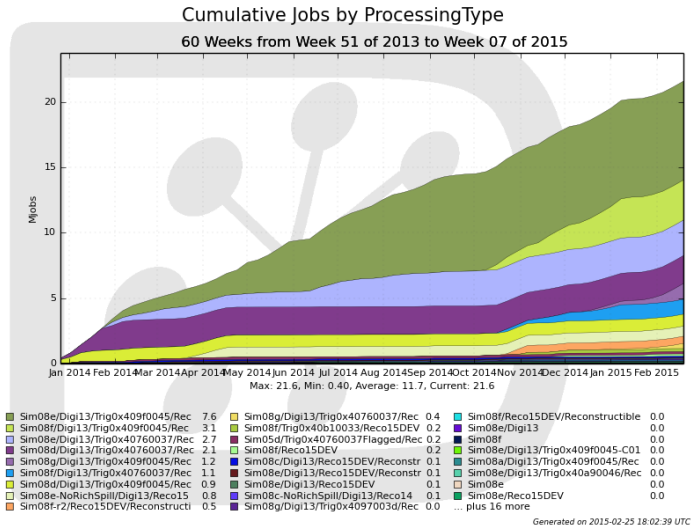

Fast Simulation – work in progress

fastMC

- Many possible solutions have been tried out
  - Parametric approach or event smearing of high level quantities
    - Different approaches to be combined into a generally useful package
  - DELPHES parametric detector package integrated in Gauss
    - Prototype available, to become a fully functional production option of Gauss
  - Event stitching inserting a signal particle gun into a real data minimum bias
    - Challenging for productions. Alternative approach of re-using underlying MC event many times with different signal decay 'easier' in production

Fast simulation: available for 2015 productions

fastMC in Sim09

- Signal Particle Gun is very useful for some studies and x10 faster to produce
  - Also smaller size samples
- Will be an alternative generator for Sim09
  - e.g. Beam4000GeV-2012-MagDown-Fix1-pGun
  - Data package with P and P_T spectra for each gun particle at each $E_{collision}$
  - To be used also to emulate CEP events
  - Productions being setup now in Sim08. Spectra only available at the moment for $D^*(2010)^+$ and $\chi_{c0,\ c1,\ c2}$

Also new in Sim09: MicroDST by default (factor 10 gain in storage)

Split HLT commissioning

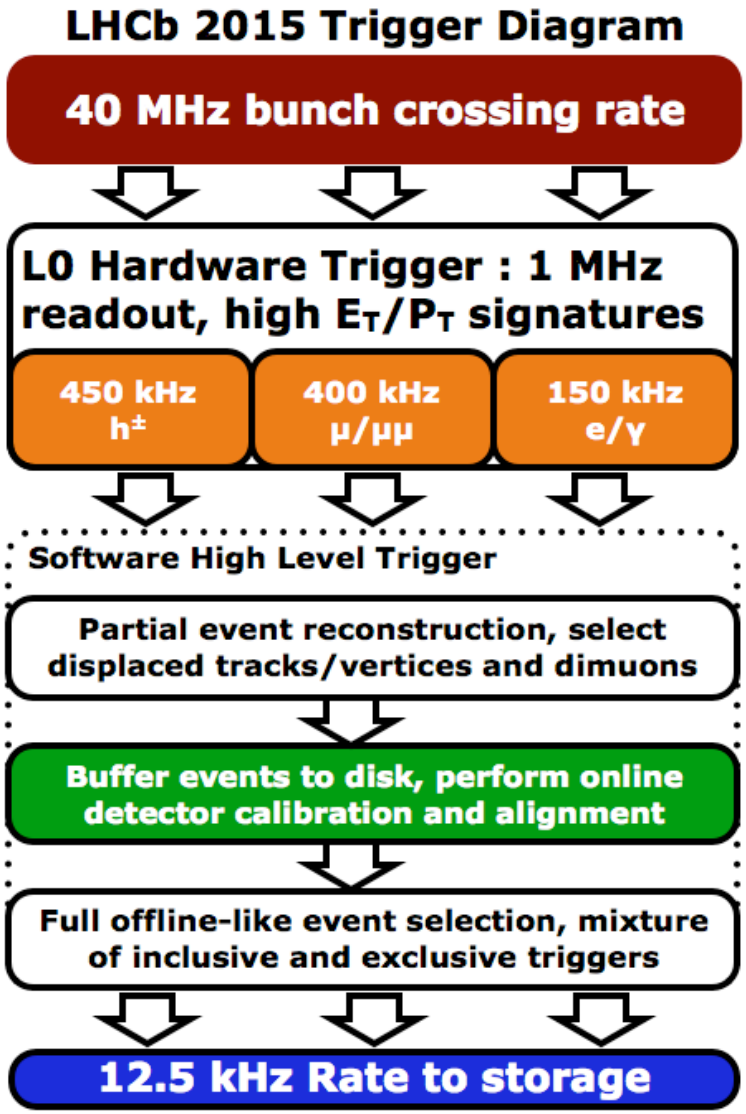

- Reminder:
  - HLT1 on prompt data, 1MHz
    - Save output to disk
  - Perform calibration and alignment
    - Automated, using HLT1 data
  - Run HLT2 fully deferred when calibration ready
    - 12.5 kHz to offline
  - Use same alignment and calibration in prompt offline reconstruction
    - No end of year reprocessing foreseen
  - Procedure tested in last commissioning week using 2012 minimum bias data at input to HLT1

Strategy for prompt calibration

Online align. and calibr. summary table

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 run	Next run	Under study (after TS in run1)	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

13

Every Cycle Counts!

callgrind "cycle count"
Moore v23r2,
Physics_September2012
Hlt1+Hlt2

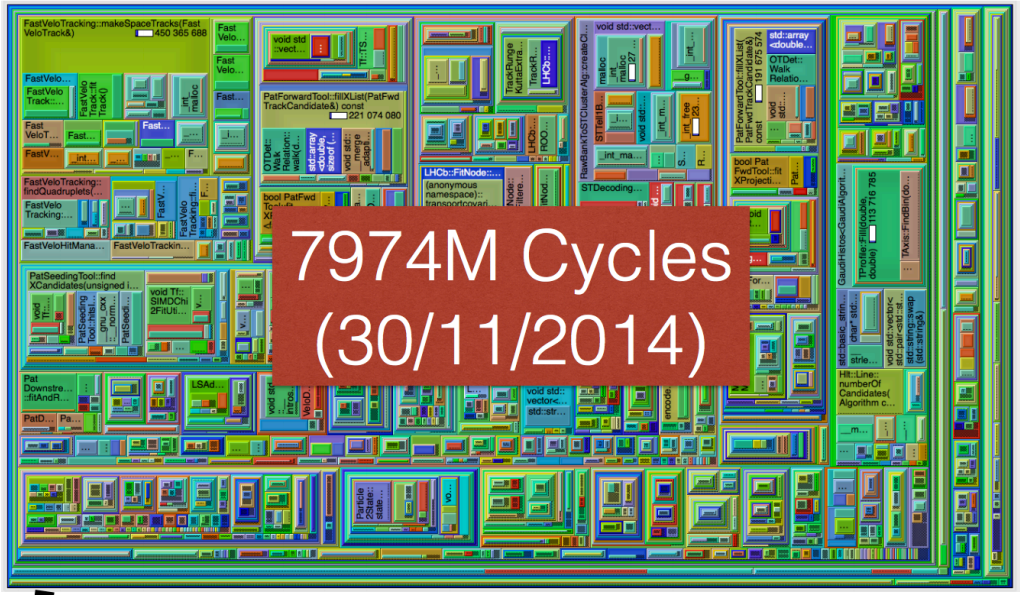

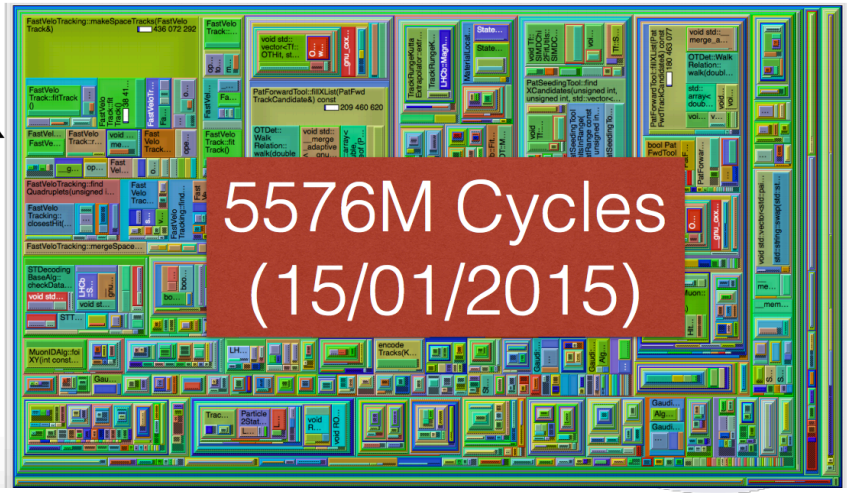

- **Major efforts to optimise the software**
  - **Targeted at HLT but of direct benefit to offline**

Gerhard Raven
Nikhef & VU Amsterdam

Timing improvements

Gerhard, Manuel, Chris et al.

Reconstruction	v23r2 [ms]	HEAD [ms]
HLT1 forward	6.2	5.4
HLT1 fit all	5.4	4.5
HLT1 fit Velo	7.9	7.8
HLT2 tracking	80.5	69.7
HLT2 RICH	81.9	63.7
HLT2 Calo (prelim.)	61.4	41.6
TisTos	0.4 /sel.	0.4 /sel.

Almost doing the full offline reconstruction in HLT2 is likely possible---we are now haggling over what "almost" in fact means. To be understood when we see 13 TeV data.

As reconstruction gets faster, HLT and offline become more similar

In the future, can envisage doing away with offline reco altogether

Not forgetting the Turbo Stream

Stripping-like selection done directly in HLT
- Selection framework ready
- Analysis framework (Tesla) ready

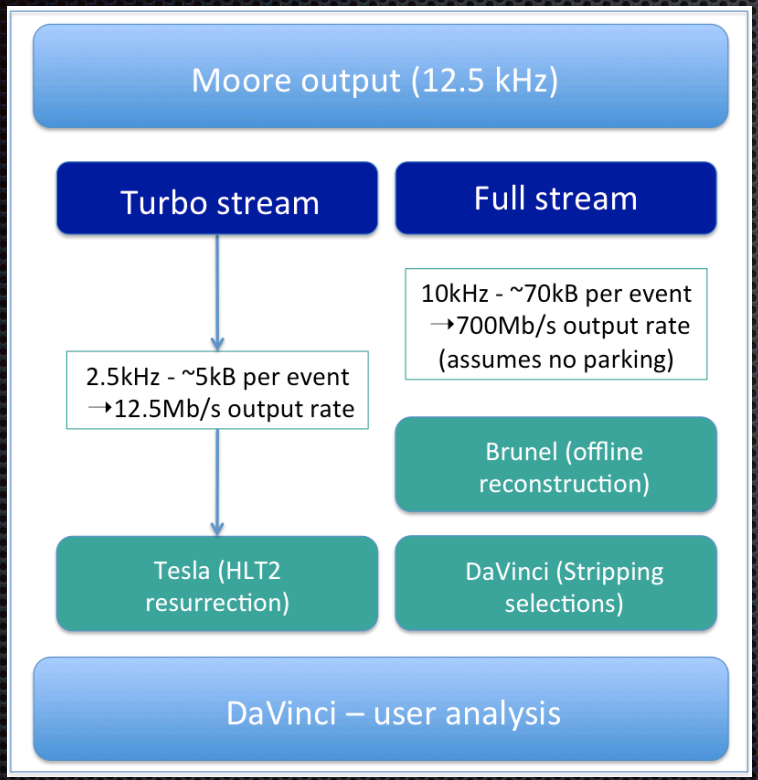

- Don't need another reconstruction (just generate the FSRs for lumi, strip away what we don't need and put the Hlt objects back in place.
- Killing what we don't need reduces the size of the event and allows a much higher rate for the same bandwidth.

2.3 kB/evt out of HLT
- RAW format,
  (in 2015 added to full event)

8.5 kB/evt after conversion to offline format
- Comparable to MDST

S. Benson Turbo update

Will be used for 2015 early measurements

DIRAC File Catalog (DFC)

- Catalog of files and replicas on the GRID
  - DIRAC component, fully integrated in DIRAC
  - Replaces LFC (WLCG service to be discontinued)
  - Designed for load of Run2
    - Major performance improvements relative to LFC
  - Migration mandatory
- LFC – DFC migration
  - Intervention last Monday
  - Change totally transparent to users
- DFC in production since last Tuesday
  - Will run in parallel to LFC for a few months, then switch off LFC

Data popularity

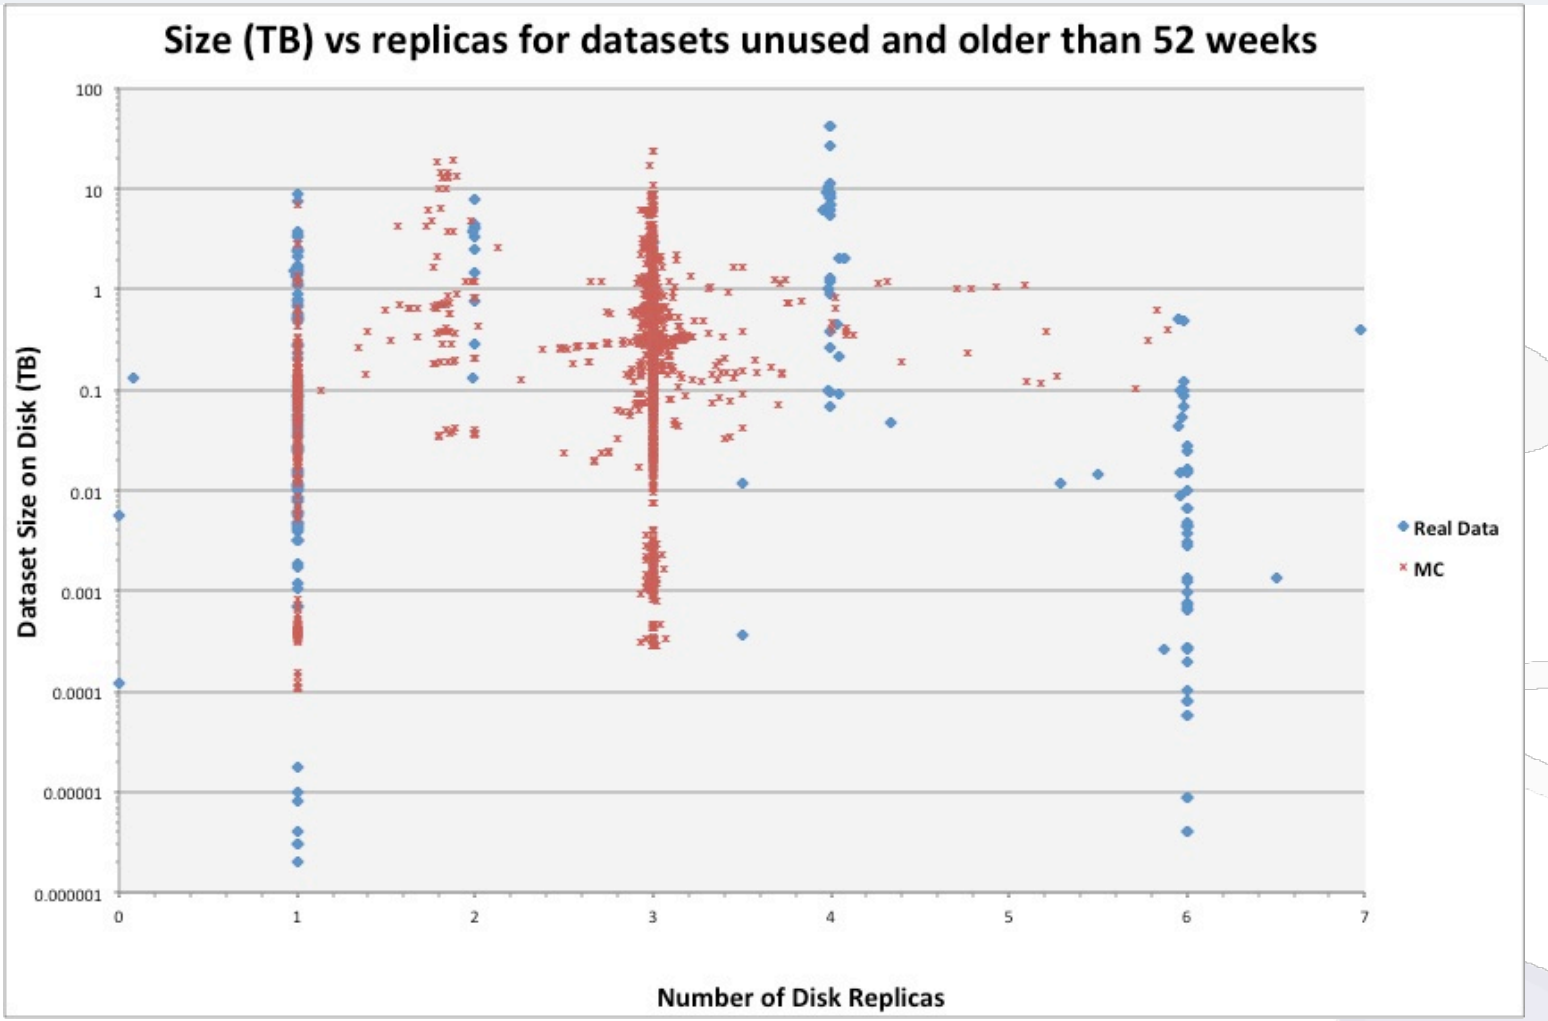

- Scope for saving disk by more dynamic placing of disk replicas
  - Including (scheduled) recall from tape for unpopular data

2015 Resources

- Pledged resources:
  - (T0+T1) ~13% above requests
    - Part of this is due to Russia now pledging a Tier1 for LHCb (RRCKI)
  - Slight (8%) underpledge for Tier2 CPU, compensated by "unpledging" sites (Yandex, OSC)
- Resources request made in 2014, based on LHC schedule available then
  - In latest schedule, 2015 running days are considerably fewer, so we probably won't use up all pledged resources
- We should be OK for 2015

2016 request

- Final request submitted on 15th February
  - Scrutiny under way, C-RSG at end April
- Assumptions:
  - 5×10^6 seconds of data-taking
  - 10kHz FULL stream (no parking in 2016), 2.5kHz Turbo
  - Event sizes and processing times comparable to Run 1
  - No end of year reprocessing
  - Two full strippings, two incremental strippings
  - Due to deferred HLT processing, HLT farm only available for offline use during long (>1 week) scheduled stoppages of LHC
- Relative to 2015 request:
  - 30% increase in CPU
  - 32% increase in disk
  - 71% increase in tape
- At the limit of what is affordable
  - Some margin due to likely underuse of 2015 storage

Computing for LHCb upgrade

- LHCb upgrade: 100kHz trigger rate
  - (c.f. 12.5 kHz in Run 2)
  - Cannot assume offline resources will scale
- Brainstorming starting on computing model for upgrade
  - Using ideas tested in Run 2
    - e.g. Turbo Stream, online reconstruction
  - Room for more ideas
- R&D programme
  - In continuity with Run 2 needs
    - e.g. software vectorisation, fast simulation
  - New technologies
    - e.g. use of co-processors
  - Collaboration with HEP Software Foundation
- Computing TDR scheduled for autumn 2017
  - Upgrade TDR working group to be announced soon
    - Discussion at next computing workshop in May