LHCb Computing

LHCC status report

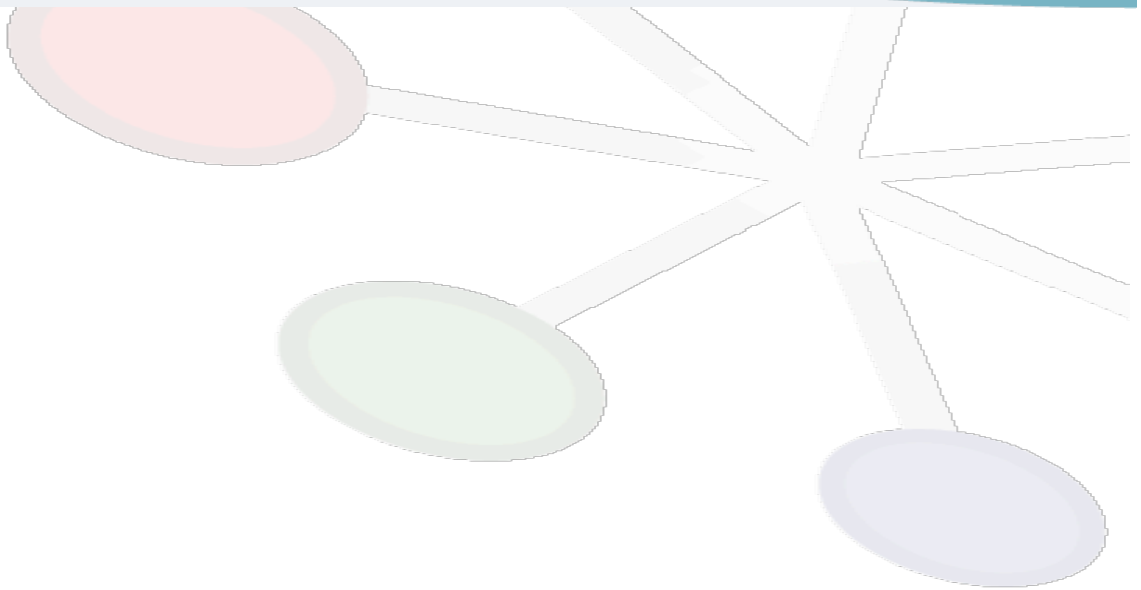

Operations June 2014 to September 2014

- Running jobs by activity
  - Montecarlo simulation continues as main activity
    - Just in time simulation driven by analysis needs
    - No large simulation campaigns this quarter
  - User jobs
  - "Swimming"
    - On request
    - Complex workflow

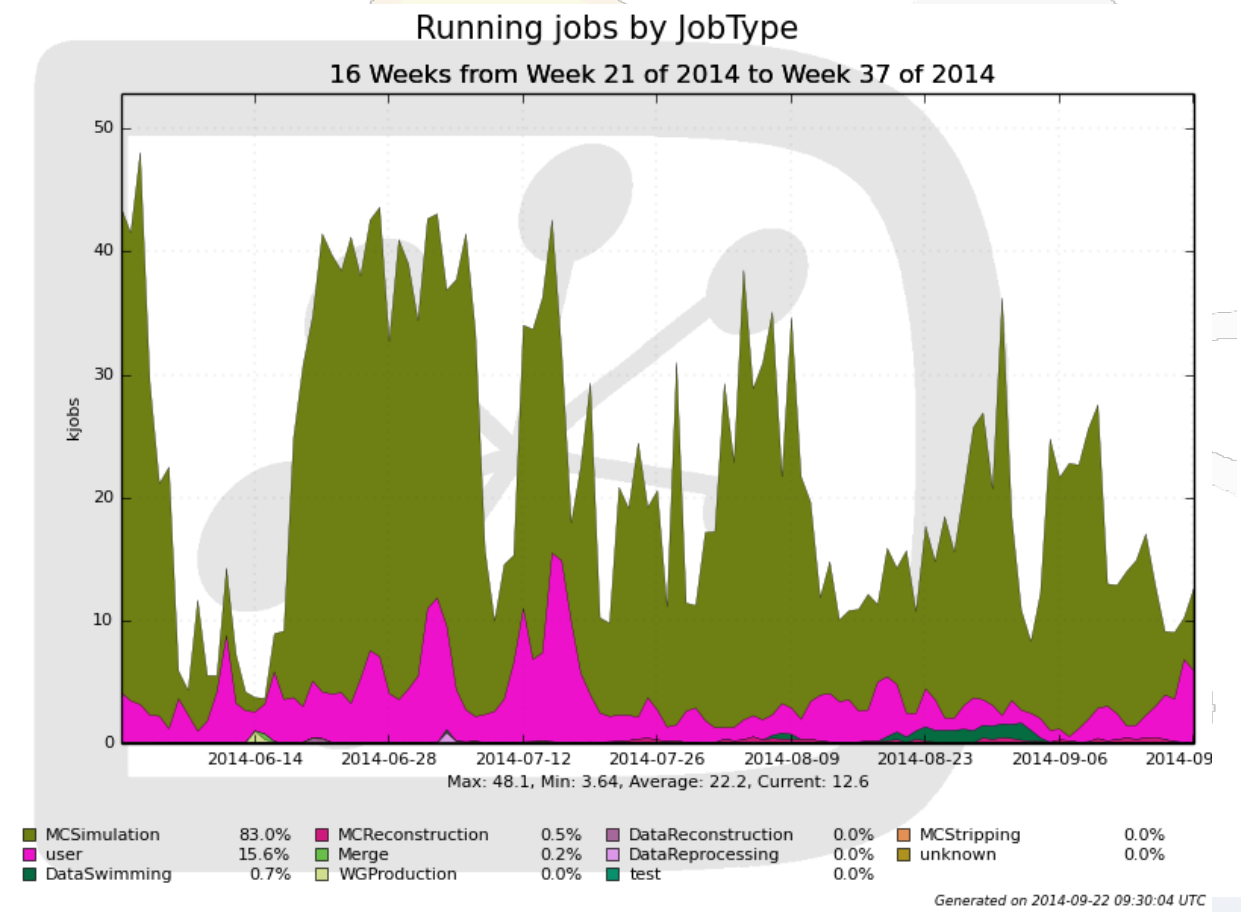

Operations June 2014 to September 2014

- CPU usage by site
  - HLT farm back in business as largest site
    - (Previous quarter affected by power maintenance and SLC6 migration)
  - Otherwise business as usual

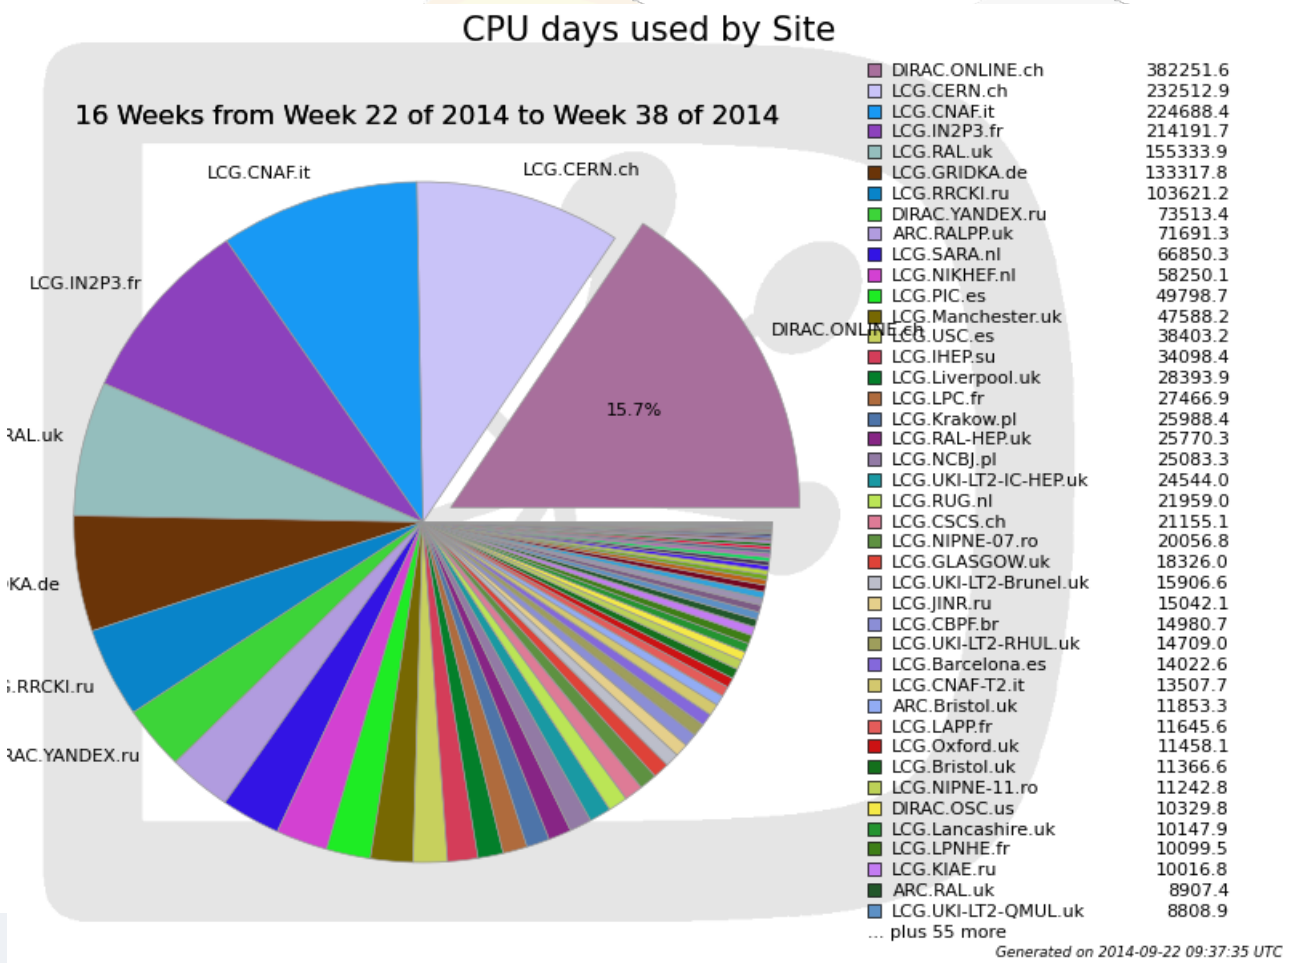

Job efficiency at CERN

- LHCb job efficiency at CERN has been dropping throughout 2014
  - Still above average though!

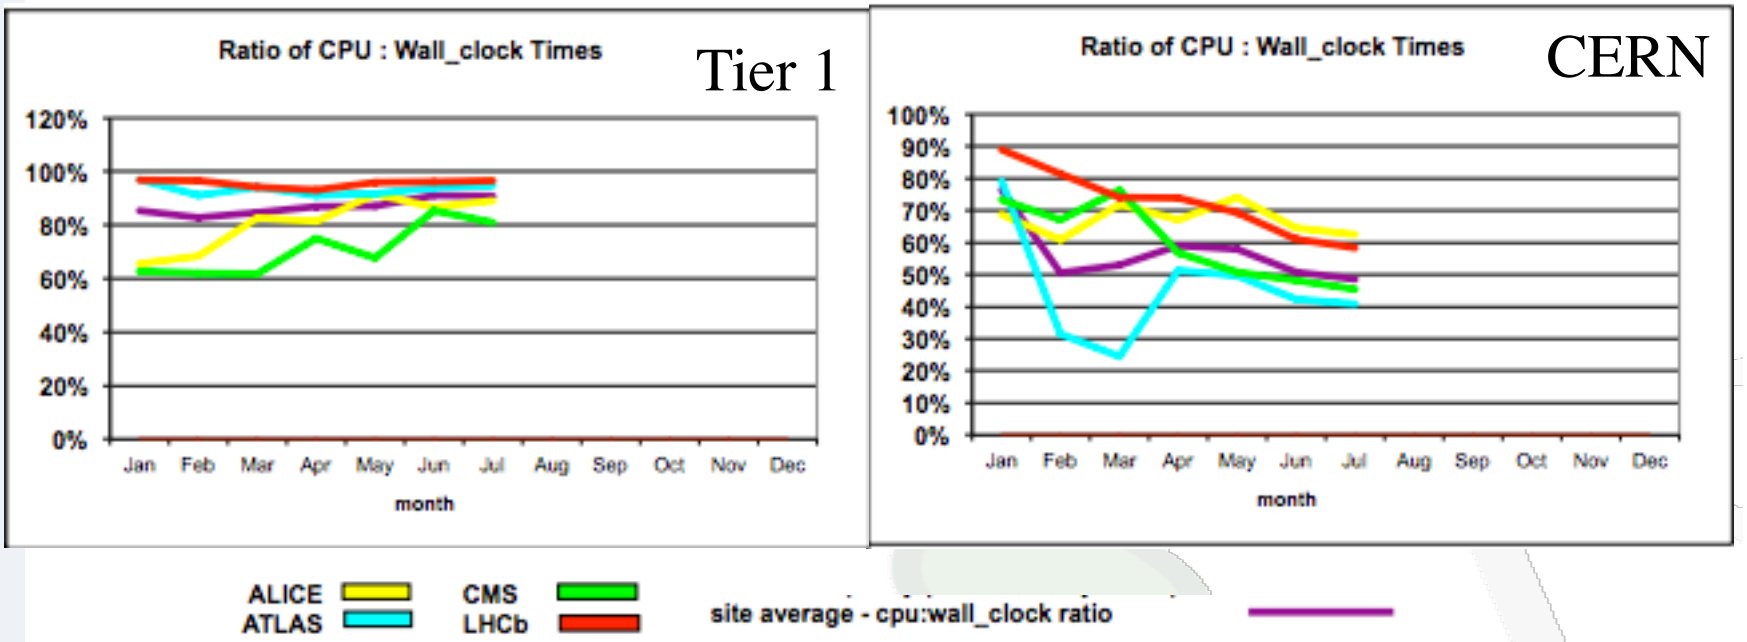

- Only seen at CERN site (same workflows at all sites)
  - MC simulations not affected
    - Related to data access?
  - Investigations ongoing

Operations June 2014 to September 2014

- Disk usage
  - Smooth increase in space used by new simulation
    - ~50 TB / week

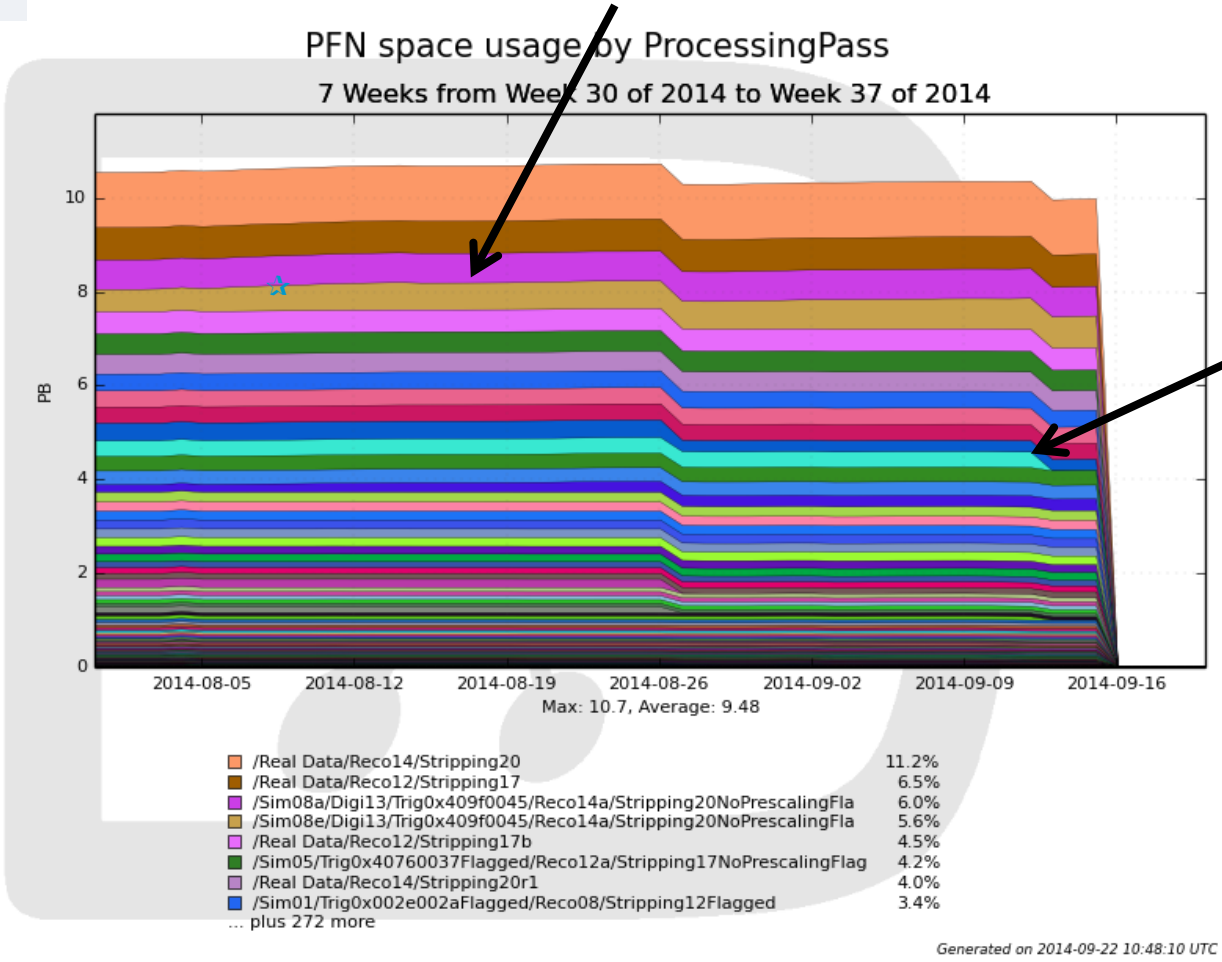

Significant reduction in space by clearing obsolete datasets

First benefits of data popularity measurements

Disk resident data usage, last 6 months

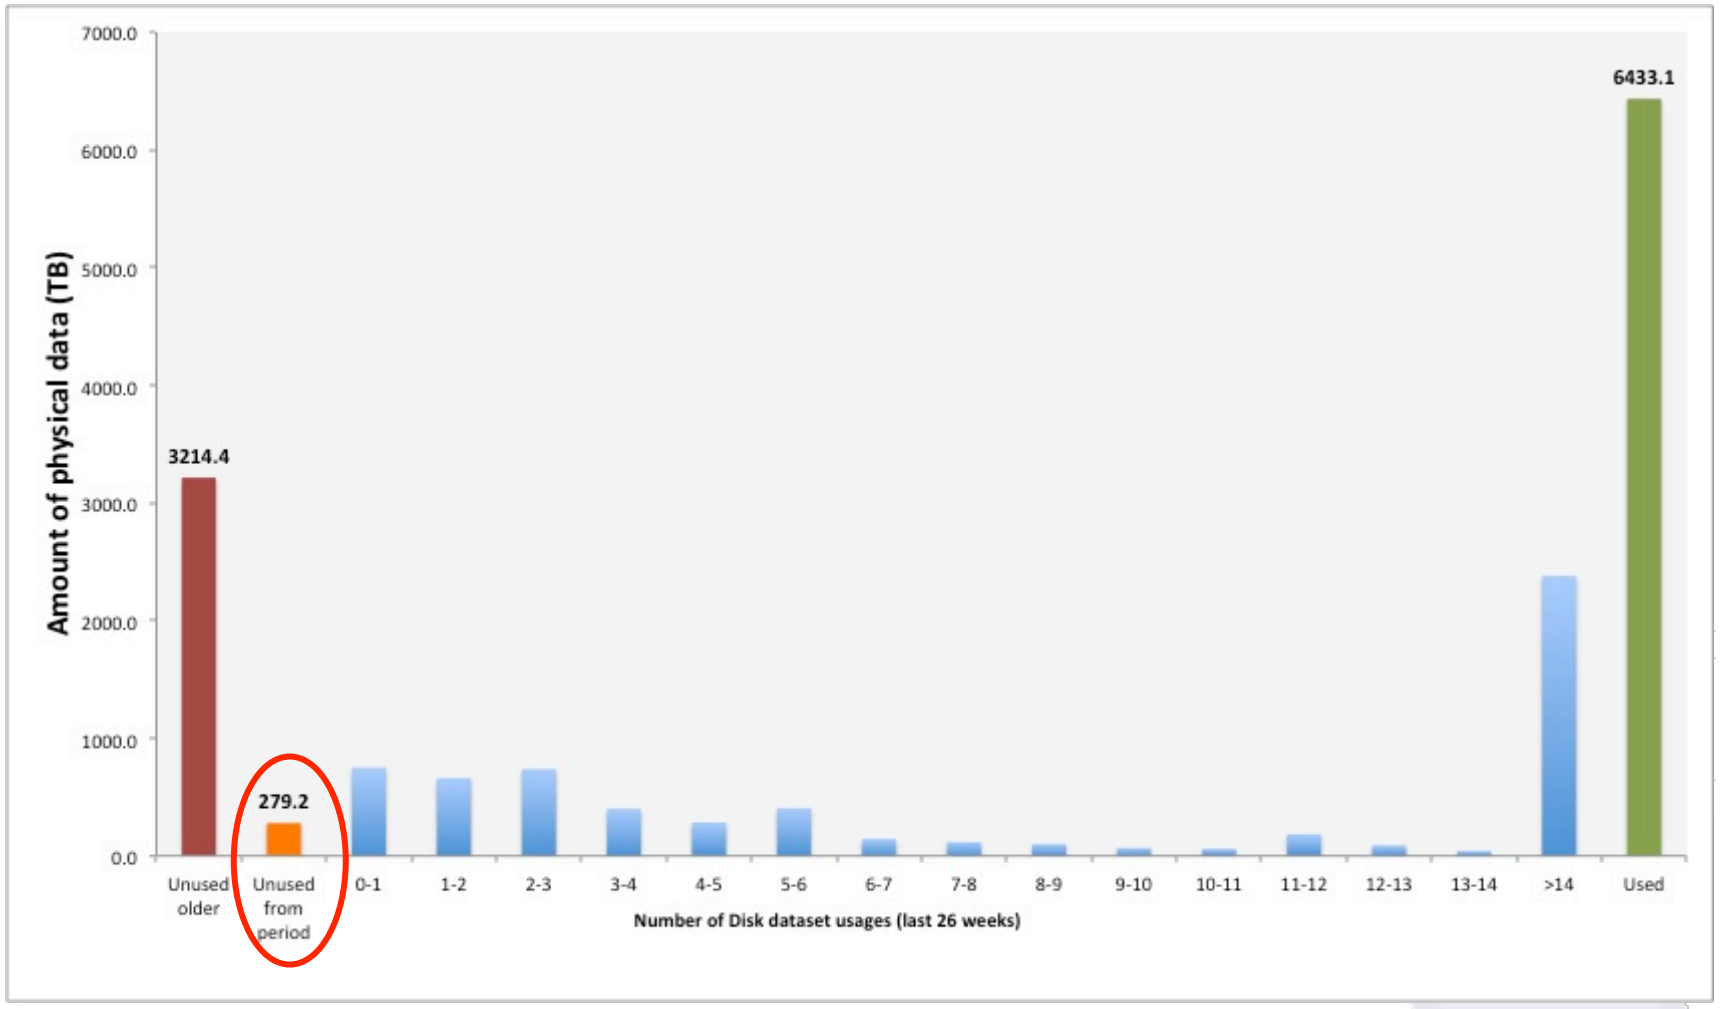

- In the last 6 months, 3.5 PB were not used, but 6.5 PB were used!

How much was used?

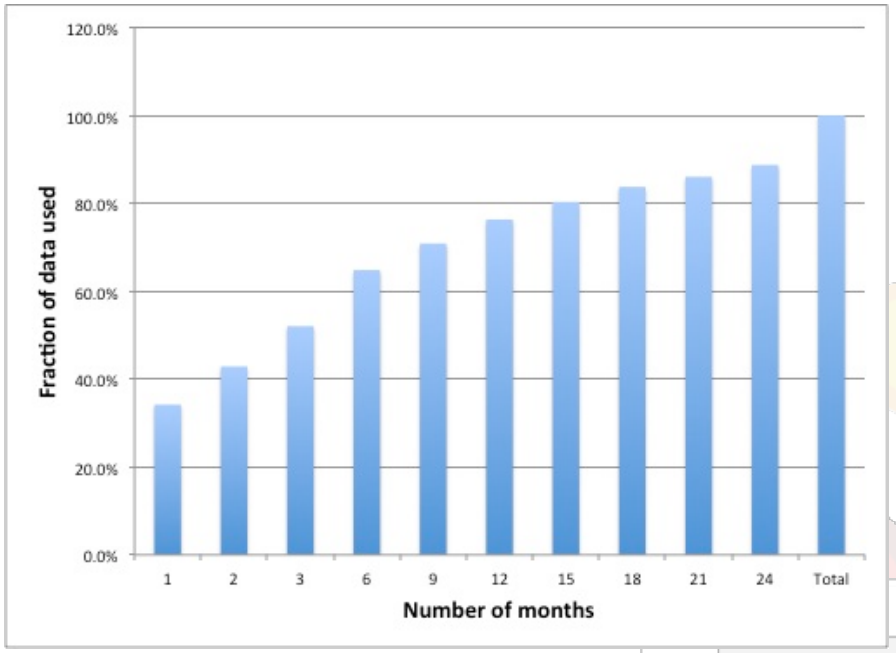

- 85% of existing datasets were used in last 2 years
- Still 15% could be safely retired!
- 1.1 PB…

- Unused datasets spread in age
- Needs more study why not used…

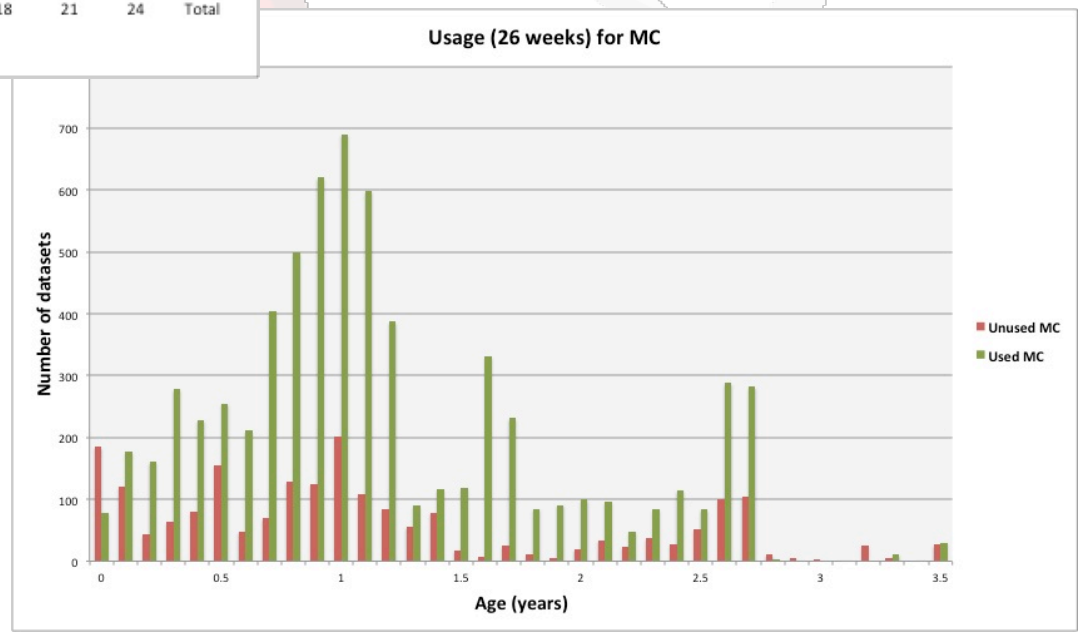

Data processing plans, next quarter

- Reprocessing of 2010 data
  - To provide legacy dataset with reconstruction consistent with 2011 and 2012
    - Considerable work to back-port alignment and calibration
  - Early data (minimum bias dataset) done
    - Do the rest as part of Stripping21 campaign
- Full restripping of 2010, 2011, 2012 data (Stripping21)
  - Applying latest calibrations
  - Intended to be legacy dataset
  - Used also to commission microDST for all stripping lines ahead of 2015 data-taking
- Simulation of 2015
  - Ready to start campaign of 2015 simulation for
    - Tuning of HLT bandwidth division
    - Systematics studies for "early measurements"

Stripping 21 status

- Legacy dataset, intended to replace all previous ones
  - ~1500 stripping lines (c.f. 900 in Stripping 20)
- Big success for microDST migration
  - Despite 70% more lines, S21 output bandwith is 90% of S20
- However this comes at a cost
  - Calculation of variables for MDST takes time
    - Code that was previously run in analysis, so overall gain, but slows down production
    - Currently ~2s/event, 3-4 times slower than S20
      - Optimisation ongoing
  - Each stripping line has a memory footprint
    - Total memory usage currently too high

Stripping 21 memory usage (first test!)

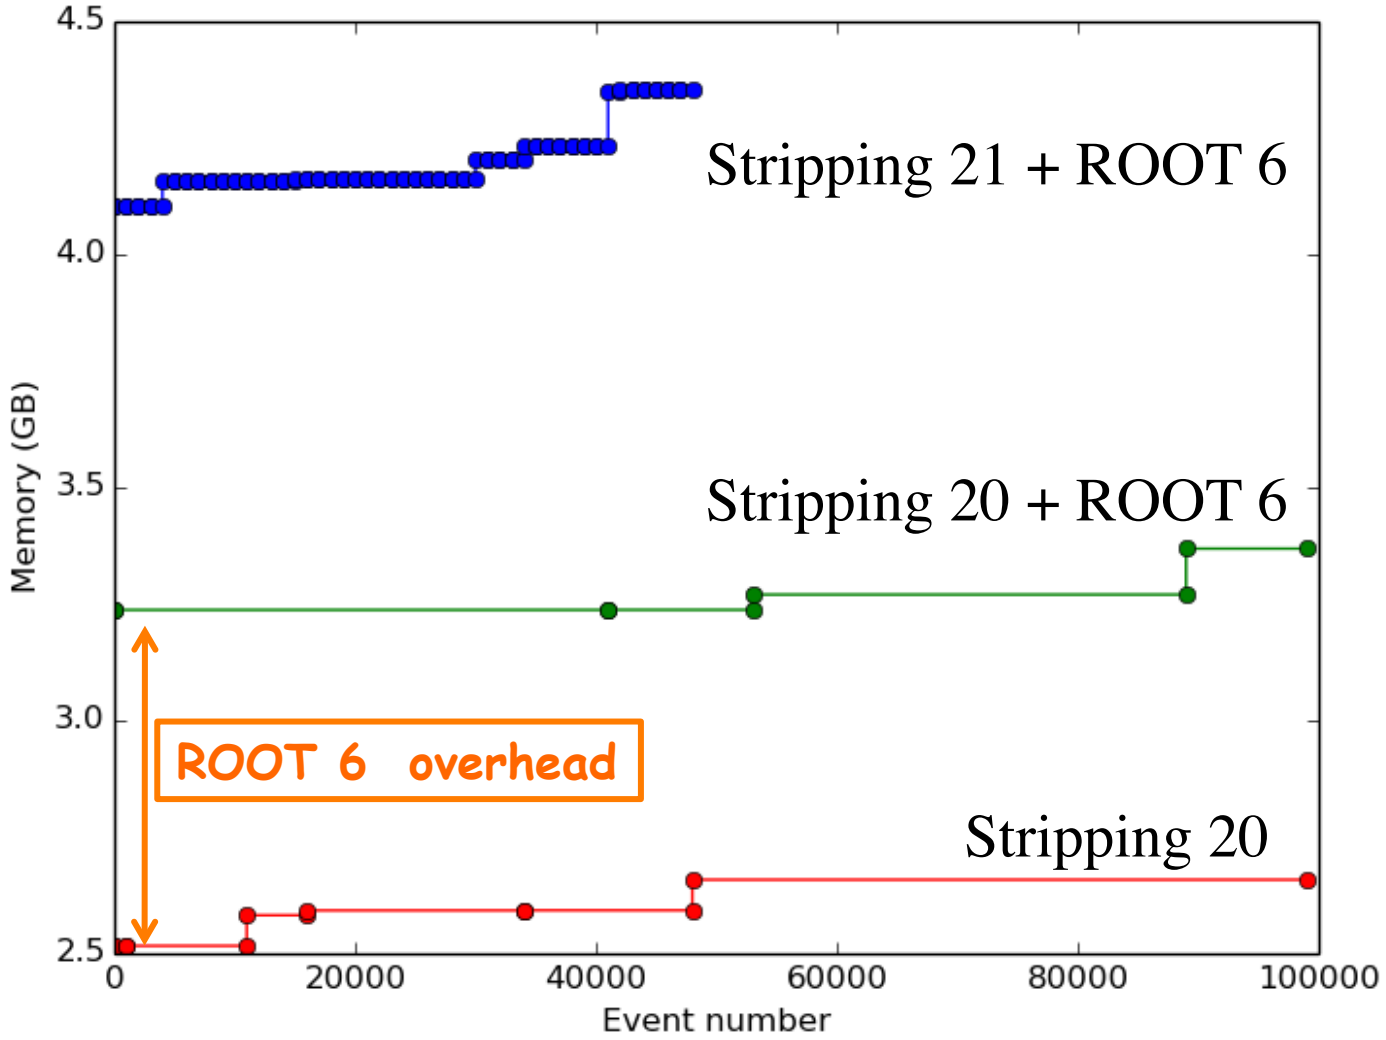

Migration to ROOT 6

- Decided in May to migrate to ROOT 6, before run 2
  - Using ROOT 6.00.00 as default since June
    - Several minor issues identified and solved promptly by development team
    - It was a good choice to expose users early, at a relatively quiet time, and get early feedback
  - Memory footprint of dictionaries is a major concern
    - Some improvements and workarounds available in next release
    - ROOT 6.02.00 scheduled for next week
- Baseline for 2015 remains ROOT 6
  - But dependent on resolution of memory issues, particularly in HLT environment
    - Benchmarking to take place with next release, early October

Software validation and deployment for 2015

- Reminder: split HLT, online calibration, single offline reconstruction pass (no reprocessing before LS2)
- Goal: perform all necessary calibrations using HLT1 data and apply to both HLT2 and offline
- Status:
  - Split HLT framework and online calibration framework implemented
  - Online calibration and monitoring procedures defined
    - Reconstruction quality demonstrated
  - Conditions database modifications for automatic transmission of new constants to HLT2 and offline implemented
- Plans
  - Monthly end to end integration tests starting in October

Run 2: HLT output streams

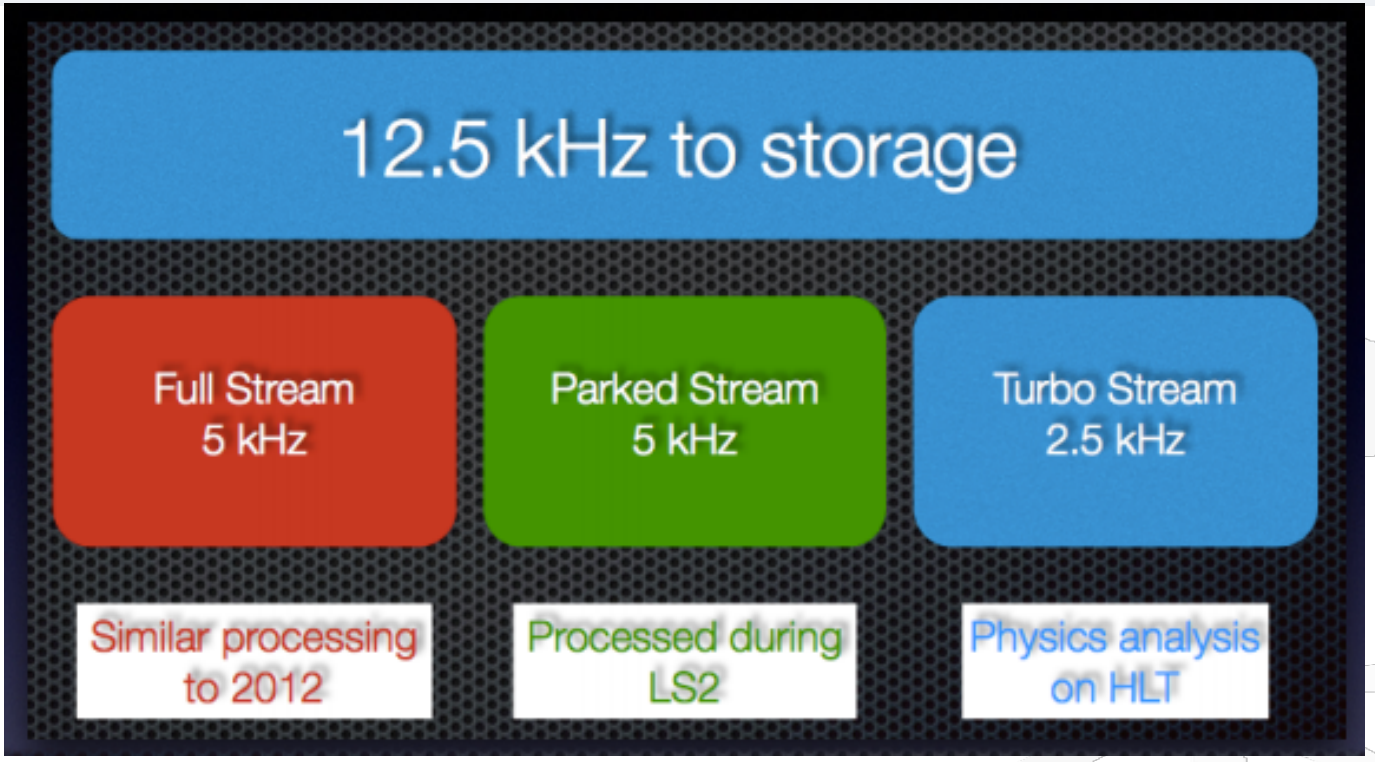

- Full stream: prompt reconstruction as soon as RAW data appears offline
- Parked stream: safety valve, probably not needed before 2017
- Turbo stream: no offline reconstruction, analysis objects produced in HLT
  - Important test for Run3

Turbo stream

- Expect similar tracking+PID quality online and offline
  - Why not do "offline" event selection already online?
  - Suitable for analyses that don't require offline 'extras' such as flavour tagging
    - e.g. charm spectroscopy
- If event reconstructed and selected online, write MDST-like output from HLT instead of the full RAW data
  - Allows larger HLT rate due to x10 smaller event size
  - To begin with (proof of principle) write also the RAW data
    - But skip offline Reconstruction and Stripping steps
- Basic framework already in place
  - Offline analysis tools running on Turbo stream output:

TupleToolKinematic

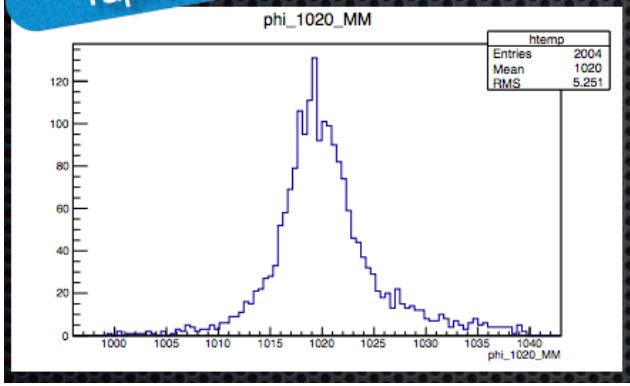

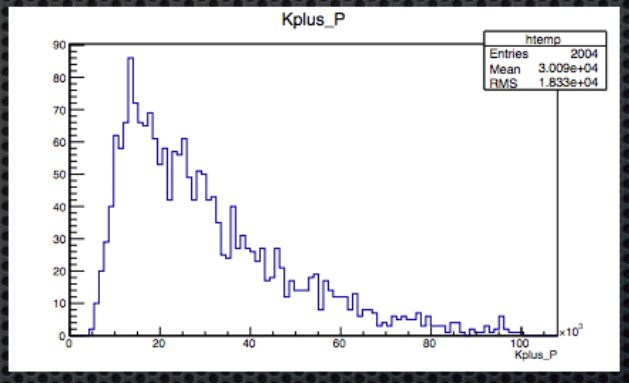

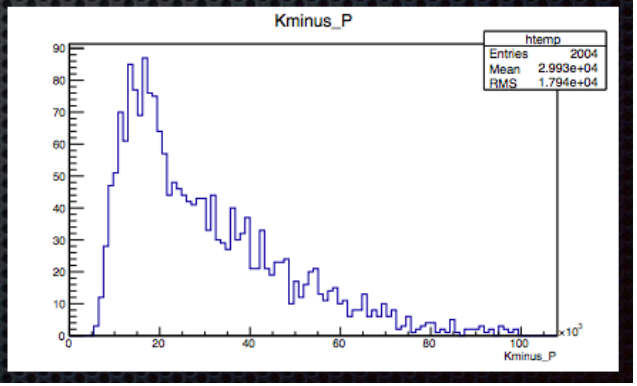

Run3: TurboDST?

- Reminder: in Run 3, full LHCb physics potential achieved with 100kHz HLT output rate
  - Order of magnitude more than in Run 2
- If Run 2 Turbo stream proof of concept successful, can it be extended to a large fraction of the Run 3 HLT rate?
  - If only MDST content is saved (no RAW data), event size is reduced by order of magnitude
- More generally, as online reconstruction approaches offline reconstruction, can we save a "DST-like" dataset out of HLT and drop the RAW data?
- Currently just ideas, but would allow a 100kHz HLT output rate without order of magnitude more computing resources
  - Run 2 can be used as a test bed for these ideas

2015 resources

- Status of 2015 pledges
  - Includes Russia Tier1 for the first time
  - For the first time in my memory, requests are fully satisfied
  - No worries for 2015 – thank you to all countries

	CPU request	CPU pledge	Disk request	Disk pledge	Tape request	Tape pledge
	kHS06	kHS06	PB	PB	PB	PB
Tier 0	36	36	5.5	5.5	11.2	11.2
Tier 1	118	139	11.7	14.2	23.7	28.0
Tier 2	66	60.5	1.9	1.9		

Comparison with "flat budget"

- Definition of flat budget: same money will buy
  - 20% more CPUs
  - 15% more disk
  - 25% more tape

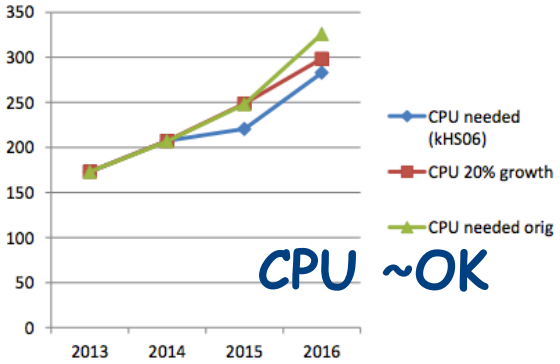

CPU ~OK

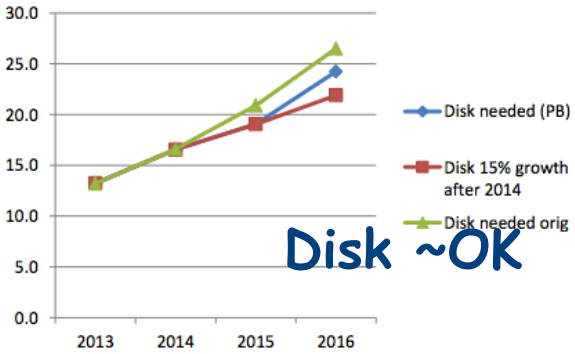

Disk ~OK

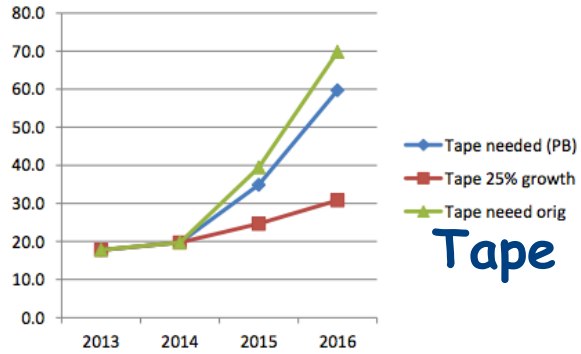

Tape explodes

- **A note on tape costs:**
  - Costing models assume a large (~10%) disk cache in front of tape
    - Drives total ownership cost of tape storage
  - LHCb computing model does NOT require such a large buffer
    - ~50% of tape storage is for archival only (write once, read ~never)
    - Bulk recalls are during planned stripping campaigns
    - Can allocate temporary storage to sustain required I/O rate

Conclusions

- Operations
  - Business as usual, no major issues
  - Data popularity is proving powerful tool for keeping disk usage under control
- Preparations for Run 2
  - Software commissioning in full swing
  - No concerns about resources for 2015
    - Need to address tape provisioning for later years
- Upgrade
  - Run 2 as a test bed for computing model changes for the future
  - Computing milestones defined:
    - Feb 2017: TDR
    - Sep 2018: Computing model for Run 3