LHCb status, plans for RunII and beyond

Raja Nandakumar

For LHCb

Science & Technology Facilities Council

Content

- Current status
- RUN-II plans
  - Dirac developments for Run-II
- RUN-III requirements
  - Upgrade thoughts contributing to these needs
- Not covering ET (Vac, ...) - covered by Andrew McNab
  - Some words on ipv6
- More information
  - 3rd LHCb computing workshop http://indico.cern.ch/event/278289/
  - WLCG workshop in Barcelona
    - https://indico.cern.ch/event/305362/session/1/contribution/14/material/slides/0.pdf
    - https://indico.cern.ch/event/305362/session/1/contribution/14/material/slides/4.pdf

Brief reminder

- LHCb computing model – what works at different sites
  - Tier-2s + non-WLCG resources
    - Any job that does not need data
      - Simulation, toy MC, …
    - Reprocessing
      - Co-processors to Tier-1s
  - Tier-2Ds
    - Tier-2 + User analysis
    - Have copy of some of the data
  - Tier-1 / 0
    - Everything

UK sites have been quite good. So used a lot.

Since GridPP 32

LHCb on the grid

Mostly simulation jobs

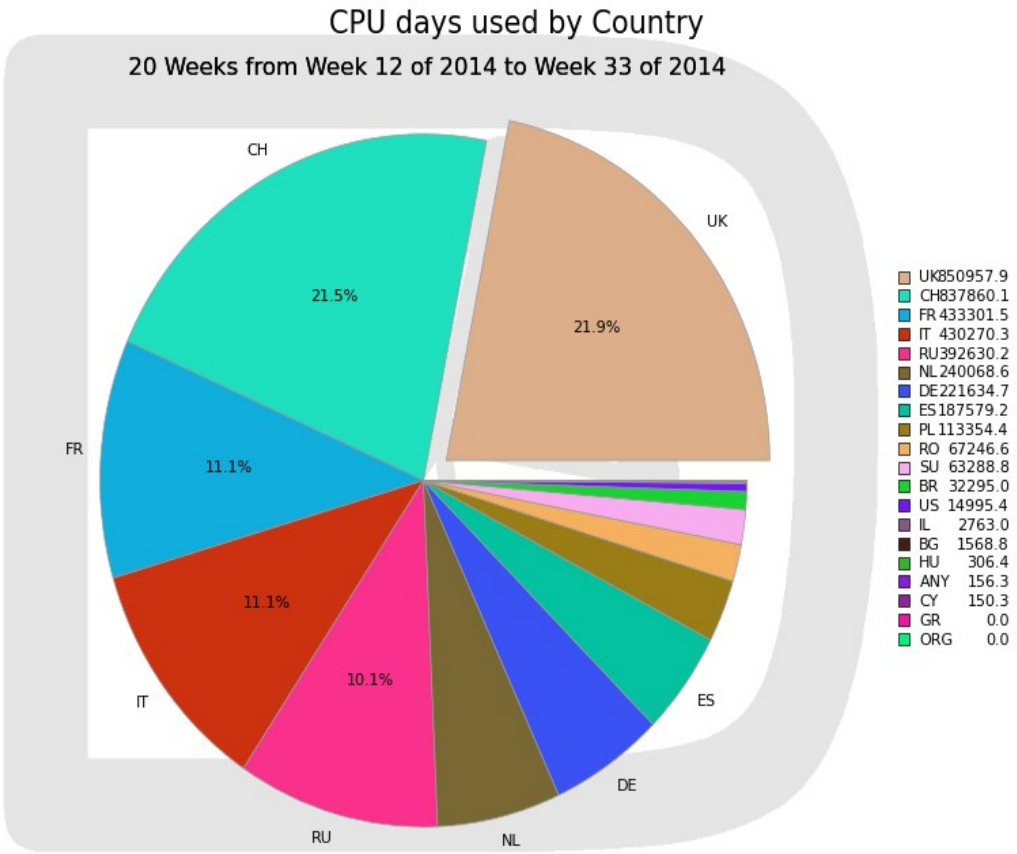

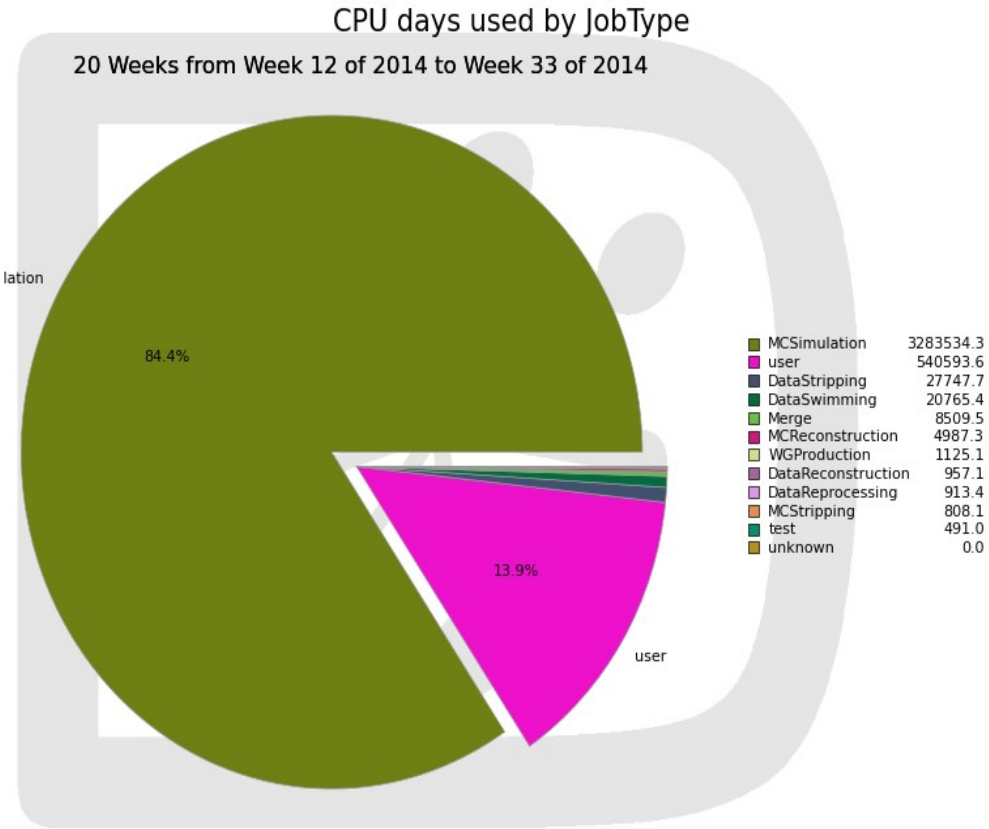

2012 Data taking

- Data rates decide most of the computing model
  - Luminosity levelling
- Pileup 2.5 int / bunch x-ing
- LHC duty cycle ~30%
- Overcommit HLT
  - Catch up during breaks

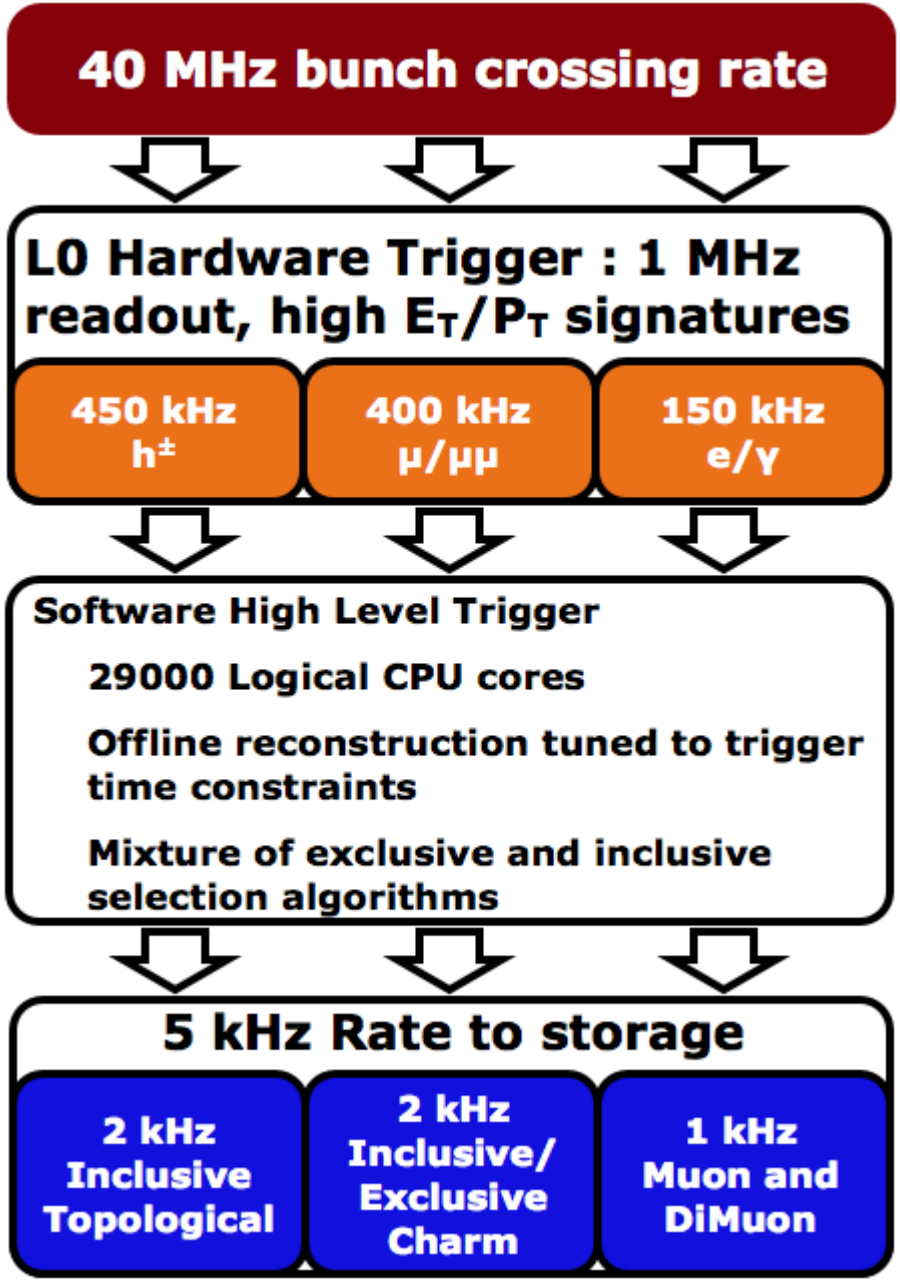

Goal

Run-II

- Extension of 2012 buffering
  - Real-time calibration run-by-run
  - "offline-quality" PID in HLT
  - Will not need full reconstruction until next break
    - Publication quality data within hours of data taking
- No reprocessing until LS2
  - No need for access to RAW on tapes
    - Process while data is still buffered

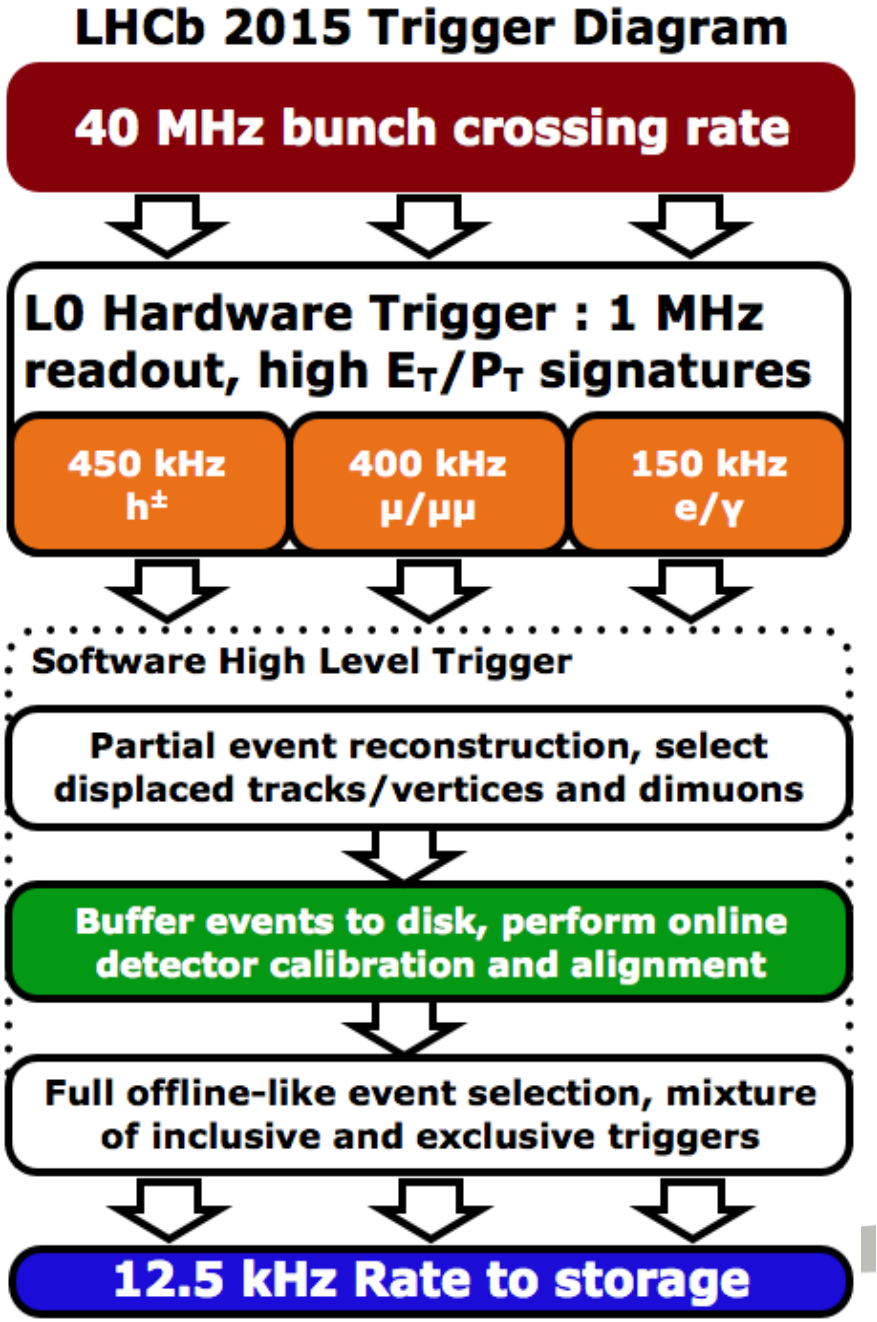

Calibrations

- 2012 experience
  - Semi automated calibration before offline reconstruction
  - 2 – 3 week delay
- Run-II plan
  - Automate execution of alignment & calibration jobs
    - Real time on online farm on HLT1 triggered data
  - Automate validation of constants and propagation to conditions dB
    - Calibrations available to HLT and offline at same time
    - No need to buffer RAW data while waiting for the calibrations

Stripping, Simulation …

- Workflows unchanged from 2012
  - Reconstruction → Stripping → Merging → Analysis
- Limit to 2 re-strippings per year
  - ~6 weeks due to tape access throughput
  - Various internal changes in formats to improve storage efficiency
- Use opportunistic resources for simulation
  - Including outside WLCG : Yandex, BOINC, …
  - Note : HLT farm is used for deferred processing

Dirac and pilots

- Pilot code being re-written
  - Modular approach
  - Only CVMFS for software
  - Same pilots for all resources
    - Batch, IaaS, Vac, Boinc, …
  - Possibility to run probes within pilot and inject to SUM
- Dirac task queues
  - Job descriptions can contain arbitrary "tag" matching a queue within DIRAC
  - First tests with multi-core jobs

Storage management

- srm is the current data access protocol
- Direct xroot access has been implemented and deployed
  - In production soon
- Storage federation
  - Jobs have local catalog file
    - Access next nearest replica if local file is inaccessible
    - Expected rarely - only as a fallback solution
  - Recent limited tests were successful
- Http planned to be implemented in future
- Note : srm will remain for a little longer
  - Space occupancy
  - Data staging for FTS

More DM within Dirac

- LFC → DFC migration
  - Testing phase finished
    - Required developments being finalised
  - Plan to have finished migration by early next year
- FTS3 implementation
  - Using FTS3 within DIRAC since late 2013
  - DIRAC client to be updated using FTS3 rest interface
- GFAL → GFAL2 updates

Other stuff

- CVMFS
  - The only way for LHCb jobs to access software
  - From 2 September : client version >= 2.1.19
  - /cvmfs/lhcb-conddb can be removed from Wns
    - dB distributed using sqldddb snapshots
- New binaries are sl6
  - Mostly – some sl5 binaries still being used
    - These sl5 binaries run fine on sl6 resources
- Tier-2Ds
  - LHCb has funding only for two sites in the UK
  - In general been running without problems
- DPM >= 1.8.8 needed for xroot access

LHCb upgrade

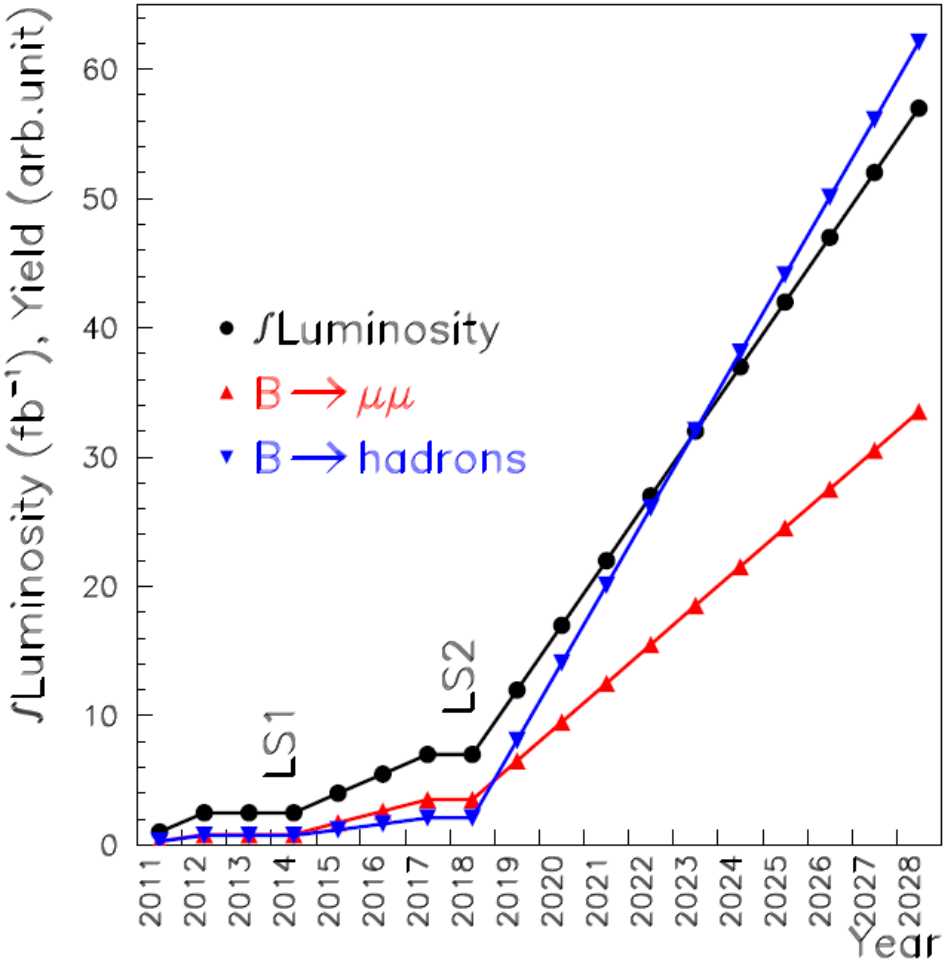

- LHC upgrade from 2018
  - Higher luminosity
    - 1-2 x 10^{33}
  - Pileup ~5 int / x-ing
- x10 or more increase in annual yields for various B final states
- Readout at bunch crossing rate
  - 40 MHz
  - No hardware trigger

Upgrade - strategy

- Read out whole detector every bunch crossing
- Fully software based trigger
  - Deferred processing / data parking as needed
- Various upgrades
  - All front-end electronics to operate at 40 MHz
  - New tracker, RICH, VELO, ...

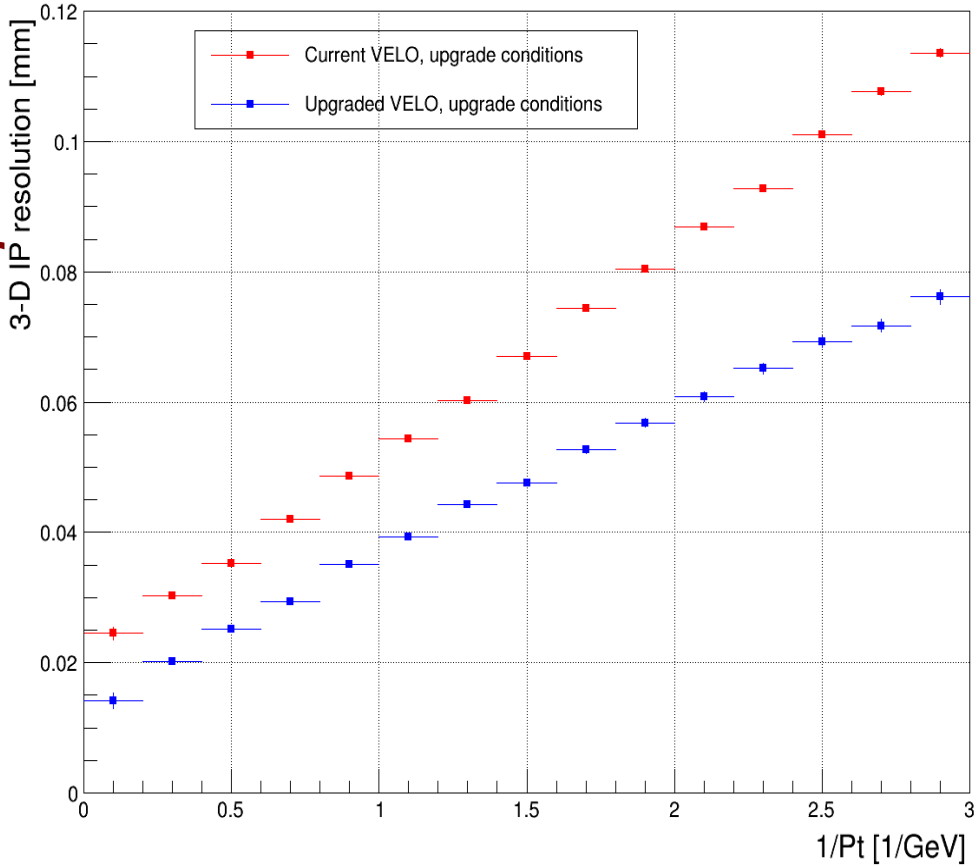

Run-III

- Three different storage rate scenarios
  - ~2GB/s (20KHz)
  - ~5GB/s (50KHz)
  - ~10GB/s (100KHz)

Selection	Output rate [kHz]		
Topo	10	25	50
LT unbiased	1	4	5
Excl. b	ϵ	1	3
Incl. di-μ	-	-	2
Charm	9	20	40
Total	20	50	100

- "Trigger" on as much signal as needed
  - Some amounts of creativity needed
  - Decide in 2018

Turbo-dst in Run-III

- Turbo–dst
  - RAW + HLT reconstruction & selection
  - = microDST from stripping
    - But produced in HLT itself
- Q : Can a complete analysis be done with just Turbo-dst?
  - No offline reconstruction, realignment, less chance for PID calibration
  - If yes, can we drop RAW data?
- Will allow ~100KHz trigger without orders of magnitude more computing resources

Thoughts

LHCb and ipv6

- Access to dual-stack machines for sometime now
  - Lxplus-ipv6
  - All tests of LHCb software run fine
    - Access to data, software area, lcg utilities
      - Simple simulation, …
    - Contacting DIRAC services, transferring data
- Next also test on pure-ipv6 machine
  - When a job lands on a worker node without ipv4 access to the outside world
  - Recently got pure ipv6 UI at Oxford
    - Many thanks to Ewan
    - Starting to test various bits and pieces on it
    - CVMFS & LHCb environment setup works "out of the box"
    - First trouble – failed voms-proxy-init

Jens' questions

- Technical aims: more xroot, more cloud, less SRM, less xroot, more http, etc?
  - More xroot, maybe moving to http
  - Would like single endpoint per site. e.g. xroot.gridpp.rl.ac.uk
  - Less srm but it will not disappear
- Scaling, will stuff scale for run 2
  - Yes – we hope. Depends on data rates
- Changes in data models
- Would you be interested in a non-SRM SE?
- And if yes, when do you think you would be able to use one?
  - Yes to points above, slide 9
- Are you getting the right kind of support from the storage group?
  - Yes :) And hope you are getting enough information from us.

Summary

- LHCb running fine (and quite quiet) at present
- Many developments coming up
  - Gearing up for Run-II processing
  - Some planning has started for Run-III
- No major bottlenecks expected in the UK
  - Good performance by UK sites from experience
    - Both cpu and data access

Backup 1 – LHCb streams

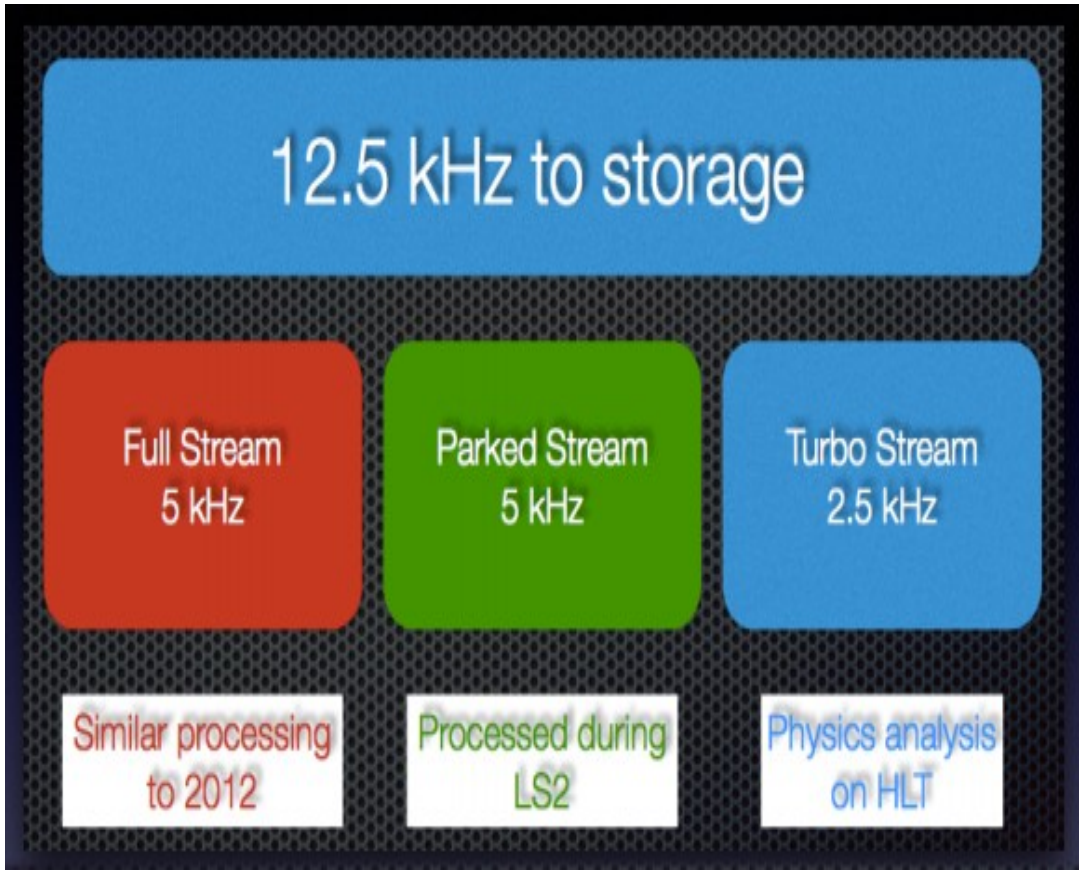

- Full stream
  - Standard data similar to 2012
- Turbo stream
  - No offline reconstruction. All analysis objects produced in HLT
  - Test for Run III
- Parked stream
  - Hope not to use it
  - Safety valve …