LHCb, Echo and more

Raja Nandakumar

STFC / RAL

Overview

- Overall LHCb status
  - ipv6
- Data management
  - Castor
  - CEPH
- Wider issues

UK and LHCb

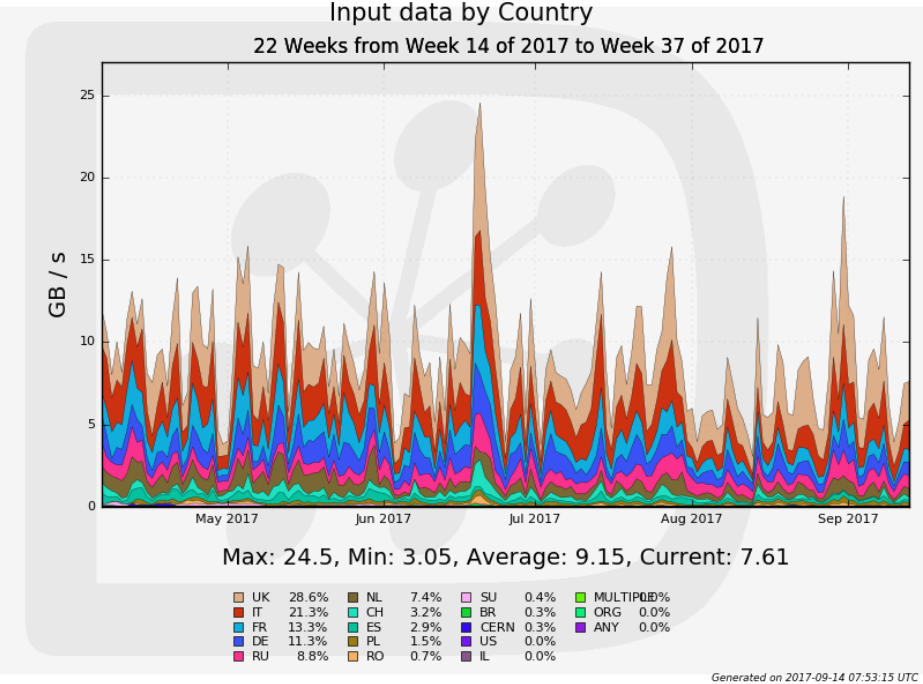

Ignoring CERN, HLT farm ...

- Run where ever allowed
  - Mostly in Europe
  - ~80 sites on the grid
- UK ~30% of LHCb
  - Both Tier-1 and Tier-2 sites are critical for LHCb operations
- Run smoothly most of the time
  - Gaudi based LHCb application stack using LHCb customised DIRAC on the grid
  - Occasional problems usually dealt with GGUS tickets

Science & Technology Facilities Council

LHCb jobs on the grid since Gridpp38

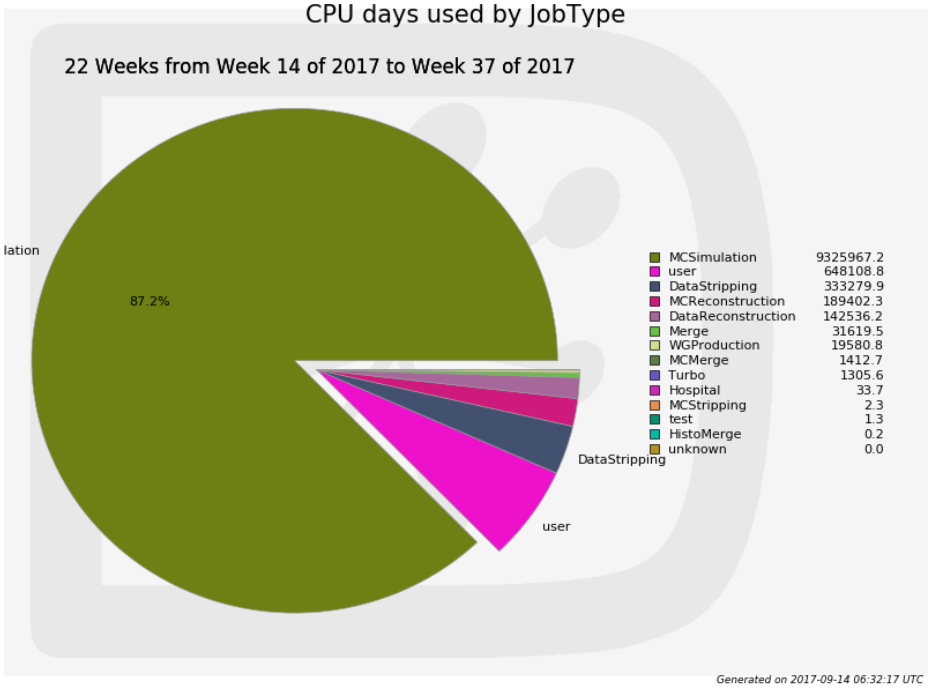

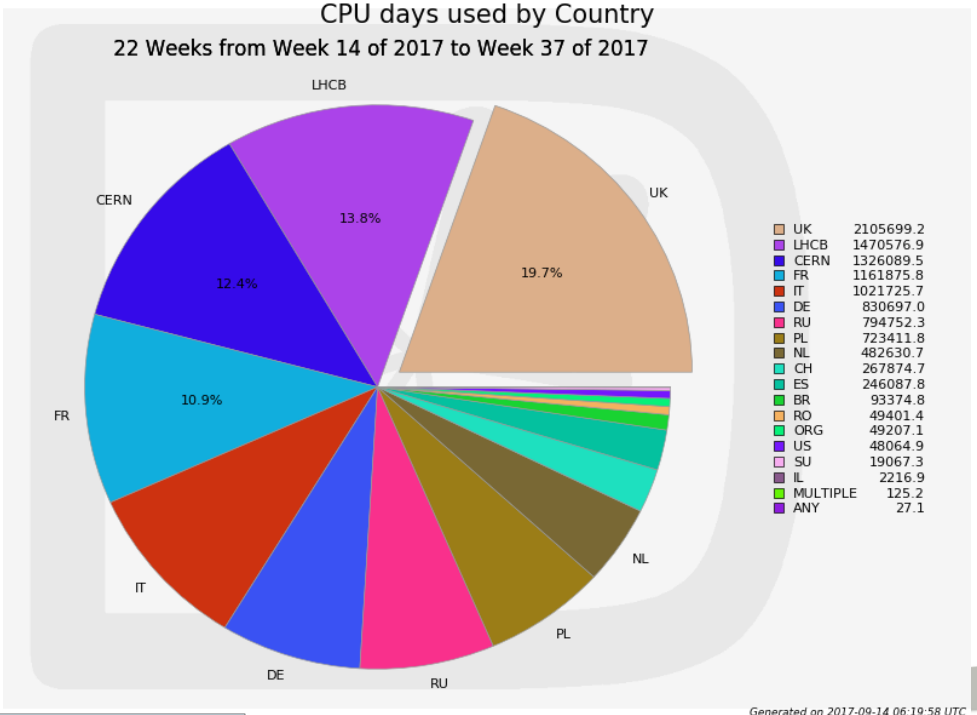

All LHCb sites

Science & Technology Facilities Council

LHCb jobs on the grid since Gridpp38

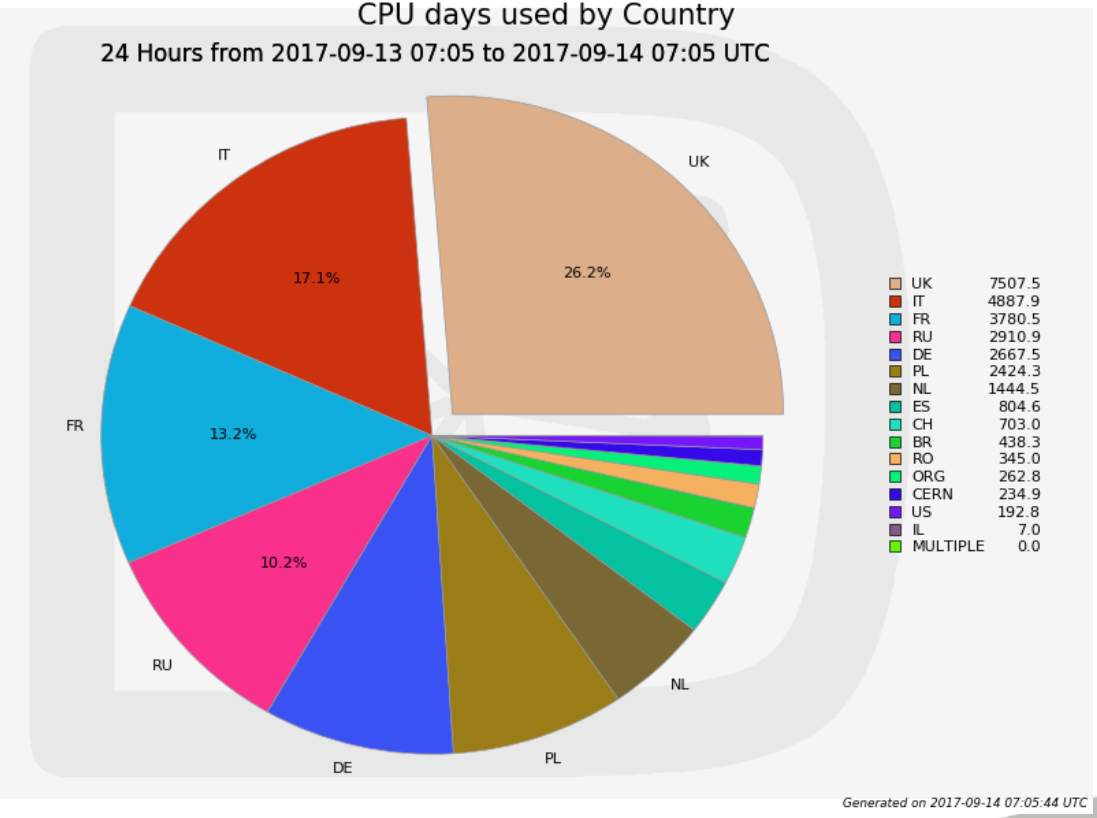

All LHCb sites
except CERN, HLT Farm

LHCb jobs on the grid since Gridpp38

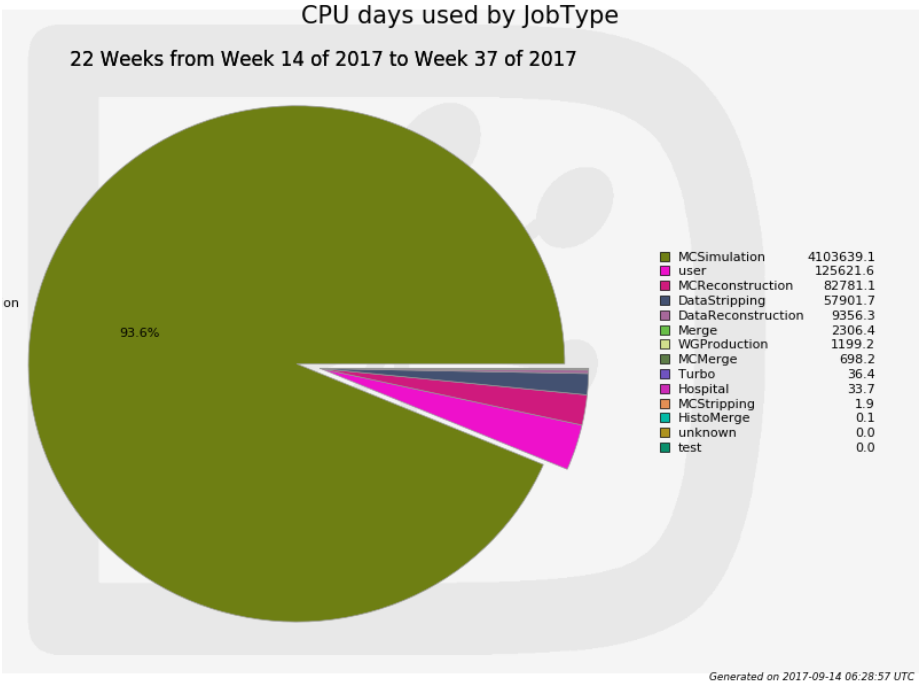

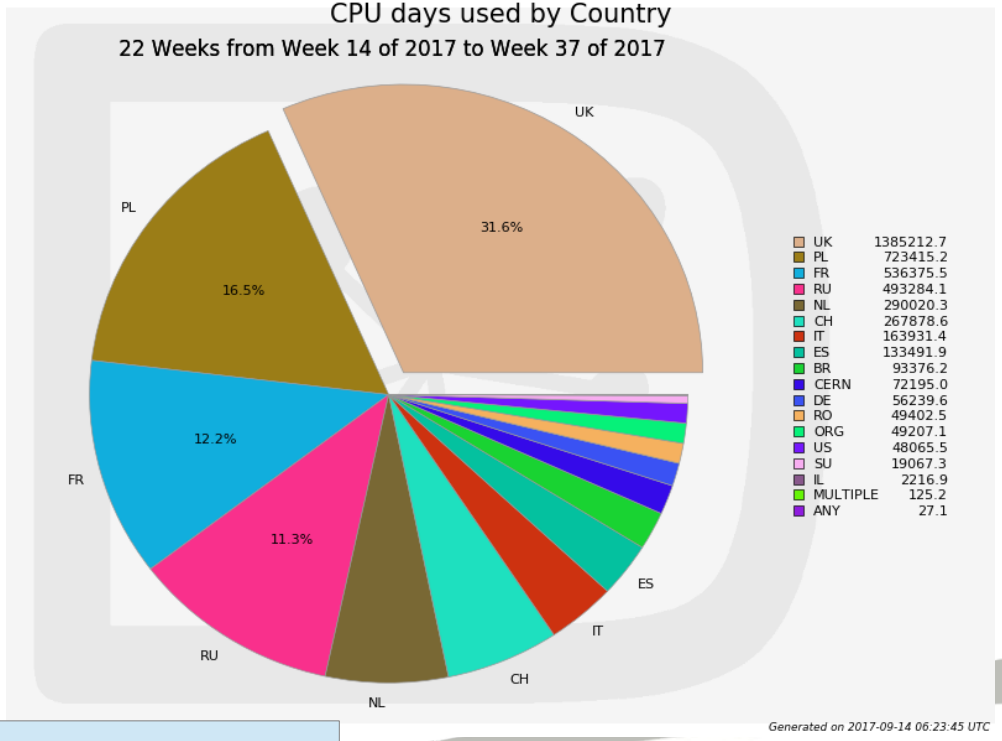

Jobs on LHCb Tier-2s only

Science & Technology Facilities Council

LHCb status

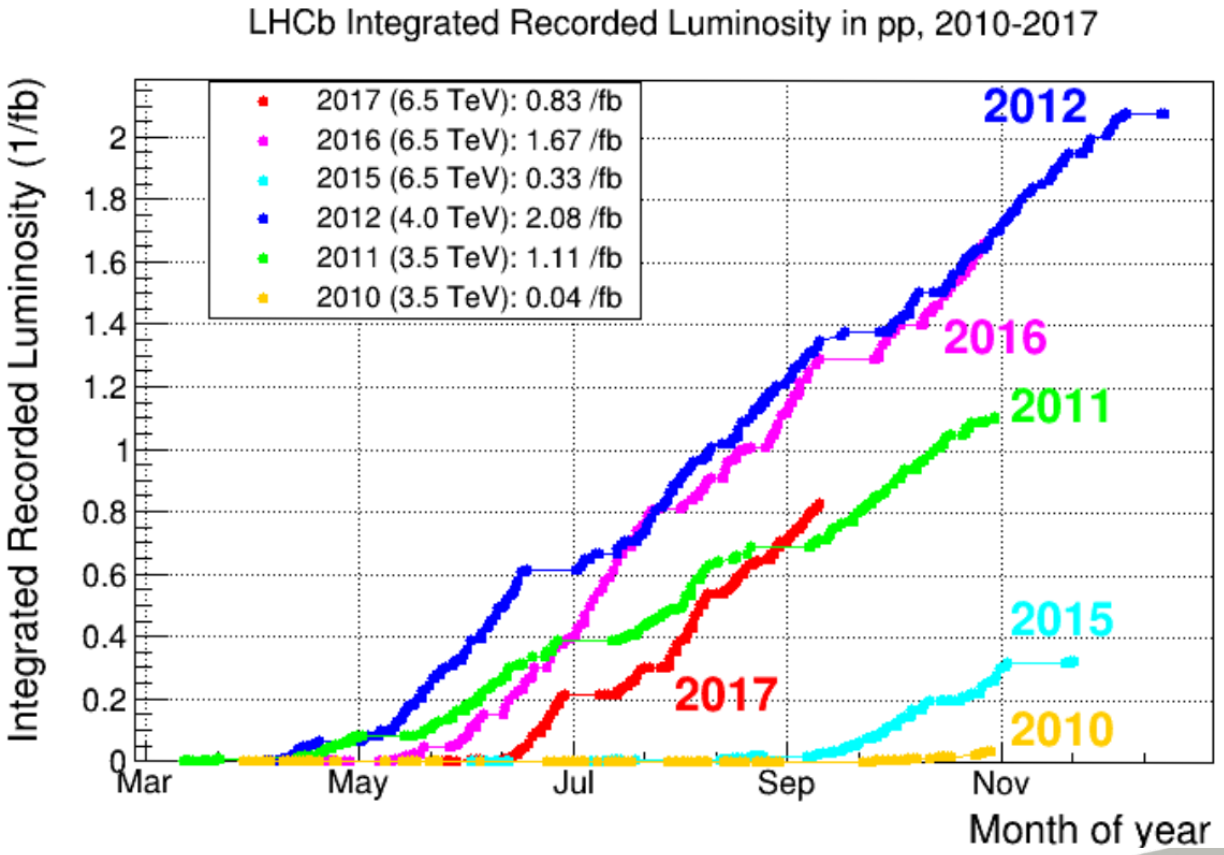

- Taking data smoothly
  - Whatever is delivered
- So far 6 fb^{-1} recorded in total
  - Out of 6.6 fb^{-1} delivered
- Final version of calibration constants in real time
  - No separate offline reprocessing immediately
    - ⇨ Unless issues are found

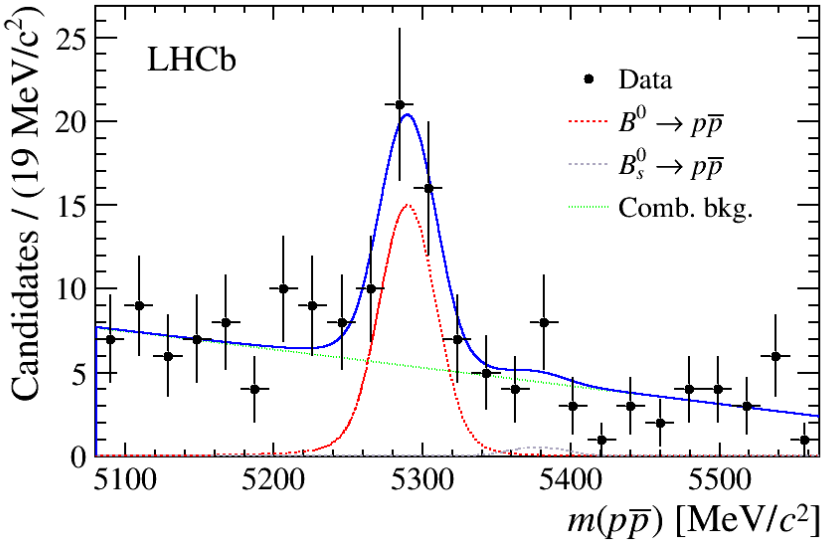

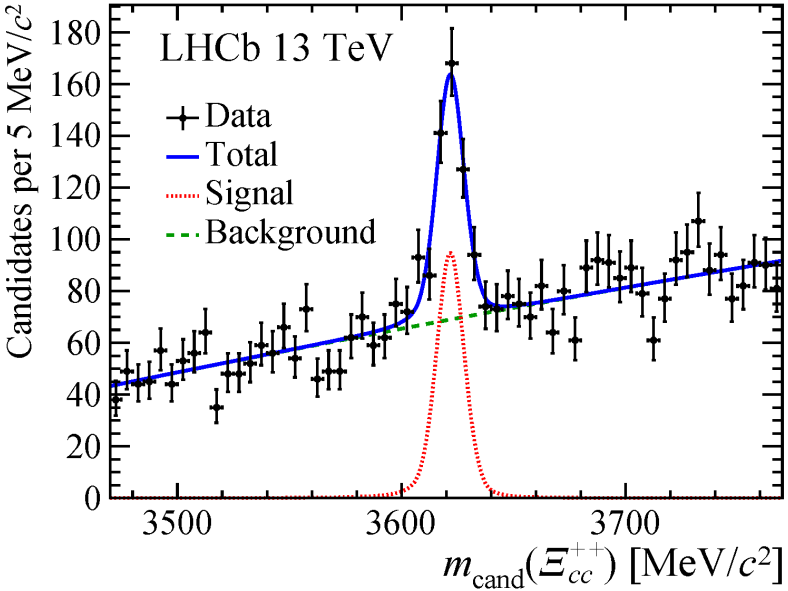

Lots of physics too

- Observation of $B^0 \to p\bar{p}$, Ξ_{cc}^{++}
  - And many other finds over the last few years!
  - So far no hard sign of new physics
- Working ongoing to probe the Standard Model in many different ways
  - Dark Energy searches, LFV, etc.

Science & Technology Facilities Council

LHCb and ipv6

- LHCb operations fully ipv6 compliant
  - Details in Duncan’s talk tomorrow

Data Management

- DIRAC to perform the actual management of data
  - Seamlessly handles DPM, dCache, Castor, StorM, EOS
  - Multiple protocols including gsiftp, xroot, file, …
    - ⇨ Not yet S3 or Swift
  - Works reliably only for SL6 / CC7
    - ⇨ Library conflicts in other cases
- Ensure a minimum of 2 copies of every file in different sites
  - Intermediate files have a single copy during processing
  - Files needed for long-term access also archived to tape
    - ⇨ Not just RAW files
- Still produce MC for 2011 / 2012 conditions
  - Half of data by size in the 2011 / 2012 dataset

Data Management -2

- Need to use srm
  - Needed for staging data from tape
  - Work in progress to remove srm necessity for non-tape accesses
    - ⇨ Get rid of srm-s for Tier-2Ds
- By default use xroot protocol for data access
  - File protocol for StorM SE
- gsiftp for file transfers
  - gfal2 libraries for data management operations
  - Working on moving to xroot for this also

LHCb and storage

- Use RAL and Tier-2 storages in the UK
  - Tier-2 storage = Tier-2D
  - Glasgow, Liverpool, RAL-HEP, QMUL, Manchester, Imperial
- RAL storage
  - Currently use only Castor
    - 12.9 PB for Tape and 4.7 PB for disk allocated
    - Also 1.5 PB disk in Echo (not yet used)
  - Fairly smooth operations in general
    - Problems with castor-srm version clash in April – May
    - Regular diskserver problems take groups of files down

ECHO @ RAL

- Block storage at RAL using CEPH
  - No directory calls possible
- Future of disk storage at RAL
  - Replacement for Castor d1t0 service class
  - Eventually also replacement for the disk in front of tape at RAL
- S3, xroot and gsiftp interfaces
- Smart management of file bits
  - Automatic bypassing of diskservers which are down
  - With erasure coding, needs << x2 storage for recovering files
- Also used for storing CE logs@RAL

LHCb and ECHO – 1

Initial (very basic) tests

- xroot access to files works
  - LHCb applications able to access input files on ECHO
  - Access also successful from outside RAL
  - The most critical test!
- Operations using gfal libraries fail
  - Bug in gfal2, fixed in upcoming version
  - DIRAC uses gsiftp for operations using gfal
- xroot commands work
  - Need a recent version of xroot

LHCb and Echo – 2

Migration from Castor to Echo

- LHCb will be the last experiment to migrate
  - Many thanks for giving LHCb the time to organise itself!
  - RAL staff extremely helpful!
- Tests of transfer of data to be done by LHCb
  - Rest to be managed by RAL
    - Including migration of diskservers as needed from Castor to Echo
  - Will take time, but should be fairly problem free
- Namespace mapping agreed upon
  - Should be fairly straightforward

Will need a xrootd redirector

- Available, but not yet in production
- Test if a file is on echo, else failover to castor
  - One way only
- Only during migration period

LHCb and Echo – 3

Migration continued : issues on the horizon

- dCache does not support third party transfers using xroot
  - Only gsiftp is supported
  - 5 out of 7 (+ CERN) LHCb Tier-1 sites run dCache
    - And a few Tier-2D sites
  - Of course, essential for file replication
- Globus dropping support for open-source toolkit
  - Likely to be supported by CERN, not yet official

LHCb and Echo – 4

- Summary
  - Usage of Echo by LHCb jobs likely to be seamless
    - Jobs okay in principle if streaming data out of Echo
      - xroot protocol
    - Data transfers likely to also be okay
      - Pending release of fixed gfal libraries
  - Early stages of discussion of migration out of Castor
    - Need the xrootd redirector for further testing
    - Fixed gfal libraries too will help
  - Timeline
    - Expect ~1 year of work on this
    - Development to be done by year end

Wider issues

- Mostly running fine around the grid
- Specifically for the UK in the last 6 months
  - RAL
    - Major issue with castor-srm incompatibility in April-May 2017
    - Also problem with location of xroot daemons at RAL in May 2017
    - Low level networking issue uploading out of RAL
      - Mostly gone away since ~ June 2017
  - Tier-2 sites
    - QMUL xroot file permissions issue
    - Various sites having ARC / HTCondor crashes
    - Publishing errors (Queue lengths, OS available, ...)
- Also many issues at other sites
  - Various downtimes and networking problems at SARA
  - Tape access speed at FZK
  - ...
- Also internal LHCb issues
  - Bugs
    - Application, DIRAC
  - Users
    - Multithreaded applications
    - Application errors
  - Library clashes

Overview

- LHCb running well on the grid
  - Many thanks for your support!
  - Lots of interesting physics results published and more coming
- LHCb / DIRAC support most technologies
  - CEs, SEs, ipv-4/6, …
    - ⇨ Even use DIRAC as a batch system (HLT farm, Yandex, ...)
  - Some blindspots still (S3, Swift)
    - ⇨ Not enough manpower
- Working on migration to Echo @ RAL
  - Basic bits in place, rest in progress

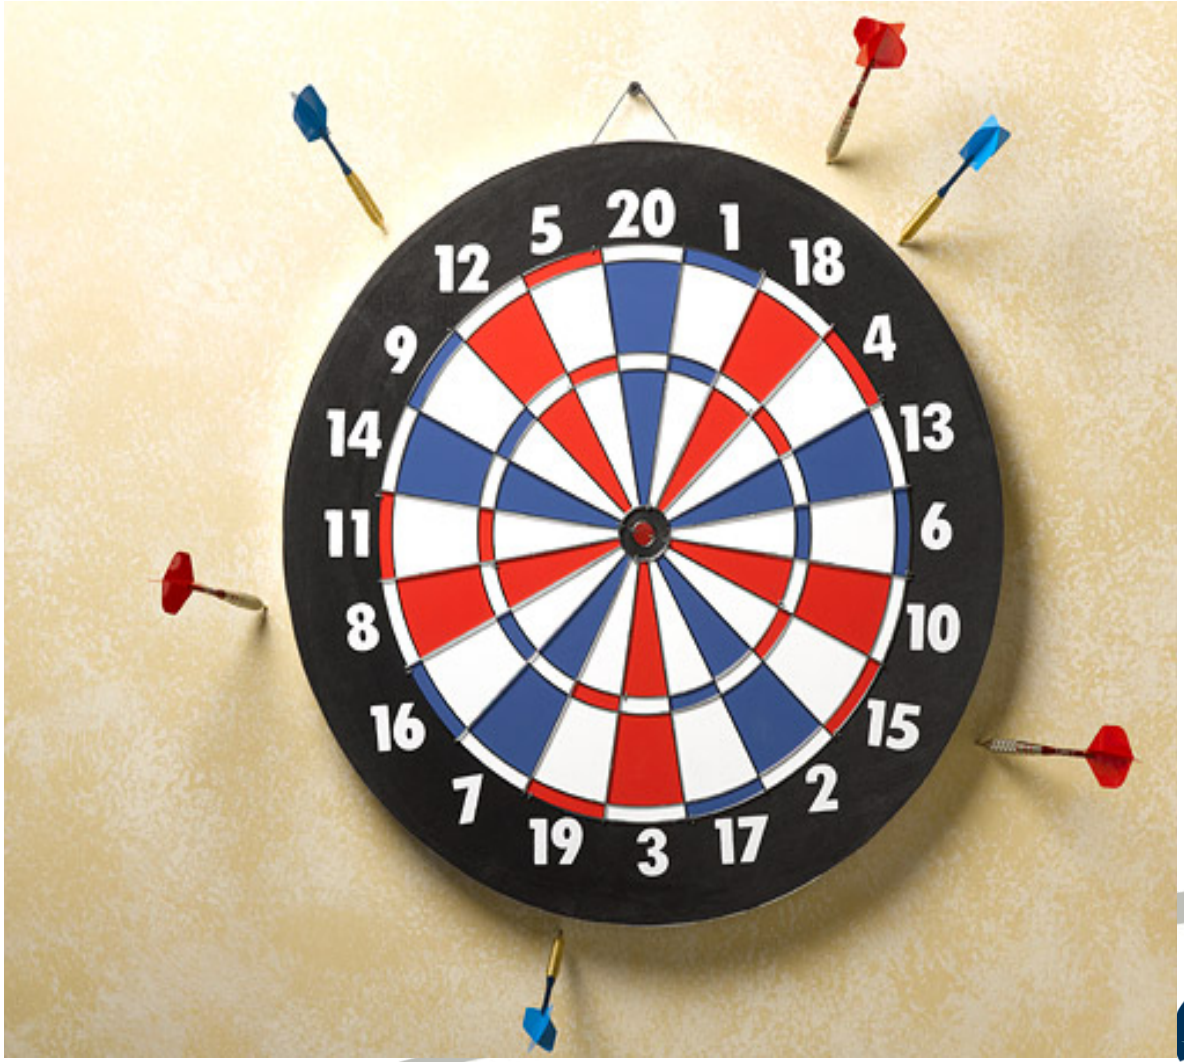

Target practice