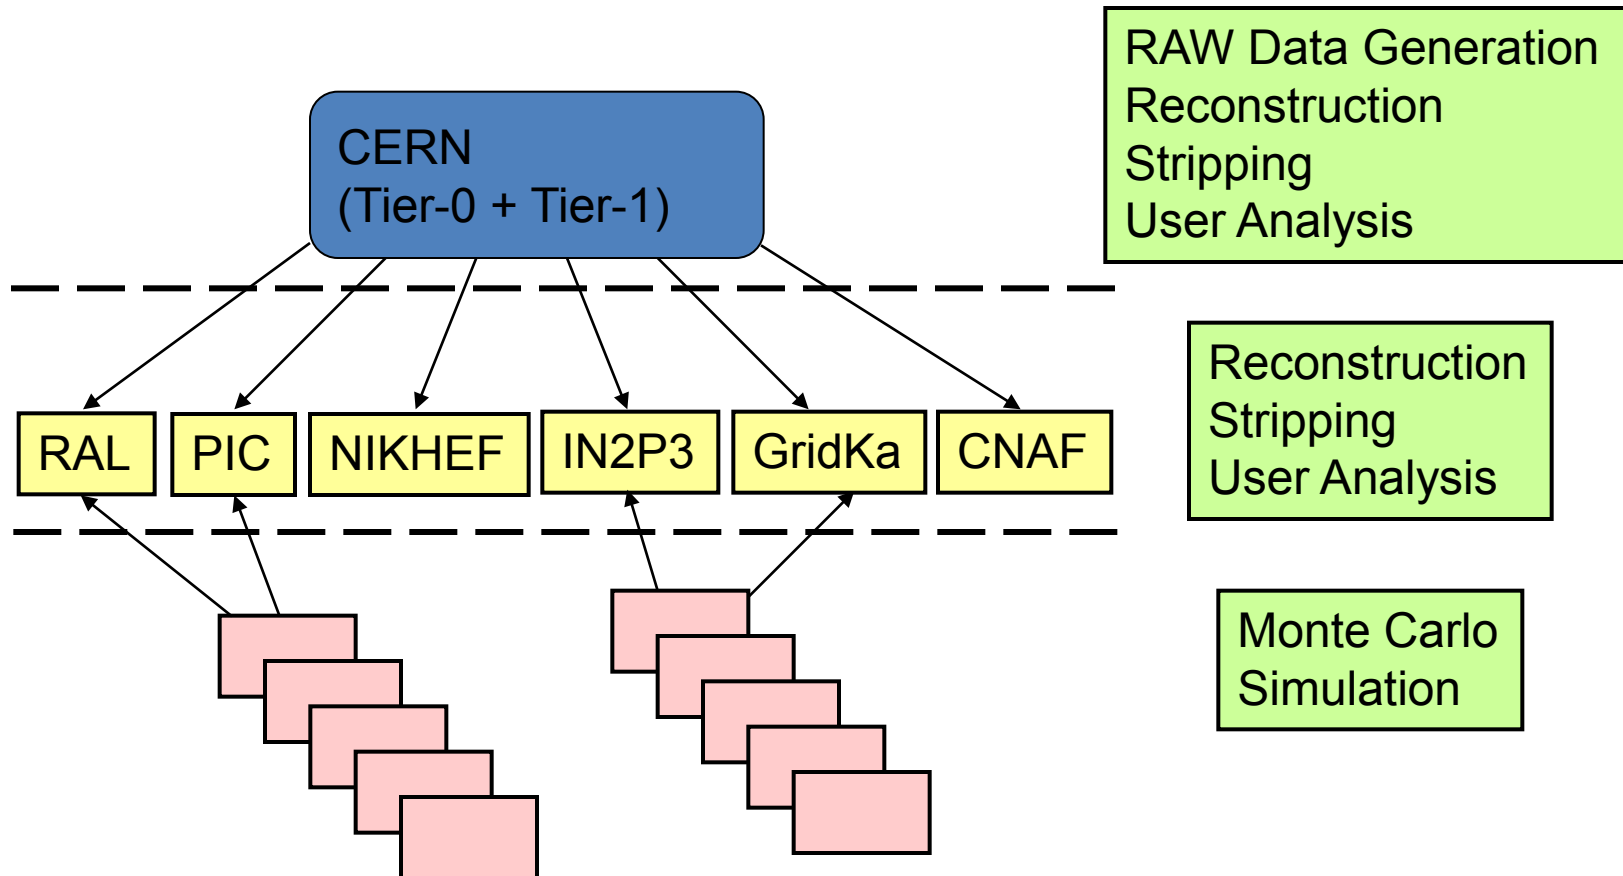


Data in LHCb

Raja Nandakumar

- ❖ Many thanks to all members of LHCb!
- ❖ Computing Model
 - ❑ And what actually happened
- ❖ What LHCb did (and is doing)
- ❖ What the sites did

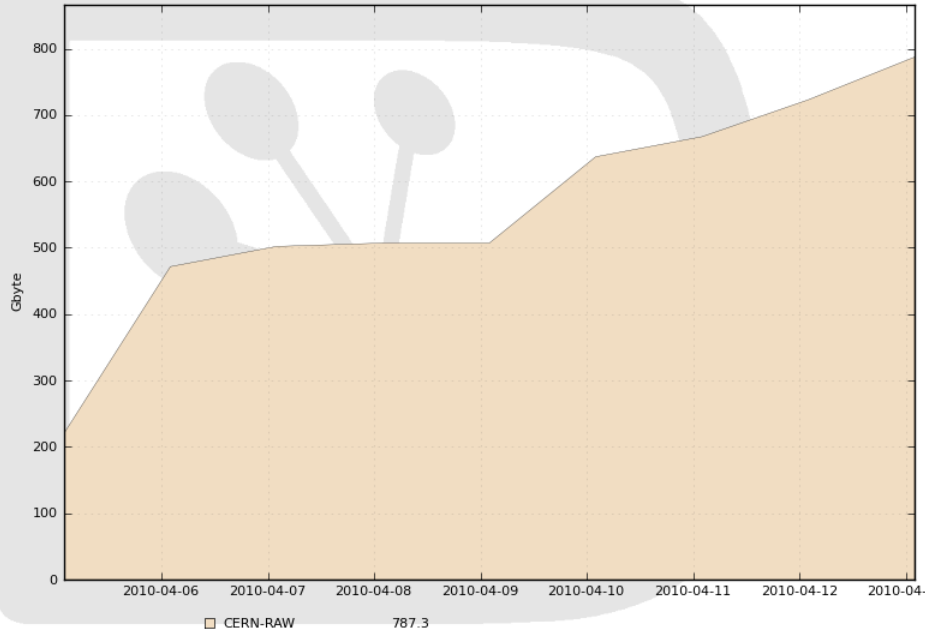
LHCb computing model




⇒ To recap : For LHCb : CERN = Tier-1++

Data distribution

Data transferred from Pit to CERN
8 Days from 2010-04-05 to 2010-04-13



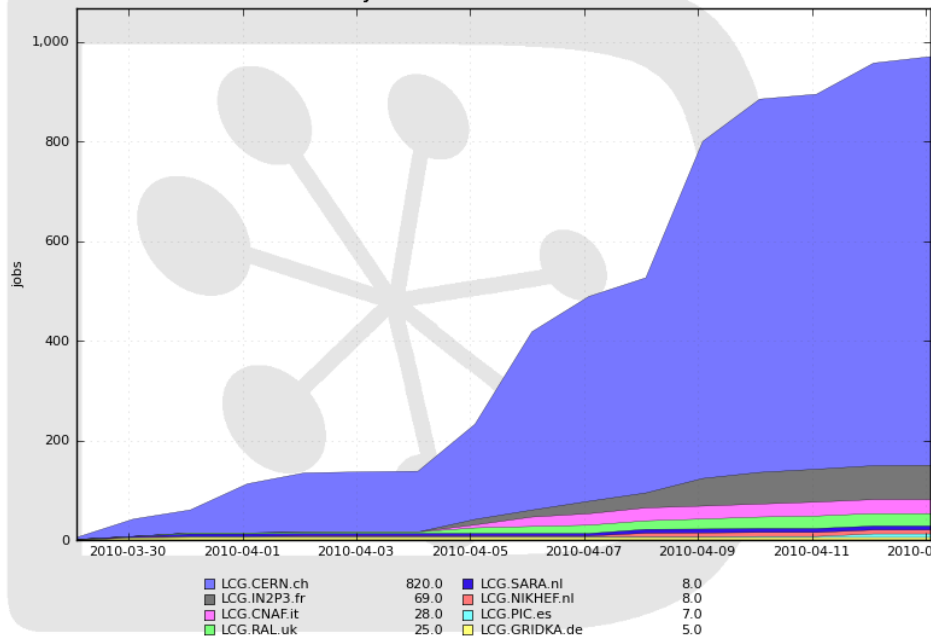
Data distribution around Tier-1s
8 Days from 2010-04-05 to 2010-04-13



- Transferred ~ 800GB data over last two weeks
 - ⊙ Mag Down : 18.7M events, 248 Files, 434GB
 - ⊙ Mag Up : 9.7M events, 86 Files, 226GB
 - ⊙ Mag Off : 3M events, 46 Files, 75GB
 - ⇒ “Usable data” with all detectors in
- RAW file size initially 3GB, now 2GB

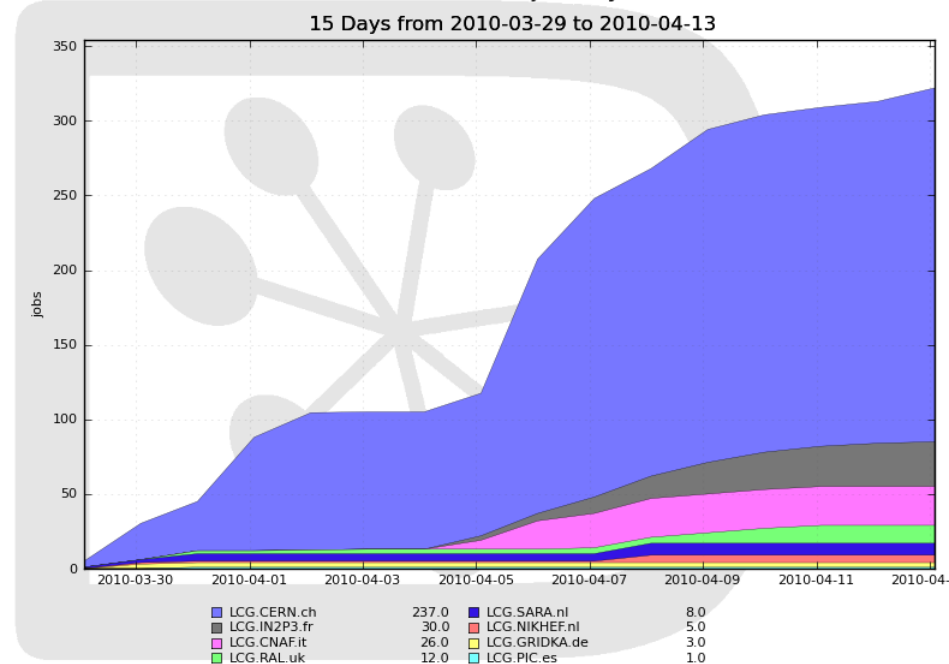
Data processing

Cumulative reconstruction jobs by site
15 Days from 2010-03-29 to 2010-04-13



Generated on 2010-04-13 11:11:17 UTC

Failed reconstruction jobs by site
15 Days from 2010-03-29 to 2010-04-13




Generated on 2010-04-13 11:12:44 UTC

- ~1000 data reconstruction jobs run so far
 - ⊙ New RAW files picked up automatically and processed
 - ⊙ CERN doing bulk of work (workflow bug)
- Currently ~0.3 - 0.5 seconds per event for reconstruction
- No stripping so far (see later)

User jobs by site

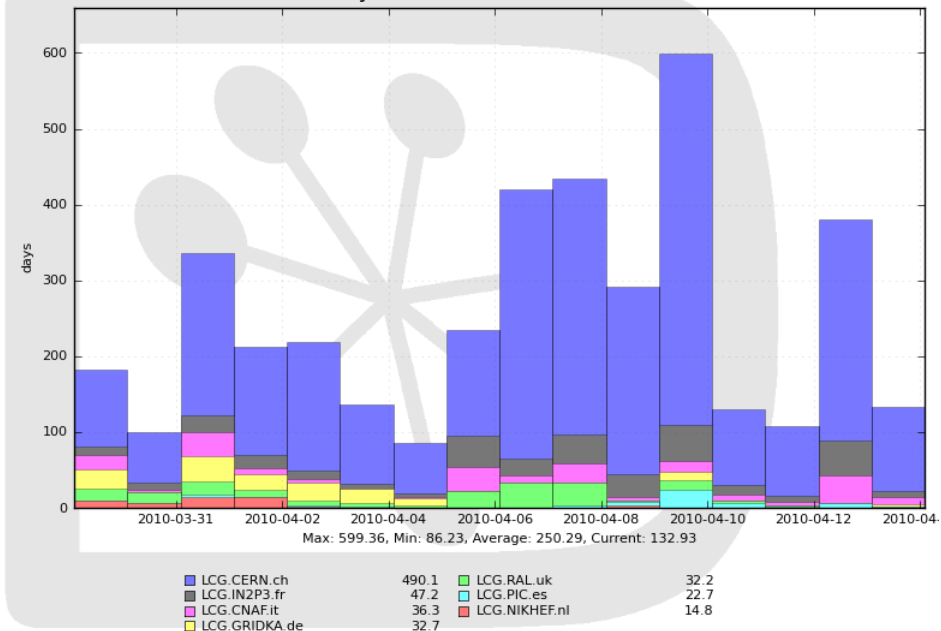
15 Days from 2010-03-29 to 2010-04-13



Generated on 2010-04-13 11:56:59 UTC

User jobs CPU used by site

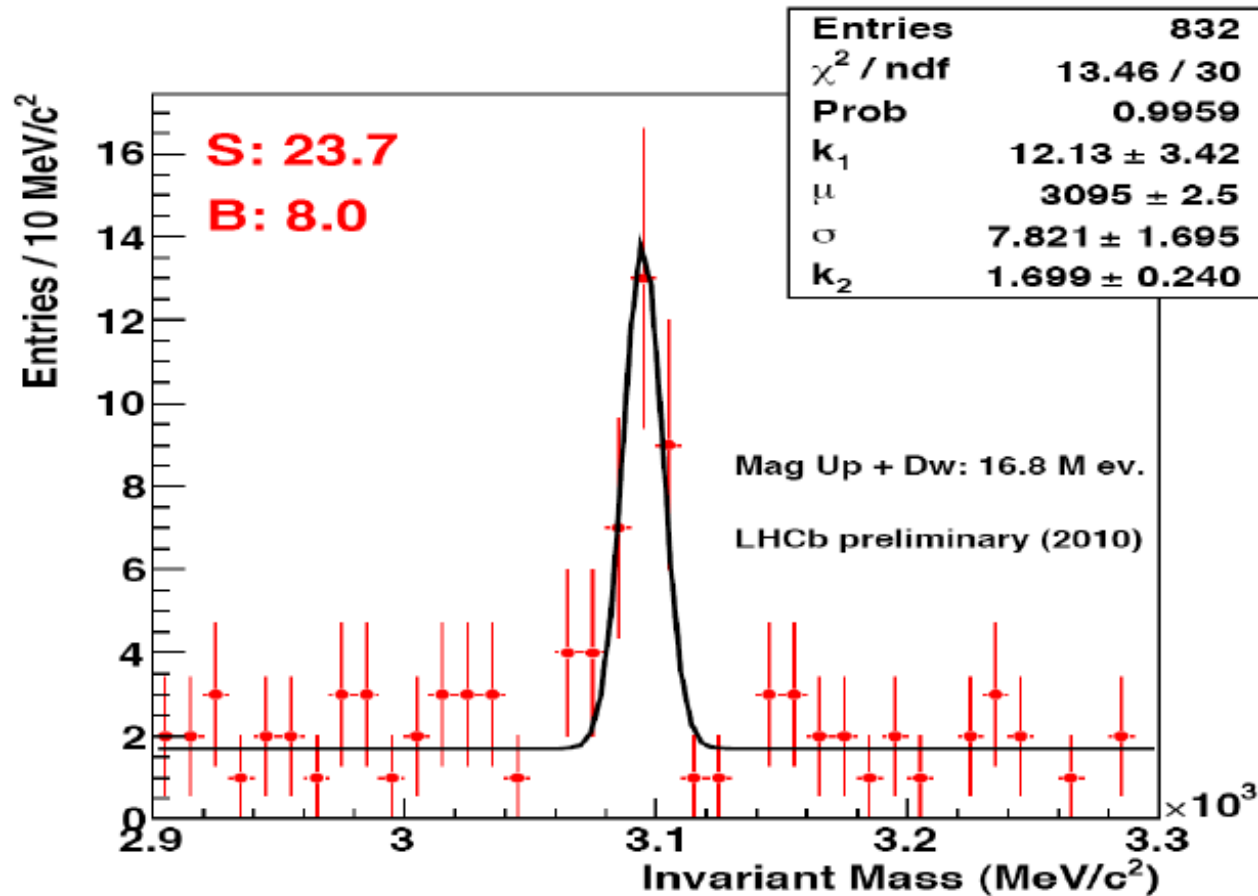
15 Days from 2010-03-29 to 2010-04-13



Generated on 2010-04-13 12:02:30 UTC

- ~90K user grid jobs in the last two weeks
 - ⊙ Many - if not most - users now processing data on the grid, mostly using Ganga
- ~1.2 hours per job on average over all users
 - ⊙ Varying (and improving) success levels
 - ⇒ Sometimes, only resubmit until successful

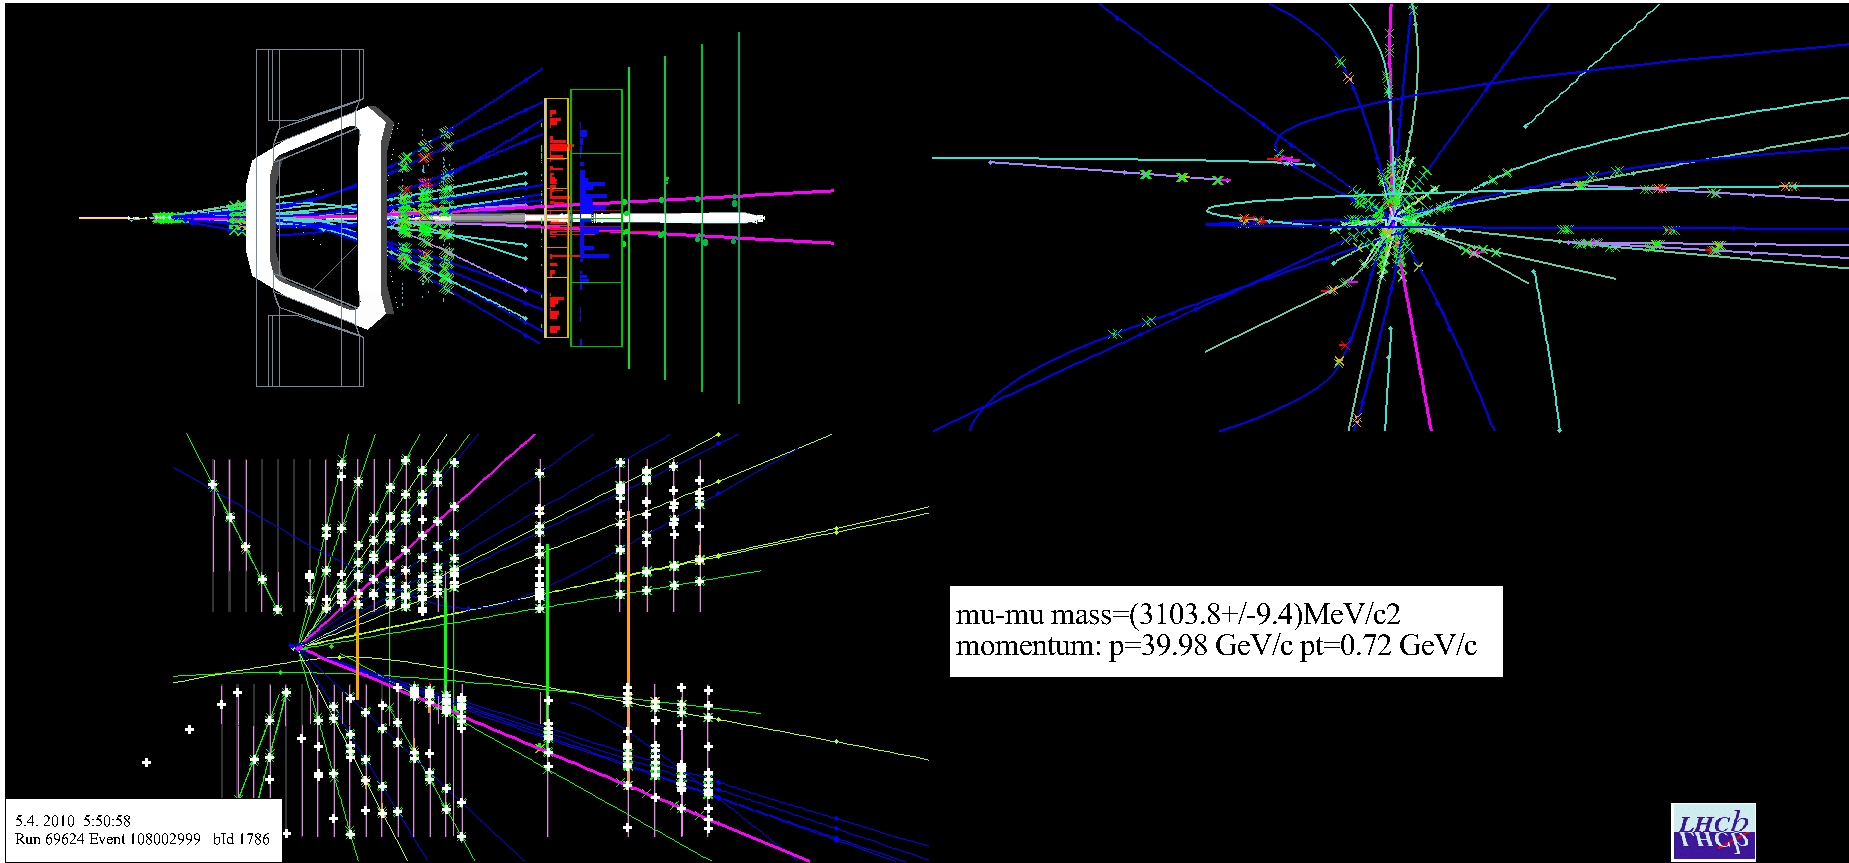
J/ ψ peak in data

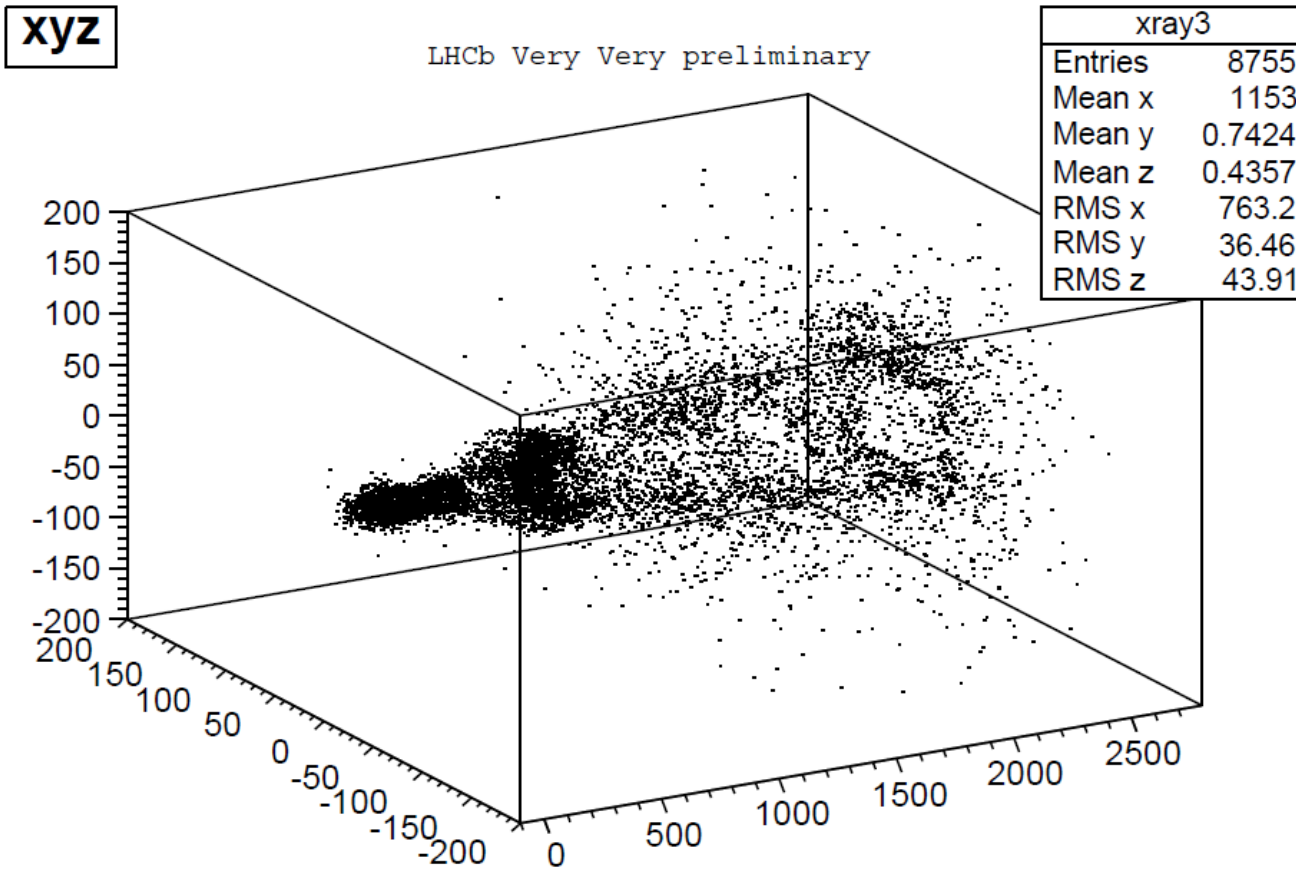


- Very very preliminary
 - ⊙ Calibration, reprocessing, ...
- Many other resonances also seen and being studied
 - ⊙ Charm physics in full steam

J/ψ event in the detector

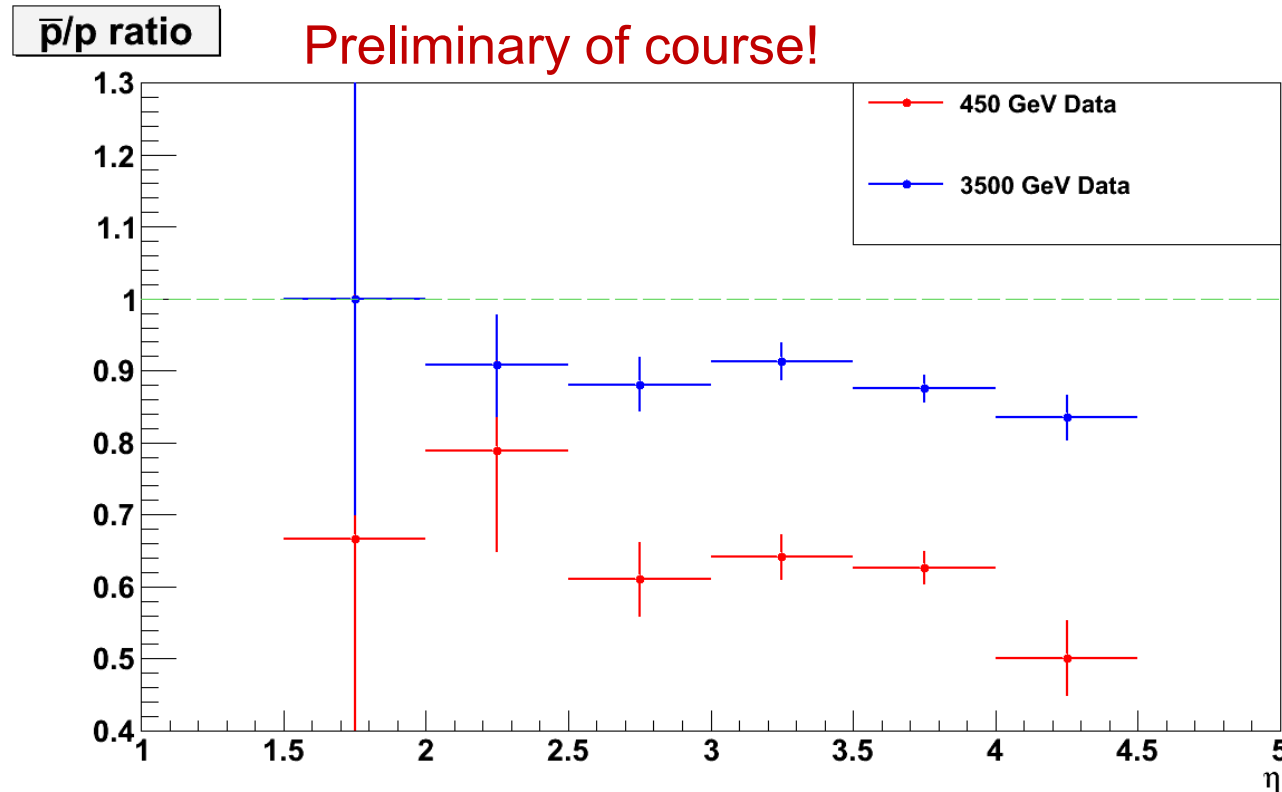
LHCb Event Display





- LHCb detector using photon conversions
 - ⦿ Of course very preliminary ...

- To compare 900GeV (2009) vs 7TeV data
 - ⊙ Look at enriched ratio of p/p vs η
 - ⊙ Asymmetry is expected at lower energy




- Most runs with all detectors in
 - ⊙ Very little “bad” data
 - ⇒ Not much to say except that it all works (mostly)
 - ⊙ Occasional (rare) event with corrupted data
- Application software
 - ⊙ Many new problems found for reconstruction
 - ⇒ Very large events, data corruption
 - ⇒ Fixes coming this week, but working for now (resubmit!)
 - ⊙ Data stripping does not work yet
 - ⇒ Awaiting fixes to be launched
 - ⇒ At present run over all reconstructed data
 - ⊙ Some “time” issues (UTC vs local times ...)

Failed user jobs

Failed user jobs by problem

15 Days from 2010-03-29 to 2010-04-13



Generated on 2010-04-13 12:32:04 UTC

Most failures due to application errors

- ⊙ Not necessarily problems with applications
 - ⇒ Many due to bugs in user code
- ⊙ Bugs in core applications (DaVinci, Brunel, ...) still being found and fixed

Many “stalled” jobs

- ⊙ Low cpu efficiency
 - ⇒ Killed by DIRAC watchdog
- ⊙ Grid middleware problems
 - ⇒ Next slide

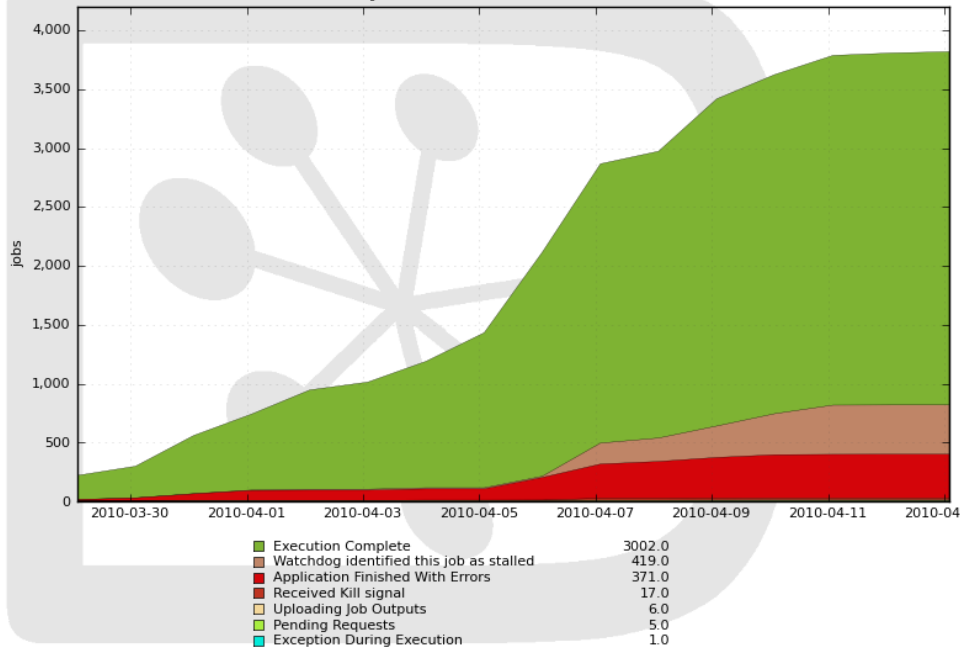
ConDB / LFC access

- ConDB used to store detector conditions
- Oracle authentication for ConDB ...
 - ⊙ Only username / password
 - ⇒ No certificates / proxies
 - ⇒ Cannot be distributed
 - ⊙ Workaround : use LFC to store Oracle endpoint and authentication
- Use CORAL software to connect to grid databases
 - ⊙ Inefficient LFC usage by CORAL
 - ⇒ Keep connection open forever
 - ⇒ LFC server threads exhausted with many jobs
 - ⊙ Temporary workaround for DIRAC jobs
 - ⇒ Still waiting for fix to CORAL
- Endpoint changes (CNAF, RAL) not properly handled
 - ⊙ Fixed release contained typo for RAL
 - ⊙ Finally now it should work
- Standard symptom : “Watchdog identified job as stalled”

At RAL by status

All jobs at RAL

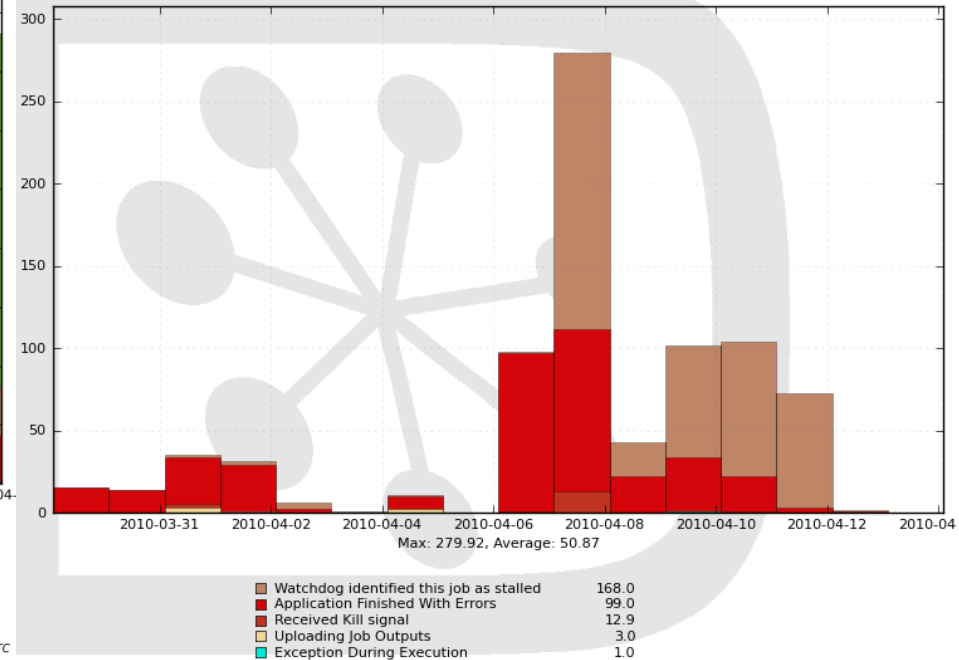
15 Days from 2010-03-29 to 2010-04-13



Generated on 2010-04-13 16:52:16 UTC

Failed jobs at RAL

15 Days from 2010-03-29 to 2010-04-13



Generated on 2010-04-13 16:41:24 UTC

➤ CERN

- ⊙ More – or – less fine and the mainstay of our computing
- ⊙ FTS channels were not configured after upgrade to FTS 2.2.3
 - ⇒ Fixed immediately
- ⊙ Some recent (possibly disk server related) problems

➤ CNAF

- ⊙ StoRM SE problems preventing output data being downloaded
- ⊙ Chose bad time to move backend hardware of LHCb ConDB

➤ GridKa, IN2P3

- ⊙ NFS problems with access to shared software area
 - ⇒ Significant problem at GridKa. Fixed with additional hardware

➤ NIKHEF / SARA

- ⊙ Strange interference between gsidcap library and ConDB
 - ⇒ Waiting for fix for this problem. Site banned till yesterday
 - ⇒ Now temporarily support dcap at SARA until fix is found

➤ RAL

- ◎ Mostly fine over the last two weeks
 - ⇒ Many thanks!
 - ⇒ Loss of connectivity twice on 4 April
- ◎ Thankfully no disk server problems so far
 - ⇒ In the last two weeks, that is
 - ⇒ There was some data loss over Christmas and New Year
 - * Mostly RAID5 disk servers which went bad
 - * Now we have replaced the RAID5 machines in d1 classes
- ◎ Stopped 1GB queue due to memory problems
 - ⇒ Request for move to limits on virtual memory rather than physical memory

➤ PIC

- ◎ Annual shutdown on 29, 30 April / CERN media day
 - ⇒ Created some issues for LHCb (see following)

- Tier-2s mostly okay within the UK
- Strange problem of data uploads at some sites
 - ⊙ Big problem in Glasgow. Banned within LHCb.
 - ⇒ ~50% of data uploads to CERN/Tier-1s fail
 - ⇒ Being investigated by Dug McNab
 - ★ 100% success for uploading data to DPM end points
 - ★ 50% failure rates for dCache, Castor and StoRM endpoints
 - ⇒ Possibly related to port changes during migration of worker nodes to sl5
 - ⇒ Not noticed before due to other problems with LHCb applications
 - ⊙ At a smaller scale at a couple of other sites – ignored for now
 - ⊙ Seen only at UK sites so far!

- DIRAC now on more resilient configuration
 - ⊙ More hardware at CERN
 - ⊙ Better configuration of software
 - ⇒ Improved separation of modules
 - ⇒ More tools for debugging problems
 - ★ Other experiments (LCD, EELA) also using DIRAC
 - ⊙ DIRAC servers able to handle load very easily so far
 - ⊙ Minor niggles which are being fixed fast ...
- Primary DIRAC web portal and configuration servers hosted at PIC
 - ⊙ Chaos on 30 March
 - ⊙ Alternatives being investigated
- More (daily) meetings of LHCb operations team with site representatives
 - ⊙ Identify and follow up problems faster for LHCb

- LHCb physics programme off to running start
 - ⊙ On the grid
 - ⊙ Users supporting each other now
 - ⇒ Faster than many experts in responding to queries
- Many challenges being seen
 - ⊙ Mostly internal to LHCb
 - ⊙ CORAL software a big problem
- RAL and UK Tier-2 sites going on fine (mostly)