LHCb CCRC'08 report
LHCC, July 08

Philippe Charpentier
Slides by courtesy of Nick Brook
(LHCb CCRC'08 co-ordinator)

Planned tasks

- May activities
  - Maintain equivalent of 1 month data taking
  - Assuming a 50% machine cycle efficiency
  - Run fake analysis activity in parallel to production type activities
    - Analysis type jobs were used for debugging throughout the period
    - GANGA testing ran for last weeks at low level

Activities across the sites

- Planned breakdown of processing activities (CPU needs) prior to CCRC08

Site	Fraction (%)
CERN	14
FZK	11
IN2P3	25
CNAF	9
NIKHEF/SARA	26
PIC	4
RAL	11

Pit -> Tier 0

- Use of rfcp to copy data from pit to CASTOR
  - rfcp is the recommended approach from IT
  - A file sent every ~30 sec
  - Data remains on online disk until CASTOR migration
  - Rate to CASTOR - ~70MB/s

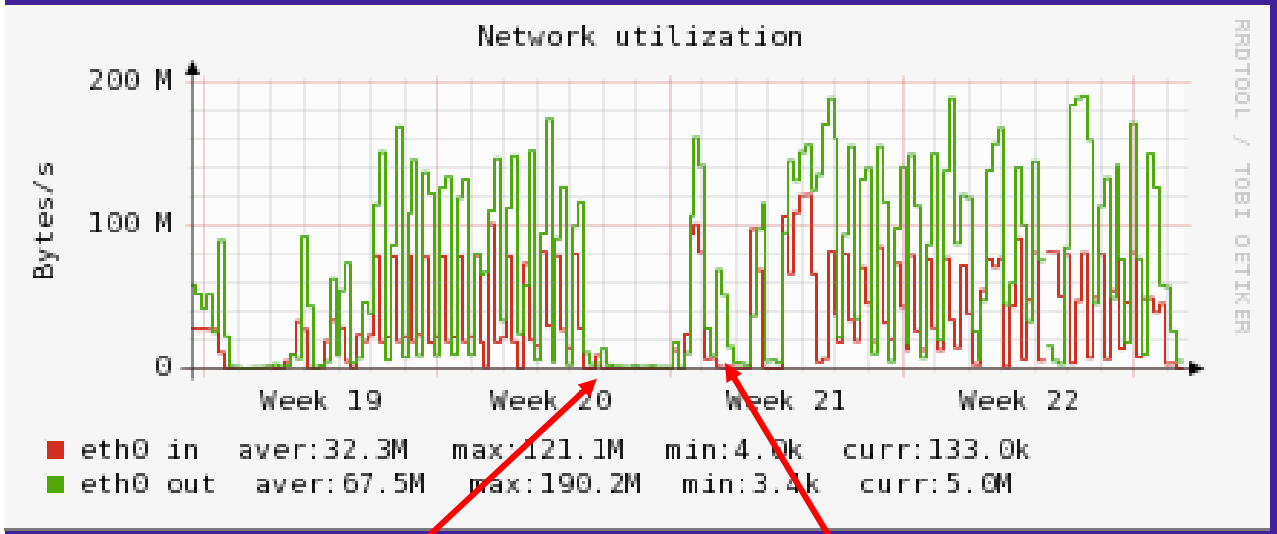

Problems with online storage area

In general ran smoothly:

- Stability problems with online storage area - solved with firmware update during CCRC

- Internal issues with sending bk-keeping info

Tier 0 -> Tier 1

- FTS from CERN to Tier-1 centres
  - Transfer of RAW will only occur once data has migrated to tape & checksum is verified
  - Rate out of CERN - ~35MB/s averaged over the period
  - Peak rate far in excess of requirement
    - In smooth running sites matched LHCb requirements

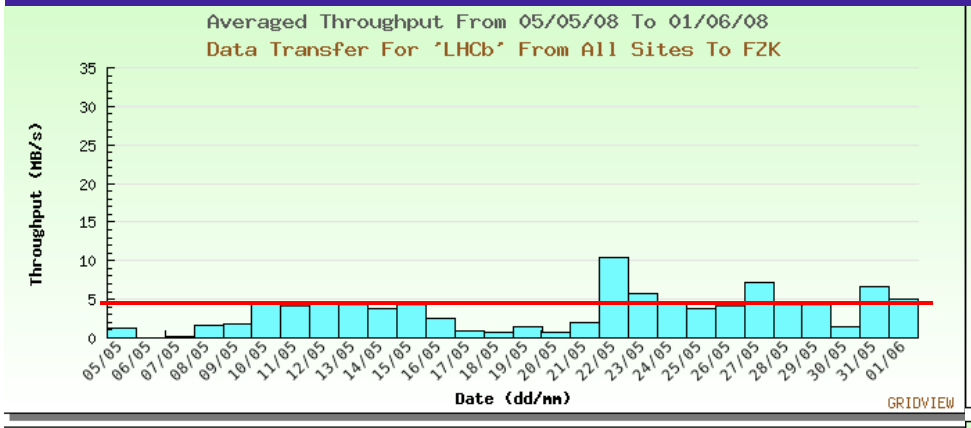
Averaged Throughput From 05/05/08 To 01/06/08
Data Transfer For 'LHCb' From All Sites To FZK
Throughput (MB/s)
Date (dd/mm)
GRIDVIEW

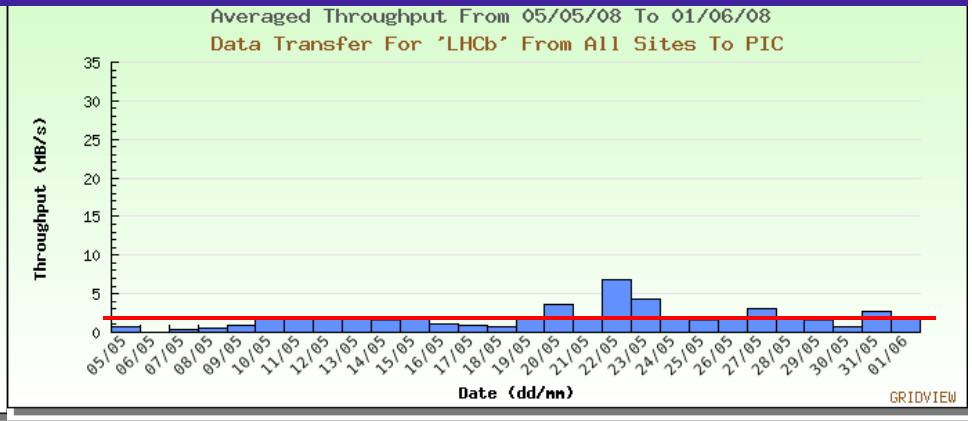
Averaged Throughput From 05/05/08 To 01/06/08
Data Transfer For 'LHCb' From All Sites To PIC
Throughput (MB/s)
Date (dd/mm)
GRIDVIEW

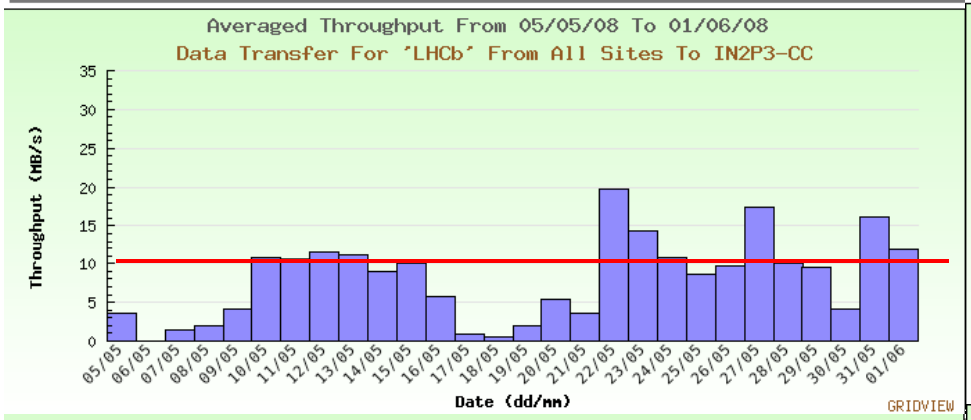
Averaged Throughput From 05/05/08 To 01/06/08
Data Transfer For 'LHCb' From All Sites To IN2P3-CC
Throughput (MB/s)
Date (dd/mm)
GRIDVIEW

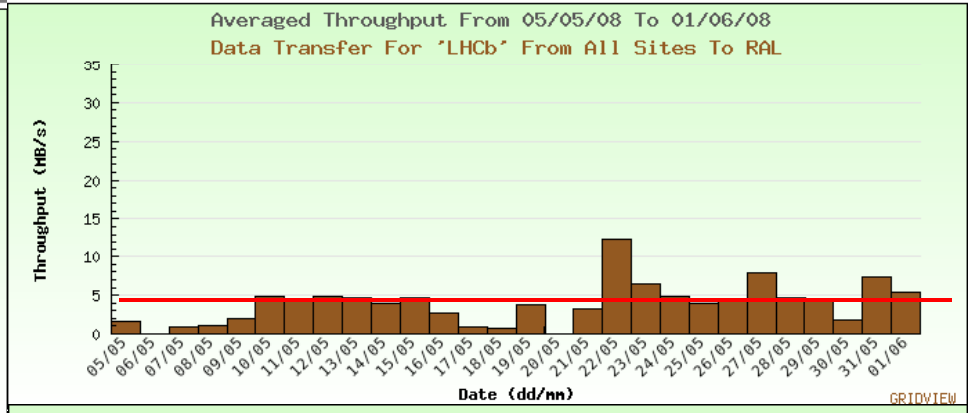
Averaged Throughput From 05/05/08 To 01/06/08
Data Transfer For 'LHCb' From All Sites To RAL
Throughput (MB/s)
Date (dd/mm)
GRIDVIEW

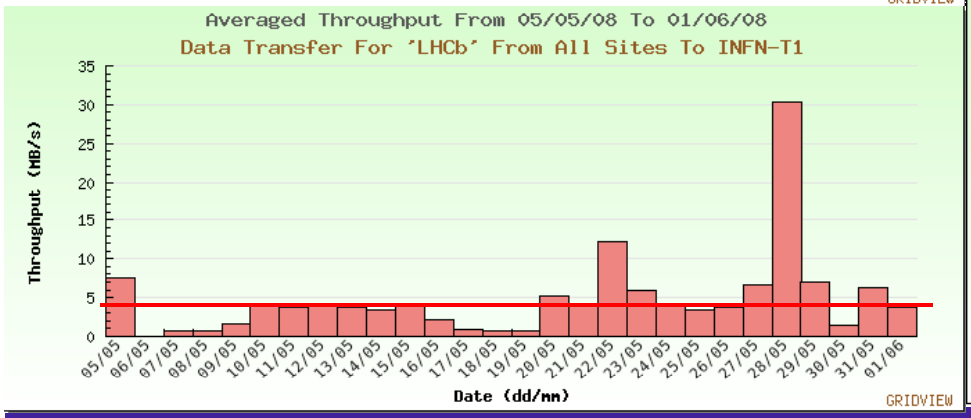
Averaged Throughput From 05/05/08 To 01/06/08
Data Transfer For 'LHCb' From All Sites To INFN-T1
Throughput (MB/s)
Date (dd/mm)
GRIDVIEW

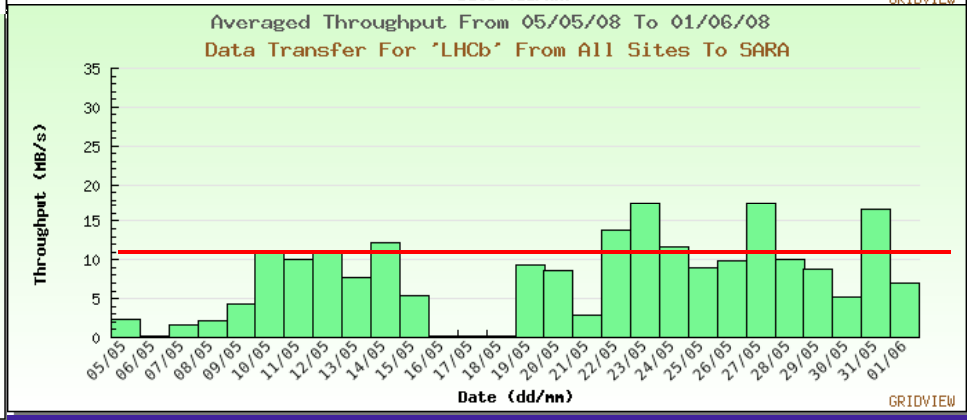
Averaged Throughput From 05/05/08 To 01/06/08
Data Transfer For 'LHCb' From All Sites To SARA
Throughput (MB/s)
Date (dd/mm)
GRIDVIEW

Tier 0 -> Tier 1

- To first order all transfers eventually succeeded
  - Plot shows efficiency on 1st attempt…

Issue with UK certificates

Restart IN2P3 SRM endpoint

CERN outage

CERN SRM endpoint problems

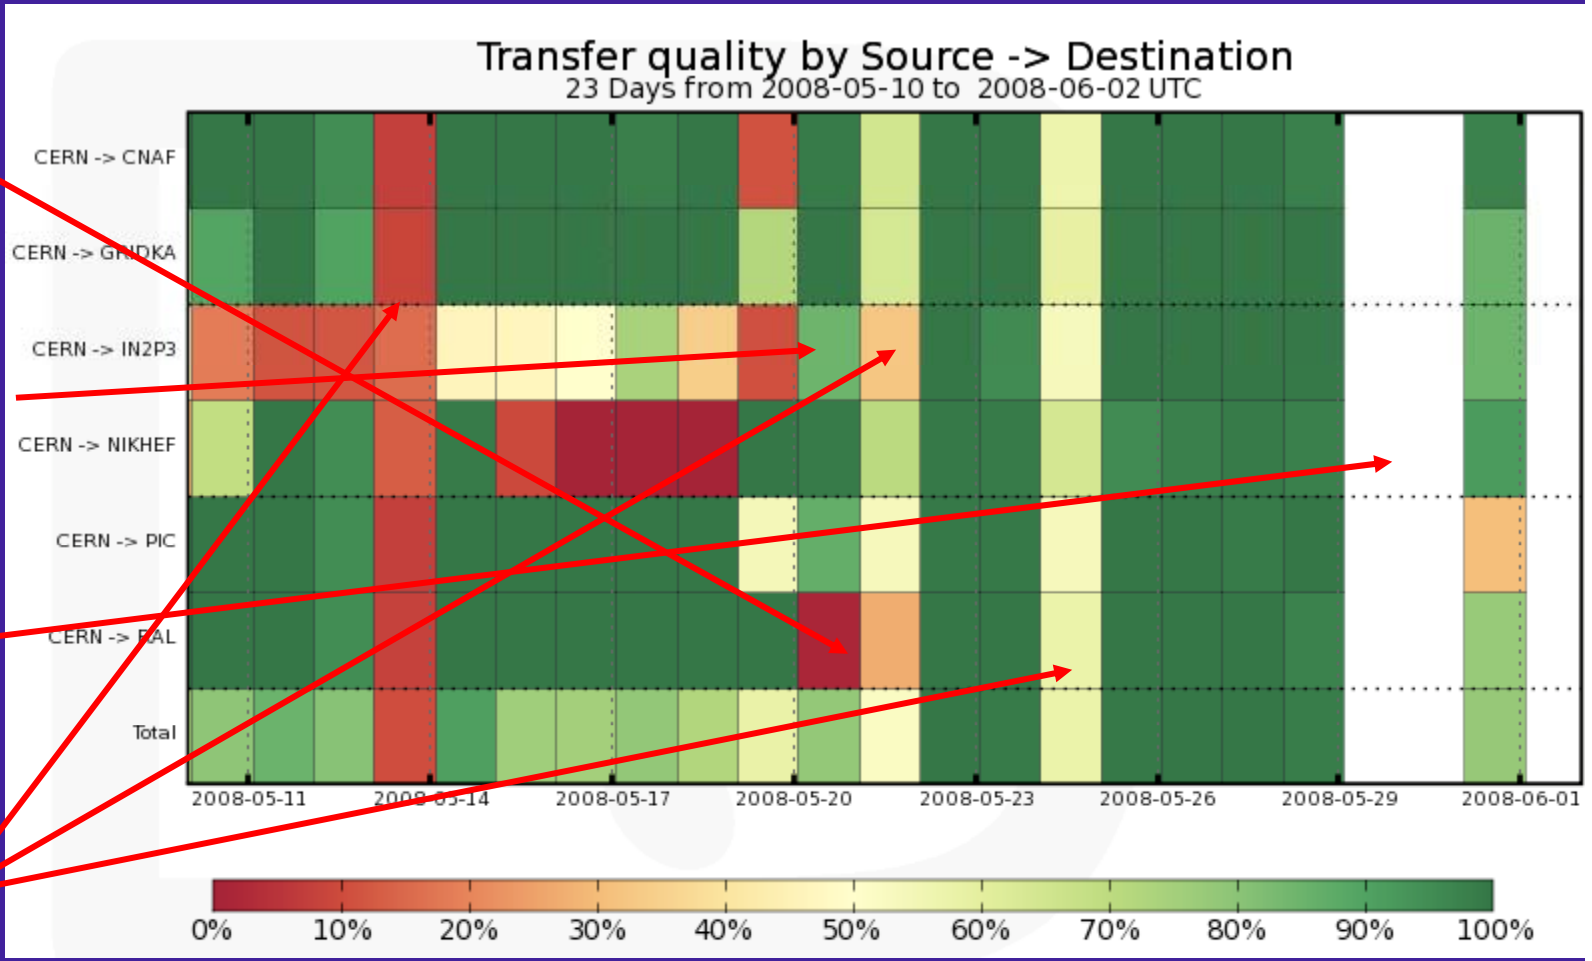

Reconstruction

- Used SRM 2.2 SE
  - LHCb space tokens are:
    - LHCb_RAW (T1D0)
    - LHCb_RDST (T1D0)
- Data shares need to be preserved
  - Important for resource planning
- Input 1 RAW file & output 1 rDST file (1.6 GB)
- Reduced nos of events per recons job from 50k to 25k (job ~12 hour duration on 2.8 kSI2k machine)
  - In order to fit within the available queues
  - Need to get queues at all sites that match our processing time
    - Alternative: reduce file size!

Reconstruction

- After data transfer file should be online, as job submitted immediately, but…
- LHCb pre-stage files & then checks on the status of the file before submitting pilot job - use gfal_ls
  - Pre-stage should ensure access availability from cache
  - Only issue at NL-T1 with reporting of file status
    - Discussed last week during Storage session (dCache version)
  - (Problem developed at IN2P3 right at end of CCRC08 - 31st May)

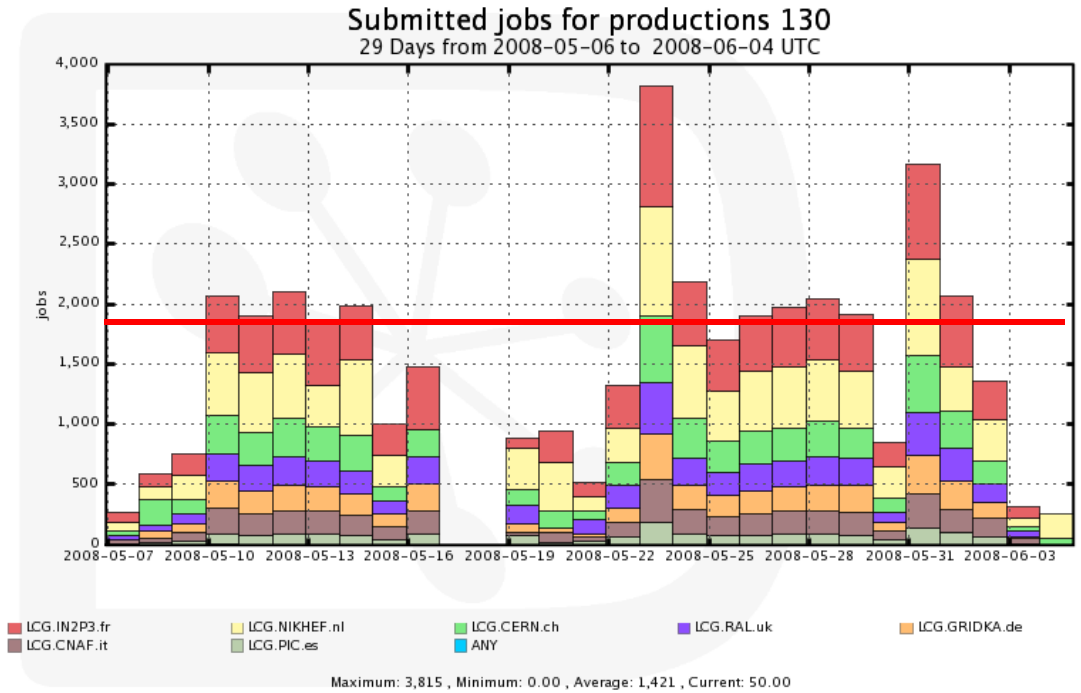

Reconstruction

41.2k reconstruction jobs submitted

27.6k jobs proceeded to done state

Done/created ~67%

CERN	6.1k (14%)	5.3k (13%)	86%
CNAF	3.9k (9%)	2.8k (7%)	72%
GridKa	4.1k (11%)	3.1k (7%)	76%
IN2P3	10.3k (25%)	6.1k (14%)	56%

	Sub jobs	Done jobs	Ratio
NIKHEF	10.3k (26%)	2.3k (6%)	23%
PIC	1.8k (4%)	1.6k (4%)	89%
RAL	4.7k (11%)	3.5k (8%)	74%

dCache Observations

Official LCG recommendation - 1.8.0-15p3

LHCb ran smoothly at half of T1 dCache sites

- PIC OK - version 1.8.0-12p6 (unsecure)
- GridKa OK - version 1.8.0-15p2 (unsecure)
- IN2P3 - problematic - version 1.8.0-12p6 (secure)
  - Seg faults - needed to ship version of GFAL to run
  - Could explain CGSI-gSOAP problem????
- NL-T1 - problematic (secure)
  - Many versions during CCRC to solve number of issues
  - 1.8.0-14 -> 1.8.0-15p3->1.8.0-15p4
  - "Failure to put data - empty file"->"missing space token" problem -> incorrect metadata returned, NEARLINE issue

Stripping

- Stripping on rDST files
  - 1 rDST files & associated RAW file
    - Space tokens: LHC_RAW & LHCb_rDST
  - DST files & ETC produced during the process stored locally on T1D1 (add storage class)
    - Space tokens: LHCb_M-DST
  - DST & ETC file then distributed to all other computing centres on T0D1 (except CERN T1D1)
    - Space tokens: LHCb_DST (LHCb_M-DST)

Stripping

CERN	2.4k	2.3k
CNAF	2.3k	2.0k
GridKa	2.0k	2.0k
IN2P3	4.5k	0.2k
NIKHEF	0.3k	<0.1k
PIC	1.1k	1.1k
RAL	2.2k	1.6k
Failed to resolve datasets	17.0k	-

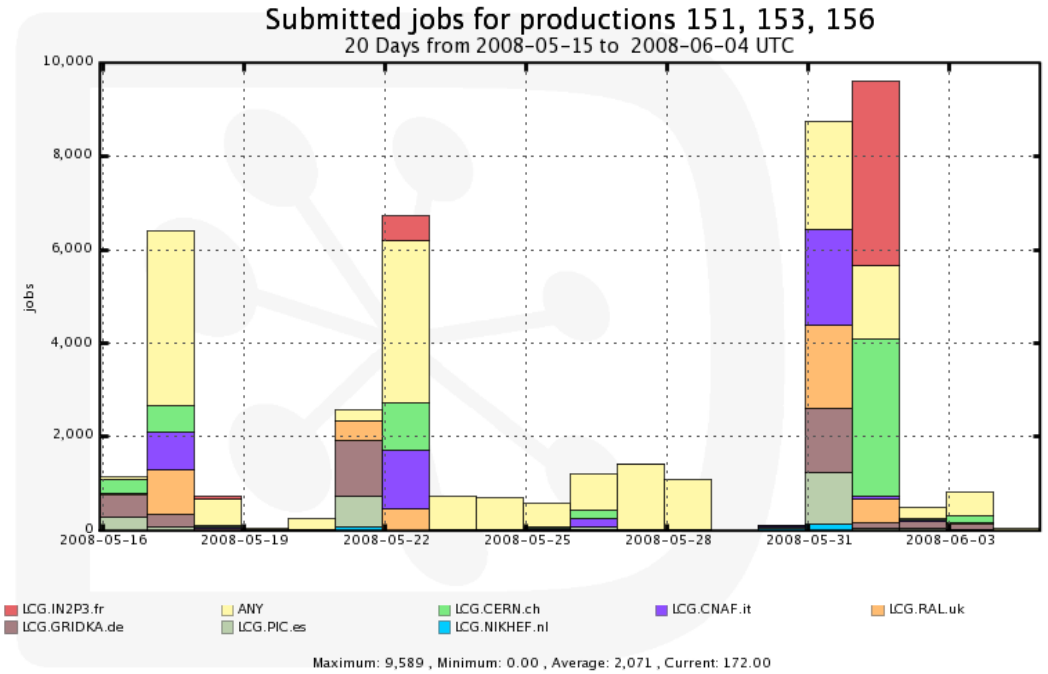

- 31.8k stripping jobs were submitted
- 9.3k jobs ran to "Done"
- Major issues with LHCb bk-keeping

Lesson Learnt for DIRAC3

- Improved error reporting in workflow & pilot logs
  - Careful checking of log files was required for detailed analysis
- Full failover mechanism is in place but not yet deployed
  - only CERN was used for CCRC08
- Alternative forms of data access
  - Minor tuning of the timeout for downloading input data was required
    - 2 timeouts needed: time of copy & activity timeout

Summary

- Data transfer of CCRC08 using FTS was successful
- Still plagued with many issues associated data access
  - Issues improved since Feb CCRC08 but…
  - 2 sites problematic for large chunks of CCRC08 - 50% of LHCb resources!!
  - Problems mainly associated with access with dCache
  - Commencing tests with xrootd
- DIRAC3 tools improved significantly from Feb
  - Still need improved reporting of problems
- LHCb bk-keeping remains a major concern
  - New version due prior to data taking
- LHCb need to implement a better interrogation of log files