

GS

CERN  
IT  
Department



LHCb CCRC'08  
February

*Roberto Santinelli IT/GS/EIS (CERN)*



*Presentation kindly provided by S. K. Paterson (CERN)*



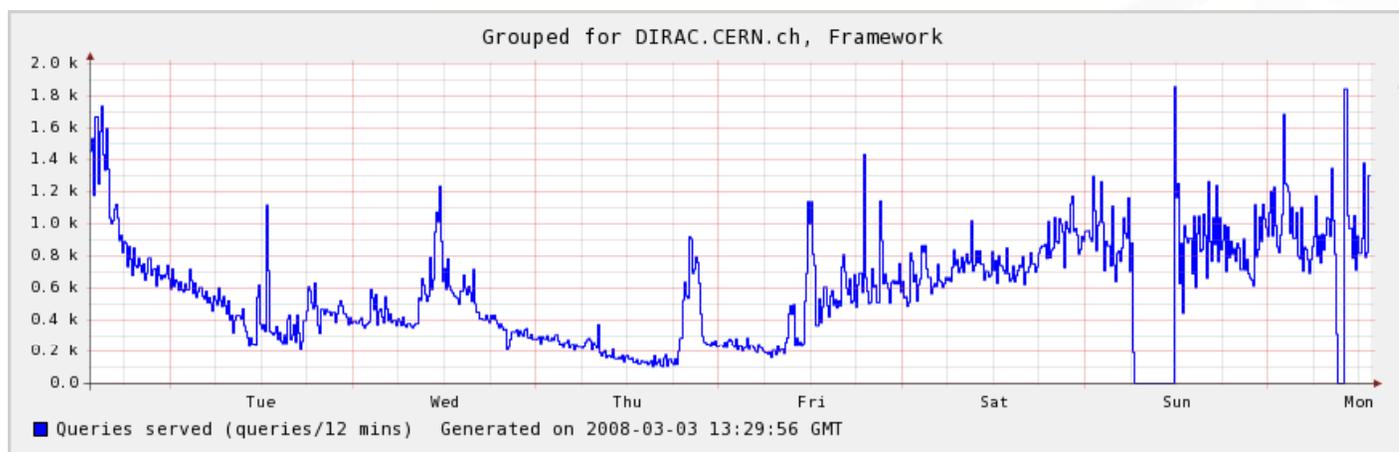
- Test of the LHCb Computing Model
  - DIRAC ported to new framework
    - Reengineered to use SRM2 via gfal python API and gLite WMS
  - CCRC'08 exercise being used to test the system
- RAW Data Upload, Online -> Tier-0 storage (CERN Castor)
  - Use DIRAC transfer framework (Castor rfc, GridFTP)
- RAW Data Distribution to Tier-1s
  - CNAF, GridKa, IN2P3, NIKHEF, PIC, RAL
- Data Reconstruction at Tier-0 + Tier-1
- Stripping of Reconstructed Data
  - Activity 'de-scoped', tests now envisaged for March
  - Will be in place for CCRC phase 2



- DIRAC ported to new framework (DIRAC3)
  - The whole software stack is completely reengineered
- New features:
  - Use of SRM2 via gfal python API and gLite WMS
  - GSI enabled web portal exposing job, production and data management functionality
  - Job peeking and killing for all jobs via 'heartbeat' mechanism
  - Job Pilot proxy delegation
  - Time Left Utility and filling mode
  - ....many others
- February CCRC'08 allowed testing of the following components
  - Data Management
  - Production Management
  - Workload Management
- Each will be elaborated upon during the course of this talk



- Some features still pending
  - Generic 'failover' mechanism for all operations
    - E.g. data upload, status transitions
  - Stager service
    - Will be put in place before stripping
- Occasional problems due to DIRAC service instability
  - Lack of mirror instances for critical services, e.g. Configuration Service
- Unexpected downtime due to (now resolved) issues in the DIRAC3 framework
  - Below plot is for the job state service
    - Downtime of this service results in 'stalled' jobs since no status transitions are sent back to the WMS





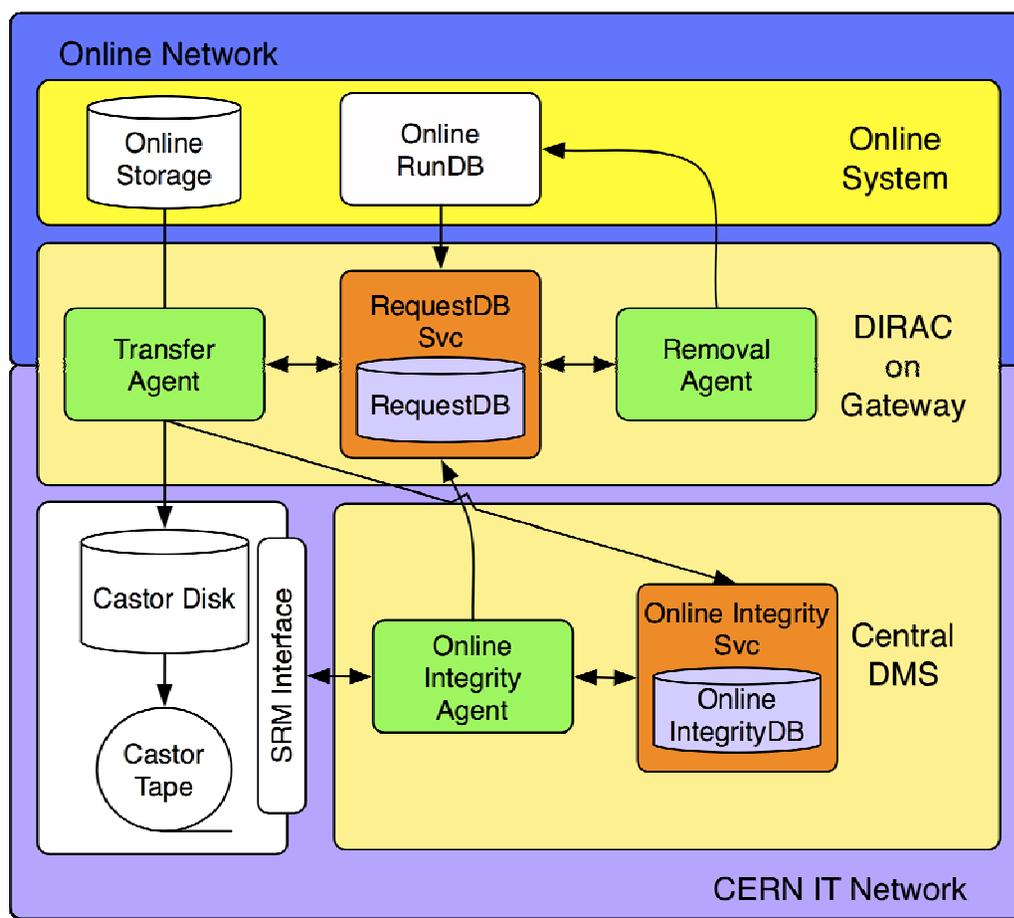
- RAW Data Upload
  - Ensure file integrity in Castor prior to removal from Online storage
- Data Driven RAW Distribution
  - After integrity verified, files become candidates for replication
  - Share amongst Tier-1s according to pledged resources
  - Replicated via gLite FTS
- Data Reconstruction at Tier-0 + Tier-1 sites
  - Files become candidates for processing after:
    - Integrity verified at Tier-0
    - Replication to a Tier-1
  - Jobs are created as soon as data is available
    - Jobs then automatically submitted to the DIRAC WMS
  - Reconstruction produces RDST file
    - Stored at the Tier-1 local SE (LHCb\_RDST)
    - Attempt to store on failover SE in case of failure (LHCb\_FAILOVER)



- SRM2 is a critical service for LHCb
  - File replication using gLite FTS
  - TURLs retrieved via gfal for access via available site protocols
    - root, rfiio, dcap, gsidcap
  - RDST output file upload to local Tier-1 SE via lcg-utils / gfal
  - File removal using gfal
- Tier-0,1 Storage Elements providing SRM2 spaces:
  - LHCb\_RAW (T1D0)
  - LHCb\_RDST (T1D0)
  - LHCb\_M-DST (T1D1)
  - LHCb\_DST (T0D1)
  - LHCb\_FAILOVER (T0D1)
    - Used for temporary upload in case of destination unavailability
- Thorough testing of production deployment of SRM2 is a key part of CCRC'08 for LHCb



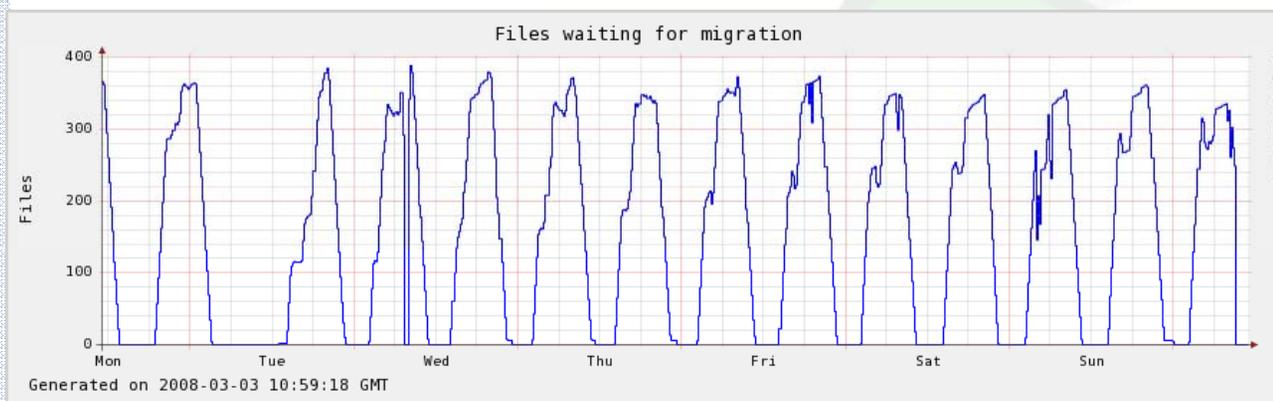
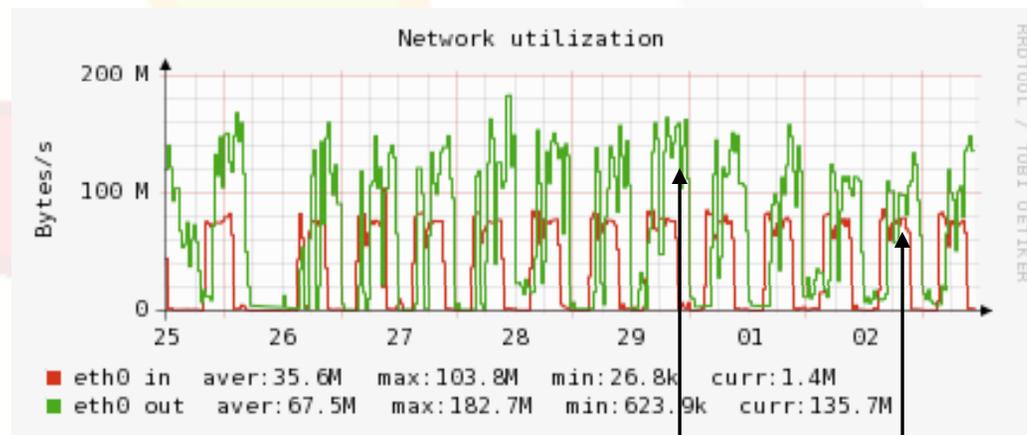
- DIRAC installed on Gateway
  - Replicate RAW data to Castor pools
  - Register data to Online Integrity DB



- Persist files Online till 'safe'
- File checksum calculated online at write time
- File checksum calculated by Castor on migration to tape
- Online Integrity Agent interrogates Castor for file checksum
- If 'safe': place removal request in DIRAC RequestDB at Gateway
- Pass removal request to Online system



- Online (Pit 8) to Tier-0 Transfers
  - Early February: continuous with low rate
  - As of February 18th proceeded at nominal rate (70MB/s) with ~50% duty cycle
    - Few pauses for software upgrades
  - Peak of 125MB/s (Feb 12th)
  - 0 checksum mismatches

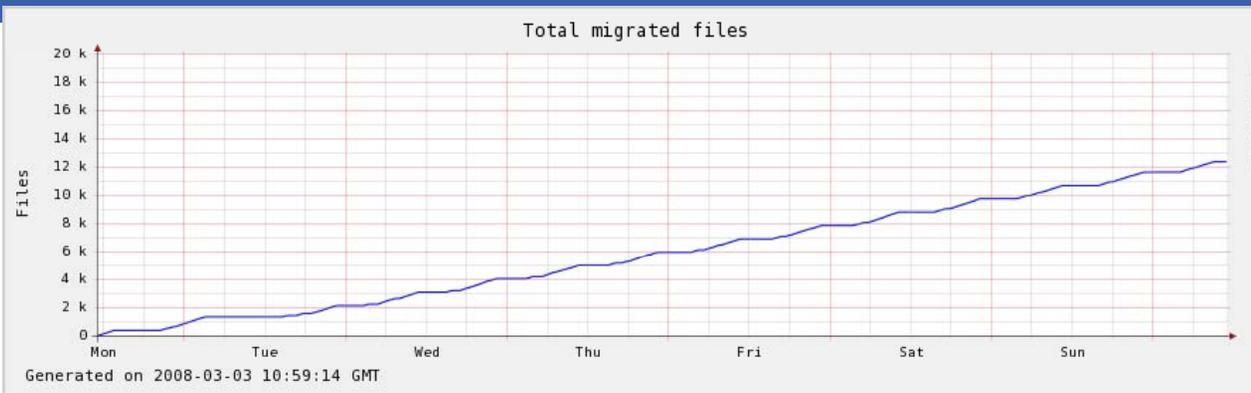


Migration

Tier-1  
transfers

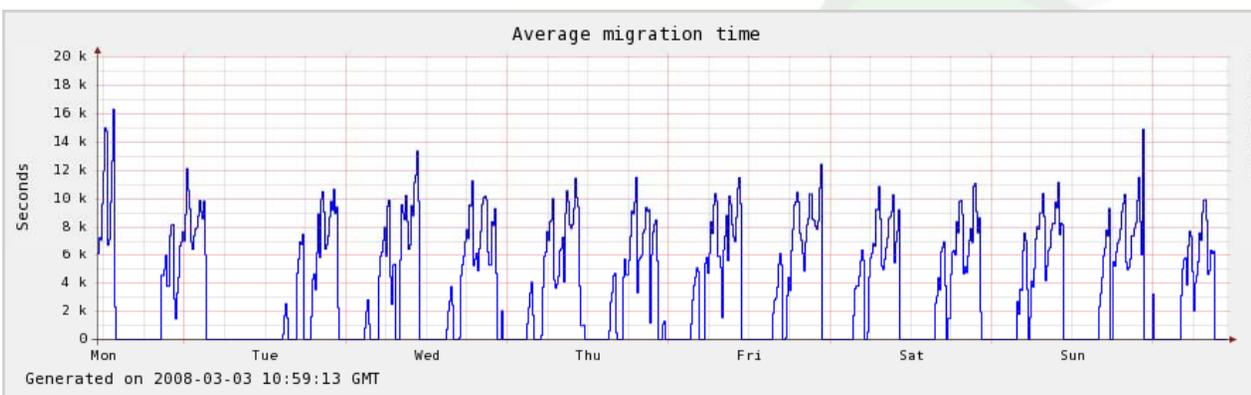
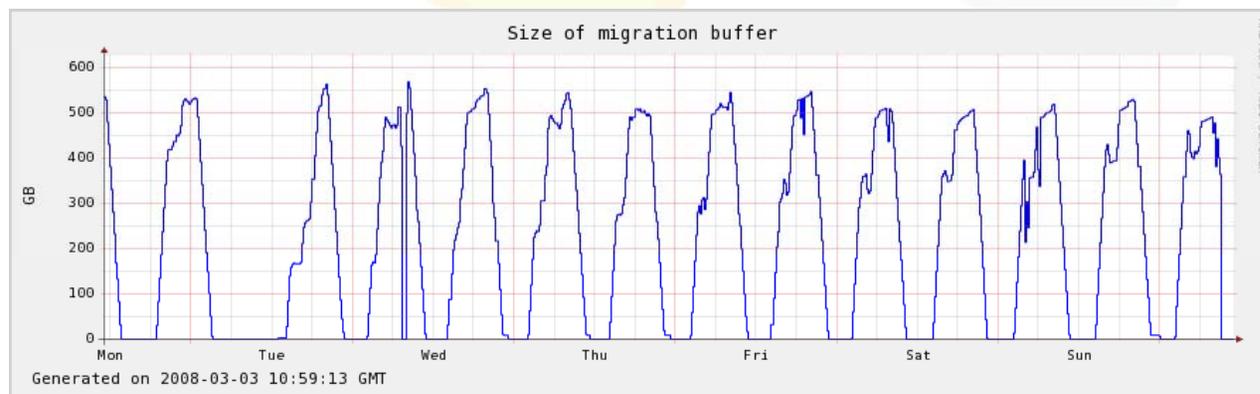


# Castor Migration (25th Feb -> 2nd March)



- Total migrated files

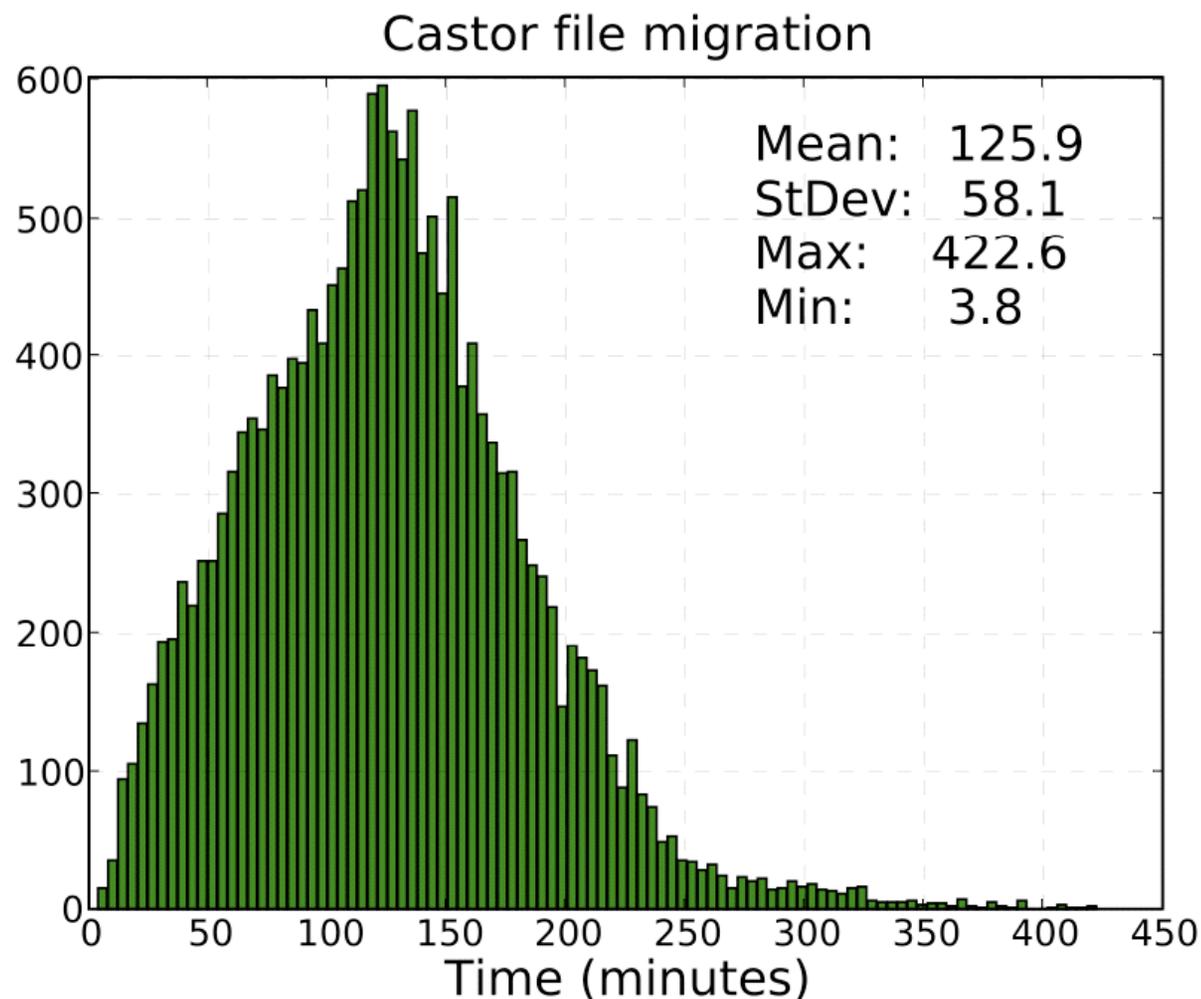
- Size of migration buffer



- Average migration time

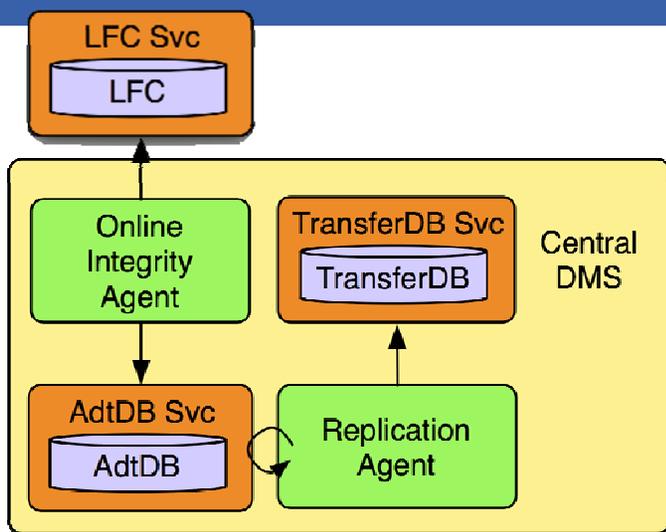


- Histogram of the Castor migration time in minutes for entire sample of LHCb CCRC'08 files



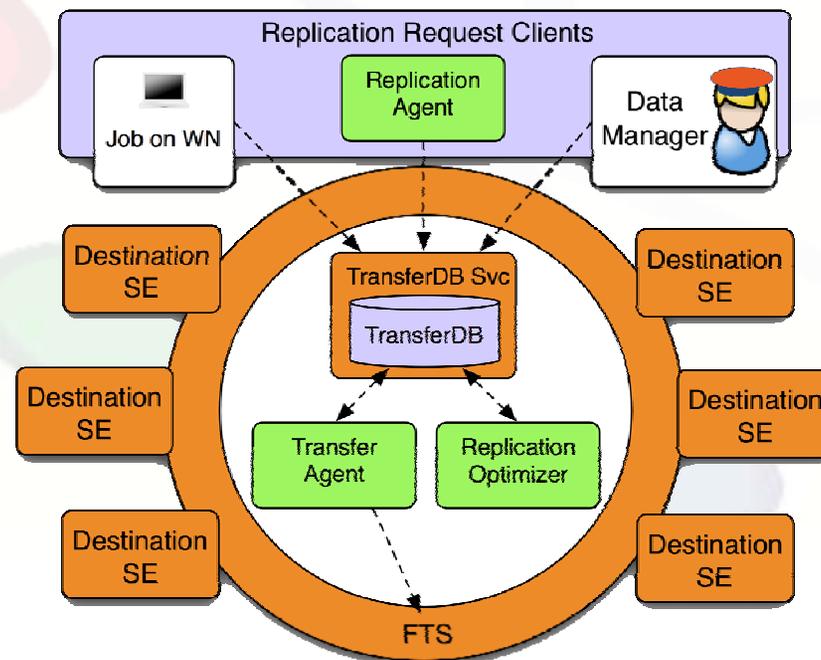


# RAW Data Distribution to Tier-1s



- Transfer Agent polls TransferDB
  - Creates bulk transfer requests
  - Submits and monitors transfers through FTS CLI
  - Requests retried in case of failure

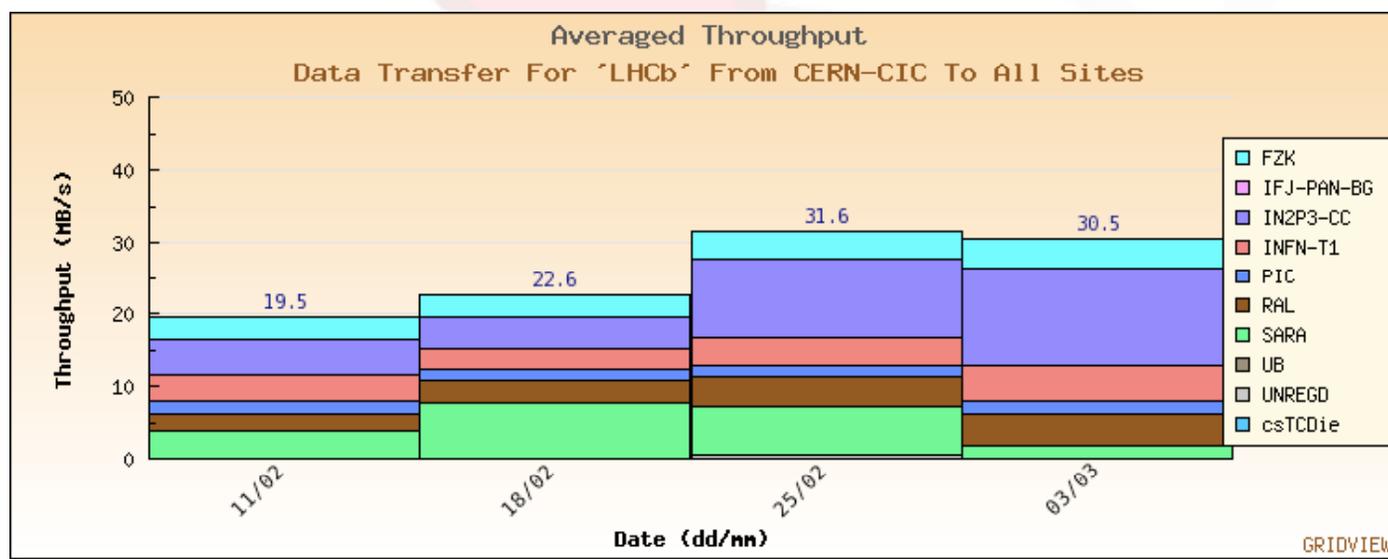
- File registered in AutoDataTransferDB when safely migrated
- Replication Agent splits files according to site shares
- Places transfer requests in TransferDB





# RAW Data Distribution to Tier-1s

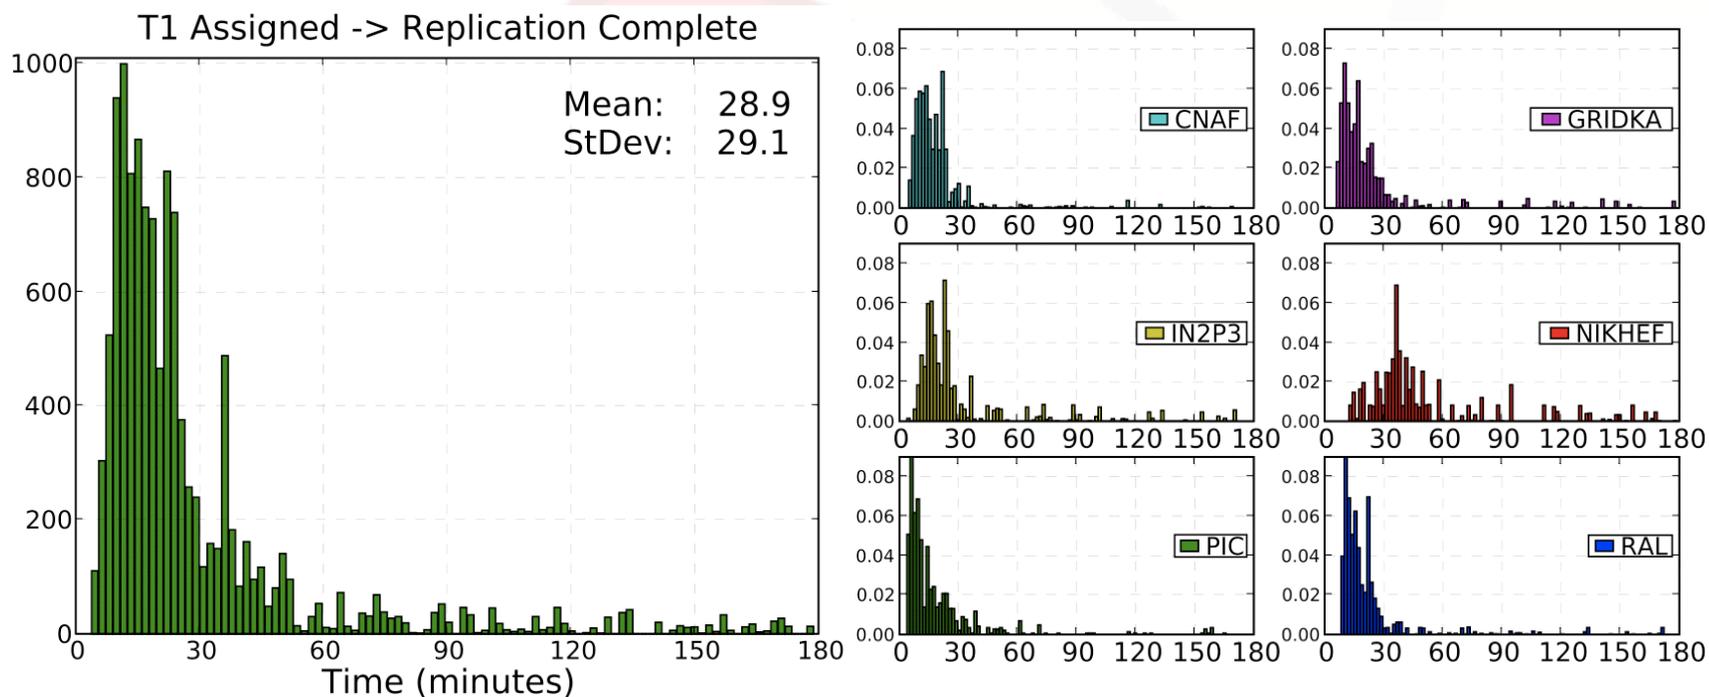
- Nominal transfer rate to the Tier-1s (70MB/s with 50% duty cycle) achieved after ramp-up on 18th February
  - Data replication according to pledged resources successful
- File removal with SRM2 works (space not recovered for dCache, see later)





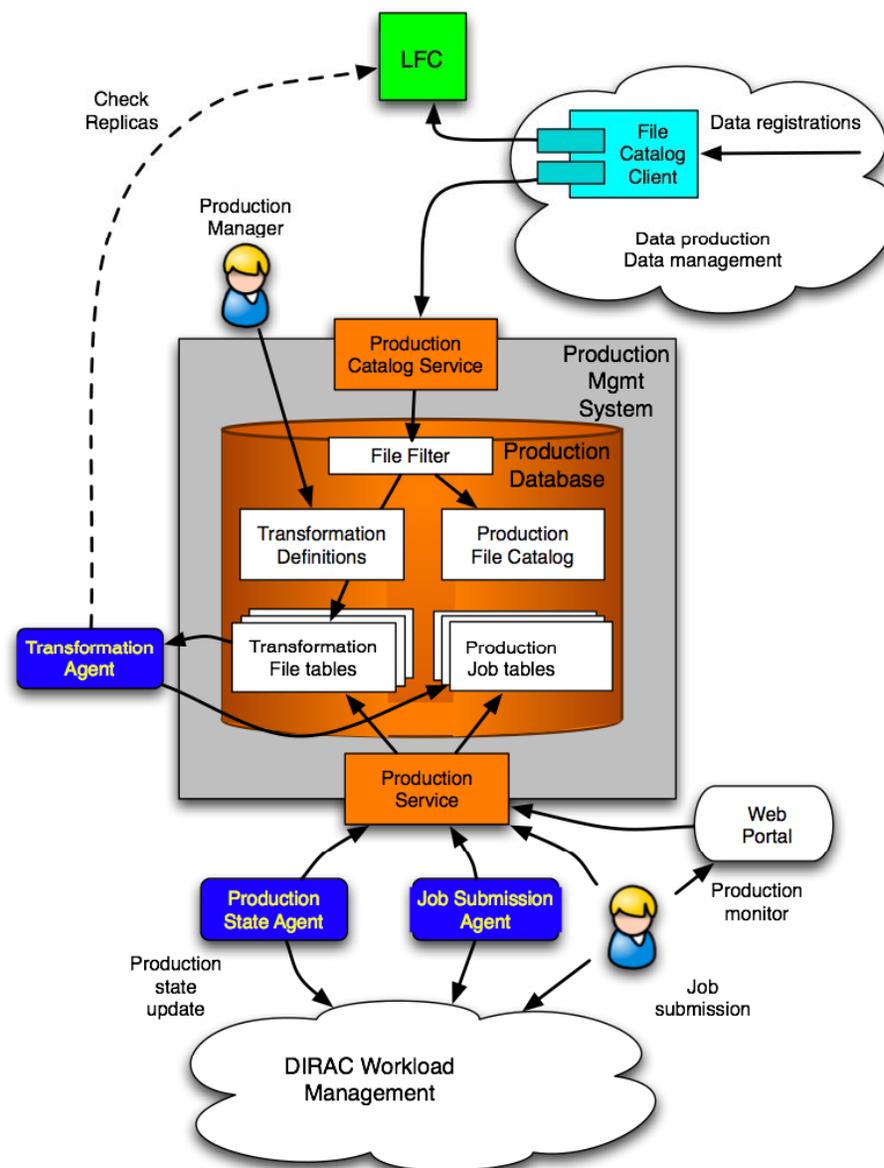
# Transfers Performance

- Histograms of time between a file being Assigned and Transferred to the LHCb Tier-1s (minutes)
  - FTS submit / monitor / done cycle
- Sites show stable behavior



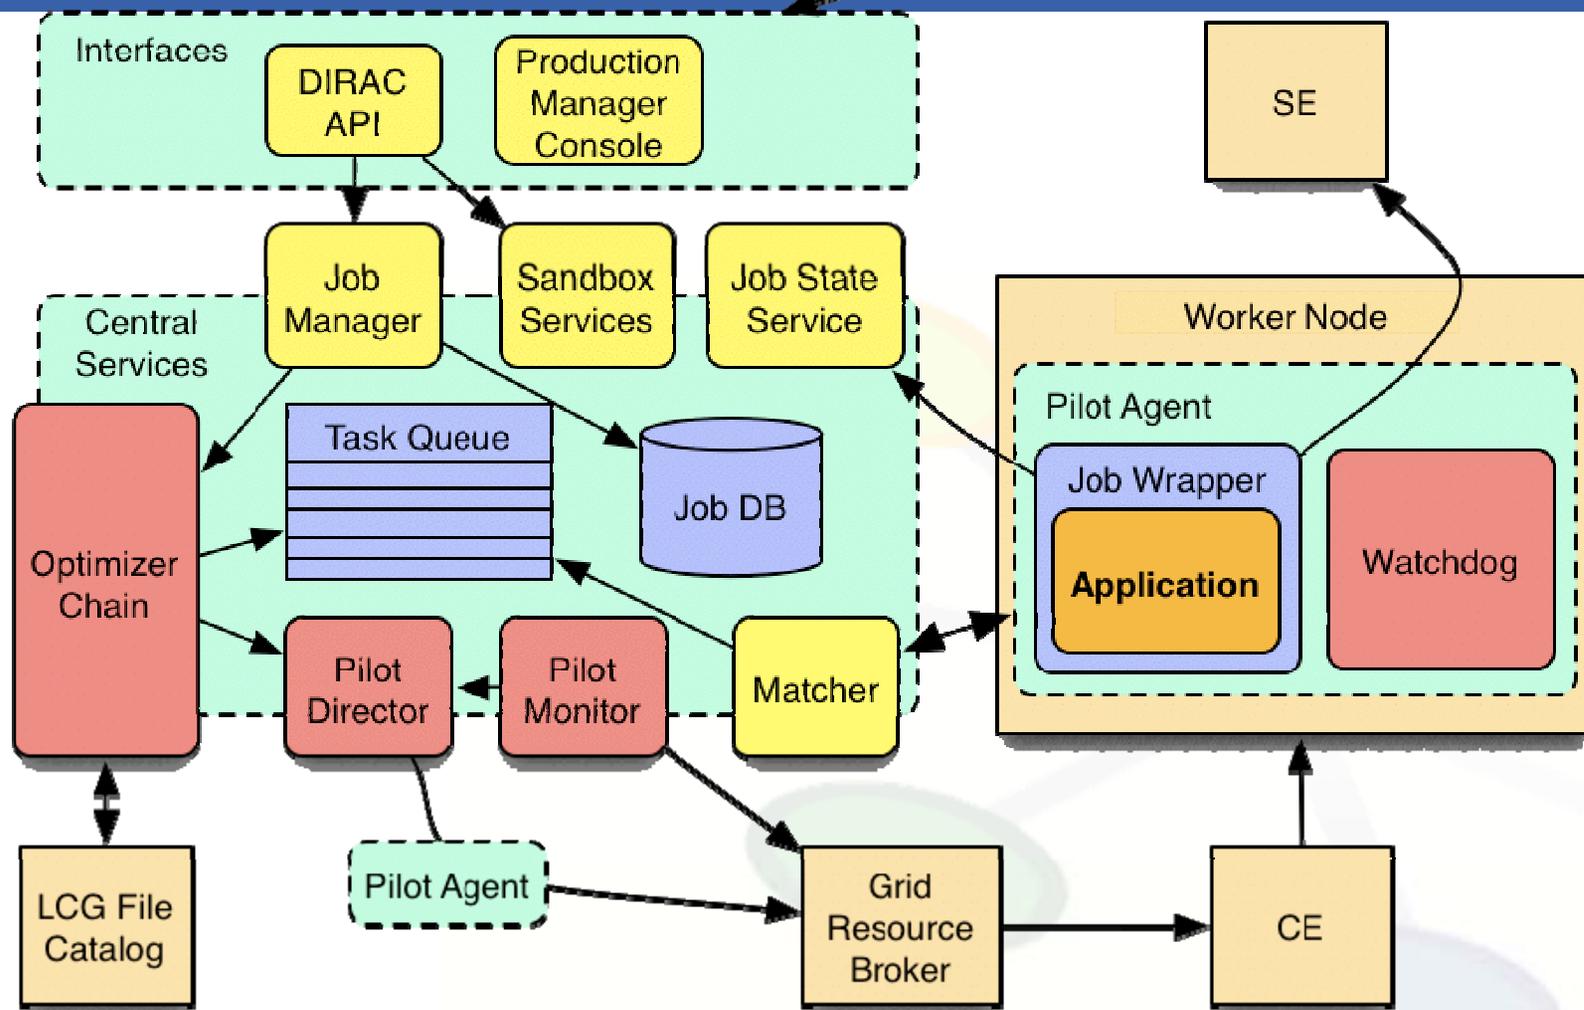


- Eligible files are entered in the Production Catalogue
  - Sorted according to the transformation definitions
  - Transformation agent creates the production jobs
- Tier-1 shares are allocated according to the quotas
  - Could be more adaptive algorithms in the future
- The Production Manager controls the job submission
  - Production service (CLI)
  - Web portal
  - Automatic submission can be enabled





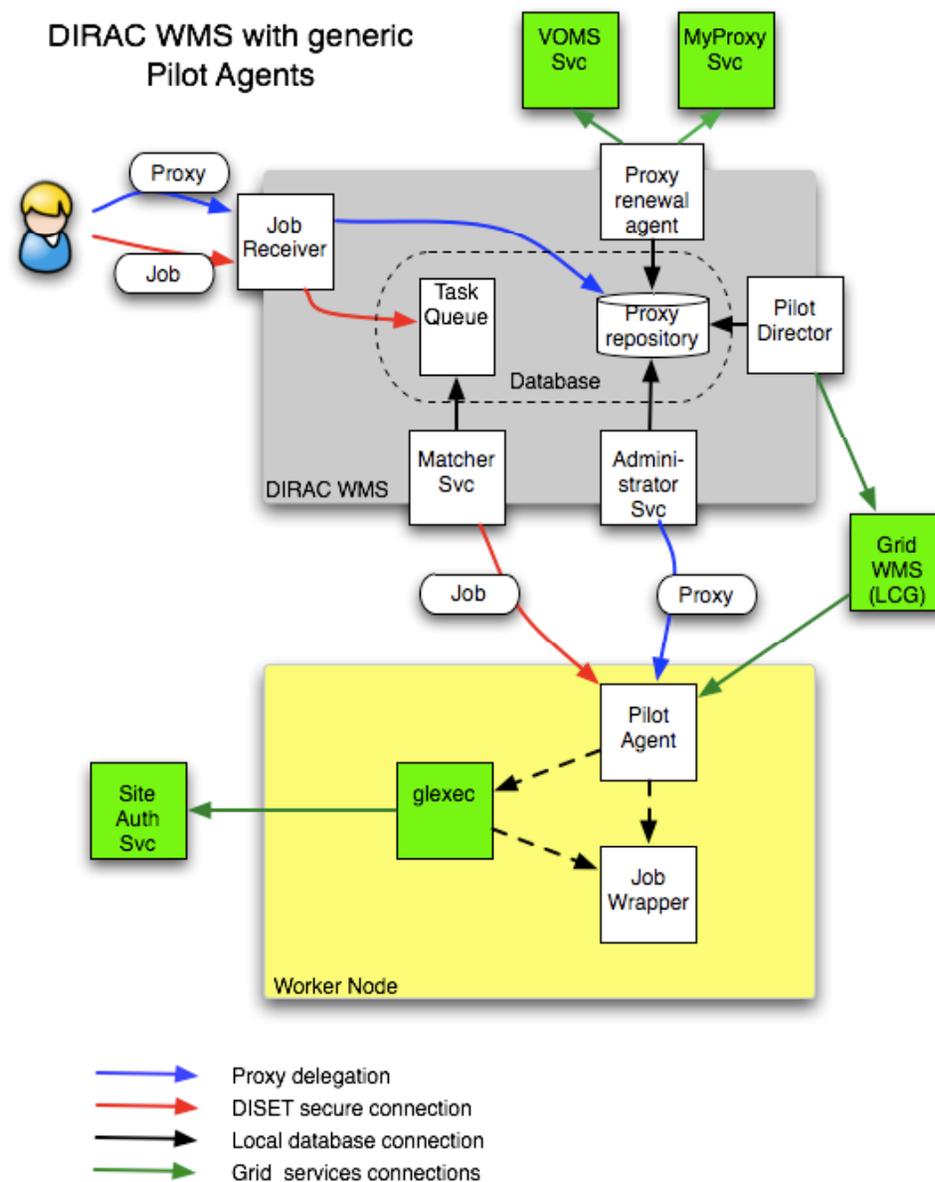
# WMS Overview





# Generic Pilots - Proxy Delegation

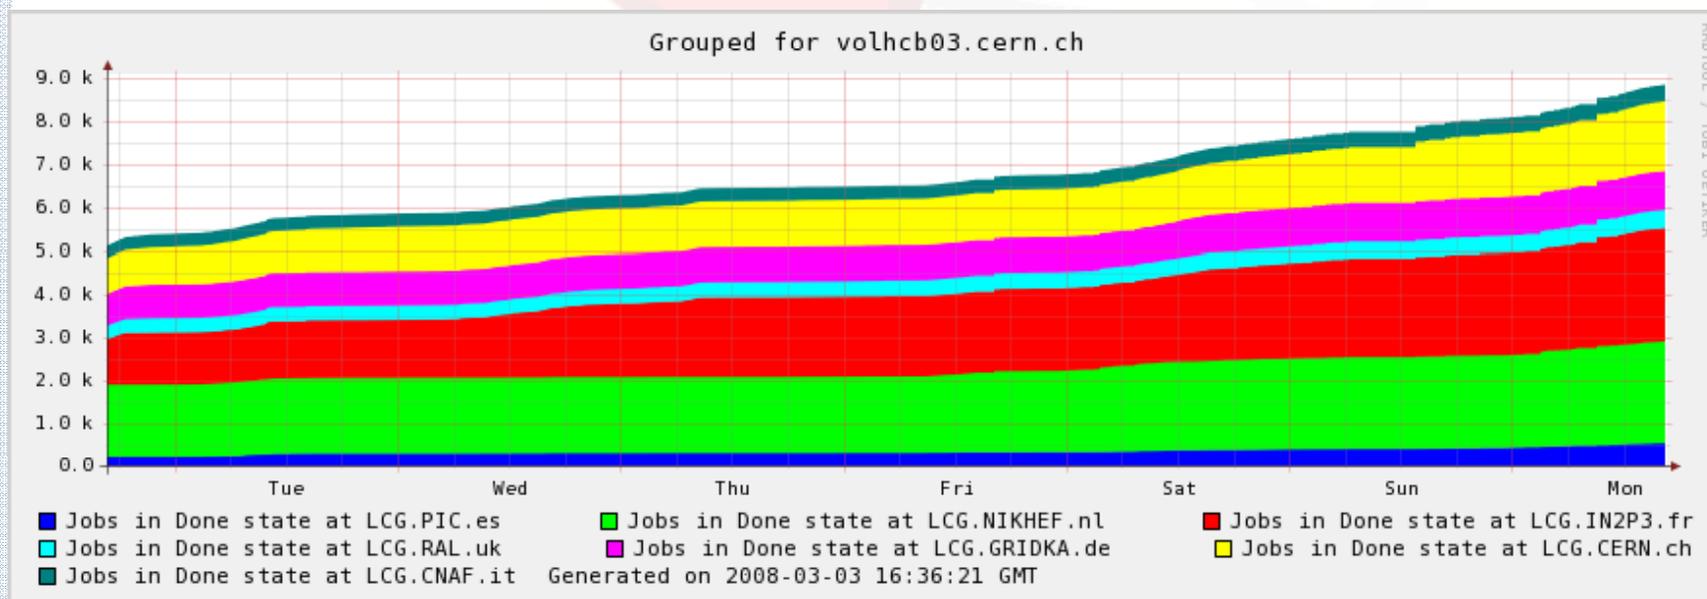
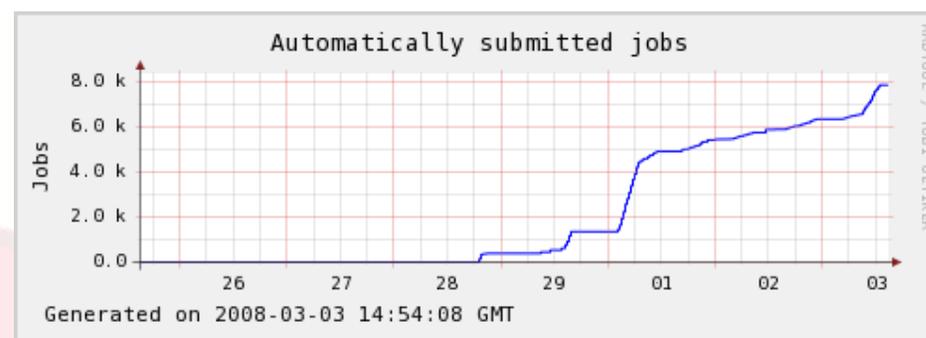
- All VO jobs submitted with generic credential
  - VOMS 'Pilot' role
    - Tested before Feb CCRC'08
  - February CCRC'08
    - Jobs submitted with 'Production' role
    - Identical 'Production' Role proxy delegated from the DIRAC WMS





# Data Reconstruction at Tier-1s (+Tier-0)

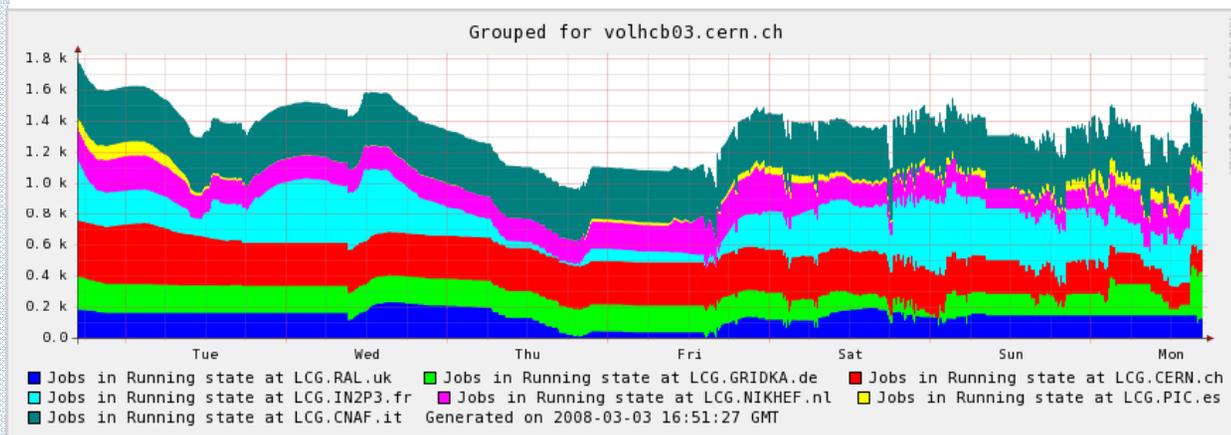
- Jobs submitted steadily and running at all Tier-1 sites
  - Mechanism for automatic job submission to DIRAC successfully demonstrated
- Looking at completed jobs during 25th Feb to 3rd March
  - Had planned to run ~23K jobs over last 2 weeks...





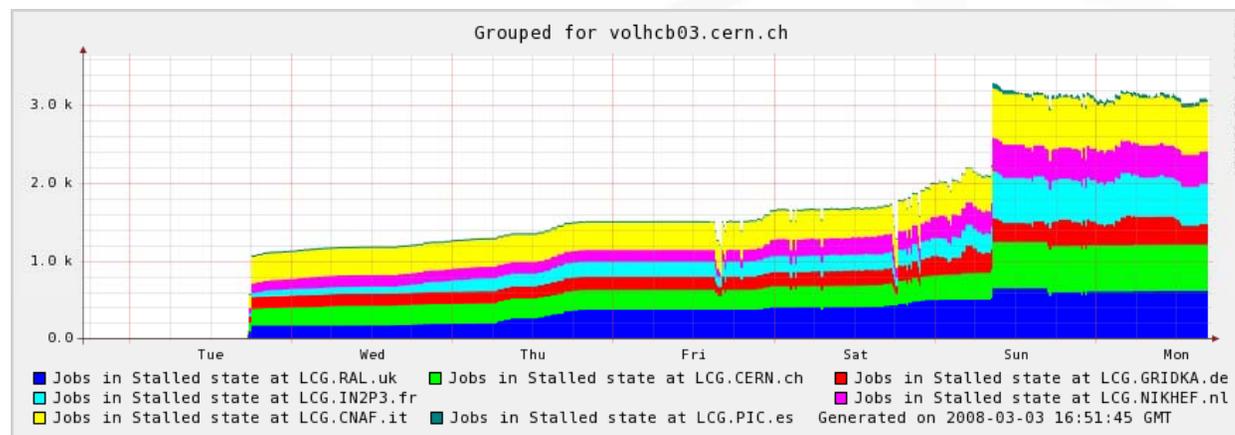
# Stalled Jobs and CPU Time Limit

- Tracking stalled jobs highlighted that the CPU requirement for our CCRC'08 jobs was too low



- CPU limit set as high as possible
  - Not all Tier-1s had sufficiently long queues

- Stalled jobs started to be identified on 26th February
  - Systematic across all Tier-1s





- Timeout set on the gsidcap ports
  - Configuration issue now resolved at affected Tier-1's (Lyon)
- Failed DESTINATION error during PREPARATION phase: [FILE EXISTS] (See Flavia's talk)
- Sites restarting gsidcap server (without downtime)
  - All open data connections dropped, running jobs fail
- dCache not releasing space reserved (even if files deleted)
  - Files in the space accounted, even on T1D0
  - Stops accepting files if maximum size is reached
    - Temporary solution is to declare a large (virtual) space size



- CERN
  - Flickering AFS behaviour caused jobs to fail accessing shared area
  - Negative file size (fixed on the latest CASTOR Gridftp+ lcg-ls)
- PIC
  - LCMAPS failing to map VOMS production FQANs
    - Correct credential is submitted to the gLite WMS
  - Some jobs were failing to upload outputs
    - No space left on device error
- CNAF
  - Queue lengths updated to accommodate LHCb jobs
  - Failures due to low number of Castor-LSF slots per diskserver
    - Issue was quickly resolved
- RAL
  - Castor issue with negative size (cfr CERN)



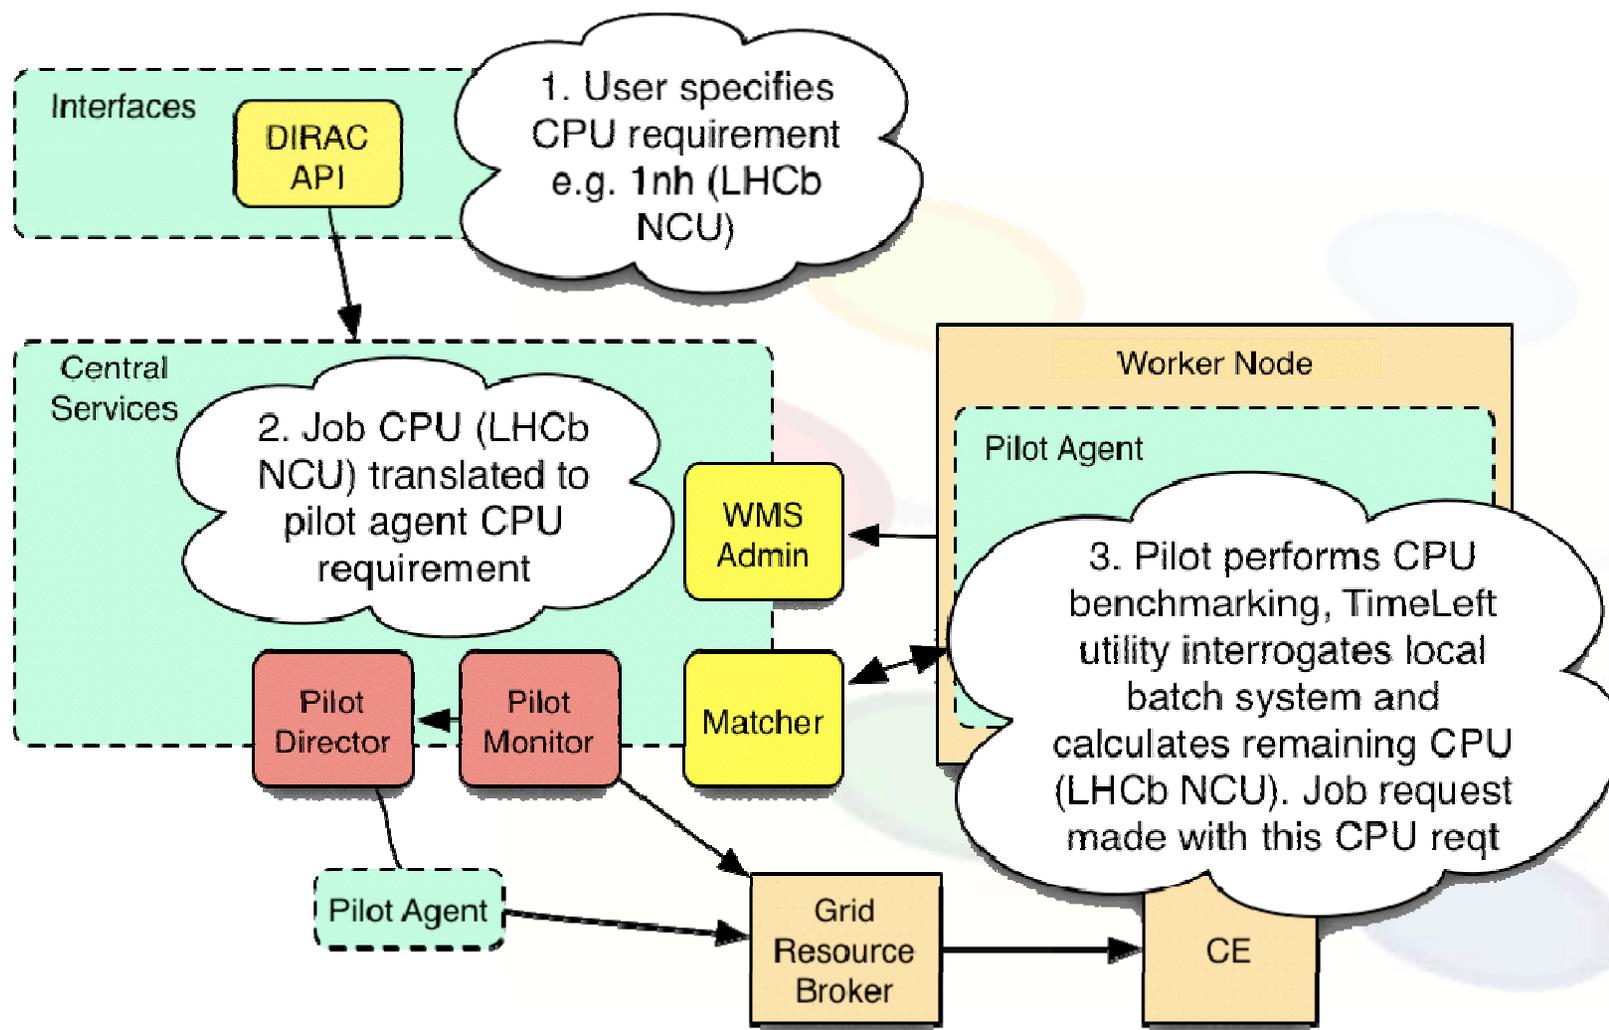
- IN2P3
  - High load on pnfs caused data access and transfer / data upload failures
- NIKHEF/SARA
  - Discovered that the longest queue is too short at NIKHEF
    - Still some successful jobs (depends on the CPU)
  - SARA problem with Sun Solaris servers affecting data access and upload
- GridKa
  - Some stalled jobs relating to dCache file access issues



- Generic pilot agent paradigm is eagerly anticipated to facilitate ‘filling’ of an available CPU slot
  - Security aspects currently being finalized by pilot working group (Maarten)
  - Can run several (different) user jobs in same slot
  - DIRAC3 is ready for ‘filling’ mode
    - Currently disabled via configuration flag
- How to ‘accurately’ determine CPU time left?
  - TimeLeft (based on LCG TimeLeft) utility improved and used in DIRAC3
    - Extracts CPU/WallClock times/limits from underlying batch system
    - Currently support LSF, PBS, BQS (ie all LHCb Tier-1s)
  - DIRAC Job Agent is instrumented to make further job requests based on this information
    - Benchmarking factor and CPU information sent back to WMS



# CPU TimeLeft Utility





## Next Steps CCRC'08 Phase 2

- Introduce Stripping workflow
- Prepare for 4 weeks steady running at nominal rate in May
- Consider alternative (and coexisting site-dependent) solutions for flickering data access via protocols
  - Download data to WNs
- If possible include analysis (using generic pilot jobs)
  - Proxy delegation already tested during CCRC'08
  - glexec will not be able for May though



## Conclusions

- Initial phase of CCRC'08 was dedicated to development and testing of DIRAC3 + interaction with SMR2
- CCRC'08 ran smoothly after initial tests
  - Online->T0 and T0-T1 transfers on the whole a success
  - Some minor issues with reconstruction activity and data upload from the WNs
  - Quick turnaround for reported problems
    - Very good responsiveness from site admins
  - Quick development and deployment of new brand of middleware or bug fixes
    - Very good responsiveness from service coordinators and software providers
- Preparation underway for 4 weeks steady running at nominal rate in May
  - Missing components (Stager/Failover/GANGA-DIRAC3 integration/ TimeLeft, Stripping workflow)