
# DIRAC status update

*A.Tsaregorodtsev,*
*CPPM-IN2P3-CNRS, Marseille*
*For the DIRAC team*

6 October 2009, LHCb Software Week, CERN

# Outline

- DIRAC subsystem highlights
  - Workload management
  - Production Management
  - Data Management
- Web portal
- Hardware configuration
- Code management, release procedures
- Conclusions

# WMS update: Priorities

- Recently introduced priorities based on user and group shares
  - Statically defined first
  - Allowed to run concurrently user and production jobs with fair shares
- Turned out to be insufficient
  - The accounting for the history of the CPU usage is necessary
  - The appropriate corrections are implemented
    - More tuning needed

# WMS: New Sandbox service

- The old database based service started to be a bottle neck with the increased load
- New Sandbox service
  - Use any Storage Element as the backend storage
    - Sandboxes stored as compressed tar files
  - Is a Storage Element itself
  - Avoids duplication of the same sandboxes
    - Will be necessary for the bulk job submission
- Migration of the clients to use the new service
  - DIRAC version prior to v4r17 is deprecated
    - Version of the 22th July
  - With the next DIRAC release old clients will get error messages while job submission
  - Old clients will not be able to retrieve output sandboxes

# WMS: further developments

- Slow entry job control
  - Single threaded optimizers that check the job sanity
  - Will be made multithreaded
- Bulk job submission
  - Bunches of jobs sharing sandboxes can be submitted in one go
  - Sandbox service is already adapted

# Prospects: new resources

- More computing resources will become available
  - DIRAC sites resurrected
    - Using pilot jobs
    - Following security rules
- OSG grid
  - If we will get access to these resources
- LHCb Online farm
  - Potentially can double the LHCb available computing resources
  - Not a standard batch system
  - Should be included seamlessly into the DIRAC production system
    - Data reconstruction and MC simulation

# Production management

- Much effort spent on the creation, monitoring and validating productions
  - Production Request is prepared by the physics coordinators
  - Productions are run and monitored by the production managers
  - Monitoring tools are available to all the users

# Production Requests

- Production Requests are defined via a Web interface
  - Sim conditions, processing path, input data, event types, etc
  - Subrequests can be defined to quickly create similar productions
- The Requests are validated by physics coordinators and production experts
- Productions are created automatically afterwards
- The Production Request progress can be followed in the Web Portal
- The recent productions were already defined using this functionality
  - The new workflows are defined (stripping, merging)

# Production Request

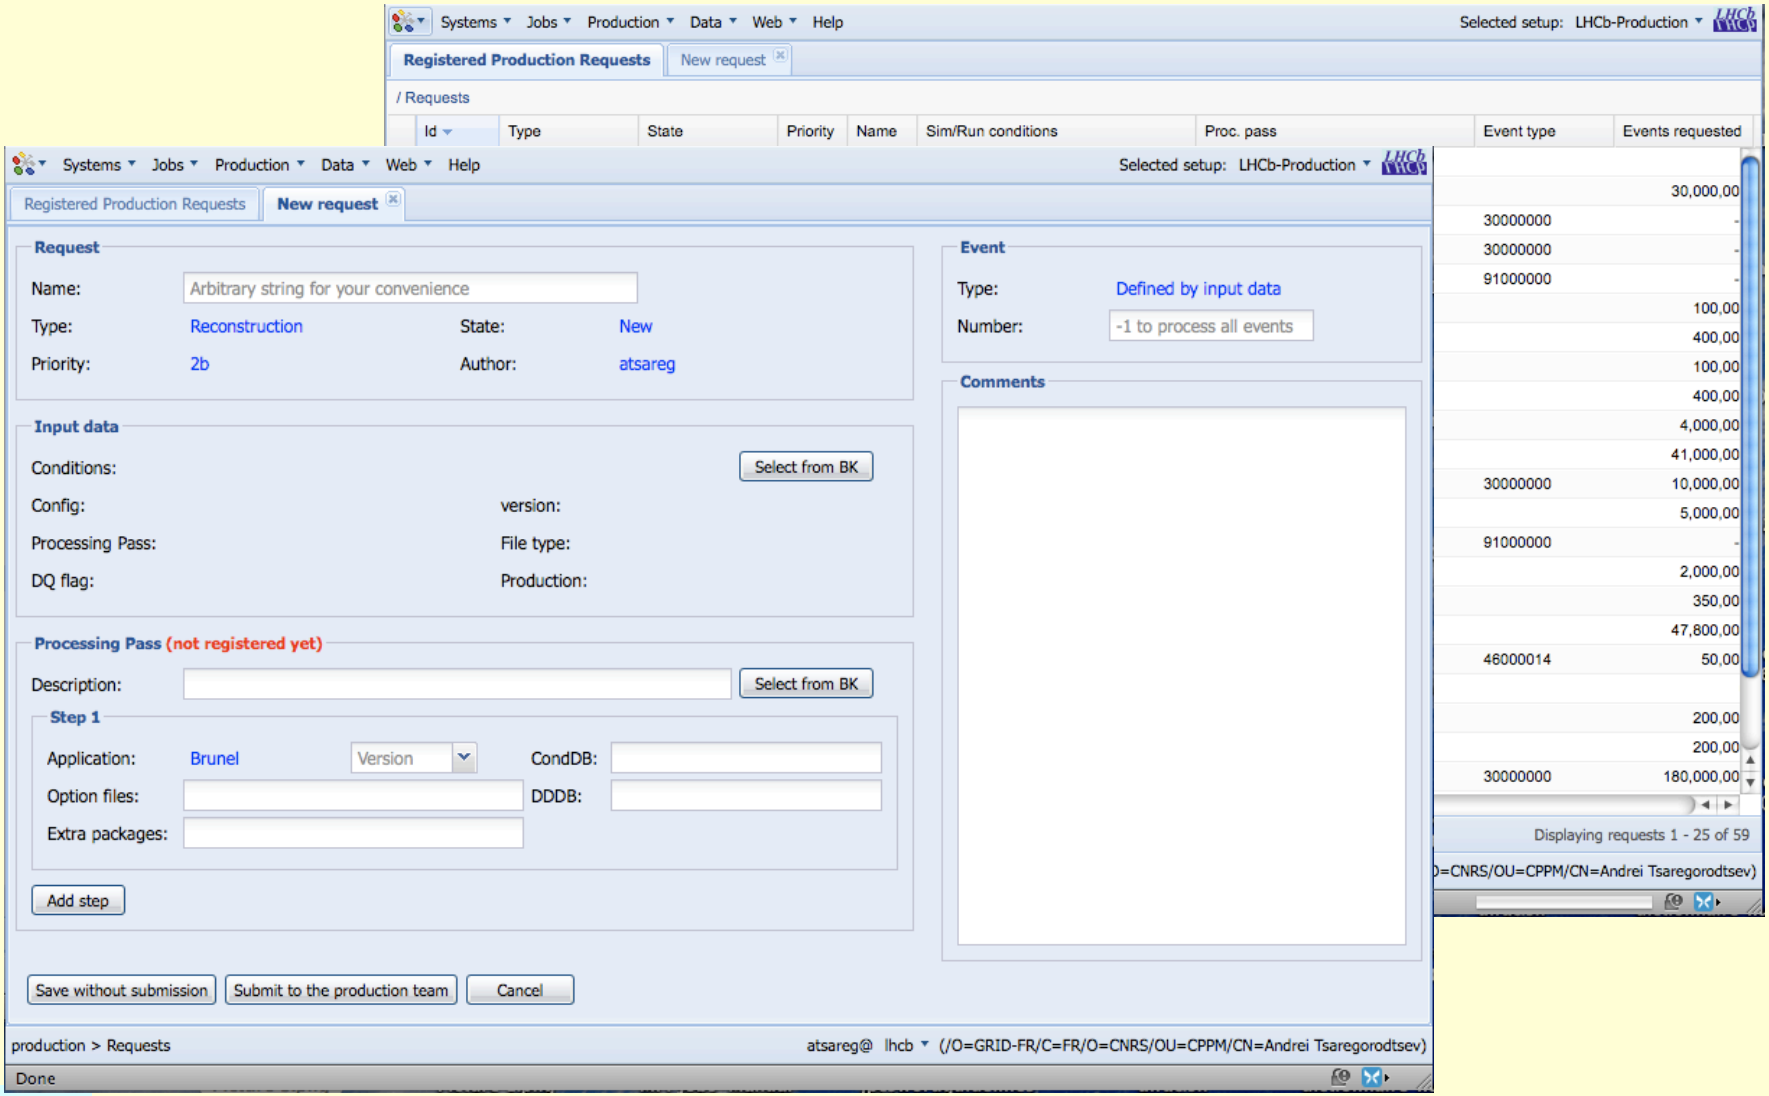

# Production Life Cycle Mgmt

- More elaborated production life cycle is introduced
- Extensive validation checks are performed after a production is completed:
  - BK information vs LFC
  - LFC vs BK, LFC vs SE
  - SE vs LFC
- Once production is validated, it is archived
  - Production definition is retained
- All the checks are fully automated

# Data Mgmt: Banning faulty SEs

- Storage Elements can be unavailalbe
  - Failures, scheduled or unscheduled shutdowns
- This should be taken into account
  - While job scheduling
    - User and production jobs
- SEs now can be declared as unavailable
  - This is equivalent to banning sites for jobs needing input data on these sites
  - For jobs without input data, the sites are still available

# Data Mgmt: Storage monitoring

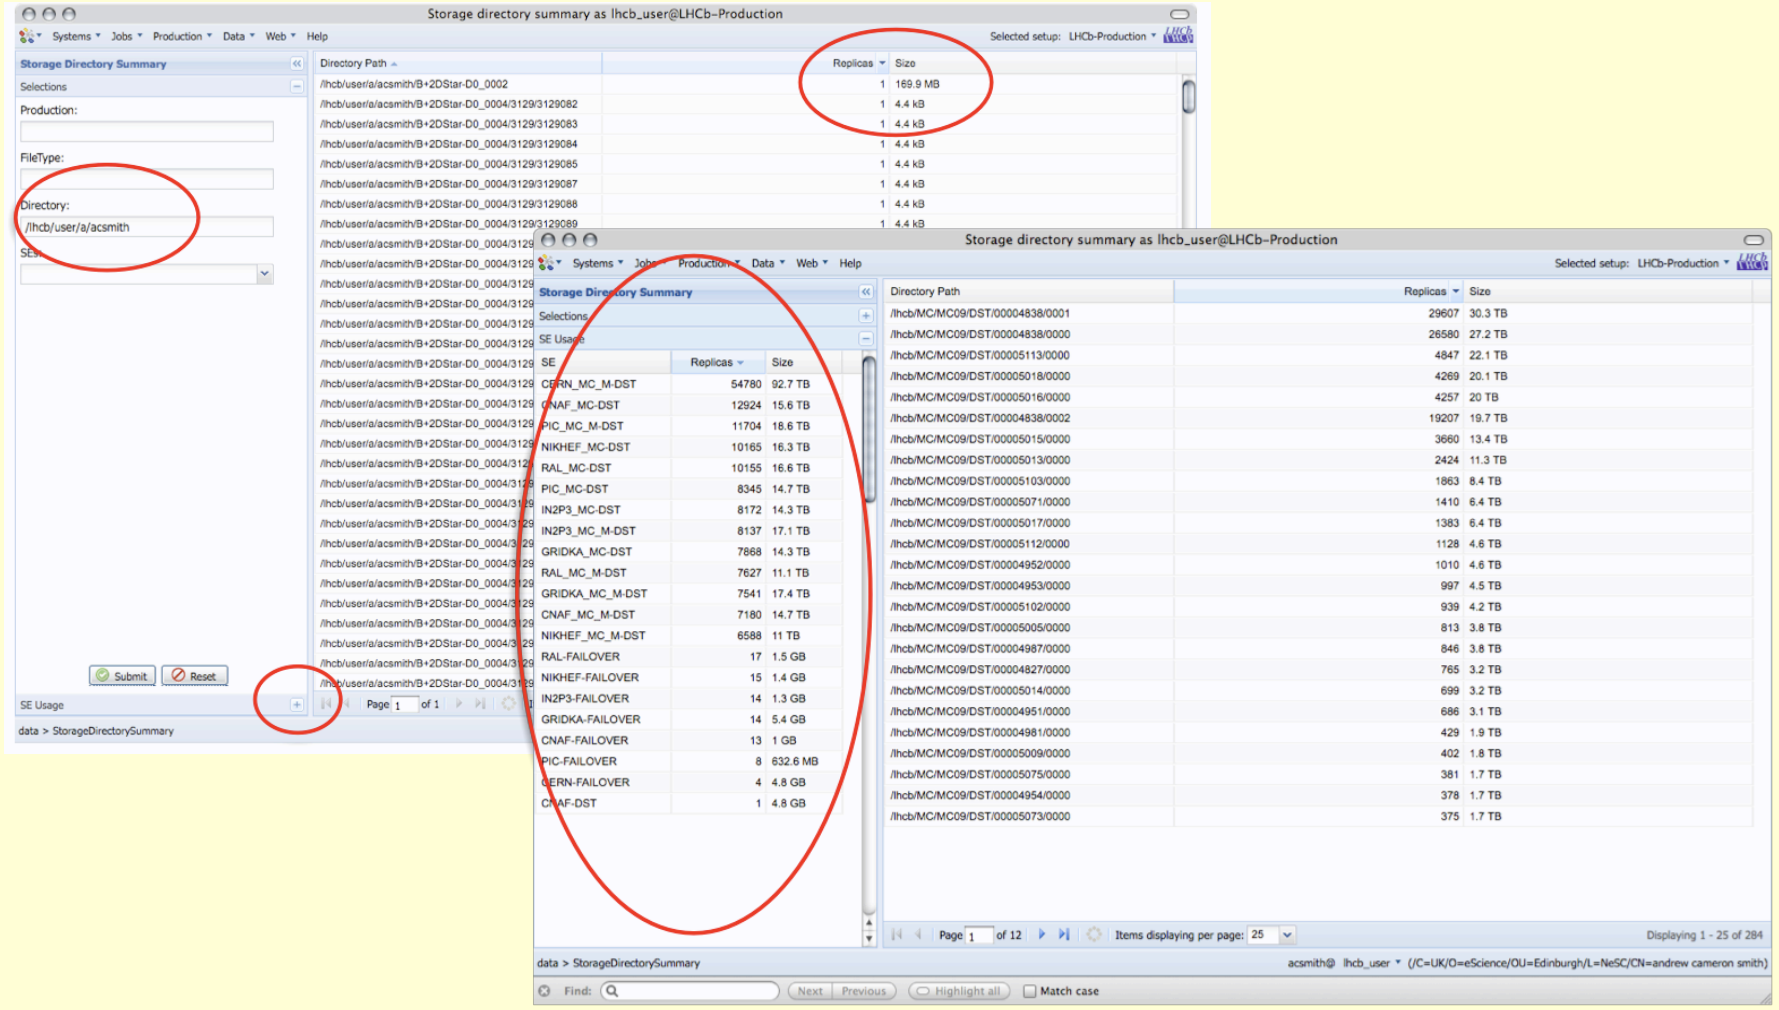

# DMS: User Storage quotas

- Storage space on the grid is not unlimited
  - Users are supposed to clean their spaces but rarely do
    - Unless they are notified about exceeding quotas
- The user storage consumption is periodically checked by a dedicated agent
  - The results are available to users
  - Currently they can be just consulted
    - Command line and API tools available
  - Eventually the user space will be locked for writing if the quotas are exceeded
    - The quotas are defined in the CS per user

# DIRAC: Secure Web Portal

- Web portal with intuitive desktop application like interface
  - Ajax, Pylons, ExtJS Javascript library
- Monitoring and control of all activities
  - User job monitoring and manipulation
  - Data production controls
  - DIRAC Systems configuration
- Secure access
  - Standard grid certificates
  - Fine grained authorization rules

# Web Portal: example interfaces

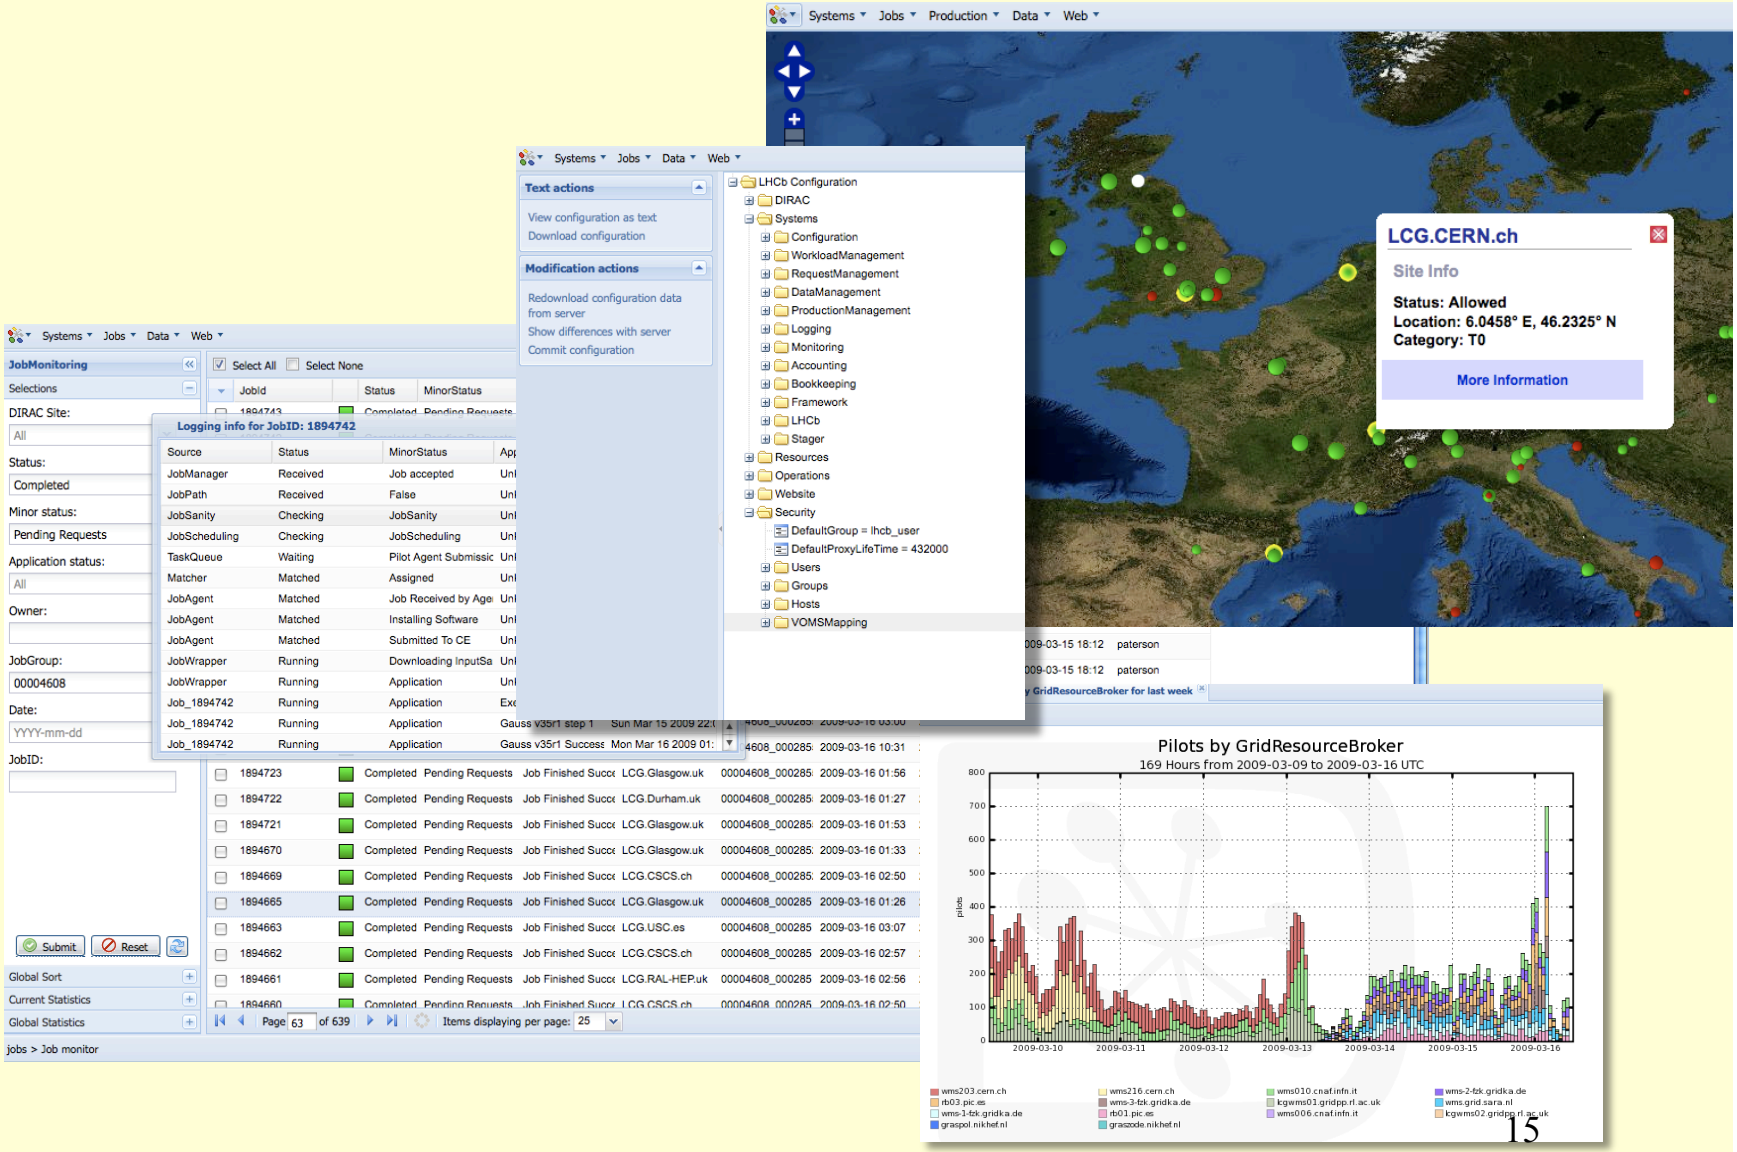

# Web Portal: new features

- Multiple developments to improve current views and to add new ones
  - Production requests, monitoring
  - Pilot monitoring
  - Etc
- Site Map resurrected
  - Mostly PR, but can be a convenient tool to navigate to the relevant site info
- Bookkeeping query page
  - Close to the Qt/GUI interface and functionality

# DIRAC: Bookkeeping page

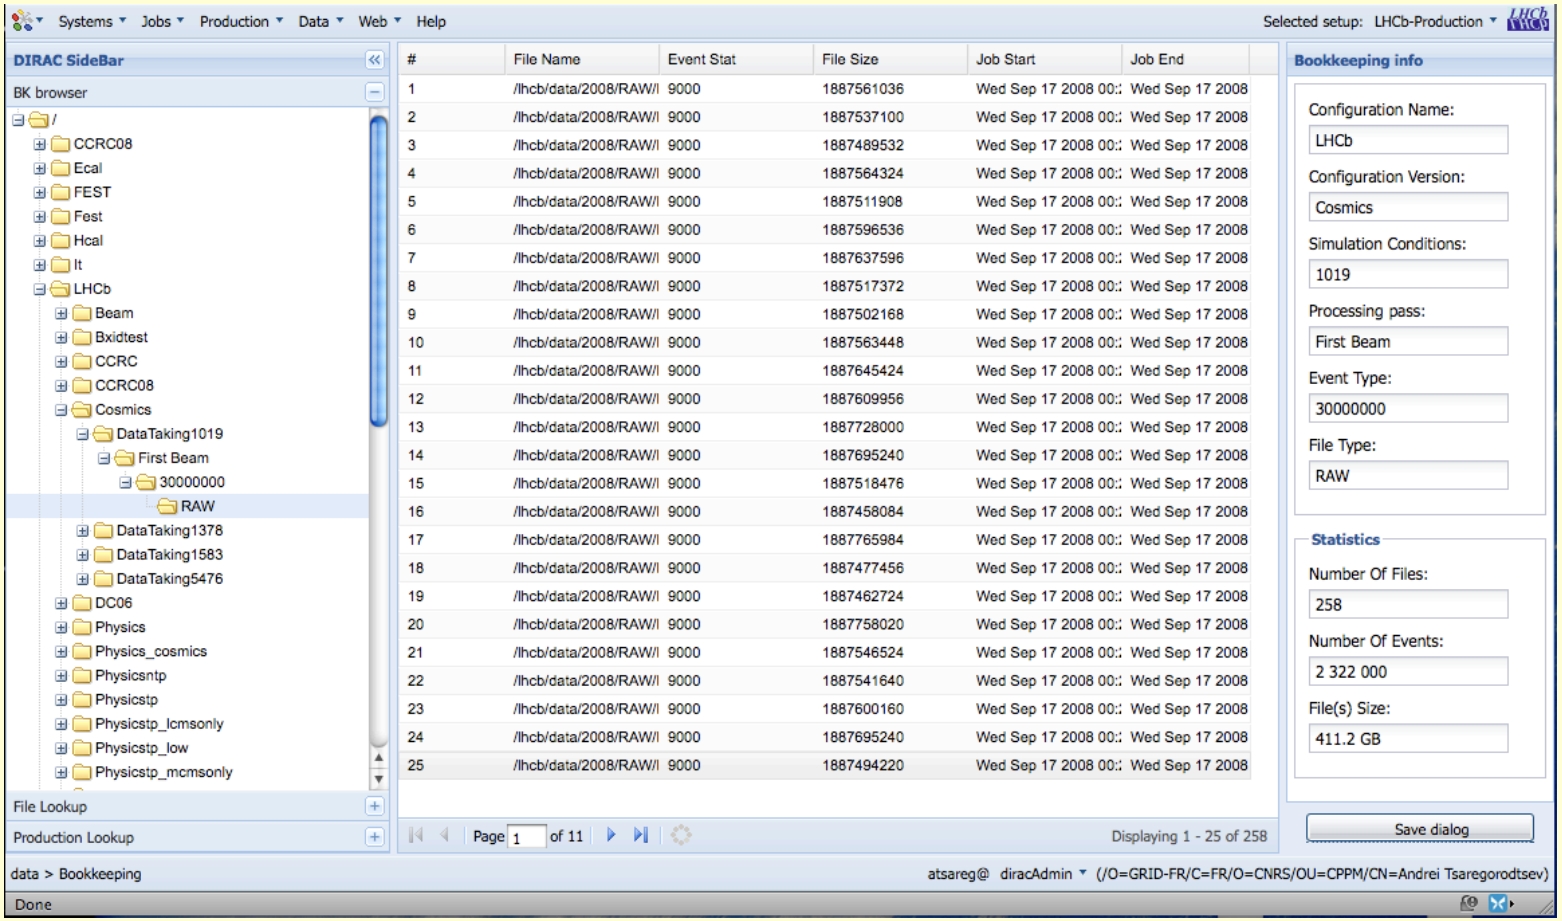


- In production this week
  - Last minute bug in the BK service communication

# Web Portal: further developments

- Web portal development will continue
- New views
  - Summaries of the site, resources, user info
- User profiles
  - Per user customizations
- Improved graphics
- More interactivity
  - Simple job submission
  - Access to the job sandboxes

# Hardware update

- The single central DIRAC WMS/Production server was split in 3:
  - volhcb04 - security sensitive services
    - Proxy repository, Security logging
  - volhcb13 - WMS components
    - More powerful with two HDDs
  - volhcb09 - Production Mgmt components
- The overall system capacity increased by a factor ~2.

# Performance

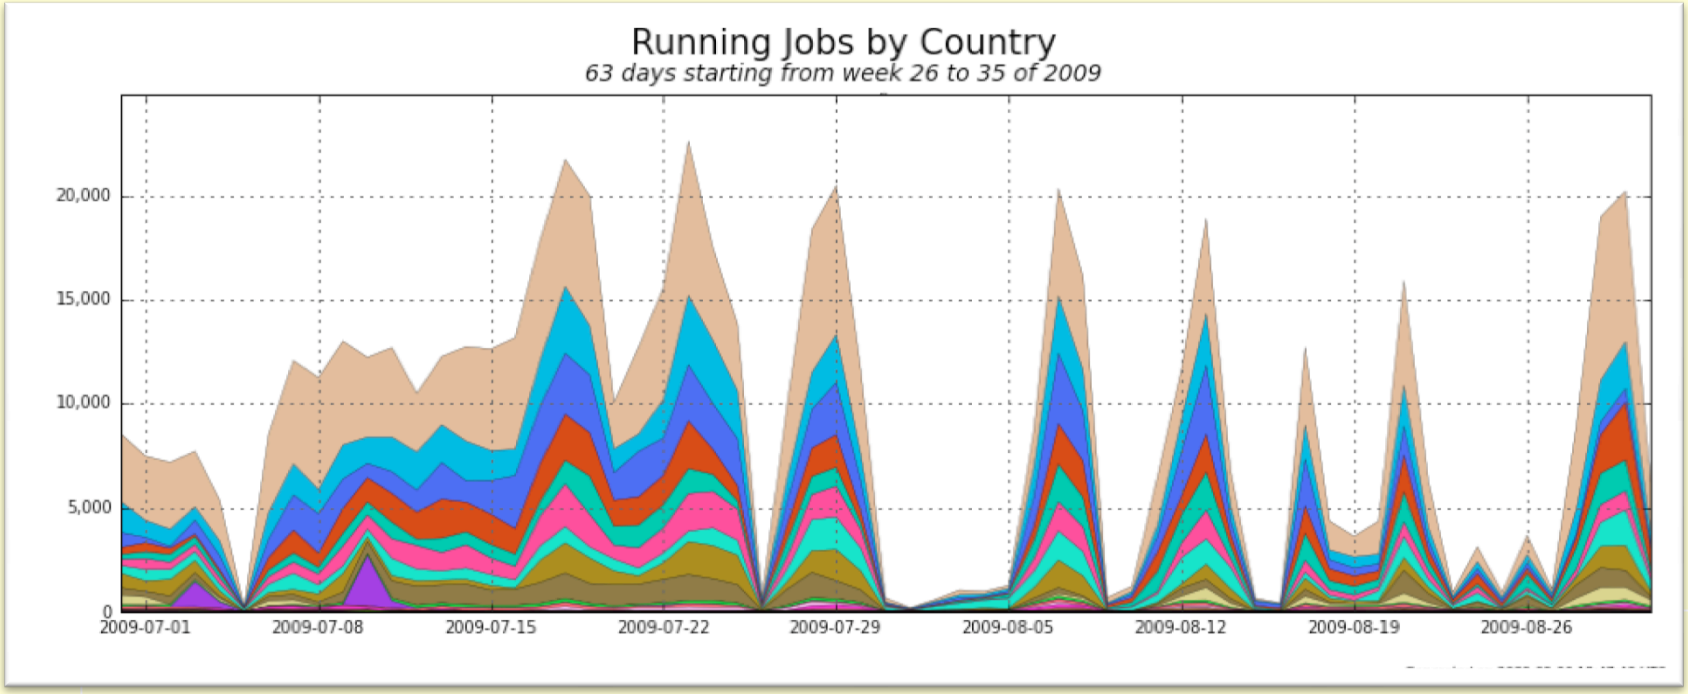


- DIRAC performance measured in the recent production and FEST'09 runs
  - Up to 25K concurrent jobs in ~120 distinct sites
    - Up to 75K jobs per day
  - Further optimizations to increase capacity are possible
    - Hardware, database optimizations, service load balancing, etc

# DIRAC Central Servers

- Front-end DIRAC servers
  - DNS or client level load balancing
- Database
  - Multiple MySQL databases
  - Real time mirroring of Master to Slave databases
  - Redundant disk arrays
  - Load balancing on the database client level
  - ORACLE solution is considered as well
- Expected operations advantages
  - Large capacity margin ( factor 2-3 )
  - Failure of any server will reduce but not disrupt operations, hot swapping of any server

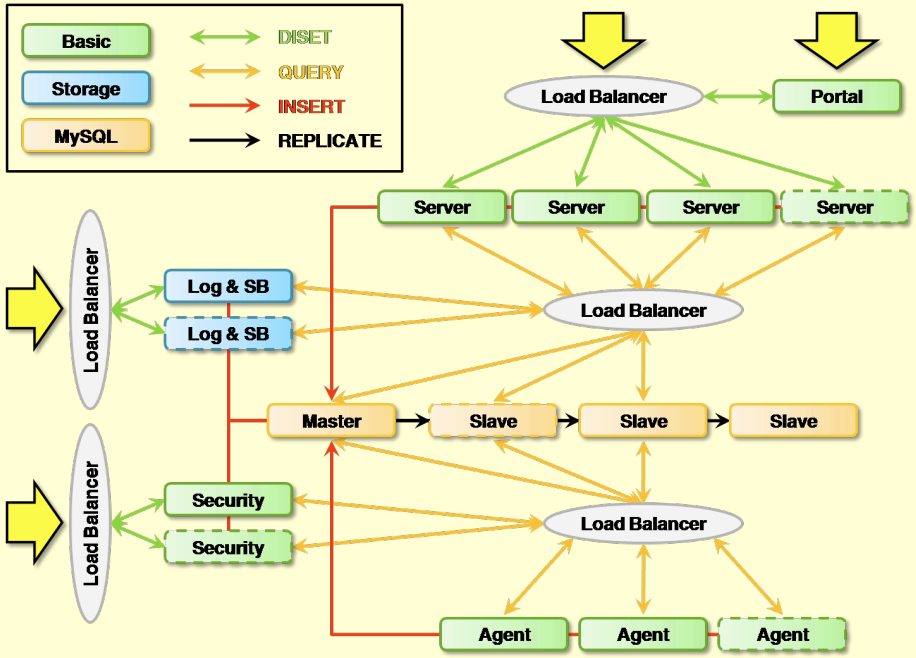

# Distribution, installation

- DIRAC is now fully available on the SLC5 platform
  - Shipped grid middleware binaries
- Use of Python 2.5 is still subject to the availability of the python grid middleware bindings
  - Memory leaks reported
- CA, CRL renewal automation
  - Command line tools
    - Useful for the user client installation
  - Automatic agent
    - For service installations

# Certification DIRAC instance

- DIRAC releases are now first tested on the Certification setup
  - Separate servers for the WMS services and for the Web Portal
  - Other systems will be added as well
- The quality of the releases should be much improved

# DIRAC v5r0 release

- A new major DIRAC v5r0 release is being prepared
- Updated binary packages
  - GSI
  - MySQL
  - Matplotlib graphics library
- Discontinued old Sandbox service
- No backward compatibility with the deprecated versions of DIRAC
  - Prior to v4r17

# Prospects: code consolidation

- Many small improvements to fix problems and improve performance and stability
- New hardware configuration will be supported in the code
  - Multiple service end-points
  - Multiple backend database instances
    - Read/write and read only
- More Database optimizations are envisaged

# Moving to SVN repositoy

- The DIRAC code is stored in the CVS repository
  - Discontinued in the fall 2009
- The code will be moved to the SVN repository
- The code will be restructured
  - Several independent projects with separate release cycle:
    - Core DIRAC, LHCbSystem, Web Portal
- The new repository structure is being prototyped now

# Conclusions

- DIRAC covers all the LHCb tasks for the distributed data processing implemented in a single coherent framework
- Many efforts going into the system polishing and code consolidation
- DIRAC demonstrates the performance adequate to the LHCb needs for real Data Taking

*http://dirac.cern.ch*