LHCb Computing & Software

Nick Brook

University of Bristol

- Organisation
- Software
- Data Challenges
- Production & analysis tools

Organisational Issues

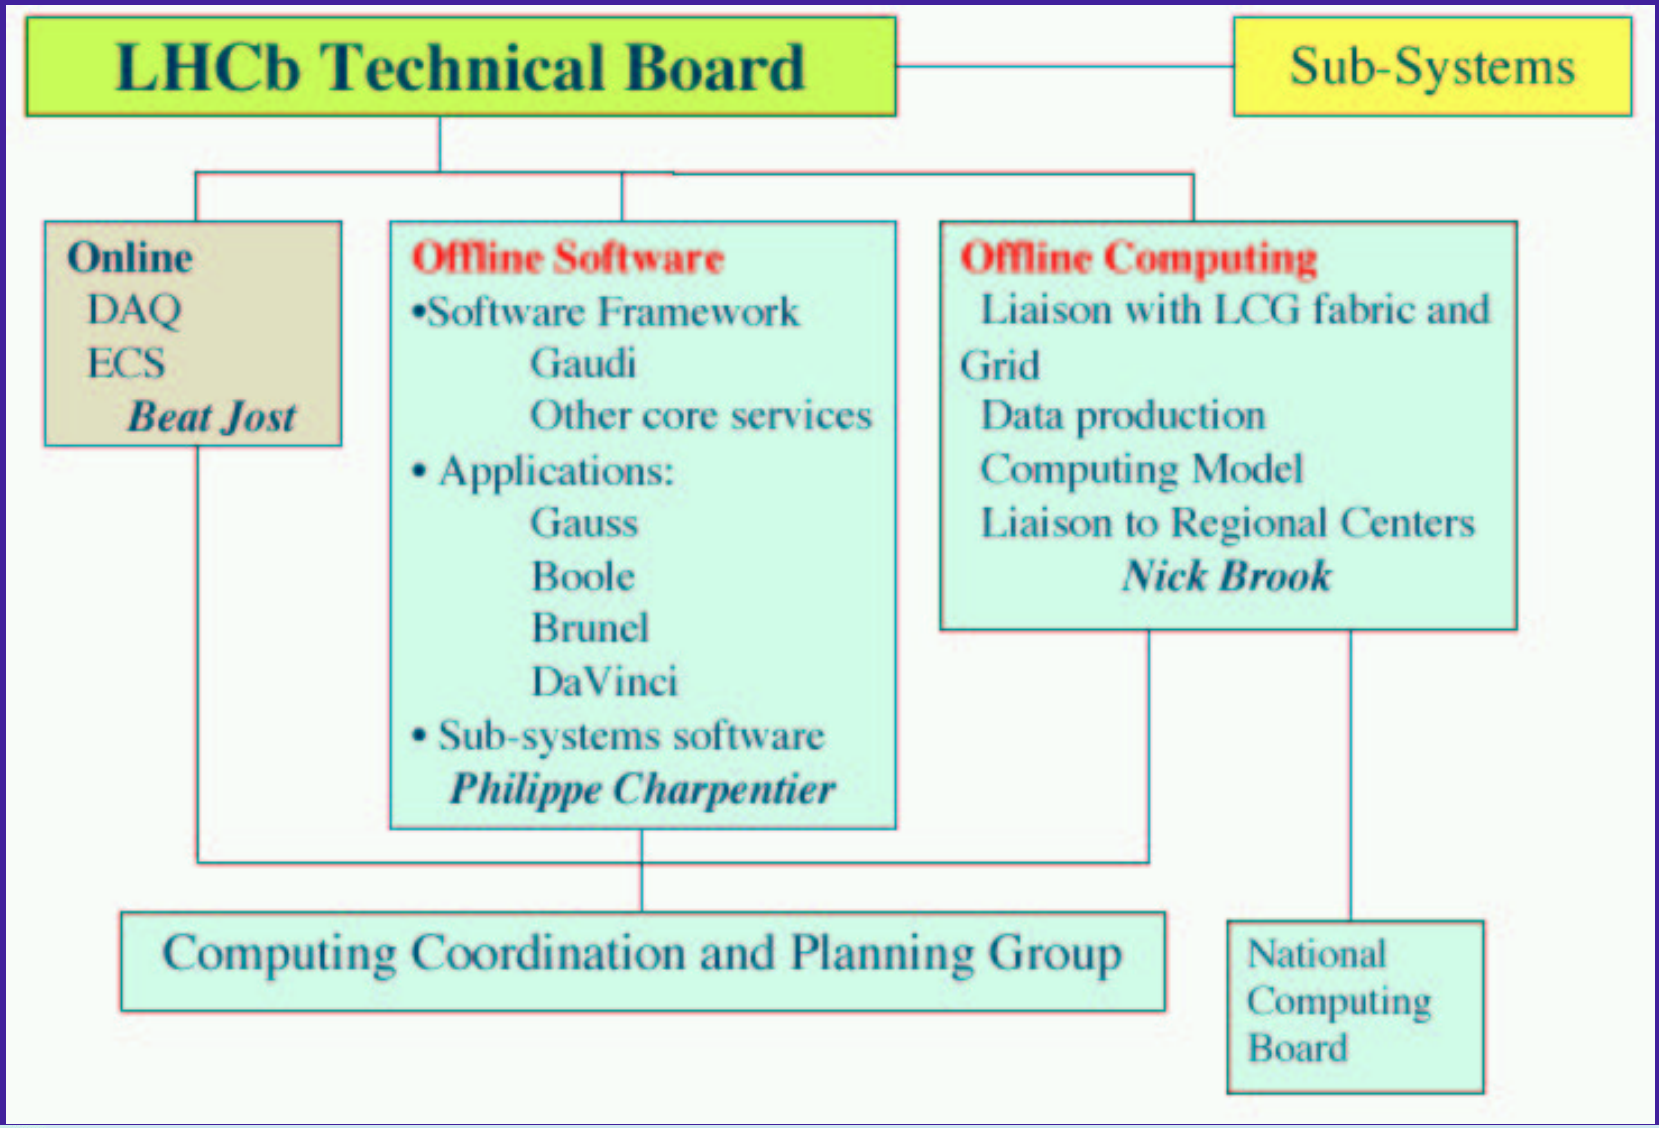

Milestones

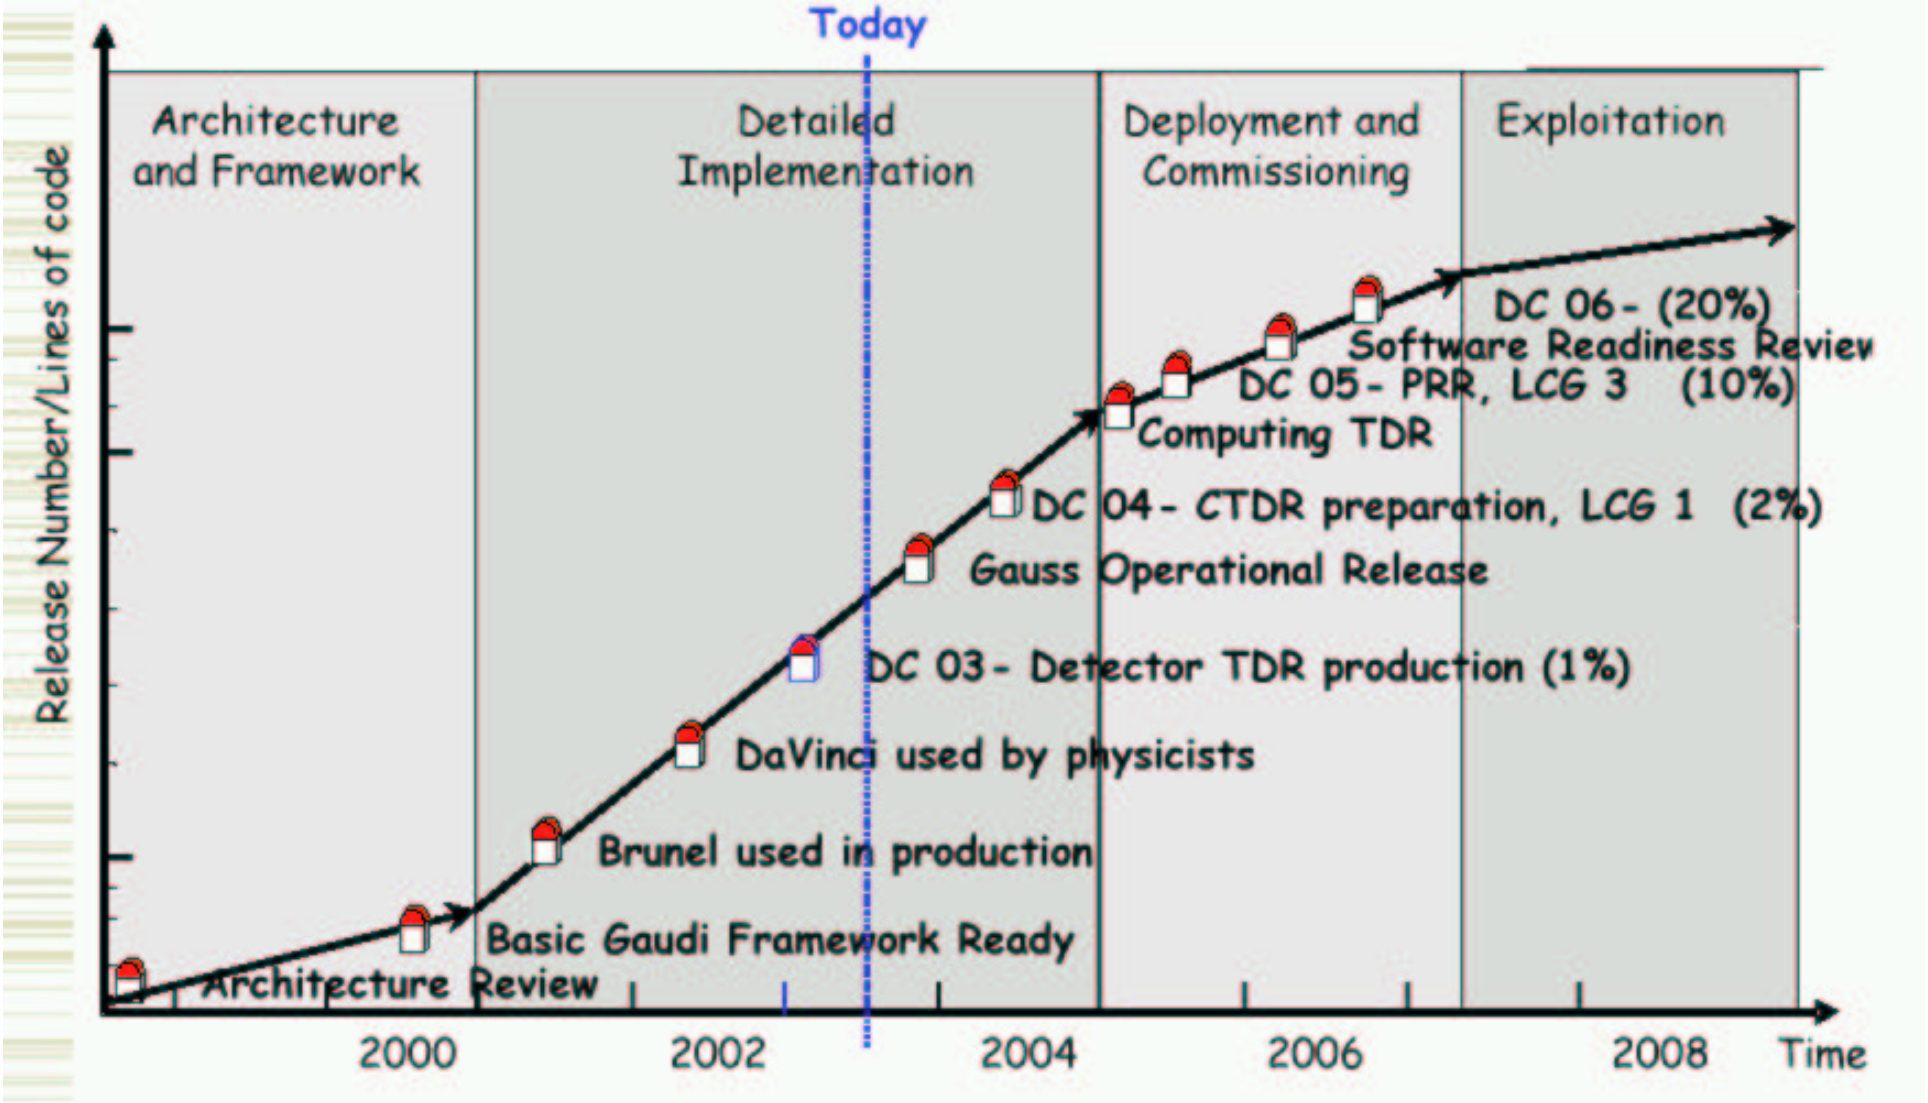
Today
Architecture and Framework
Detailed Implementation
Deployment and Commissioning
Exploitation
Release Number/Lines of code
Architecture Review
Basic Gaudi Framework Ready
Brunel used in production
DaVinci used by physicists
DC 03 - Detector TDR production (1%)
Gauss Operational Release
DC 04 - CTDR preparation, LCG 1 (2%)
Computing TDR
DC 05 - PRR, LCG 3 (10%)
Software Readiness Review
DC 06 - (20%)
2000
2002
2004
2006
2008
Time

Software

- Develop an Architecture ('blueprint') and a Framework (real code) to be used at all stages of LHCb data processing
  - high level triggers, simulation, reconstruction, analysis
  - a single framework used by all members of the collaboration
- Avoid fragmentation and duplication of computing efforts
  - common vocabulary, better understanding of the system
  - better specifications of what needs to be done
  - identify and build common components
    - guidelines and coordination for sub-detector groups
- Transparent use of third-party components wherever possible
  - GUI, persistency, simulation....
- Applications are developed by customizing the Framework

Software Organisation

Applications built on top of frameworks and implementing the required algorithms.

One framework for basic services + various specialized frameworks:
detector description, visualization, persistency, interactivity, simulation, etc.

A series of basic libraries widely used: STL, CLHEP, GSL etc.

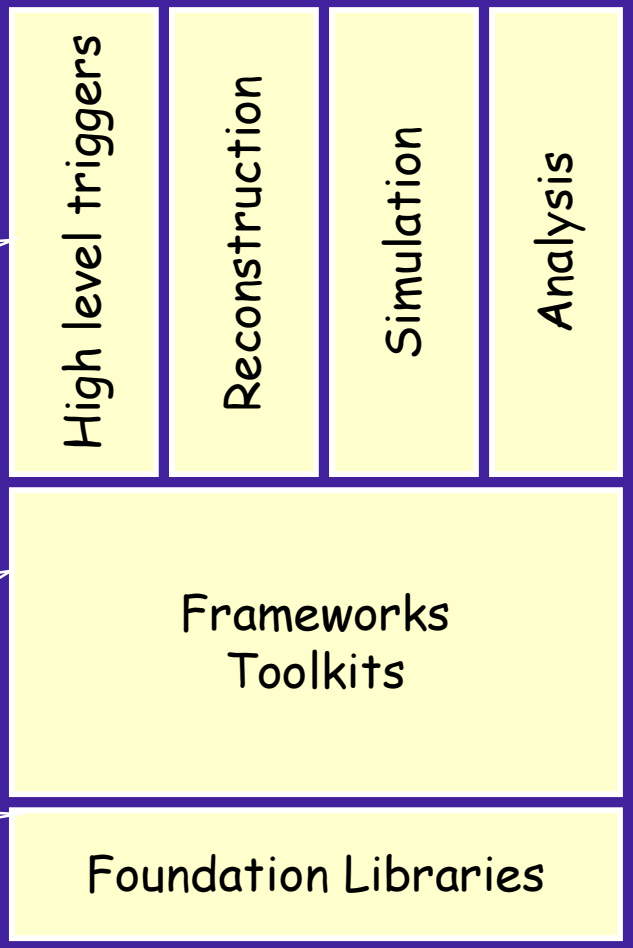

Offline Software

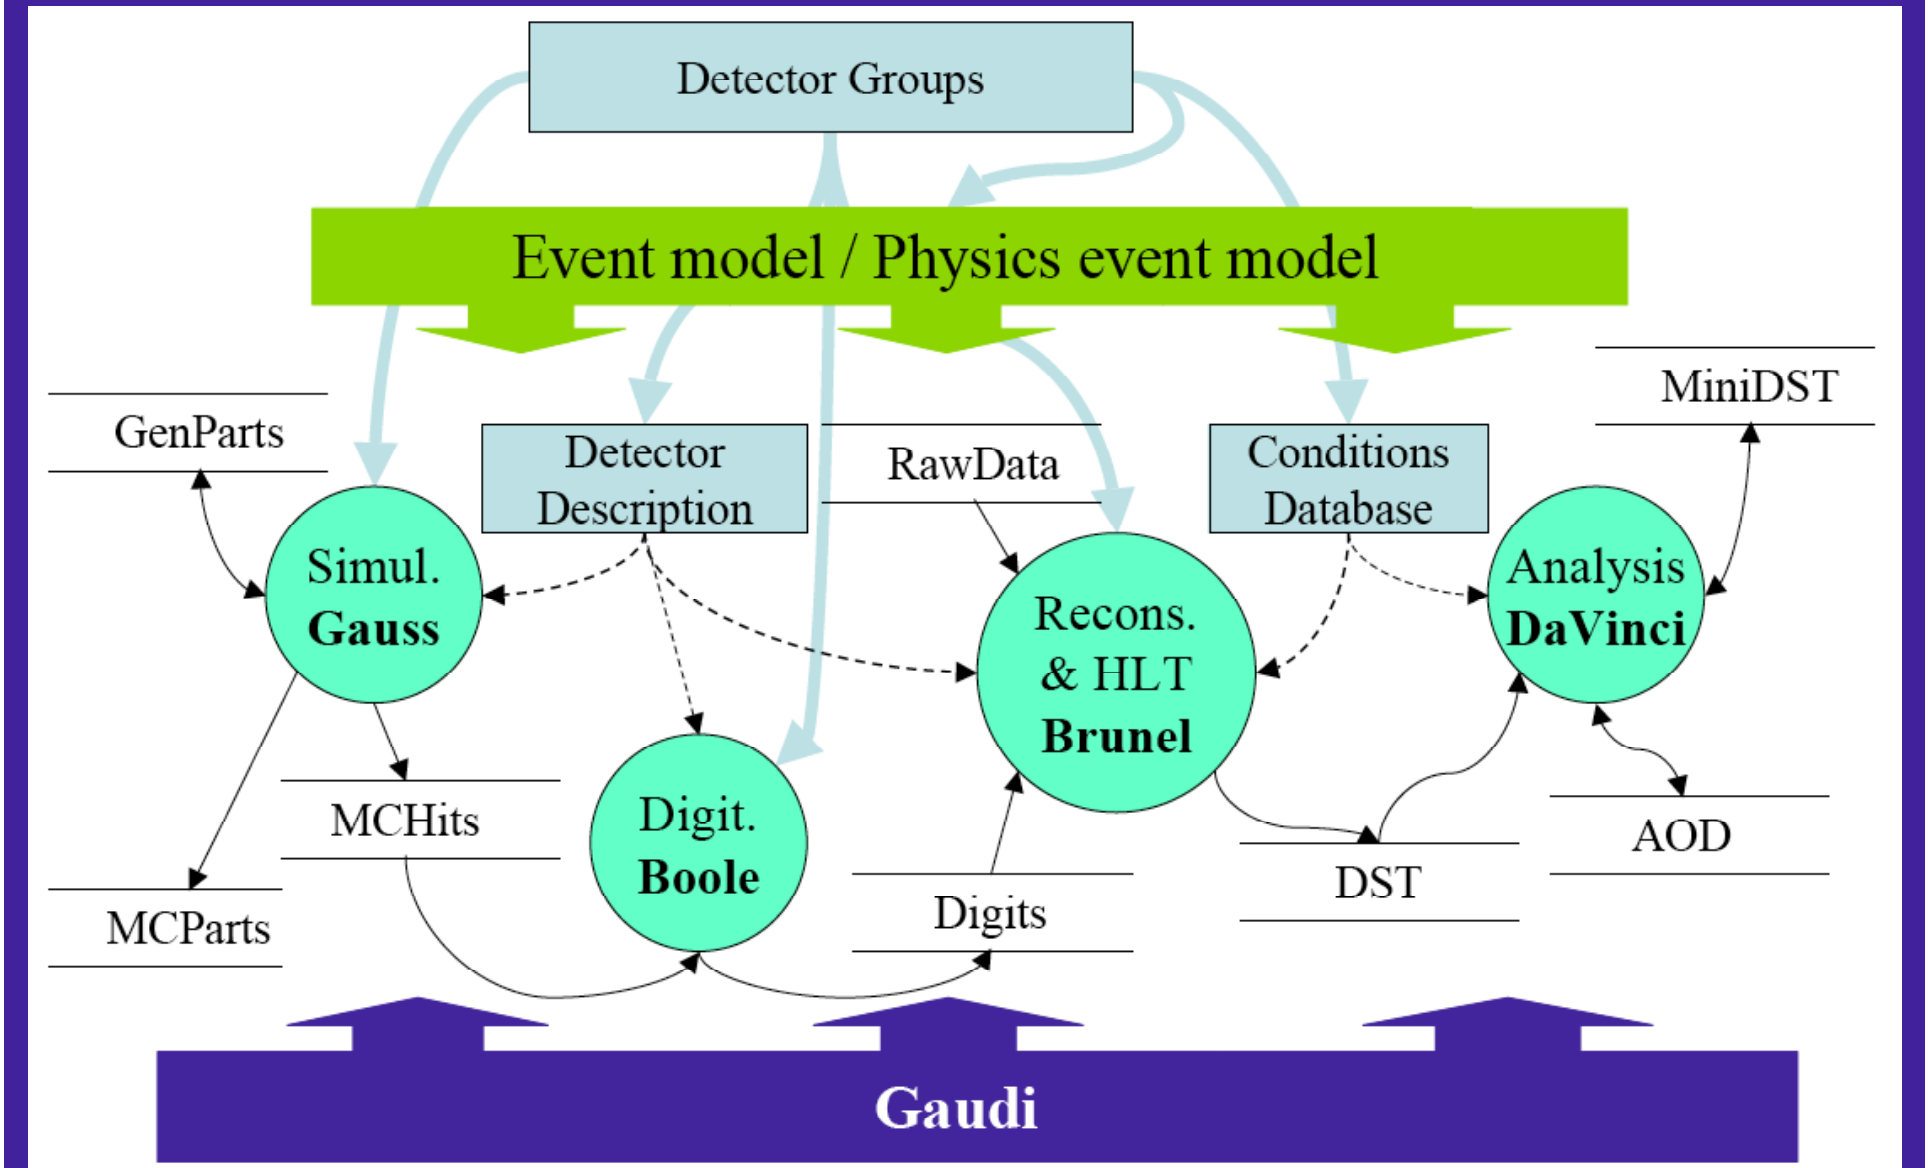

Gaudi Framework

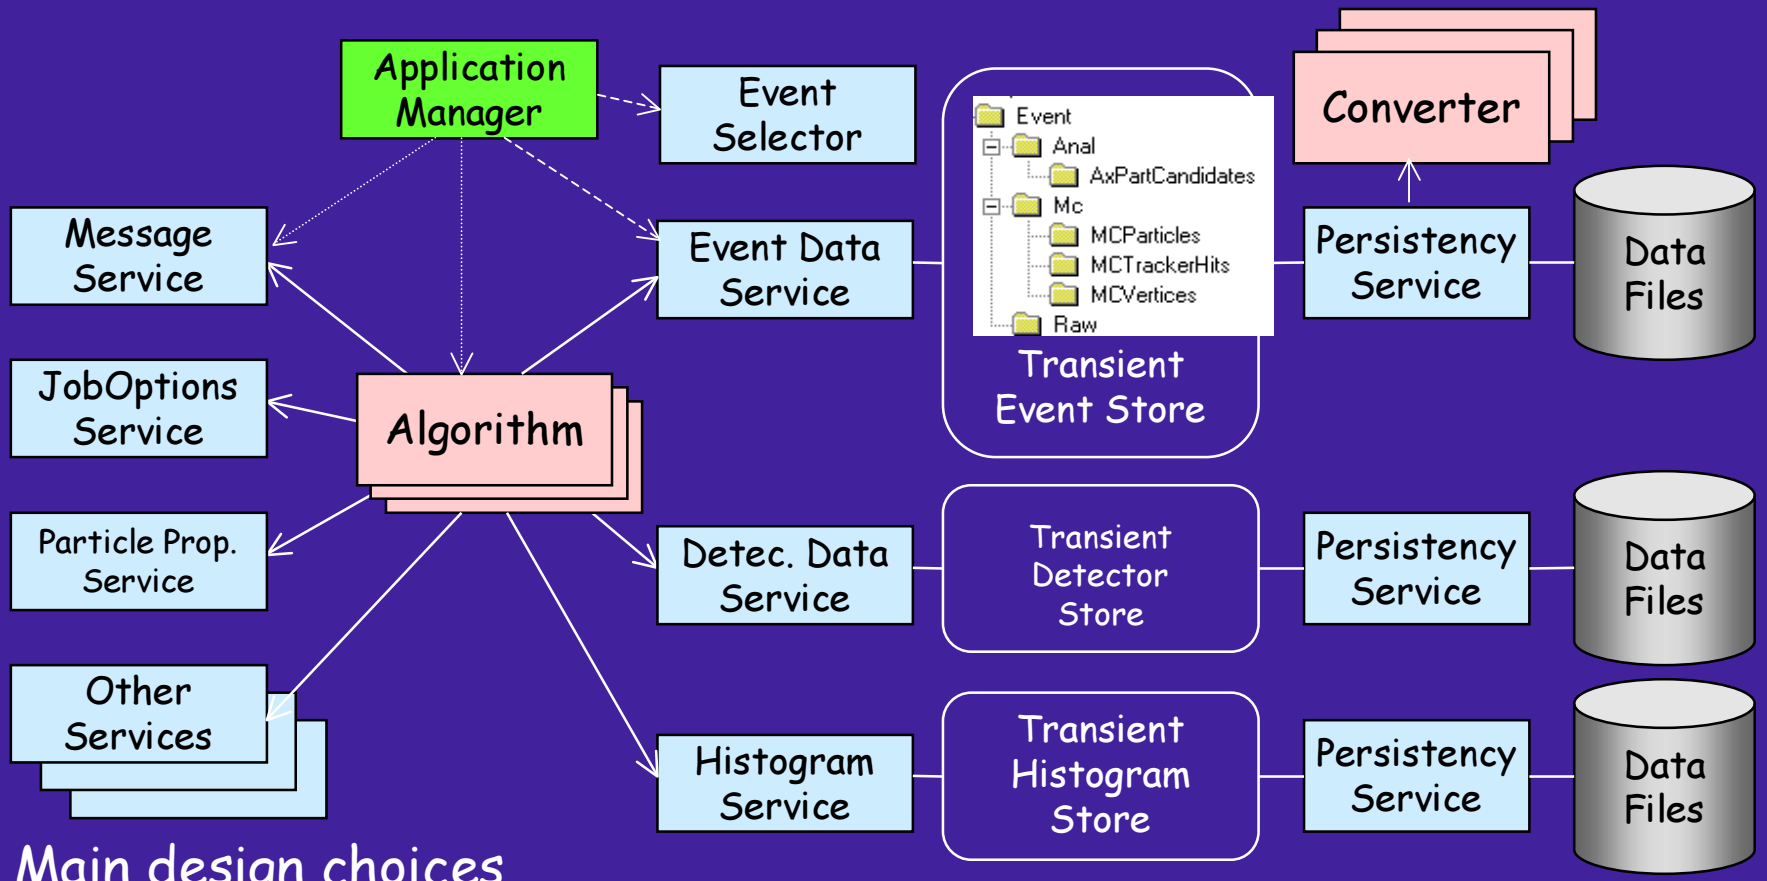

- Main design choices
  - Separation between "data" and "algorithms"
  - Separation between "transient" and "persistent" representations of data
  - "Physicist code" encapsulated in few specific places (*Algorithms*, *Tools*, *Converters*)
  - Well defined component "interfaces"

GAUSS - simulation

- Currently the LHCb Simulation program SICBMC is a Fortran Geant3-based program
  - DC'03 performed with SICBMC
  - SICBMC retired at end of 2003
  - DC'04 will be performed with Gauss
- Event Generation and Detector simulation as a Gaudi based application is provided by Gauss.
  - Specialized External software (e.g. Pythia, GEANT4)
  - To be used for DC 04
  - Undergoing final testing for production release
  - Pileup in the same bunch crossing is handled at event generation time

GAUSS - simulation

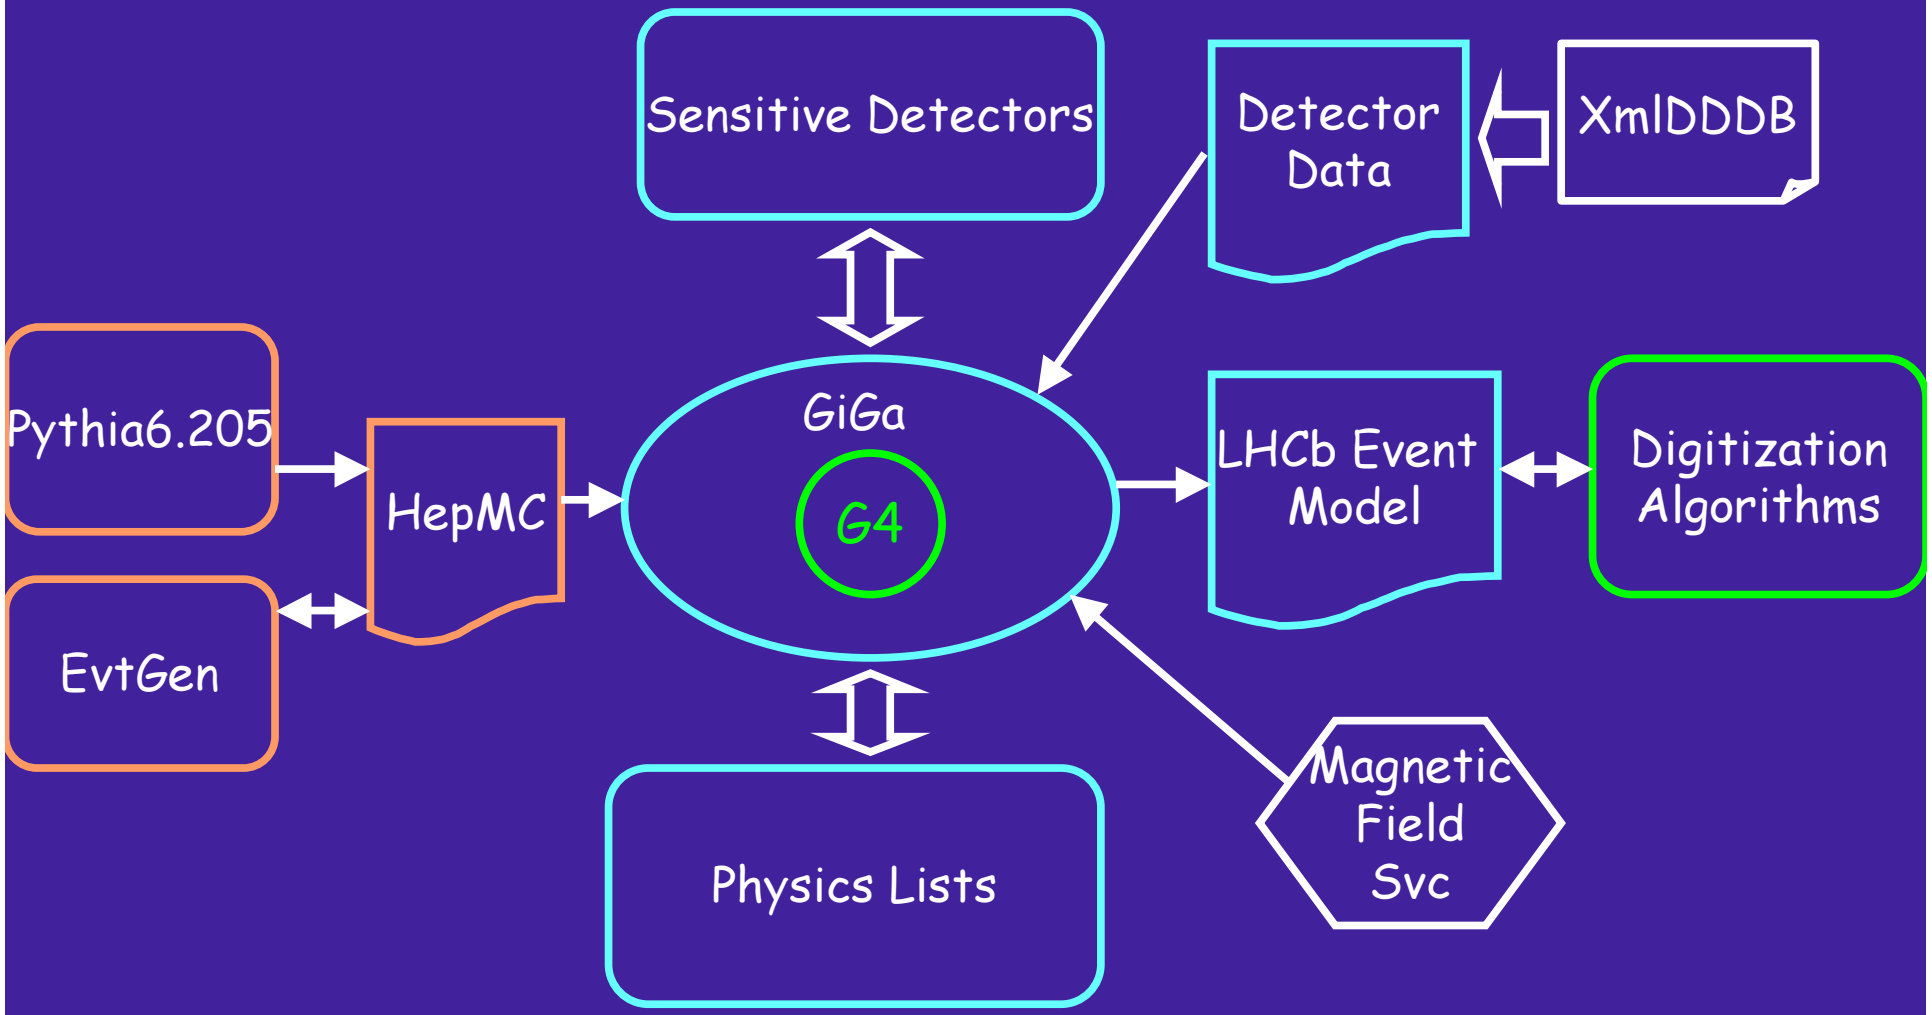

BOOLE - digistisation

- Simulation of detector response transforming hits in sensitive detectors to produce digitized data and provide them in DAQ-like format is provided by Boole
  - Handling of spill-over effects (data from adjacent beam crossings)
  - Simulation of detector response
  - General backgrounds and detector imperfections
  - L0 and L1 trigger as last step
    - Large data reduction before reconstruction
- Separate application from Gauss
  - Convenience
  - Flexibility (process GEANT3 hits)
- Up to mid-2003 a single application with Reconstruction
  - Split completed and validated
  - Interfacing to Gauss under validation
  - Output of digitization in DAQ format recently started
    - General DAQ-Event

BRUNEL - reconstruction

- Reconstruction of both Monte Carlo data and DAQ data is provided by Brunel
  - Many components that require careful coordination
    - Clustering, Pattern recognition, Track reconstruction, Particle ID
  - Provides global reconstructions objects
  - Output are DST file containing final results of sub-systems and global reconstruction
  - High Level Triggers
- Up to mid-2003 included digitization of SICBMC data
  - Digitization and Triggers no longer part of Brunel
  - Process output of Boole
    - Completely independent from Simulation
- Brunel operational and used in production since August 2002
- Brunel has been used in DC03 to produce the Re-optimization and Trigger TDRs

DaVinci - analysis

- Group and individual analysis proceeds from the output of the reconstruction. The framework for selections of events and analysis is called DaVinci.
  - "Particles" are produced from reconstruction objects with a specified particle ID hypothesis, interesting decays are "built" via selection and combination of particles
  - Tools of general utilities are provided:
    - for vertexing and particle manipulation
    - to compare with MC truth for Monte Carlo data and evaluate physics performance
  - Integration of event selection algorithms
    - Stripping data
    - Contents are responsibility of physics groups and outside scope of this Review
- Output can be reduced miniDSTs or AODs
- DaVinci used for Physics studies since July 2002
  - Most users interact with this application

Data Challenge'03 - Goals

- Large statistic of simulated events for the Trigger TDR and the Re-optimisation TDR
- 30 Million events of Minimum Bias
- 10 Million events of generic b-decays
- Between 50 K and 150 K events per channel for ~30 decay channels

Data Challenge'03 - Results

- *A total of 47 Million events have been produced in two months*
- 18 centres participate and 80% of CPU outside CERN
- 36 600 jobs have been run and each job:
  - producing between 250 and 500 events
  - using from 32 to 56 hours on 1GHz PC

- *It would have taken more than 170 years on a single PC*

Data Challenge'03 - Results

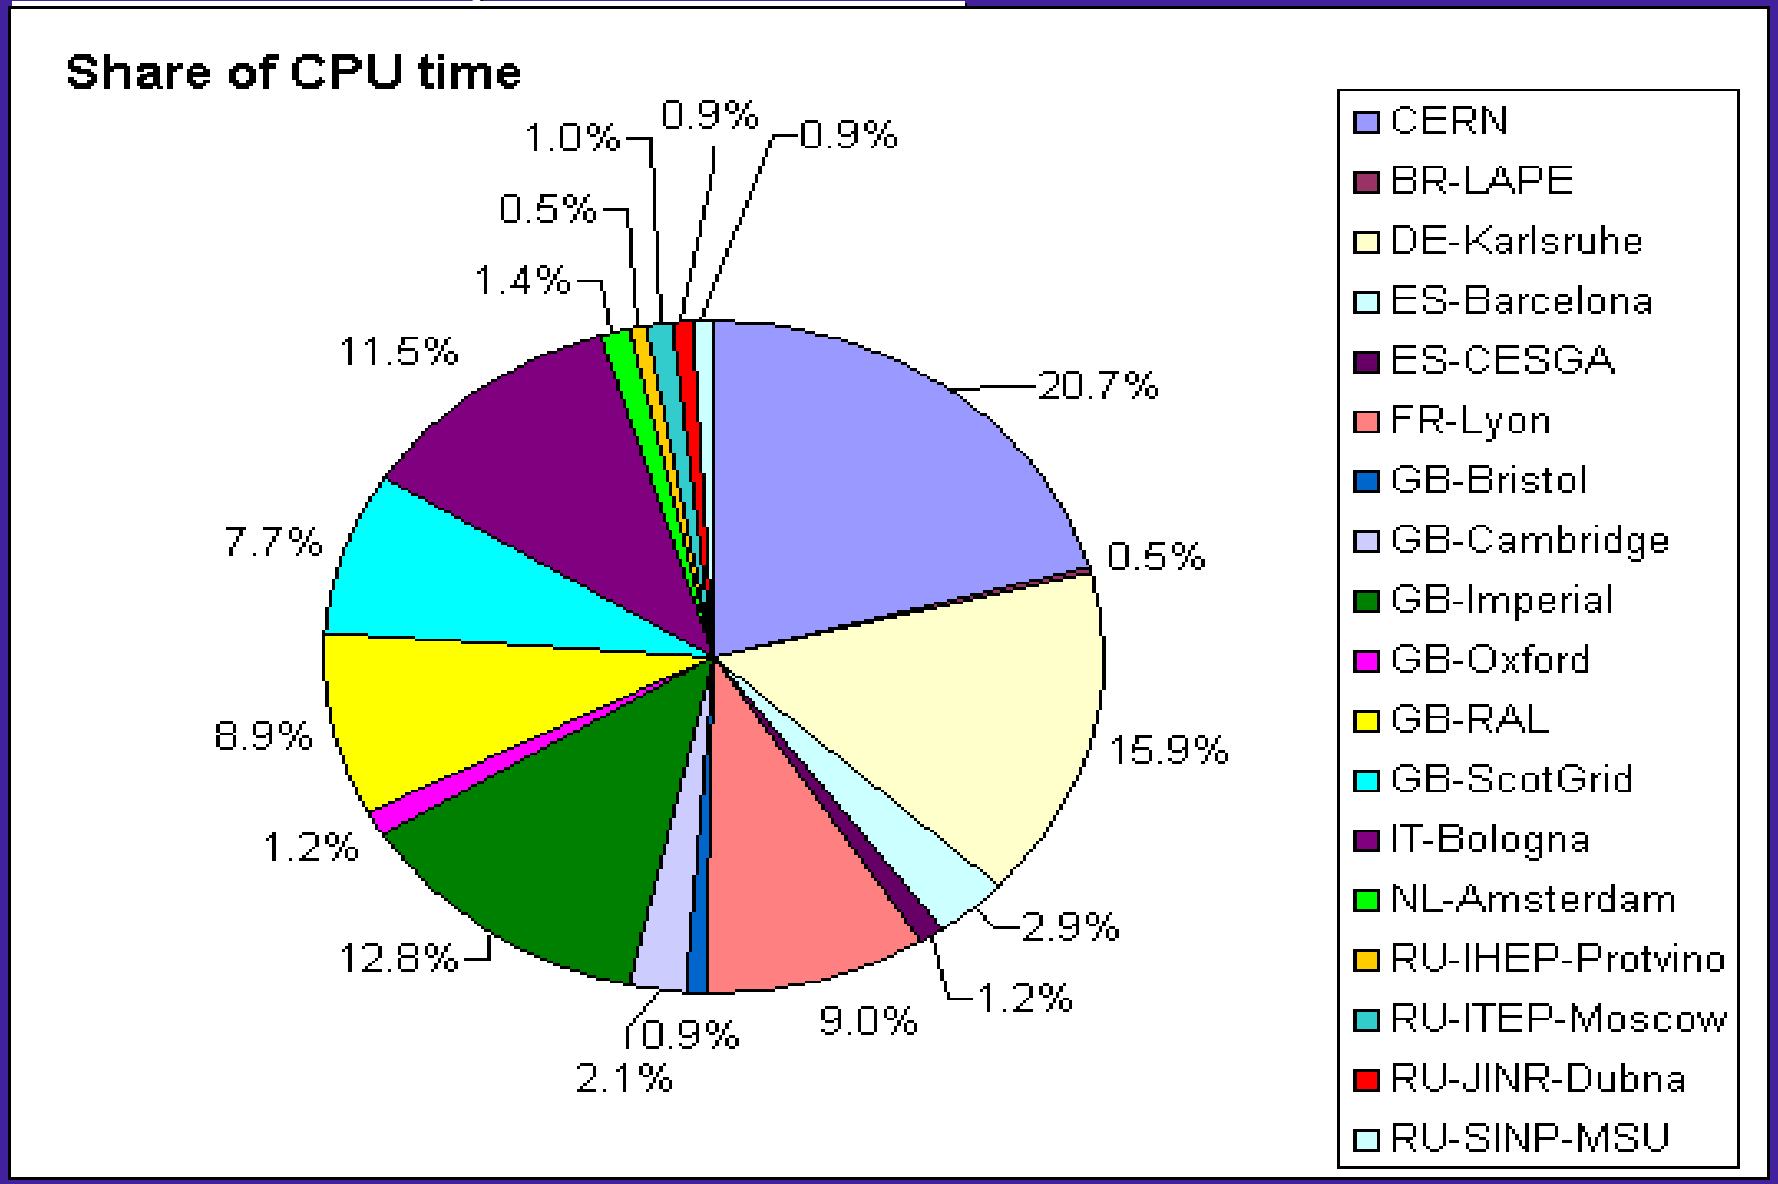

Data Challenge'03 - Results

- Data Challenge co-ordinated from CERN
- None of the 18 centres are LHCb dedicated and to first order there is no dedicated LHCb manpower to operate the computing centres
- External manpower needed to perform DC'03 - integrates to about ~1FTE
- Large external centres are not even LHC dedicated like IN2P3, RAL or Karlsruhe.
- Some external centre are not even Particles Physics dedicated like Imperial College

Data Challenge'04 - Goals

- Robustness test of the LHCb software and production system
- Test of the LHCb computing model
- Perform distributed analyses
- Incorporation of the LCG application area software into the LHCb production environment
- Use of LCG resources/environment as a substantial fraction (>50%) of the production

Data Challenge'04 - Numbers

- Production of :
  - 150 Million events of Minimum bias
  - 50 Million events inclusive b decays
  - 20 Million exclusive b decays in the channels of interest
- All events will be digitised
- Minimum Bias and inclusive b decays passing the L0 and L1 trigger will be reconstructed and the entire exclusive b decays as well
- The Data Challenge will be run by the production manager at CERN in close collaboration with the LHCb production site managers

Data Challenge'04 – Robustness tests

- First use of the simulation program **Gauss** based on Geant4
- Introduction of the new digitisation program, **Boole**
- Robustness of the reconstruction program, **Brunel**
- All the application will use the POOL persistency mechanism to store event data and SEAL services.

Data Challenge'04 – computing model tests

- The output of Brunel (DST) will be stored
- The output of Gauss (RAW) and Boole (DIGITS) will be stored if they pass the trigger.
- Pre-selection of events based on physics criteria will be performed after the Reconstruction and multiple DST output will be produced
- The master copies of RAW/DIGITS will be about 14TB and 7TB for the DST - all DSTs will be copied at CERN.
- A factor of 10 less storage is expected for pre-selected events that will be replicated to associated Tiers 1 to test the distributed analysis

Data Challenge'04 – resource needs

- The CPU requirements is 10 times what was needed for the Data Challenge 2003.
- Data Challenge will be over 3 months
- Assumption: Gauss is twice slower than SICBMC
  - Not necessarily GEANT4 performance but still need to tune GAUSS to SICBMC level
- We expect that minimum of 50% of the CPU resources needed will be accessed via LCG1 Grid prototype

Future Data Challenges

- The 2005 data Challenge
  - mimicking LHC data taking,
  - testing the performance of High Level Trigger, reconstruction and streaming software
  - Evaluation of the feasibility of a full reconstruction at 200 Hz in the Online filter farm
- Early 2006, large statistics simulation/reconstruction for physics studies will be needed to further assess the Physics performance of LHCb

DIRAC – Distributed Infrastructure with Remote Agent Control

- Distributed MC production system for LHCb
  - Production tasks definition and steering
  - Software installation on production sites
  - Job scheduling and monitoring
  - Data transfers
- Automates most of the production tasks
  - minimum participation of local production managers
- *PULL* rather than *PUSH* concept for tasks scheduling
  - Different from the DataGRID architecture

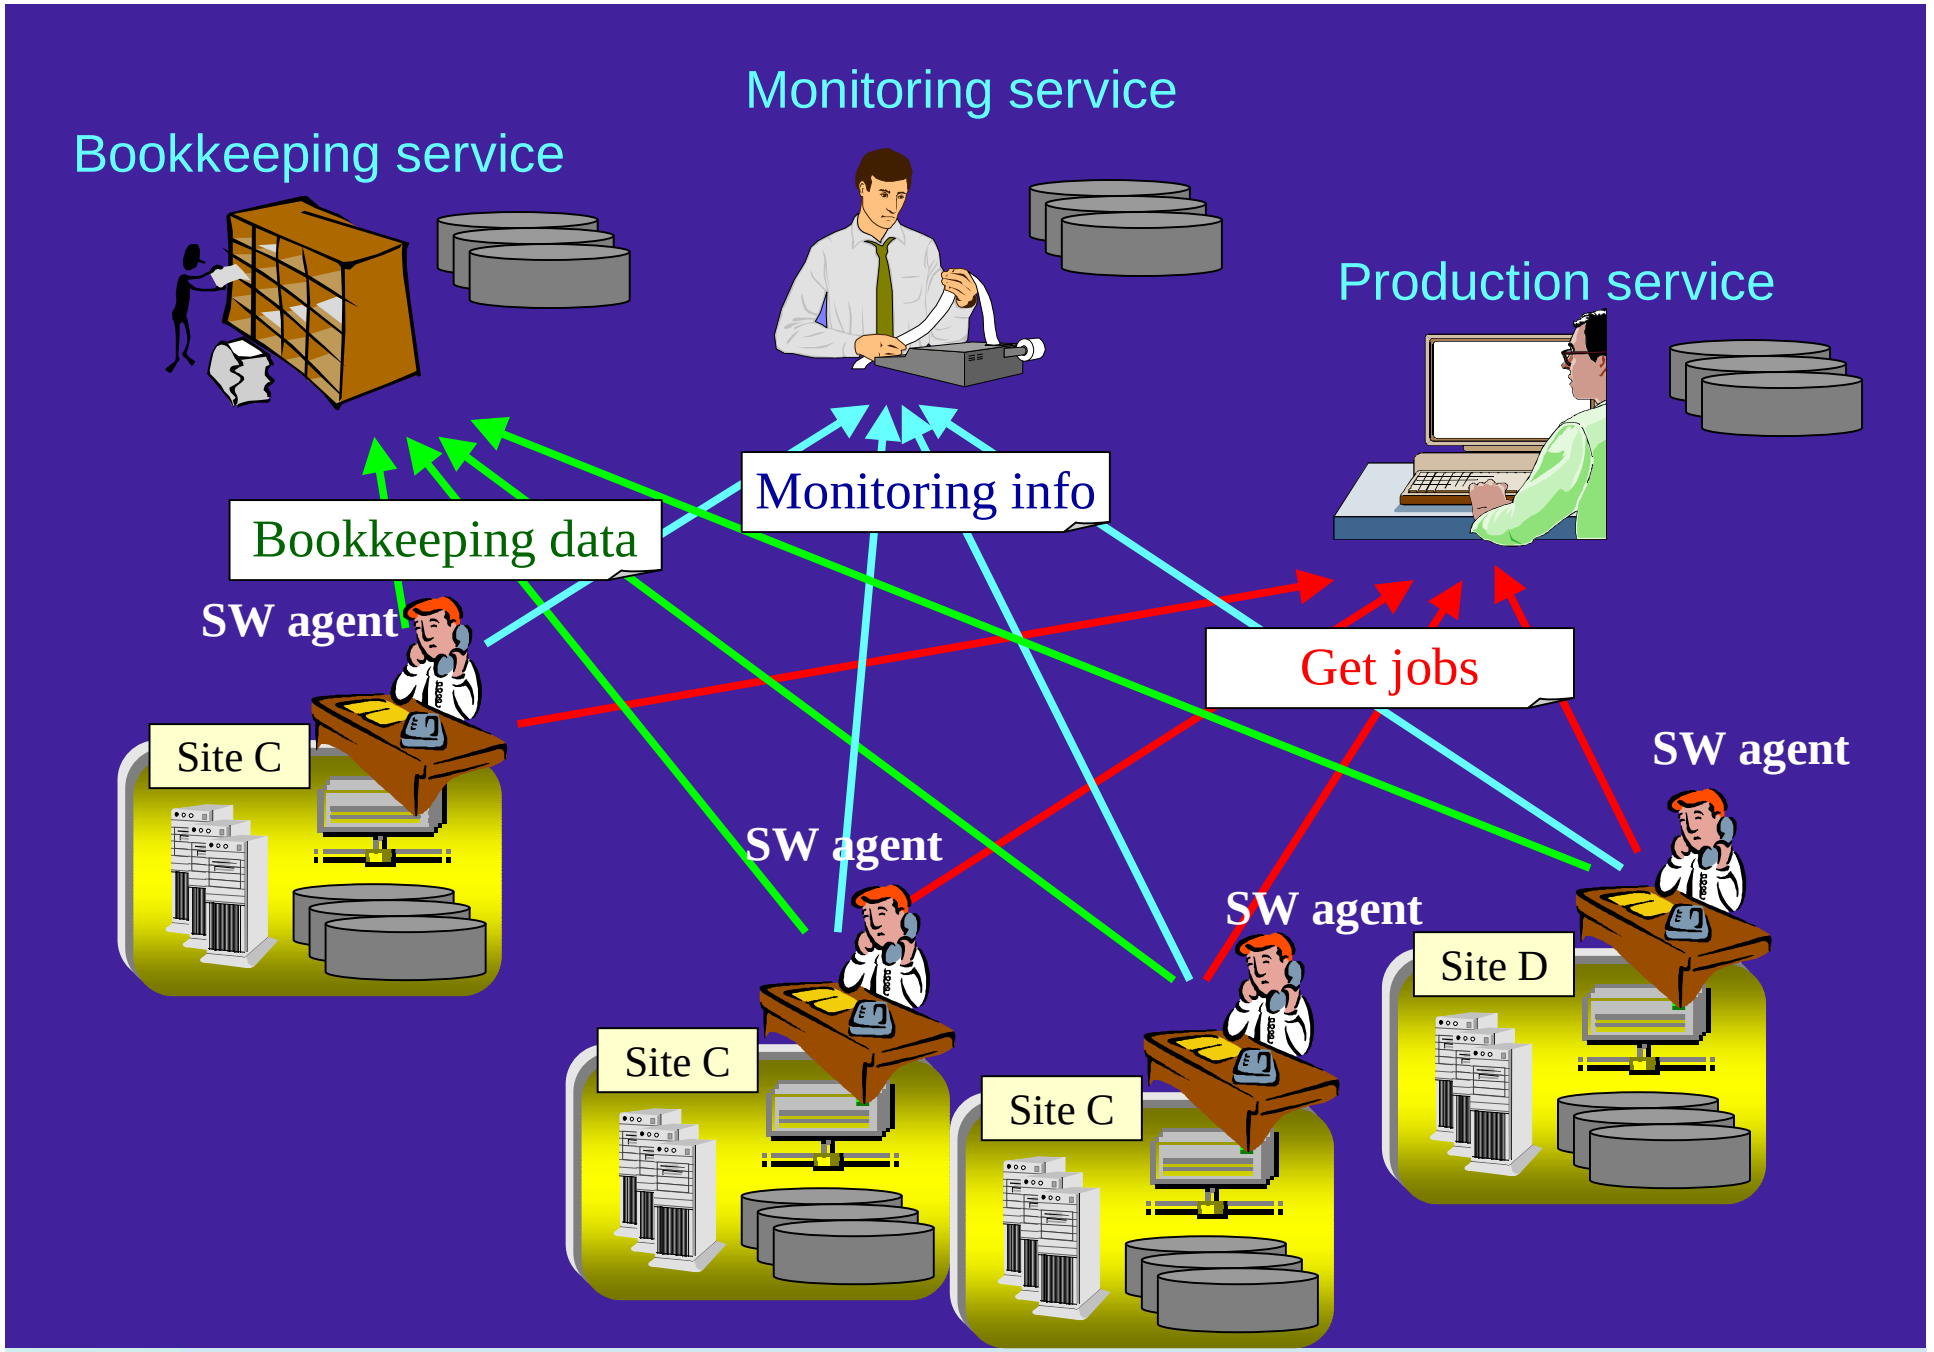
Monitoring service
Bookkeeping service
Production service
Monitoring info
Bookkeeping data
SW agent
Get jobs
Site C
SW agent
SW agent
SW agent
Site D
Site C
Site C

Job Description

Workflow description

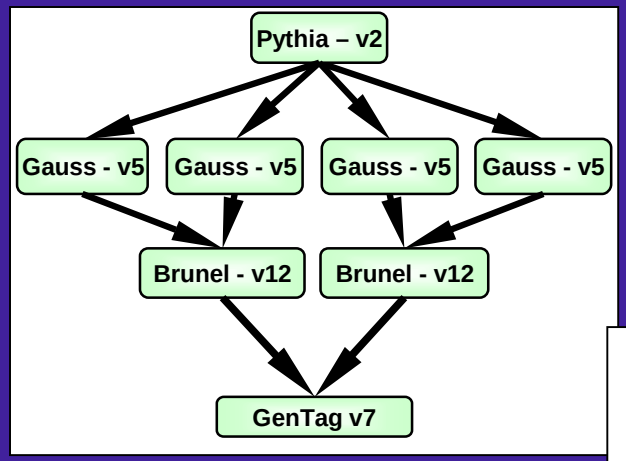

Production manager

Web based editors

Production run description

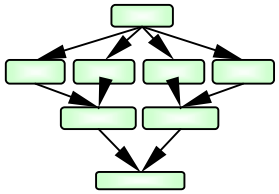

+

- Event type
- Application options
- Number of events
- Execution mode
- Destination site …

XML job descriptions

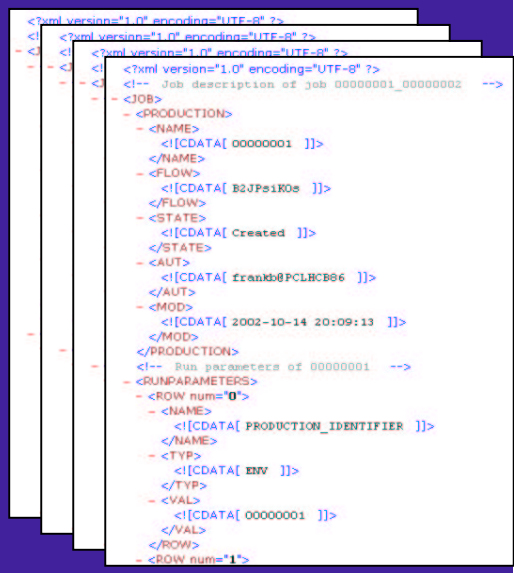

Production DB

Role of Agent

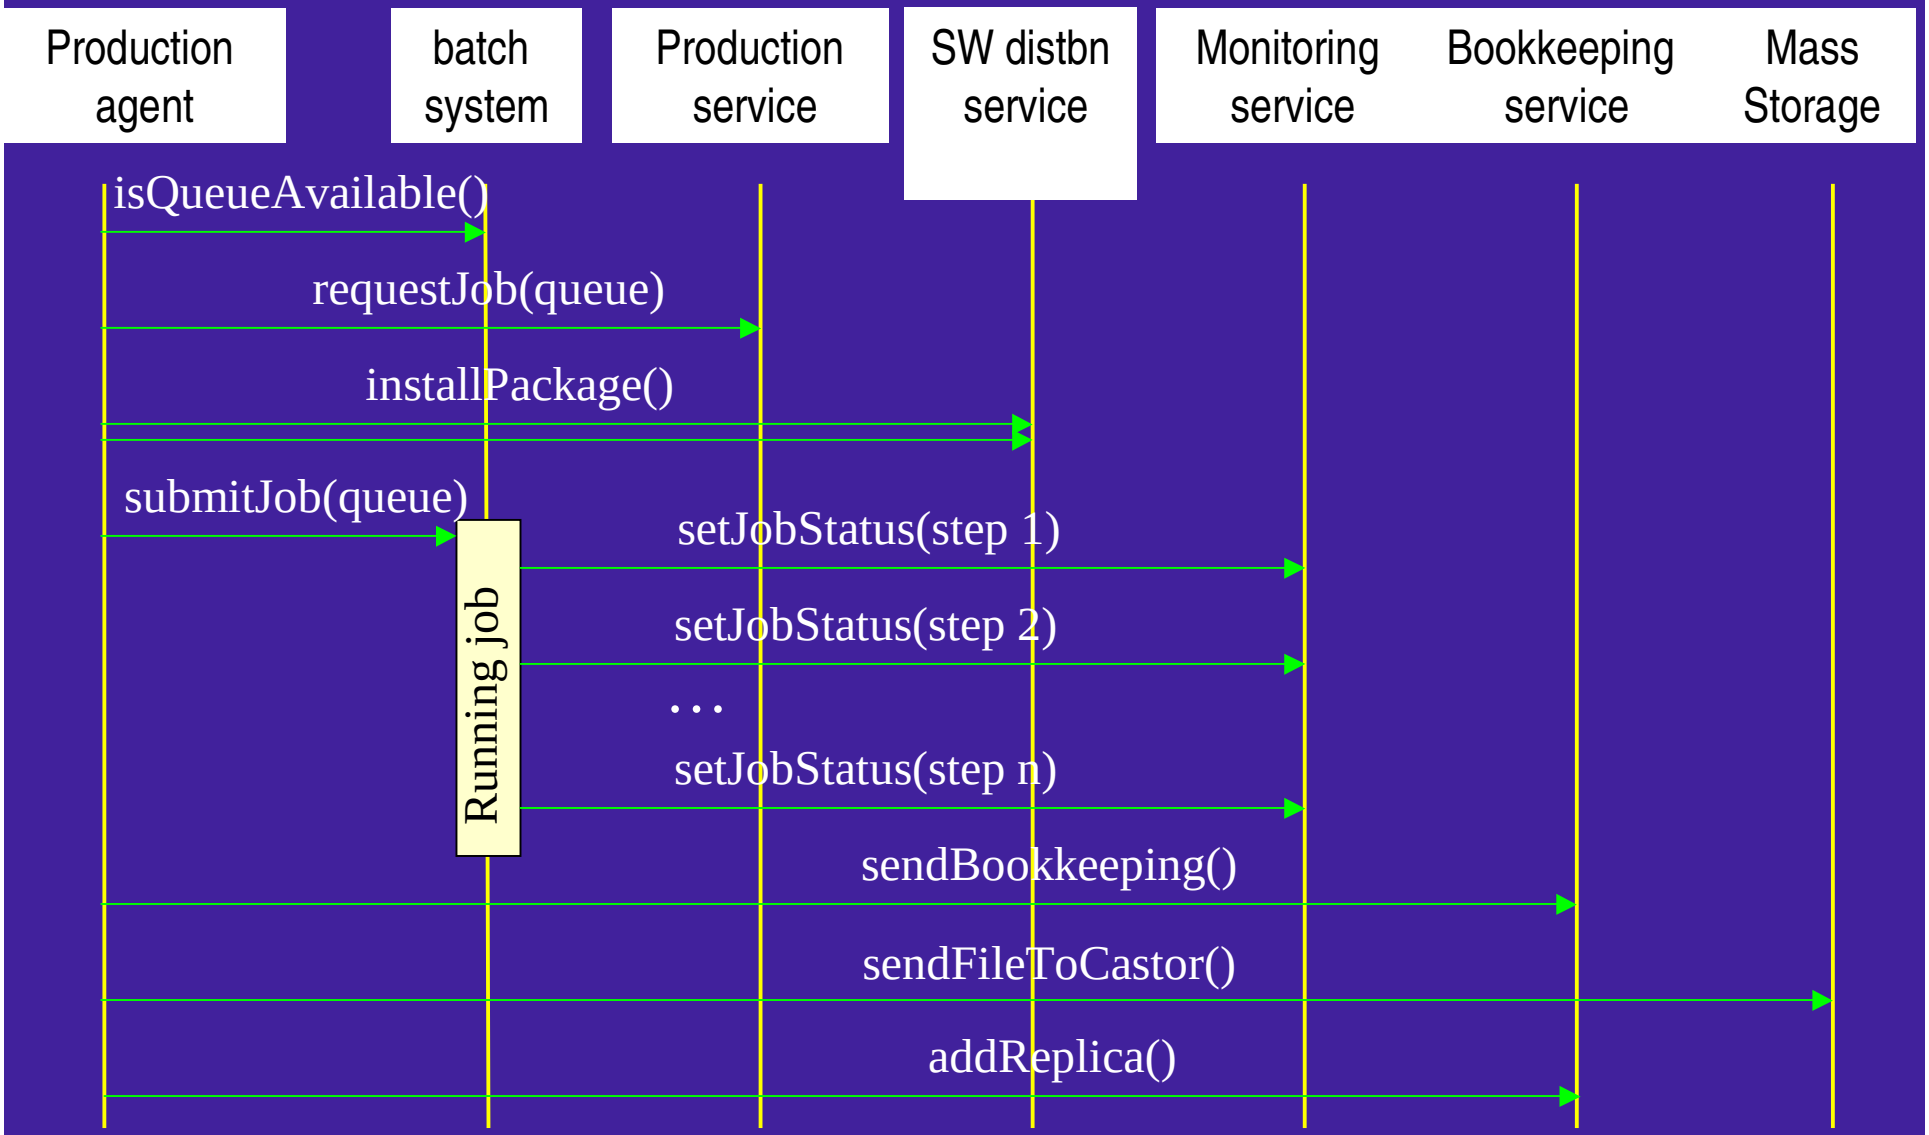

DIRAC & the DataGrid

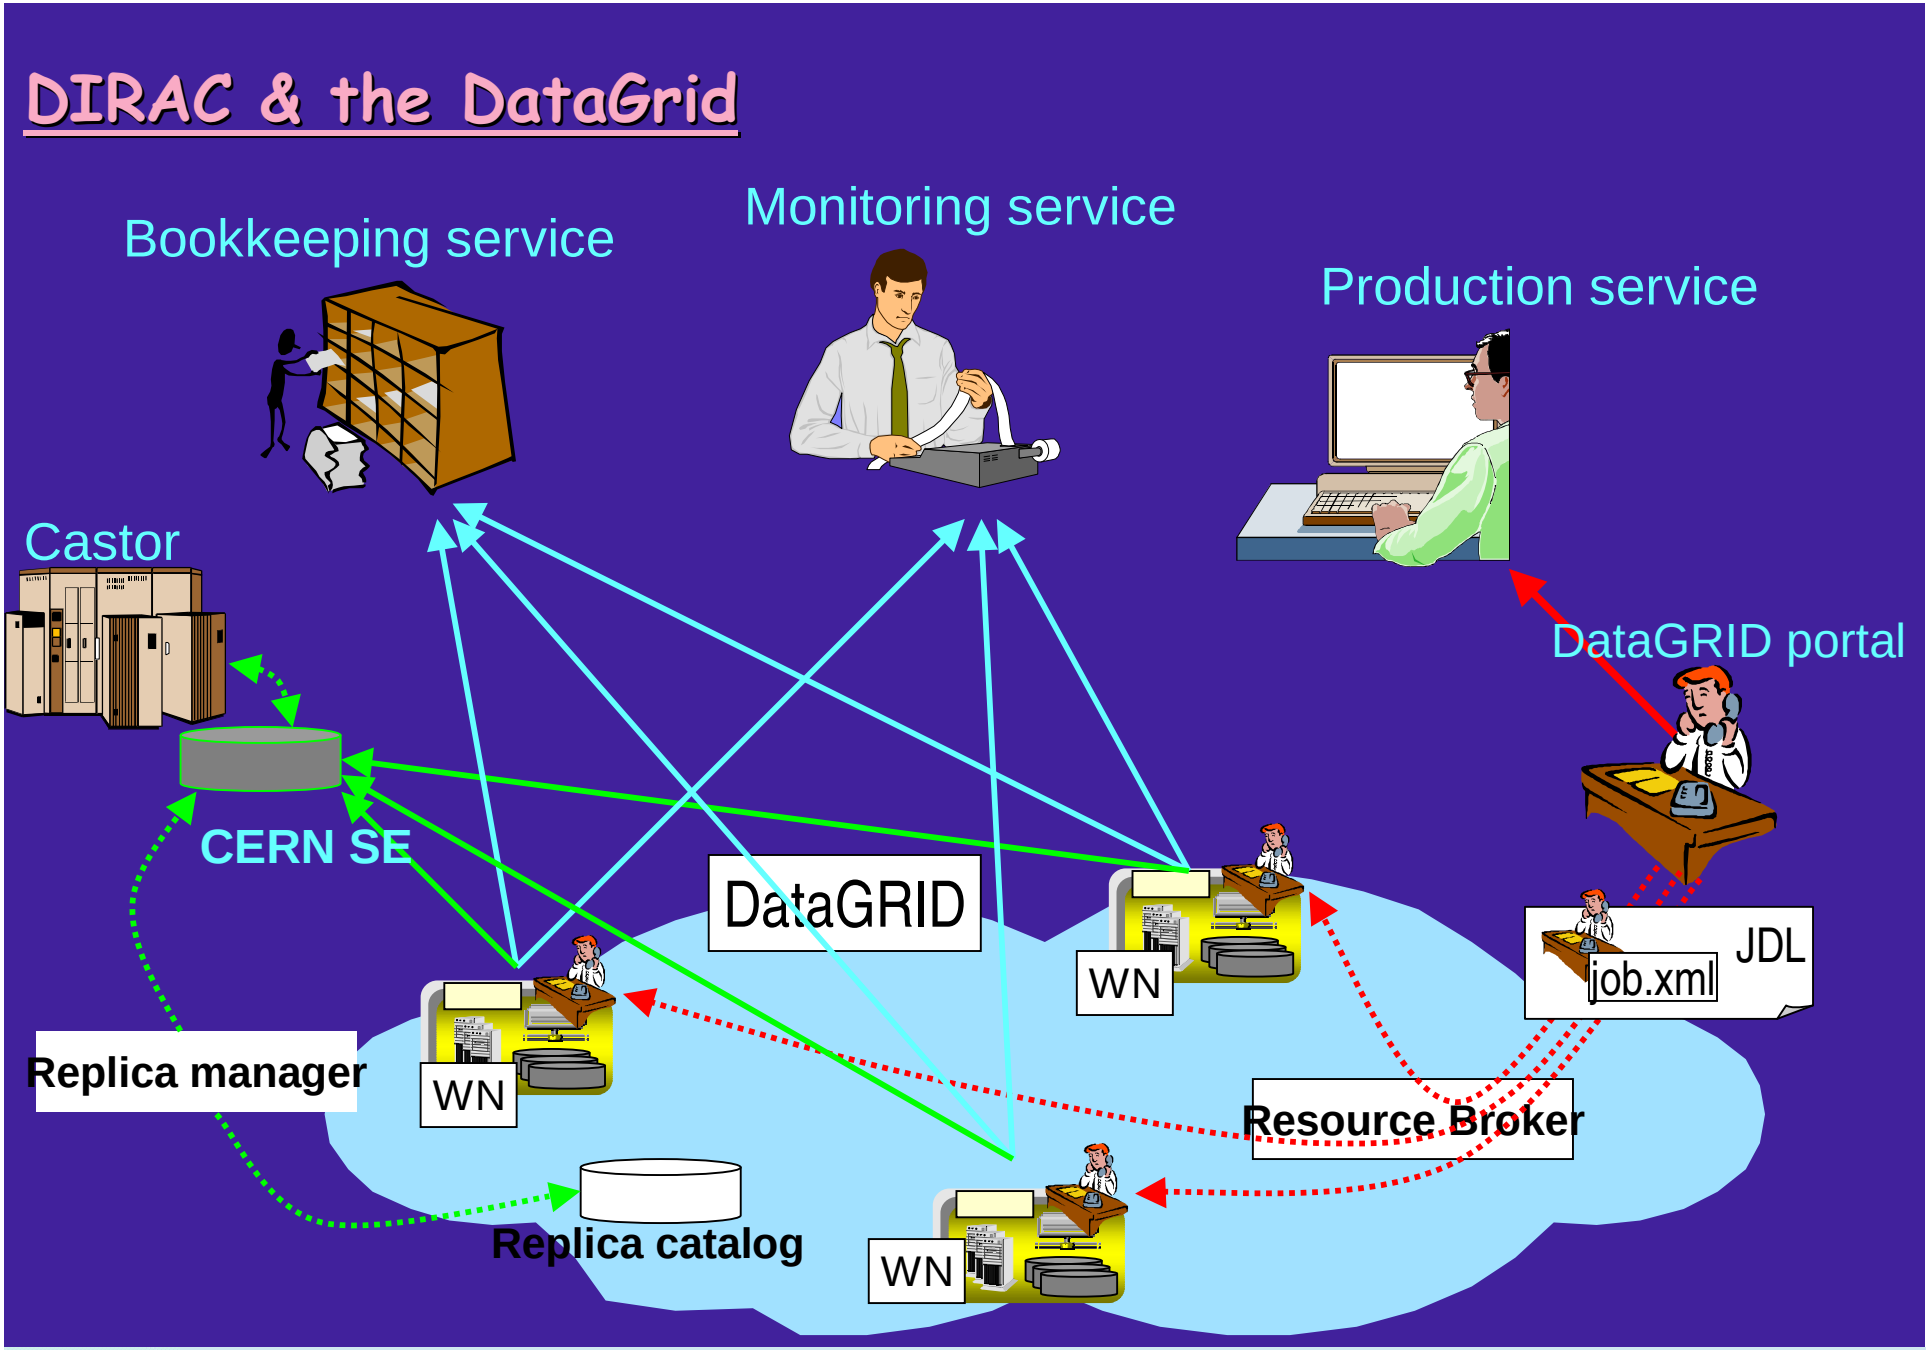

The "DataGrid" Experience

- Standard LHCb production jobs were used for the tests
  - Jobs of different statistics with 8 steps workflow.
- Jobs submitted to 4 EDG testbed Resource Brokers
  - keeping ~50 jobs per broker
- Software installed for each job

Job type (hours)	Total	Success	Success rate
Mini (0.2)	190	113	59%
Short (6)	171	102	59%
Medium (24)	1195	346	29%
Total	1556	561	36%

Total of ~300K events produced - EDG testbed already a "competitive" LHCb production site

Future of DIRAC

- Grid Services based architecture;
- Most of the services are coming from the LCG development;
- Some custom services tuned for the LHCb needs

Input to the ARDA RTAG - looking at distributed analysis

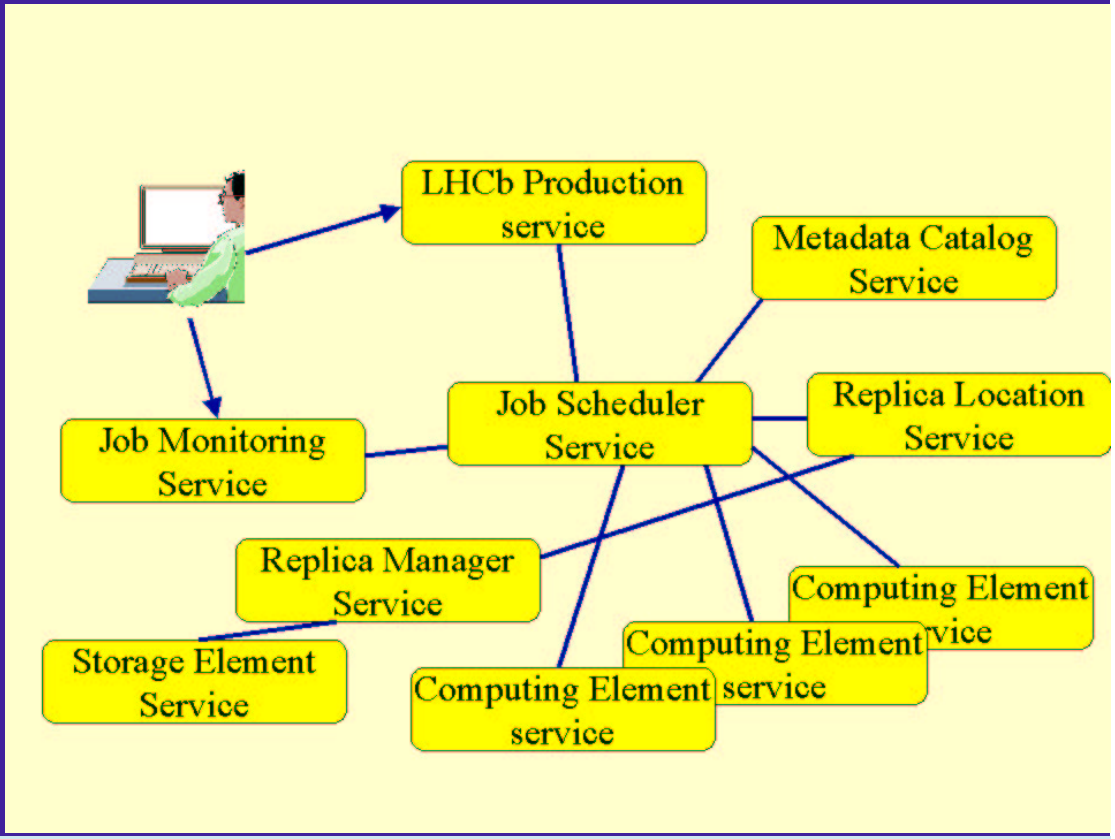

DIRAC development & needs

- Computing Element as a service with:
  - Information provider interface;
  - Job manipulation interface;
- Job Proxy service:
  - Allows to communicate with jobs running on WN's without IP connections to outside world.
- Grid wide File catalog:
  - Passing input and output sandboxes, log files, stdout/stderr, etc;
  - Grid File Access Library can be a solution.
    - Wraps variety of storage access protocols behind a file system like interface.

DIRAC development & needs

- Replica management service;
  - Managing datasets transfers, bookkeeping.
- LHCb VO production service:
  - LHCb users jobs executed on the grid should pass by the production service;
  - Accounting, policies, priorities, etc
- Scheduler service;
- Job monitoring service;
- …

Motivation and background

- ATLAS and LHCb develop applications within a common framework: Gaudi/Athena
- Both collaborations aim to exploit potential of Grid for large-scale, data-intensive distributed computing

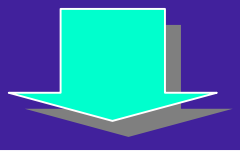

- Simplify management of analysis and production jobs for end-user physicists by developing tools for accessing Grid services with built-in knowledge of how Gaudi/Athena works: **Gaudi/Athena and Grid Alliance (GANGA)**
- Generic tools interfacing to the Grid can be used at other experiments.

Schematic on GANGA-Gaudi-Grid services

Existing tools will be used wherever possible

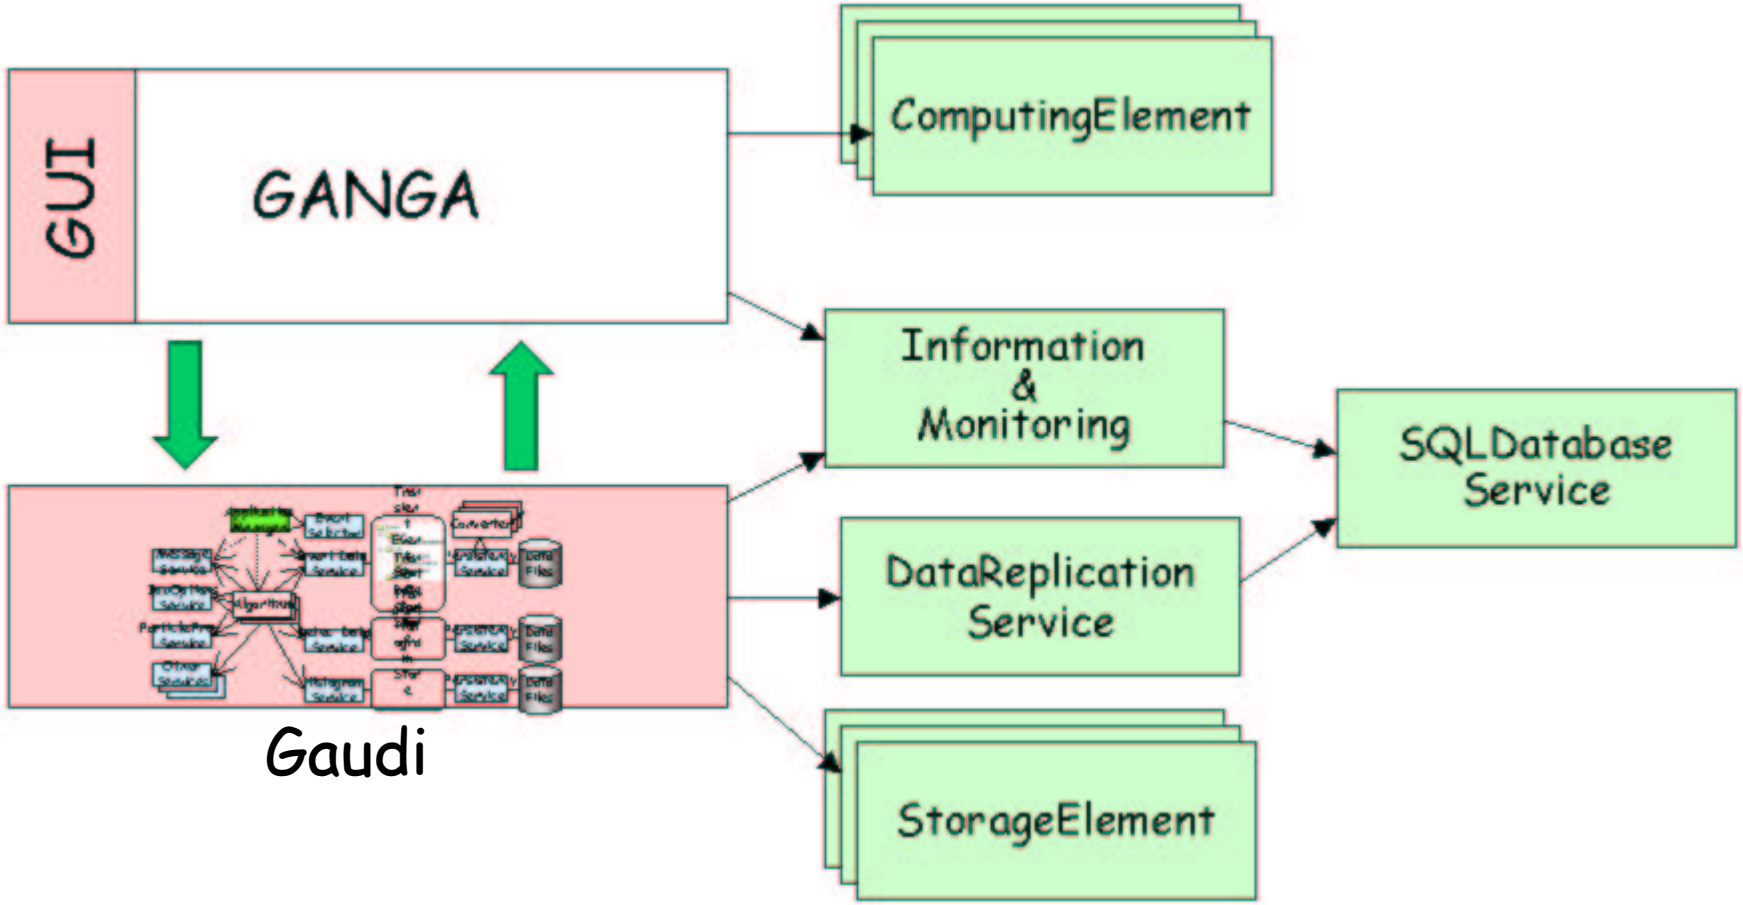

Design Issues

GANGA accessible through CLI and GUI via PYTHON s/w bus

3 GANGA components:

- general functionality
- Gaudi specific
- External components

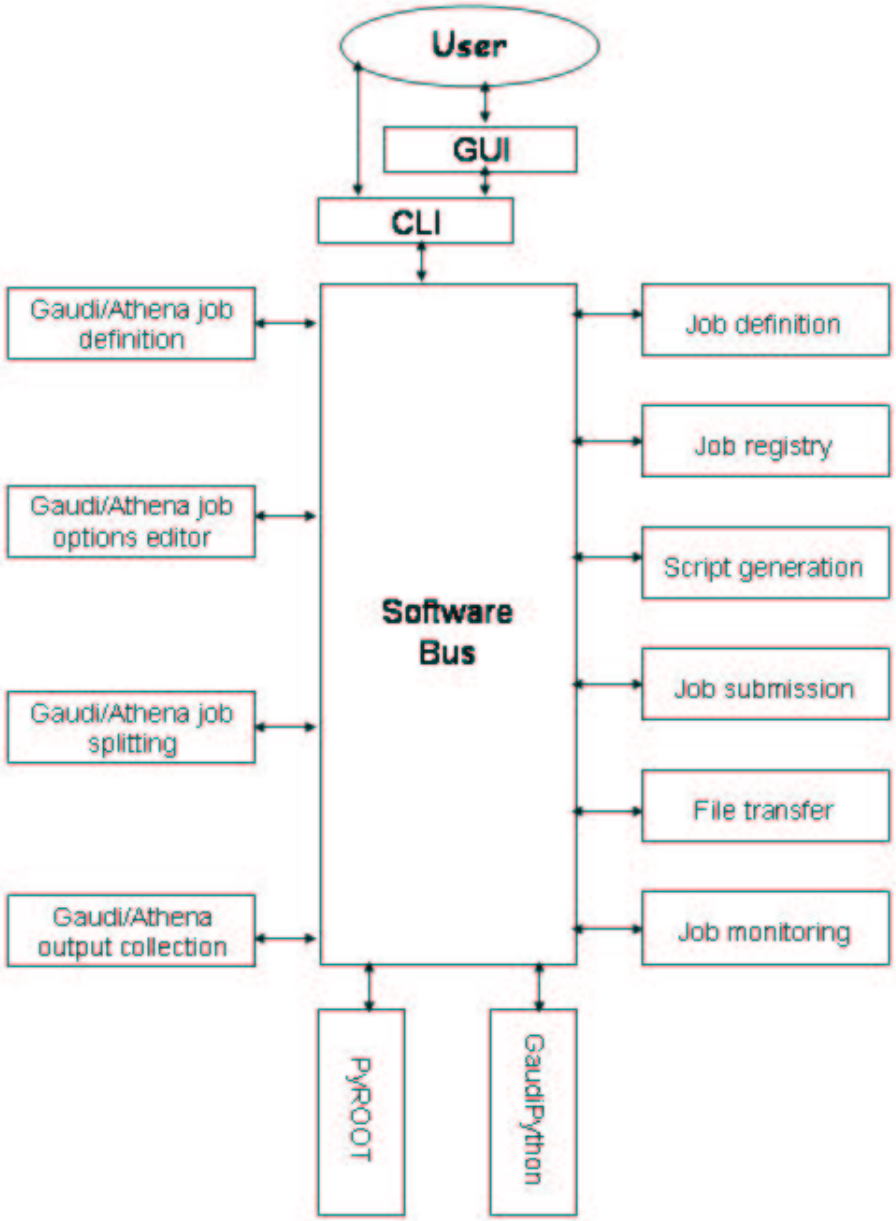

GUI Snapshot

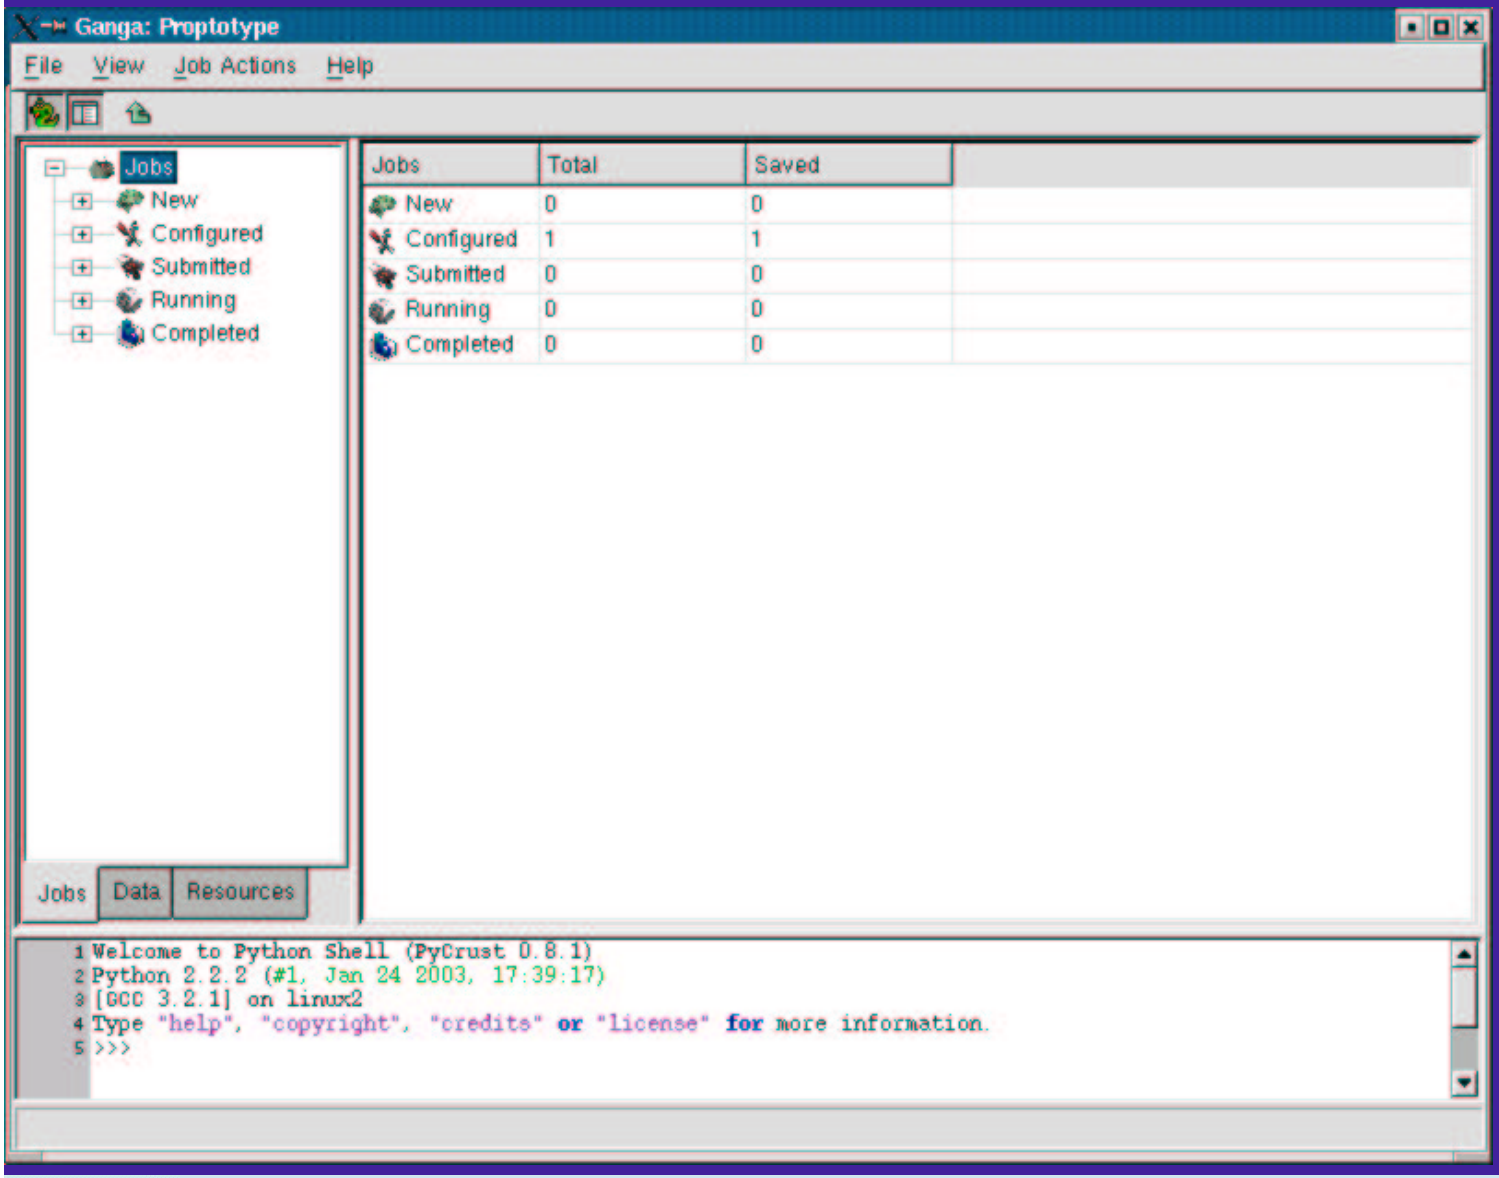

Short term plans

- GUI & CLI for working with DaVinci analysis job
- GANGA to workflow definition mechanism of its own
- Job-options editor
- Splitting cloning procedure
- Catalogue of job and status
- Pull info from jobs – allowing automatic updates of status
- Single procedure for submitting jobs to different "batch" systems (EDG, LSF, PBS, …) from local machine

Long term plans

- Integrate production jobs into GANGA; interface with DIRAC
- More sophisticated workflow definition e.g. in context of DIRAC
- Job-option editor with sensible defaults
- More sophisticated Splitting cloning procedure
- Increased functionality for Catalogue of job and status
- Jobs monitroing based on "push" system; integrate with NetLogger for Grid
- Introduce concept of "gatekeeper" to allow job submission to any machine where users has access; software installation

DIRAC architecture with GANGA user interface

- GANGA should be useful for defining and executing complex production and analysis tasks:
  - Job configuration
  - Job splitting and submission
  - Job monitoring
  - Checking the results

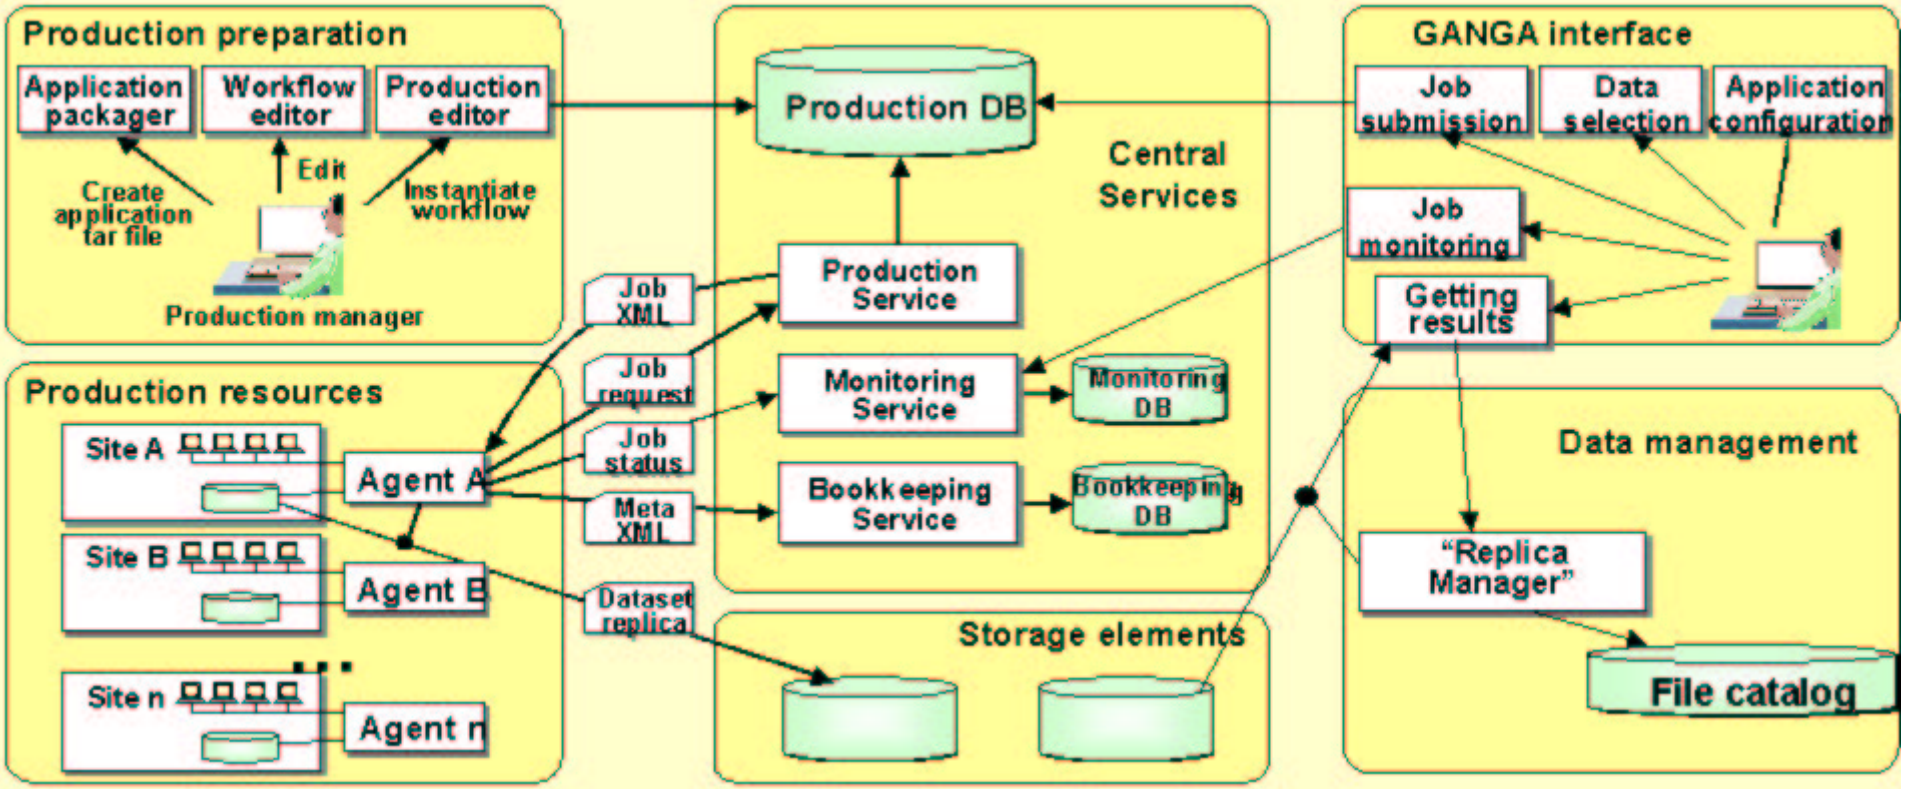

Summary

- Expt s/w framework is in a mature state
  - Iterative & incremental releases
  - All LHCb event data processing applications based in on it
  - Currently integrating POOL, SEAL, …
  - New developments: conditions DB, data management, …
- The data processing Gaudi-based applications are in good status
- DIRAC production system established and used very successfully
  - Integrated into EDG testbed
  - Needs to be developed & implemented into LCG
- Prototype GANGA exists now
  - Development and integration into DIRAC