LHCb Computing

Nick Brook

- Organisation
- LHCb software
- Distributed Computing
- Computing Model
- LHCb & LCG
- Milestones

Organisation

Software framework & distributed computing

- provision of the software framework
  - Core s/w, conditions DB, s/w engineering, …
- tools for distributed computing
  - Production system, user analysis interface, …

Computing Resource

- coordination of the computing resources
- organisation of the event processing of both real and simulated data

Physics Applications

- integration of algorithms (both global and sub-system specific) in the software framework
- global reconstruction algorithms that will run in the online & offline environment
- coordination of the sub-detector software

LHCb software framework

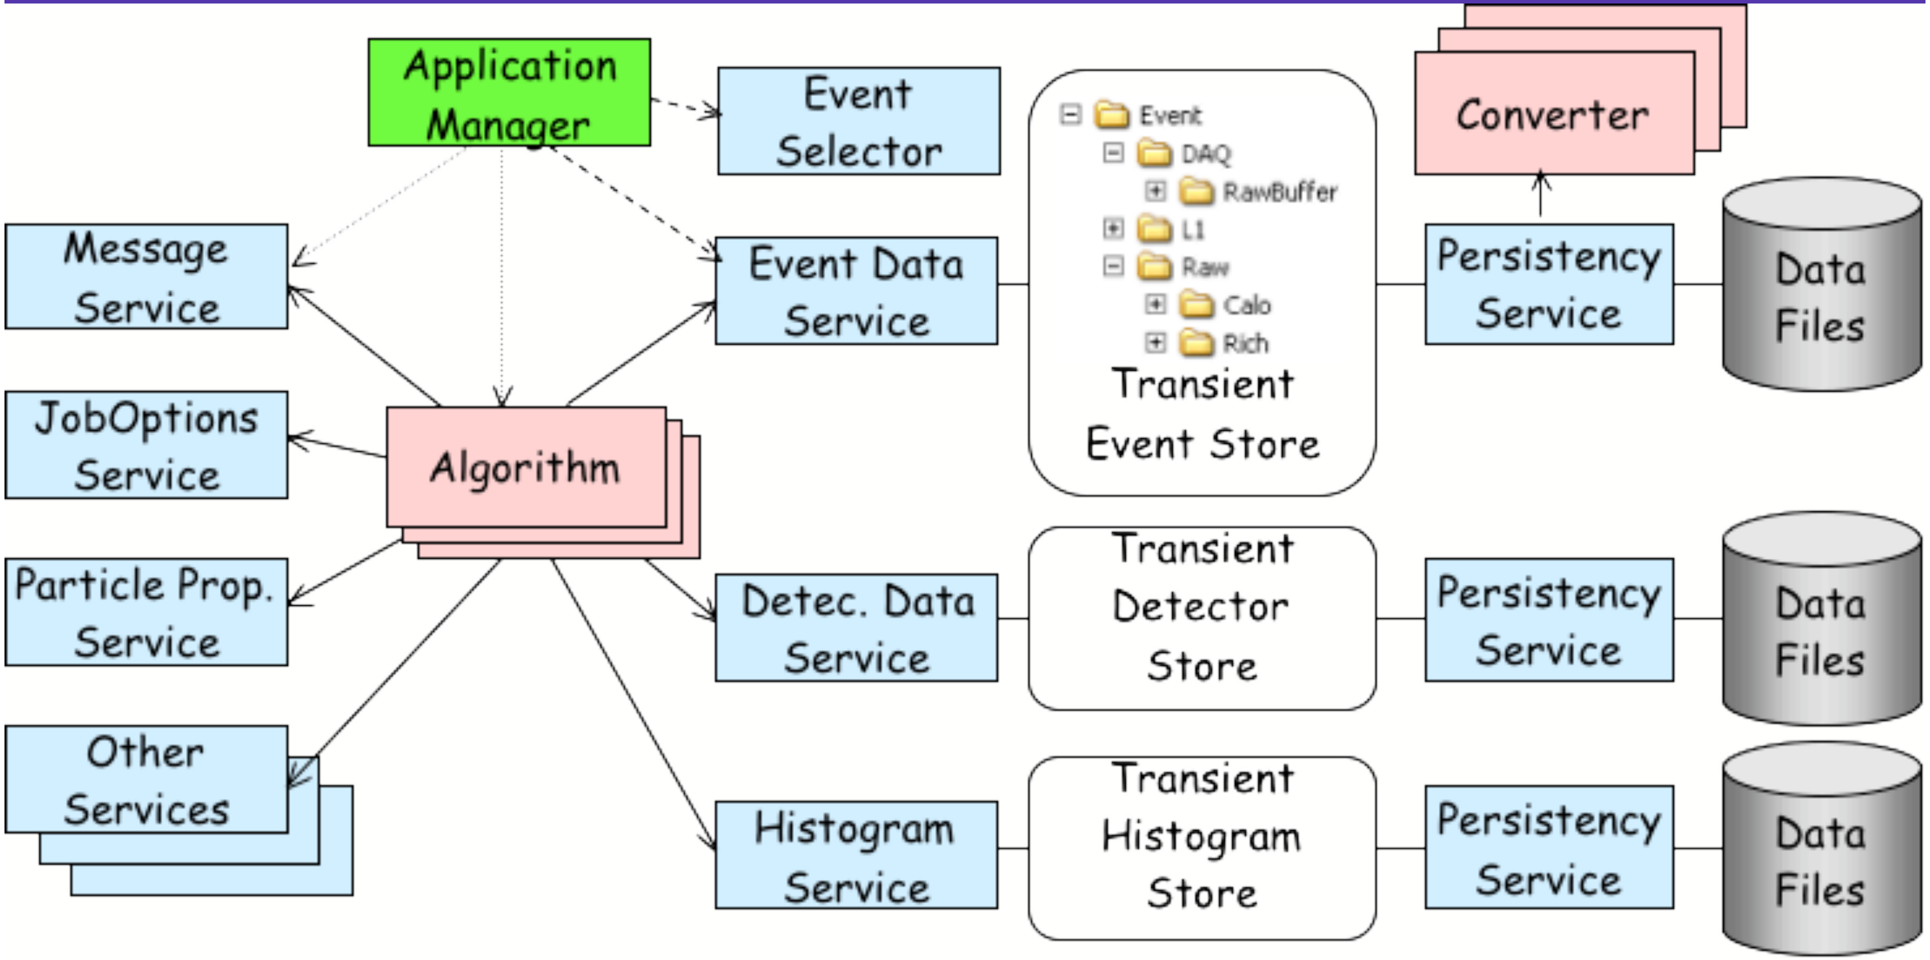

Object diagram of the Gaudi architecture

LHCb software framework

- Gaudi is architecture-centric, requirements-driven framework
  - Adopted by ATLAS; used by GLAST & HARP
  - Same framework used both online & offline
- algorithmic part of data processing as a set of OO objects
  - decoupling between the objects describing the data and the algorithms allows programmers to concentrate separately on both.
  - allows a longer stability for the data objects (the LHCb event model) as algorithms evolve much more rapidly
- An important design choice has been to distinguish between a transient and a persistent representation of the data objects
  - changed from ZEBRA to ROOT to LCG POOL without the algorithms being affected.
- Event Model classes only contain enough basic internal functionality for giving algorithms access to their content and derived information
  - Algorithms and tools perform the actual data transformations

LHCb software

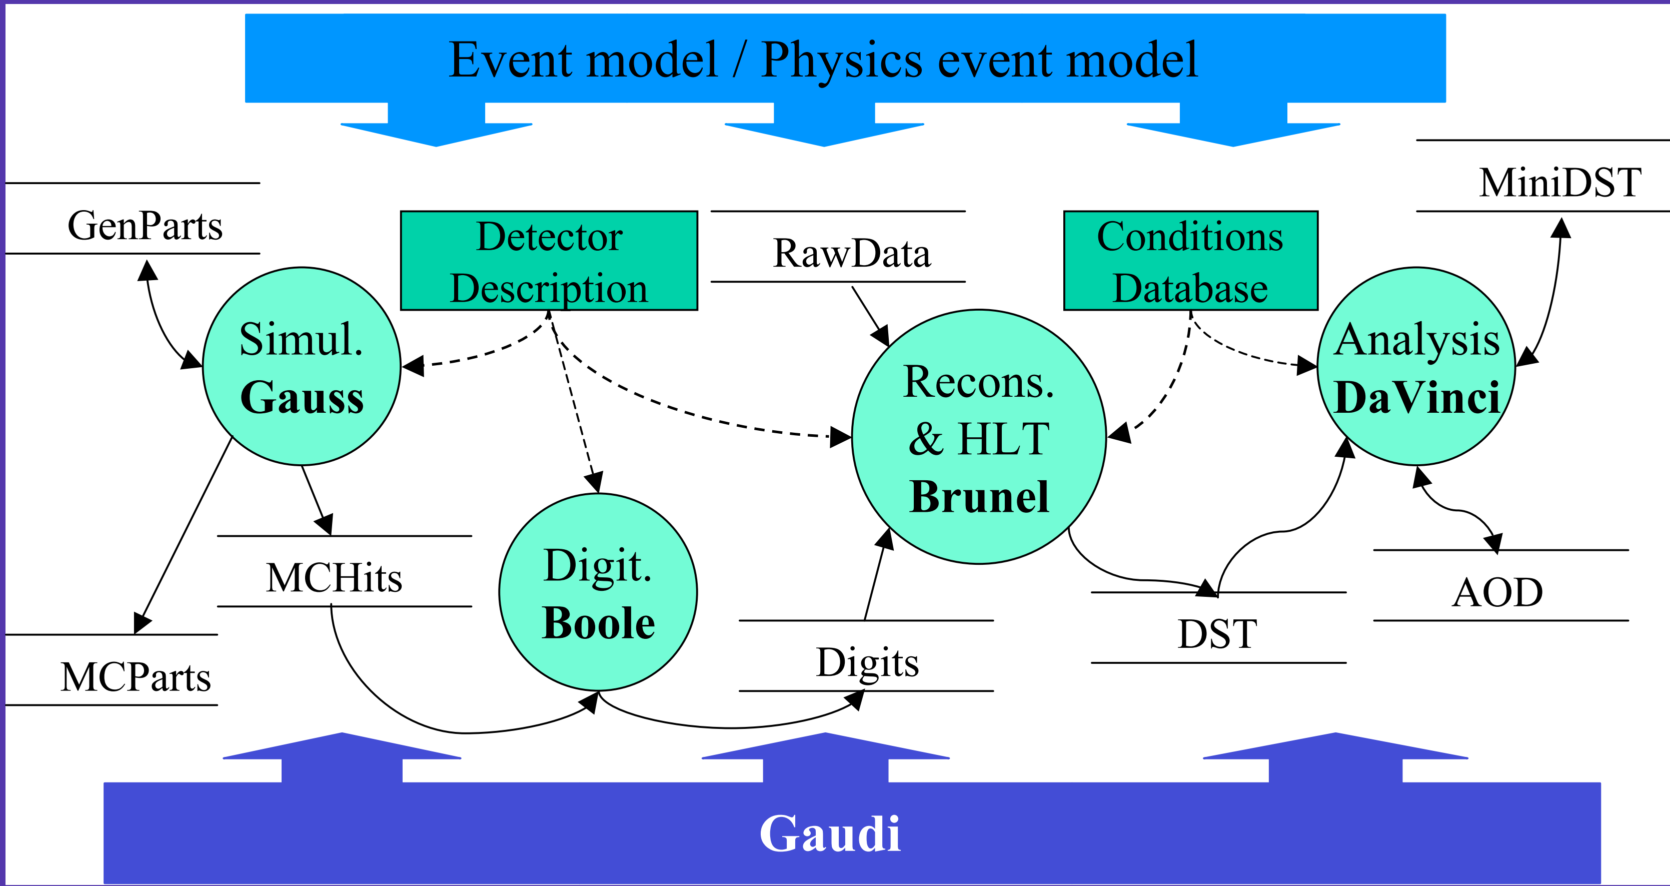

LHCb data processing applications and data flow

LHCb software

- Each application is a producer and/or consumer of data for the other applications
- The applications are all based on the Gaudi framework
  - communicate via the LHCb Event model and make use of the LHCb unique Detector Description
  - ensures consistency between the applications and allows algorithms to migrate from one application to another as necessary
- subdivision between the different applications has been driven by their different scopes as well as CPU consumption and repetitiveness of the tasks performed

Event sizes & processing requirements

	Aim	Current
Event Size	**kB**	
RAW	25	35
rDST	25	8
DST	75	58
Event processing	**kSI2k.s per event**	
Reconstruction	2.4	2.7
Stripping	0.2	0.6
Analysis	0.3	??
Simulation (bb-incl)	50	50

Conditions DB

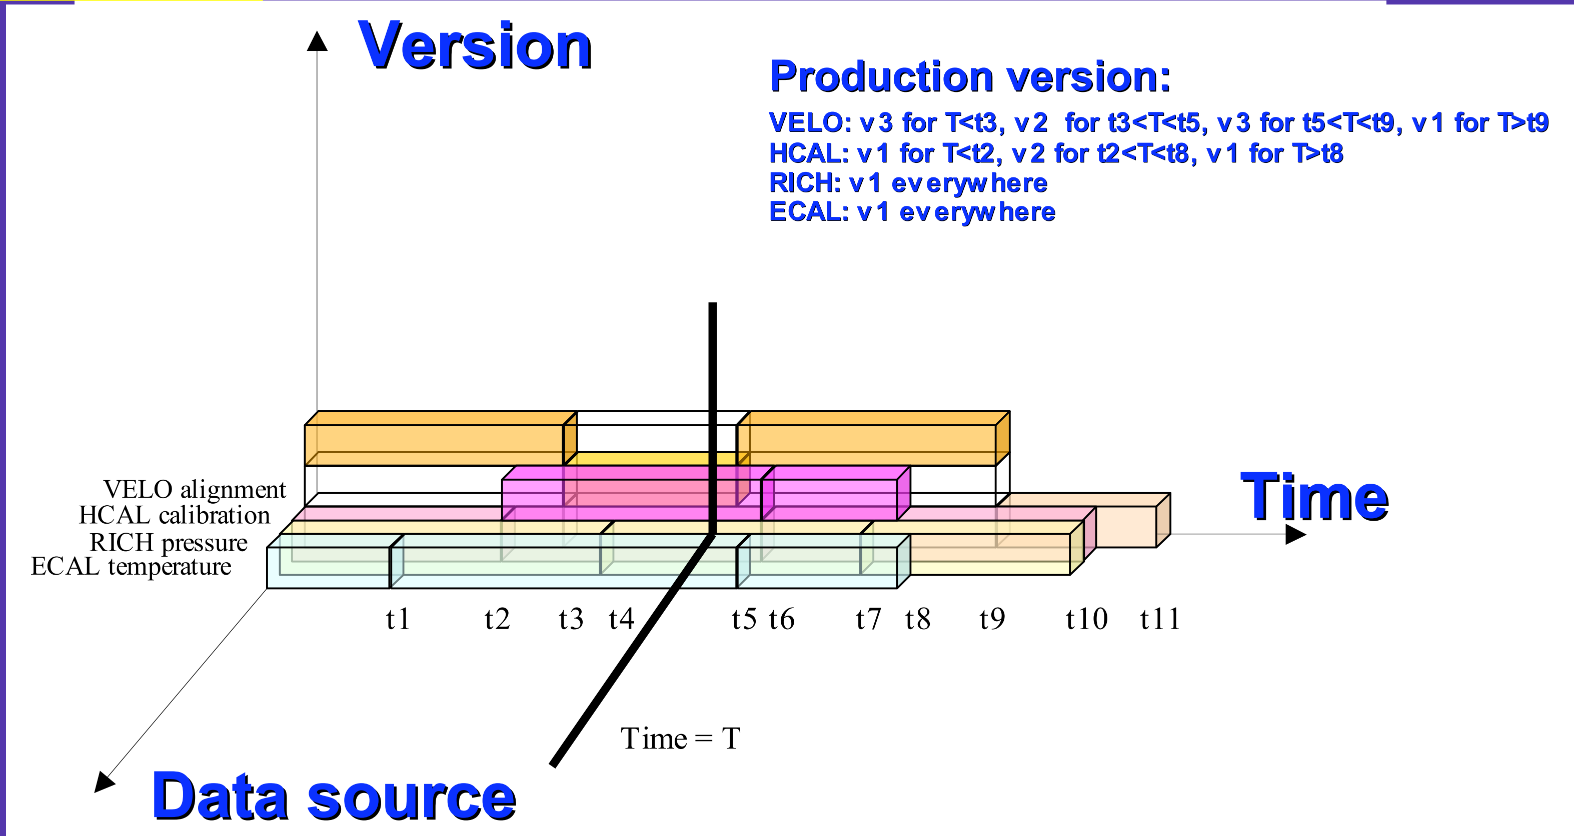

Tools and framework to deal with conditions DB and non-perfect detector geometry is in place

LCG COOL project is providing the underlying infrastructure for conditions DB

Distributed computing - production with DIRAC

DIRAC uses the paradigm of a Services Oriented Architecture (SOA).

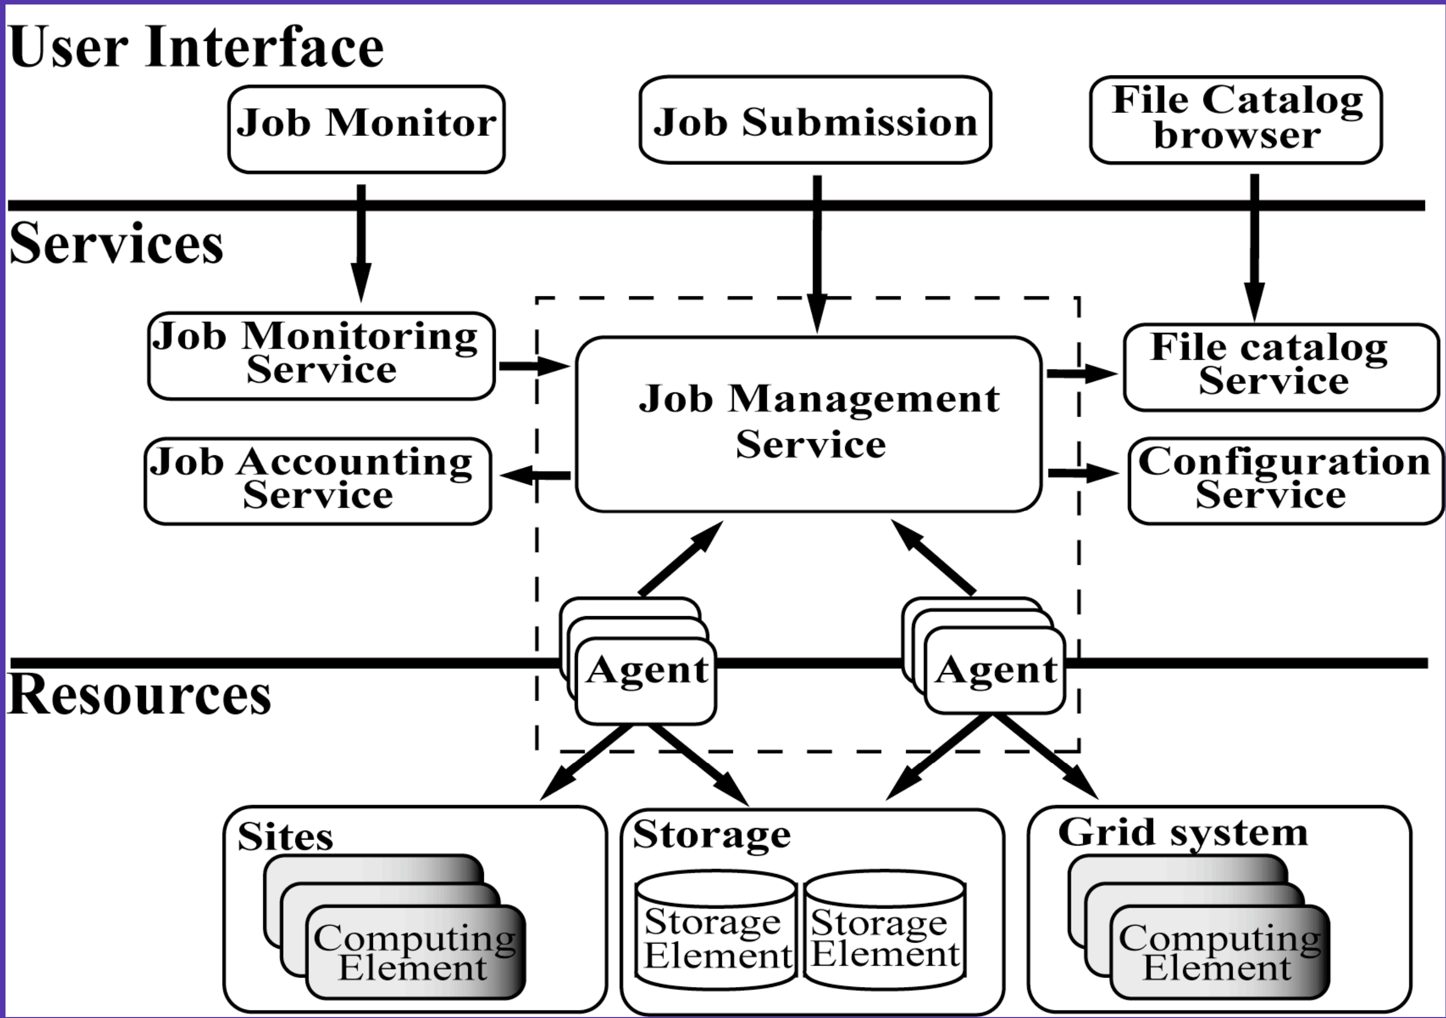

Distributed computing - production with DIRAC

- The DIRAC overlay network paradigm is first of all there to abstract heterogeneous resources and present them as single pool to a user :
  - LCG or DIRAC sites or individual PC's (or other Grids)
  - Single central Task Queue is foreseen both for production and user analysis jobs
- The overlay network is dynamically established
  - No user workload is sent until the verified LHCb environment is in place

GANGA - user interface to the Grid

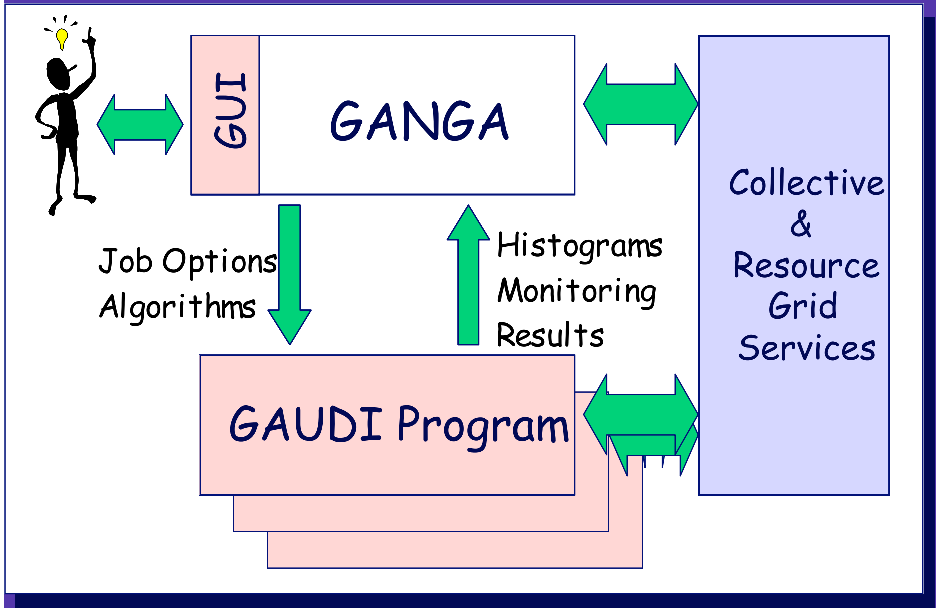

- Goal
  - Simplify the management of analysis for end-user physicists by developing a tool for accessing Grid services with built-in knowledge of how Gaudi works
- Required user functionality
  - Job preparation and configuration
  - Job submission, monitoring and control
  - Resource browsing, booking, etc.
- Done in collaboration with ATLAS
- Use Grid middleware services
  - interface to the Grid via Dirac and create synergy between the two projects

Computing Model

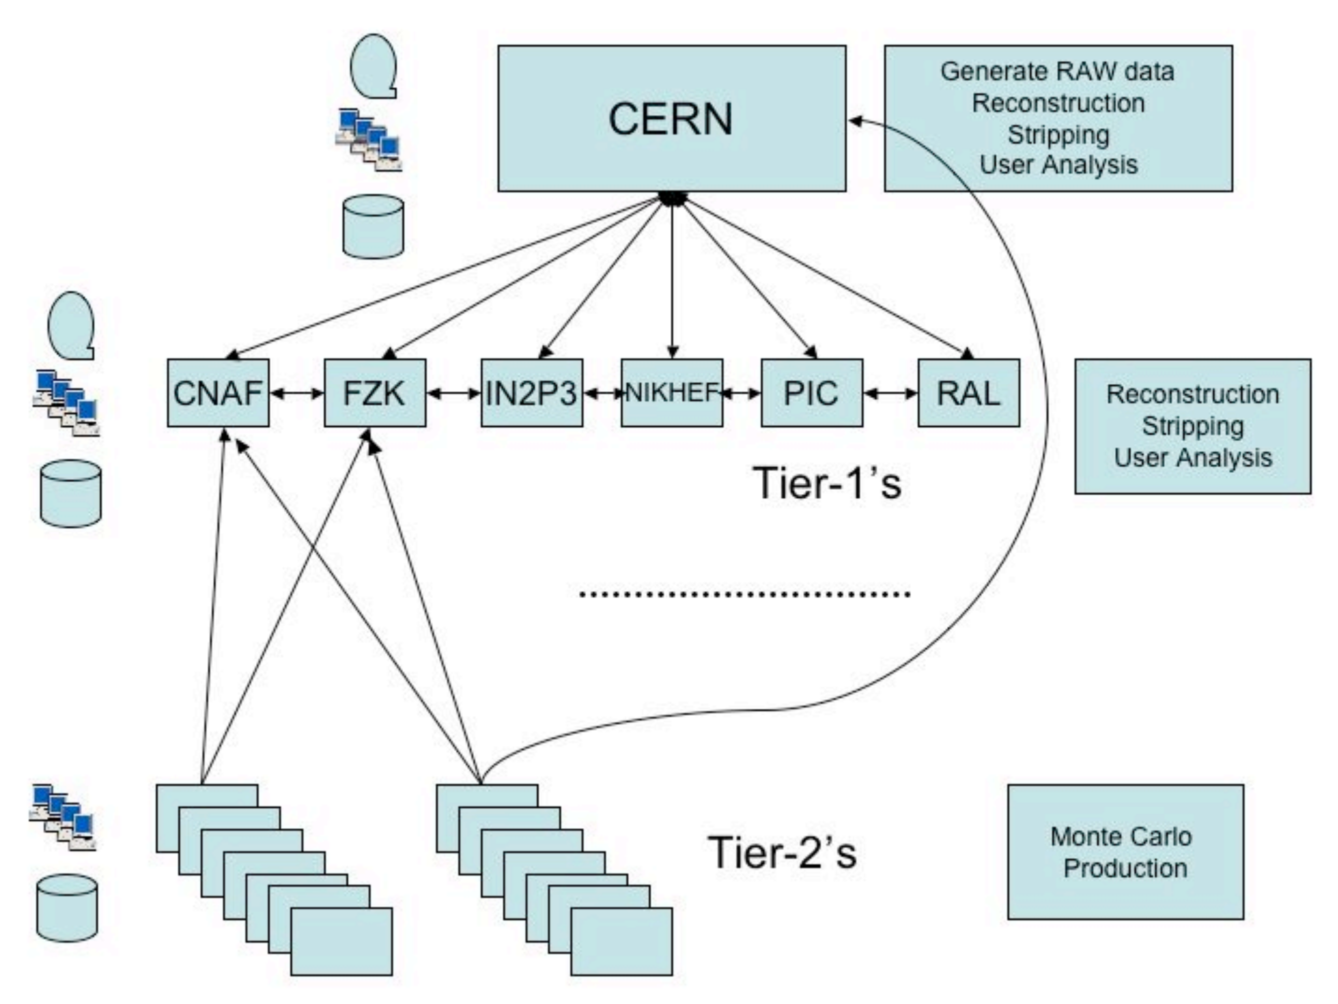

CERN Tier-1 centre will be essential for accessing the "hot stream" data to:

i. First alignment & calibration

ii. First high level analysis

Computing Model - resource summary

Nos. of CPUs (2.4GHz PIV)	2006	2007	2008	2009	2010
CERN	312	624	1040	1445	2173
Tier-1's	1537	3063	5109	6416	9653
Tier-2's	2647	5306	8843	8843	8843
Total	**4497**	**8994**	**14994**	**16705**	**20670**

Computing Model - resource profiles

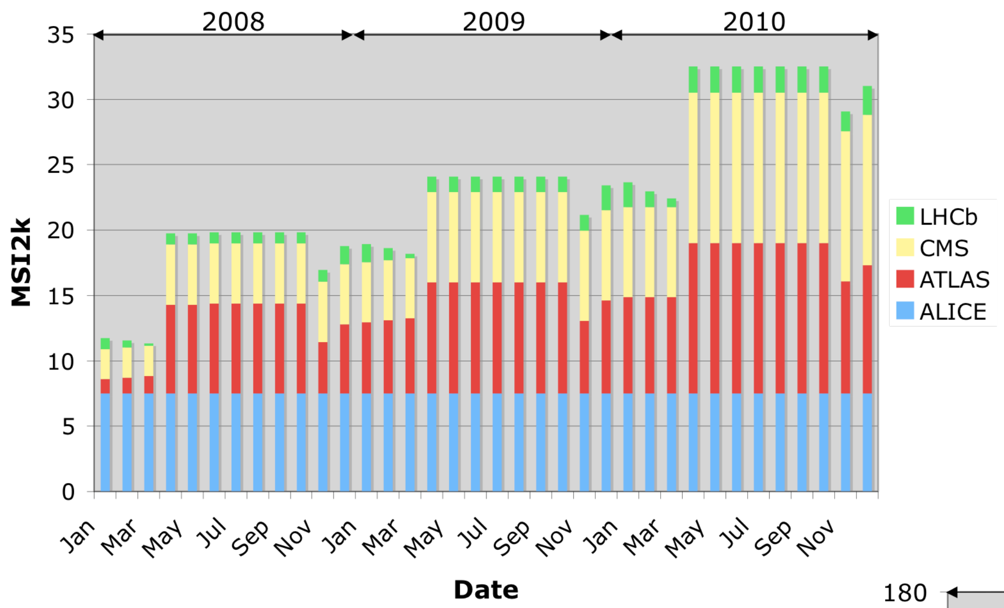

CERN CPU

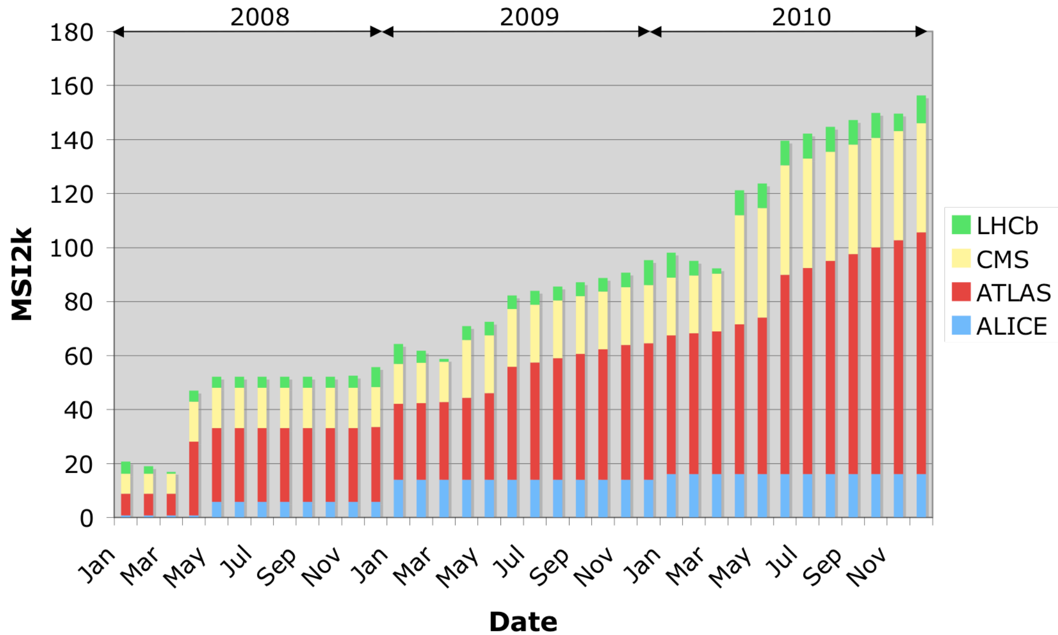

Tier-1 CPU

Computing Model - resource summary

Disk(TB)					
CERN	248	496	826	1095	1363
Tier-1's	730	1459	2432	2897	3363
Tier-2's	7	14	23	23	23
Total	**984**	**1969**	**3281**	**4015**	**4749**
MSS (TB)					
CERN	408	825	1359	2857	4566
Tier-1's	622	1244	2074	4285	7066
Total	**1030**	**2069**	**3433**	**7144**	**11632**

LHCb & LCG

- DC04 (May-August 2004)
  - 187 Mevts simulated and reconstructed
  - 61 Tbytes of data produced
  - 43 LCG sites used
  - 50% using LCG resources (61% efficiency)
- DC04v2 (December 2004)
  - 100 Mevts simulated and reconstructed
- DC04 stripping
  - Helped in debugging CASTOR-SRM functionality
  - CASTOR-SRM now functional (at CERN, CNAF, PIC)
- RTTC production (May 2005)
  - 200 Mevts simulated (minimum bias) in 3 weeks (up to 5500 jobs simultaneously)

LHCb & LCG - Data Challenge 2004

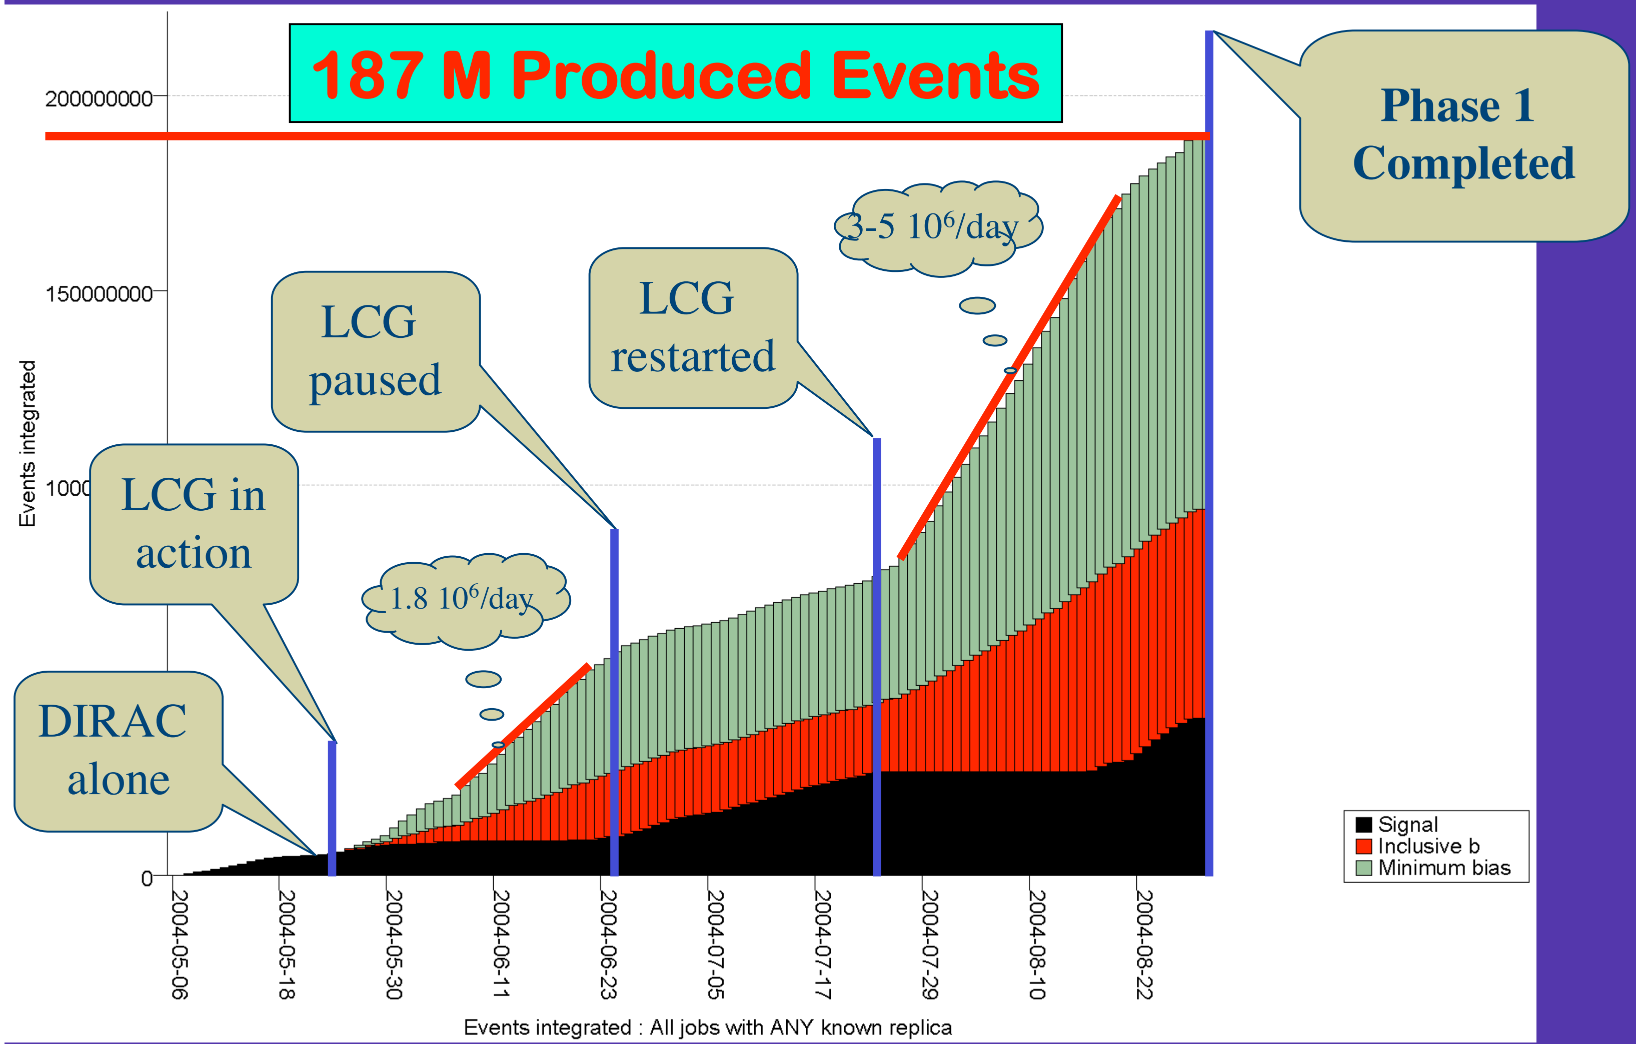

DC04 production

424 CPU years

Events: 185.55 M

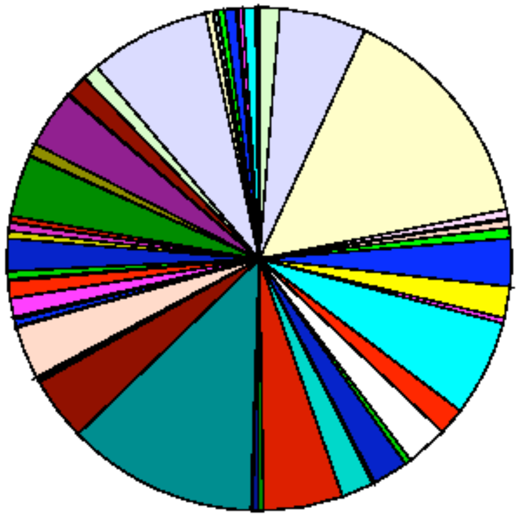

DIRAC.Barcelona.es	1.305%
DIRAC.Bologna.it	5.560%
DIRAC.CERN.ch	14.96%
DIRAC.CracowAgu.pl	0.532%
DIRAC.IF-UFRJ.br	0.124%
DIRAC.IHEP-Protvino.ru	0.504%
DIRAC.IHEP2-Protvino.ru	0.691%
DIRAC.ITEP-Moscow.ru	3.066%
DIRAC.Imperial.uk	2.017%
DIRAC.JINR-Dubna.ru	0.353%
DIRAC.Karlsruhe.de	6.181%
DIRAC.LHCBONLINE.ch	1.752%
DIRAC.Liverpool.uk	2.674%
DIRAC.Lpool.uk	0.405%
DIRAC.Lyon.fr	2.273%
DIRAC.Manno.ch	0.035%
DIRAC.Oxford.uk	0.137%
DIRAC.Santiago.es	2.053%
DIRAC.ScotGrid.uk	5.078%
DIRAC.Zurich.ch	0.355%
LCG.BHAM-HEP.uk	0.341%
LCG.Barcelona.es	0.106%
LCG.Bari.it	0.010%
LCG.CERN.ch	12.22%
LCG.CNAF.it	4.097%
LCG.Cagliari.it	0.049%
LCG.Cambridge.uk	0.128%
LCG.Carleton.ca	0.146%
LCG.Catania.it	0.031%
LCG.FNAL.us	0.017%
LCG.FZK.de	3.375%
LCG.Ferrara.it	0.094%
LCG.IN2P3.fr	0.419%
LCG.ITEP.ru	0.176%
LCG.Imperial.uk	1.188%
LCG.JINR.ru	0.021%
LCG.KFKI.hu	1.077%
LCG.Krakow.pl	0.127%
LCG.Lancashire.uk	0.515%
LCG.Legnaro.it	2.076%
LCG.Manchester.uk	0.473%
LCG.Milano.it	0.527%
LCG.Montreal.ca	0.041%
LCG.NCU.tw	0.408%
LCG.NIKHEF.nl	3.963%
LCG.Napoli.it	0.062%
LCG.Oxford.uk	0.791%
LCG.PIC.es	3.716%
LCG.Padova.it	0.099%
LCG.QMUL.uk	1.417%
LCG.RAL-HEP.uk	1.042%
LCG.RAL.uk	7.726%
LCG.RHUL.uk	0.463%
LCG.Roma.it	0.052%
LCG.SARA.nl	0.246%
LCG.SINP.ru	0.034%
LCG.Sheffield.uk	0.420%
LCG.Torino.it	0.722%
LCG.Toronto.ca	0.143%
LCG.Triumf.ca	0.317%
LCG.UCL-CCC.uk	0.795%
LCG.USC.es	0.193%
LCG.WEIZMANN.il	0.034%

@2004-09-06 Between 2004-05-03 – 2004-08-31

20 non-LCG Sites

43 LCG Sites

Both production environments under the control of DIRAC

424 CPU years

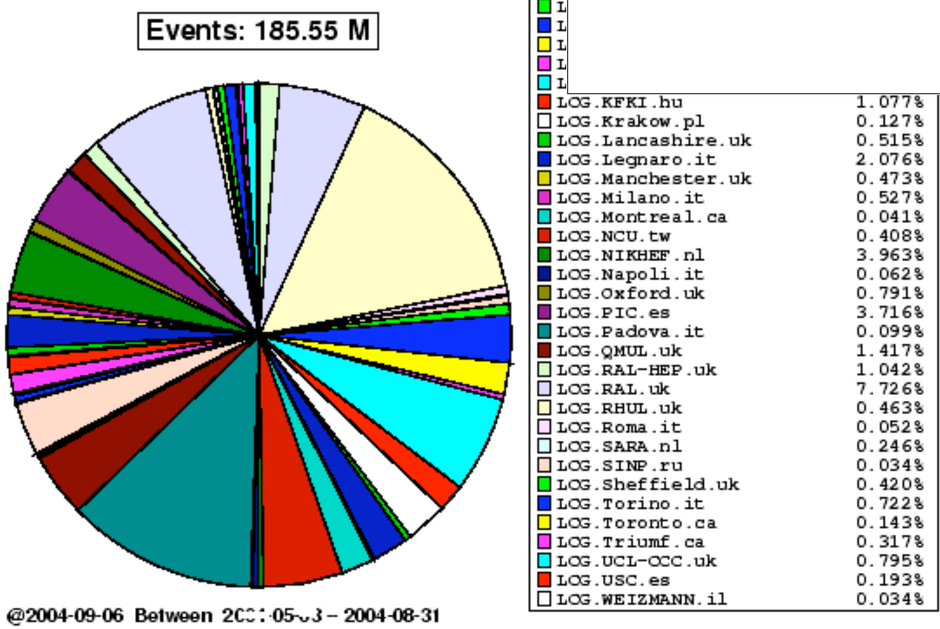

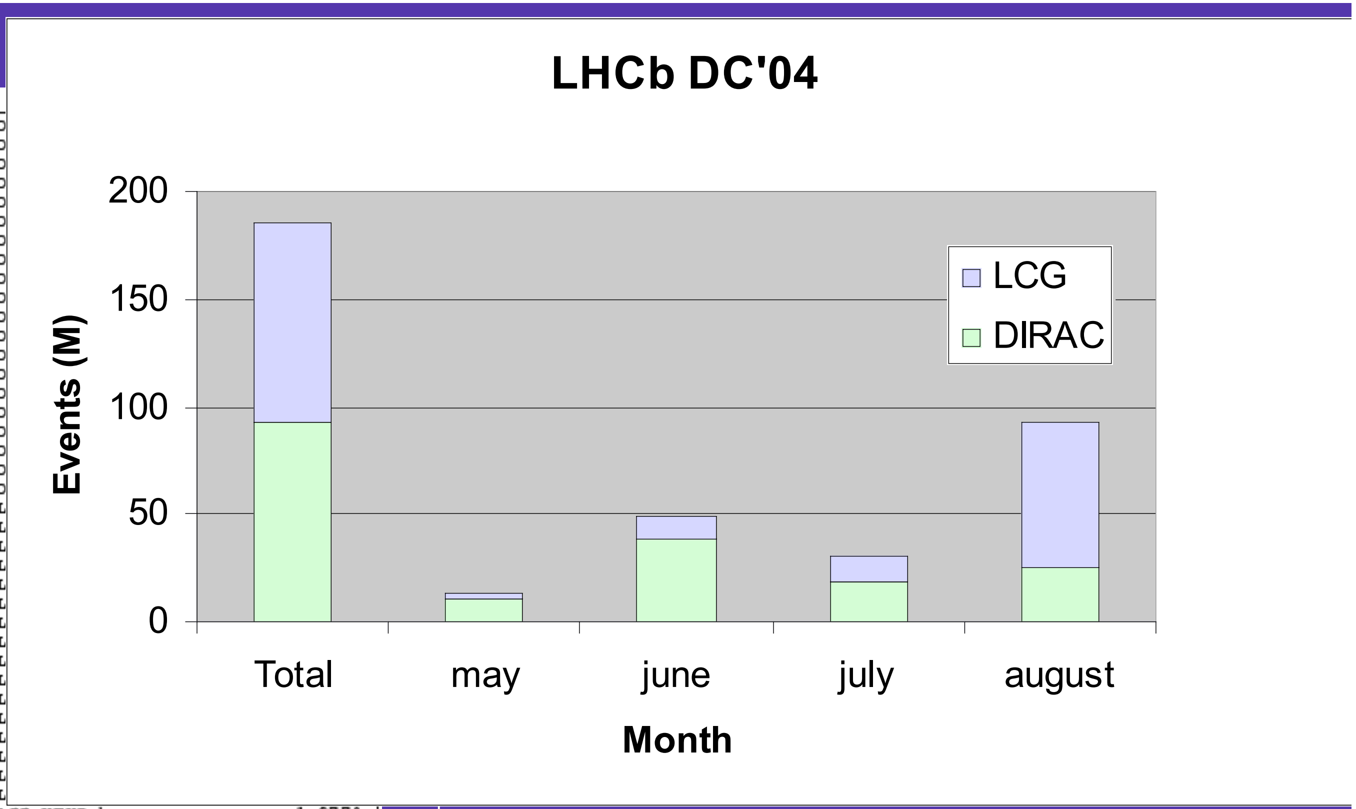

43 LCG Sites

Both production environments under the control of DIRAC

DC04 produced data

TIER 0	Nb of Events	Size (TB)
CERN	187.6M	62

Tier 1	Nb of Events	Size (TB)
CNAF	37.1M	12.6
RAL	19 .5M	6.5
PIC	16.5M	5.4
Karlsruhe	12.5M	4
Lyon	4.4M	1.5

Large scale production in 2005 on the Grid

- The RTTC production lasted just 20 days
- The startup was very fast
  - In a few days almost all available sites were in production
  - system was able to run with 4000 CPUs over 3 weeks, with a peak of over 5500 CPUs - improvement with respect to DC04 data challenge.
- 168 M events produced (11 M events as final output after L0)

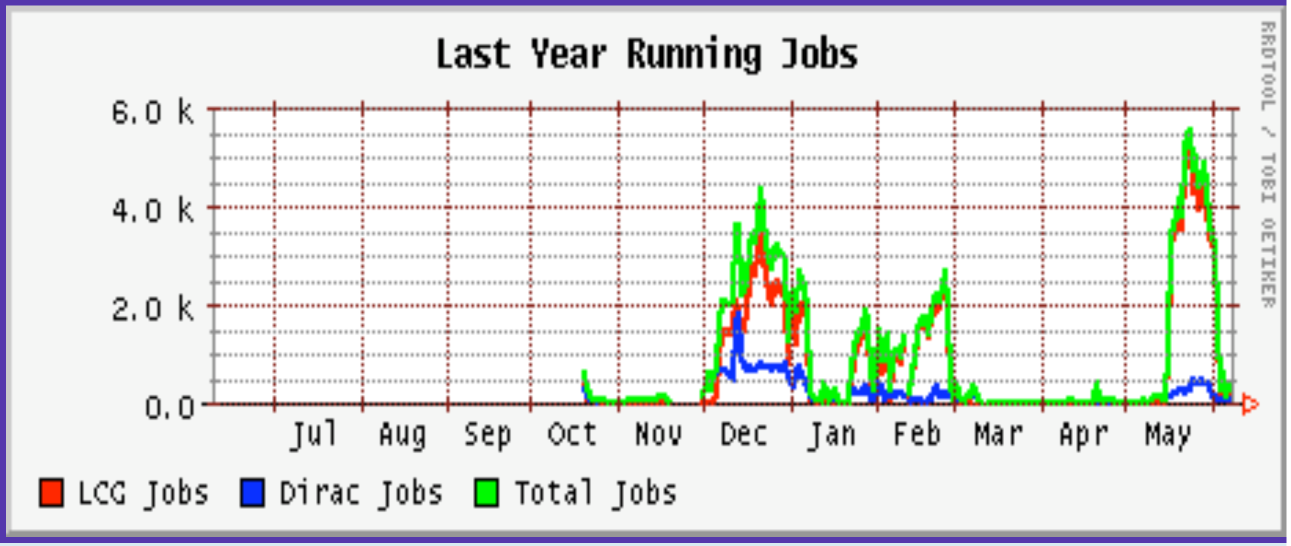

LHCb & LCG - SC3 & beyond

- Data Management
  - Storage Elements for permanent storage should have a common S(torage) R(esource) M(anagement) interface
    - Supports the LCG requirements for SRM (v2.1)
  - Evaluating for transfer gLite-FTS in Service Challenge 3 (SC3)
  - Evaluating LCG File Catalog in SC3
    - Previously used AliEn FC and LHCb bookkeeping DB
  - Uses its own "metadata" catalogue (LHCb Bookkeeping DB)
    - Implementation based on ARDA metadata interface being tested
- Computing resources
  - Requires a standard Computing Element (front-end to local resource management system) interface to which Dirac agents could submit jobs and query status and monitoring information
  - Requires a framework for deploying LHCb-specific agents at major sites
    - Resources (CPU, disk, database) to be defined with sites

LHCb Computing Milestones

Analysis at all Tier-1's - November 2005

Start data processing phase of DC'06 - May 2006

i. Distribution of RAW data from CERN

ii. Reconstruction/stripping at Tier-1's including CERN

iii. DST distribution to CERN & other Tier-1's

Alignment/calibration challenge – October 2006

i. Align/calibrate detector

ii. Distribute DB slice – synchronize remote DB's

iii. Reconstruct data

Production system and software ready for data taking - April 2007

Summary

- LHCb has in place a robust s/w framework
- Grid computing can be successfully exploited for production-like tasks
- Next steps:
  - Realistic Grid user analyses
  - Prepare reconstruction to deal with real data
    - particularly calibration, alignment, …
  - Stress testing of the computing model