# Usecases of LCG-2 at GridKa

Johan Blouw

Max Planck Institute for Nuclear Physics

# Overview

- Introduction
- LHC
- LHCb
- Distributed Computing
- GridKa
- DC03
- DC04
- Conclusions

Max-Planck-Institut für Kernphysik

# Introduction

B-physics in the 21-st century:

- Use hadron machines (CERN, DESY, Fermilab. . . )
- high cross section and high interaction rate $\rightarrow$ lot of data
- Measure cp-violation in $B_s$ meson system
- Search for rare decays
- over-constrain CKM-triangle
- Search for physics beyond the Standard Model

Max-Planck-Institut für Kernphysik

# Introduction

- $b\bar{b}$ mesons produced at forward/backward angles

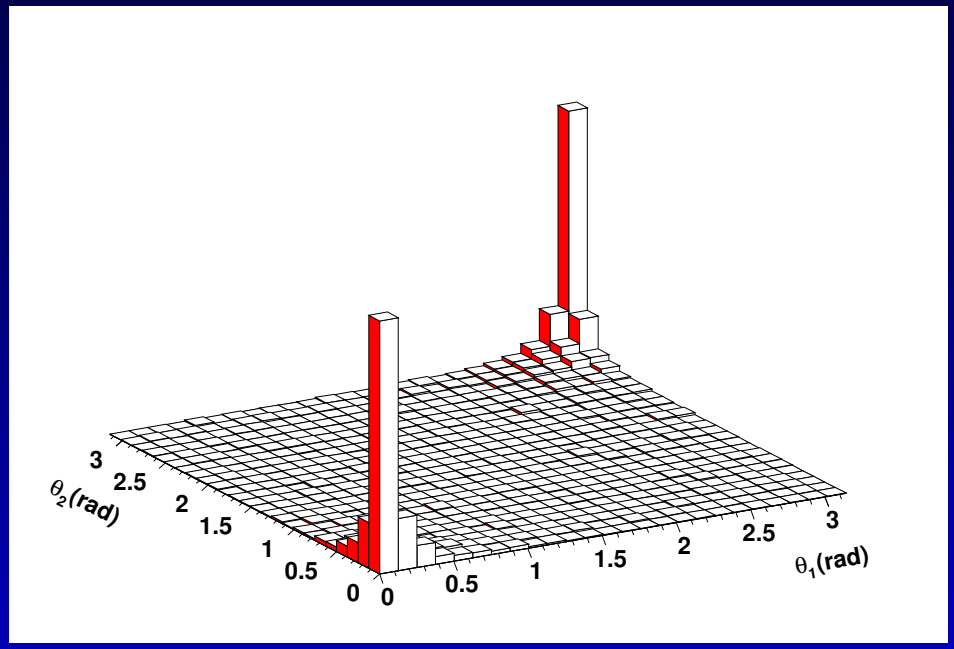


- large boost $\longrightarrow$ may resolve $\mathrm{B}_s - \bar{\mathrm{B}}_s$ oscillations...
- simple forward spectrometer sufficient (HERAB, BTEV, LHCb)

Max-Planck-Institut für Kernphysik

# LHC

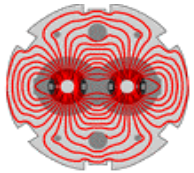


CMS

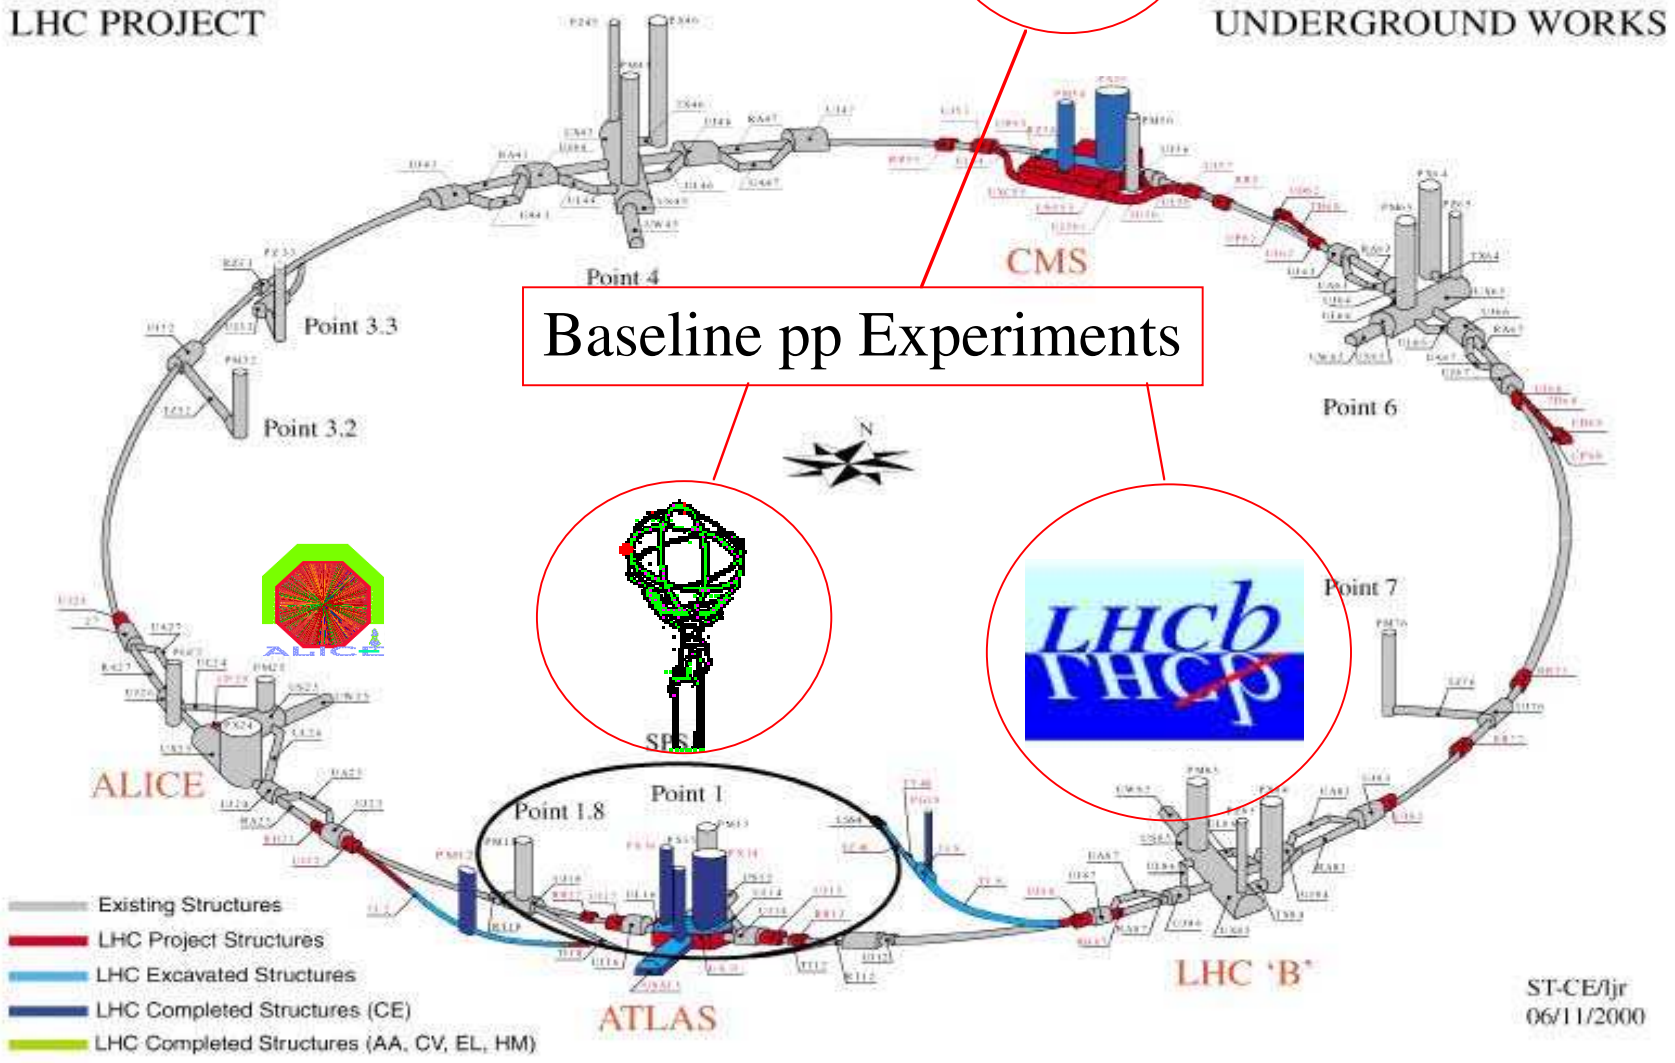
LHC PROJECT
UNDERGROUND WORKS
Point 4
CMS
Point 3.3
Point 3.2
Point 6
Baseline pp Experiments
N
Point 7
LHCb
ALICE
Point 1.8
Point 1
ATLAS
LHC 'B'
Existing Structures
LHC Project Structures
LHC Excavated Structures
LHC Completed Structures (CE)
LHC Completed Structures (AA, CV, EL, HM)
ST-CE/ljr
06/11/2000

# LHC at CERN

- $\mathrm{p\bar{p}}$ collisions at $\sqrt{s} = 14$ TeV
- 40 MHz bunch crossing rate
- $\mathcal{L} \sim 10^{32}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$

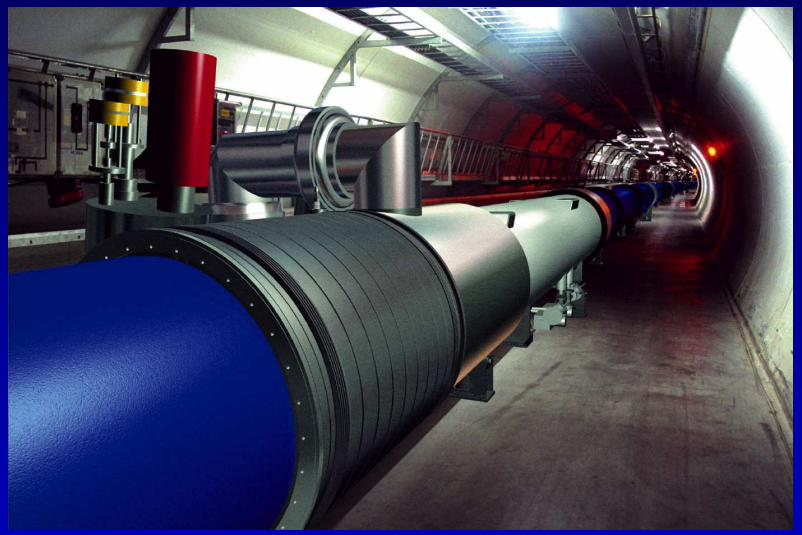

Max-Planck-Institut für Kernphysik

# LHC at CERN

- large $b\bar{b}$ cross section: $\sigma \sim 500\ \mu\mathrm{b}$
- signal/noise ratio: $\sigma_{b\bar{b}}/\sigma_{\mathrm{inelastic}} \sim 5 \times 10^{-3}$
- $10^{12}\ \mathrm{b\bar{b}}$ pairs per year

Max-Planck-Institut für Kernphysik

# LHCb challenges

LHCb detector located in DELPHI pit:

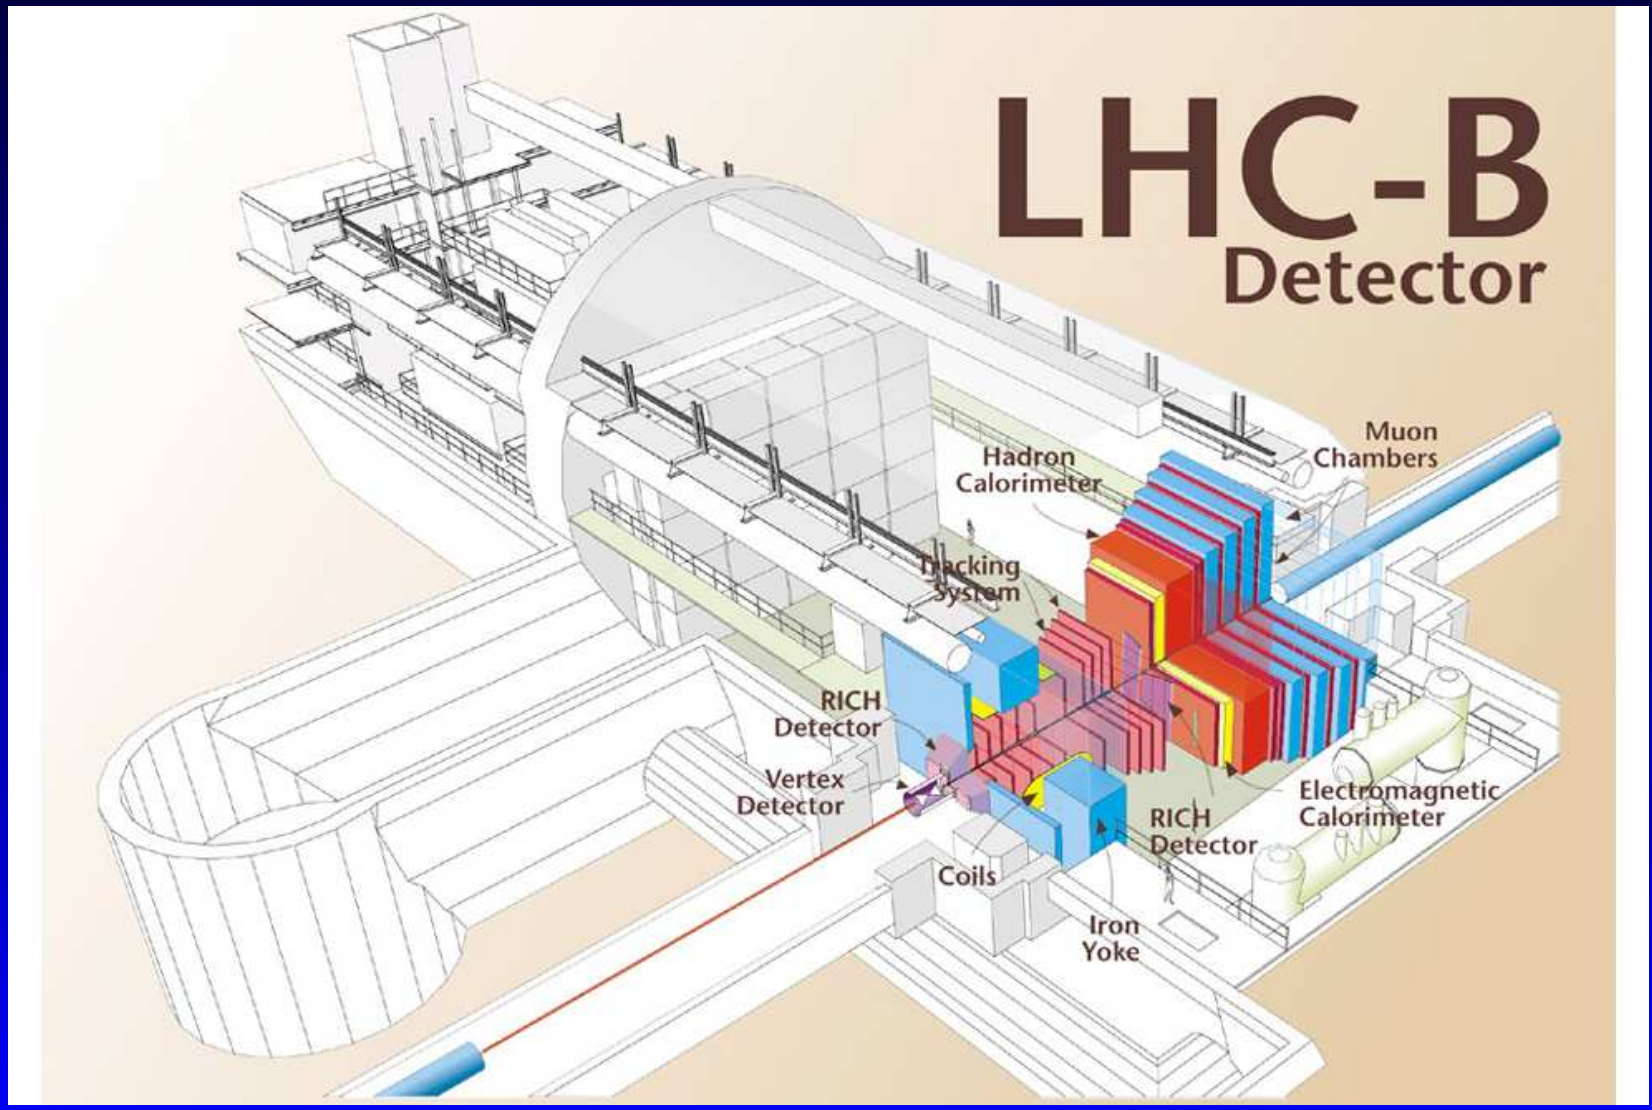


Max-Planck-Institut für Kernphysik

# LHCb challenges

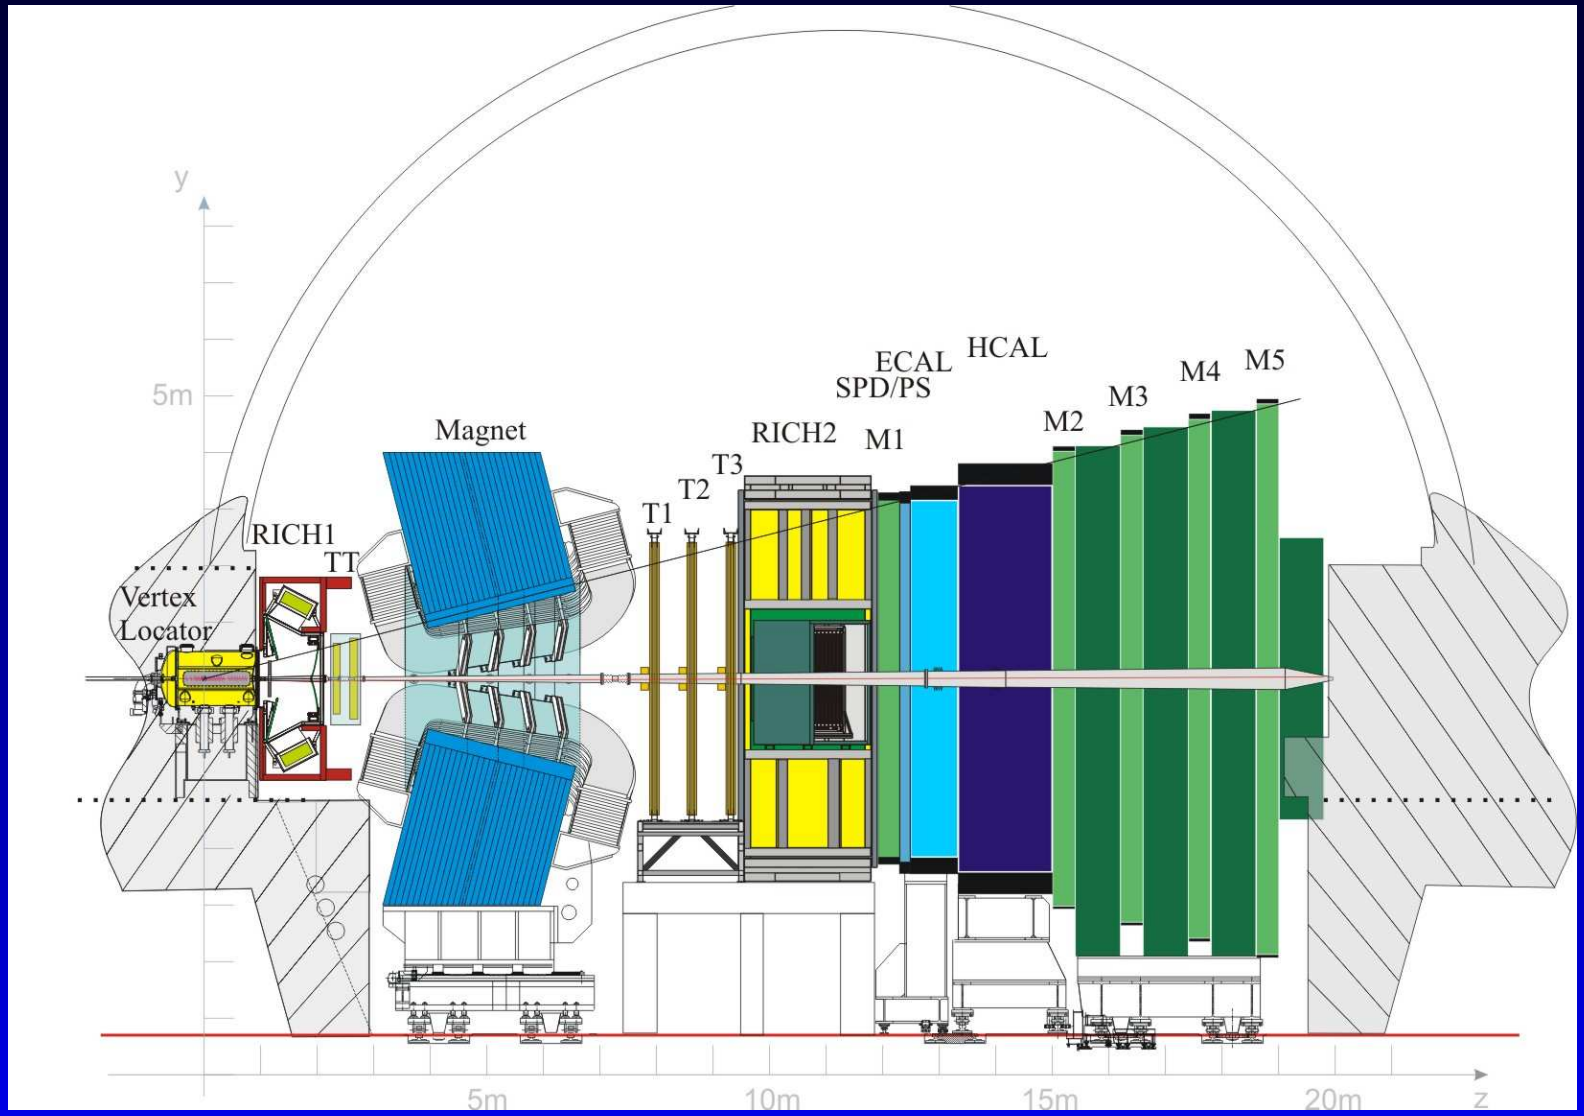

# LHCb challenges

- Study physics performance

depends on

- trigger performance (efficiency, event rejection. . . )
- tracking performance (efficiency, ghosts, momentum resolution)
- material budget

$\Rightarrow$ use detector and event simulation

Max-Planck-Institut für Kernphysik

# LHCb challenges

predicted rates for 1 Snowmass year:

- $10^{12} \mathrm{b\bar{b}}$ events per year
- at 10 MHz of visible $\mathrm{p\bar{p}}$ crossings
- Level 0: 40 MHz input, reduce to 1 MHz
- Level 1: reduce 1 MHz to 40 KHz
- HLT: reduce to 200 Hz
- Data volume:
  $200\ \mathrm{Hz} \times 1\ \mathrm{Snowmass\ year} = 2 \times 10^9$ events
- at 50 KByte/event $\rightarrow \mathcal{O}(10^4)$ TByte/year

Max-Planck-Institut für Kernphysik

# LHCb challenges

- enormous amount of simulated data needed
- detector simulation: $\sim$ minute/event/1GHz
- trigger studies
- background studies
- mass reconstruction efficiency studies
- cp-reach, rare decays, etc. etc. . .

$\Rightarrow$ too much data for CERN alone

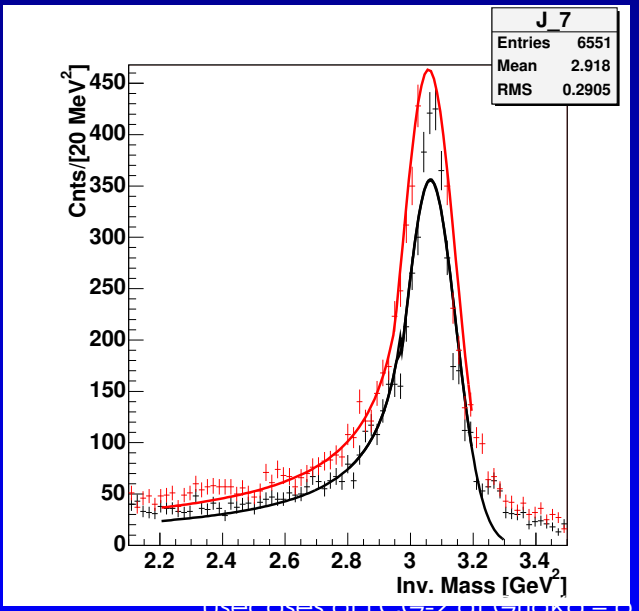


Max-Planck-Institut für Kernphysik

# Computing Challenge

- use distributed computing resources (cpu, storage)
- need high bandwidth connections between computing centers

$\Rightarrow$ birth of GRID…

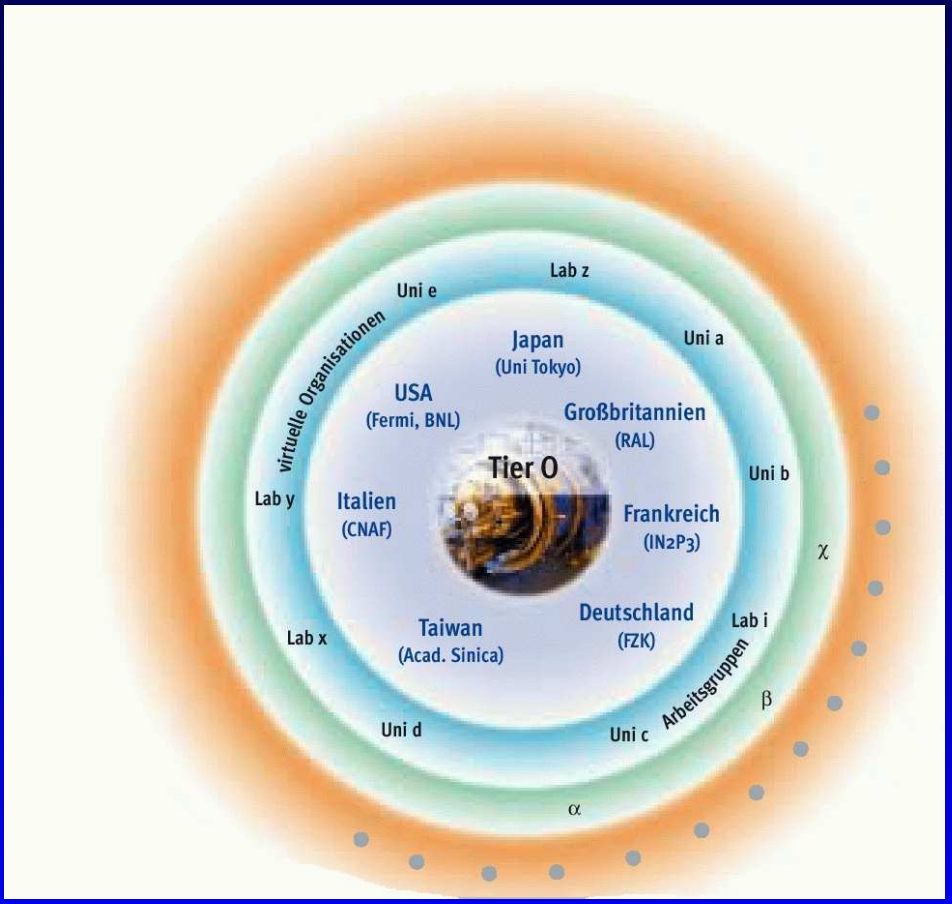


Max-Planck-Institut für Kernphysik

# Computing Challenge

GRID: nice idea, but no easy meat. . .

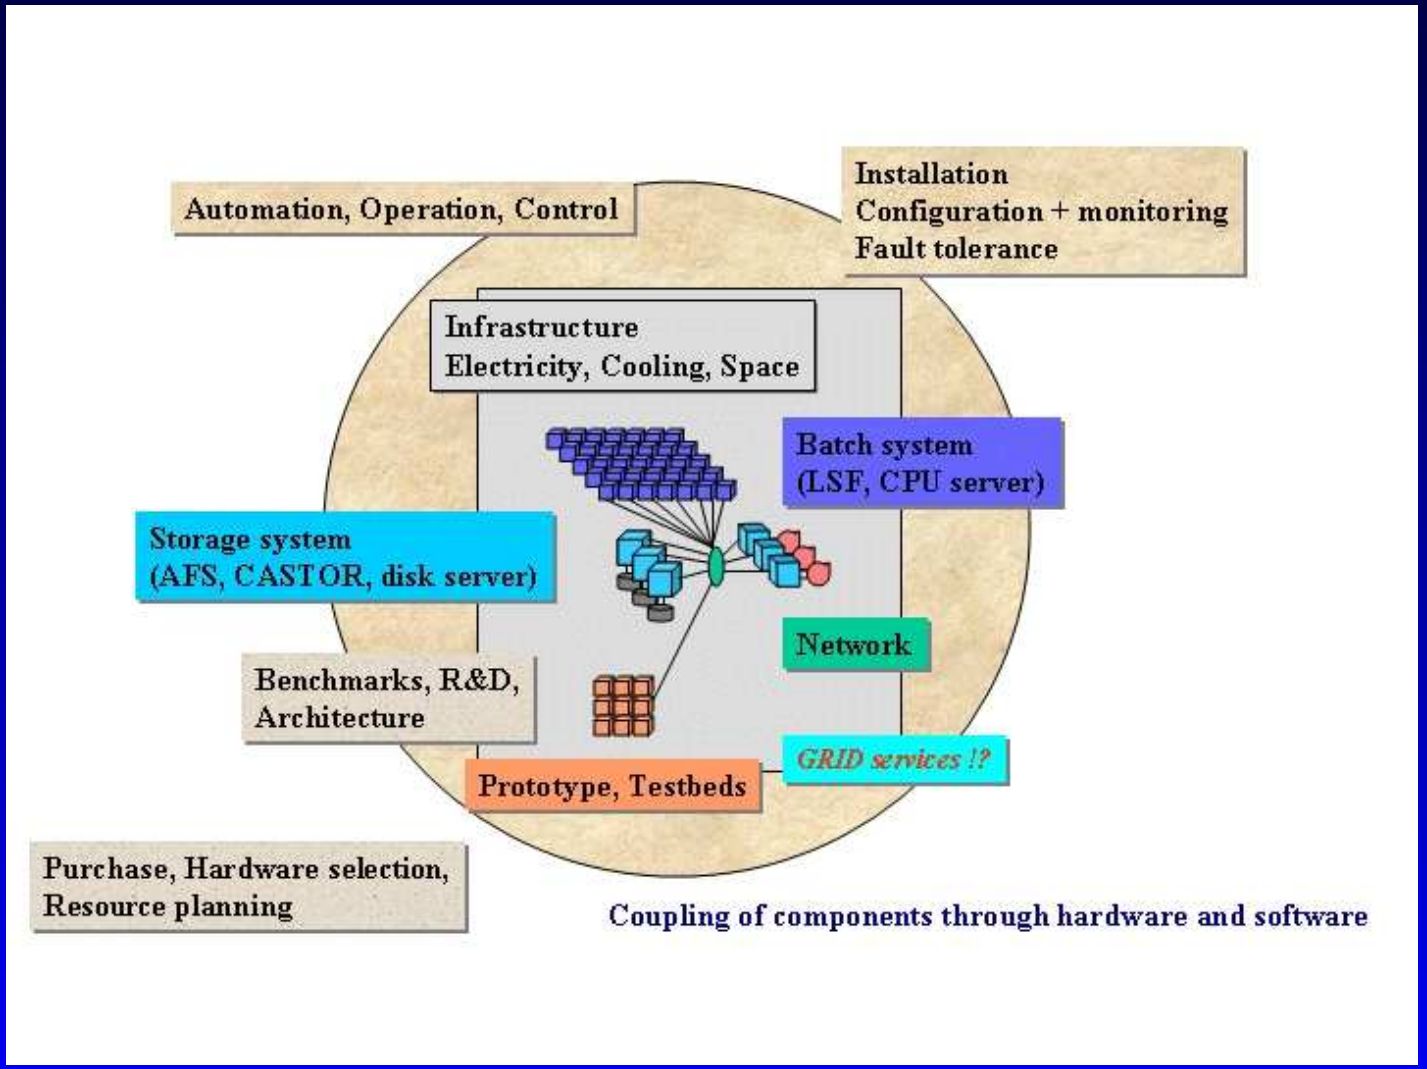


Max-Planck-Institut für Kernphysik

# DIRAC '03 (No LCG usage...)

Distributed Infrastructure with Remote Agent Control...
Different from DataGRID philosophy:

- jobs not time critical
- many more computing jobs than computing resources

! use "PULL" concept instead of "PUSH".

- run job-agents at each participating site
- job-agents request jobs from central job-server

Max-Planck-Institut für Kernphysik

# DIRAC '03 (No LCG usage...)

- Production task definition
- Software installation on production sites
- Production steering
- Job scheduling
- Job monitoring
- data transfer
- bookkeeping

Max-Planck-Institut für Kernphysik

# DIRAC '03 (No LCG usage...)

- use simple system of job submission

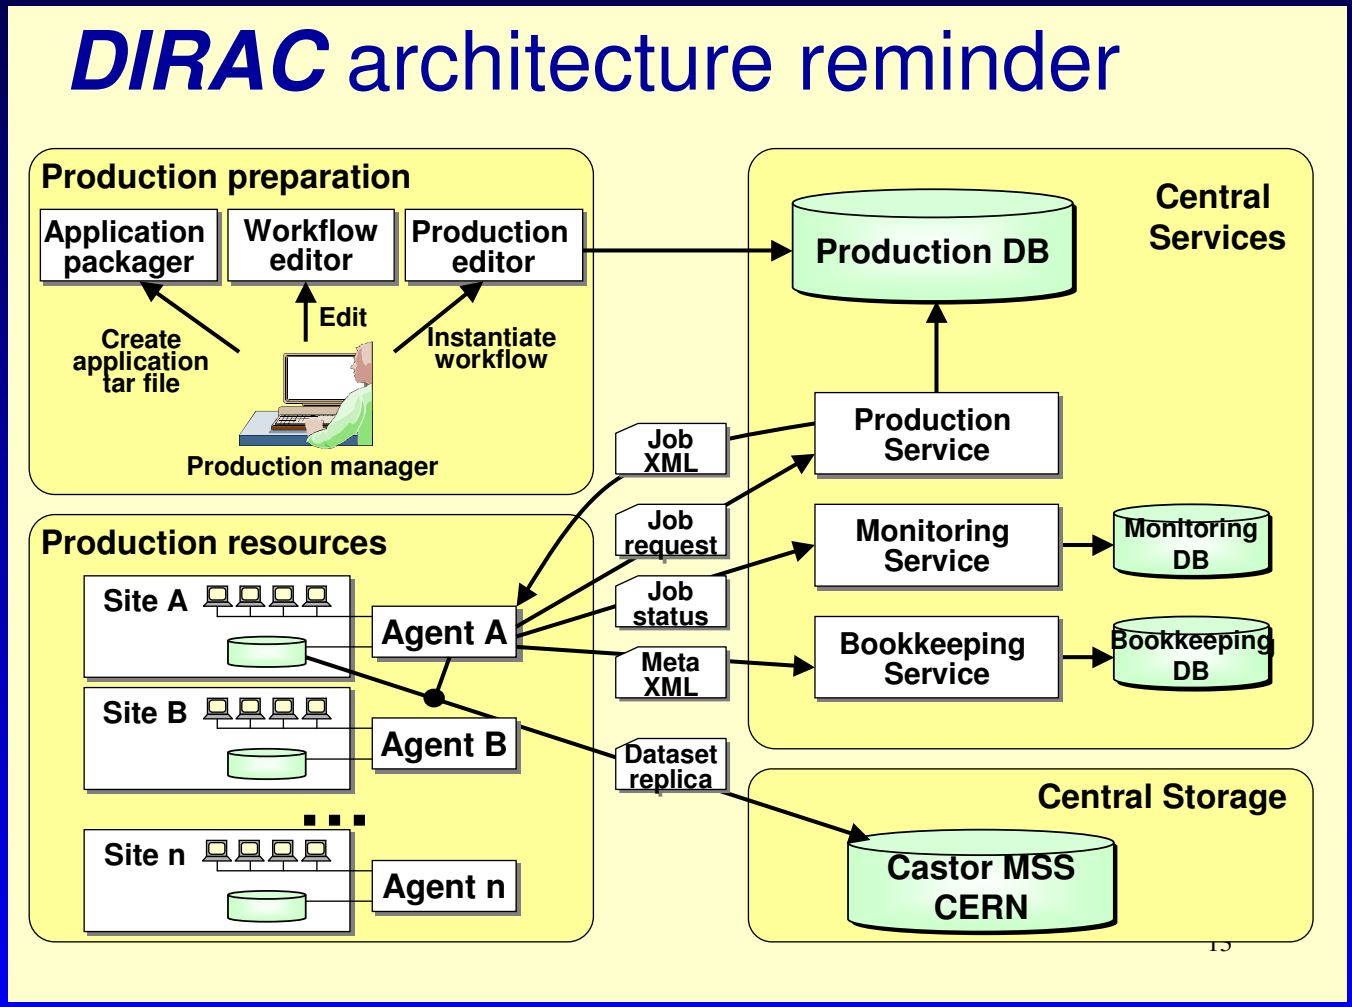

# DIRAC ’03 (No LCG usage…)

Job description in ’sandbox’:

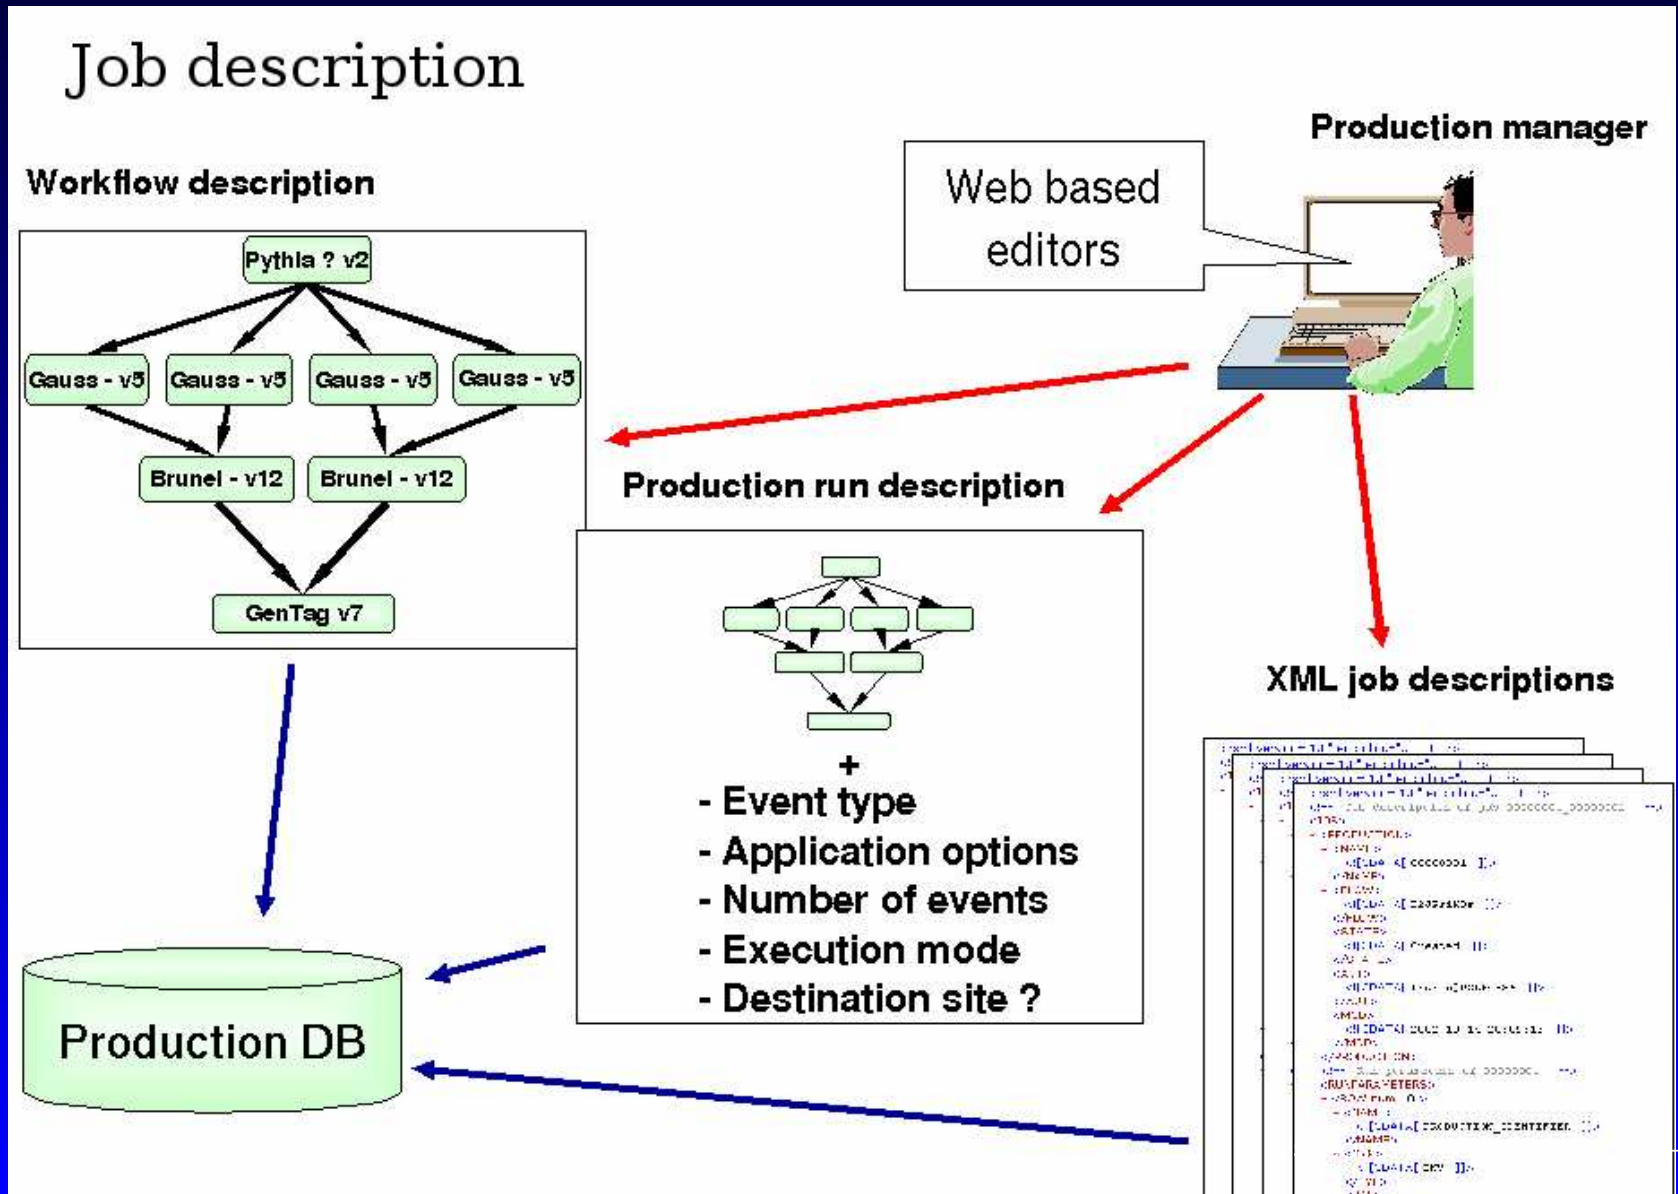

# DIRAC ’03 (No LCG usage…)

- use the “PULL” approach:

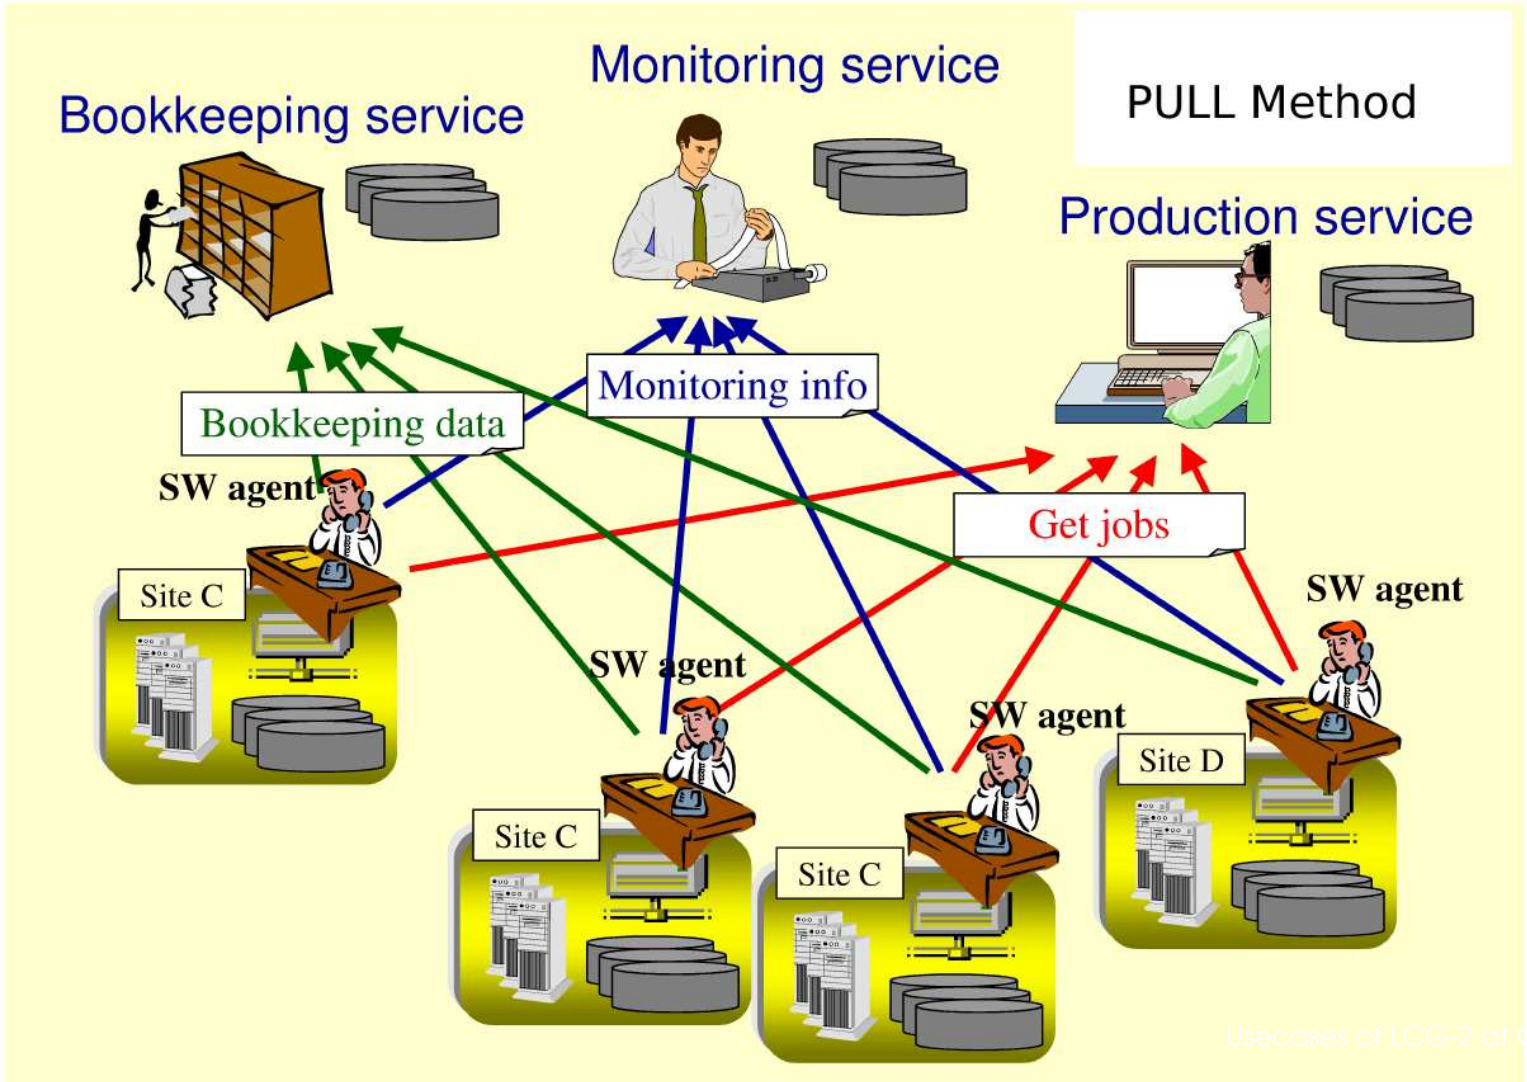

# LHCb DC03 Goals

Goals:

- Produce 30 million minimum bias events
- 10 million $B\bar{B}$ events
- 38 different signal event channels (50K each)
- use data for “re-optimized TDR”
- physics reach studies

Max-Planck-Institut für Kernphysik

# LHCb DC03 Goals

Participating sites:

| | | |
|---|---|---|
| CERN | Imperial | ScotGrid |
| Bologna | Lyon | RAL |
| Karlsruhe | Barcelona | Rio |
| Cambridge | CESGA | Bristol |
| Oxford | VU Amsterdam | . . . |

Max-Planck-Institut für Kernphysik

# Intermezzo: GridKa

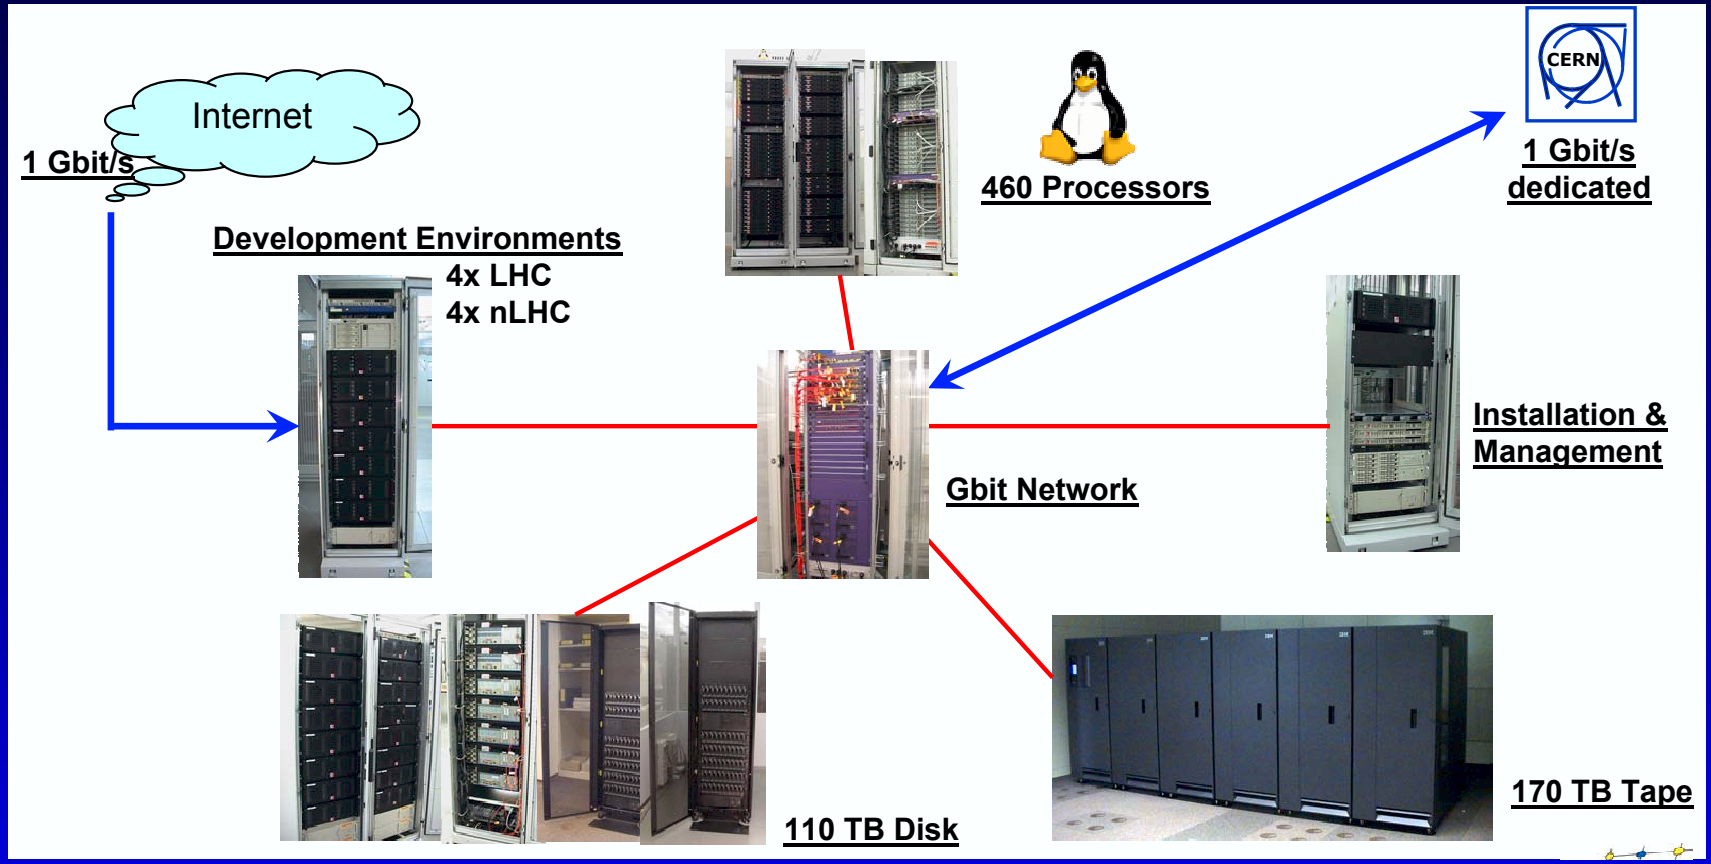

# Intermezzo: GridKa

- But software did not use hardware optimally

→ run $< 50$ jobs because of limited disk i/o

- major debugging necessary to solve bottleneck

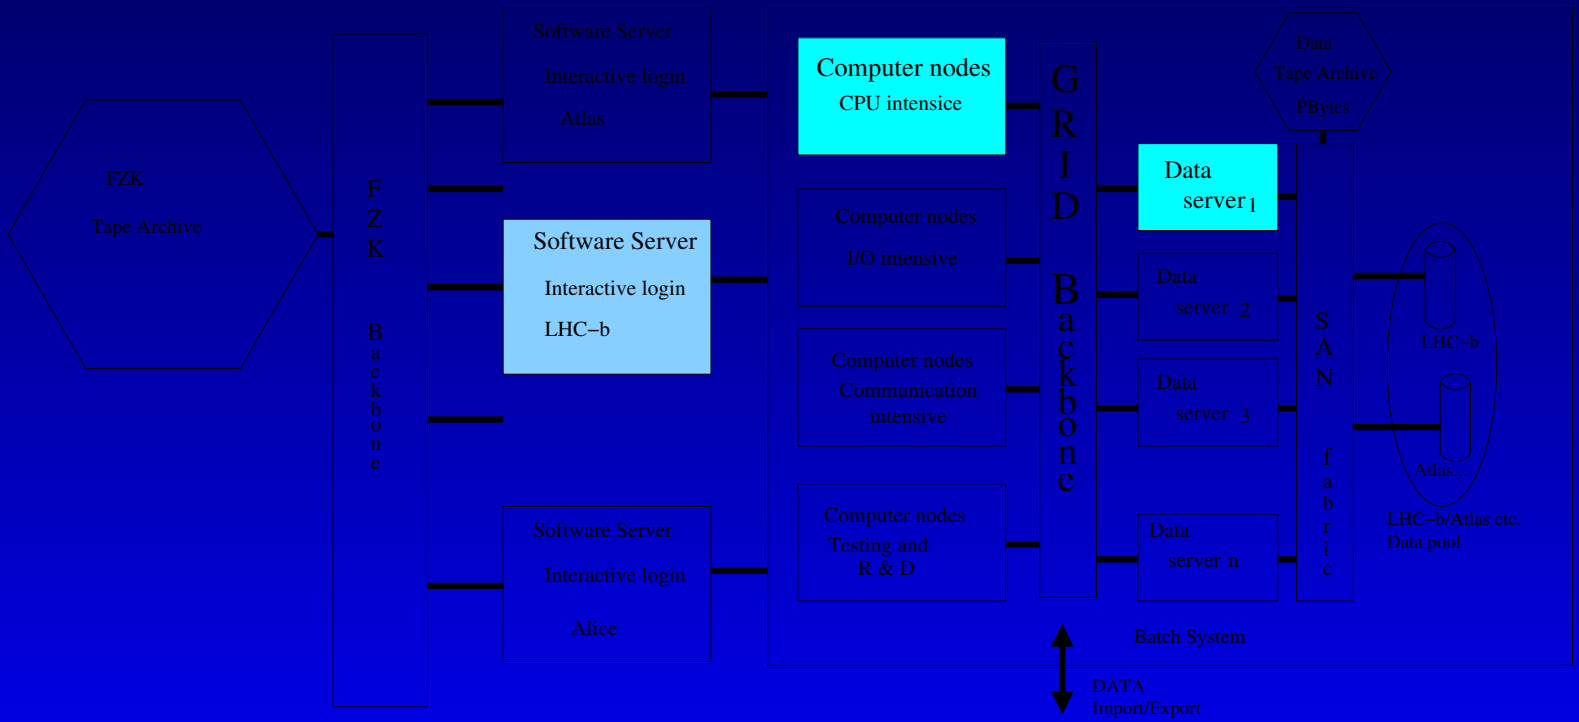

# Intermezzo: GridKa

- Each WN accessed FS simultaneously for read/write
- data through-put to FS limited by FS

$\Rightarrow$ write data to local disk first

! copy data to FS at end of job

! mainly hacking of job-submission software. . .

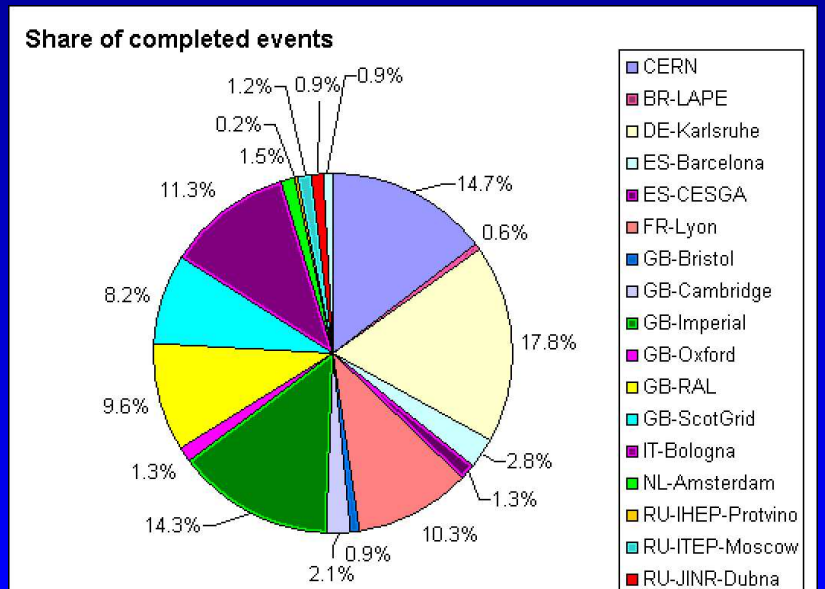

# Intermezzo: GridKa

But:

- excellent cooperation with GridKa:

→ use three different FS (1 big slow, 2 small fast servers)

  - while WNs write to fast FS
  - move data from 2nd fast FS to slow FS
  - alternate writing to fast FS1 and FS2

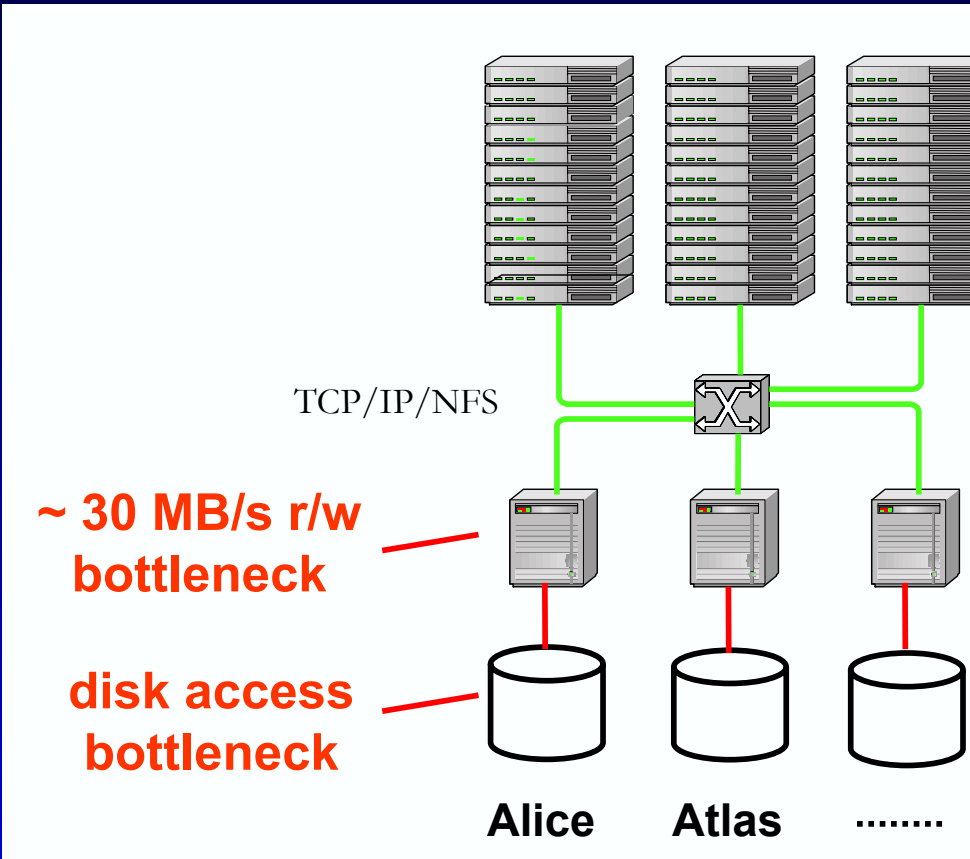


Max-Planck-Institut für Kernphysik

# DC03 Results

## Usage Example: LHCb 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*

Max-Planck-Institut für Kernphysik

# DC03 Results

- GridKa contributed $\sim 18\%$ of produced events

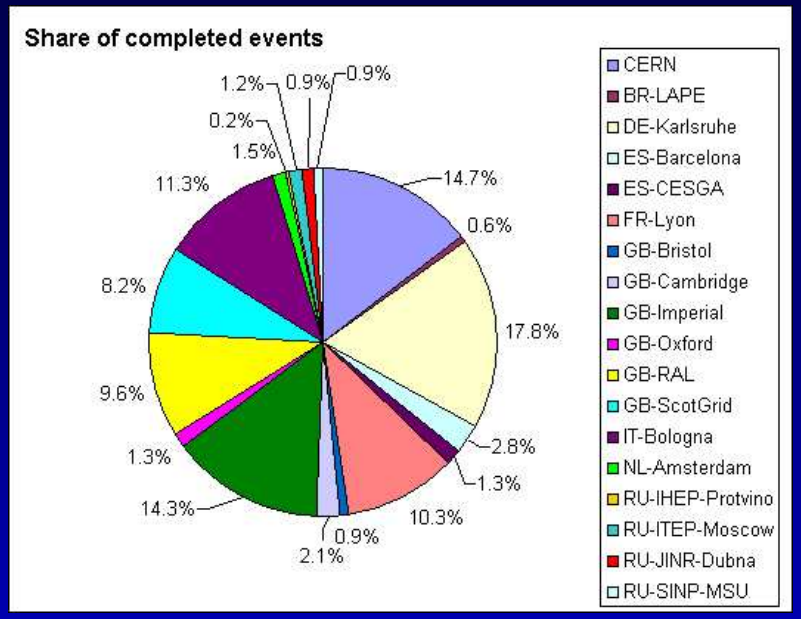


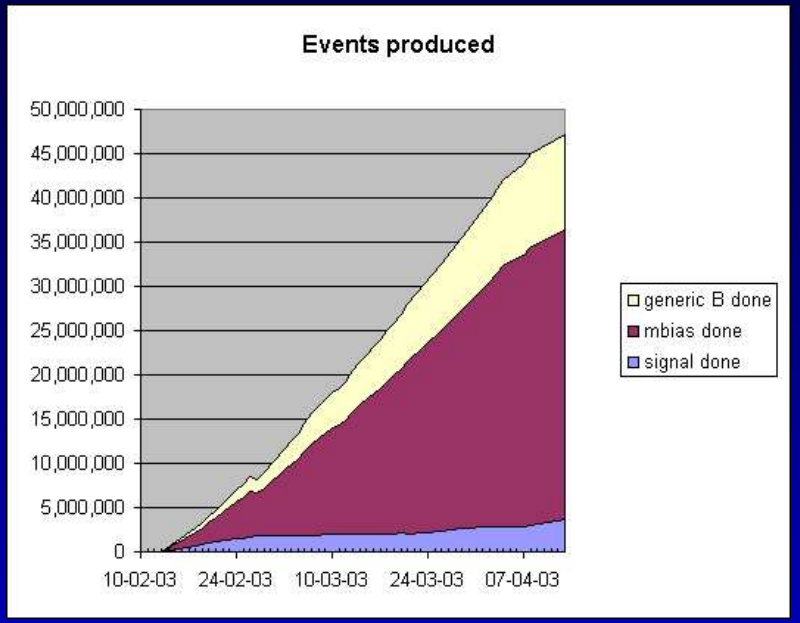


Very successful!

Max-Planck-Institut für Kernphysik

# DC03 Results

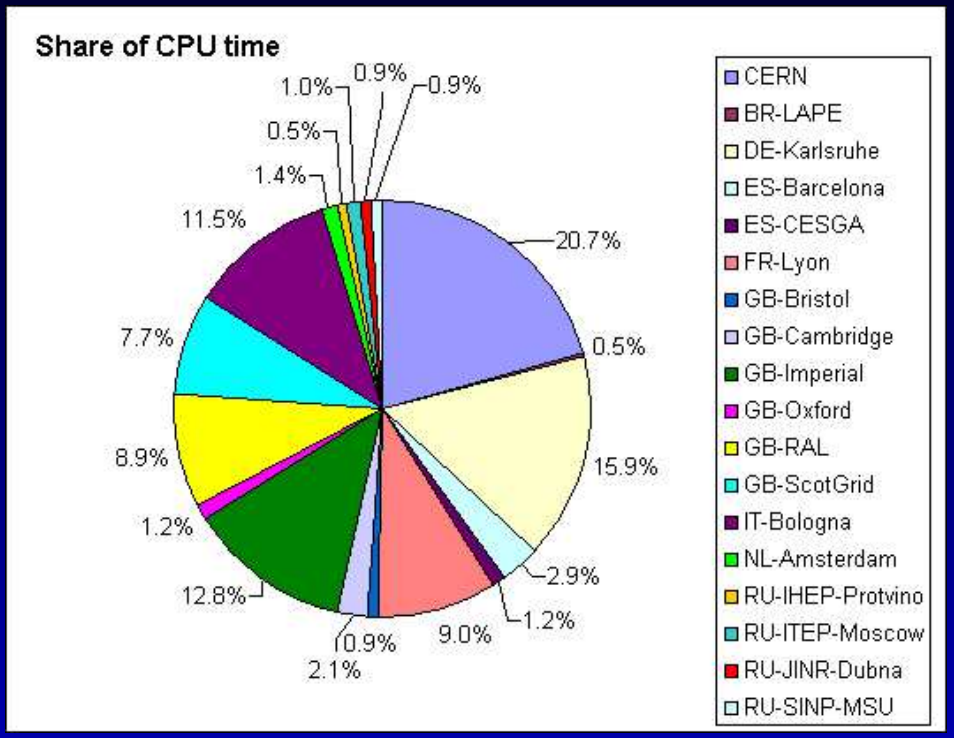


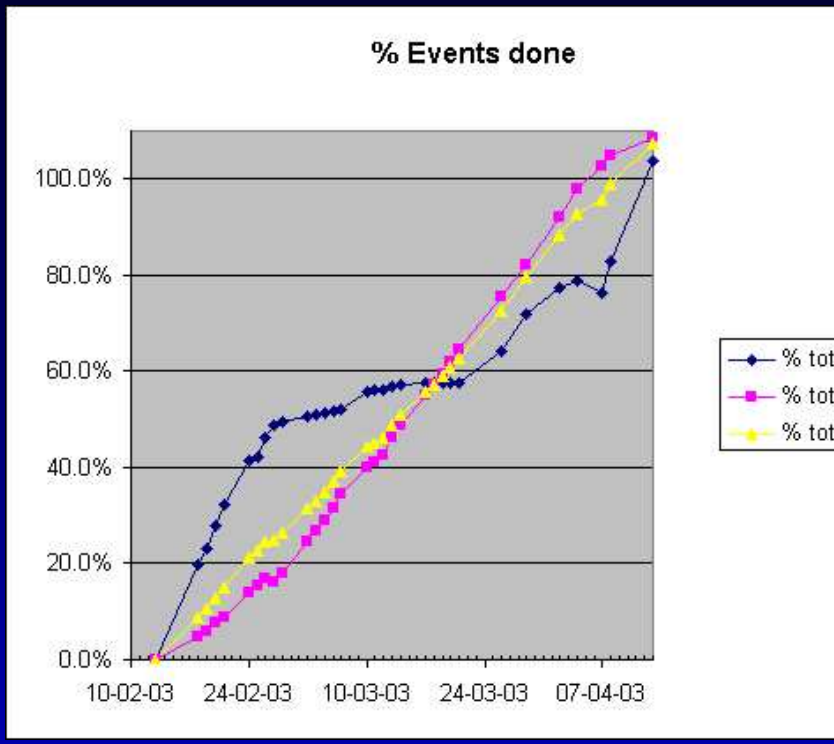


Max-Planck-Institut für Kernphysik

# Data Challenge '04

GridKa improved internal data through-put:

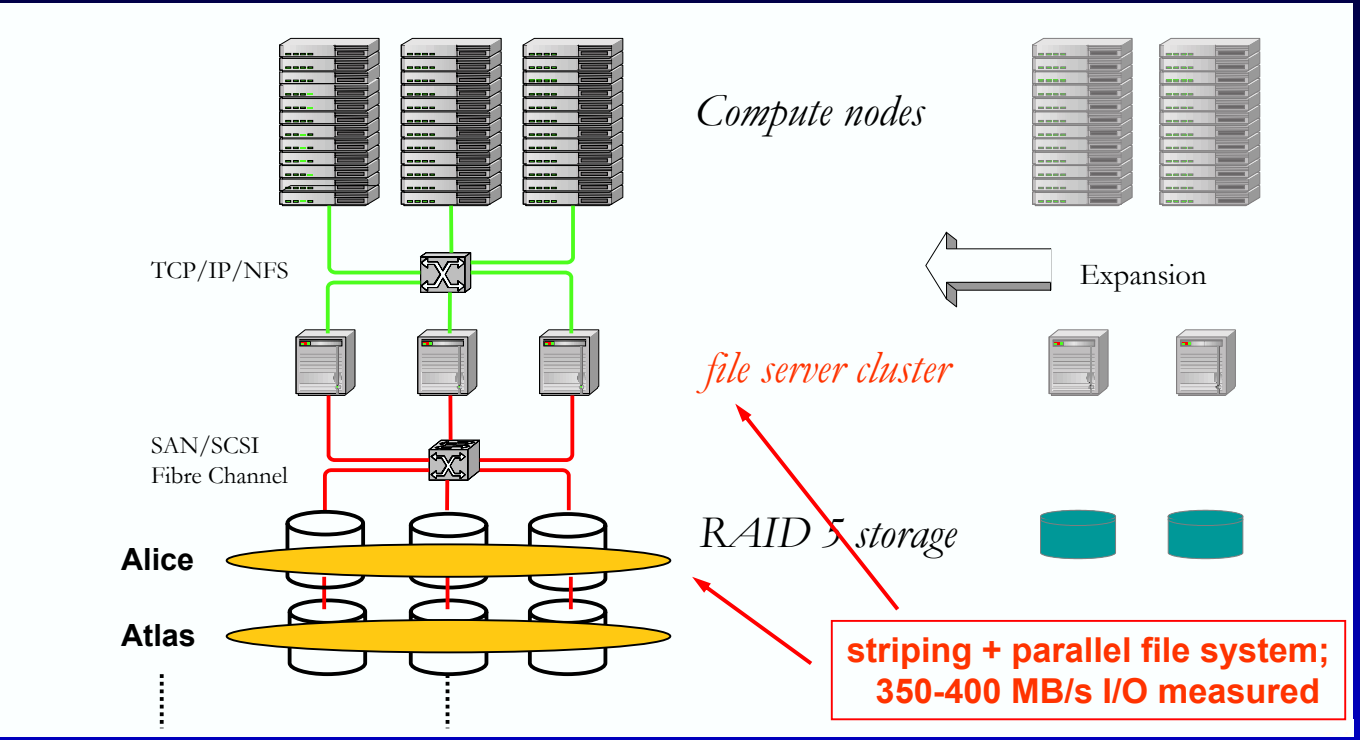


Max-Planck-Institut für Kernphysik

# Data Challenge '04

LHCb modified DIRAC to include job submission through LCG

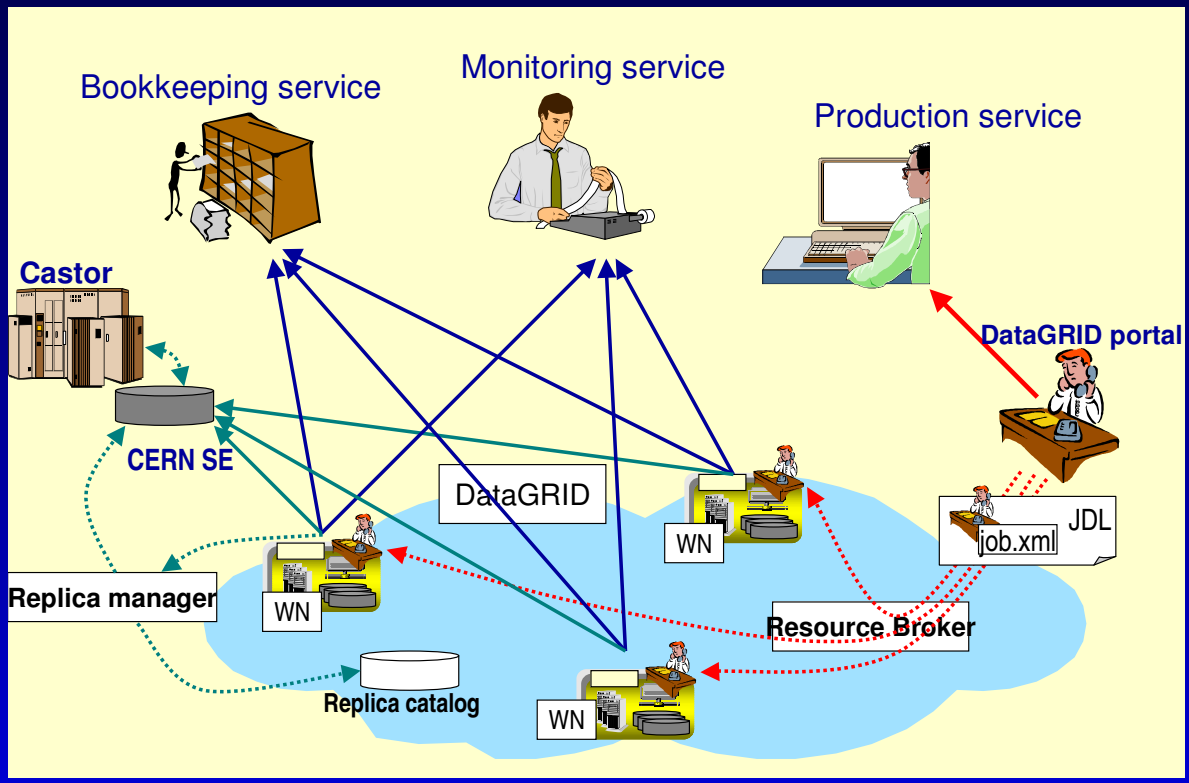

# Data Challenge '04

- accomodate LHCb specific software with
- "sandbox" technology:
  - job description
  - launch script
- use Grid software for transferring data to CERN
- but: need sandboxes for log files...
- DataGrid not very stable
  - Resource Broker unavailable
  - Proxy expires before job is over
  - large variation in available CPU power
  - some WN have no outbound IP connection

Max-Planck-Institut für Kernphysik

# Data Challenge '04

## Circumvent GridPortal...

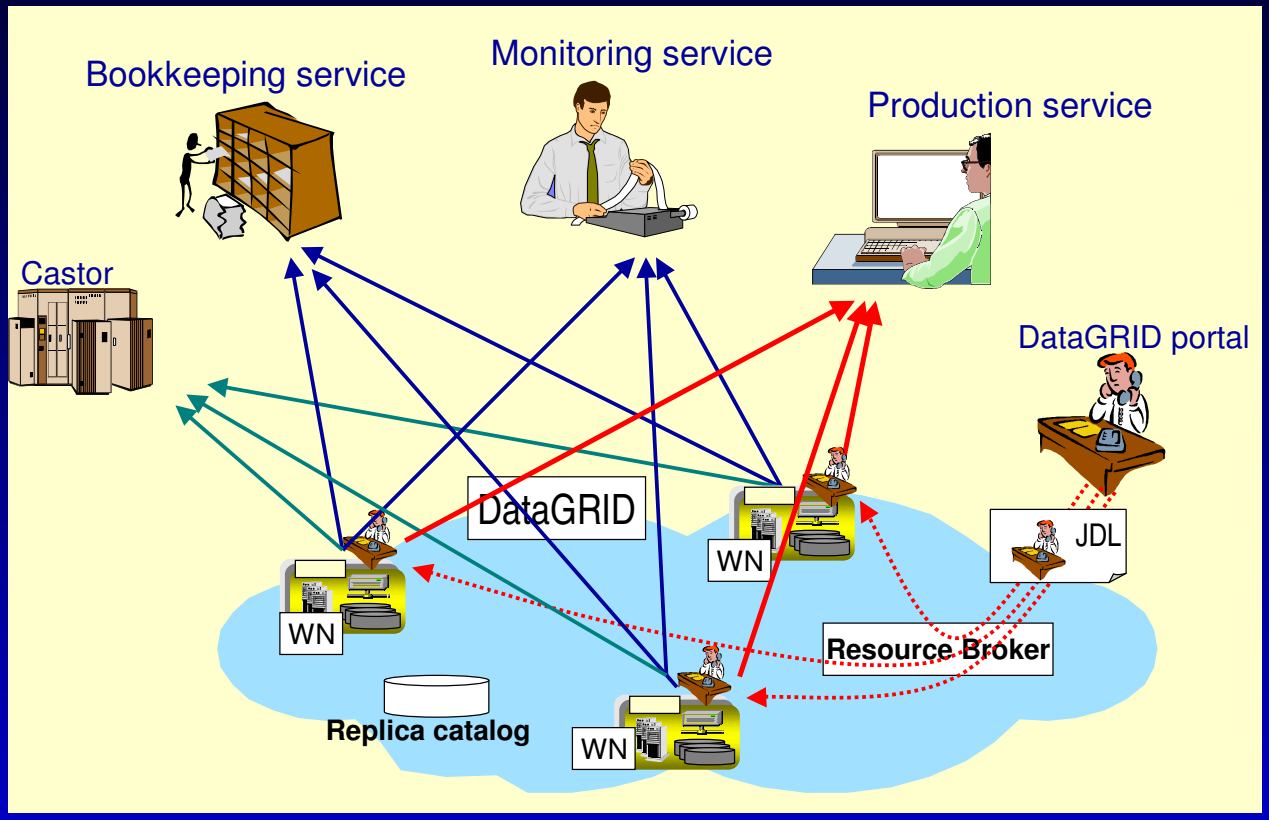

# Data Challenge '04

DIRAC $\Longleftrightarrow$ DataGrid

PULL $\Longleftrightarrow$ PUSH

decentralized job $\Longleftrightarrow$ centralized job

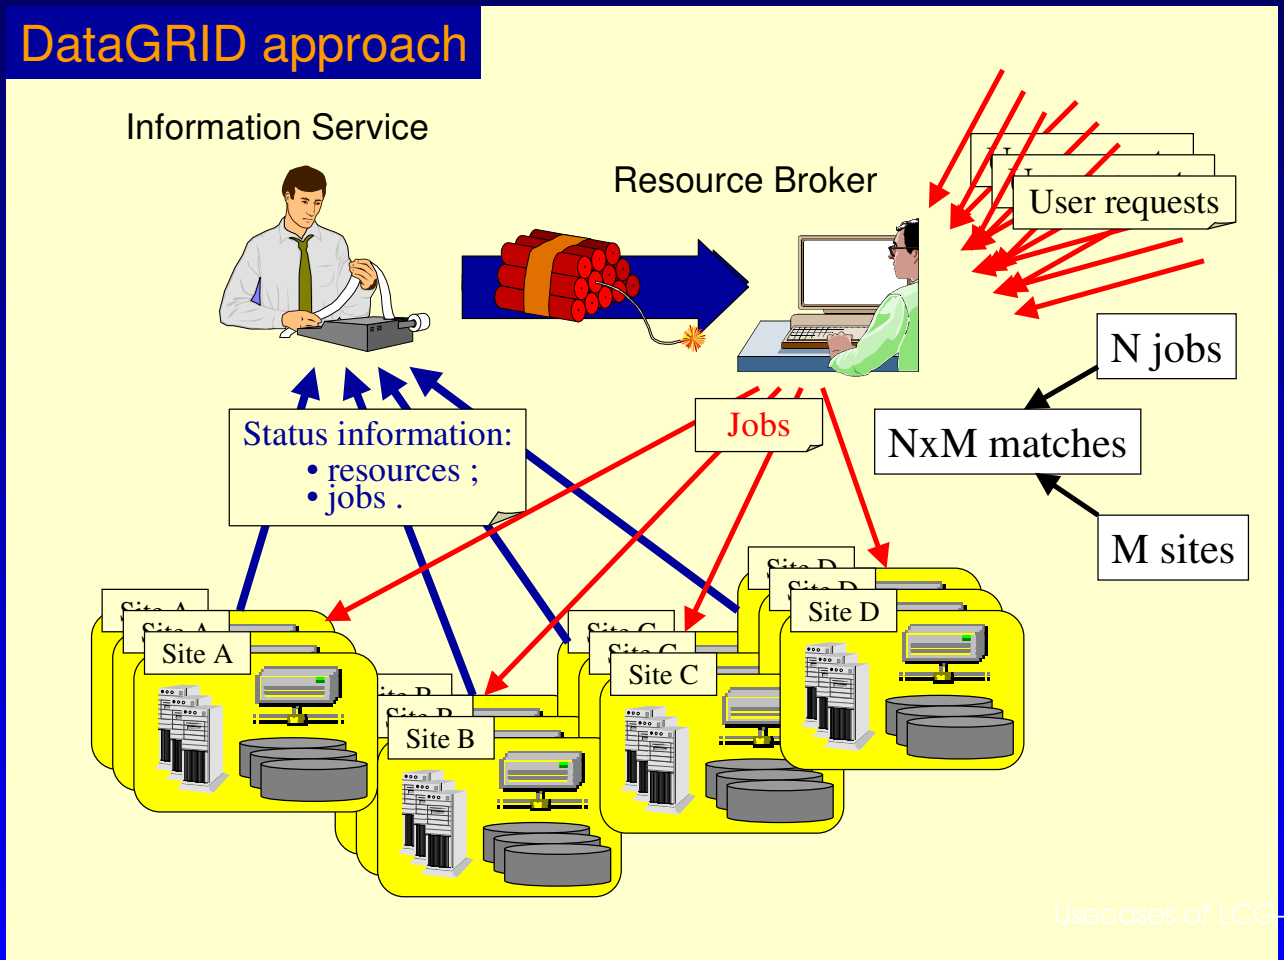

# DC04 Experiences...

- Job submission software more complicated...
- site configuration much easier
- DIRAC also on DataGrid middleware

→ non-trivial exercise...

Max-Planck-Institut für Kernphysik

# DC04 Experiences...

By using LCG, many more sites may contribute to data challenge

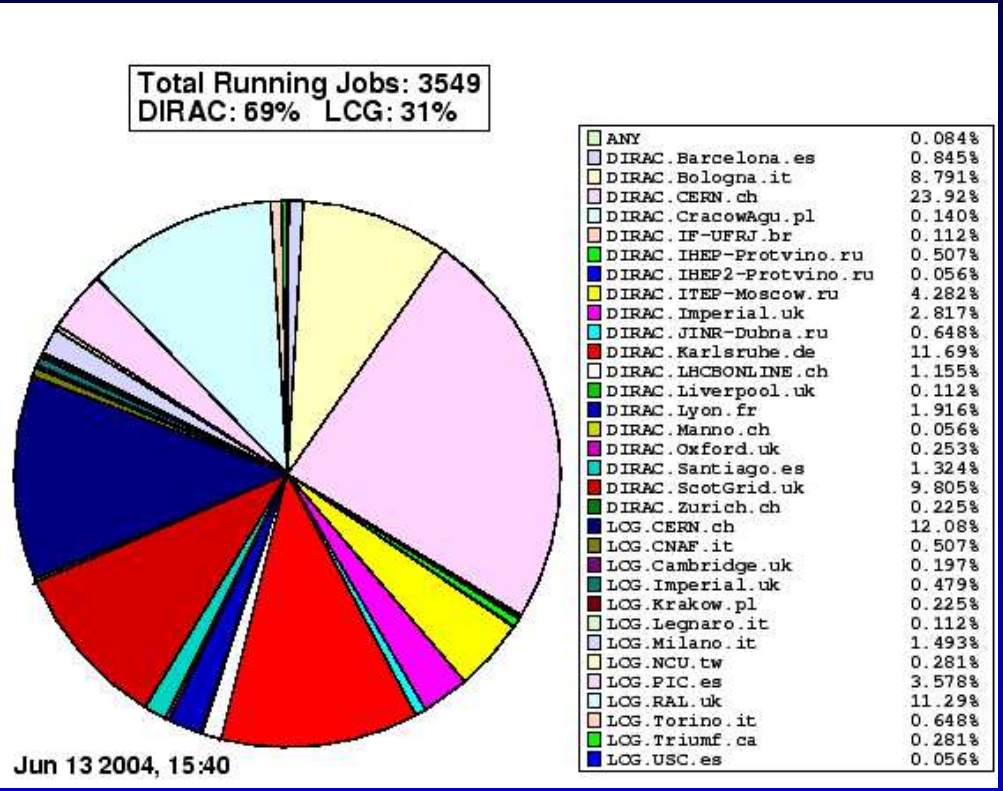


Max-Planck-Institut für Kernphysik

# DC04 Experiences...

- many updates of DIRAC software
- hiccups in data transfer to CERN

→ $\sim$ 1300 files waiting to be transferred from FZK to CERN...

→ transfer can not keep up with data production!

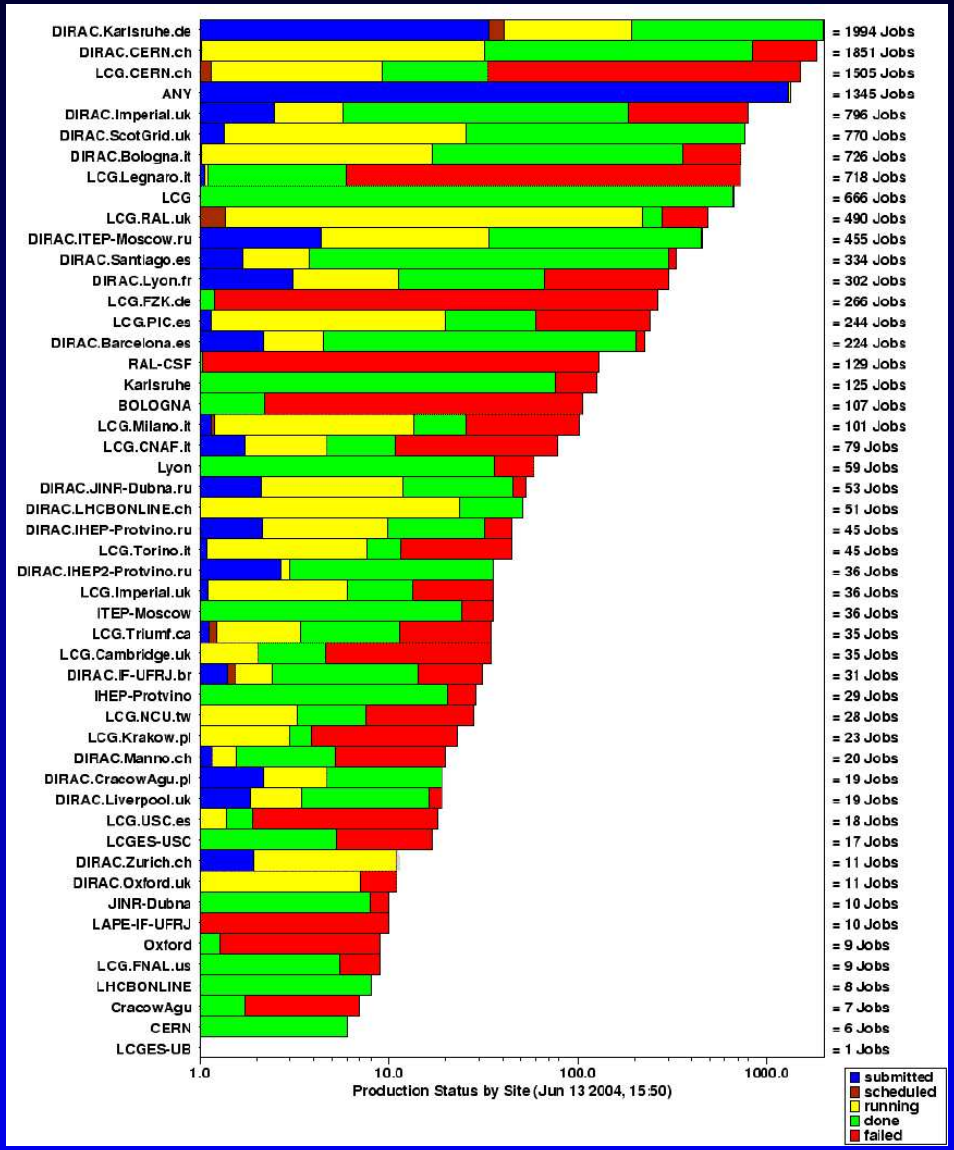


Max-Planck-Institut für Kernphysik

# DC04 Experiences...

LCG related problems:

- Resource Broker bottle-neck
- Determine Ranking (which site gets the LCG jobs?)
- Faulty WM sucks up all LCG jobs
- Software installation
- Site configuration (not always according to LCG-2 specs)
- Proxy expiration
- 10% of jobs aborted at 'job-finilization' stage

Max-Planck-Institut für Kernphysik

# DC04 Experiences...

Hardware problems:

- RB server machine at CERN suffer from high-load problems
- log-file server sshd problems
- Bookkeeping dbase problems
- HW problems at local sites

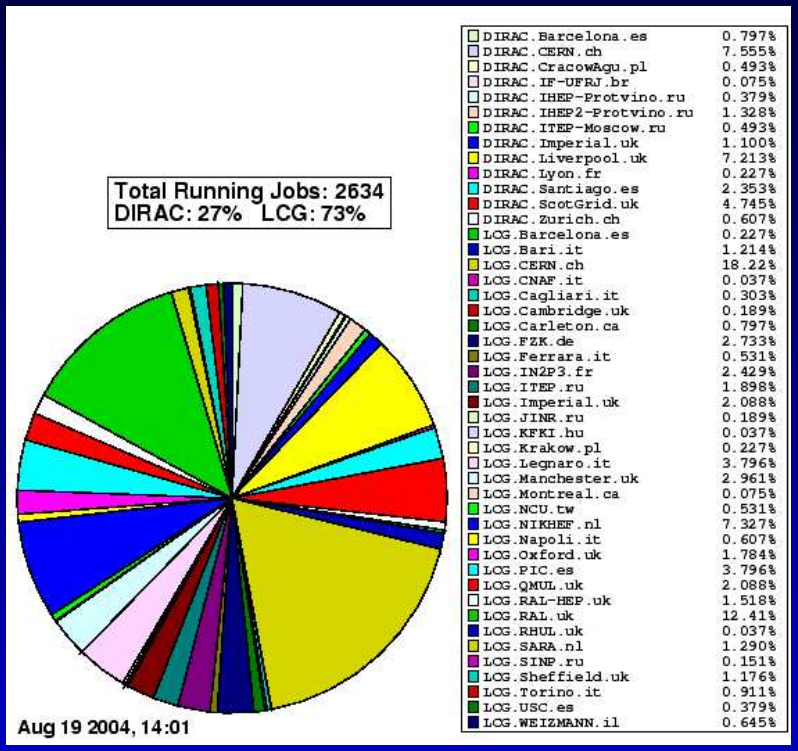


Max-Planck-Institut für Kernphysik

# DC04 Experiences...

DIRAC & LCG contributions:

DIRAC LCG

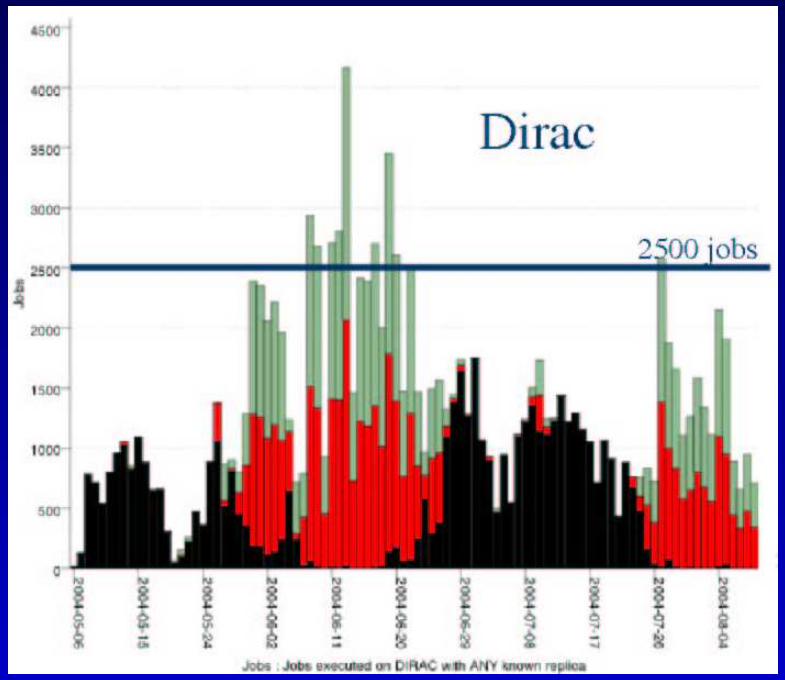


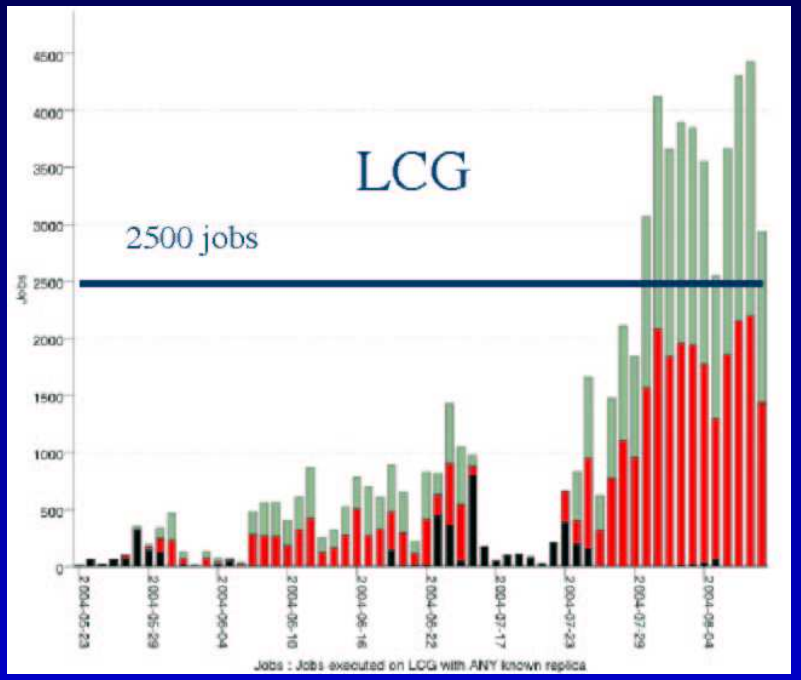


Max-Planck-Institut für Kernphysik

# DC04 Summary

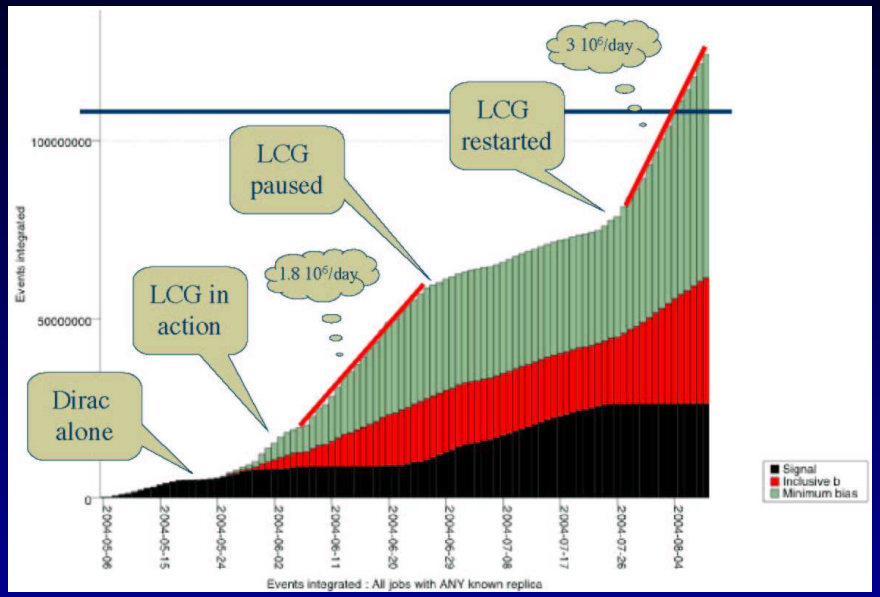


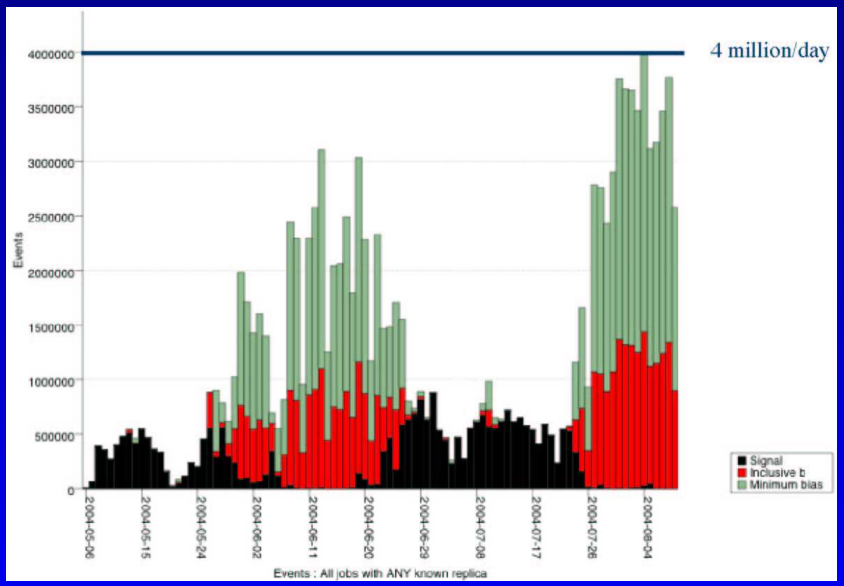


Max-Planck-Institut für Kernphysik

# DC04 Summary

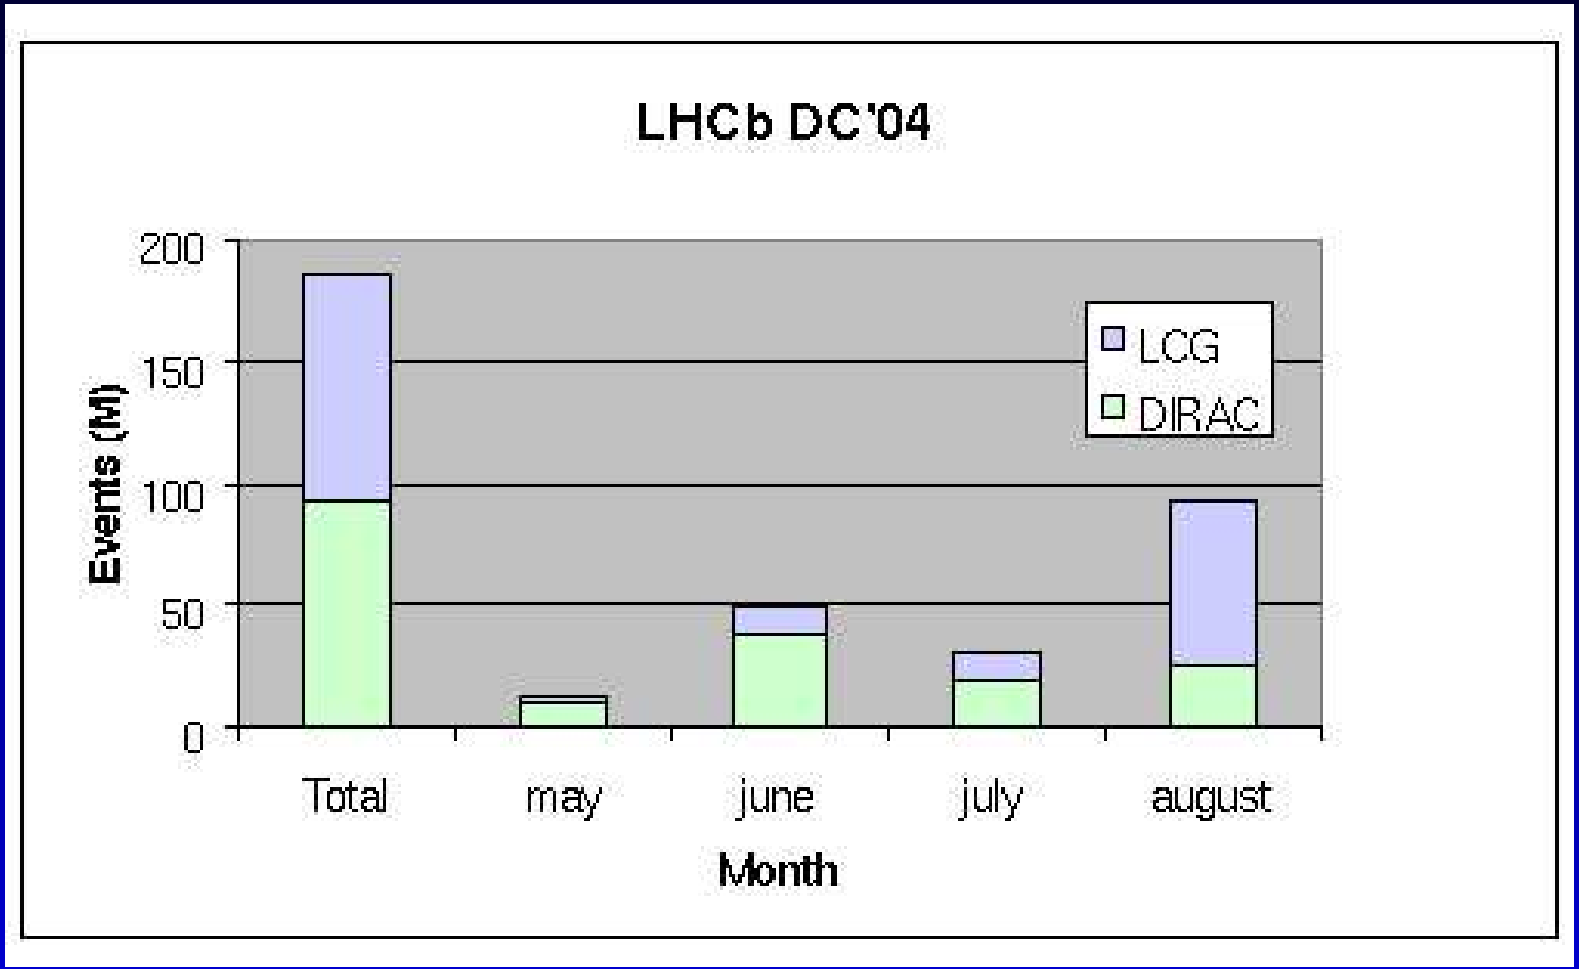

# DC04 Summary

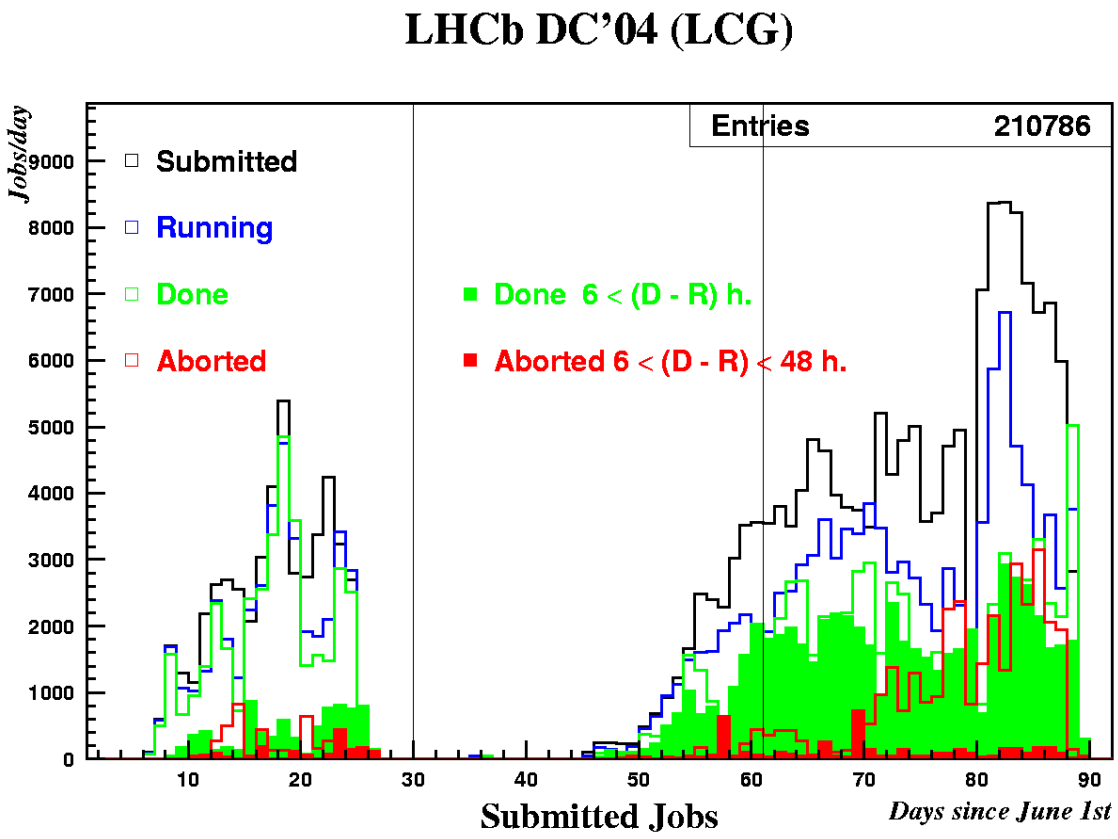

# DC04 Summary

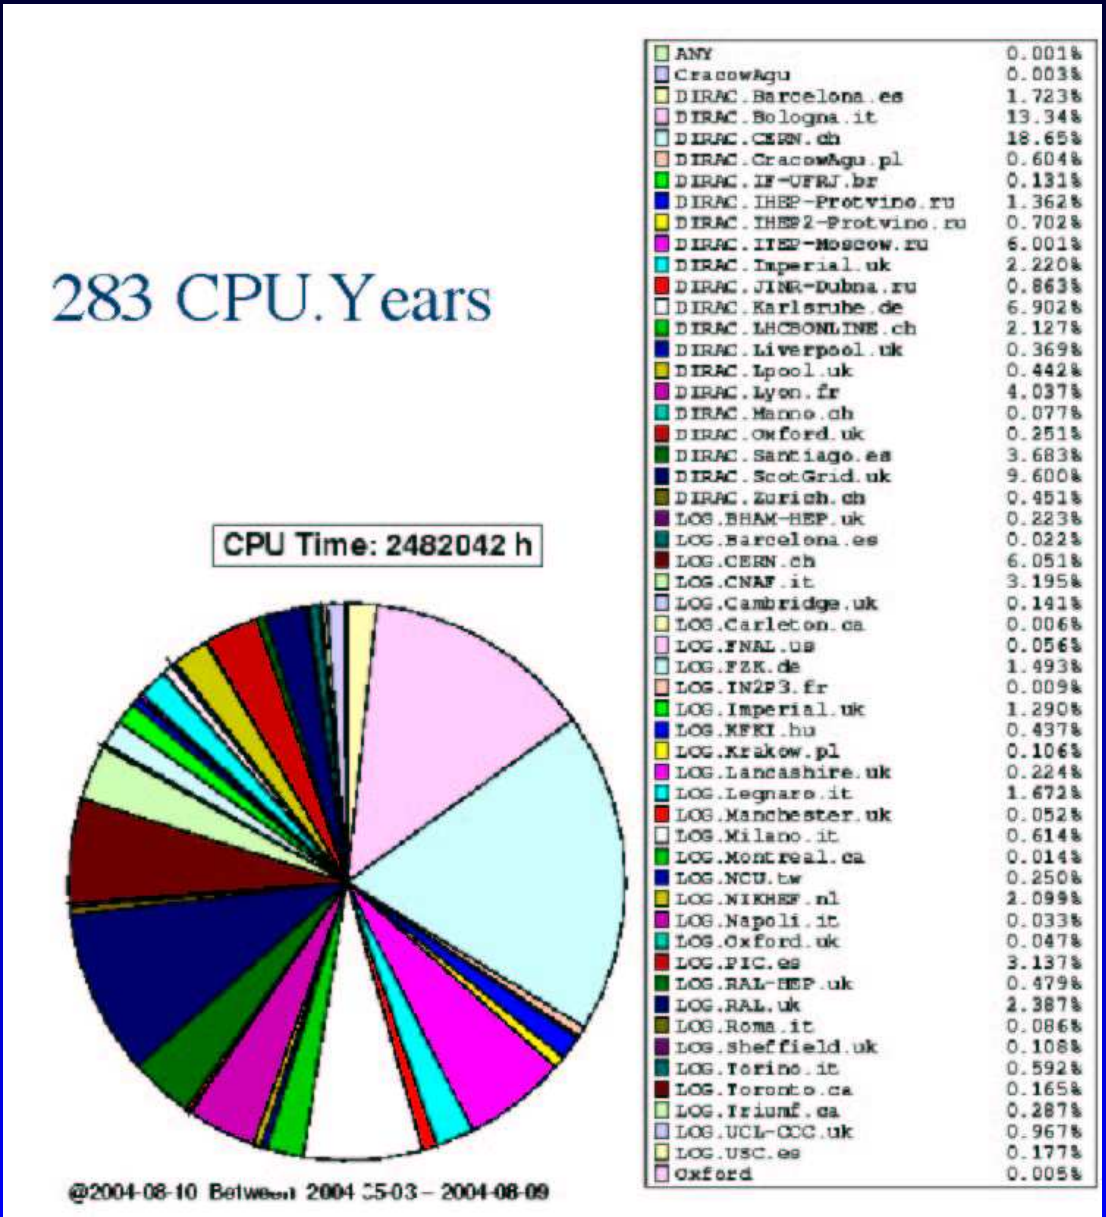


Max-Planck-Institut für Kernphysik

# DC04 Summary

LCG Job Statistics

| | Jobs | Jobs(k) | %Sub | %Remain |
|---|---|---|---|---|
| Submitted | 210786 | 211 | 100.00% | |
| Cancelled | 25669 | 26 | 12.18% | |
| Remaining | 185117 | 185 | 87.82% | 100.00% |
| Aborted (not Run) | 37157 | 37 | 17.63% | 20.07% |
| Running | 147596 | 148 | 70.02% | 79.73% |
| Aborted (Run) | 34228 | 34 | 16.24% | 18.49% |
| Done | 113368 | 113 | 53.78% | 61.24% |
| Retrieved | 113331 | 113 | 53.77% | 61.22% |

| | Jobs | Jobs(k) | %Sub | %Remain |
|---|---|---|---|---|
| Done | 113368 | 113 | 53.78% | 61.24% |
| Done (CPU > 6h) | 69296 | 69 | 32.88% | 37.43% |
| Aborted (CPU > 6h) | 9971 | 10 | 4.73% | 5.39% |

Output Sandbox Analysis: 69K Successful jobs

LHCb Accounting: 81K Successful jobs

Max-Planck-Institut für Kernphysik

# Conclusions

- DC03 was very successful
- DIRAC works! (“PULL” vs. “PUSH”)
- DC04: many hiccups/problems solved
- include running on LCG (non-trivial)

To do:

- skimming of data
- incorporate distributed analyses
- employ more gridtools

Max-Planck-Institut für Kernphysik