# Overview of Grid Computing within the BaBar Experiment

T. Adye, C.A.J. Brew, G. Castelli, F.F. Wilson
(Rutherford Appleton Laboratory – United Kingdom)

D. Andreotti, E. Luppi
(INFN Sezione di Ferrara - Italy)

A. Khan
(Brunel University – United Kingdom)

D. Bailey, R. Barlow, J. Werner
(Manchester University – United Kingdom)

ISGC 2007 - International Symposium on Grid Computing 2007
March 26 - 29, 2007, Taipei, Taiwan
"Application – High Energy Physics II" session

BABAR




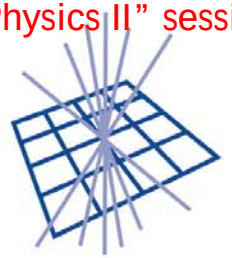

GridPP
UK Computing for Particle Physics

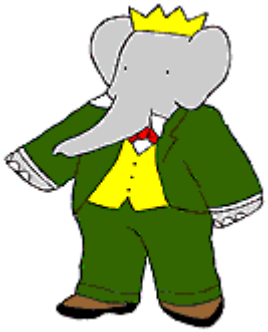




# Outline

- BaBar experiment.
- Computational activities.
- LCG-Grid issues and solutions.
- Conclusions.

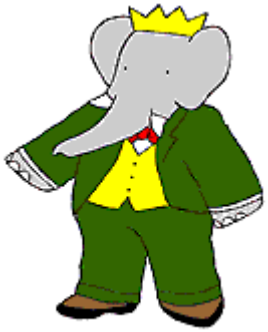
TM & © Nelvana

# In the beginning…


…and not like this, where matter and antimatter have annihilated?

Perhaps matter and anti-matter behave differently…
No anti-matter in universe? Where has it gone?
BaBar at SLAC (California) and Belle at KEK (Japan) are looking at the differences between matter and antimatter.

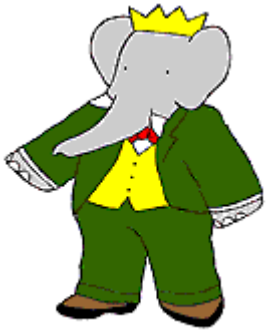




# The BaBar detector at SLAC

10 Countries 77 Institutions ~522 Physicists

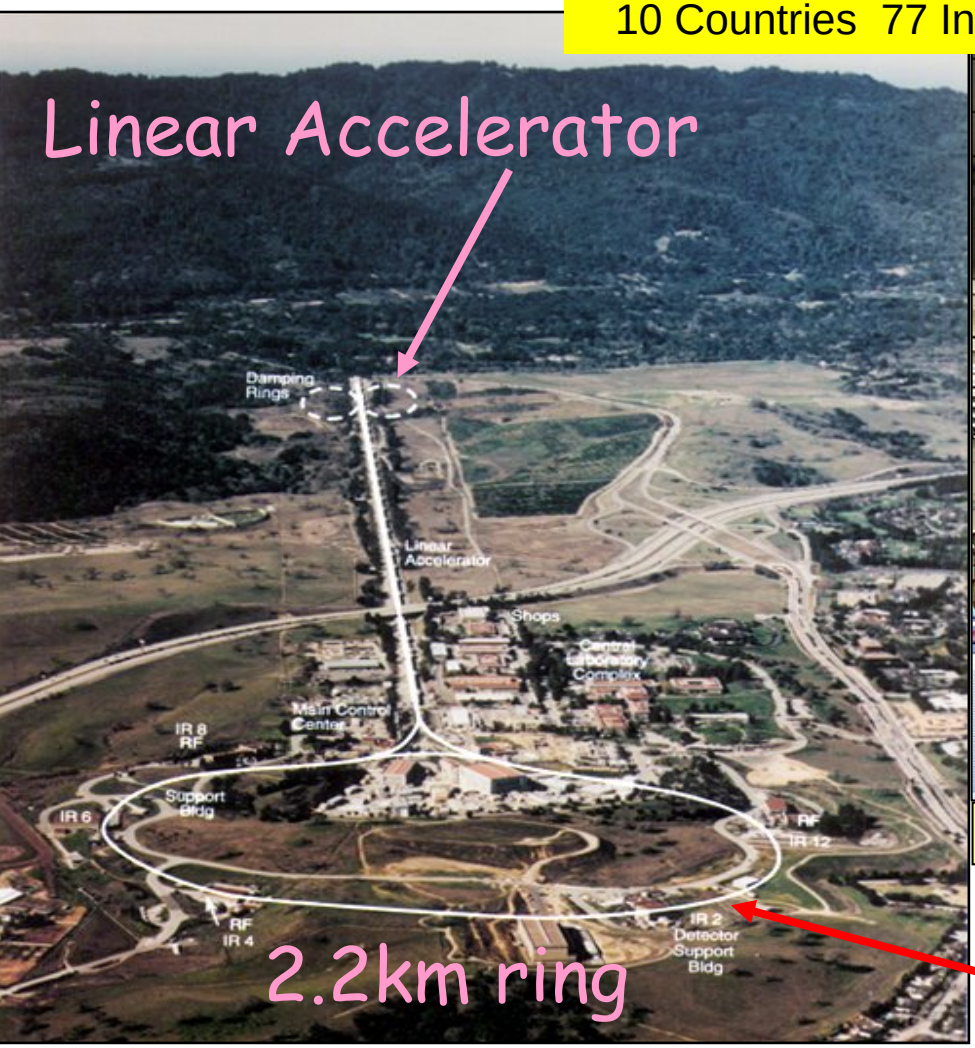


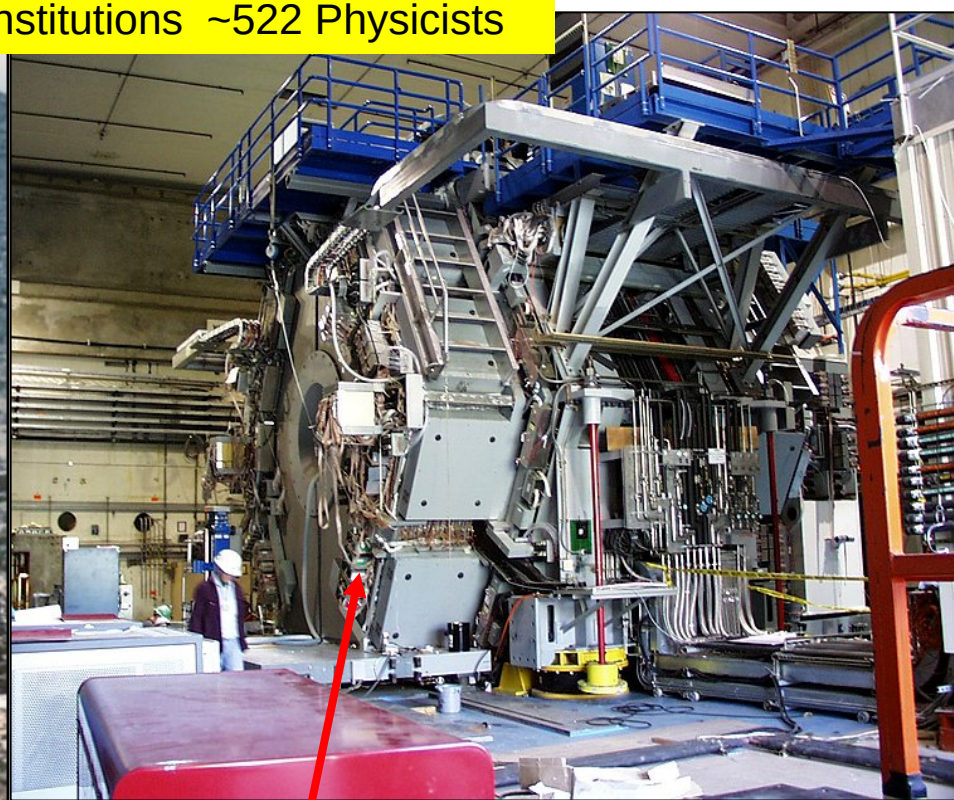

May 1999: first events recorded by BaBar

BaBar

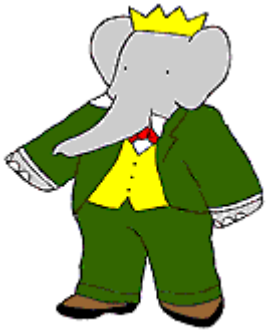


# What happens?

- Bunches of $10^{10}$ electrons and positrons collide together 250 million times per second.
- The collision energy is tuned to produce events: pairs of B and anti-B-meson particles.
- They decay in 1.6 picoseconds ($10^{-12}$ sec).
- We measure the results of their decays and look for differences between B and anti-B-meson.
- We store on disk about 5 events per second.

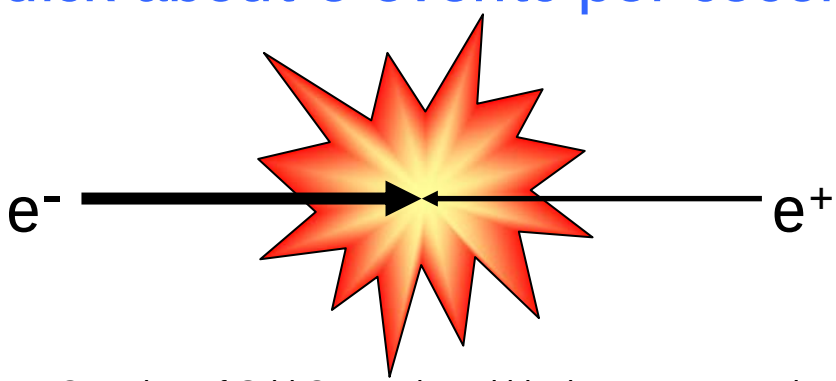

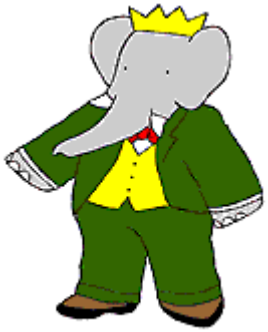


# Data

- Since starting in May 1999, the experiment has recorded ~ 4.5 billion events (about 20 times previous experiments).

- We currently have 971 TB of reconstructed data (~60% simulation).

- In 2007/8 we aim to take as much data as in the previous 7 years.

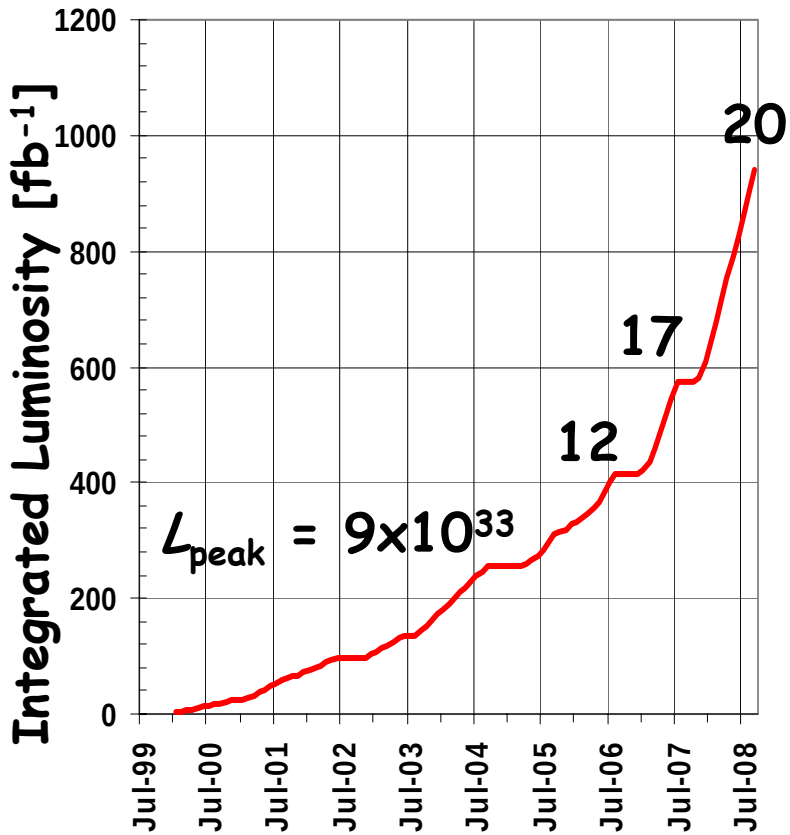

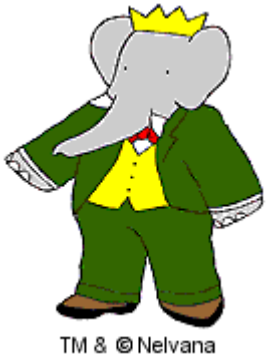


# Computing issues

Four compute intensive steps:

1. Reconstruction of real data taken at SLAC.
   - Sent to dedicated farms in Padova.
   - Results available at SLAC within 48 hours.
2. Production and reconstruction of simulated data.
   - Various farms at external institutes.
   - Simulation is gridified and in production.
   - To analyse the events real data are compared with simulated data, we need much more simulated than real data (~4 times).
3. Production of pre-selected analysis dataset.
   - Dedicated farms at SLAC, INFN Padova and Karlsruhe.
   - Gridified, on-going in production now.
   - Process real and simulated data to select ~200 sub-samples defined by the BaBar physics analysis working groups.
     - Much quicker to run over pre-selected analysis dataset than full data sample.
     - Pre-selected analysis includes physics analysis code and saves the results, so CPU time spent in it is regained many times over.
4. User analysis.
   - Farms at SLAC and large “Tier 1” sites.

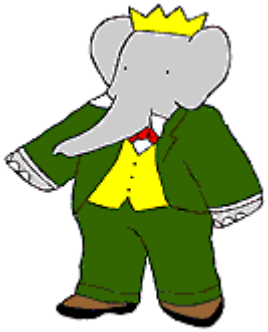


# Distributed data flow

- What all this means in terms of computing?

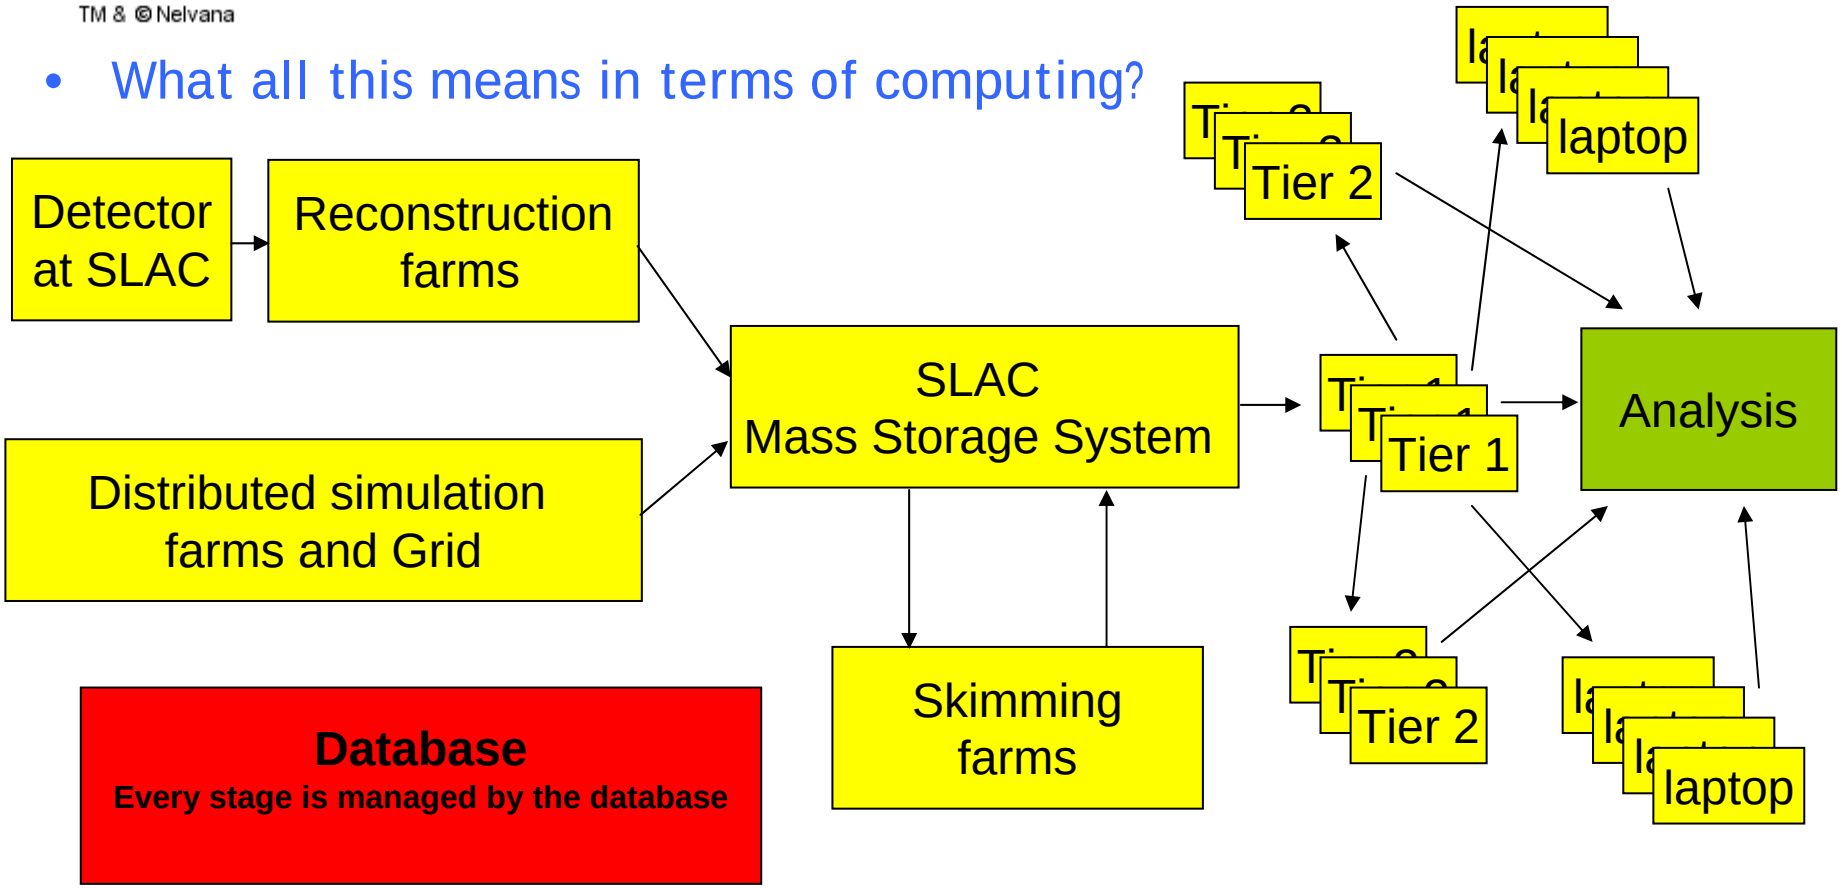

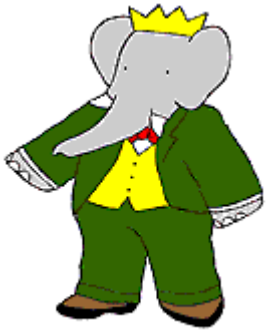


# Job statistics

| | Simulation data production | Production of pre-selected analysis dataset | Data analysis |
|---|---|---|---|
| **Data input rates per job** | negligible | 15.3 kB/s | 50 kB/s |
| **Data output rates per job** | 1 kB/s | 80 kB/s | negligible |
| **Events per second per job** | 0.185 | 2.11 | 1-10 |
| **Events per job** | 2000 events | 100000 events | 100000 events |
| **Memory requirements (RAM) per job** | 0.7 GB | 2 GB | 0.2 GB |
| **Number of open files** | 10 | 220 | 5 |
| **Transfer rates T2->T1** | 0.5 MB/s | 10 MB/s | not applicable |
| **Failures rates on the Grid** | 7% | 11% | not yet tested |
| **Simultaneous jobs in UK** | 1000 | 1000 | 800 |
| **Executable** | defined | defined | variable, undefined |
| **Required input dataset** | 0.3 TB | 5 TB | 100 TB |

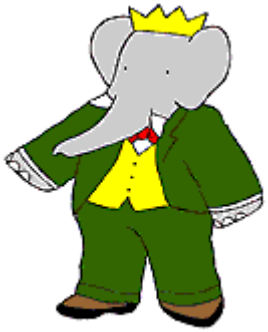


GridPP
UK Computing for Particle Physics

# Simulation data production

**To analyse the events real data are compared with simulated data, we need much more simulated than real data (~4 times).**

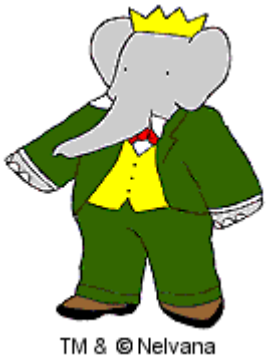


# Life of a BaBar Grid job

Resource Broker

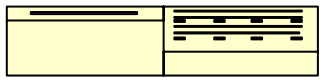

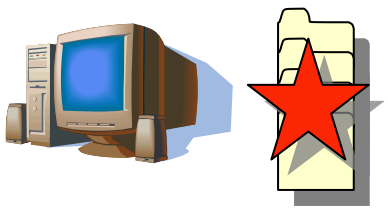

- Job preparation
- Submit
  - Pack
  - Dispatch
- Execute
  - Unpack
  - Run
  - Copy to SE
- Recover/Unpack
- Merge
- Export

Xrootd data server

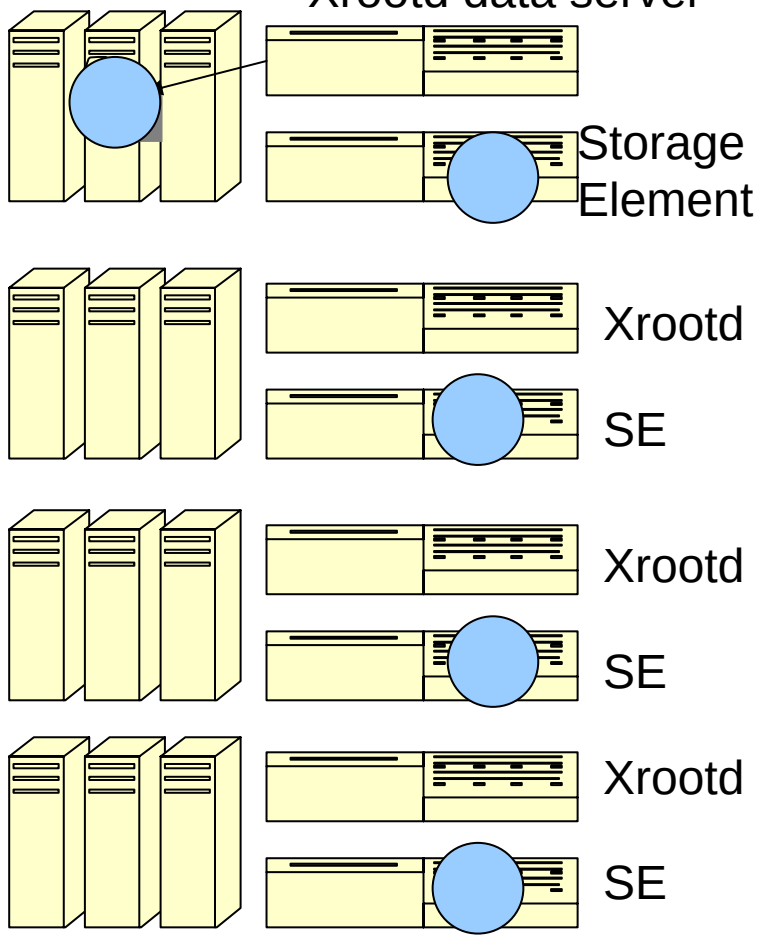

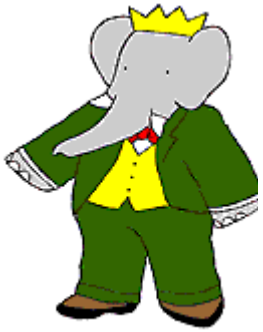


# Grid simulation production rates

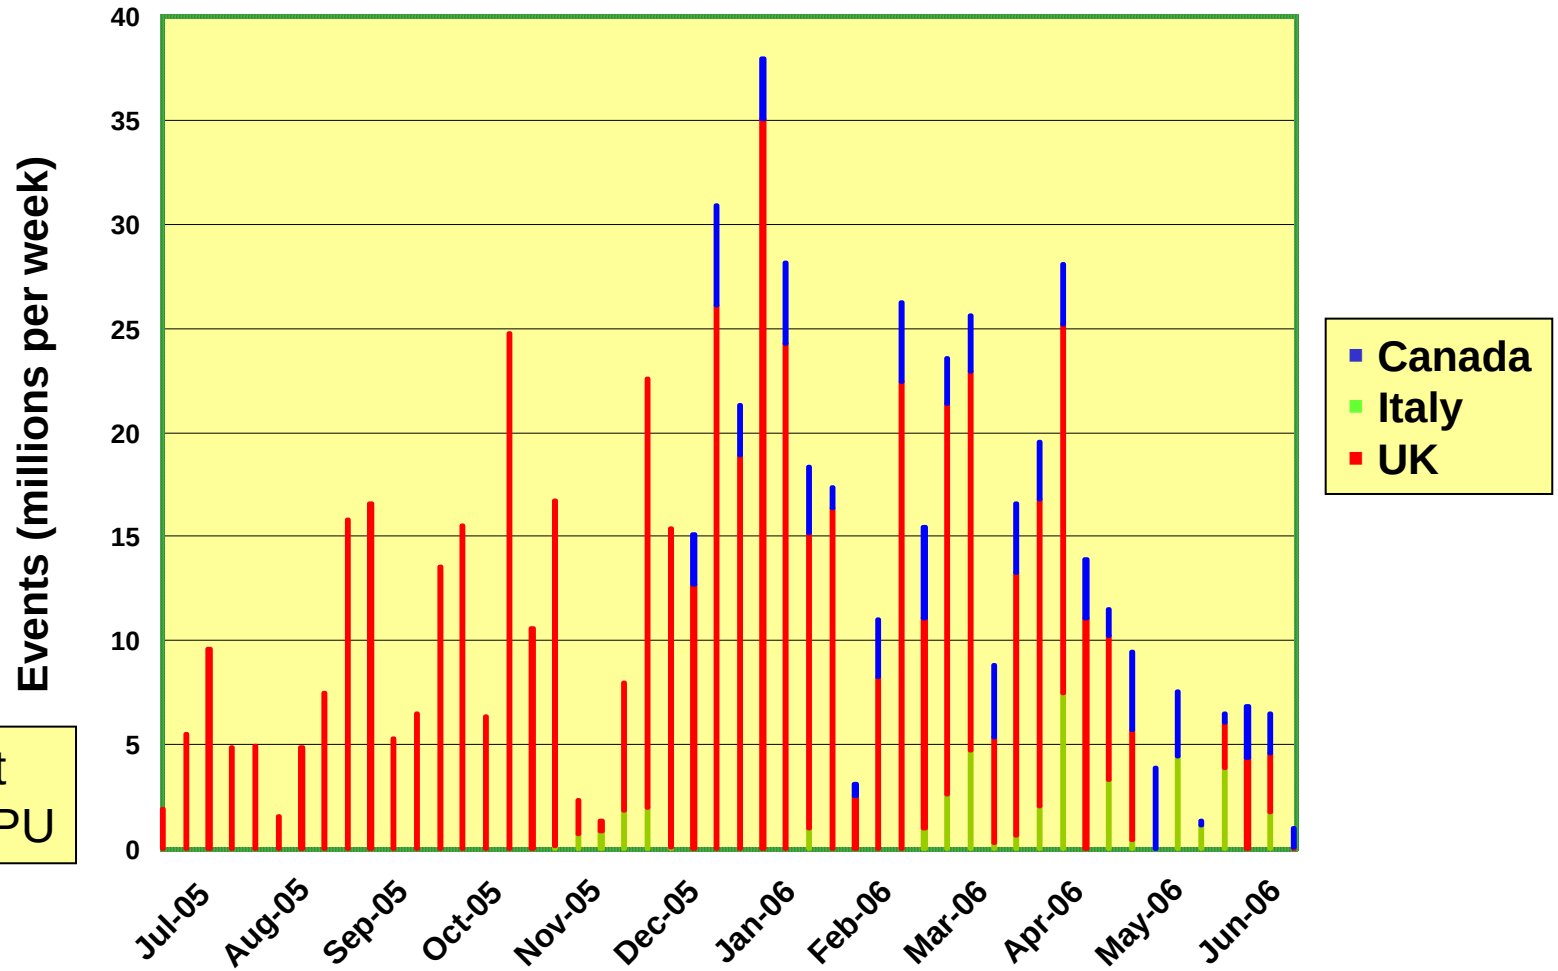


1 event
= 5s CPU

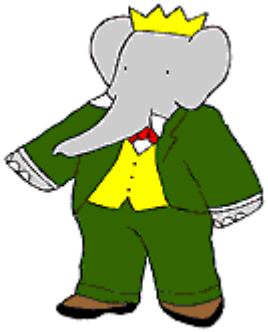


# Production of pre-selected analysis dataset

**Process real and simulated data to select ~200 sub-samples defined by the BaBar physics analysis working groups.**

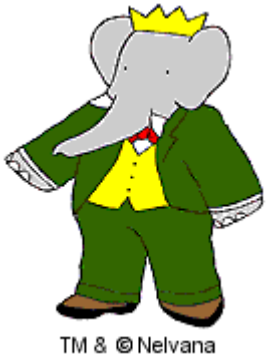


# Life of a BaBar Grid job

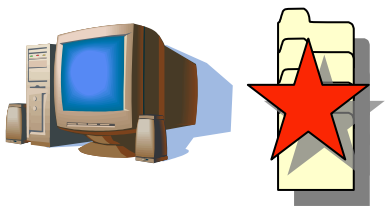

Resource Broker

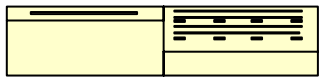

- Job preparation
- Submit
  - Pack
  - Dispatch
- Execute
  - Unpack
  - Run
  - Copy to SE
- Recover/Unpack
- Merge
- Export

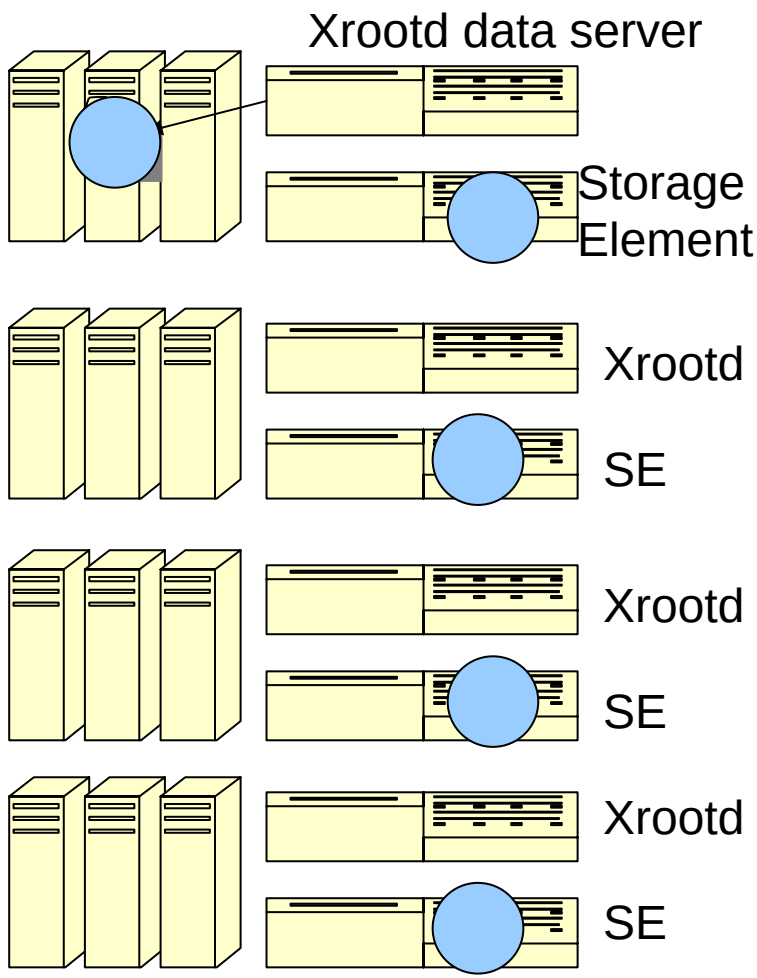

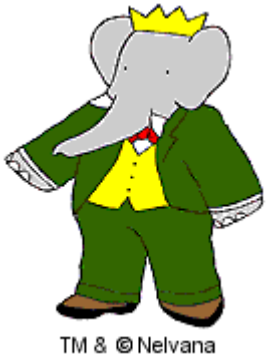


# Production of pre-selected analysis dataset: overview

- We can use the tools and experience gained by porting the production of simulated data to the grid to enable BaBar to use Grid resources for one of the other compute intensive steps: production of pre-selected analysis dataset.
- Remember table (page 9):
  - input dataset 5 TB (it was 0.3 TB for simulation data production).
  - output is 5-6 times the input (it is 80 times less for simulation data production).
- Greater care in the control system and in the matching together of all the pieces
  - we need to know 100% what's happening
  - we need quality assurance
- New version of the control software has been developed with the Grid as batch system in mind.
- Plan is to start running production at Manchester Tier 2 (2000 CPUs, 500 TB, fair share with other experiments).

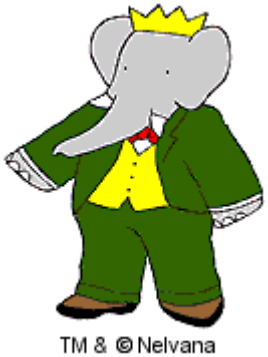


# Status of pre-selected analysis dataset production

- The BaBar software has been adapted to run on the Grid.
- Thousands of validation jobs have been submitted and run successfully and their output recovered, merged and sent to SLAC.
- 100 parallel jobs for the moment, target 1000.
- We will move into production in the next two weeks.

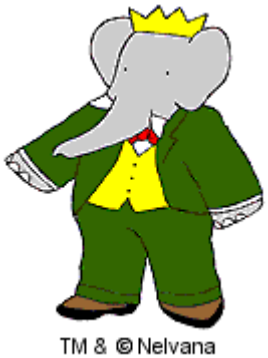


# Past and present Grid issues

- Unique experiment services:
  - BaBar jobs needed access to an Objectivity database.
    - It needed to be deployed at each site we used.
  - This dependency has now been removed.

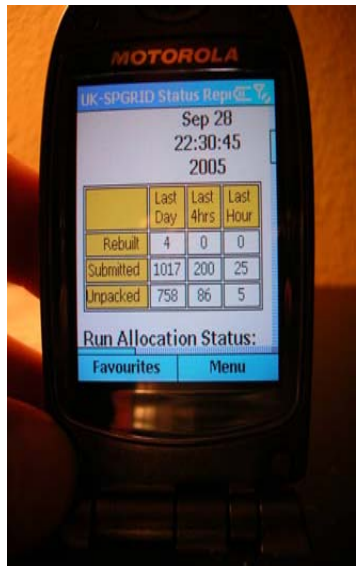


- Monitoring:
  - Unreliability and slowness.
  - Improved when job status information published via R-GMA.
- Grid service reliability:
  - Not possible to monitor them, i.e. RB status: you cannot find in advance if it didn't work.
  - Workarounds, i.e. using multiple RBs.
- High load testing:
  - System didn't scale very well.
  - Adapted the number of jobs to the particular farm.

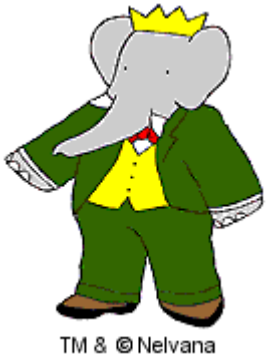


# and more issues

- Site support:
  - Separation of experiments from site organization did cause:
    - Miscommunication.
    - Increased delays in response to questions.
- Increased points of failure:
  - Many more necessary services means more things to go wrong.
- Software releases:
  - Incompatibilities between middleware versions and patchy upgrading by sites caused some problems.
- Documentation:
  - It was initially limited.

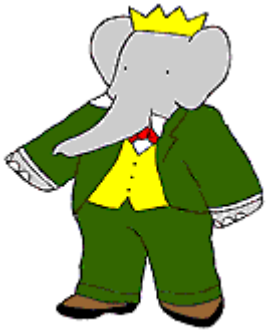


# Conclusions

- This Grid-based model is how BaBar UK will produce its simulated data events in the future, and will be the major part of the pre-selected analysis dataset production in 2007.
- We hope to exploit US Science Grid as well.
- Many of the problems we encountered have improved as the Grid has matured.
- We have demonstrated that Grid technologies are able to aid BaBar in meeting its distributed computing needs.