The *grid* for LHC Data Analysis

ICAP 2006 Conference
Chamonix
5 October 2006

Les Robertson - CERN
LHC Computing Grid Project Leader

The LHC Computing Challenges

1. Data
   - After reduction by triggers and data acquisition filter the raw data will be recorded at 100 – 1,000 MBytes/sec
   - With processed datasets, across the 4 experiments,
     ➔ 15 PetaBytes of new data each year
2. Computation
   - A few thousand users and about 100K of today's fastest PC processors
3. Funding
   - Most of the computing resources will not be at CERN – distributed across ~100 scientific computing centres around the world



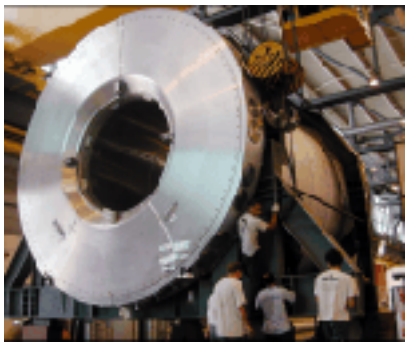
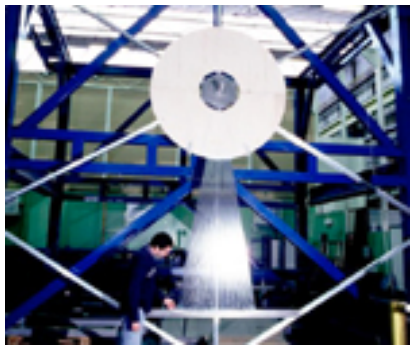

LCG - The LHC Computing Grid

- Purpose
  - Develop, build and maintain a distributed computing environment for the storage and analysis of data from the four LHC experiments
    - Ensure the computing service
    - … and common application libraries and tools

- Collaboration – experiments and ~100 regional computing centres
- Phase I – 2002-05 - Development & planning
- Phase II – 2006-2008 – Deployment & commissioning of the initial services

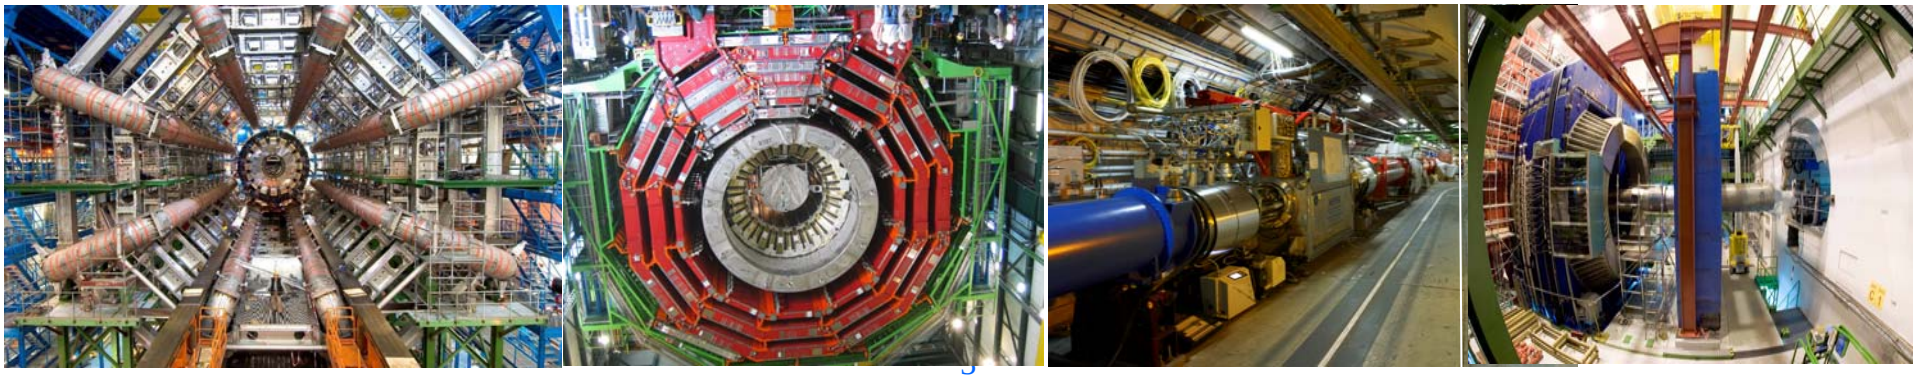

Distribution of Computing Services

Summary of Computing Resource Requirements

All experiments - First full year

From LCG TDR - June 2005

	CERN	*All Tier-1s*	*All Tier-2s*	*Total*
CPU (MSPECint2000s)	25	56	61	142
Disk (PetaBytes)	7	31	19	57
Tape (PetaBytes)	18	35		53

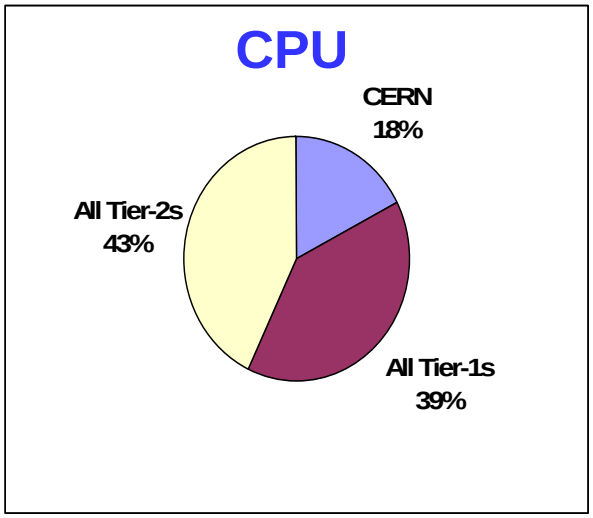

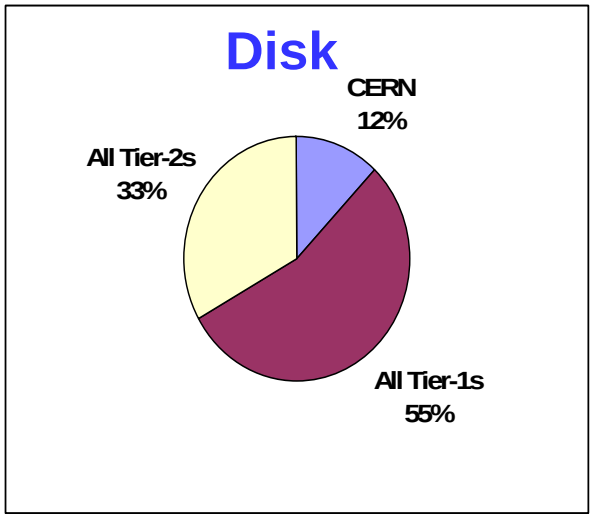

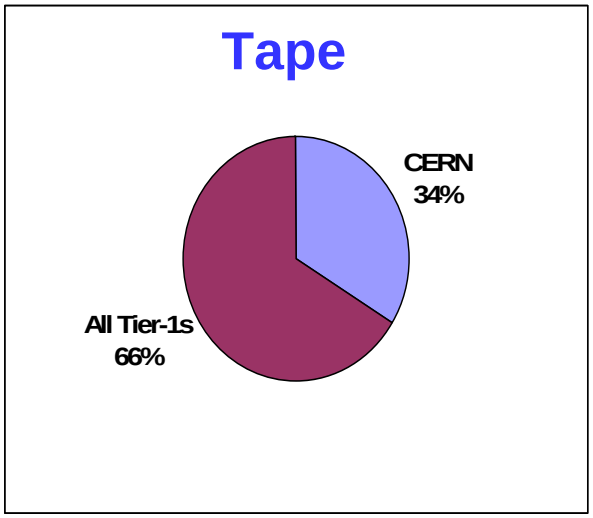

LCG Service Hierarchy

Tier-0 – the accelerator centre

- Data acquisition & initial processing
- Long-term data curation
- Distribution of data → Tier-1 centres

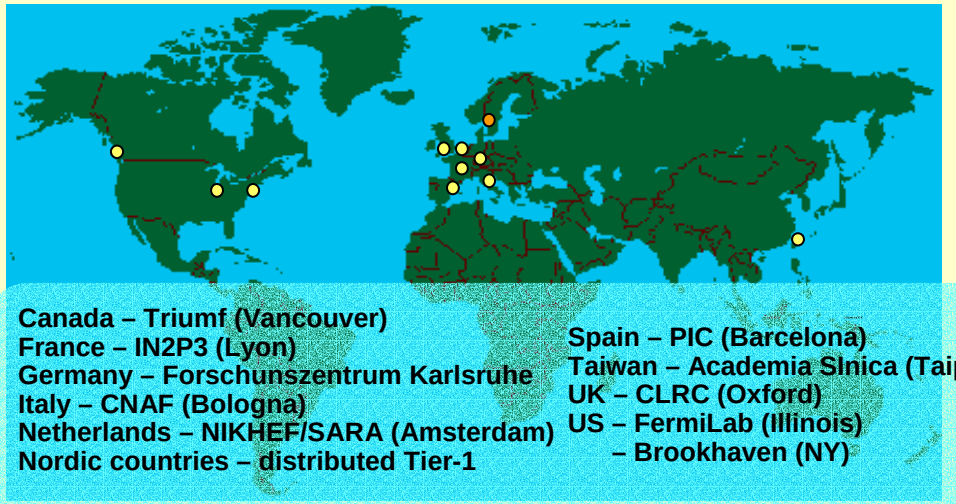

Canada – Triumf (Vancouver)
France – IN2P3 (Lyon)
Germany – Forschunszentrum Karlsruhe
Italy – CNAF (Bologna)
Netherlands – NIKHEF/SARA (Amsterdam)
Nordic countries – distributed Tier-1
Spain – PIC (Barcelona)
Taiwan – Academia SInica (Taipei)
UK – CLRC (Oxford)
US – FermiLab (Illinois)
– Brookhaven (NY)

Tier-1 – "online" to the data acquisition process → high availability

- Managed Mass Storage – → grid-enabled data service
- Data-heavy analysis
- National, regional support

Tier-2 – ~100 centres in ~40 countries

- Simulation
- End-user analysis – batch and interactive

Using the *grid*

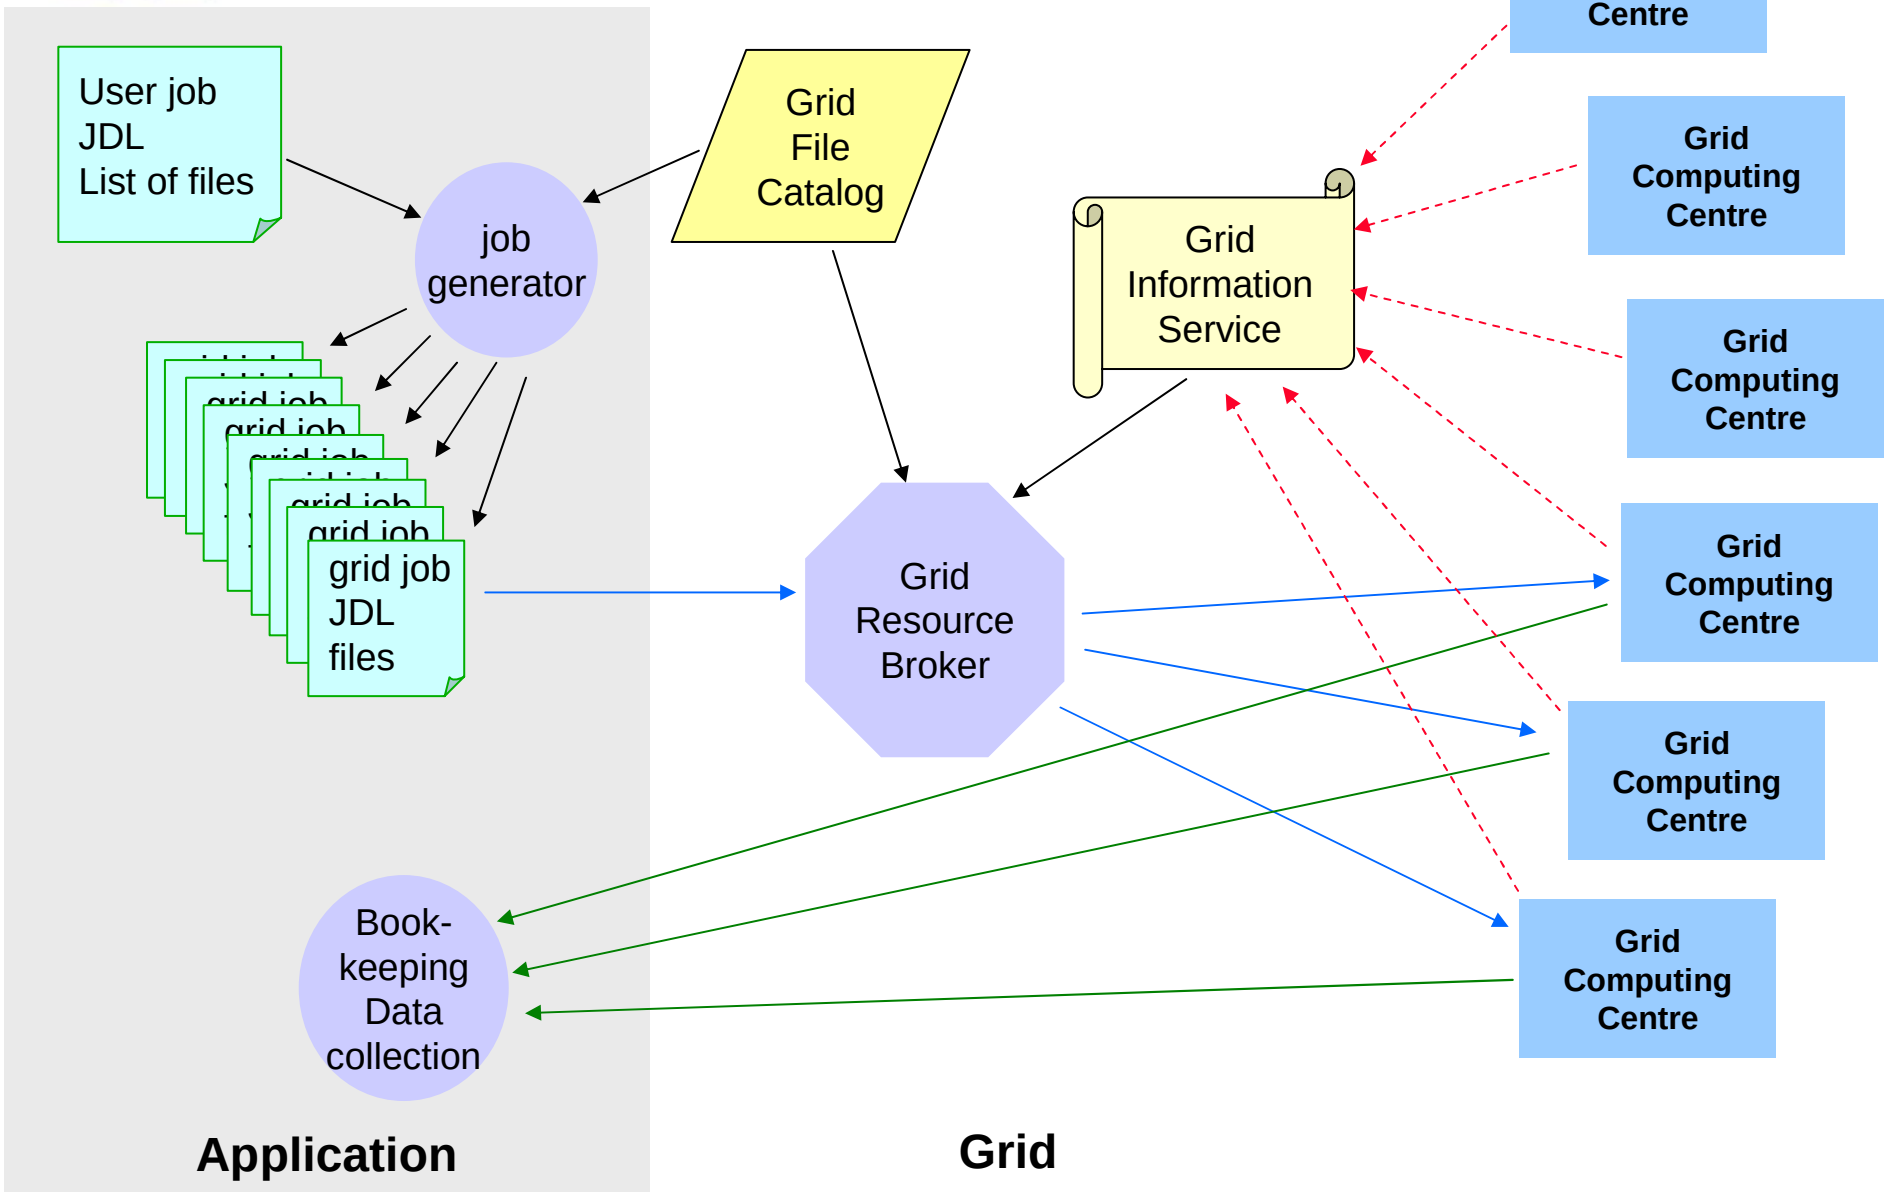

LCG Baseline Grid Services

- Information system
- Security framework
  - Security model
  - Certificate infrastructure
  - Policy and processes

- Virtual Organisation management services
  - User registration
  - Authentication & authorisation

- Grid operations
  - Duty operator responsibility cycles around 5 large sites
  - Monitoring tools
  - Availability/reliability test framework
  - Accounting
  - User support

- Resource Broker
- Grid Catalog
- Applications software installation

grid

- Compute Resource Services
  - Interface to site
  - Job submission, monitoring & control

- Information publisher
- Authentication and authorisation services

- Storage management services
  - Standard interface (SRM)
- Data transfer services
  - Basic file transfer
  - Scheduling and recovery service

- Grid catalogue services
- Database services
- Data management toolkit
- POSIX-like I/O service

site

HEP analysis using the grid

- Quite different from large simulation
  - Emphasis on low latency, variable work-load
  - Multi-user concurrent activities
  - Lots of data
  - Still an activity for *power users*
- But:
  - Lots of end-users (physicists in the collaborations)
  - Ultimately most analysis will need the grid
- User environment for the *grid*:
  - All collaborations are building their own interfaces/environments
  - Purpose:
    - *Facilitate*: job preparation, job generation, job monitor/control, bookkeeping and data collection
    - *Hide*: error recovery and details of the computing environments (local batch system, grid)
  - Most of the experiments use at least 2 grid infrastructures

Example: ATLAS and LHCb Ganga

GridPP
UK Computing for Particle Physics

Ganga GUI

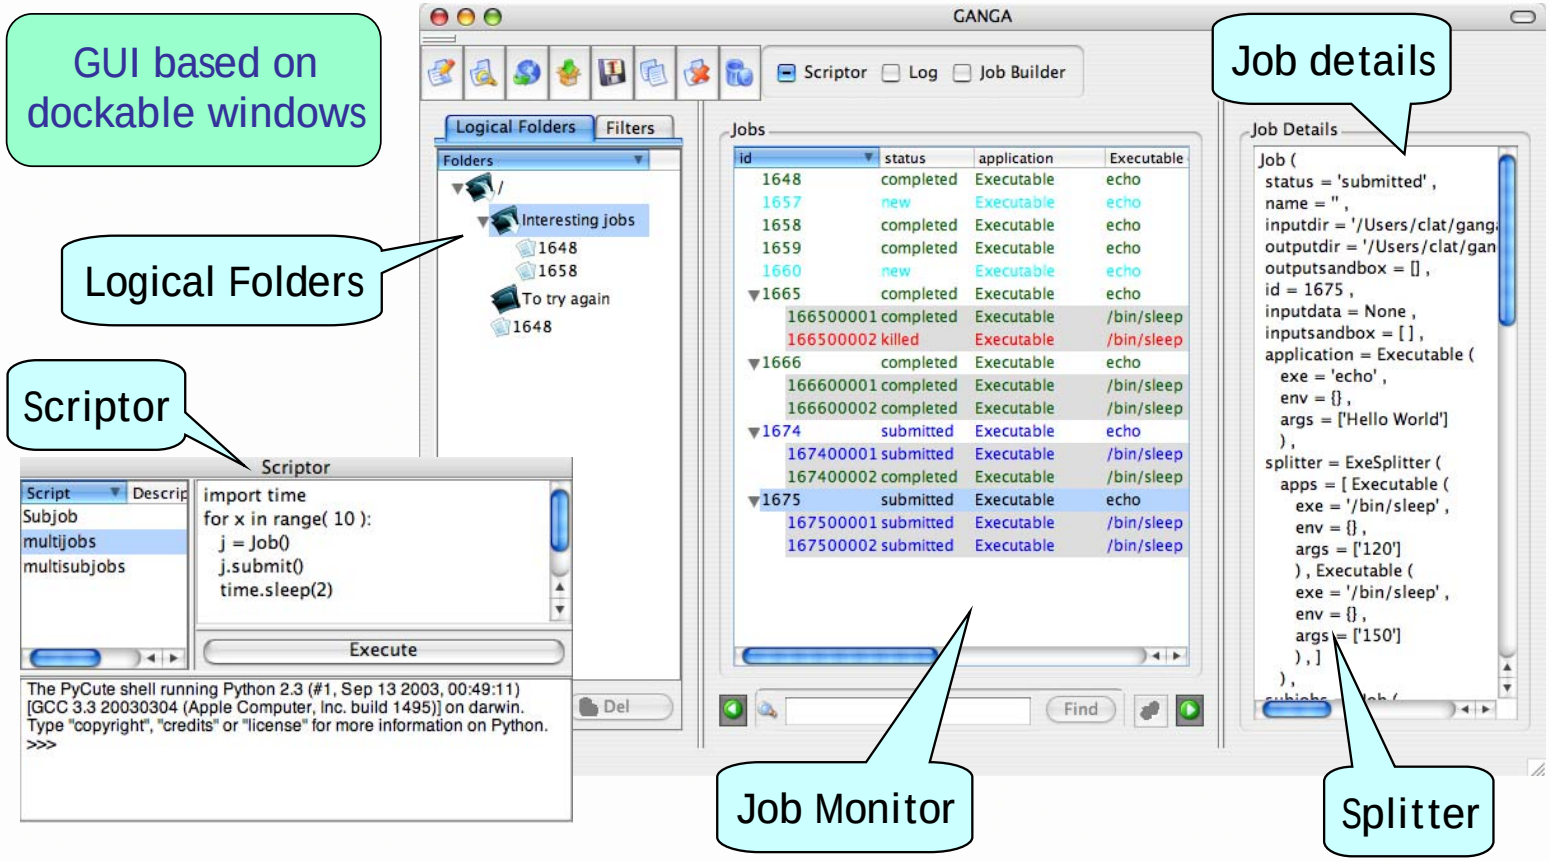

WLCG depends on two major science grid infrastructures ….

EGEE - Enabling Grids for E-Science
OSG - US Open Science Grid

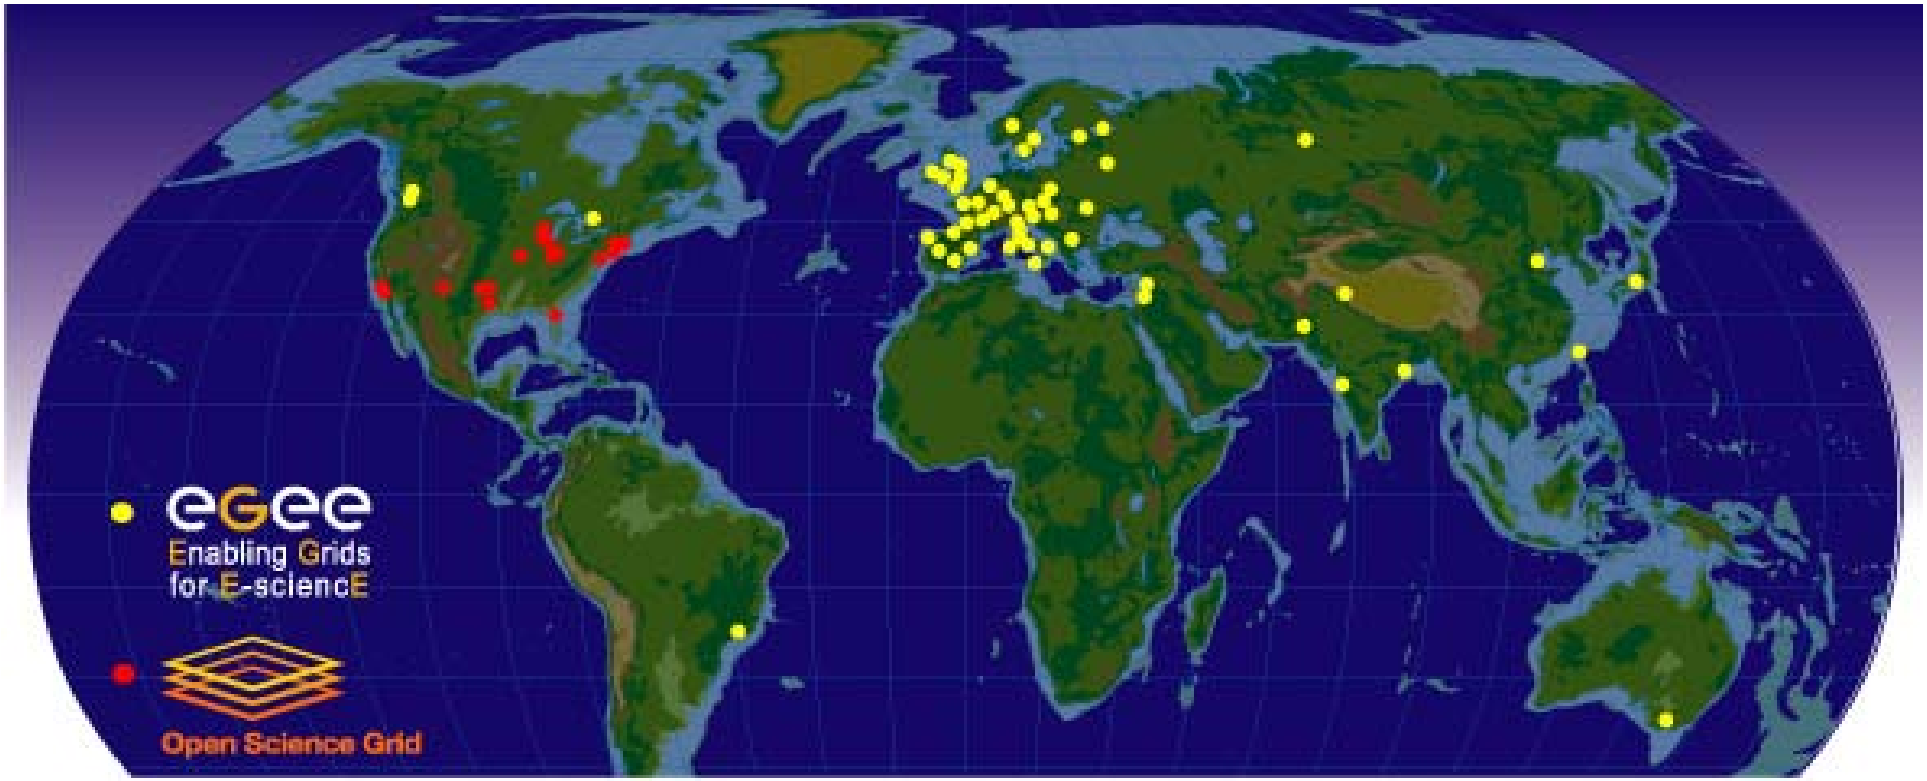

A map of the worldwide LCG infrastructure operated by EGEE and OSG.

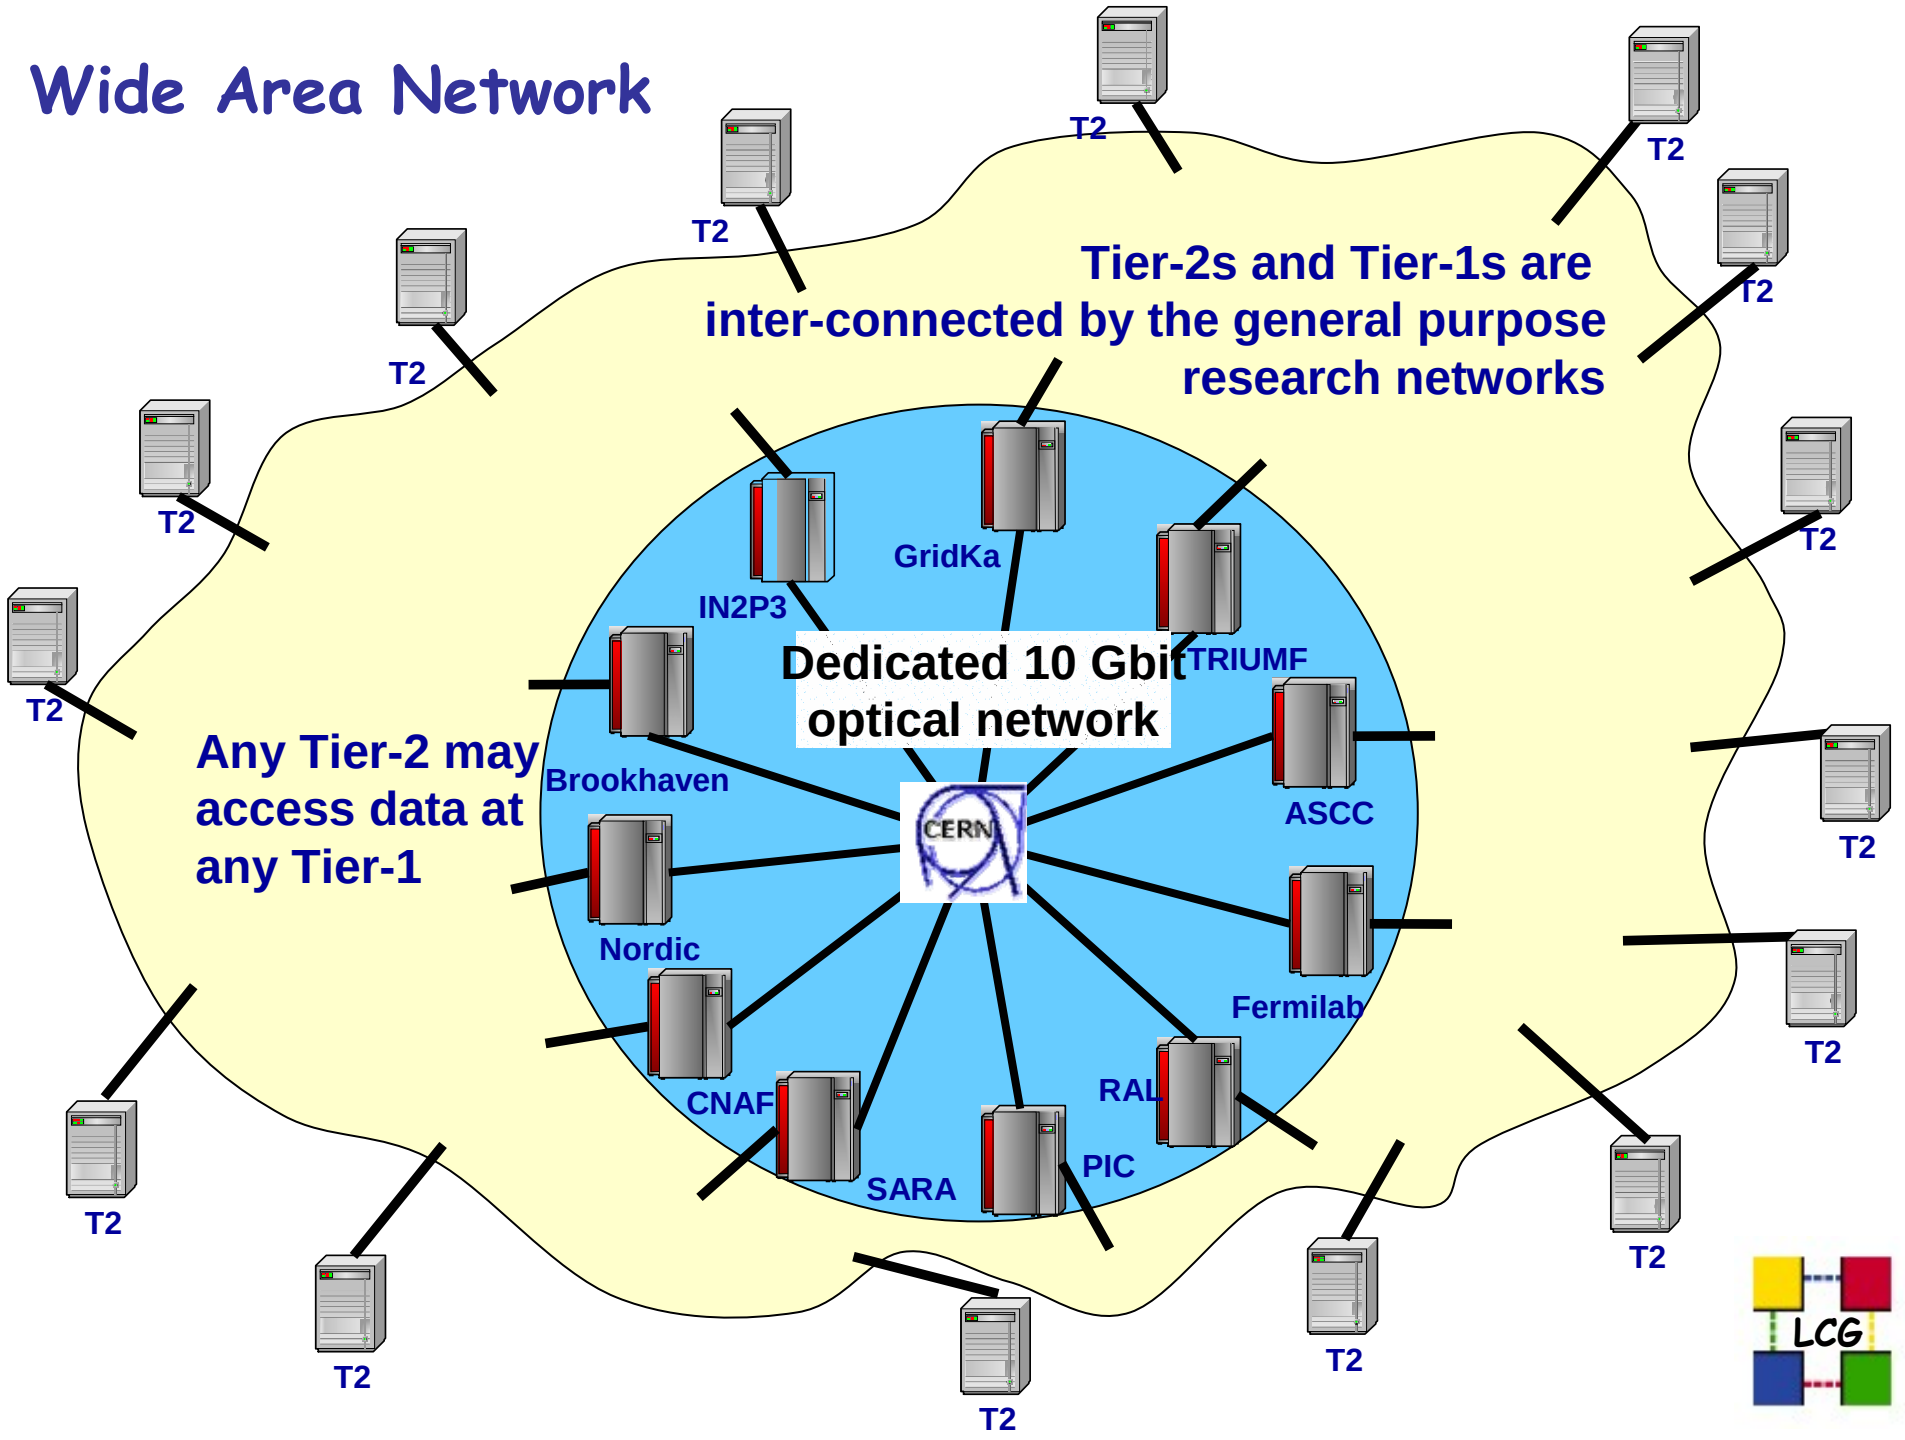
Wide Area Network
Tier-2s and Tier-1s are inter-connected by the general purpose research networks
Any Tier-2 may access data at any Tier-1
Dedicated 10 Gbit optical network
CERN
IN2P3
GridKa
TRIUMF
ASCC
Fermilab
RAL
PIC
SARA
CNAF
Nordic
Brookhaven
T2
LCG

Production Grids for LHC

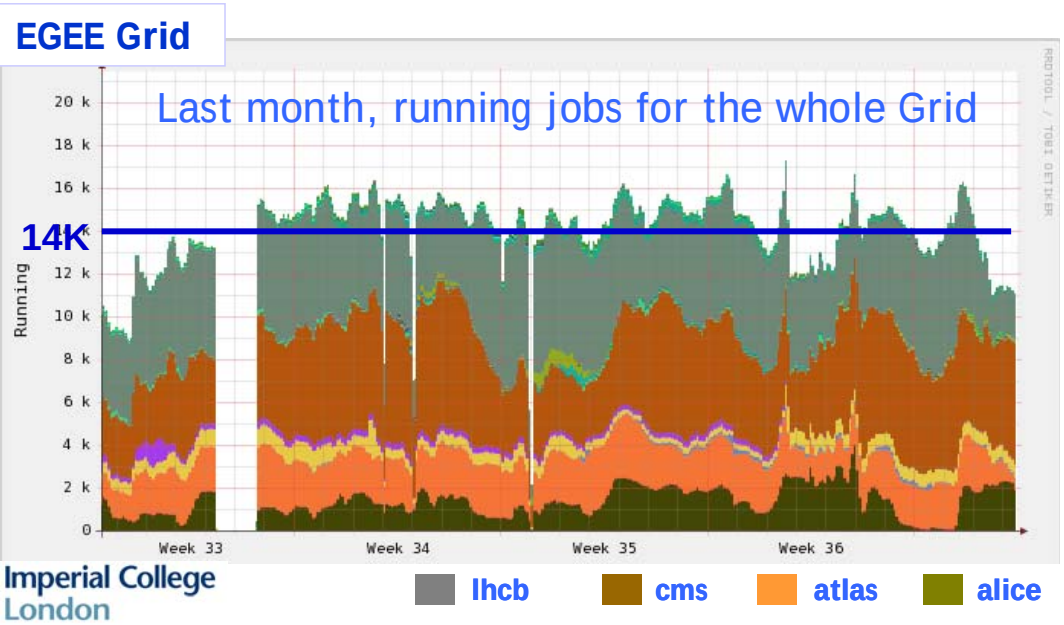

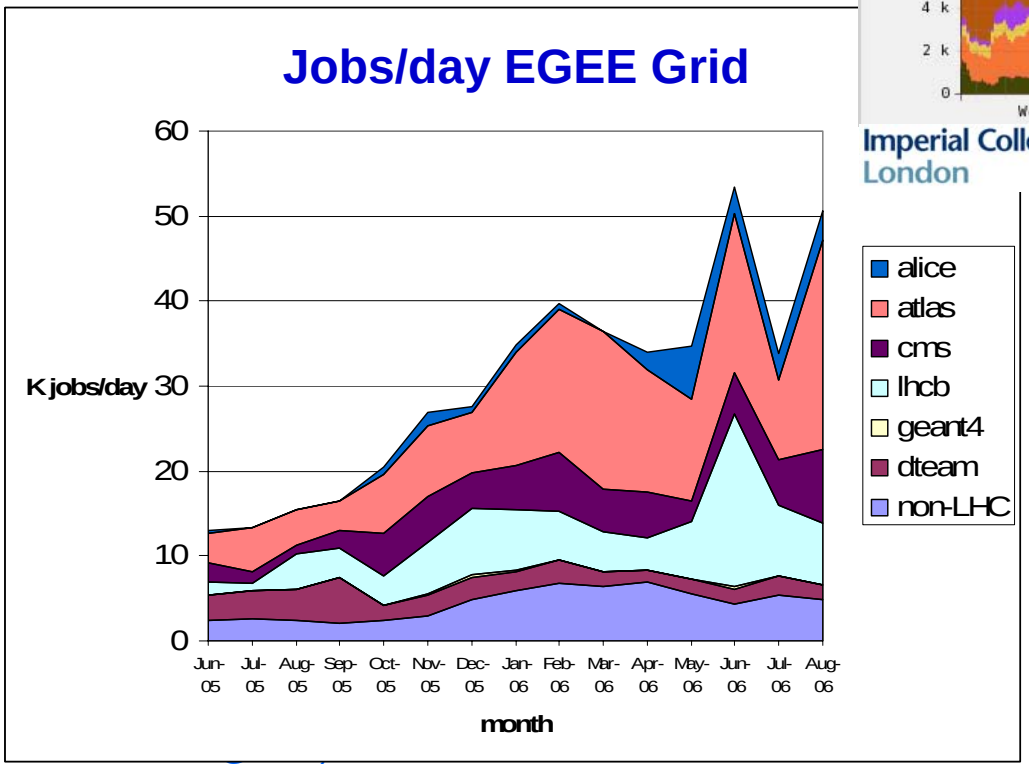

EGEE Grid

- ~50K jobs/day
  ~14K *simultaneous* jobs during prolonged periods

→ ~20% of the 2008 requirement

Data Distribution

- CERN →T1s – April 06 test period
  - "nominal" rate when LHC is operating - 1.6 GB/s – reached – but only for one day
  - Sustained data rate 80% of the target

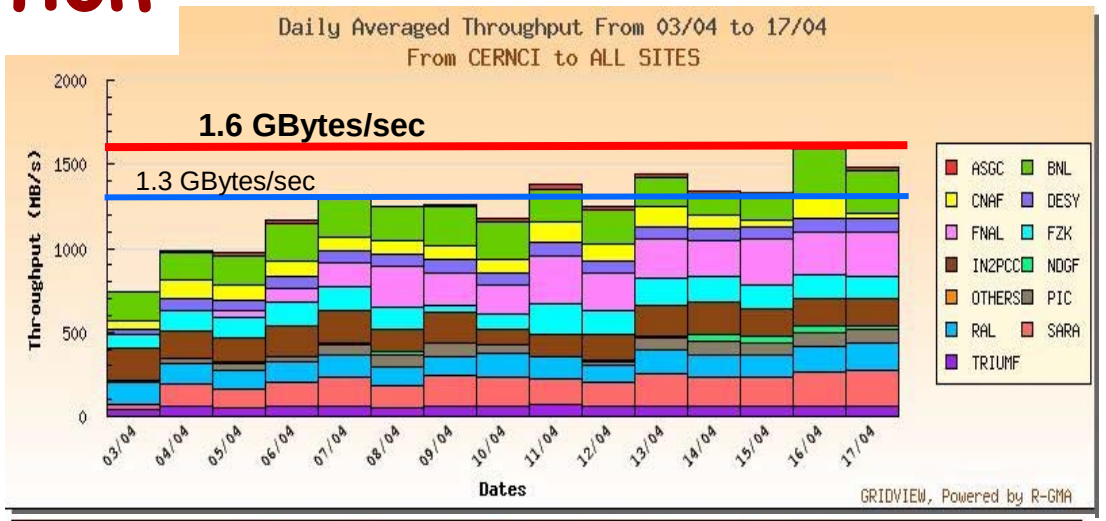

- August experiment-driven transfers (ATLAS and CMS) sustained 50% of the SC4 target under *much more realistic conditions*

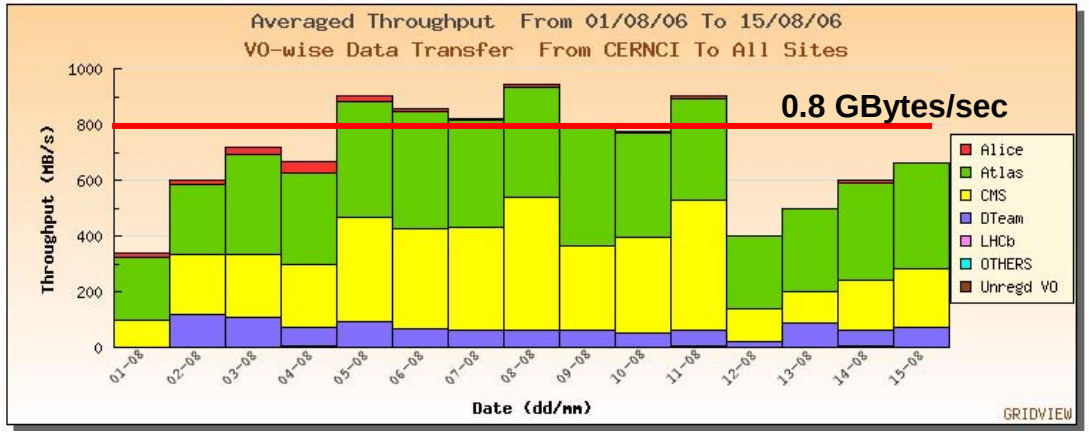

- CMS transferred a steady 1 PByte/month between Tier-1s & Tier-2s during a 90 day period
- ATLAS distributed 1.25 PBytes from CERN during a 6-week period

Commissioning Schedule

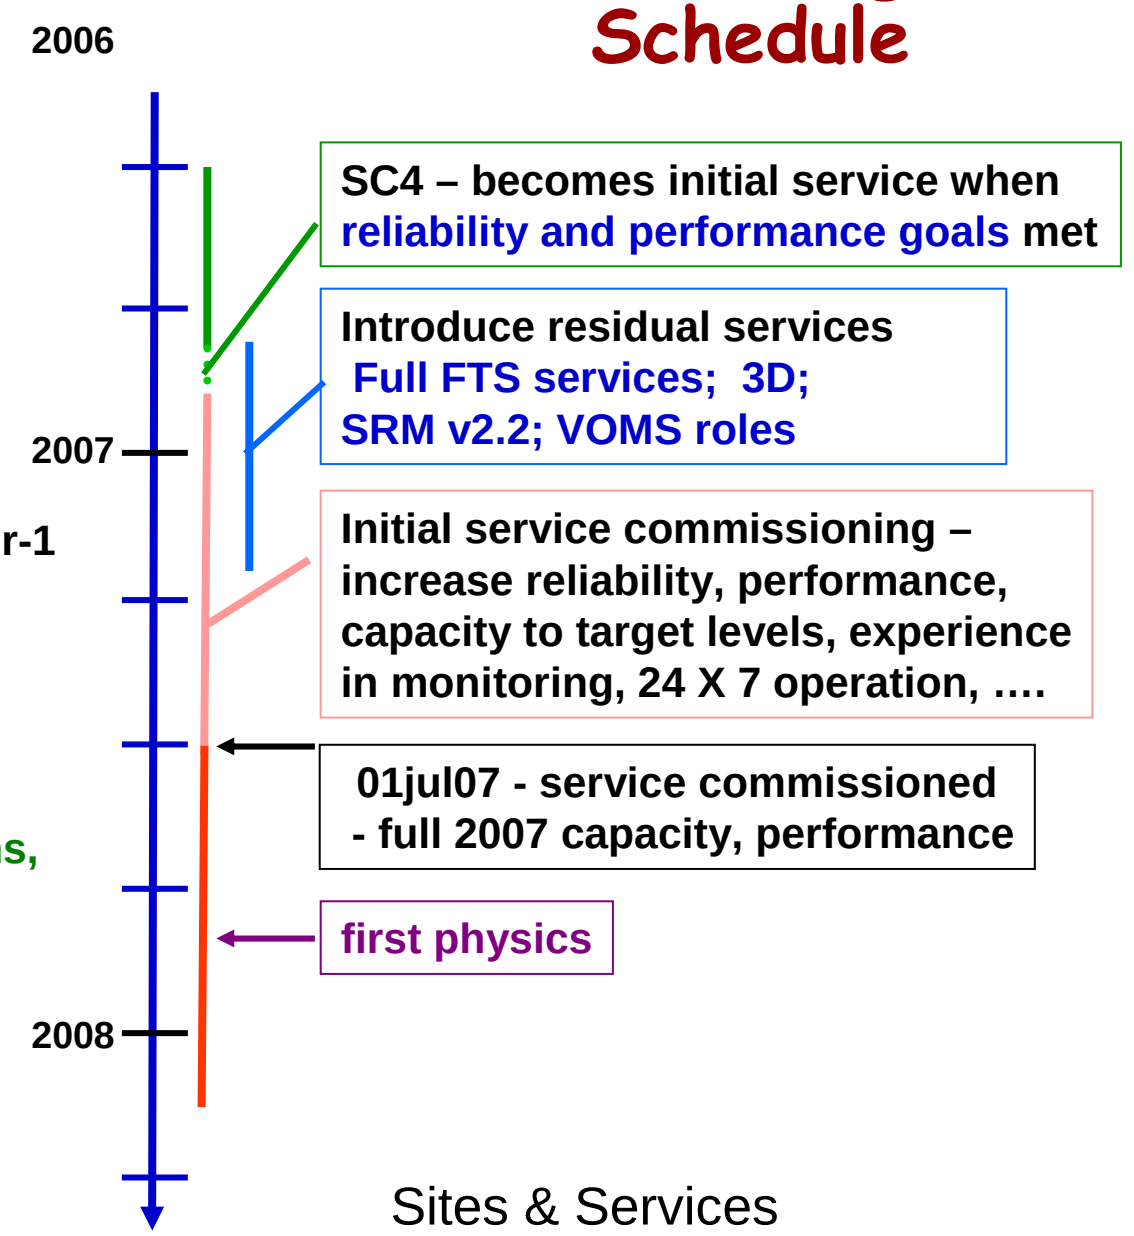

Continued testing of computing models, basic services

Testing DAQ→Tier-0 (??) & integrating into DAQ→Tier-0→Tier-1 data flow

Building up end-user analysis support

Exercising the computing systems, ramping up job rates, data management performance, ….

Experiments

Sites & Services

Applications Running

- **More than 20 applications from 7 domains**
  - High Energy Physics (Pilot domain)
    - 4 LHC experiments
    - Other HEP (DESY, Fermilab, etc.)
  - Biomedicine (Pilot domain)
    - Bioinformatics
    - Medical imaging
  - Earth Sciences
    - Earth Observation
    - Solid Earth Physics
    - Hydrology
    - Climate
  - Computational Chemistry
  - Fusion
  - Astronomy
    - Cosmic microwave background
    - Gamma ray astronomy
  - Geophysics
    - Industrial applications

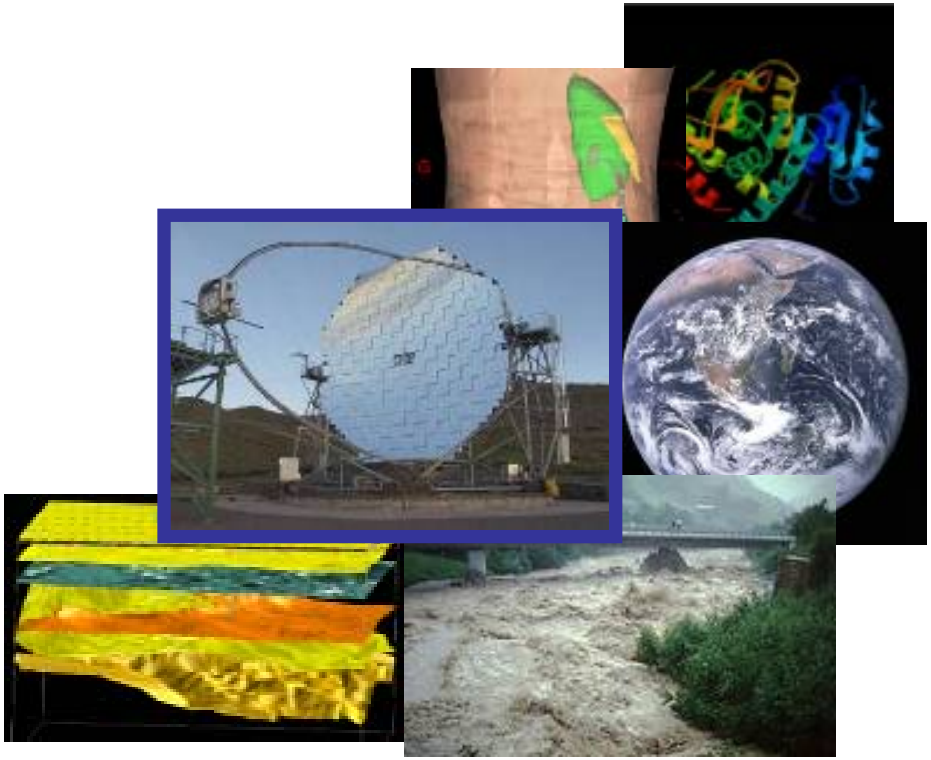

Example from Biomedical applications: Avian Flu

Millions of chemical compounds available in laboratories

High Throughput Screening
2$/compound, nearly impossible

300,000 Chemical compounds: ZINC
Chemical combinatorial library

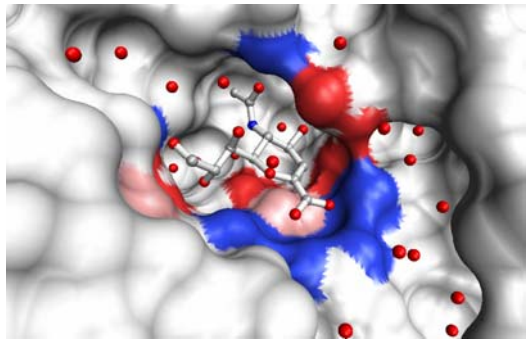

Target (PDB) :
Neuraminidase (8 structures)

Molecular docking (Autodock)
~100 CPU years, 600 GB data

Data challenge on EGEE grid
~6 weeks on ~2000 computers

Hits sorting and refining

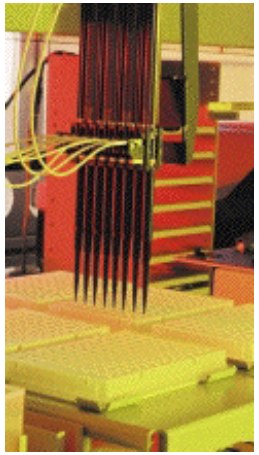

In-vitro screening of 100 hits

Academia Sinica
Genomics Research Center

Another example: the gridCC project (www.gridcc.org)

- Using Grid technology to operate distributed instruments
  - And couple them with grid resources for simulation and analysis
- Pilot applications:
  - Remote operation of an accelerator
  - Meteorology (Ensemble Limited Area Forecasting)
  - Device Farm for the Support of Cooperative Distributed Measurements in Telecommunications and Networking Laboratories
  - Geo-hazards: Remote Operation of Geophysical Monitoring Network

Main GridCC Pilot Applications:
Remote Operation of an Accelerator

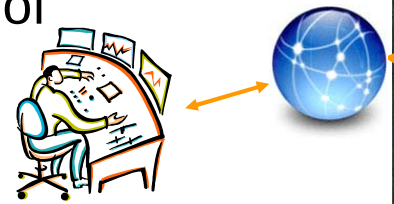

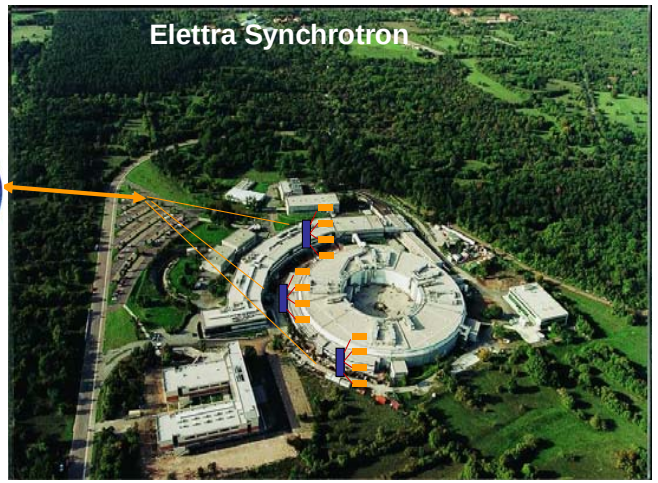

- Different applications
  - LHC Monte Carlo simulation
  - Fusion
  - Drug discovery (like avian flu and malaria)
- Similar characteristics
  - Jobs are CPU-intensive
  - Large number of independent jobs
  - Run by few (expert) users
  - Small input; large output
- & needs
  - Massive CPU needs
  - Data collections for further analyses

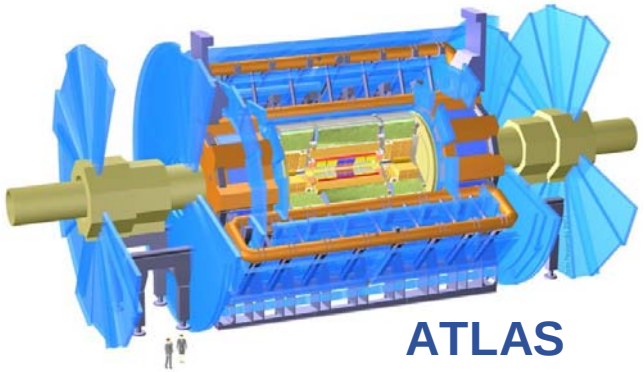

ATLAS

ITER

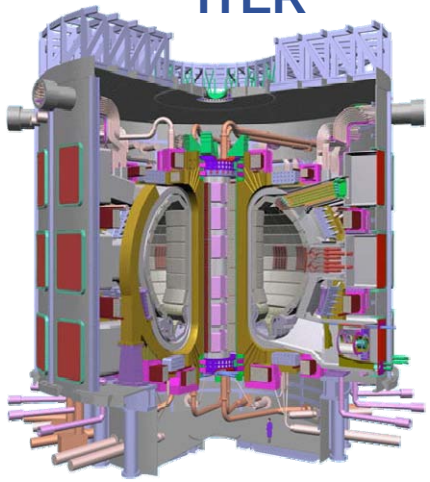

Drug discovery

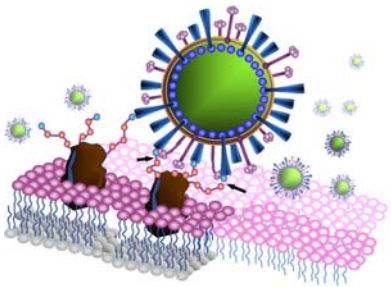

Towards General Science Infrastructures

- EGEE
  - Funded 50% by the EU - 90 partners - 34 countries
  - Funding for people - middleware development, operations, pilot applications, outreach
  - No funding for hardware
  - Many sites outside Europe
  - Current (second) phase ends March 2008!
  - Expectation that there will be a third phase of 2-3 years
- Post EGEE –
  - Early discussions on a long term science grid infrastructure (cf. the GEANT research and education network backbone)
- OSG
  - Organised as a consortium – with some funding ($6M/year for next 5 years) from DoE and NSF for facility development and operations
  - No middleware activity, no funding for hardware

How to get going

- NA4 website (http://egeena4.lal.in2p3.fr/)

- EGEE Conferences and Users' Forums
  - Share your expertise, learn from other users.
  - Be open to collaboration with others
  - Next event: Manchester May 2007

- EGEE website www.eu-egee.org
- OSG website www.opensciencegrid.org

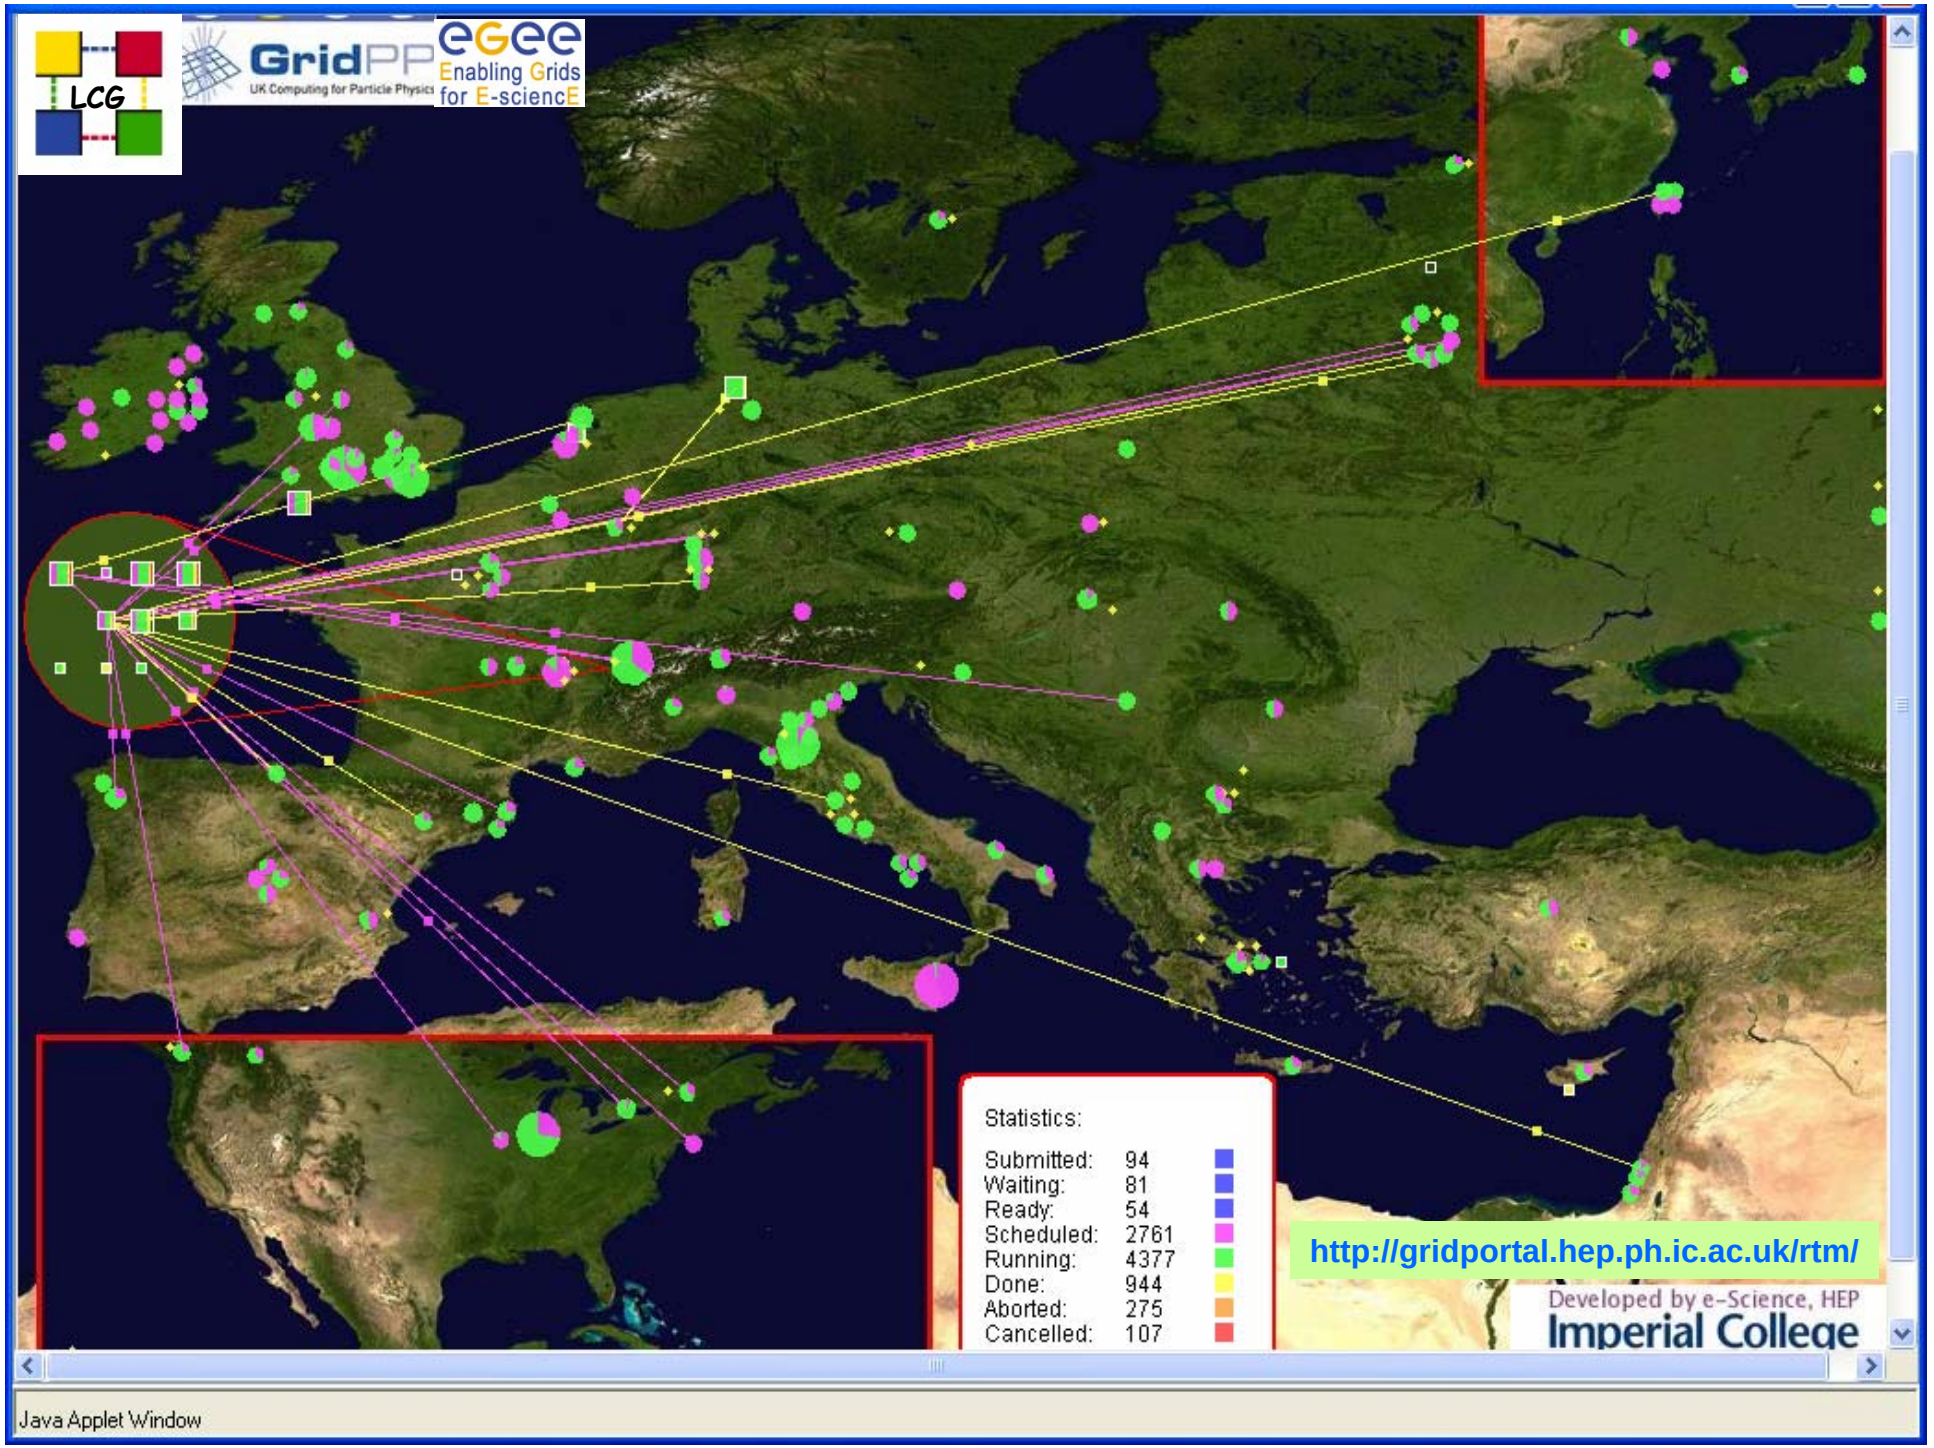
LCG
GridPP
UK Computing for Particle Physics
eGee
Enabling Grids for E-sciencE
Statistics:
Submitted: 94
Waiting: 81
Ready: 54
Scheduled: 2761
Running: 4377
Done: 944
Aborted: 275
Cancelled: 107
http://gridportal.hep.ph.ic.ac.uk/rtm/
Developed by e-Science, HEP
Imperial College
Java Applet Window