EGEE – a worldwide Grid infrastructure

Fabrizio Gagliardi
Project Director EGEE
CERN, Switzerland

IFIC, 6 October 2005

www.eu-egee.org

Information Society

Presentation overview

- **Data intensive science and the rationale for Grid computing**
- **General description of the EGEE project**
- **EGEE operates a production infrastructure:**
  - Operations
  - Middleware
  - Applications
- **Promote and enable international collaboration**
- **The next phase: EGEE-II**

Computing intensive science

- **Science is becoming increasingly digital and needs to deal with increasing amounts of data**

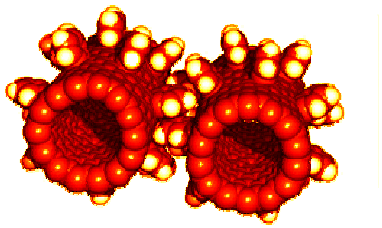

- **Simulations get ever more detailed**
  - Nanotechnology – design of new materials from the molecular scale
  - Modelling and predicting complex systems (weather forecasting, river floods, earthquake)
  - Decoding the human genome

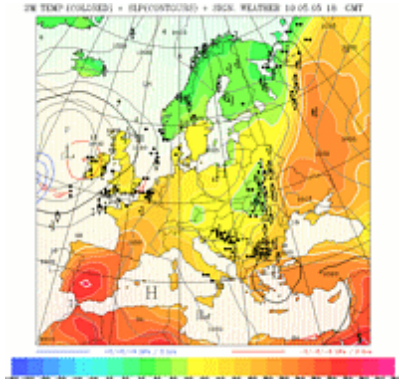

- **Experimental Science uses ever more sophisticated sensors to make precise measurements**
  - → Need high statistics
  - → Huge amounts of data
  - → Serves user communities around the world

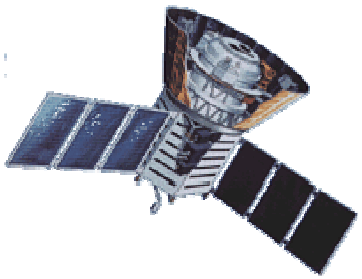

A good example: Particle Physics

- **Large amount of data produced in a few places**
  (CERN, FNAL, KEK…)
- **Worldwide collaborations of computer-savvy scientists**
  (e.g. LHC CERN experiments)
- **Distributed Computing and storage resources**
  owned and managed by many different entities worldwide
- **Large Hadron Collider (LHC) at CERN in Geneva, Switzerland:**
  - One of the most powerful instruments ever built to investigate matter

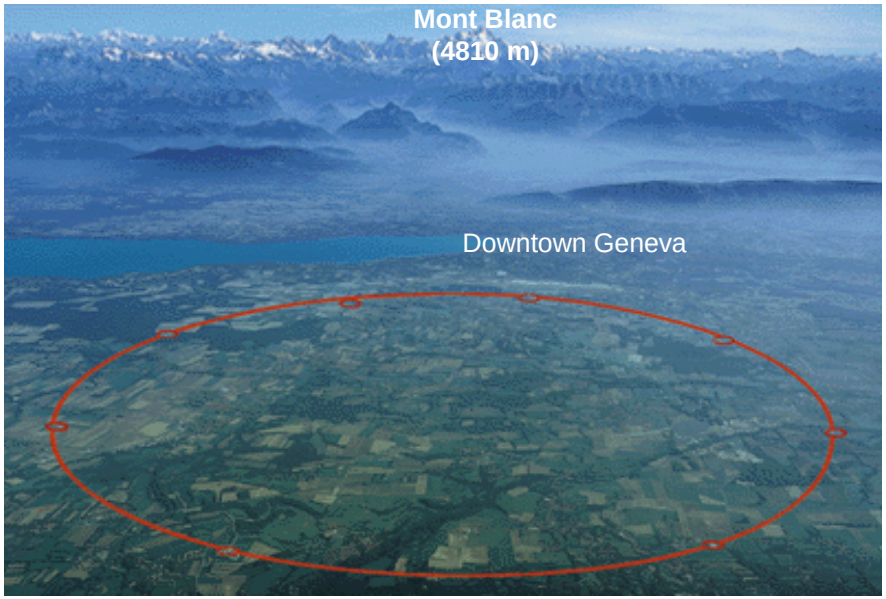

The LHC Experiments

- **Large Hadron Collider (LHC):**
  - four experiments:
    - ALICE
    - ATLAS
    - CMS
    - LHCb
  - 27 km tunnel
  - Start-up in 2007

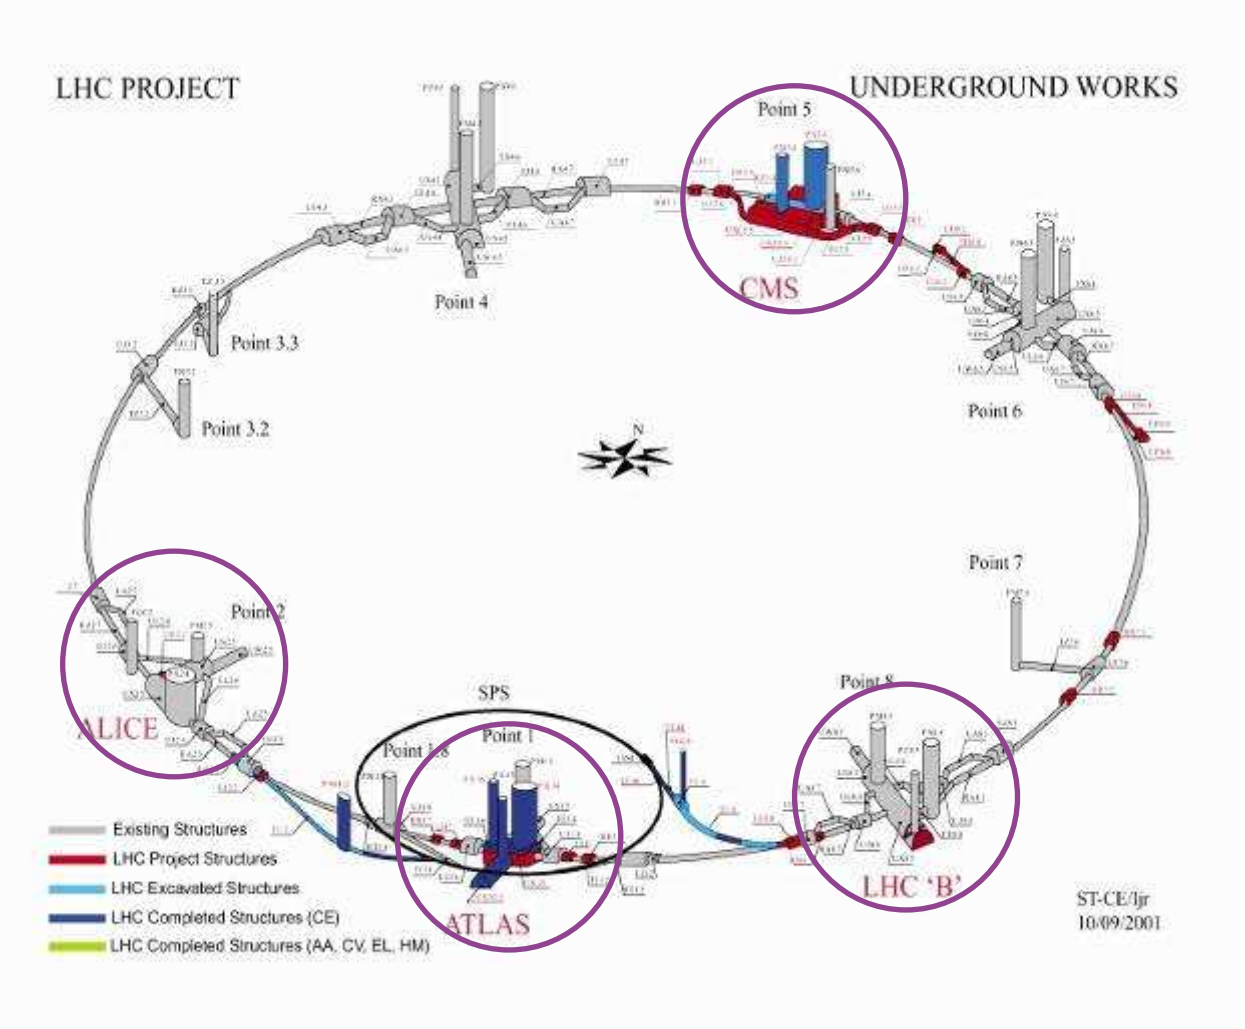

The LHC Data Challenge

Starting from this event

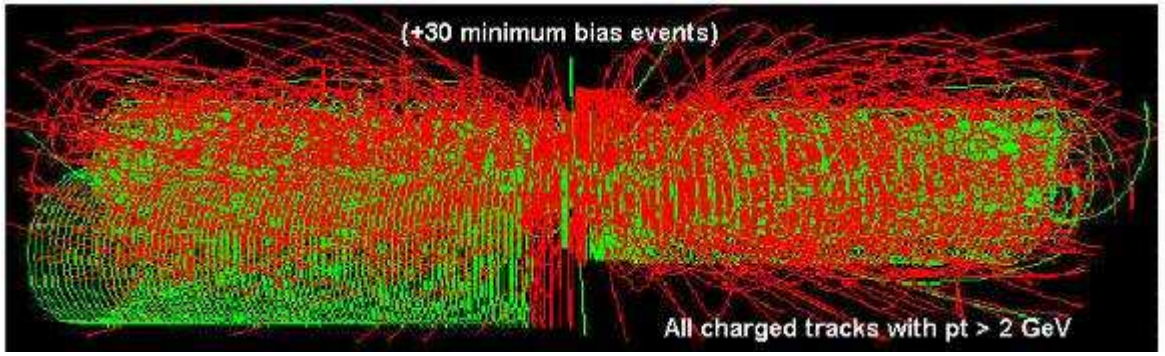

Looking for this "signature"

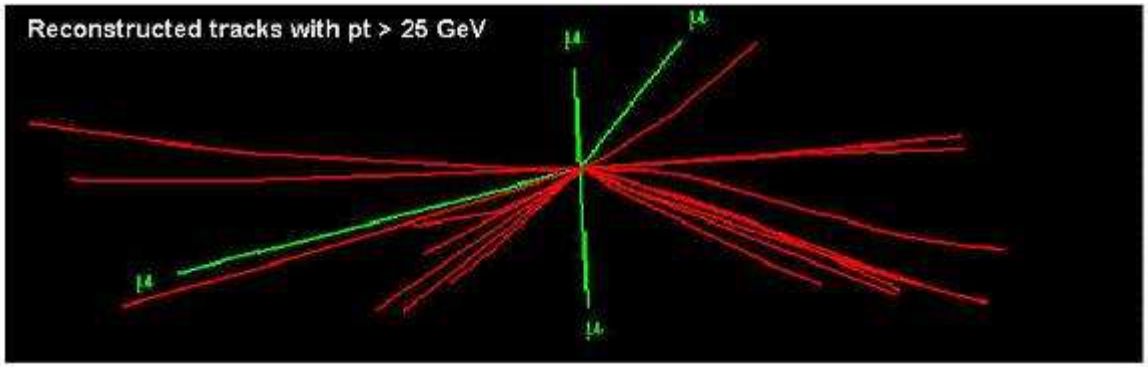

→ **Selectivity: 1 in 10^{13}**

(Like looking for a needle in 20 million haystacks)

LHC Data

- 40 million collisions per second
- After filtering, 100 collisions of interest per second
- A Megabyte of data for each collision = recording rate of 0.1 Gigabytes/sec
- 10^{10} collisions recorded each year
- ~ 10 Petabytes/year of data
- LHC data correspond to about 20 million CDs each year!
- ~ 100,000 of today's fastest PC processors

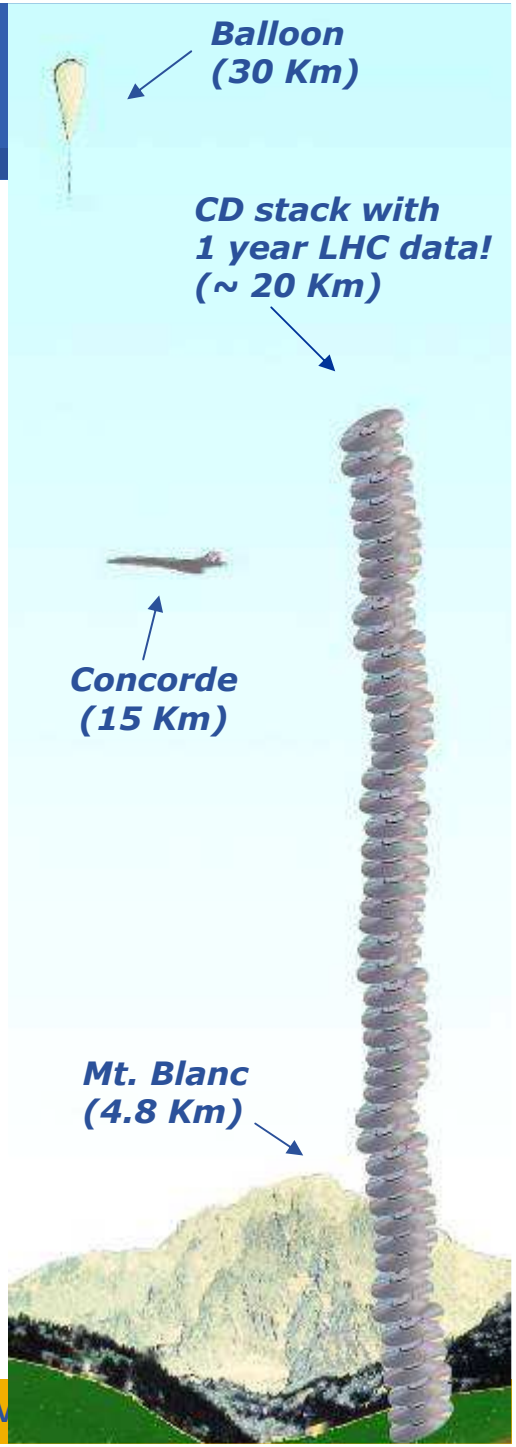

The solution: the Grid

- Integrating computing and storage capacities at major computer centres
- 24/7 access, independent of geographic location

→ Effective and seamless collaboration of dispersed communities, both scientific and commercial

→ Ability to use thousands of computers for a wide range of applications

- Best cost effective solution for HEP LHC Computing Grid project (LCG) and from this the close integration of LCG and EGEE projects

- **Objectives**
  - consistent, robust and secure service grid infrastructure
  - improving and maintaining the middleware
  - attracting new resources and users from industry as well as science

- **Structure**
  - 71 leading institutions in 27 countries, federated in regional Grids
  - leveraging national and regional grid activities worldwide
  - funded by the EU with ~32 M Euros for first 2 years starting 1st April 2004

EGEE Activities

- 48 % service activities (Grid Operations, Support and Management, Network Resource Provision)
- 24 % middleware re-engineering (Quality Assurance, Security, Network Services Development)
- 28 % networking (Management, Dissemination and Outreach, User Training and Education, Application Identification and Support, Policy and International Cooperation)

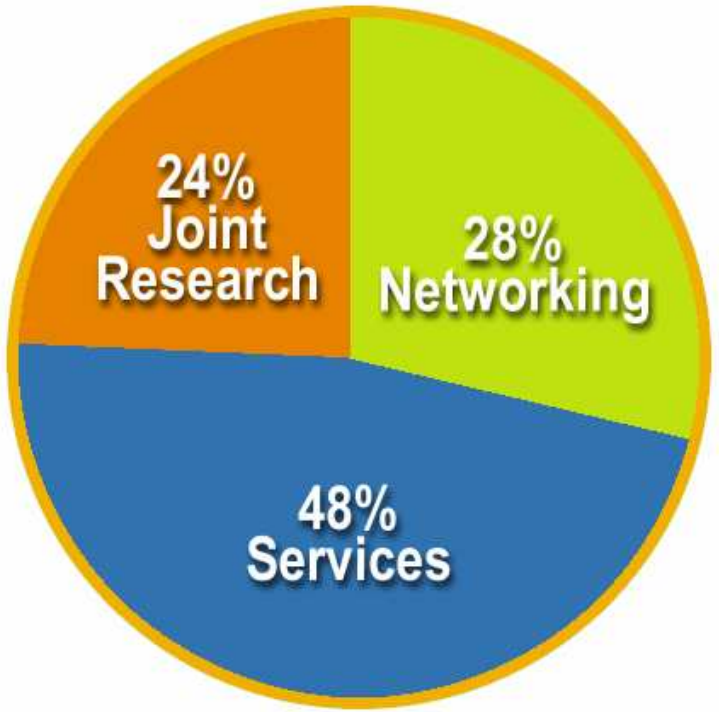

Emphasis in EGEE is on operating a production grid and supporting the end-users

EGEE Infrastructure

In collaboration with LCG

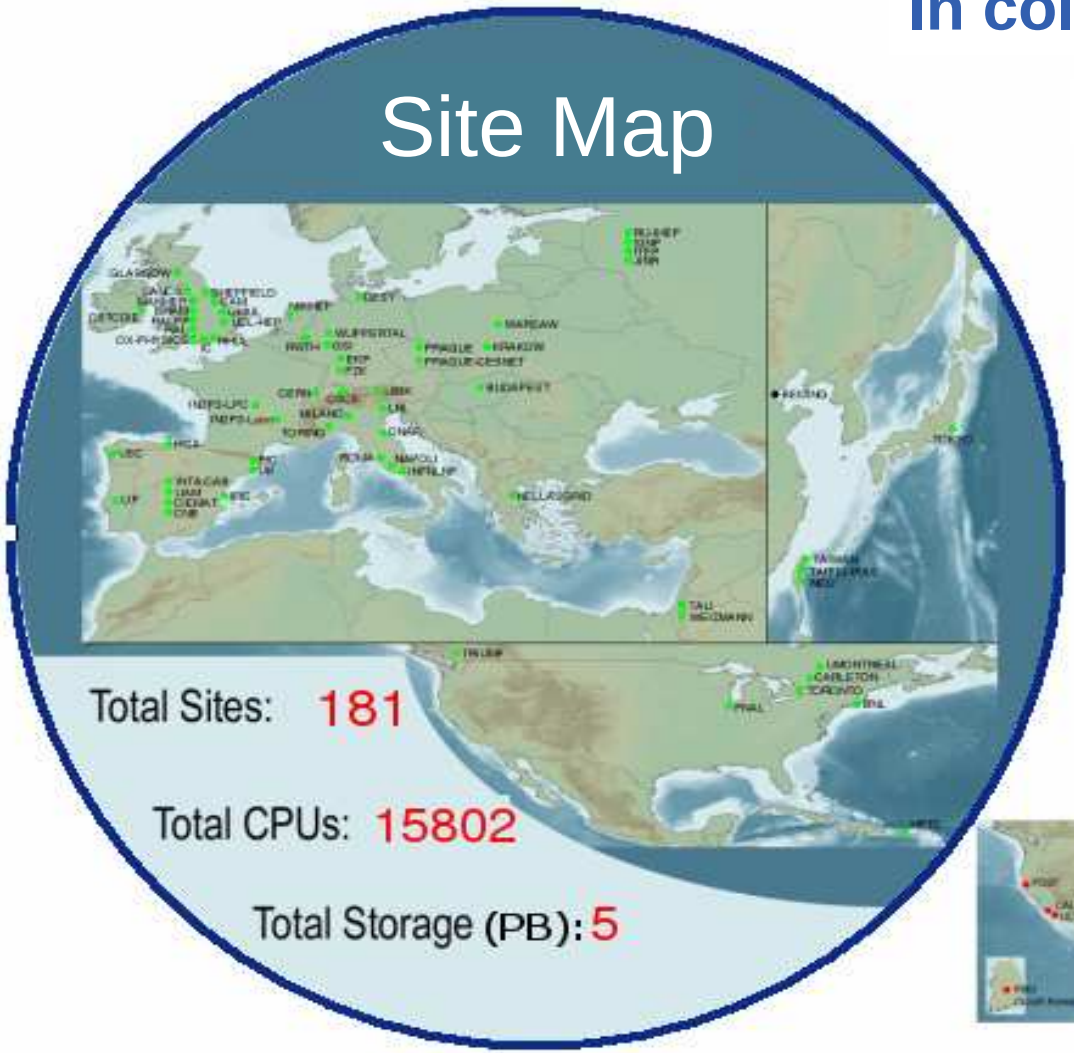

NorduGrid

Grid3/OSG

Status 25 July 2005

EGEE Production Service

- >180 sites
- >15 000 CPUs (with peaks >18 000 CPUs)
- ~14 000 jobs successfully completed per day
- 20 VOs
- >800 registered users, representing thousands of scientists

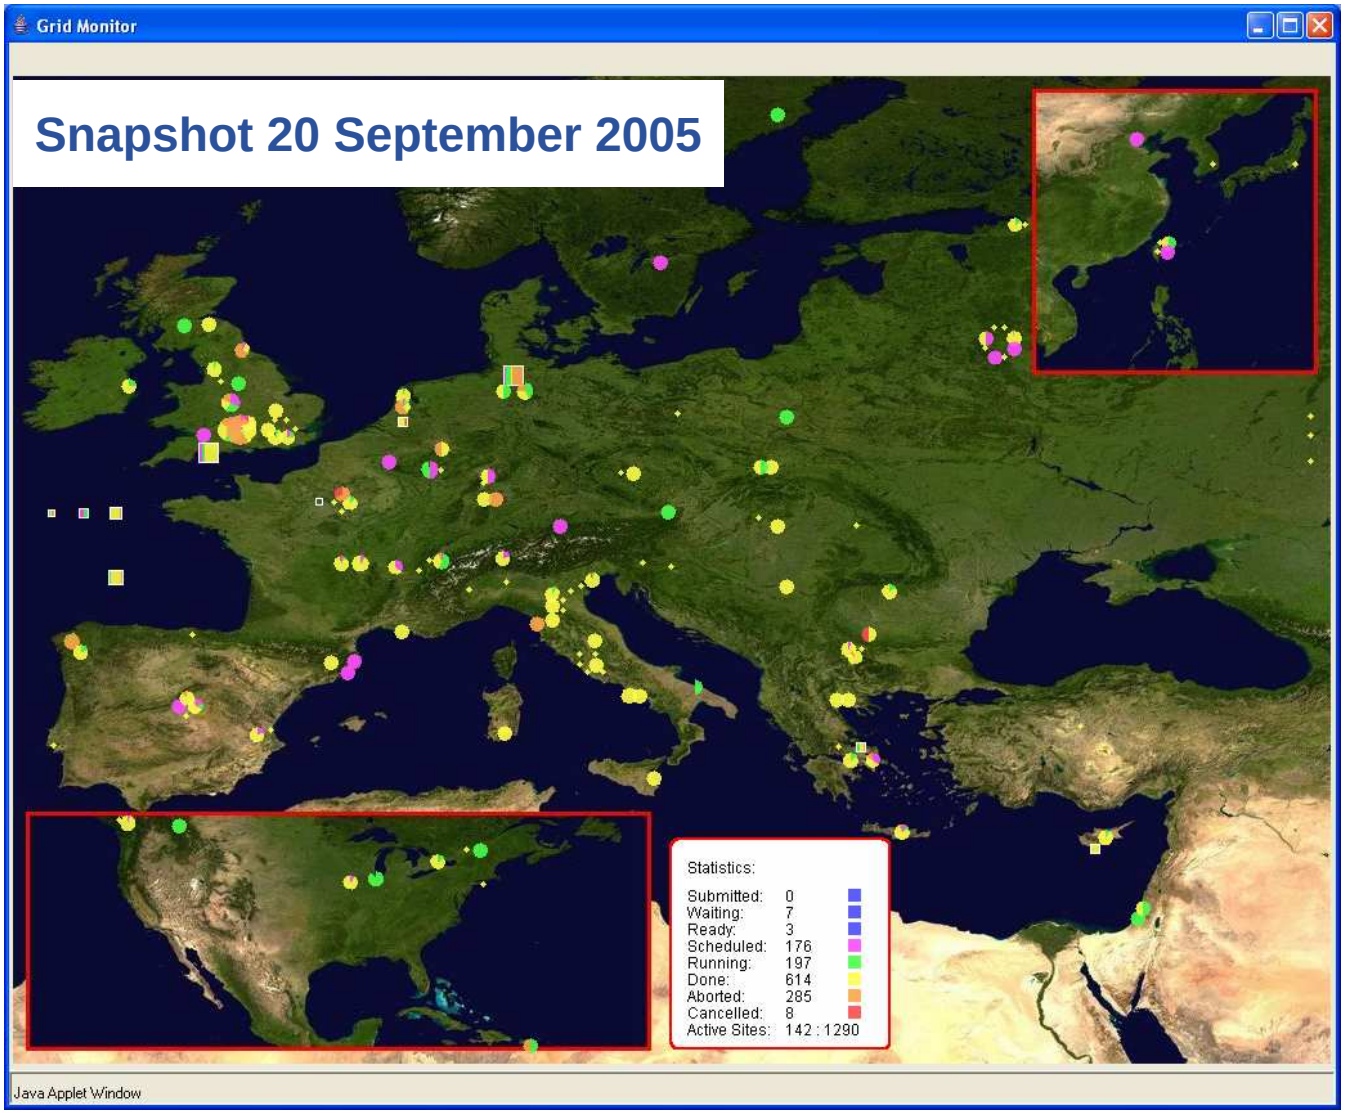

http://gridportal.hep.ph.ic.ac.uk/rtm/

EGEE services

- **Production service**
  - Based on the LCG-2 service
  - With new resource centres and new applications encouraged to participate
  - Stable, well-supported infrastructure, running only well-tested and reliable middleware

- **Pre-production service**
  - Run in parallel with the production service
  - First deployment of new versions of the middleware
  - Applications test-bed

- **GILDA testbed**
  - https://gilda.ct.infn.it/testbed.html
  - Complete suite of Grid elements and applications
    - Testbed, CA, VO, monitoring
  - Everyone can register and use GILDA for training and testing

Grid monitoring

- **Operation of Production Service: real-time display of grid operations**
- **Accounting Information**
- **Selection of Monitoring tools:**
  - GIIS Monitor + Monitor Graphs
  - Sites Functional Tests
  - GOC Data Base
  - Scheduled Downtimes
  - Live Job Monitor
  - GridIce – VO + Fabric View
  - Certificate Lifetime Monitor

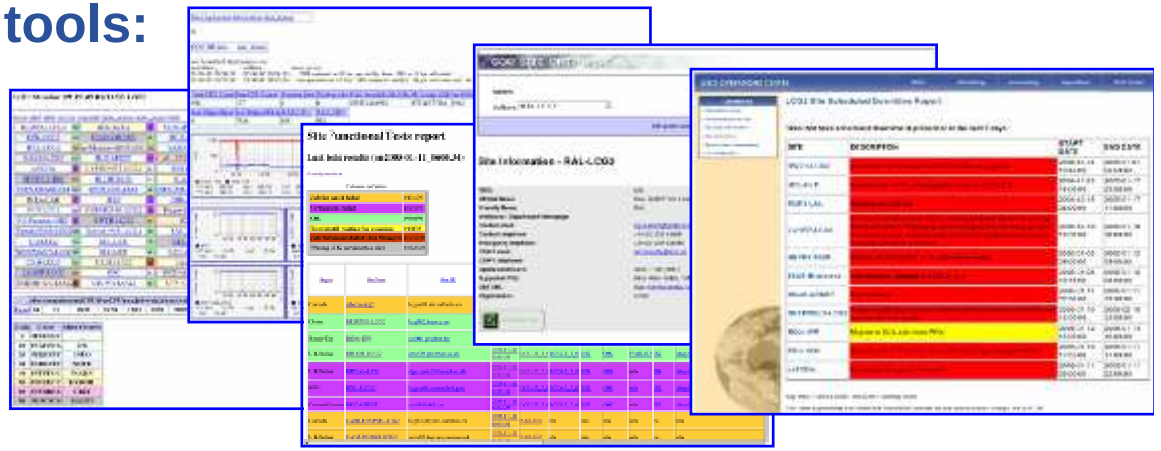

Grid middleware

- **The Grid relies on advanced software, called middleware, which interfaces between resources and the applications**

- **The GRID middleware:**
  - Finds convenient places for the application to be run
  - Optimises use of resources
  - Organises efficient access to data
  - Deals with authentication to the different sites that are used
  - Runs the job & monitors progress
  - Recovers from problems
  - Transfers the result back to the scientist

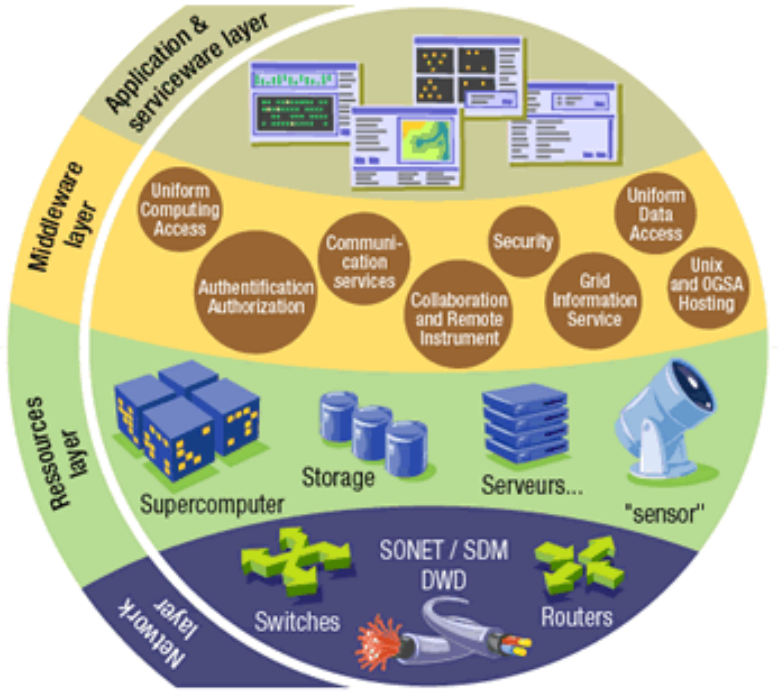

EGEE Middleware

- **First release of gLite end of March 2005** www.gLite.org
  - Now testing gLite v1.3 in pre-production

- **Guiding principles**
  - Lightweight services
  - Interoperability & Co-existence with deployed infrastructures
  - Robust: Performance & Fault Tolerance
  - Portable
  - Service oriented architecture
  - Site autonomy
  - Open source license

- **Development / Integration / Testing**
  - Workload Management
  - Information Systems
  - Security
  - Data Management

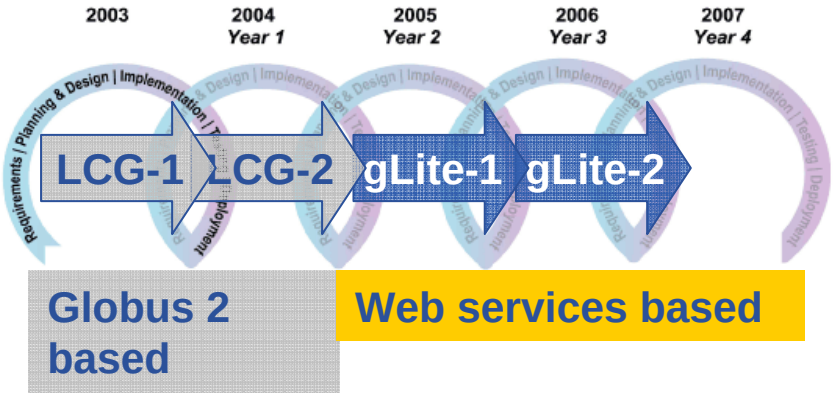

→ **gLite offers a complete data management solution in a distributed environment building on existing technology**

Architecture & Design

- **Design team includes**
  - Representatives from middleware providers (AliEn, Condor, EDG, Globus,…)
  - Colleagues from the Operations activity
  - Partners from related projects (e.g. OSG)

- **gLite development takes into account input and experiences from applications, operations, related projects**
  - Effective exchange of ideas, requirements, solutions and technologies
  - Coordinated development of new capabilities
  - Open communication channels
  - Joint deployment and testing of middleware
  - Early detection of differences and disagreements

gLite is not "just" a software stack, it is a "new" framework for international collaborative middleware development

Software stack and origin of services in release 1 (simplified)

- **Computing Element**
  - Gatekeeper *(Globus)*
  - Condor-C *(Condor)*
  - CE Monitor *(EGEE)*
  - Local batch system *(PBS, LSF, Condor)*
- **Workload Management**
  - WMS *(EDG)*
  - Logging and bookkeeping *(EDG)*
  - Condor-C *(Condor)*
- **Information and Monitoring**
  - R-GMA *(EDG)*
- **Storage Element**
  - gLite-I/O *(AliEn)*
  - Reliable File Transfer *(EGEE)*
  - GridFTP *(Globus)*
  - SRM: Castor *(CERN)*, dCache *(FNAL, DESY)*, other SRMs
- **Catalog**
  - File/Replica & Metadata Catalogs *(EGEE)*
- **Security**
  - GSI *(Globus)*
  - VOMS *(DataTAG/EDG)*
  - Authentication for C and Java based (web) services *(EDG)*

Now doing rigorous scalability and performance tests on pre-production service

Statistics

- **The core of gLite consists of** www.gLite.org
  - ~200 packages in 23 logical deployment modules
    → easy mechanism to deploy the gLite services in a flexible way
  - > 700.000 LOCs, mostly in Java, C, and C++
    (mostly re-engineered from predecessor projects, e.g EDG, DataTag)

- **… and depends on 3rd party software**
  - by partner projects
    - Condor, Globus, and LCG
  - from the open source community
    - Apache, Java, mySQL, etc.
  - distributed with the gLite package

Applications Running

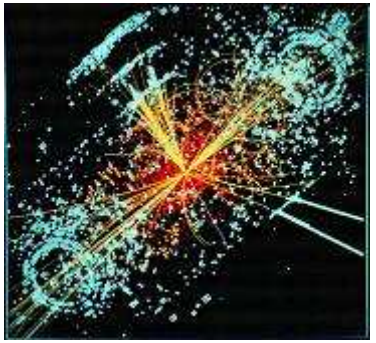

- **More than 20 applications from 6 domains**
  - High Energy Physics
    - 4 LHC experiments (Alice, ATLAS, CMS, LHCb)
    - BaBar, CDF, DØ, ZEUS
  - Biomedicine
    - Bioinformatics (Drug Discovery, GPS@, Xmipp_MLrefine, etc.)
    - Medical imaging (GATE, CDSS, gPTM3D, SiMRI 3D, etc.)
  - Earth Sciences
    - Earth Observation, Solid Earth Physics, Hydrology, Climate
  - Computational Chemistry
  - Astronomy
    - MAGIC
    - Planck
  - Geo-Physics
    - EGEODE

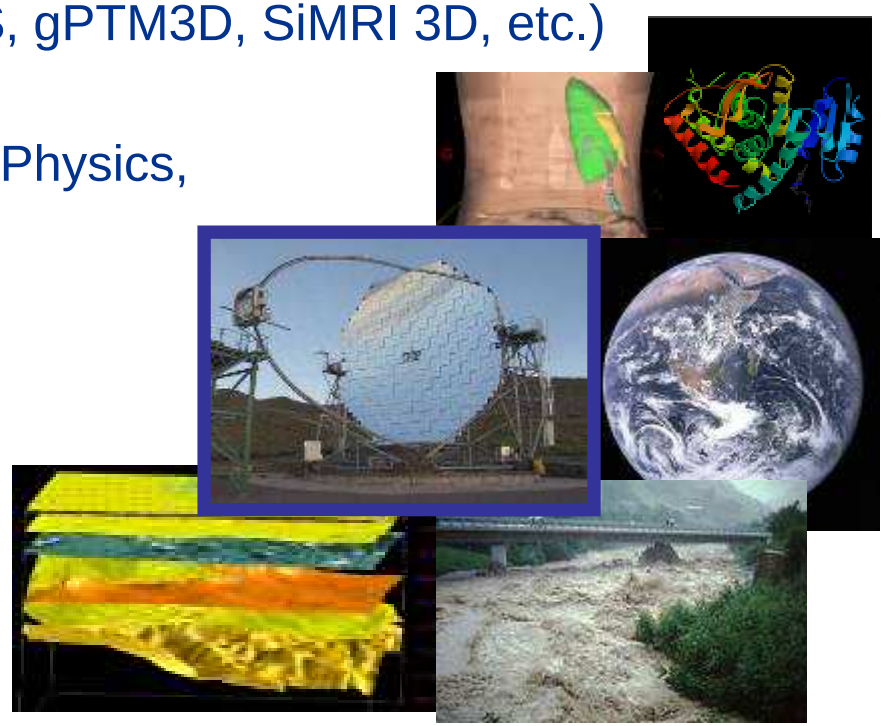

Service Usage

- **More than 800 registered users in the VOs on the production service**
  - + Many local VOs, supported by their ROCs
- **Scale of work performed:**
  - – LHC Data challenges 2004:
    - ▪ >1 M SI2K years of CPU time (~1000 CPU years)
    - ▪ 400 TB of data generated, moved and stored
    - ▪ 1 VO achieved ~4000 simultaneous job (~4 times CERN grid capacity)

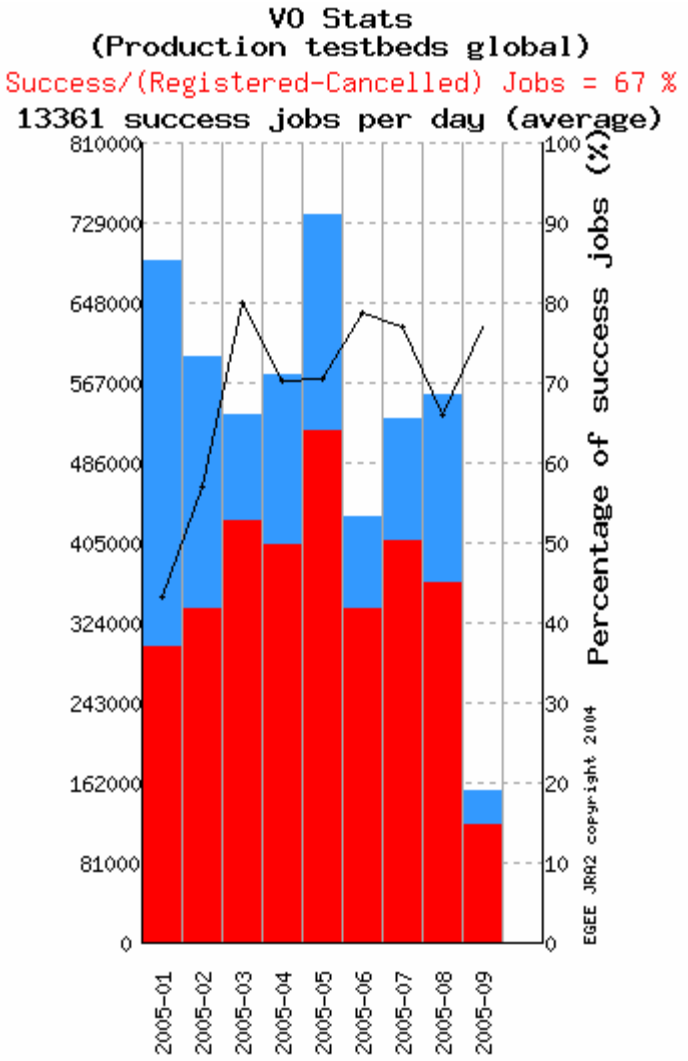

Number of jobs processed/month

LCG2 testbed usage

- **Average job duration January 2005 – June 2005 for the main VOs**

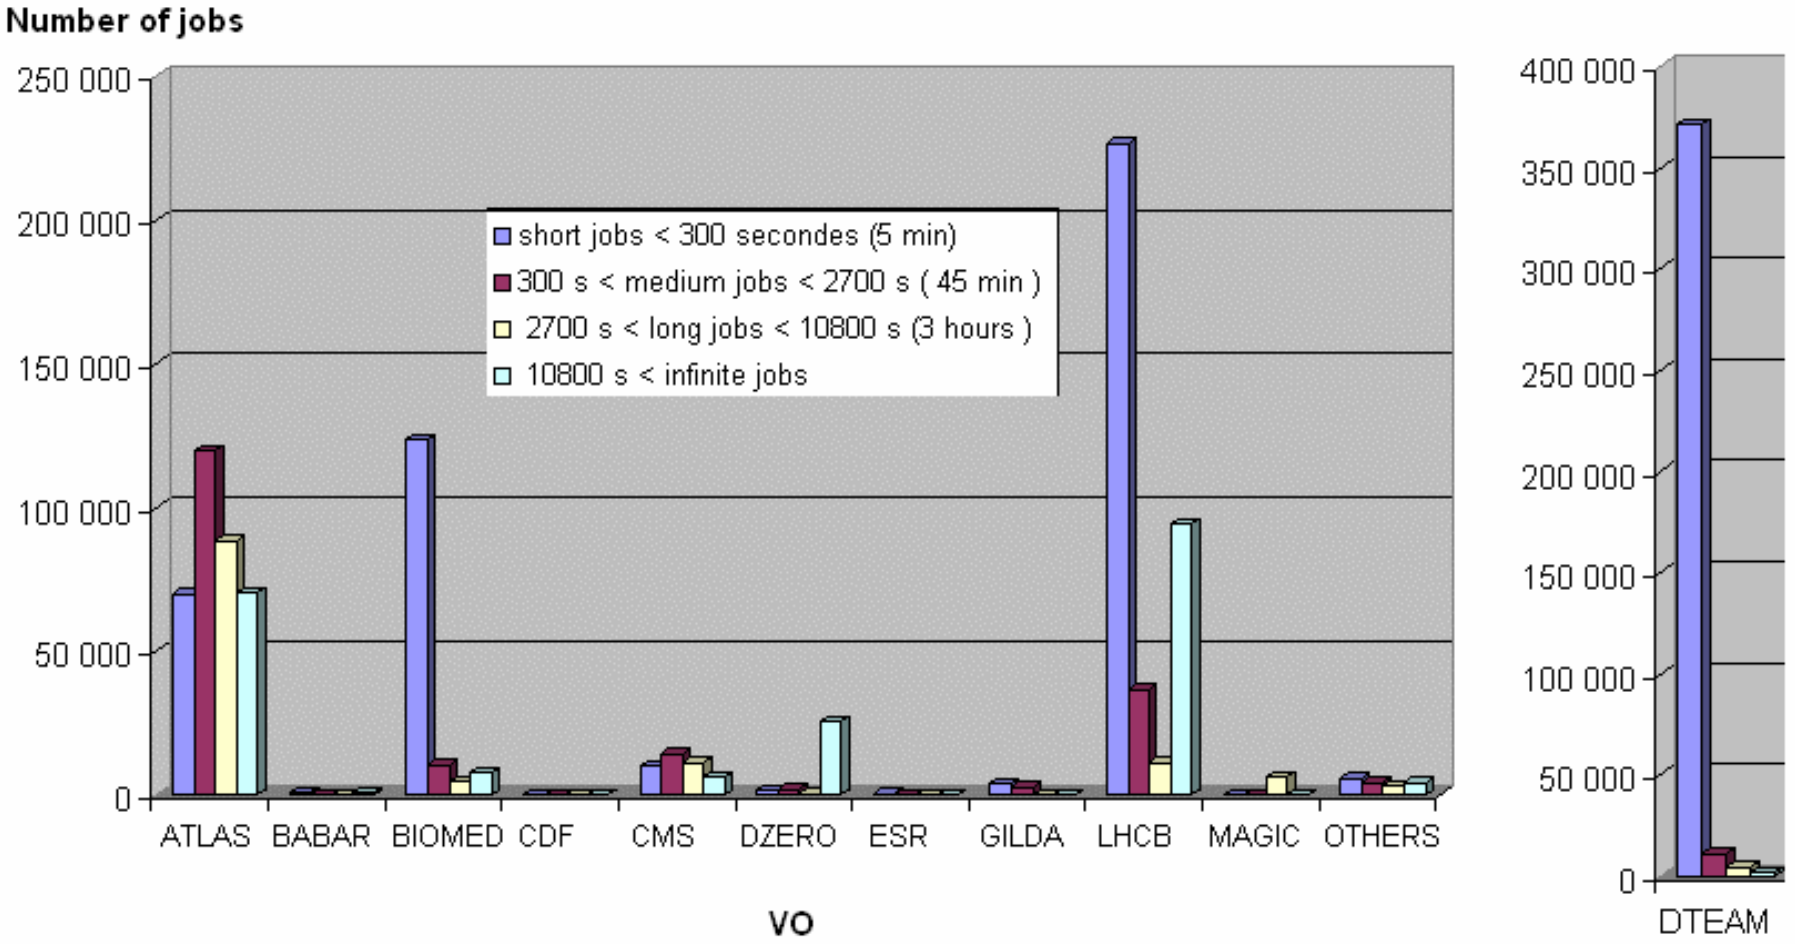

EGEE pilot applications

- **High-Energy Physics (HEP)**
  - Provides computing infrastructure (LCG)
  - Challenging:
    - thousands of processors world-wide
    - generating petabytes of data
    - ‘chaotic’ use of grid with individual user analysis (thousands of users interactively operating within experiment VOs)

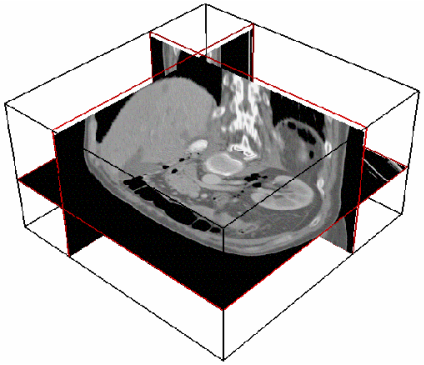

- **Biomedical Applications**
  - Similar computing and data storage requirements
  - Major challenge: security

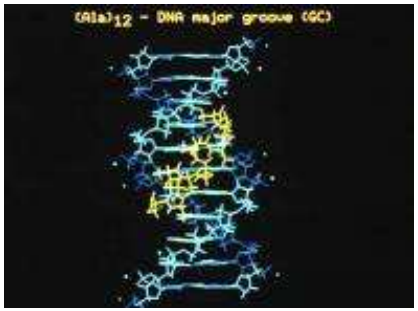

BioMed Overview

- **Infrastructure**
  - ~2.000 CPUs
  - ~21 TB of disk
  - in 12 countries

- **>50 users in 7 countries working with 12 applications**
- **18 research labs**

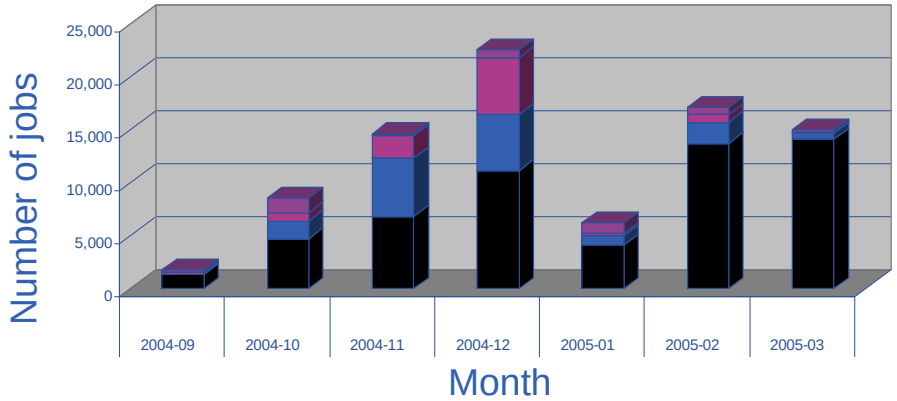

Bioinformatics

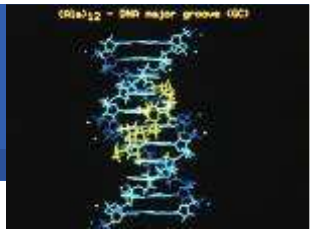

- **GPS@: Grid Protein Sequence Analysis**
  - **Gridified version of NPSA web portal**
    - Offering proteins databases and sequence analysis algorithms to the bioinformaticians (3000 hits per day)
    - Need for large databases and big number of short jobs
  - **Objective**: increased computing power
  - **Status**: 9 bioinformatic softwares gridified
  - **Grid added value**: open to a wider community with larger bioinformatic computations

- **xmipp_MLrefine**
  - **3D structure analysis of macromolecules**
    - From (very noisy) electron microscopy images
    - Maximum likelihood approach to find the optimal model
  - **Objective**: study molecule interaction and chem. properties
  - **Status**: algorithm being optimised and ported to 3D
  - **Grid added value**: parallel computation on different resources of independent jobs

Drug Discovery

- **Grid-enabled drug discovery process for neglected diseases**
  - *In silico* docking: compute probability that potential drugs will dock with a target protein
  - To speed up and reduce cost required to develop new drugs

- **WISDOM (Wide *In Silico* Docking On Malaria)**
  - Drug Discovery Data Challenge
  - 11 July – 19 August
  - 46 million docked ligands produced (typical for computer clusters: 100 000 ligands)
  - Equivalent to 80 CPU years
  - 1000 computers in 15 countries used simultaneously
  - Millions of files (adding up to a few TB of data)

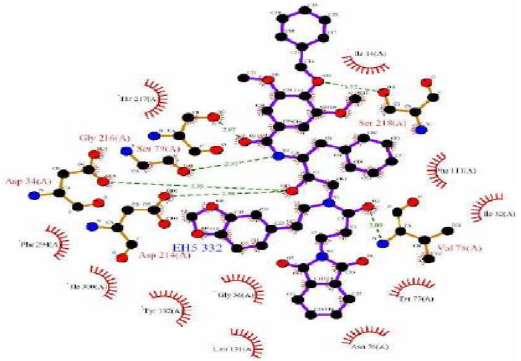

→ Never done on a large scale production infrastructure
→ Never done for a neglected disease

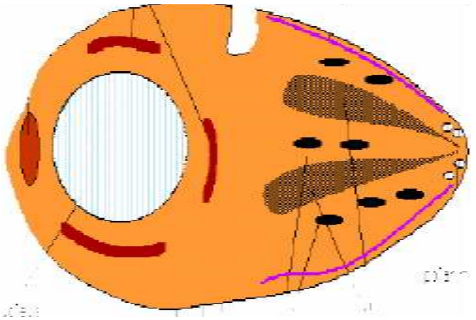

- **Next steps**
  - Sort through data to identify potential drugs
  - Develop the next steps of the process (molecular dynamics)

Medical imaging

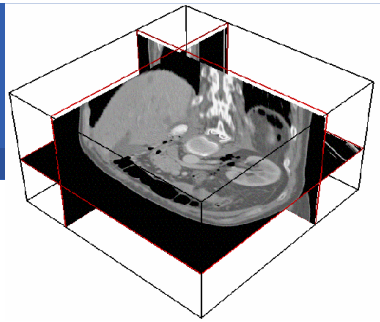

- **GATE**
  - **Radiotherapy planning**
    - Improvement of precision by Monte Carlo simulation
    - Processing of DICOM medical images
  - **Objective**: very short computation time compatible with clinical practice
  - **Status**: development and performance testing
  - **Grid Added Value**: parallelisation reduces computing time

- **CDSS**
  - **Clinical Decision Support System**
    - Assembling knowledge databases
    - Using image classification engines
  - **Objective:** access to knowledge databases from hospitals
  - **Status**: from development to deployment, some medical end users
  - **Grid Added Value**: ubiquitous, managed access to distributed databases and engines

Medical imaging

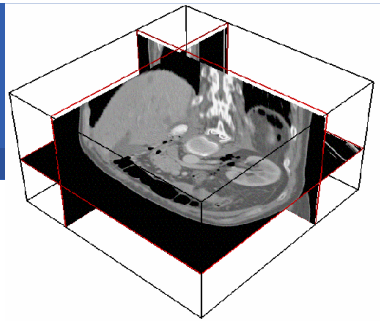

- **SiMRI3D**
  - **3D Magnetic Resonance Image Simulator**
    - MRI physics simulation, parallel implementation
    - Very compute intensive

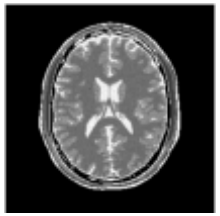

  - **Objective**: offering an image simulator service to the research community
  - **Status**: parallelised and now running on EGEE resources
  - **Grid Added Value**: enables simulation of high-res images

- **gPTM3D**
  - **Interactive tool to segment and analyse medical images**
    - A non gridified version is distributed in several hospitals
    - Need for very fast scheduling of interactive tasks

  - **Objectives**: shorten computation time using the grid
    - Interactive reconstruction time: < 2min and scalable
  - **Status**: development of the gridified version being finalized
  - **Grid Added Value**: permanent availability of resources

Generic Applications

- **EGEE Generic Applications Advisory Panel (EGAAP)**
  - UNIQUE entry point for “external” applications
  - Reviews proposals and make recommendations to EGEE management
    - Deals with “scientific” aspects, not with technical details
    - Generic Applications group in charge of introducing selected applications to the EGEE infrastructure
  - 6 applications selected so far:
    - Earth sciences (I and II)
    - MAGIC
    - Computational Chemistry
    - PLANCK
    - Drug Discovery
    - GRACE (end Feb 2005)

Earth sciences applications

- **Earth Observations by Satellite**
  - Ozone profiles
- **Solid Earth Physics**
  - Fast Determination of mechanisms of important earthquakes
- **Hydrology**
  - Management of water resources in Mediterranean area (SWIMED)
- **Geology**
  - Geocluster: R&D initiative of the Compagnie Générale de Géophysique

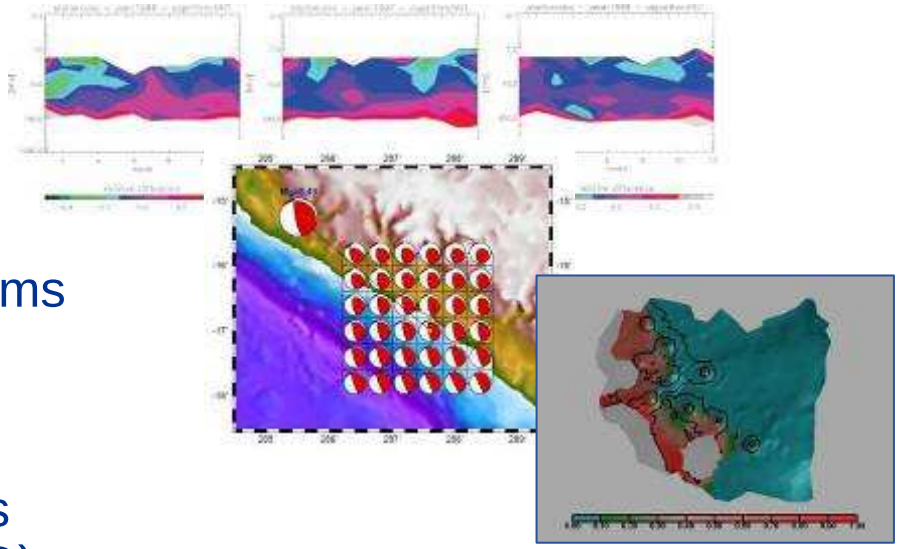

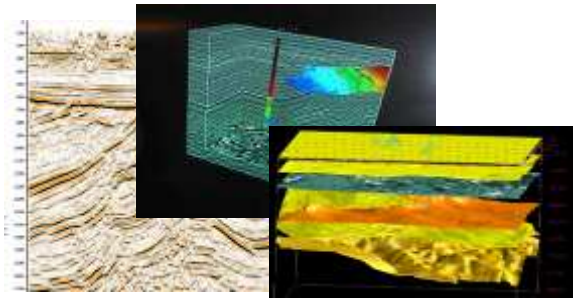

- **A large variety of applications ported on EGEE which incites new users**
- **Interactive Collaboration of the teams around a project**

MAGIC

- **Ground based Air Cerenkov Telescope 17 m diameter**
- **Physics Goals:**
  - Origin of VHE Gamma rays
  - Active Galactic Nuclei
  - Supernova Remnants
  - Unidentified EGRET sources
  - Gamma Ray Burst
- **MAGIC II will come 2007**
- **Grid added value**
  - Enable “(e-)scientific“ collaboration between partners
  - Enable the cooperation between different experiments
  - Enable the participation on Virtual Observatories

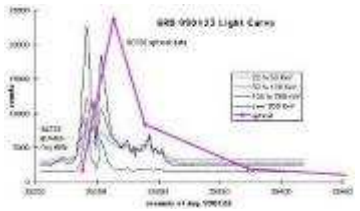

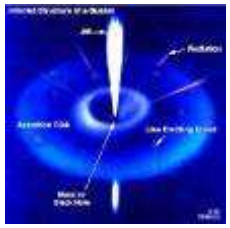

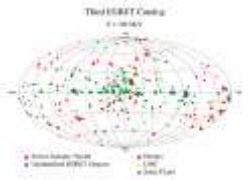

Computational Chemistry

- **The Grid Enabled Molecular Simulator (GEMS)**
  - Motivation:
    - Modern computer simulations of biomolecular systems produce an abundance of data, which could be reused several times by different researchers.
      → data must be catalogued and searchable
  - GEMS database and toolkit:
    - autonomous storage resources
    - metadata specification
    - automatic storage allocation and replication policies
    - interface for distributed computation

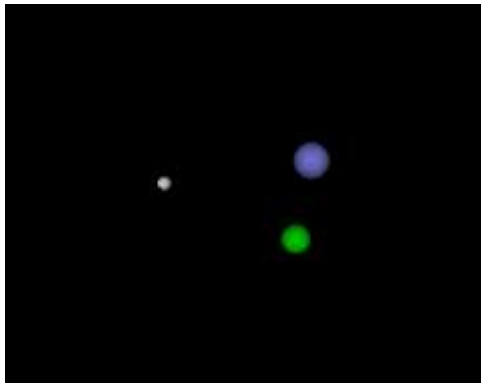

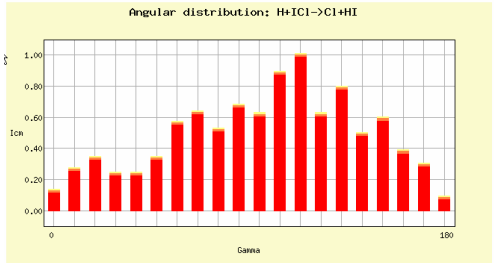

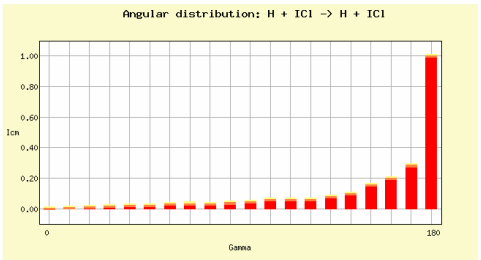

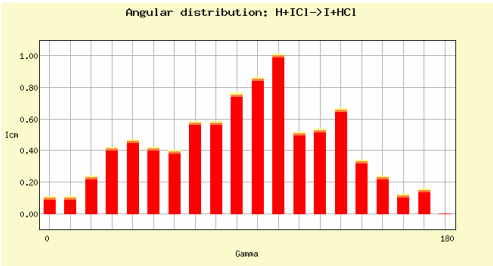

Planck

- **ESA satellite mission to map the cosmic microwave background**
  - Three steps:
    - Imaging sky emission at many frequencies
    - Peeling back the layers of different sources (detector noise, galactic, etc.)
    - Recovering cosmological information
  - Complex data structure
    → data handling important
- **On the Grid:**
  > 12 times faster (but ~5% failures)
- **The Grid as**
  - collaboration tool
  - common user-interface
  - flexible environment
  - new approach to data and S/W sharing

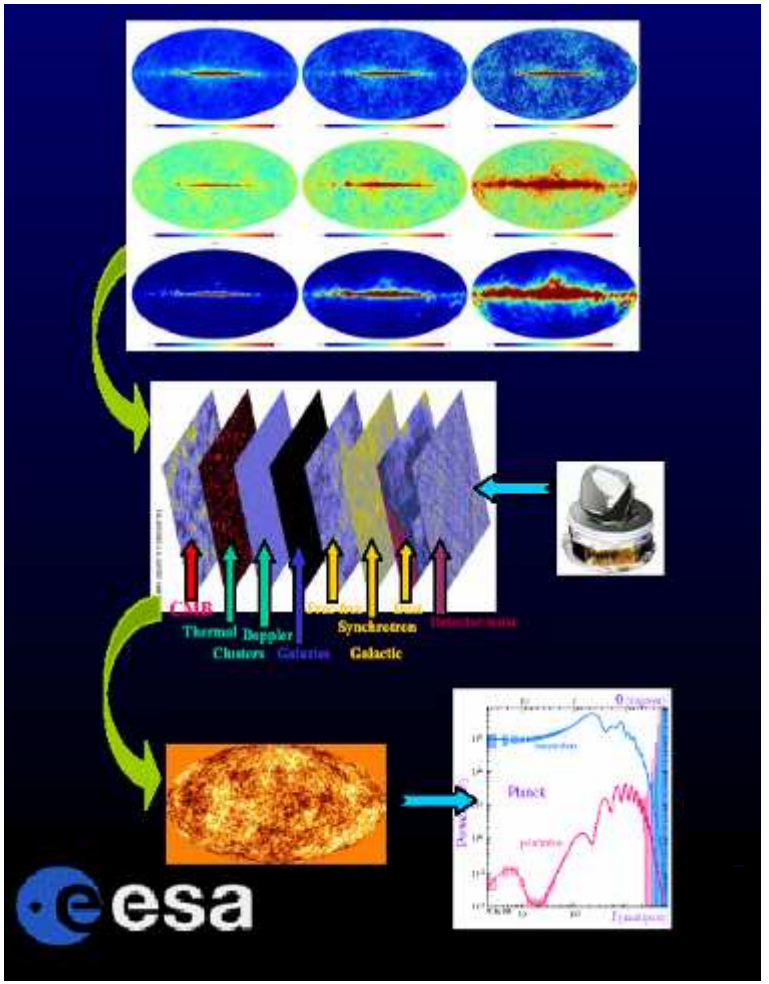

User information & support

- **More than 140 training events (including the ISSGC school) across many countries**
  - >1200 people trained
    induction; application developer; advanced; retreats
  - Material archive coming online with ~200 presentations

- **Public and technical websites constantly evolving to expand information available and keep it up to date**

- **3 conferences organized**
  - ~ 300 @ Cork
  - ~ 400 @ Den Haag
  - ~ 450 @ Athens

- **Pisa: 4th project conference 24-28 October '05**

Collaborations

- **EGEE closely collaborates with other projects, e.g.**

- **Flooding Crisis (CrossGrid) demonstrated at 3rd EGEE conference in Athens**
  - Simulation of flooding scenarios
  - Display in Virtual Reality
  - Optimize data transport

  → **won prize for "best demo"**

CrossGrid

Collaboration with Slowak Academy of Sciences

How to collaborate with EGEE

- Contact us and consider to participate in or organize a EGEE tutorial
- Make first experience and tests with our t-Infrastructure Gilda and our Genius portal
- Consider installing the EGEE middleware gLite on your own infrastructure or on the general EGEE infrastructure
- Alternatively apply to the EGAAP to port and deploy your application on the present EGEE infrastructure
- Consider joining forces with EGEE to submit a EGEE related project

EGEE as partner

- Ongoing collaborations
  - with non EU partners in EGEE: US, Israel, Russia, Korea, Taiwan…
  - with other European infrastructure projects:
    - GÉANT
    - DEISA
    - SEE-GRID
  - with other European grid projects
    - DILIGENT
    - GRIDCC
  - with non-European projects:
    - OSG: OpenScienceGrid (USA)
    - NAREGI

- EGEE as incubator
  - >10 related projects have been created

Related projects

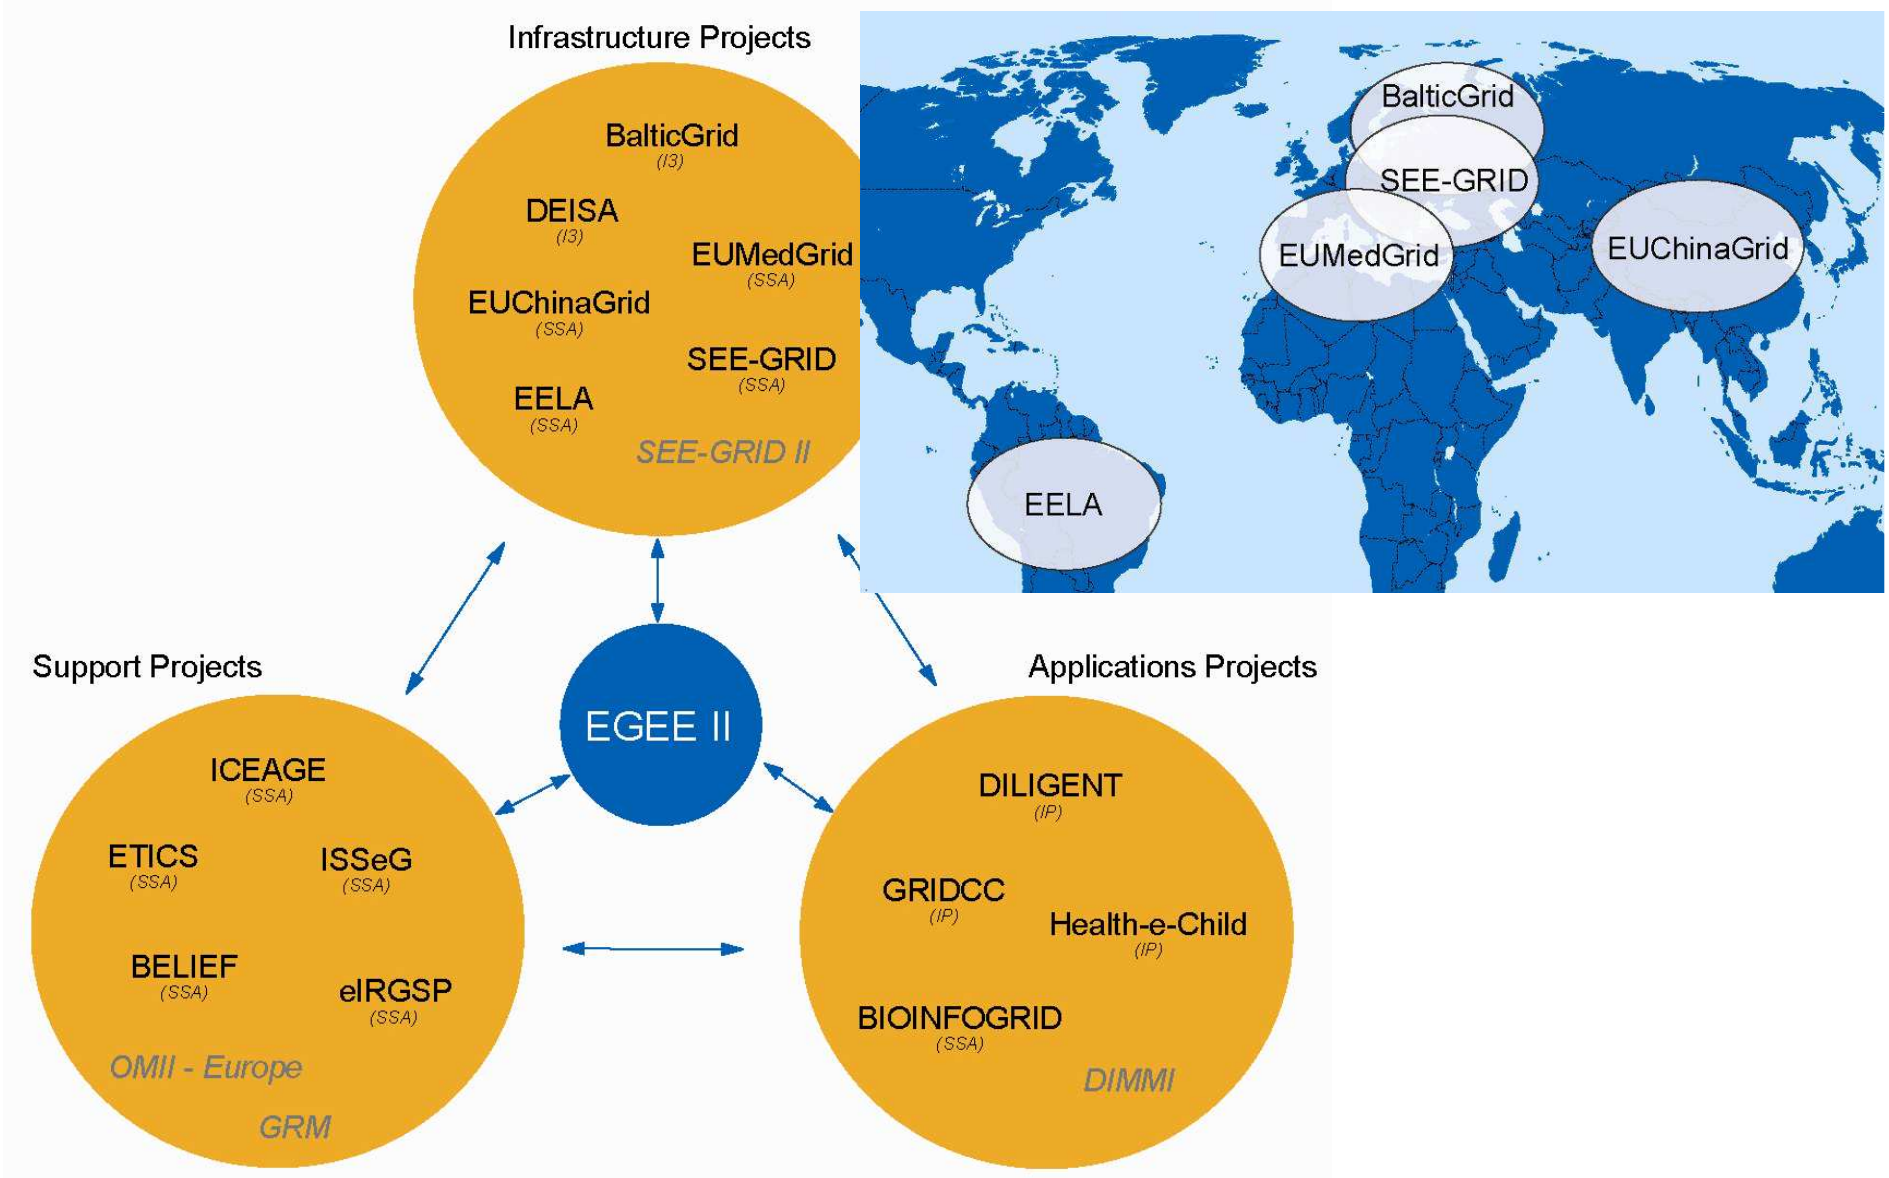

Toward EGEE-II

- **EGEE-II proposal submitted to the EU**
  - On 8 September 2005
  - Proposed start 1 April 2006

- **Natural continuation of EGEE**
  - Emphasis on providing an infrastructure for e-Science
    - → increased support for applications
    - → increased multidisciplinary Grid infrastructure
    - → more involvement from Industry
  - Expanded consortium
    - > 90 partners in 32 countries (Non-European partners in USA, Korea and Taiwan)
    - related projects

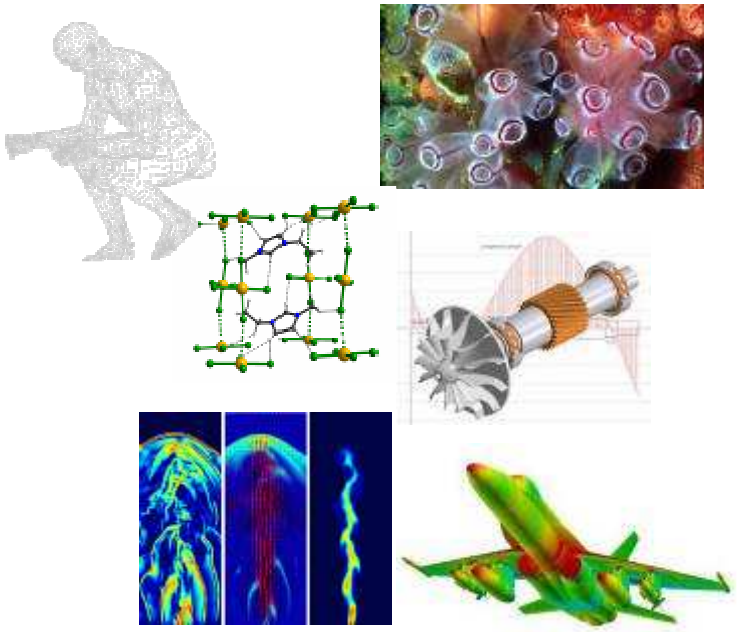

→ **world-wide Grid infrastructure**
→ **increased international collaboration**

Security & Intellectual Property

- **The existing EGEE grid middleware is distributed under an Open Source License developed by EU DataGrid**
  - No restriction on usage (scientific or commercial) beyond acknowledgement
  - Same approach for new middleware
- **Application software maintains its own licensing scheme**
  - Sites must obtain appropriate licenses before installation

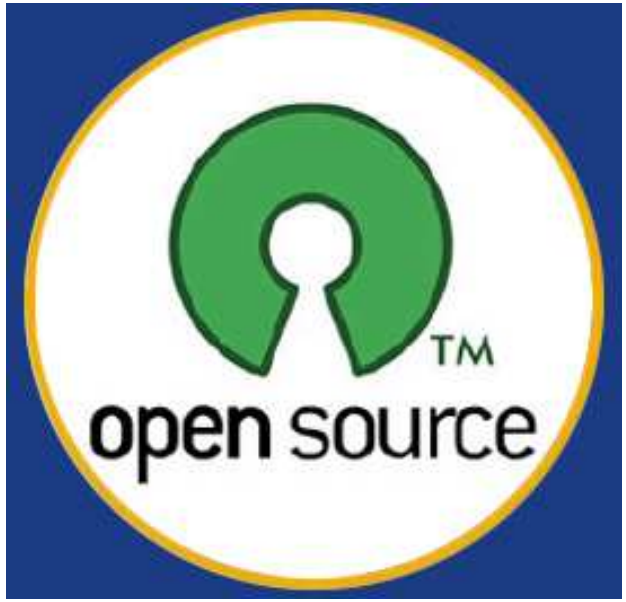

Conclusions I

- Grids are a powerful new tool for science – as well as other fields
- Grid computing has been chosen by CERN and HEP as the most cost effective computing model
- Several other applications are already benefiting from Grid technologies
- Investments in Grid projects are growing world-wide
- Europe is strong in the development of Grids also thanks to the success of EGEE and related projects

Conclusions II

- **Collaboration across national and international programmes is very important:**
  - Grids are above all about collaboration at a large scale
  - Science is international and therefore requires an international computing infrastructure

- **Thanks for the invitation to come and speak in Valencia and for this opportunity to explore any further collaboration in this area**

Contacts

- **EGEE Website**

  http://www.eu-egee.org

- **How to join**

  http://public.eu-egee.org/join/

- **EGEE Project Office**

  project-eu-egee-po@cern.ch