GRID: from HEP to Science

Federico Ruggieri – INFN

La Thuile 6-11 March 2006

Outline

- What is a Grid ?
- A bit of history (how we got involved)
- GRID services and Infrastructure (EGEE)
- Applications and Communities
- Social Impact and Digital Divide
- Conclusions

What Grid is not

- GRID is not an acronym (like Good Resources In Disbanding)
- GRID is not a larger bandwidth network (although having an High Bandwidth network is mandatory for most of the services).
- GRID is not a proprietary solution (even if there are many products with GRID word in their name).
- GRID is not a new High Performance Computer, but it's made up by very many (10^4 and more) computers (mostly off-the-shelf servers).
- GRID is not magic: will not work without a very efficient infrastructure organization.

What GRID is supposed to be

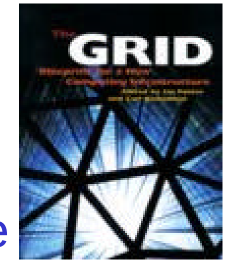

"A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities." I. Foster & K. Kesselman - The Grid: Blueprint for a New Computing Infrastructure – Morgan Kaufman 1998.

- A dependable infrastructure that can facilitate the usage of distributed resources by many groups of distributed persons or Virtual Organizations.
- The GRID paradigm is an extension of the WEB one, which was originally limited to distributed access to distributed information and documents.
- The classical example is the Power GRID: you plug in and receive power; you don't know (and you don't care) where it comes from.

A Grid Checklist

- Ian Foster more recently suggested that GRID is a system that:
  - 1) ***coordinates resources that are not subject to centralized control*** … (A Grid integrates and coordinates resources and users that live within different control domains—for example, the user's desktop vs. central computing; different administrative units of the same company; or different companies; and addresses the issues of security, policy, payment, membership, and so forth that arise in these settings. Otherwise, we are dealing with a local management system.)
  - 2) *…* ***using standard, open, general-purpose protocols and interfaces***… (A Grid is built from multi-purpose protocols and interfaces that address such fundamental issues as authentication, authorization, resource discovery, and resource access. As I discuss further below, it is important that these protocols and interfaces be *standard* and *open*. Otherwise, we are dealing with an applicationspecific system.)
  - 3) *…* ***to deliver nontrivial qualities of service***. (A Grid allows its constituent resources to be used in a coordinated fashion to deliver various qualities of service, relating for example to response time, throughput, availability, and security, and/or co-allocation of multiple resource types to meet complex user demands, so that the utility of the combined system is significantly greater than that of the sum of its parts.)

A bit of (My) short history in Grids

- In the 80’ and early 90’ the accent was on client-server and meta-computing.
- In 1998 I. Foster & K. Kesselman - The Grid: Blueprint for a New Computing Infrastructure & Globus project (www.globus.org).
- First GRID presentation in CHEP’98 – Chicago.
- 1999 – 2000 INFN-GRID Project started based on Globus, GridPP in UK.
- 2000 - 2003 First EU Project DataGRID and PPDG & GRIPHYN in US.
- 2003 – 2006 EGEE Project in EU and OSG in US
- Many other projects in many countries (Japan, China, etc.)

A GRID for LHC and HEP

- We got involved in Grids to solve the huge LHC computational problem which was, at that time, starting to be investigated (after an initial under-evaluation).
- In the late 90' client-server and meta-computing were the frontier and Computer Farms were just started (Beowulf).
- The largest problem anyway was the huge amount of data expected to be produced and analyzed (PB).
- The "social" challenge was to allow thousands of physicists to access those data easily from tens of countries in different continents.

LHC Computational Problem

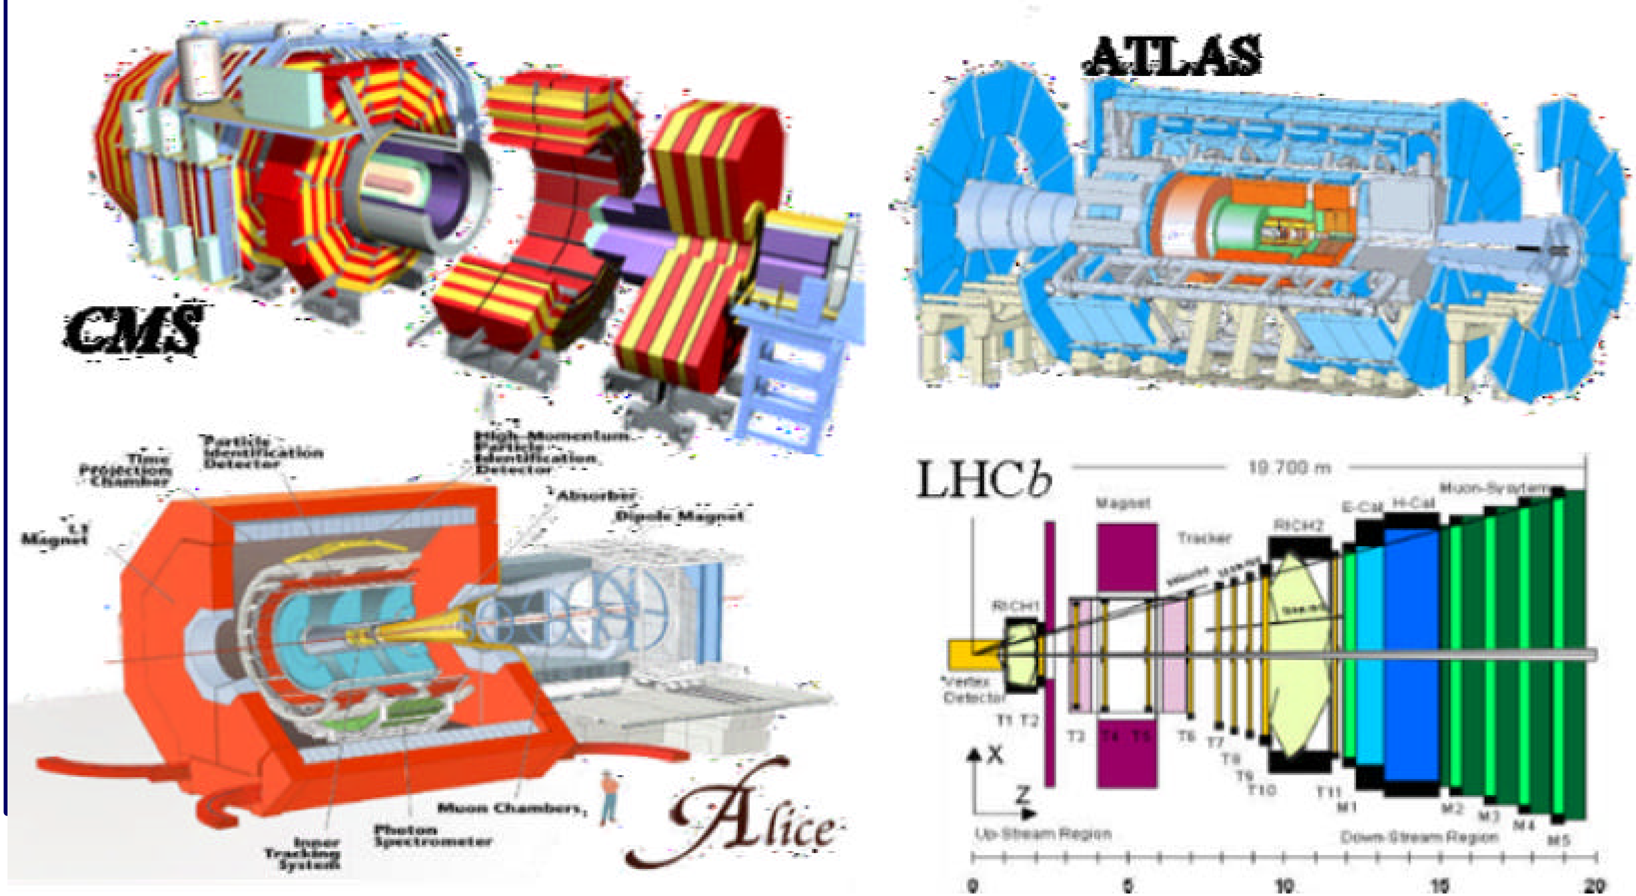

One Experiment

PetaBytes of data every year produced by Montecarlo simulation and real data.

> Thousand of Physicists and many hundreds thousands CPU.

CMS
Compact Muon Solenoid

Estimate of Computing needs at CERN for LHC (2000)

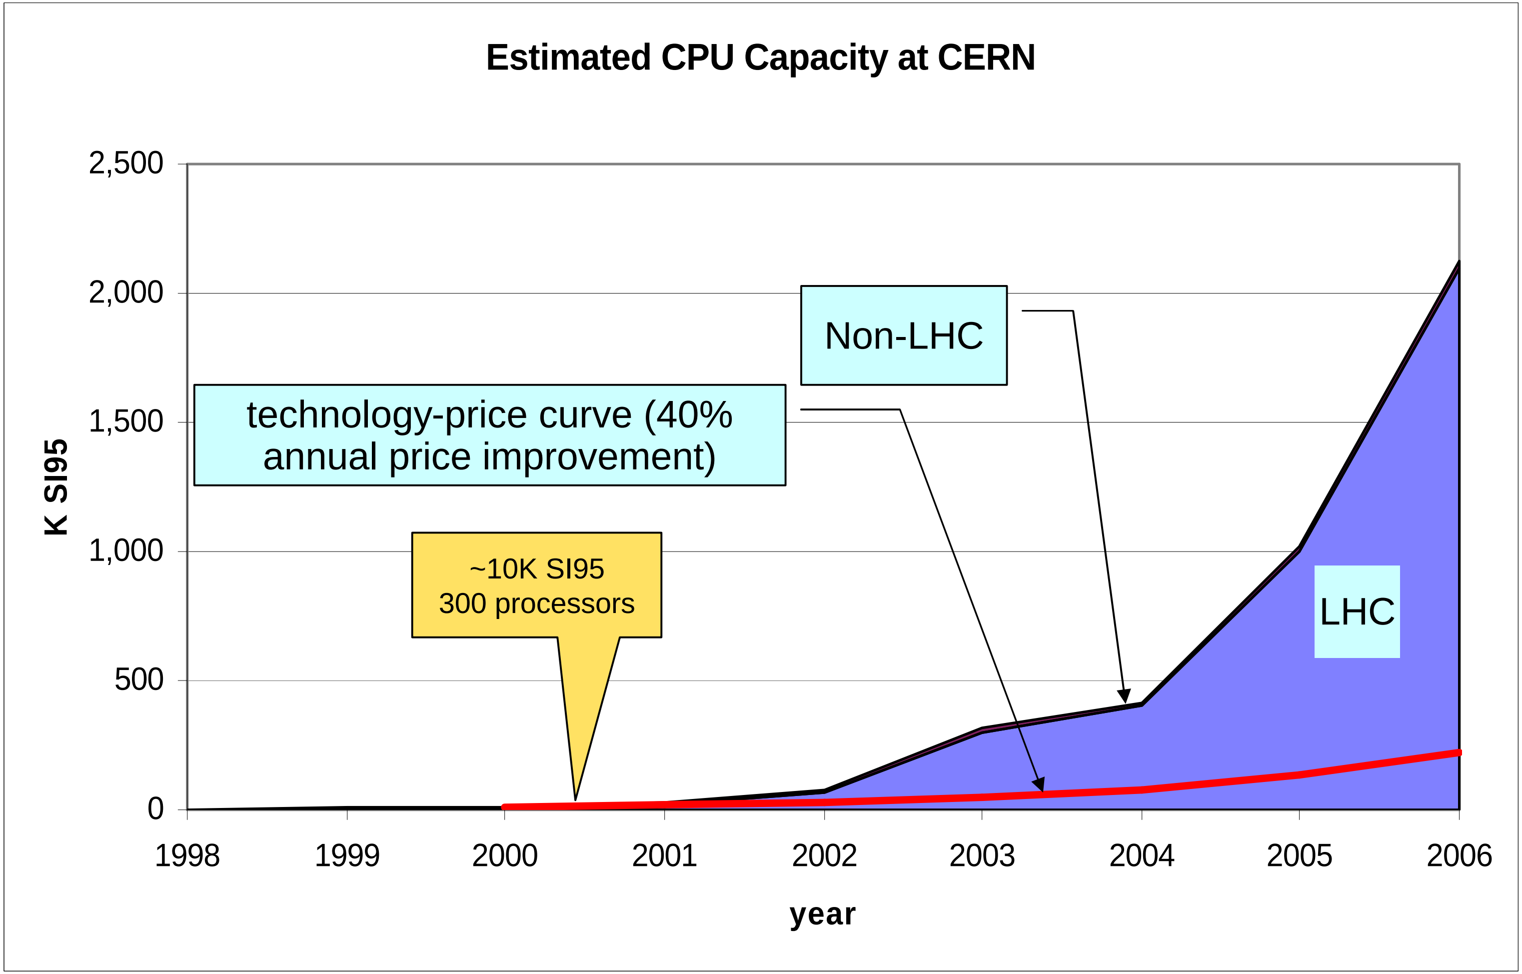

Extension of Web Paradigm

Web: Uniform Access to Information and Documents

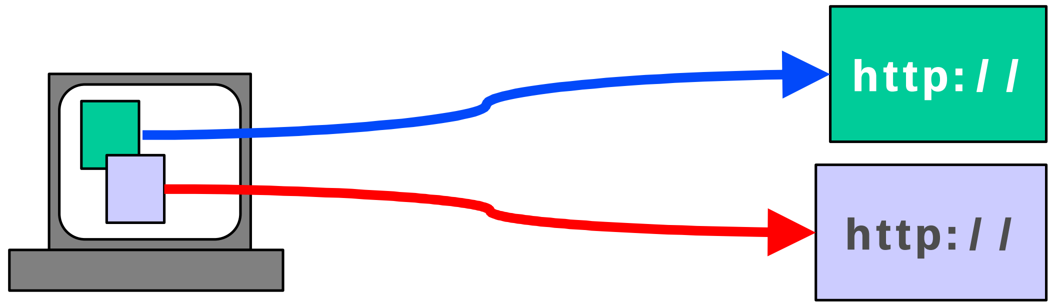

Grid: Flexible and High Performance access to (any kind of) resources

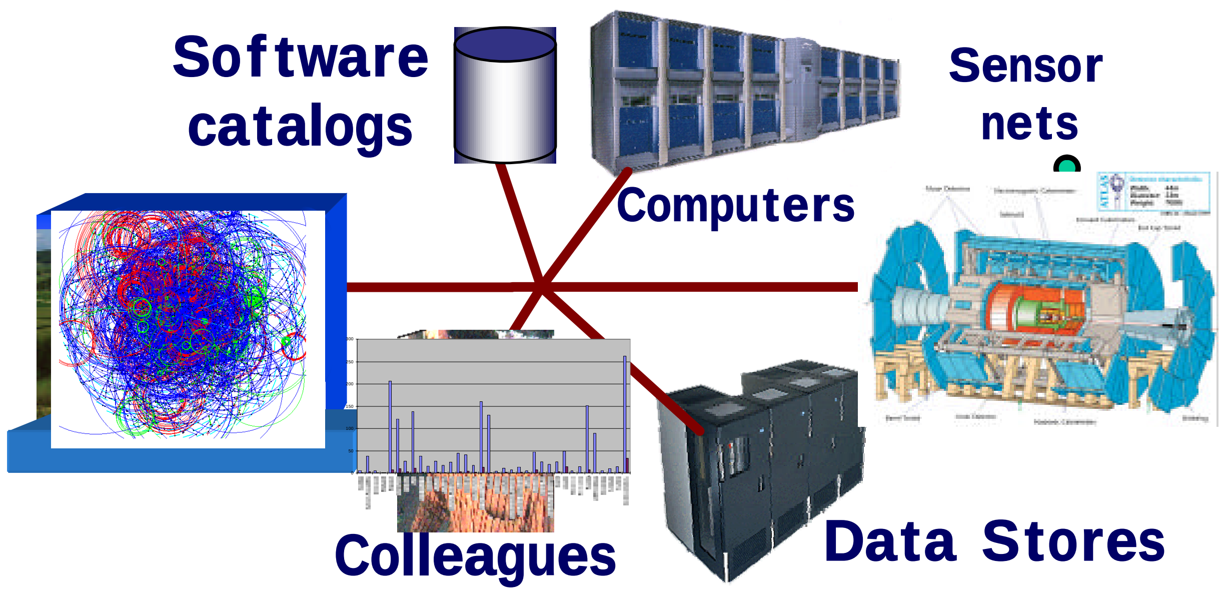

On-demand creation of powerful virtual computing and data systems

A Power GRID for Computing

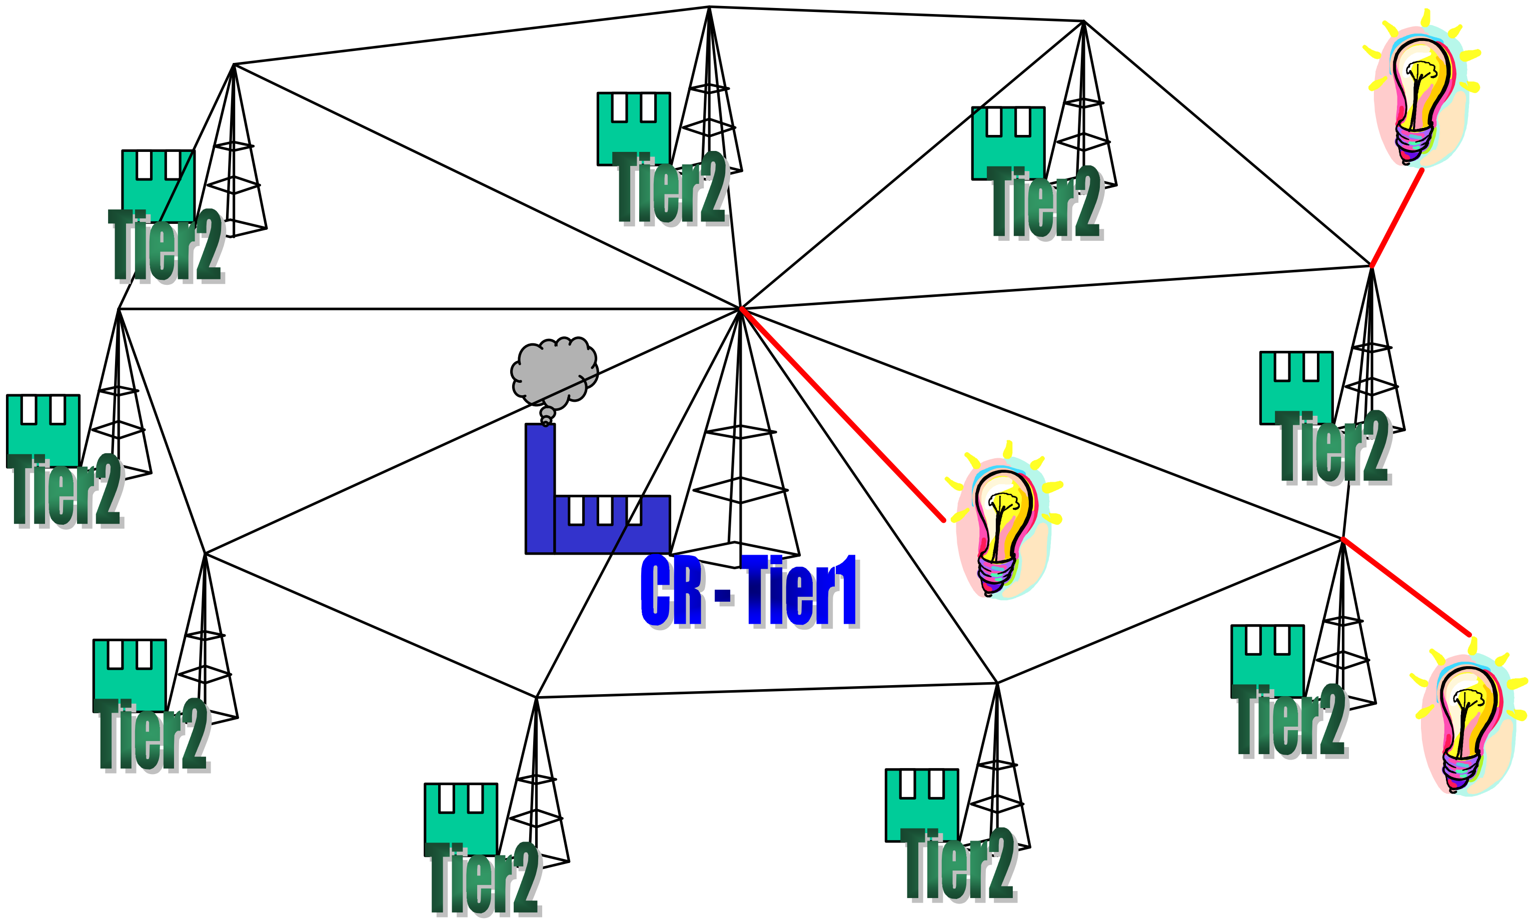

DataGRID Layered structure (2000)

Applications	Chemistry; Cosmology; Environment; Biology; High Energy Physics
Application Toolkits	Distributed Computing Toolkit; Data-Intensive Applications Toolkit; Collaborative Applications Toolkit; Remote Visualization Applications Toolkit; Problem Solving Applications Toolkit; Remote Instrumentation Applications Toolkit
Grid Services (Middleware)	<u>Resource-independent and application-independent services</u> authentication, authorization, resource location, resource allocation, events, accounting, remote data access, information, policy, fault detection
Grid Fabric (Resources)	<u>Resource-specific implementations of basic services</u> E.g., Transport protocols, name servers, differentiated services, CPU schedulers, public key infrastructure, site accounting, directory service, OS bypass

Natural HEP Farming
High Throughput Computing

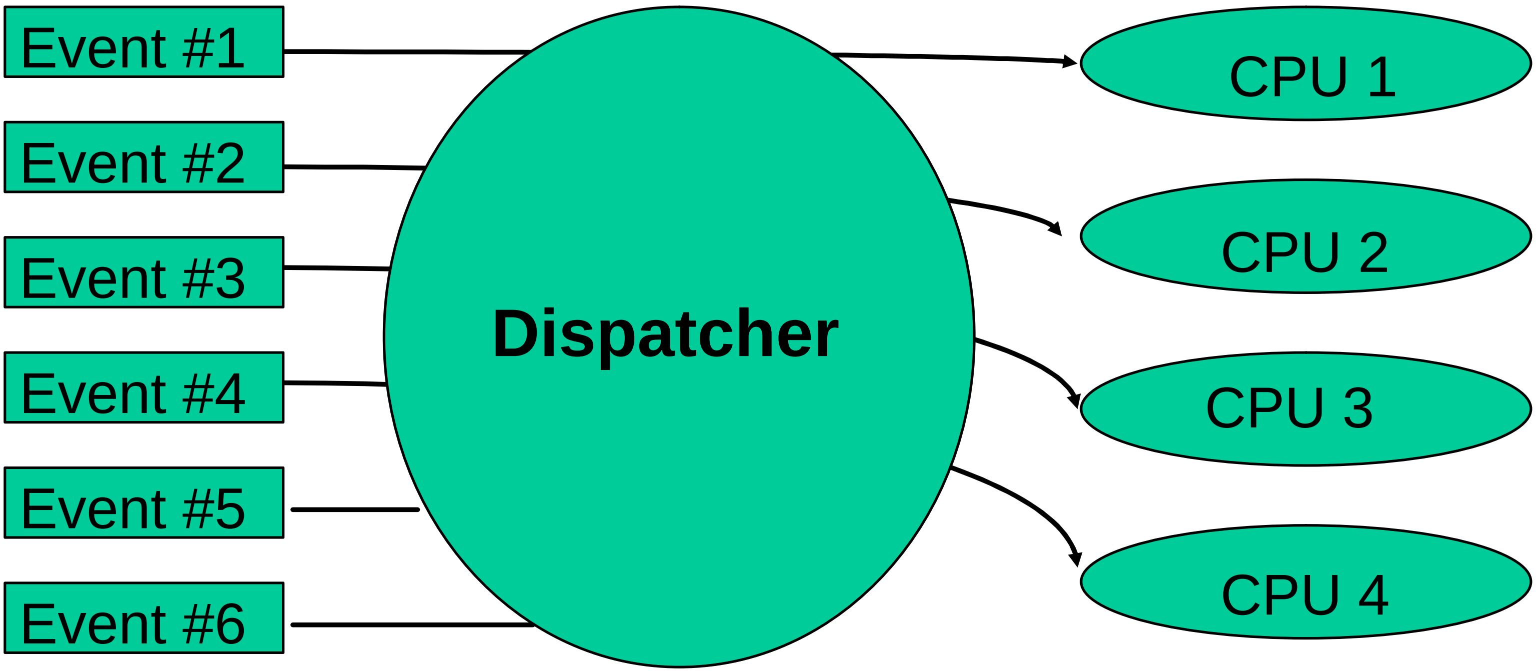

GRID Computing Farm

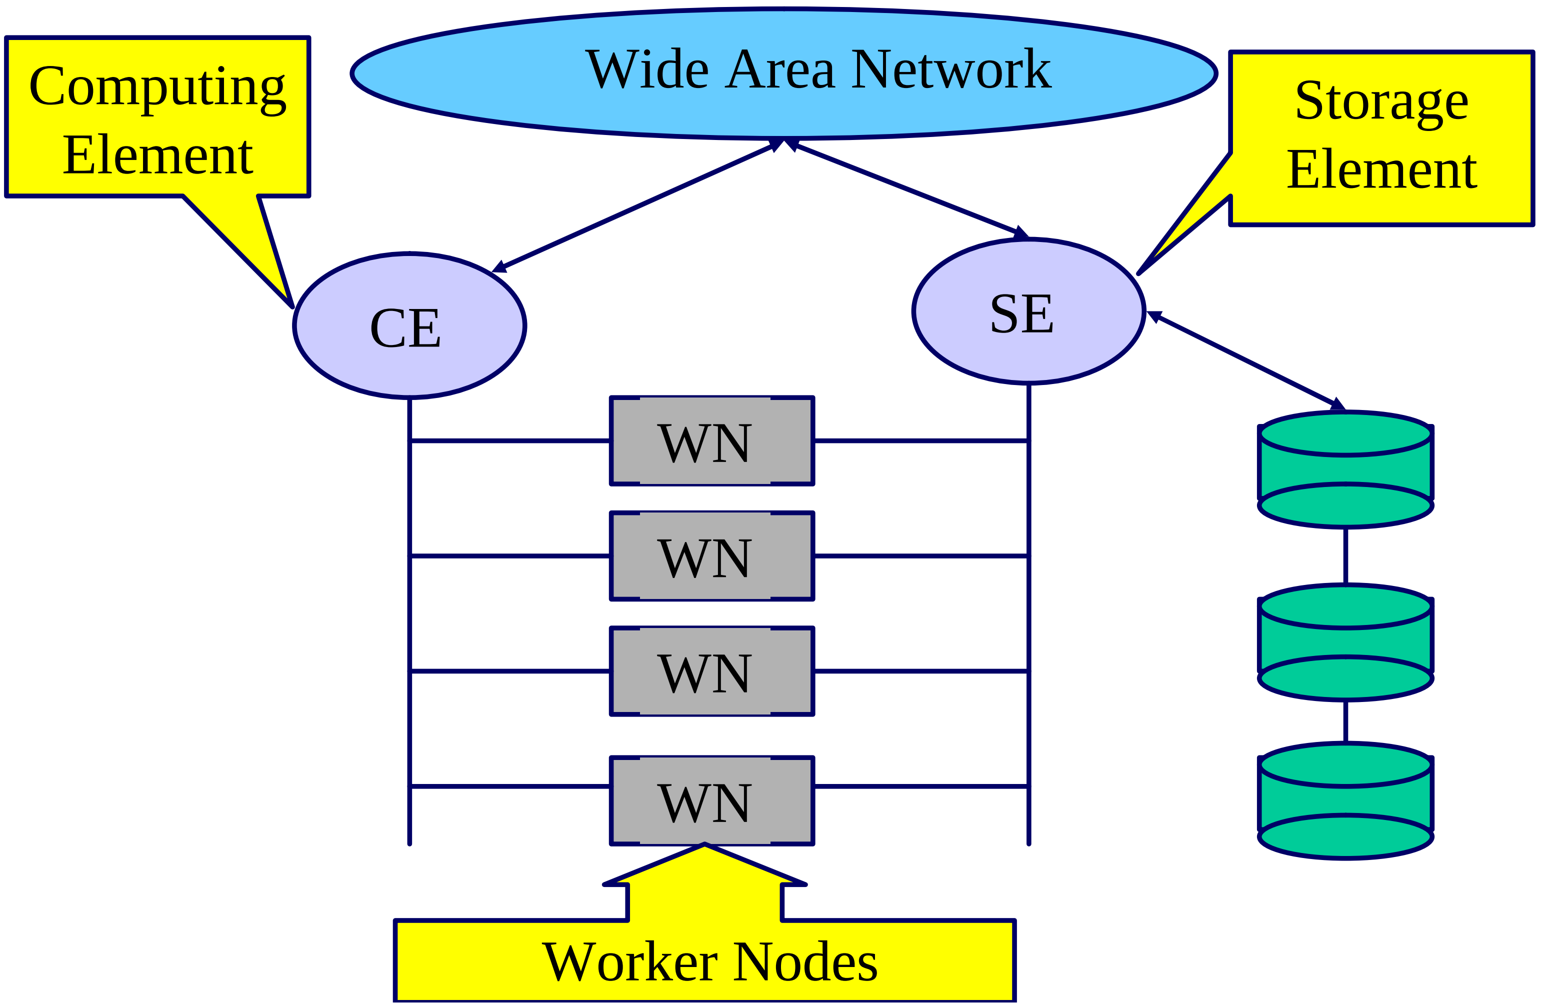

(Some) GRID Services

- Workload Management System (Resource Broker) chooses the best resources matching the user requirements.
- Virtual Organization Management System allows to map User Certficates with VO's describing rights and roles of the users.
- Data Oriented Services: Data & Meta-data Catalogs, Data Mover, Replica Manager, etc.
- Information & Monitoring Services which allow to know which resources and services are available and where.
- Accounting services to extract resouce usage level related to users or group of users and VO's.

- **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 started 1st April 2004 -> EGEE II on 1 April 2006

EGEE is the major Grid project in Europe and operates the biggest infrastructure for e-Science world-wide

EGEE Project

- CERN
- Central Europe (Austria, Czech Republic, Hungary, Poland, Slovakia, Slovenia)
- France
- Germany and Switzerland
- Ireland and UK
- Italy
- Northern Europe (Belgium, Denmark, Estonia, Finland, The Netherlands, Norway, Sweden)
- NRENs
- Russia
- South-East Europe (Bulgaria, Cyprus, Greece, Israel, Romania)
- South-West Europe (Portugal, Spain)
- USA

EGEE Grid Sites : November 2005

- Many more sites than anticipated for this stage of project
  - 179 actual, cf. 50 proposed for end of year 2
  - ~2000 CPU in sites outside of EGEE federations (7 countries)
- Includes industrial partner sites (HP in Puerto Rico and UK)
- Exposes full complexity of grid operations

EGEE:

179 sites, 39 countries
>17,000 processors,
~5 PB storage

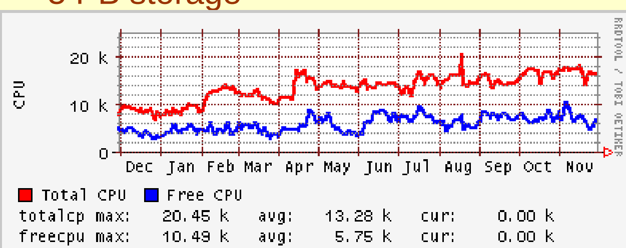

country	sites	country	sites	country	sites
Austria	2	India	2	Russia	12
Belgium	3	Ireland	15	Serbia	1
Bulgaria	4	Israel	3	Singapore	1
Canada	7	Italy	25	Slovakia	4
China	3	Japan	1	Slovenia	1
Croatia	1	Korea	1	Spain	13
Cyprus	1	Netherlands	3	Sweden	4
Czech Republic	2	Macedonia	1	Switzerland	1
Denmark	1	Pakistan	2	Taipei	4
France	8	Poland	5	Turkey	1
Germany	10	Portugal	1	UK	22
Greece	6	Puerto Rico	1	USA	4
Hungary	1	Romania	1	CERN	1

Operations Process

- **CIC – on – duty (grid operator on duty)**
  - Started November 2004
  - SFTs
  - 6 teams working in weekly rotation
    - CERN, IN2P3, INFN, UK/I, Ru,Taipei
  - Crucial in improving site stability and management
- **Operations coordination**
  - Regular ROC, CIC managers meetings
  - Series of EGEE Operations Workshops
    - Last one was a joint workshop with Open Science Grid
  - These have been extremely useful
    - Bring in related infrastructure projects – coordination point
- **Geographically distributed responsibility for operations:**
  - There is no “central” operation
  - Tools are developed/hosted at different sites:
    - GOC DB (RAL), SFT (CERN), GStat (Taipei), CIC Portal (Lyon)

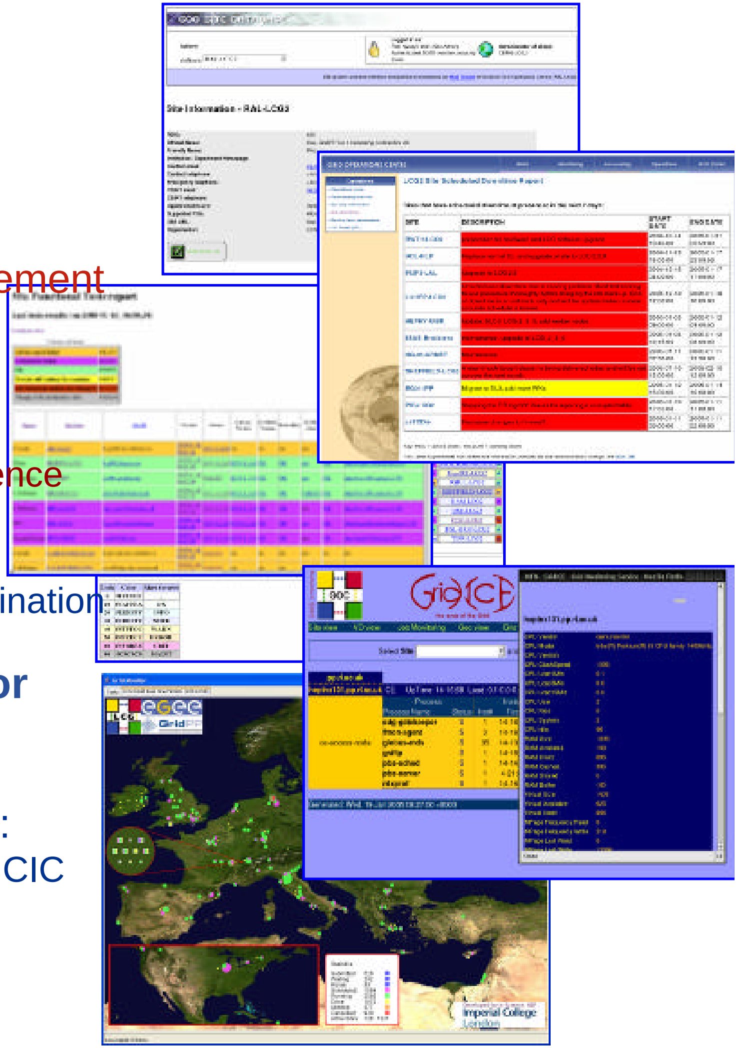

EGEE Highlights - Applications

>20 applications from 7 domains

- High Energy Physics
- Biomedicine
- Earth Sciences
- Computational Chemistry
- Astronomy
- Geo-Physics
- Financial Simulation

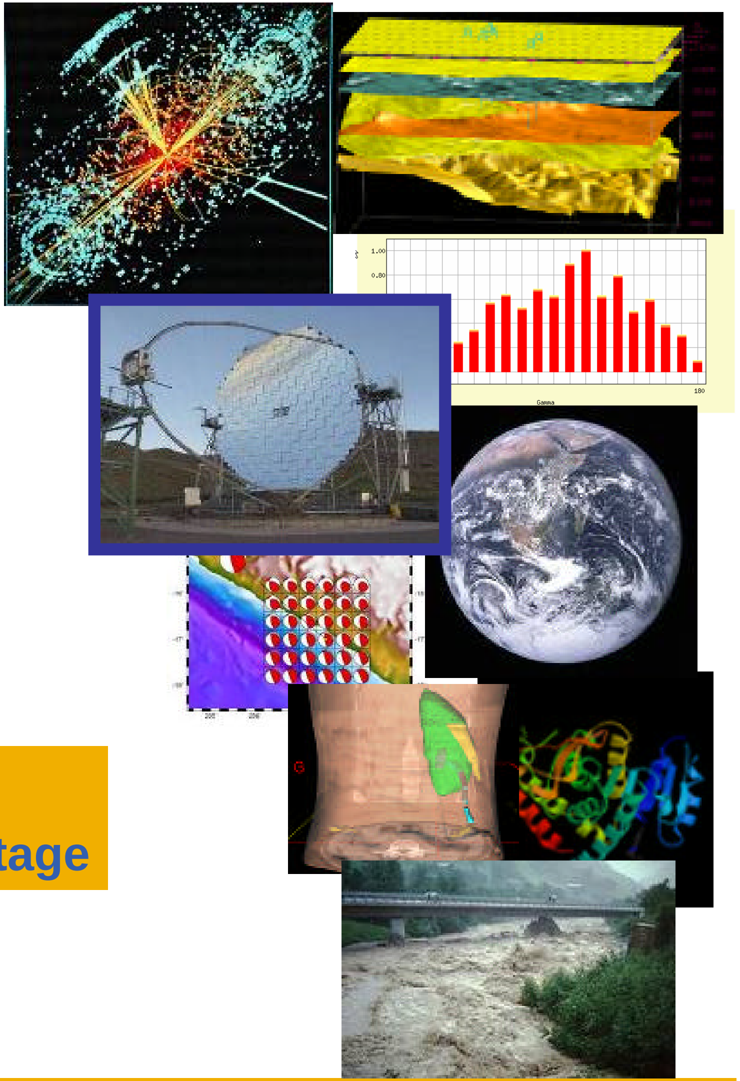

Another 8 applications from 4 domains are in evaluation stage

MAGIC, cosmic physics

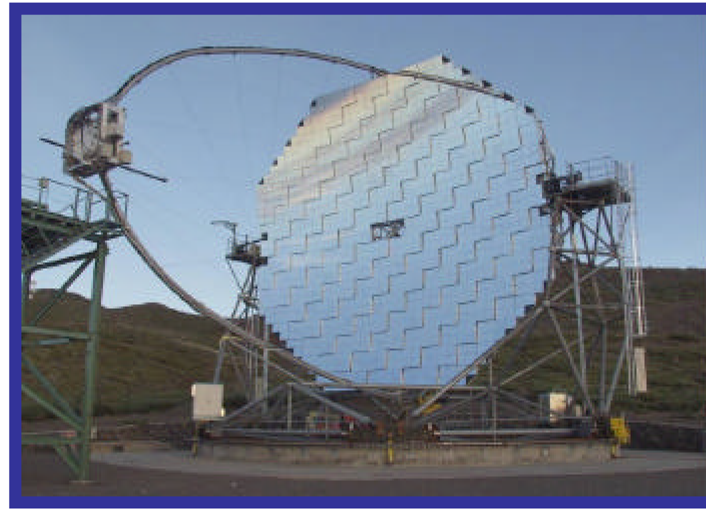

- Ground based Air Cerenkov Telescope 17 m diameter
- LaPalma, Canary Islands (28° North, 18° West)
- Physics Goals:
  - Origin of VHE Gamma rays
  - Active Galactic Nuclei
  - Supernova Remnants
  - Unidentified EGRET sources
  - Gamma Ray Burst

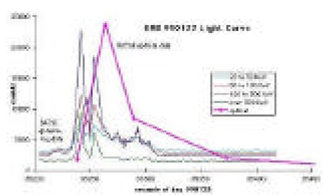

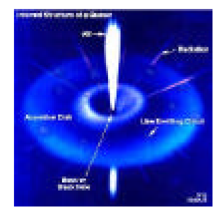

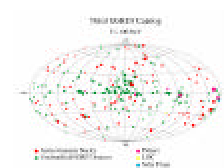

- MAGIC II will come 2007
- Grid added value
  - Enable "(e-)scientific" collaboration between partners
  - Enable the cooperation between different experiments
  - Enable the participation on Virtual Observatories

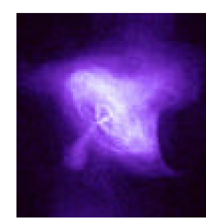

MAGIC – Why the Grid?

MAGIC is an international collaboration

- Partners distributed all over Europe
- Amount of data can NOT be handled by one partner only (up to 200 GB per night)
- Access to data and computing needs to be more efficient
- MAGIC will build a second telescope

Analysis is based on Monte Carlo simulations

- CORSIKA code
- **CPU consuming**
  - 1 night of hadronic background needs 20000 days on 70 computer
- Lowering the threshold of MAGIC telescope requires new methods based on MC simulations
- More CPU power needed!

The Planck Mission

- **Measure cosmic microwave background (CMB)**
  - succeeds COBE, Boomerang & WMAP missions
  - aims at even higher resolution
- **Timeline**
  - launch August 2007
  - start of observations 2008
  - duration >1 year
- **Characteristics**
  - continuous data stream (TOD)
  - large datasets
    (a TOD of ~7 TB for the whole LFI mission)
  - changing calibration
    (parameters configuration)
  - high-performance computing for data analysis

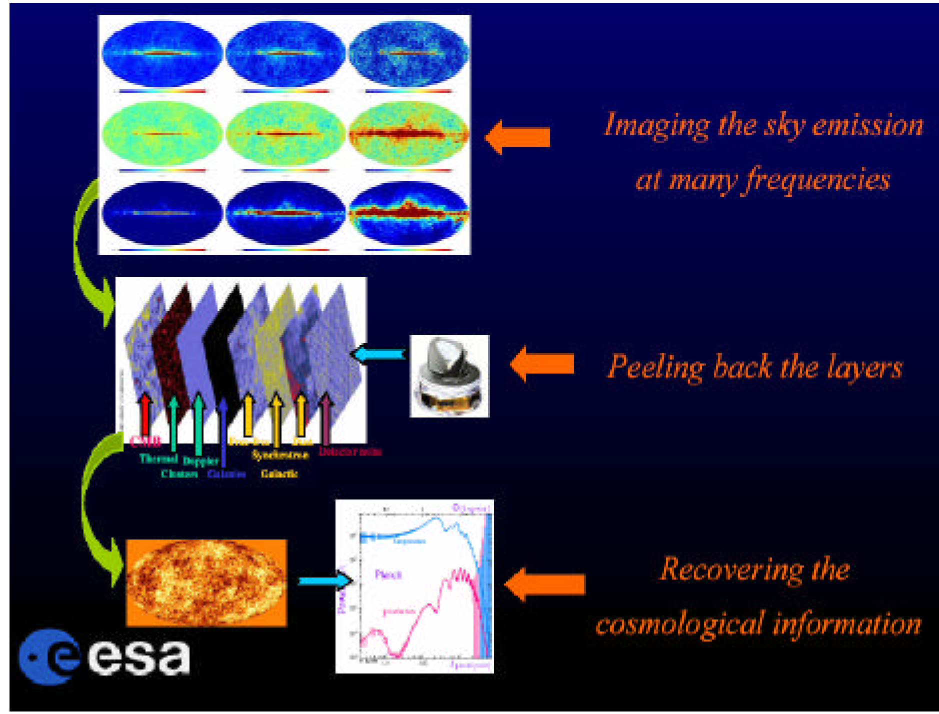

Grid added values to Plank

- **CPU power:**
  - E-computing lab;
  - Production burst;
  - Efficient CPU usage/sharing.
- **Data storing/sharing:**
  - Distributed data for distributed users;
  - Replica and security;
  - Common interface to software and data.

Planck simulations are highly computing demanding and produce a huge amount of data. Such resources cannot be usually afforded by a single research institute, both in terms of computing power and data storage space.

Worldwide Ozone distribution

Earth Observation by satellite: large community

– Space Agencies
– Scientists in many fields
– Organization in charge of natural risk alerts
– SMEs for elaborated products

- Grid Technology seems very well adapted to a lot of applications related to Earth Observations
- ESA/ESRIN, IPSL and KNMI have started to port applications in DataGrid

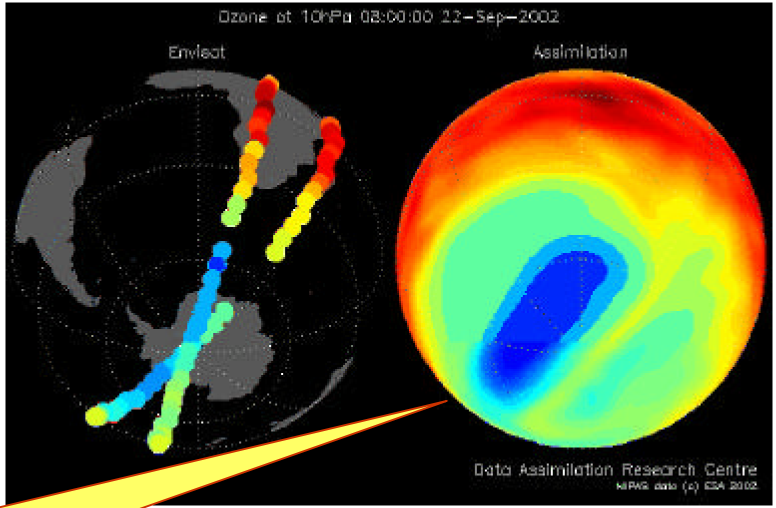

GOME total ozone assimilation

- Figure from ESRIN

SCAI
Institut
Algorithmen und Wissen-
schaftliches Rechnen

In silico Drug Discovery

- **Scientific objectives**
  - **Provide docking information helping in search for new drugs.**
  - **Biological goal: propose new inhibitors (drug candidates) addressed to neglected diseases.**
  - **Bioinformatics goal: *in silico* virtual screening of drug candidate DBs.**
  - **Grid goal : demonstrate to the research communities active in the area of drug discovery the relevance of grid infrastructures through the deployment of a compute intensive application.**
- **Method**
  - **Large scale molecular docking on malaria** to compute million of potential drugs with some software and parameters settings.
  - Docking is about computing the binding energy of a protein target to a library of potential drugs using a scoring algorithm.

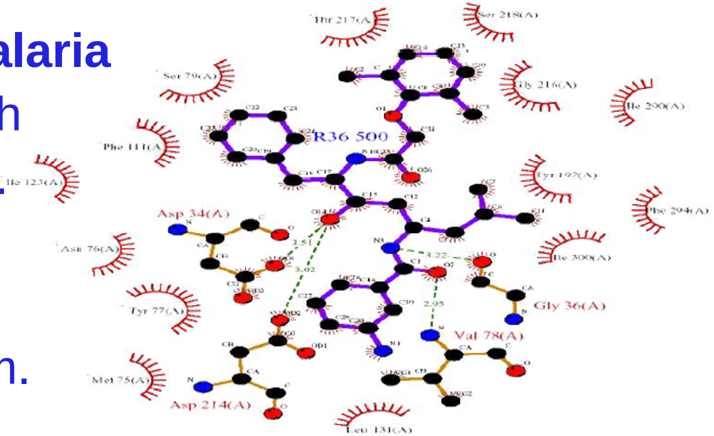

In silico Drug Discovery

- **Grid added value**

  Drug discovery lead by pharmaceutical companies takes up to 12 years to complete. **Molecular docking** has the potential to drastically **speed-up** this process but considering large databases yield to **heavy computations**.

  The **computations** involved can be **distributed** on grid nodes by splitting the candidate drug input on different grid resources. The **data management** services will facilitate the storage and the post-processing of the output files

- **Results and perspectives**

  First experiments have shown that a limited size computation (10^5 candidate drugs tested against 1 protein target) are achievable in 2 days using the EGEE infrastructure compared to 6 months of CPU time involved.

  A full data challenge is planed that should involve $3x10^6$ candidate drugs to be tested against 5 protein target structures. The total computing time is expected to reach 80 years of CPU and 6 TB of storage.

Biology Applications: The “never born” proteins

- Natural proteins are only a tiny fraction of the possible ones
  - Approx. 10^{13} natural proteins *vs* 20^{100} possible proteins with a chain length of 100 amino acids !
- Does the subset of natural proteins has particular properties?
- Do exist in principle protein scaffolds with novel structure and/or activity not yet exploited by Nature?
- GRID technology allows to tackle the problem through high throughput prediction of protein structure of a large library of “never born proteins”

Social Impact

- A large part of the Globe has not advanced digital infrastructures yet.
- The European Research Area program wants to set Europe as the most advanced region in eInfrastructures and promote the take-up to speed of other less advanced countries to alleviate as much as possible the so called Digital Divide.
- eInfrastructures support wide geographically distributed communities which share problems and resources to work towards common goals -> enhance international collaboration of scientists -> promote collaboration in other fields.
- Problems too big to be handled with conventional local computer clusters and time sharing computing centers can be attacked with GRIDs.
- eInfrastructures are leveraging international network interconnectivity -> High Bandwidth connections will improve exchange of knowledge and be the basis for GRID Infrastructures.
- Based on safe AAA (Authentication, Authorization and Accounting) architecture -> secure and dependable infrastructures.
- Need of persistent software & middleware -> Software is integral part of the infrastructure.

Digital Divide

http://maps.maplecroft.com/

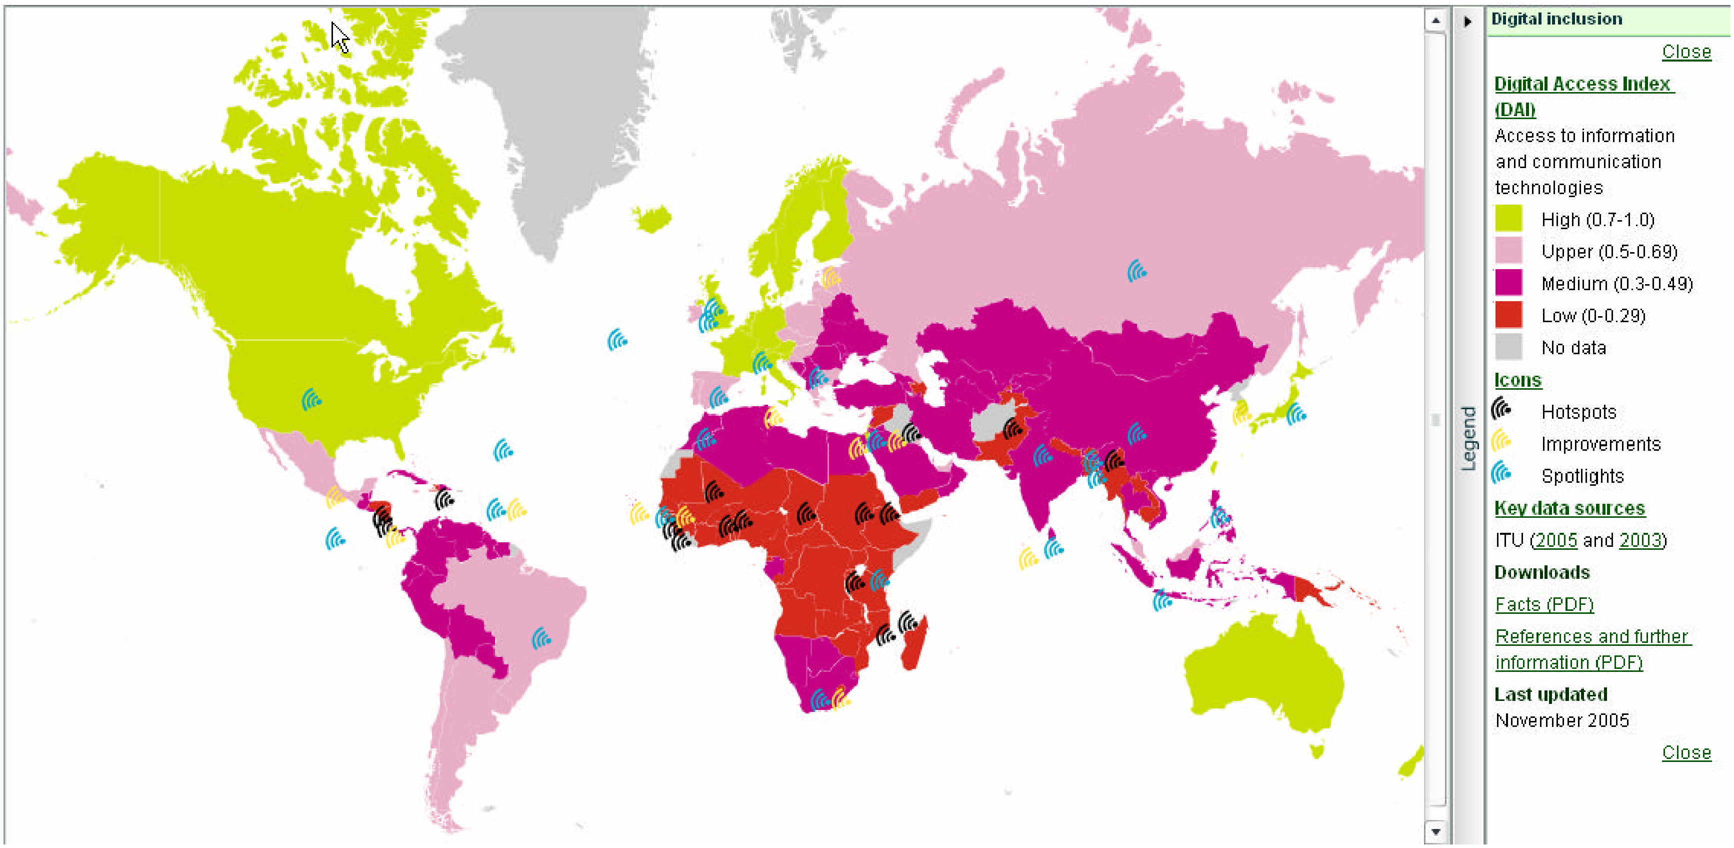

Docking on Malaria

- Of a large interest for many developing countries.
- Based on Grid-enabled drug discovery process.
- Data challenge proposal never done on a large scale production infrastructure and for a neglected disease
  - 5 different structures of the most promising target
  - Output Data: 16,5 million results, ~10 TB

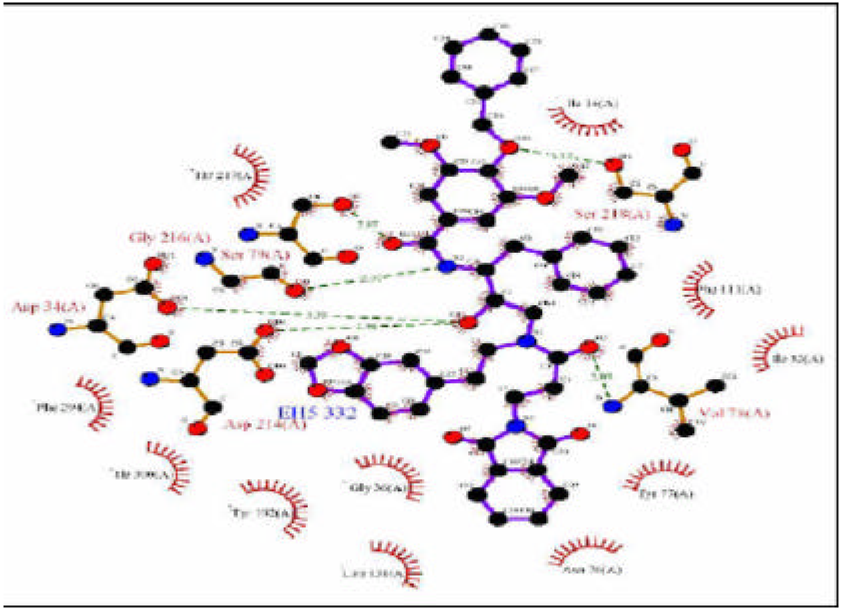

FP6- 2004- Infrastructures- 6-SSA-026024

EUMEDGRID Geography

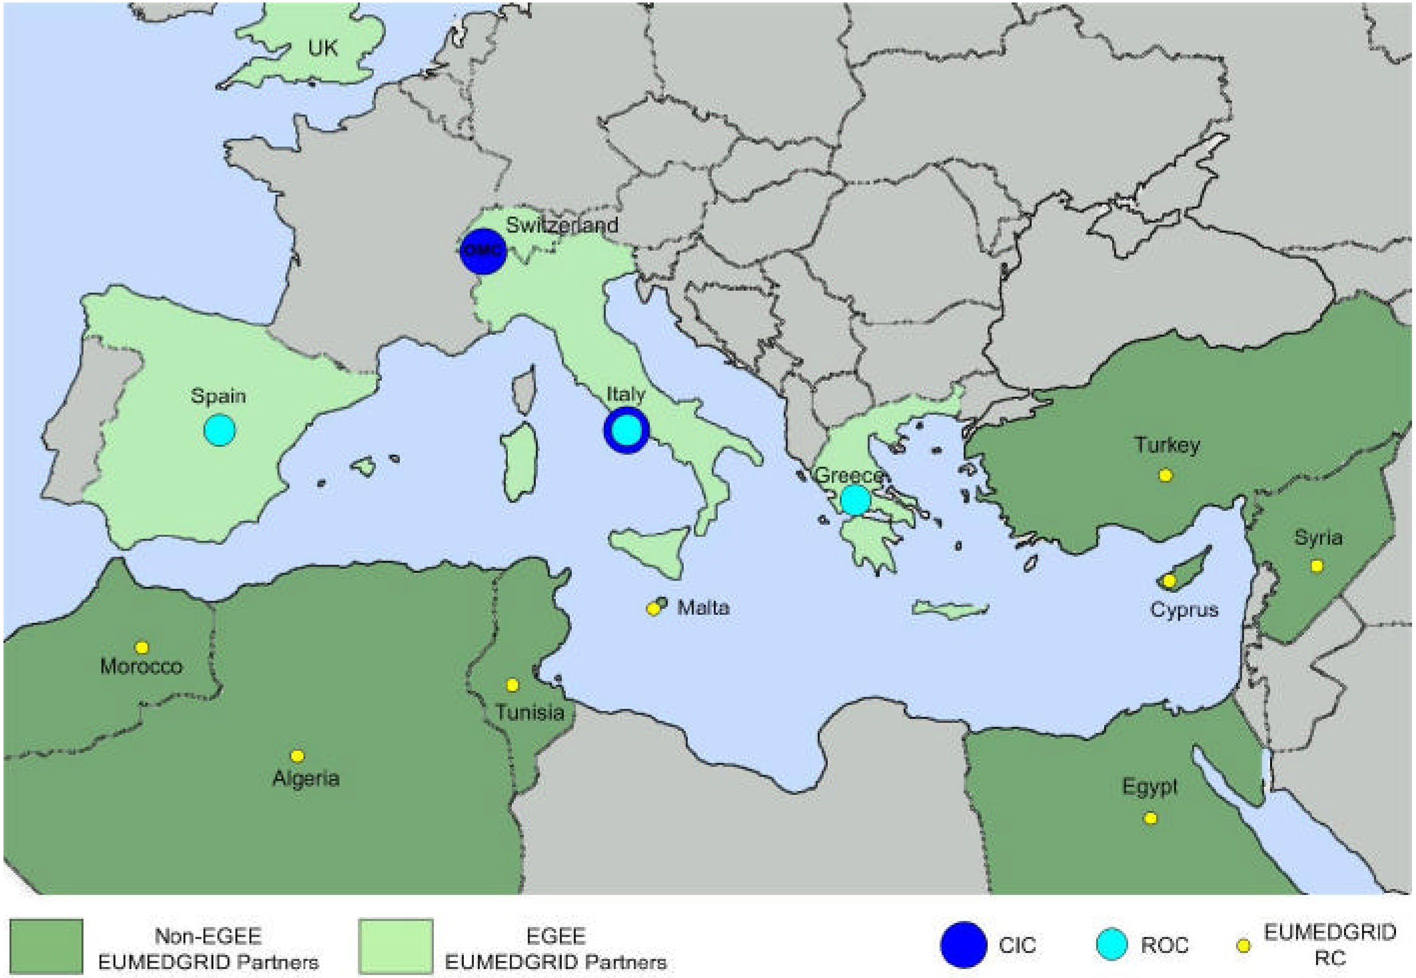

Linking Mediterranean research and educational communities to Europe

INFN

EUMED connect

www.eumedconnect.net

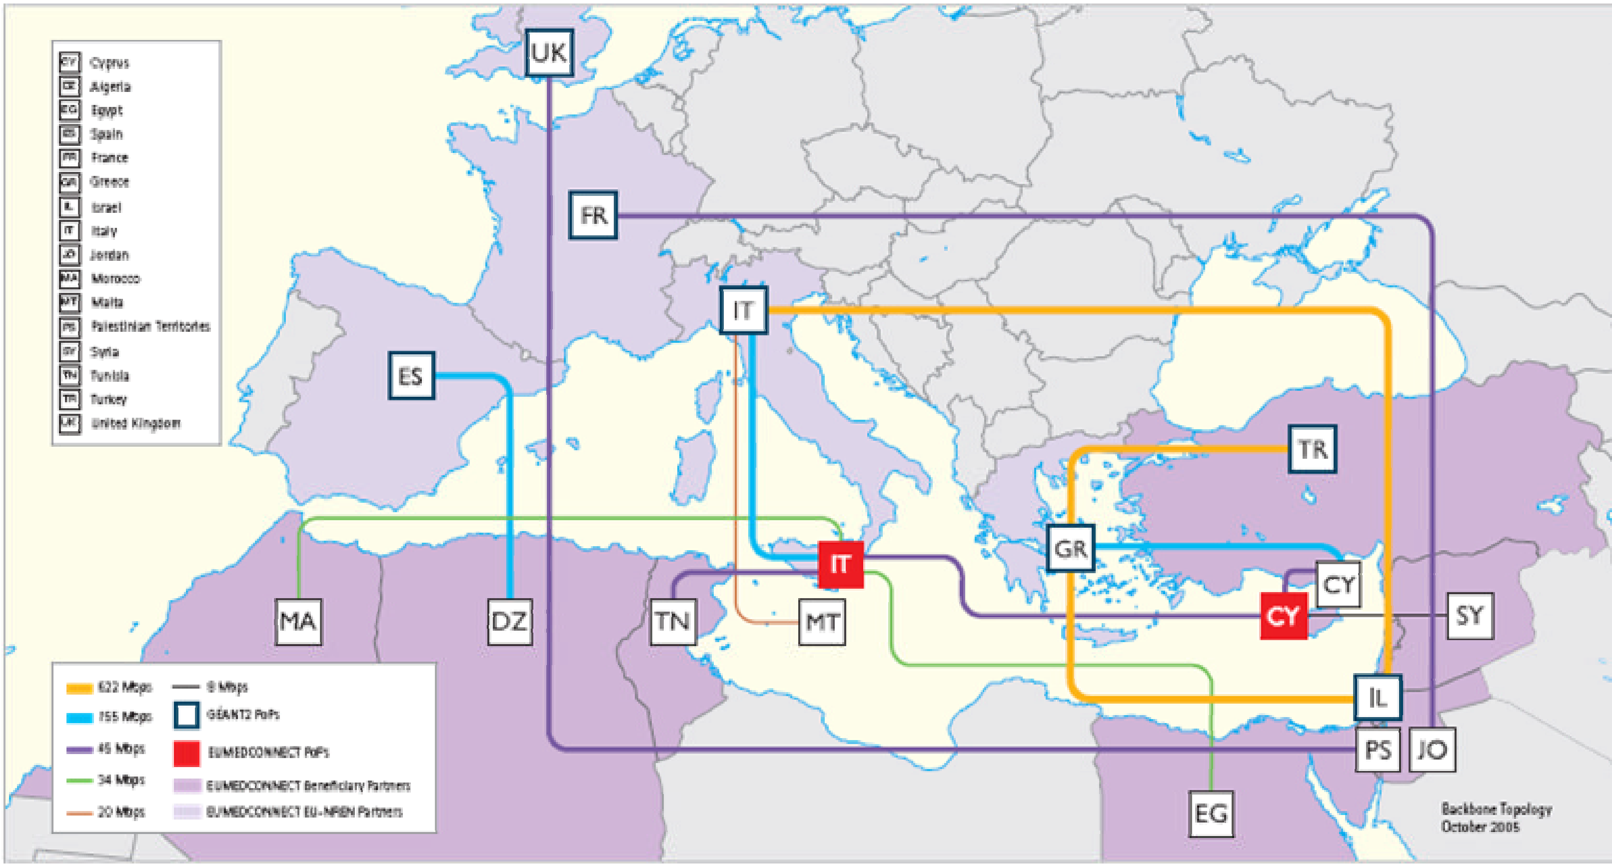

DANTE www.dante.net

Consortium GARR

GRNET

Renater

Creating a research and education network for the Mediterranean

GÉANT2 www.geant2.net

EUMEDCONNECT is co-funded by the European Commission through the EUMEDIS programme

eumedis

EuropeAid

EELA Project

Conclusions

- Grids infrastructure are an expanding reality.
- They can stimulate new aggregations of scientists working together on new challenges which are now made affordable.
- They are the basis of eInfrastructures which can promote high bandwidth networks and make a little step forward to fight Digital Divide in the developing countries.
- But nothing comes for free; you need to know who is using the (your) resources and for which purpose. You need security, accounting and, eventually, billing systems.
- Long Term Sustainability of such a huge investment needs Governmental Priorities and a strong Industrial Uptake.

Thanks for your kind invitation and attention…