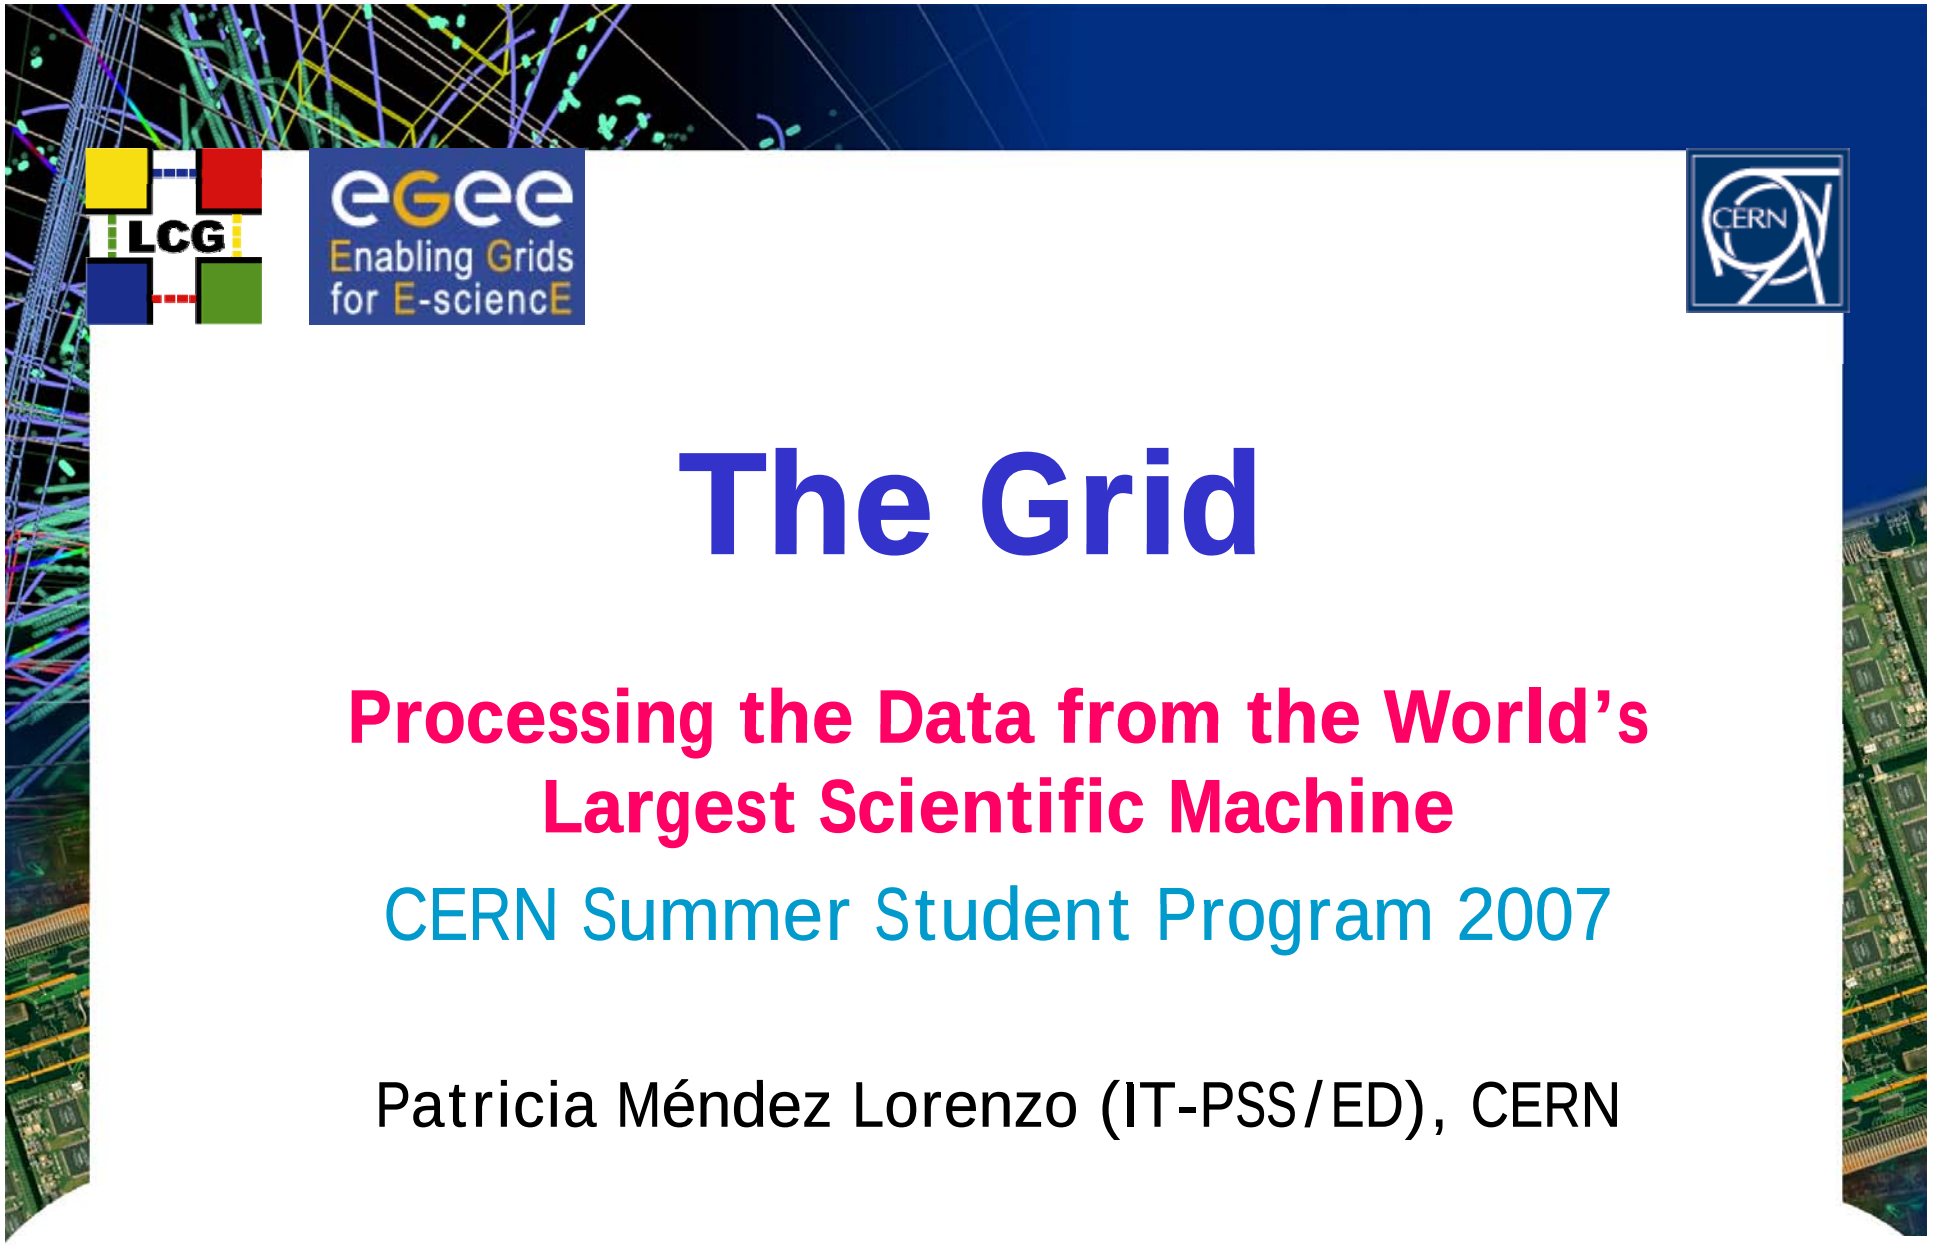

# The Grid

## Processing the Data from the World's Largest Scientific Machine

CERN Summer Student Program 2007

Patricia Méndez Lorenzo (IT-PSS/ED), CERN

Worldwide LHC Computing Grid

Distributed Production Environment for Physics data Processing

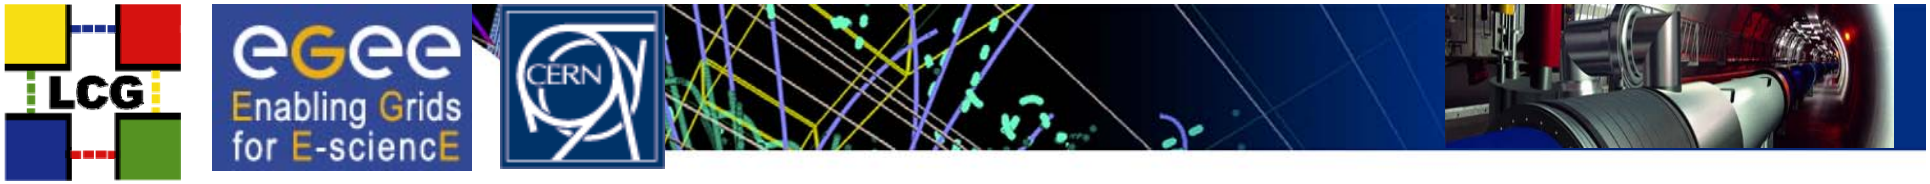


# Abstract


- The world's **largest scientific machine** will enter production about one year from the time of this talk
- In order to exploit its scientific potential, computational resources way beyond those needed for previous accelerators are required
- Based on these requirements, a distributed solution based on **Grid technologies** has been developed
- This talk describes the overall **Grid infrastructure** offered to the experiments, the state of **deployment** of the production services, and the applications of the Grid beyond the HEP

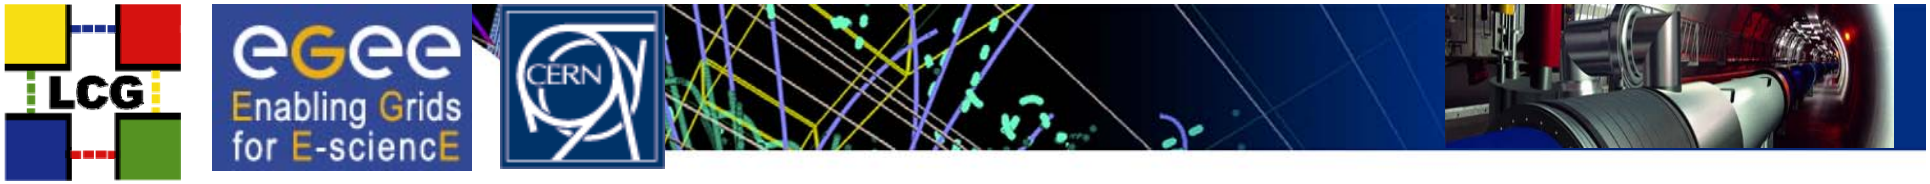


# Several considerations before beginning…

- I assume that you have never submitted a job to the Grid
  - No large knowledge of computing is required
- **Announcement to the Physicist!**
  - "…This is a computer related talk, no interest for me…" **(WRONG!)**
  - If you do not know the software infrastructure of your experiment and you do not know computing, you will have big problems in your future work as researcher
- This presentation does not pretend to be a tutorial
  - I will avoid as much as possible the use of code
- It will show you the infrastructure that many of you will have to face as researchers and it begins to be expanded in many fields of research (not only HEP)

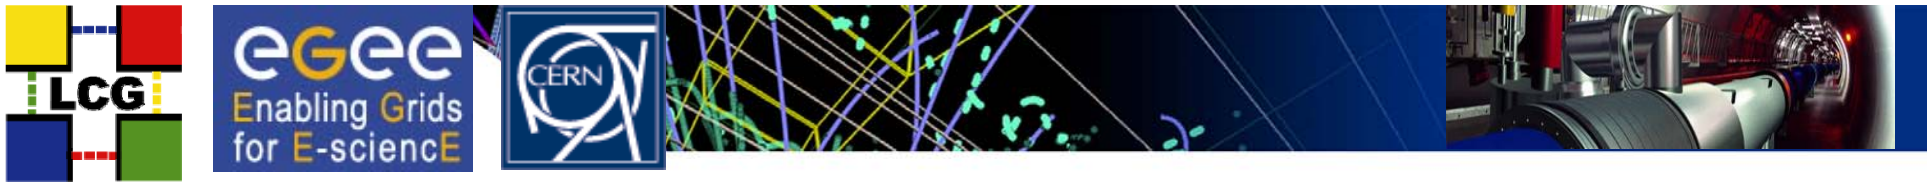


# Overview

- First Part: The general concept of the Grid
- Second Part: The WLCG Grid Computing
- Third Part: The Infrastructure and the Services
- Fourth Part: Current Status
- Fifth Part: Applications beyond HEP

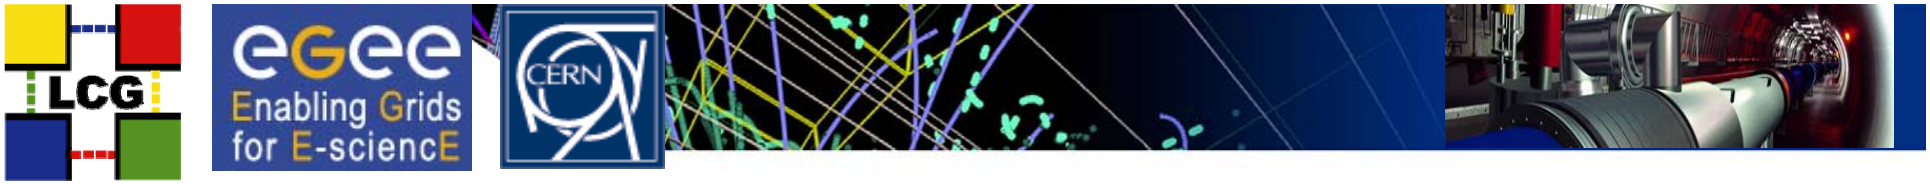


# 1st Part

- General Concept of the GRID

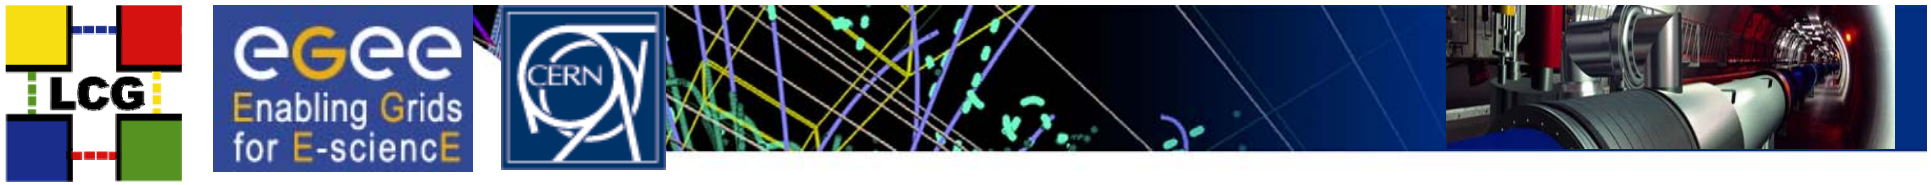


# 1st Part: What is a Grid

- A genuine new concept in distributed computing
  - Could bring radical changes in the way people do computing
  - **Share** the computing power of many countries for your needs
  - **Decentralized** the placement of the computing resources
- A hype that many are willing to be involved in
  - Many researchers/companies work on "grids"
    - More than 100 projects (some of them commercial)
    - Only few large scale deployments aimed at production
  - Very confusing (hundreds of projects named Grid something)
- Interesting links: Ian Foster, Karl Kesselman
  - Or have a look at http://globus.org
    - Not the only grid toolkit, but one of the first and most widely used
  - Web pages of: EGEE, WLCG (we will talk about them)

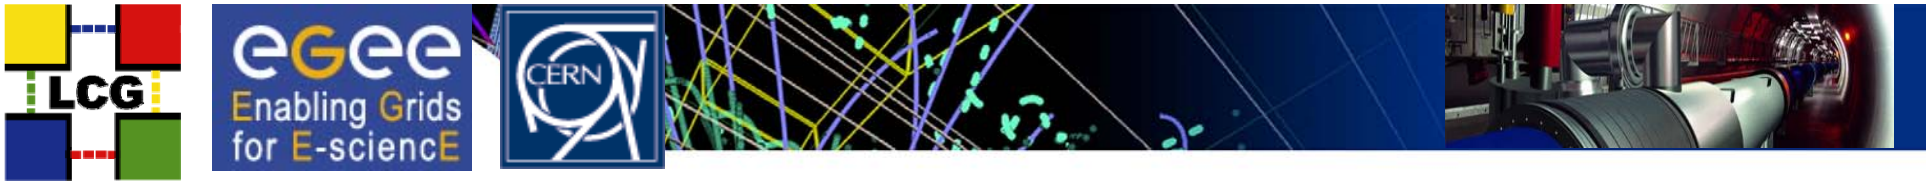


# What is a Grid (cont.)

- Basic concept is simple:
  - I.Foster: "coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations. "
  - Or:"On-demand, ubiquitous access to computing, data, and services"
  - From the user's perspective:
    - I want to be able to use computing resources as I need
    - I don't care who owns resources, or where they are
      - Have to be secure
      - My programs have to run there
  - The owners of computing resources (CPU cycles, storage, bandwidth)
    - My resources can be used by any authorized person (not for free)
    - Authorization is not tied to my administrative organization
  - **NO centralized control of resources or users**

**Basic Philosophy**

# How to see the Grid

- The Grid connects Instruments, Computer Centers and Scientists

- If the WEB is able to share information, **the Grid is intended to share computing power and storage**

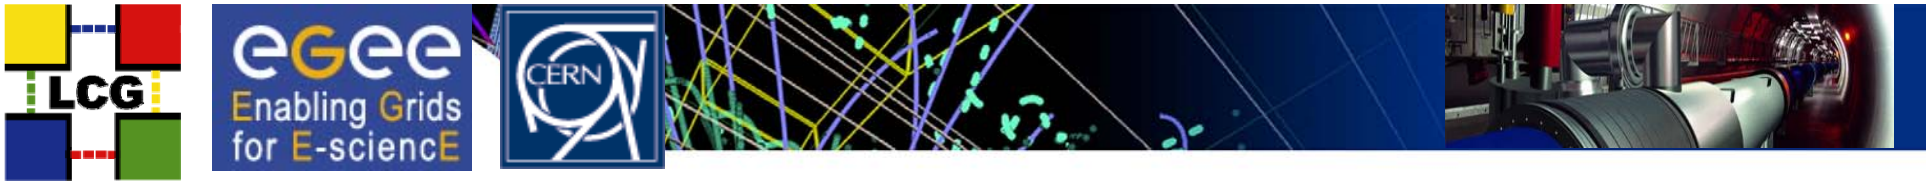


# 2nd Part

- The WLCG GRID Computing

# The LHC Grid Computing

- LHC will be completed in 2008 and run for the next 10-15 years
- Experiments will produce about **15 Millions Gigabytes per year** of data (about 20 million CDs!)
- LHC data analysis requires a computing power equivalent to around **100000 of today's fastest PC processors**
- R
  c
  o
  c

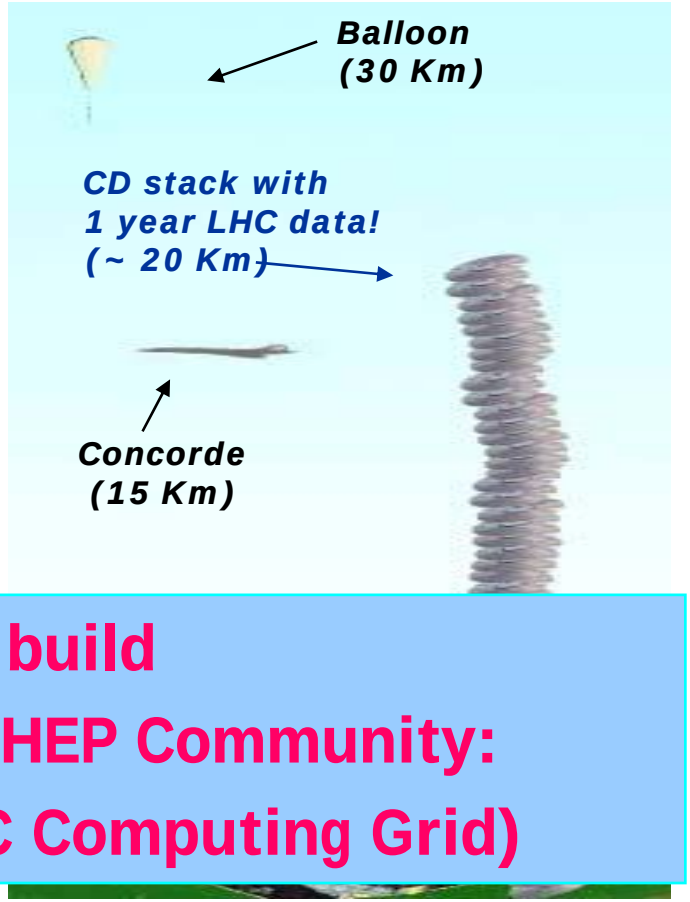


**Therefore we build**

**a Computing Grid for the HEP Community:**

**The WLCG (Worldwide LHC Computing Grid)**

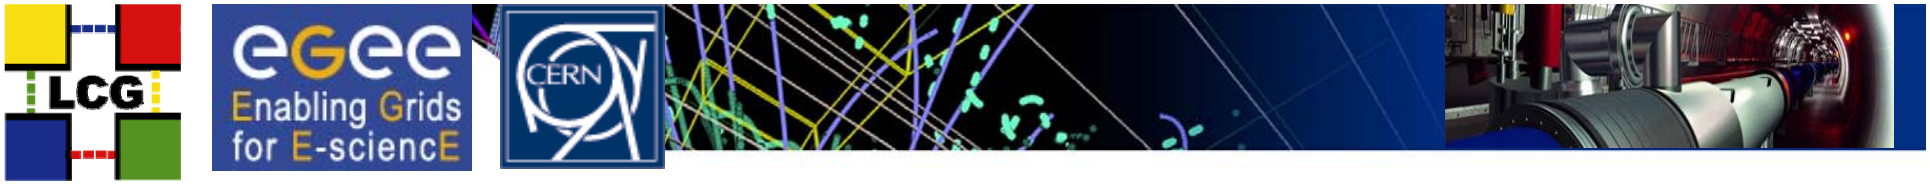


# Many reasons to use the GRID

The LCG Technical Design Report lists:

1. Significant costs of maintaining and upgrading the necessary resources … more easily handled in a distributed environment, where individual institutes and organizations can fund local resources whilst contributing to the global goal
2. … no single points of failure. Multiple copies of the data, automatic reassigning of tasks to resources … facilities access to data for all scientists independent of location. … round the clock monitoring and support.

# Projects we use for LHC Grid

- WLCG depends on 2 major science grid infrastructures provided by:
  - EGEE – Enabling Grid for E-SciencE (160 communities)
  - OSG – US Open Science Grid (also supporting other communities beyond HEP)
  - Both infrastructures are federations of independent GRIDs
  - WLCG uses both OSG resources and many (but not all) from EGEE

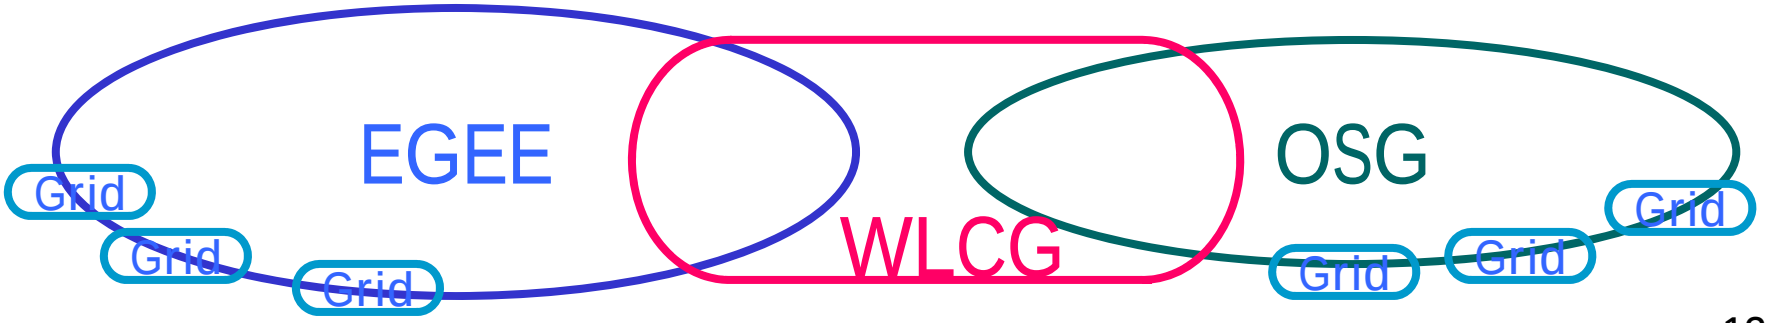

**EGEE:**
Steady growth over the lifetime of the project

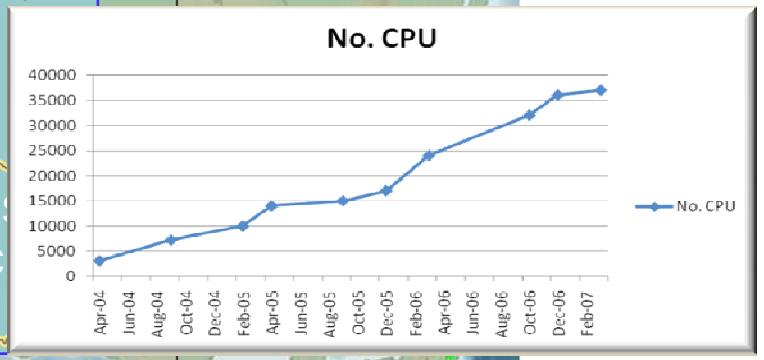


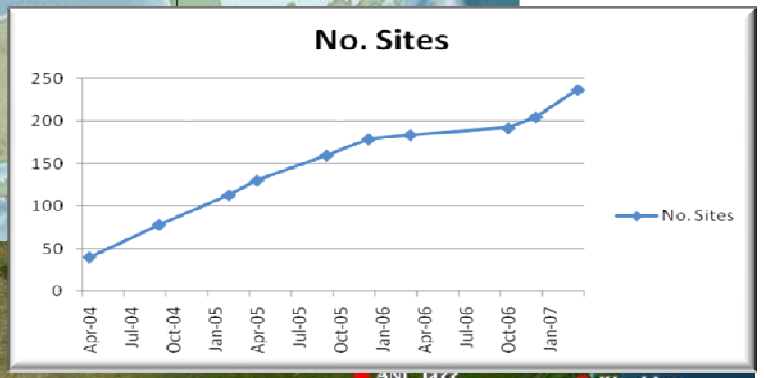


**EGEE:**
> 237 sites, 45 countries
> 35,000 processors,
~ 10 PB storage

Patricia Méndez Lorenzo

# The WLCG Distribution of Resources

## Tier-0 – the accelerator centre

- Data acquisition and initial Processing of raw data
- Distribution of data to the different Tier's

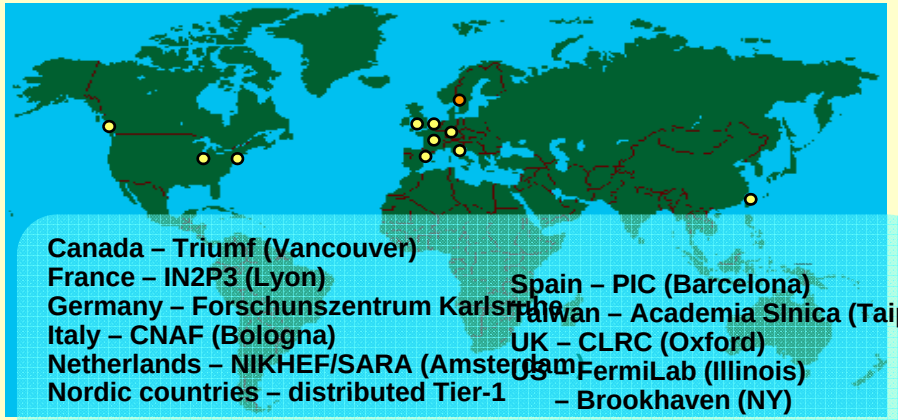


## Tier-1 (11 centers ) – "online" to the data acquisition process → high availability

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

## Tier-2 – ~200 centres in ~40 countries

- Simulation
- End-user analysis – batch and interactive

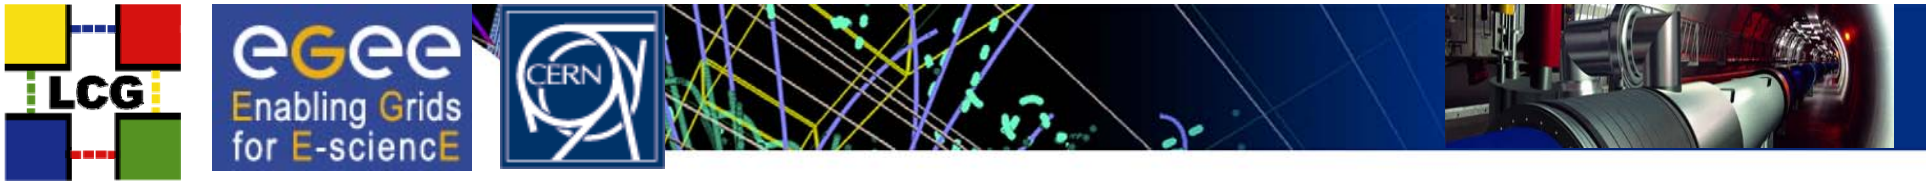


# From the raw data… to the paper

- **Reconstruction**: transform signals from the detector to physical quantities
  - energy, charge, tracks, momentum, particle id.
  - this task is computational intensive and has modest I/O requirements
  - structured activity (production manager)
- **Simulation**: start from the theory and compute the response of the detector
  - very computational intensive
  - structured activity, but larger number of parallel activities
- **Analysis**: complex algorithms, search for similar structures to extract physics
  - very I/O intensive, large number of files involved
  - access to data cannot be effectively coordinated
  - iterative, parallel activities of hundreds of physicists

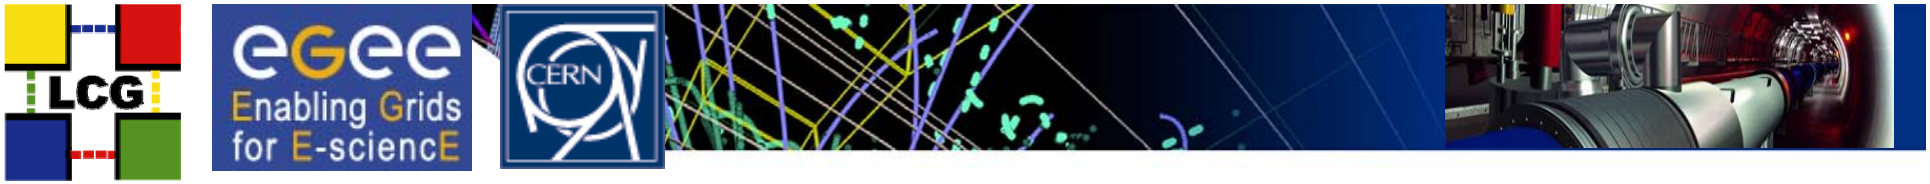


# 3rd Part

- The Infrastructure and the Services

# 3rd Part: Who is who in the Grid

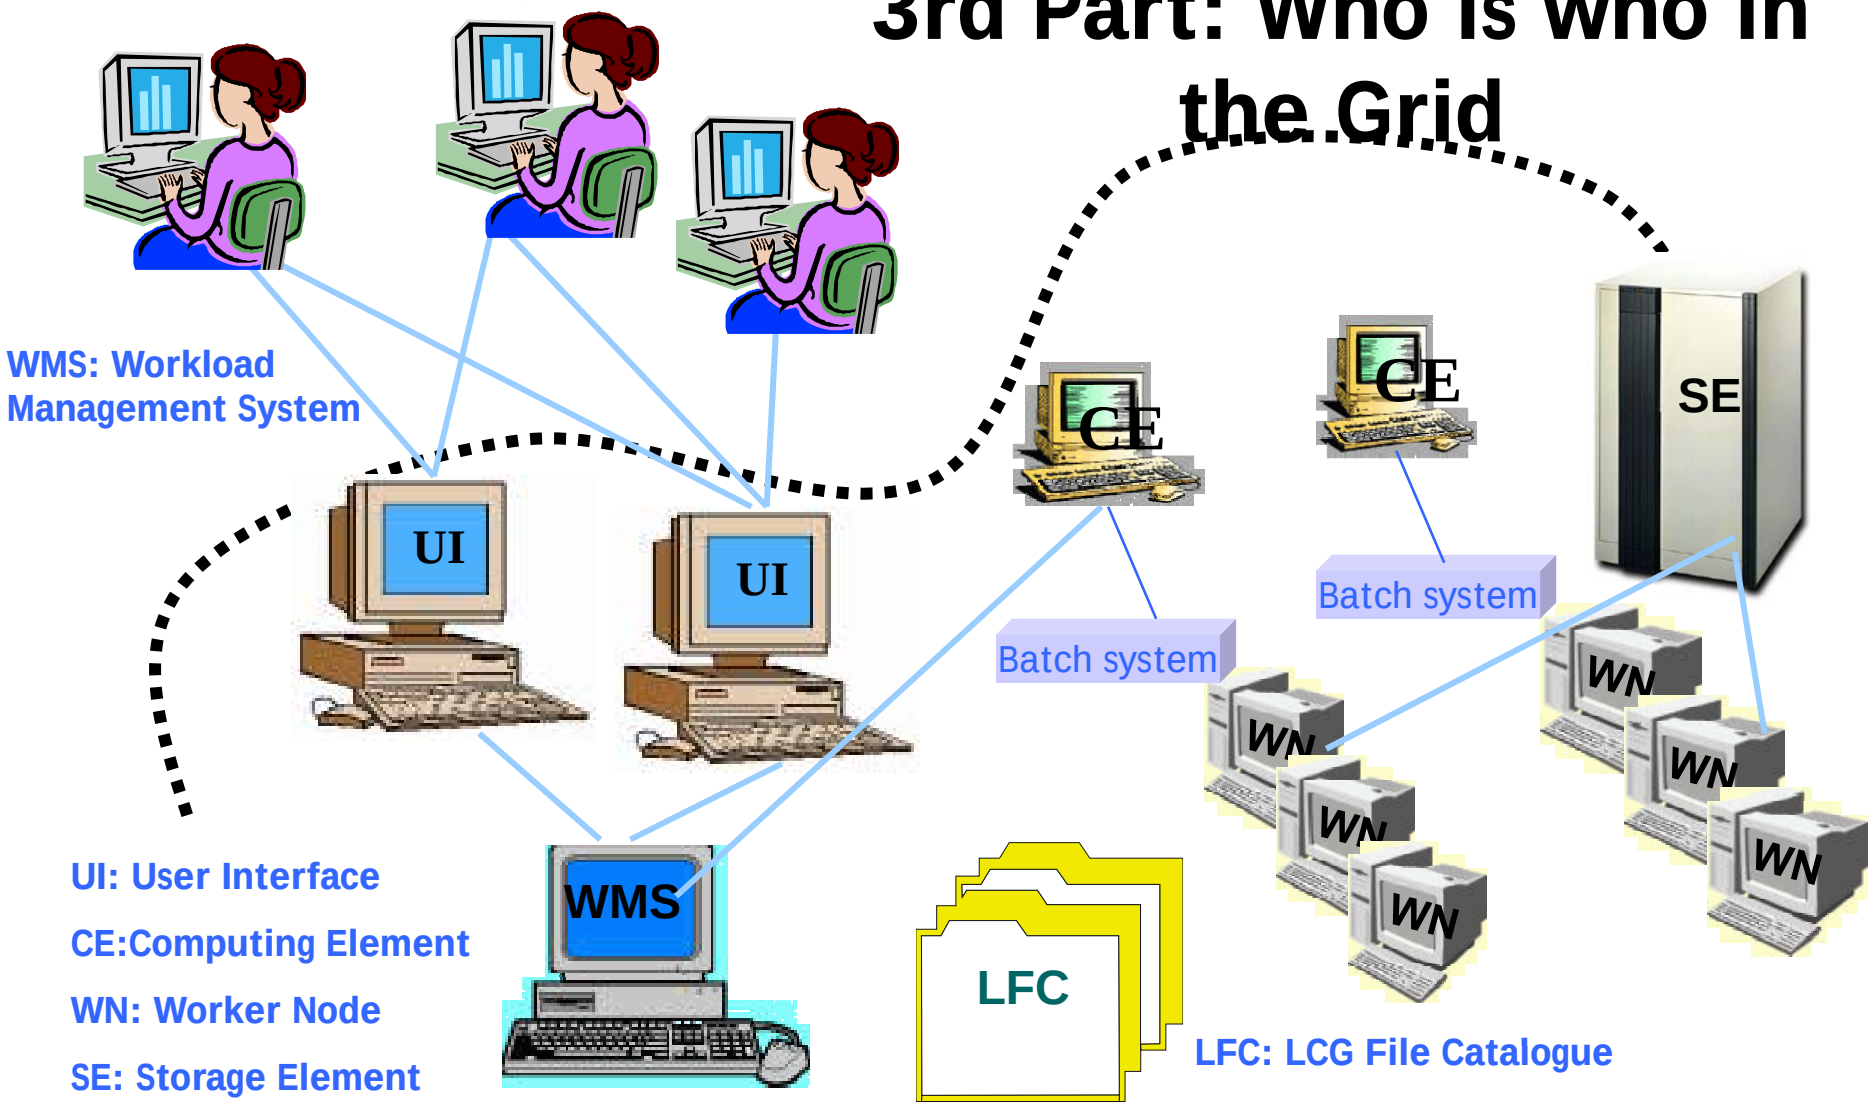

# Middleware: What it is

- All these services and their tools should talk the same language to make the whole mechanism able to work … in Grid terms this is called: MIDDLEWARE
  - The middleware is software and services that sit between the user application and the underlying computing and storage resources, to provide an uniform access to those resources
- The Grid middleware services: should
  - Find convenient places for the application to be run
  - Optimize the use of the resources
  - Organize efficient access to data
  - Deal with security
  - Run the job and monitor its progress
  - Recover from problems
  - Transfer the result back to the scientist

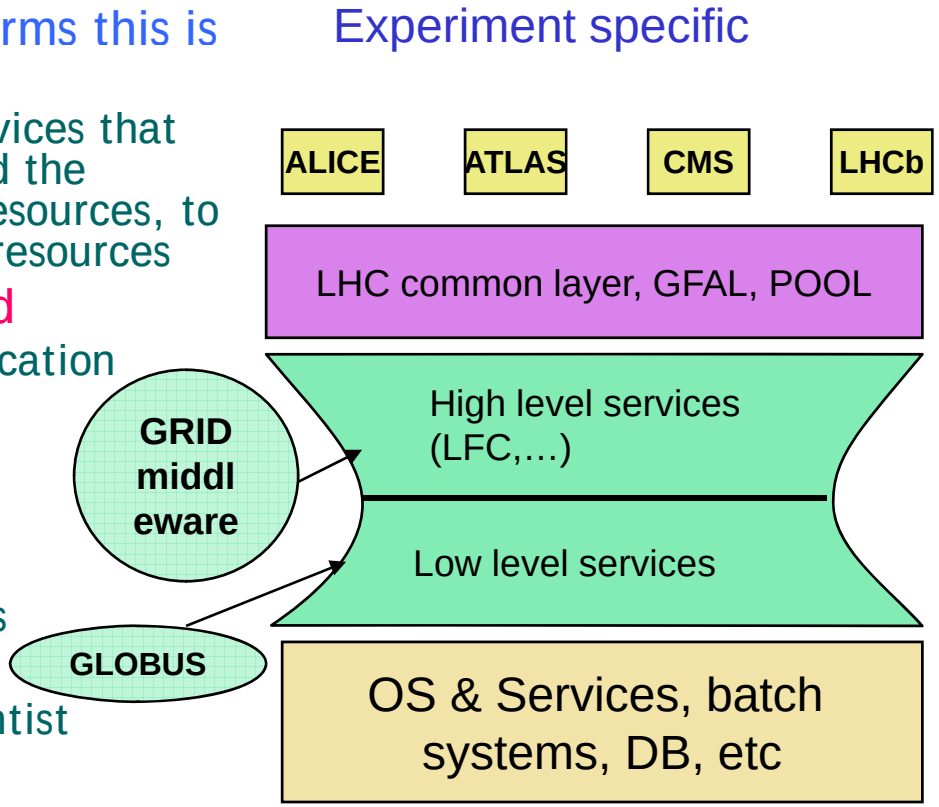

# Blocks of Middleware (I)

## Security

- GSI: Globus Security Infrastructure
  - Based on PKI (public/private key)
- Authentication: *Certificates*
- You are who you say to be
  - Provided by personal certificates delivered by Certification authorities
- Authorization: *Virtual Organization*
  - What can you do in the Grid?
  - VO: homogeneous groups of people with common goals
  - VOMS: Rules within the VO

## Info System

- Used by middleware (e.g. WMS) middleware developers and Grid users
- Published and provided in a hierarchical way:

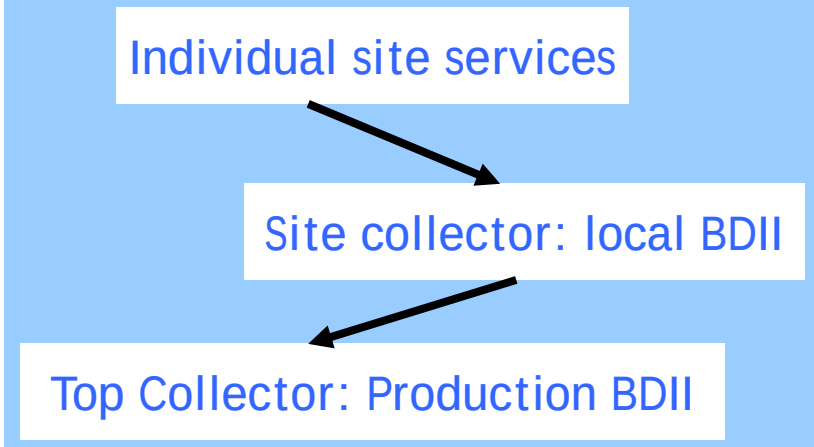

# Blocks of Middleware (II)

## WMS

- Accepts requirements from clients for jobs submission and management
- Support Requests:
  - Individual Job submission
  - DAG: direct acyclic graph of dependent jobs
  - Bulk submission: set of independent jobs
- Keeps track of the status of the job
- Matches jobs requirements with available resources
- Resubmission if needed

## Data Managt.

- Data access becoming the bottleneck
- How a good SE should look like
  - Good access to storage spaces
  - Services for the access
- Supported SE:
  - MSS (Castor, dCache) T1
  - DPM (disk arrays) T2
- SE managed by SRM
  - It hides Storage implementat.
  - File pinning
  - Space reservation

# Monitoring Systems

- 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
  - Dashboard
  - Certificate Lifetime Monitor

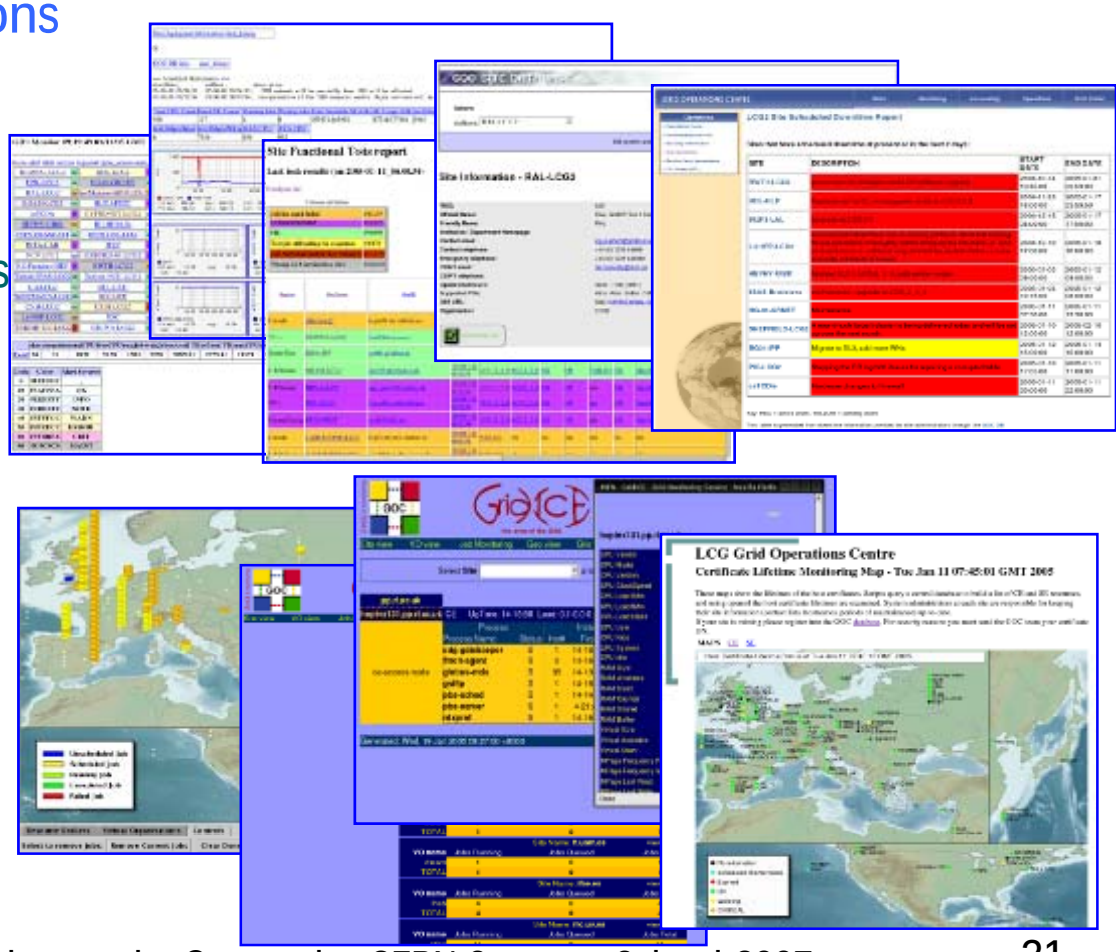

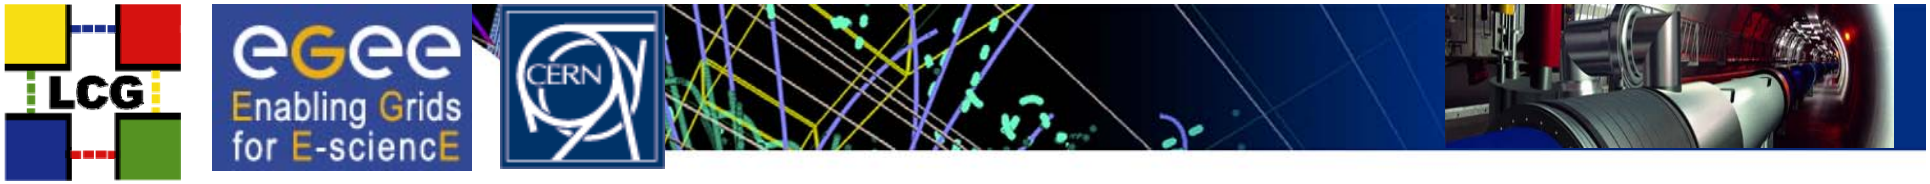


# 4th Part

- What are we doing now

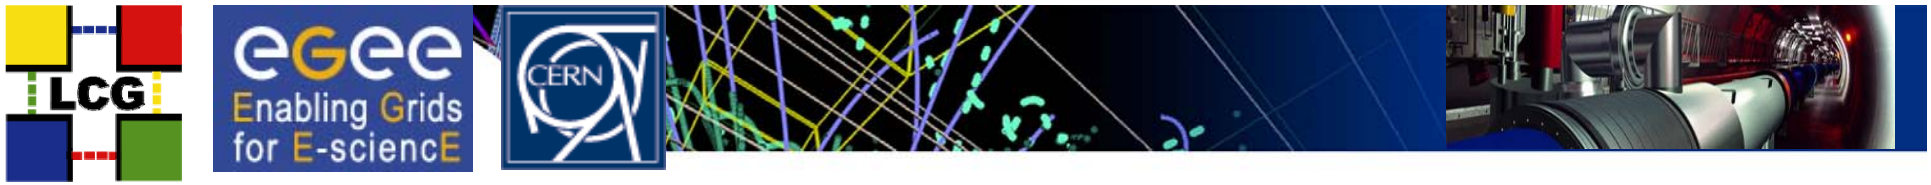


# Current Approach

- Grid Services
  - The LCG Service has been validated over 2 years (2004-05) via a series of dedicated "Service Challenges", designed to test the readiness of the service infrastructure
  - These are complementary to tests by the experiments of the offline Computing Models
    - Experiments participated during the last Services Challenge
  - Full support to experiments
    - Each experiment has a WLCG specific contact person + contact with the sites + support of the services
- Experiments
  - Testing of their own software using Grid services by Data Challenges
    - Continuous procedure by several experiments
  - Testing of the full system: Experiments Dress Rehearsals
    - From raw data to analysis: A "real" data taking

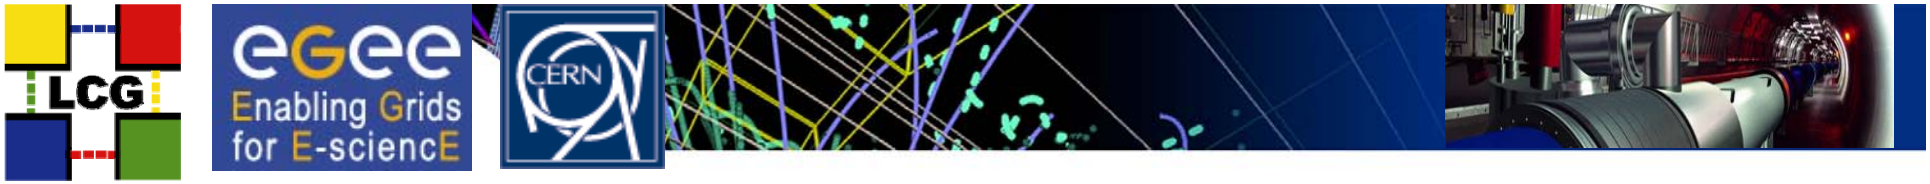


# 5th Part

- Applications beyond HEP

# Beyond HEP

- Due to the computational power of the EGEE new communities are requiring services for different research fields
  - In many cases, their productions are shorter and well defined in the year
  - The amount of CPU required is much lower and also the Storage capabilities

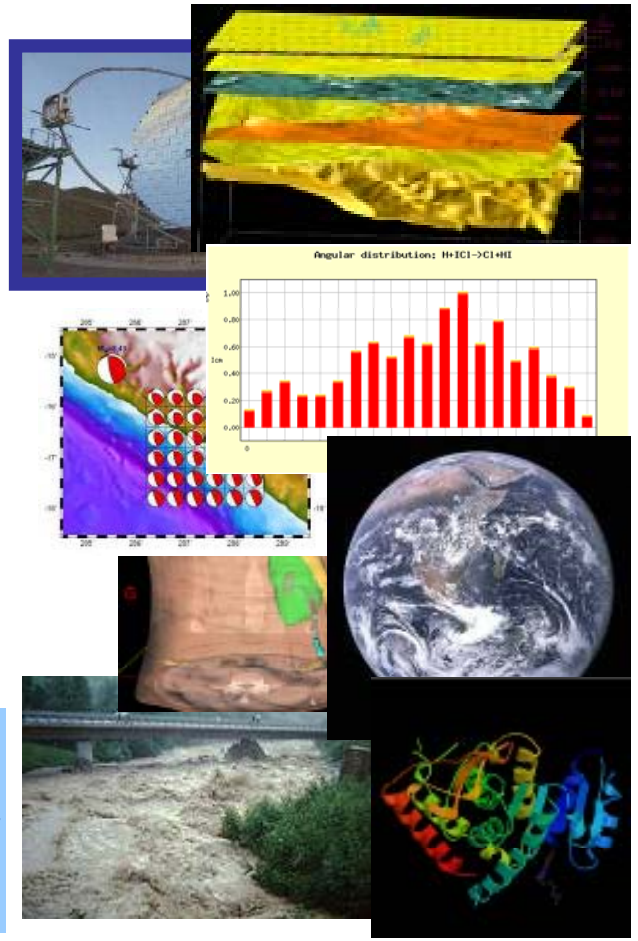


**20 applications from 7 domains**

High Energy Physic, Biomedicine, Earth Sciences, Computational Chemistry

Astronomy, Geo-physics and financial simulation

# UNOSAT

- Objectives
  - United Nations Initiative
  - Provide the humanitarian community with access to satellite imagery and Geographic Information System services
    - Reduce disasters and plan sustainable development
  - Ensure cost-effective and timely products
- Core Services
  - Humanitarian services
  - image processing

VEGETATION – 1 Km

IKONOS – 1m

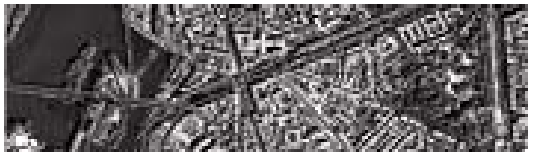

- The UNOSAT Gridification project
  - Bringing the Grid to mobile devices
    - Via GPS the user can know the coordinates of his physical position
    - Using this information as input the user obtains an image of his placement previously stored in the Grid
  - Hiding completely the Grid using web services

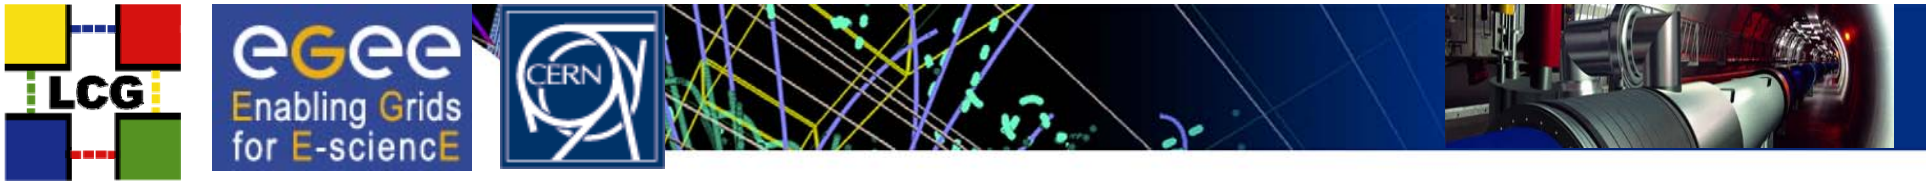


# Production with the ITU

- May 2006: The **International Telecommunication Union** organized a world conference to establish establish a new frequency plan for the introduction of digital broadcasting in Europe, Africa, Arab States and former Russian Federation States
- The software developed by the European Broadcasting Union (EBU), performs compatibility analysis between digital requirements and existing analogue broadcasting stations
- The peculiarity is that the CPU of each job is not easily predictable, since it depends on the details of the input data set although the total one could be calculated
- The duration of the jobs variates from few seconds until several hours
- The GRID was used during the whole production to executed all the required jobs (several thousands per weekend)

# Final Consideration

- Users have to be aware that GRID is not the magic bullet
  - Be sure what you want to do with the system
  - If you want to run few short jobs, maybe GRID is not your solution
- You will have to face situations as:
  - CPU at CERN, data in Taipei, RB in Italy….

```
[lxplus094] ~ > ping adc0018.cern.ch
PING adc0018.cern.ch (137.138.225.48) from 137.138.4.103 : 56(84) bytes of data.
--- adc0018.cern.ch ping statistics ---
11 packets transmitted, 11 received, 0% loss, time 10099ms
rtt min/avg/max/mdev = 0.204/0.405/1.332/0.332 ms
[lxplus094] ~ > ping lcg00105.grid.sinica.edu.tw
PING lcg00105.grid.sinica.edu.tw (140.109.98.135) from 137.138.4.103 : 56(84) bytes of data.
--- lcg00105.grid.sinica.edu.tw ping statistics ---
10 packets transmitted, 10 received, 0% loss, time 9086ms
rtt min/avg/max/mdev = 301.246/301.342/301.837/0.714 ms
```

O,4ms

301ms

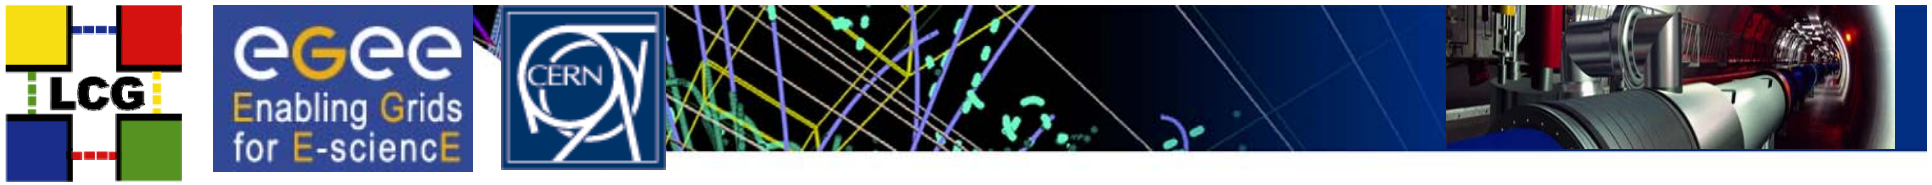


# Summary

- WLCG has been born to cover the high computer and storage demand of the experiments from 2008
- We count today more than 200 sites in 34 countries
- The whole project is the results of a big collaboration between experiments, developers and sites
- It is a reality, GRID is being used in production
- Covering a wide range of research fields, from HEP applications until humanitarian actions
- If we were able to share information with the WEB, GRID will allow us to share computational and storage power all over the world