Experimental techniques in high-energy nuclear and particle physics

"Dottorato di Ricerca in Ingegneria dell'Informazione"

LECTURE 9.

Prof. Rino Castaldi
INFN-Pisa
rino.castaldi@pi.infn.it

GridCafé

Data, data, data!!

- 40 million collisions per second
- After filtering, 100 collisions of interest per second
- > 1 Megabyte of data digitised per collision
  recording rate > 1 Gigabyte / sec
- 10^{10} collisions recorded each year
  stored data > 15 Petabytes / year

  This is a MASSIVE data-handling challenge!

1 Megabyte (1MB)
A digital photo

1 Gigabyte (1GB)
= 1000MB
5GB = A DVD movie

1 Terabyte (1TB)
= 1000GB
World annual book production

1 Petabyte (1PB)
= 1000TB
Annual production of one LHC experiment

1 Exabyte (1EB)
= 1000 PB
3EB = World annual information production

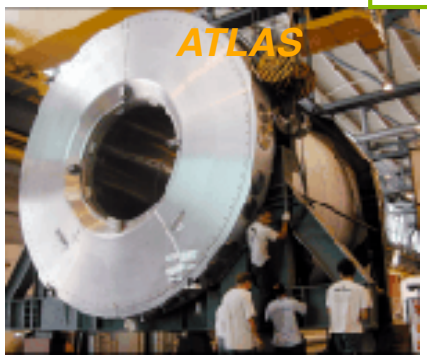

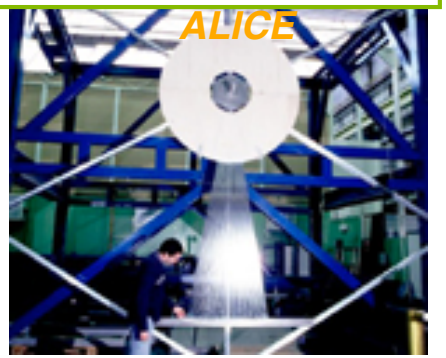

LHC Data every year

The LHC will produce around

15 Petabytes

of data every year!

That's about 20 million CDs each year!

Where will the experiments store all of these data?

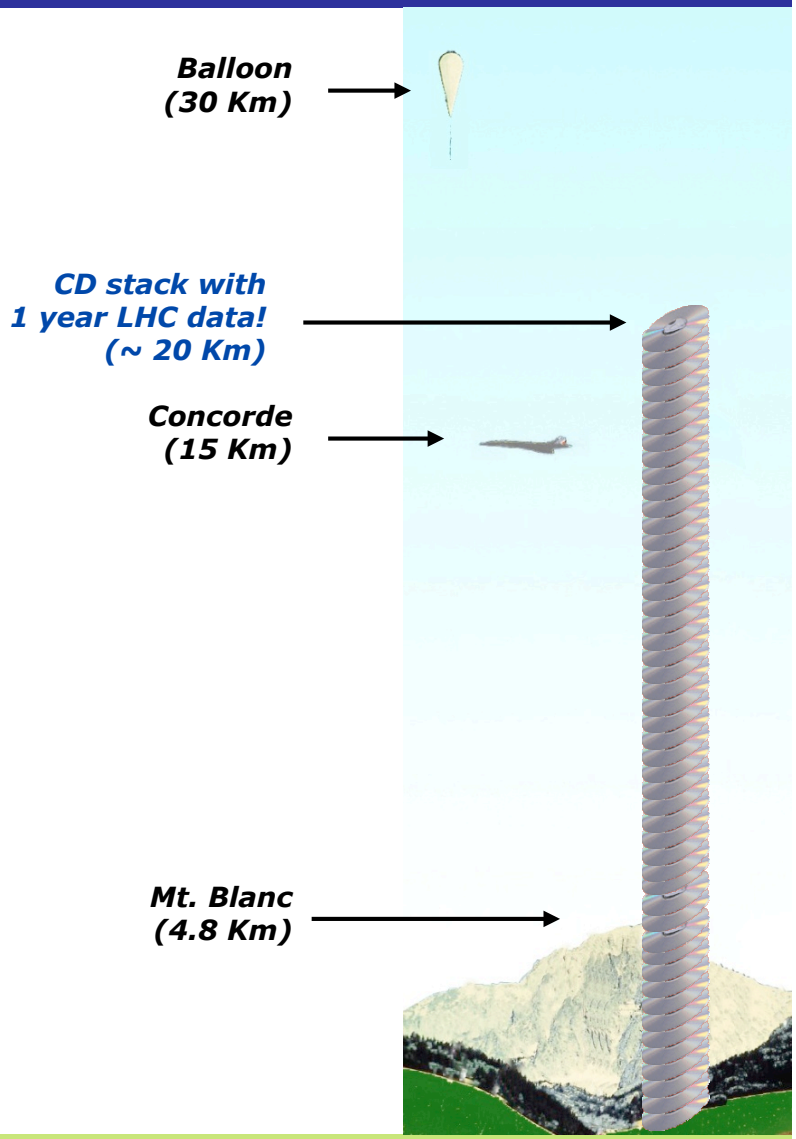

LHC processing

• Simulation: start from theory and detector characteristics and compute what detector should have seen

• Reconstruction: transform signals from the detector to physical properties (energies, charge of particles, ..)

• Analysis: Find collisions with similar features, use of complex algorithms to extract physics…

LHC Data Processing

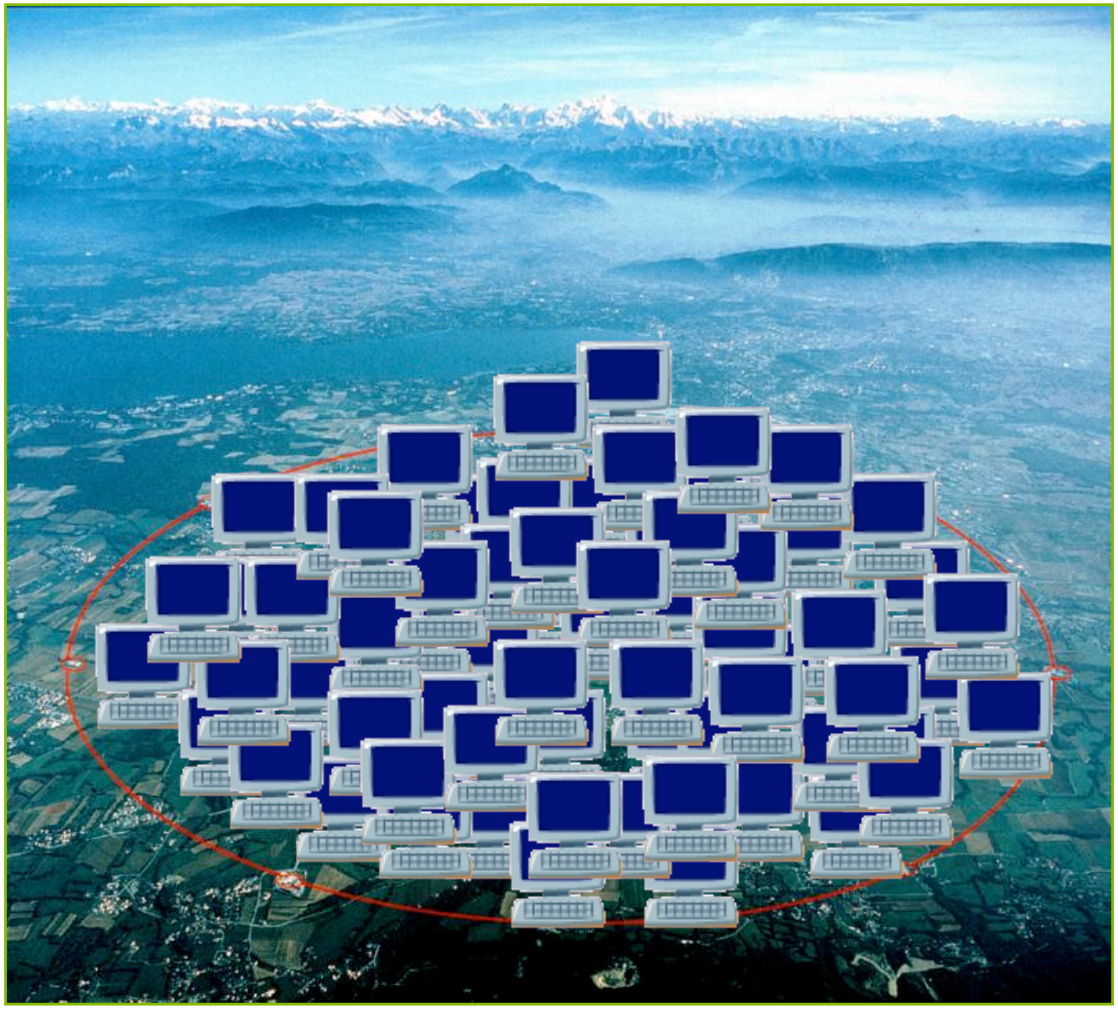

LHC data analysis requires a computing power equivalent to ~ 100,000 of today's fastest PC processors

Where will the experiments find such a computing power?

Computing power available at CERN

- High-throughput computing based on reliable "commodity" technology
- More than 35'000 CPUs in about 6000 boxes (Linux)
- 14 Petabytes on 14'000 drives NAS Disk storage)
- 34 Petabytes on 45'000 tape slots with 170 high speed drives

Nowhere near enough!

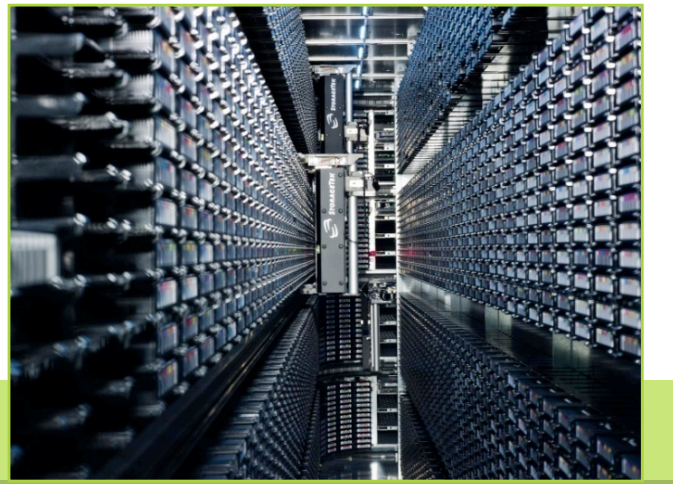

Solution: The Internet?

The network which physically connects the worlds computers allowing them to communicate as done in the World Wide Web

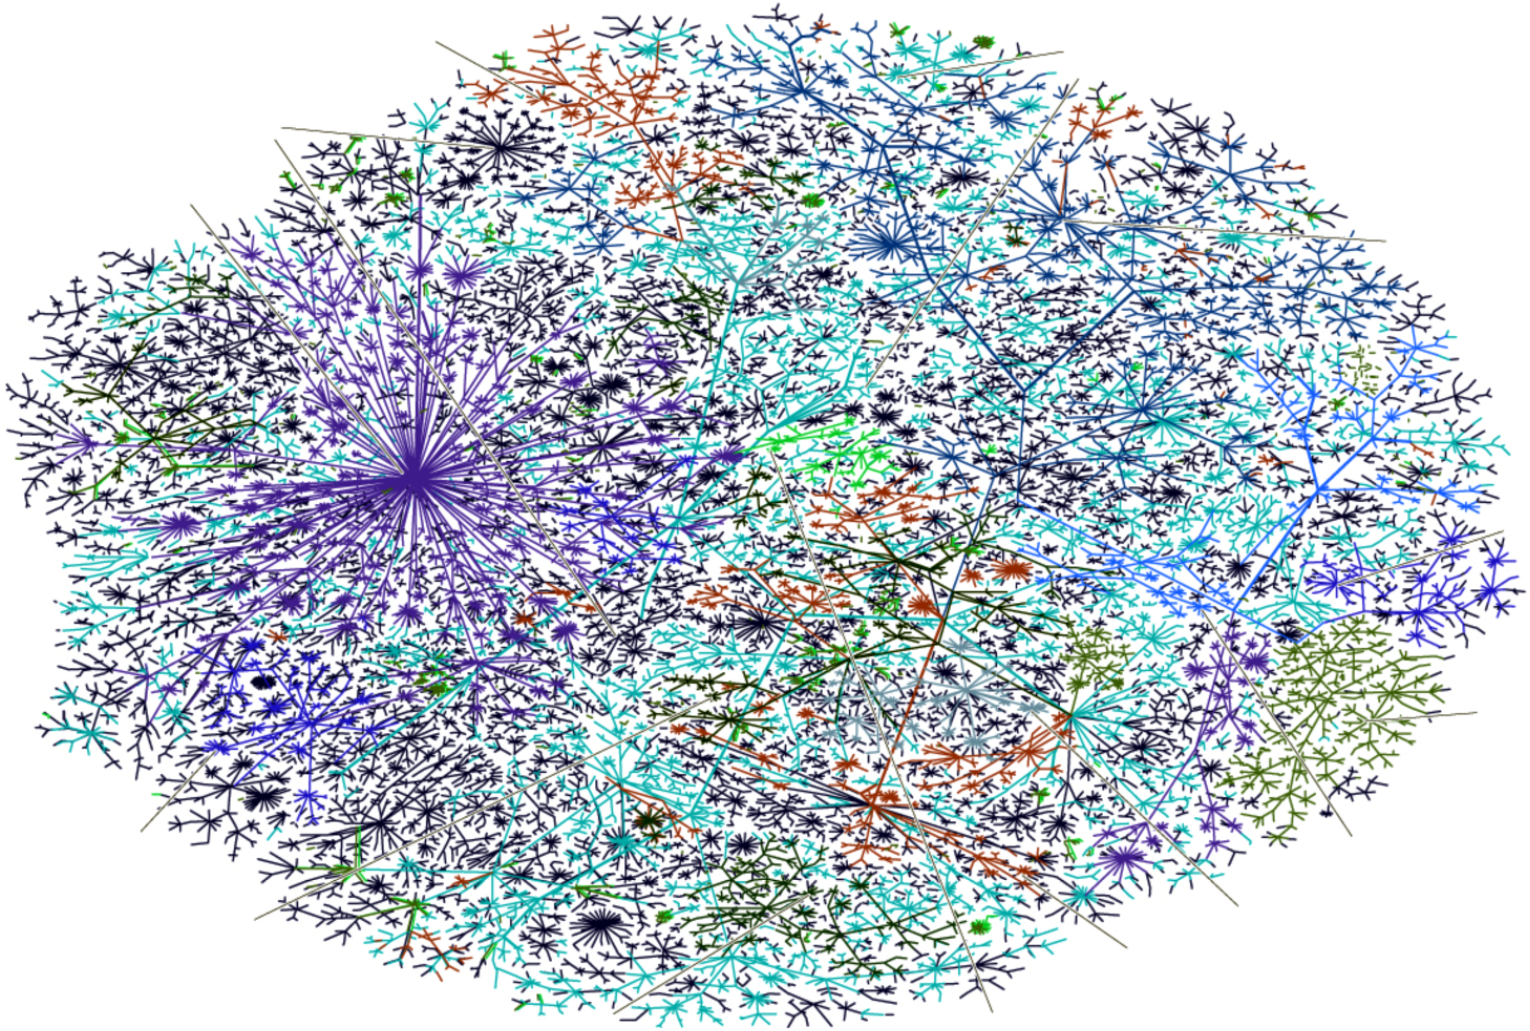

What is the Web?

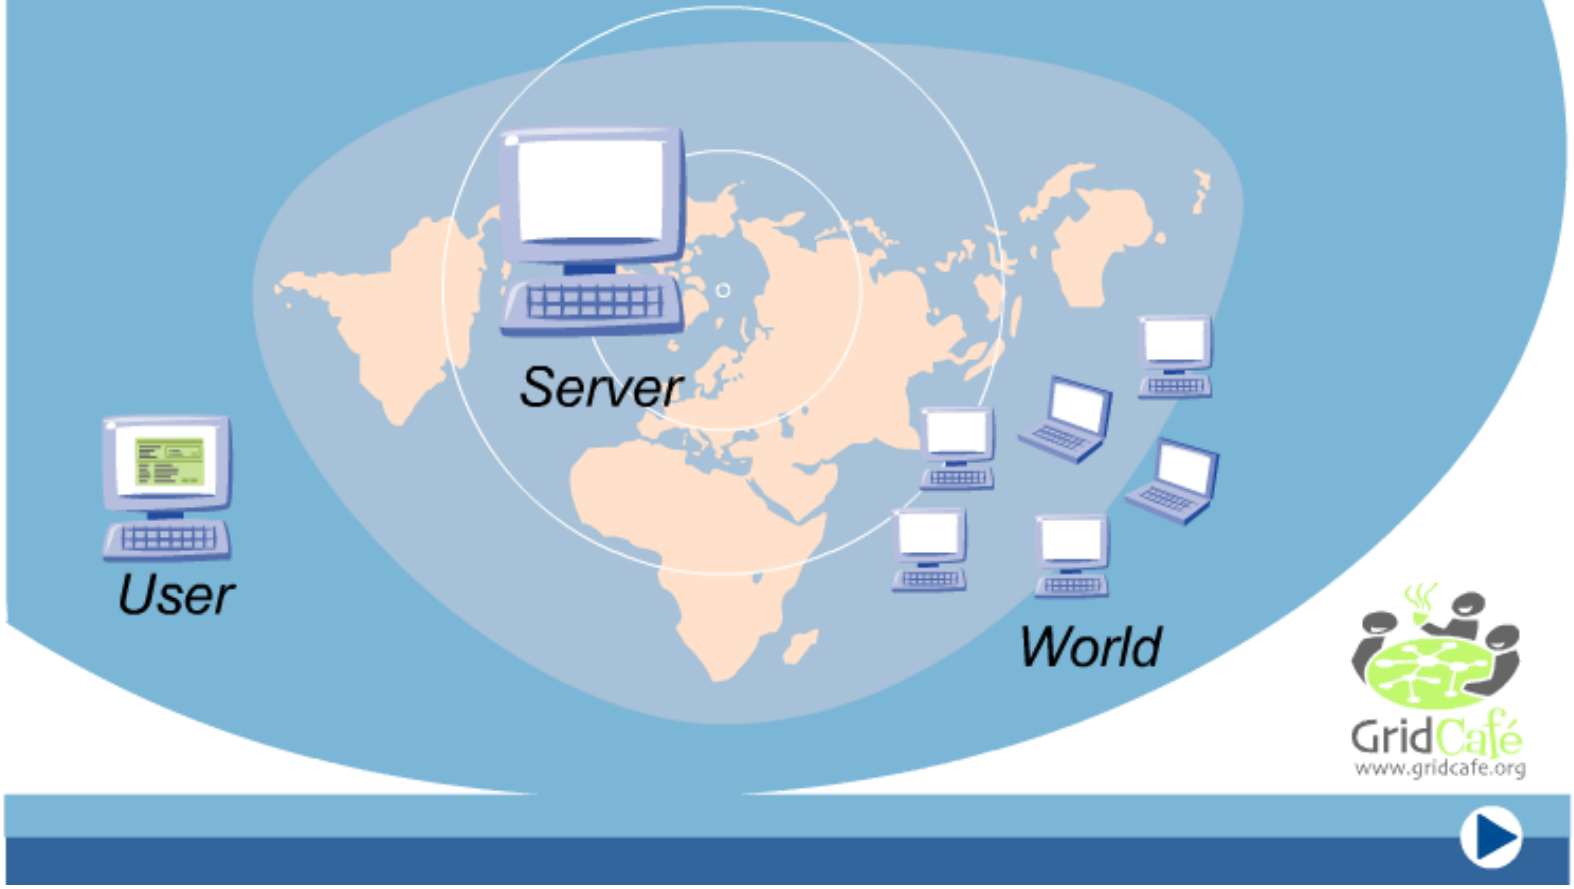

Why was the Web invented at CERN?

- Science depends on free **access to information** and exchange of ideas. CERN is the hub of a worldwide community of **6500 scientists** in **80 countries.**
- CERN has a long history of being at the forefront of **scientific computing** and **networking** (first lab on Internet outside the US).
- During the preparation of the previous large project LEP, the need to share documents in a global way became vital.

How did it start?

- 1989: **Tim Berners-Lee** circulates “Information Management: A proposal” to help with future Large Hadron Collider project.
- 1991: Early **www system released** to high energy physics via the CERN program library. First web servers located in European physics laboratories.
- 1993: First Mosaic browser; web reaches 500 servers and 1% of Internet traffic; CERN places **WWW in the public domain**.

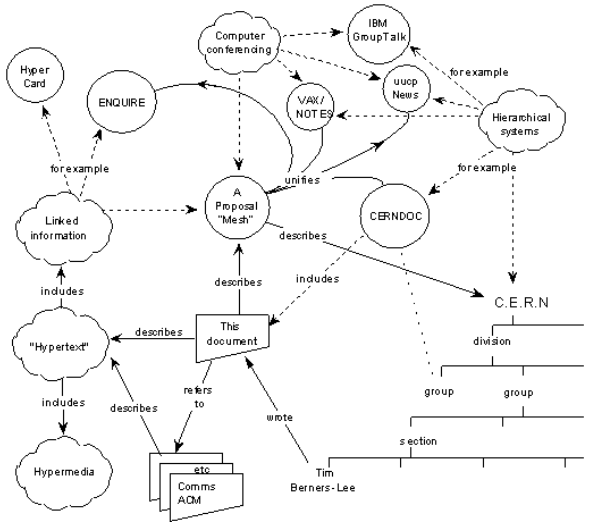

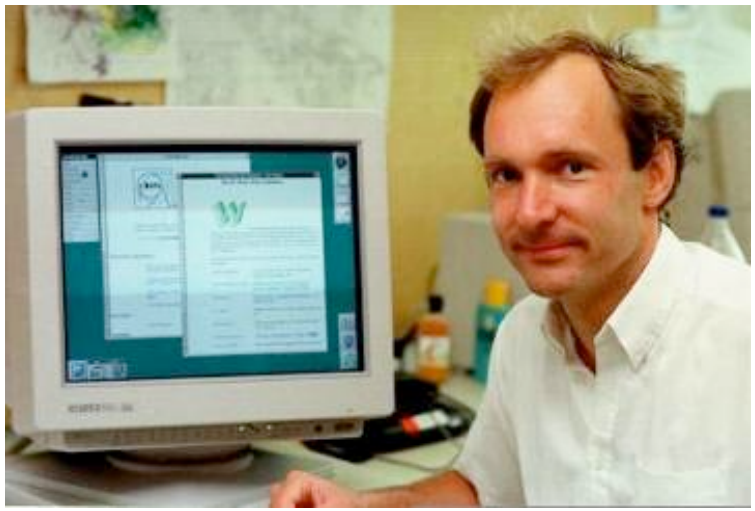

Web of Today!


Mittwoch, 10.9.2008
Schlagzeilen | TV-Programm | RSS | Newsletter | Mobil | Dienste | Startseite | Wetter | DER SPIEGEL
Le Matin Online
Mon Matin Online Inscription | Login
Rechercher
Shopping | Le Matin Mobile | Services abonnés | Edition électronique | Aide | Contact | Plan
Le Matin Video | BleuBlog | Dossiers | Photos | Dessins | Concours | Jeux | Annonces
Accueil | Actu | Sports | People | Tendances | Multimédia | Loisirs | Dépêches | Femina | Bilan
Imiter le chef
Toutes les vidéos et les recettes sur le blog
Golfe du Mexique: l'ouragan Ike menace les USA
L´ouragan Ike, balayant les eaux du Golfe du Mexique, poursuivait lentement sa route mercredi vers les Etats-Unis et les côtes du Texas. A Cuba, les intempéries ont entraîné l´évacuation de 2,6 millions de personnes et fait quatre morts. | Suite
Edipresse: bénéfice en hausse de 10% au 1er semestre
Edipresse a vu son bénéfice net progresser de 10% à 24 millions de francs au premier semestre. En revanche, les revenus de l´éditeur vaudois ont reculé dans les mêmes proportions, à 385,3 millions de francs. | Suite
Vidéo
© keystone
COLLISIONNEUR DU CERN
RSS
Mise en route réussie du LHC: ce n'est pas la fin du monde!
Le CERN a procédé avec succès ce matin au démarrage du plus grand accélérateur de particules au monde, le Grand collisionneur de hadrons (LHC). A 09h33, le faisceau de protons était injecté dans l´anneau de 27 kilomètres et à 10h25, le tour complet était réalisé pour la première fois. Suivez l'expérience en direct sur le matin.ch | Suite
Le Matin Online 69 commentaires
Dossier
› Le CERN et la fin du monde

What is the Grid?

• The **World Wide Web** provides seamless access to information that is stored in many millions of different geographical locations

• In contrast, the **Grid** is an emerging infrastructure that provides seamless access to computing power and data storage capacity distributed over the globe

Computing for LHC

- **Problem:** even with Computer Centre upgrade, CERN can provide only a fraction of the necessary resources

- **Solution:** Computing centers, which were isolated in the past, will be connected, **uniting the computing resources of particle physicists worldwide**

Users of CERN

Europe:
267 institutes
4603 users
Out of Europe:
208 institutes
1632 users

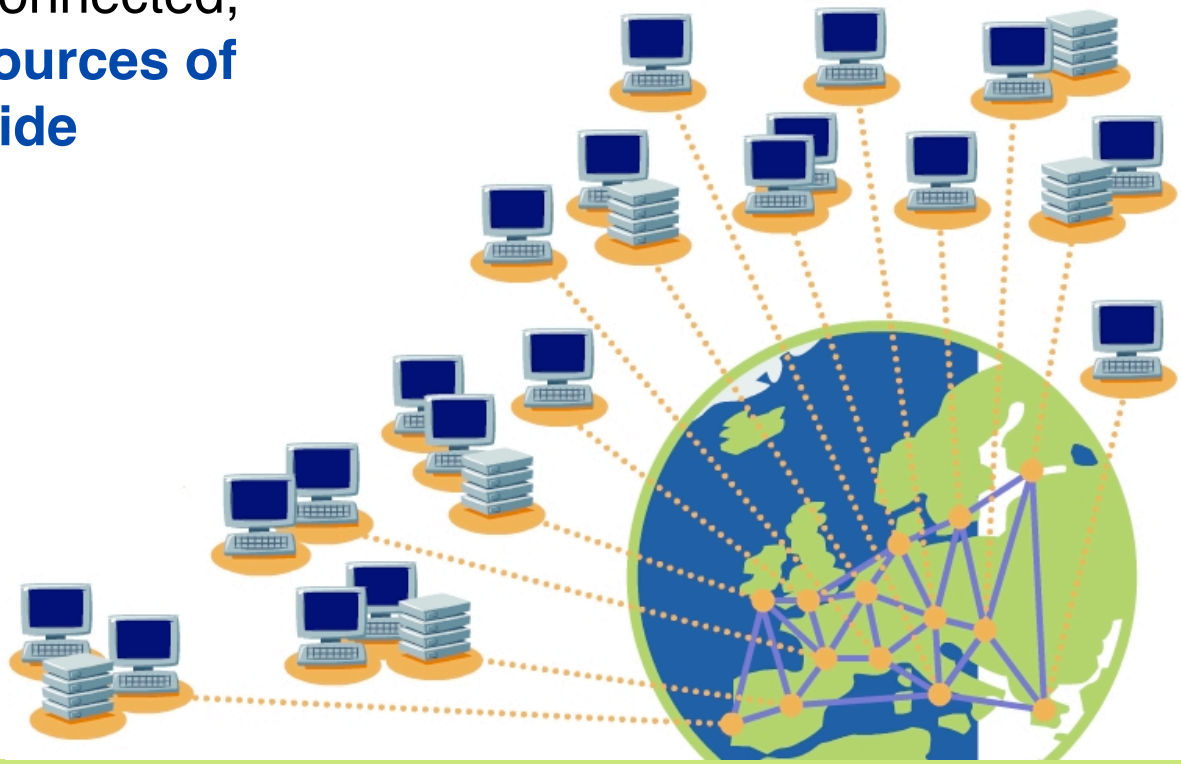

What is the Grid?

• The term Grid was coined by Ian Foster and Carl Kesselman (Grid bible “The Grid: blueprint for a new computing infrastructure”).

• The name Grid is chosen by analogy with the electric power grid: plug-in to computing power without worrying where it comes from, like a toaster.

• The idea has been around under other names for a while (distributed computing, metacomputing, …).

•This time, technology is in place to realise the dream on a global scale.

Computing Grid

Computing and Data Centres

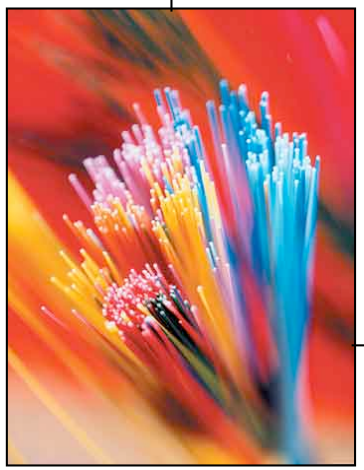

Fibre Optics of
the Internet

Analogy with the Electricity Power Grid

Power Stations

Distribution Infrastructure

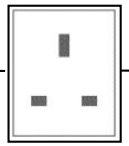

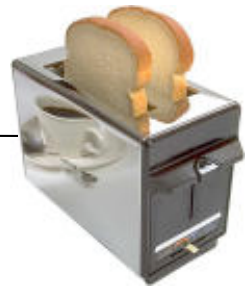

'Standard Interface'

4 Essential Utilities and Delivery Networks

(1) Electricity

Power Grid

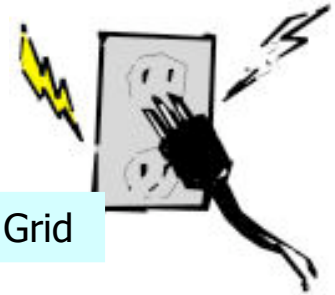

(2) Water

Water Distribution Network

(3) Gas

(4) Telephone

Telecom Networks

How will it work?

• The Grid relies on advanced software, called middleware, which ensures seamless communication between different computers and different parts of the world

• The Grid search engine will not only find the data the scientist needs, but also the data processing techniques and the computing power to carry them out

• It will distribute the computing task to wherever in the world there is spare capacity, and send the result to the scientist

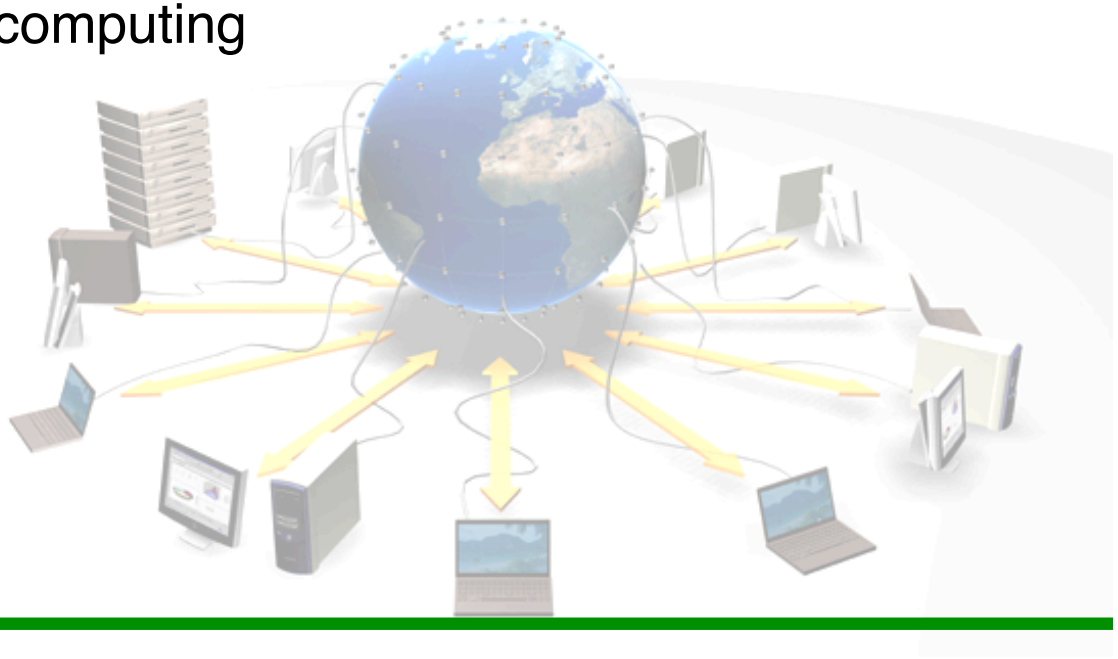

Building a massive distributed computer system - From Web to Grid

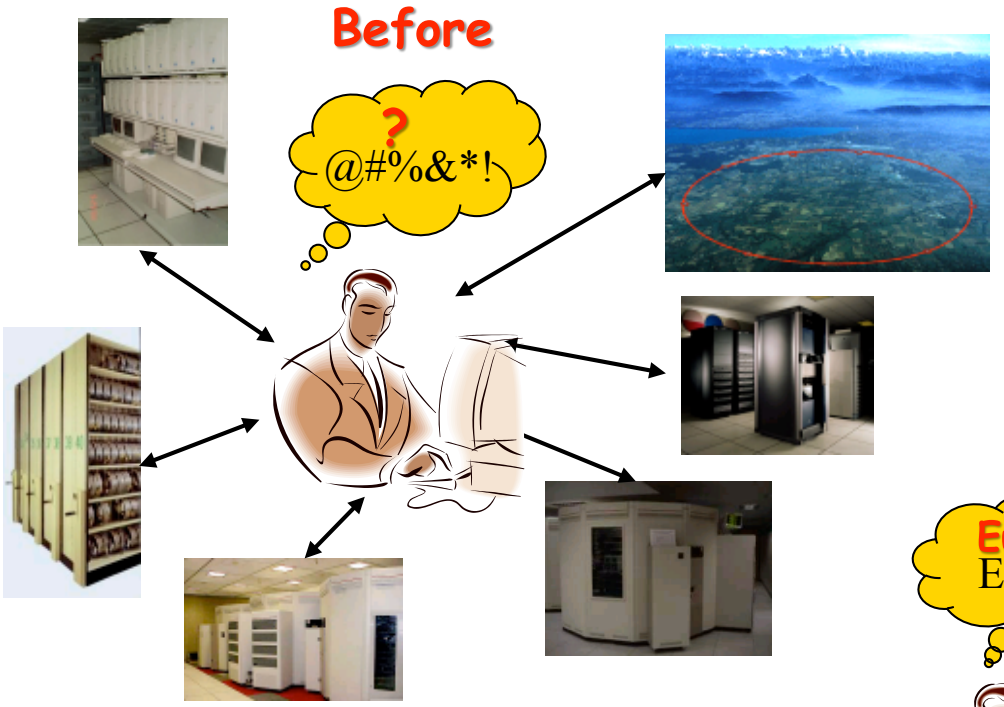

Distributed computing has been available to scientists for some time but:

- The use of different sites has to be negotiated by each scientist individually.
- They need a separate account on each system.
- Jobs have to be submitted and results collected back by hand.

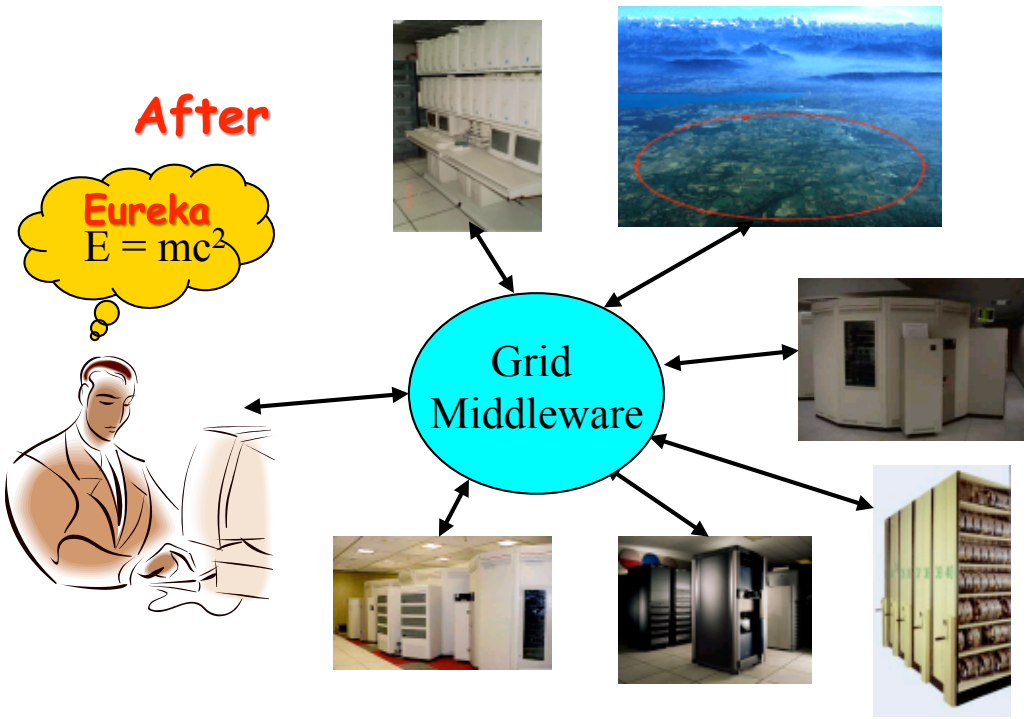

Middleware lets users simply submit jobs to the Grid without having to know where the data is or where the jobs will run.

The software can run the job where the data is, or move the data to where there is CPU power available.

Virtual Organization

High Energy Physics, Earth Observation and Biology are examples of communities made up of several institutions and individuals sharing the same interests and the same scientific goals. They greatly benefit from putting together their computing resources, data and scientific instruments.

Such distributed communities are called *Virtual Organisations*.

Step 1: Join Virtual Organization

To have access to the Grid facility, you need to join a *Virtual Organisation*.

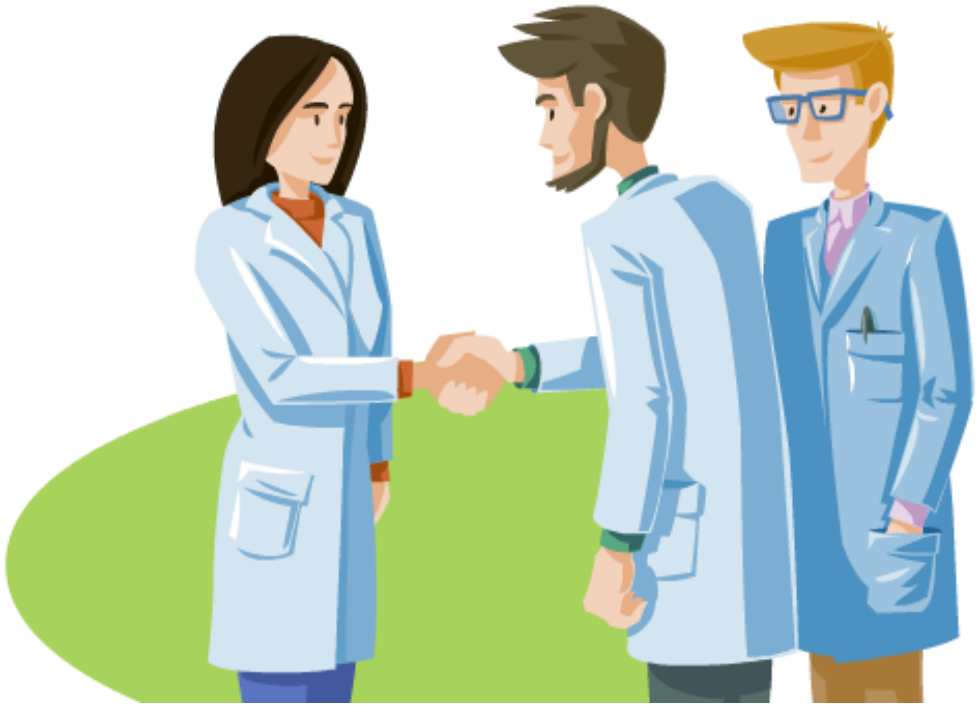

Step 3: Get a permission

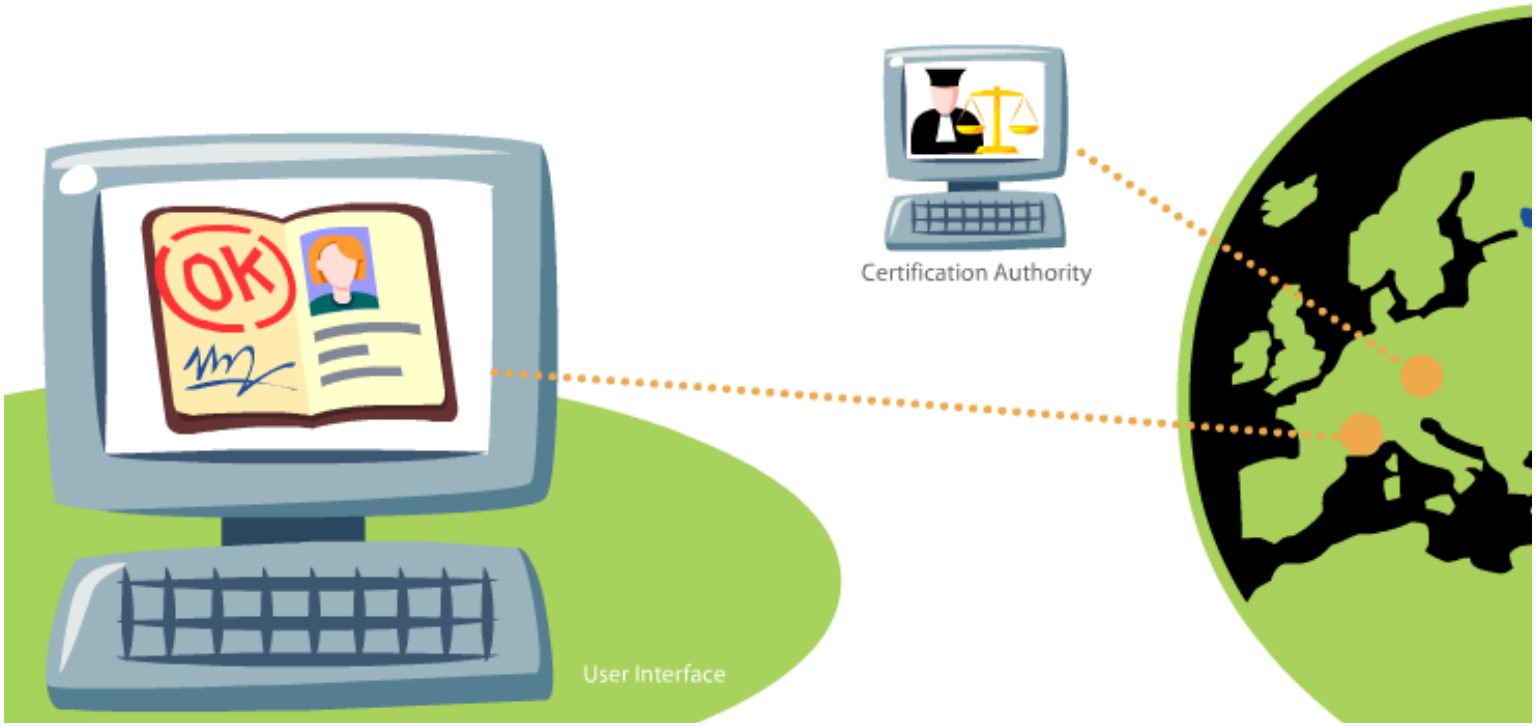

Step 4: Write a file describing your job

To run a job on the LCG/EGEE Grid Facility, you have to describe it in the *Job Description Language* (*JDL*).

JDL specifies job characteristics such as the application to use, the input data, the required resources, etc.

Step 5: Submit your job to the resource broker

Once you have the jdl file for your job, you can submit it to the ***Resource Broker***.

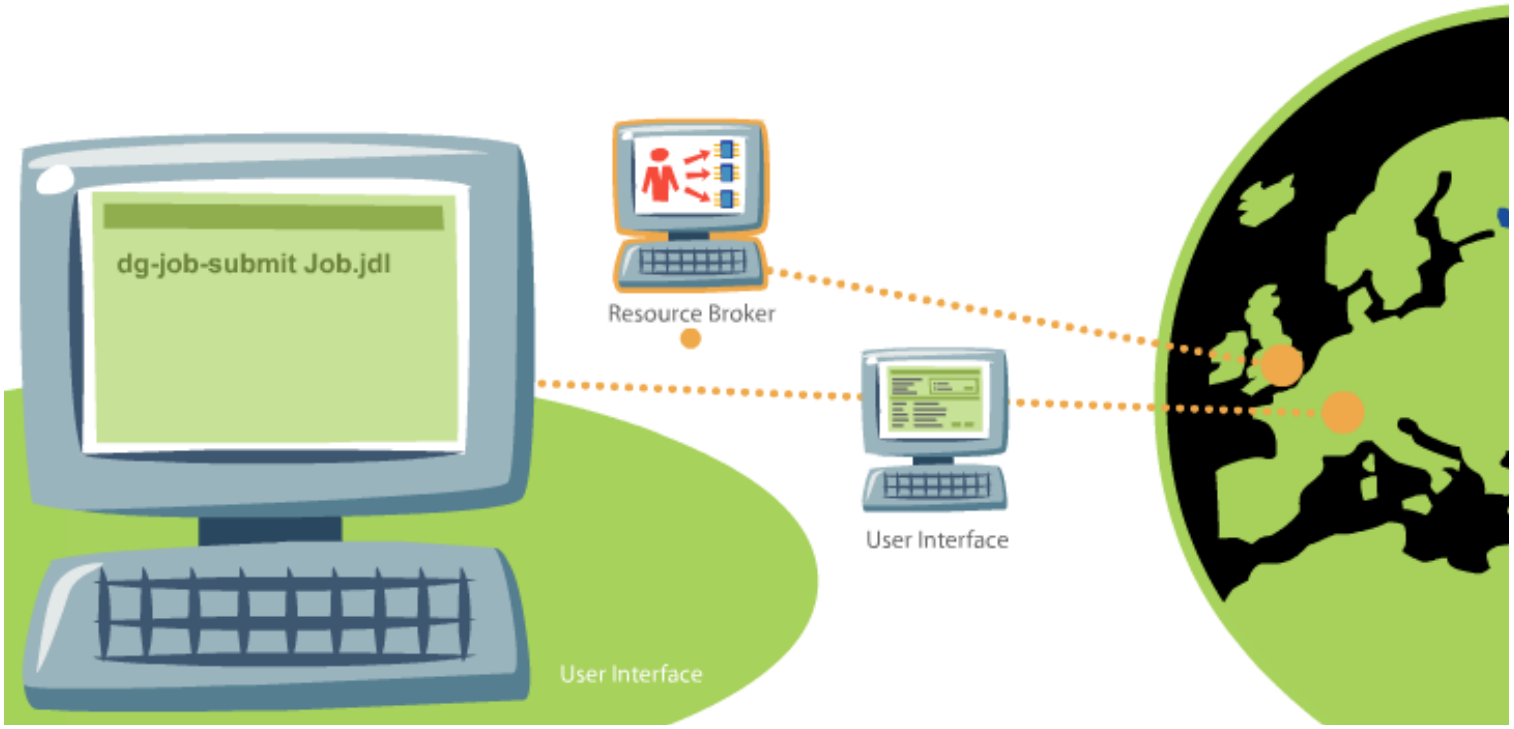

Step 5: Submit your job to the resource broker

Based on the information given in the JDL file, the *Resource Broker* queries the *Information Service* and the *Replica Catalog* to check resources.

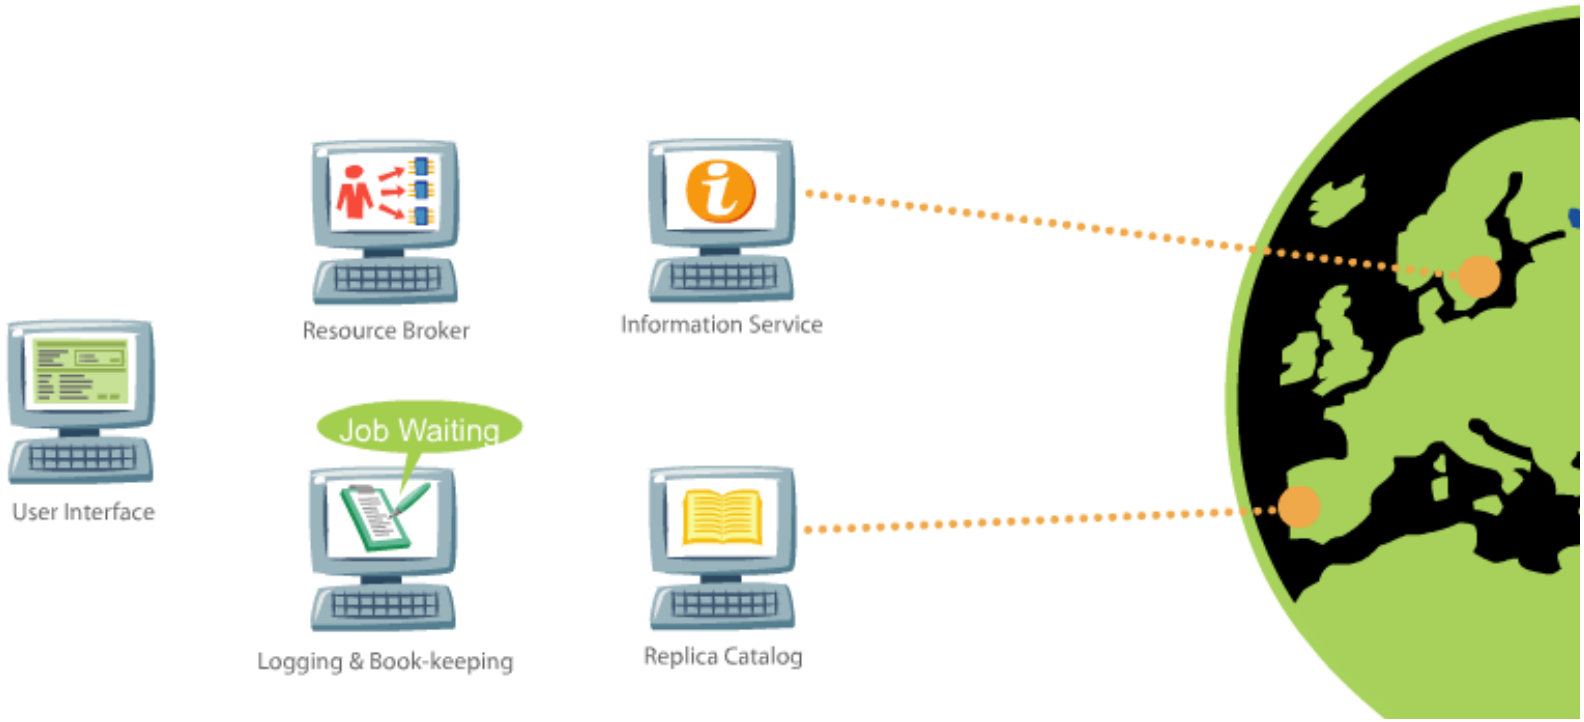

Step 5: Submit your job to the resource broker

The *Replica Catalog* and the *Information Service* hold information on the current status of all the sites. The Resource Broker uses this info to match the job to a suitable *Computing Element*. During this phase, the job is in a WAITING status.

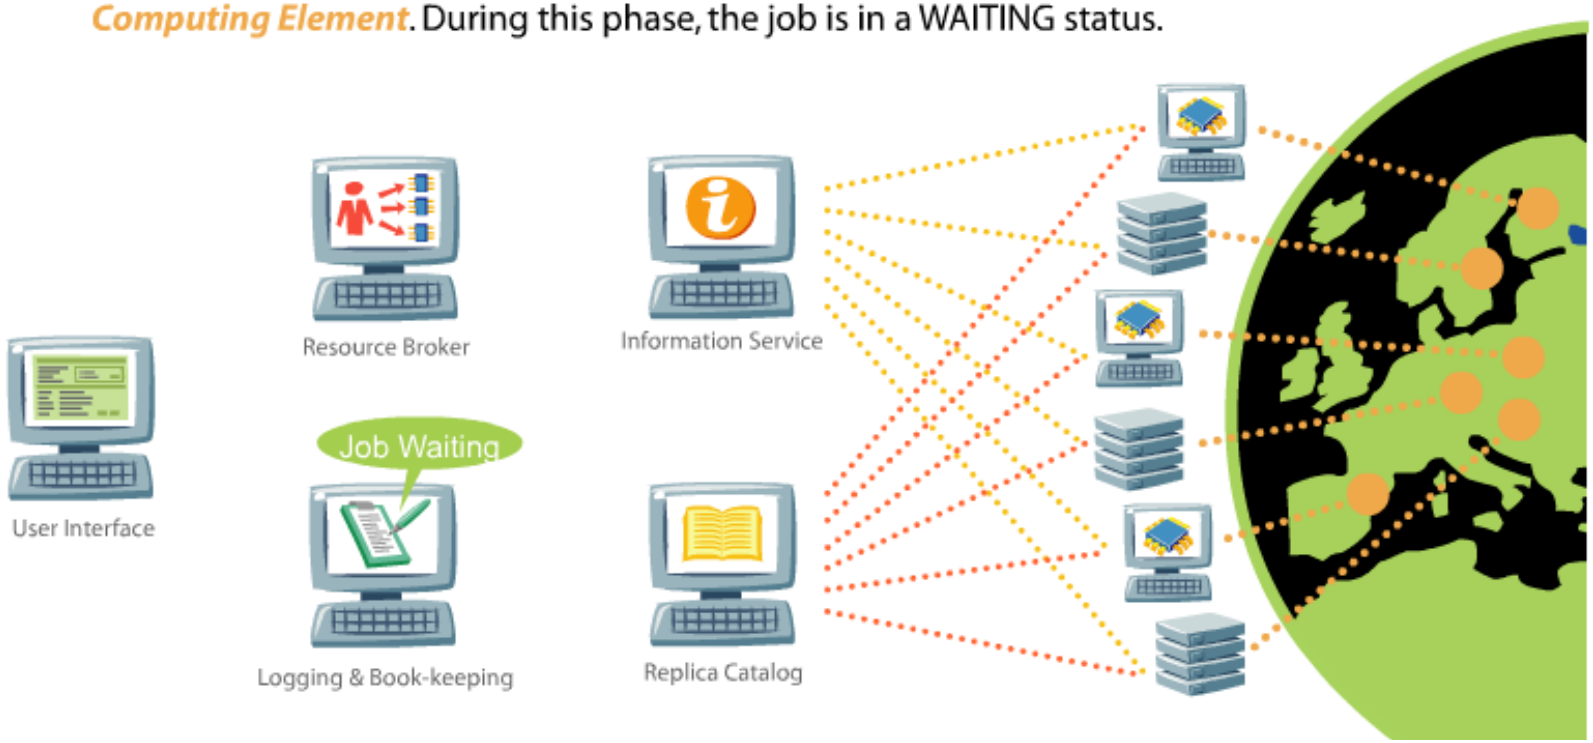

Step 5: Submit your job to the resource broker

The Resource Broker *makes its choice*. It has found a suitable Computing Element and the Storage Element with the necessary data.
It informs the logging & Book-keeping service of its decision.

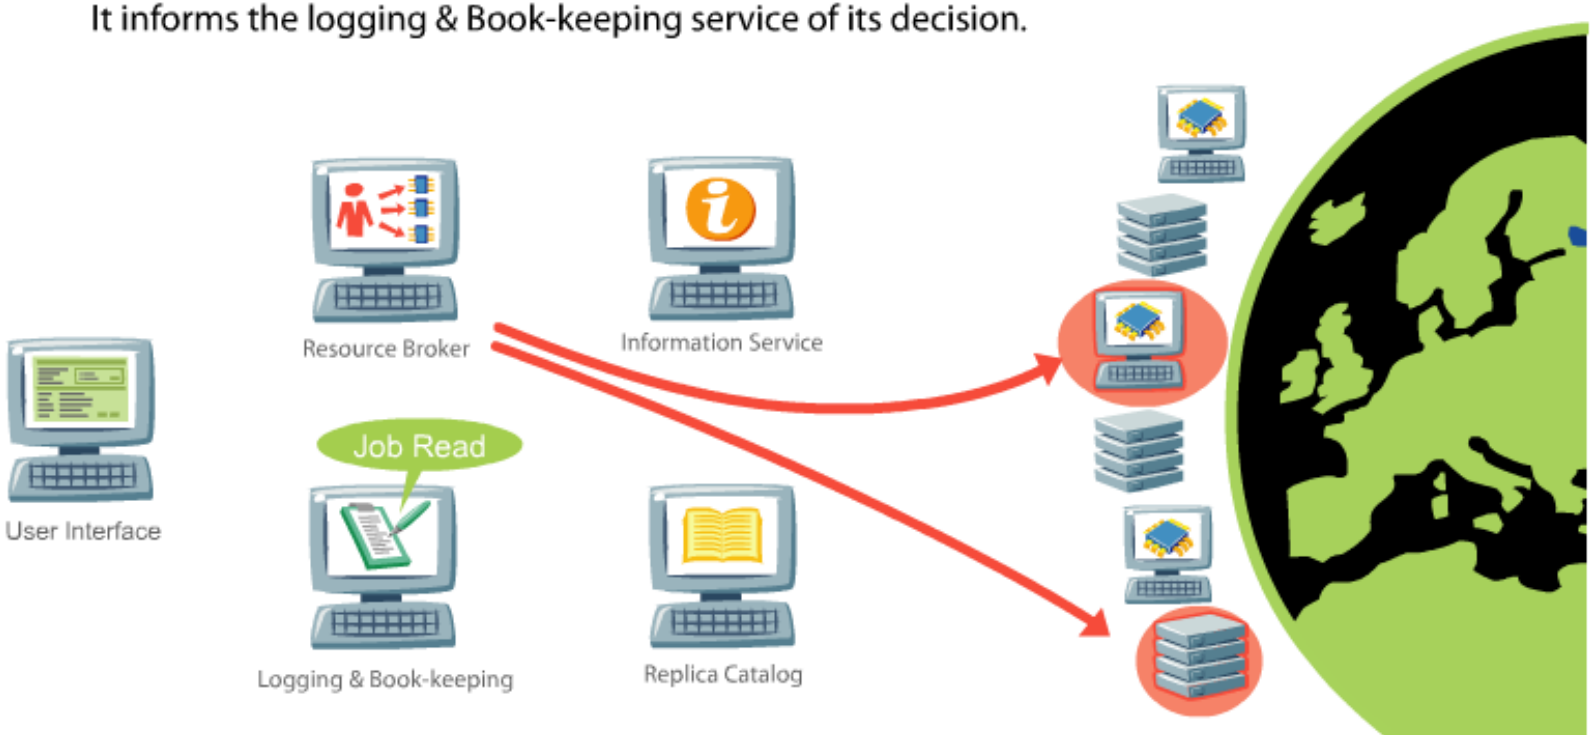

Step 6: Check the status of your job

During all this process, you can check the status of your job by contacting the *Logging & Book-keeping Service*.

Step 6: Check the status of your job

The execution of the job has completed on the Computing Element.
The *Computing Element transfers the output to the Resurce Broker*.

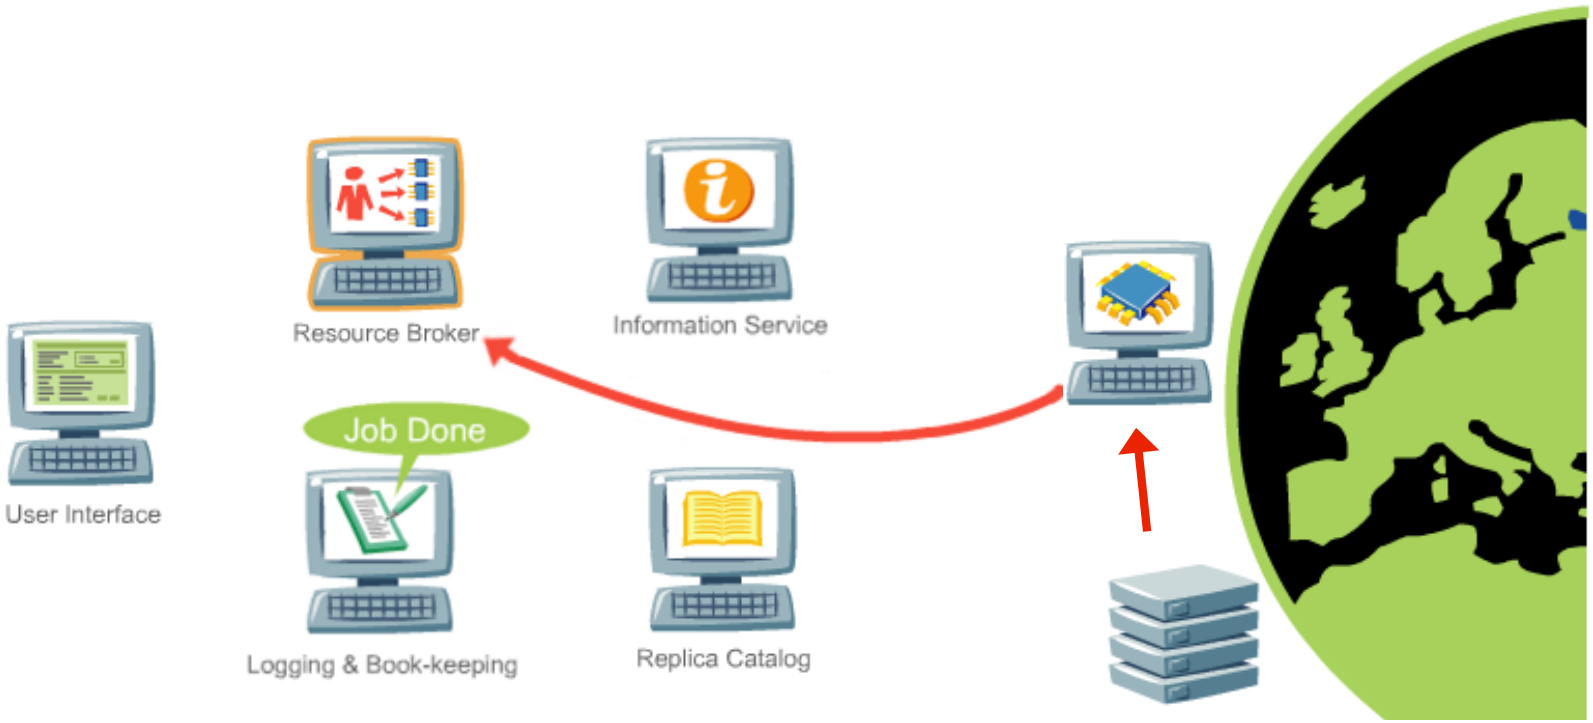

Step 7: Get the results

You can now retrieve your Output from the Resource Broker.
When finished, book-keeping *information is purged*.

How will it work?

The GRID middleware:

- **Finds convenient places for the scientists "job" (computing task) to be run**
- **Optimises use of the widely dispersed resources**
- **Organises efficient access to scientific data**
- **Deals with authentication to the different sites that the scientists will be using**
- **Interfaces to local site authorisation and resource allocation policies**
- **Runs the jobs**
- **Monitors progress**
- **Recovers from problems**

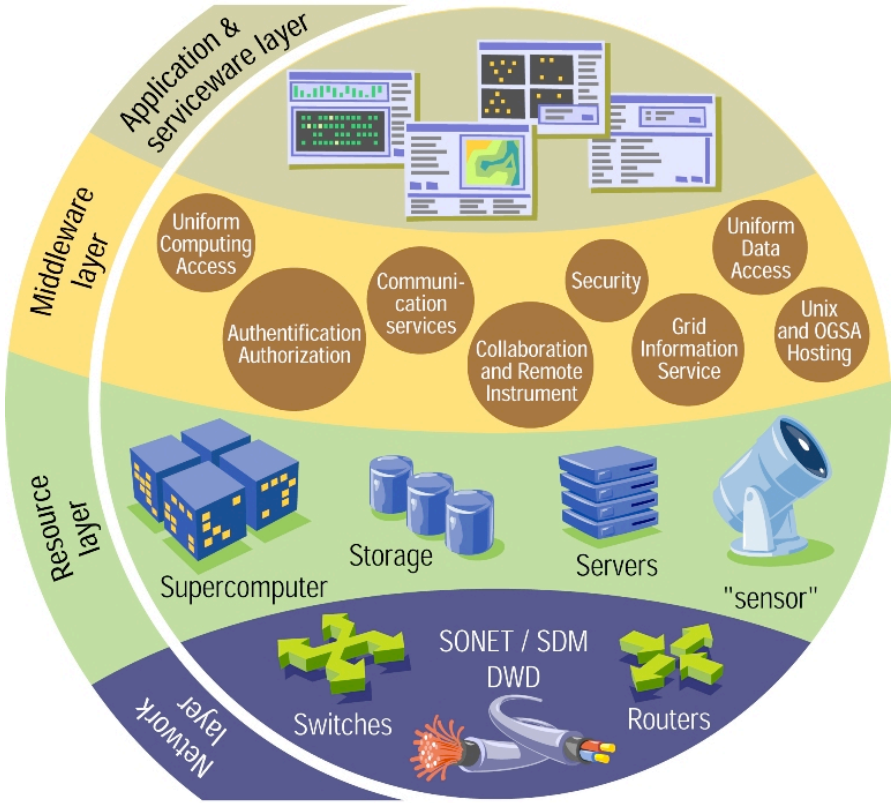

… and ….

Tells you when the work is complete and transfers the result back!

The Grid at CERN?

HEP Grid Computing Model (since the end of 90's)

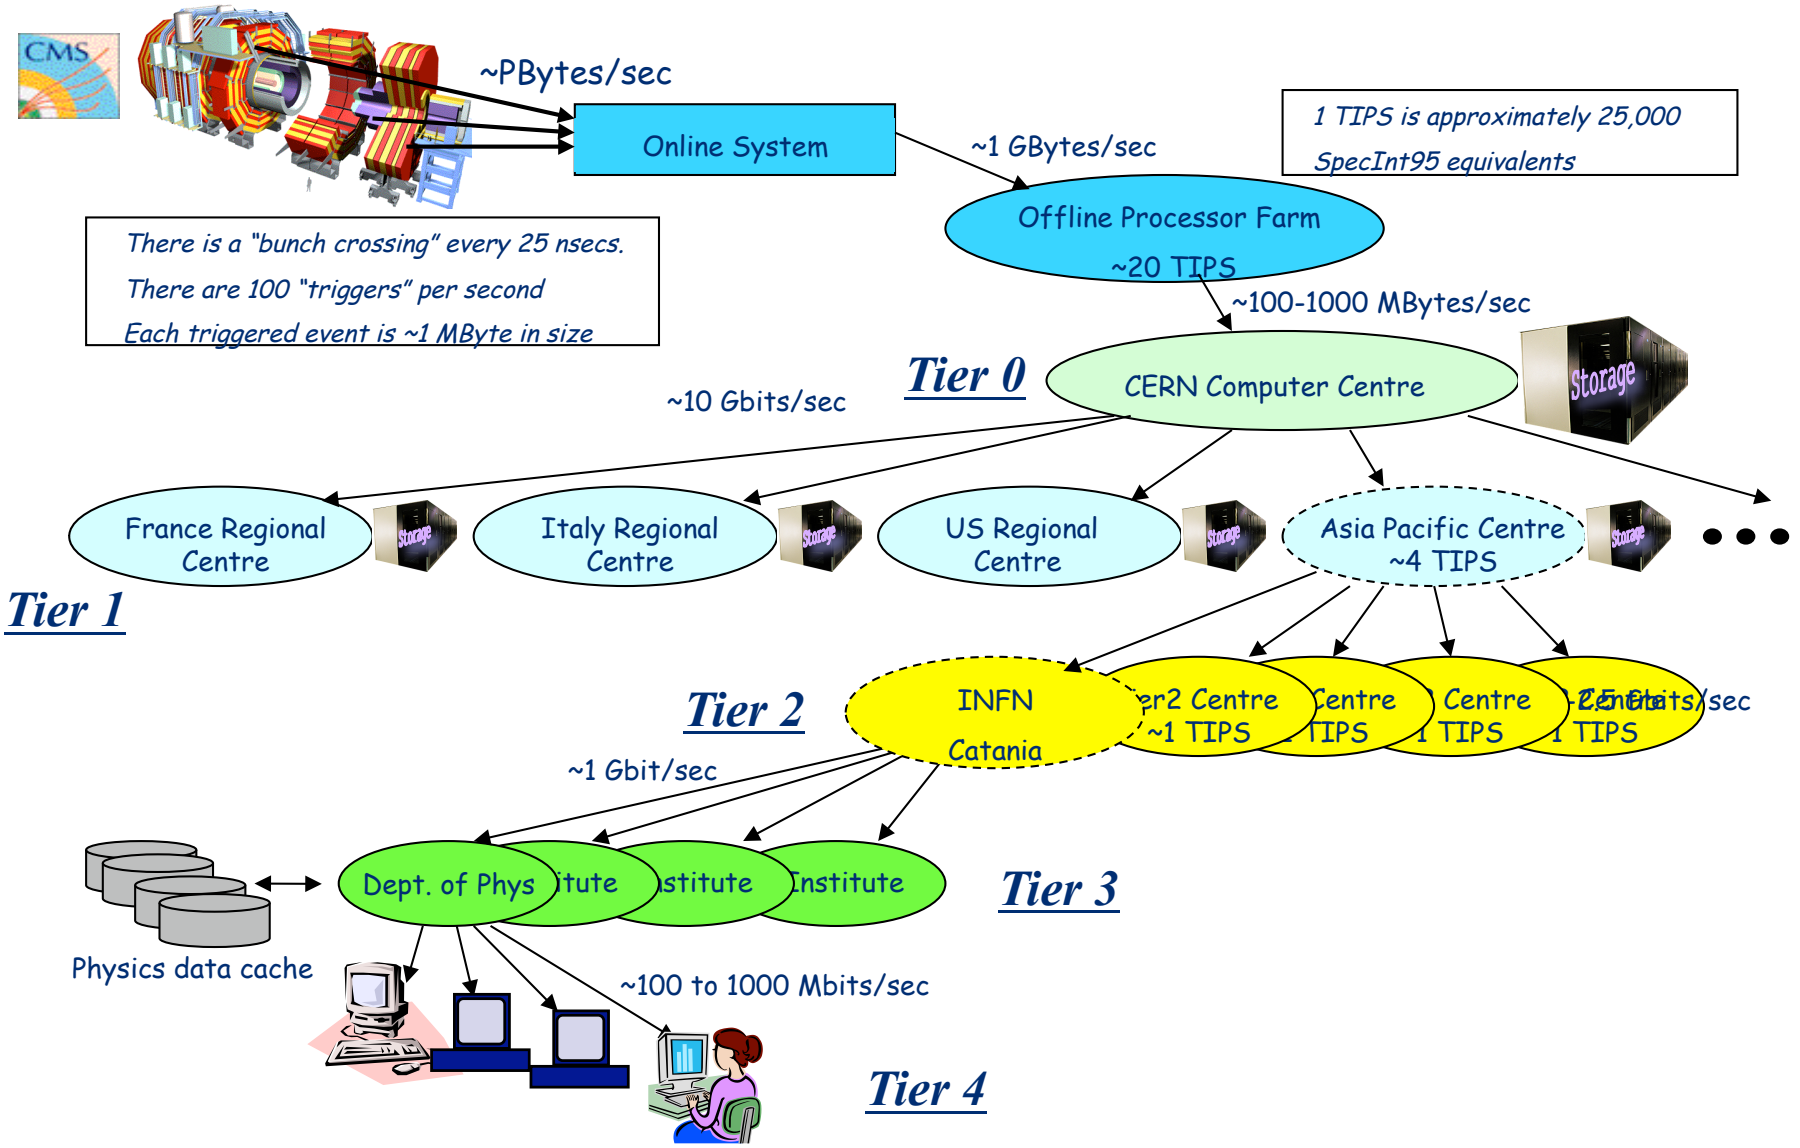

Grid @ CERN

- CERN projects:
  LHC Computing Grid (LCG)

- EC funded projects led by CERN:
  European DataGrid (EDG)
  + others

- Industry funded projects:
  CERN openlab for DataGrid applications

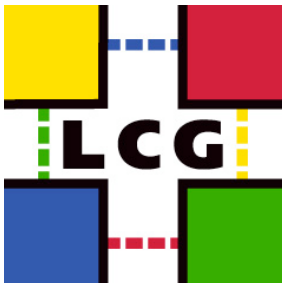

European DataGrid (EDG)

Mission:

• Develop the necessary middleware to run a Grid on a “testbed” involving computer centers in Europe

Key features:

• Largest software development project ever funded by the EU (9.8 million euros)

• Three year phased developments & demos (2001-2003)

• Three application fields: High Energy Physics, Earth Observation and Genomic Exploration

European DataGrid (EDG)

People:

• Total of 21 partners, over 150 programmers from research and academic institutes as well as industrial companies

Status:

• Testbed including approximately 1000 CPUs at 15 sites

• Several improved versions of middleware software (final release end 2003)

• Several components of software integrated in LCG

• Software used by partner projects: DataTAG, CROSSGRID

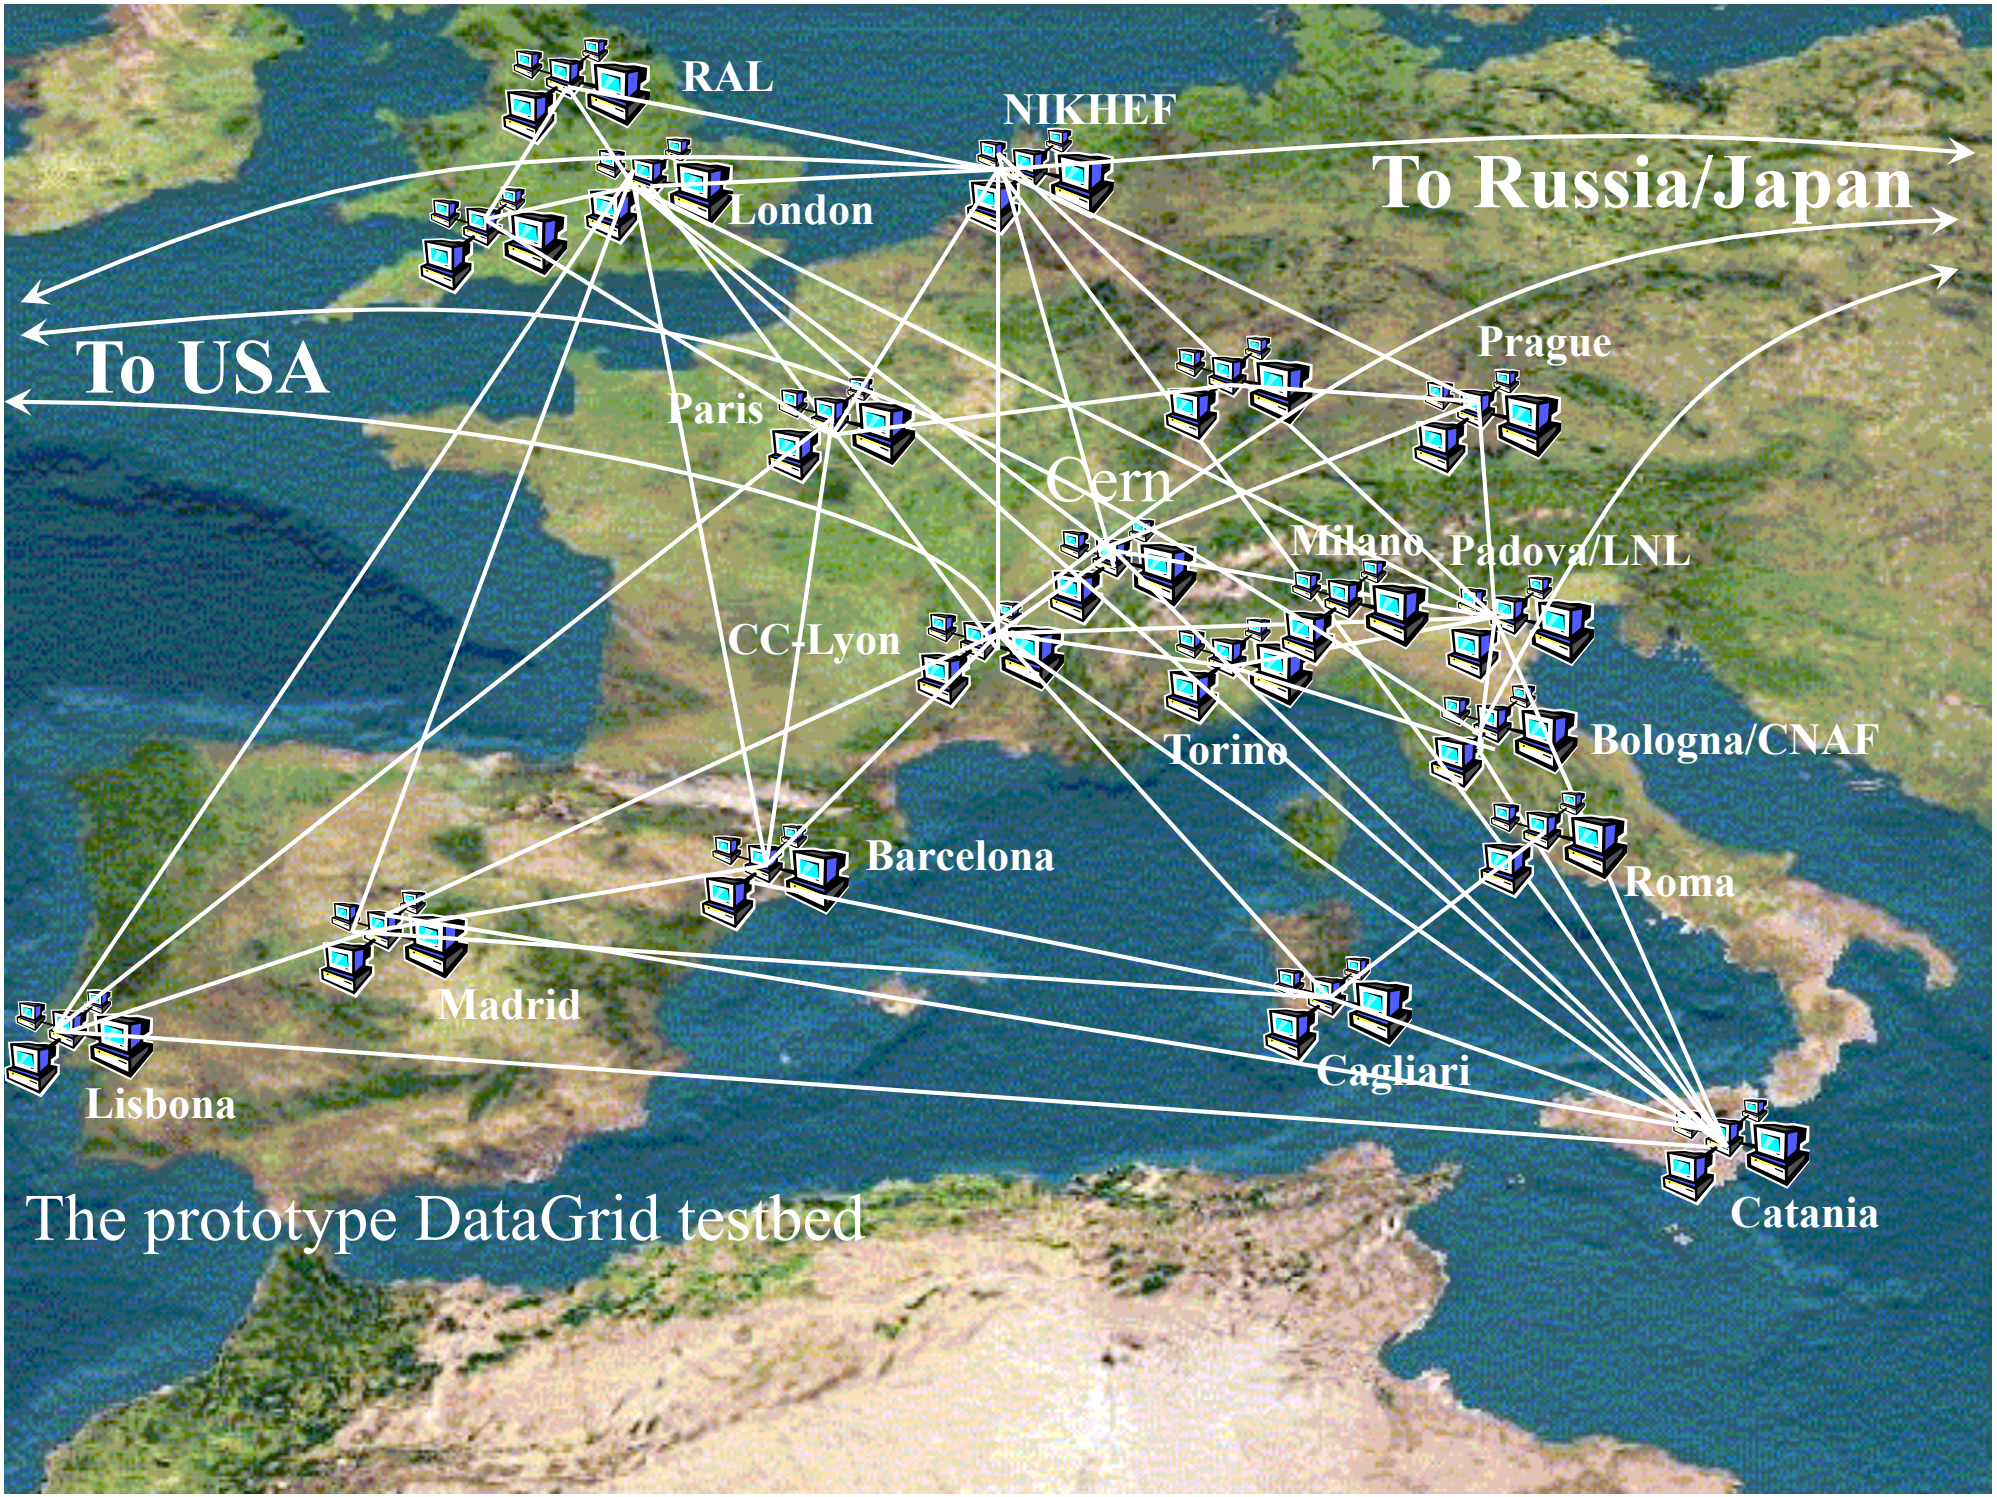
RAL
NIKHEF
To Russia/Japan
London
To USA
Prague
Paris
Cern
Milano
Padova/LNL
CC-Lyon
Torino
Bologna/CNAF
Barcelona
Roma
Madrid
Cagliari
Lisbona
Catania
The prototype DataGrid testbed

CERN openlab for DataGrid applications

Mission:

• Testbed for cutting edge Grid software and hardware

• Industry consortium for Grid-related technologies of common interest

• Training ground for a new generation of engineers to learn about Grid

Partners:

• CERN
• ENTERASYS
• HP
• IBM
• INTEL

hp invent

ENTERASYS NETWORKS™

IBM®

intel®

CERN openlab for DataGrid applications

CERN openlab for DataGrid applications

CERN opencluster:

- Build an ultrahigh performance computer cluster
- Link it to the DataGrid and test its performance
- Evaluate potential of future commodity technology for LHC

Student Program:

- student teams get hands-on experience with some of the latest hardware and software technologies for the Grid
- learn about how CERN and its partners are developing Grid technology for scientific and industrial purposes
- external lab visits and special invited talks

CERN
openlab for DataGrid applications

CERN Openlab III

Framework for evaluating and integrating cutting edge IT technologies or services in partnership with industry, focusing on future versions of the WLCG.

Phase III focuses on

- Security
- Automation tools
- Data replication and monitoring
- Behaviour of a network of thousand machines
- Platform: thermal optimisation, applications...

LHC Computing Grid (LCG)

Mission:

• Grid deployment project aimed at installing a functioning Grid to help the LHC experiments collect and analyse the data coming from the detectors

Strategy:

• Integrate thousands of computers at dozens of participating institutes worldwide into a global computing resource

• Rely on software being developed in advanced grid technology projects, both in Europe and in the USA

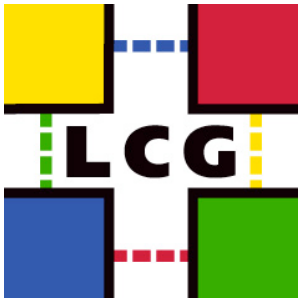

LHC Computing Grid (LCG)

People:

• Over 150 physicists, computer scientists and engineers from partner research centres around the world

Timeline:

• 2002: start project

• 2003: service opened (Sept)

• 2002 - 2005: prepare and deploy the environment for LHC computing

• 2006 – 2008: acquire, build and operate the LHC computing service

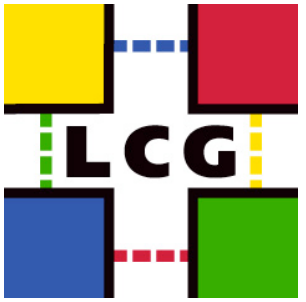

GridCafé

LHC Computing Grid

Mission:

• Install a functioning grid to help the LHC experiments collect and analyse data coming from the detectors

Strategy:

• Integrate thousands of computers at hundreds of participating institutes worldwide into a global computing resource.

Results:

The Worldwide LCG launched in October 2008 with more than 100,000 processors from 140 institutions in 33 countries, producing a massive distributed supercomputer that will provide more than 7000 physicists around the world with near real-time access to LHC data, and the power to process it.

Data acquisition and storage for LHC experiments

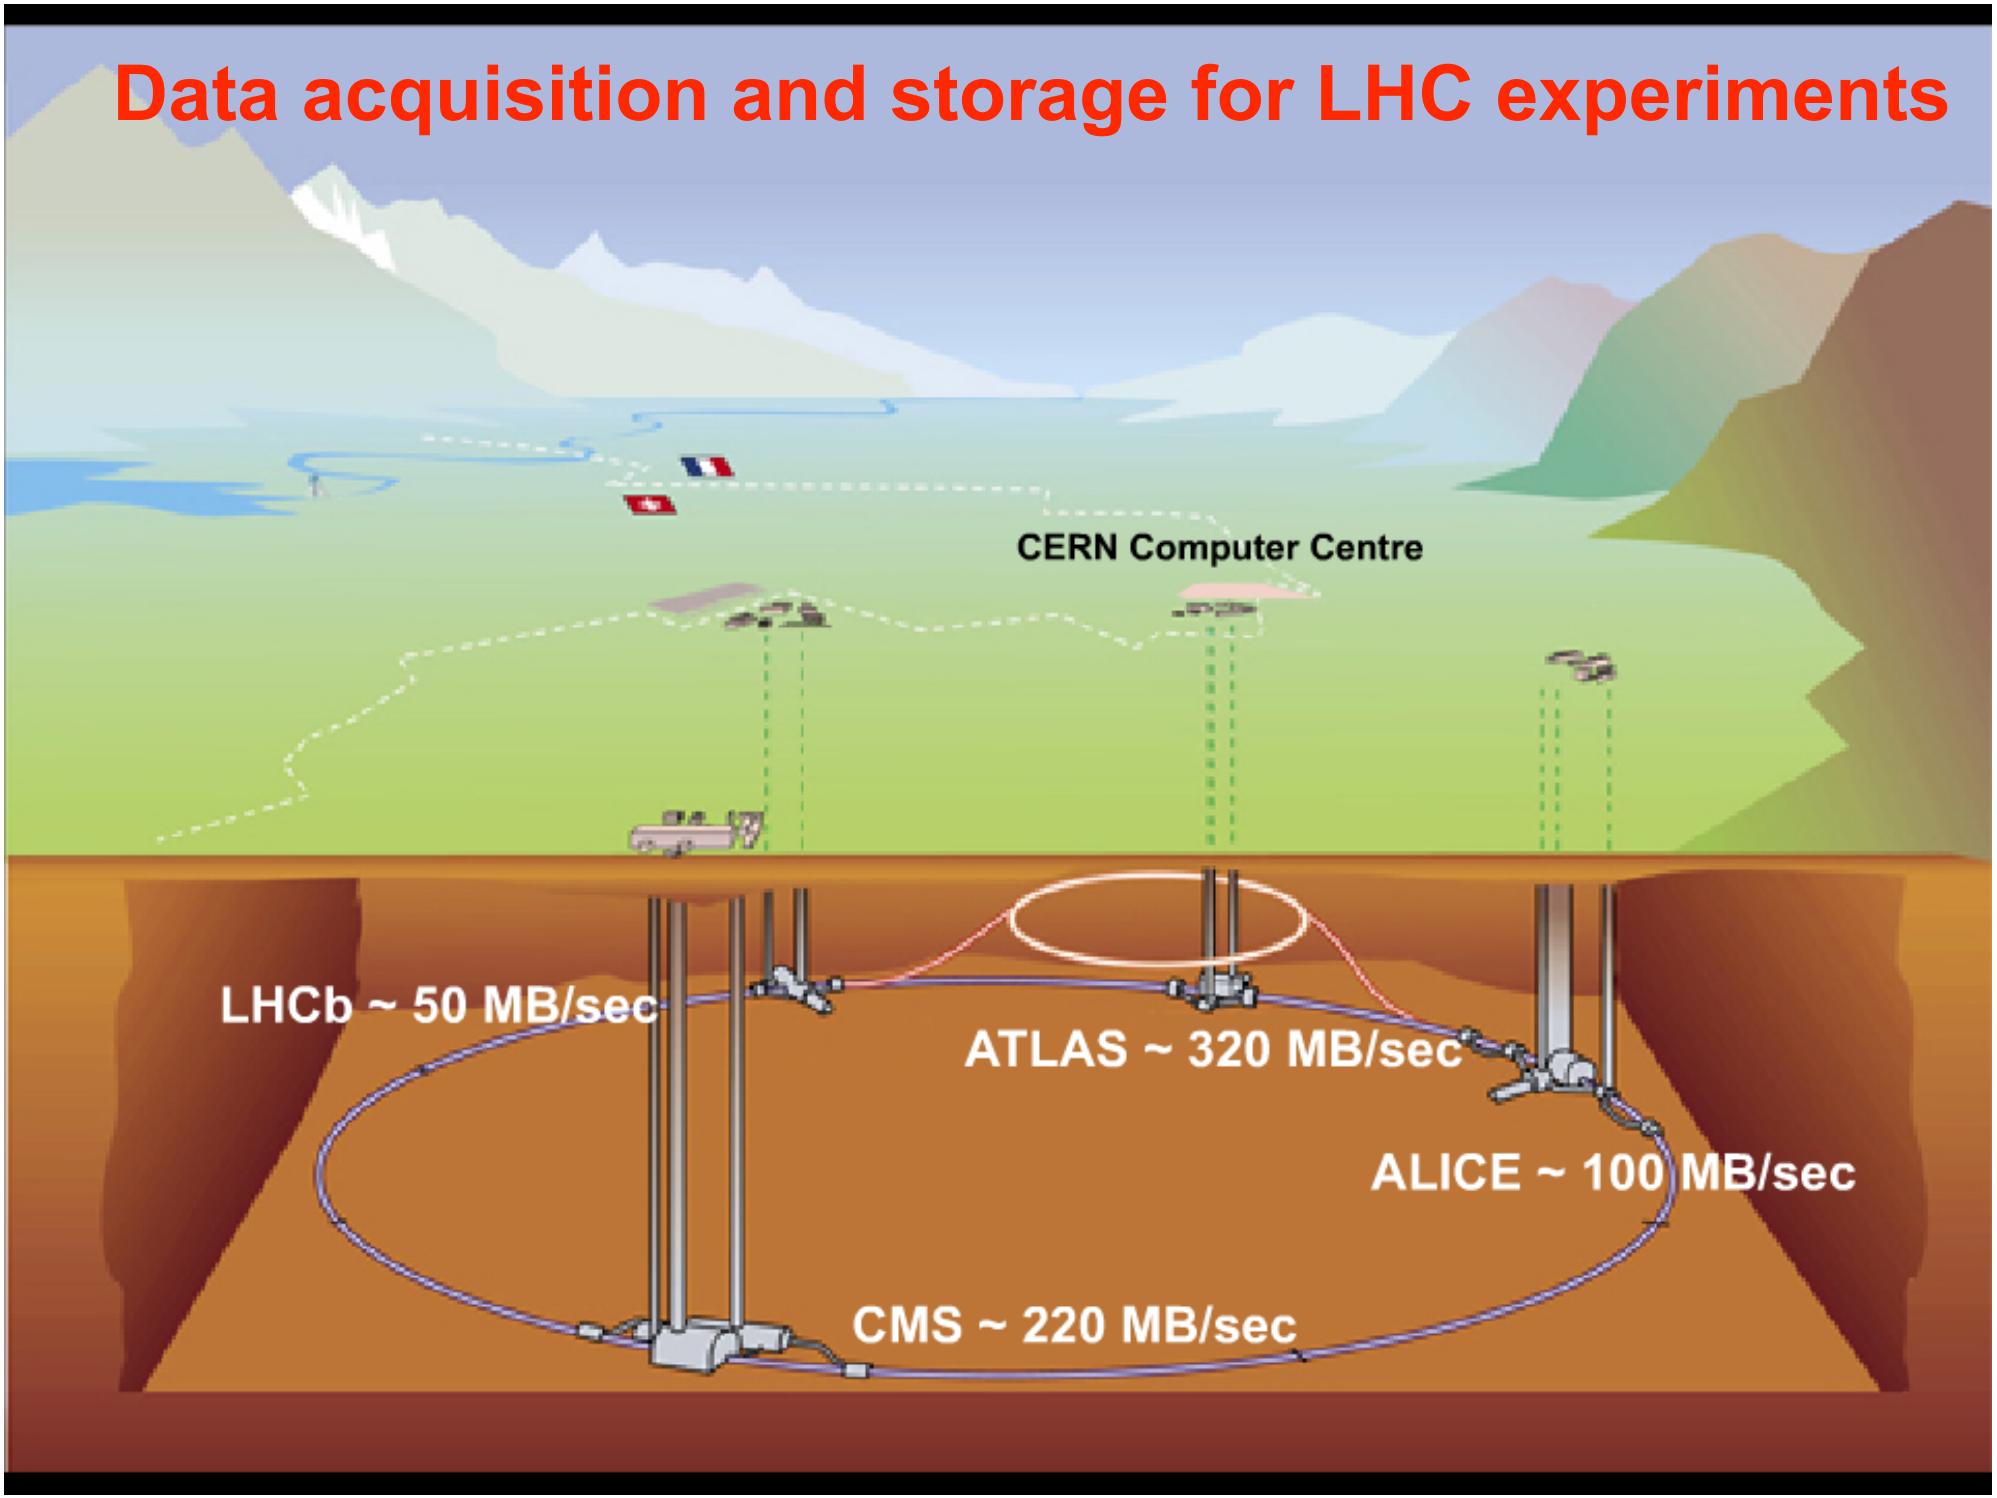

Experiment dataflow

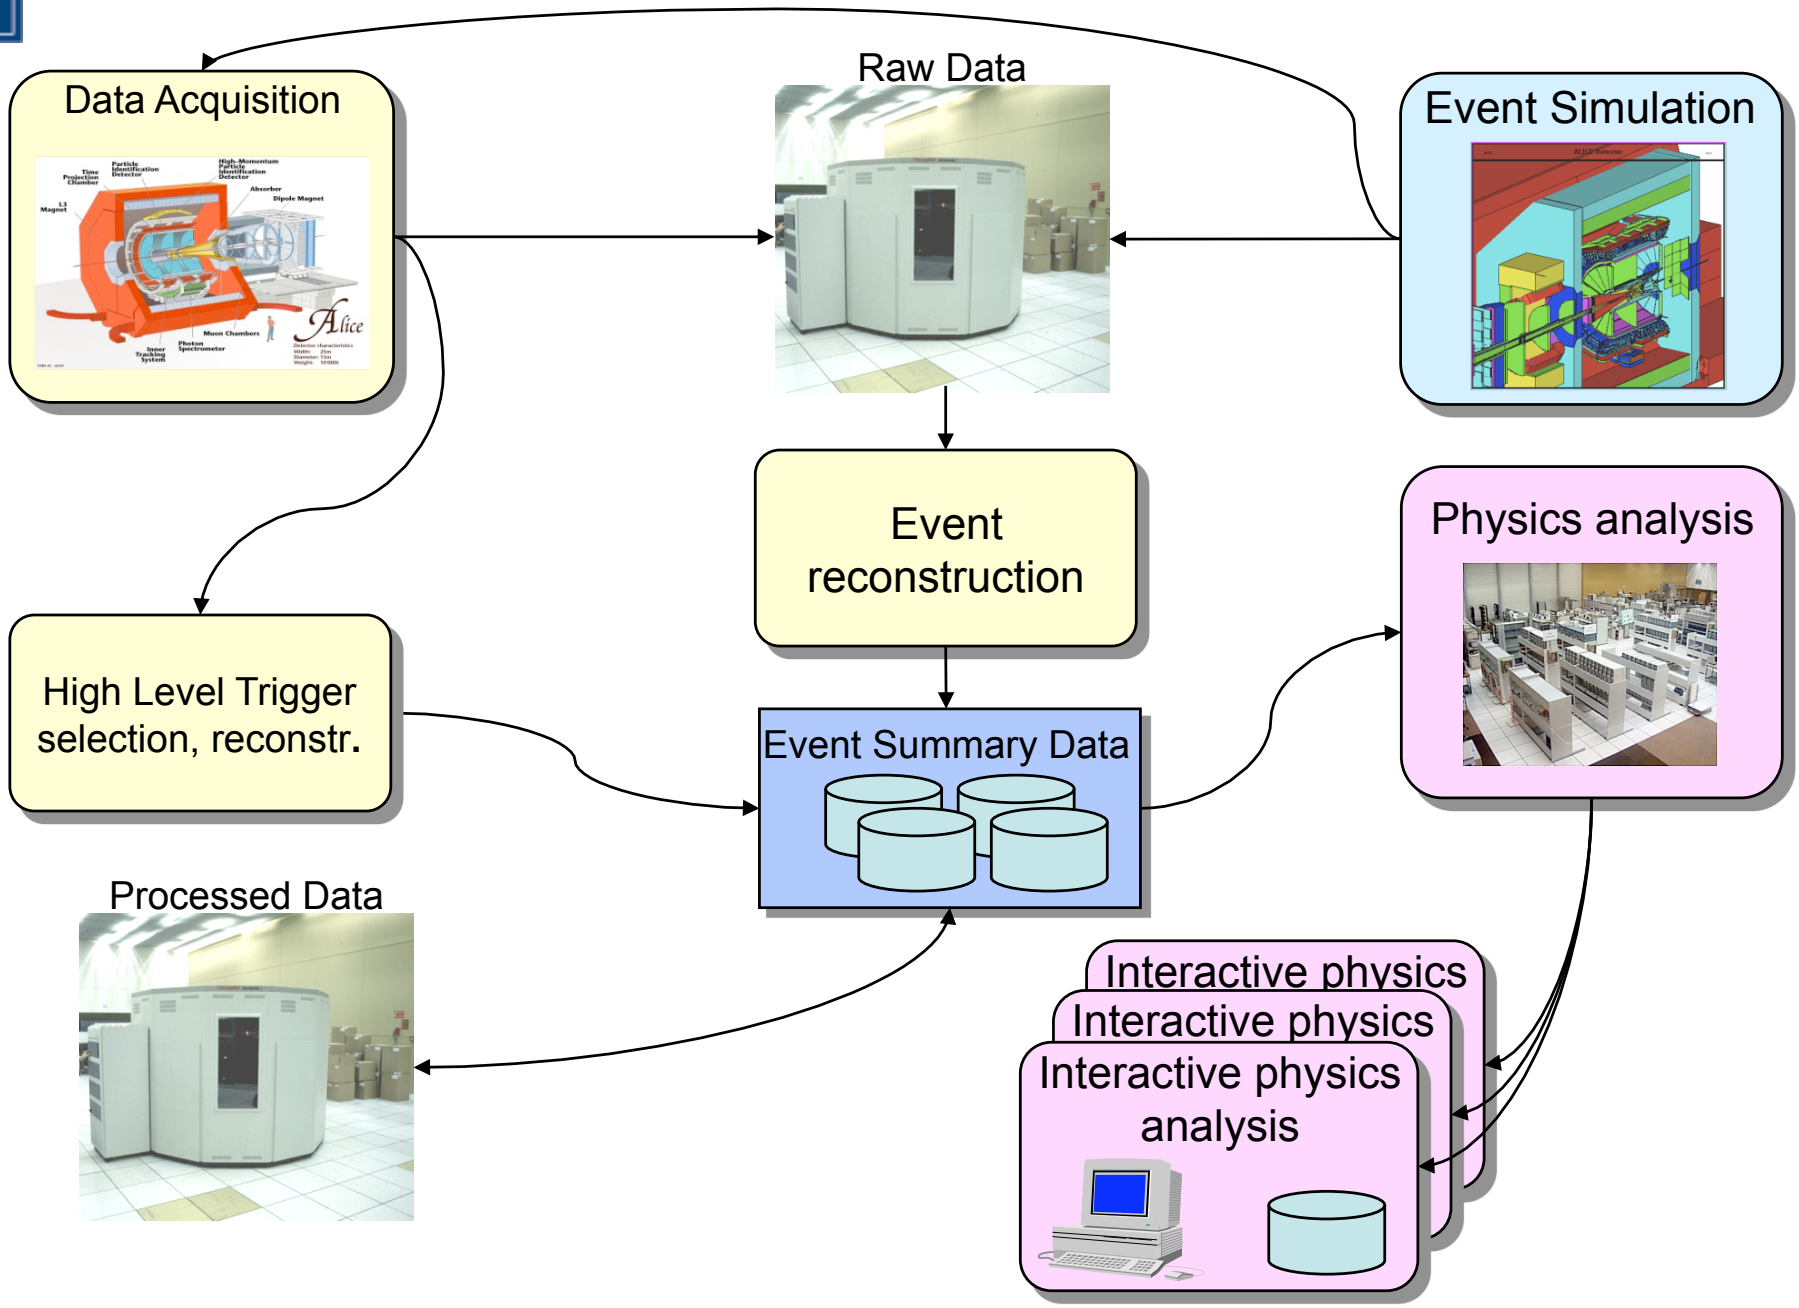

Resources needed for the LHC computing

Summary of Computing Resource Requirements

All experiments - 2008

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**

CPU

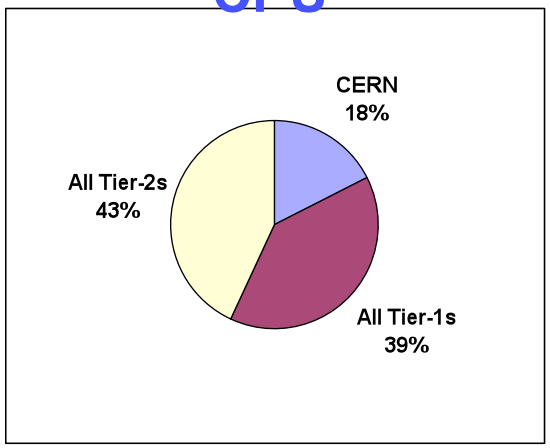

Disk

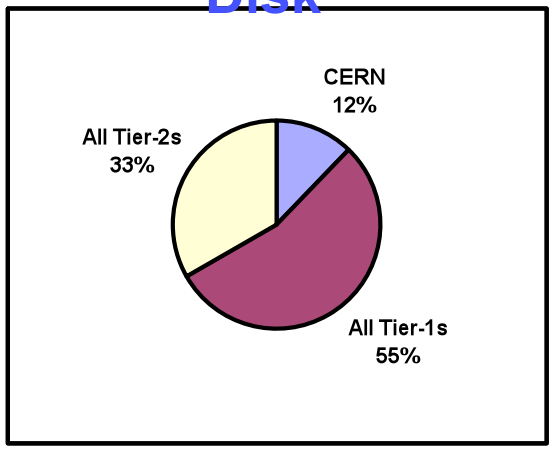

Tape

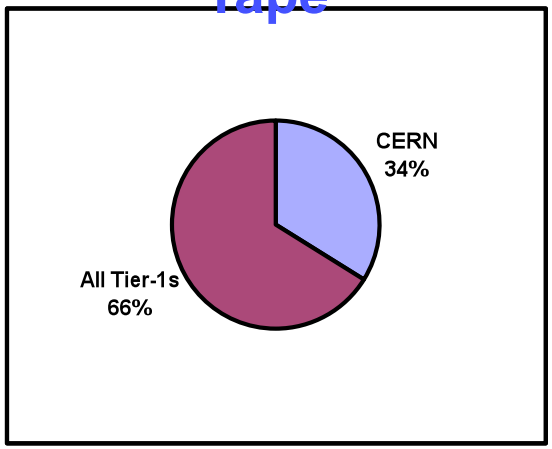

WLCG: Worldwide LHC Computing Grid

- Biggest scientific Grid project in the world
- More than 140 sites in 35 countries
- 90'000 processors by end of 2009
- 12 large centres for primary data management: CERN (Tier-0) and eleven Tier-1s
- 38 federations of smaller Tier-2 centres
- All Tiers 1 must store all the data produced by LHC
- Should run 100 millions programmes in 2009
- Used by 5000 scientists in 500 institutes

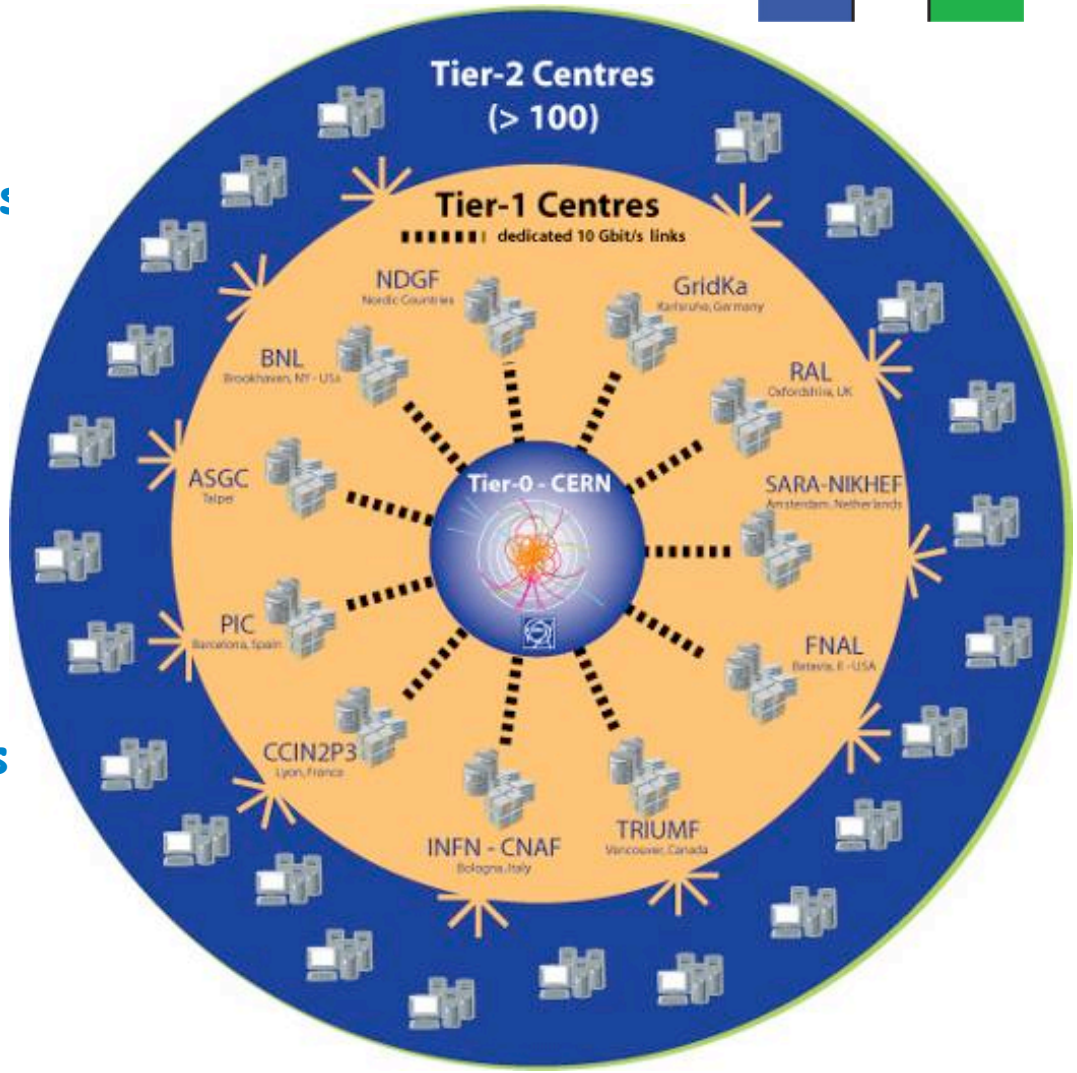

One Tier0 Centre (CERN)

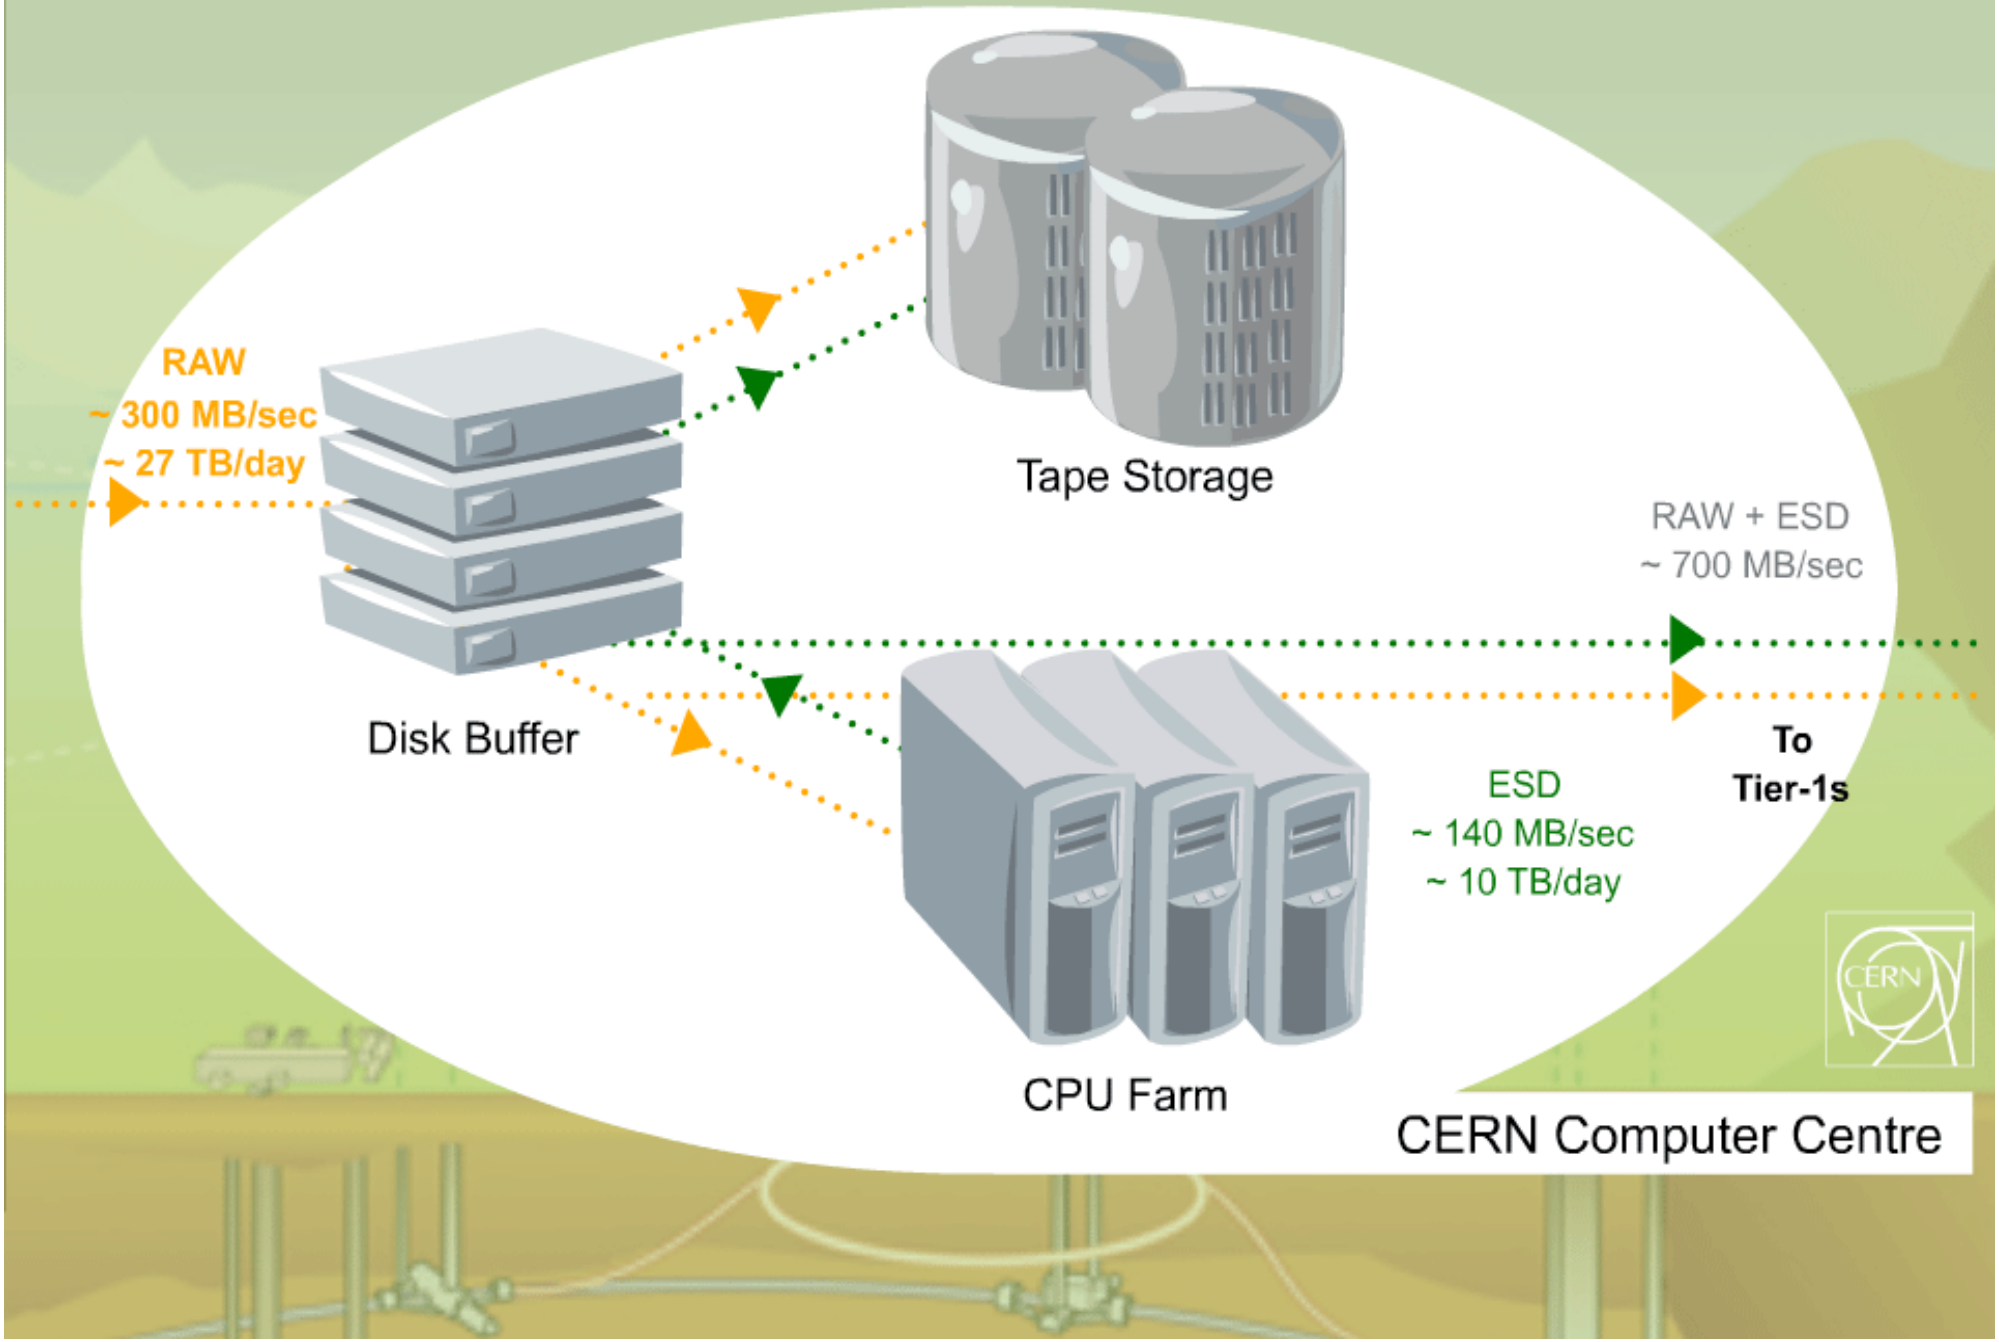

Eleven Tier1 Centres

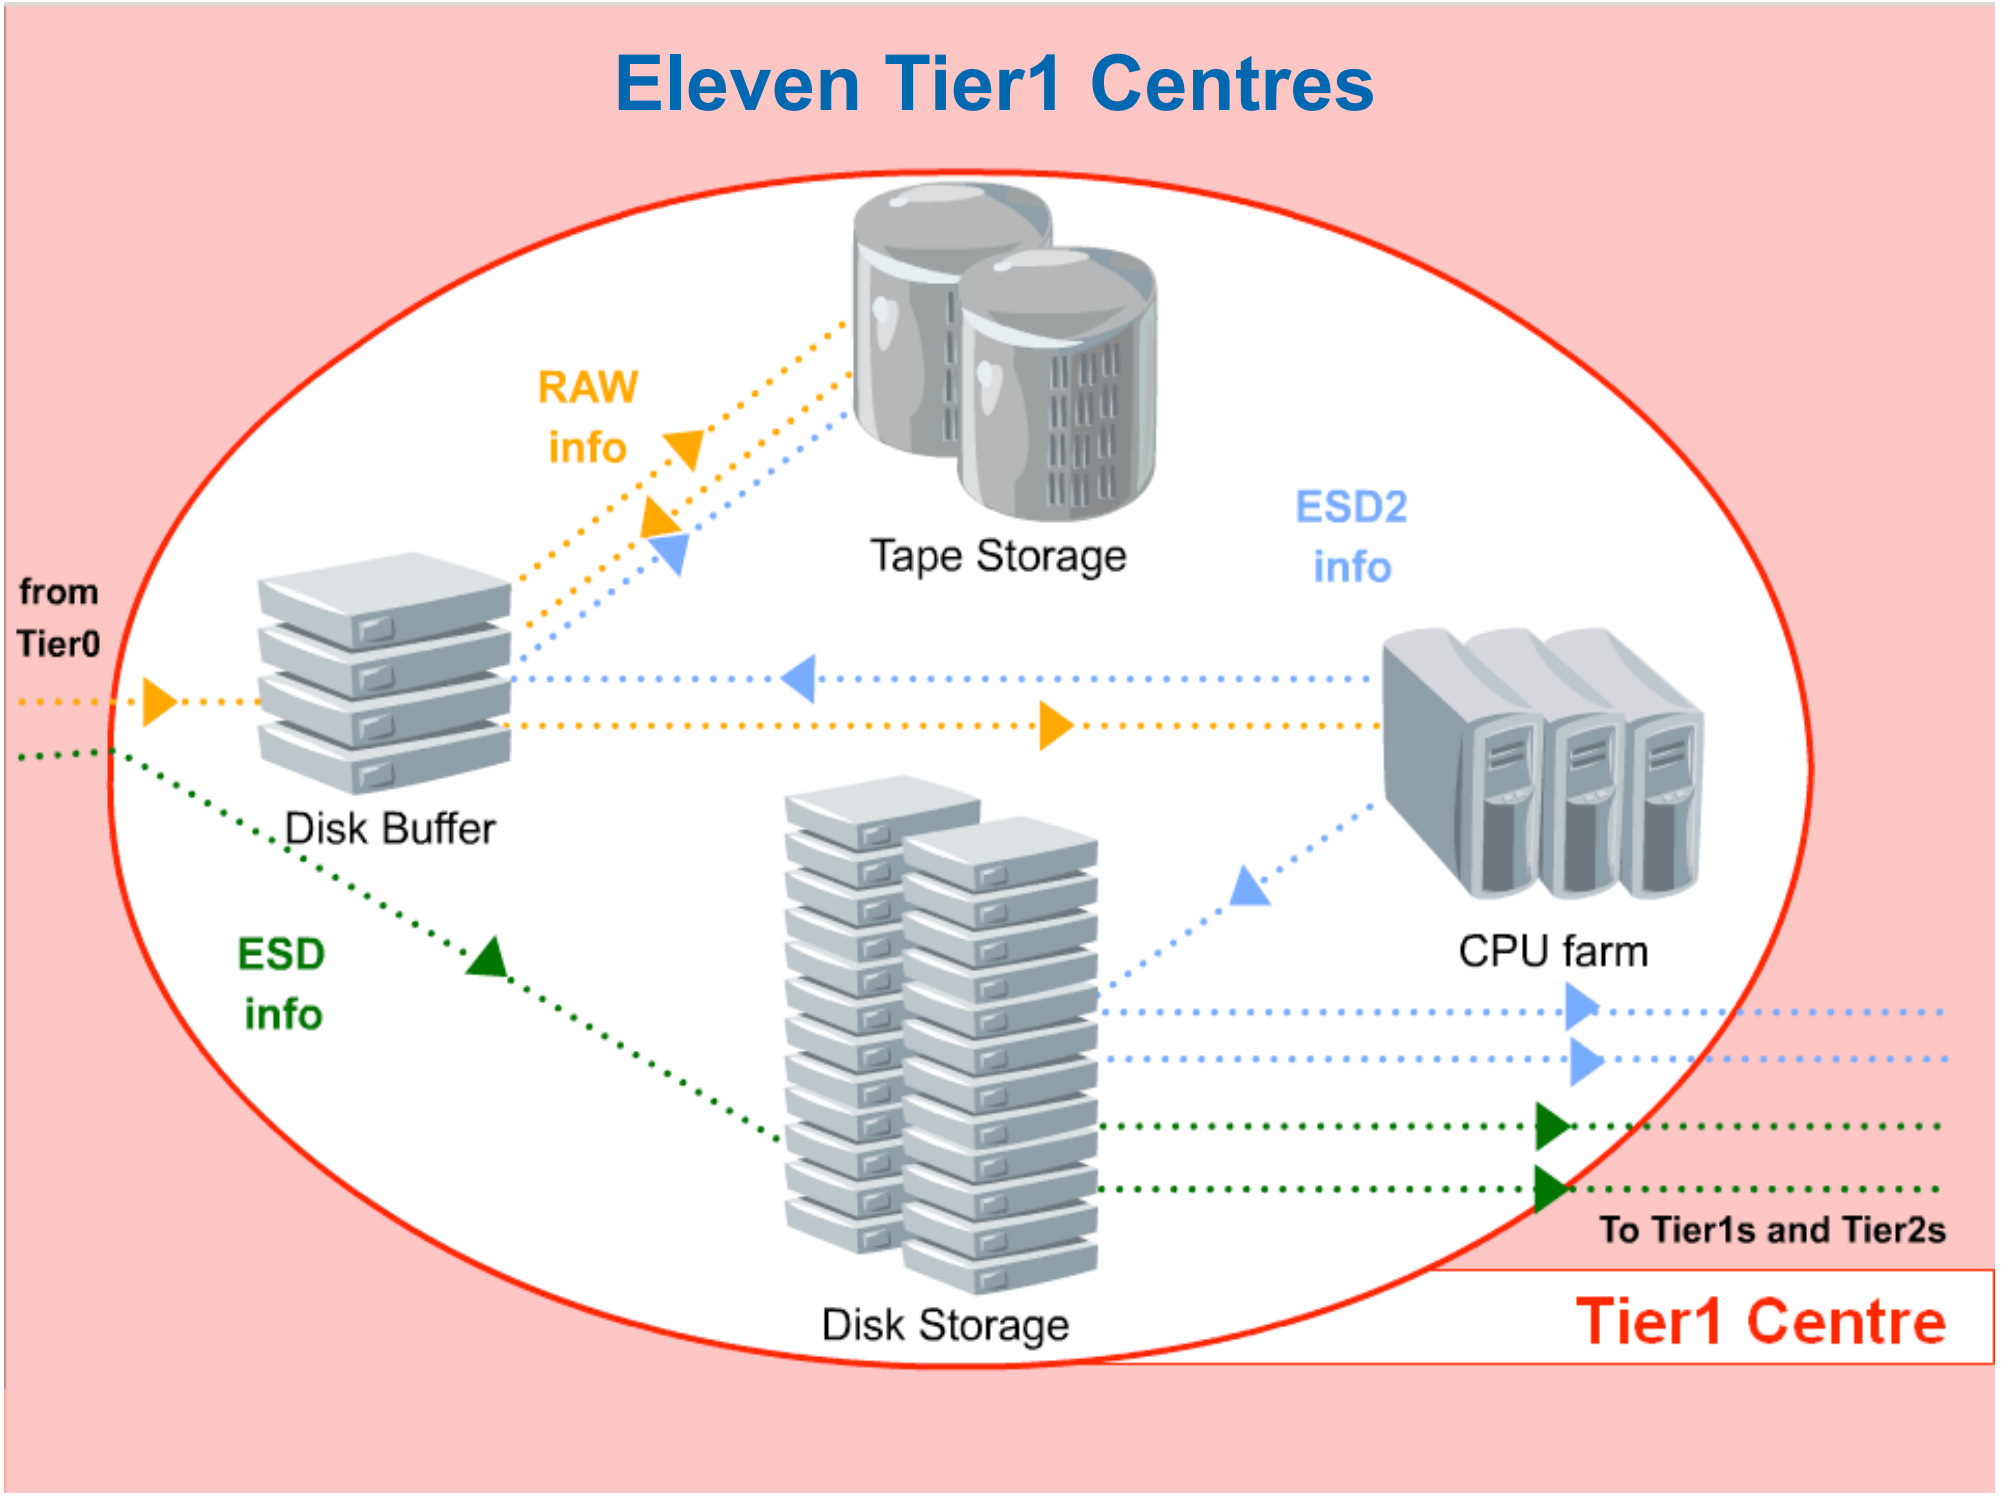

>100 Tier2 Centres

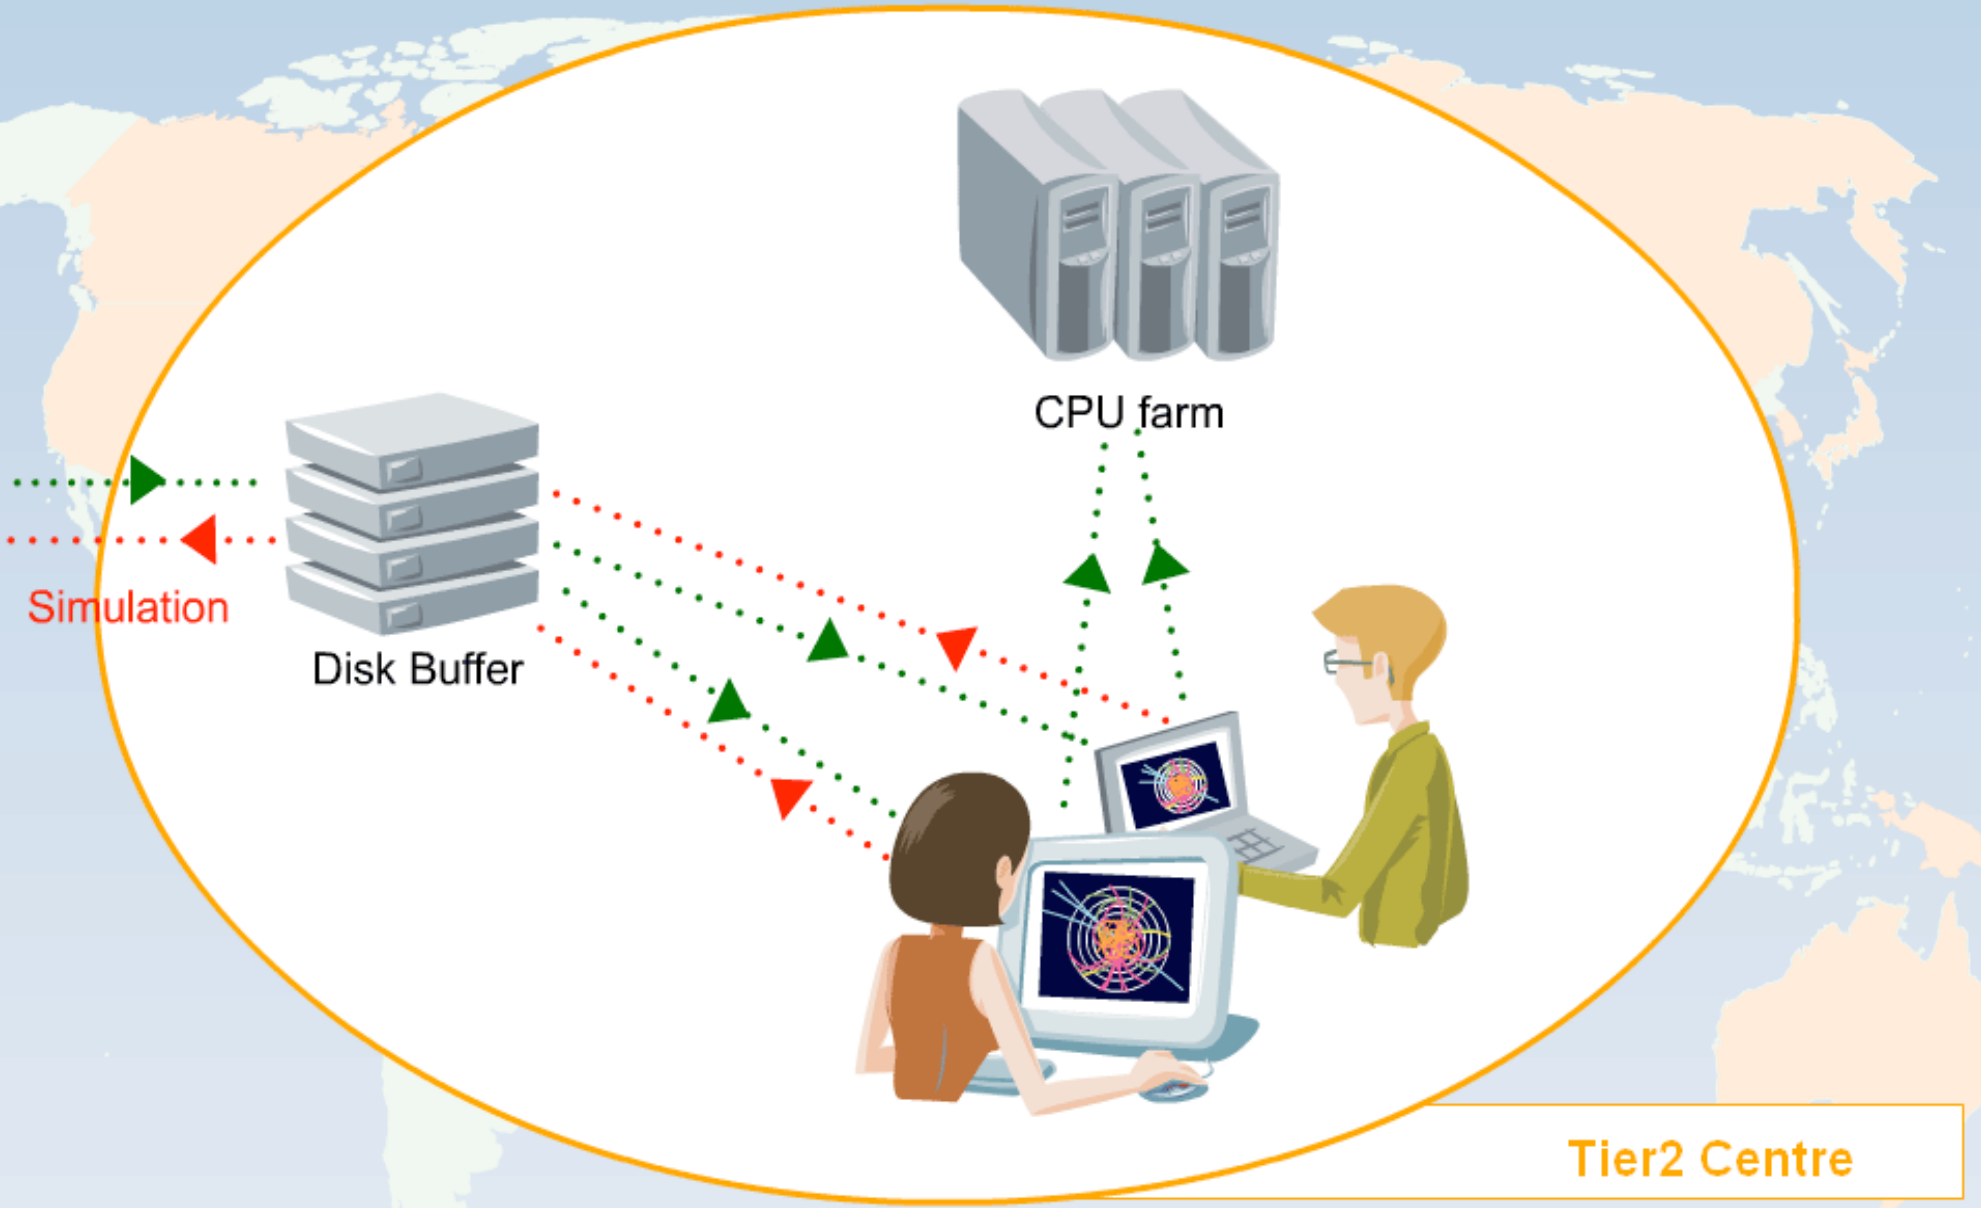

WLCG: Worldwide LHC Computing Grid

- Project to build and maintain data storage and computing infrastructure for LHC
- Uses infrastructure of several Grid organizations where 2 the biggest ones are
  - EGEE (founded by EC)
  - OSG (founded by US)

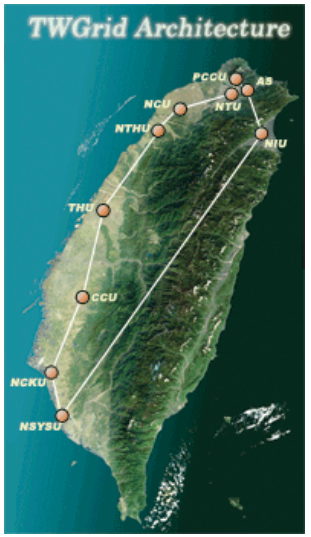

GridPP
UK Computing for Particle Physics

NORDUGRID
Grid Solution for Wide Area Computing and Data Handling

Computing @ CERN

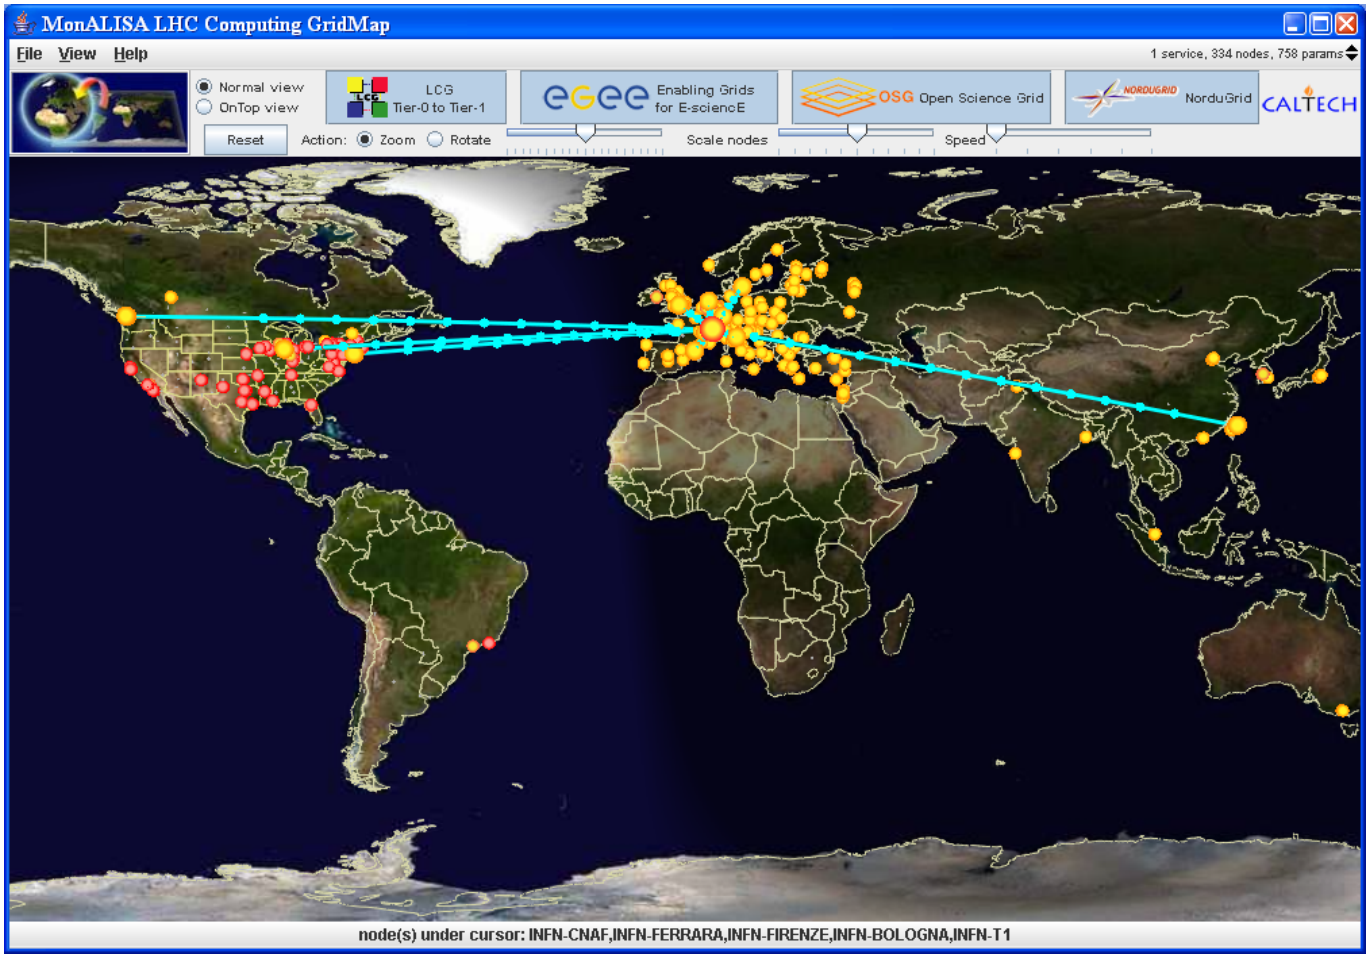

Grid@CERN

- CERN projects:
  Worldwide LHC Computing Grid (WLCG)

- External projects with CERN participation:
  Enabling Grids for E-SciencE (EGEE)
  European Grid Initiative (EGI)

- Industry funded projects:
  CERN Openlab

GridCafé

Enabling Grids for E-Science

Creating a global grid for global e-science

Mission:

• Deliver 24/7 grid service to European science; re-engineer grid middleware for production; market grid solutions to different scientific communities

Results:

•EGEE currently involves more than 240 institutions in 45 countries, supporting science in more than 20 disciplines, including bioinformatics, medical imaging, education, climate change, energy, agriculture and more.

Enabling Grids for e-Science (EGEE)

- EGEE Scope : ALL-Inclusive for academic applications (open to industrial and socio-economic world as well)
- LCG has been the driving force for the European multi-science Grid EGEE (Enabling Grids for E-sciencE)
- EGEE is now a global effort, and the largest Grid infrastructure worldwide
  - Used by 10,000 users from at least 15 disciplines
  - Provided by 259 sites from 52 countries
  - over 72,000 CPU's
  - over 20 Petabytes (10^{15}) HDD storage
  - massive data transfer: 1.5 GB/s
  - between 30K and 150K jobs/day

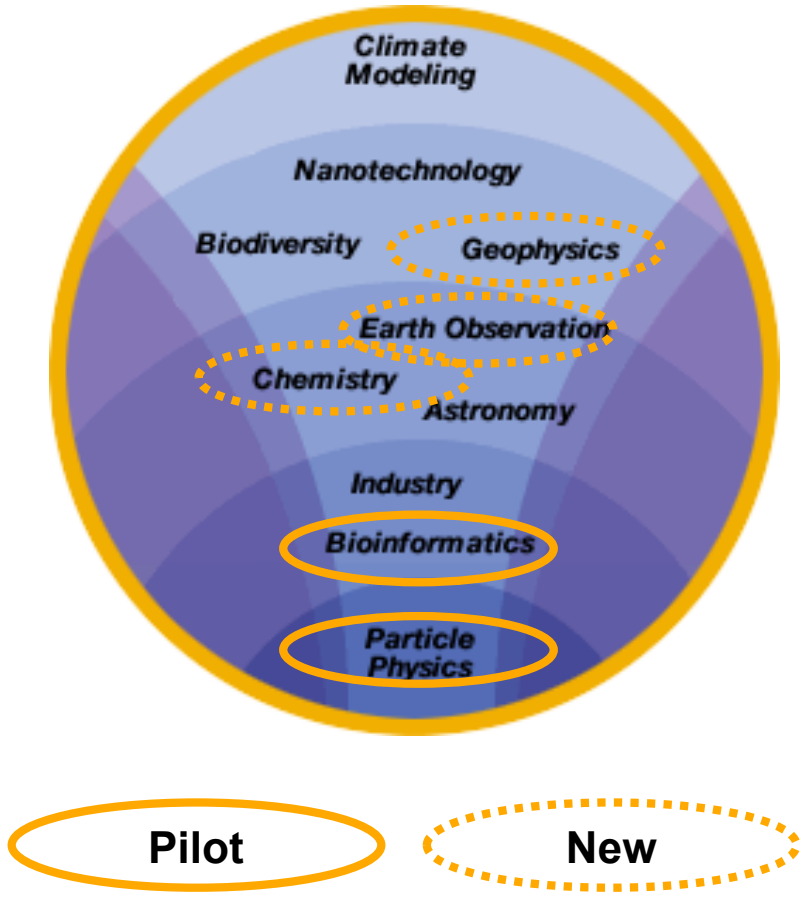

The Grid Vision for e-Science

Researchers perform their activities regardless geographical location, interact with colleagues, share and access data

The Grid: networked data processing centres and ”middleware” software as the “glue” of resources.

Grids are both a “dream” and a tool for realizing even larger “dreams”.
(http://access.ncsa.uiuc.edu/witg/)

Scientific instruments and experiments provide huge amount of data

GridCafé

The EGEE Vision

Grid computing is already changing the way science and much else is done around the world. What will the future hold?

An international network of scientists will be able to model a new flood of the Danube in real time, using meteorological and geological data from several centers across Europe.

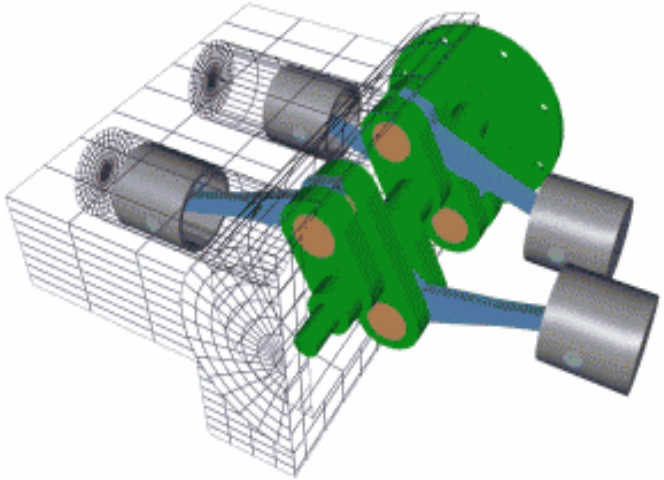

A team of engineering students will be able to run the latest 3D rendering programs from their laptops using the Grid.

A geneticist at a conference, inspired by a talk she hears, will be able to launch a complex biomolecular simulation from her mobile phone.

EGEE Status July 2008

- **Infrastructure**
  - Number of sites connected to the EGEE infrastructure: 259
  - Number of countries connected to the EGEE infrastructure: 52
  - Number of CPUs available to users 24/7: ~ 72,000
  - Storage capacity available: ~ 20 PB disk + tape MSS
- **Users**
  - Number of Virtual Organisations using the EGEE infrastructure: > 200 (this number is according to gstat)
  - Number of registered Virtual Organisations: >130
  - Number of registered users: >7500
  - Number of people benefiting from the existence of the EGEE infrastructure: ~14000
  - Number of jobs: >150k jobs/day
  - Number of application domains making use of the EGEE infrastructure: more than 15
    - Archeology
    - Astronomy & Astrophysics
    - Civil Protection
    - Computational Chemistry
    - Computational Fluid Dynamics
    - Computer Science/Tools
    - Condensed Matter Physics
    - Earth Sciences
    - Finance (through the Industry Task Force)
    - Fusion
    - Geophysics
    - High-Energy Physics
    - Life Sciences
    - Multimedia
    - Material Sciences
    - ... ***Further applications under evaluation***

Project

Duration: 24 months
European Commission contribution: 32,000,000 euro
Total budget: cca. 47,150,000 euro, with a further estimated 50,000,000 euro worth of computing resources contributed by the partners.
Total manpower: 9,010 Person Months, of which over 4,500 Person Months contributed by the partners from their own funding sources.

For the overall project metrics, please also refer to the Quality Assurance section.

GridCafé

Grid benefits

- More effective and seamless collaboration of dispersed communities, both scientific and commercial
- Ability to run large-scale applications comprising thousands of computers, for wide range of applications
- Transparent access to distributed resources from your desktop, or even your mobile phone
- The term "e-Science" has been coined to express these benefits

One Web but many Grids

Grid development has been initiated by the academic, scientific and research community, but industry is also interested.

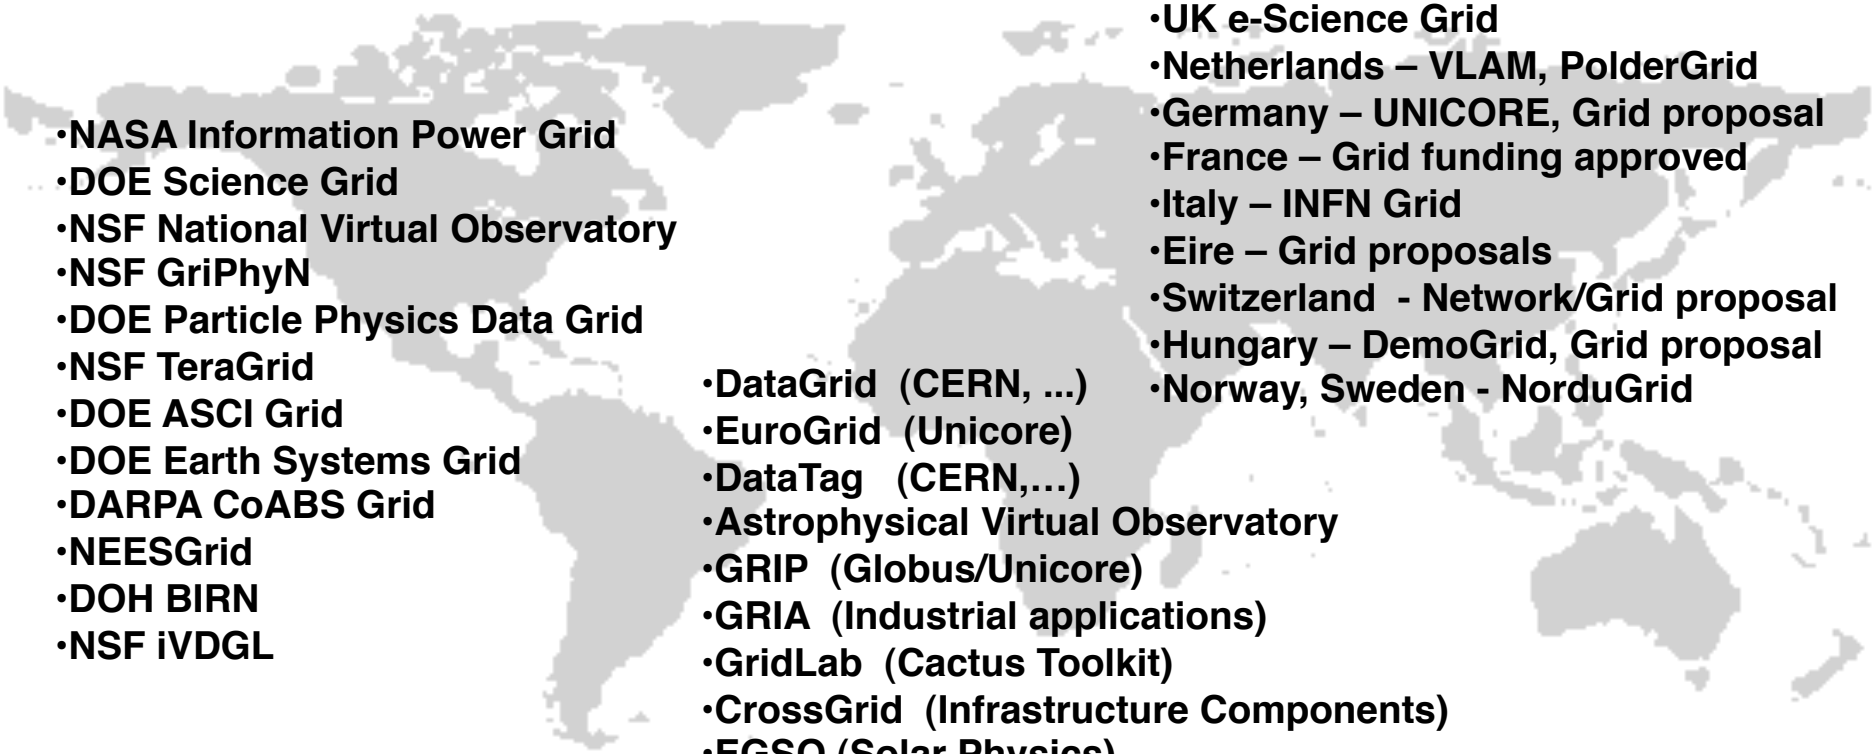

- NASA Information Power Grid
- DOE Science Grid
- NSF National Virtual Observatory
- NSF GriPhyN
- DOE Particle Physics Data Grid
- NSF TeraGrid
- DOE ASCI Grid
- DOE Earth Systems Grid
- DARPA CoABS Grid
- NEESGrid
- DOH BIRN
- NSF iVDGL

- DataGrid (CERN, ...)
- EuroGrid (Unicore)
- DataTag (CERN,…)
- Astrophysical Virtual Observatory
- GRIP (Globus/Unicore)
- GRIA (Industrial applications)
- GridLab (Cactus Toolkit)
- CrossGrid (Infrastructure Components)
- EGSO (Solar Physics)

- UK e-Science Grid
- Netherlands – VLAM, PolderGrid
- Germany – UNICORE, Grid proposal
- France – Grid funding approved
- Italy – INFN Grid
- Eire – Grid proposals
- Switzerland - Network/Grid proposal
- Hungary – DemoGrid, Grid proposal
- Norway, Sweden - NorduGrid

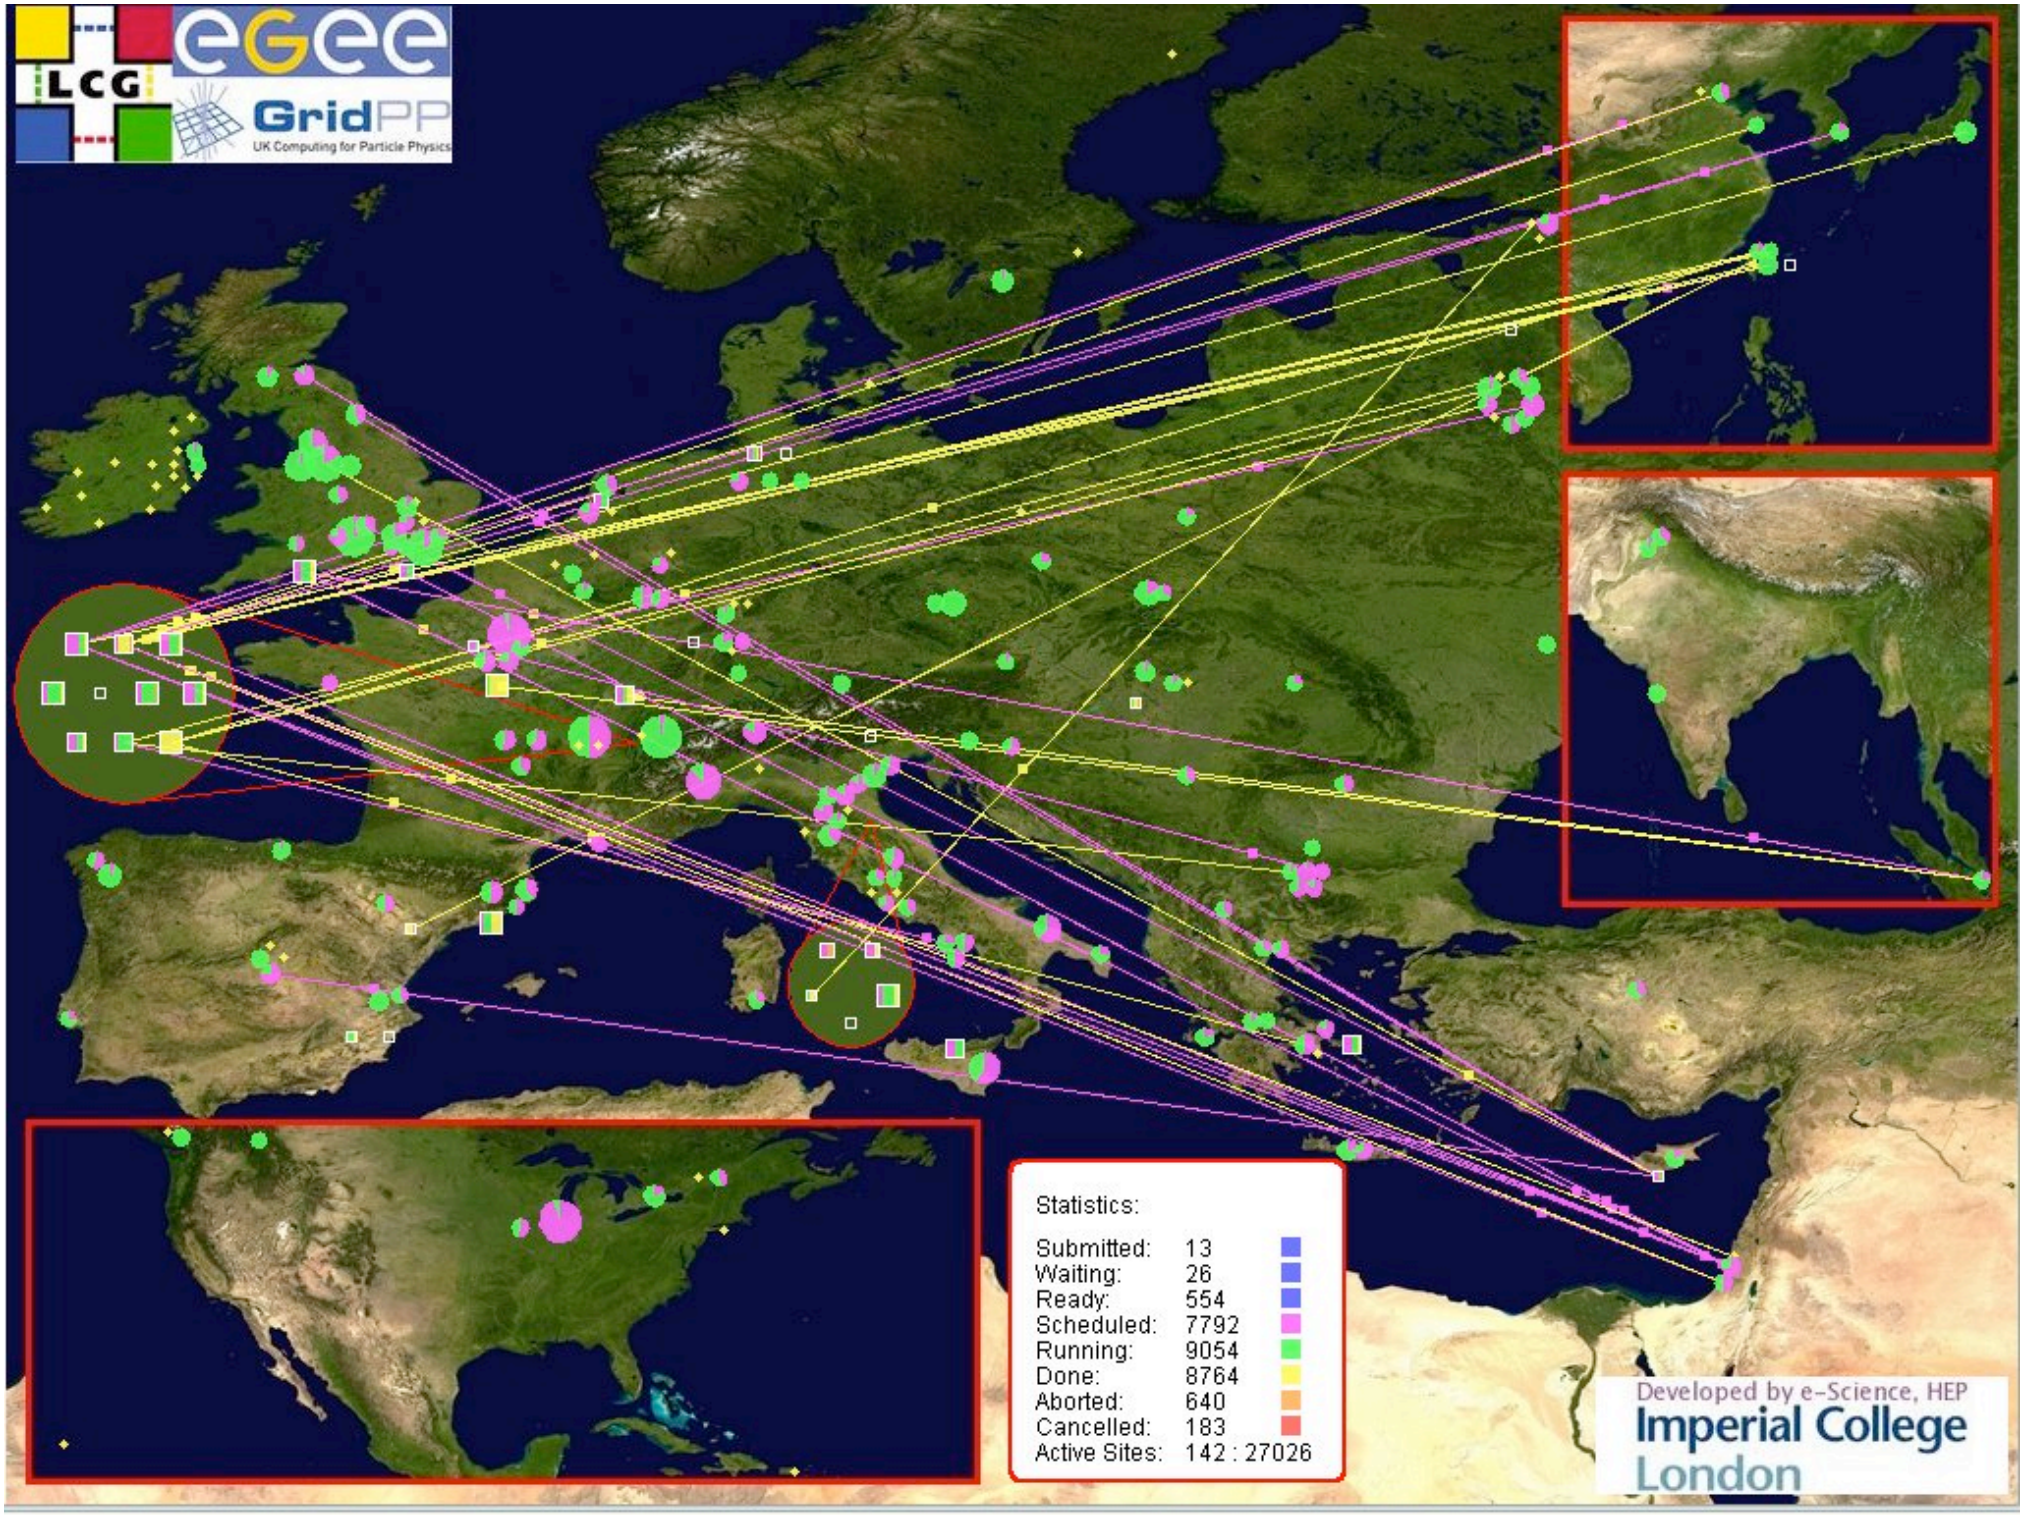
LCG
eGee
GridPP
UK Computing for Particle Physics
Statistics:
Submitted: 13
Waiting: 26
Ready: 554
Scheduled: 7792
Running: 9054
Done: 8764
Aborted: 640
Cancelled: 183
Active Sites: 142 : 27026
Developed by e-Science, HEP
Imperial College London

Central European Grid

- Central European (CE) Federation is one of 12 federations in EGEE project
- seven countries: Austria, Croatia, Czech Republic, Hungary, Poland, Slovakia and Slovenia.
- 10 EGEE-approved virtual organizations (VOs)
- Sites, around 10 000 CPUs
- 850 TBs

Source: http://goc.grid.sinica.edu.tw/gsta

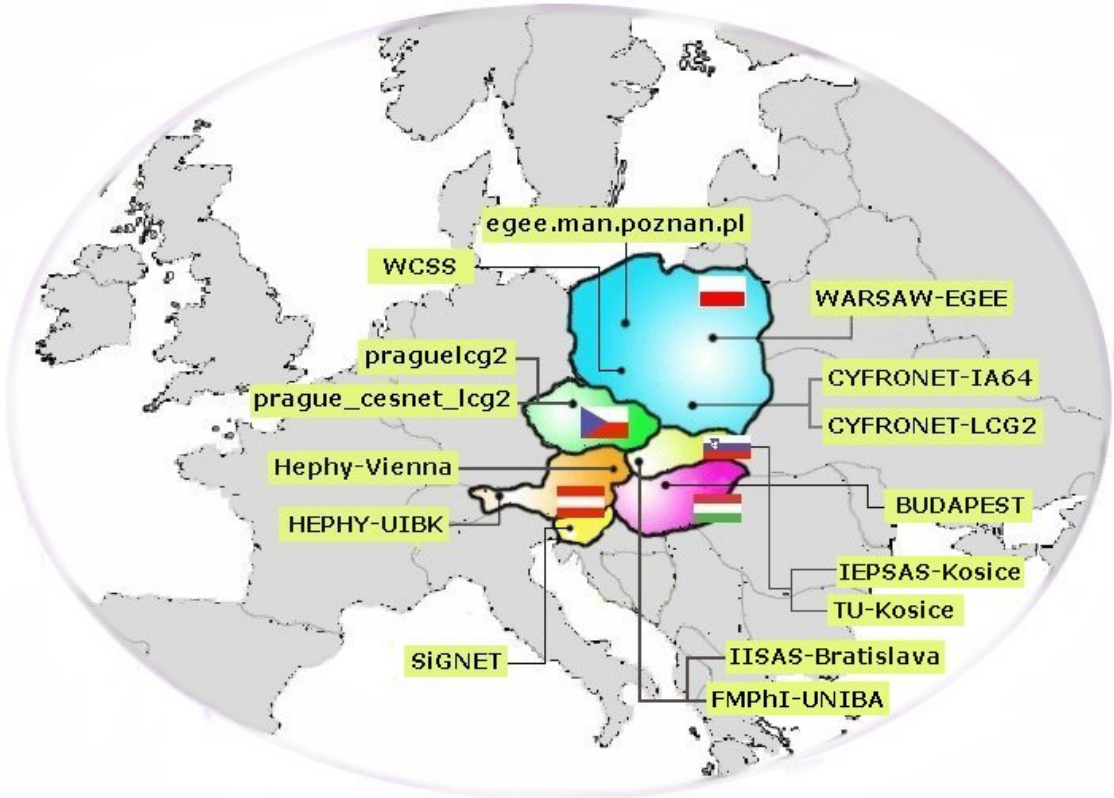

GridCafé

Grids in science

• There's so much that grids can already do! Here are some examples in:

•Medicine *(imaging, diagnosis and treatment)*

• Bioinformatics *(study of the human genome and proteome to understand genetic diseases)*

• Nanotechnology *(design of new materials from the molecular scale)*

• The environment
(weather forecasting, earth observation, modeling and prediction of complex systems)

GridCafé

Medicine

- Digital image archives
- Collaborative virtual environments
- On-line clinical conferences

"Grids will enable a standardized, distributed digital mammography resource for improving diagnostic confidence"

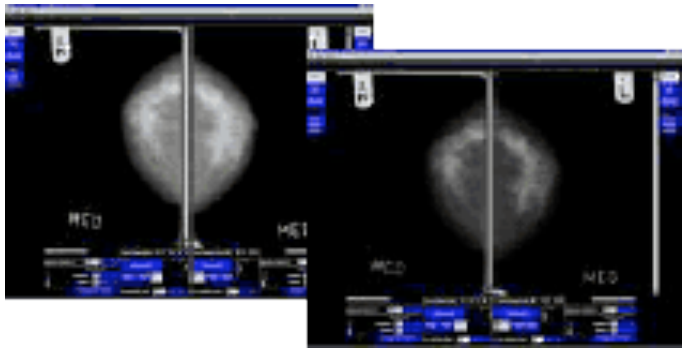

"Grids make it possible to use large collections of images in new, dynamic ways, including medical diagnosis."

"The ability to visualise 3D medical images is key to the diagnosis of pathologies and pre-surgical planning"

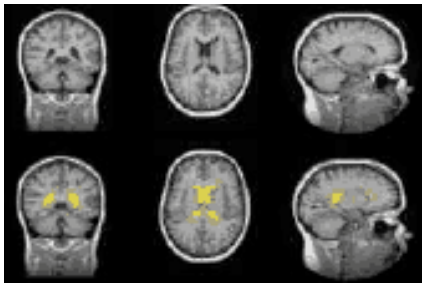

eDiaMoND: Screening for Breast Cancer

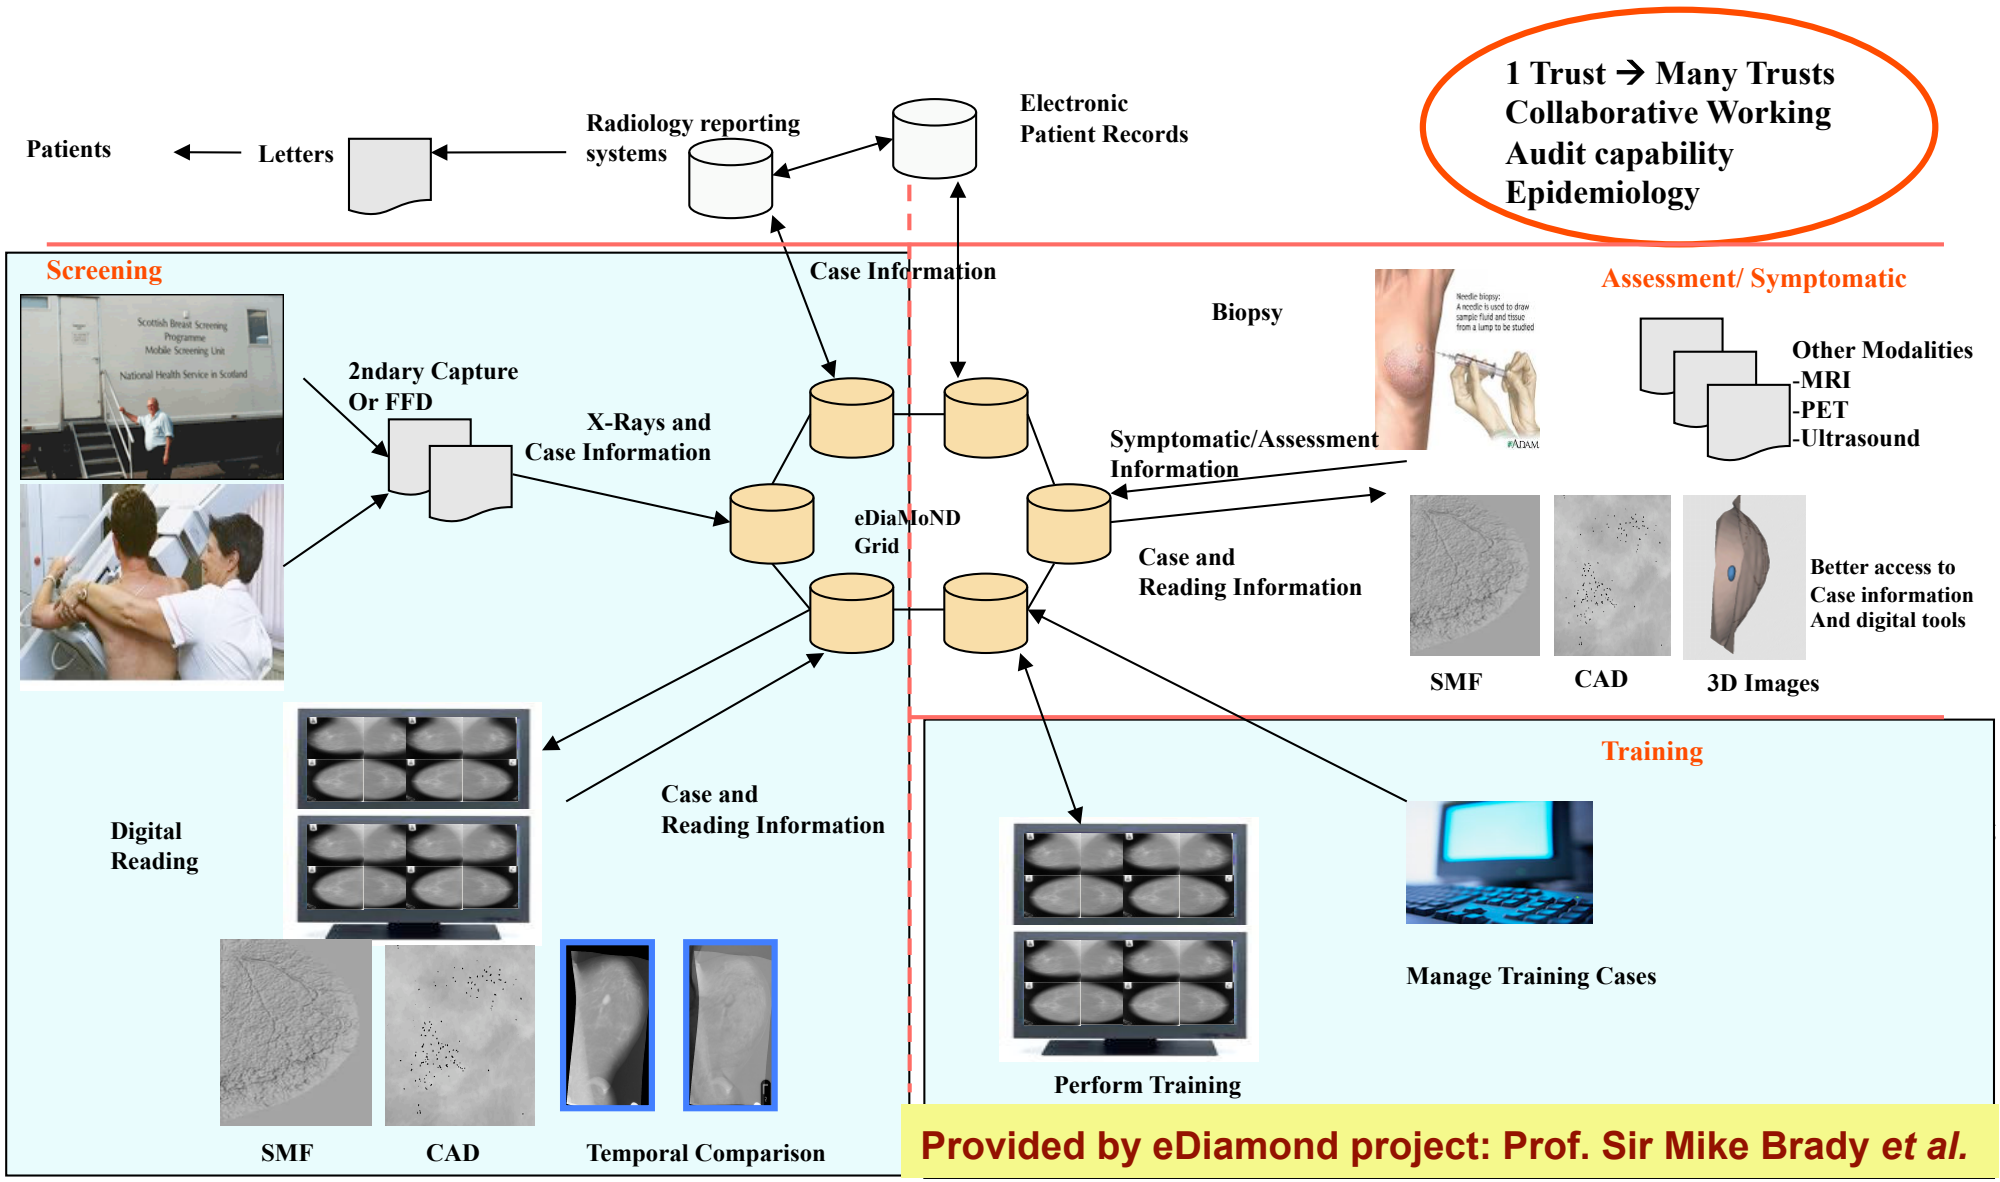

Provided by eDiamond project: Prof. Sir Mike Brady *et al.*

GridCafé

Bioinformatics

•Capturing the complex and evolving patterns of genetic information, determining the development of an embryo

• Understanding the genetic interactions that underlie the processes of life-form development, disease and evolution.

"Every time a new genome is sequenced the result is compared in a variety of ways with other genomes. Each code is made of 3.5 billion pairs of chemicals…"

GridCafé

Nanotechnology

- New and 'better' materials
- Benefits in pharmaceuticals, agrochemicals, food production, electronics manufacture from the faster, cheaper discovery of new catalysts, metals, polymers, organic and inorganic materials

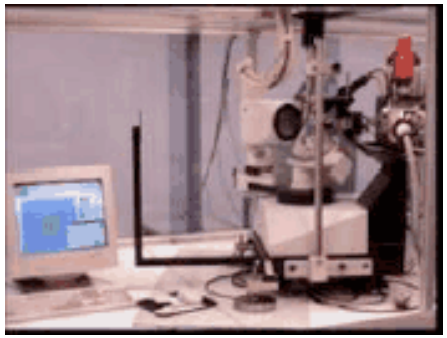

"Grids have the potential to store and analyze data on a scale that will support faster, cheaper synthesis of a whole range of new materials."

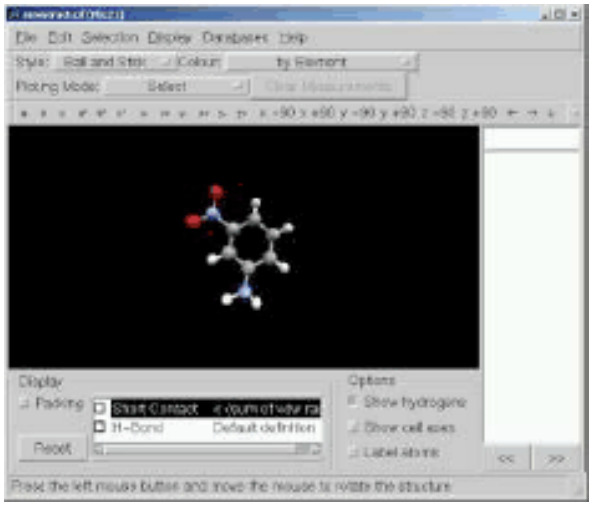

GridCafé

Environment

- Modeling and prediction of earthquakes
- Climate change studies and weather forecast
- Pollution control
- Socio-economic growth planning, financial modeling and performance optimization

"Federations of heterogeneous databases can be exploited through grid technology to solve complex questions about global issues such as biodiversity."

climate*prediction*.net and GENIE

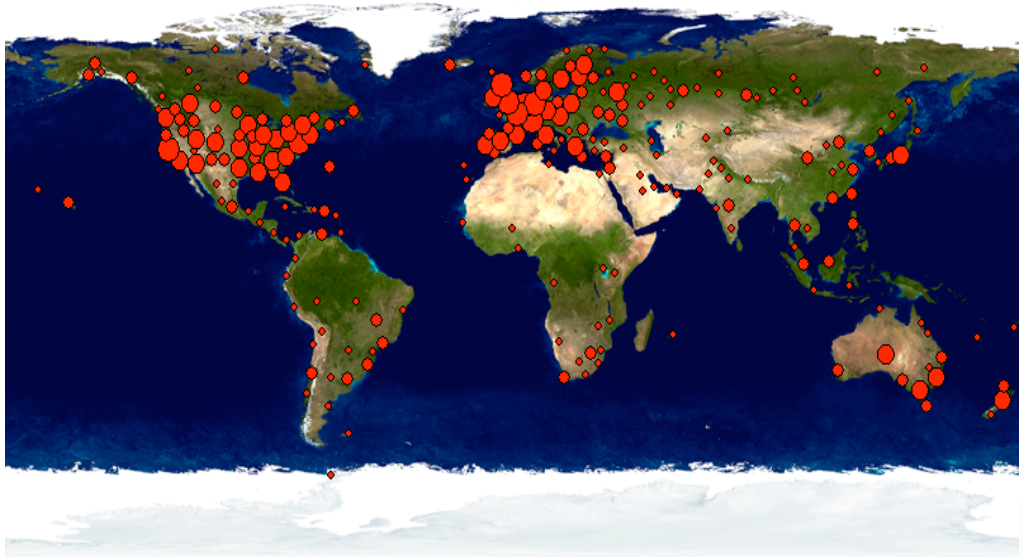

- Largest climate model ensemble
- >45,000 users, >1,000,000 model years

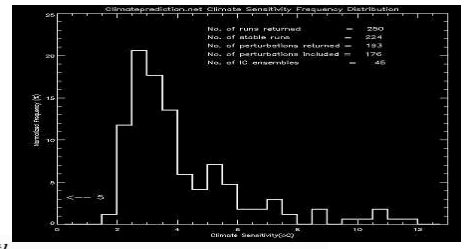

Response of Atlantic circulation to freshwater forcing

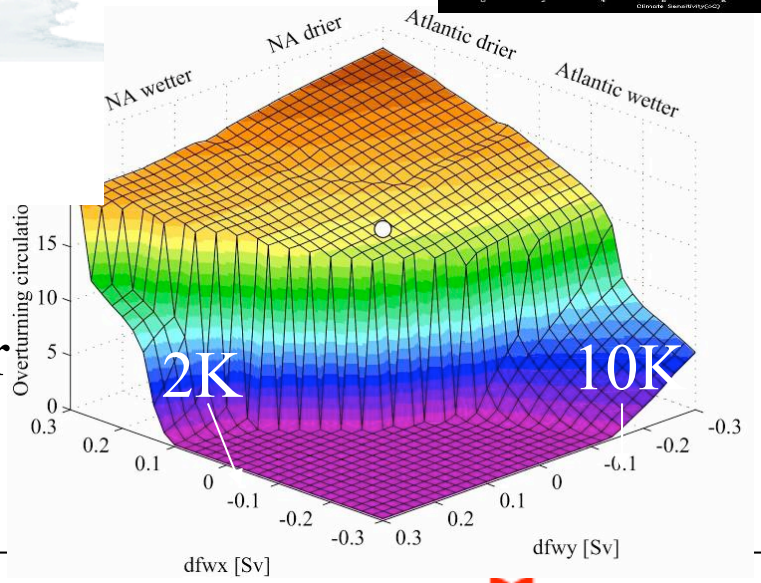

e-Science dti

NATURAL ENVIRONMENT RESEARCH COUNCIL

Other Applications

- Astronomy

- Bioinformatics

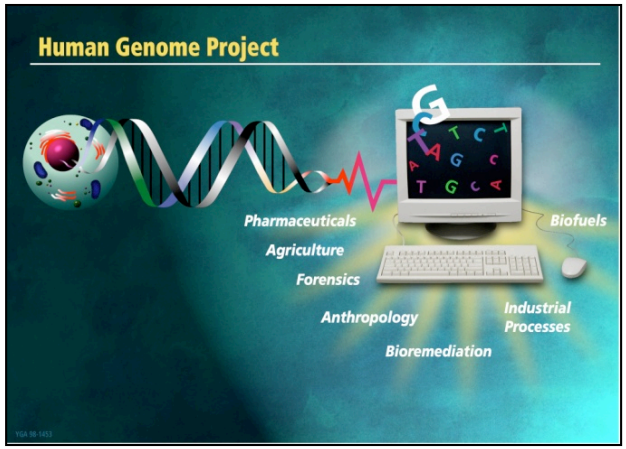

- Engineering

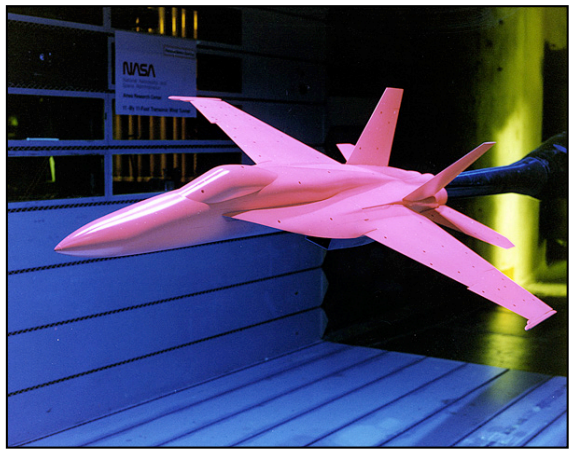

- Healthcare

- Commerce

- Gaming


Computer Science & Engineering Department
Global Grid Community
TERAGRID
UNICORE
Distributed European Infrastructure for Supercomputing Applications
D-GRID
Open Science Grid
INFN GRID
EUROGRID
National Grid NG SINGAPORE
Grid3 2003
PPDG
iVDgL
K*GRID
Data GRID
ACCESS GRID
grangenet
GRIP UK Particle Physics
DataTAG
NEESgrid
Building the National Virtual Collaboratory for Earthquake Engineering.
LCG
WestGrid
GriPhyN
Data Intensive Science
CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider
NORDUGRID
Grid Solution for Wide Area Computing and Data Handling
Versió en Català
English Version
Versión en Castellano
tota la informació sobre el projecte
presentació
descripció tècnica
els serveis oferts
participants
documentació pública
GridCat
com puc participar?
accés al Portal
Intranet per a investigadors
estat dels serveis actius
dintre del grid
APGrid
Asia-Pacific Grid
rid Consortium Japan
GridLab
eGee
Enabling Grids for E-science in Europe
NAREGI
超高速コンピュータ網形成プロジェクト
National Research Grid Initiative
Grid Applications
Grid Middleware
Networking
国立情報学研究所グリッド研究開発推進拠点 NII -The National Institute of Informatics
grid.it project
Grid-Ireland

Summary

- World Wide Web – born at CERN - elaborate set of interlinked documents accessible over the Internet
- Grid computing
  - infrastructure that provides seamless access to computing power and data storage
  - Suitable for problems that need huge/shared data, must run in parallel or need to be run frequently
- EGEE – largest existing Grid infrastructure (academic)
- Many existing applications besides high energy physics

To know more about the Grid…