Introduction to Grid Computing and the EGEE initiative

Dr. Rüdiger Berlich,
Forschungszentrum Karlsruhe / Germany

Singapore, 09.02.06

Slides contributed by FZK, INFN and EGEE Team

Forschungszentrum Karlsruhe Germany

Institute for Scientific Computing (IWR)

- Part of the „Helmholtz Gemeinschaft“
- One of the largest independent German research institutions
- Many different research areas ranging from environmental studies over nano technology to Grid Computing

Many different aspects of computing and research justify a move towards a standardised, distributed computing infrastructure.

The technical viewpoint

Exponential Growth ...

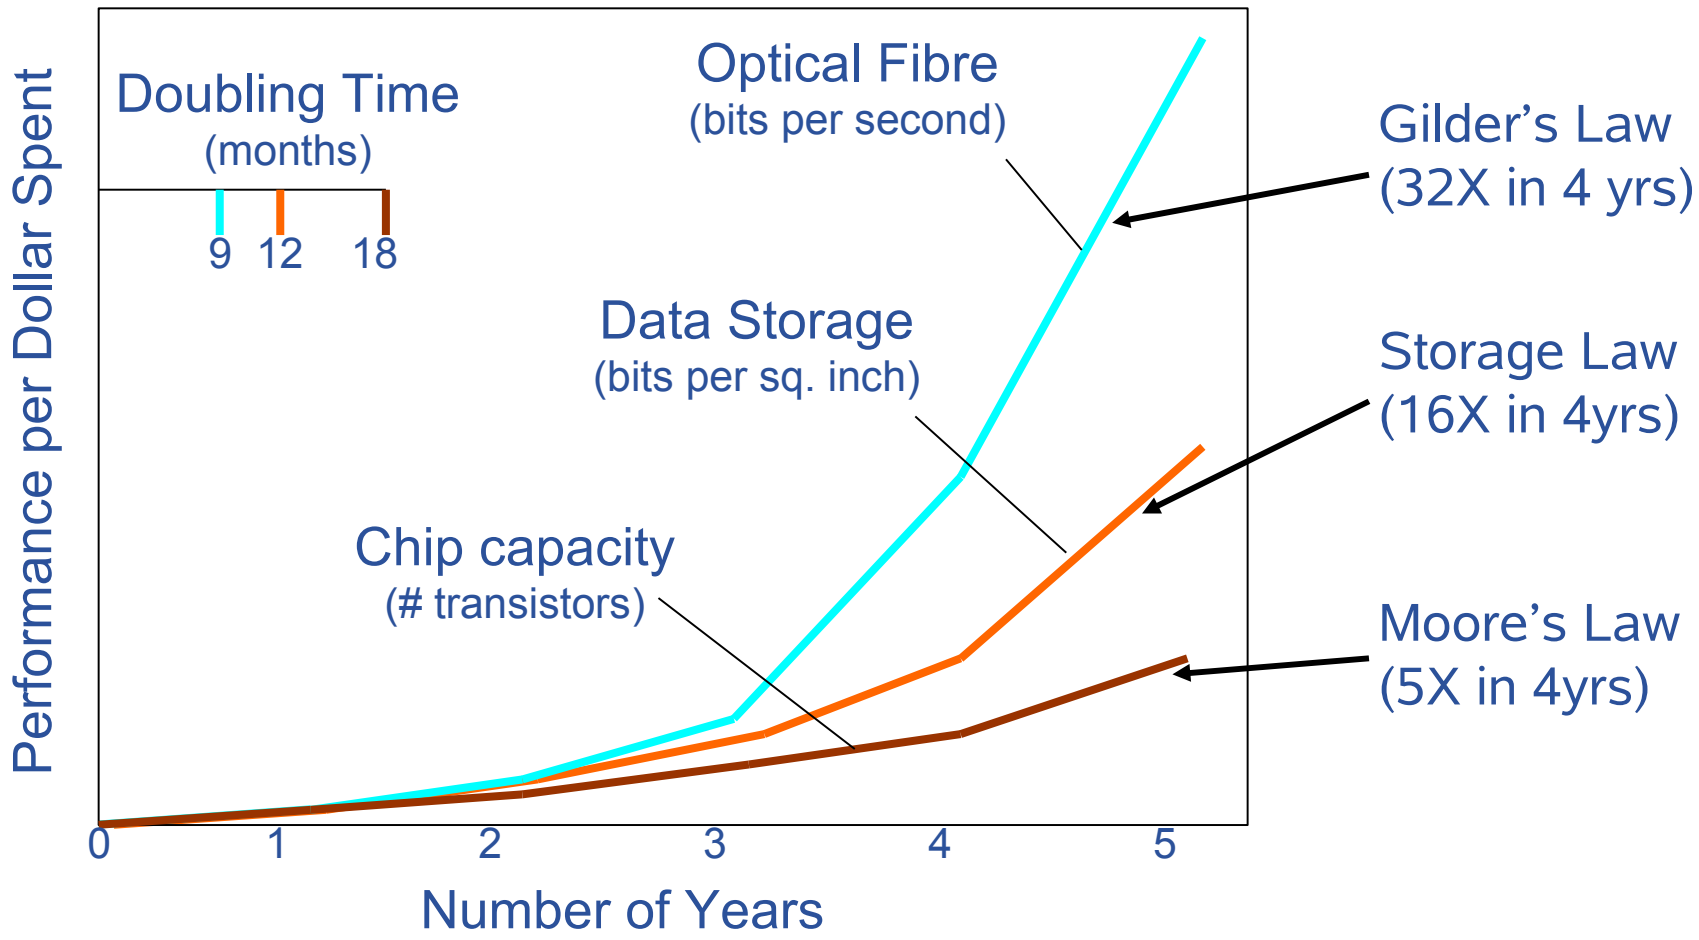

Triumph of Light – *Scientific American*. George Stix, January 2001

„When the network is as fast as the computer's internal links, the machine dis-integrates across the net into a set of special purpose appliances“

(Gilder Technology Report, June 2001)

Technical Limitations

- „Speed“ of a network consists of two components
- Bandwidth (scales to any number)
- Latency (doesn't scale)
- Possible application types in a distributed environment are limited by latency

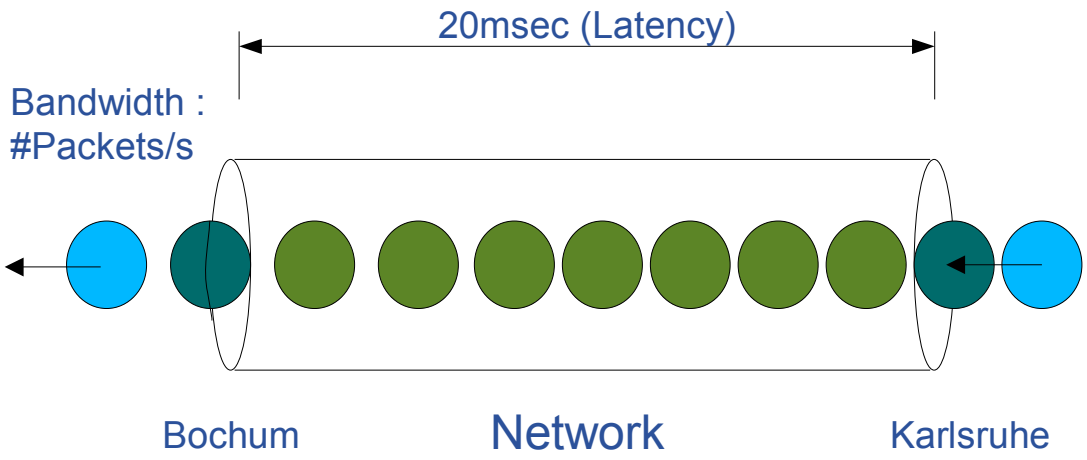

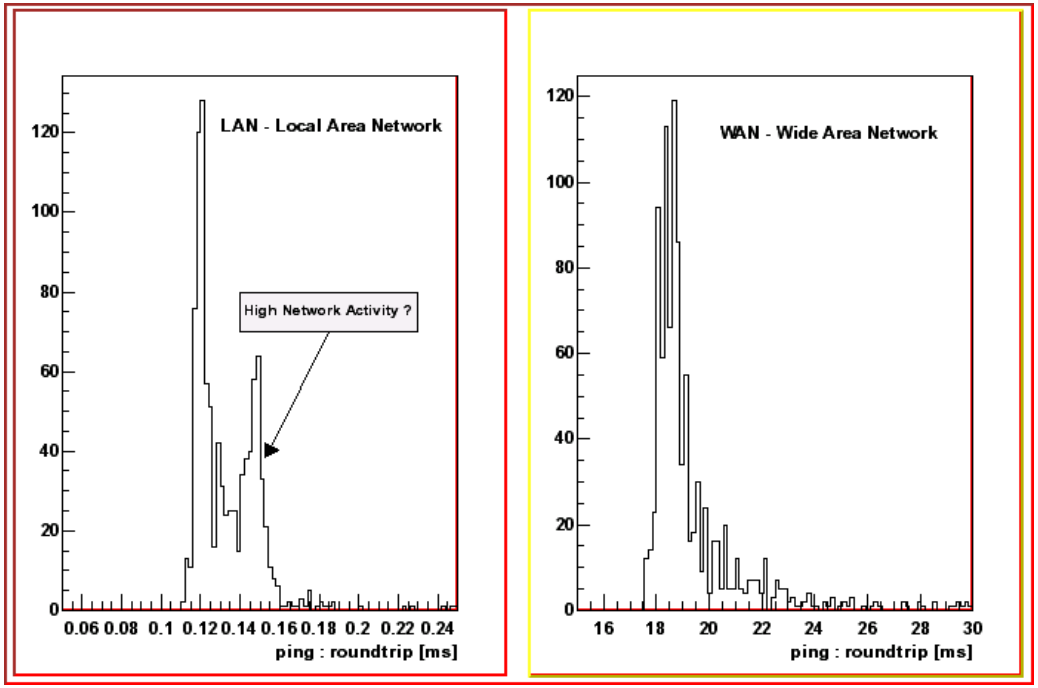

The application viewpoint

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

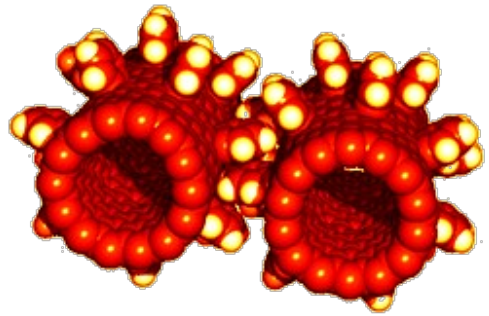

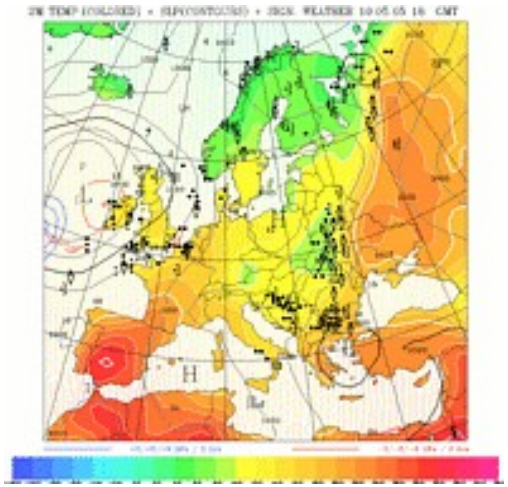

- **Experimental Science uses ever more sophisticated sensors to make precise measurements**
  - → Need high statistics
  - → Huge amounts of data
  - → Serves user communities around the world

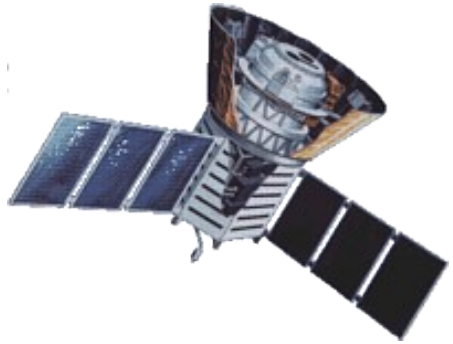

- **A good example: Particle Physics**

BaBar

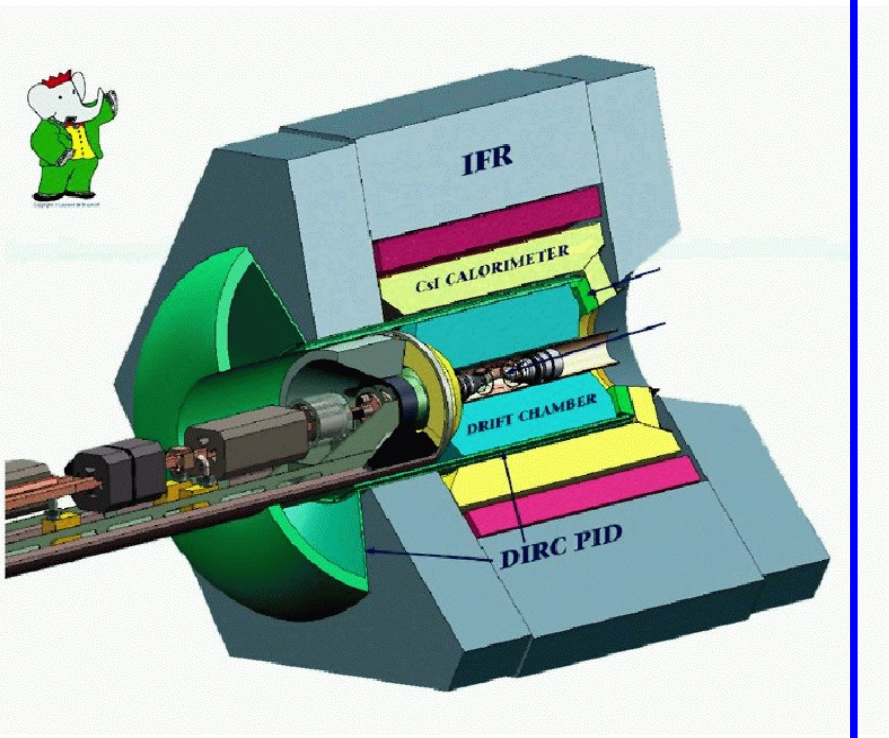

BaBar Detektor / SLAC, Stanford / USA

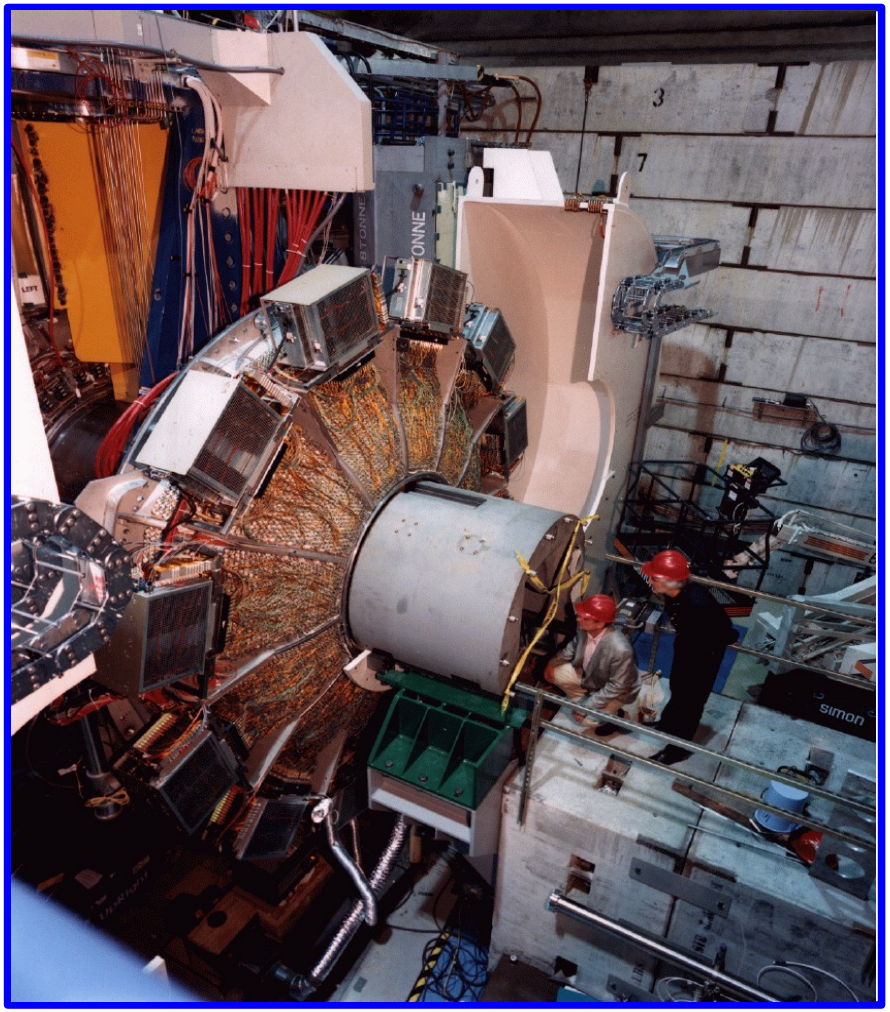

BaBar

Particle Physics / BaBar

- Collision / Annihilation of electrons and positrons
- $E=mc^2$: Energy is transformed to matter. Matter can be detected.
- One Collision ~ 1 Event
- 1 Event ~ KByte – MByte
- Several thousand events/ s (but: „trigger“)
- Need to store and analyse all data

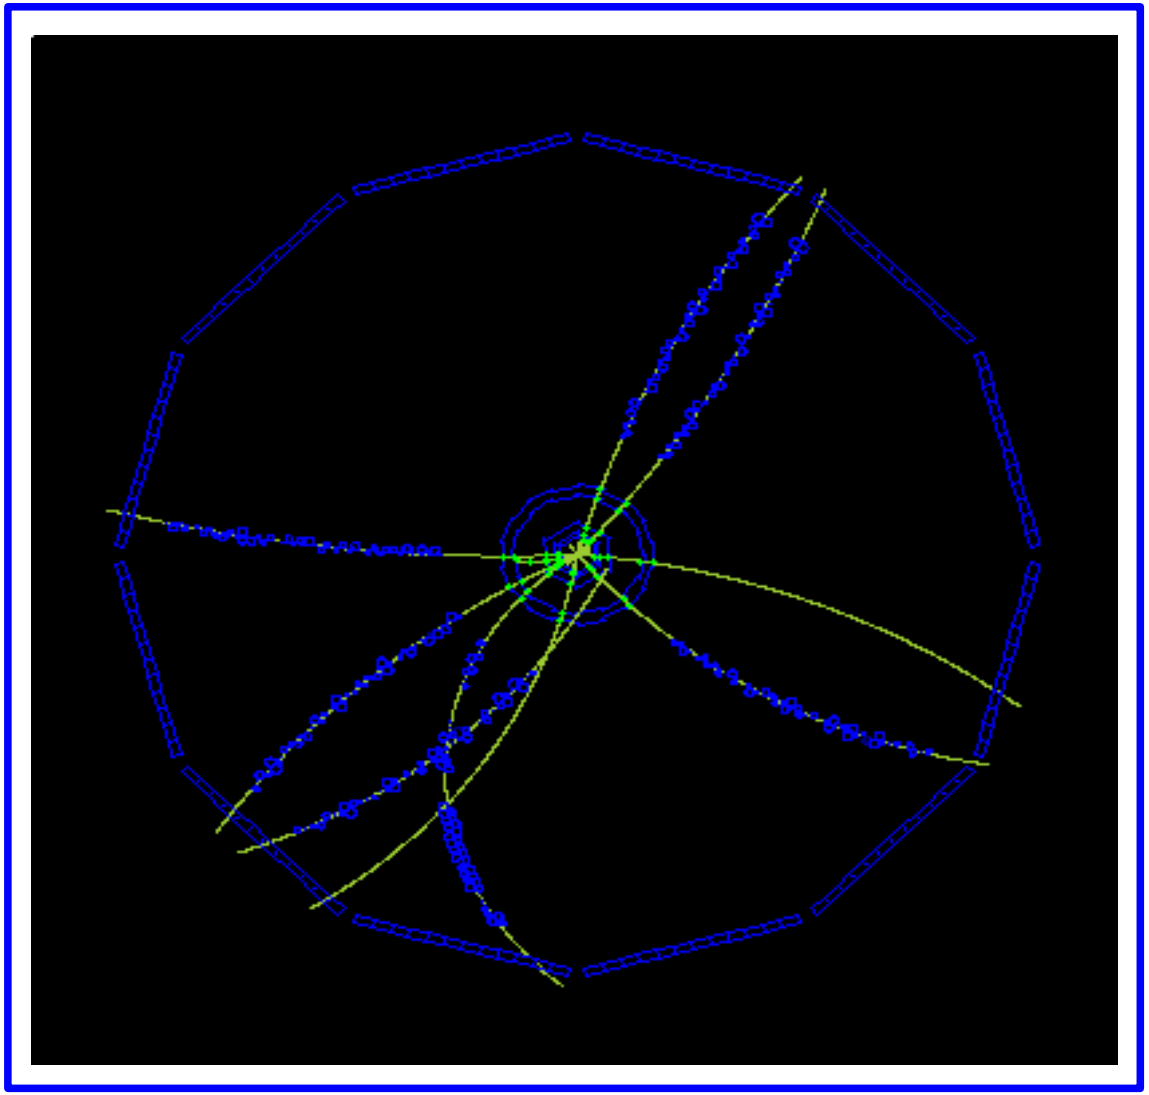

The Large Hadron Collider

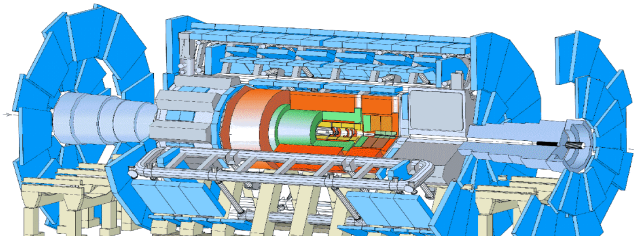

Atlas

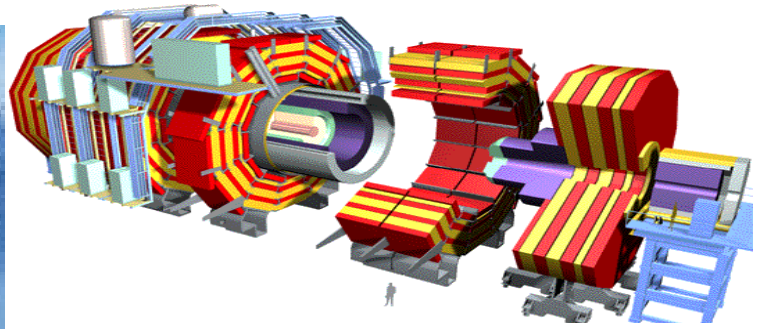

CMS

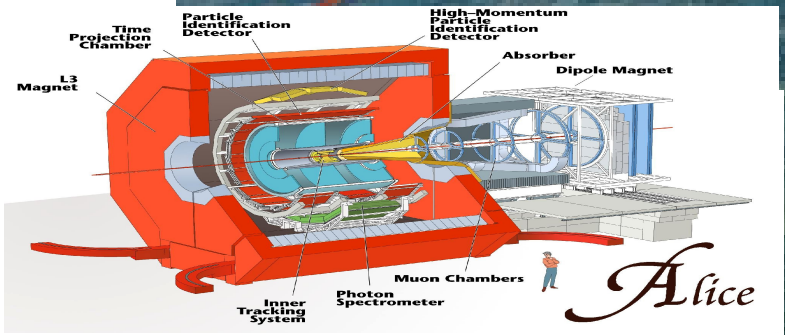

LHC

LHC*b*

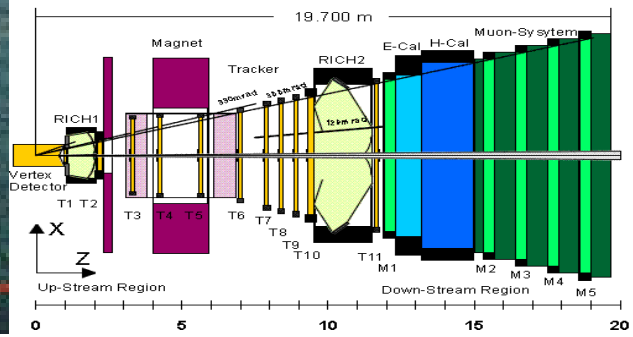

LHC / CMS

In LHC:
Expect data rates of
10 - 40 Petabytes
for all experiments
per year -
after triggering ...

But: trivial to run in parallel ...

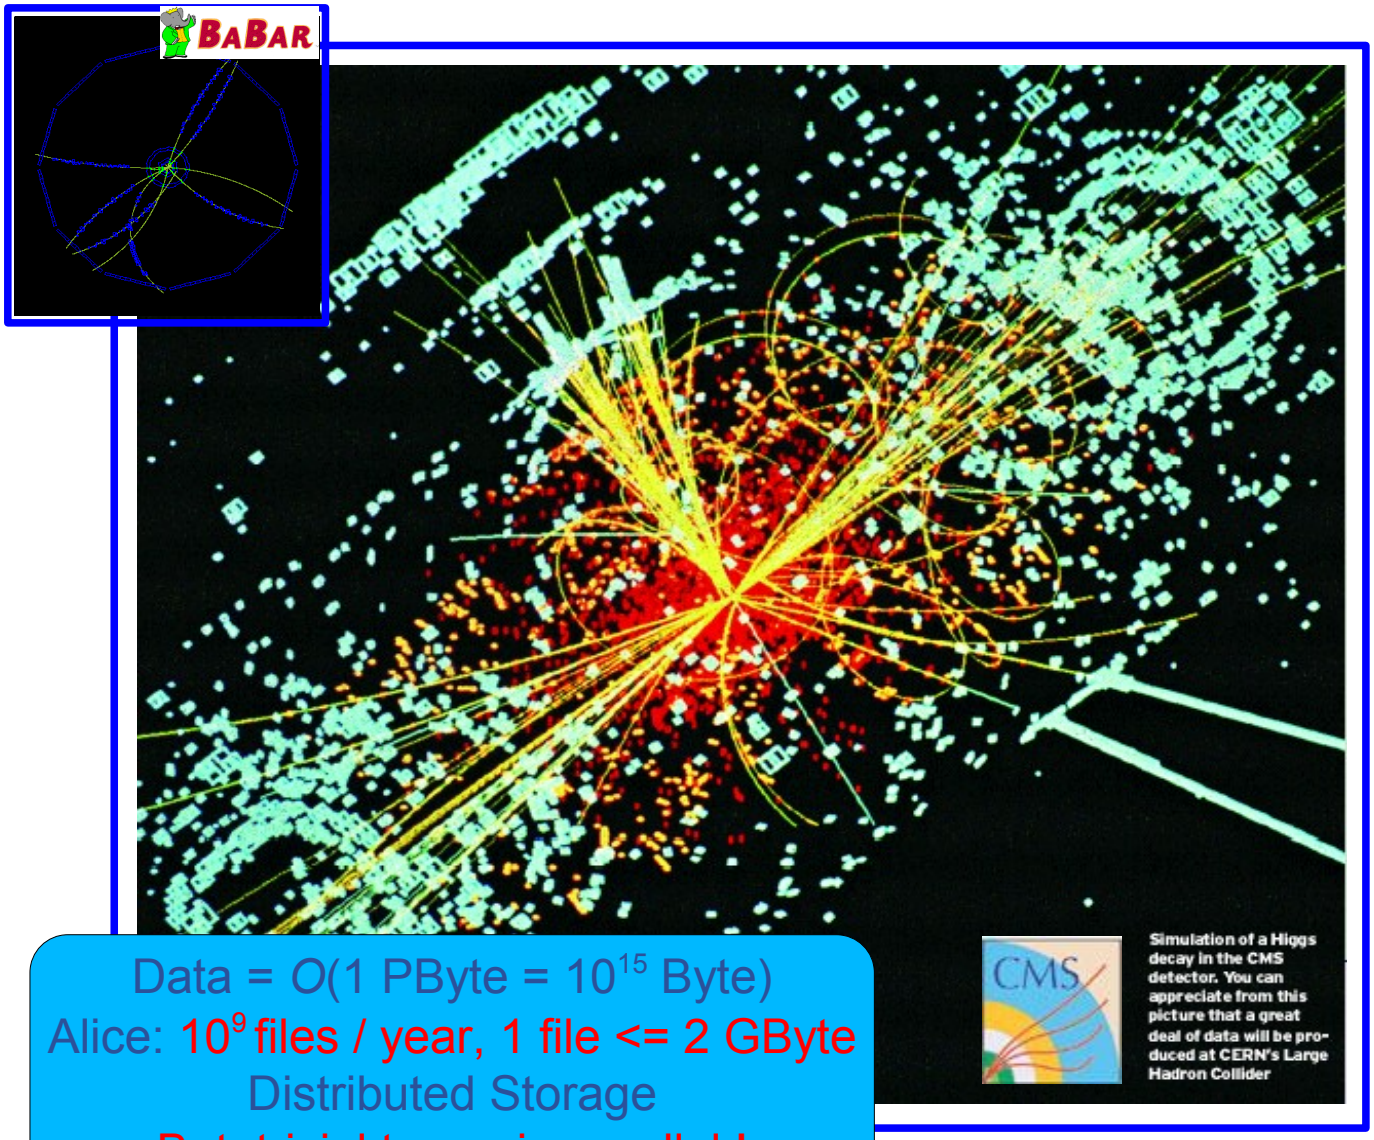

Data = O(1 PByte = 10^{15} Byte)
Alice: 10^9 files / year, 1 file <= 2 GByte
Distributed Storage
But: trivial to run in parallel !

The needle and the haystacks

Starting from this event

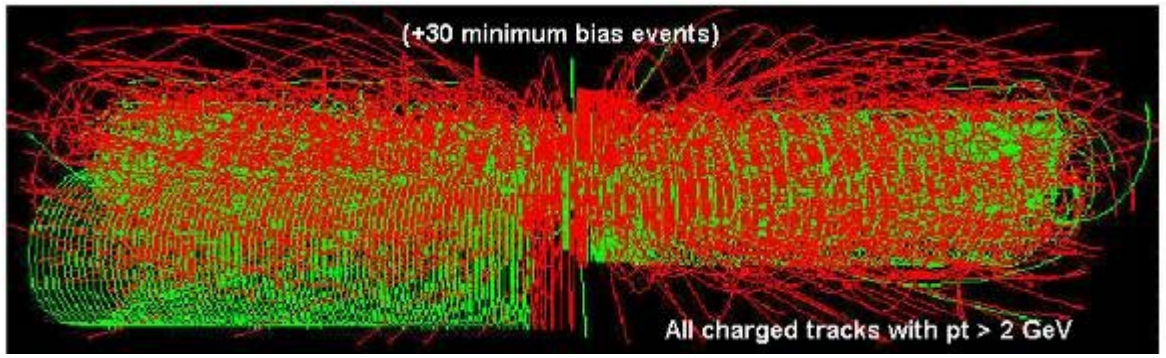

Looking for this "signature"

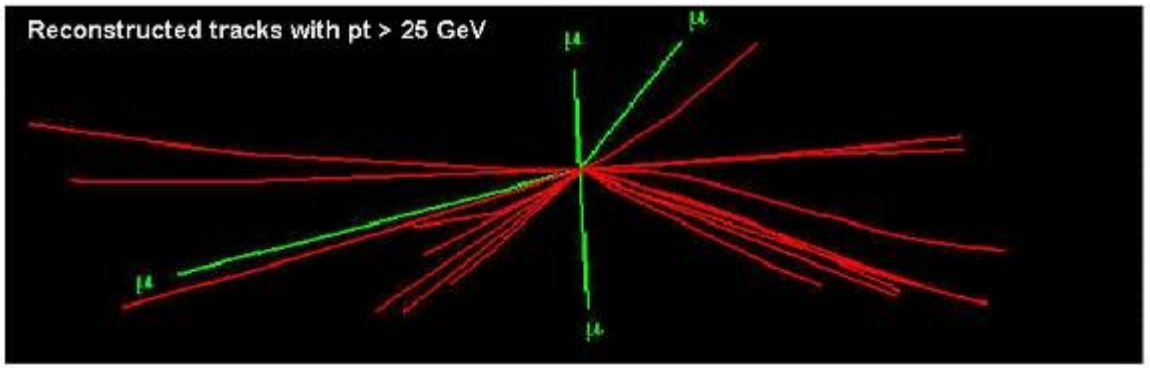

→ **Selectivity: 1 in 10^{13}**

(Like looking for a needle in 20 million haystacks)

Over 6000 LHC Scientists world wide

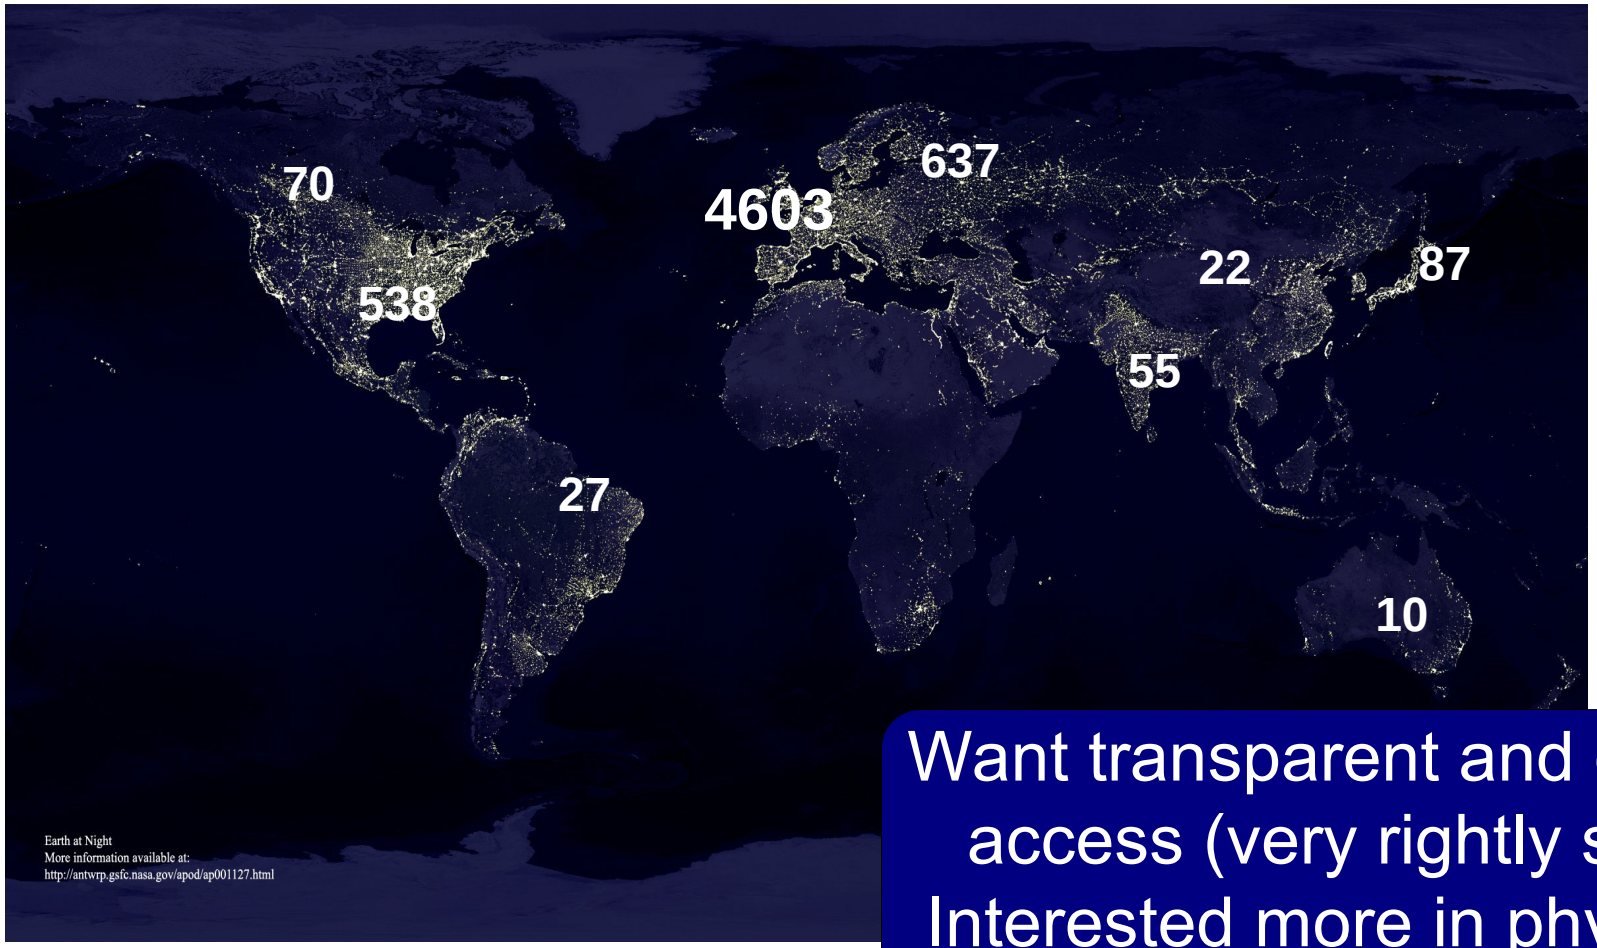

Europe: 267 Institutes, 4603 Users
Other: 208 Institutes, 1632 Users

Want transparent and quick access (very rightly so). Interested more in physics results, than computing revolutions

LHC / CMS

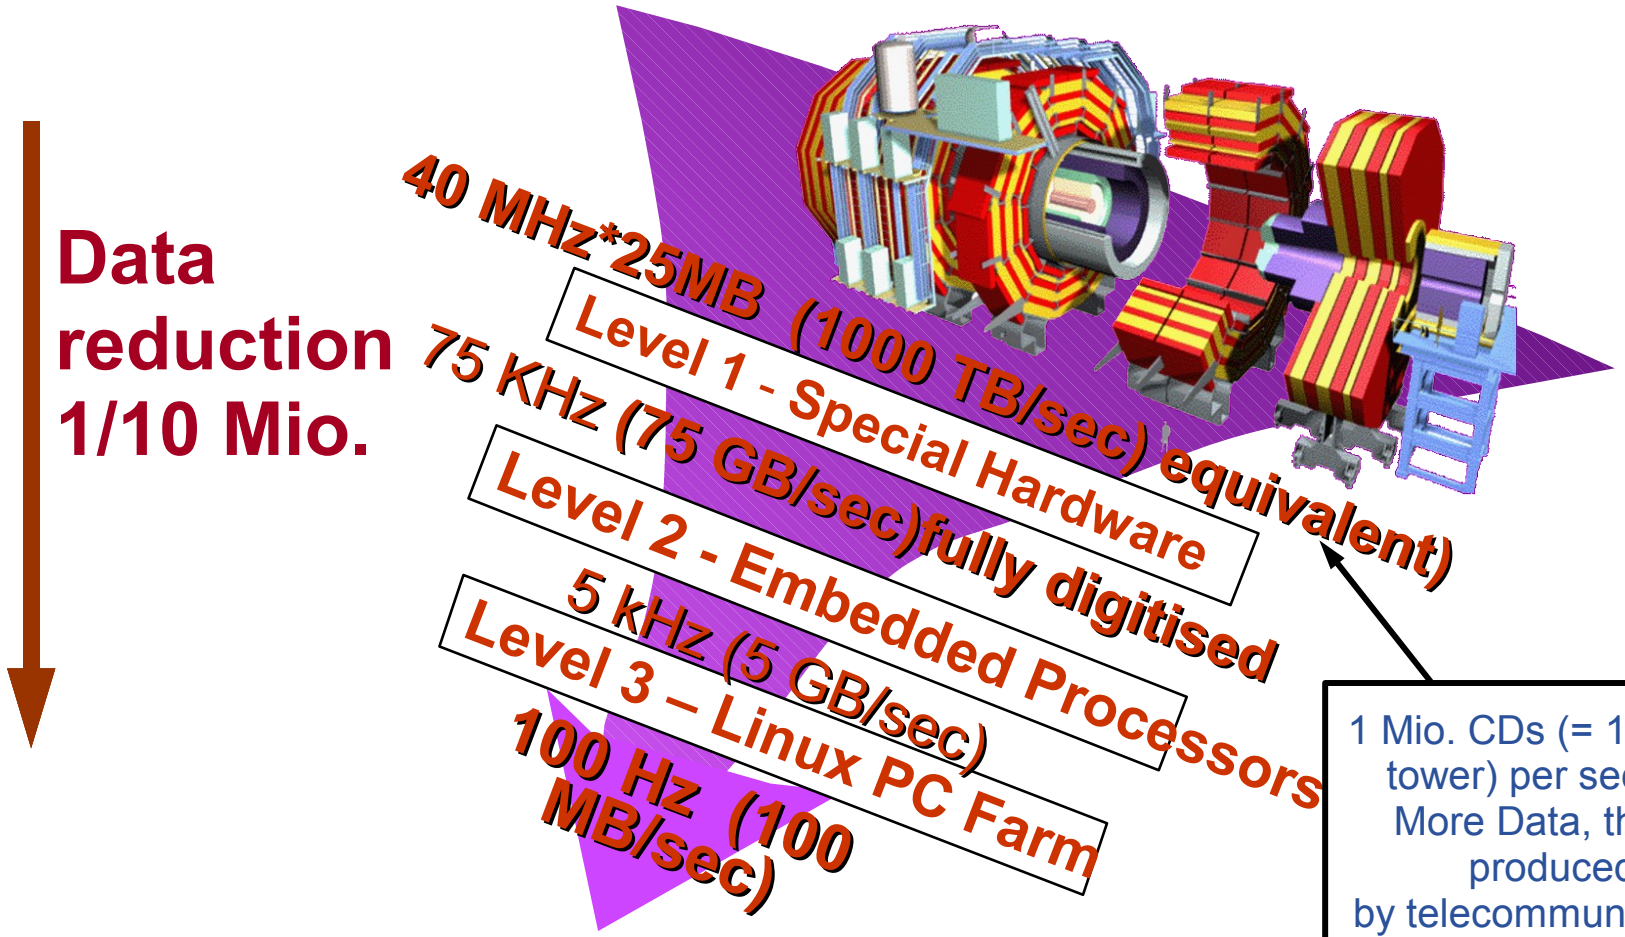

1 PB of data per year and experiment

... and 6000 physicist that want to access it !

The MONARC* study and Tier-1 centers

- Basic idea: hierarchical distribution of tasks
- Idea accepted by the LHC Computing Grid (responsible for planning and management of LHC computing)
- Tier-0: Initial reconstruction and storage of raw eve distribution to Tier-
- Tier-1: Data-heavy analysis, reprocessing of data, regional support
- Tier-2: Managed disk storage, simulation of PP events, computing

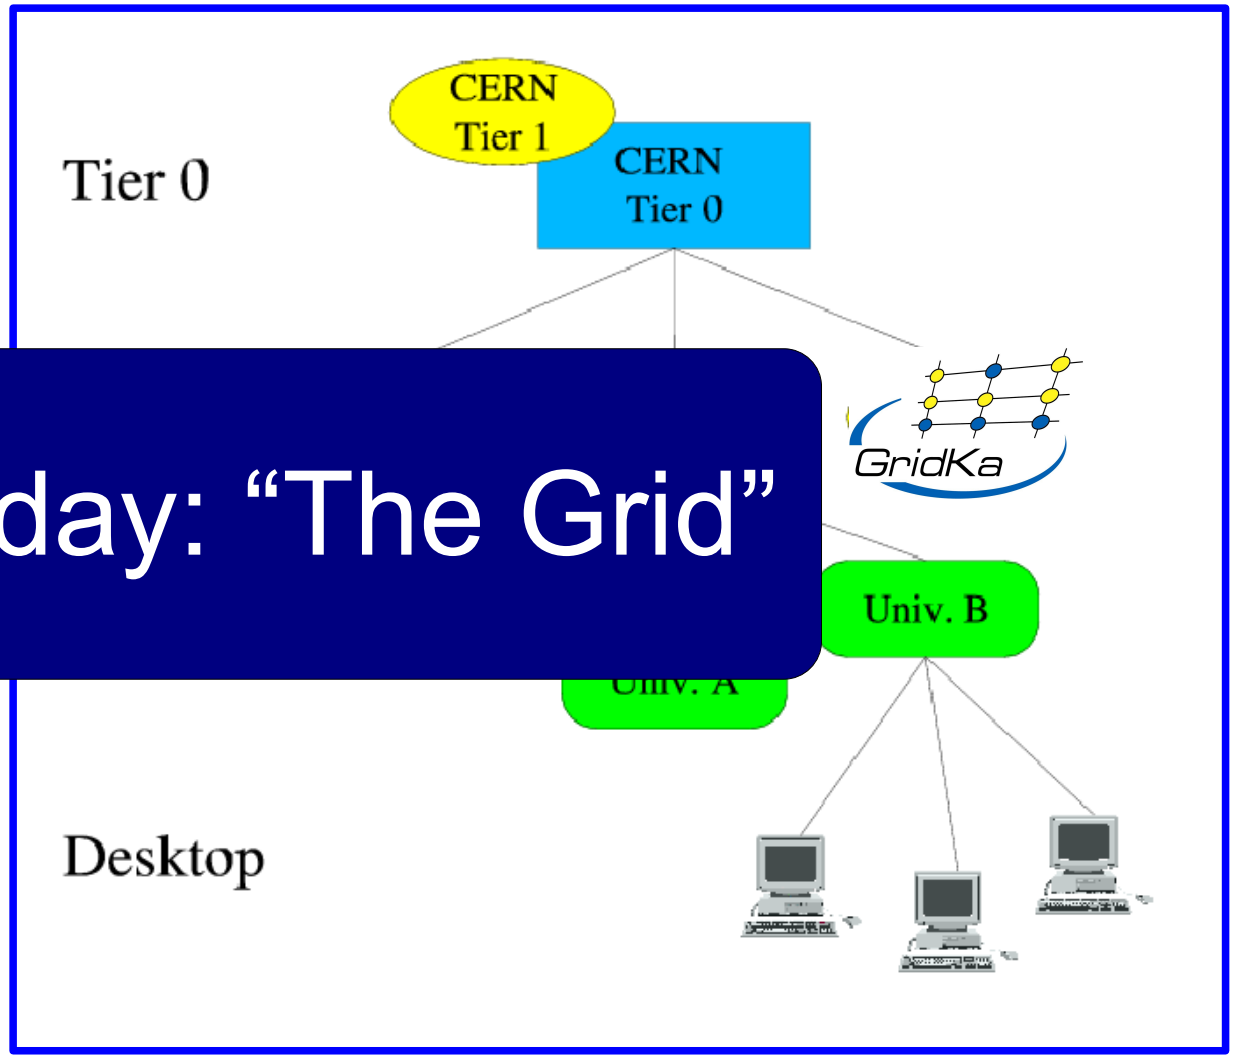

Today: "The Grid"

* **MONARC** == **M**odels **O**f **N**etworked **A**nalysis at **R**egional **C**enters

The GridKa Cluster: A Tier-1 centre in LCG

ca. 1500 CPUs
(approx. 4500 in 2007)

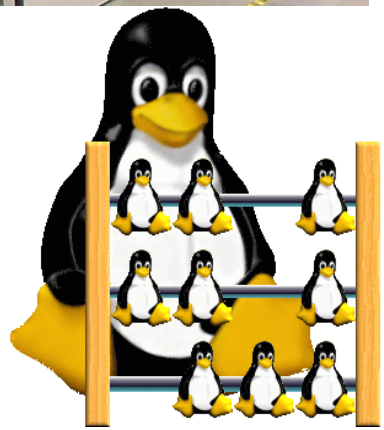

Scientific Linux

GridKa

The GridKa Cluster in the Mainstream

Practical definition in some areas

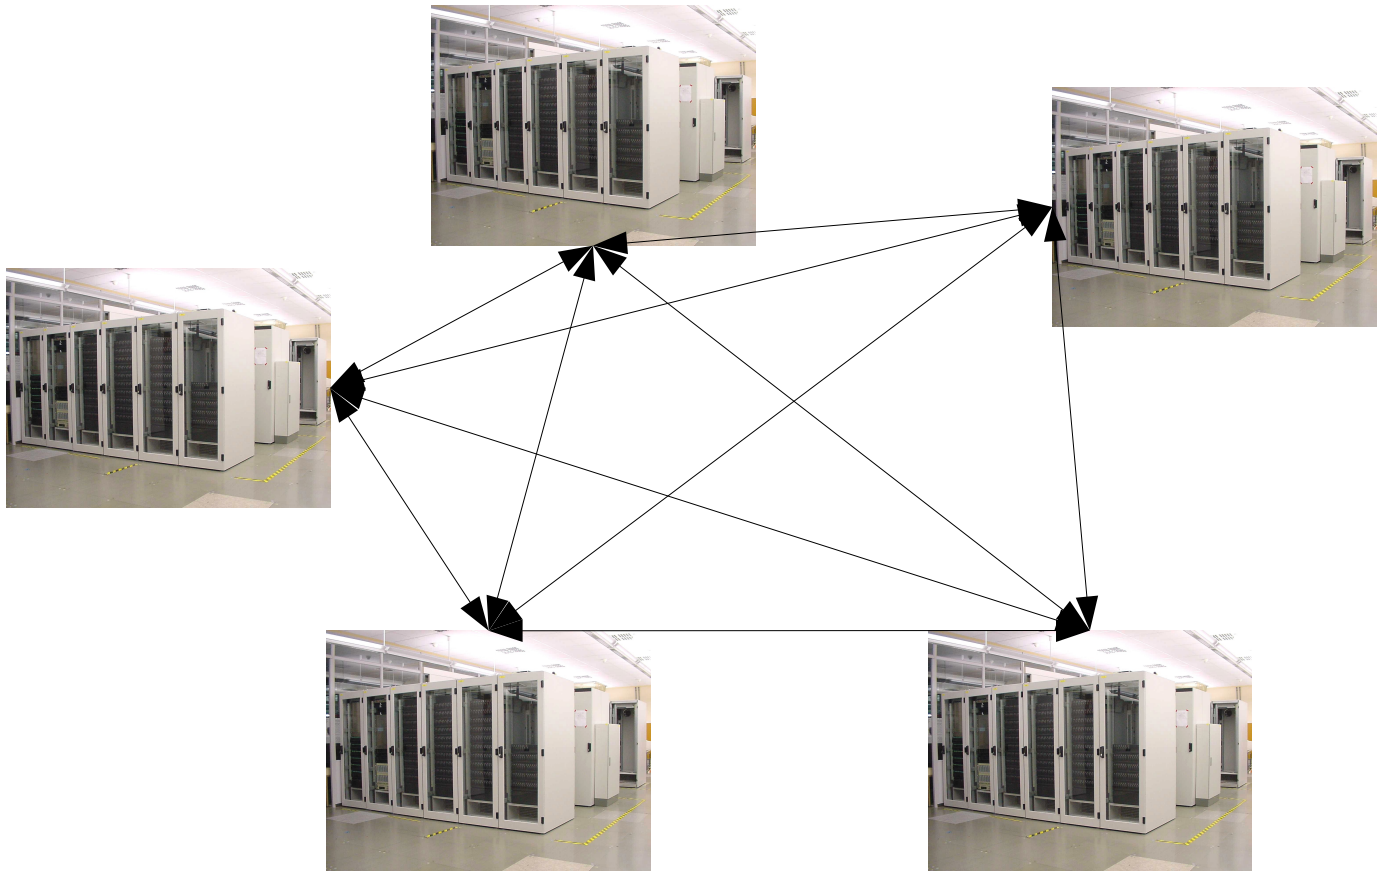

Grid Computing == Clustering Clusters; Building a global batch submission system.

(Please note that this definition neglects many advanced aspects of Grid Computing and does *not* work in a number of application areas)

The collaboration viewpoint

e-Science and the Grid

- **What is e-Science? Collaborative science that is made possible by the sharing across the Internet of resources (data, instruments, computation, people's expertise...)**
  - Often very compute intensive
  - Often very data intensive (both creating new data and accessing very large data collections) – data deluges from new technologies
  - Always: crosses administrative boundaries
- **More precisely "e-Research": arts, humanities, social science,… are engaging with distributed environments.**
- **A Grid can orchestrate these resources**
- **e-Science is a primary driver for Grid computing**

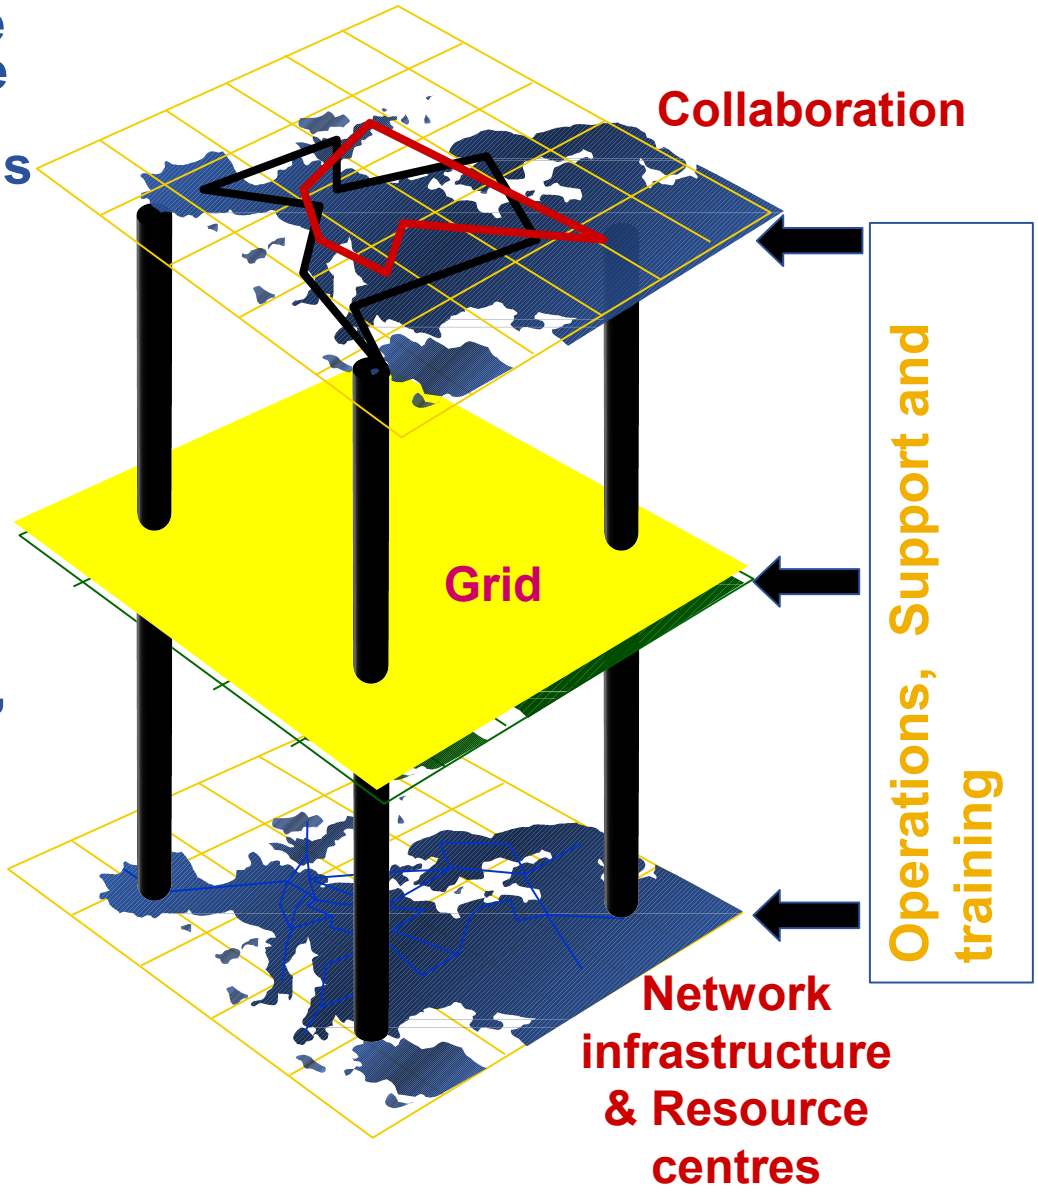

Defining Grid Computing

Genealogy: The term "Grid Computing" comes from the analogy to the electrical power grid - "computing power from a plug in the wall"

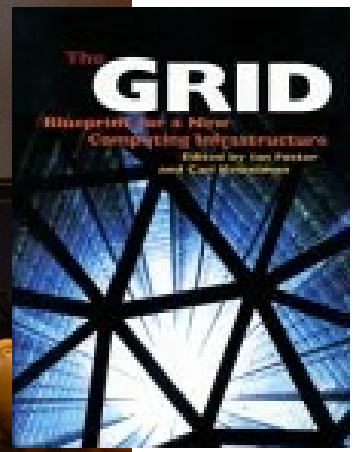

Definition, Ian Foster

> *"Grid computing is coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations"* (I.Foster)

A *Virtual Organisation* is:

- **People from different institutions working to solve a common goal**
- **Sharing distributed processing and data resources**
- **Technically not too different from Unix rights management (access control)**

Focus: Wide area, collaboration, virtual organisations

From an Oracle talk (September 2005)

"What is Grid Computing ?
In basic terms, grids are clusters of interconnected servers, enabling shared computing resources utilization"

(originally from "Defining Grid Computing",
Giga Research, August 2002)

Focus: Clustering ?

Definition, Sun

"[A grid] is a collection of computing resources that perform tasks. It appears to users as a large system, providing a single point of access to powerful distributed resources. Users treat the grid as a single computational resource. Resource management software, such as Sun Grid Engine, accepts jobs submitted by users and schedules them for execution on appropriate systems in the grid based on resource management policies. Users can literally submit thousands of jobs at a time without being concerned where they run." -- Grid Engine 5.3 Administration and User's Guide, p 23

found on http://sysadmin.cs.caltech.edu/docs/help/software/gridengine_help

Focus: distributed computing, but Sun Grid Engine comparable to PBS, LSF ?

Grid computing enables the virtualization of distributed computing and data resources such as processing, network bandwidth and storage capacity to create a single system image, granting users and applications seamless access to vast IT capabilities. Just as an Internet user views a unified instance of content via the Web, a grid user essentially sees a single, large virtual computer.

taken from http://www-1.ibm.com/grid/about_grid/what_is.shtml

Focus: Virtualisation, distributed computing, comparison with the Internet and the Web

Definition, HP

"Grid is a software environment based on open standards and protocols that make it possible to share disparate, loosely coupled IT resources across organizations and geographies. IT resources are freed from their physical boundaries and offered as services. They can potentially include almost any IT component -- computer cycles, storage spaces, databases, applications, files, sensors or scientific instruments."

From http://h71028.www7.hp.com/enterprise/cache/125371-0-0-225-121.html

Focus: Wide area, Open Standards, Webservices ...

Xgrid

- Recycling of “spare” compute power
- similar to Condor ??
- automated configuration
- plug & play

http://www.apple.com/server/macosx/features/xgrid.html

Joining Forces: The EGEE Initiative

Grid Projects

http://www.cordis.lu/ist/grids/projects.htm

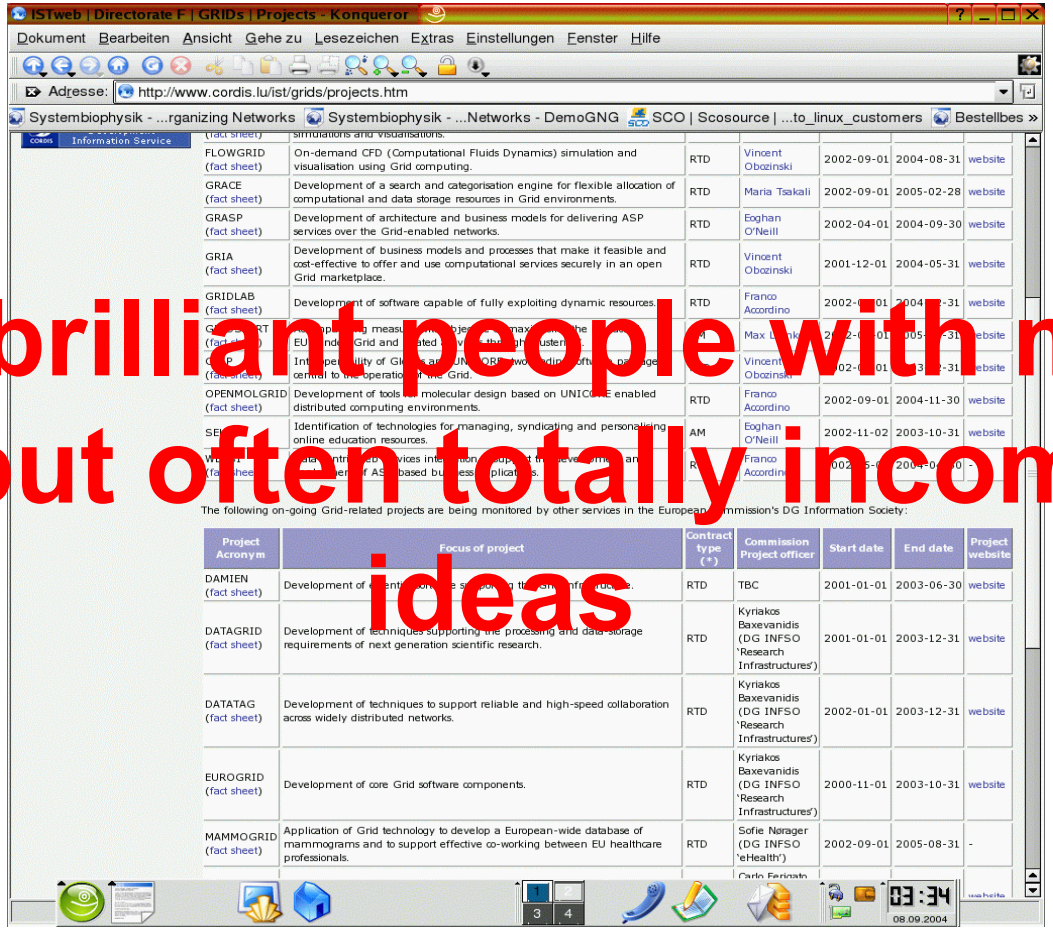

Many brilliant people with many brilliant (but often totally incompatible) ideas

The largest Grid Infrastructure: EGEE

- **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
  - Second project phase starts April 2006

EGEE ==
Enabling Grids
for E-SciencE

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)

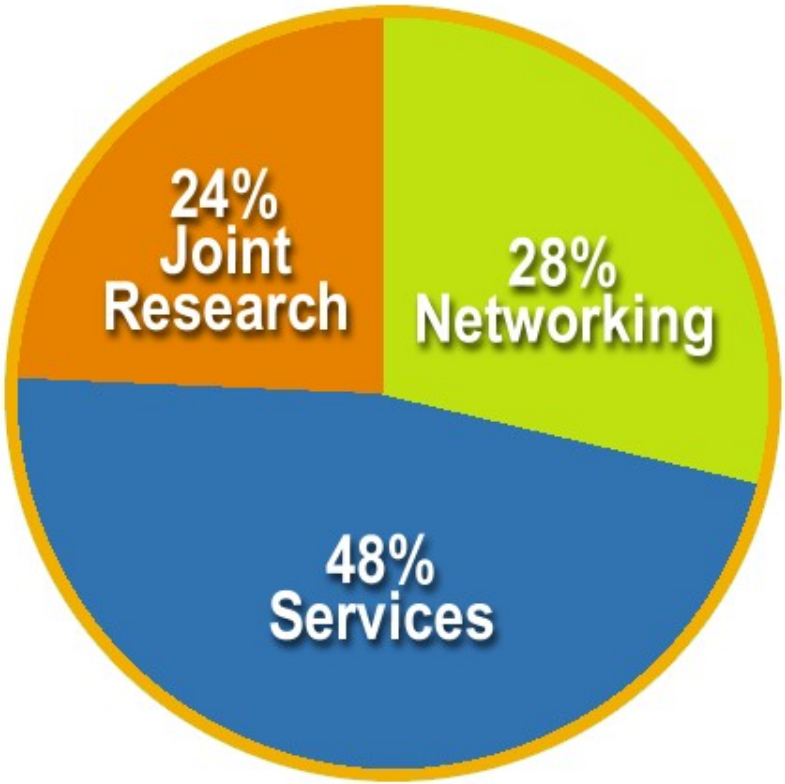

Emphasis in EGEE is on operating a production grid and supporting the end-users

Grid Operations

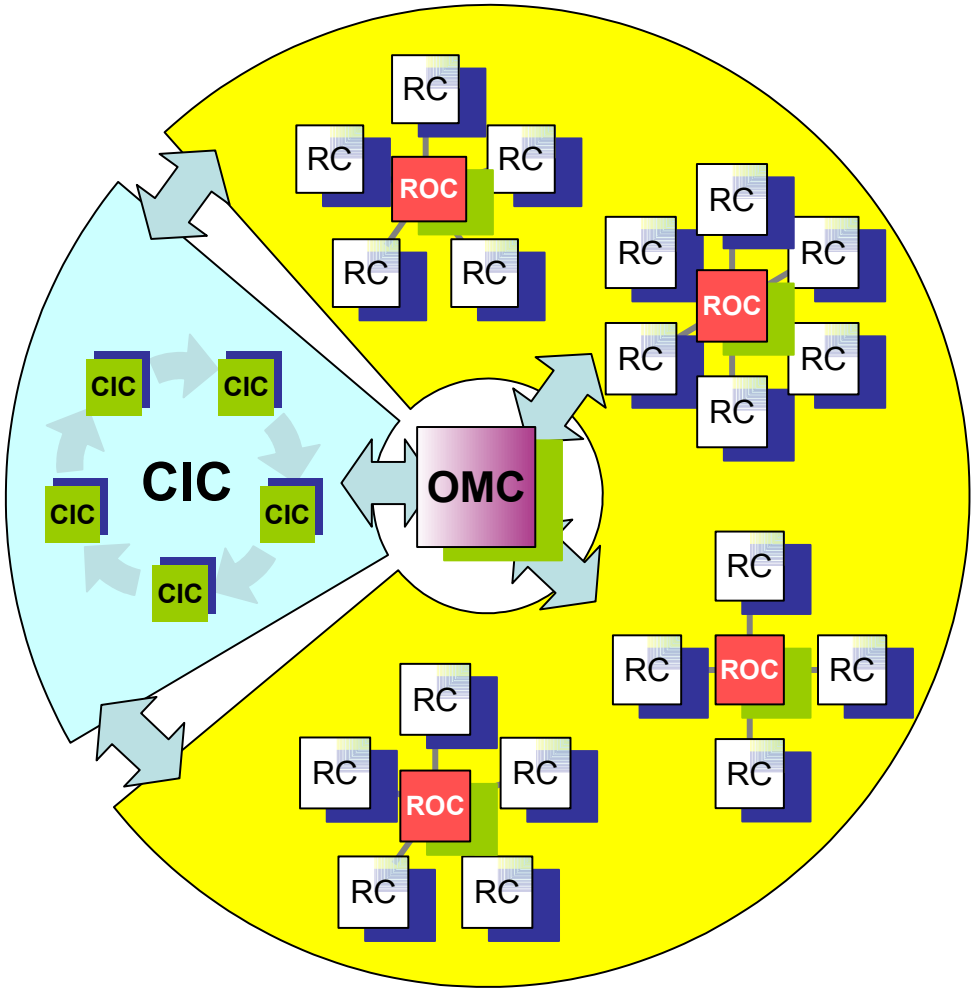

RC = Resource Centre
ROC = Regional Operations Centre
CIC = Core Infrastructure Centre
OMC = Operations Management Centre

- **The *grid* is flat, but**
- ***Hierarchy* of responsibility**
  - Essential to scale the operation
- **CICs act as a single Operations Centre**
  - Operational oversight ***(grid operator)*** responsibility
  - rotates weekly between CICs
  - Report problems to ROC/RC
  - ROC is *responsible* for ensuring problem is resolved
  - ROC oversees regional RCs
- **ROCs responsible for organising the operations in a region**
  - Coordinate deployment of middleware, etc
- **CERN coordinates sites not associated with a ROC**

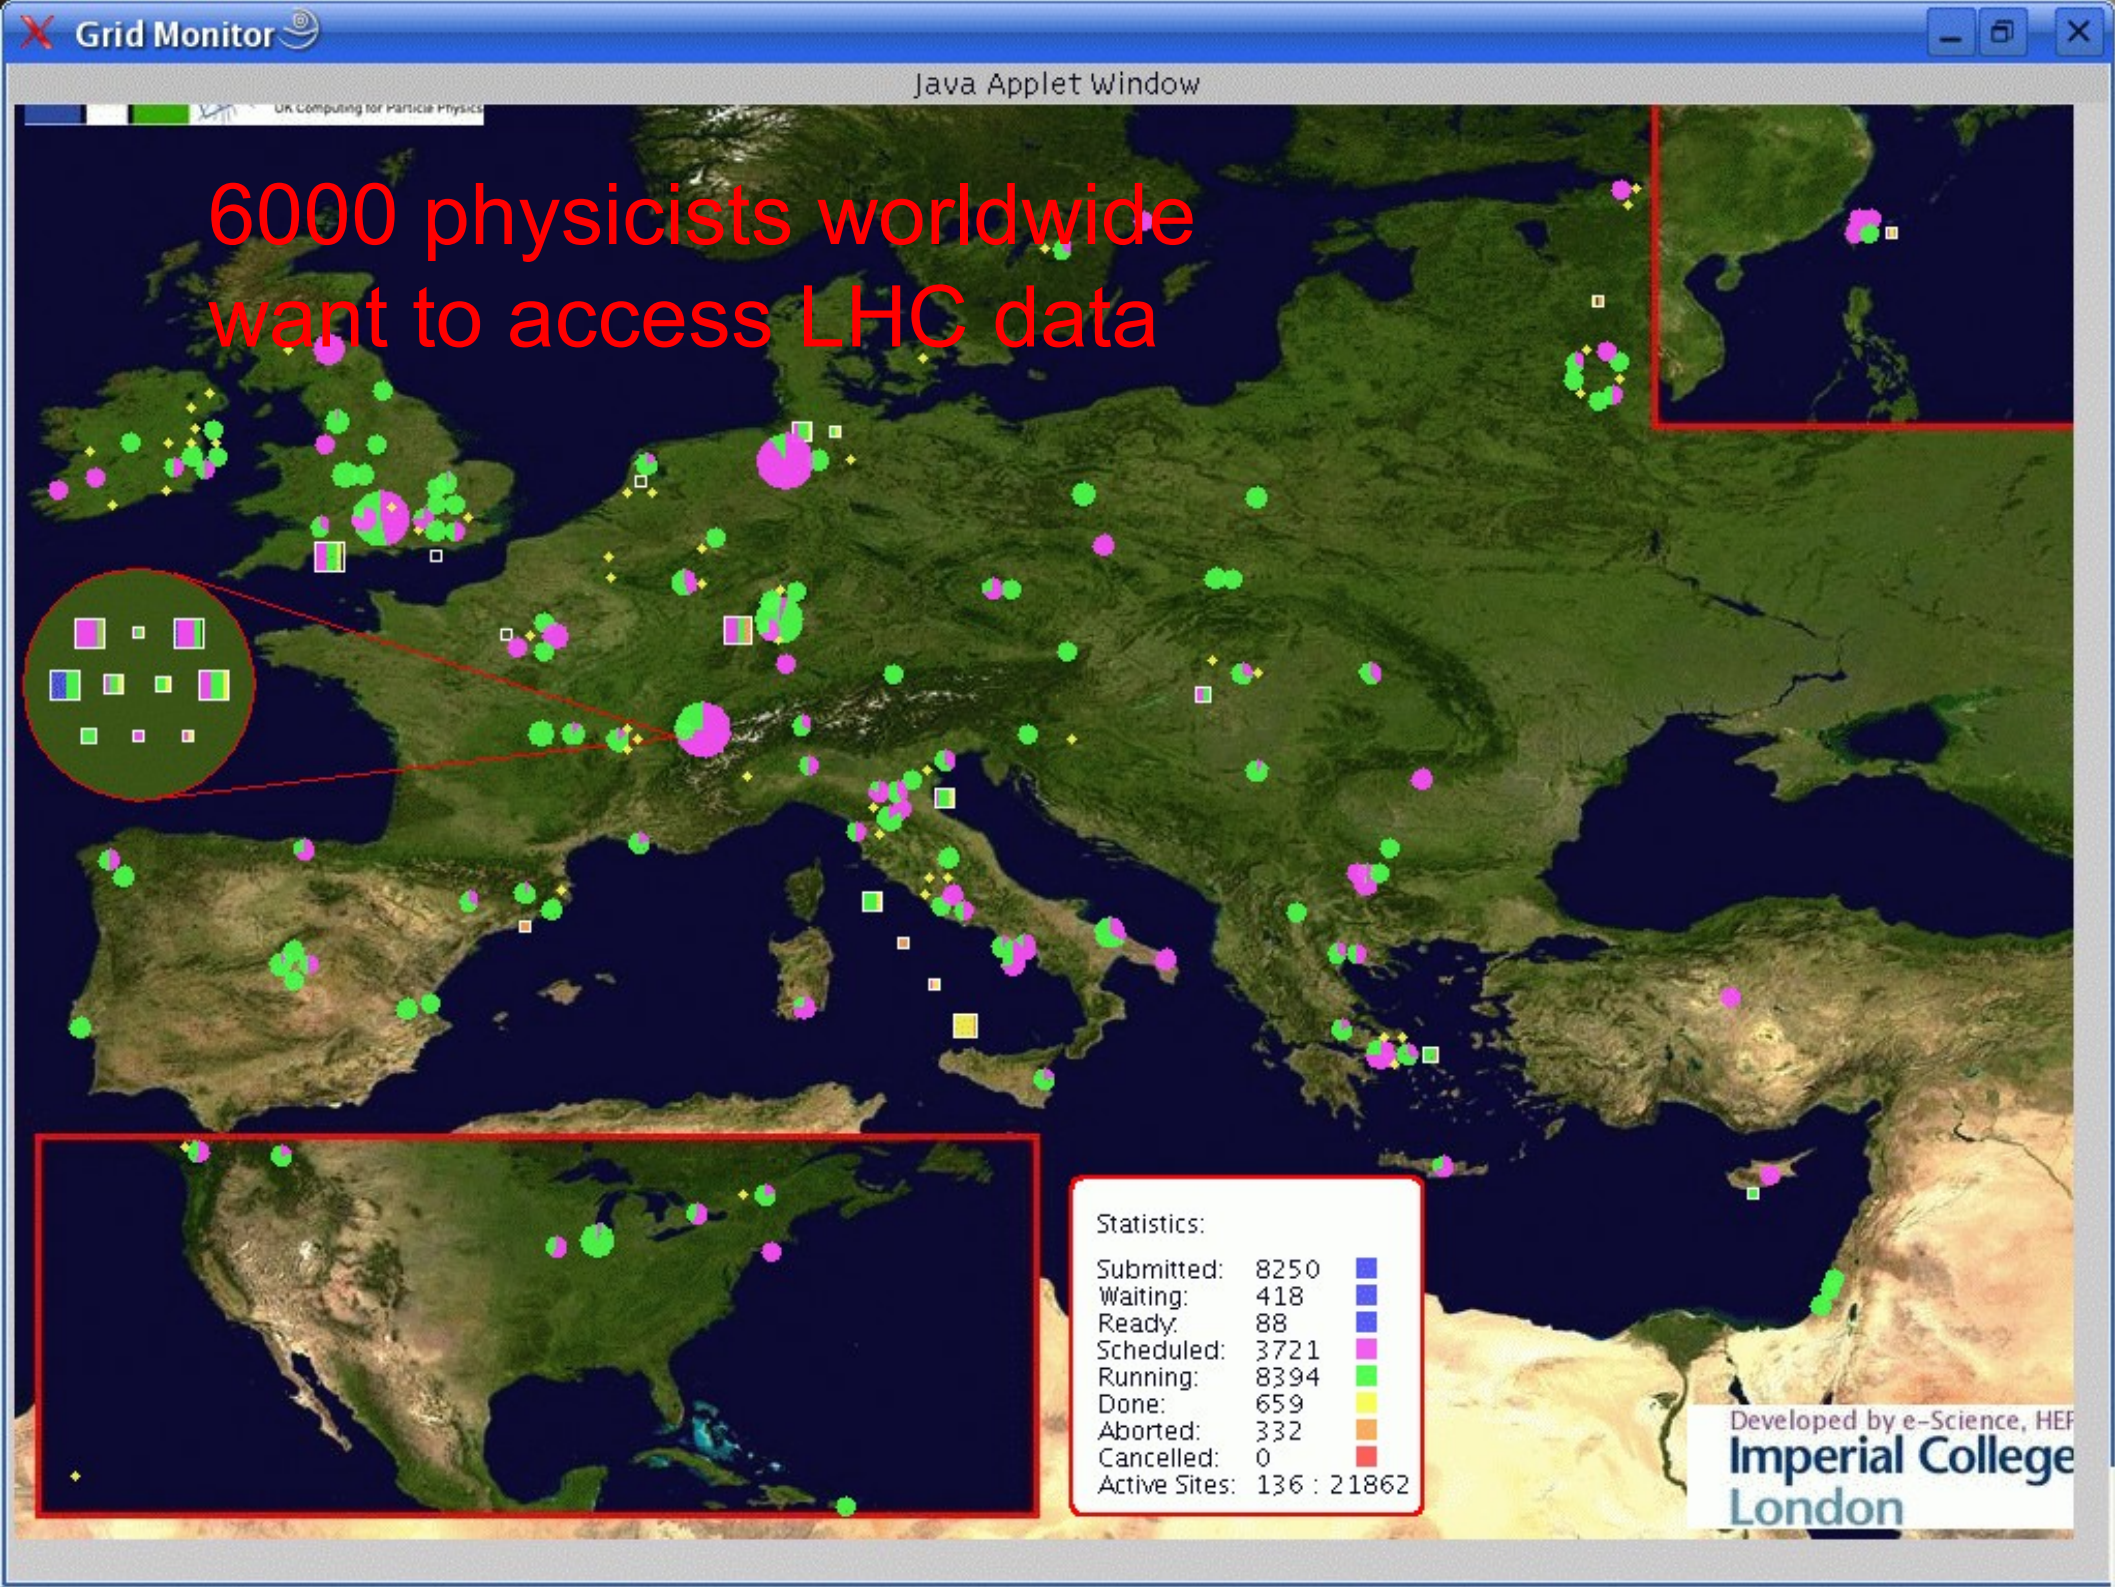
Grid Monitor
Java Applet Window
6000 physicists worldwide
want to access LHC data
Statistics:
Submitted: 8250
Waiting: 418
Ready: 88
Scheduled: 3721
Running: 8394
Done: 659
Aborted: 332
Cancelled: 0
Active Sites: 136 : 21862
Developed by e-Science, HEP
Imperial College
London

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 additional challenge:

    security & privacy
- **Chemistry, Earth Observation, Astronomy, Geophysics, ...**

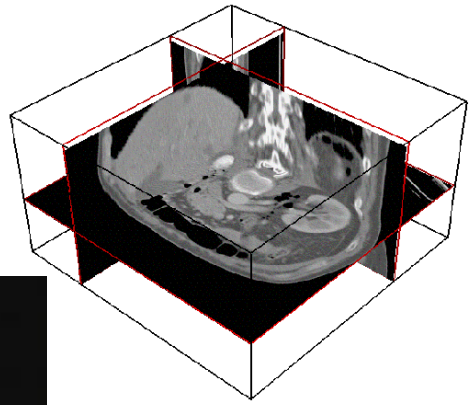

Grid middleware

- **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 scientis

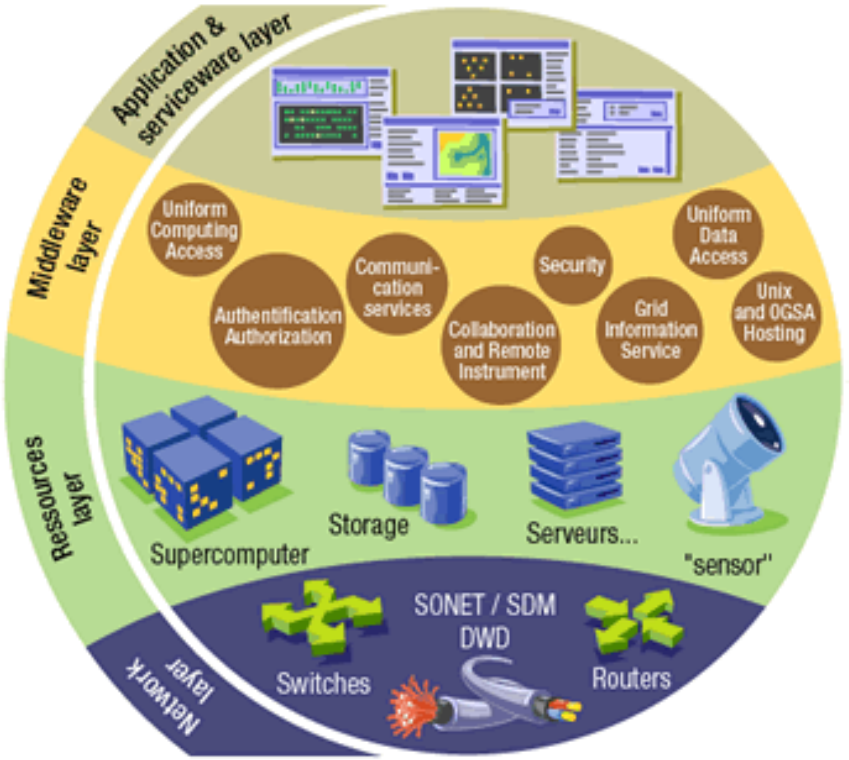

EGEE Middleware gLite

- First release of gLite end of March 2005
  - Release 1.1 in May 05
  - Release 1.2 in July 05
  - Release 1.5 in January 06
  - see www.gLite.org

- Interoperability & Co-existence with deployed infrastructure
- Robust: Performance & Fault Tolerance
- Service oriented approach
- Open source license
- Webservice focus
- Likely successor to LCG-2 / Merge

User information, training & support

- Over 2000 people trained
- Induction, application developer, advanced courses, retreats, ...
- Material archive online > 200 presentations
- Press work
- Grid summer schools
- Public and technical websites constantly evolving to expand information available and keep it up to date
- 4 conferences organized
  - ~ 300 @ Cork
  - ~ 400 @ Den Haag
  - ~ 450 @ Athens
  - ~ 450 @ Pisa (10/05)

EGEE operations: OMC and CIC

- **Operation Management Centre**
  - located at CERN, coordinates operations and management
  - coordinates with other grid projects
- **Core Infrastructure Centres**
  - behave as single organisations
  - operate core services (VO specific and general Grid services)
  - develop new management tools
  - provide support to the Regional Operations Centres

EGEE operations: ROC

- **Regional Operations Centre responsibilities and roles:**
  - Testing (certification) of new middleware on a variety of platforms before deployment
  - Deployment of middleware releases + coordination + distribution inside the region
  - integration of ‘Local’ VO
  - Development of procedures and capabilities to operate the resources
  - First-line user support (GGUS)
  - Bring new resources into the infrastructure and support their operation
  - Coordination of integration of national grid infrastructures Provide resources for pre-production service
  - Germany/Switzerland: Distributed ROC

- **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"**

Collaboration with Slowak Academy of Sciences

EGEE as partner

- Ongoing collaborations
  - with non-EU partners: US, Israel, Russia, Korea, Taiwan…
    - MoU with the Chonnam–Kangnung–Sejong–Collaboration project (CKSC)
    - Strong relationship KISTI (Korea Institute of Science and Technology Information), developing into partnership for EGEE II
  - with other European projects, in particular:
    - GÉANT
    - DEISA
    - SEE-GRID
  - with non-European projects:
    - OSG: OpenScienceGrid (USA)
    - NAREGI (Japan)
    - International Grid Trust Federation
      - *EU-GridPMA joining with Asia-Pacific and American counterparts*

- EGEE as incubator
  - 18 recently submitted EU proposals supported
  - More proposals in next calls and national funding programmes

From Phase I to II

- **From 1st EGEE EU Review in February 2005:**
  - "The reviewers found the overall performance of the project very good."
  - "… remarkable achievement to set up this consortium, to realize appropriate structures to provide the necessary leadership, and to cope with changing requirements."

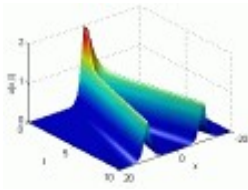

- **EGEE I**
  - Large scale deployment of EGEE infrastructure to deliver production level Grid services with selected number of applications

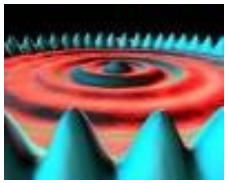

- **EGEE II**
  - approved !
  - Natural continuation of the project's first phase
  - Emphasis on providing an infrastructure for e-Science
    - → increased support for applications
    - → increased multidisciplinary Grid infrastructure
    - → more involvement from Industry
  - **Extending the Grid infrastructure world-wide**
    - → **increased international collaboration**
      (Asia-Pacific is already a partner!)

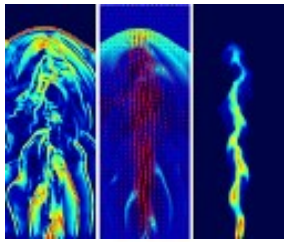

Grid Applications

Use Case in Science: 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

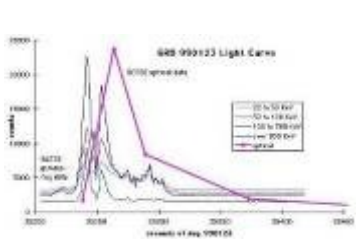

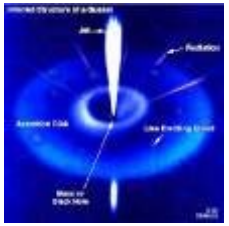

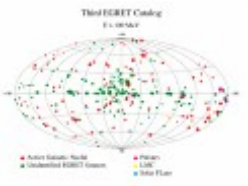

Use Case in Science: Bioinformatics

- **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

Bio-medicine applications

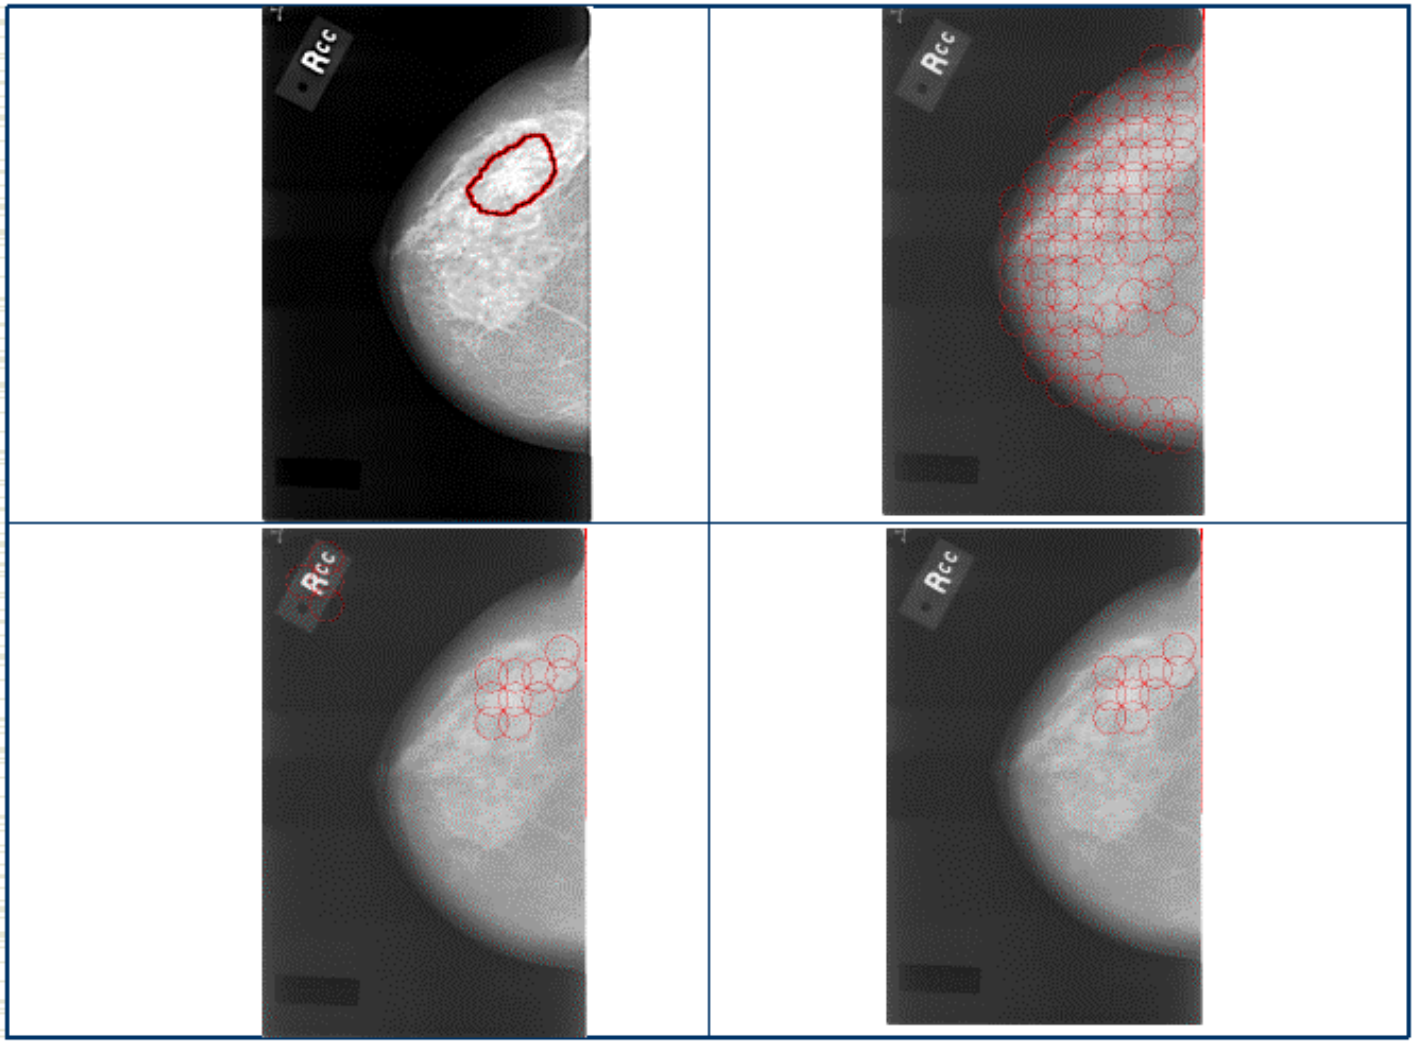

September 21st, 2001 *Datagrid Meeting, Lyon* 15

Mammogrid -> AliEn

Similarity search

Similarity computation

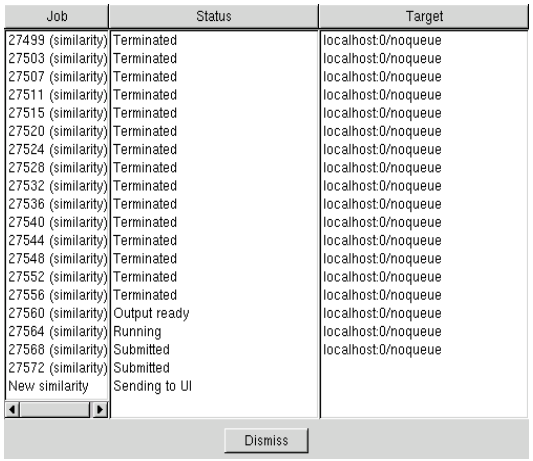

Job monitoring

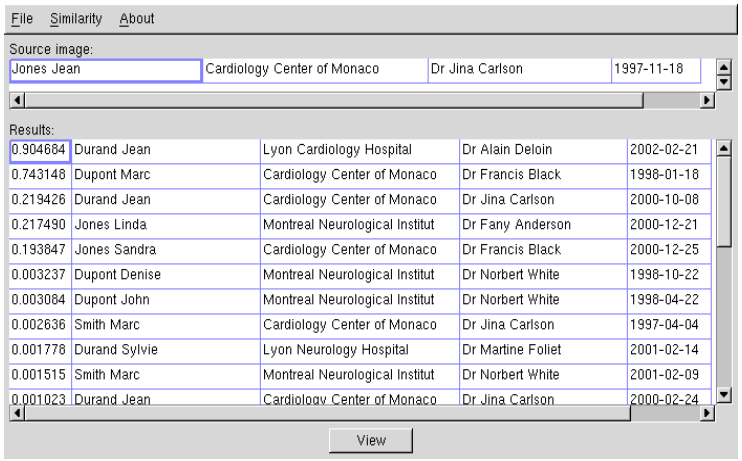

Ranked list of images

Results visualization

Source image

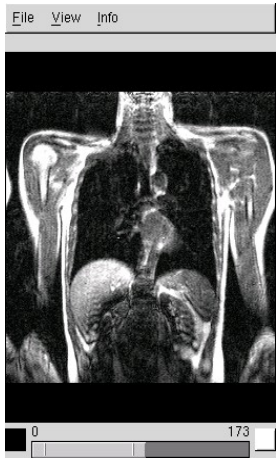

Most similar images

Low score images

3.3 Heart Modeling

- Objectives: modeling heart anatomy, dynamics and physiology for heart image processing
  - bio-mecanical model
  - electrical model
  - very complex structure
  - biological scale out of range

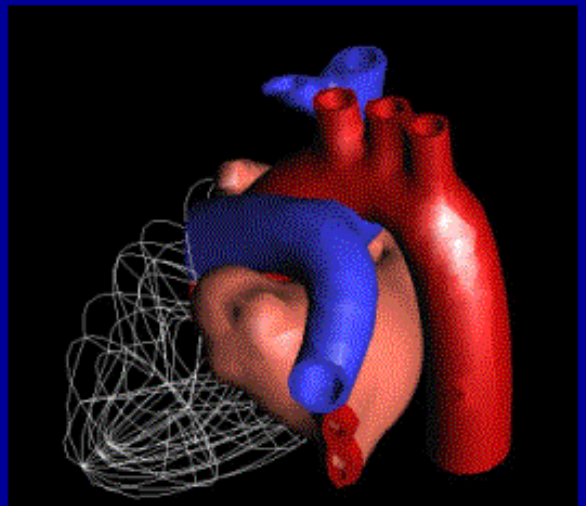

Bioengineering research group, Auckland

- Finite Element modeling
  - elements oriented in heart fibers direction: fine resolution
  - electrical propagation model based on bidomain theory
  - 4D model (3D+T)

Creatis

GGF1 - DataGrid WP 10 - March 2001

Earth observation applications

Virtual Observatory

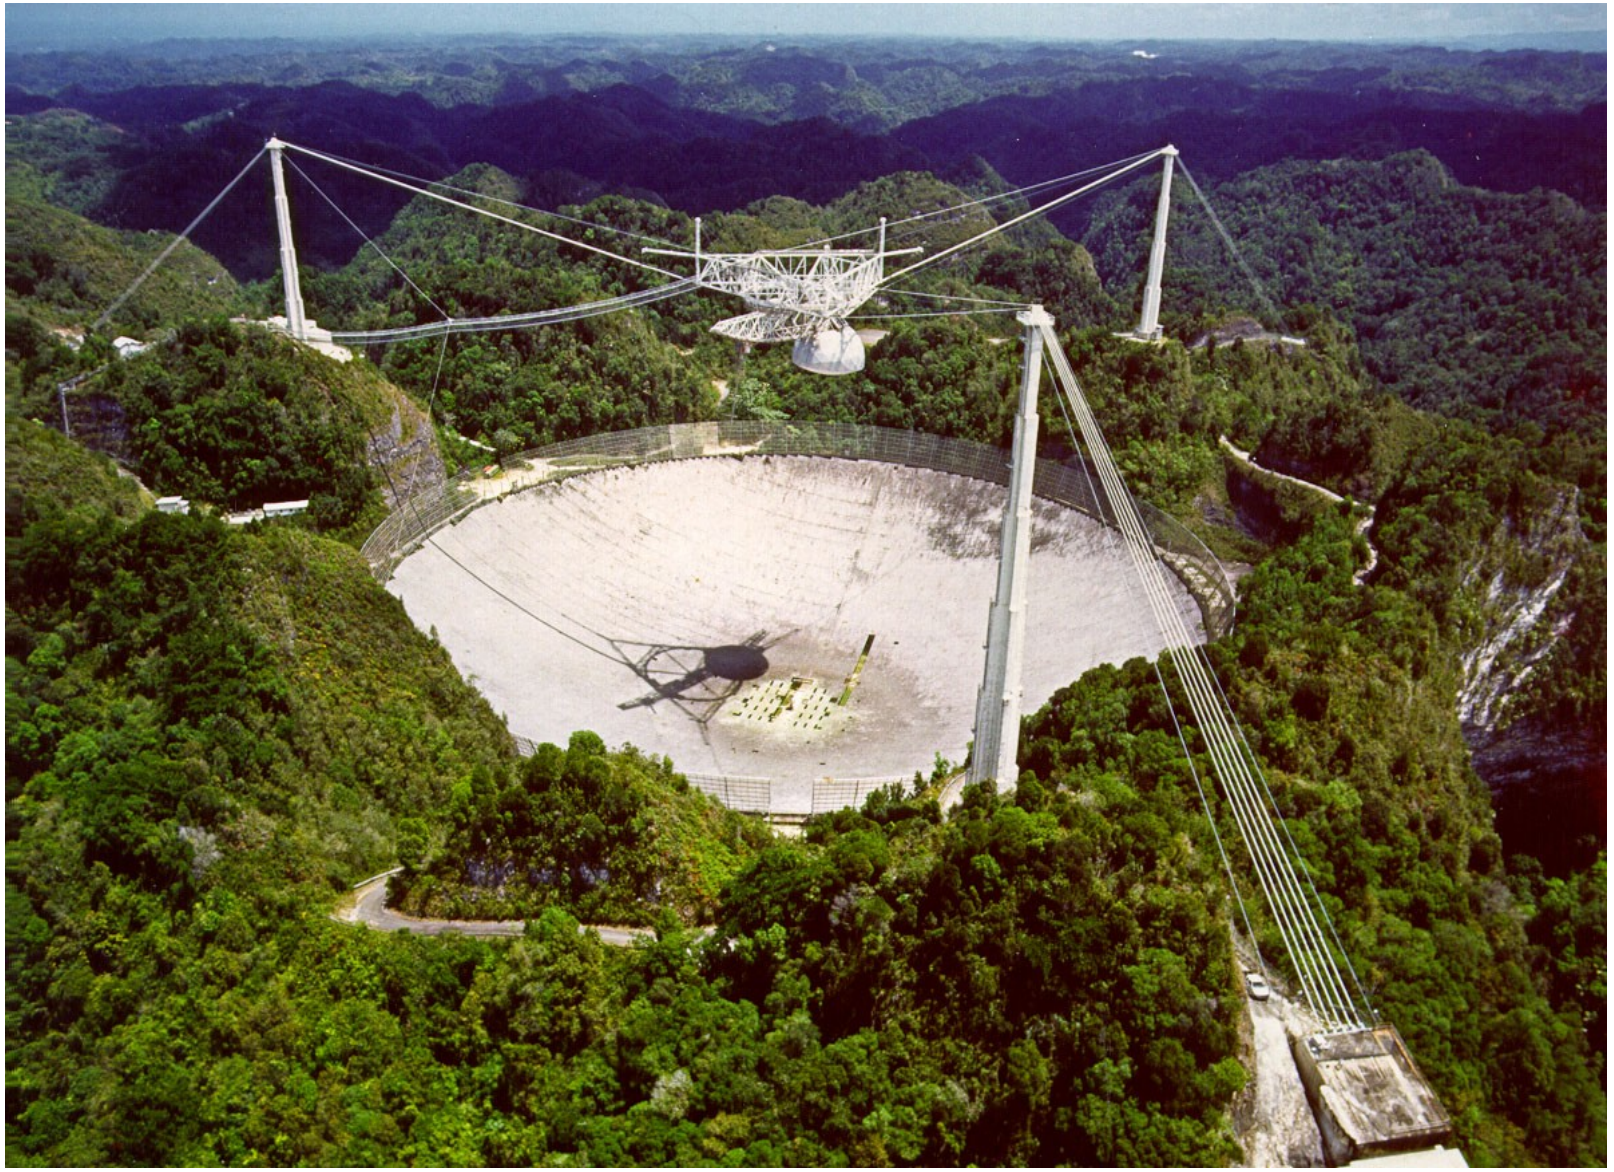

Flood simulation

Sample
Vah river

Computer vision

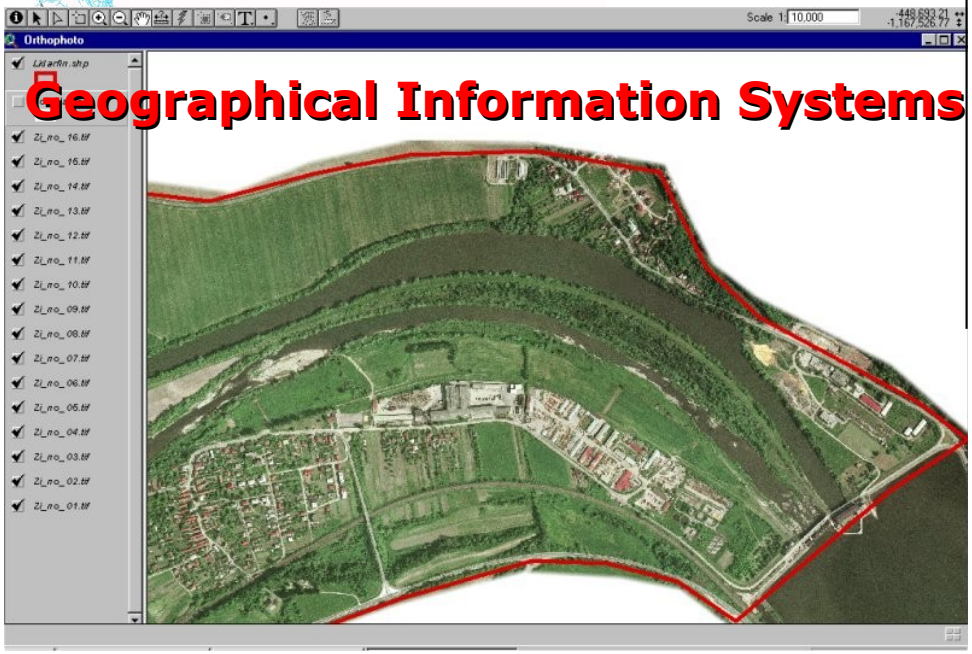

Geographical Information Systems

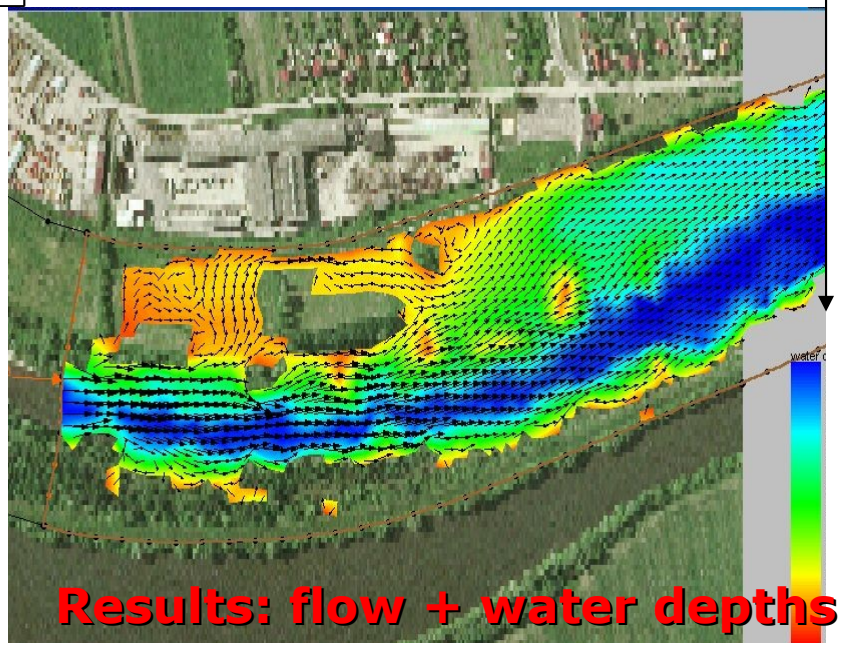

Results: flow + water depths

Network for Earthquake Engineering Simulation

- NEESgrid: national infrastructure to couple earthquake engineers with experimental facilities, databases, computers, & each other
- On-demand access to experiments, data streams, computing, archives, collaboration

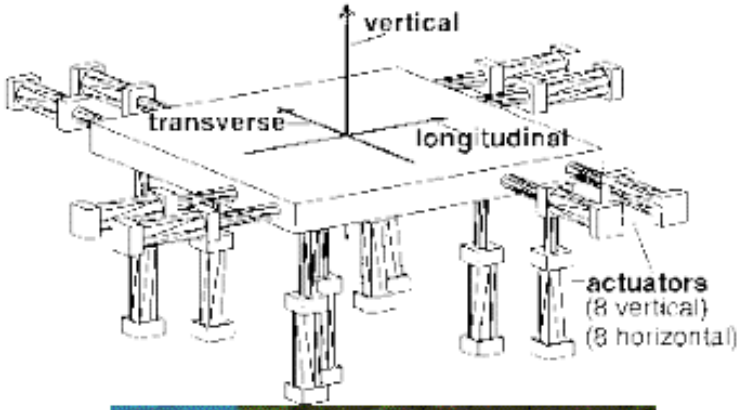

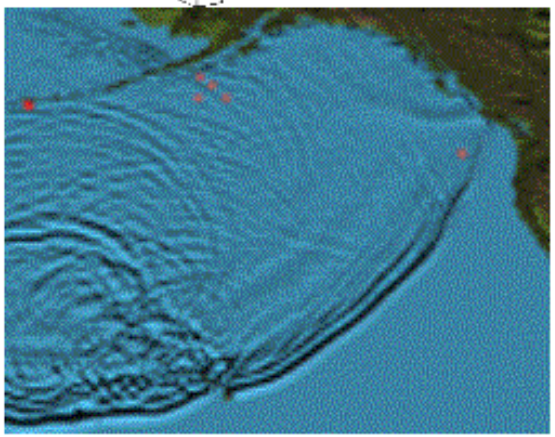

NEESgrid: Argonne, Michigan, NCSA, UIUC, USC

Grid Applications: Art

Paintings are being scanned in at

30 GB each
in the EU CRISATEL Project

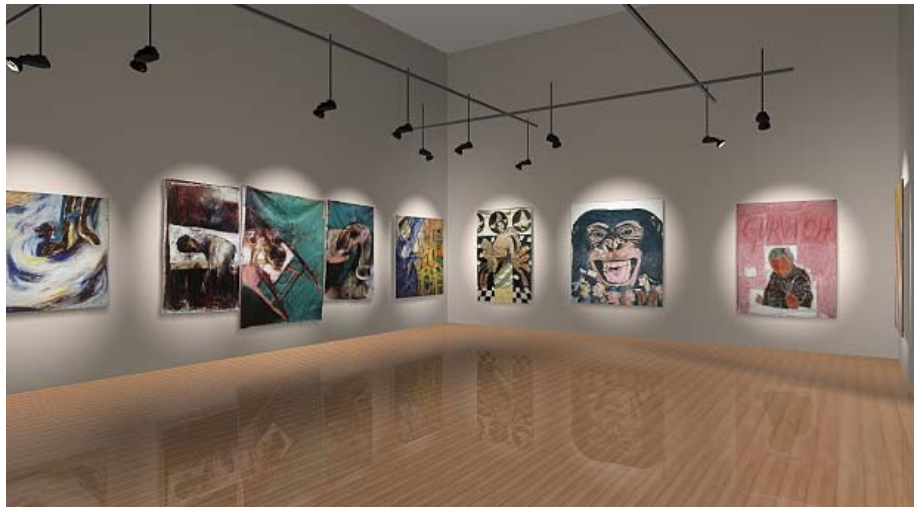

Museo Virtual de Artes El Pais
(MUVA)

Books are being scanned in at
767 MB per page
1/2 Terabyte for Gutenberg
Bible

Connecting *People*: Access Grid

Use Case: Business (1) – CPU Cycles

- Selling compute cycles
- Sun: 1$ / CPU-hour

Use Case: Business (2) – DAME

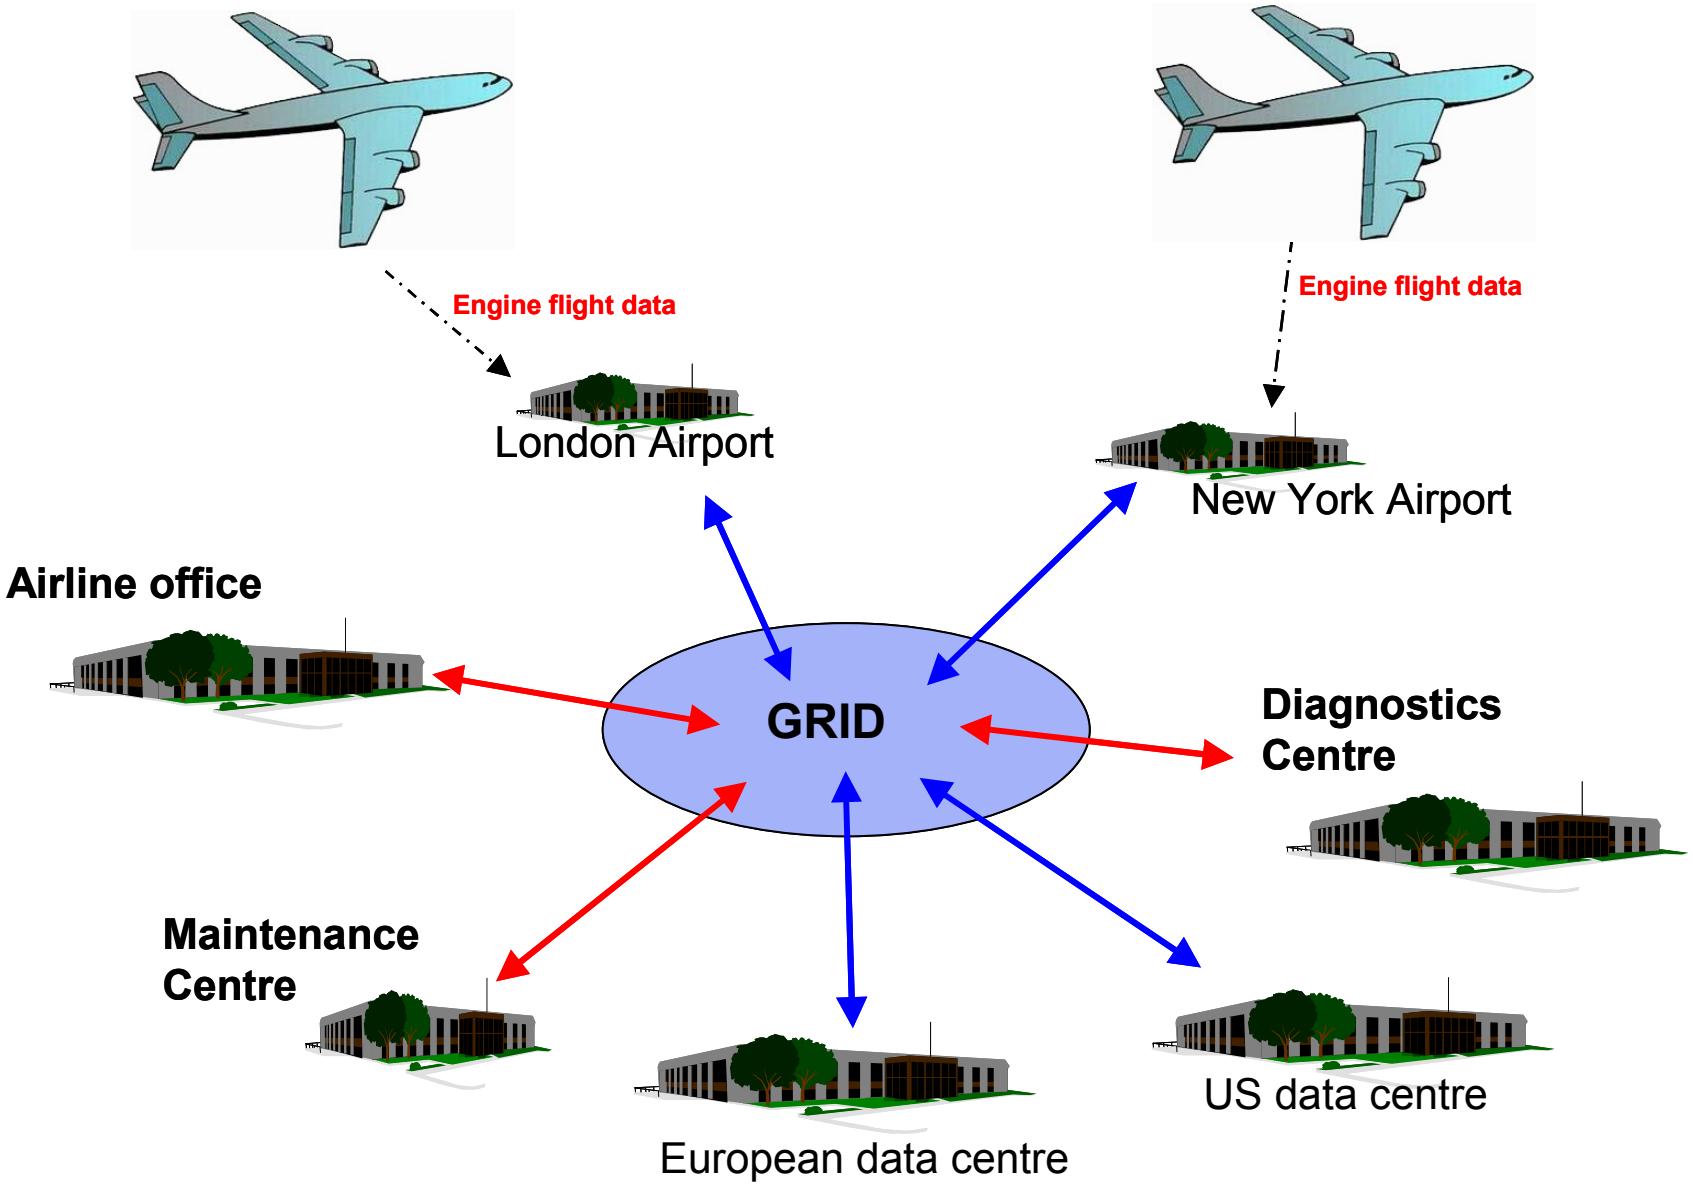

Slide taken from talk given at GridKa School by Dr. Tom Jackson, University of York

DAME Objectives

- **Building a demonstration system as proof of concept for Grid technology in the aerospace diagnostic domain.**
- **Three primary Grid challenges:**
  - Management of large, distributed and heterogeneous data repositories;
  - Rapid data mining and analysis of fault data;
  - Information management and data fusion for diagnosis/prognosis applications;
- **Other key (commercial) issues:**
  - Remote, secure access to flight data and other operational data and resources;
  - Management of distributed users and resources;
  - Quality of Service issues (and Service Level Agreements)

Slide taken from talk given at GridKa School by Dr. Tom Jackson, University of York

Conclusion

- There is no such thing as “The Grid”
- But everyone is working hard on Grid solutions (albeit different ones)
- EGEE plays a special role as a unifying force
- Some Grids can and *will* grow together, they might one day be called the “World Wide Grid”.
- As a deployment project, EGEE is instrumental in this movement
- Application types are limited by latency
- It will be applications that define the ultimate meaning of “The Grid”

Please do ask questions!

I'd like to thank the audience, the organisers of this event, the European Comission, the German Federal Ministry of Education and Research (BMB+F), the EGEE project and its representatives as well as Forschungszentrum Karlsruhe!

bmb+f - Förderschwerpunkt
Hadronen - und Kernphysik
Großgeräte der physikalischen Grundlagenforschung