Applications and the Grid

EDG Tutorial @ CERN 19.03.2003

Ingo Augustin

CERN

DataGrid HEP Applications

Introduction

You've heard much about WHAT the Grid is, but not much about WHY the Grid is, or will be, or should be or whatever….

The Rationale behind the Grid *)

- **Size:**
  - The Large Hadron Collider Experiments
- **Geographical Distribution:**
  - The Monarc Computing Model
- **Complexity:**
  - Earth Observation Applications
- **User Community:**
  - Biomedical Applications

*) **I am a physicist! All mistakes in EO & Bio applications are due to my ignorance.**

Electrical Power Grid Metaphor

- Power on demand
  - User unaware of actual provider
- Resilience
  - Re-routing
  - Redundancy
- Simple interface
  - Wall socket
- Standardised protocols
  - 230 V, 50 Hz

LHC Experiments

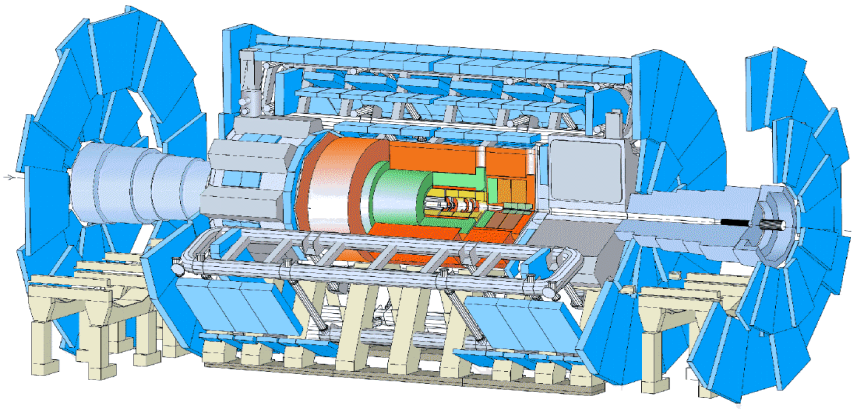

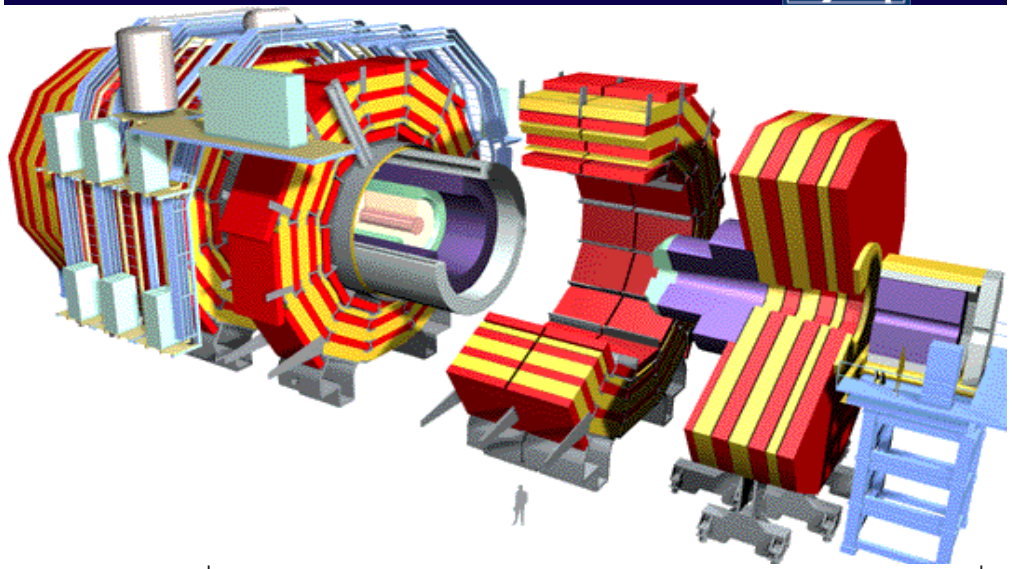

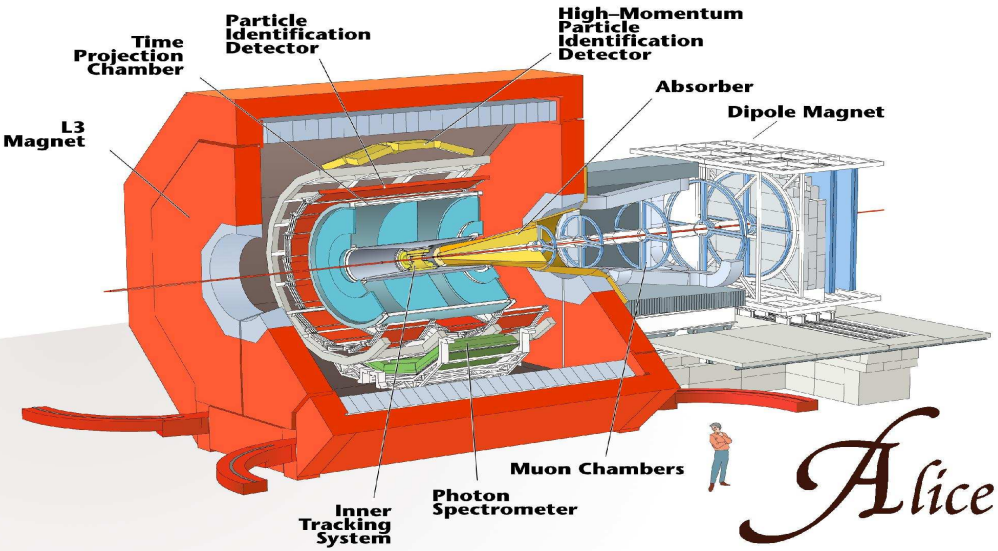

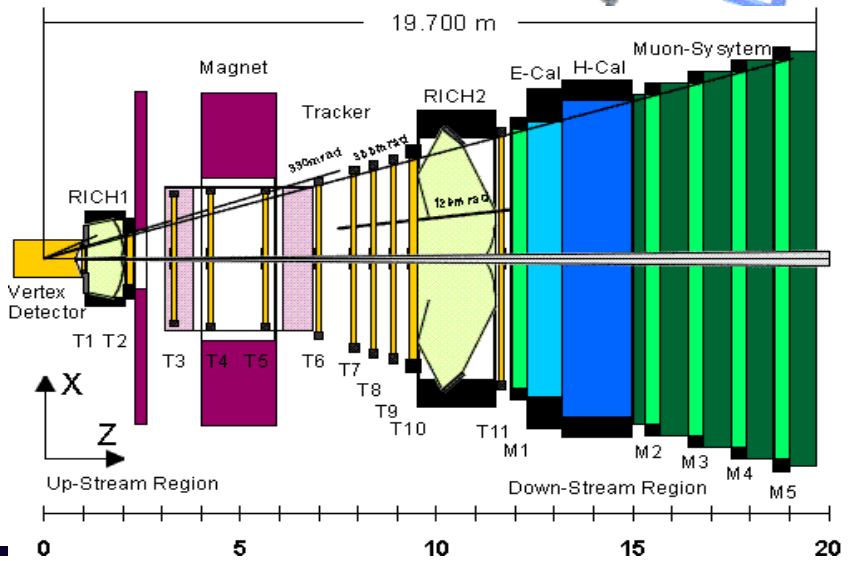

More Complex Events

< 2000

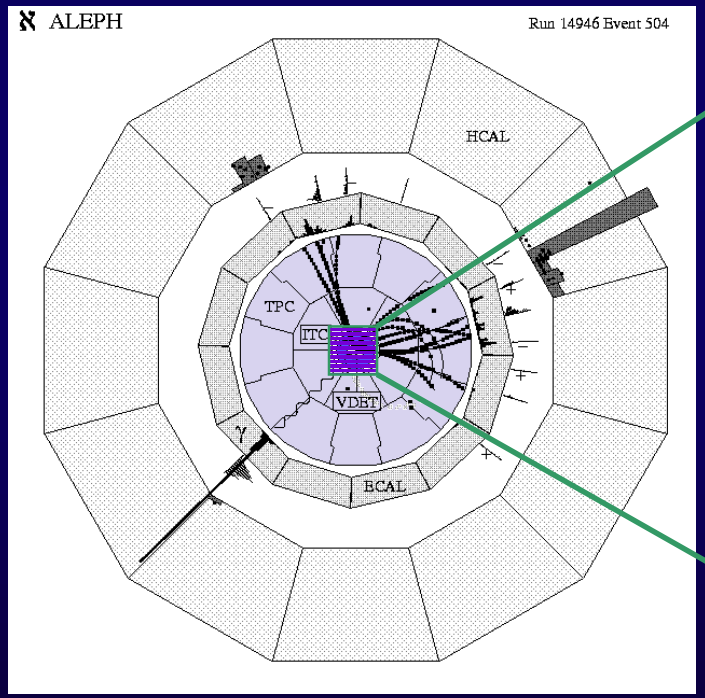

> 2007

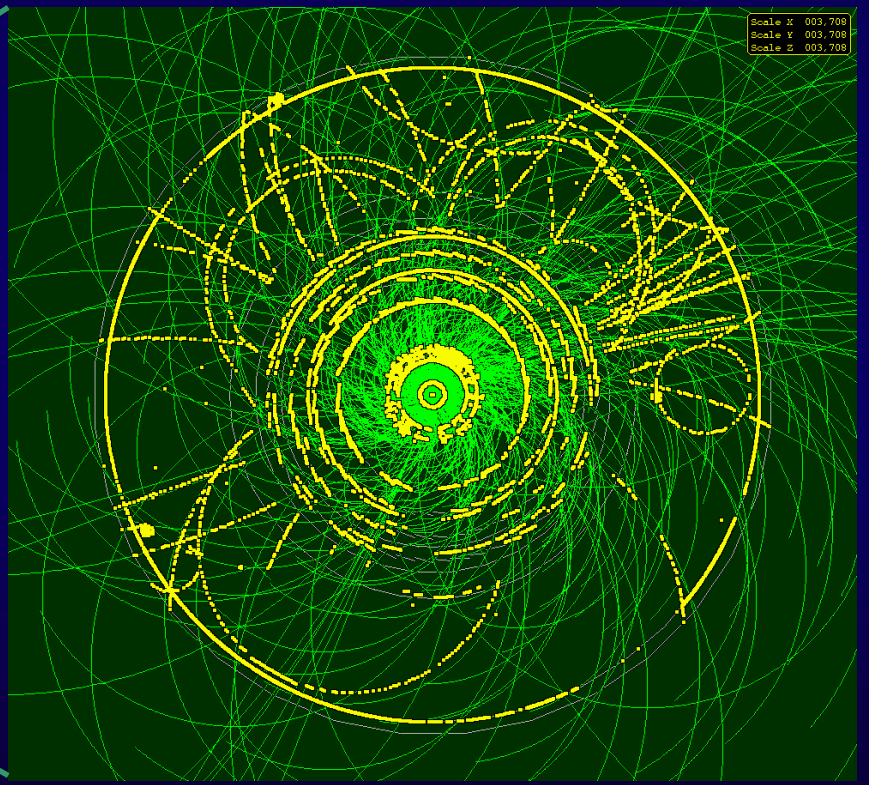

Typical HEP Software Scheme

Detector description

Detector alignment

Detector calibration

Reconstruction parameters

Generate Events

Build Simulation Geometry

Build Reconstruction Geometry

Reconstuction geometry

Reconstruct Events

Physics

Analyze Events

Simulation geometry

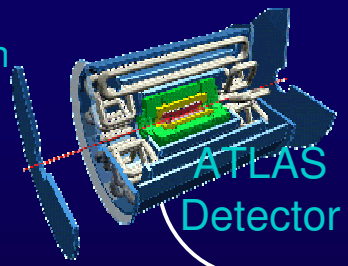

Simulate Events

Raw Data

ESD AOD

Characteristics of HEP computing

Event independence

- Data from each collision is processed independently
- Mass of independent problems with no information exchange

Massive data storage

- Modest event size: 1-25 MB
- Total is very large - Petabytes for each experiment.

Mostly read only

- Data never changed after recording to tertiary storage
- But is read often ! cf.. magnetic tape as an archive medium

Modest floating point needs

- HEP computations involve decision making rather than calculation
- Computational requirements in SPECint95 secs

Typical Layout of a Computing Farm (up to several hundred nodes)

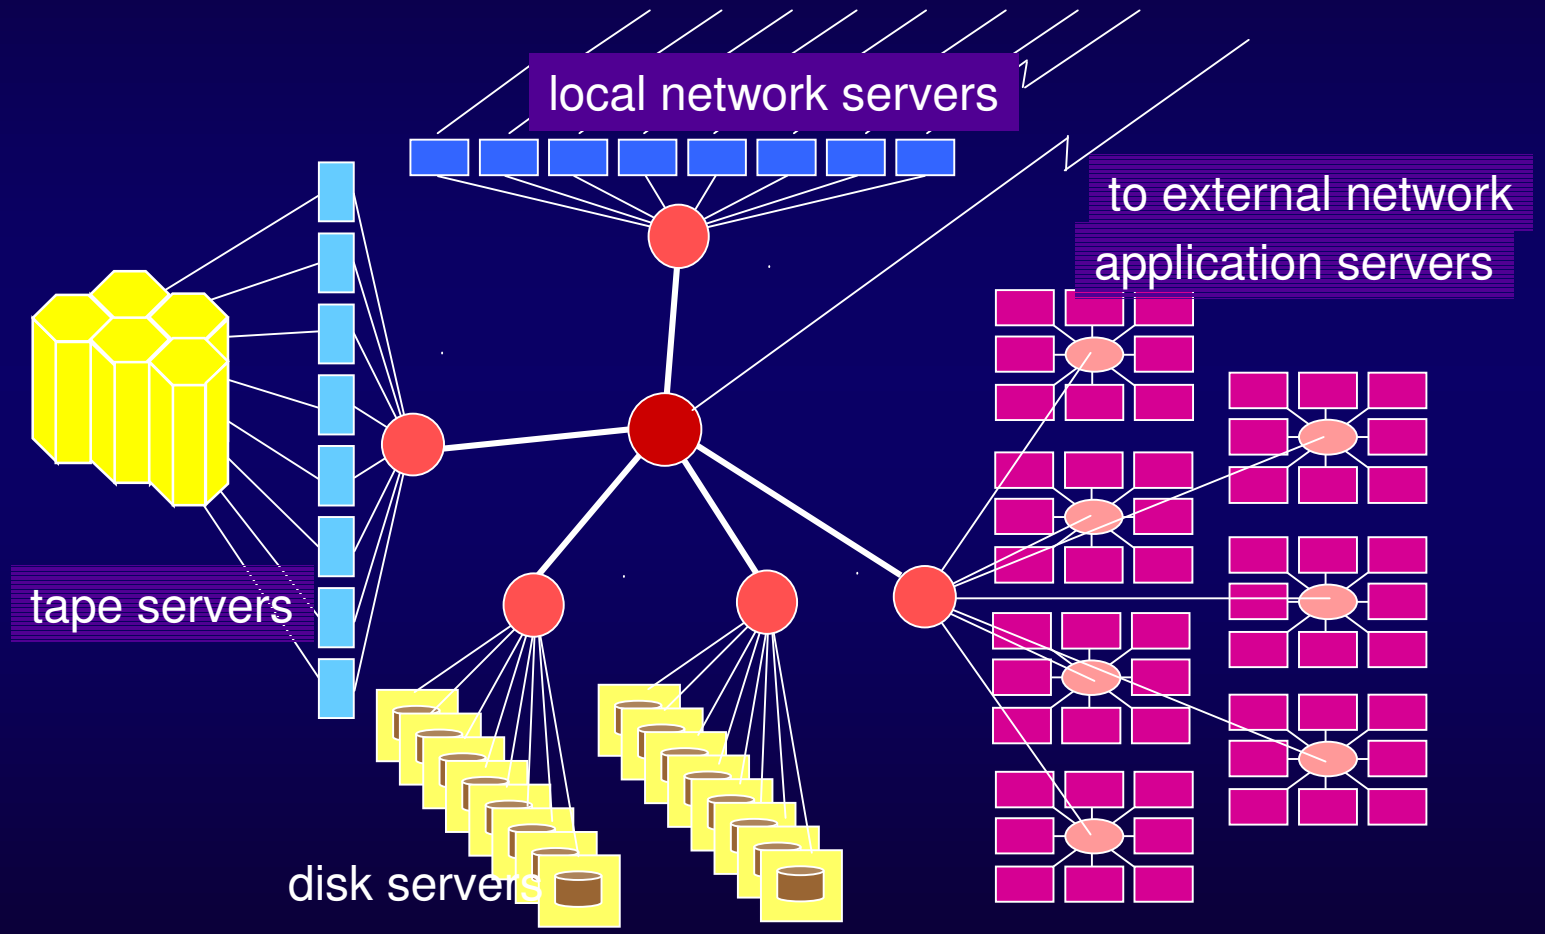

The Constraints

Needed during a year of LHC operations		
Tape	Disk	CPU
29'400 TB	9'600 TB	6.2 * 10^6 SI95
In today's units:		
60 STK Silos	160'000 60GB disks	150'000 800 MHz CPUs

Taken from: LHC Computing Review, CERN/LHCC/2001-004

World Wide Collaboration

⇒ distributed computing & storage capacity

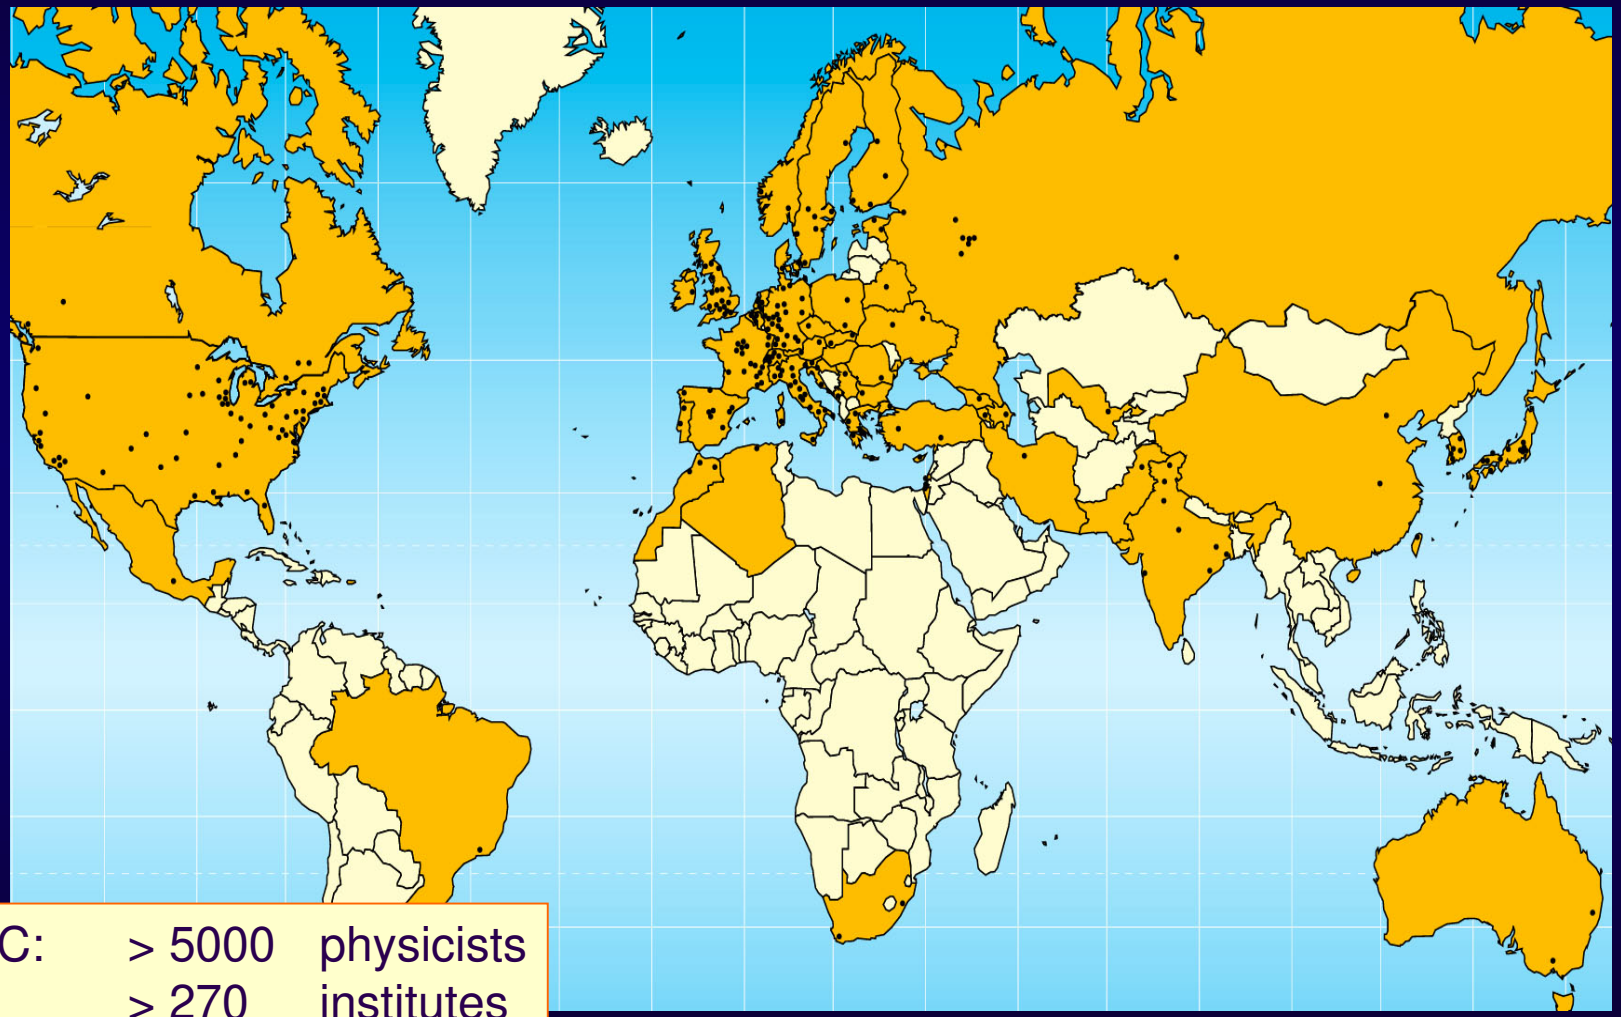

LHC:	> 5000	physicists
	> 270	institutes
	> 60	countries

World-wide computing

Two problems:

- **Funding**
  - will funding bodies place all their investment at CERN?

No

- **Geography**
  - does a geographically distributed model better serve the needs of the world-wide distributed community?

Maybe –

if it is **reliable** and **easy to use**

Need to provide physicists with the best possible access to LHC data irrespective of location

Present LHC Computing Model

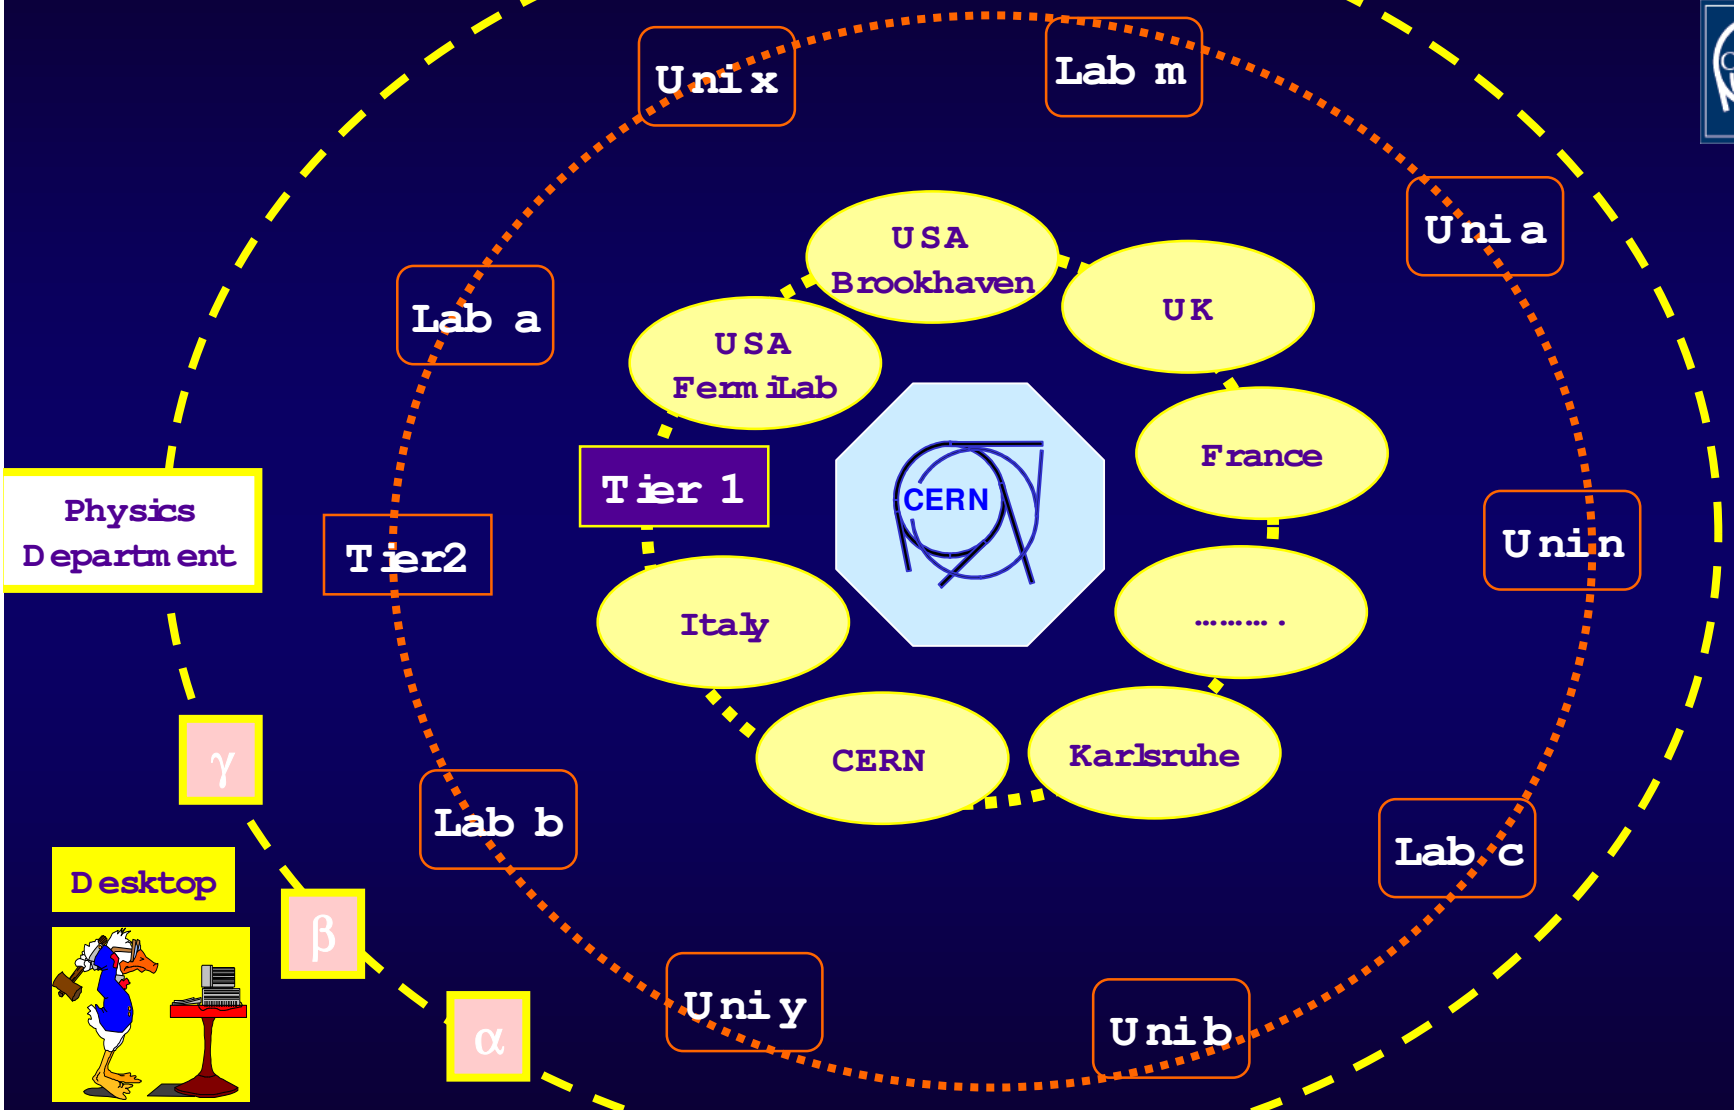

Regional Center

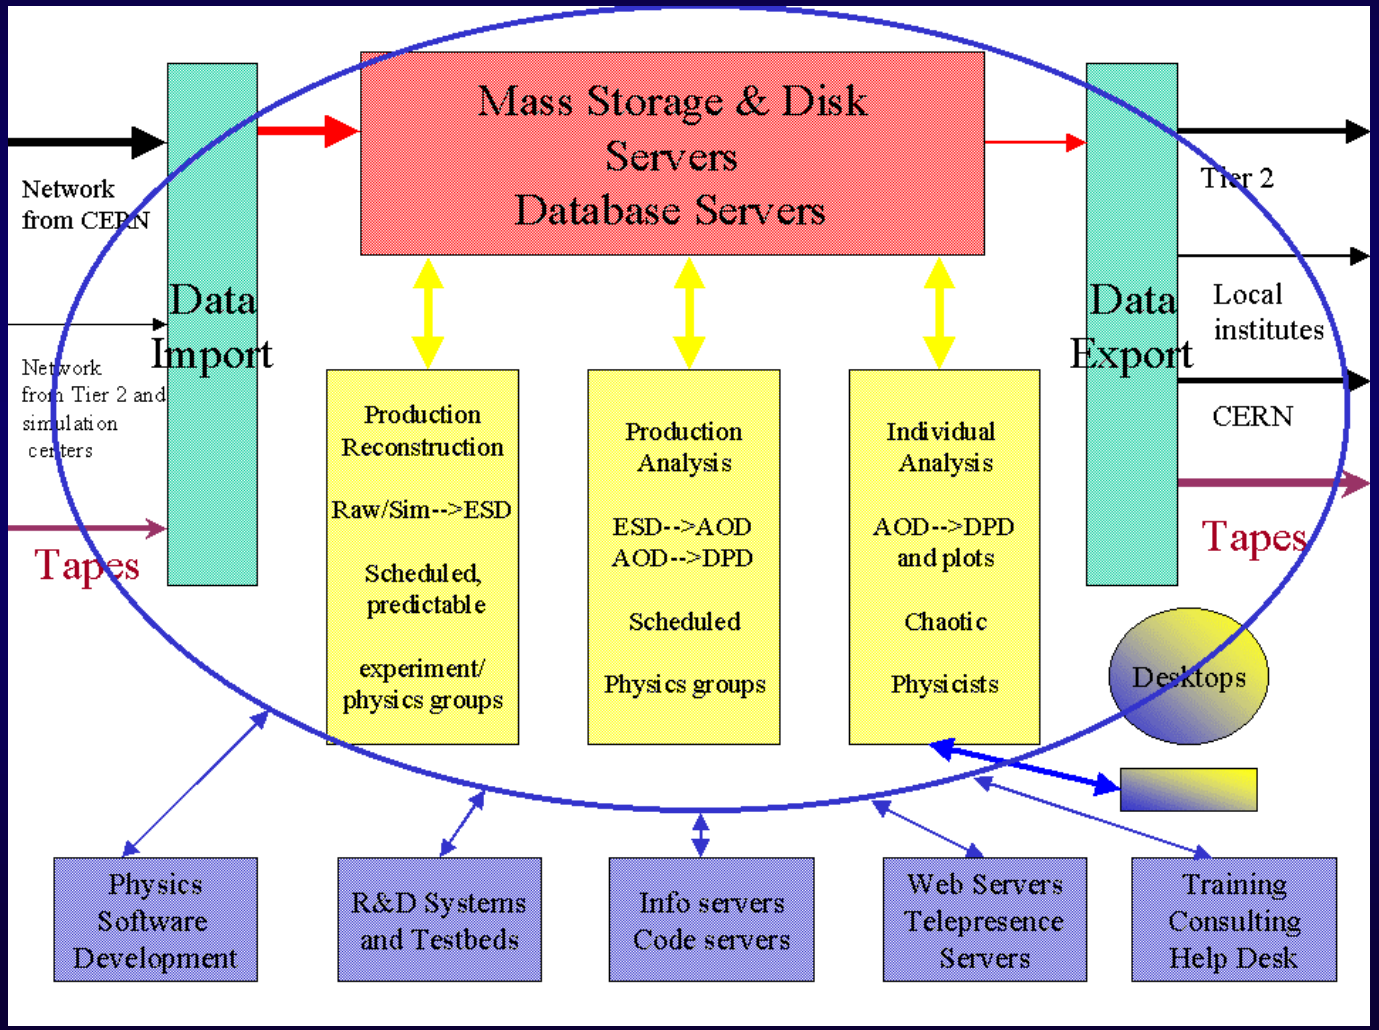

The Dungeon

- **Pain (administration)**
  - **money**
  - **manpower**
    - reduction by ~ 30% before start of LHC
  - **commodity**
    - PC & Network & ...
- **Torture (users & history)**
  - **anarchic user community**
  - **legacy (software & structures)**
    - evolution instead of projects
- **Execution (deadline)**
  - **2006/7 start-up of LHC**

Earth Observation (WP9)

- Global Ozone (GOME) Satellite Data Processing and Validation by KNMI, IPSL and ESA
- **The** DataGrid testbed **provides a** collaborative processing environment **for 3 geographically distributed** EO **sites (Holland, France, Italy)**

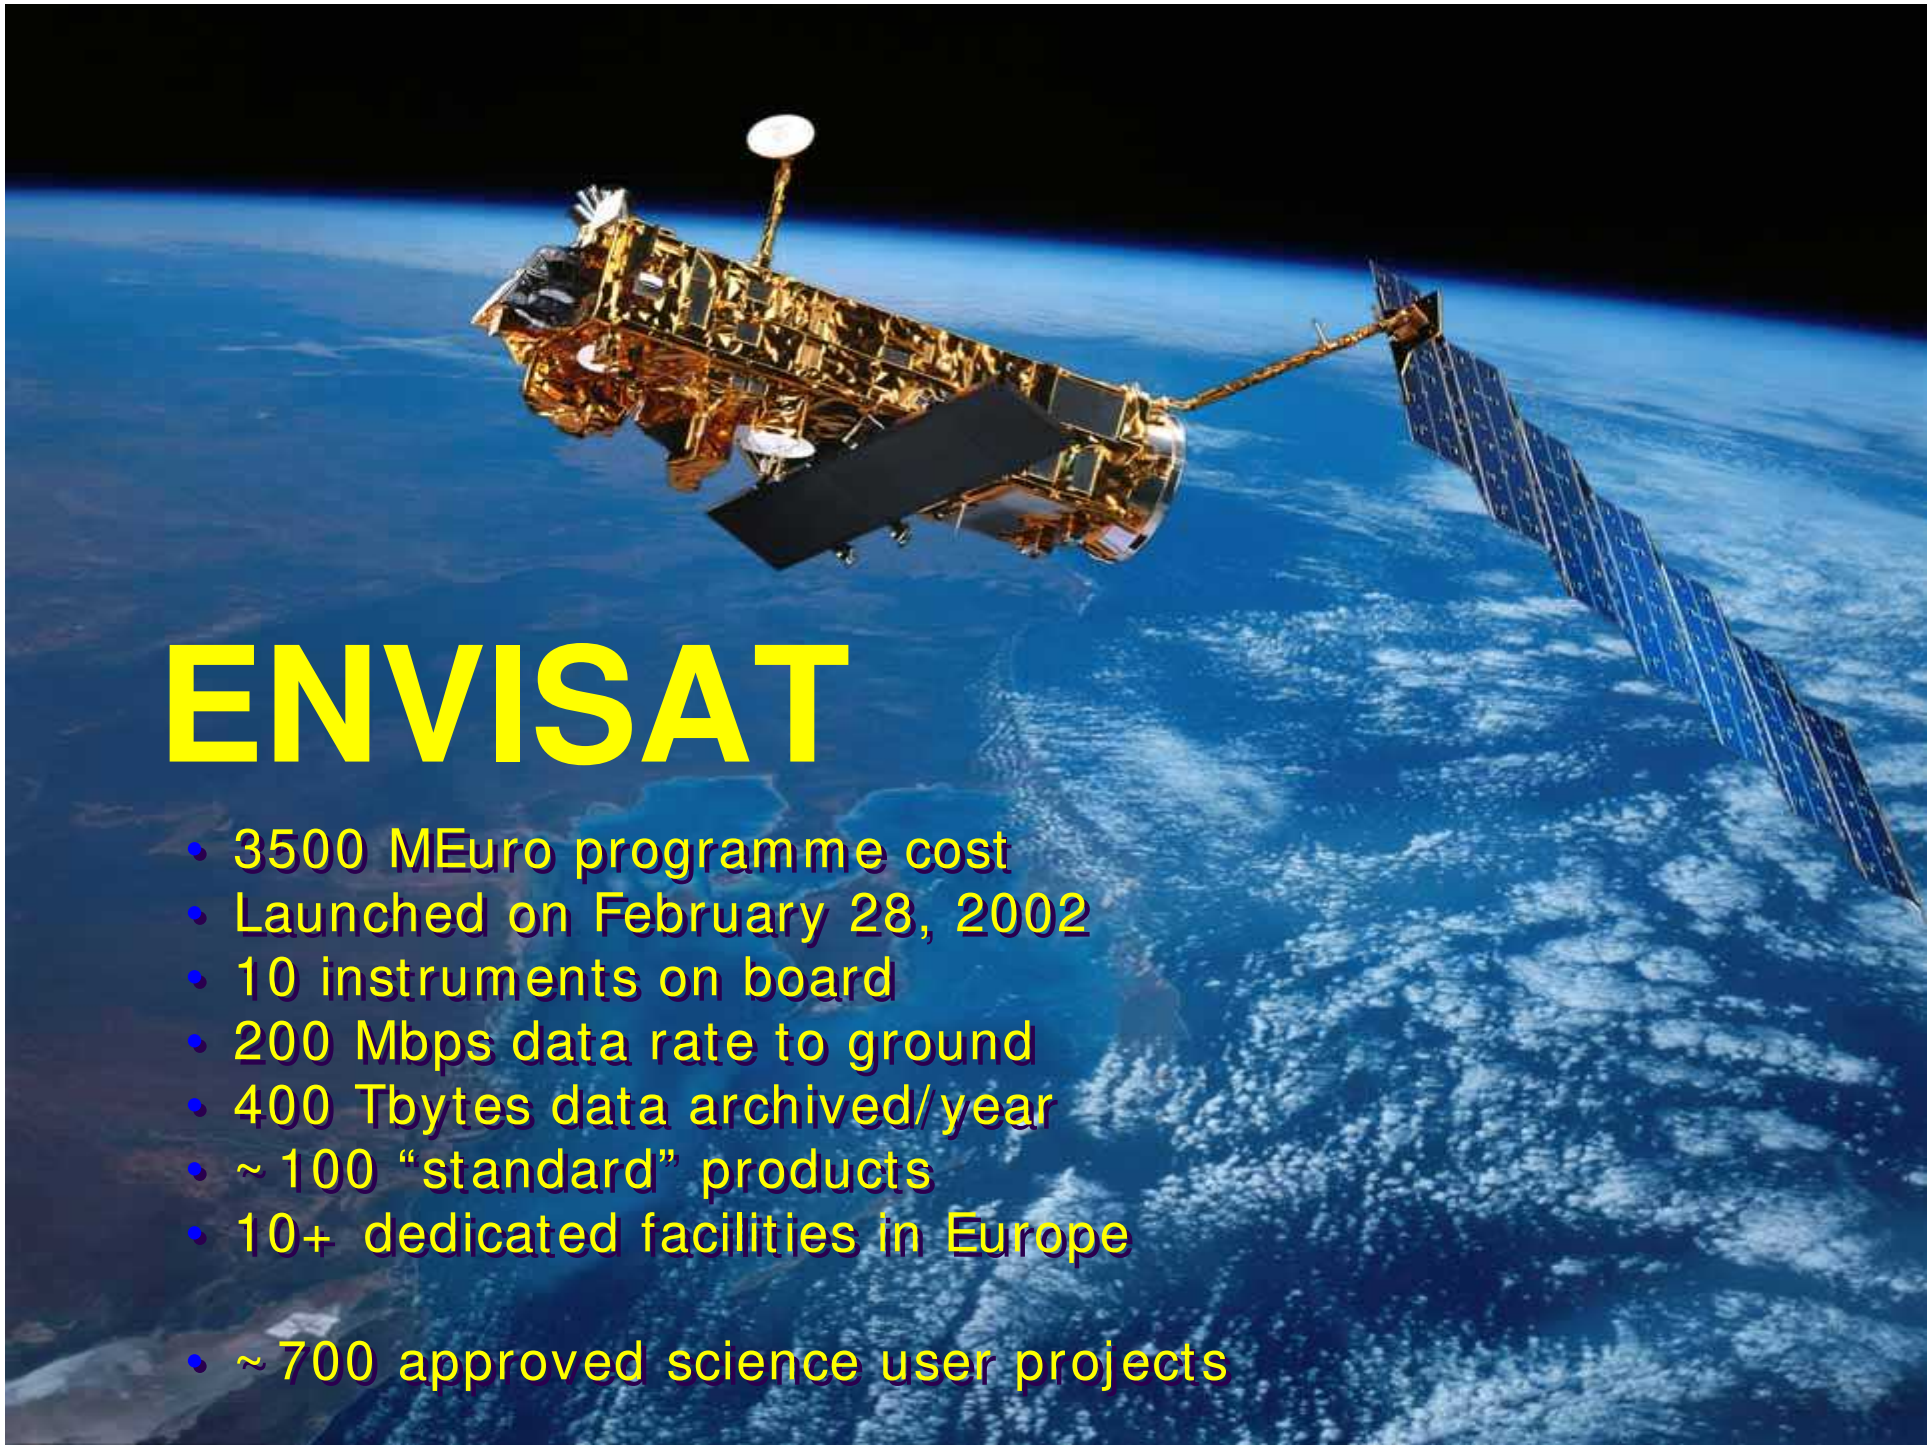
ENVISAT
• 3500 MEuro programme cost
• Launched on February 28, 2002
• 10 instruments on board
• 200 Mbps data rate to ground
• 400 Tbytes data archived/year
• ~100 “standard” products
• 10+ dedicated facilities in Europe
• ~700 approved science user projects

Earth Observation

- **Two different GOME processing techniques will be investigated**
  - **OPERA (Holland) - Tightly coupled - using MPI**
  - **NOPREGO (Italy) - Loosely coupled - using Neural Networks**
- **The results are checked by VALIDATION (France). Satellite Observations are compared against ground-based LIDAR measurements coincident in area and time.**

GOME OZONE Data Processing Model

- Level-1 **data (raw satellite measurements) are analysed to retrieve actual physical quantities : Level-2 data**
- Level-2 **data provides measurements of OZONE within a vertical column of atmosphere at a given lat/lon location above the Earth's surface**
- Coincident **data consists of Level-2 data co-registered with LIDAR data (ground-based observations) and compared using statistical methods**

The EO Data challenge: Processing and validation of 1 year of GOME data

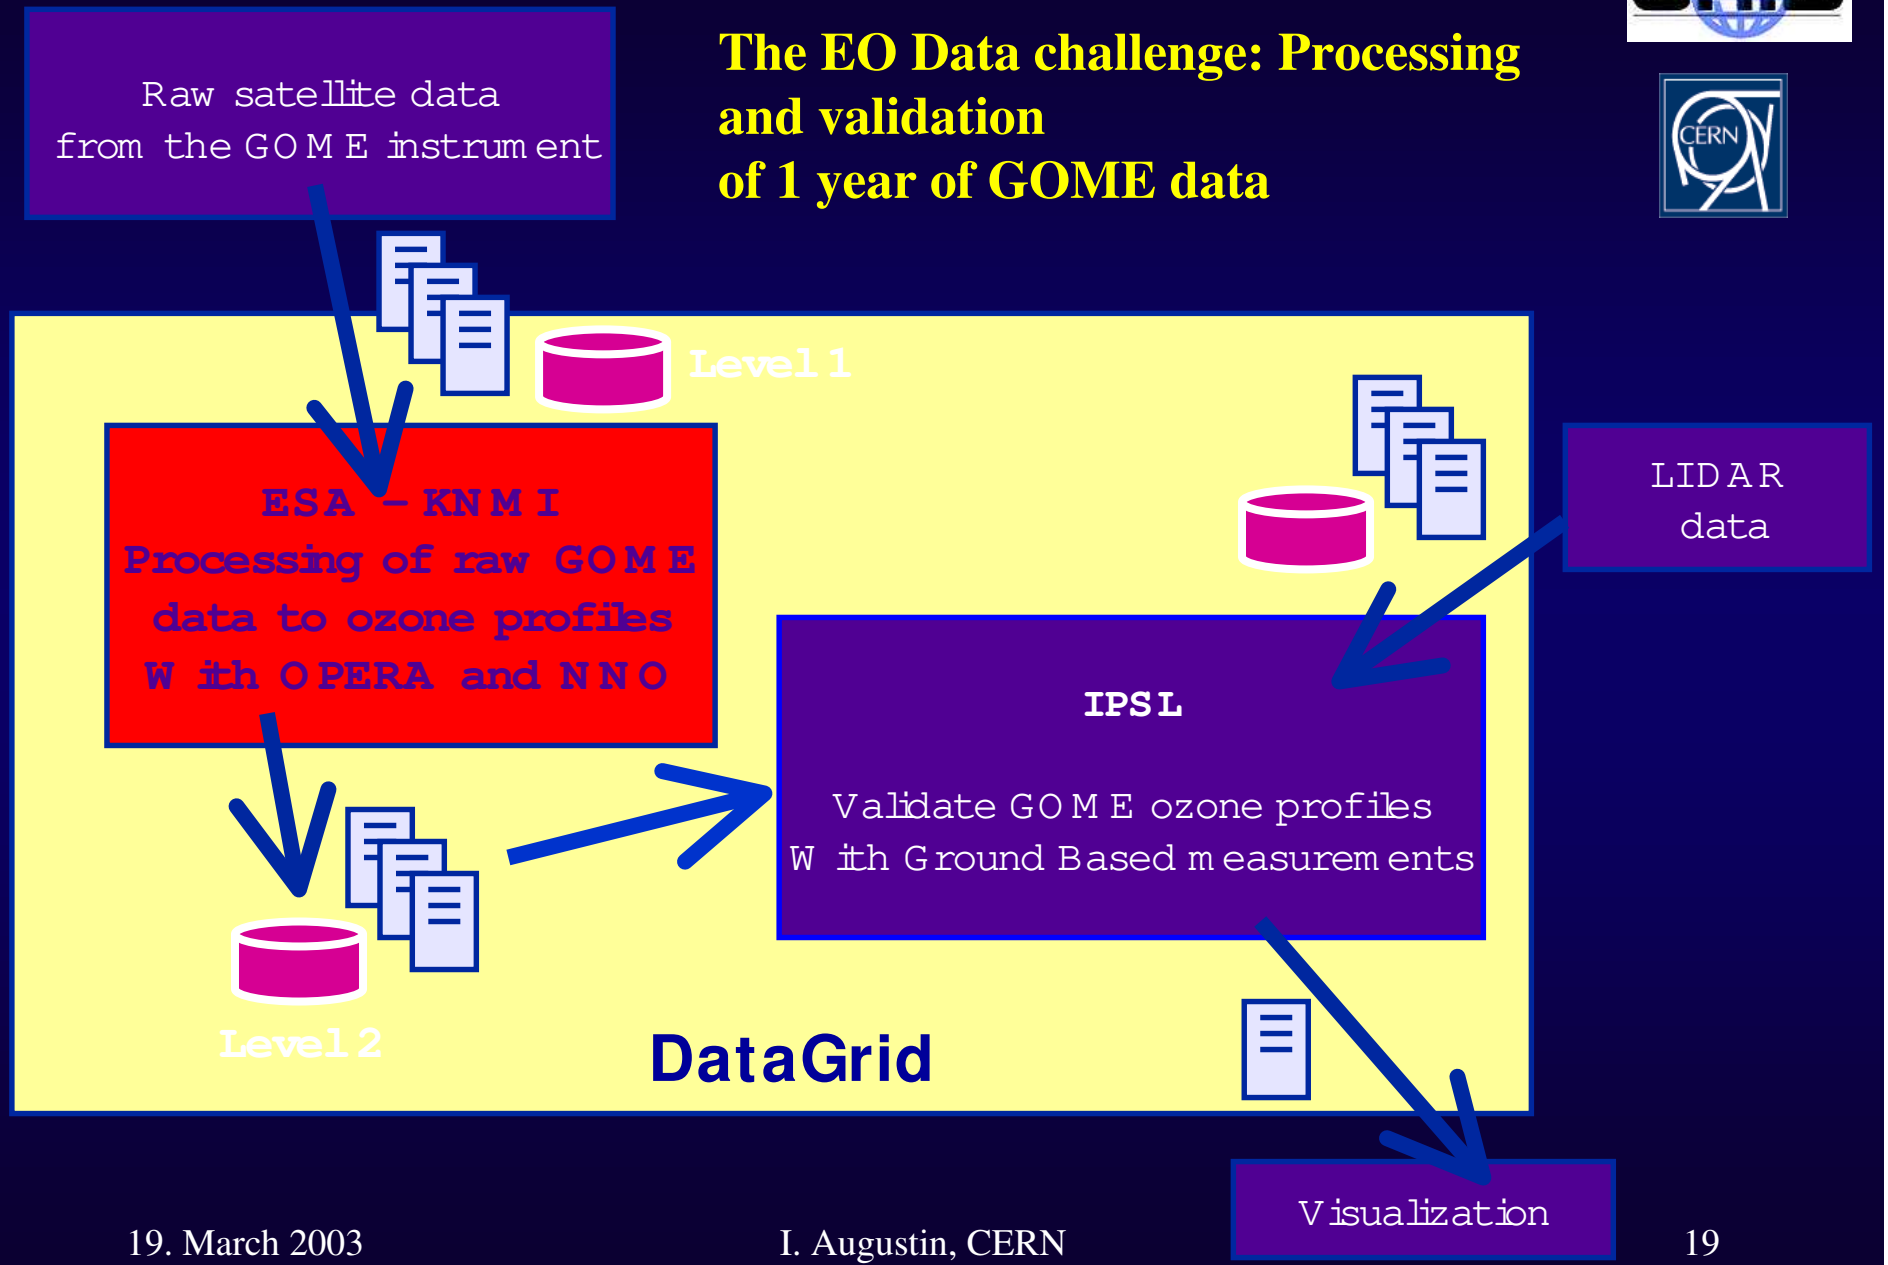

EO Use-Case File Numbers

1 Year of GOME data

Data	Number of files to be processed and replicated	Size
Level 1 (Satellite data)	4,724	15 Mb
Level 2 (NNO)	9,448,000	10 kb
Level 2 (Opera)	9,448,000	12 kb
Coincident (Validation)	12	2.5 Mb
Total:	18,900,736 files	267 Gbyte

Part of a 5-year, global dataset

GOME Processing Steps (1-2)

Step 1: Transfer **Level1** data to the Grid Storage Element

Step 2: Register **Level1** data with the ReplicaManager

Replicate to other SEs if necessary

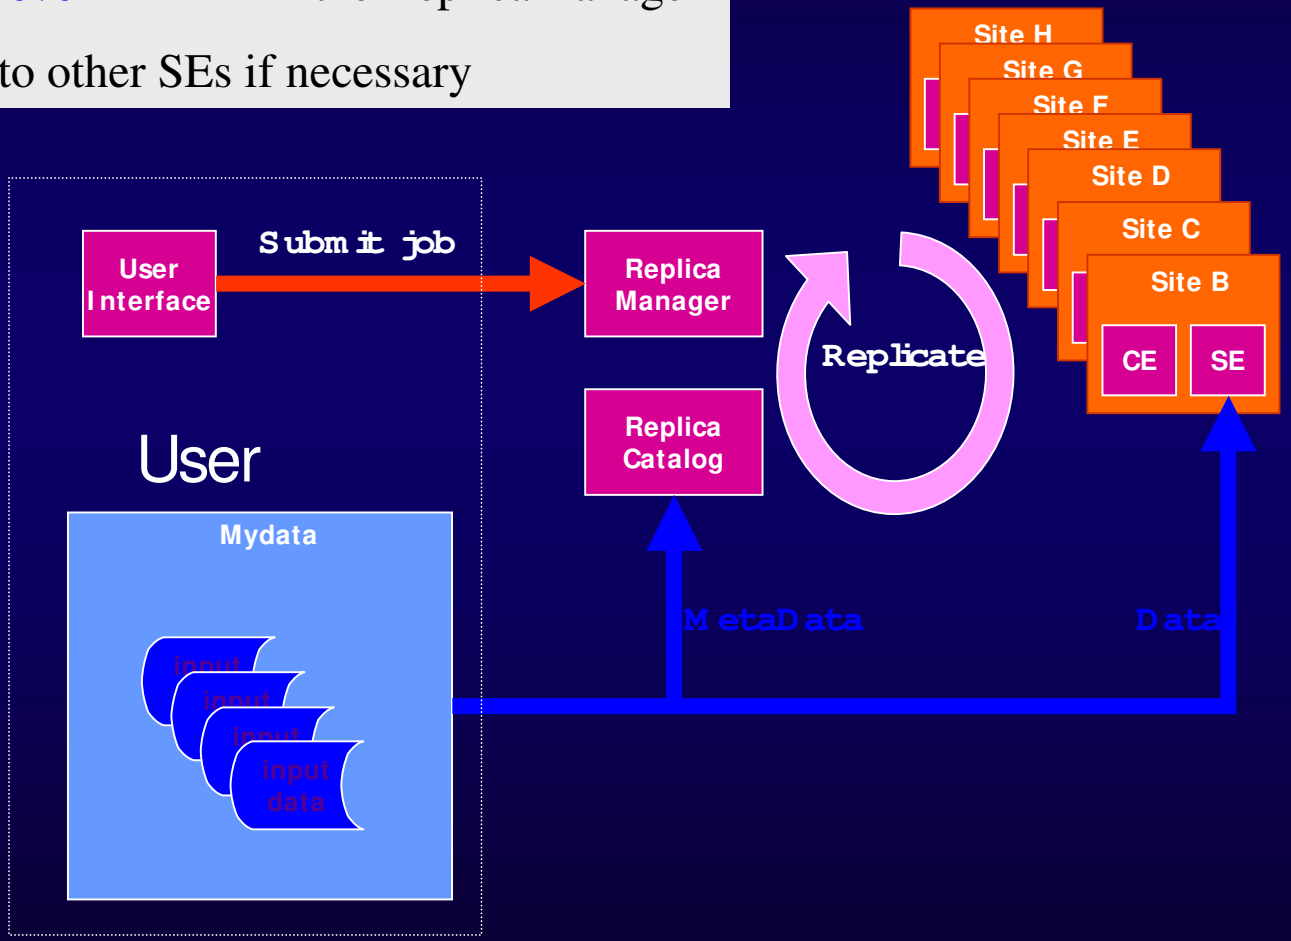

GOME Processing Steps (3-4)

Step 3: Submit jobs to process **Level1** data, produce Level2 data

Step 4: Transfer **Level2** data products to the Storage Element

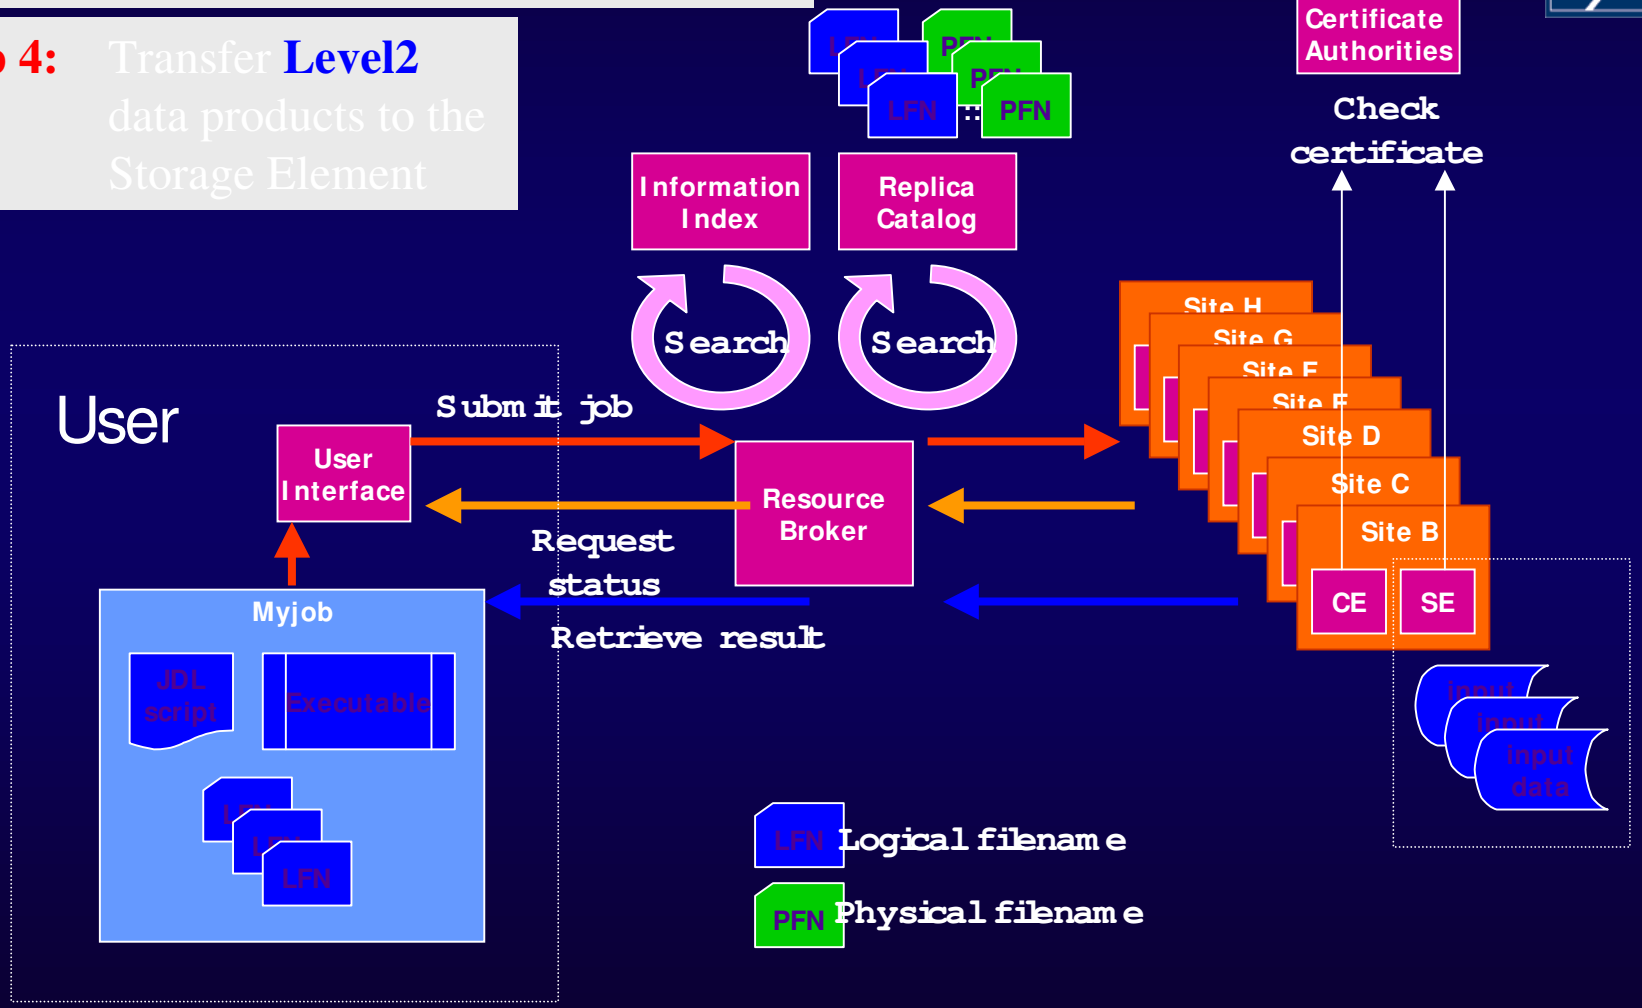

GOME Processing Steps (5-6)

Step 5: Produce Level-2 / LIDAR **Coincident** data perform **VALIDATION**

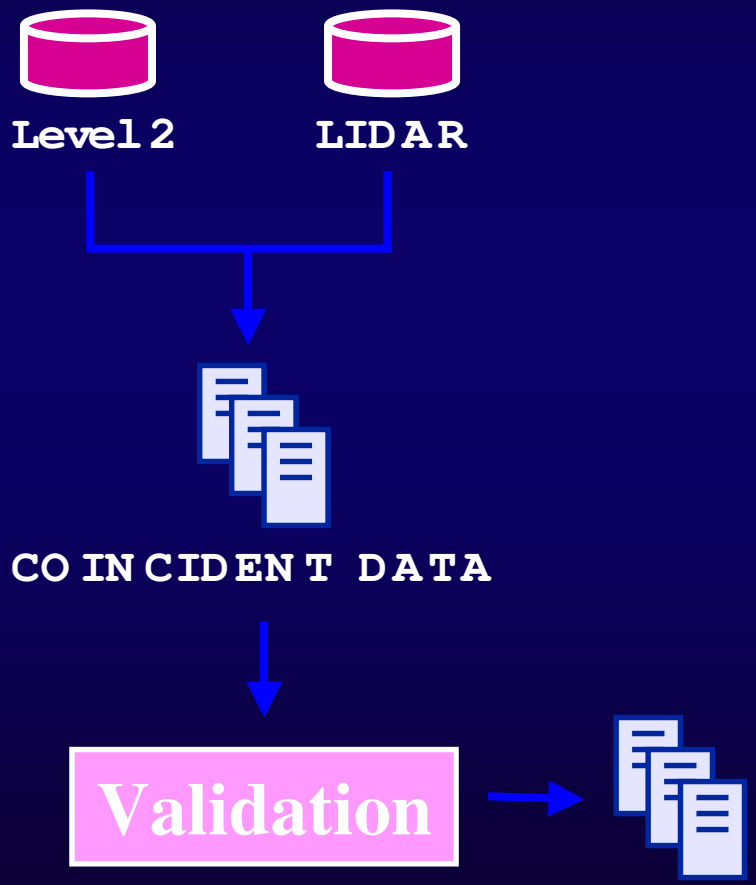

Step 6: Visualize Results

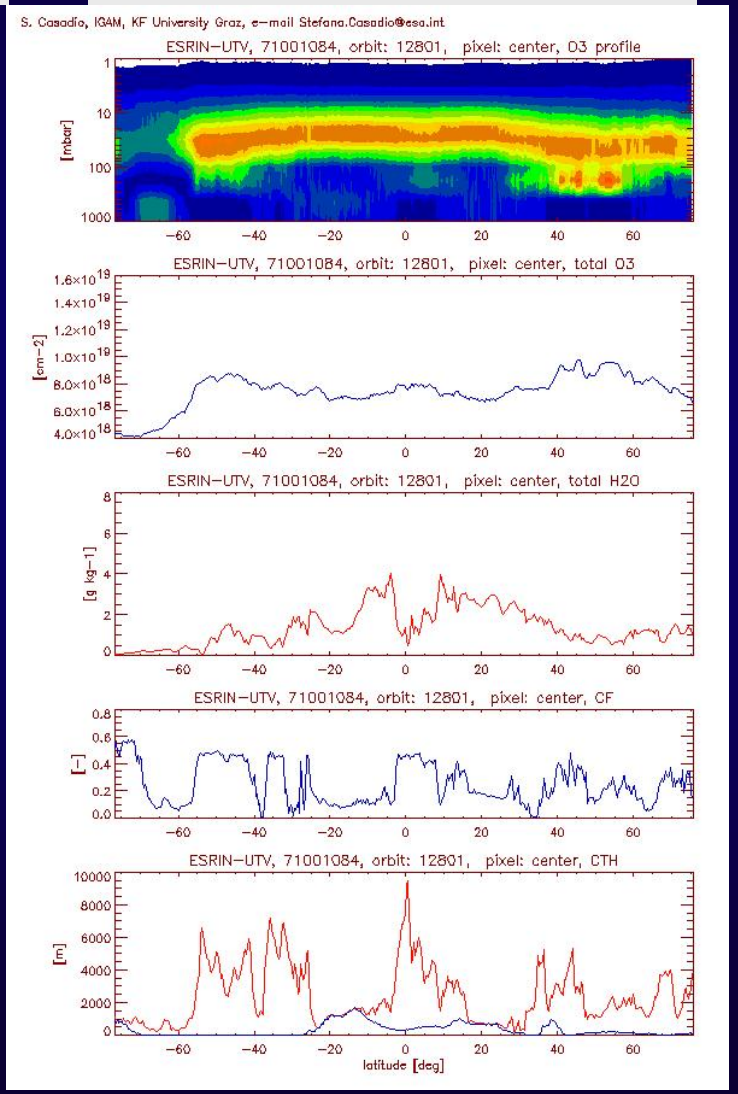

Genomics and Bioinformatics (WP10)

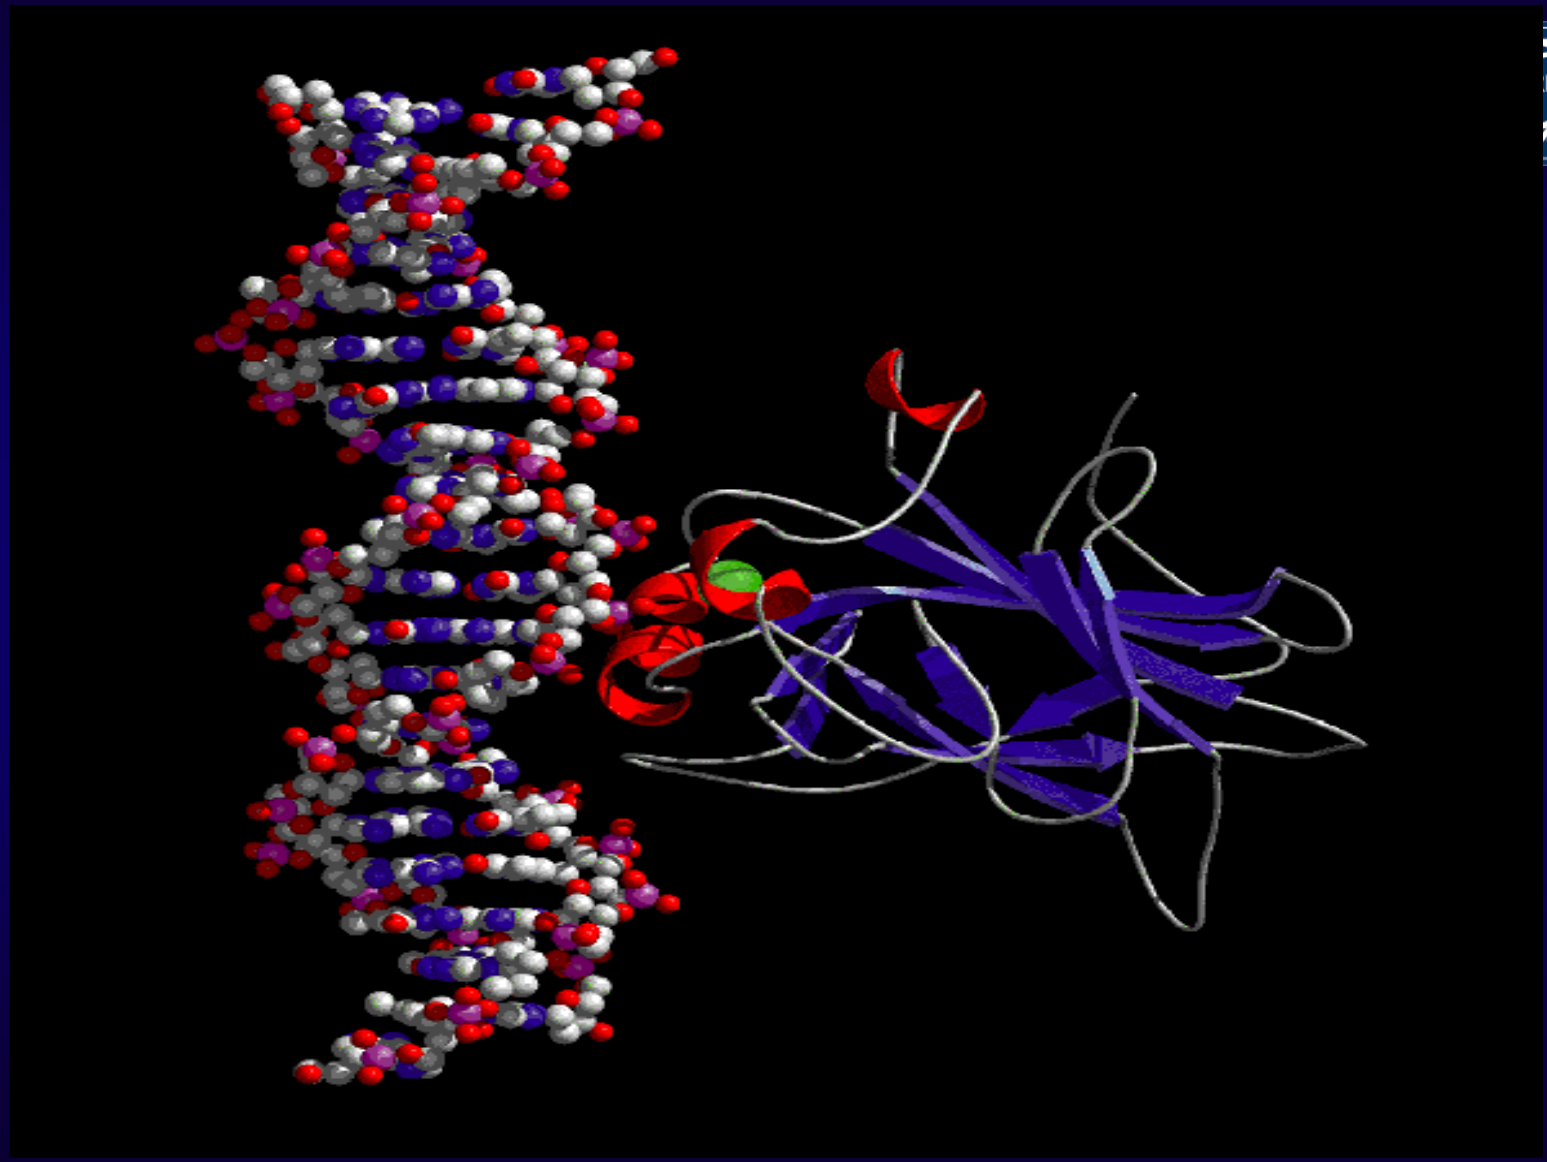

Challenges for a biomedical grid

- The biomedical community has NO strong center of gravity in Europe
  - No equivalent of CERN (High-Energy Physics) or ESA (Earth Observation)
  - Many high-level laboratories of comparable size and influence without a practical activity backbone (EMB-net, national centers,…) leading to:
    - **Little awareness of common needs**
    - **Few common standards**
    - **Small common long-term investment**
- The biomedical community is very large (tens of thousands of potential users)
- The biomedical community is often distant from computer science issues

Biomedical requirements

- Large user community(thousands of users)
  - anonymous/group login
- Data management
  - data updates and data versioning
  - Large volume management (a hospital can accumulate TBs of images in a year)
- Security
  - disk / network encryption
- Limited response time
  - fast queues
- High priority jobs
  - privileged users
- Interactivity
  - communication between user interface and computation
- Parallelization
  - MPI site-wide / grid-wide
  - Thousands of images
  - Operated on by 10's of algorithms
- Pipeline processing
  - pipeline description language / scheduling

The grid impact on data handling

- DataGrid will allow mirroring of databases
  - An alternative to the current costly replication mechanism
  - Allowing web portals on the grid to access updated databases

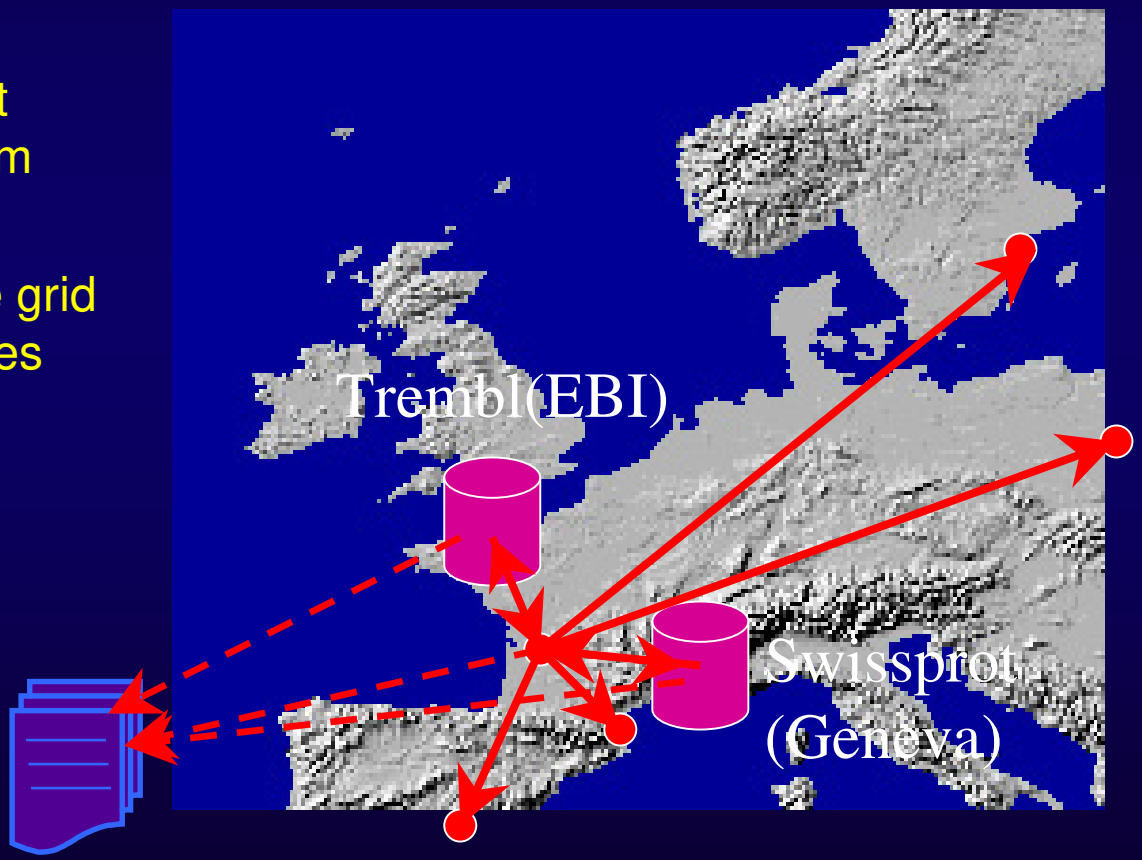

Biomedical Replica Catalog

Web portals for biologists

- Biologist enters sequences through web interface
- Pipelined execution of bio-informatics algorithms
  - Genomics comparative analysis (thousands of files of ~Gbyte)
    - Genome comparison takes days of CPU (~n**2)
  - Phylogenetics
  - 2D, 3D molecular structure of proteins…
- The algorithms are currently executed on a local cluster
  - Big labs have big clusters …
  - But growing pressure on resources – Grid will help
    - More and more biologists
    - compare larger and larger sequences (whole genomes)…
    - to more and more genomes…
    - with fancier and fancier algorithms !!

The Visual DataGrid Blast

- A graphical interface to enter query sequences and select the reference database
- A script to execute the BLAST algorithm on the grid
- A graphical interface to analyze result
- Accessible from the web portal genius.ct.infn.it

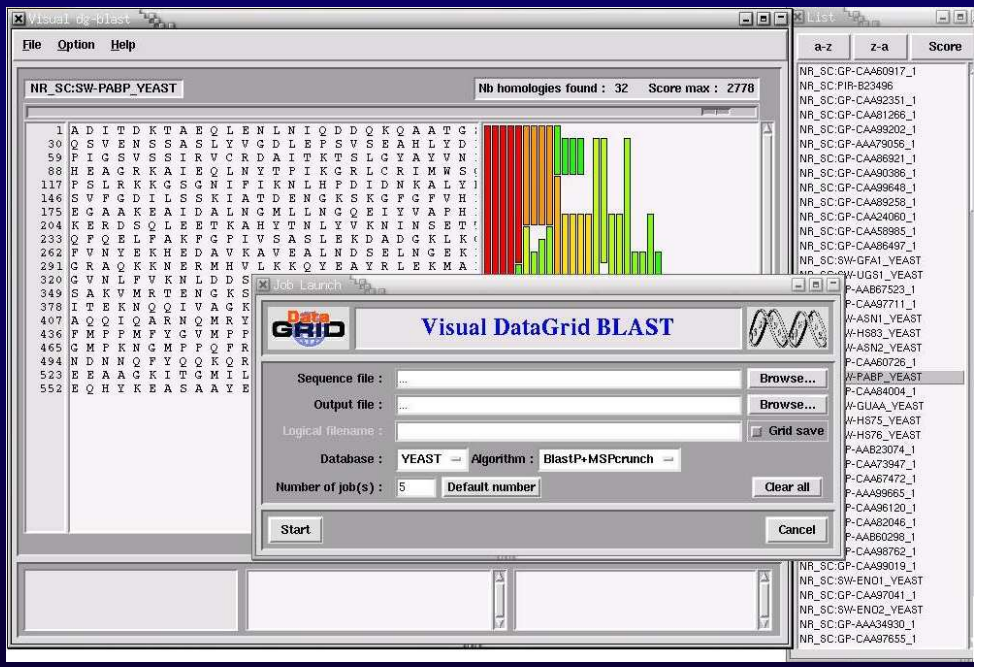

Summary of added value provided by Grid for BioMed applications

- Data mining on genomics databases (exponential growth).
- Indexing of medical databases (Tb/hospital/year).
- Collaborative framework for large scale experiments (e.g. epidemiological studies).
- Parallel processing for
  - Databases analysis
  - Complex 3D modelling

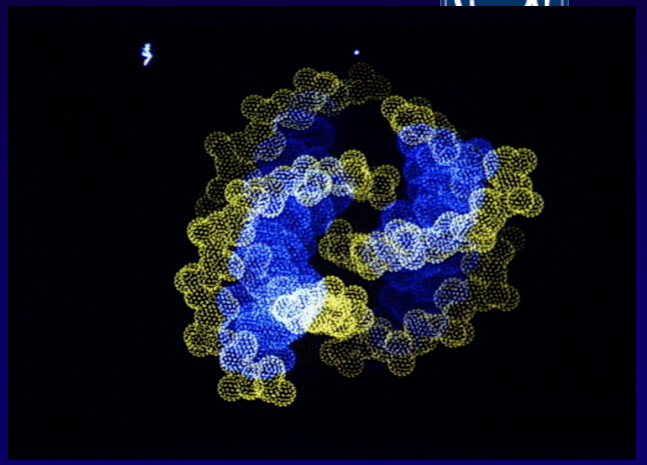

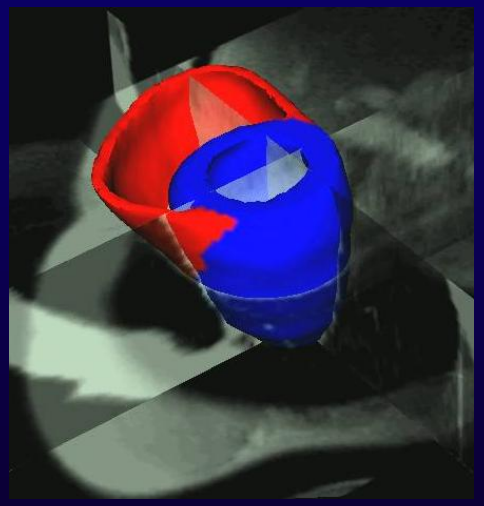

Conclusions

- Grid or Grid-like systems are clearly needed
- EDG is a start that has to be followed up
- EDG is nowhere near to be the “real thing”
- Currently key focus is resilience and scalability

References

- Some interesting WEB sites and documents
  - LHC Review http://lhc-computing-review-public.web.cern.ch/lhc-computing-review-public/Public/Report_final.PDF (LHC Computing Review)
  - LCG http://lcg.web.cern.ch/LCG
    http://lcg.web.cern.ch/LCG/SC2/RTAG6 (model for regional centres)
    http://lcg.web.cern.ch/LCG/SC2/RTAG4 (HEPCAL Grid use cases)
  - GEANT http://www.dante.net/geant/ (European Research Networks)
  - POOL http://lcgapp.cern.ch/project/persist/
  - WP8 http://datagrid-wp8.web.cern.ch/DataGrid-WP8/
    http://edmsoraweb.cern.ch:8001/cedar/doc.info?document_id=332409 (Requirements)
  - WP9 http://styx.srin.esa.it/grid
    http://edmsoraweb.cern.ch:8001/cedar/doc.info?document_id=332411 (Reqts)
  - WP10 http://marianne.in2p3.fr/datagrid/wp10/
    http://www.healthgrid.org
    http://www.creatis.insa-lyon.fr/MEDIGRID/
    http://edmsoraweb.cern.ch:8001/cedar/doc.info?document_id=332412 (Reqts)

The End
(finally)