Applications and the Grid

F Harris (Oxford/CERN) WP8

- **An applications view of the the Grid**
- **Current models for use of the Grid in**
  - **High Energy Physics (WP8)**
  - **Biomedical Applications (WP10)**
  - **Earth Observation Applications (WP9)**
- **Summary and a forward look for applications**
- **Acknowledgments and references**

GRID Services: The Overview

Chemistry

Cosmology

Environment

Biology

High Energy Physics

- Distributed computing toolkit
- Data-intensive applications toolkit
- Collaborative applications toolkit
- Remote Visualisation applications toolkit
- Problem solving applications toolkit
- Remote instrumentation applications toolkit

Resource-independent and application-independent services

authentication, authorisation, resource location, resource allocation, events, accounting, **remote data access**, information, policy, fault detection

Resource-specific implementations of basic services

E.g., transport protocols, name servers, differentiated services, CPU schedulers, public key infrastructure, site accounting, directory service, OS bypass

What all applications want from the Grid (the basics)

- A homogeneous way of looking at a ‘virtual computing lab’ made up of heterogeneous resources as part of a VO(Virtual Organisation) which manages the allocation of resources to authenticated and authorised users
  - A uniform way of ‘logging on’ to the Grid
  - Basic functions for job submission, data management and monitoring
    - Ability to obtain resources (services) satisfying user requirements for data, CPU, software, turnaround……

LHC Computing (a hierachical view of grid…this has evolved to a ‘cloud’ view)

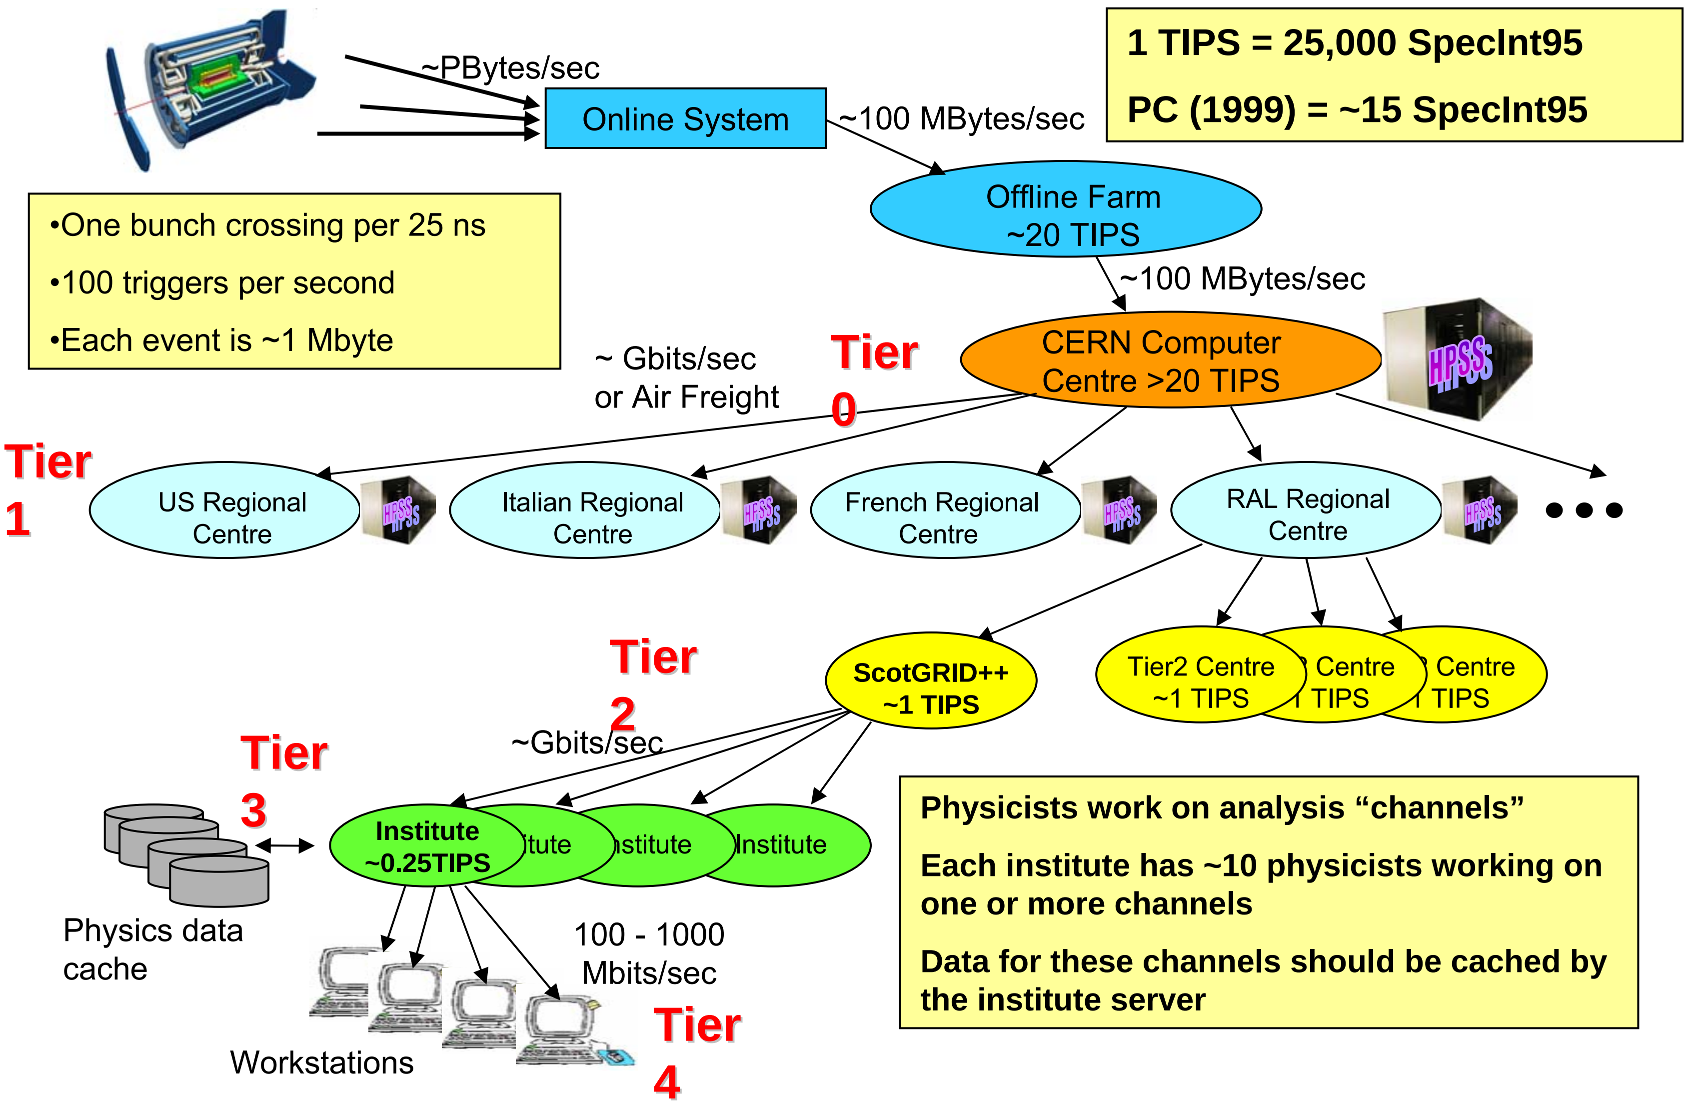

Data Handling and Computation for Physics Analysis

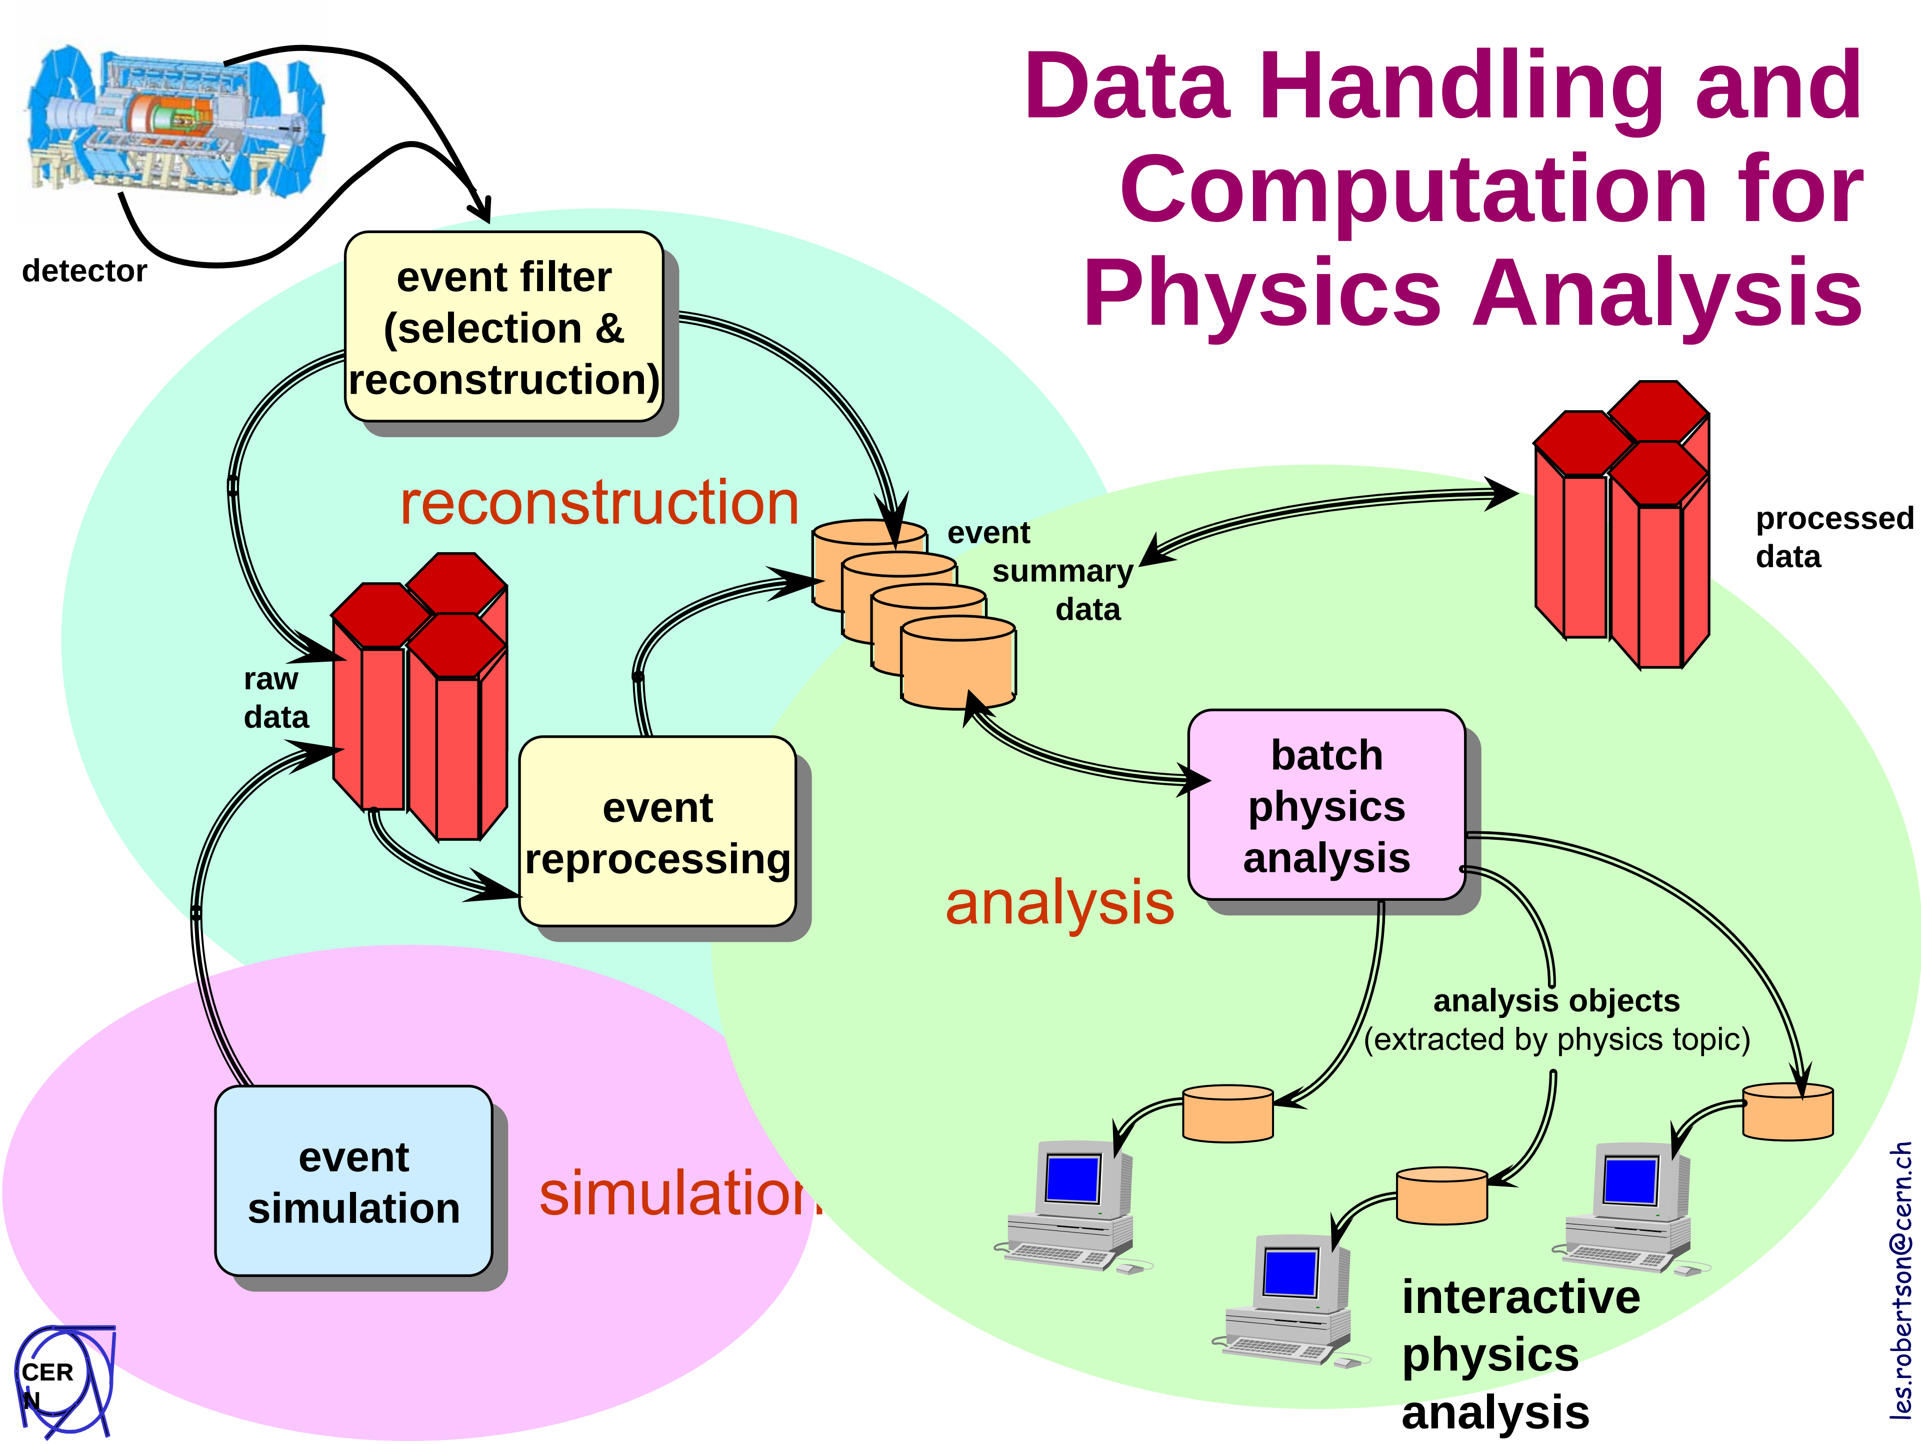

les.robertson@cern.ch

HEP Data Analysis and Datasets

- **Raw data (RAW)** **~ 1 MByte**
  - **hits, pulse heights**
- **Reconstructed data (ESD) ~ 100 kByte**
  - **tracks, clusters…**
- **Analysis Objects (AOD)** **~ 10 kByte**
  - **Physics Objects**
  - **Summarized**
  - **Organized by physics topic**
- **Reduced AODs(TAGs)** **~1 kByte**
- **histograms, statistical data on collections of events**

ATLAS Barrel Inner Detector

$H \rightarrow b\bar{b}$

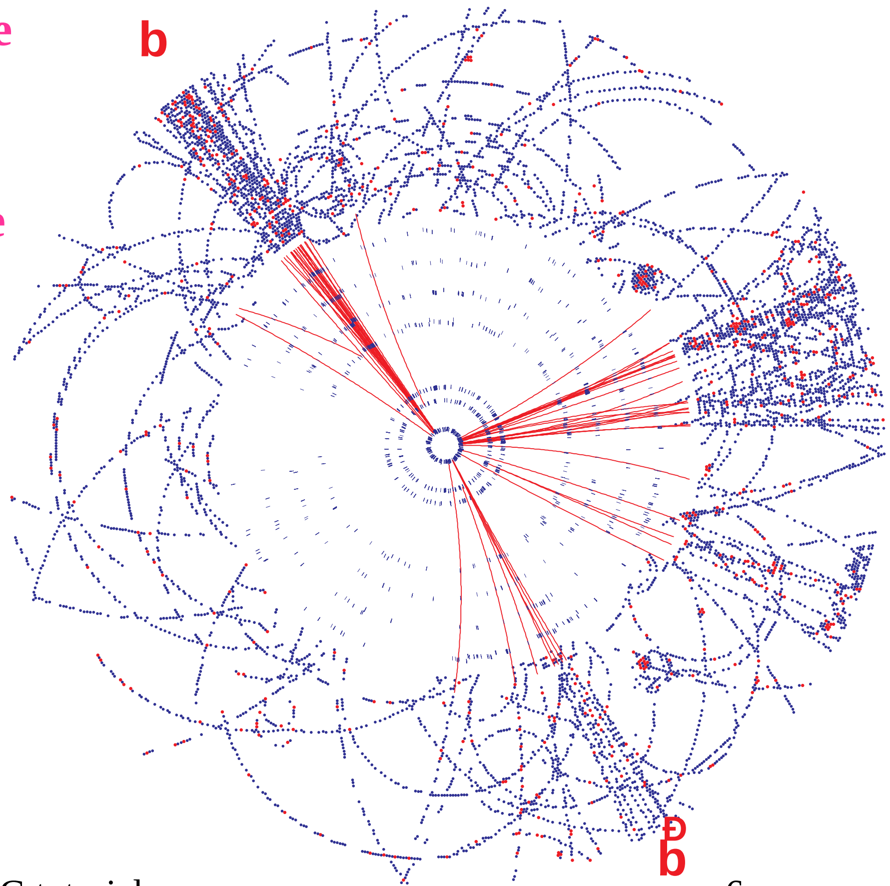

HEP Data Analysis –processing patterns

- **Processing fundamentally parallel** due to independent nature of ‘events’
  - So have concepts of splitting and merging
  - Processing organised into ‘jobs’ which process N events
    - (e.g. simulation job organised in groups of ~500 events which takes ~ day to complete on one node)
      - A processing for 10**6 events would then involve 2,000 jobs merging into total set of 2 Tbyte
- **Production processing is planned** by experiment and physics group data managers(this will vary from expt to expt)
  - Reconstruction processing (1-3 times a year of 10**9 events)
  - Physics group processing (? 1/month). Produce ~10**7 AOD+TAG
  - This may be distributed in several centres

Processing Patterns(2)

- **Individual physics analysis - by definition 'chaotic' (according to work patterns of individuals)**
  - **Hundreds of physicists distributed in expt may each want to access central AOD+TAG and run their own selections . Will need very selective access to ESD+RAW data (for tuning algorithms, checking occasional events)**
- **Will need replication of AOD+TAG in experiment, and selective replication of RAW+ESD**
  - **This will be a function of processing and physics group organisation in the experiment**

A Logical View of Event Data for physics analysis

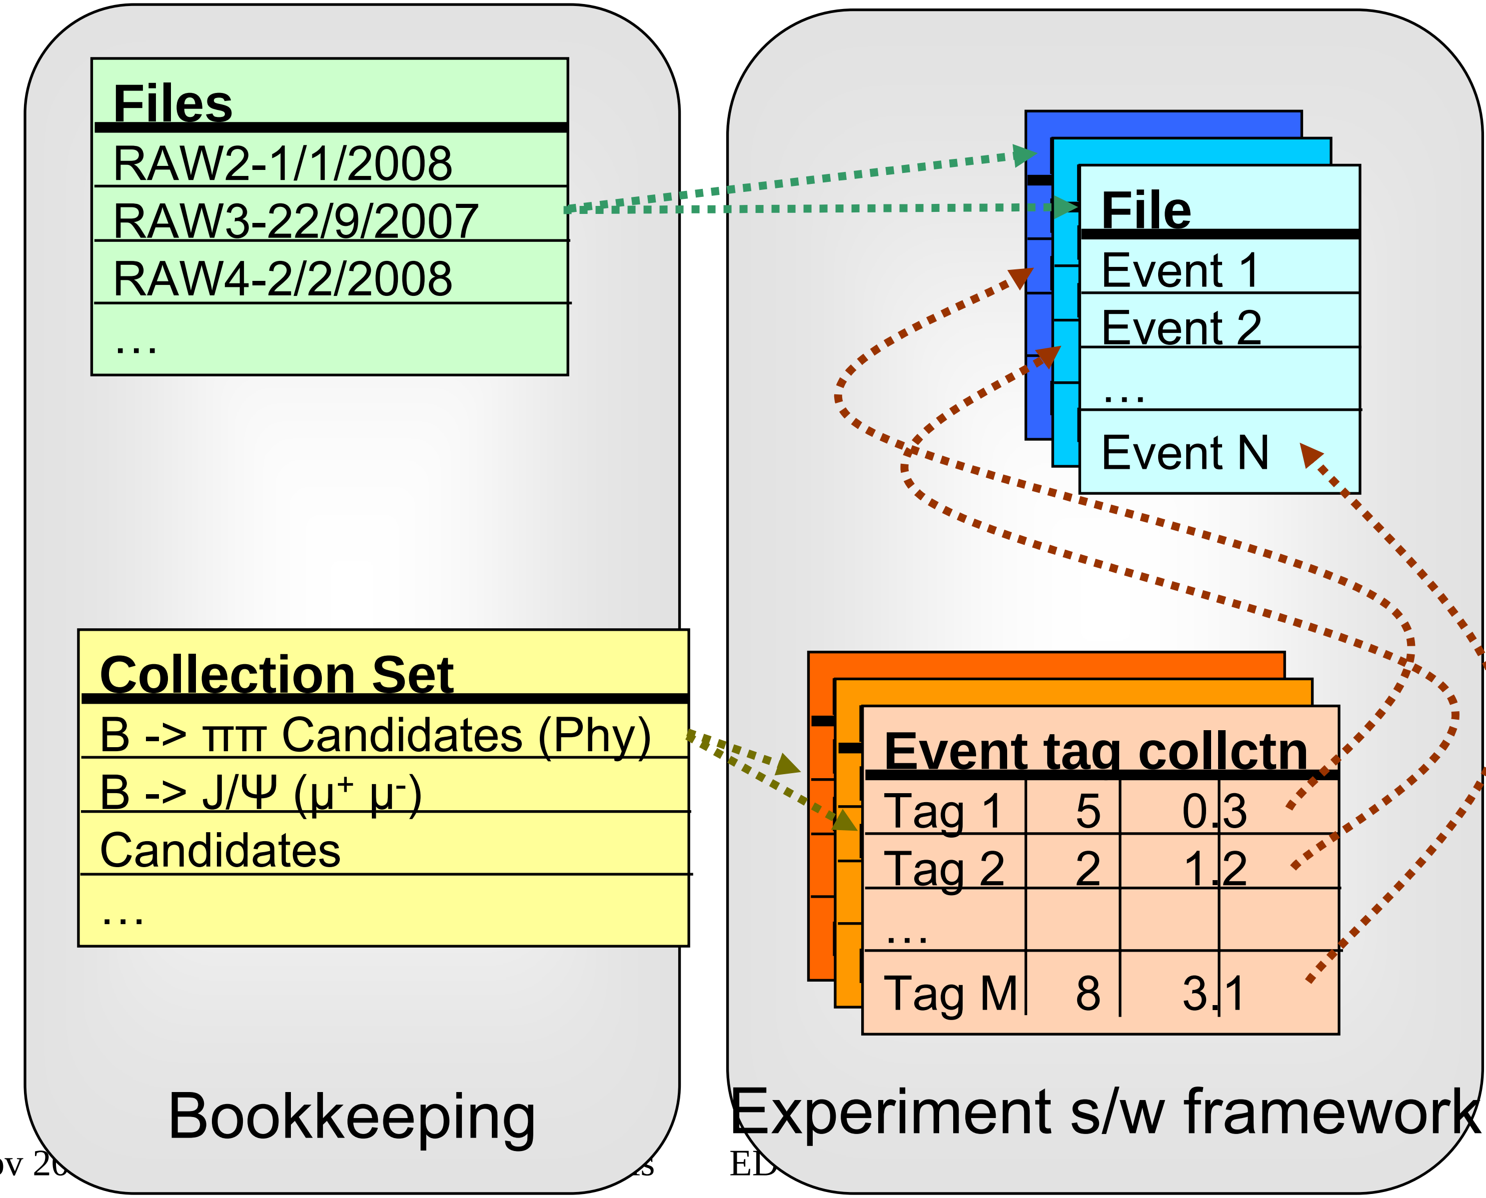

LCG/Pool on the Grid

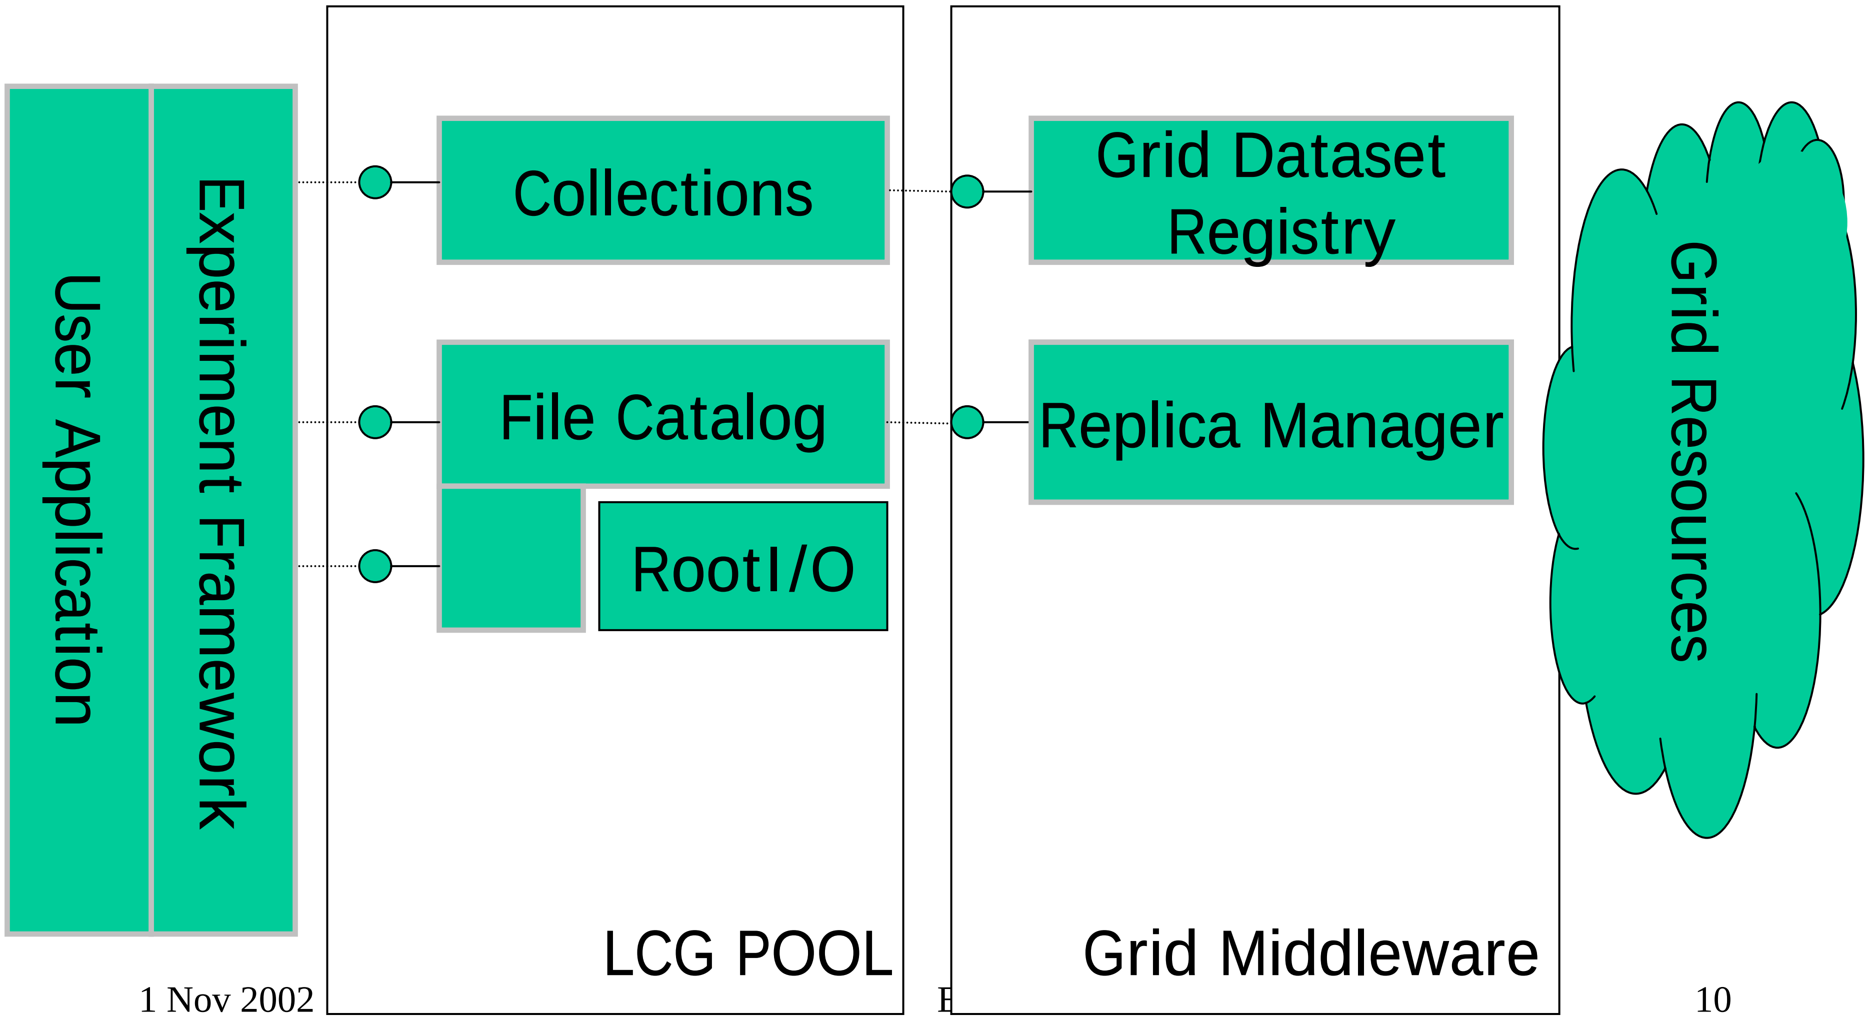

An implementation of distributed analysis in ALICE using natural parallelism of processing

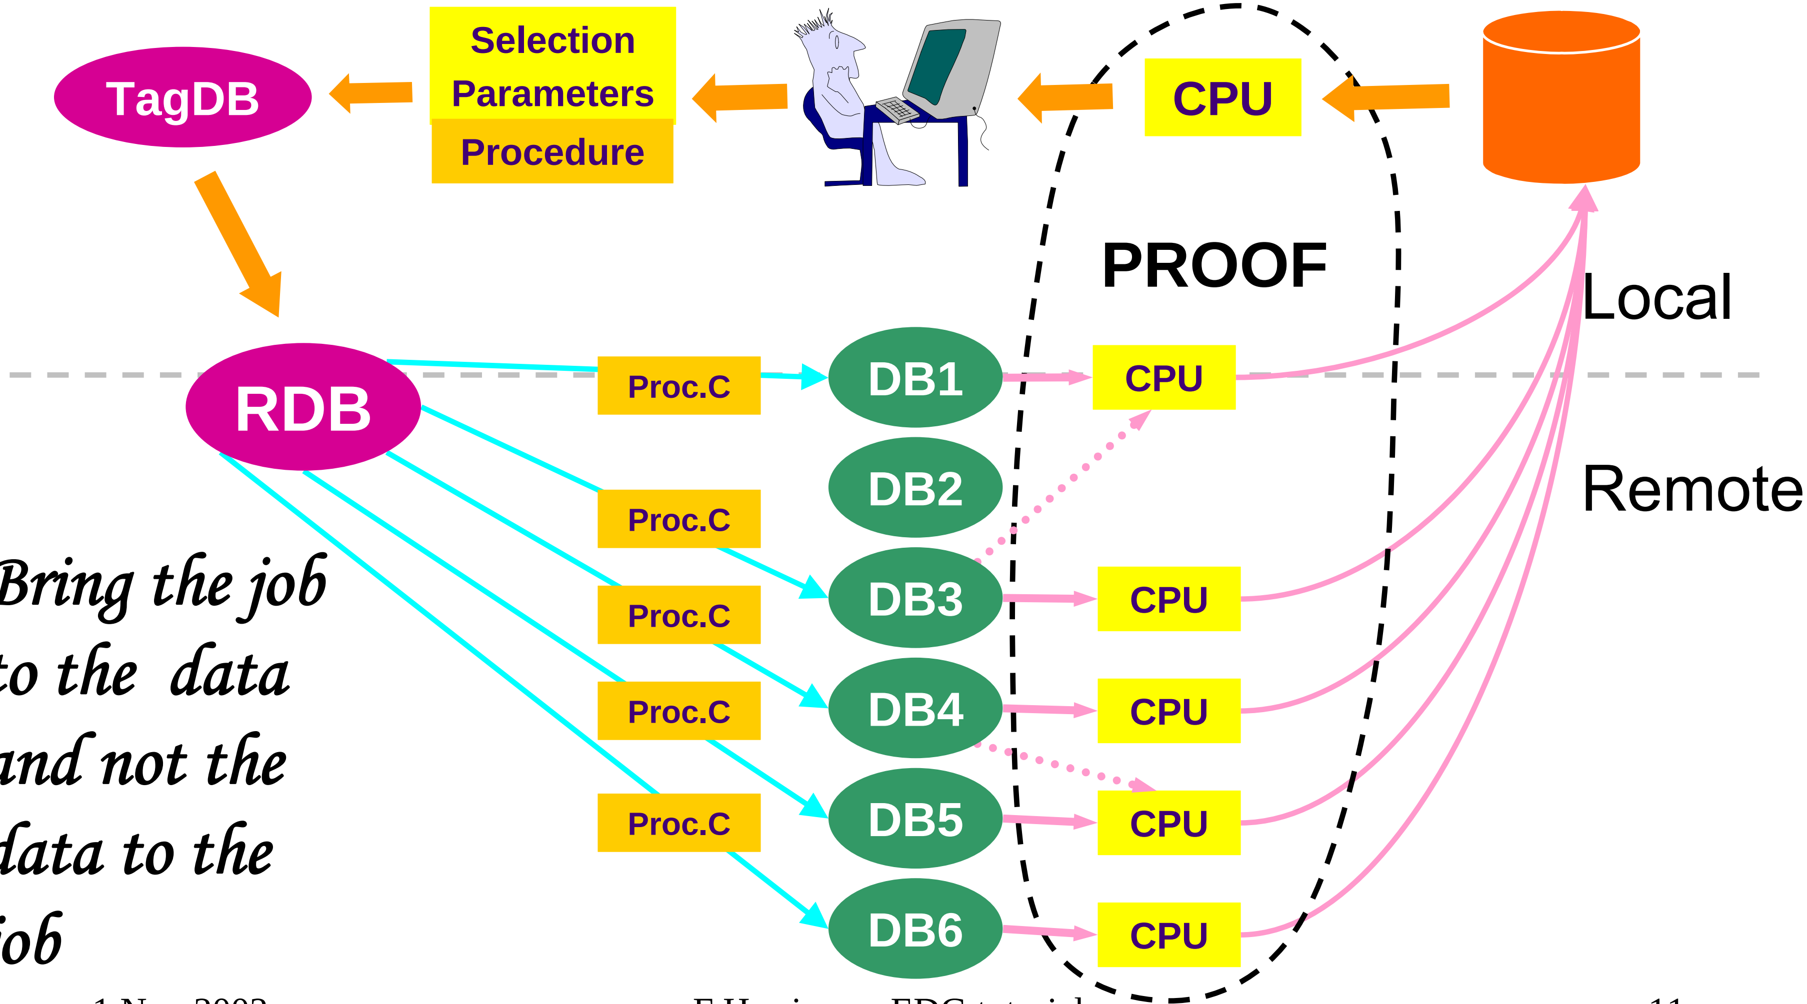

Bring the job to the data and not the data to the job

ALICE production distributed Environment

- **ALICE production distributed Environment**
  - **Entirely ALICE developed**
- **File Catalogue as a global file system on a RDB**
- **TAG Catalogue, as extension**
- **Secure Authentication**
  - **Interface to Globus available**
- **Central Queue Manager ("pull" vs "push" model)**
  - **Interface to EDG Resource Broker available**
- **Monitoring infrastructure**

The CORE GRID functionality

- **Automatic software installation with AliKit**
- **Being interfaced to EDG and iVDGL(US Testbed)**
- http://alien.cern.ch

ATLAS/LHCb Software Framework (Based on Services)

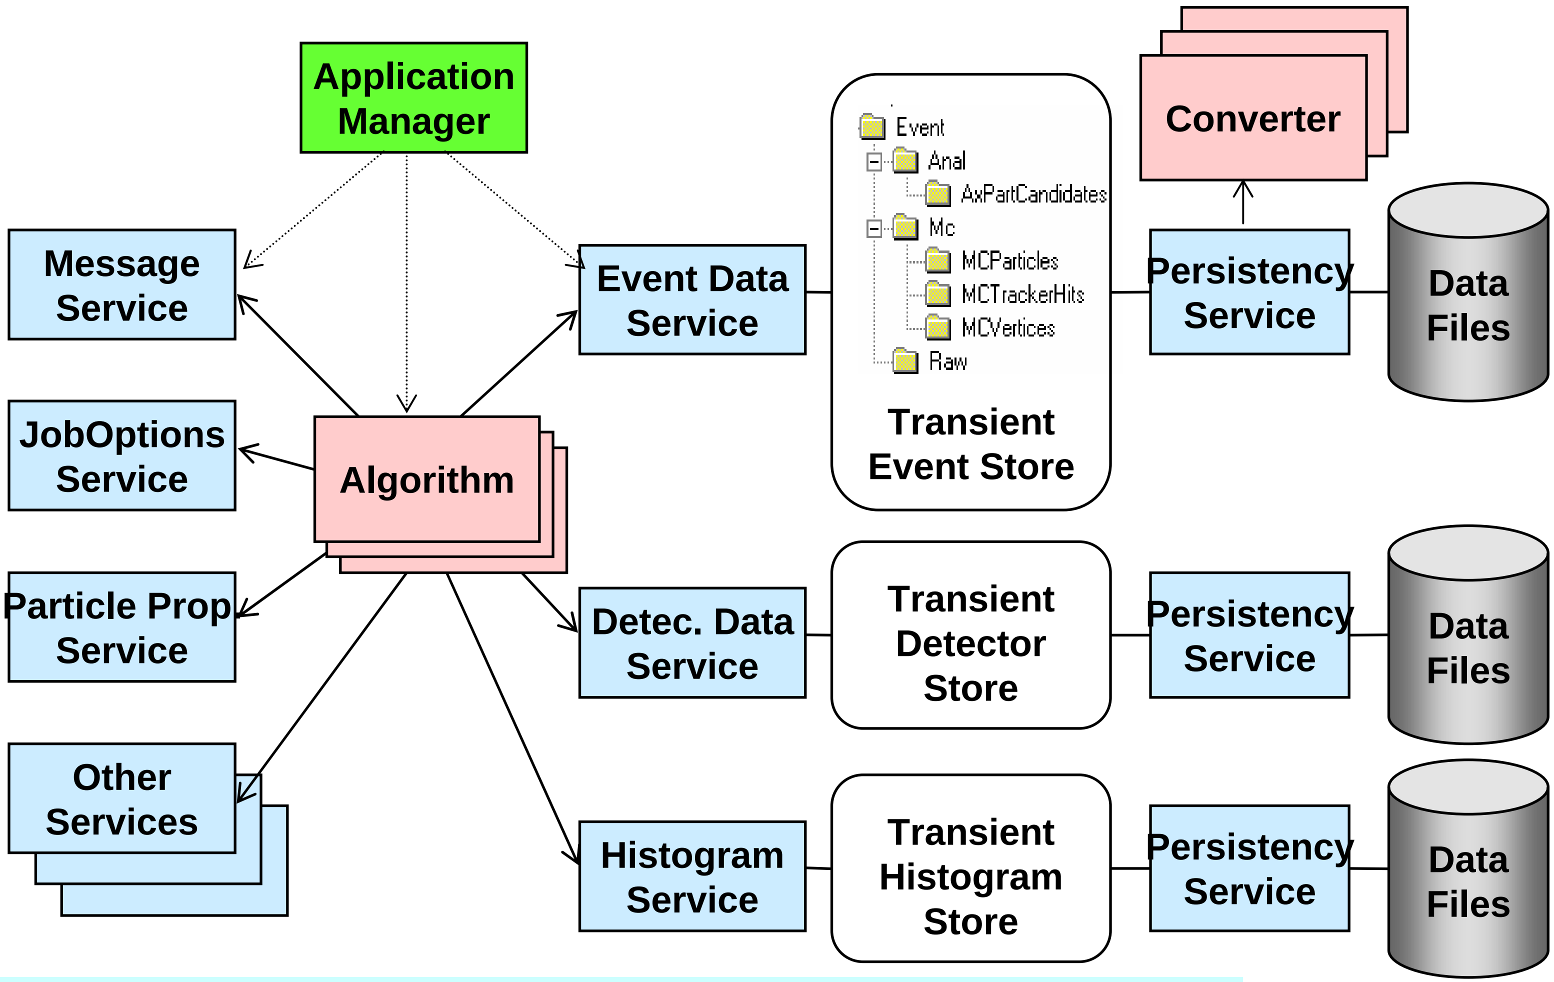

The Gaudi/Athena Framework – Services will interface to Grid (e.g. Persistency)

GANGA: Gaudi ANd Grid Alliance

Joint Atlas/LHCb project

- Application facilitating end-user physicists and production managers the use of Grid services for running Gaudi/Athena jobs.

- a GUI based application that should help for the complete job life-time:
  - job preparation and configuration
  - resource booking
  - job submission
  - job monitoring and control

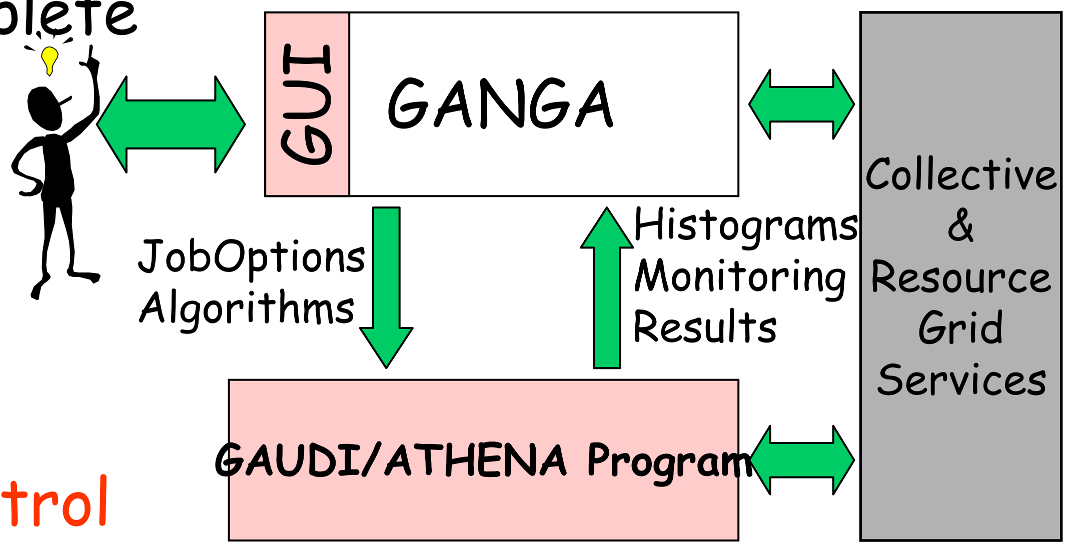

A CMS Data Grid Job
The vision for 2003

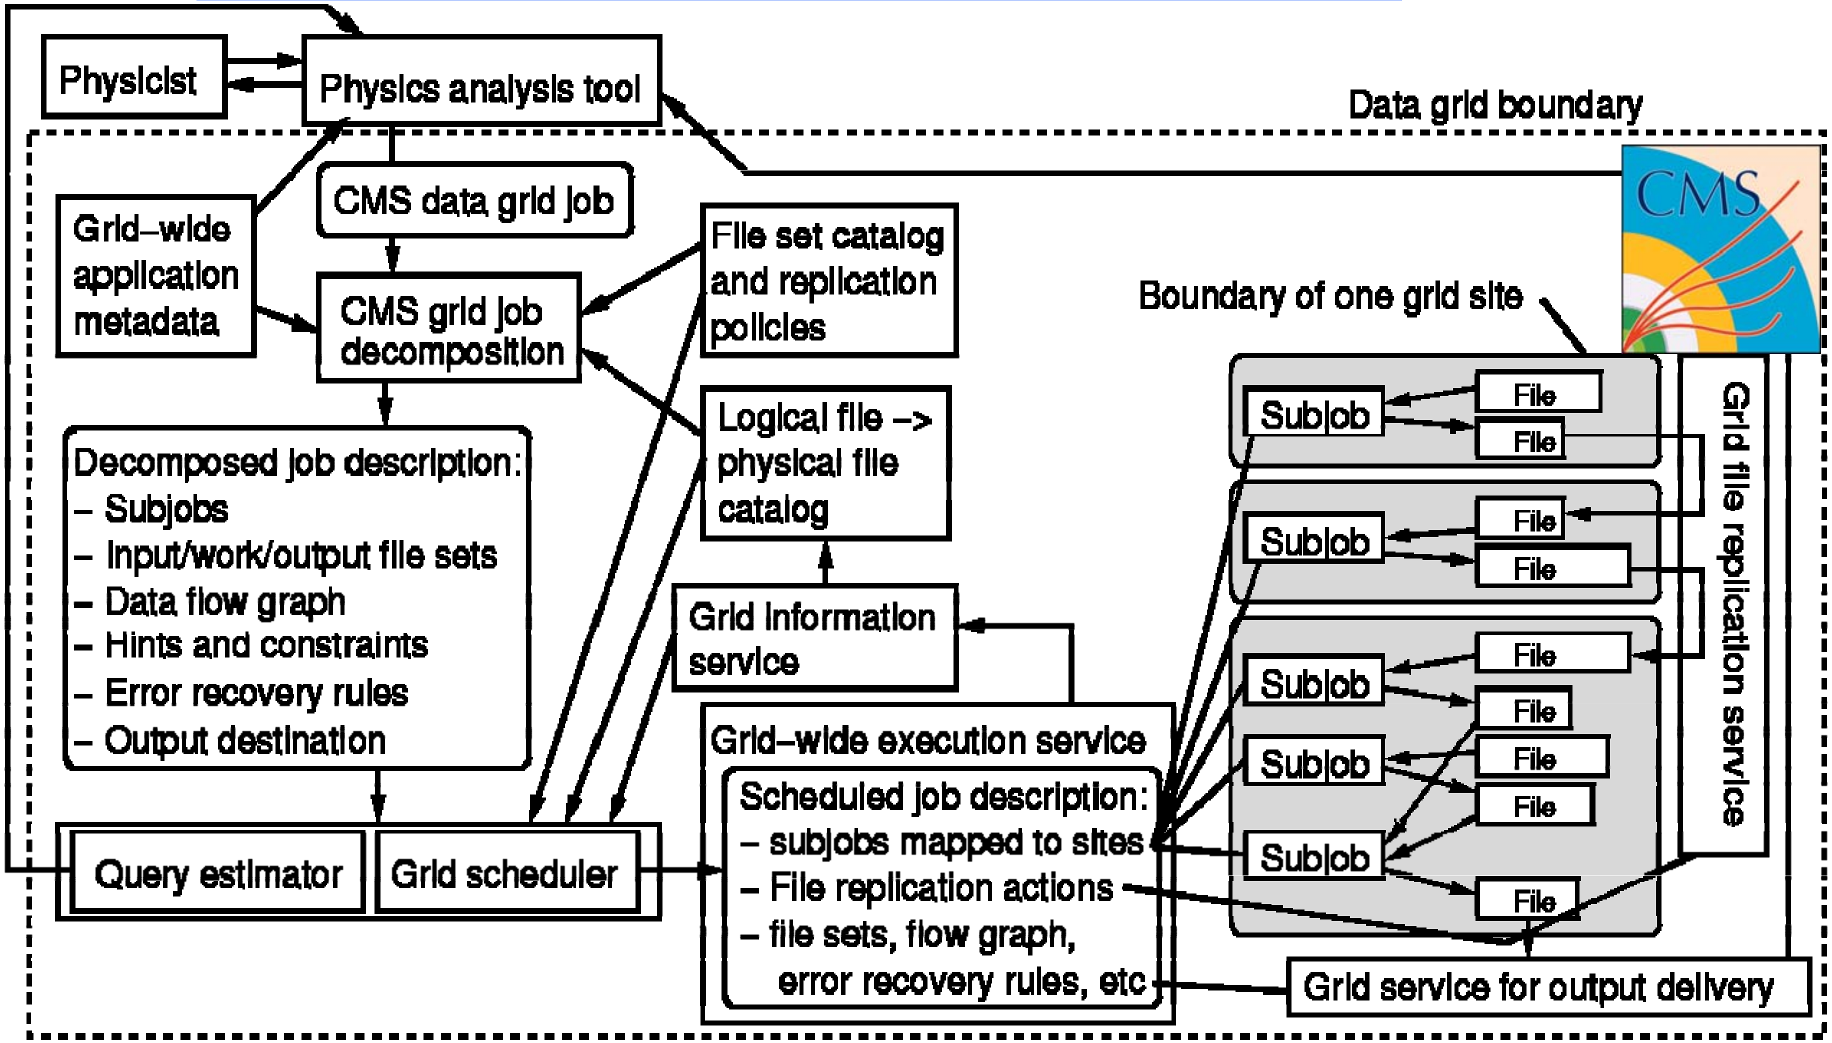

Deploying the LHC Global Grid Service

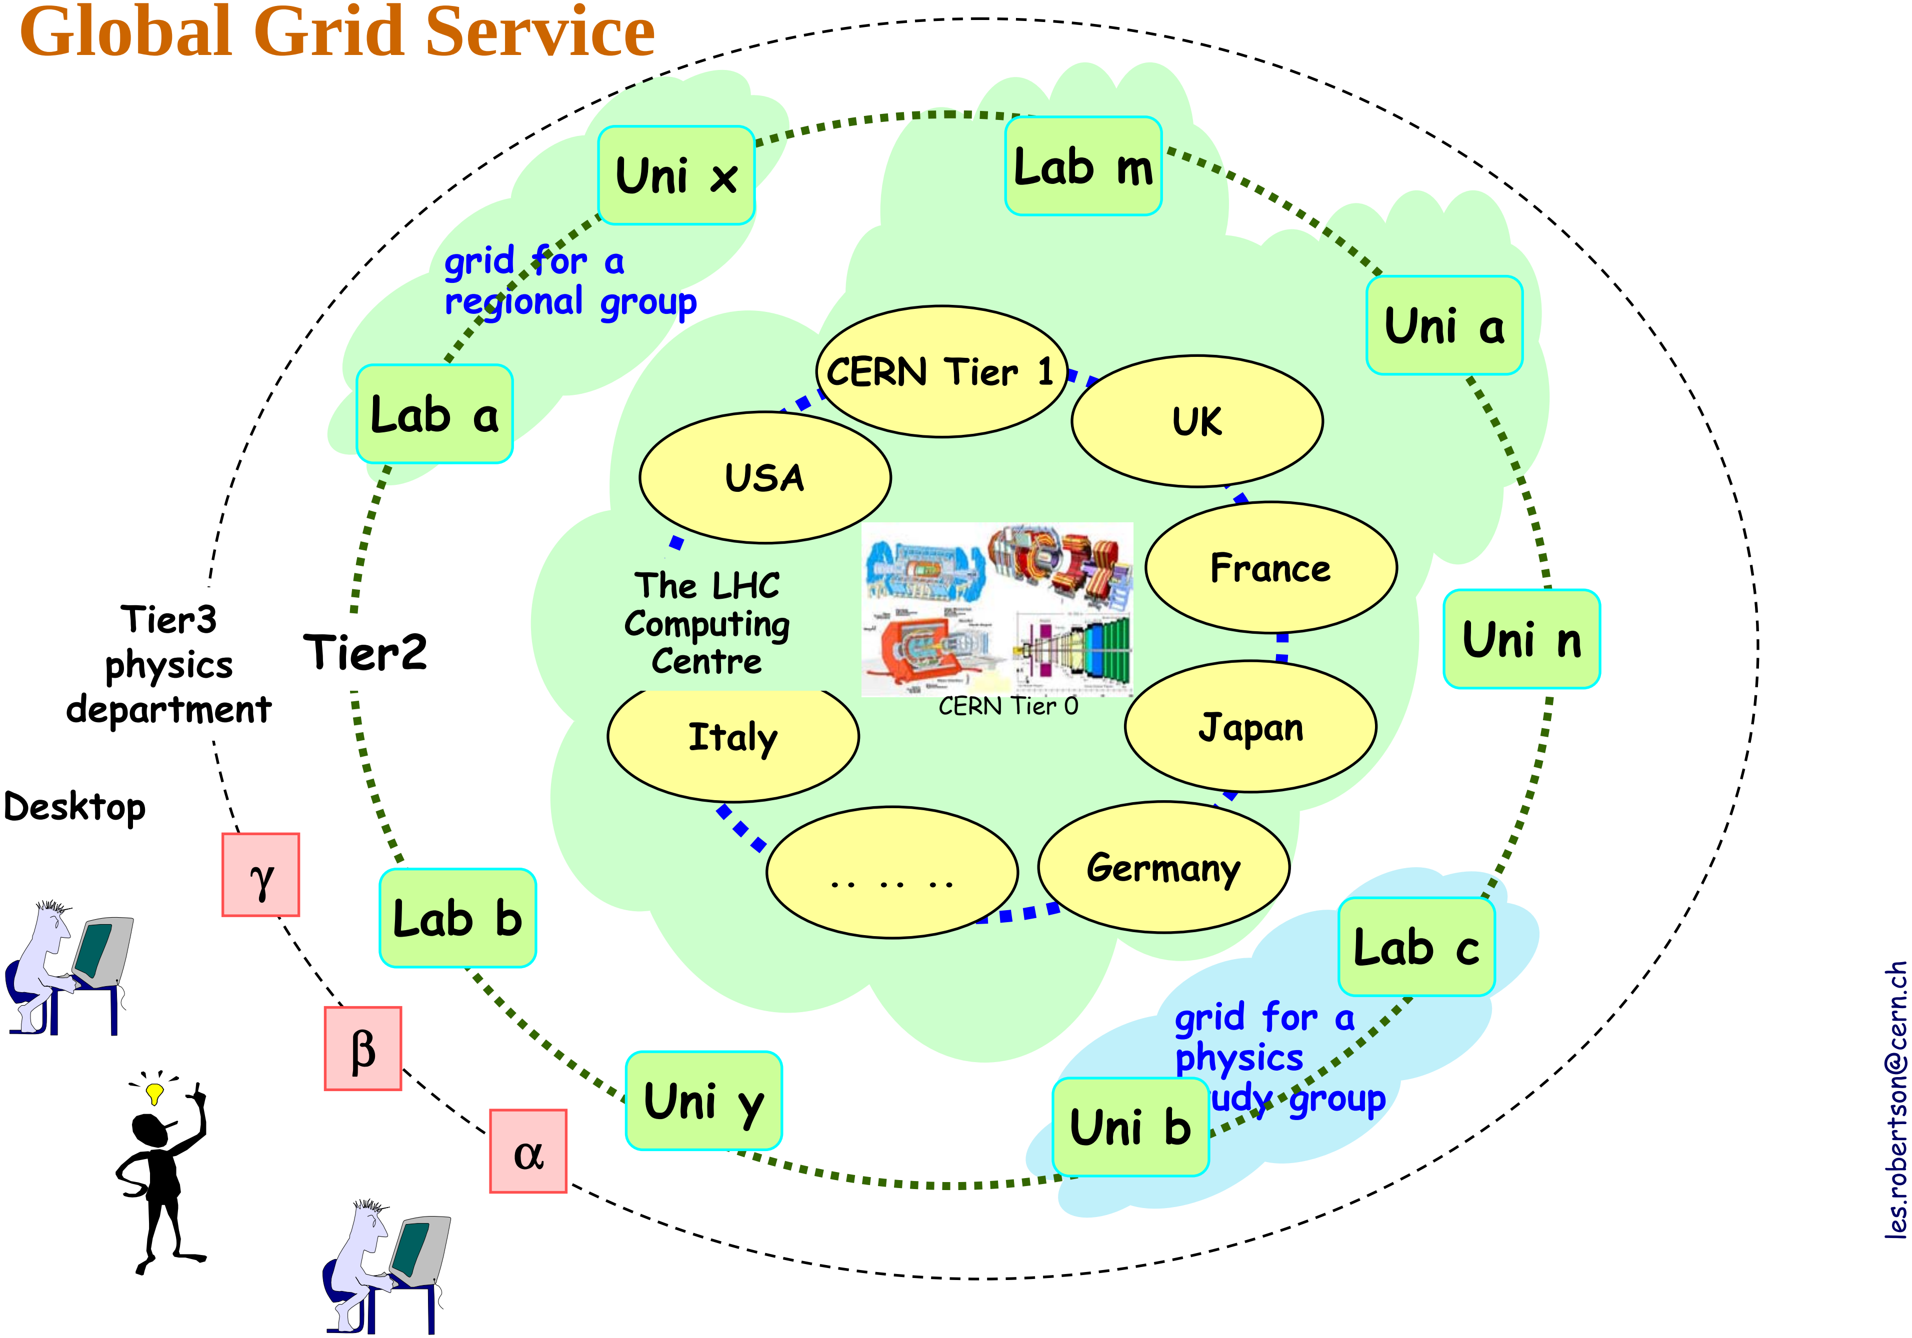

DataGrid Biomedical work package 10

- Grid technology opens the perspective of large computational power and easy access to heterogeneous data sources.
- A grid for health would provide a framework for sharing disk and computing resources, for promoting standards and fostering synergy between bio-informatics and medical informatics
- A first biomedical grid is being deployed by the DataGrid IST project

http://dbs.cordis.lu/fep-cgi/srchidadb?ACTION=D&SESSION=221592002-10-18&DOC=27&TBL=EN_PROJ&RCN=EP_RCN_A:63345&CALLER=PROJ_IST

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**

Biomedical projects in DataGrid

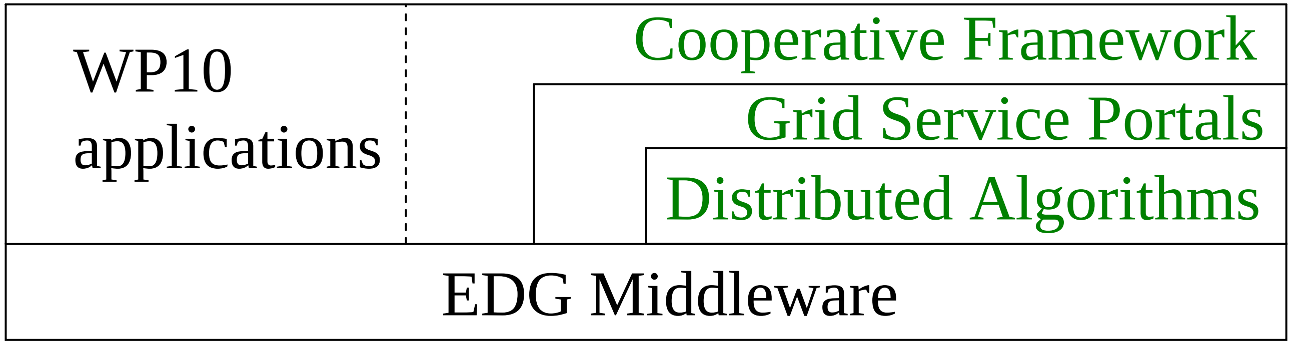

- **Distributed Algorithms**. New distributed "grid-aware" algorithms (bio-info algorithms, data mining, …)
- **Grid Service Portals**. Service providers taking advantage of the DataGrid computational power and storage capacity.
- **Cooperative Framework**. Use the DataGrid as a cooperative framework for sharing resources, algorithms, and organize experiments in a cooperative manner.

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**

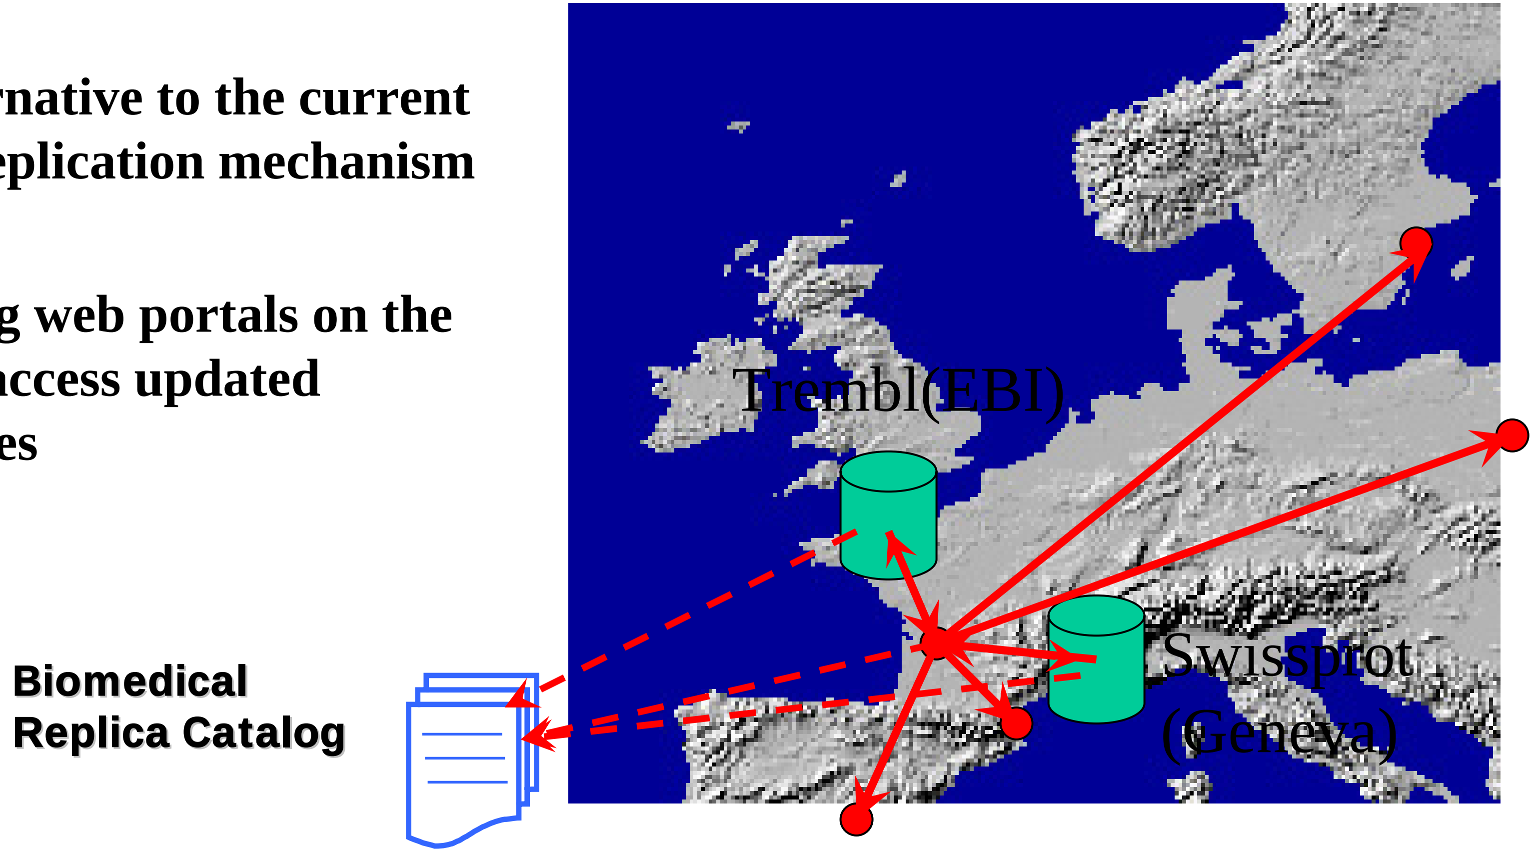

Web portals for biologists

- 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 first genomics application on DataGrid

- 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

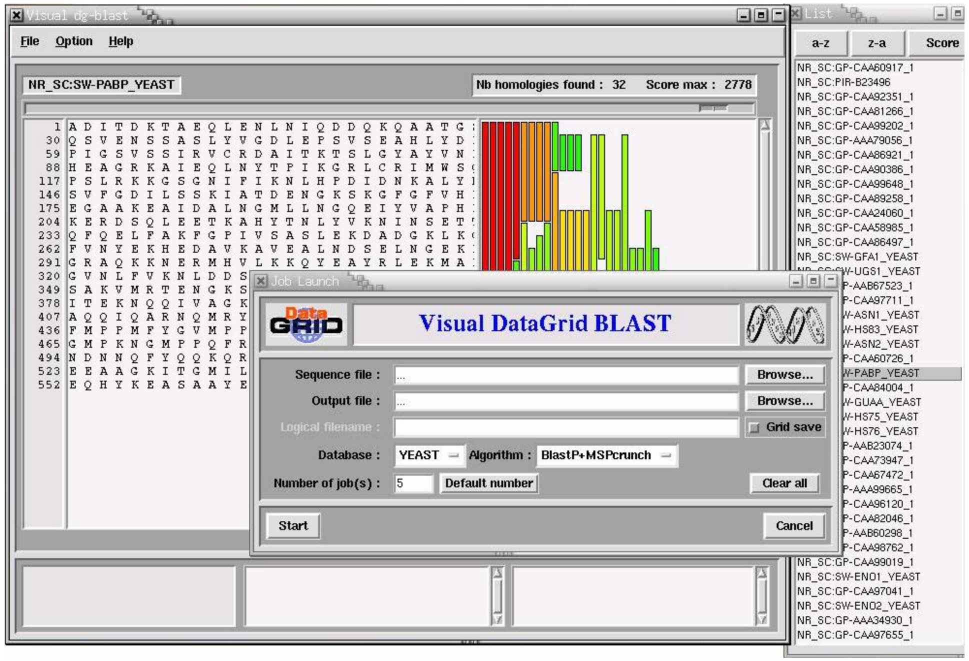

Summary of added value provided by Grid for BioMed applications

- 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

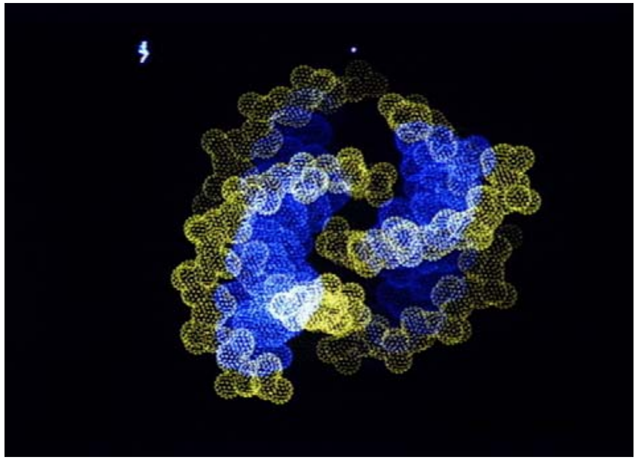

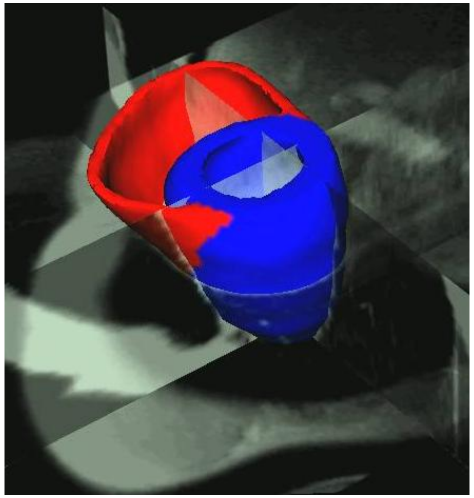

Creatis

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)

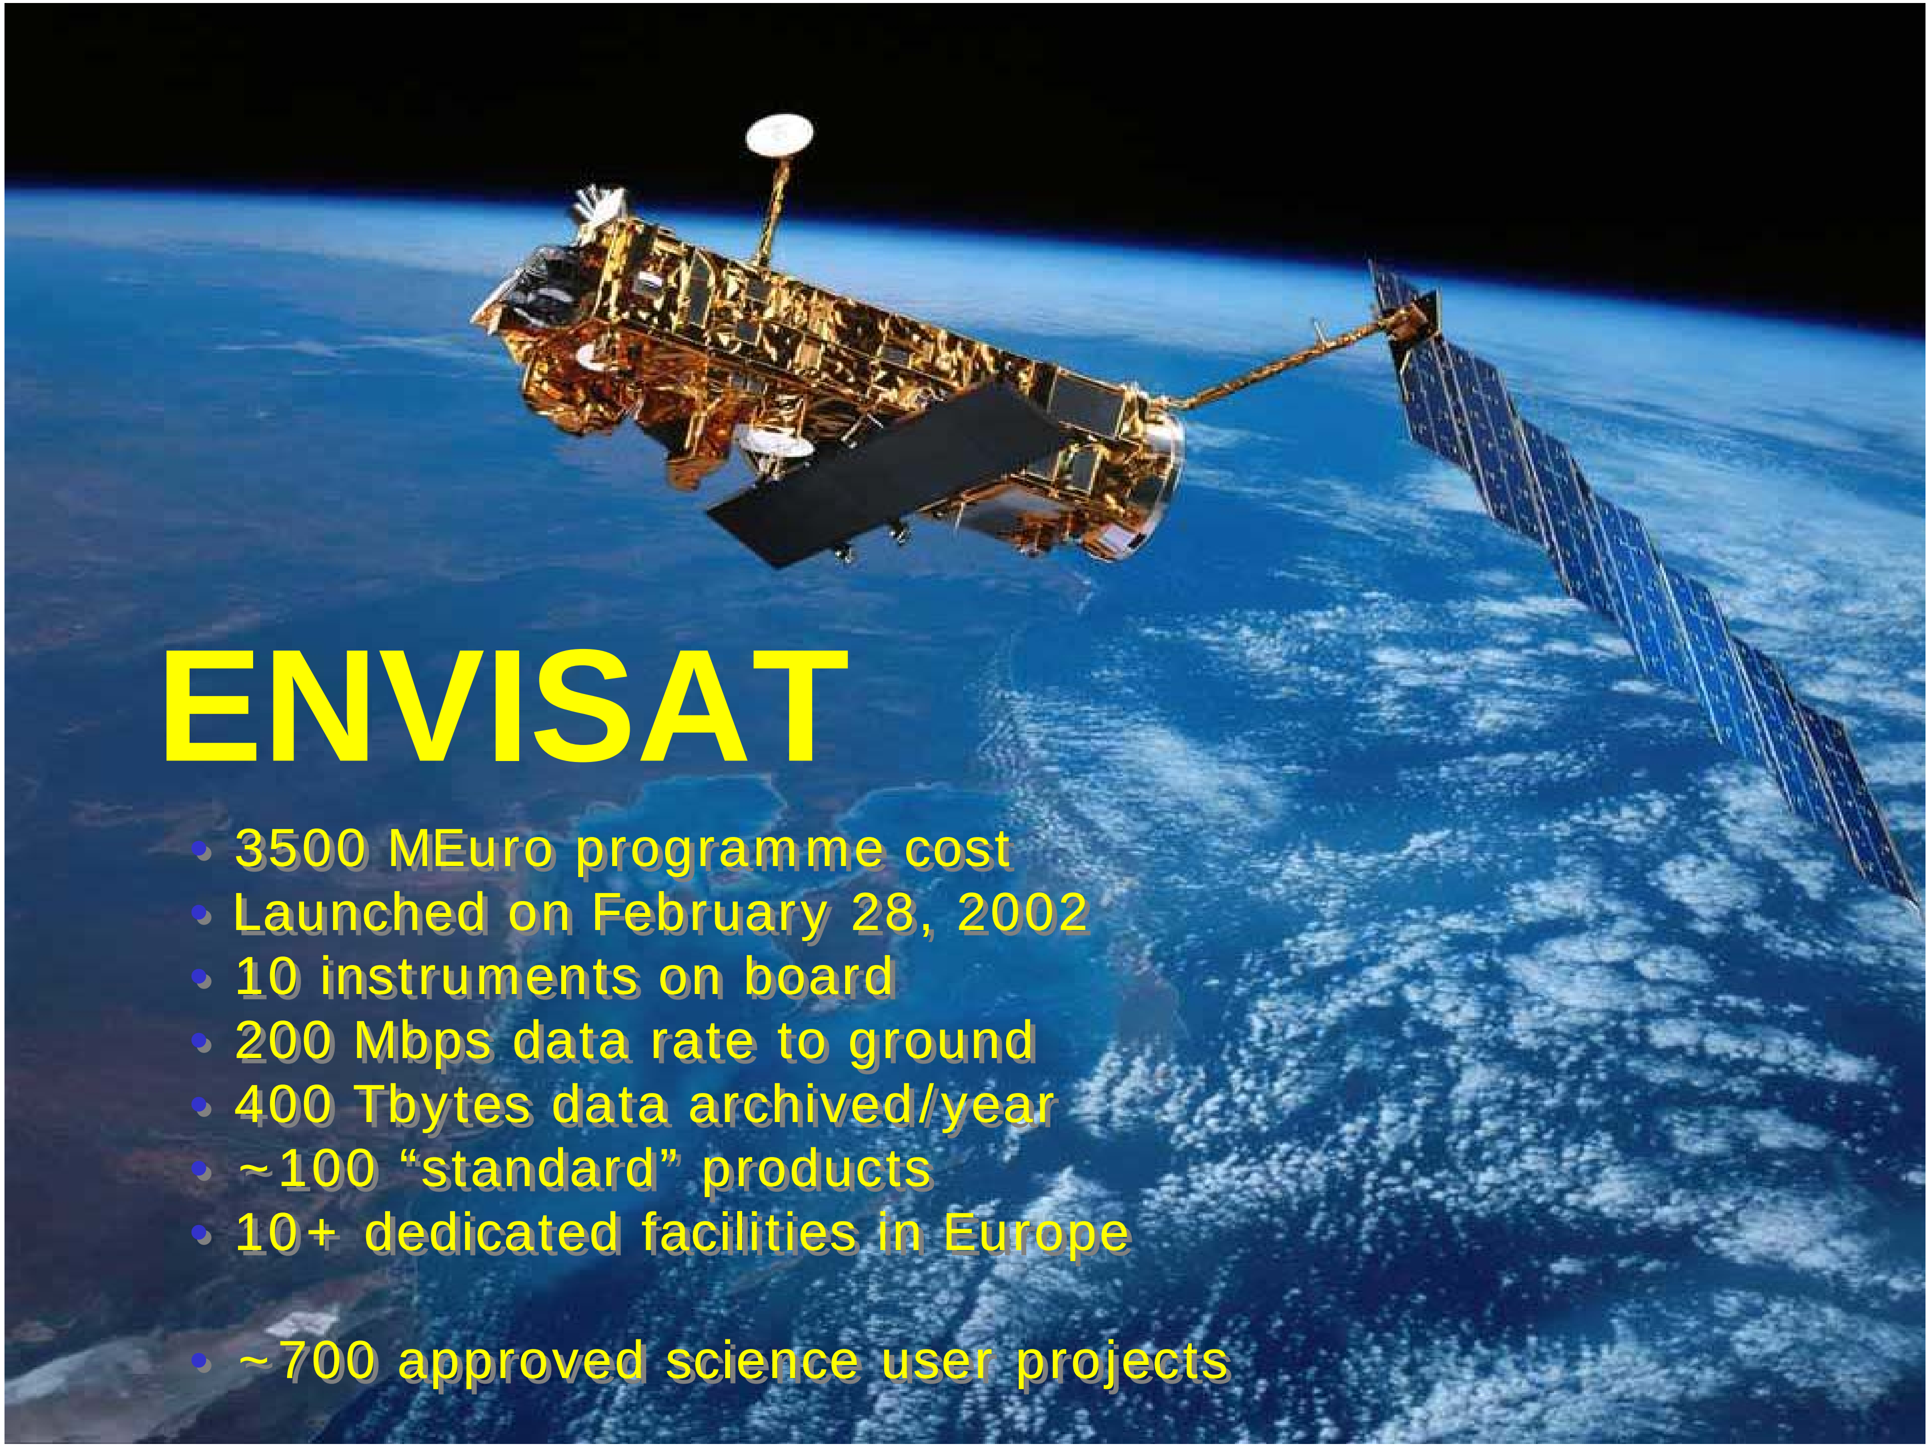
ENVISAT
3500 MEuro programme cost
Launched on February 28, 2002
10 instruments on board
200 Mbps data rate to ground
400 Tbytes data archived/year
~100 “standard” products
10+ dedicated facilities in Europe
~700 approved science user projects

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

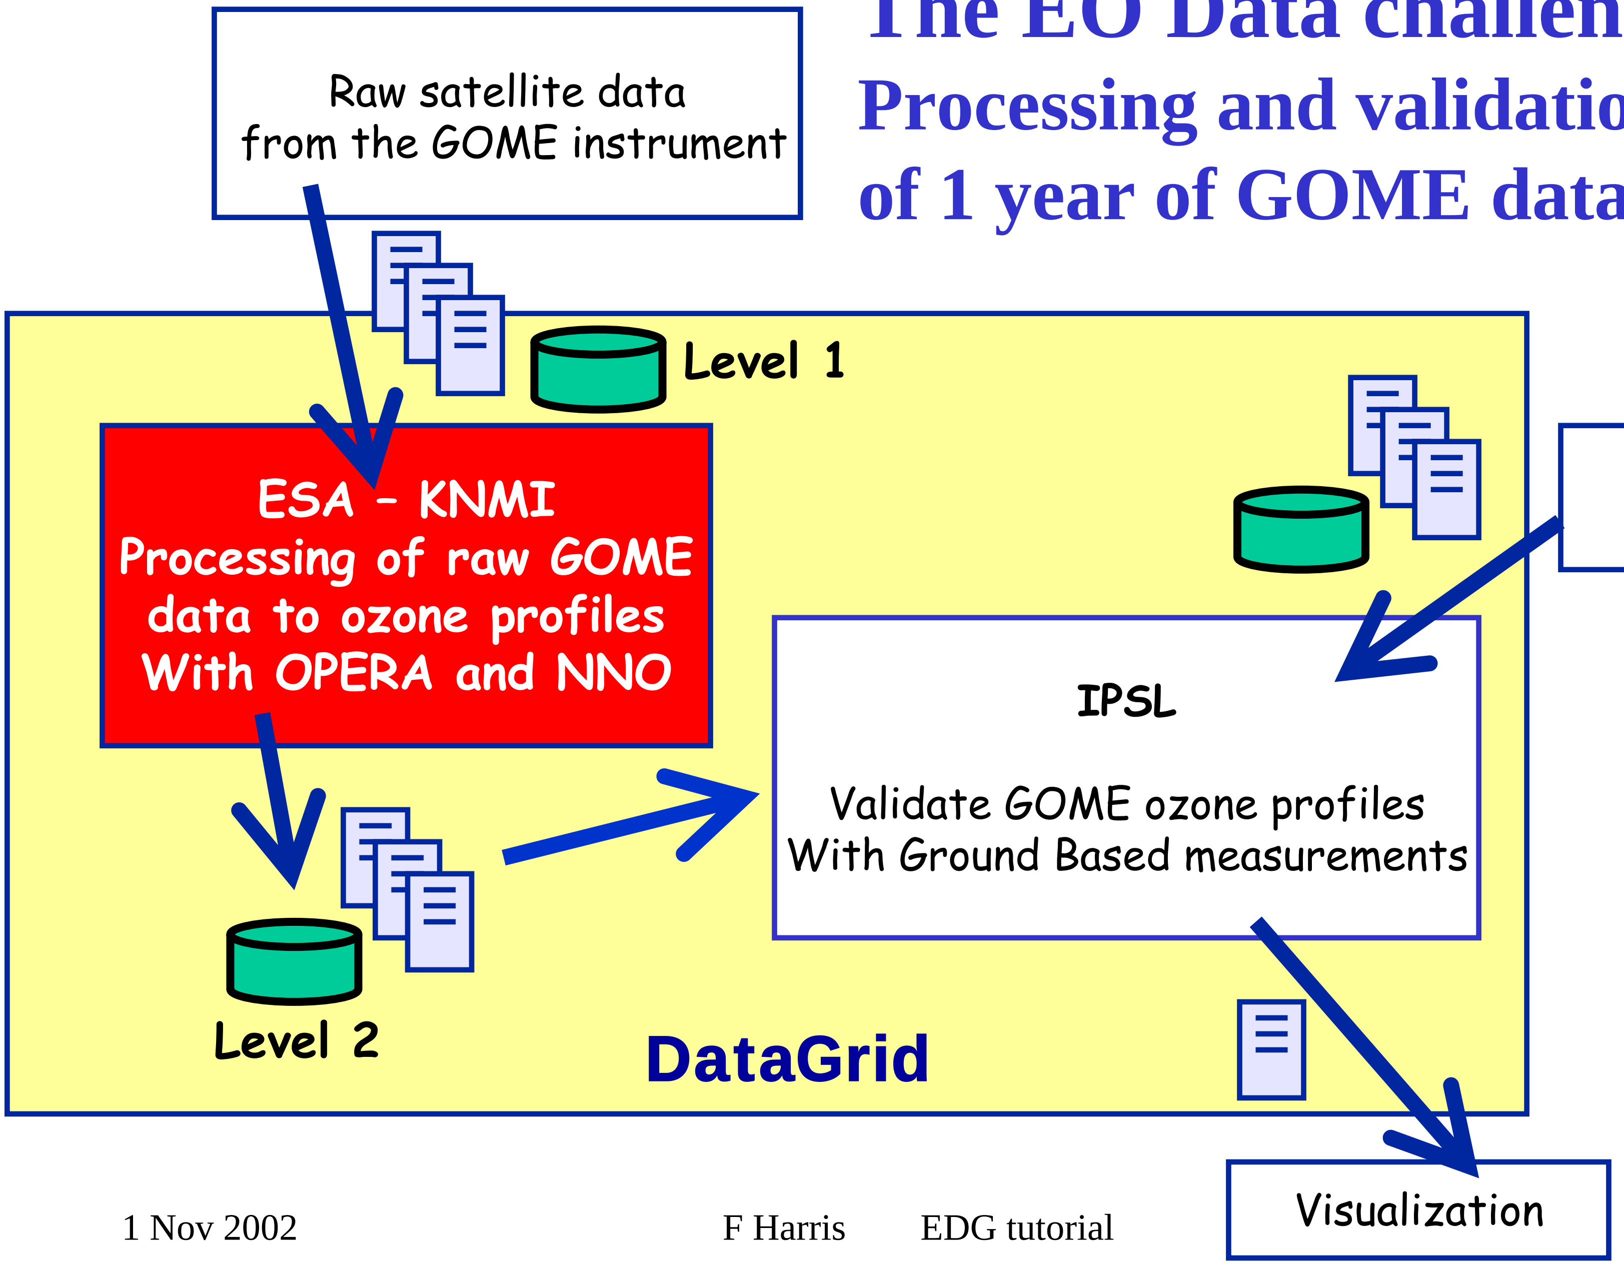

EO Use-Case File Numbers

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

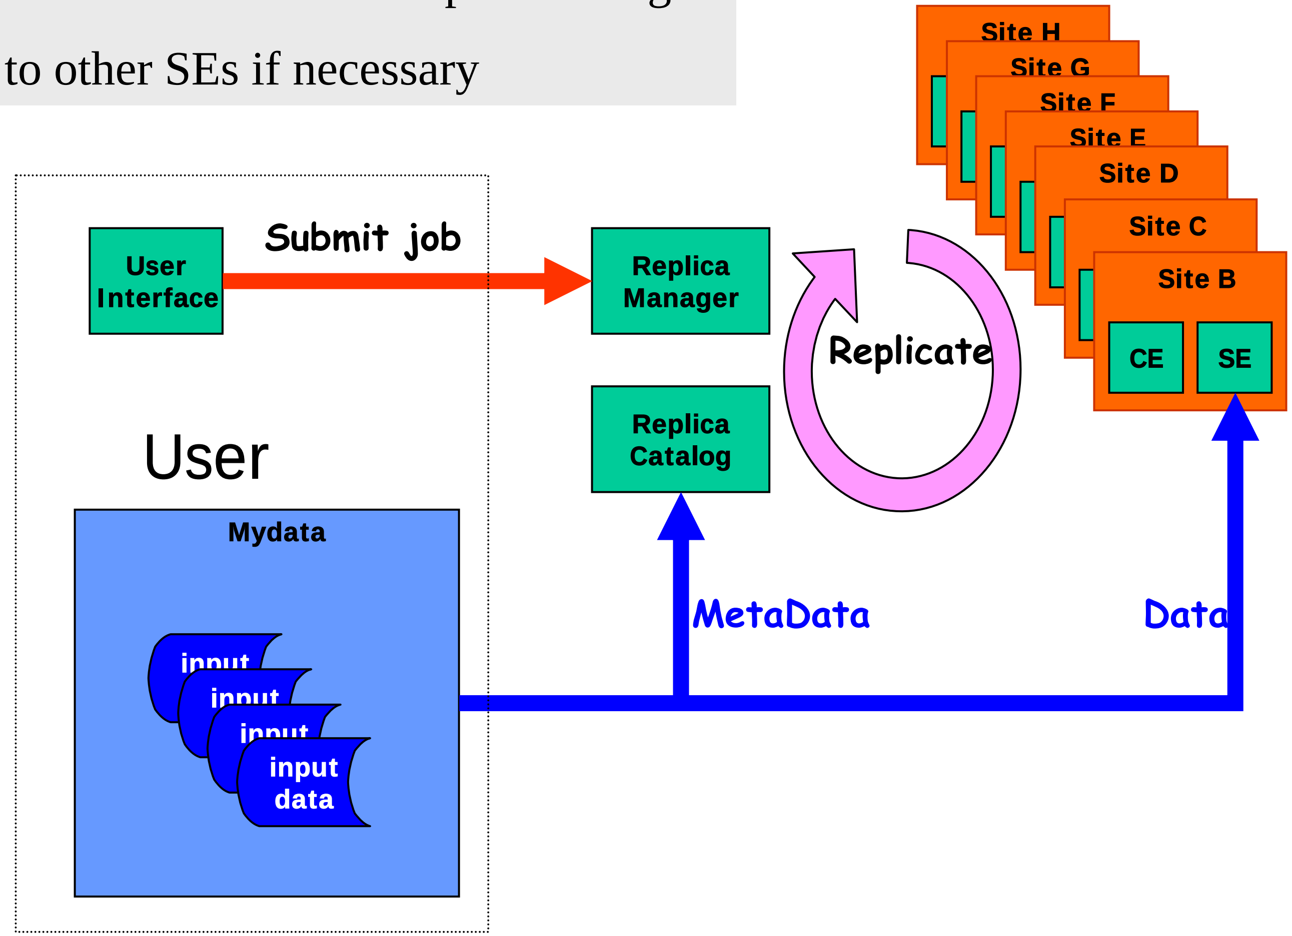

GOME Processing Steps (3-4)

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

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

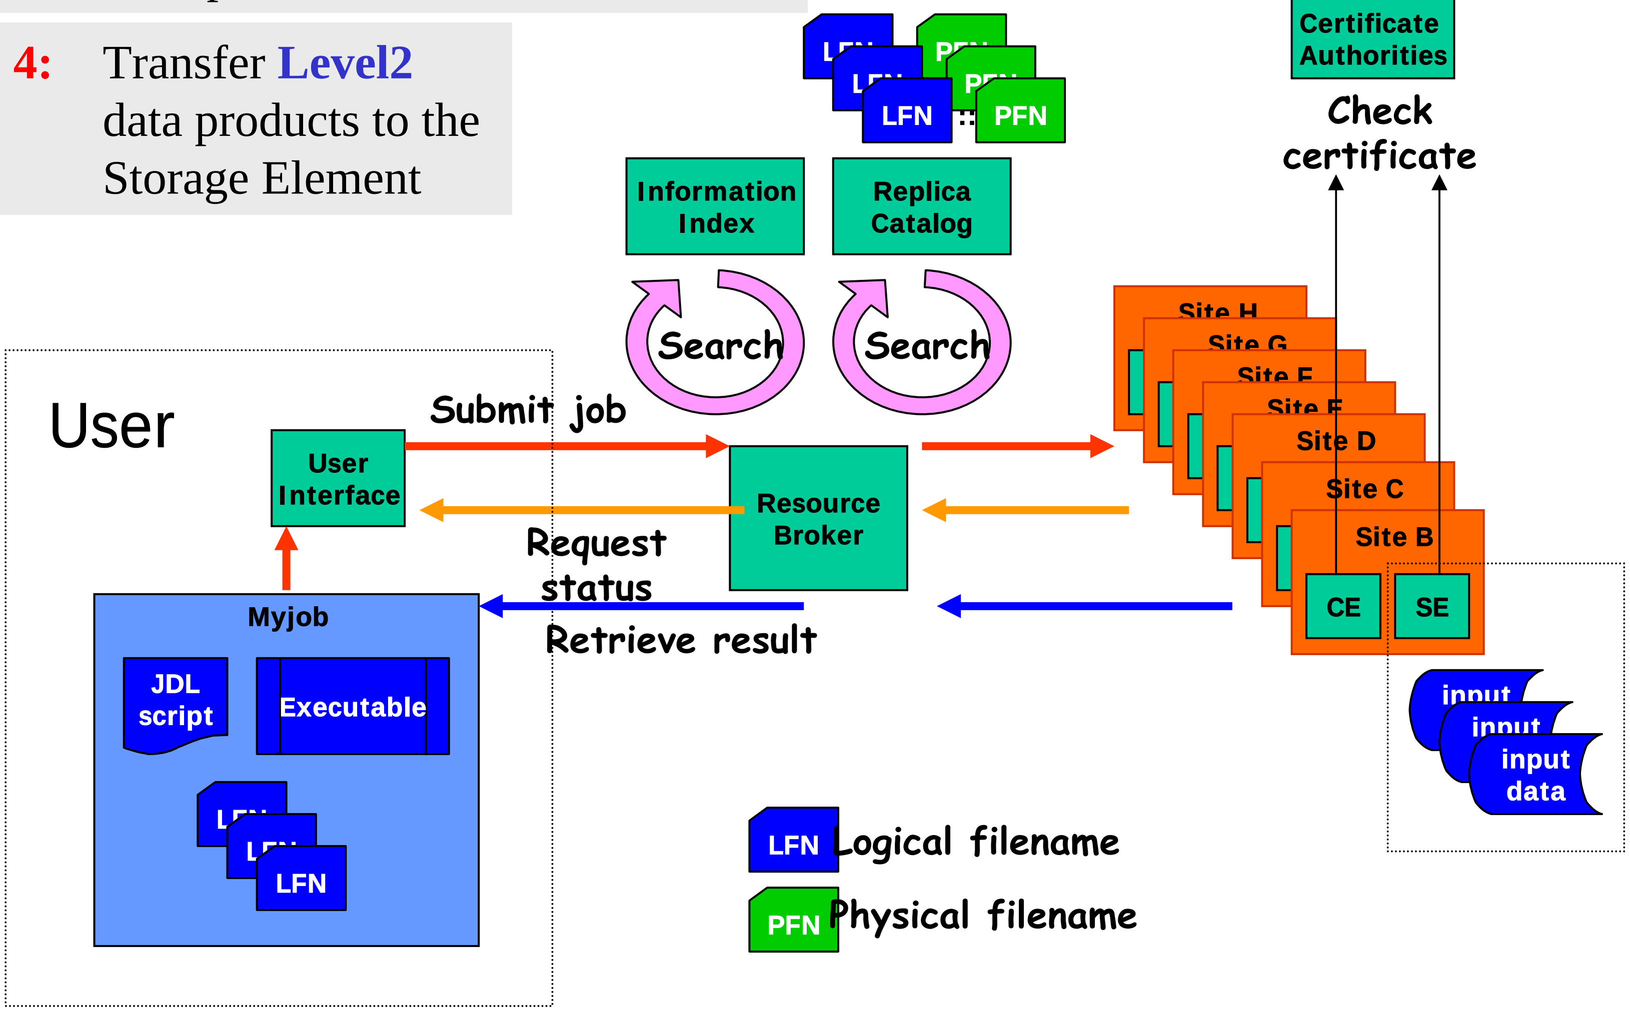

GOME Processing Steps (5-6)

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

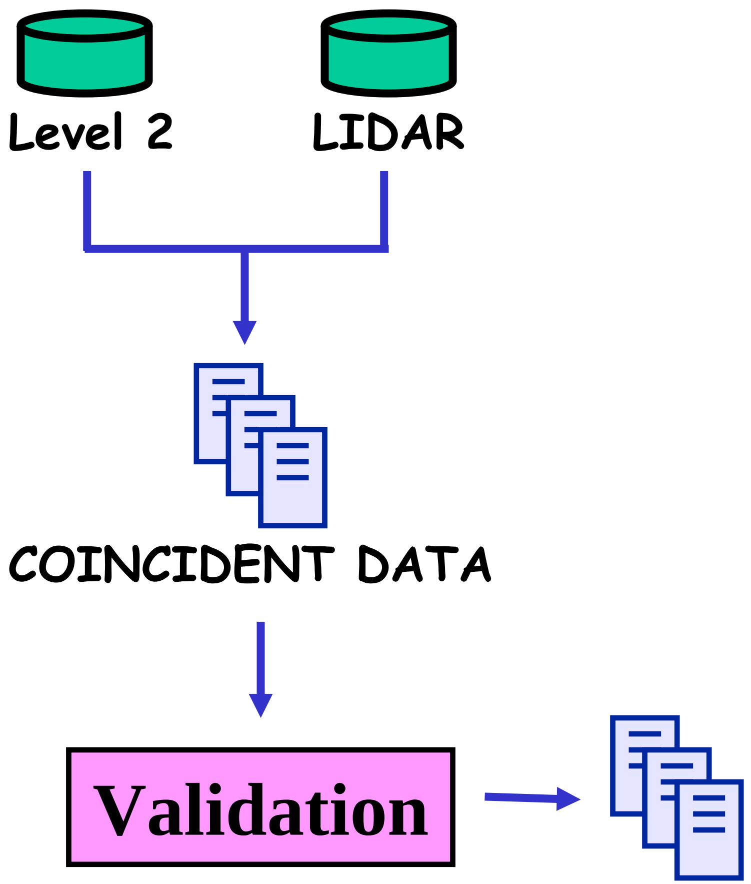

Step 6: Visualize Results

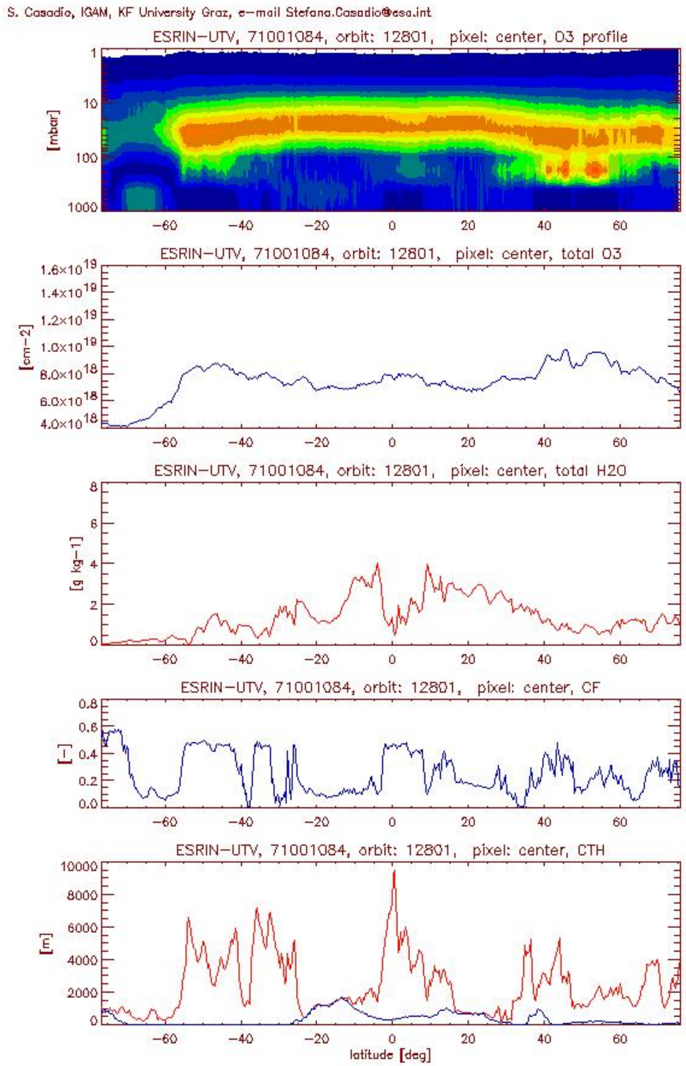

Common Applications Work

- Several discussions between application WPMs and technical coordination to consider the common needs of all applications

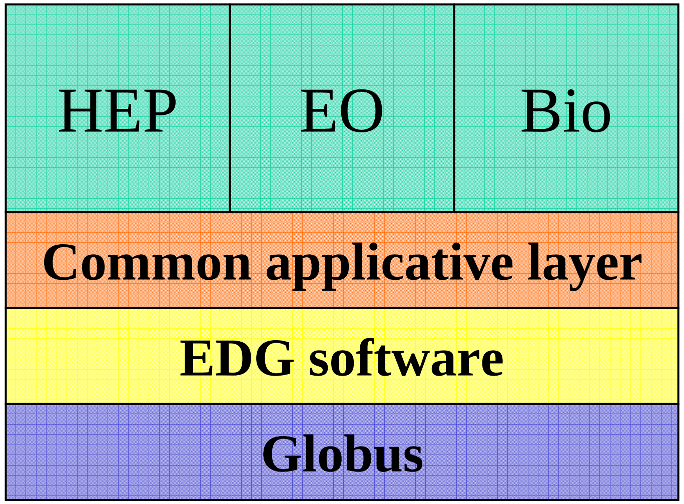

Summary and a forward look for applications work within EDG

- **Currently evaluating the basic functionality of the tools and their integration into data processing schemes. Will move onto areas of interactive analysis, and more detailed interfacing via APIs**
  - **Hopefully experiments will do common work in interfacing applications to GRID under the umbrella of LCG**
  - **HEPCAL (Common Use Cases for a HEP Common Application Layer) work will be used as a basis for the integration of Grid tools into the LHC prototype**

    http://lcg.web.cern.ch/LCG/SC2/RTAG4

- **There are many grid projects in the world and we must work together with them**
  - **e.g. in HEP we have DataTag,Crossgrid,Nordugrid + US Projects(GryPhyn,PPDG,iVDGL)**

- **Perhaps we can define shared project between HEP,Bio-med and ESA for applications layer interfacing to basic Grid functions.**

Acknowlegements and references

- **Thanks to the following who provided material and advice**
  - **J Linford(WP9),V Breton(WP10),J Montagnat(WP10),F Carminati(Alice),JJ Blaising(Atlas),C Grandi(CMS),M Frank(LHCb),L Robertson(LCG),D Duellmann(LCG/POOL) ,T Doyle(UK GridPP),M Reale(WP8)**
- **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)**