

Life sciences production on the EGEE grid infrastructure

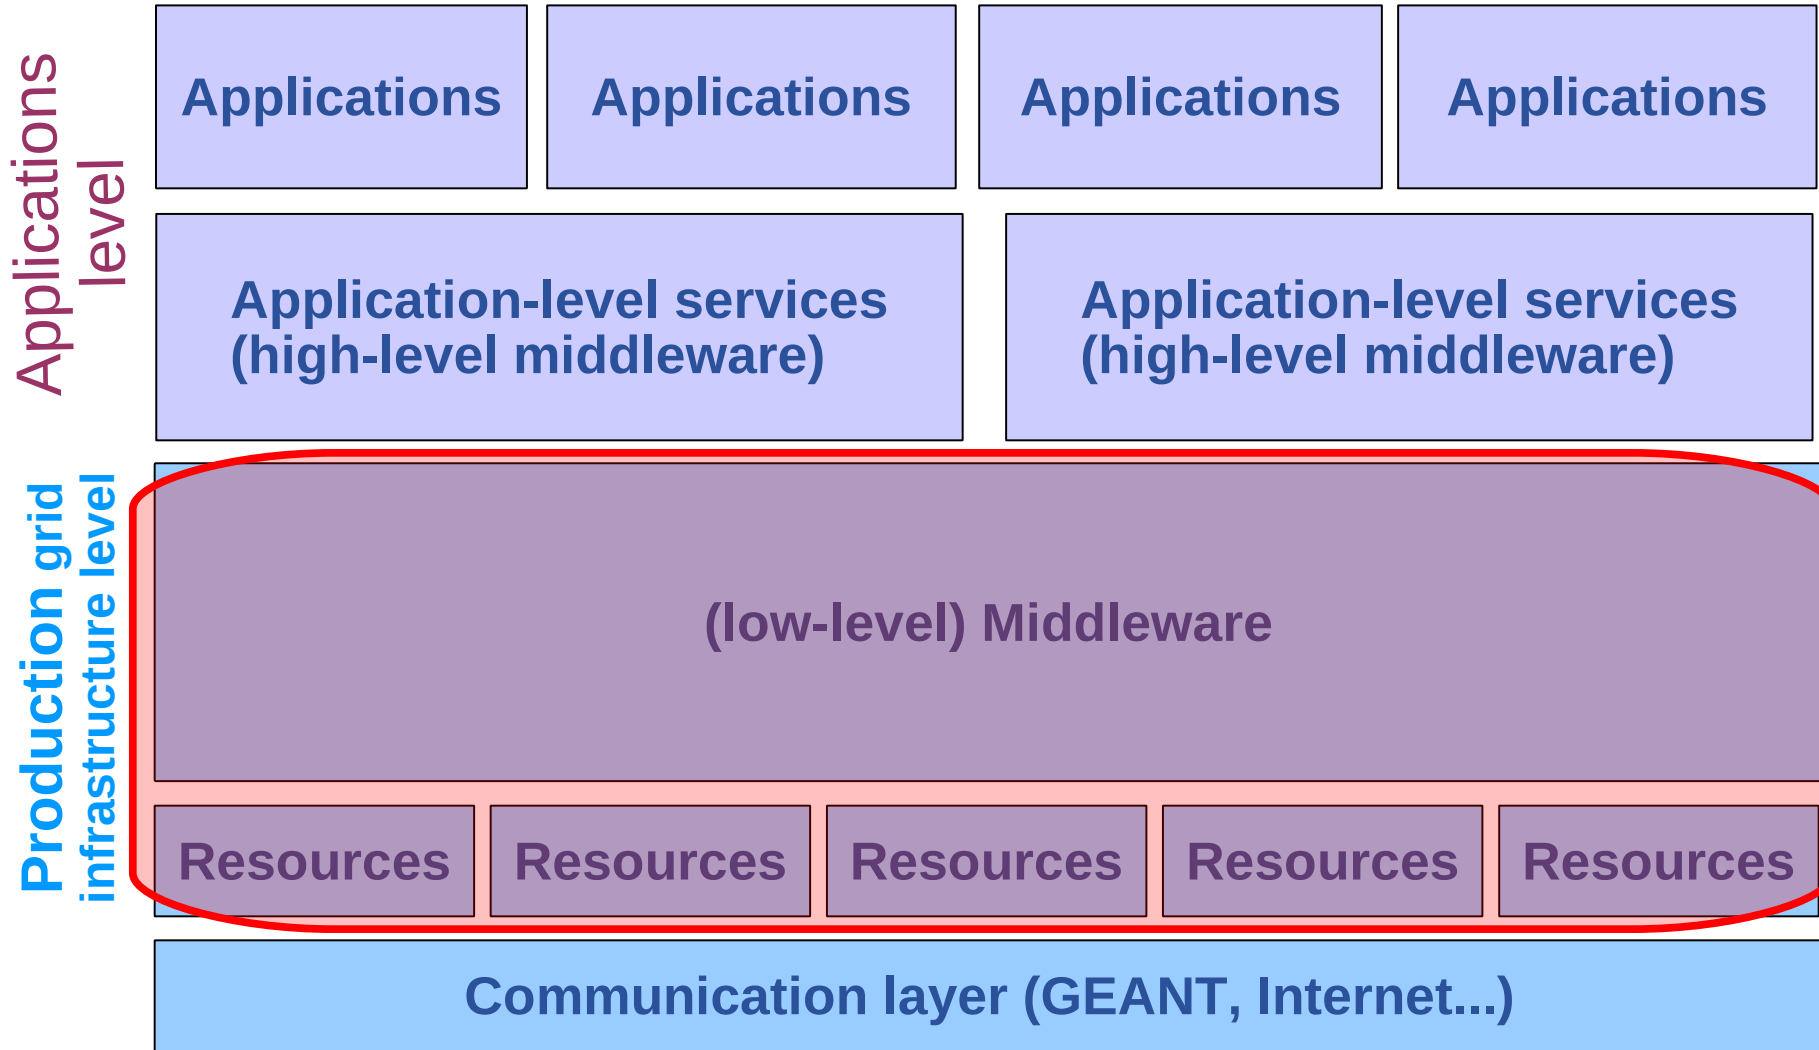
CSIR

Pretoria, May 12, 2008

Johan Montagnat

CNRS, I3S laboratory

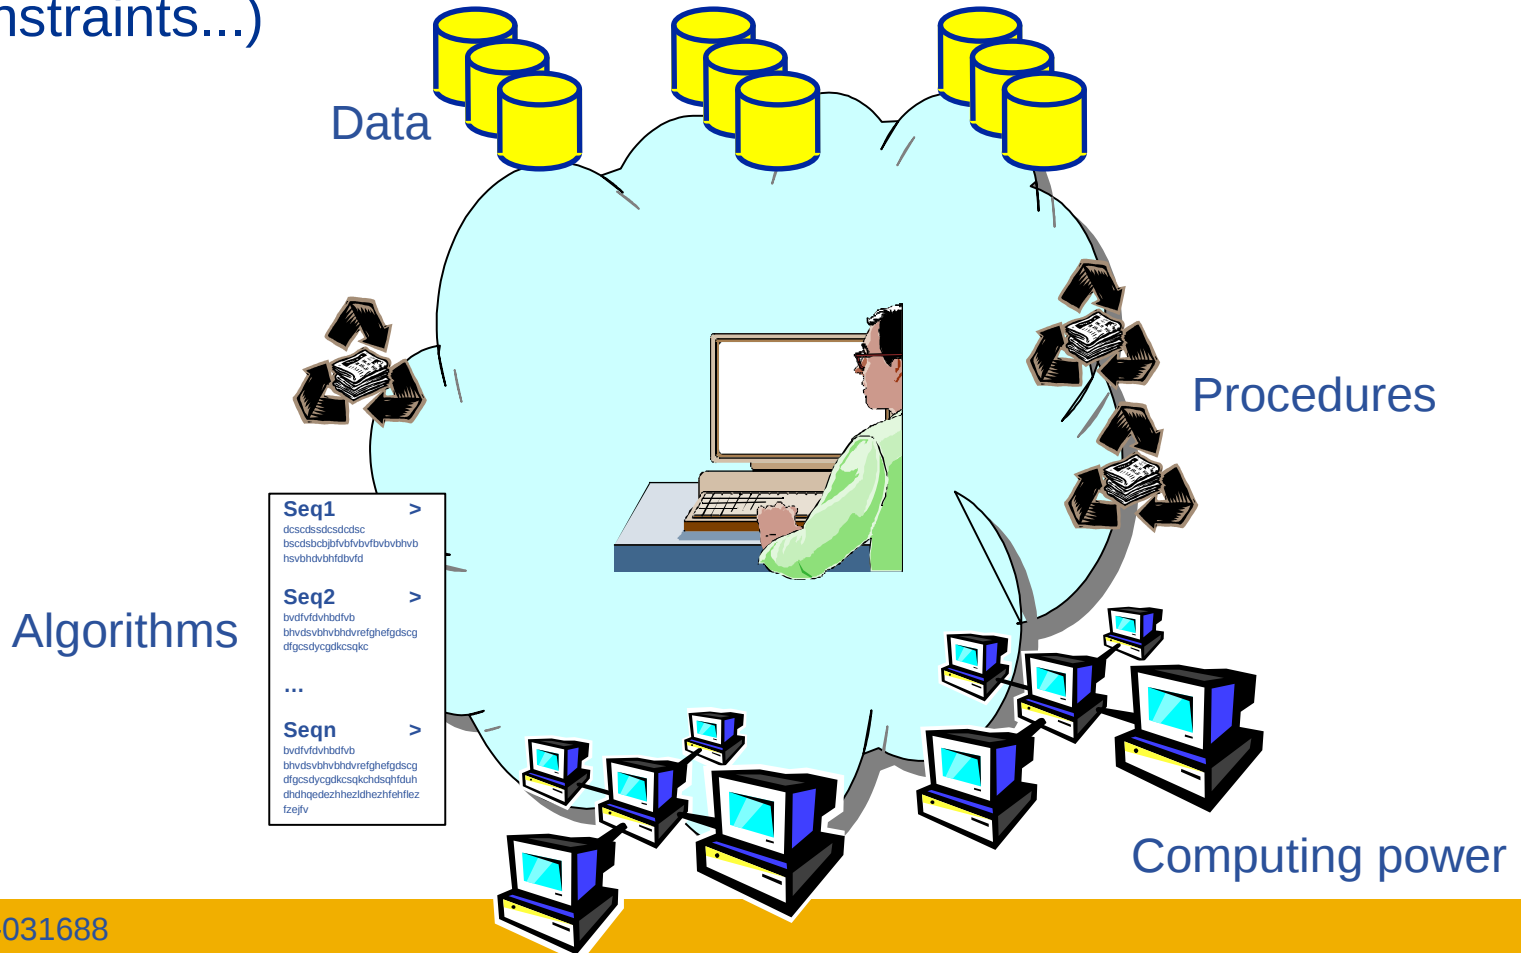




- **Complex middleware stack**
 - On top of Globus Toolkit 2, VDT, CONDOR, batch schedulers...
- **Fundational services**
 - User authentication and authorization
 - Information Index: registration of Computing Elements (CEs), Storage Elements (SEs), etc.
 - Workload Management System: workload distribution over sites
 - Data Management System: virtual file hierarchy, standard (SRM) interface to storage resources
- **Batch oriented**
 - EGEE is a federation of computing centers
 - and File Servers oriented




- **Sharing computing resources and algorithms**
 - Research (populations studies, models design, validation, statistics)
 - Complex analysis (compute intensive image processing, time constraints...)



- **Fostering scientific communities**
 - EGEE is an international, multi-disciplinary research infrastructure
 - Scientific communities are expanding beyond administrative boundaries
- **Virtual Organizations**
 - Authentication and Authorization management
 - Share resources (computers, data, algorithms...) inside a VO
 - Application areas identification and support unit
- **The Biomed VO now divided in 3 sectors**
 - Medical imaging
 - Bioinformatics
 - Drug discovery

Flagship European grid infrastructure project
Now in 2nd phase with 91 partners in 32 countries

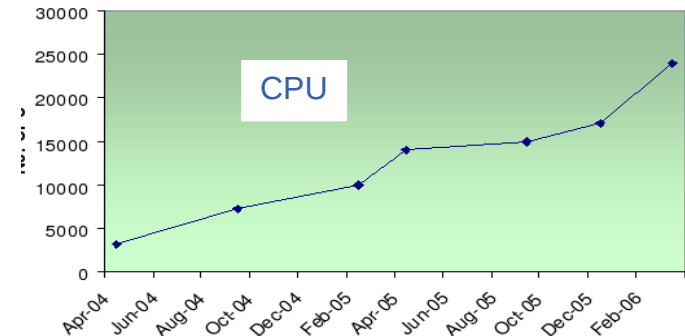


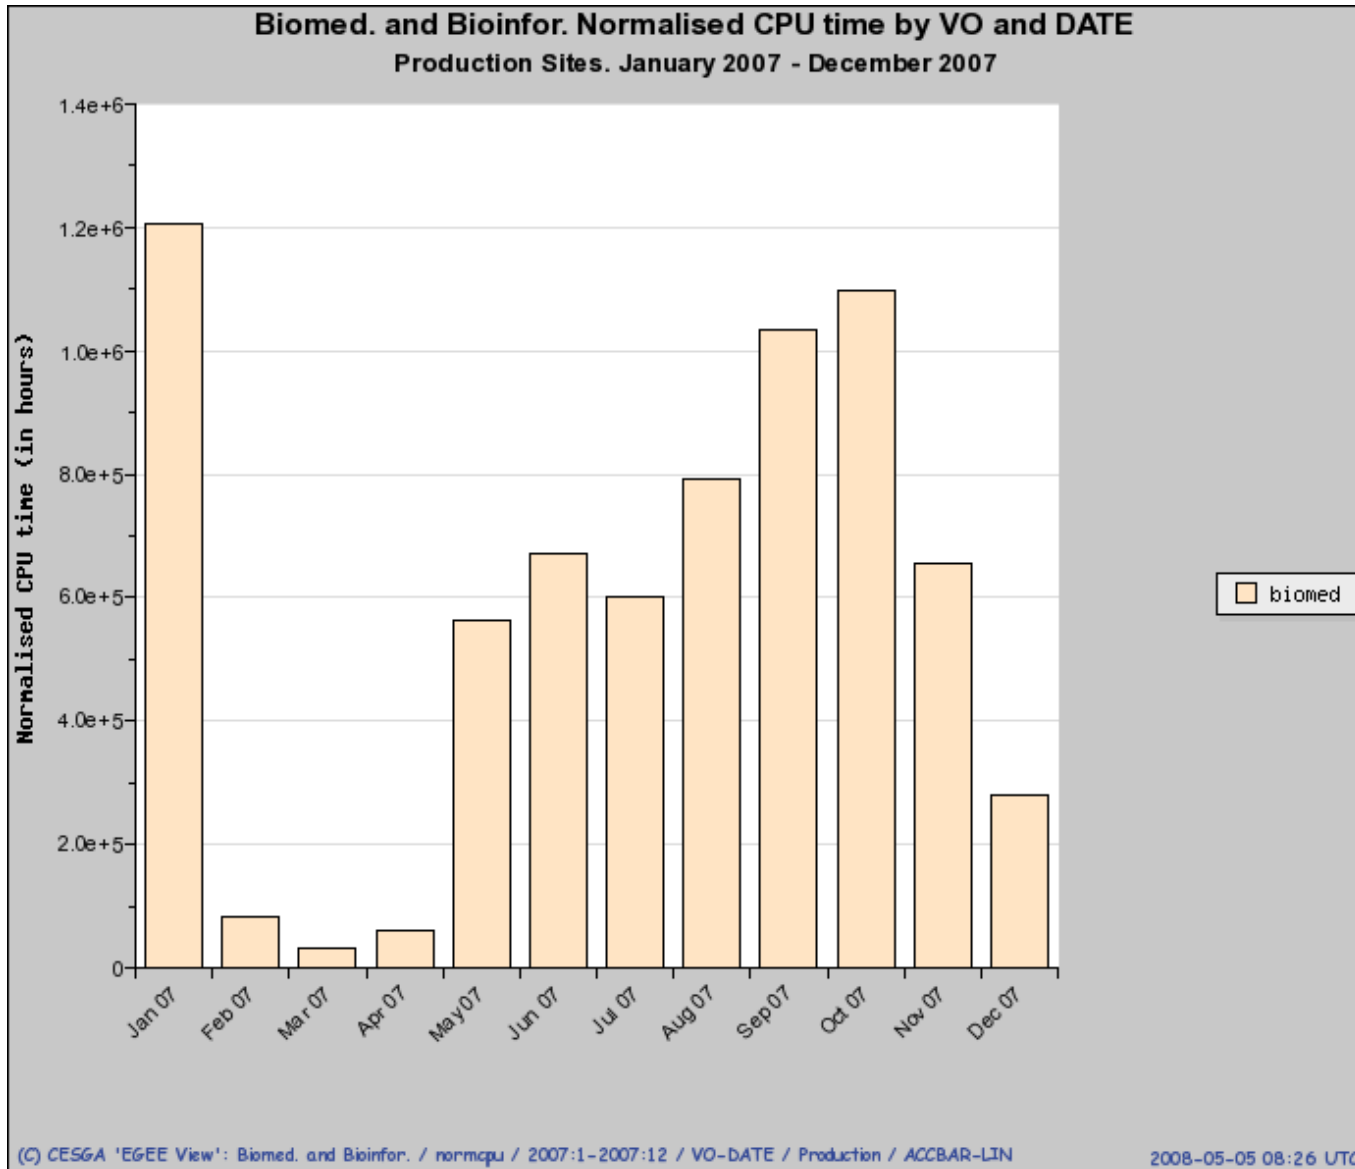
Size of the infrastructure today:

- > 240 sites in 40 countries
- > 50 000 CPU
- ~ 5 PB disk + tape MSS
- > 100 000 concurrent jobs/day
- > 8000 registered users

Out of which, Biomed VO:


- > 100 sites in 30 countries
 (170 CEs, 130 SEs)
- ~ 17 000 CPU
- > 150 registered users





Year 2007 statistics collected per month


Biomed. and Bioinfor. Cumulative Normalised CPU time by VO and DATE
Production Sites. January 2007 - December 2007



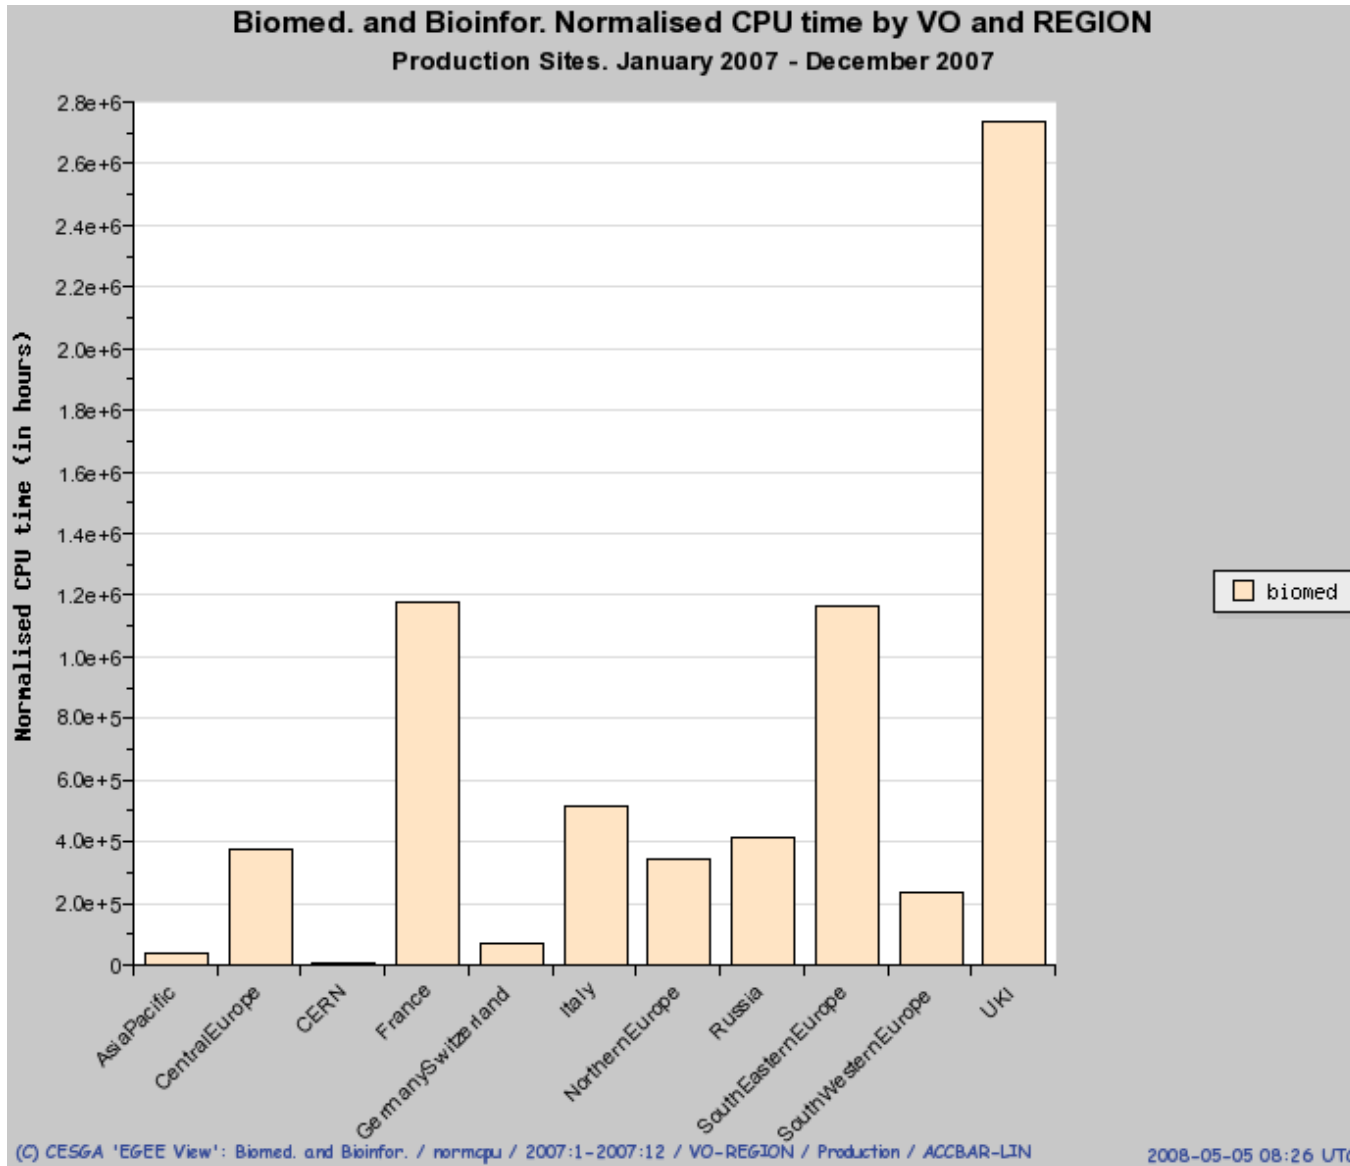
~800 CPU years
 (1000 SpecInt 2000
 Normalised CPU
 time)

biomed

Biomed. and Bioinfor. Cumulative Total number of jobs by VO and DATE
Production Sites. January 2007 - December 2007

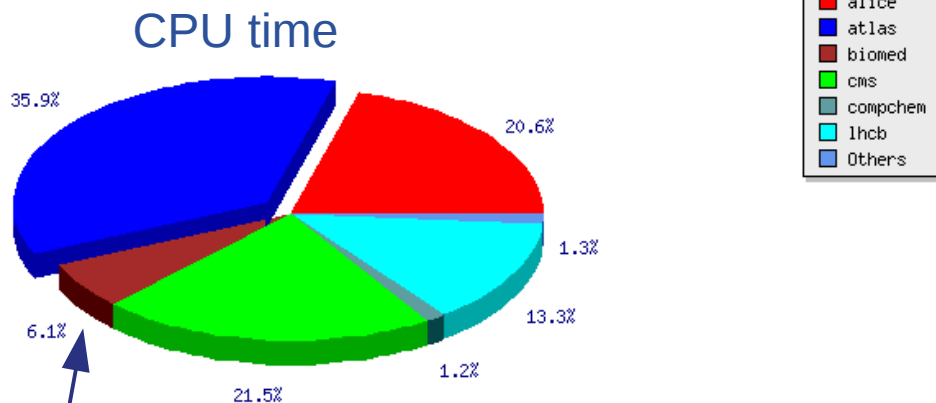


Computing hours per EGEE federation in 2007



• Biomed VO share

PRODUCTION Normalised CPU time per VO (Excluded dteam and ops VOs)
EGEE VOs. January 2007 - December 2007



PRODUCTION Total number of jobs per VO (Excluded dteam and ops VOs)
EGEE VOs. January 2007 - December 2007

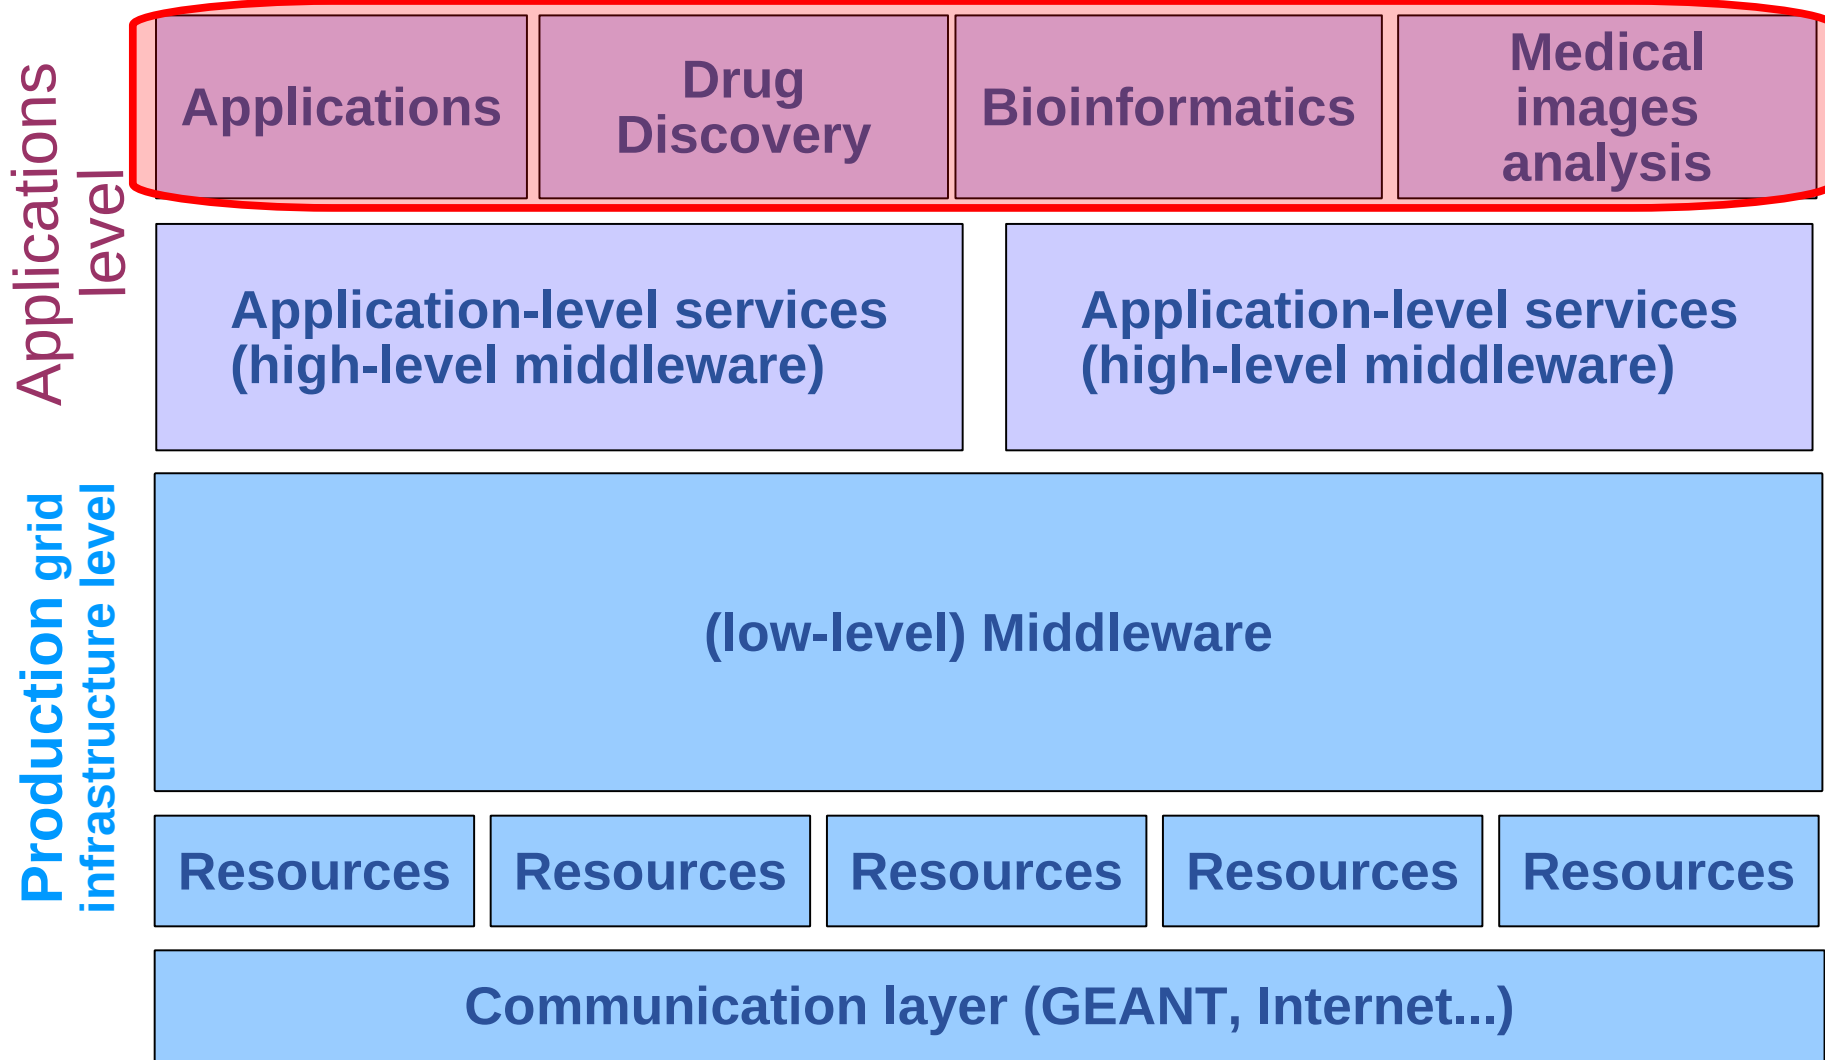


(C) CESSGA 'EGEE View': PRODUCTION / normcpu / 2007:1-2007:12 / VO-REGION / egge (x) /

Biomed VO

(C) CESSGA 'EGEE View': PRODUCTION / njobs / 2007:1-2007:12 / VO-REGION / egge (x) / ACCBAR-LIN / x

2008-05-05 08:26 UTC



- **WISDOM:** <http://wisdom.healthgrid.org/>
- **Goal: find new drugs for neglected and emerging diseases**
 - Neglected diseases lack R&D
 - Emerging diseases require very rapid response time
- **Need for an optimized environment**
 - To achieve production in a limited time
 - To optimize performances
- **Method: grid-enabled virtual docking**
 - Cheaper than in vitro tests
 - Faster than in vitro tests



Grid-enabled virtual docking

Millions of potential drugs to test against interesting proteins!




High Throughput Screening
 ~10\$/compound, several hours

Too costly for neglected disease!

Compounds:

ZINC: 4.3M


Chembridge: 500 000



Molecular docking (FlexX, Autodock)
 ~1 to 15 minutes

Targets:


PDB: 3D structures



Data challenge on EGEE
 ~ 2 to 30 days on ~5000 computers

Cheap and fast!

Selection of the best hits

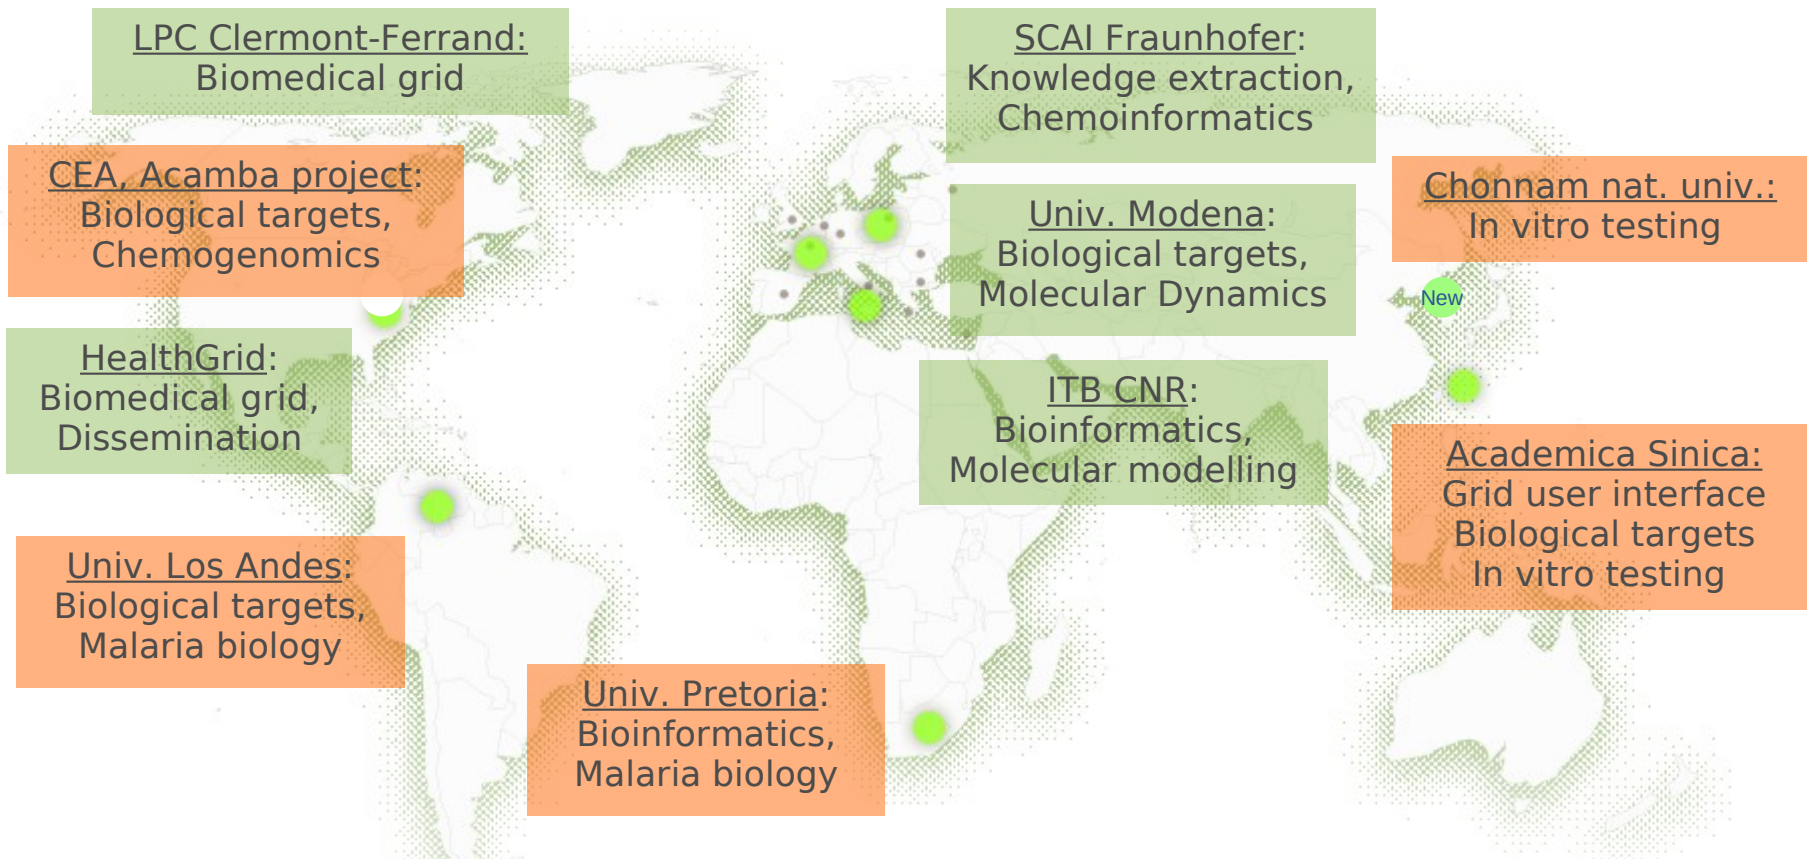


Hits screening using assays performed on living cells



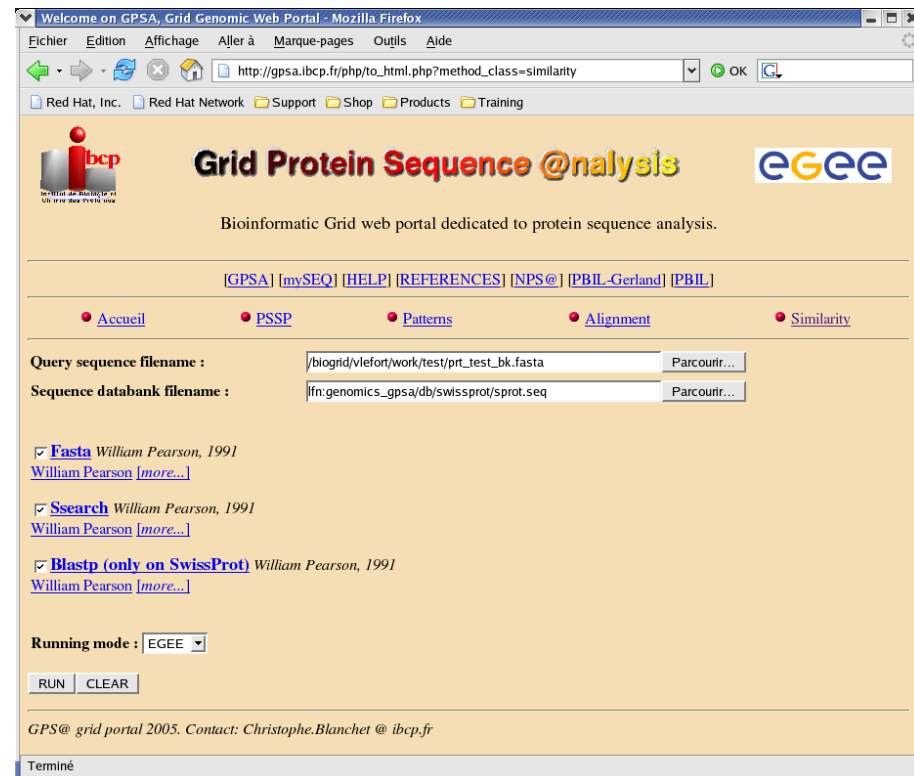
Leads
 Clinical testing
 Drug


- The grid provides the centuries of CPU cycles required on demand
- The grid provides the reliable and secure data management services to store and replicate the biochemical inputs and outputs
- The grid offers a collaborative environment for the sharing of data in the research community on avian flu and malaria




- **First Data Challenge: July 1st - August 15th 2005**
 - Target: **malaria**
 - 80 CPU years
 - 1 TB of data produced
 - 1700 CPUs used in parallel
 - **1st large scale docking deployment world-wide on a e-infrastructure**
- **Second Data Challenge: April 15th - June 30th 2006**
 - Target: **avian flu**
 - 100 CPU years
 - 800 GB of data produced
 - 1700 CPUs used in parallel
 - **Collaboration initiated on March 1st: deployment preparation achieved in 45 days**
- **Third Data Challenge: October 1st - 15th December 2006**
 - Target: **malaria**
 - 400 CPU years
 - 1,6 TB of data produced
 - Up to 5000 CPUs used in parallel
 - **Very high docking throughput: > 100.000 compounds per hour**

- GPS@: bioinformatics portal
 - <http://gpsa.ibcp.fr/> web portal
 - Relaces existing (but overloaded) NPSA portal
 - Tens of bioinformatics legacy code
 - Thousands of potential users
 - Large input databases





- > 7.000.000 analyses
- > 5000 analyses/day



Courtesy of C. Combet

- **Protein analysis methods**
 - Sequence similarity: BLAST, SSEARCH, FASTA
 - Multiple Alignment: ClustalW, MultAlin
 - Protein site/signatures: PattInProt
 - Secondary Struct. Pred: Predator, Gor4, Simpa96

- **Biological data (clear and encrypted)**
 - Sequence Bank: SWISSPROT, TrEMBL
 - Pattern/Profile bank: PROSITE

Enabling Grids for E-science


Welcome on GPSA, Grid Genomic Web Portal - Mozilla Firefox

Fichier Edition Affichage Aller à Marque-pages Outils Aide


http://gpsa.ibcp.fr/php/to_html.php?method_class=similarity

Red Hat, Inc. Red Hat Network Support Shop Products Training

Google Welcome on GPSA, Grid Genomi...



Grid Protein Sequence @analysis



Bioinformatic Grid web portal dedicated to protein sequence analysis.

[GPSA] [mySEQ] [HELP] [REFERENCES] [NPS@] [PBIL-Gerland] [PBIL]

● Accueil ● PSSP ● Patterns ● Alignment ● Similarity

Query sequence filename :

Sequence databank filename :

[Fasta](#) William Pearson, 1991
[William Pearson](#) [more...]

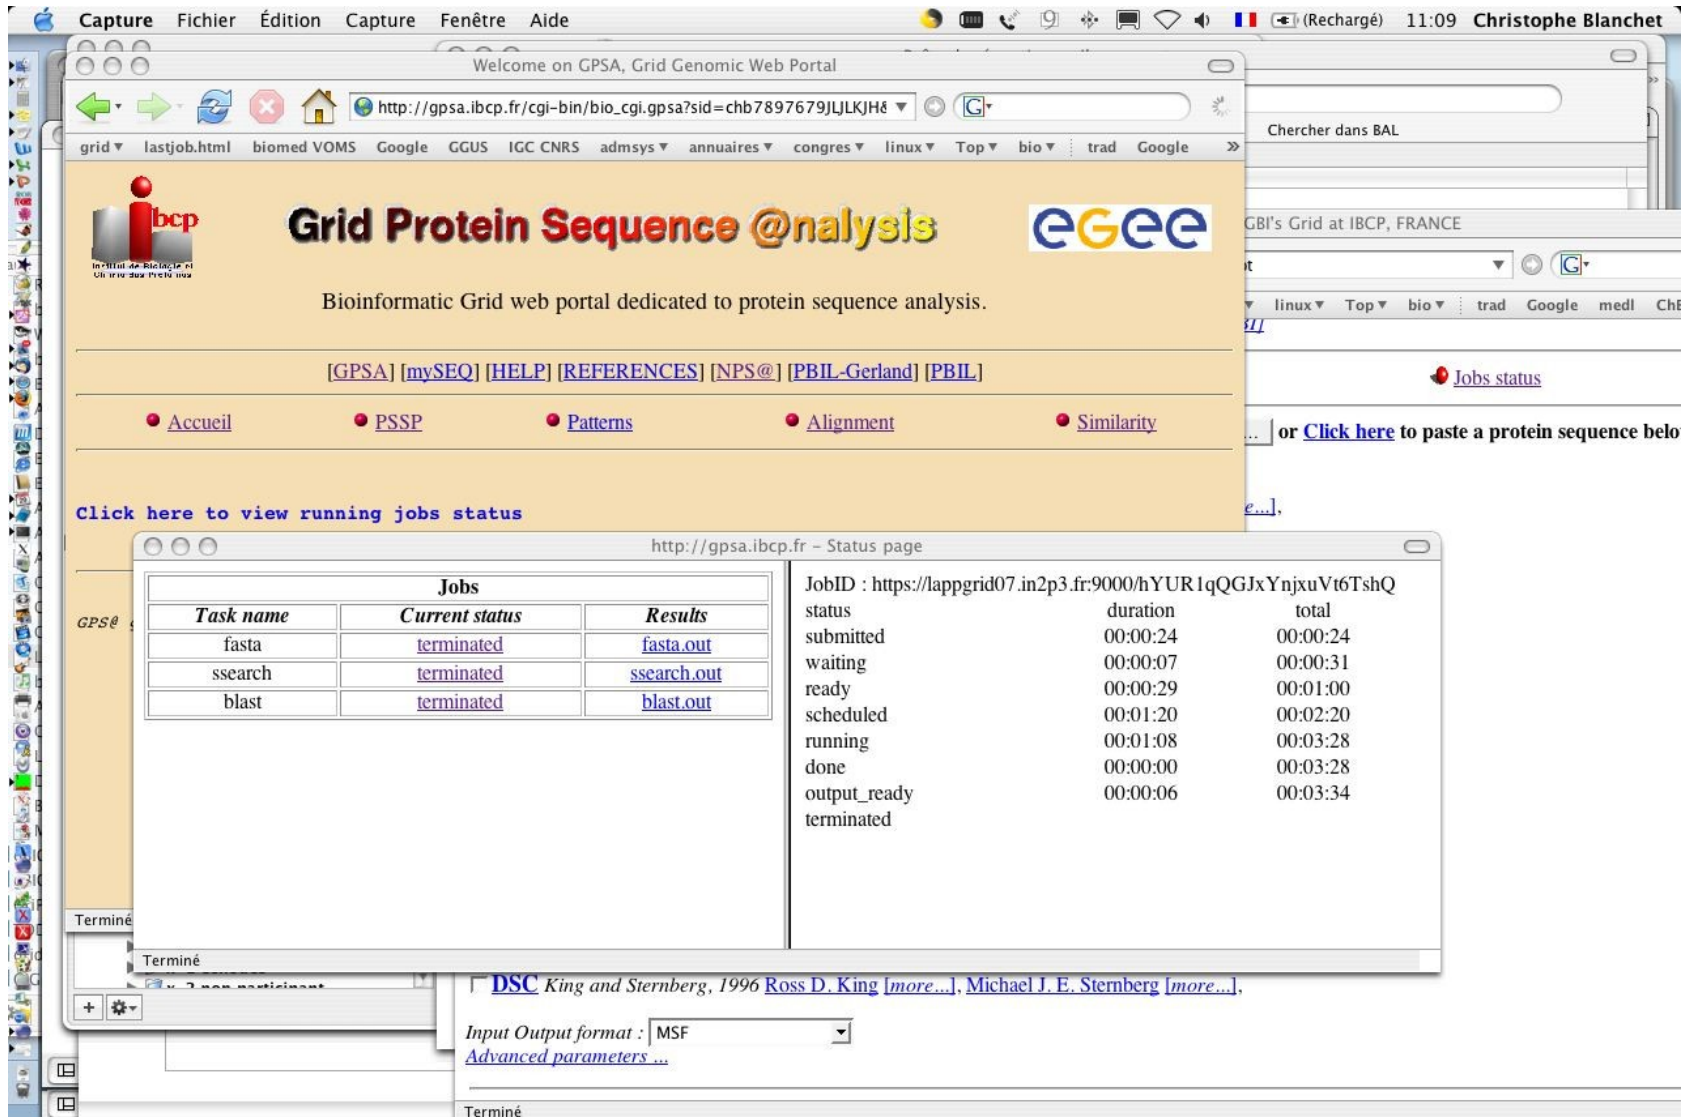
[Ssearch](#) William Pearson, 1991
[William Pearson](#) [more...]

[Blastp](#) William Pearson, 1991
[William Pearson](#) [more...]

Running mode :

GPS@ grid portal 2005. Contact: Christophe.Blanchet @ ibcp.fr


Terminé



Welcome on GPSA, Grid Genomic Web Portal

http://gpsa.ibcp.fr/cgi-bin/bio CGI.gpsa?sid=chb7897679JLJKJHé

grid ▾ lastjob.html biomed VOMS Google GGUS IGC CNRS admsys ▾ annuaires ▾ congres ▾ linux ▾ Top ▾ bio ▾ trad Google >>

Grid Protein Sequence @analysis 

Bioinformatic Grid web portal dedicated to protein sequence analysis.

[GPSA] [mySEQ] [HELP] [REFERENCES] [NPS@] [PBIL-Gerland] [PBIL]

● [Accueil](#) ● [PSSP](#) ● [Patterns](#) ● [Alignment](#) ● [Similarity](#)

[Click here to view running jobs status](#)

Jobs status

... or [Click here](#) to paste a protein sequence below.

http://gpsa.ibcp.fr - Status page

Task name	Current status	Results
fasta	terminated	fasta.out
ssearch	terminated	ssearch.out
blast	terminated	blast.out

JobID : https://lappgrid07.in2p3.fr:9000/hYUR1qQGJxYnjxvT6TshQ

status	duration	total
submitted	00:00:24	00:00:24
waiting	00:00:07	00:00:31
ready	00:00:29	00:01:00
scheduled	00:01:20	00:02:20
running	00:01:08	00:03:28
done	00:00:00	00:03:28
output_ready	00:00:06	00:03:34
terminated		

Terminé




Terminé

[DSC King and Sternberg, 1996](#) [Ross D. King \[more...\]](#), [Michael J. E. Sternberg \[more...\]](#)


Input Output format :

[Advanced parameters ...](#)

Terminé




Unregistered



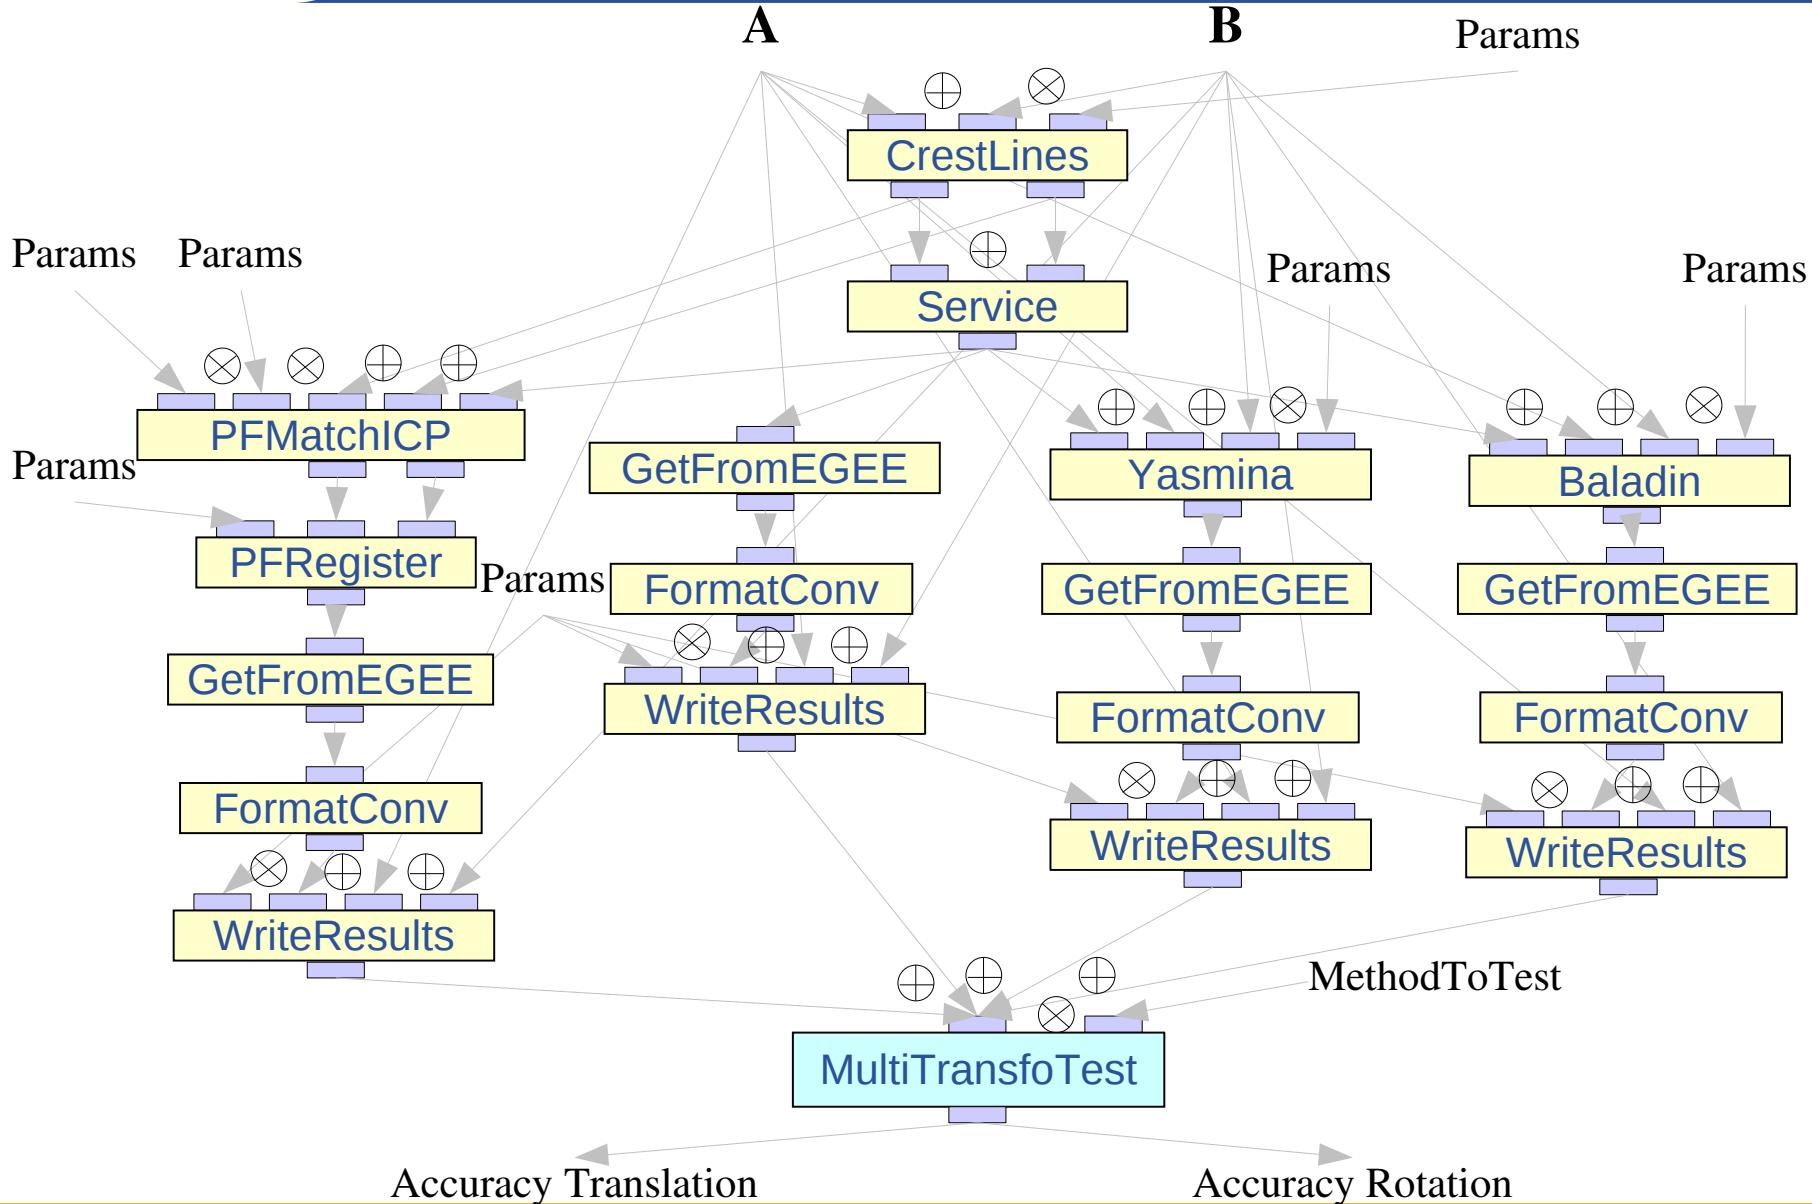
Registered

- N images, m algorithms
- $N \cdot (N-1) \cdot m$ transformations measured
- N-1 transformations to estimate


} Redundancy



- **Exploit redundancy to compute**
 - Mean transformations \overline{T}_{ij} (Bronze standard)
 - Variances on the transformations (Accuracy)



- 29 patients
- 2 time points minimum
- Gadolinium injected T1 MRIs
- Example for one patient (3 time points):



t1

t2

t3

- Mean error on the transformations:

$$\sigma_r = 0.130 \text{ deg} ; \sigma_t = 0.345 \text{ mm}$$

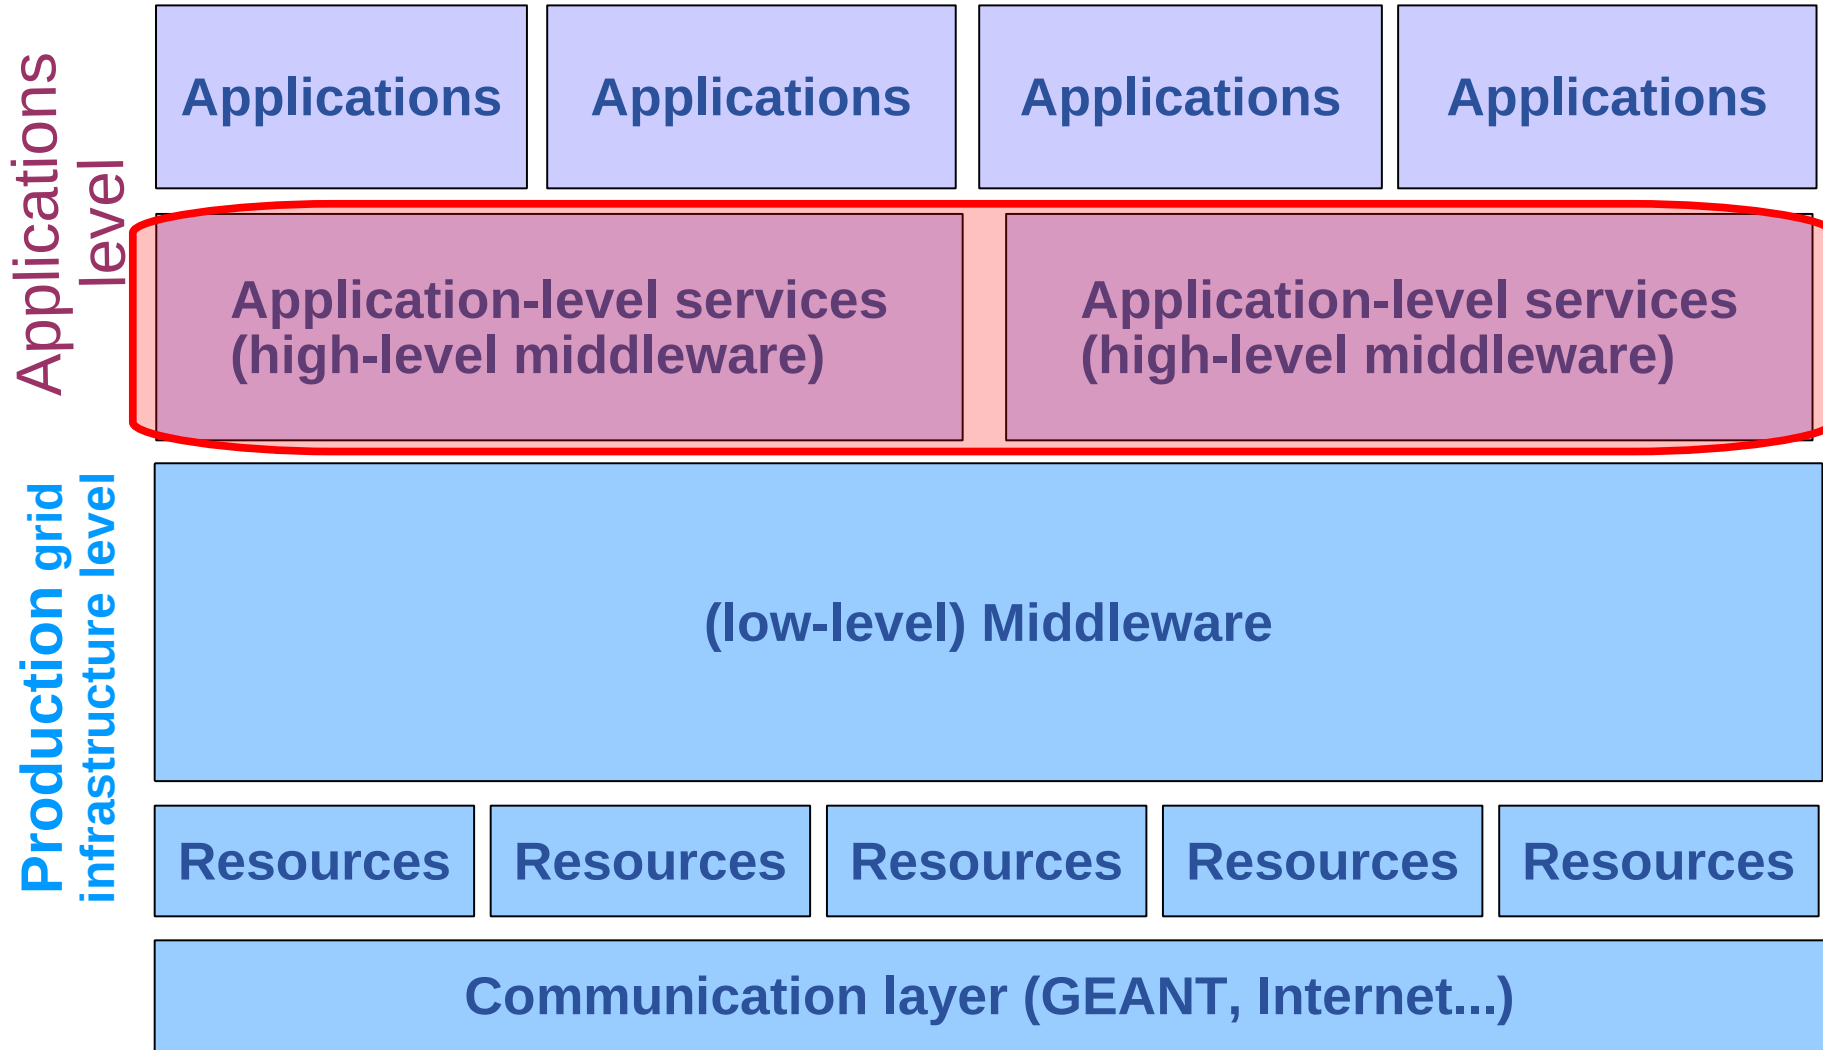
- Error on the bronze standard:

$$\sigma_r = 0.05 \text{ deg} ; \sigma_t = 0.148 \text{ mm}$$

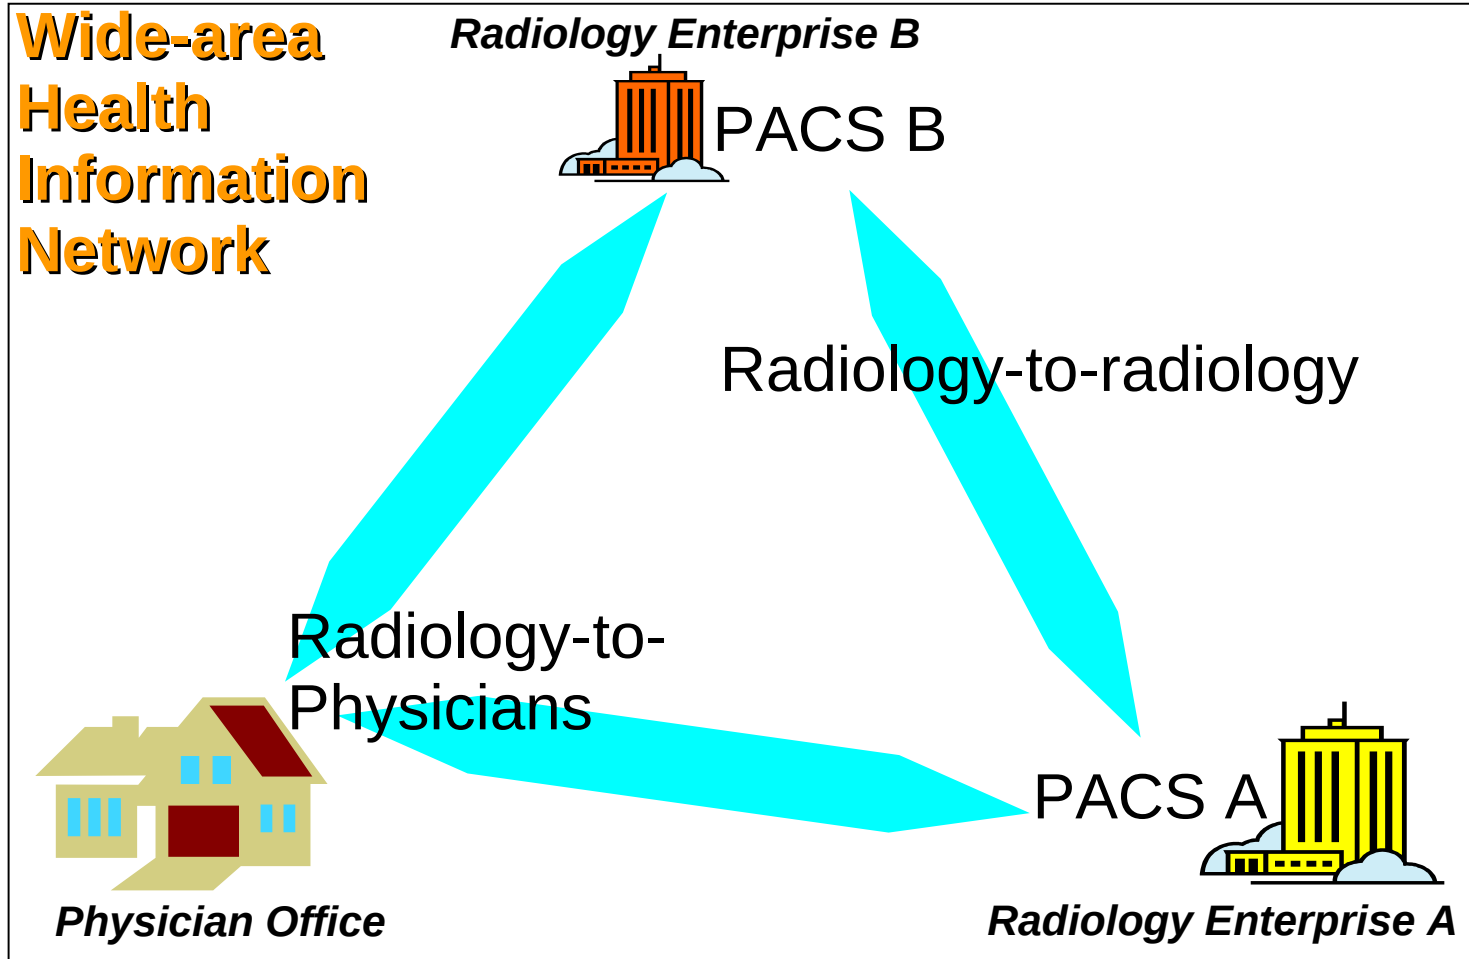
- Accuracy of the algorithms:

Algorithm	σ_r (deg)	σ_t (mm)
CrestMatch	0.150	0.424
PFRegister	0.180	0.416
Baladin	0.139	0.395
Yasmina	0.137	0.445

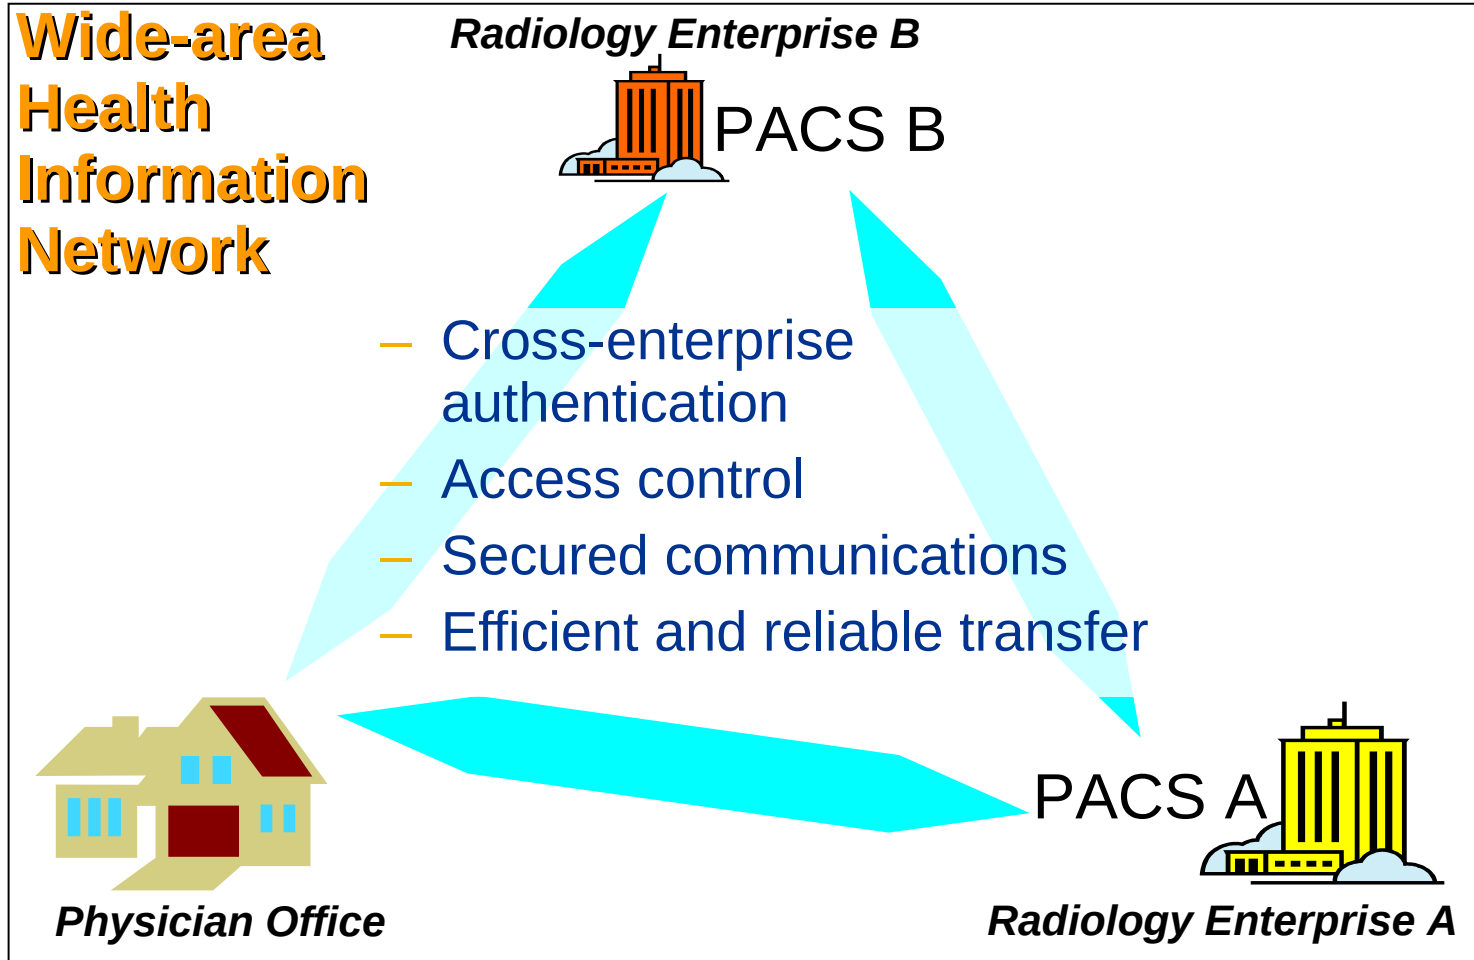
- **“Embarassingly” parallel problem**
- **Fast turn over of short jobs**
- **Data confidentiality**
- **Interactivity**
- **Fine grain parallelism (MPI)**
- **Workflow-based**
- **Portals / user interface**



- Cross-enterprise exchange of radiology reports and images

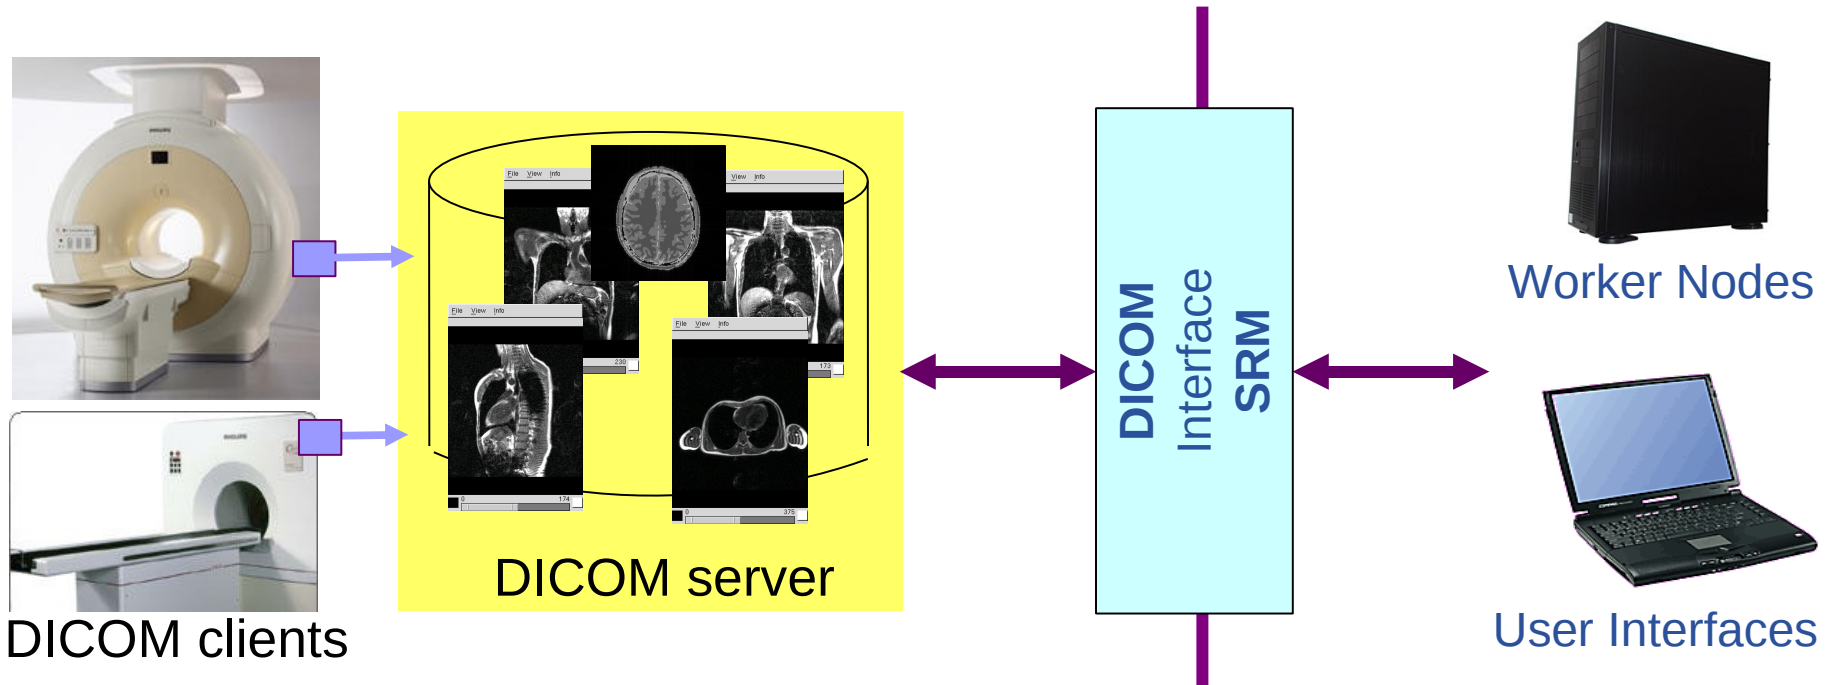


- Grid technologies




Objectives

- Expose a **standard grid interface (SRM)** for **medical image servers (DICOM)**
- Use native DICOM storage format
- Fulfill medical applications security requirements
- Do not interfere with clinical practice






1. Image is acquired



2. Image is stored in DICOM server



DICOM server


3. lcg client

**GFAL
API**

3a. Image is registered (a GUID is associated)

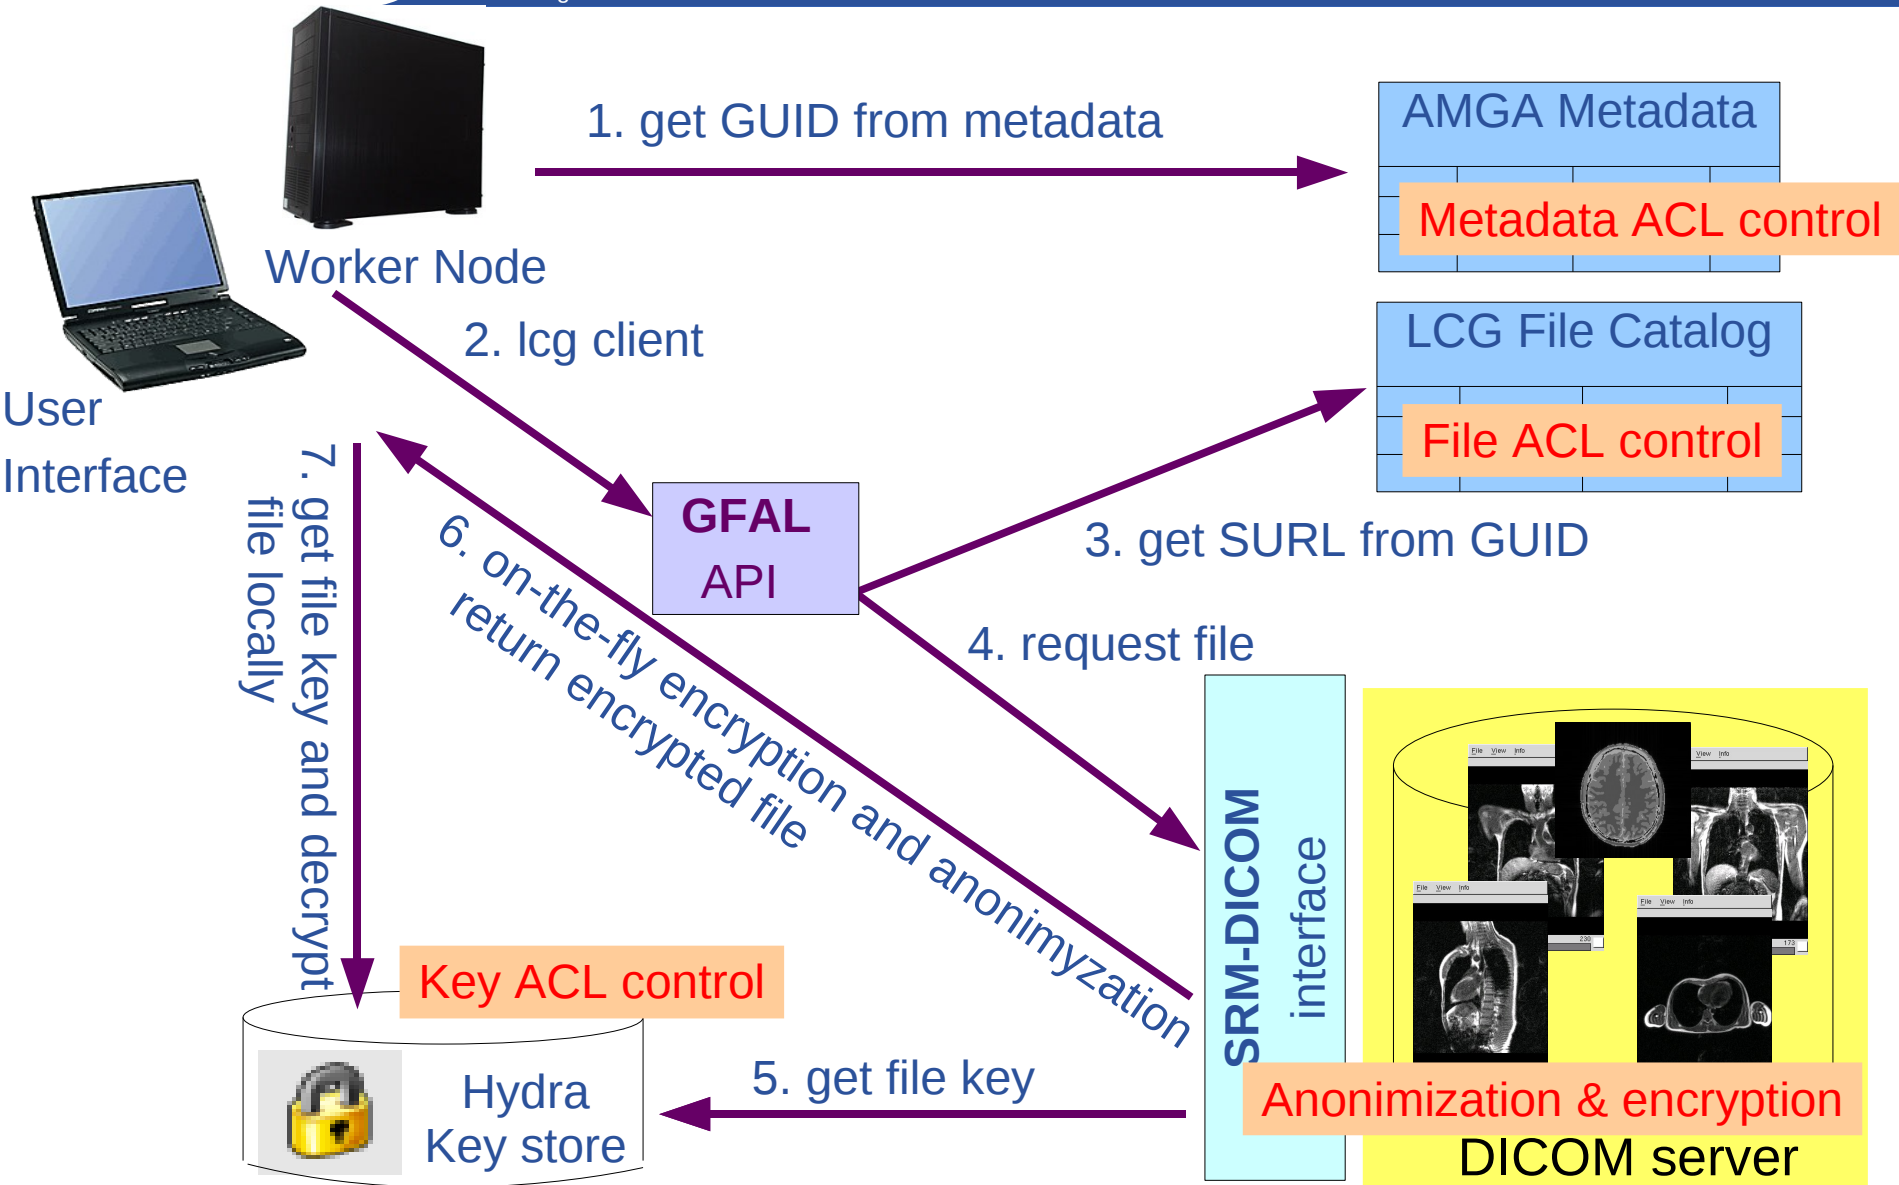
LCG File Catalog

3b. Image key
is produced and
registered



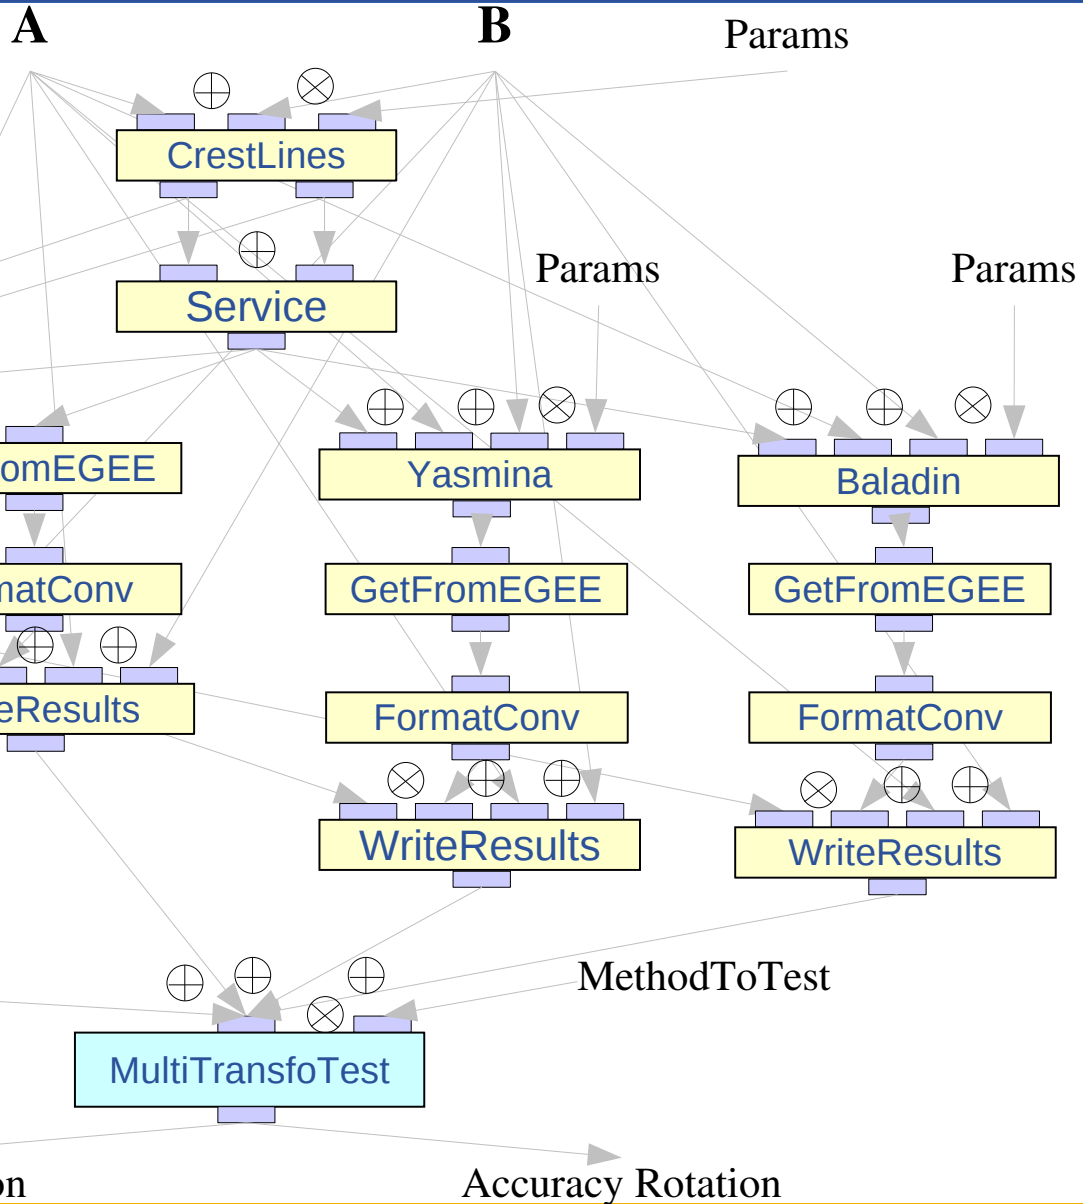

4. image metadata
are registered

AMGA Metadata




~100 image pairs

~800 EGEE jobs




- **MOTEUR, I3S laboratory, CNRS**
 - <http://egee1.unice.fr/MOTEUR>
- **High level interface**
 - Hides grid complexity to the user
- **Service-based approach**
 - Legacy code service wrapper
- **Scufl language (myGrid / Taverna)**
 - Pure data flow approach
- **Grid submission interfaces**
 - EGEE (LCG2, gLite)
 - Grid5000 (OAR, DIET)
- **Transparent parallelism exploitation**
 - Code and data parallelism




- A workflow naturally provides application parallelization
- MOTEUR transparently exploits 3 kinds of parallelism


Workflow parallelism



Data parallelism



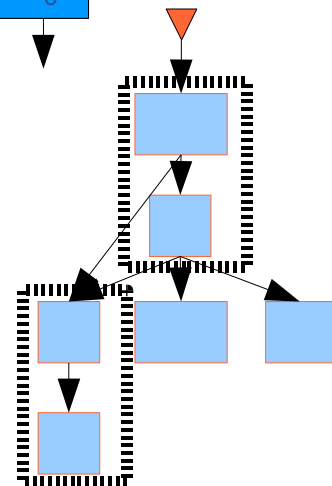
Service parallelism



- Jobs grouping strategy in sequential branches in order to reduce grid latency



<http://egee1.unice.fr/MOTEUR>



- **Torque + MAUI batch scheduler**
 - Used in most EGEE sites
 - specific configuration (virtual processors allocation)
 - Does not interfere with normal batch scheduling (shared processor time)
 - Enables efficient processing of short tasks on the production infrastructure
 - Alternative to jobs prioritization
- **Special submission queues**
 - Three SDJ queues deployed on biomed-compliant sites
 - Time-limited queues
- **Submit-or-reject paradigm**
 - Jobs are immediately executed or rejected if a too high number of short jobs are already executing.

- **Grid used for sustained production in life sciences**

- Regular usage in drug discovery, bioinformatics and medical imaging community
- Closer to non computing-specialized end-users
 - RSNA'08 demonstration (EGEE-MEDICUS)
 - Bioinformatics web portal
 - ...



- **Application-domain high level services being developed**

- Data protection
- Data semantics
- User interfaces

- **Demonstrated direct interest for uses community**

- Computing power... but not only computing power
- Data, algorithms and procedures federation