Port d'Informació Científica

Generalitat de Catalunya
Departament d'Universitats, Recerca i Societat de la Informació

IFAE

Universitat Autònoma de Barcelona

The LHC Computing Grid Project in Spain (LCG-ES)

Presentation to RECFA

Prof. Manuel Delfino
Coordinating Principal Investigator, LCG-ES
Physics Department, Universitat Autònoma de Barcelona
Director, Port d'Informació Científica (PIC)*

* Collaboration of the Generalitat de Catalunya, IFAE and Universitat Autònoma de Barcelona

LHC is multi-Petabyte, Tera-object data-intensive scientific computing

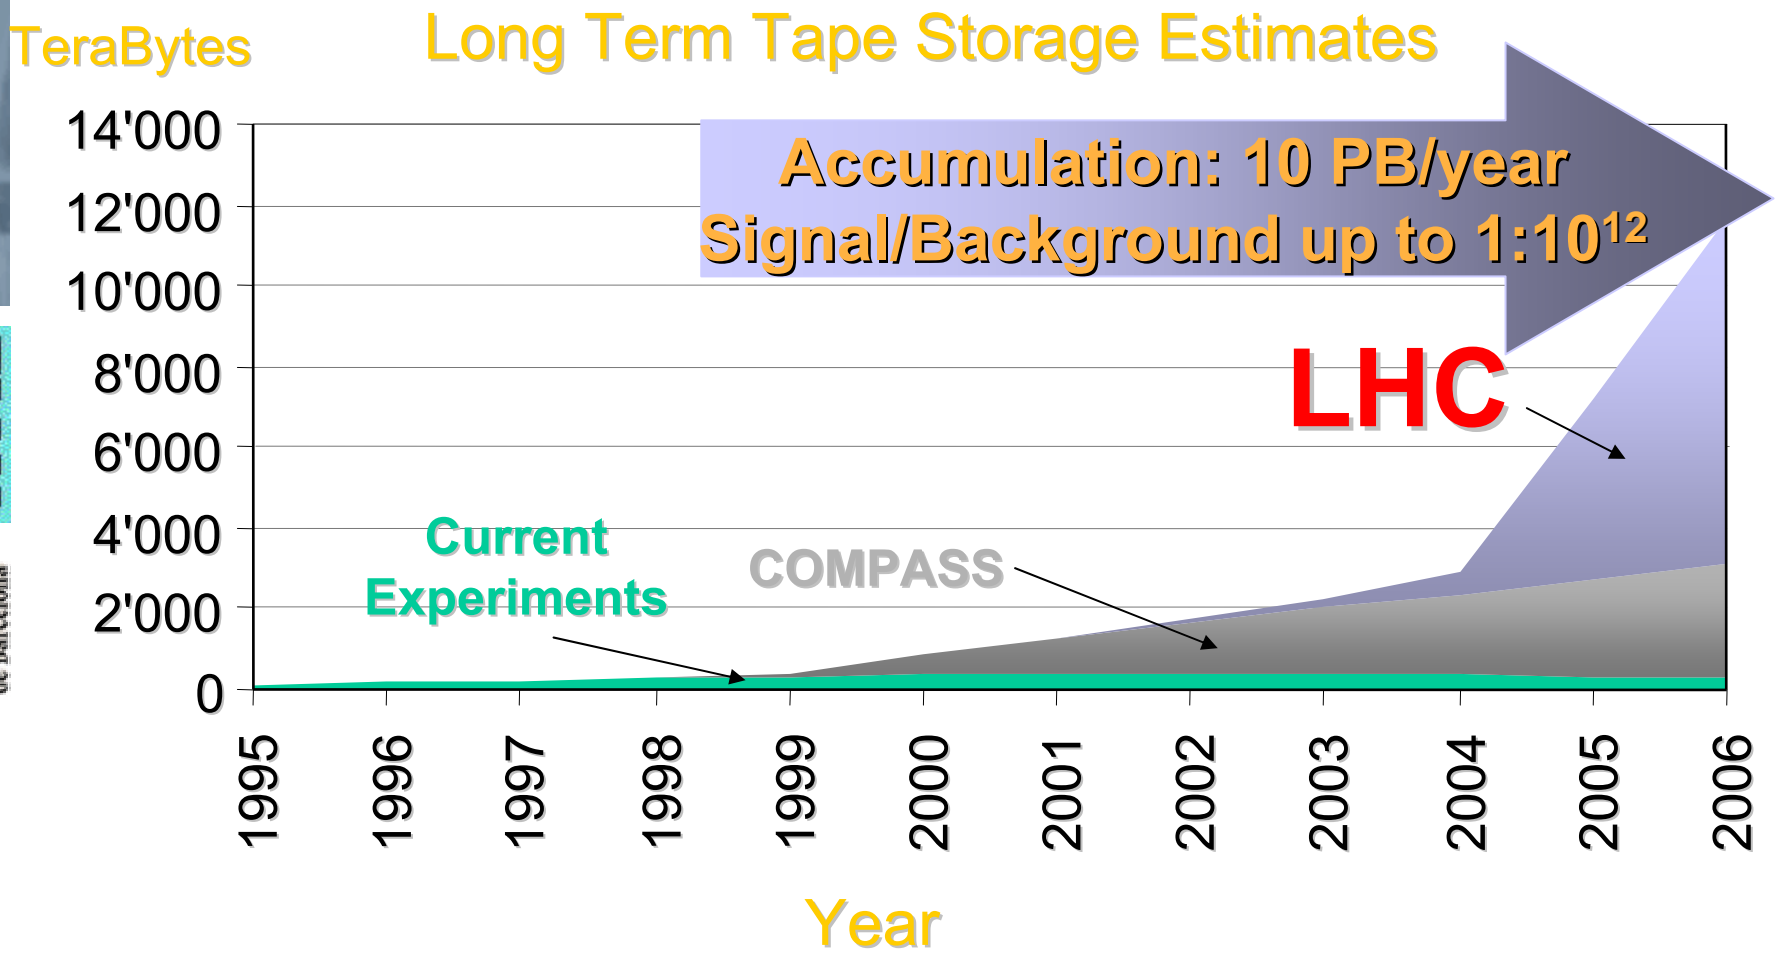

Complex LHC data = more CPU/byte

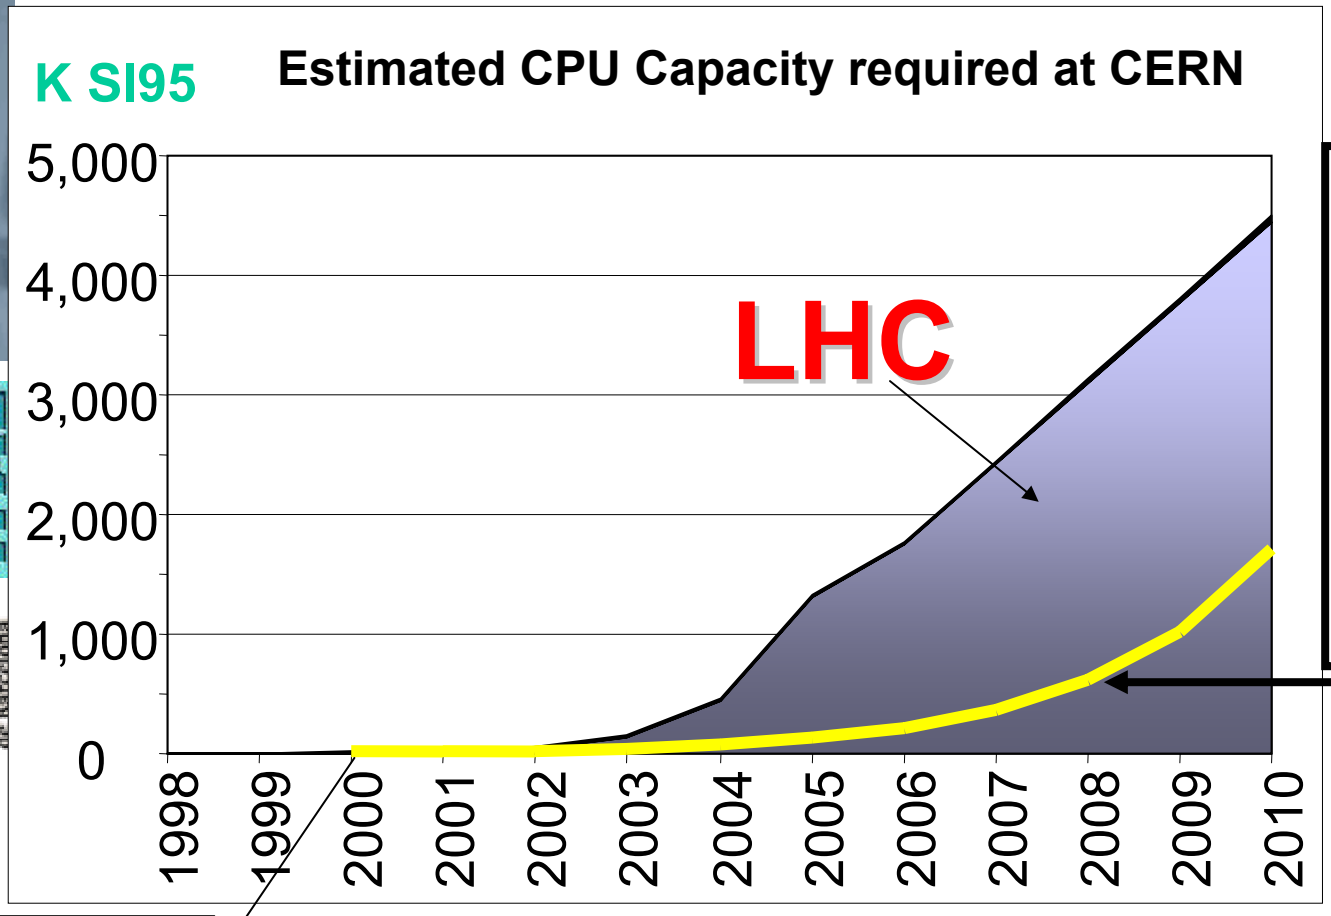

Moore's law – some measure of the capacity technology advances provide for a constant number of processors or investment

Jan 2000: 3.5K SI95

Port d'Informació Científica

LHC analysis will be driven by the demands of ~6000 physicists worldwide

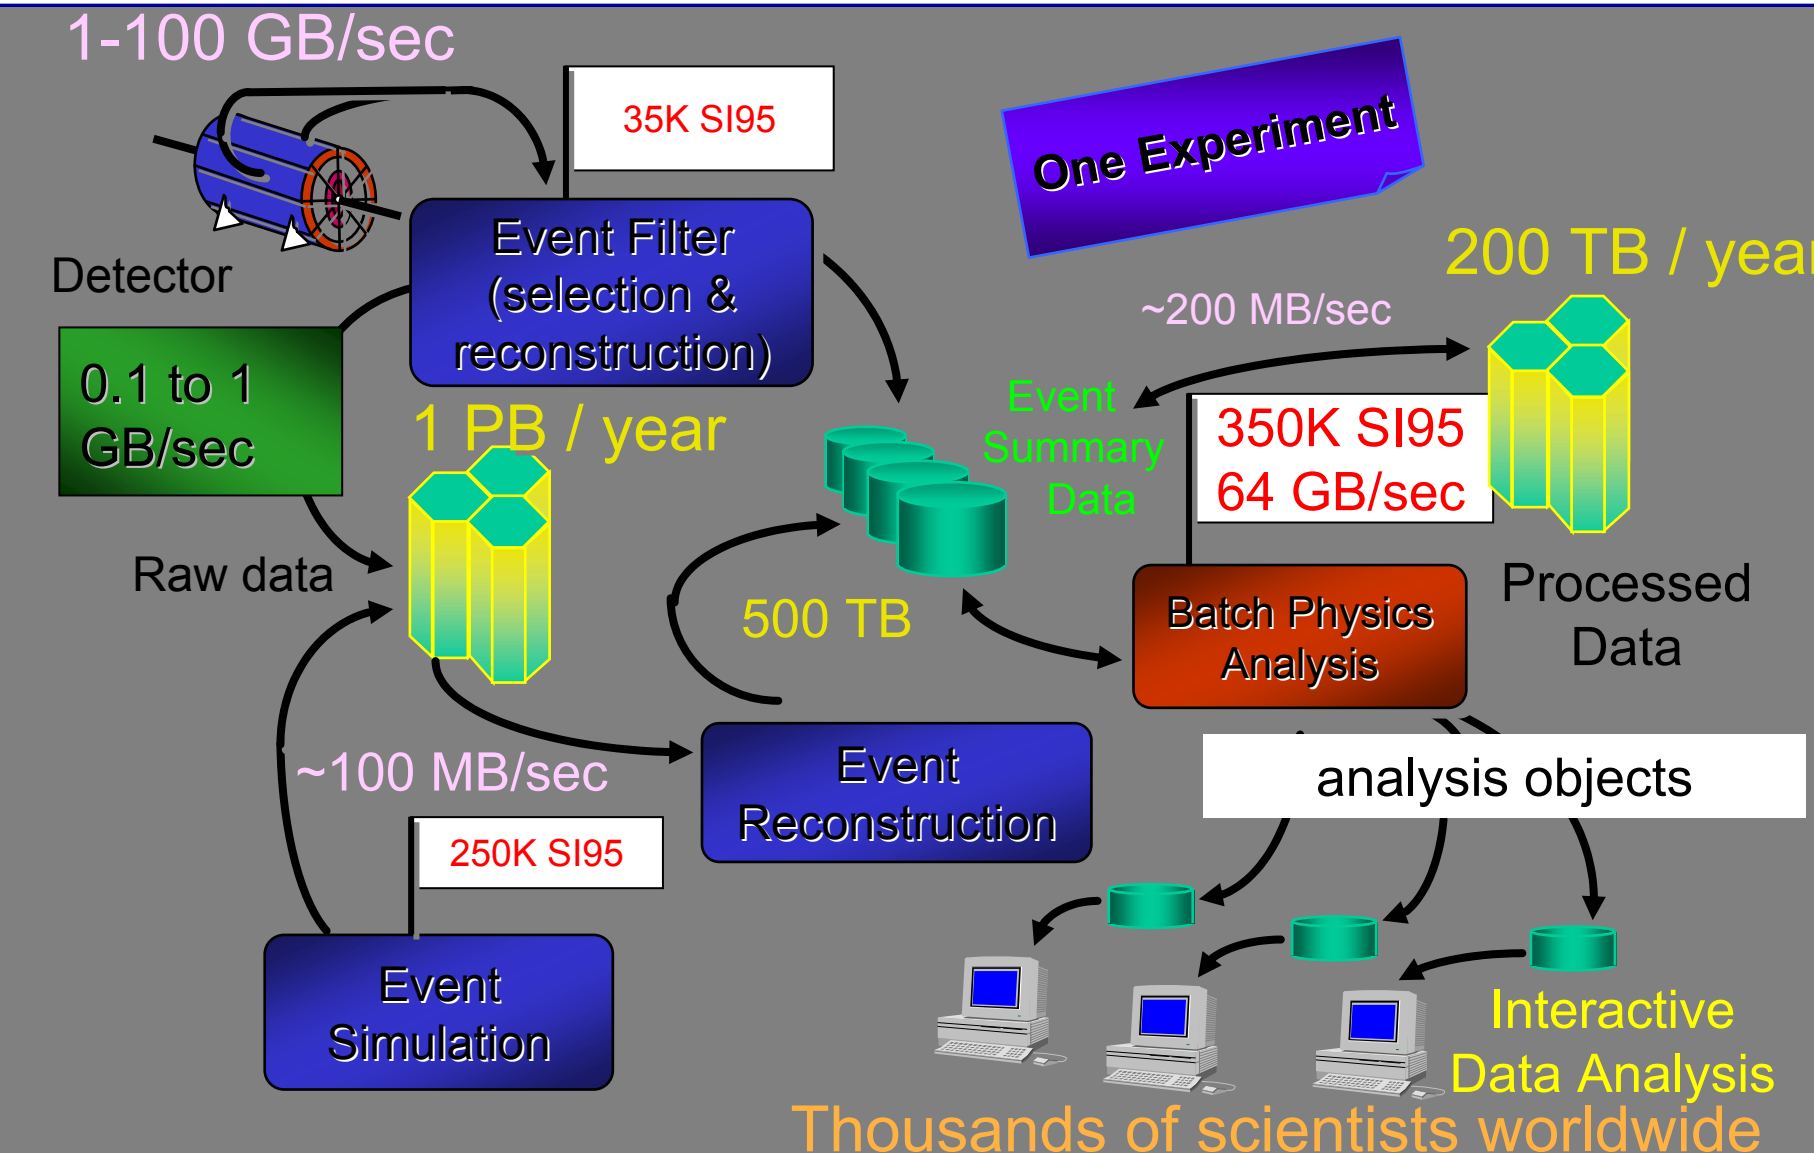

Technology Domains for developing solutions for LHC Data Analysis

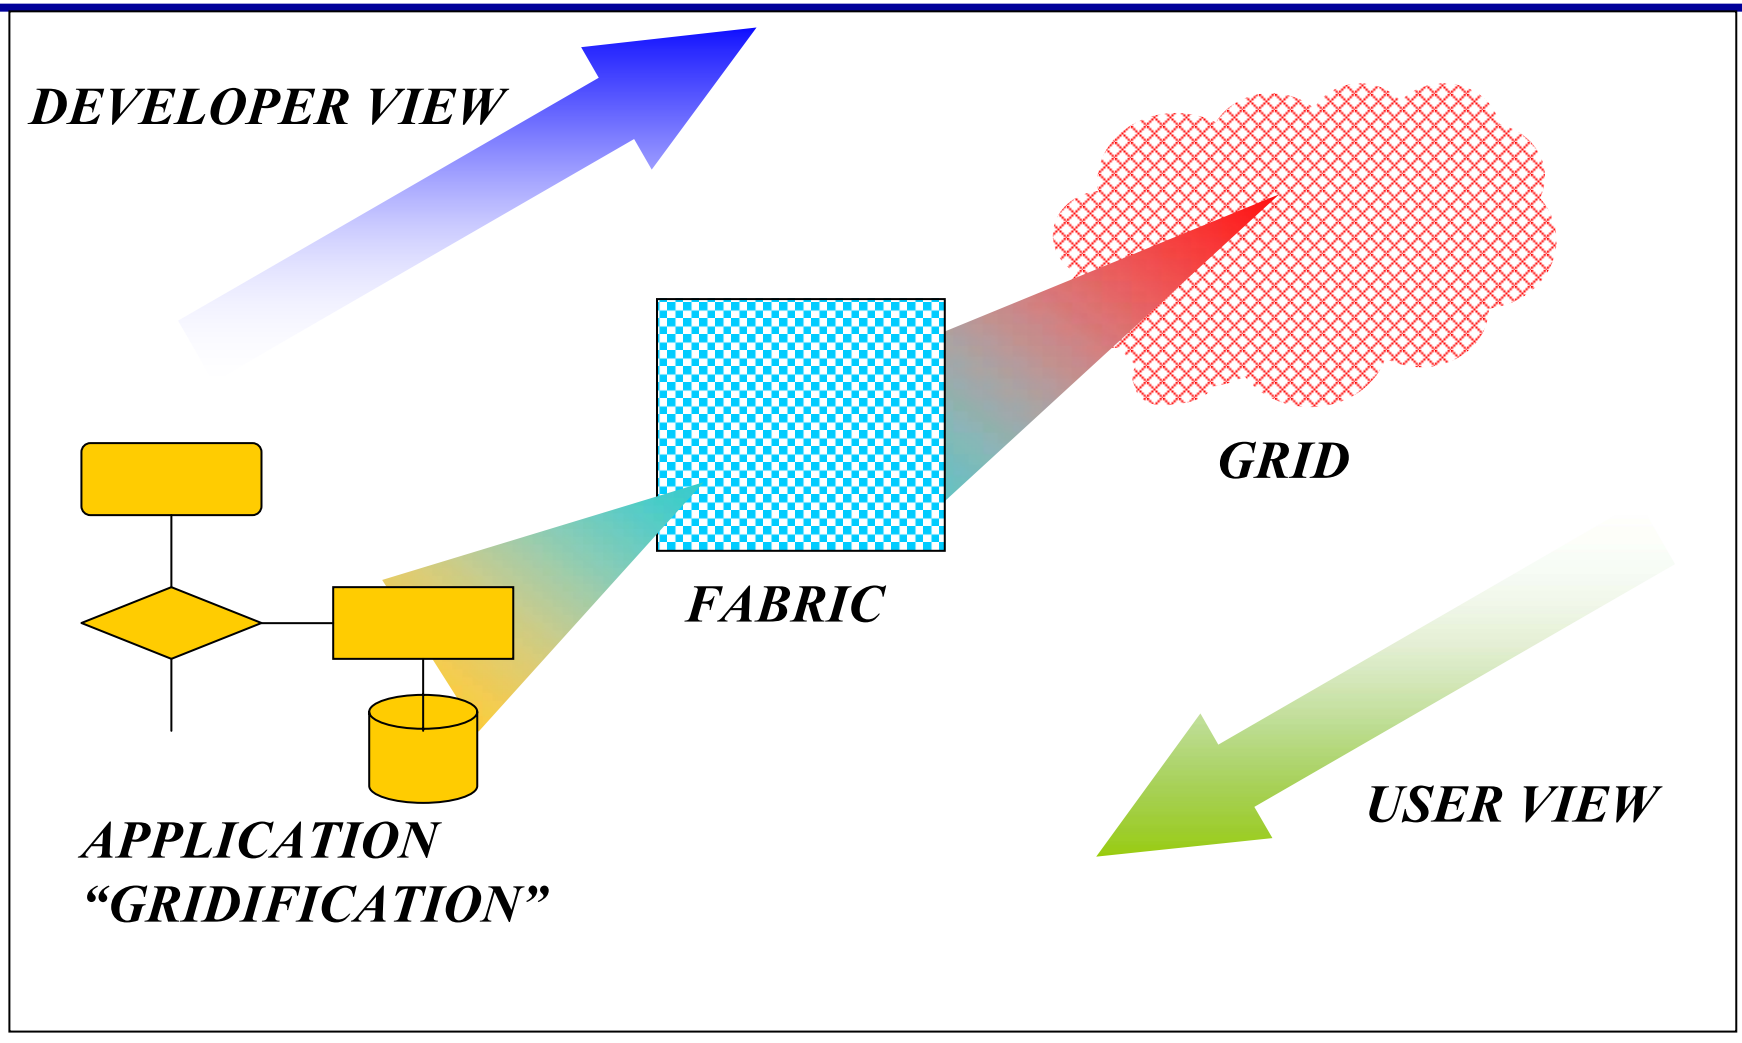

Port
'Informació
Científica

From physical to virtual communities

User point of view:
Virtual analysis communities

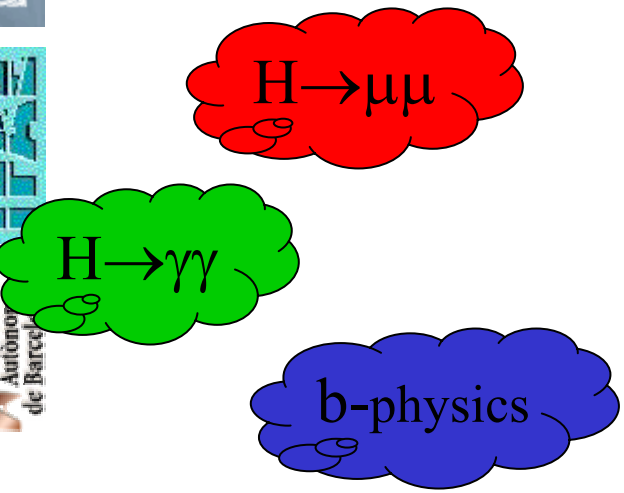

Worldwide multi-cultural collaboration

Physical setup

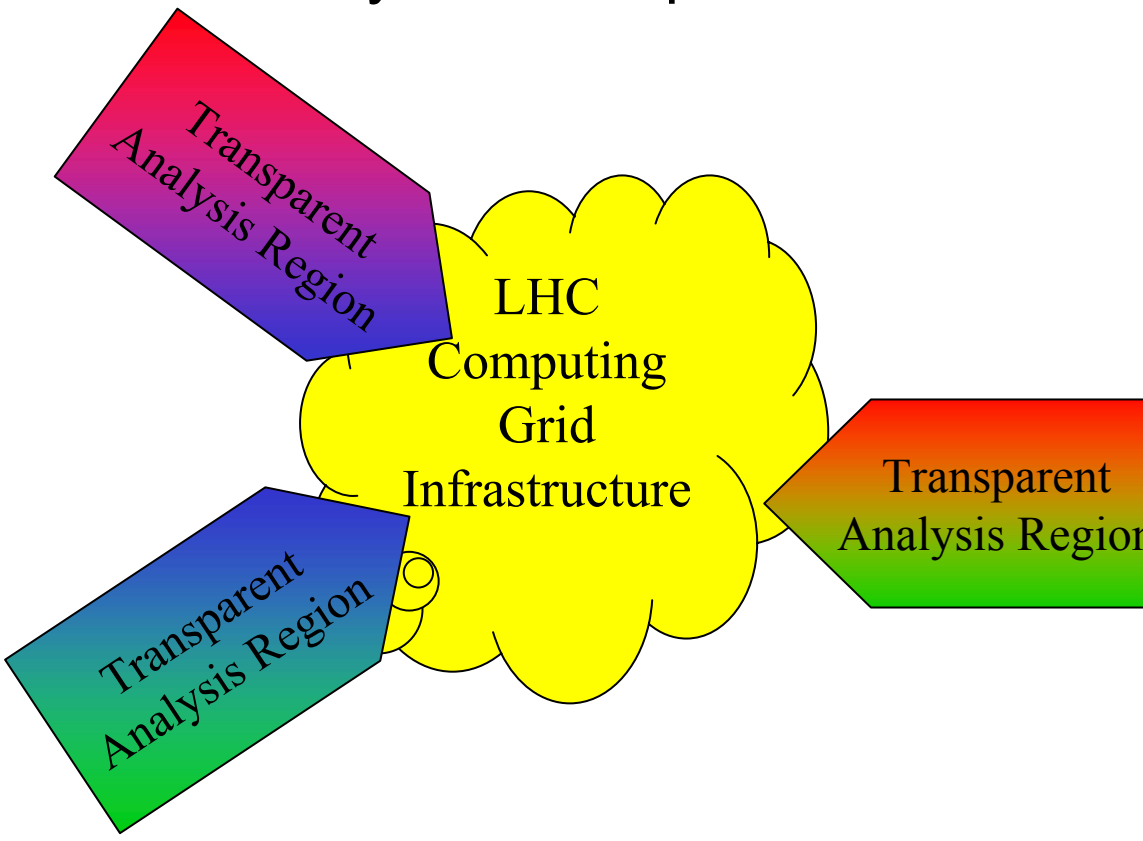

Institutional, National, Geographical Setups

Benefits of Grid Infrastructure: User-centric view

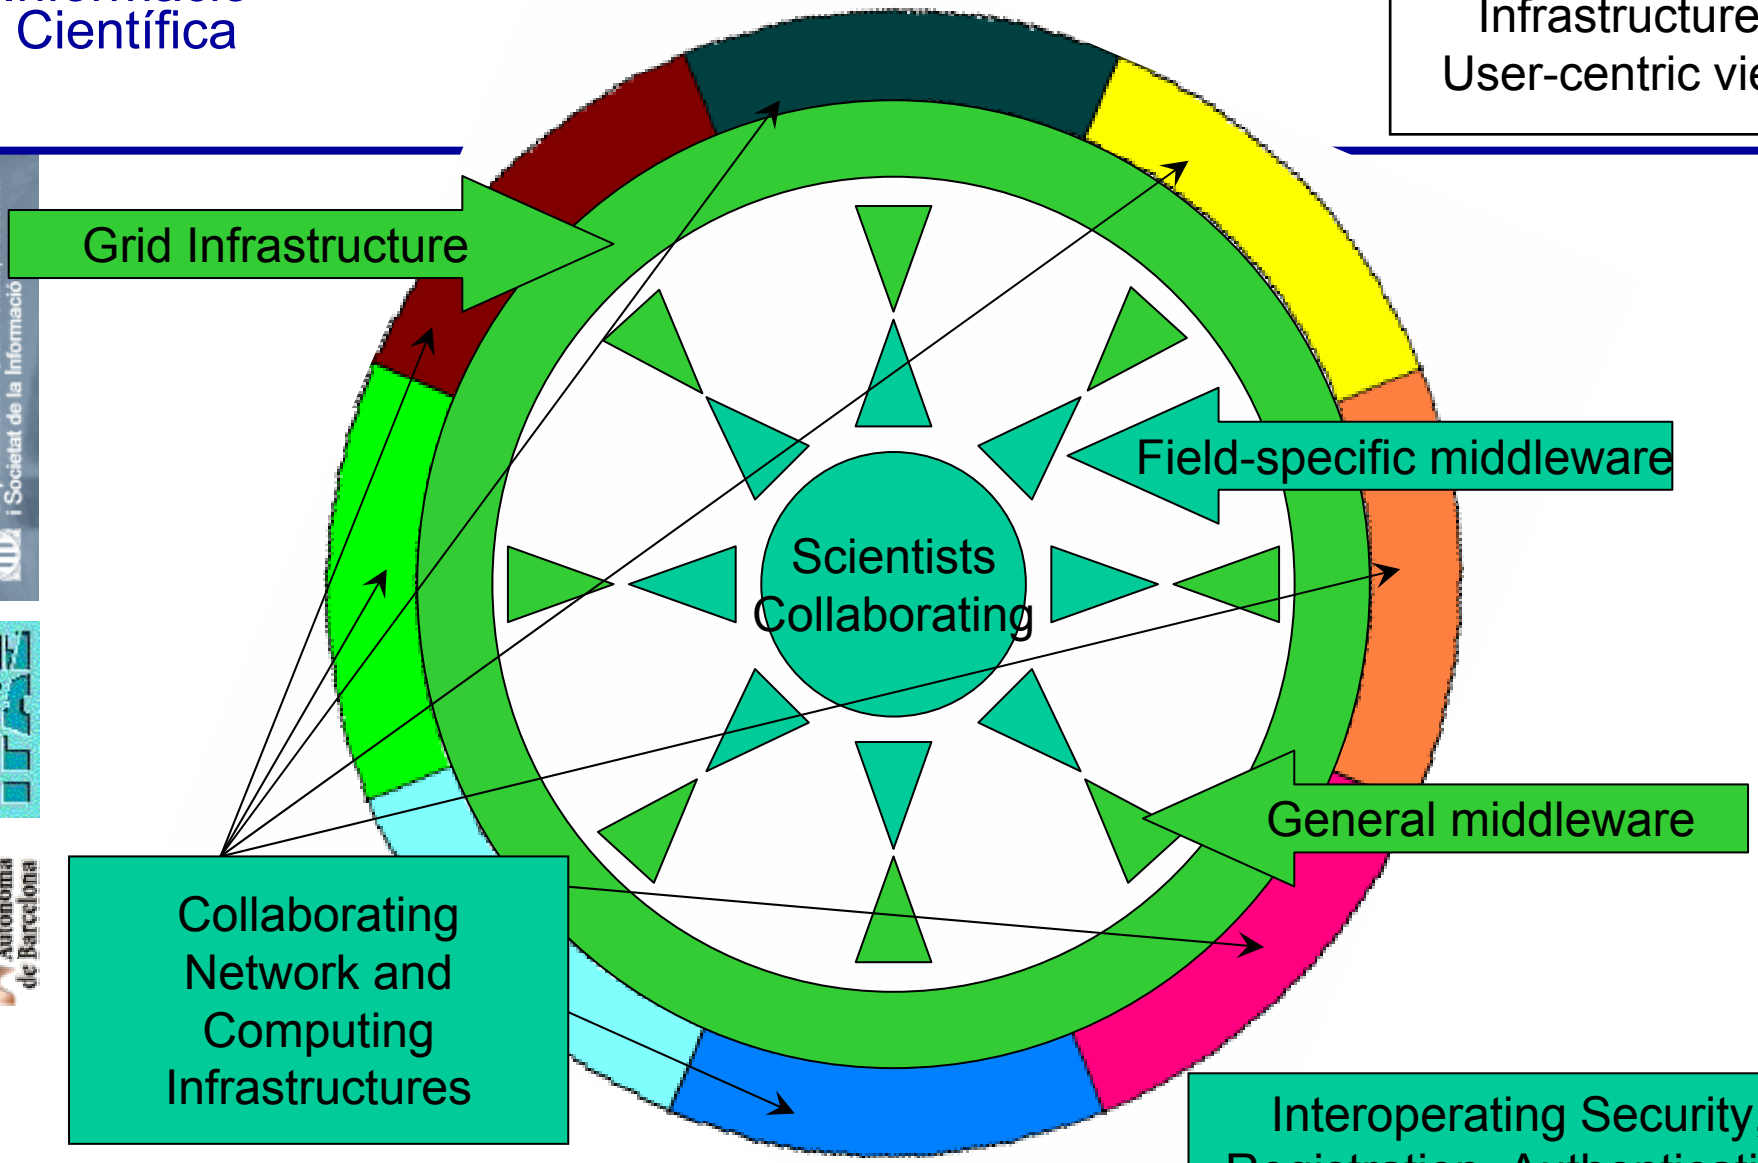

Port
d'Informació
Científica

The Internet connects Computers
The Grid enables Virtual Communities

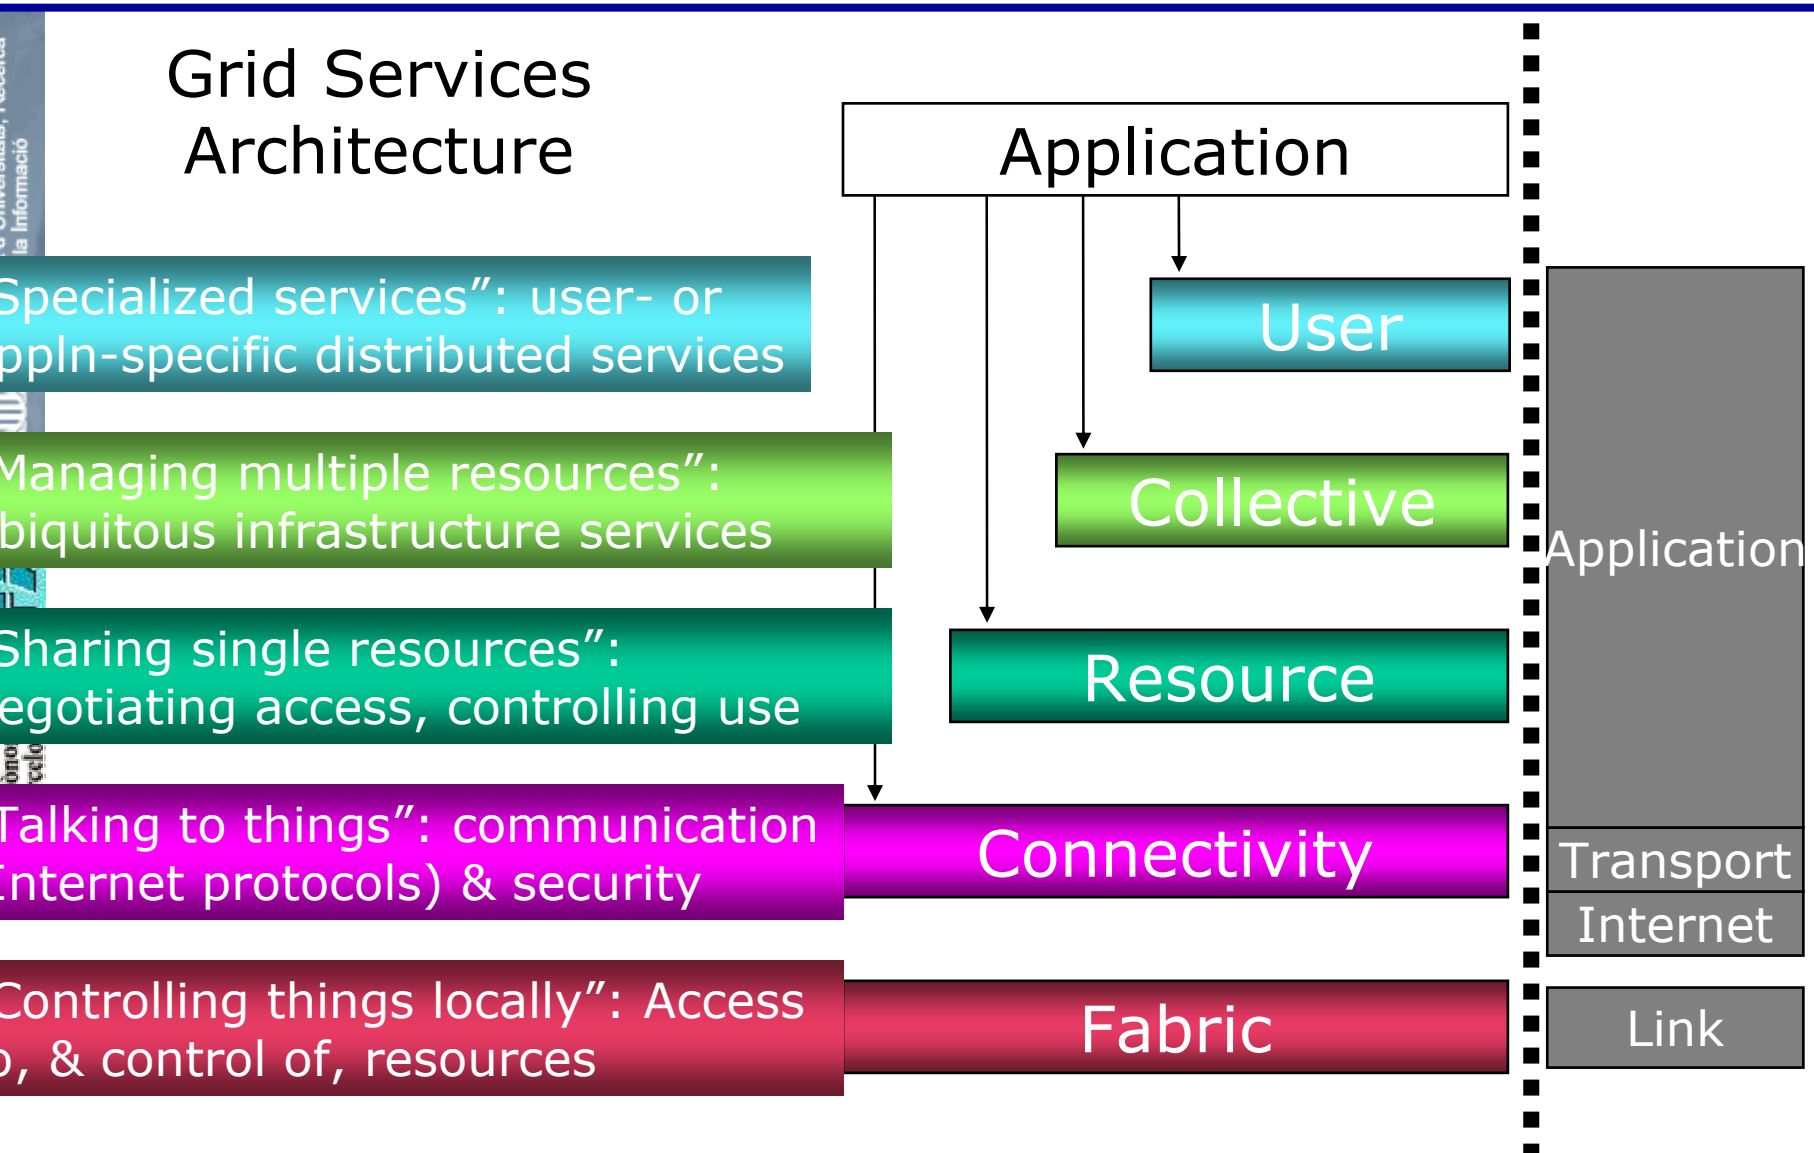

From “farms” to “computing fabrics”

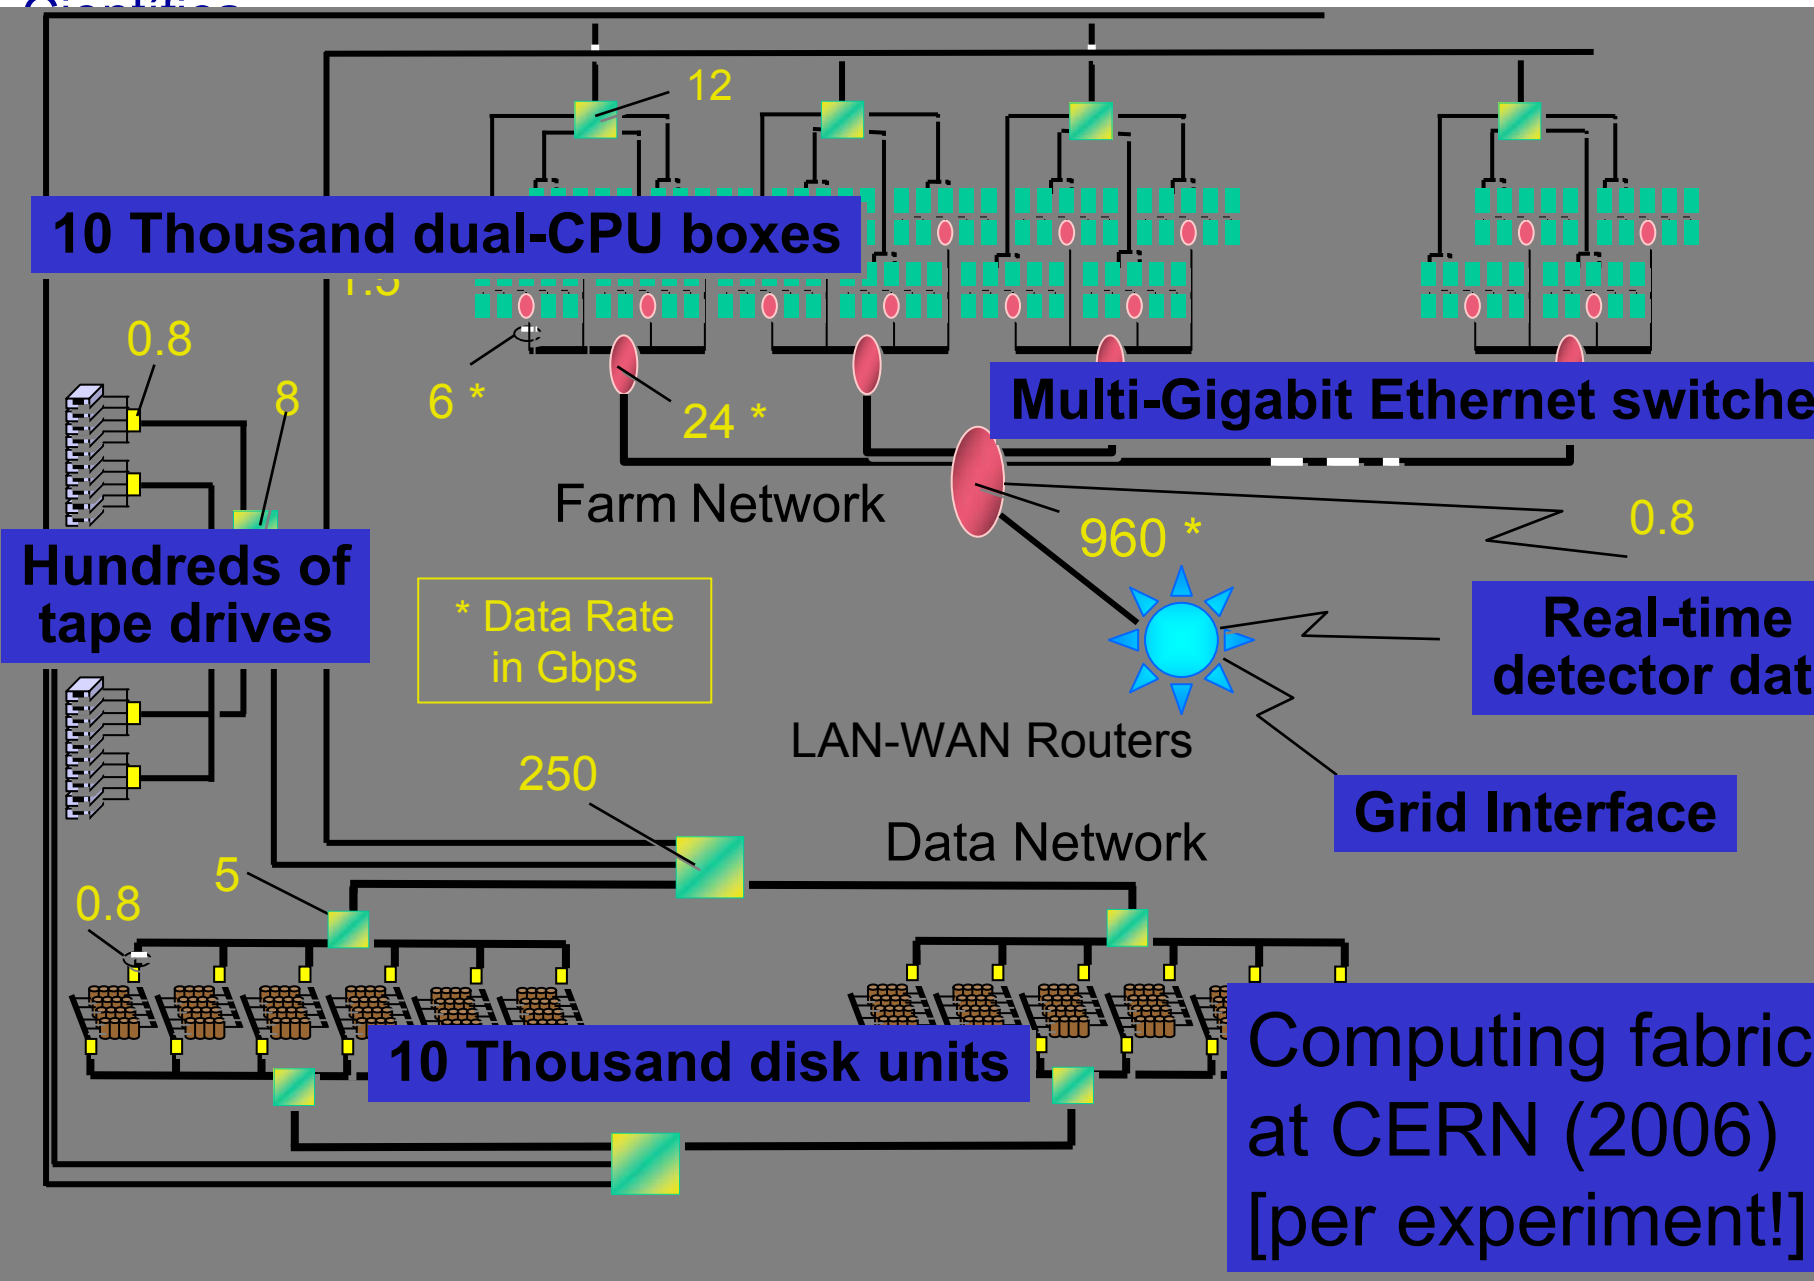

For a fixed Quality of Service for the Users:

- Costs will be dominated by
  the level of Automation and
  the reliability of Certification procedures.
- This in turn is related to the Architecture of the solution.

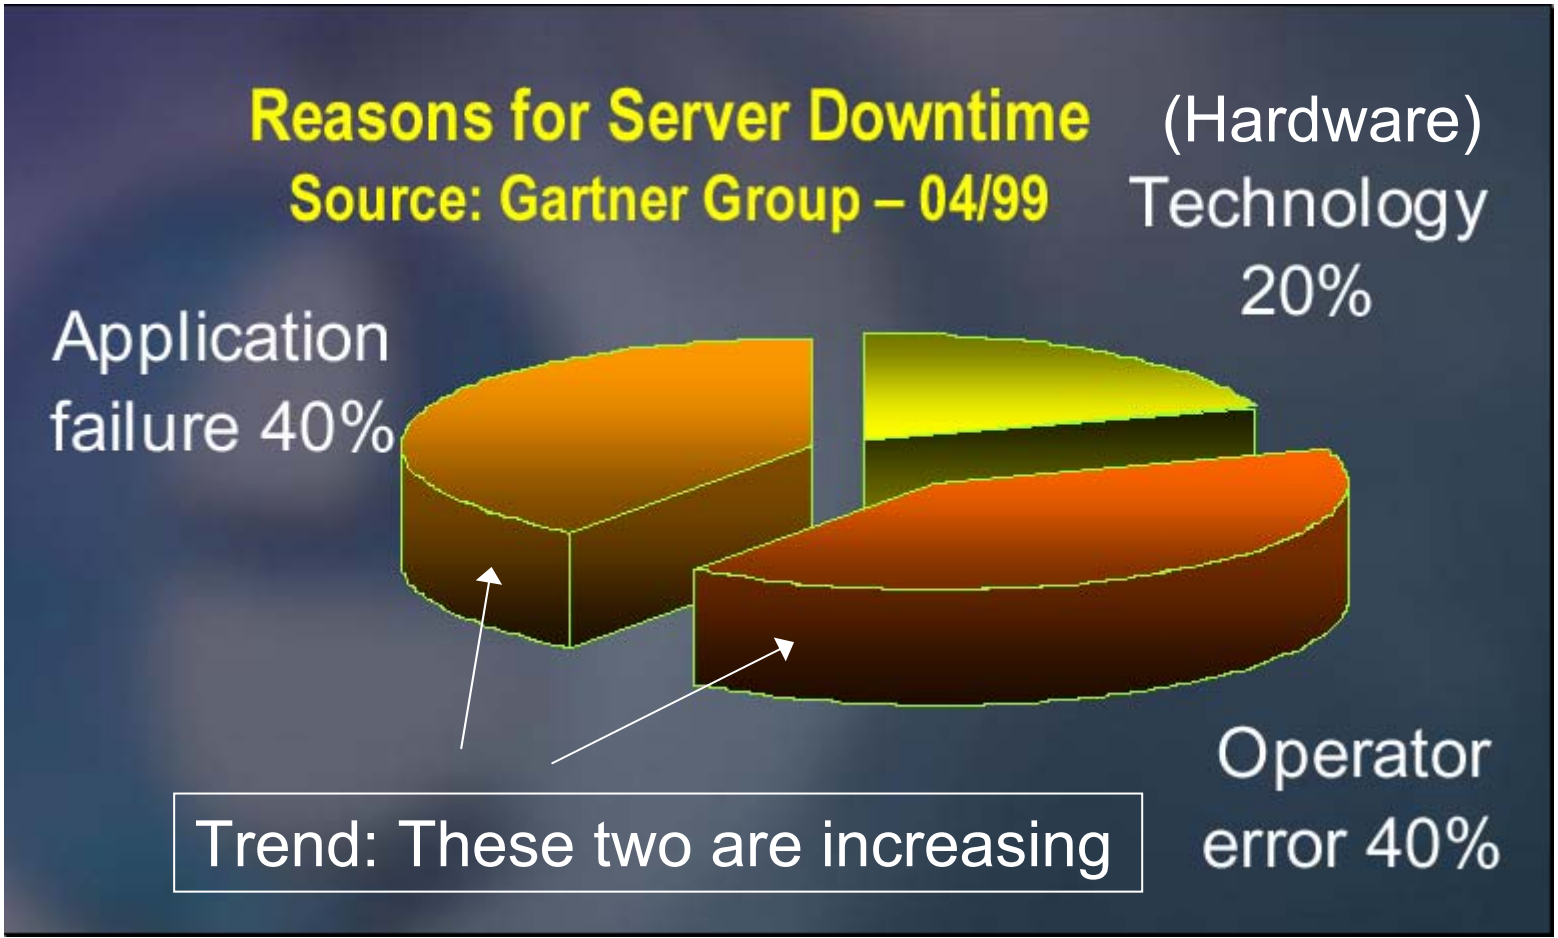

Port
d'Informació
Científica

- LCG-ES is a Coordinated Project of 7 Groups involved in 3 LHC experiments.

  It is a development and deployment project, not an R&D project.

  It is managed similar to an EU project.

  You can think of it as the overlapped "pre-production" , "pre-assembly" and "development of Engineering Design Report" of the Computing Subdetector in 36 months.

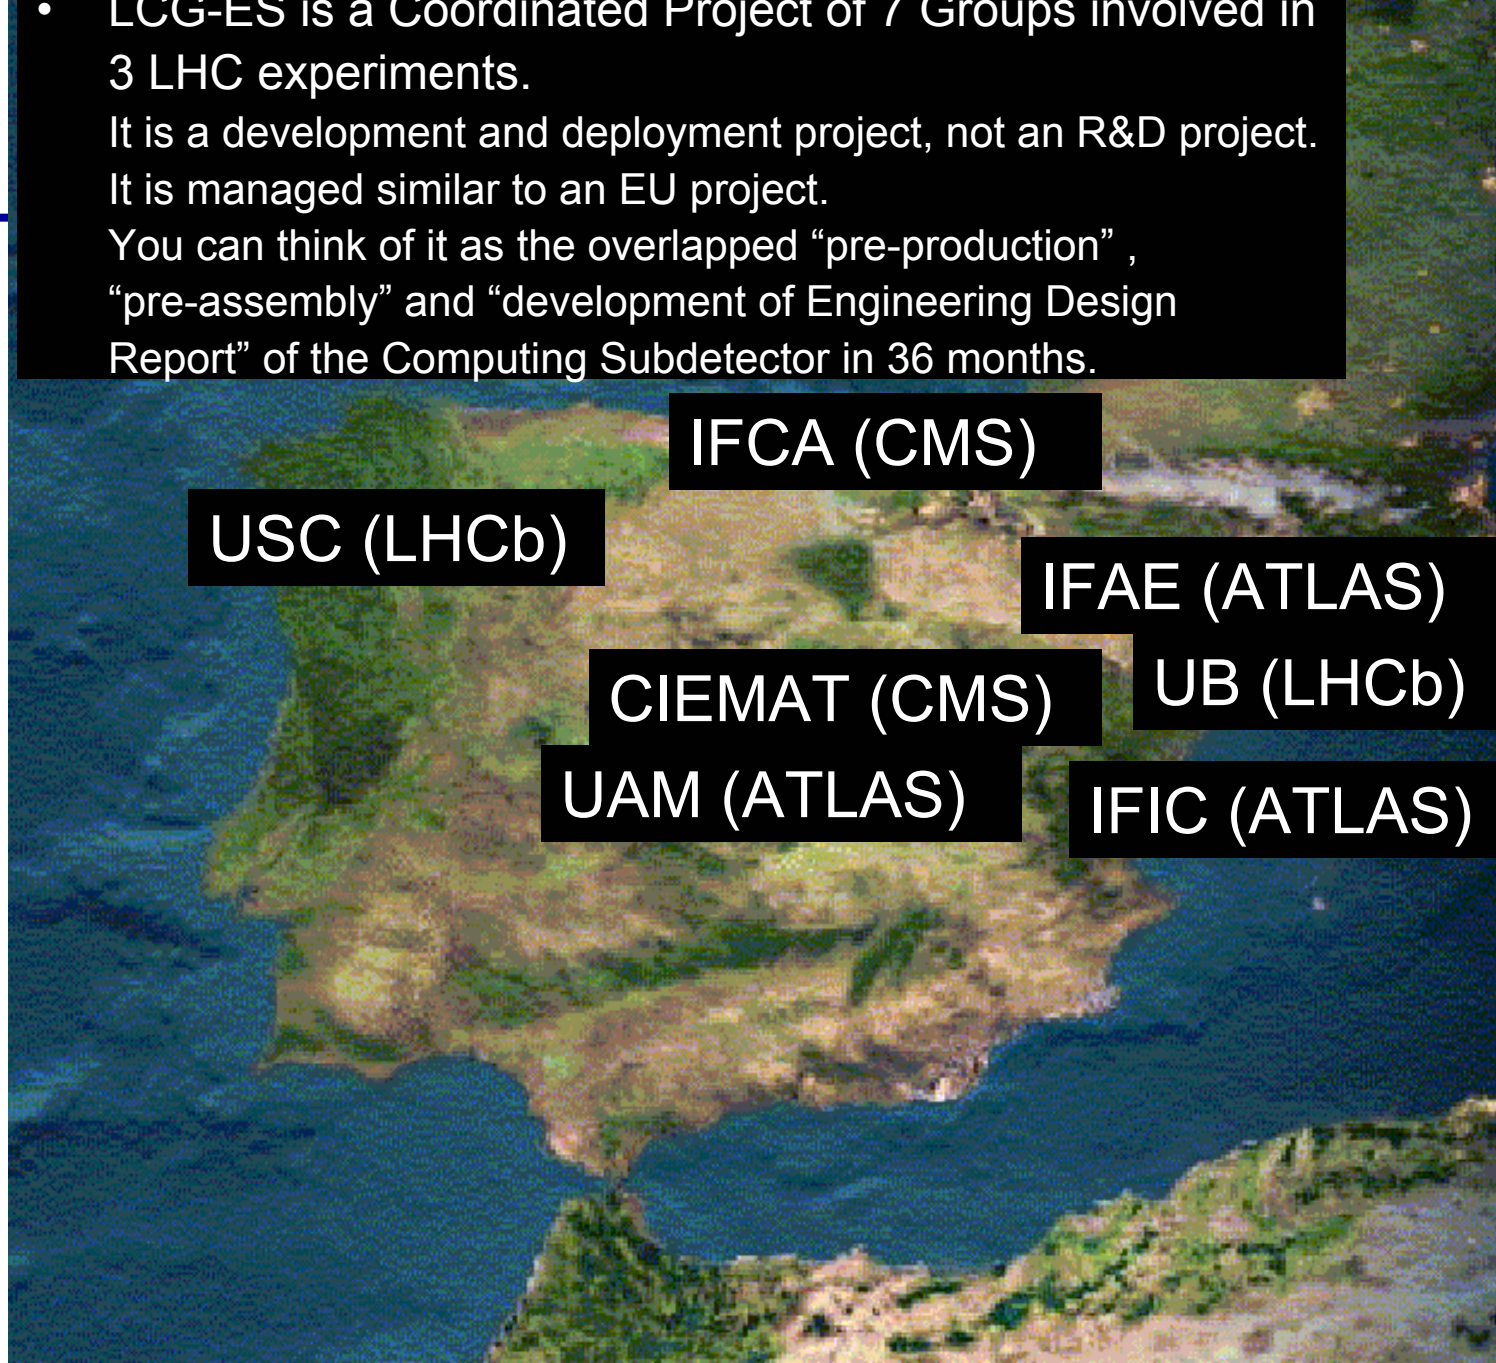

Background leading to LCG-ES

- Overall situation of Spanish groups in computing was analyzed and found to be too weak to make "just-in-time LHC ramp-up"
  - Analysis for LEP mostly done at CERN towards the end
  - General lack of technicians, system admins and sw professionals
  - Work on LHC dispersed, not fully recognized by LHC collaborations (in the shadows of "US ATLAS/CMS", "France", etc…)
- On the other hand, we had a small string of high quality activities
  - Various contributions to online and slow controls.
  - FALCON quasi-online farm for ALEPH (+ Florida reprocess farm)
  - CERN RD-47 (HEP Data Processing using Commodity Components)
  - R&D in Object Databases for HEP Data
  - Participation in EU DataGrid and EU CrossGrid
- Initial boosts from "ad-hoc" mechanisms, lobbying and FPA
  - 2 multi-disciplinary farms (~200 nodes) at IFIC and IFCA
  - "Special Action" for 1 year of 500 K€ (about 50% materials, 50% personnel & travel) loosely coordinated by Jesus Marco (IFCA)

Port
d'Informació
Científica

LCG-ES: High-Level Objectives

- Develop and test under realistic conditions a series of prototypes, increasingly more powerful and complex, of data processing systems for the LHC utilizing Grid techniques.
- Research the possibilities of sharing complex data processing infrastructures between several communities (those of ATLAS, CMS and LHCb in our case).

Coordinated within the global LCG project plan and the Data Challenges of the experiments

Finer-grain objectives (1 of 3)

1. Establish a solid base of knowledge and experience in Grid distributed computing in order to ensure participation of Spanish groups in LHC data analysis through a shared infrastructure.
2. Establish a solid base of knowledge and experience in management and processing of the large data volumes (Petabyte level) that will be generated by the LHC experiments.
3. Actively participate in Grid computing projects in order to ensure the deployment in Spain of a Grid and DataGrid infrastructure which is integrated internationally.

Finer-grain objectives (2 of 3)

4. Take active part in development of Grid computing within the LHC experiments and contribute to the LHC Computing Grid Technical Design Report.
5. Create Monte Carlo production centers such that groups can contribute to the official productions of the LHC experiments. In particular:
   - Generate official MC consisted with Spanish participation in the LHC program and the objectives of groups within their experiment.
   - Explore efficient and affordable methods to deliver the data of these official productions to the international collaborators in the experiments.

Finer-grain objectives (3 of 3)

6. Deliver to the groups a local software development environment (for simulation, reconstruction and analysis) appropriate for their experiment.
7. Facilitate the transfer of the experience acquired fostering the possible re-use of this knowledge in other HEP projects or in other disciplines.

Work Methodology

- Major agreements by groups
  - Commit to give services to each other
  - Allow jobs from any experiment to run at institute's facilities (with some safeguards)
- Opened the possibility to weave deliverables together into a "Spanish Grid Economy" linked to the global LCG Grid
- Forced to use build-to-cost, as cost estimates are difficult with fast-moving technologies and demand seems to far exceed reasonable funding expectations in 2003-2005
- Allow a reasonable amount of personnel "scaling" in order to study farm and data management automation to minimize personnel needs in the production system

Risk Analysis

- Major agreements by groups (R1: Management support)
  - Commit to give services to each other
  - Allow jobs from any experiment to run at institute's facilities (with some safeguards)
- Opened the possibility to weave deliverables together into a "Spanish Grid Economy" linked to the global LCG Grid
  (R2: Will Grid techniques deliver. Alternatives?)
- Forced to use build-to-cost, as cost estimates are difficult with fast-moving technologies and demand seems to far exceed reasonable funding expectations in 2003-2005
- (R3: Fulfilling official MC. R4: Production scale-up)
- Allow a reasonable amount of personnel "scaling" in order to study farm and data management automation to minimize personnel needs in the production system
  (R5: "Ratchet" effect, local vs national personnel policy)

LCG-ES Budget summary

Concepto	Personal	Equipamientos	Viajes, Dietas, Otros gastos	P/(P+M+O)
Contribución LCG-CERN Fase 1 Años 2002-2004	0 €	300,000 €	0 €	0%
PIC (CIEMAT Arquitecto de Seguridad)	115,000 €	0 €	0 €	100%
PIC (IFAE)	0 €	250,000 €	50,000 €	0%
Grupo CIEMAT	115,000 €	143,750 €	44,850 €	38%
Grupo IFAE	230,000 €	101,280 €	20,000 €	65%
Grupo IFCA	230,000 €	136,000 €	48,000 €	56%
Grupo IFIC	230,000 €	135,700 €	50,140 €	55%
Grupo UAM	184,000 €	74,761 €	9,235 €	69%
Grupo UB	92,000 €	28,750 €	7,130 €	72%
Grupo USC	115,000 €	51,560 €	30,320 €	58%
Subtotal (excl. LCG-CERN)	1,311,000 €	921,801 €	259,675 €	53%
Original Request (excl. LCG-CERN)	1,447,500 €	1,951,500 €	0 €	43%
Request/Granted	91%	61%		
LCG-CERN/Total	11% (plus 5 CDTI fellows plus possibly 5 more this year)			

Notes:

1. Spanish contribution to LCG-CERN (300 K€) handled administratively through IFAE
   (Additional 100 K€ for LCG-CERN Phase 1 to be funded in 2005)
2. Project budget is 50% human resources. If you include human resources funded by institutes and universities, it is approximately 2/3 people. Consistent with Gartner Group????

LCG-ES Personnel details (estimates)

Personnel per Work Package

IP Total	0.80
EAD Total	4.00
EDS Total	3.55
RSG Total	1.25
GSW Total	5.38
MCF Total	5.45
GVM Total	2.90
CTS Total	1.50
ETD Total	3.05
GDC Total	1.30
GRD Total	1.00
SEG Total	1.45
PIC Total	0.90
CDC Total	1.40
Grand Total	33.93

Personnel per Group and Funding Source

Group	Other	LCGES	Total
CIEMAT	3.00	3.00	6.00
IFAE	1.00	2.00	3.00
IFAE-PIC	2.50	2.00	4.50
IFCA	3.50	2.00	5.50
IFIC	3.50	2.50	6.00
UAM	1.20	2.00	3.20
UB	1.00	1.00	2.00
USC	0.83	1.00	1.83
Grand Total	16.53	15.50	32.03

Port
'Informació
Científica

Deliverables to fulfill the objectives

EAD, MCF, GSW, CDC

IFCA

EAD, GSW, CDC

USC

EAD, ETD, PIC, EDS, GSW

IFAE

EAD, MCF, EDS, RSG, SEG, CTS

CIEMAT

UB

EAD,EDS, CTS

UAM

EAD, GVM

IFIC

EAD, MCF, CTS, CDC

- Stay away from glitz.
- Concentrate on deployment, MC & analysis
- Use local Univ. for TT to other disciplines
- 600 KCHF materials contribution to LCG-CERN

AD = Analysis Farm
DS = SW Dev Platform
SG = SW repository
SW = SW Gridification
CF = MC Fabric
VM = Virtual MC farm
TD = Data Transform
IC = Gridified Data Store
EG = Security Architect
TS = Tech MC Support
DC = Data Chal. Coord.

Port
'Informació
Científica

Use Grid technology to generate a common infrastructure

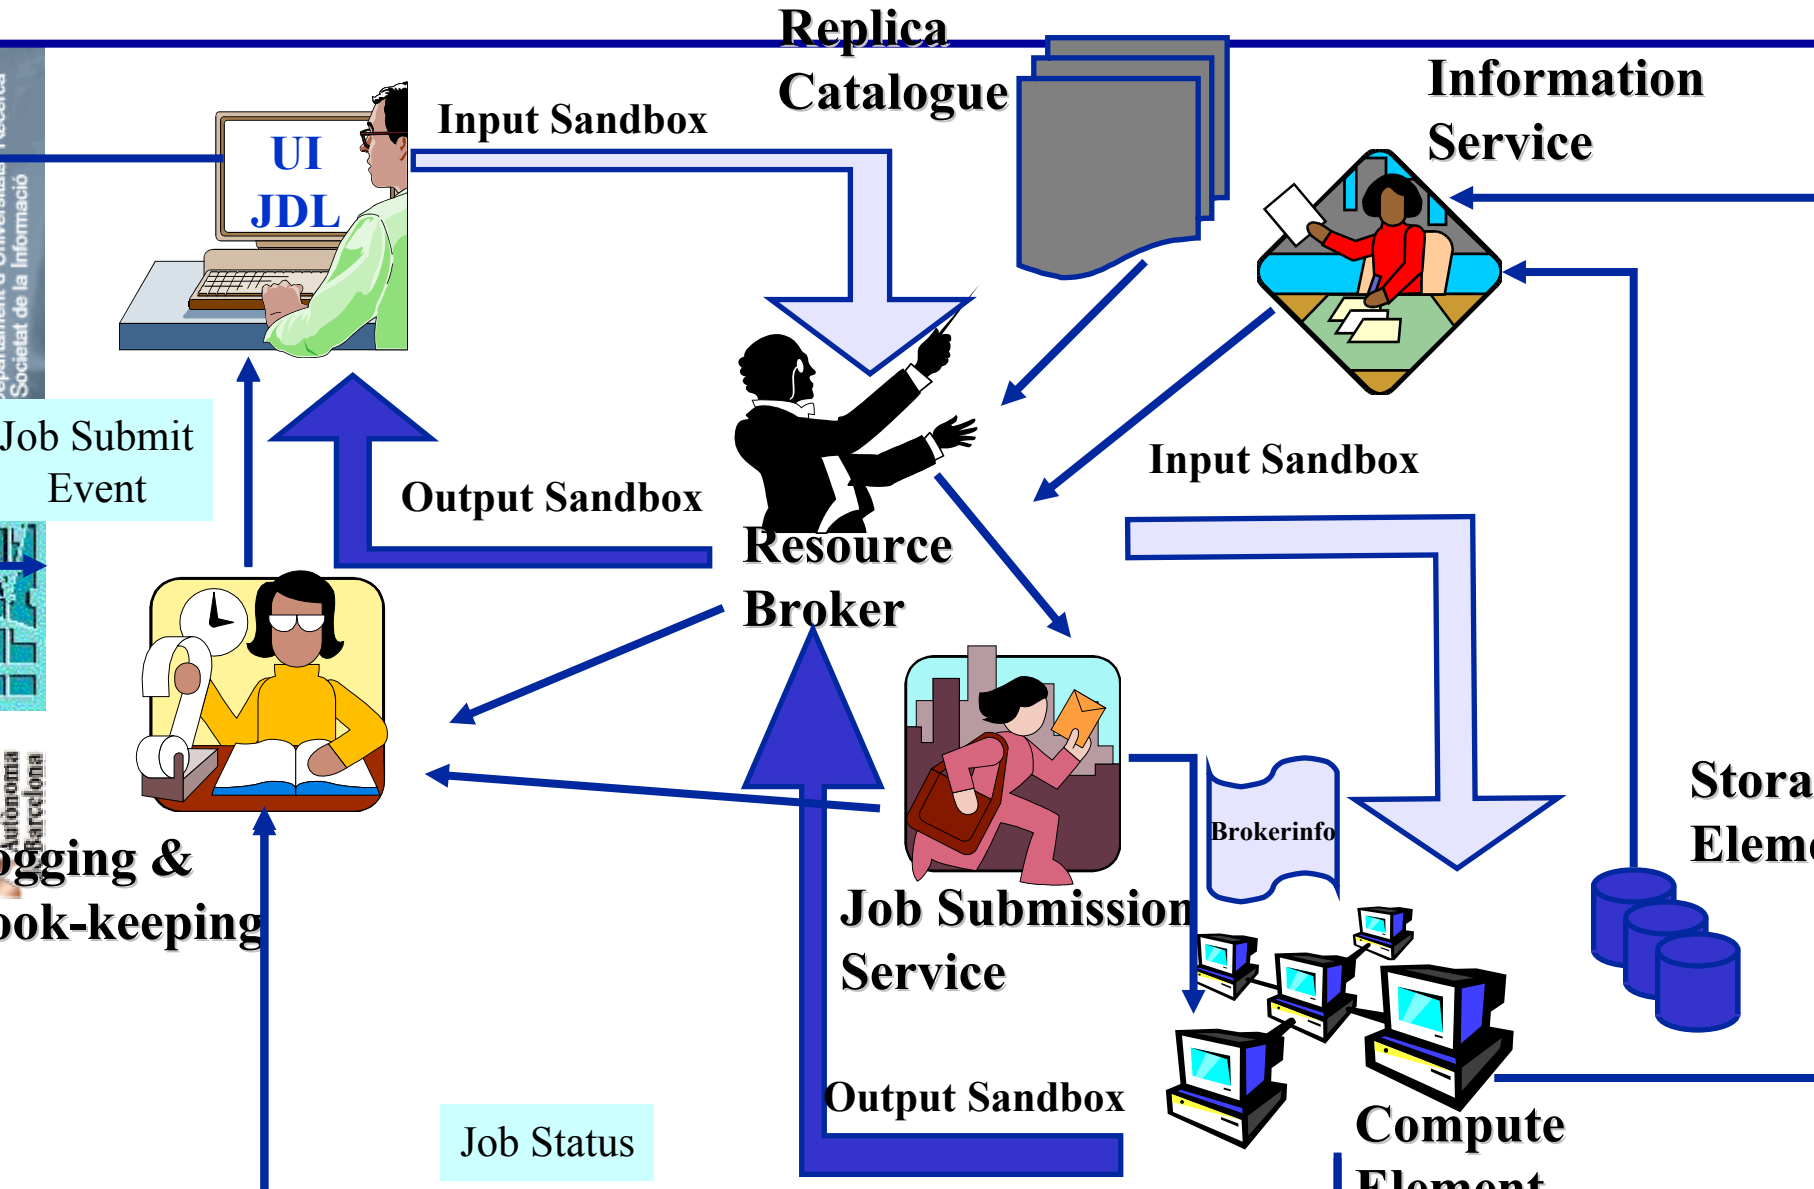

Port d'Informació Científica

Functionality developed by EDG, Globus, etc.
brought to high performance level
by LCG + Data Challenges

EU DataGrid Monitor

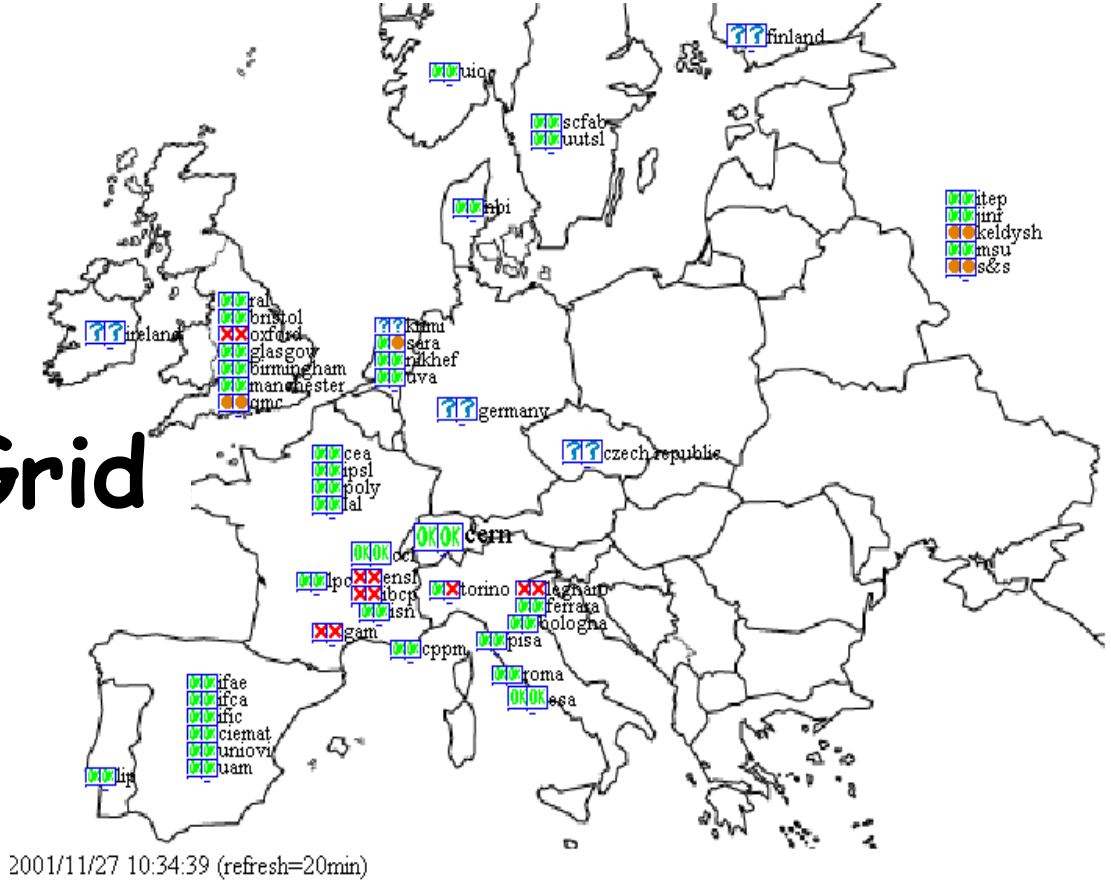

http://ccwp7.in2p3.fr/mapcenter/

Port
d'Informació
Científica

The LHC Computing Model: 1999-2002 and still evolving

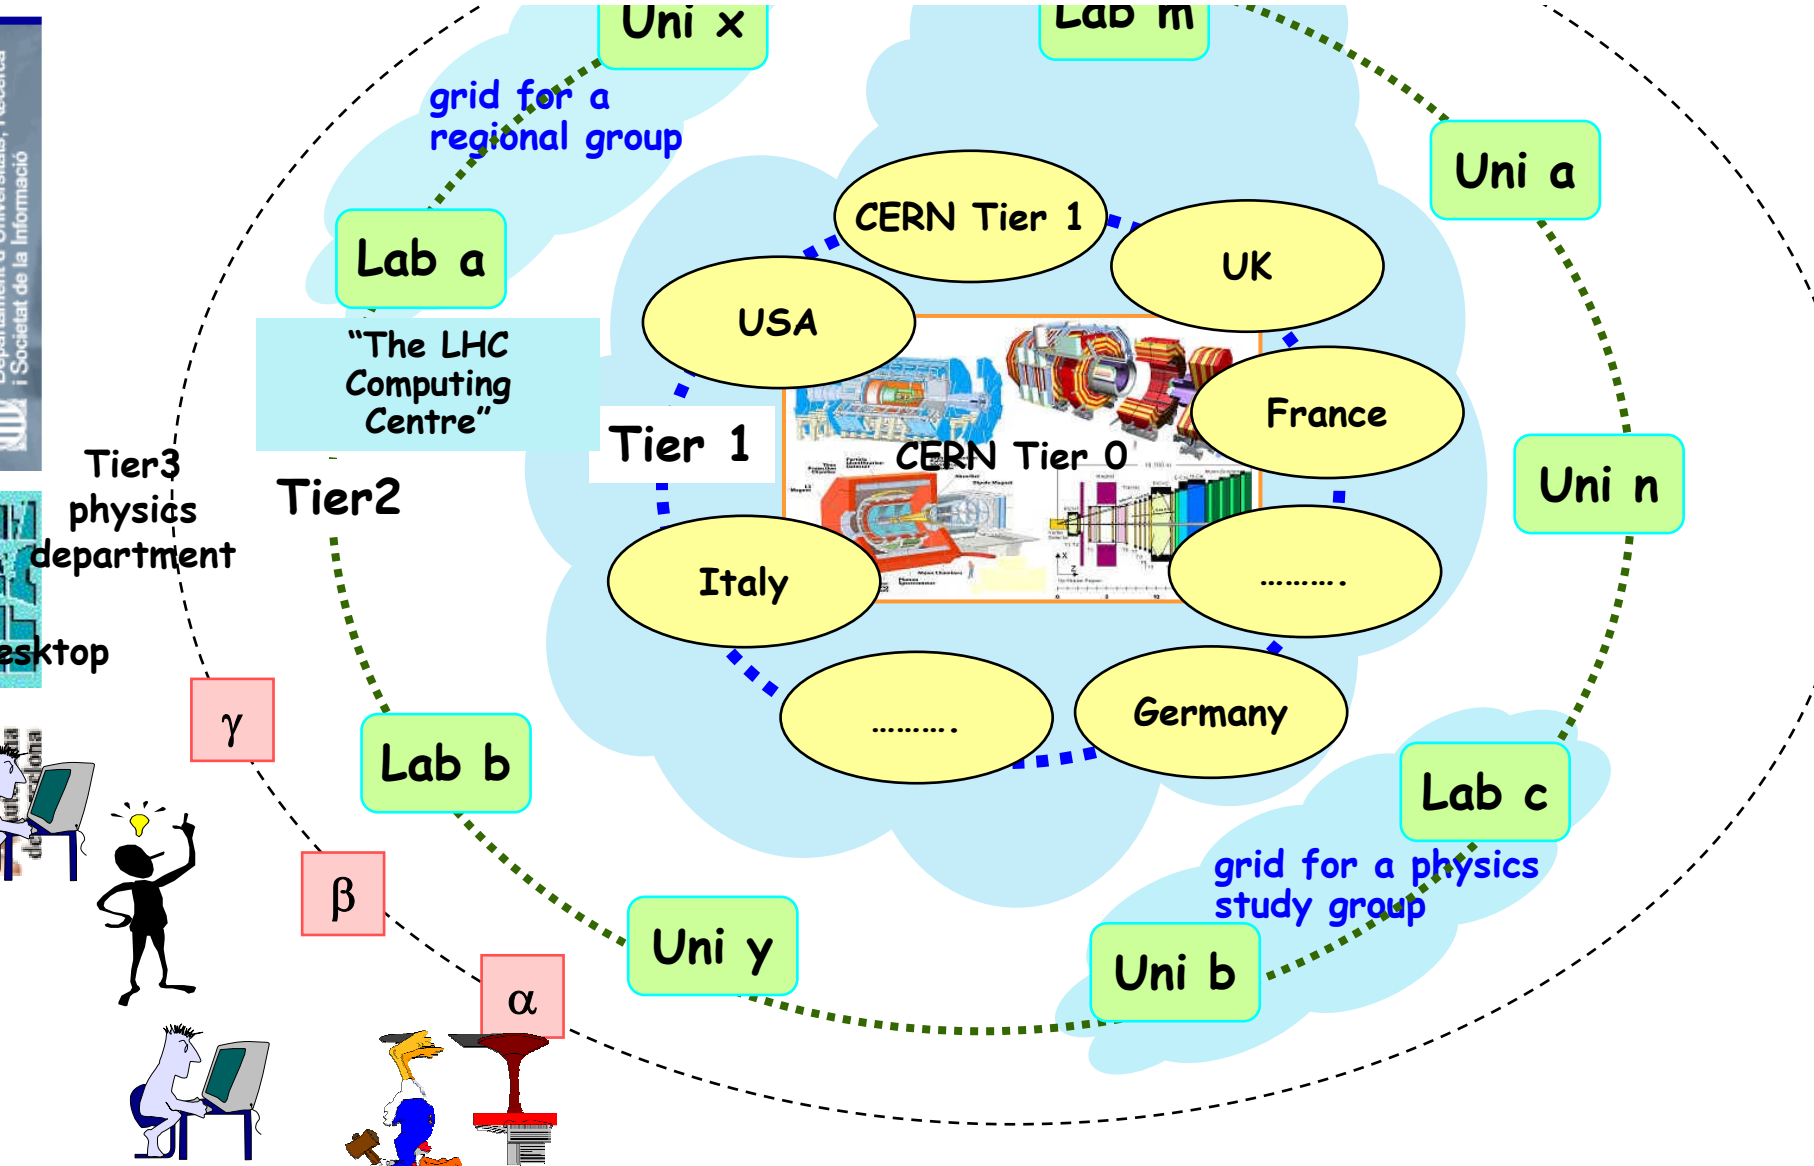

Port
d'Informació
Científica

The "Tier-1 Dilemma"

Classic T1:

Serve 10% of a global experiment

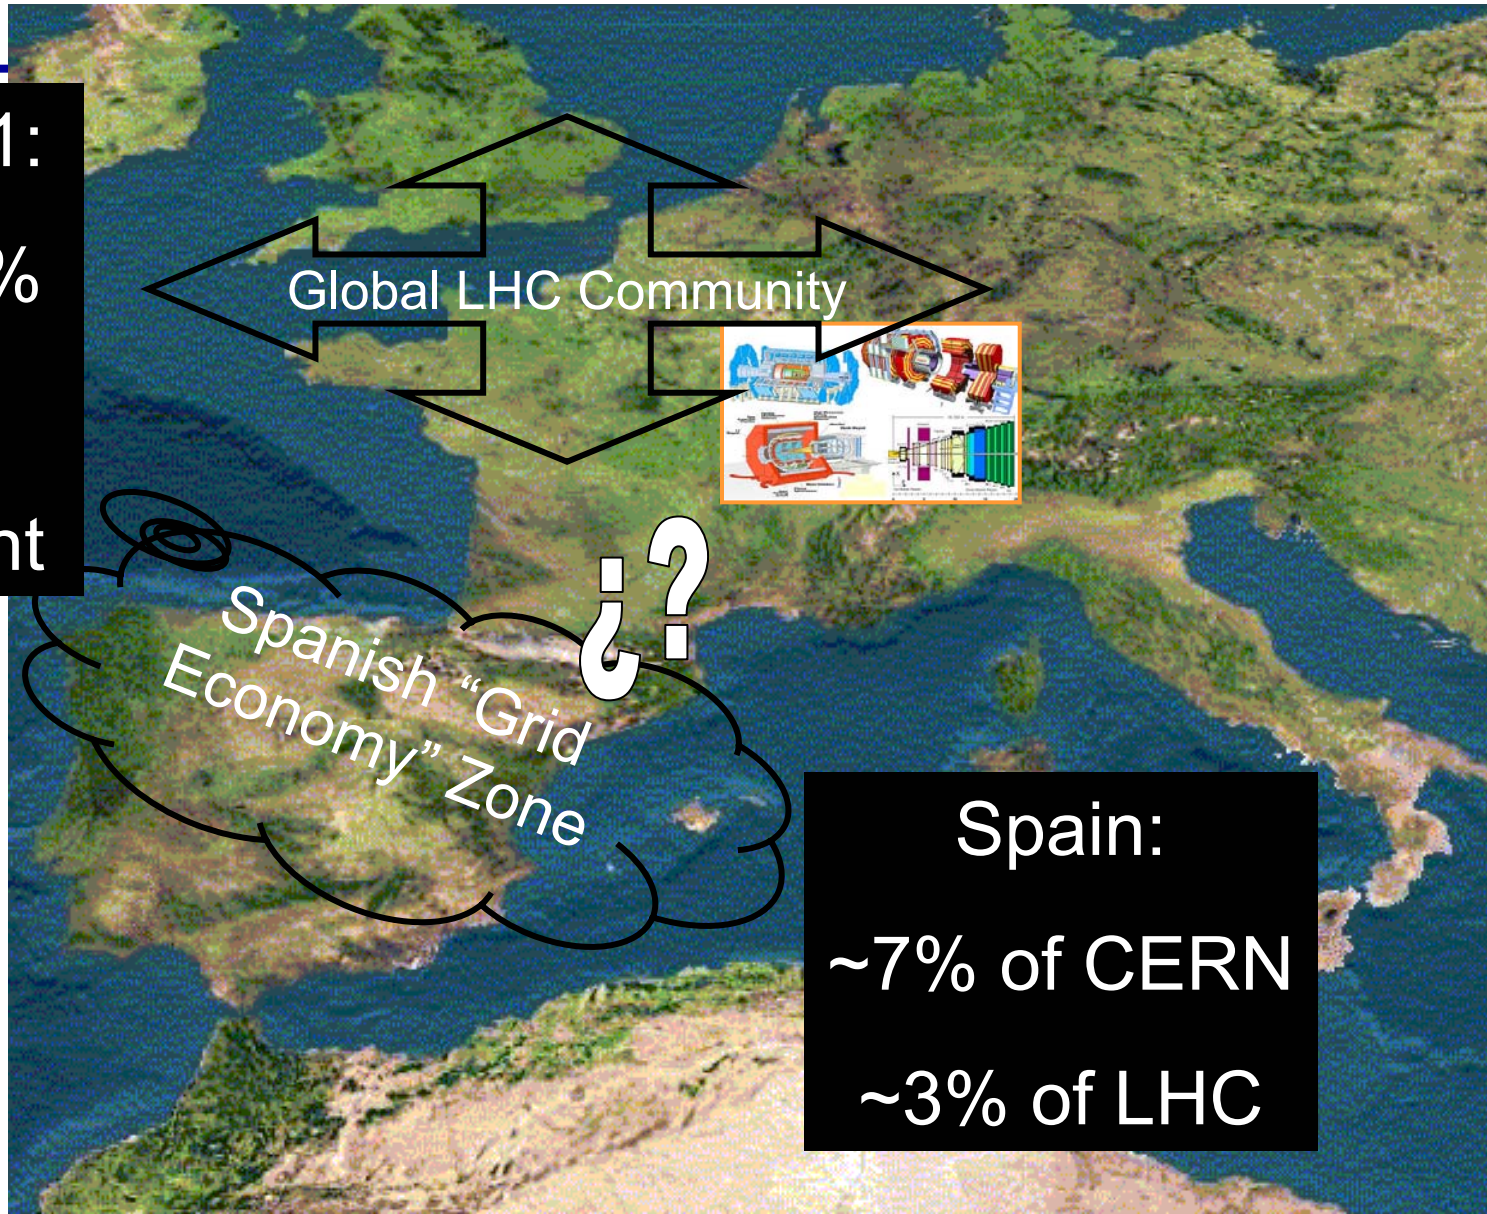

Autònoma de Barcelona

Port
d'Informació
Científica

LCG Requirements Ctte (SC2) revisited Tier 1/2/3 issue. New document focuses more on services than on capacity

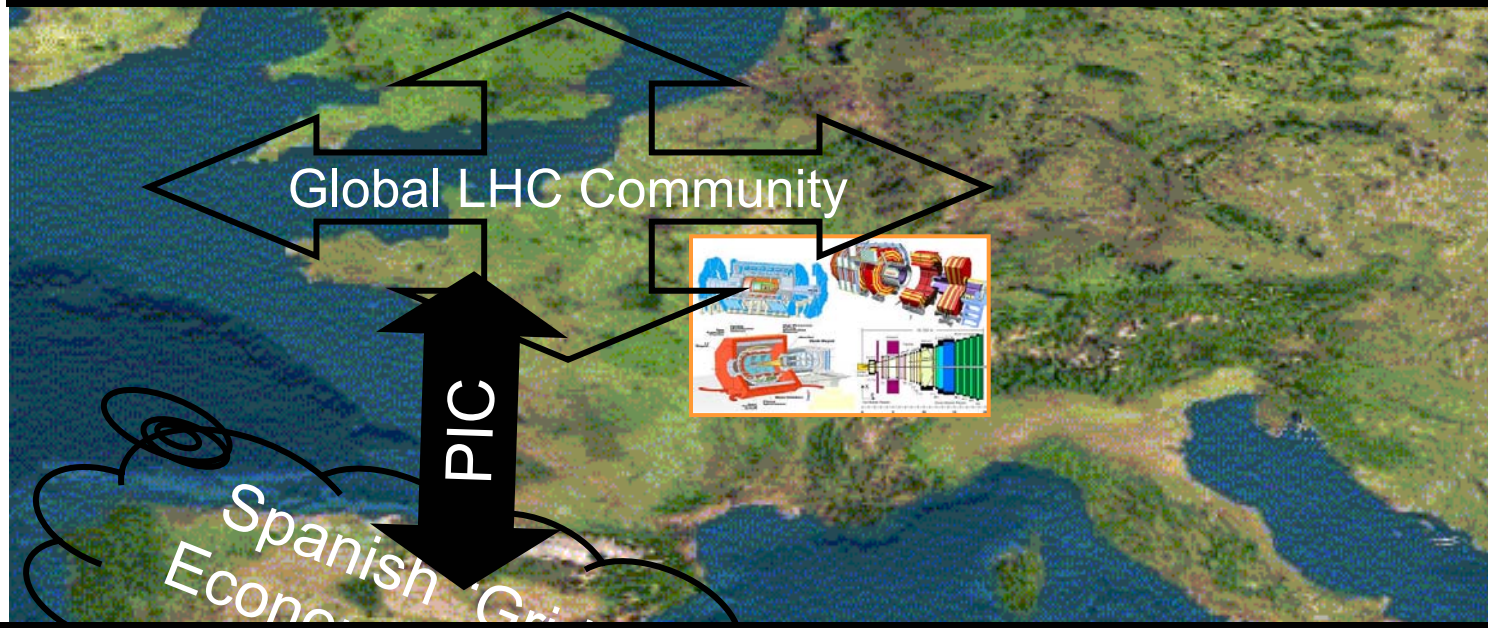

Prototype a "Grid Services Exchange" focused on LHC data

PIC = Port d'Informació Científica

- "Legacy-free" center for data-intensive scientific computing
- Exploit commodity hardware and automation to lower costs
- Store and offer produced datasets to global LCG securely
- Cache datasets from global LCG for secure use by ES physicists

Port d'Informació Científica

Objectives of the PIC

- Coordination centre for data-intensive scientific computing for users in Spain and open to other users in the EU and other nations.
- Coordination centre for LHC computing of Spanish HEP, maintaining contacts particularly with CERN.
- Reference centre in data-intensive scientific computation techniques, including transfer of these techniques to scientific disciplines which may benefit from them.
- Establishment of collaborations with other institutes involved in scientific computation, in particular supercomputing and massively parallel computing.
- Development of techniques and general services in the context of the future Global Information Grid.
- Development of a Centre of Excellence to enable participation in European projects for the development of a future International Grid Infrastructure.

Port d'Informació Científica

PIC Base Funding Plan
(Note: CIEMAT agreement not signed yet)

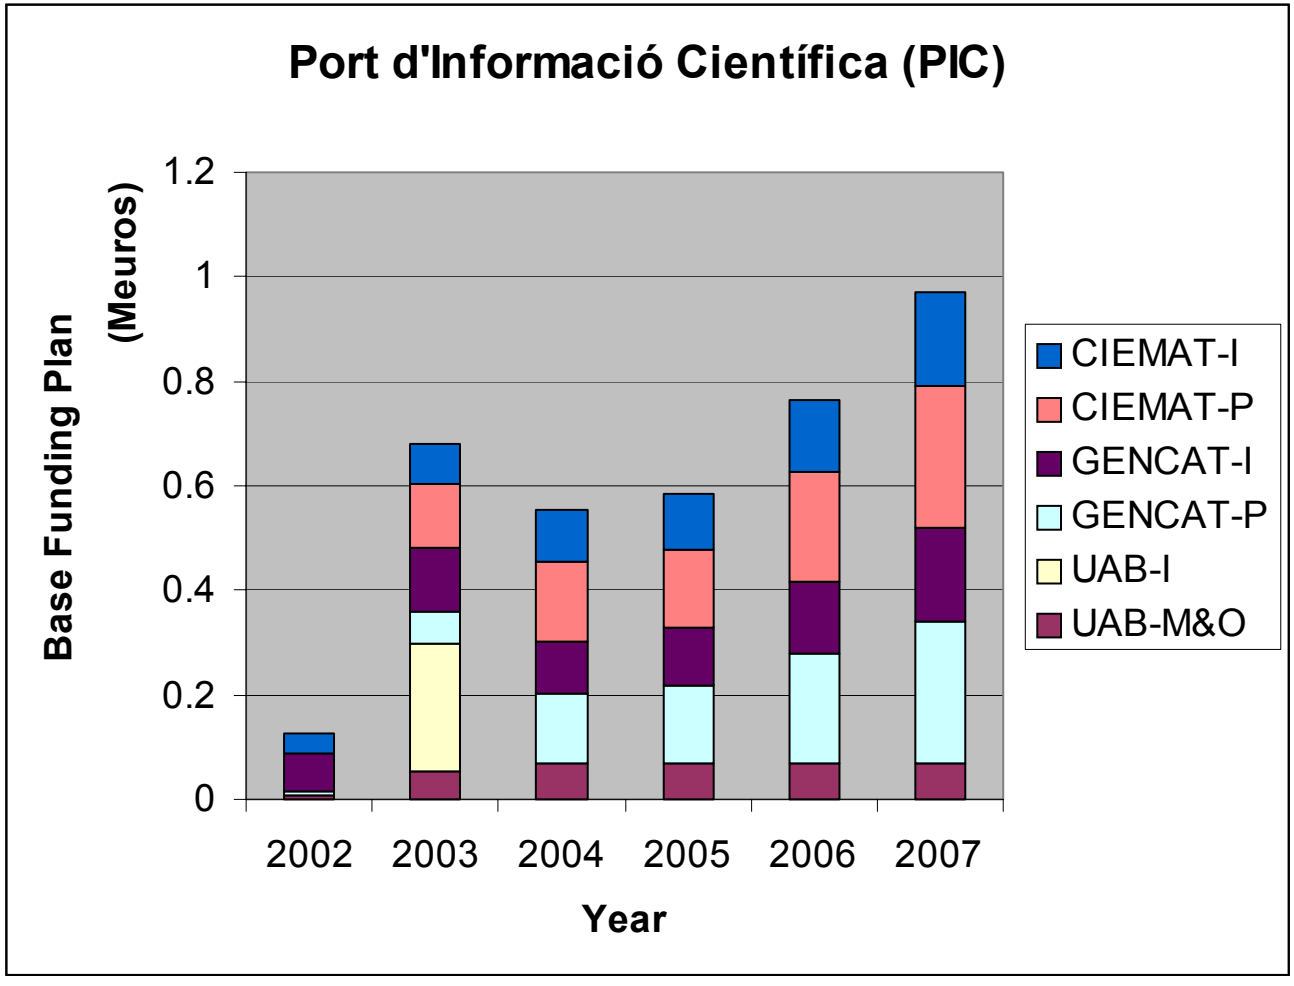

Port d'Informació Científica

PIC infrastructure at UAB

- UAB is turning "Edifici D" into a shared infrastructure to house computationally oriented scientific research activities.
  The largest "user" of the building will be the PIC.
- PIC will have ~150 m^2 office space there (mostly "open-plant" type, we already have 90 m^2). Some space reserved for visitors.
- The PIC tape robots and data servers are housed in a partition of the computer room of Edifici D.
  - 150 m^2 of high quality machine room with false floor
  - 200 KVA electricity with battery backup capilarized to about 70 16-Amp circuits each with smoke detector
  - 300 KW of triply redundant air conditioning
  - 500 KVA diesel generator (Rolls-Royce!!)
- The Catalan and Spanish A&R Network are being boosted to the Gbps range. GÈANT is already multi-Gbps.

Port d'Informació Científica

A few pictures of where we are with PIC

Overall view of machine room

Space for:

2 STK cylindrical robots

35 19" racks

Future:

Double tape space by using additional room

Port d'Informació Científica

A few pictures of where we are with PIC

Latest Generation L5500 Storagetek Tape Robot (6000 slots, >1.4 PBytes capacity)

2 Latest Generation 9940B drives (200 GB cartridge, 30 MB/sec)

Port
'Informació
Científica

One Grid infrastructure makes many Grids possible

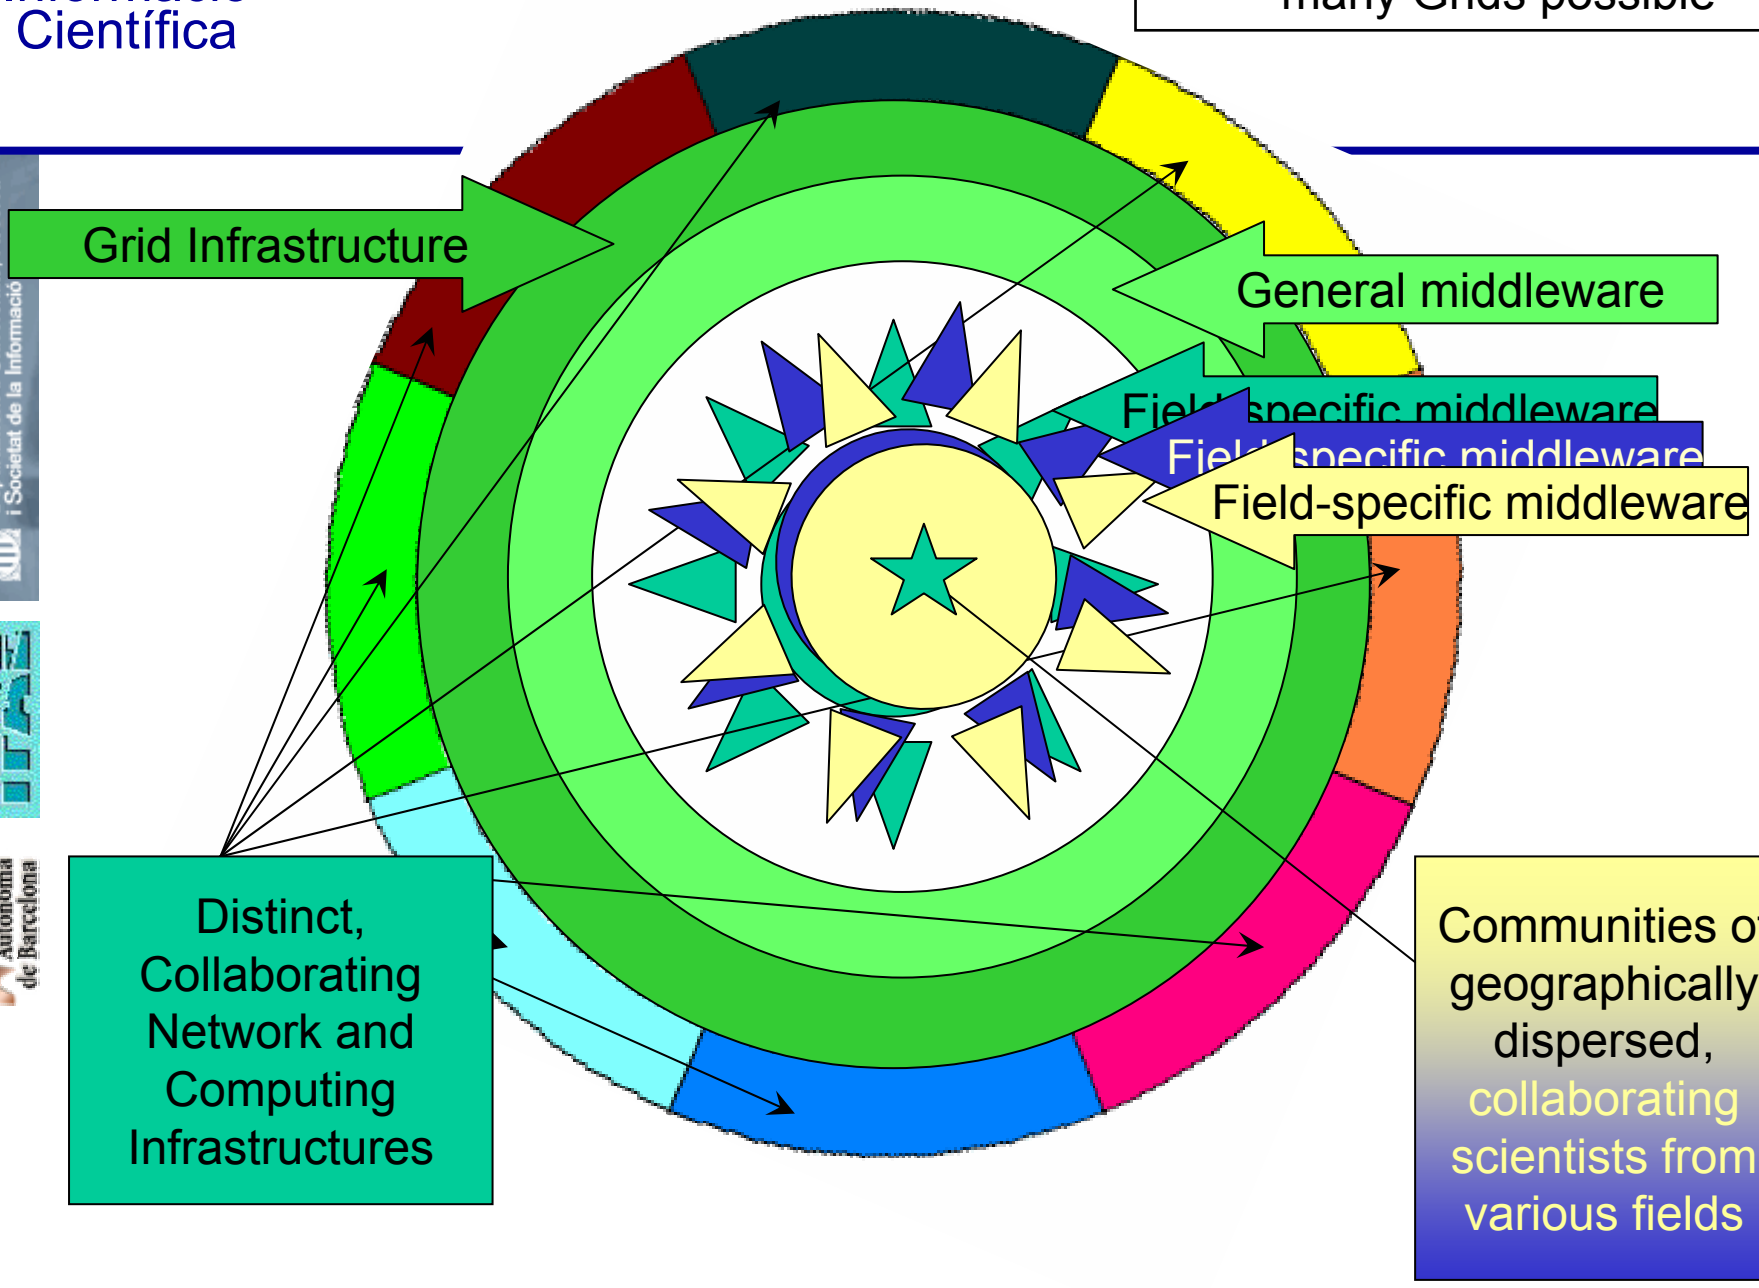

Grids and Grid Infrastructure:
Provider-centric view

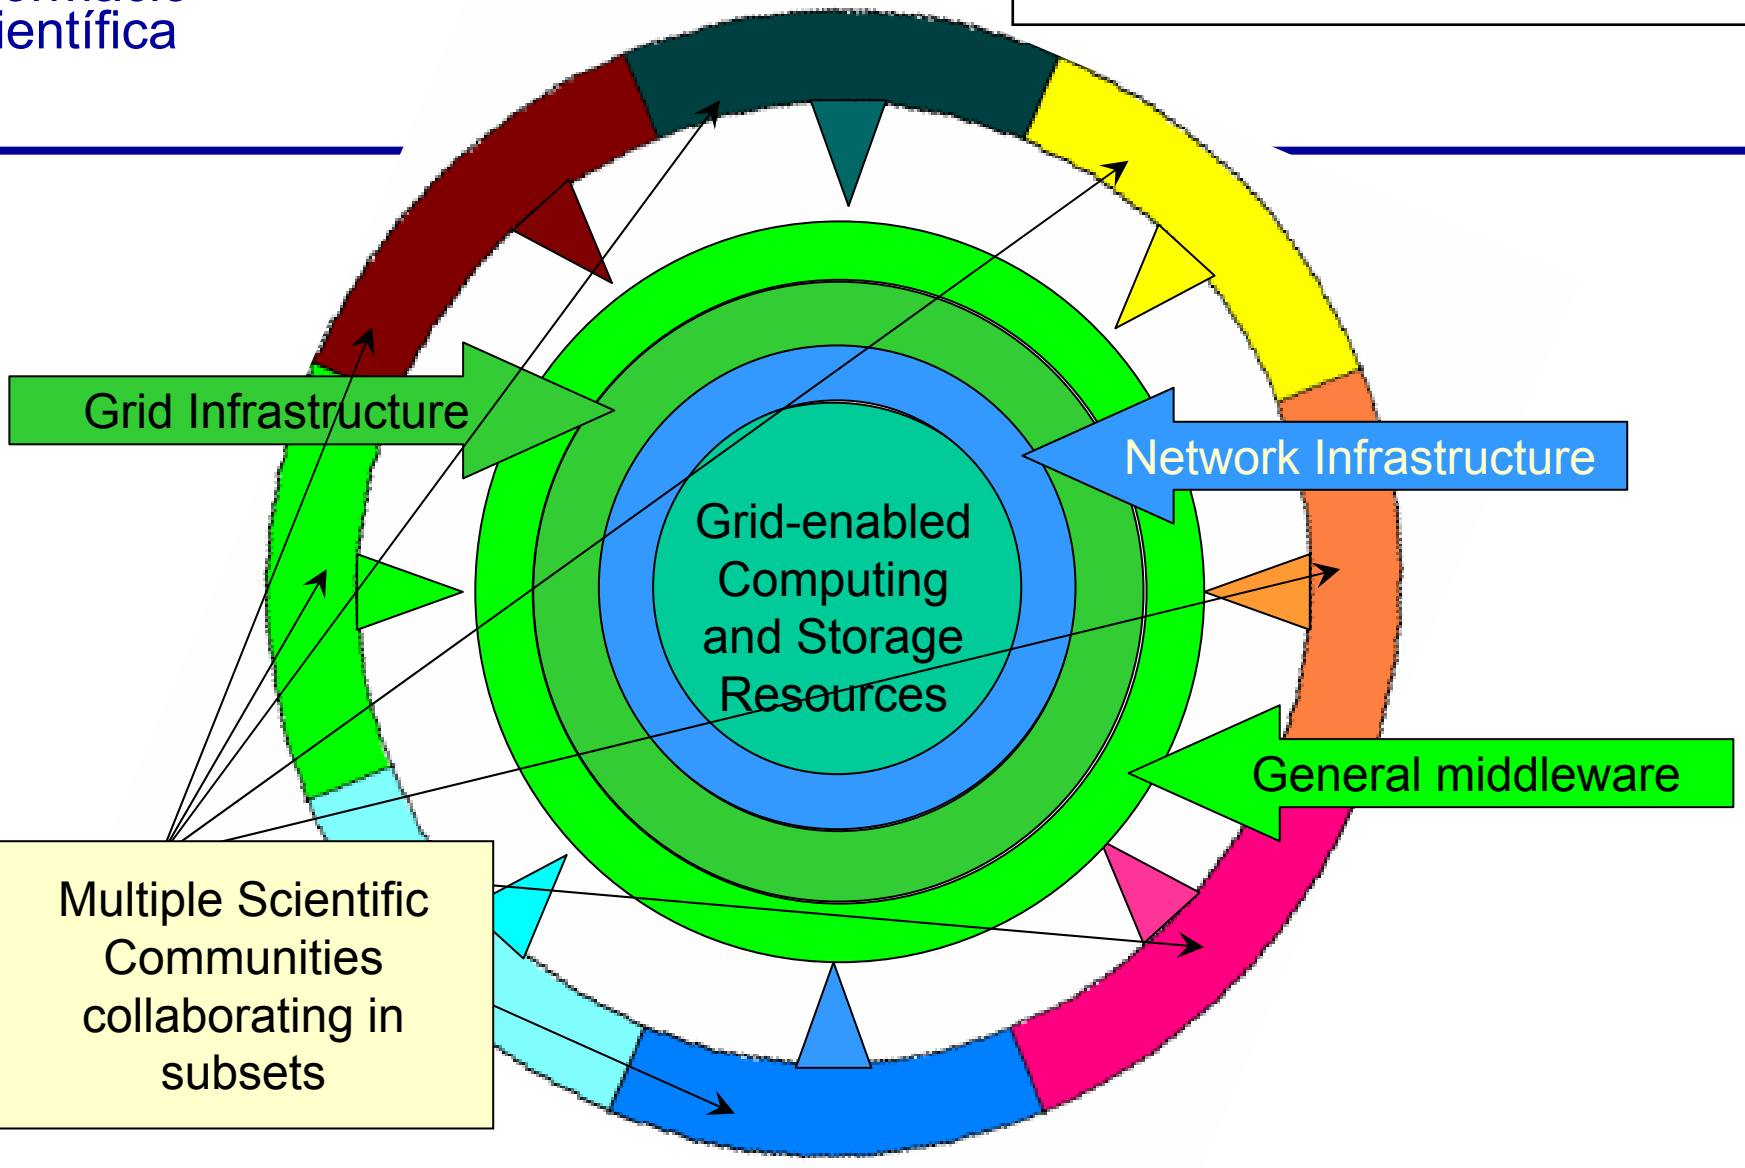

Port
d'Informació
Científica

Questions of Terminology

"e-Science"	Enhanced Science possible by continuous collaboration o mobile scientists supported by digital resources
"Grids"	The set of all instances of "a Grid"
"a Grid"	The set of digital resources that a Community is using at any given time to support their e-activity
"Grid Infrastructure"	The set of persistent Services and Processes which perm construction and reconfiguration of Grids
"Research Network Infrastructure and/or Internet"	The 'net that we know and which will grow influenced by and influencing Grid Infrastructure and Grids

An illustration: GÉANT and GRIDs (as seen by EU IST programme)

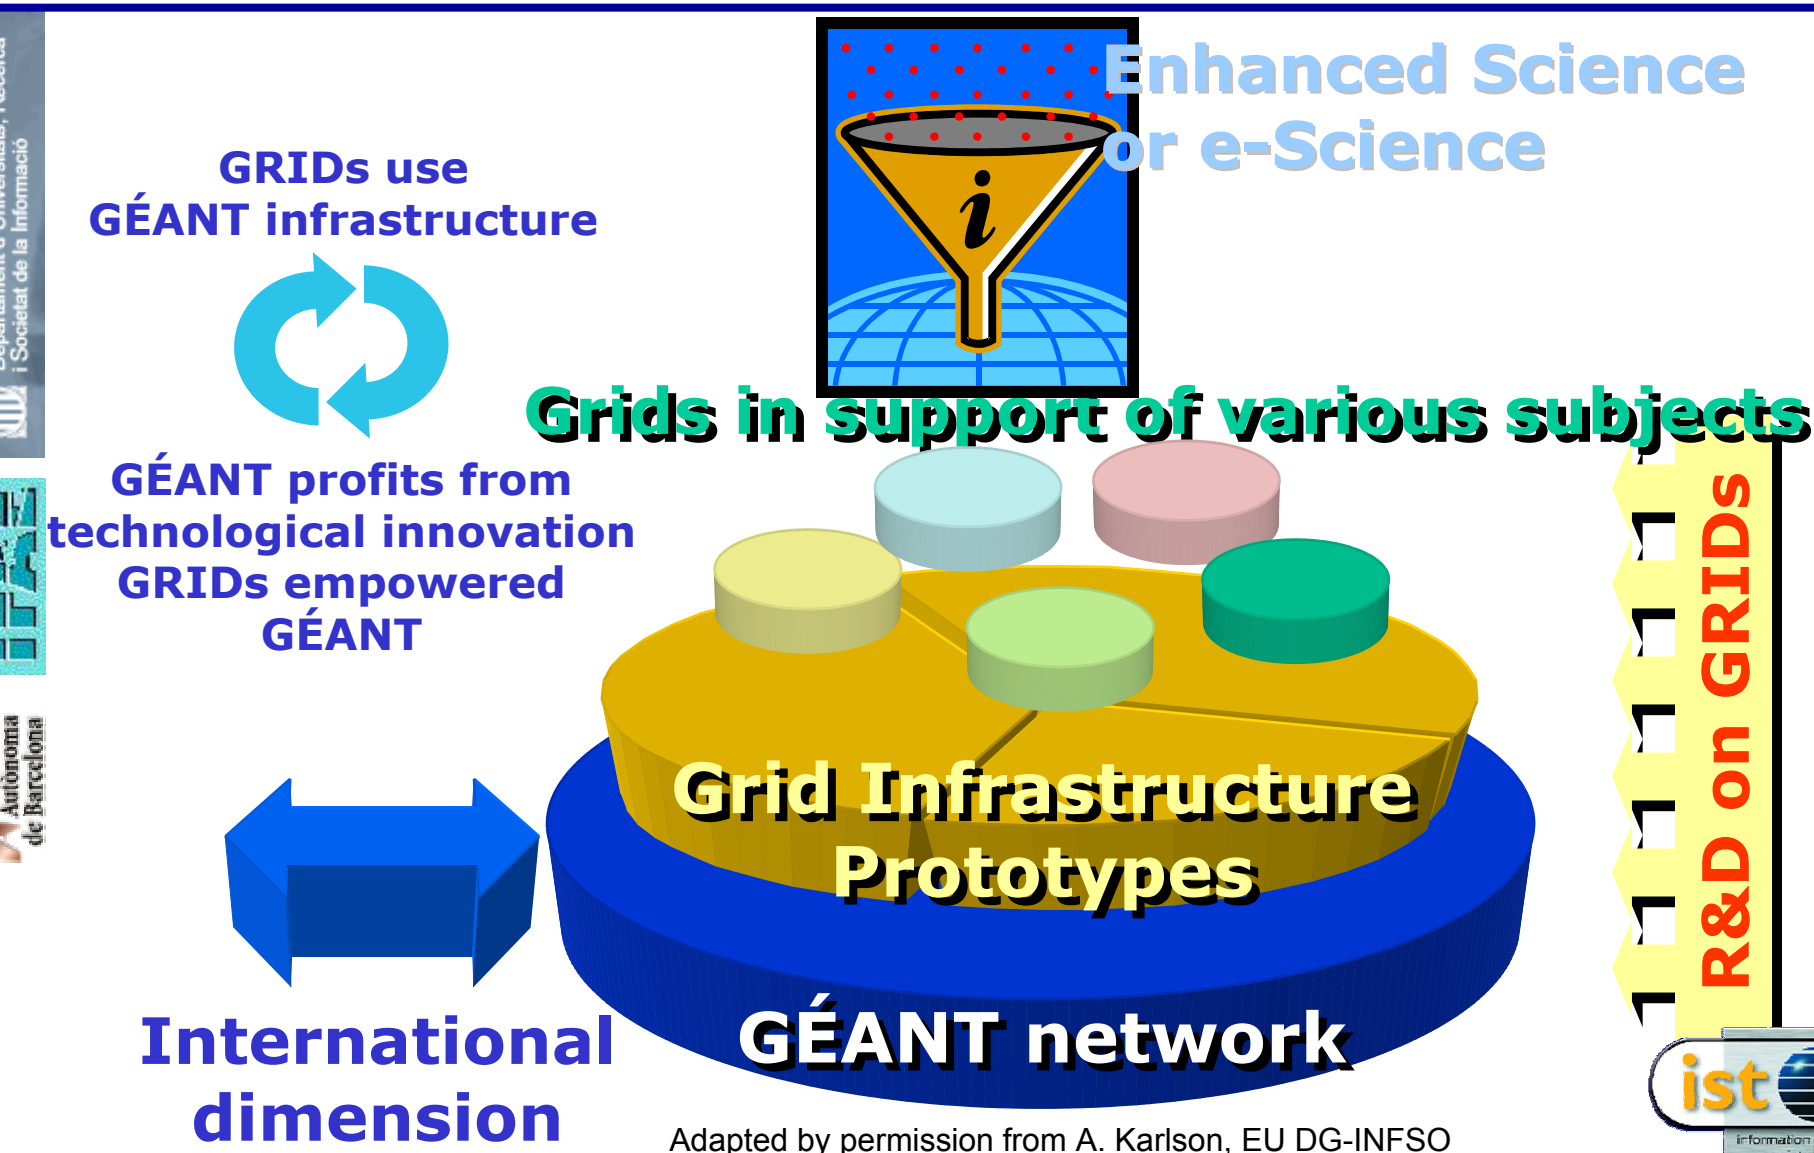

Adapted by permission from A. Karlson, EU DG-INFSO

Port d'Informació Científica

Lead-times and momentum in R+D
Internet – Web – Web Services – Grid

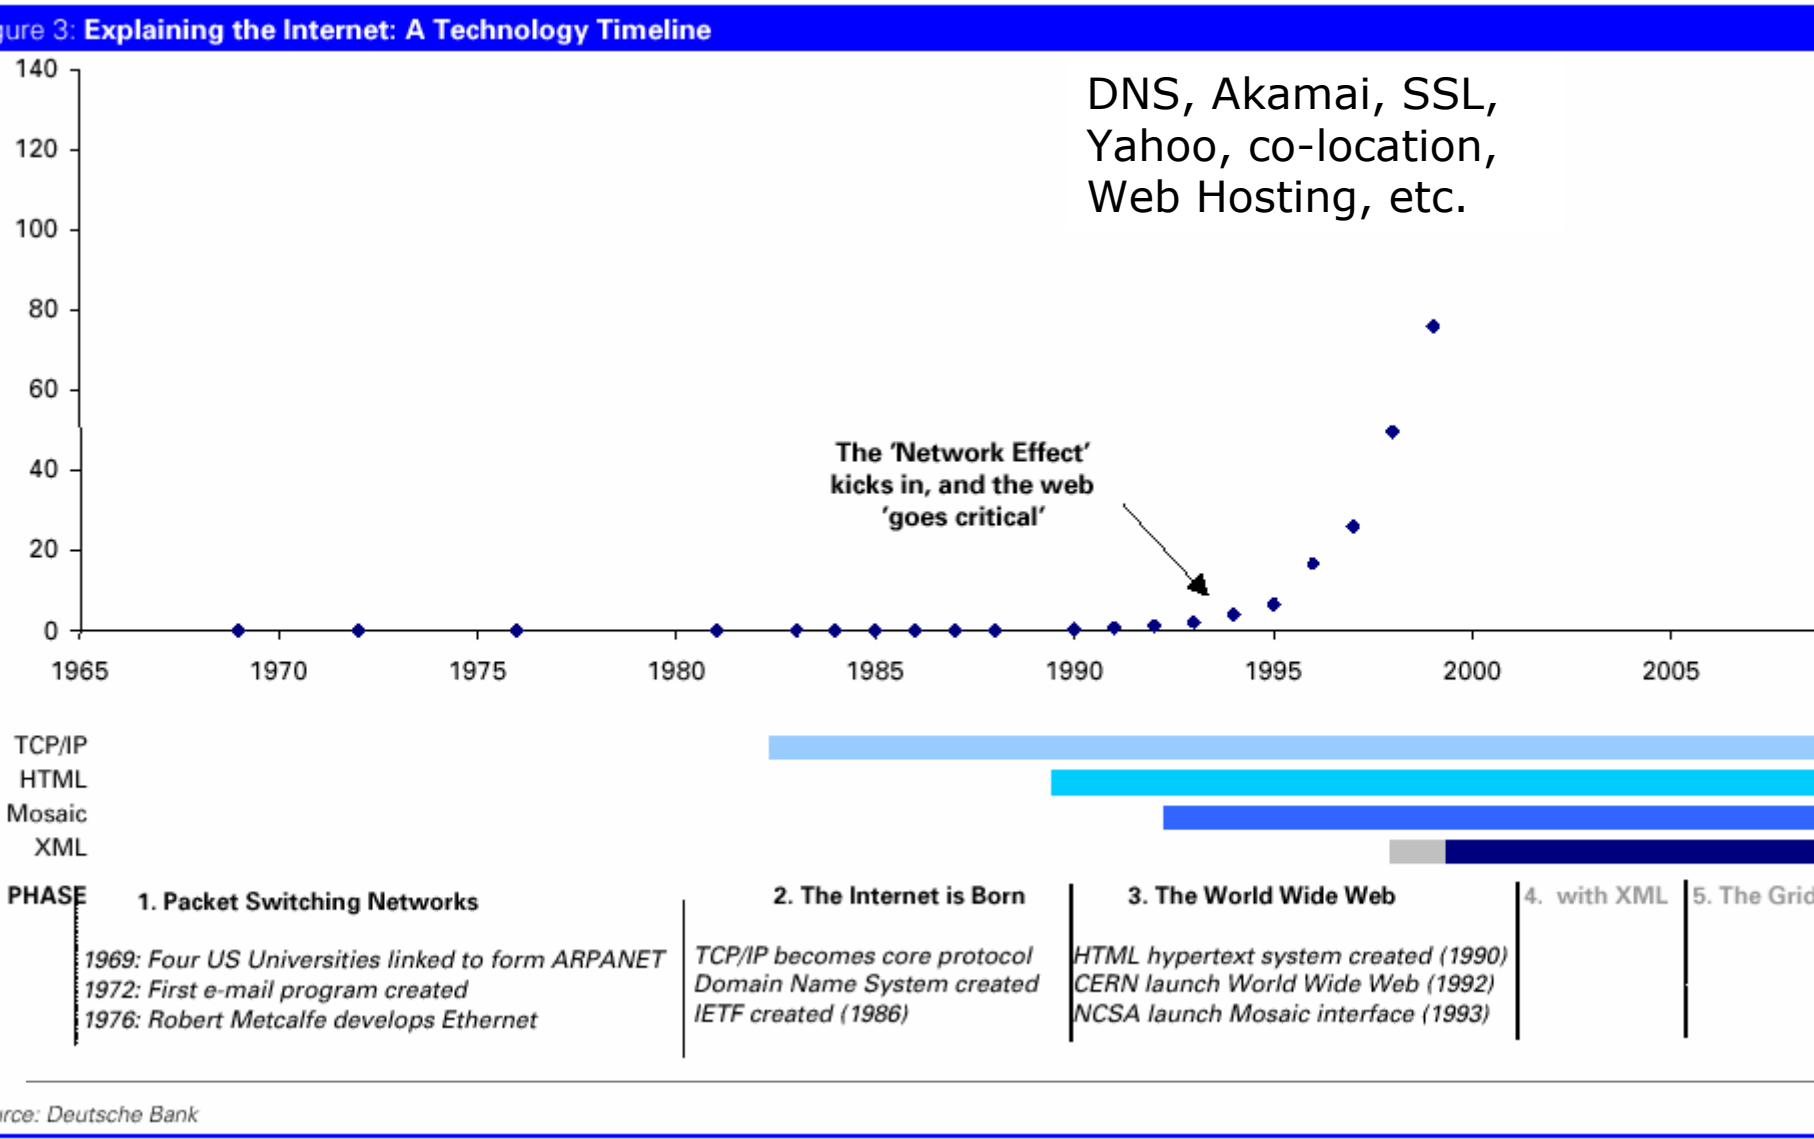

Evolution of data-processing environments

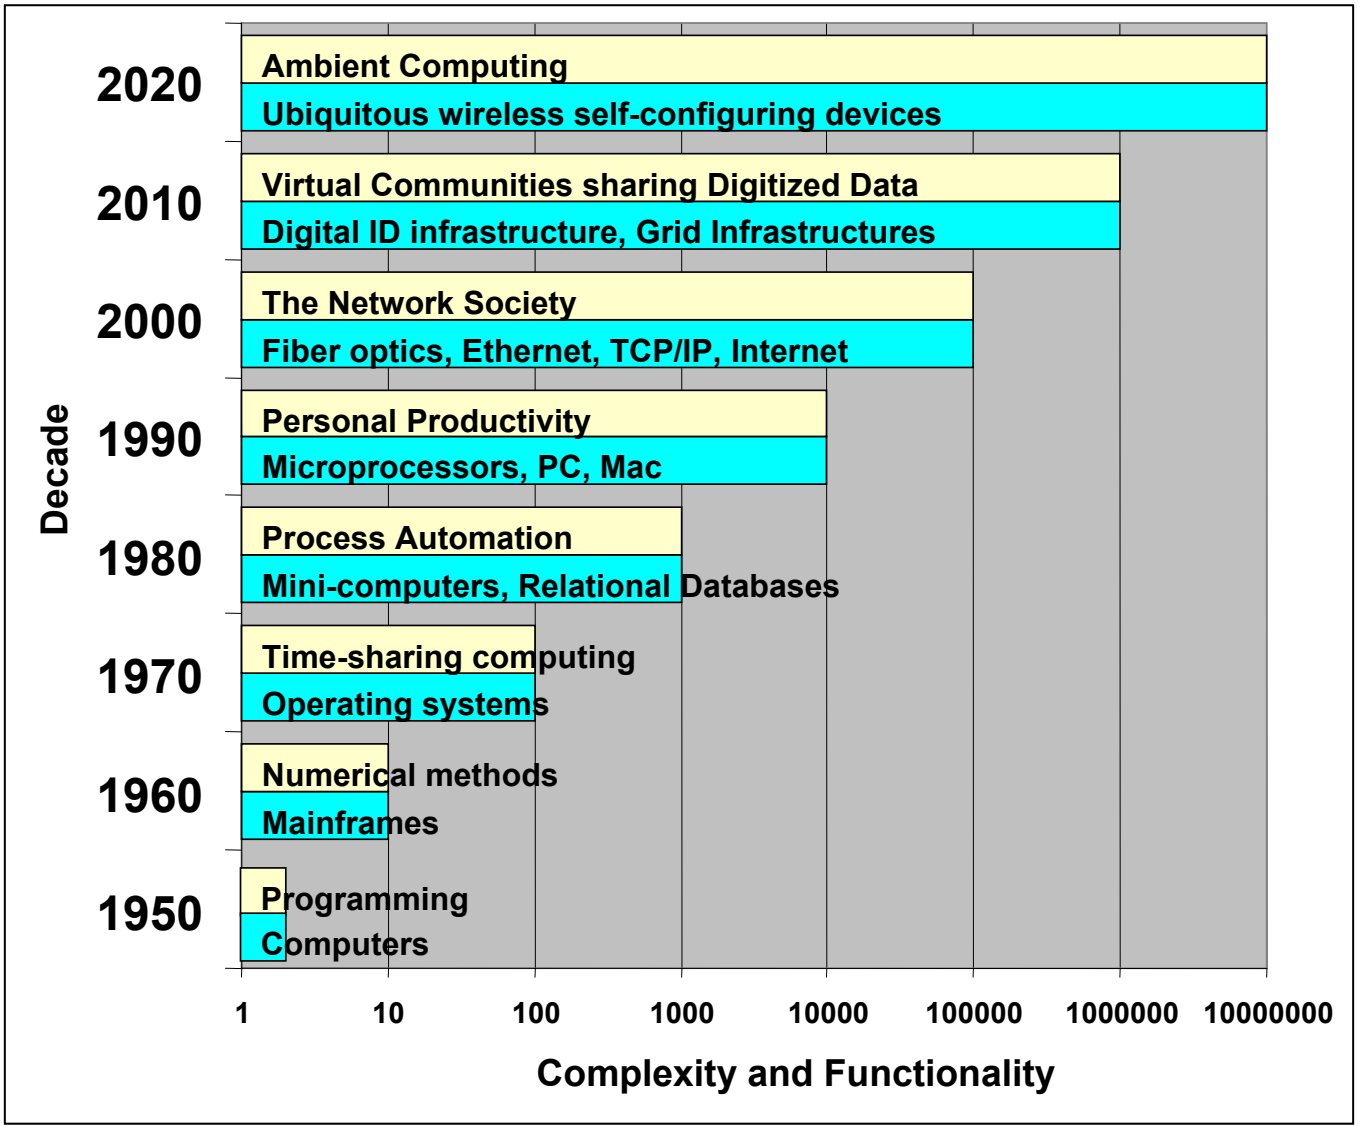

Evolution of data-processing environments

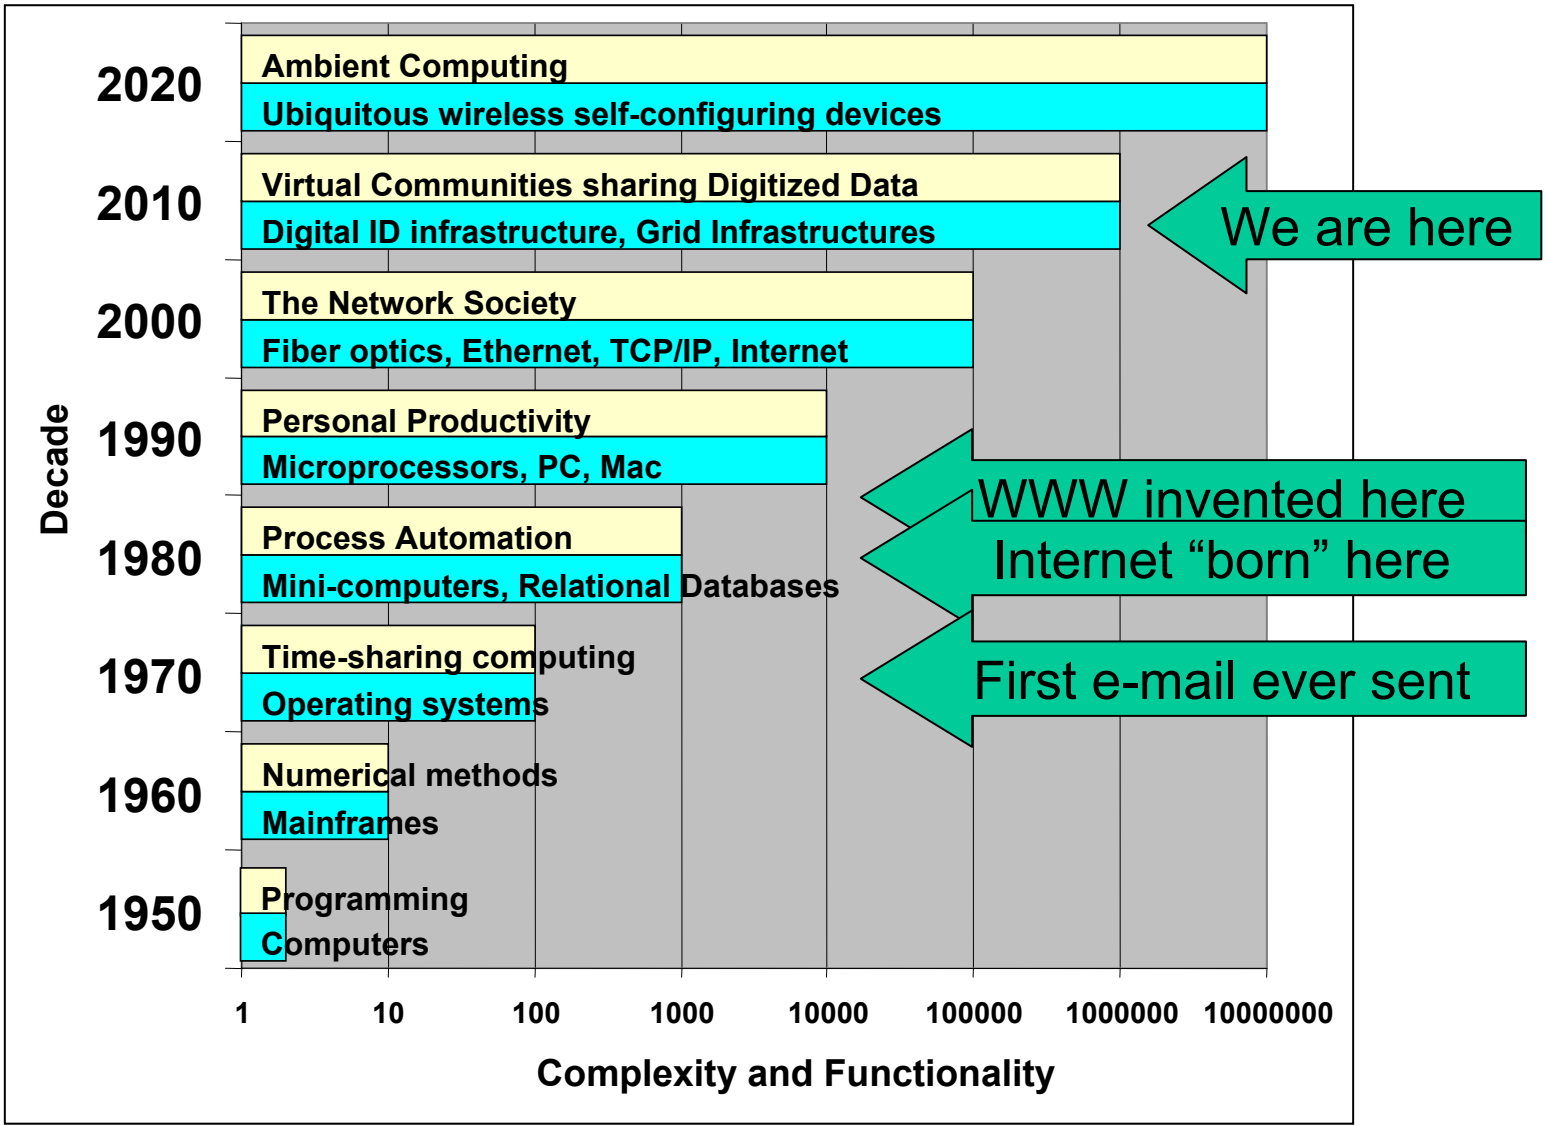

Port
'Informació
Científica

Deliverables to fulfill the objectives

EAD, MCF, GSW, CDC

IFCA

EAD, GSW, CDC

USC

EAD, ETD, PIC, EDS, GSW

IFAE

EAD, MCF, EDS, RSG, SEG, CTS

CIEMAT

UB

EAD,EDS, CTS

UAM

EAD, GVM

IFIC

EAD, MCF, CTS, CDC

- Stay away from glitz.
- Concentrate on deployment, MC & analysis
- Use local Univ. for TT to other disciplines
- 600 KCHF materials contribution to LCG-CERN

AD = Analysis Farm
DS = SW Dev Platform
SG = SW repository
SW = SW Gridification
CF = MC Fabric
VM = Virtual MC farm
TD = Data Transform
IC = Gridified Data Store
EG = Security Architect
TS = Tech MC Support
DC = Data Chal. Coord.