Grids in Europe and the LCG Project

Ian Bird
LCG Deployment Manager
Information Technology Division, CERN
Geneva, Switzerland

Lepton-Photon Symposium 2003
Fermilab
14 August 2003

Outline

- Introduction
  - Why are grids relevant to HENP?
- European grid R&D program
  - Existing projects
  - New project - EGEE
- LCG project
  - Deploying the LHC computing environment
  - Using grid technology to address LHC computing
- Outlook
  - Interoperability and standardisation
  - Federating grids – what does it mean?

Introduction

Why is particle physics involved with grid development?

The Large Hadron Collider Project
4 detectors

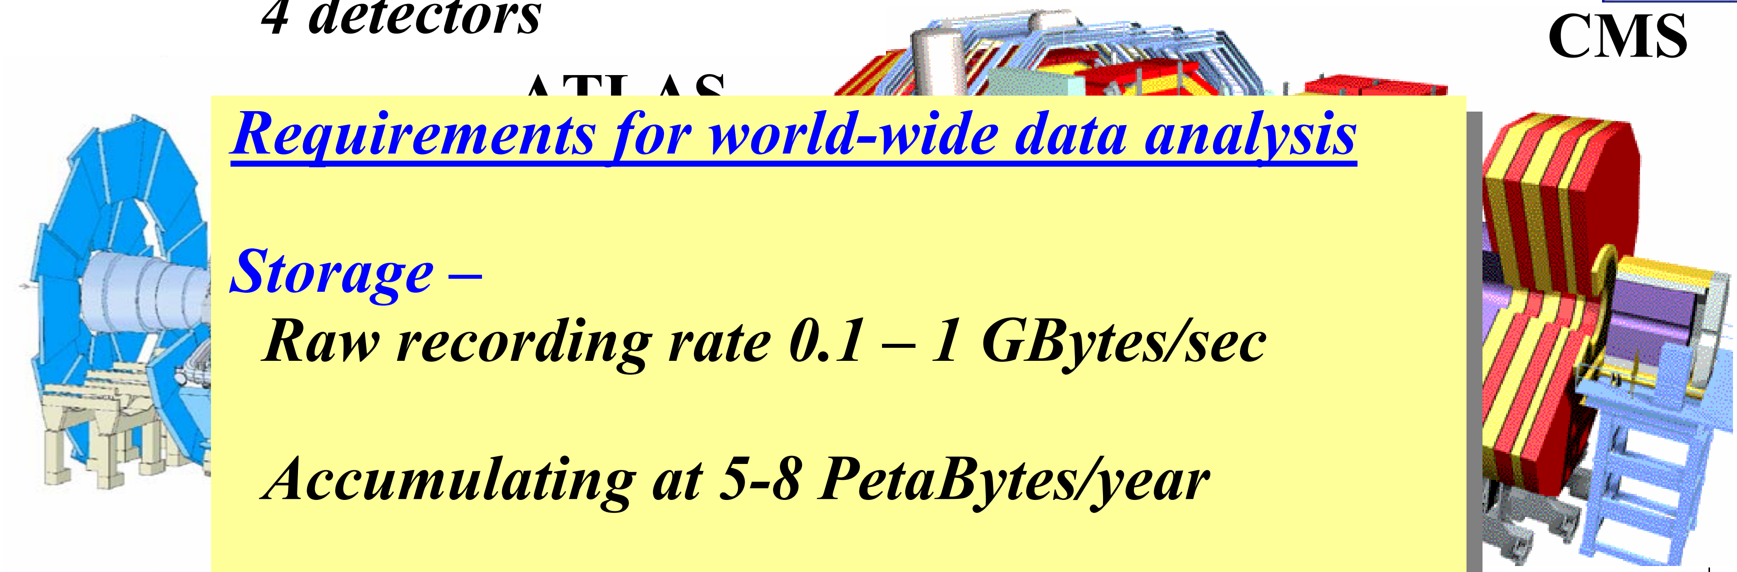

[*Requirements for world-wide data analysis*]

Storage –

Raw recording rate 0.1 – 1 GBytes/sec

Accumulating at 5-8 PetaBytes/year

10 PetaBytes of disk

Processing –

100,000 of today's fastest PCs

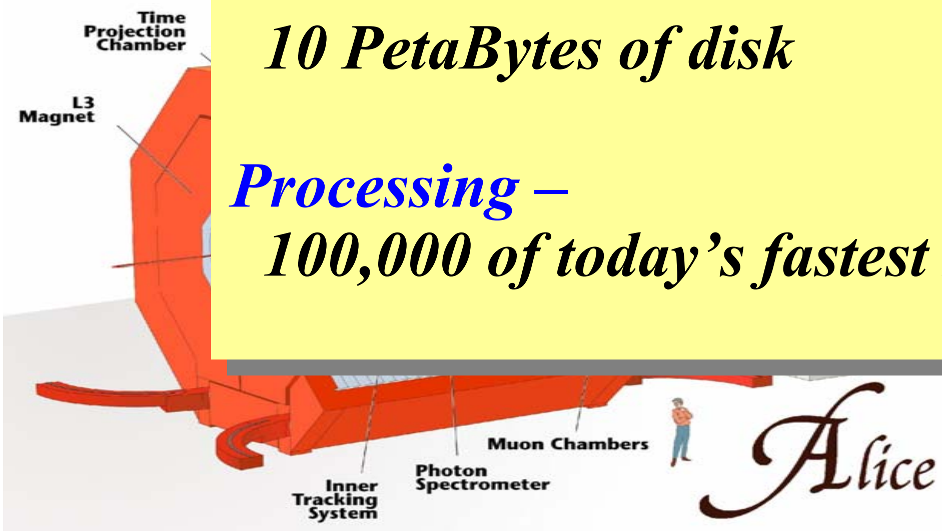

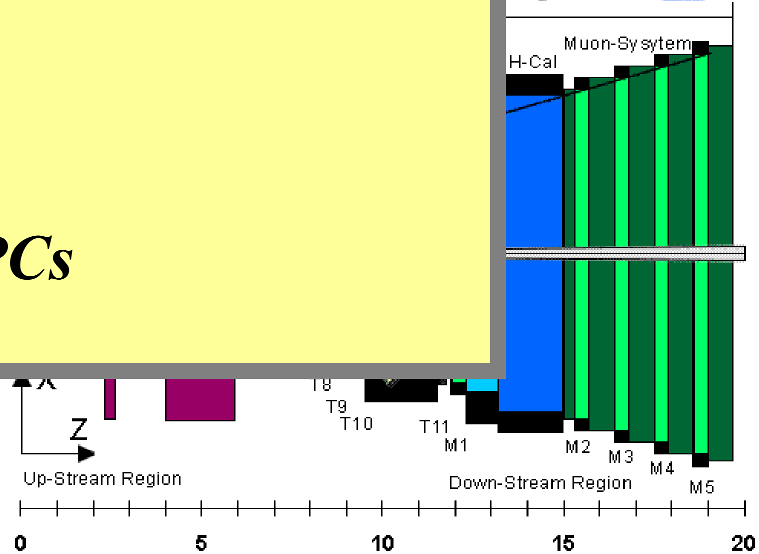

p-p collisions at LHC

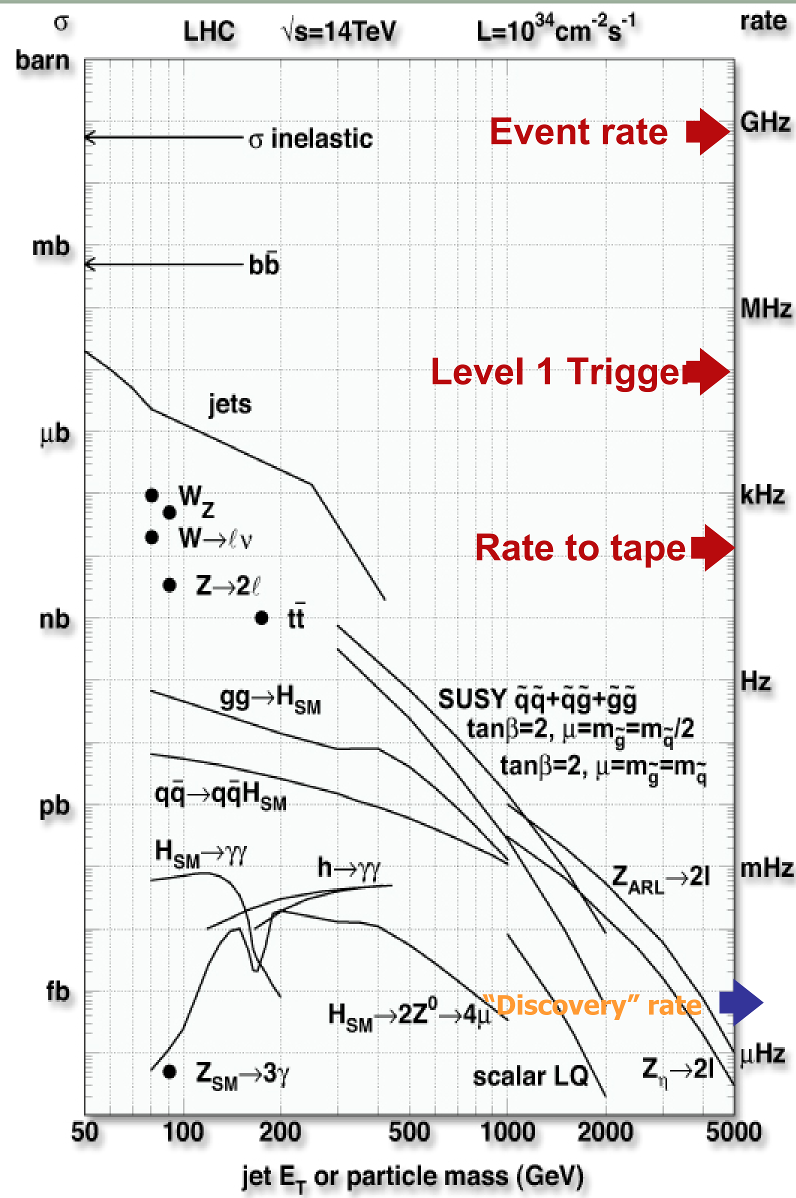

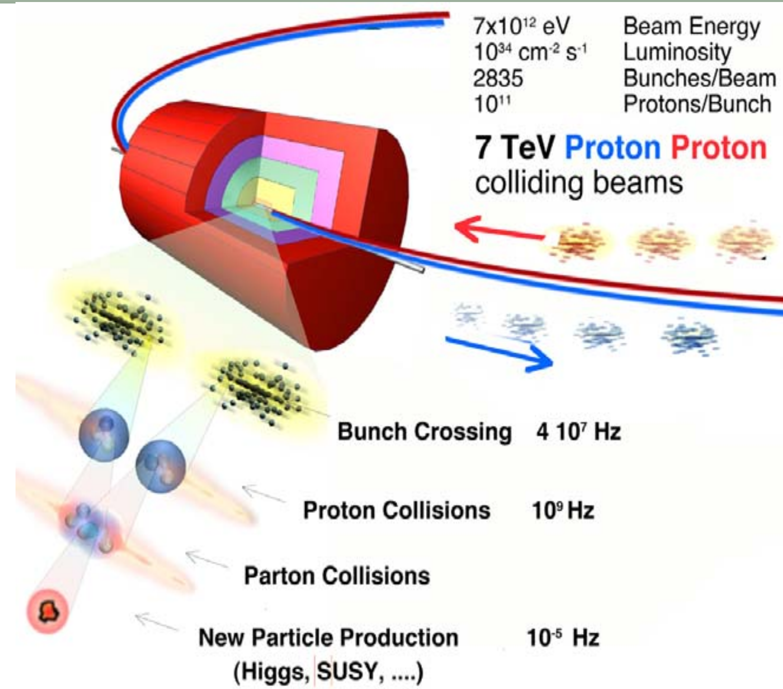

Crossing rate 40 MHz
Event Rates: ~10^9 Hz

Max LV1 Trigger 100 kHz
Event size ~1 Mbyte
Readout network 1 Terabit/s
Filter Farm ~10^7 Si2K
Trigger levels 2
Online rejection 99.9997% (100 Hz from 50 MHz)
System dead time ~ %
Event Selection: ~$1/10^{13}$

Luminosity
Low 2×10^{33} cm^{-2} s^{-1}
High 10^{34} cm^{-2} s^{-1}

LHC Computing Hierarchy

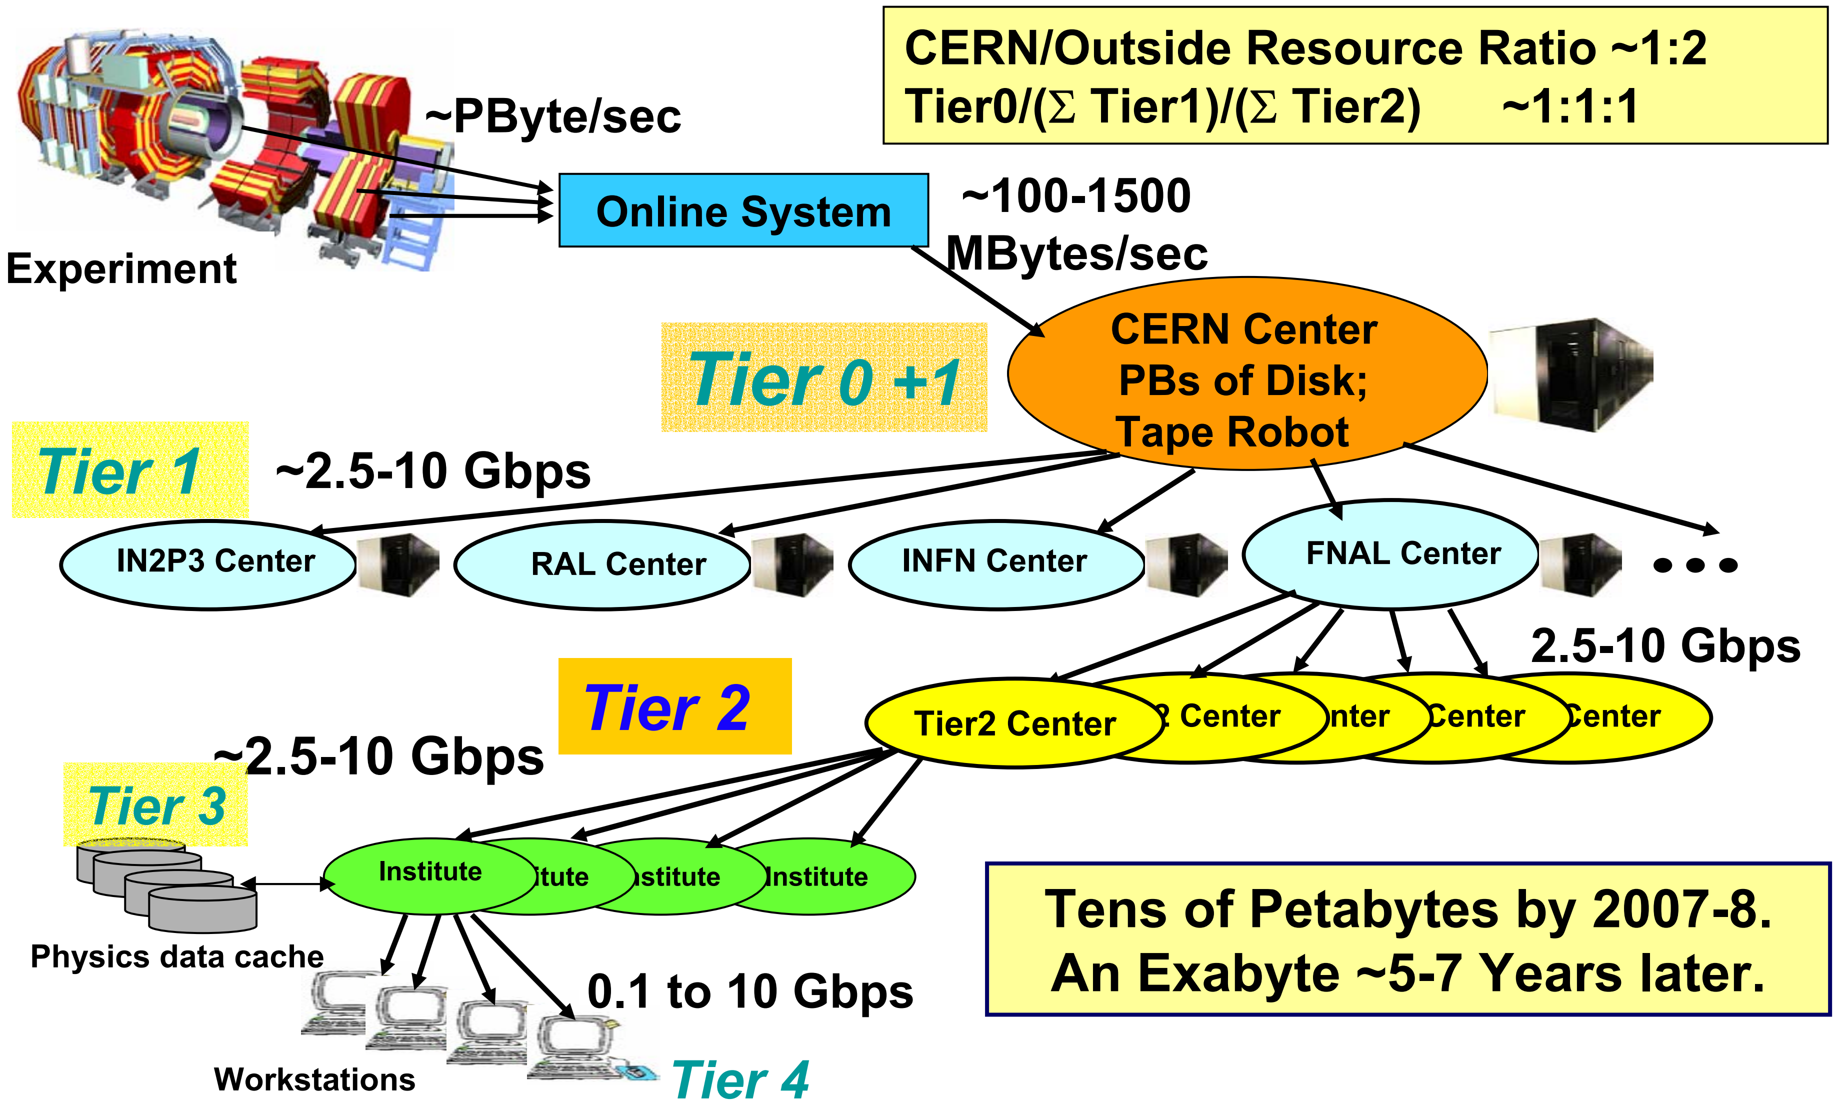

Emerging Vision: A Richly Structured, Global Dynamic System

Summary – HEP/LHC Computing Characteristics

- ☺ independent *events* (collisions)
  - easy parallel processing
- ☺ bulk of the data is read-only
  - versions rather than updates
- ☺ meta-data (few %) in databases
- ☺ good fit to simple PCs
  - modest floating point
  - modest per-processor I/O rates

- ☹ very large aggregate requirements – computation, data, i/o
  - more than we can afford to install at the accelerator centre
- ☹ chaotic workload –
  - batch & interactive
  - research environment - physics extracted by iterative analysis, collaborating groups of physicists

→ unpredictable
→ unlimited demand

Grids as a solution

- LHC computing is of unprecedented scale
  - Requirements are larger than could feasibly install in one place
- Computing must be distributed for many reasons
  - Political, economic, staffing
  - Enable access to resources for all collaborators

- Many projects over the last few years have addressed aspects of the LHC computing problem
  - In the US and Europe
- In 2002 LCG was proposed to set up the LHC computing environment (assumed to be based on grid technology)
  - Using the results of EU and US projects to deploy and operate a real production-level service for the experiments
  - As a validation of the LHC computing models

tem

- Note:
  - Other HENP experiments currently running (Babar, CDF/DO, STAR/PHENIX), with significant data and computing requirements
    - Have already started to deploy solutions based on grid technology
  - We can learn from the running experiments

European Grid projects

European grid projects

CrossGrid

- Many grid research efforts, either
  - Nationally funded – including regional collaborations, or
  - EU funded
- Most with particle physics as a major (but not the only) application
- Address different aspects of grids:
  - Middleware
  - Networking, cross-Atlantic interoperation
- Some are running services at some level

- In this talk I will address some of the major EU funded projects:
  - Existing projects: DataGrid and DataTAG
  - New project: EGEE

European DataGrid (EDG)

http://www.eu-datagrid.org

The EU DataGrid Project

- 9.8 M Euros EU funding over 3 years
- 90% for middleware and applications (Physics, Earth Observation, Biomedical)
- 3 year phased developments & demos
- Total of 21 partners
  - Research and Academic institutes as well as industrial companies
- Extensions (time and funds) on the basis of first successful results:
  - DataTAG (2002-2003) www.datatag.org
  - CrossG… www.c…
  - GridSta… www.g…

- Project started on Jan. 2001
- Testbed 0 (early 2001)
  - International test bed 0 infrastructure deployed
    - Globus 1 only - no EDG middleware
- Testbed 1 (early 2002)
  - First release of EU DataGrid software to defined users within the project
- Testbed 2 (end 2002)
  - Builds on Testbed 1 to extend facilities of DataGrid
  - Focus on stability
- Passed 2nd annual EU review Feb. 2003
  - …alability

Built on Globus and Condor for the underlying framework, and, since 2003 provided via the Virtual Data Toolkit (VDT)

DataGrid in Numbers

People

>350 registered users

12 Virtual Organisations

19 Certificate Authorities

>300 people trained

278 man-years of effort

100 years funded





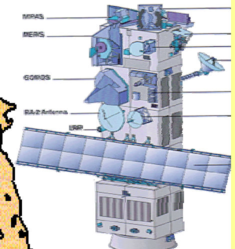

Testbeds

>15 regular sites

>40 sites using EDG sw

>10'000s jobs submitted

>1000 CPUs

>15 TeraBytes disk

3 Mass Storage Systems

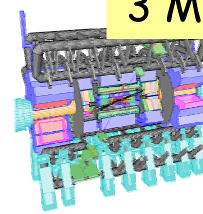

Software

50 use cases

18 software releases

Current release 1.4

>300K lines of code

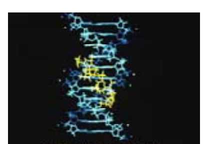
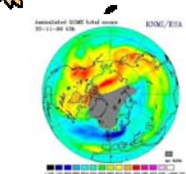


Scientific applications

5 Earth Obs institutes

9 bio-informatics apps

6 HEP experiments

GEANT

DataGrid Status
Applications & Testbeds

- Intense usage of application testbed (**release 1.3 and 1.4**) in 2002 and early 2003
  - **WP8**: 5 HEP experiments have used the testbed
    - ATLAS and CMS task forces very active and successful
      - Several hundred ATLAS simulation jobs of length 4-24 hours were executed & data was replicated using grid tools
      - CMS Generated ~250K events for physics with ~10,000 jobs in 3 week period
      - Since project review: ALICE and LHCb have been generating physics events
    - Results were obtained from focused ***task-forces***. Instability prevented the use of the testbed for standard production
  - **WP9**: EarthObs level-1 and 2 data processing and storage performed
  - **WP10**: Four biomedical groups able to deploy their applications
  - First Earth Obs site joined the testbed (Biomedical on-going)
- Steady increase in the size of the testbed until a peak of approx **1000 CPUs at 15 sites**
- **The EDG 1.4 software is frozen**
  - The testbed is supported and security patches deployed but **effort has been concentrated on producing EDG 2.0**
  - Application groups were warned that the **application testbed will be** upgrade on short notice sometime after June 15th.

DataTAG Project

DataTAG: Research and Technological Development for a Trans-Atlantic GRID

- EU ⇔ US Grid Interoperability
- EU ⇔ US Grid network research
  - High Performance Transport protocols
  - Inter-domain QoS
  - Advance bandwidth reservation
- Two years project started on 1/1/2002
  - extension until 1Q04 under consideration
- 3.9 MEUROs
  - 50% Circuit cost, hardware
  - Manpower

Interoperability: Objectives

- Address issues of middleware interoperability between the European and US Grid domains to enable a selected set of applications to run on the transatlantic Grid test bed
- Produce an assessment of interoperability solutions
- Provide test environment to applications
- Provide input to a common Grid LHC middleware projects

Interoperability issues

- Information System: demonstrate the ability to discover the existence and use grid services offered by the testbed; define minimal requirements on information services: glue information schema.
- Authentication / Authorisation : demonstrate the ability to perform cross-organizational authentication / test common user authorization Services based on VO.
- Data movement and access infrastructure: demonstrate the ability to move data from storage services operated by one site to another and to access them.
- LHC Experiments, distributed around the world, need to integrate their applications with interoperable GRID domains services.
- Demo test-bed demonstrating the validity of the solutions

DataTAG WP4 GLUE testbed

Grid Computing and Storage elements in

- INFN Bologna, Padova, Milan
- CERN
- FNAL
- Indiana University

Middleware

- INFN Bologna, Padova, Milan
  - EDG 1.4/GLUE
- CERN
  - LCG-0
- FNAL - Indiana University
  - VDT 1.1.X

Grid Services in Bologna/INFN:

- RB/Glue aware based on EDG1.4
- GIIS GLUE testbed top level
- VOMS
- Monitoring Server

Network Research Testbed

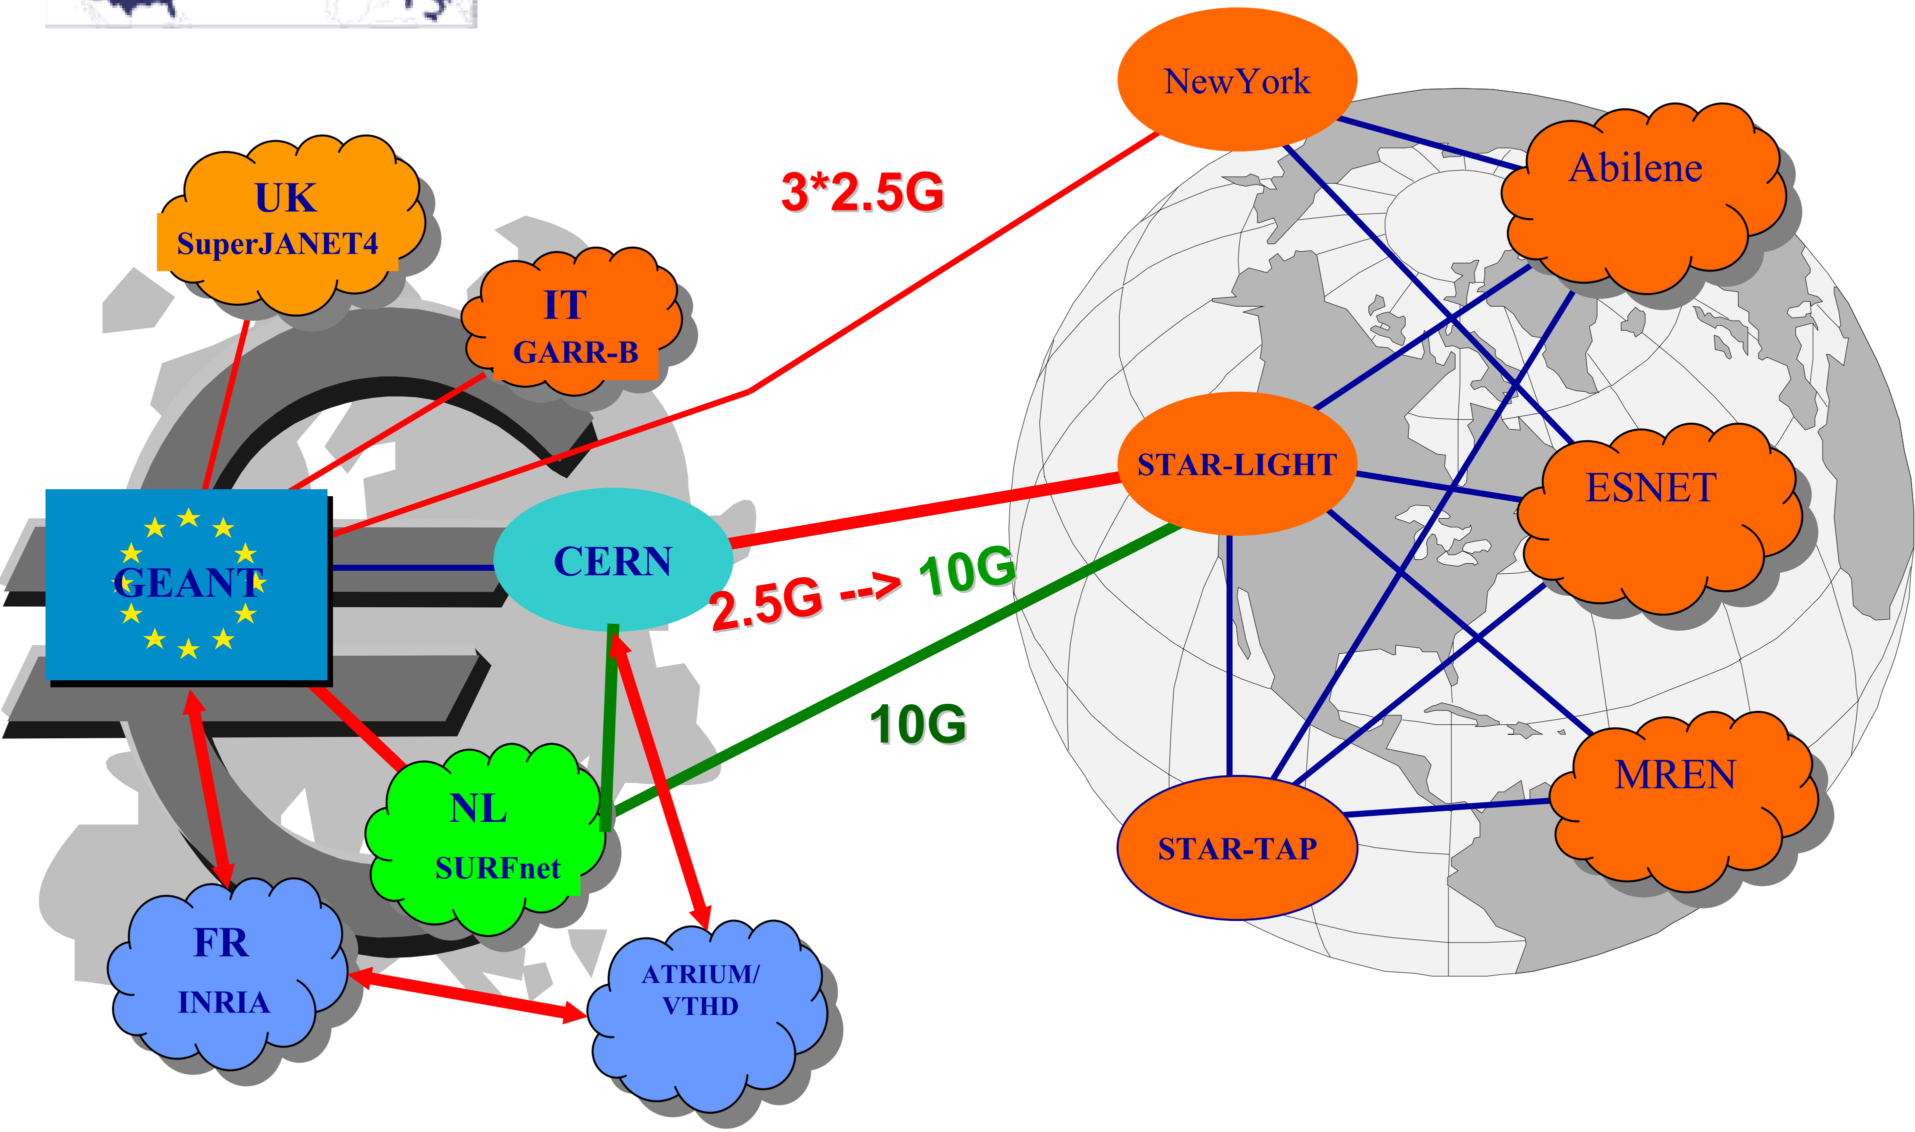

Land Speed Record

DataTAG

Information Society Technologies

California Institute of Technology 1891

Stanford Linear Accelerator Center

TERAGRID

Los Alamos National Laboratory

Argonne National Laboratory

European Commission

On February 27-28, a Terabyte of data was transferred by S. Ravot of Caltech between the Level3 PoP in Sunnyvale near SLAC and CERN through the TeraGrid router at StarLight from memory to memory as a single TCP/IP stream with 9KB Jumbo frames at a rate of 2.38 Gbps for 3700 seconds. This beat the former record by a factor of approximately 2.5, and used the US-CERN link at 96% efficiency. This is equivalent to:

- Transferring a full CD in 2.3 seconds (i.e. 1565 CDs/hour)
- Transferring 200 full length DVD movies in one hour (i.e. 1 DVD in 18 seconds)

Deutsche Telekom

Level(3) Communications — The Network Partner You Can Rely On.

STARLIGHT℠ The Optical STAR TAP℠

intel.

Cisco Systems

DataTAG Summary

- First year review successfully passed
  - GRID interoperability demo during the review:
    - Glue information system/EDG infoprovoders/EDG RB-glue
    - VOMS
    - GRID monitoring
    - LHC experiment applicaton using interoperable GRID
- Demonstration of applications running across heterogeneous Grid domains: EDG/VDT/LCG
- Comprehensive Transatlantic testbed built
- Advances in very high rate data transport

Enabling Grids for E-science in Europe

A seamless international Grid infrastructure to provide researchers in academia and industry with a distributed computing facility

STRATEGY

- Leverage current and planned national and regional Grid programmes
- Build on existing investments in Grid Technology by EU and US
- Exploit the international dimensions of the HEP-LCG programme
- Make the most of planned collaboration with NSF CyberInfrastructure initiative

ACTIVITY AREAS

SERVICES

- Deliver "production level" grid services (manageable, robust, resilient to failure)
- Ensure security and scalability

MIDDLEWARE

- Professional Grid middleware re-engineering activity in support of the production services

NETWORKING

- Proactively market Grid services to new research communities in academia and industry
- Provide necessary education

PARTNERS

70 partners organized in nine regional federations

Coordinating and Lead Partner: CERN

CENTRAL EUROPE – FRANCE - GERMANY & SWITZERLAND – ITALY - IRELAND & UK - NORTHERN EUROPE - SOUTH-EAST EUROPE - SOUTH-WEST EUROPE – RUSSIA - USA

May 2003: proposal submitted

July 2003: proposal accepted

Sept - Oct 2003: EU contract negotiation

April 2004: planned start of project

CERN

Information Society Technologies

European Commission

EGEE: Enabling Grids for E-science in Europe

Goals

- **Create a European-wide Grid Infrastructure for the support of research in all scientific areas, on top of the EU Reseach Network infrastructure**
- **Establish the EU part of a world-wide Grid infrastructure for research**

Strategy

- Leverage current and planned national and regional Grid programmes (e.g. LCG)
- Build on EU and EU member states major investments in Grid Technology
- Work with relevant industrial Grid developers and National Reseach Networks
- Take advantage of pioneering prototype results from previous Grid projects
- Exploit International collaboration (US and Asian/Pacific)
- Become the natural EU counterpart of the US NSF Cyber-infrastructure

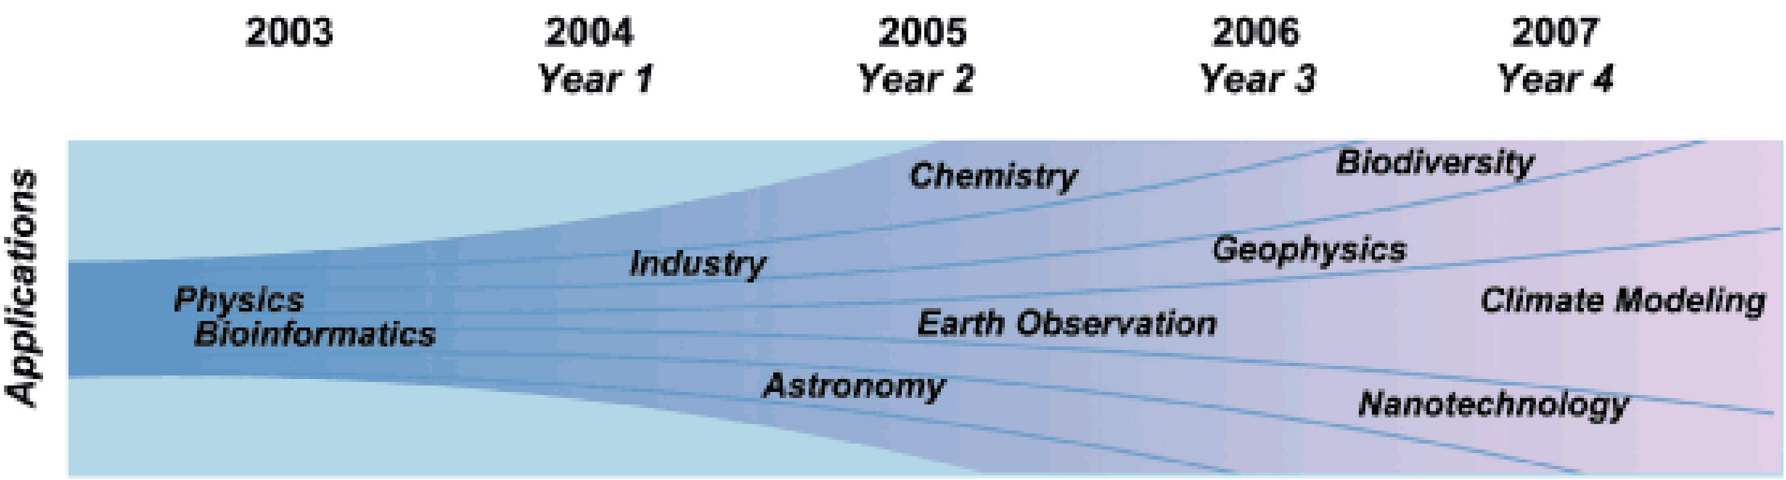

EGEE: partner federations

- Integrate regional grid efforts
- Represent leading grid activities in Europe

9 regional federations covering 70 partners in 26 countries

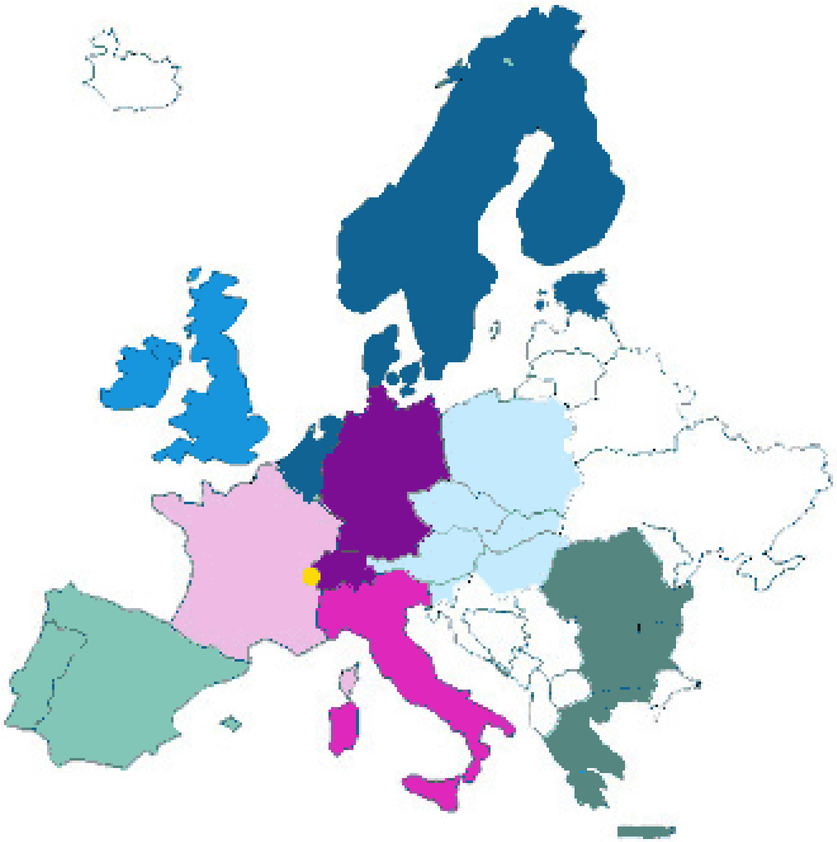

- CERN
- Central Europe (Austria, Czech Republic, Hungary, Poland, Slovakia, Slovenia)
- France
- Germany and Switzerland
- Ireland and UK
- Italy
- Northern Europe (Belgium, Denmark, Estonia, Finland, The Netherlands, Norway, Sweden)
- South-East Europe (Bulgaria, Cyprus, Greece, Israel, Romania)
- South-West Europe (Portugal, Spain)

■ GÉANT (plus NRENs*)

The “super-vehicle” for information transfer

- World leading Research Network
- Connecting more than 3100 Universities and R&D centers
- Over 32 countries across Europe
- Connectivity to NA, Japan, …
- Speeds of up to 10 Gbps
- Focus on the needs of “very demanding” user communities (PoC radio astronomers)

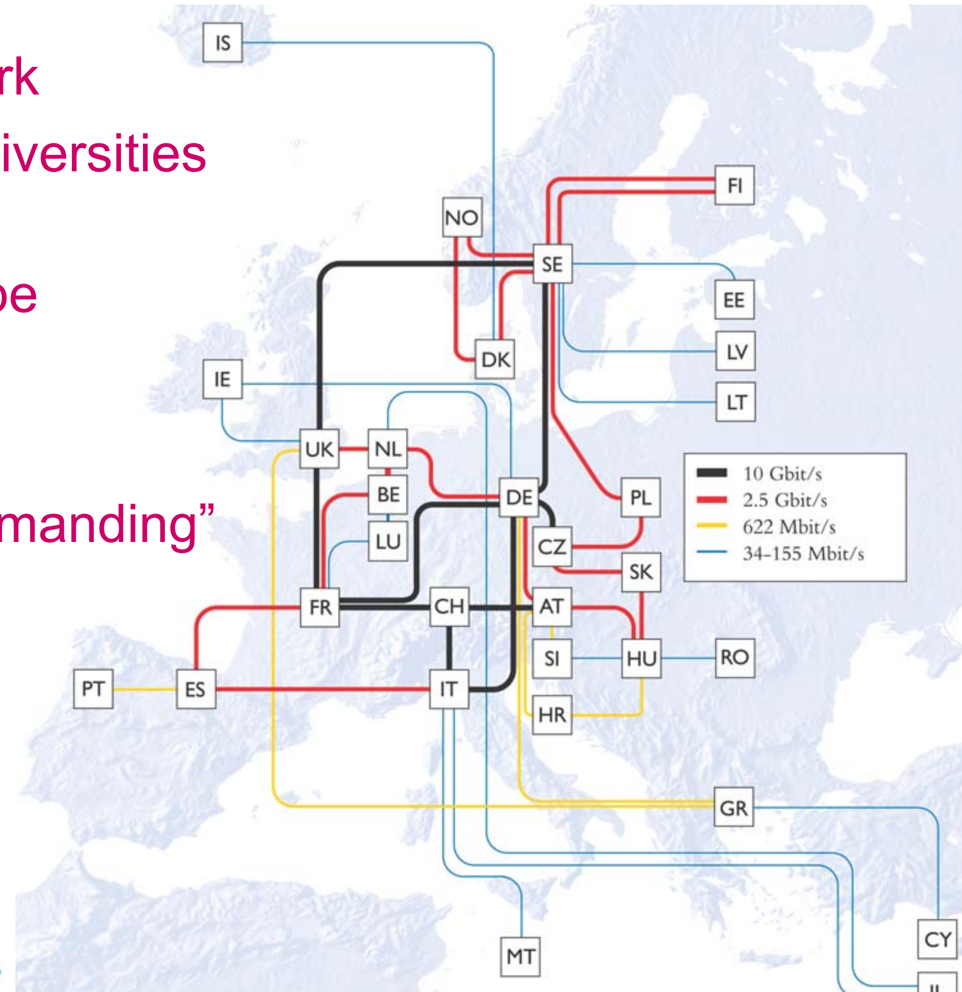

* National Research and Education Networks

European Commission

Information Society

■ GÉANT - a world of opportunities

- EU-North America: 7,5 Gbps (soon + 10Gbps)
- EU-Japan: 155Mbps + 2,5Gbps
- EUMEDIS: Evaluation of results of CfT in progress
- ALIS: Contract signed, CfT launched and CLARA in place
- TEIN: 20 Mbps (new developments in the course of the year)
- South Africa: Good prospects for connectivity

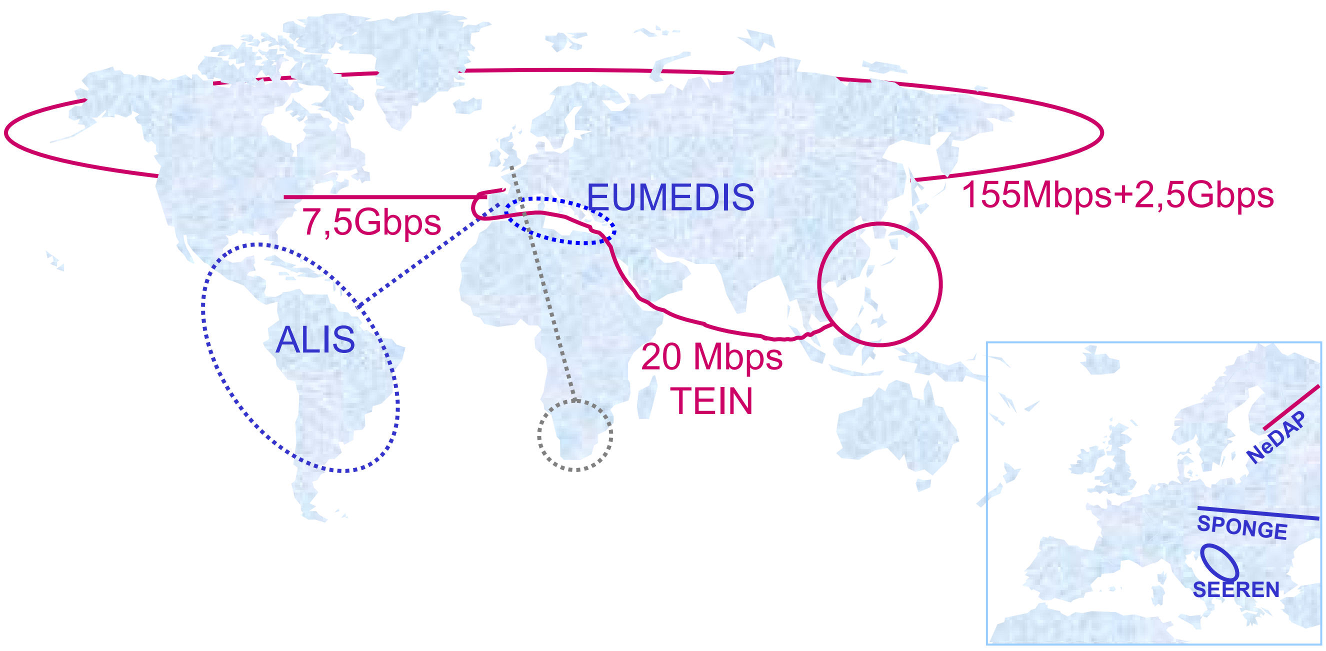

EGEE Proposal

- Proposal submitted to EU IST 6th framework call on 6th May 2003
  - Executive summary (exec summary: 10 pages; full proposal: 276 pages)
    - http://agenda.cern.ch/askArchive.php?base=agenda&categ=a03816&id=a03816s5%2Fdocuments%2FEGEE-executive-summary.pdf

- Activities
  - Deployment of Grid Infrastructure
    - Provide a grid service for science research
    - Initial service will be based on LCG-1
    - Aim to deploy re-engineered middleware at the end of year 1
  - Re-Engineering of grid middleware
    - OGSA environment – well defined services, interfaces, protocols
    - In collaboration with US and Asia-Pacific developments
    - Using LCG and HEP experiments to drive US-EU interoperability and common solutions
    - A common design activity should start now
  - Dissemination, Training and Applications
    - Initially HEP & Bio

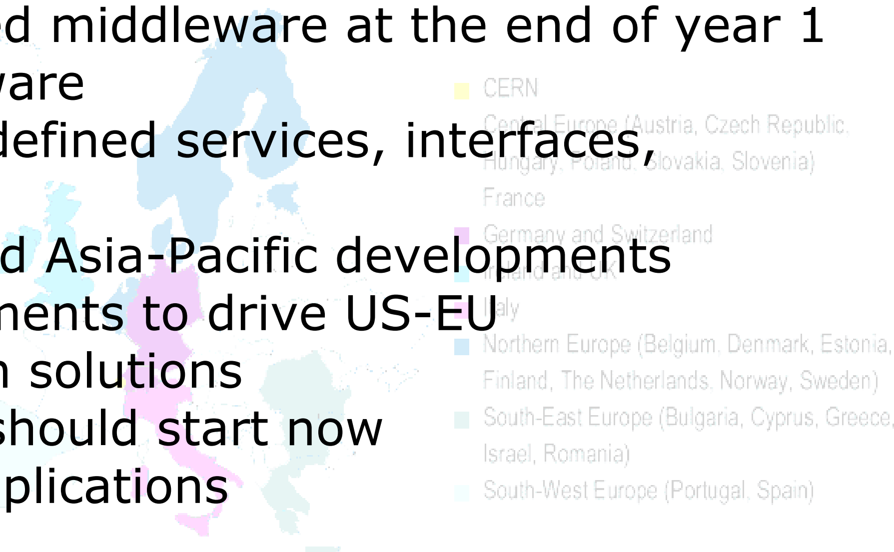

EGEE: timeline

- **May 2003**
  proposal submitted
- **July 2003**
  positive EU reaction
- **September 2003**
  start negotiation
  approx 32 M€ over 2 years
- **December 2003**
  sign EU contract
- **April 2004**
  start project

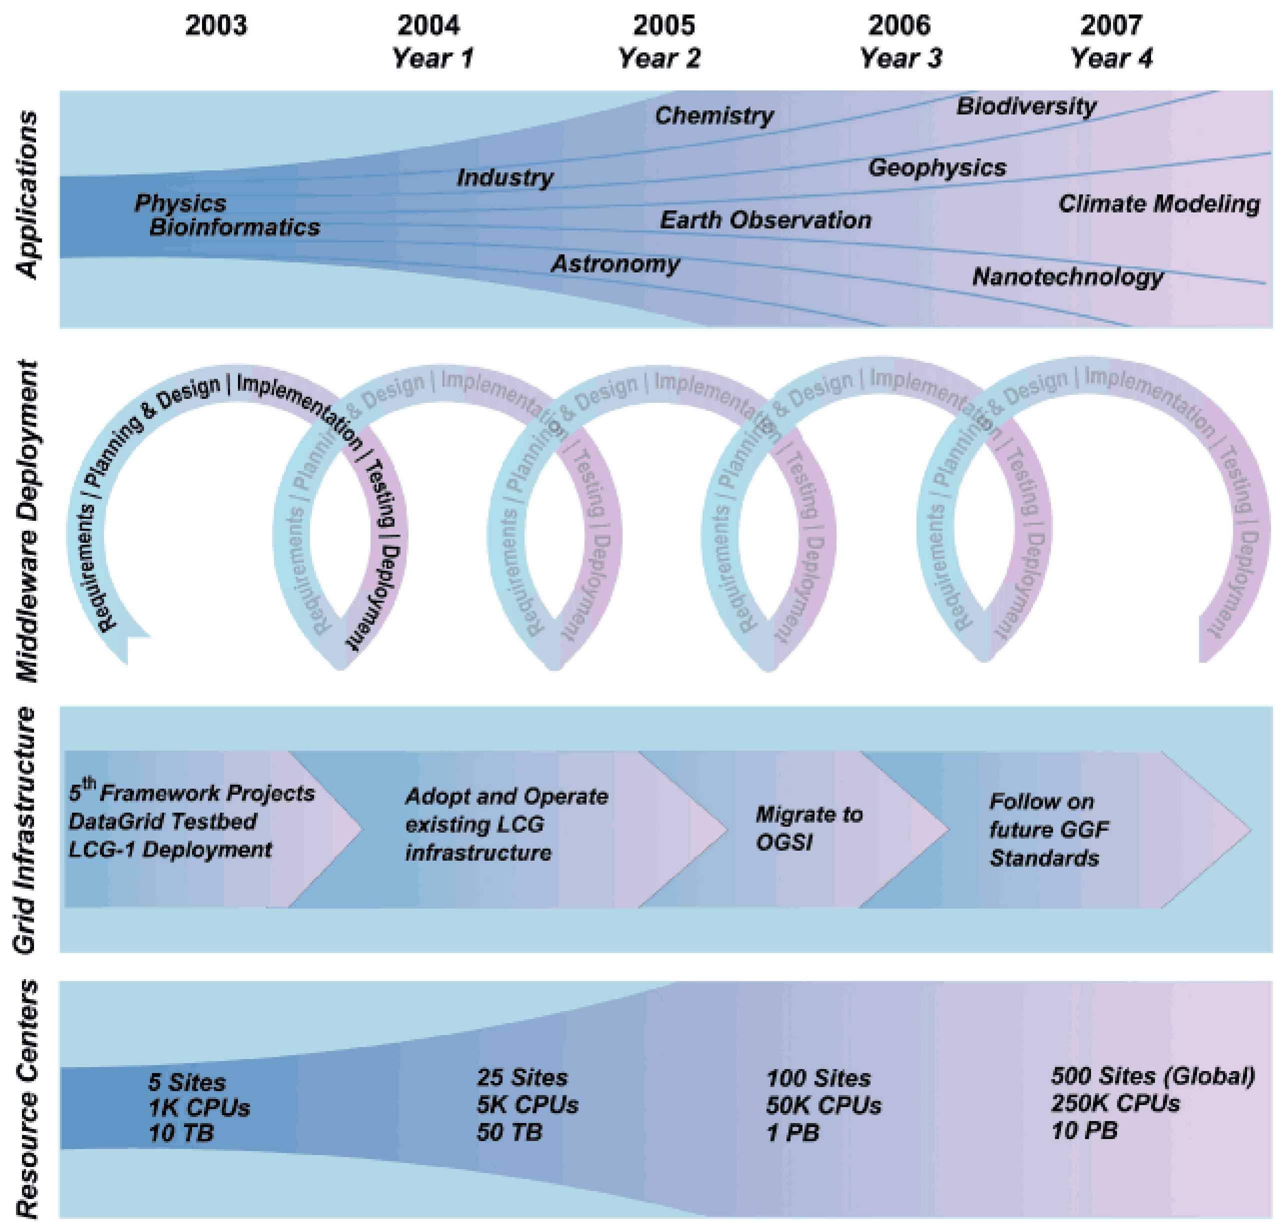

The LHC Computing Grid (LCG) Project

LCG - Goals

- The goal of the LCG project is to prototype and deploy the computing environment for the LHC experiments

- Two phases:
  - **Phase 1: 2002 – 2005**
  - Build a service prototype, based on existing grid middleware
  - Gain experience in running a production grid service
  - Produce the TDR for the final system

  - **Phase 2: 2006 – 2008**
  - Build and commission the initial LHC computing environment

LCG is not a development project – it relies on other grid projects for grid middleware development and support

LHC Computing Grid Project

- The LCG Project is a **collaboration** of –
  - The LHC experiments
  - The Regional Computing Centres
  - Physics institutes

  .. working together to prepare and deploy the computing environment that will be used by the experiments to analyse the LHC data

- This includes support for **applications**
  - provision of common tools, frameworks, environment, data persistency
- .. and the development and operation of a **computing service**
  - exploiting the resources available to LHC experiments in computing centres, physics institutes and universities around the world
  - presenting this as a reliable, coherent environment for the experiments
  - the goal is to enable the physicist to concentrate on science, unaware of the details and complexity of the environment they are exploiting

Deployment Goals for LCG-1

- Production service for Data Challenges in 2H03 & 2004
  - Initially focused on batch production work
  - But '04 data challenges have (as yet undefined) interactive analysis
- Experience in close collaboration between the Regional Centres
  - Must have wide enough participation to understand the issues
- Learn how to maintain and operate a global grid
- Focus on a production-quality service
  - Robustness, fault-tolerance, predictability, and supportability take precedence; additional functionality gets prioritized
- LCG should be integrated into the sites' physics computing services – should not be something apart
  - This requires coordination between participating sites in:
    - Policies and collaborative agreements
    - Resource planning and scheduling
    - Operations and Support

2003 – 2004 Targets

Project Deployment milestones for 2003:

- Summer: Introduce the initial publicly available LCG-1 global grid service
  - With 10 Tier 1 centres in 3 continents
- End of year: Expanded LCG-1 service with resources and functionality sufficient for the 2004 Computing Data Challenges
  - Additional Tier 1 centres, several Tier 2 centres – more countries
  - Expanded resources at Tier 1s (e.g. at CERN make the LXBatch service grid-accessible)
  - Agreed performance and reliability targets

Resource commitments for 2004:

	CPU (kSI2K)	Disk TB	Support FTE	Tape TB
CERN	700	160	10.0	1000
Czech Rep.	60	5	2.5	5
France	420	81	10.2	540
Germany	207	40	9.0	62
Holland	124	3	4.0	12
Italy	507	60	16.0	100
Japan	220	45	5.0	100
Poland	86	9	5.0	28
Russia	120	30	10.0	40
Taiwan	220	30	4.0	120
Spain	150	30	4.0	100
Sweden	179	40	2.0	40
Switzerland	26	5	2.0	40
UK	1656	226	17.3	295
USA	801	176	15.5	1741
Total	5600	1169	120.0	4223

LHC Computing Grid Service

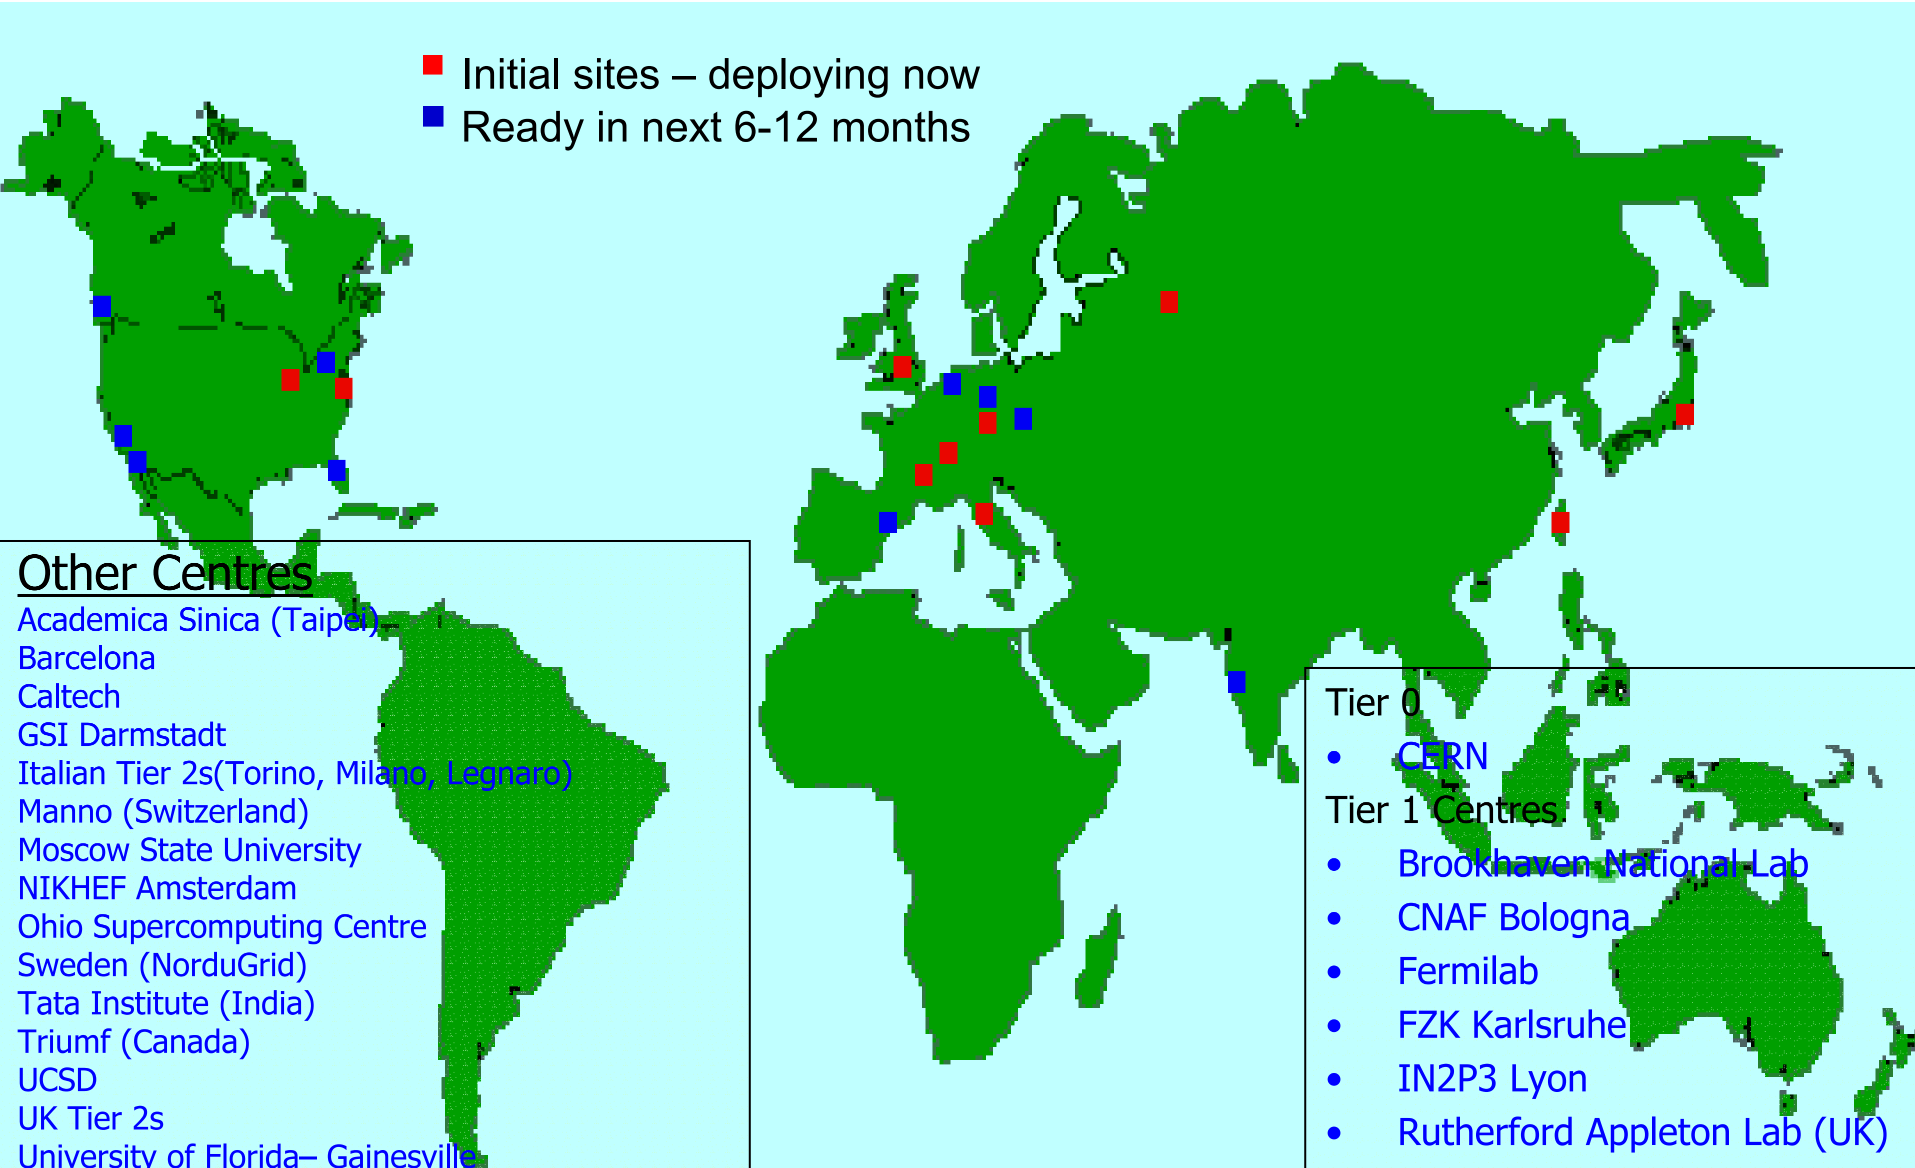

Other Centres

Academica Sinica (Taipei)
Barcelona
Caltech
GSI Darmstadt
Italian Tier 2s(Torino, Milano, Legnaro)
Manno (Switzerland)
Moscow State University
NIKHEF Amsterdam
Ohio Supercomputing Centre
Sweden (NorduGrid)
Tata Institute (India)
Triumf (Canada)
UCSD
UK Tier 2s
University of Florida– Gainesville
University of Prague
……

Tier 0

- CERN

Tier 1 Centres

- Brookhaven National Lab
- CNAF Bologna
- Fermilab
- FZK Karlsruhe
- IN2P3 Lyon
- Rutherford Appleton Lab (UK)
- University of Tokyo
- CERN

Elements of a Production LCG Service

- Middleware:
  - Testing and certification
  - Packaging, configuration, distribution and site validation
  - Support – problem determination and resolution; feedback to middleware developers
- Operations:
  - Grid infrastructure services
  - Site fabrics run as production services
  - Operations centres – trouble and performance monitoring, problem resolution – 24x7 globally
  - RAL is leading sub-project on developing operations services
    - Initial prototype –
      - Basic monitoring tools
      - Mail lists and rapid communications/coordination for problem resolution
- Support:
  - Experiment integration – ensure optimal use of system
  - User support – call centres/helpdesk – global coverage; documentation; training
  - FZK leading sub-project to develop user support services
    - Initial prototype –
      - Web portal for problem reporting
      - Expectation that initially experiments will triage problems and experts will submit LCG problems to the support service

Timeline for the LCG services

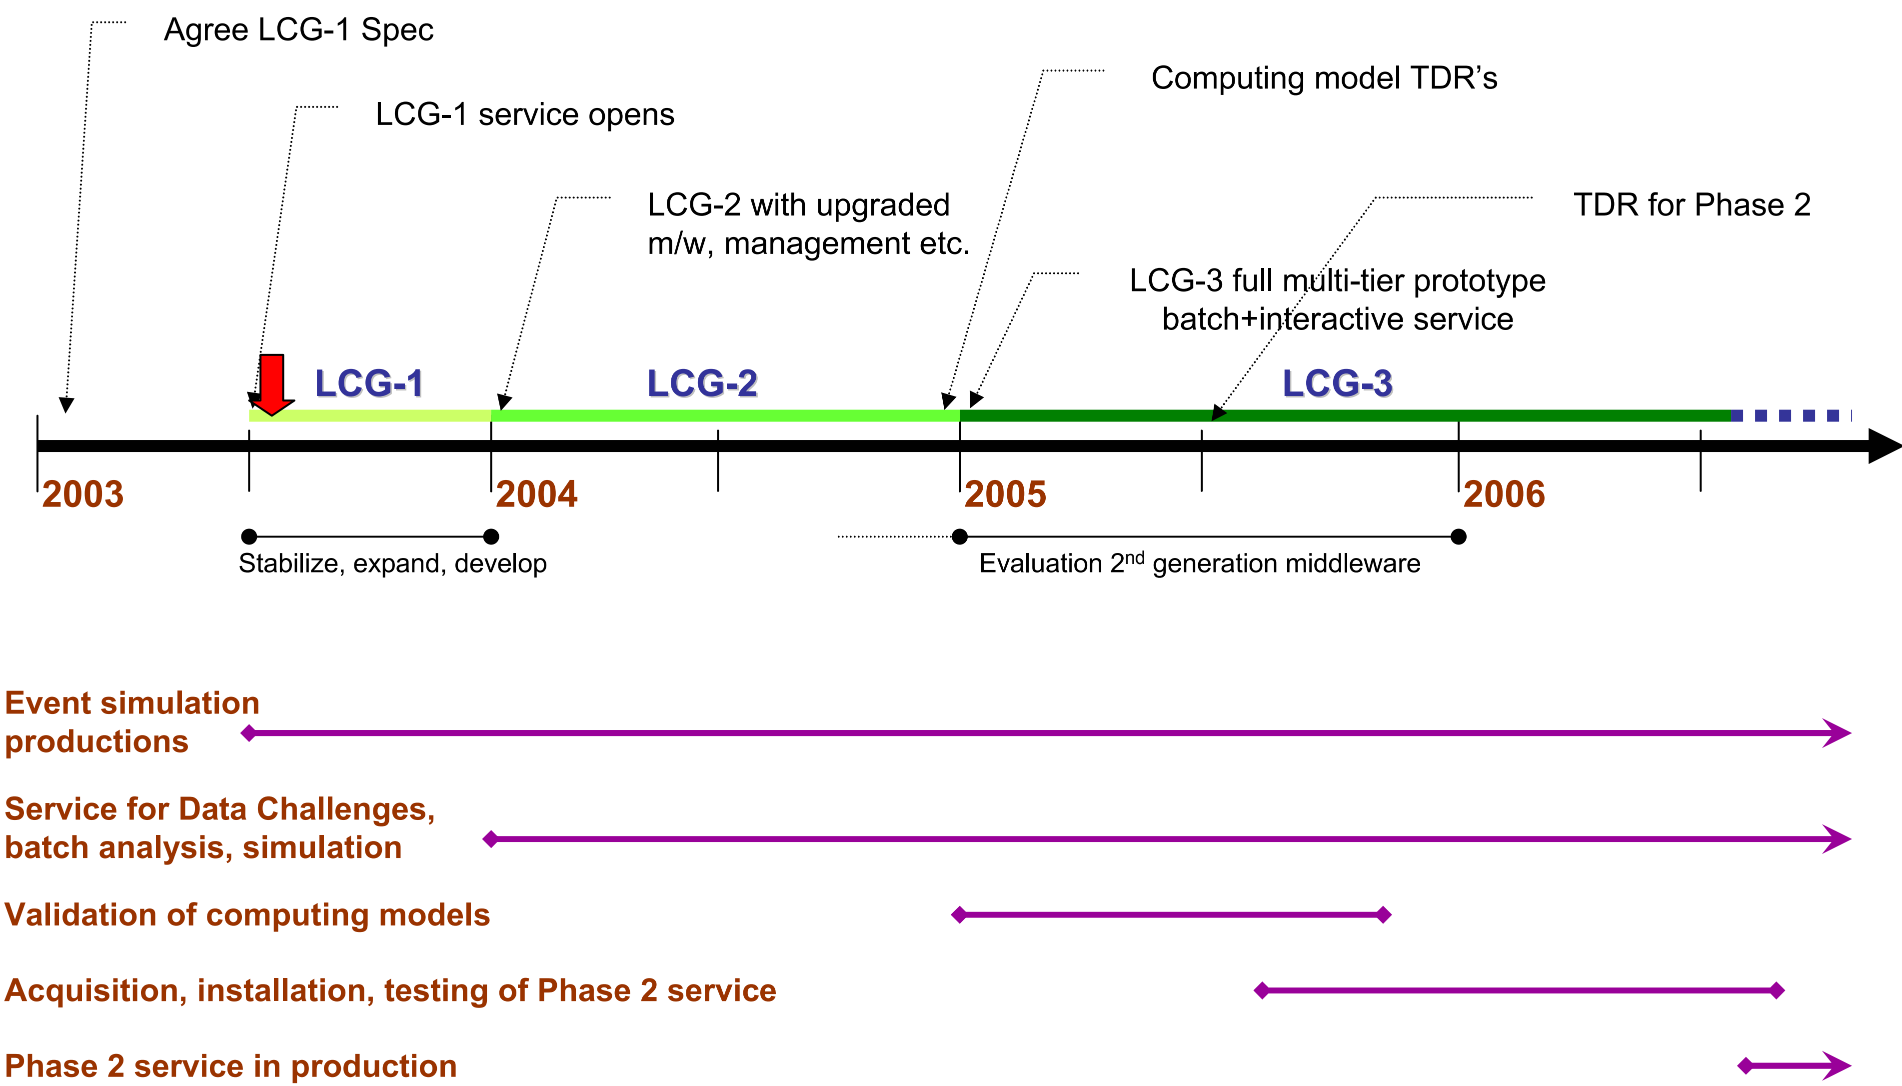

LCG-1 components

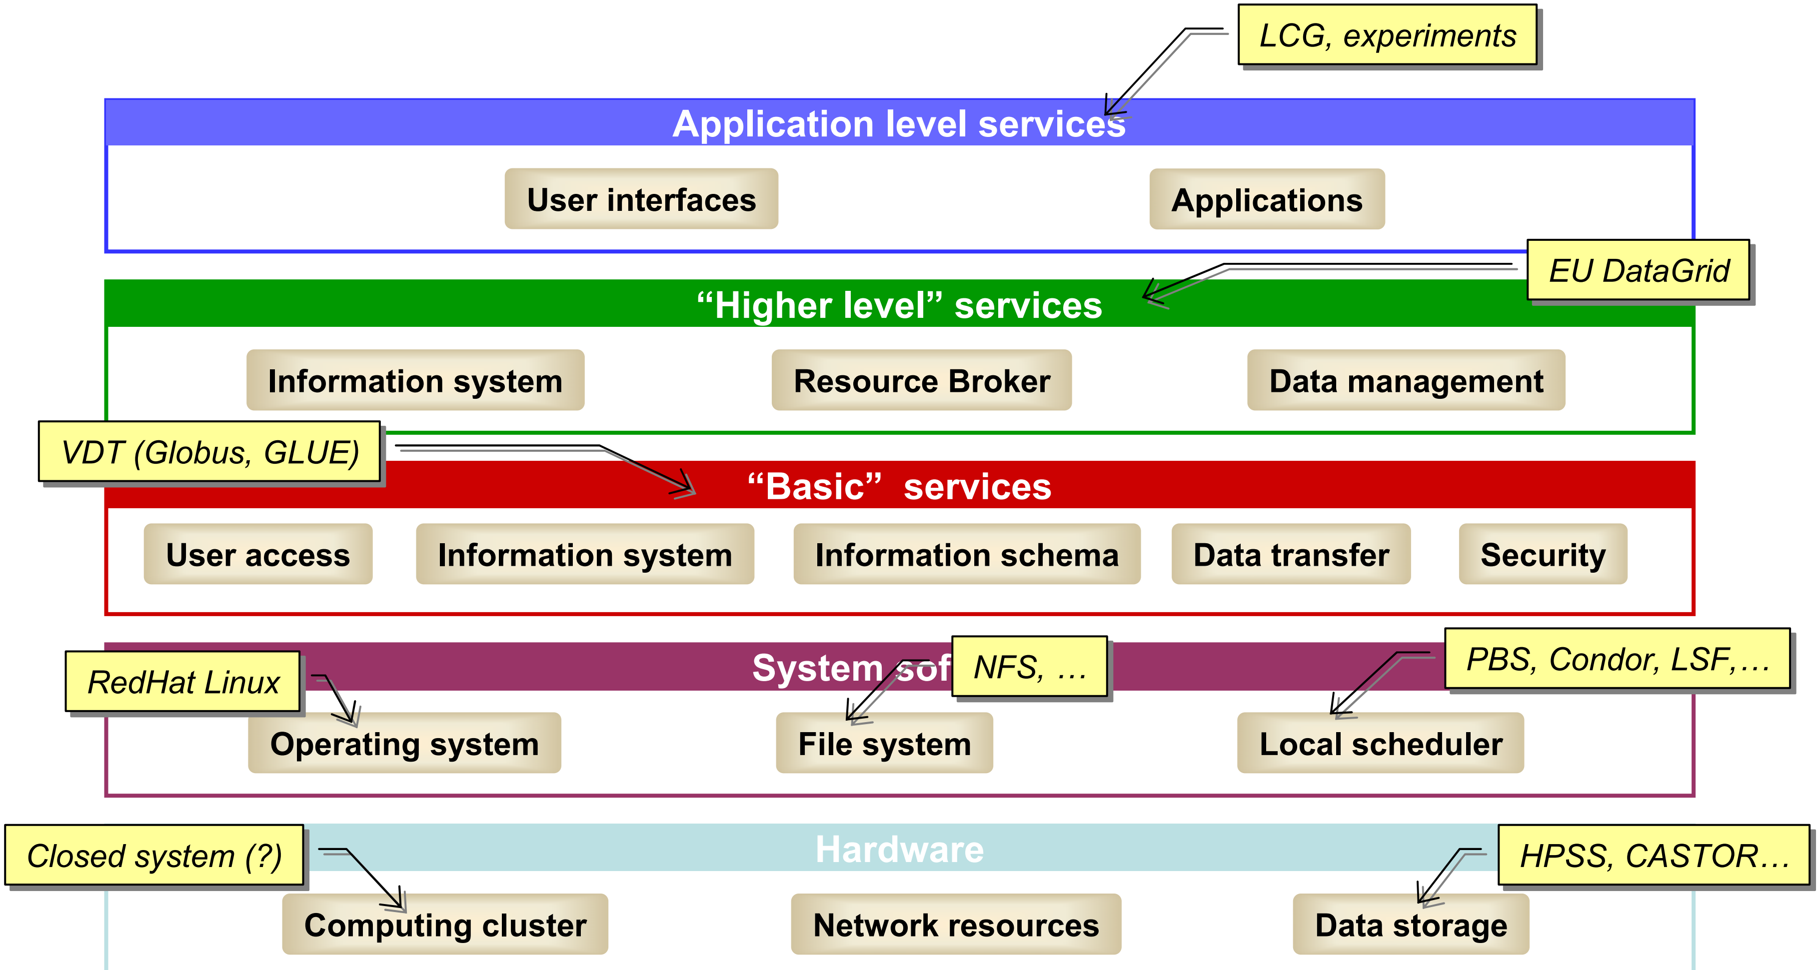

LCG – summary

- LHC data analysis has enormous requirements for storage and computation
- HEP
  - large global collaborations
  - good track record of innovative computing solutions
  - that do real work
- Grid technology offers a solution for LHC - to unite the facilities available in different countries in a virtual computing facility
- The technology is immature – but we need reliable solutions that can be operated round the clock, round the world
- The next three years work –
  - set up a pilot service – and use it to do physics
  - encourage the technology suppliers to work on the quality as well as the functionality of their software
  - learn how to operate a global grid

Outlook

LCG (and particle physics) as a major driving force to build interoperation and standardization

EU Vision of E-infrastructure in Europe

Moving towards an e-infrastructure

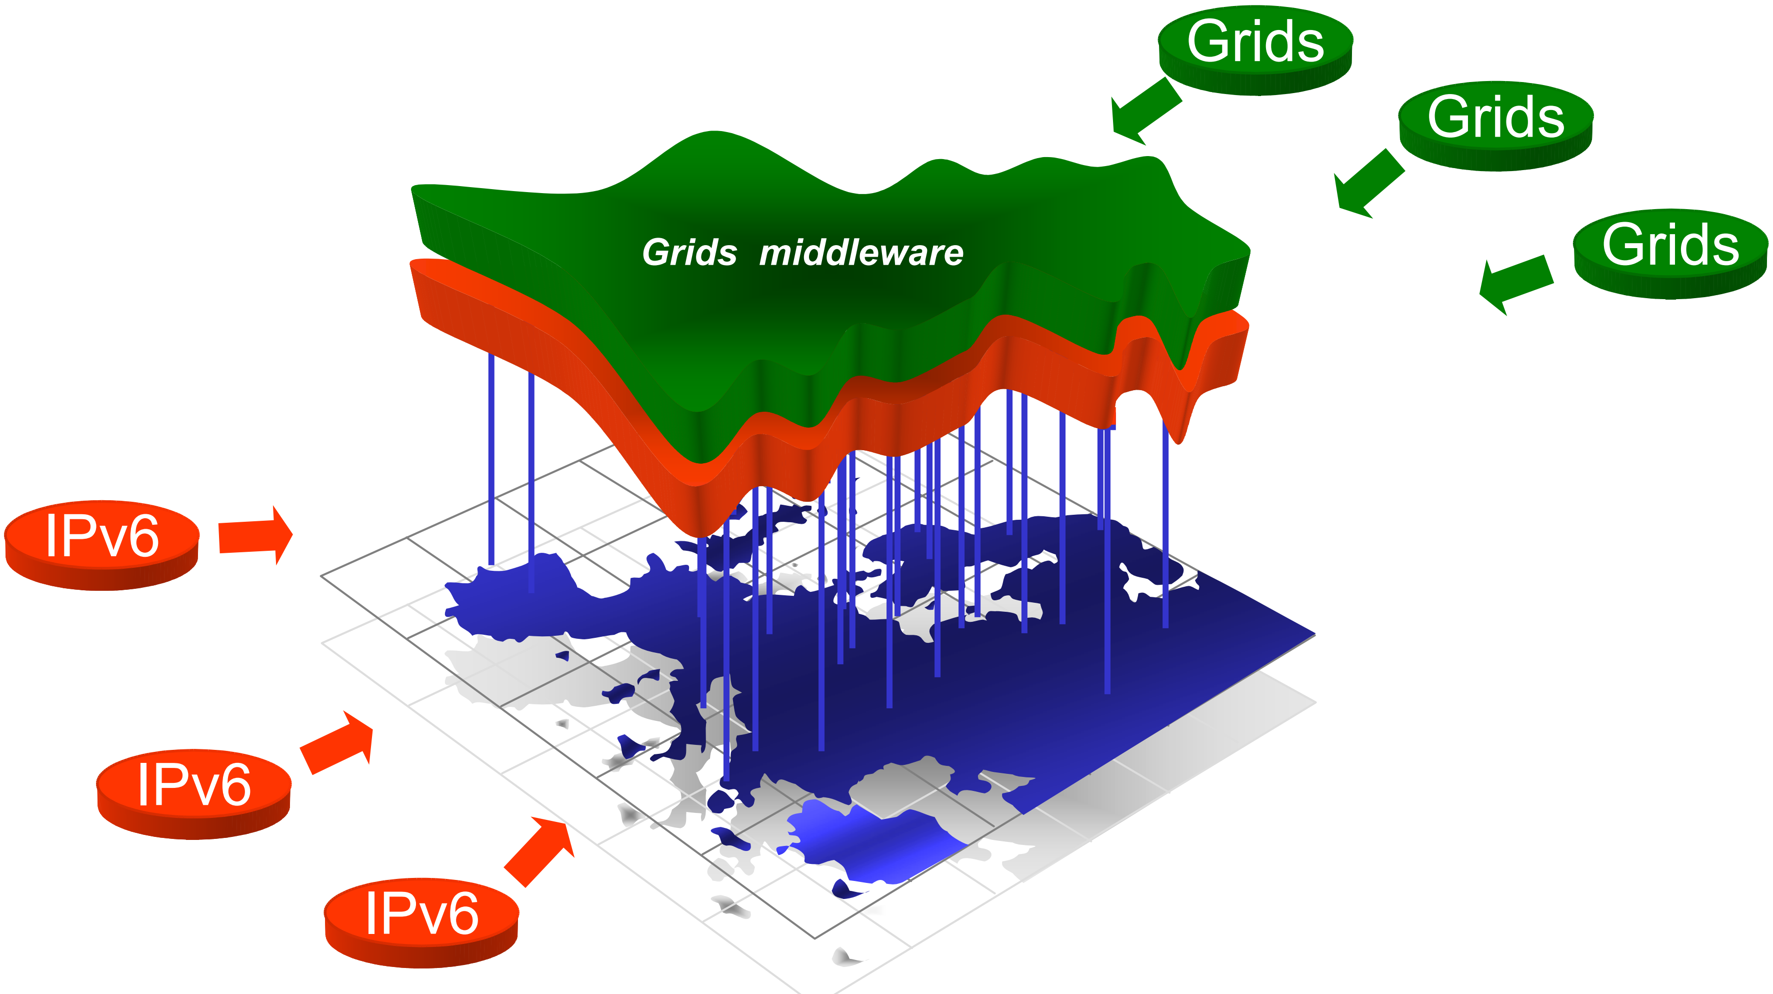

■ Moving towards an *e-infrastructure*

In FP6, deployment of a Grid empowered *e*-infrastructure

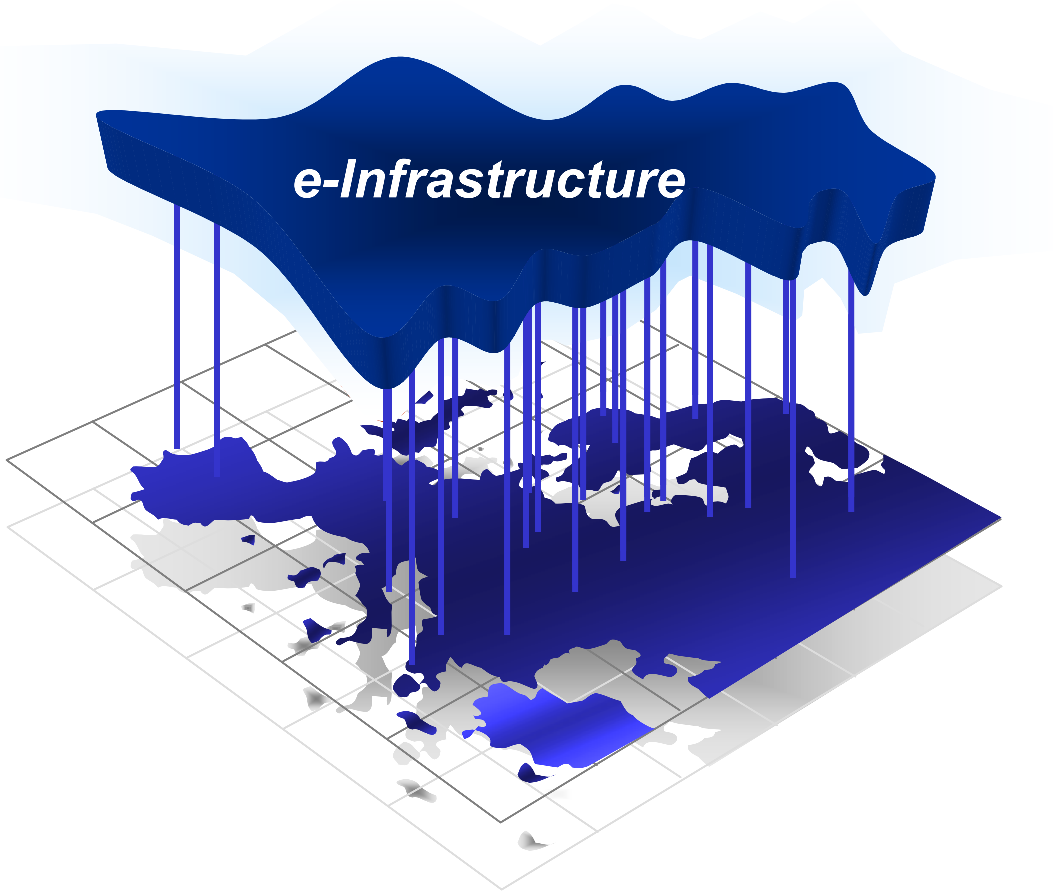

■ e-infrastructure - initial prospects (2004)

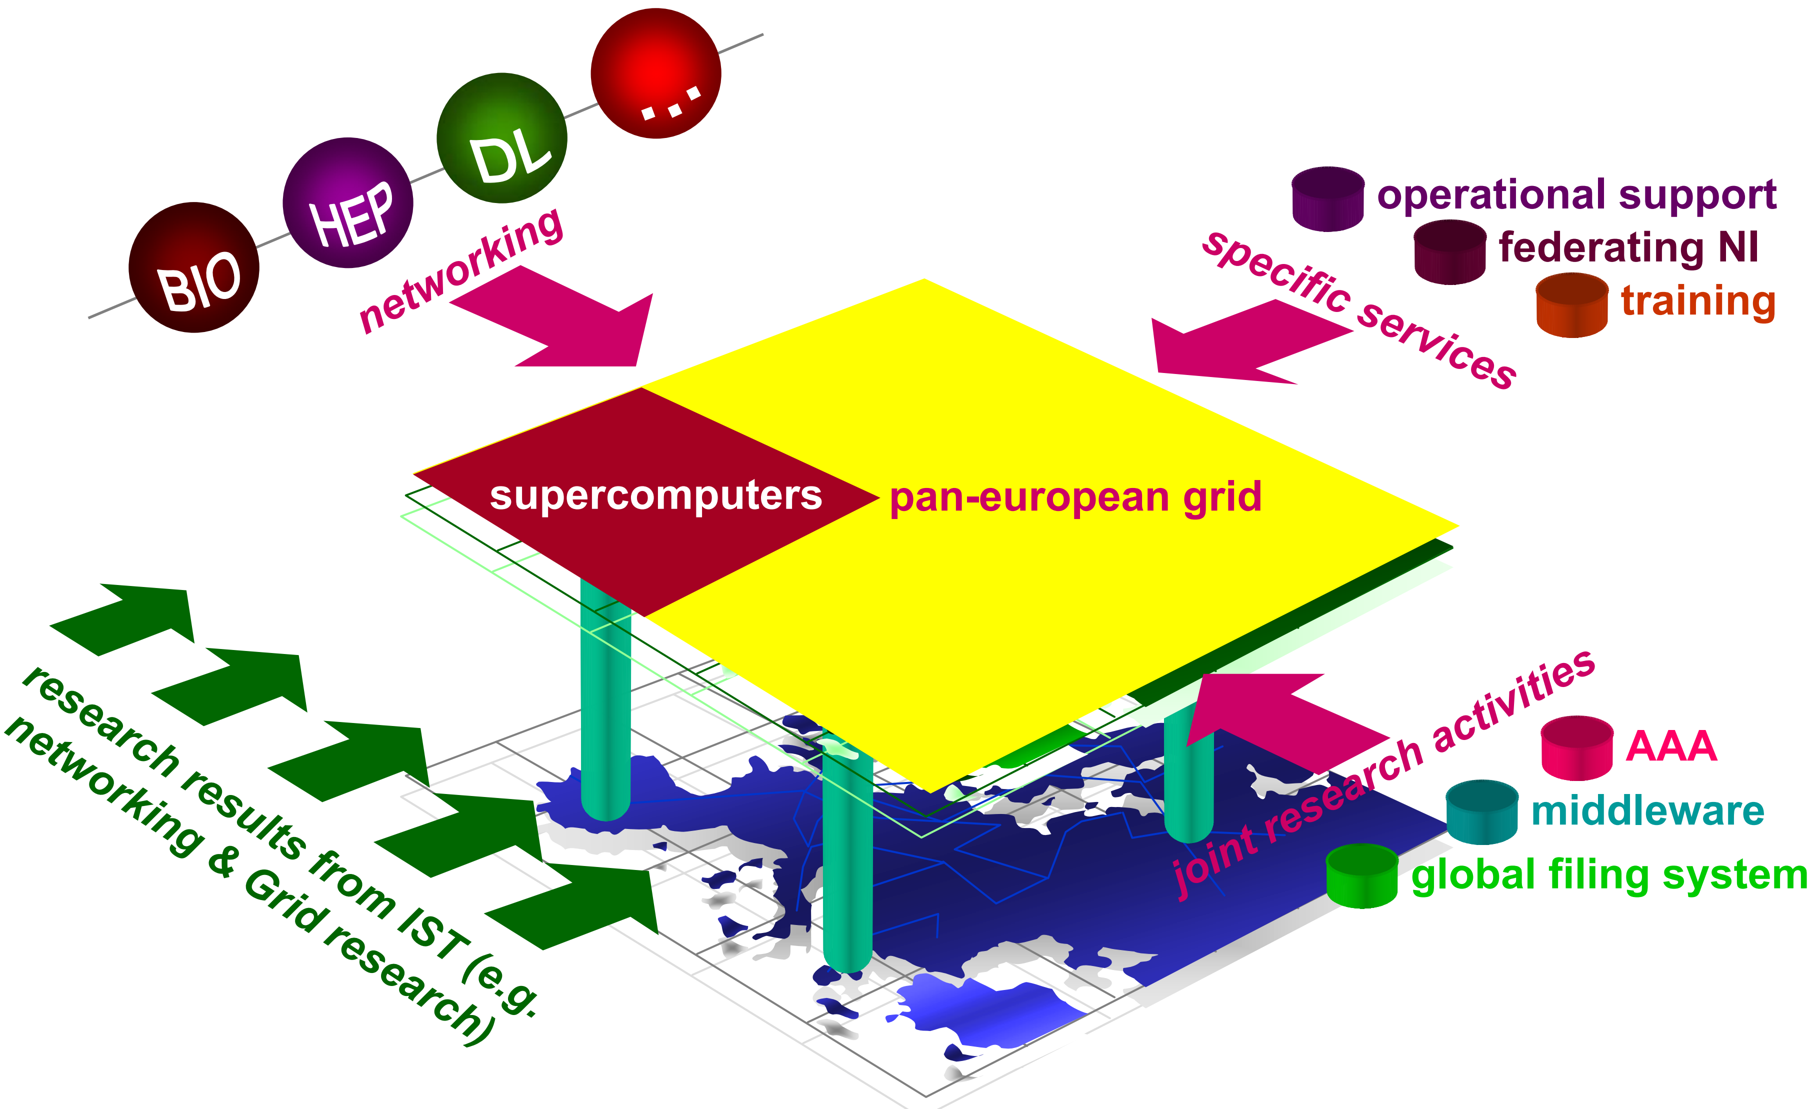

(international dimension to be taken from the start - cyberinfrastructure/Teragrid)

Interoperability for HEP

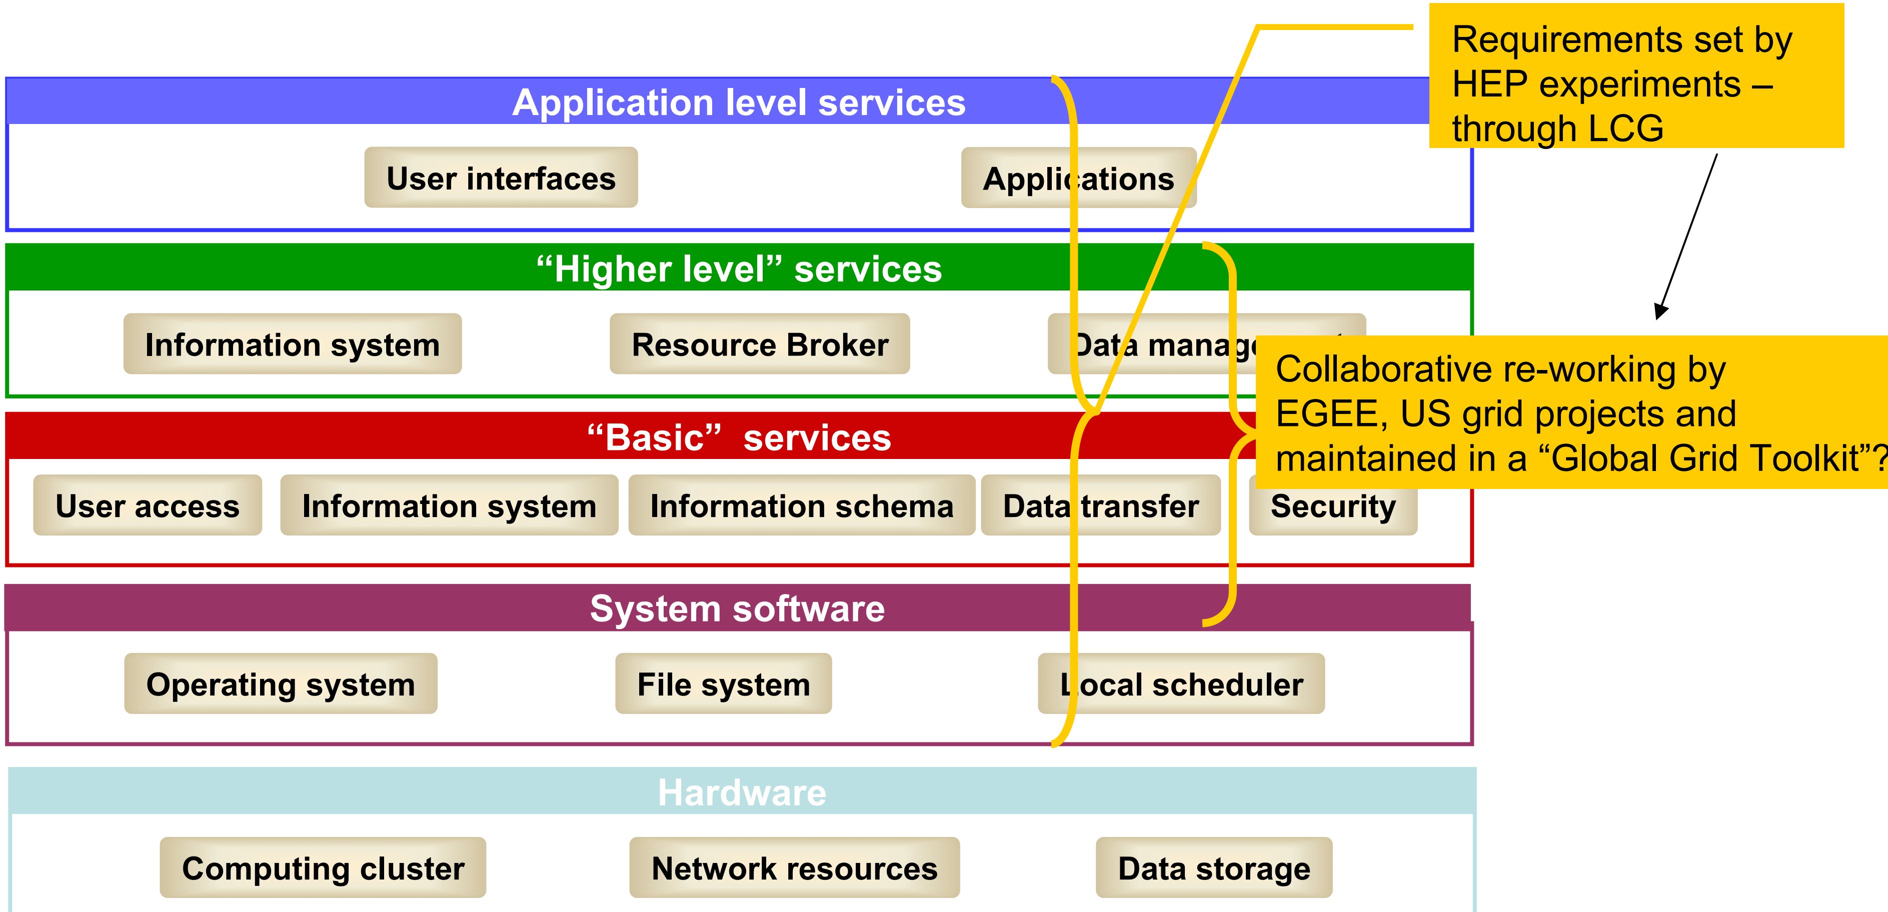

Relationship between LCG and grid projects

- LCG is collaboration representing the interests of the LHC experiments
- Negotiate with EGEE, US grid infrastructure, etc for services on behalf of the experiments

- Not just LHC experiments – other HENP communities exploring similar solutions
- Huge overlap of computing centres used by various experiments
  - Cannot have different grid solutions for each experiment
  - Must co-exist and inter-operate
  - Only way to inter-operate is through agreed standards and consistent implementations
  - Standards:
    - Service granularity
    - Service interfaces
    - Protocols

Standardization and interoperation

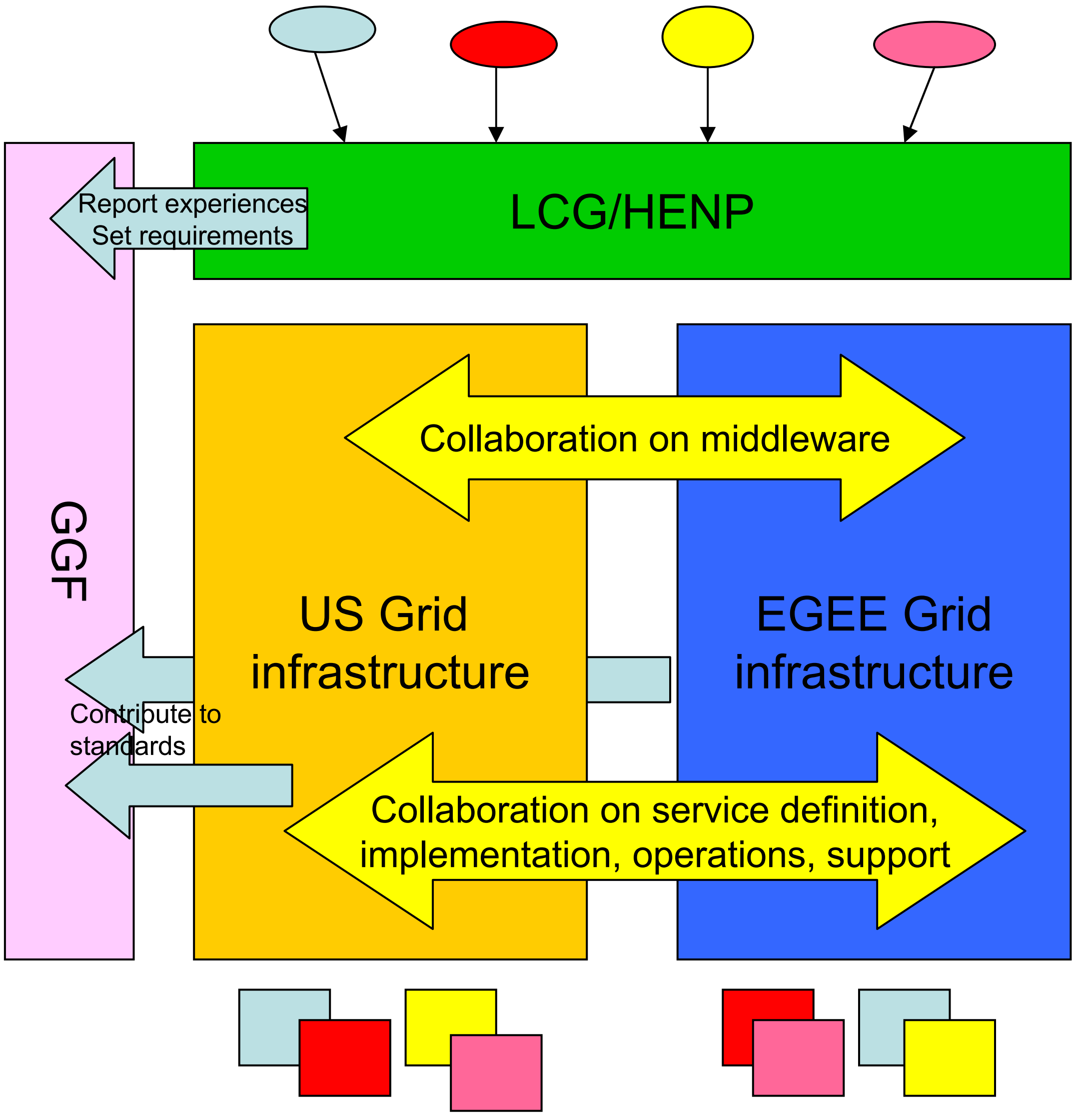

Experiment VOs

Drives common projects to ensure common solutions:
- Agreed service definitions
- Agreed interfaces
- Common protocols

Operate grid services on behalf of the customers (LCG, other sciences), Including support, problem resolution etc.
Implement policies set by the VOs for the use of resources

Resources owned by VOs

Summary

- Huge investment in e-science and grids in Europe
  - National and cross-national funded
  - EU funded
- Emerging vision of European-wide e-science infrastructure for research
  - Building upon and federating the existing national infrastructures
  - Peer with equivalent infrastructure initiatives in the US, Asia-Pacific
- High Energy Physics and LCG is a major application that needs this infrastructure today and is pushing the limits of the technology
  - Provides the international (global) dimension
- We must understand how to federate and use these infrastructures
  - A significant challenge – technology is not yet stable – there is no such thing today as a production-quality grid with the functionality we need
  - but we know already that we must make these interoperate