Tony Doyle
a.doyle@physics.gla.ac.uk

What is The Grid?

IOP HEPP Annual Conference 2004,
Birmingham, 7 April 2004

GRIDPP UK Particle Physics

THE UNIVERSITY OF BIRMINGHAM

University of Bristol

BRUNEL UNIVERSITY WEST LONDON

CERN

UNIVERSITY OF CAMBRIDGE

CCLRC

University of Durham

THE UNIVERSITY OF EDINBURGH

UNIVERSITY of GLASGOW

Imperial College London

LANCASTER UNIVERSITY

THE UNIVERSITY of LIVERPOOL

THE UNIVERSITY of MANCHESTER

UNIVERSITY OF OXFORD

Queen Mary University of London

Royal Holloway University of London

THE UNIVERSITY OF SHEFFIELD

THE UNIVERSITY OF SUSSEX

PRIFYSGOL CYMRU ABERTAWE
UNIVERSITY OF WALES SWANSEA

UCL

Contents

- **What is the Grid?**
  1. Hype
  2. Demo
  3. Prototype
  4. A Standard
  5. Evolving
  6. Irrelevant
- **GridPP summary**
  - Brief..
  - Middleware Hour Glass
  - Software Process
  - Hardware Resources
  - Is GridPP a Grid?
  - How does the Grid Work?
  - (Security) Interlude
  - The Challenges Ahead..
  - Philosophy
  - Historical Perspective
- **Answer?**

GridPP – the UK Computing Grid for Particle Physics

- **The UK's contribution to LHC computing**
  - **19 UK Institutes and CERN**
  - **collaboration of particle physicists and computer scientists**
  - **PPARC funded £33m over 6 years (2001-2007)**

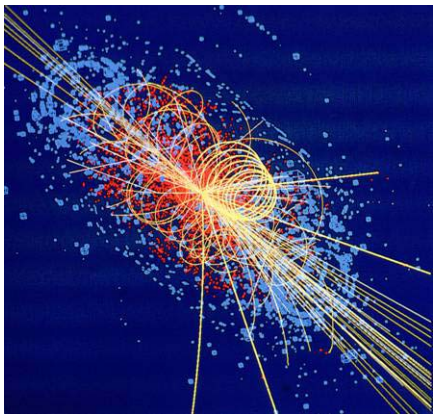

THE UNIVERSITY OF BIRMINGHAM
University of Bristol
BRUNEL UNIVERSITY WEST LONDON
CERN
CCLRC
UNIVERSITY OF CAMBRIDGE
University of Durham
UNIVERSITY of GLASGOW
Imperial College London
LANCASTER UNIVERSITY
THE UNIVERSITY of MANCHESTER
THE UNIVERSITY of LIVERPOOL
Queen Mary University of London
Royal Holloway University of London
THE UNIVERSITY OF SHEFFIELD
THE UNIVERSITY OF SUSSEX
PRIFYSGOL CYMRU ABERTAWE UNIVERSITY OF WALES SWANSEA
UCL

GridPP International Collaboration

- **CERN IT – PPARC funded £5.7m for staff and hardware**
  - prototype LHC Computing Grid went live in September 2003 in 12 countries
  - UK sites among the first to join
  - Grid Operations Centre based at RAL, UK
- **EU DataGrid – PPARC one of 6 main partners**
  - GridPP software central to the project
  - testbed at more than 40 sites across Europe with 1,000+ computer processors
  - around 100,000 jobs successfully submitted via the UK to the prototype EU-wide grid
  - 3 main areas
    - high energy physics
    - biology and medical image processing
    - Earth observation

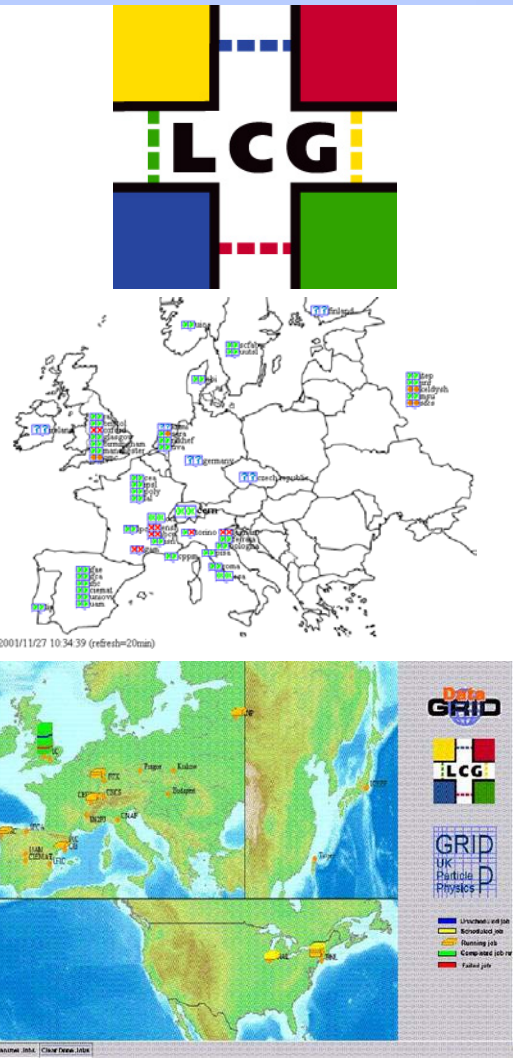

What is the Grid? 1: Hype

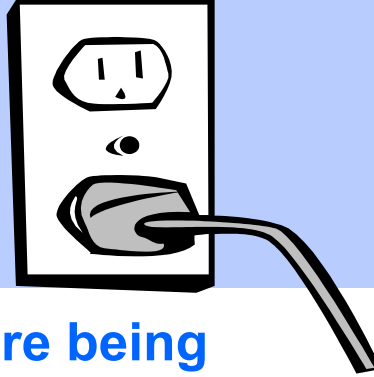

EITHER The new Grid computing technologies are set to revolutionise the way scientists use the world's computing resources?

"The technology now being deployed for particle physics will ultimately change the way that science and business are undertaken in the years to come. This will have a profound effect on the way society uses information technology, much as the World Wide Web did.

Grid technology will extend to fields like bioinformatics, digital archive and biodiversity informatics."

"'The Grid', as it is provisionally known, will work far more quickly and reliably than today's internet. It should eventually enable computer users to receive exactly the information they want from anywhere in the world within seconds – and without having to go through a tortuous search process."

"The Grid will link super-computers on every continent to create the first 'self-conscious' and virtual computer system, capable of exhibiting artificial intelligence"

"A scientist anywhere could access the world's most powerful computers. The Grid itself will decide where best to obtain data and necessary processing power"

"The end is in sight, even for the internet. Personal computers and the web are outdated technologies"

"When the network is as fast as the computer's internal links, the machine disintegrates across the net into a set of special purpose appliances"

"e-Science means science increasingly done through distributed global collaborations enabled by the Internet, using very large data collections, terascale computing resources and high performance visualisation"

"e-Science will change the dynamic of the way science is undertaken"

"We are very excited to be able to participate in such a revolutionary global collaboration."

OR Grids exist and are being routinely deployed now?

1,865,989 years of cpu time

4,935,832 users

- **PLEASE NOTE:**
- **Due to potential security problems user email addresses are not being shown - you can, however, now have your user name link to a URL that you specify. See the Account Change page to do so.**

SETI@home
The Search for Extraterrestrial Intelligence

313,868 years of cpu time

1,144,738 users

grid.org

Answer in the small print..

Tony Doyle - University of Glasgow

GRIP UK Particle Physics

Gartner Hype Cycle
Big Bang (or Big Crunch)

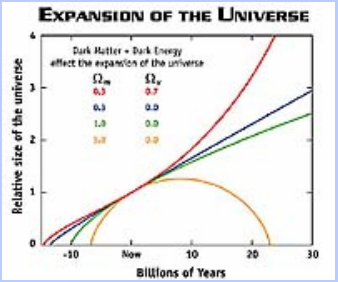

Business Development perspective:

Grid in the early stages of visibility

Research Development perspective?:

Grid in the trough of disillusionment

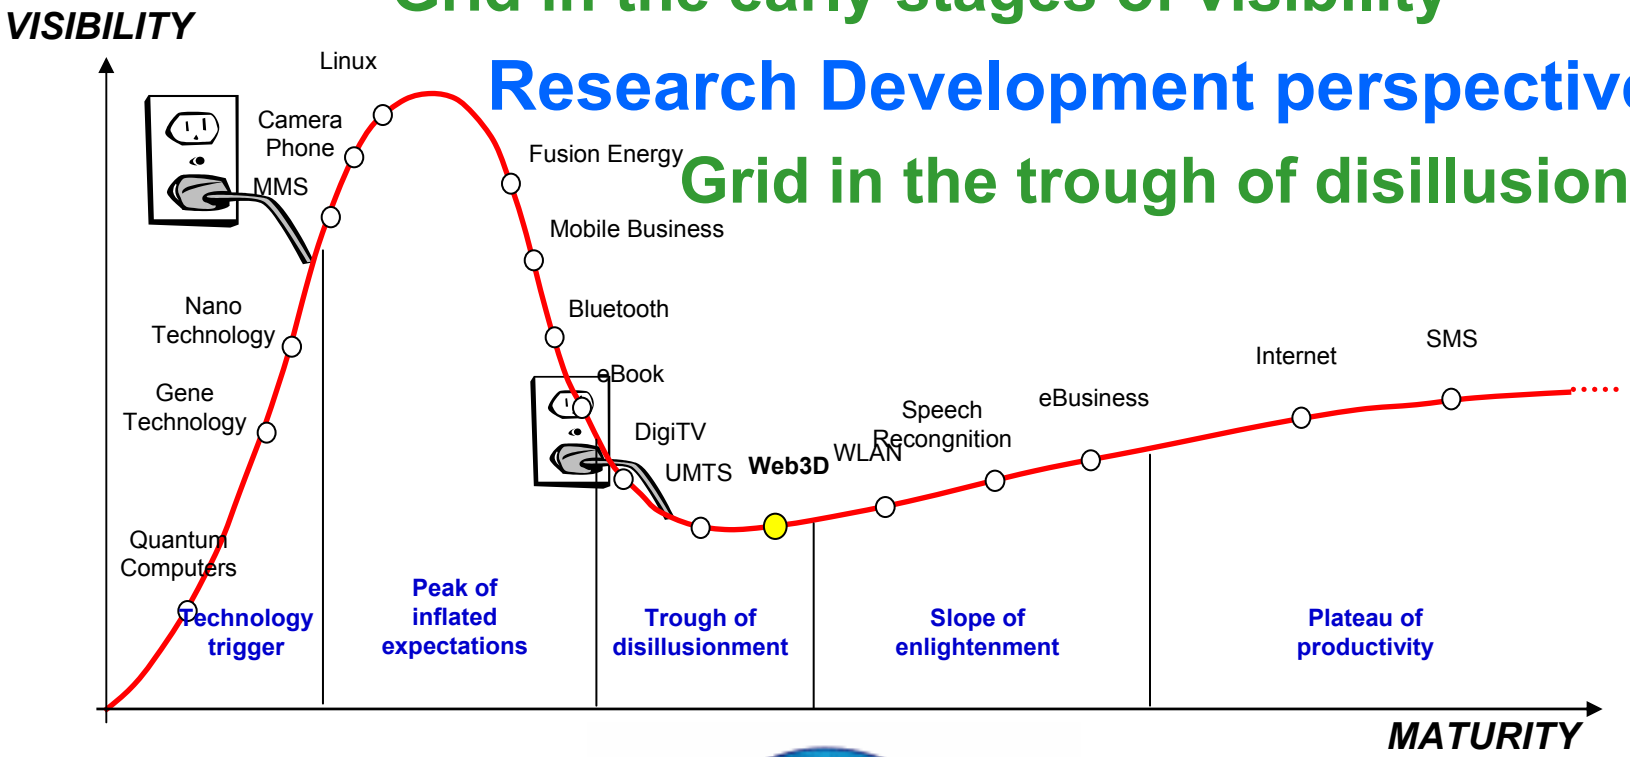

Take back the web

Tony Doyle - University of Glasgow

What is the Grid? 1: Hype?

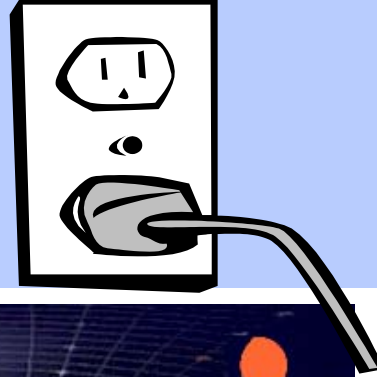

PPARC April 1st News:

PPARC web site was down..

"Grid Computing Steps up a Gear" Enabling Grids for E-Science in Europe launch

Substance behind the hype?

The Physics of the Universe

PPARC

TEXT VERSION WHAT'S NEW COMMENT CONTACTS SEARCH

Current phase of moon

Contents

RESEARCH PROGRAMMES

STUDENTS & FELLOWSHIPS

KNOWLEDGE TRANSFER

PUBLIC AND SCHOOLS

PRESS AND MEDIA

ABOUT PPARC

Welcome

The **Particle Physics and Astronomy Research Council [PPARC]** is the UK's strategic science investment agency. By directing, coordinating and funding research, education and training in particle physics and astronomy, PPARC delivers world-leading science, technologies and people for the UK. Based in Swindon, along with several other research councils, PPARC operates three scientific sites: the UK Astronomy Technology Centre in Edinburgh, the Isaac Newton Group in La Palma and the Joint Astronomy Centre in Hawaii.

Latest news

The PPARC web site will be unavailable between 19:00 GMT on Friday 2 April 2004 and 12:00 GMT on Sunday 4 April 2004 due to system maintenance. Apologies for any inconvenience.

01 April 2004

Grid Computing Steps up a Gear

UK plans for Grid computing changed gear this week. The pioneering European DataGrid (EDG) project came to a successful conclusion at the end of March, and on 1 April a new project, known as Enabling Grids for E-Science in Europe (EGEE), begins.

Quick links

Using this site
Including site map, quickfinder and access keys

CERN 50th Anniversary
2004 marks the 50th Anniversary of CERN

Diary
PPARC events and meetings

Funding
How to apply, schemes, forms, handbook and adminstration

Je-S 1 login
Prepare and submit research grant proposals to PPARC electronically using the Je-S 1 system

Frontiers
Frontiers, PPARC's flagship magazine

Grants on the Web
Summaries of all current PPARC research grants

Publications
Free publications, posters

The Grid: Demonstrations (QED?)

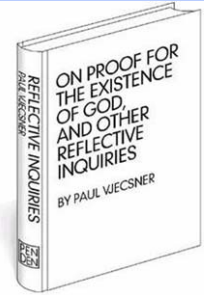

- Demos at this meeting established that the two LHC multi-purpose detector collaboartions
- can run jobs on an International Grid
- Use common Grid infrastructure with secure Grid access
- But doesn't mean that the Grid works in production mode
- (yet)
- This is however

signi**ficant**

- Thanks esp. to Frederic Brochu (ATLAS), Dave Colling (CMS)++

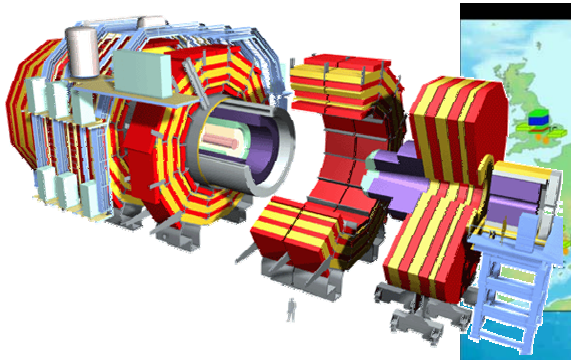

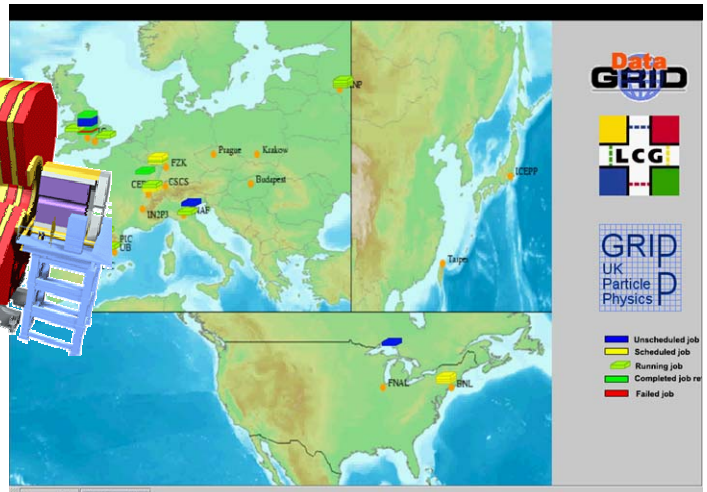

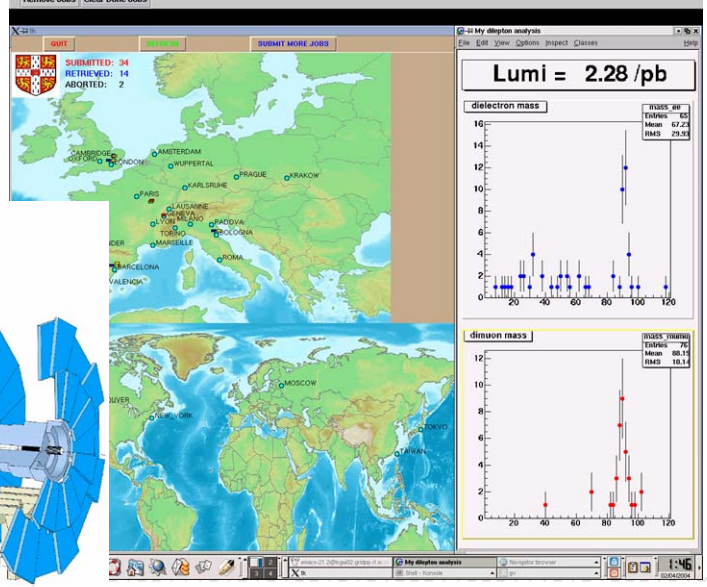

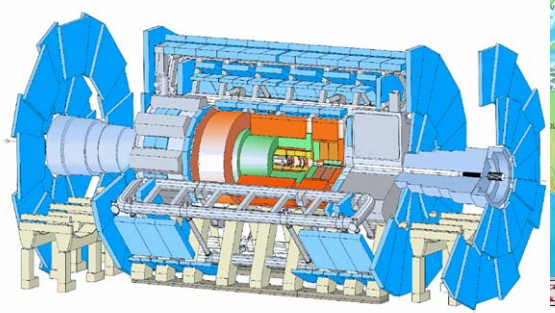

What is “The GRID”?
2. Demo

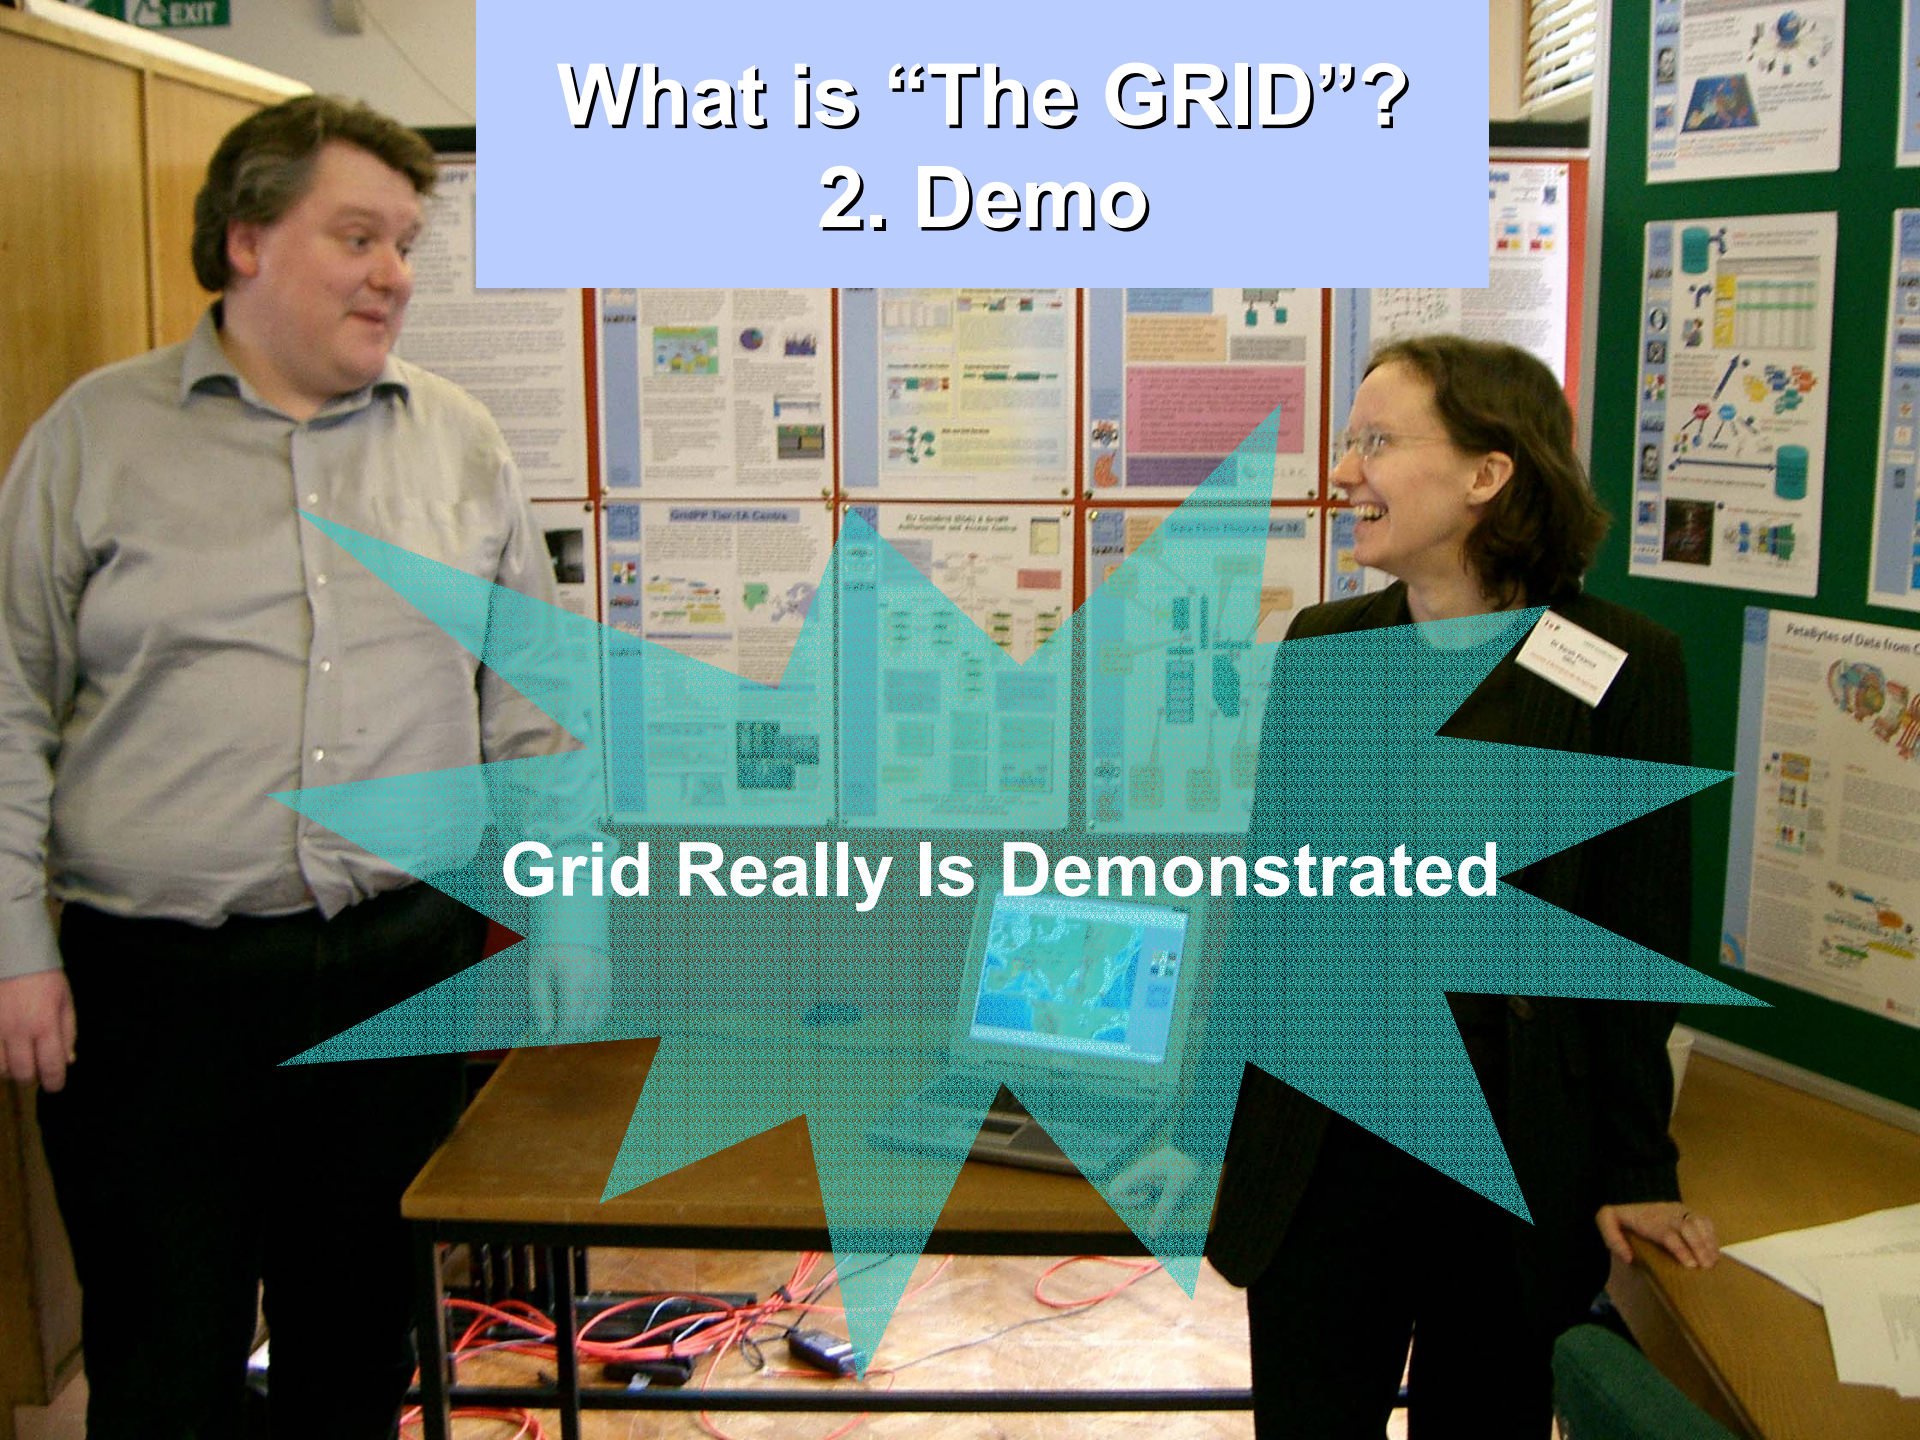

What is the Grid? 3. Prototype Performance indicators (as measured by end users)

$$E_{crude} = \frac{\text{Number of jobs successfully completed}}{\text{Total Number of jobs submitted}}$$

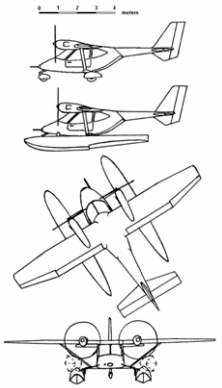

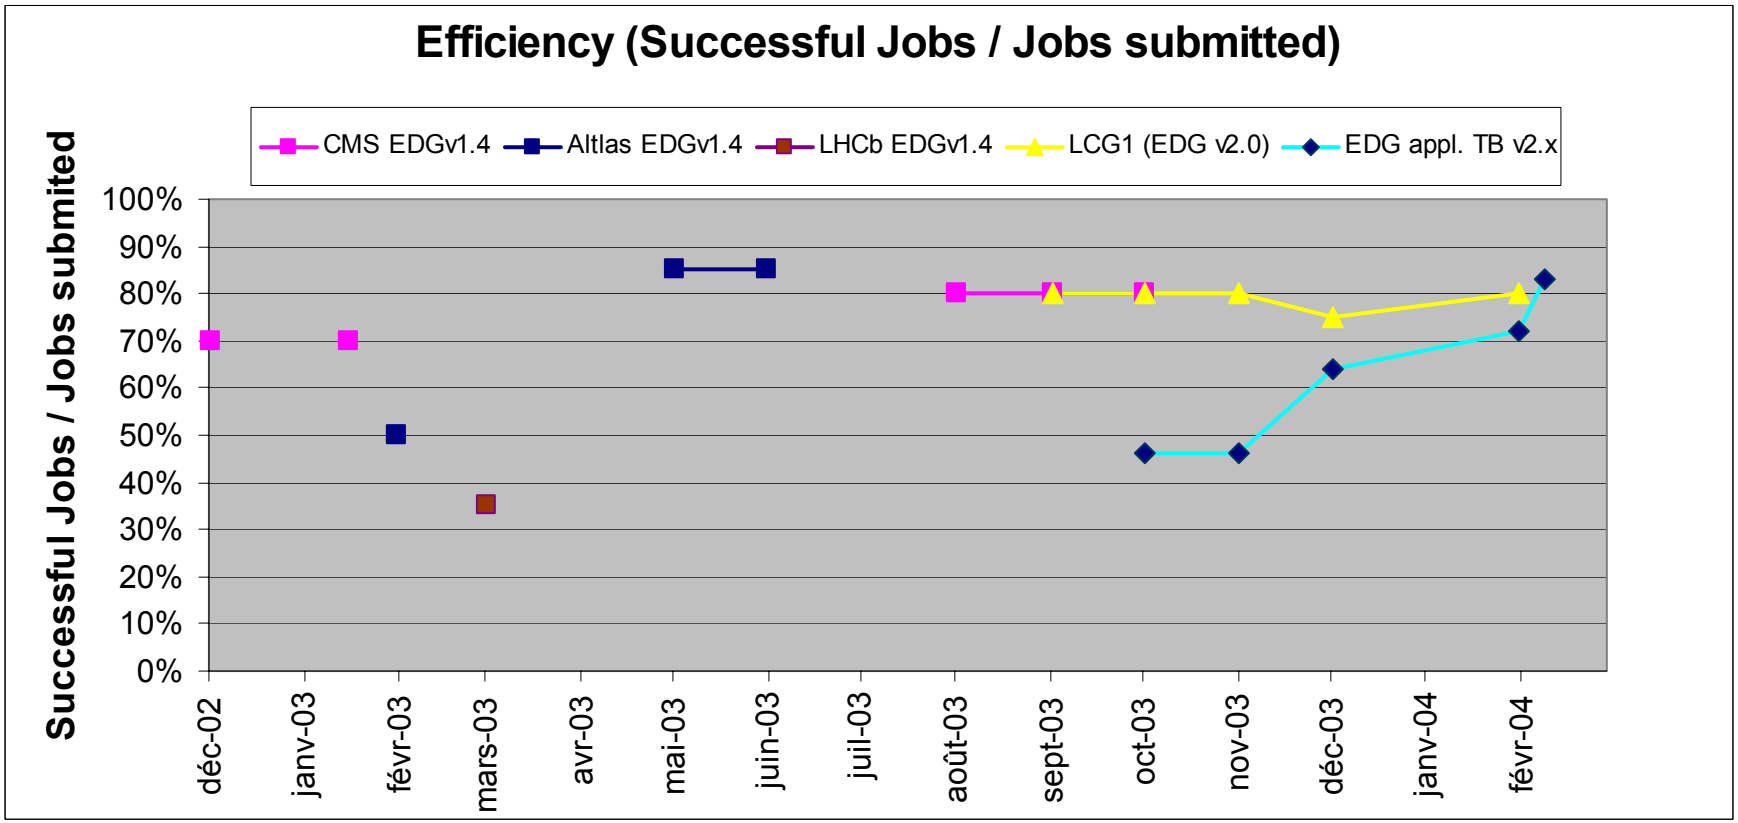

Conclusion: prototype performance, but with quality assurance mechanisms built-in

What is the Grid? Hour Glass

I. Experiment Layer

e.g. Portals

II. Application Middleware

e.g. Metadata

III. Grid Middleware

e.g. Information Services

IV. Facilities and Fabrics

e.g. Storage Services

Application Middleware Development

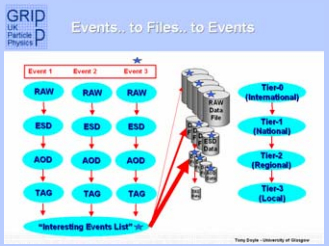

LHCb

BaBar

Job submission

Today	Tomorrow
• Locate data at sites • Submit jobs remotely, with automatic authorisation • Retrieve output	• Specify data • Submit jobs – sent automatically to best site • Output sent back

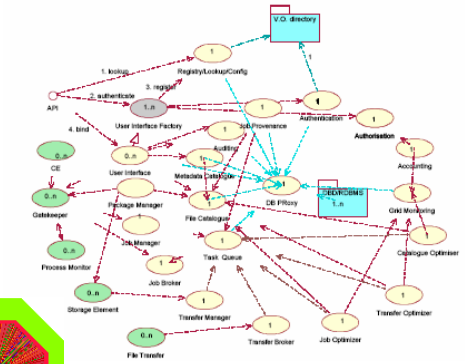

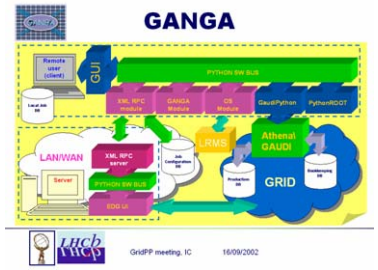

Data GRID

LCG

ALICE

UKQCD collaboration

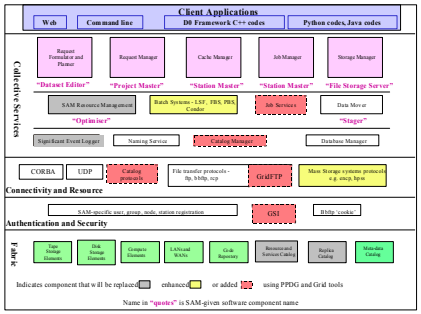

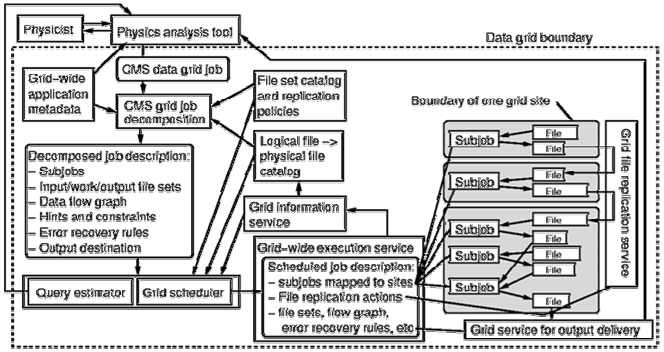

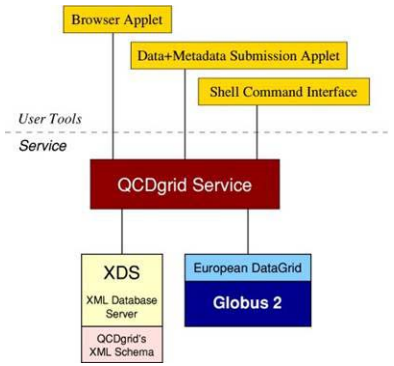

What is the Grid?
4. Standard 5. Evolving
Grid Middleware Development

Grid Data Management

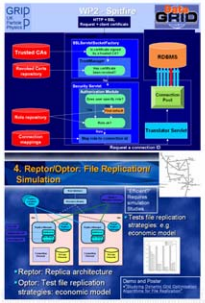

- Secure access to metadata
  - metadata: where are the files on the grid?
  - database client interface
  - "grid service" using standard web services
  - develop with UK e-science programme
- Optimised file replication
  - simulations required
  - economic models using CPU, disk, network inputs

Tony Doyle - University of Glasgow

Network Monitoring

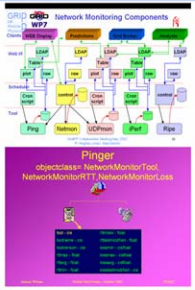

- Network monitoring essential
  - user writes sensor code to collect or measure networking information
- Format of metadata requires (International) agreement
  - emerging standards

Tony Doyle - University of Glasgow

Configuration Management

- software updates and maintenance
- Need to configure every component automatically

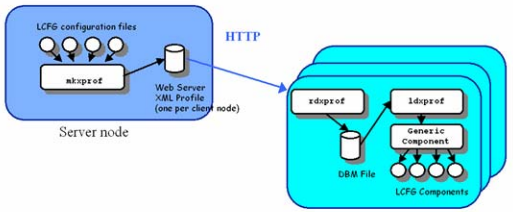

Tony Doyle - University of Glasgow

Information Service

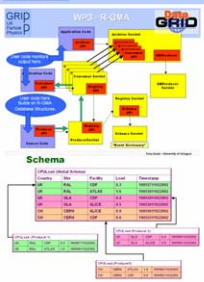

- Information service essential
  - user writes sensor code to collect or measure information
  - information is stored in buffers of appropriate length, duration
  - efficient access required for all metadata
  - relational model natural for e.g. network monitoring information
- Format of metadata requires (International) agreement
  - diplomacy required

Tony Doyle - University of Glasgow

Security

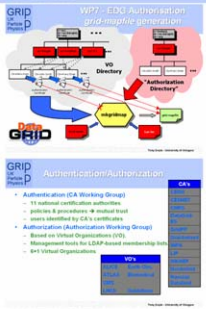

- User authentication required
  - who am I?
- User authorisation required
  - What experiment am I in?
  - What am I allowed to do?
- Automated tools
- International trust relationships required
  - Matrix of trust relationships
  - Agreed arrangements when security is compromised = certificate revocation

Tony Doyle - University of Glasgow

Storage Interfaces

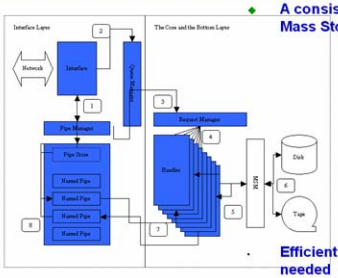

- A consistent interface to Mass Storage Systems
  - MSS
    - Castor
    - HPSS
    - RAID arrays
    - SRM
    - DMF
    - Enstore
  - Interfaces
    - GridFTP
    - GridRFIO
    - /grid
    - OGSA
- Efficient access standard needed

Tony Doyle - University of Glasgow

The Software Process: Extreme Programming

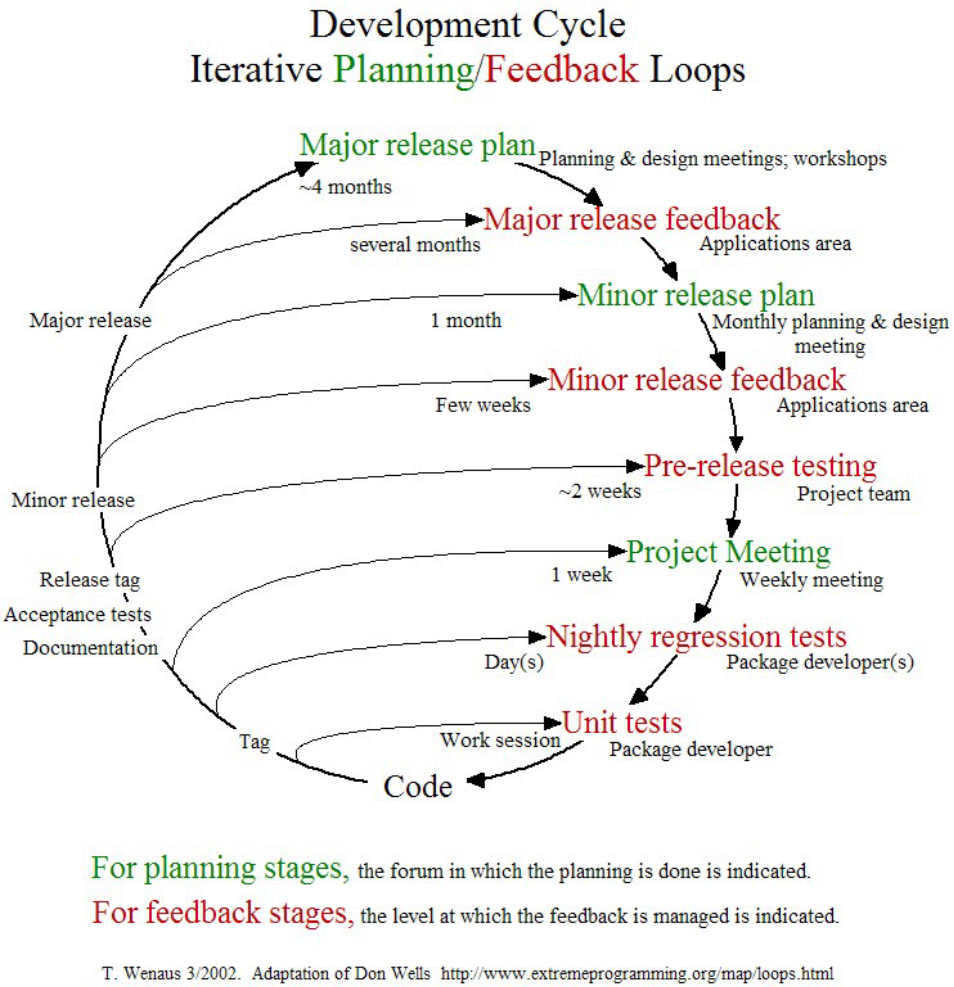

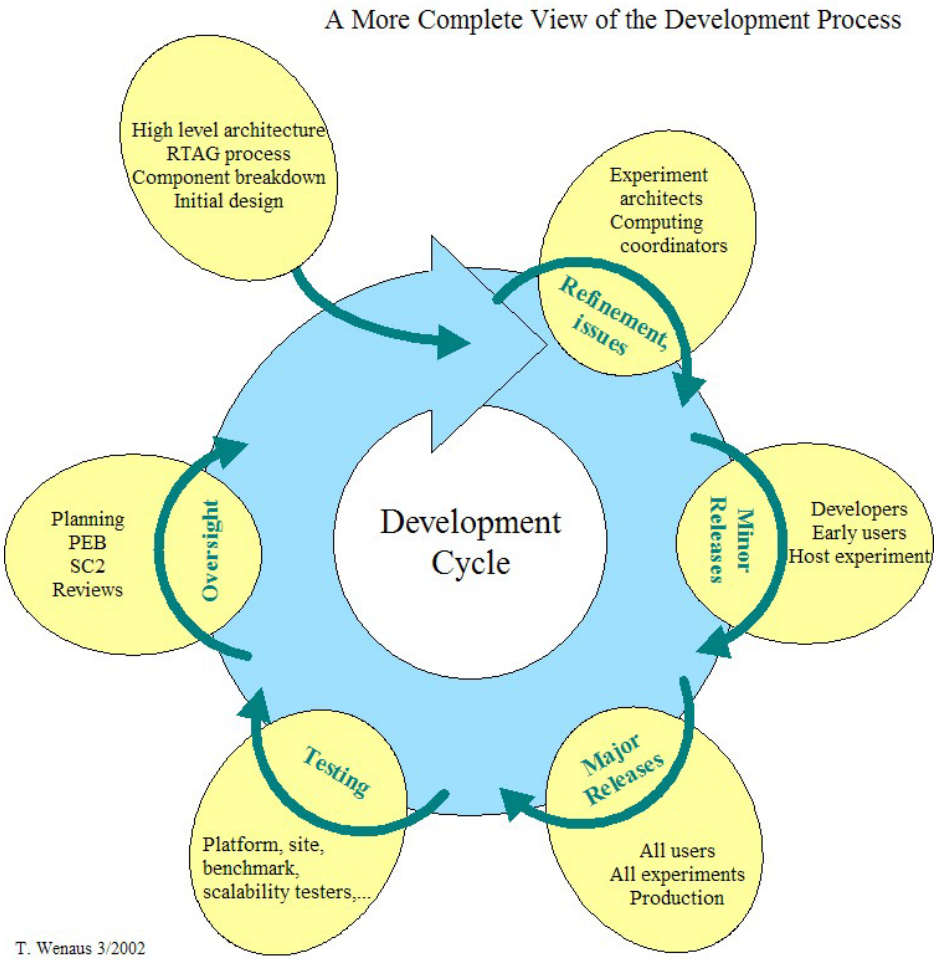

How Is the process applied?
Middleware Validation: From Testbed to Production

GRID UK Particle Physics

Build System | **Development Testbed ~15CPU** | **Certification Testbed ~40CPU** | **Application Testbed ~1000CPU**

Unit Test

add unit tested code to repository

WPs

Build

Run nightly build & auto. tests

Build system

Tagged package

Integration

Individual WP tests

Integration Team

Overall release tests

Tagged Releases

Tagged release selected for certification

Certification

Grid certification

Test Group

Application Certification

Apps. Representatives

Certified Releases

Certified release selected for deployment

Production

Certified public release for use by apps.

Users

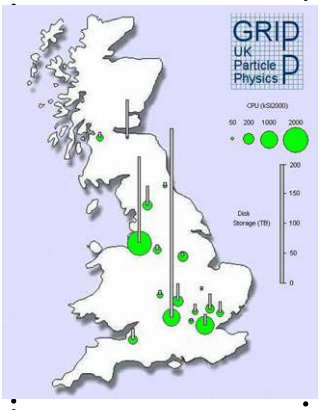

24x7

Fix problems

Problem reports

DataGrid Release Milestones

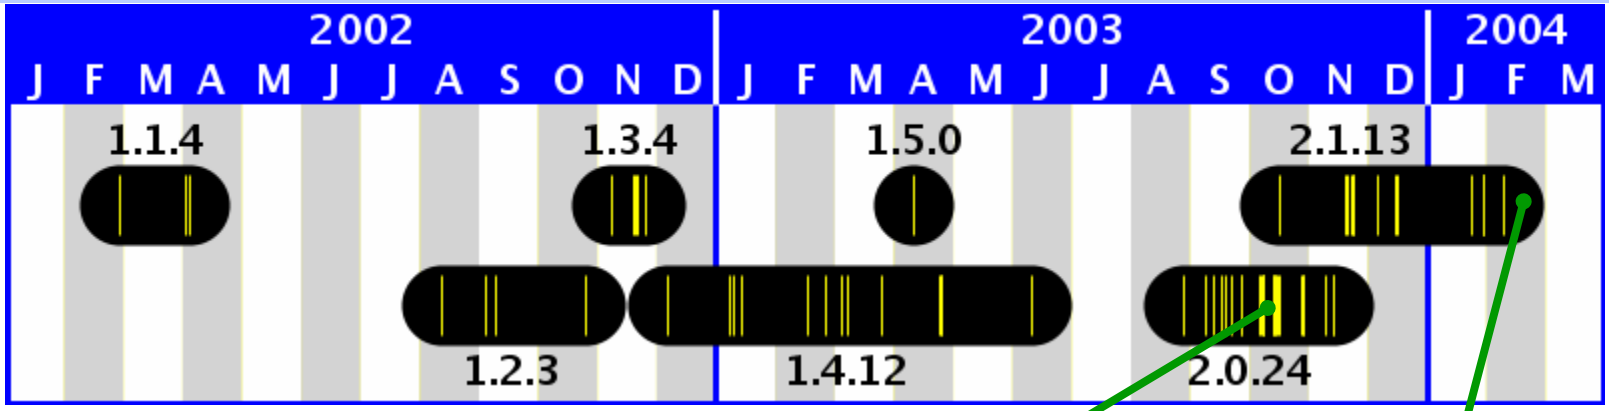

Evaluations (2.0.12)

- **Features (2.0.12)**
  - R-GMA replaced MDS
  - Refactored workload mgt.
  - Interactive, MPI, chkpt. jobs
  - Replica Location Service
  - Web Service SE

EU Review (2.1.13)

- **Features (2.1.13) [0.5Mloc]**
  - Reasonable stability, reliability
  - VOMS incorporated
  - Bug fixes for all services.

- **Stabilisation time on application testbed typically a few months**

Testbed: UK Tier-2 Centres

NorthGrid
Daresbury, Lancaster, Liverpool, Manchester, Sheffield

SouthGrid
Birmingham, Bristol, Cambridge, Oxford, RAL PPD

ScotGrid
Durham, Edinburgh, Glasgow

LondonGrid
Brunel, Imperial, QMUL, RHUL, UCL

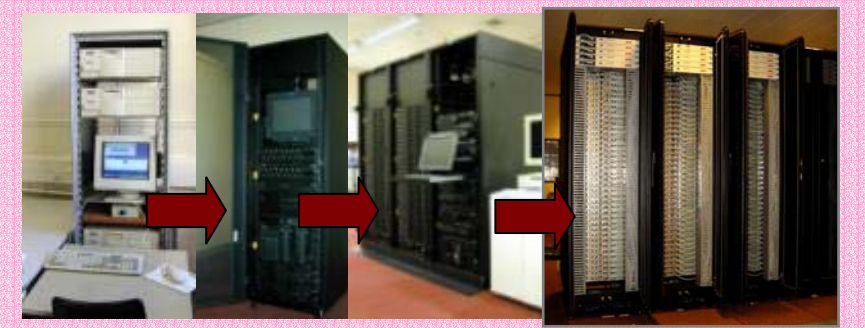

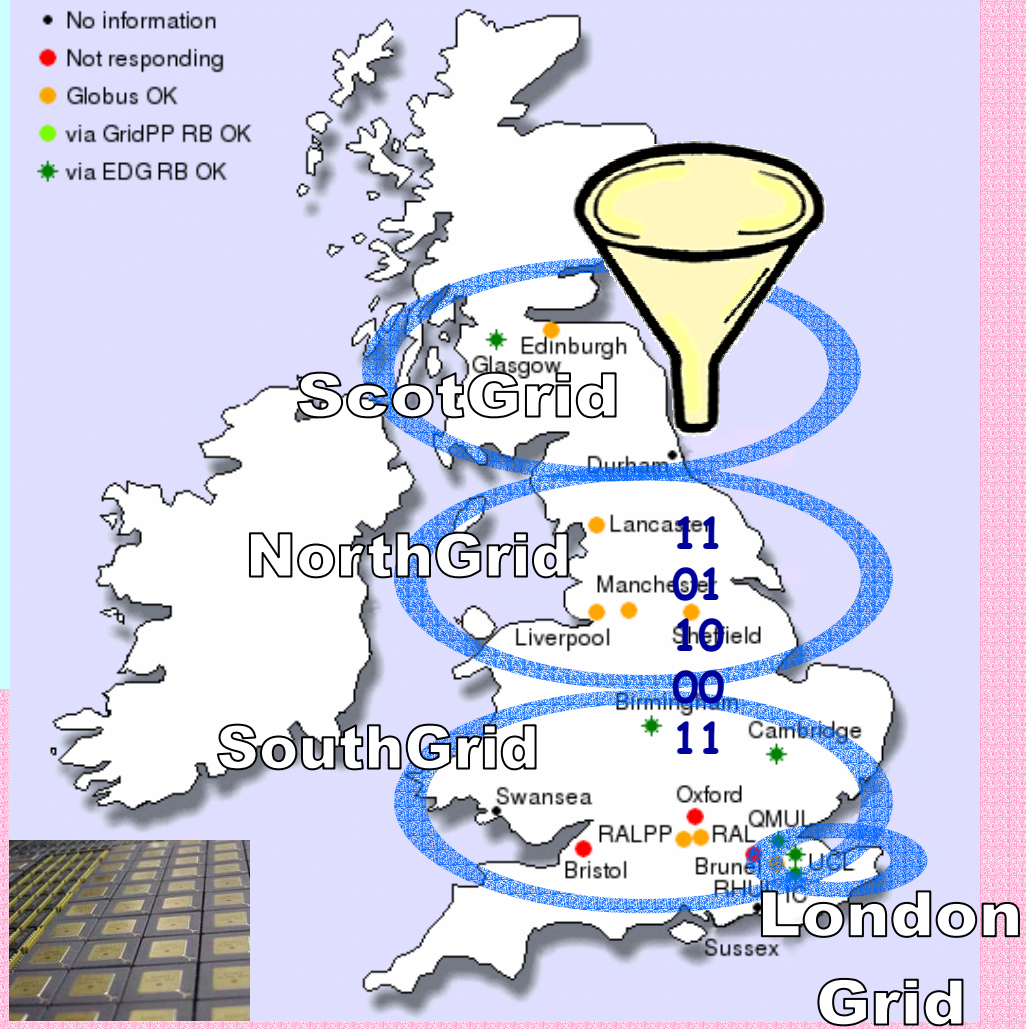

Tier-1/A Services

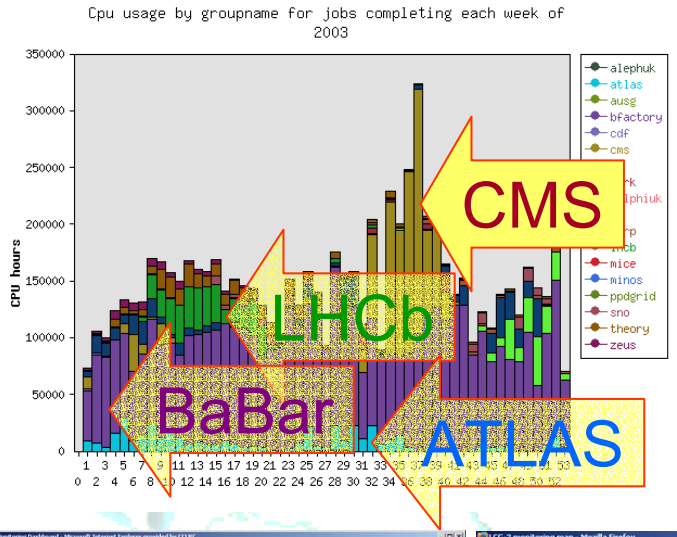

- **High quality data services**
- **National and International Role**
- **UK focus for International Grid development**

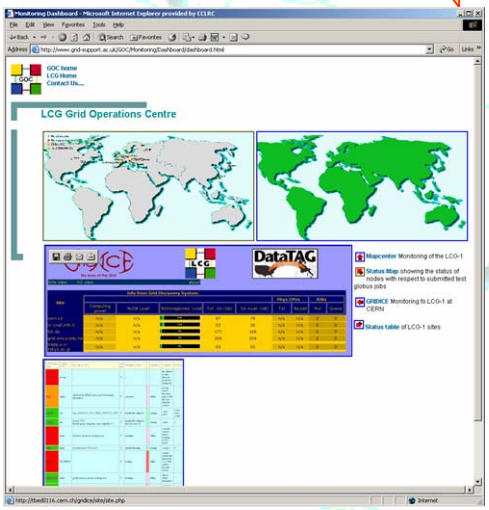

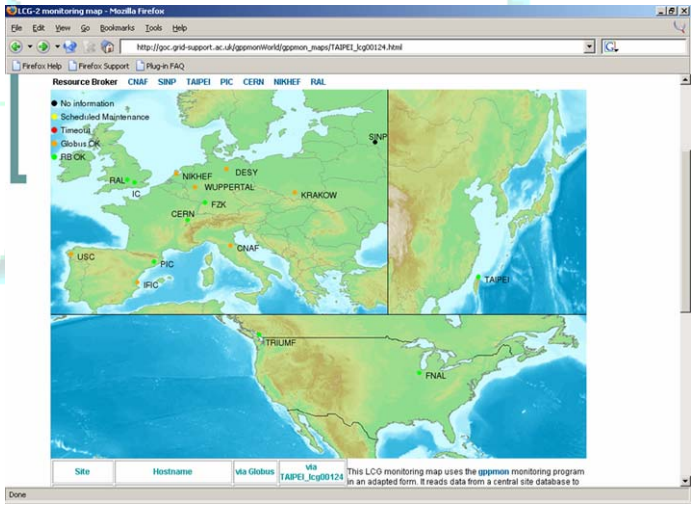

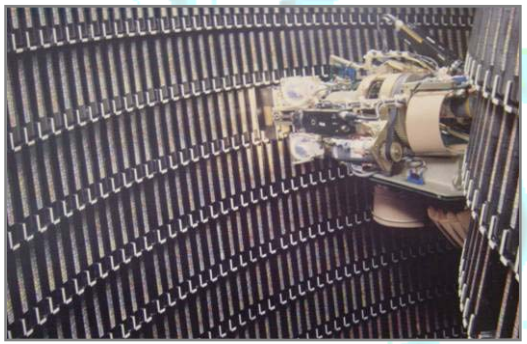

The UK Testbed

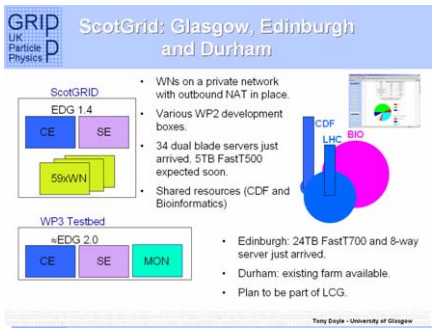

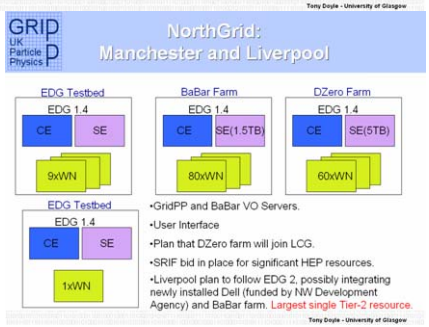

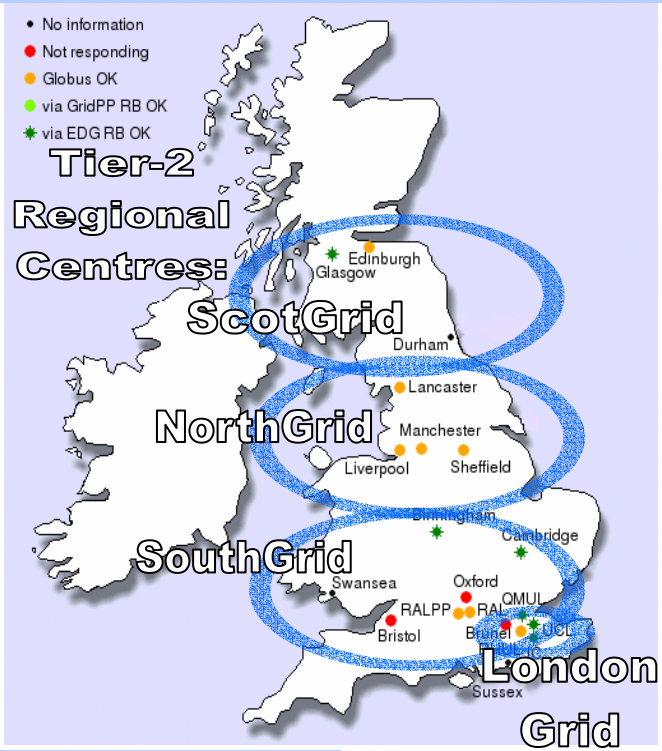

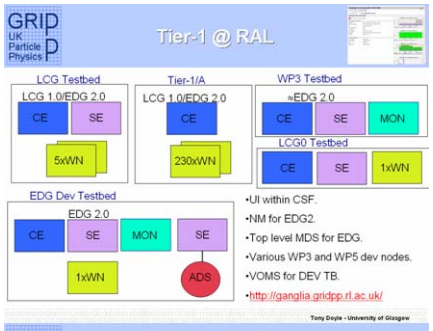

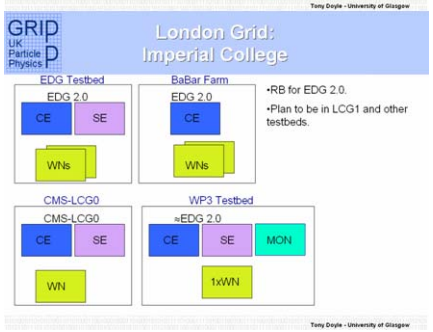

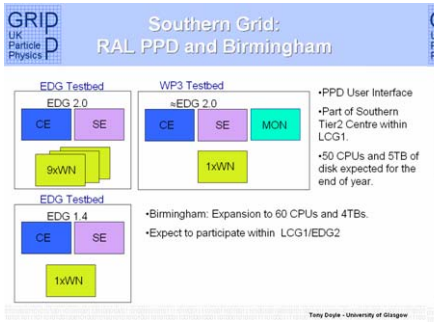

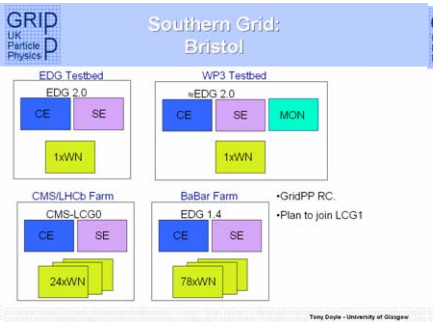

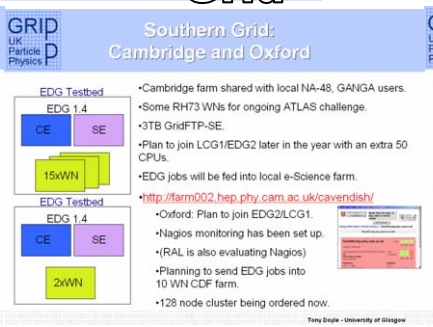

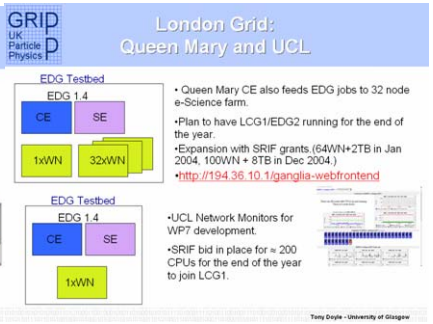

What is GridPP?

- A team that have built a working prototype grid
  - hundreds of computers across 16 UK institutions
  - Allows scientists to access data and processing power seamlessly, wherever they are:
  - No need for accounts at many different universities
  - No need for lots of passwords
  - ‘Middleware’ detects where the computing power and data is available and deals with security
- A project where 76% of the 190 tasks for the first three years have been completed
- Working with other disciplines
  - e.g. UK e-Science (Grid deployment)
    medical researchers (Grid security)
    BT (Grid Info. Services, networking)

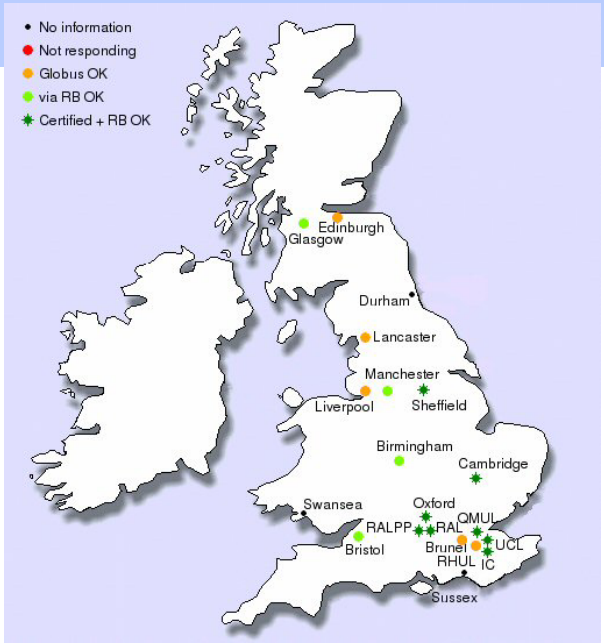

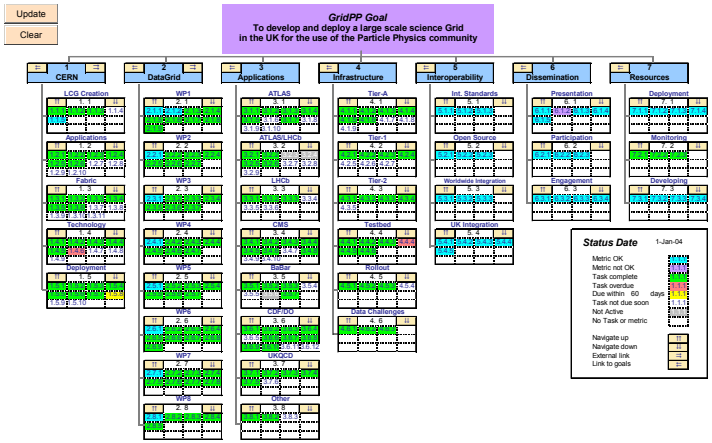

What is "The Grid" Is GridPP a Grid? Anyway?

http://www-fp.mcs.anl.gov/~foster/Articles/WhatIsTheGrid.pdf

1. Coordinates resources that are not subject to centralized control
2. … using standard, open, general-purpose protocols and interfaces
3. … to deliver nontrivial qualities of service

1. YES. This is why development and maintenance of a UK-EU-US testbed is important
2. YES... Globus/CondorG/EDG meet this requirement. Common experiment application layers are also important here.
3. NO(T YET)… Experiments define whether this is true - currently only ~100,000 jobs submitted via the testbed c.f. internal component tests of up 10,000 jobs per day. Next step: LCG-2 deployment outcome… this year

How Does the Work?

0. Web User Interface… or CLI

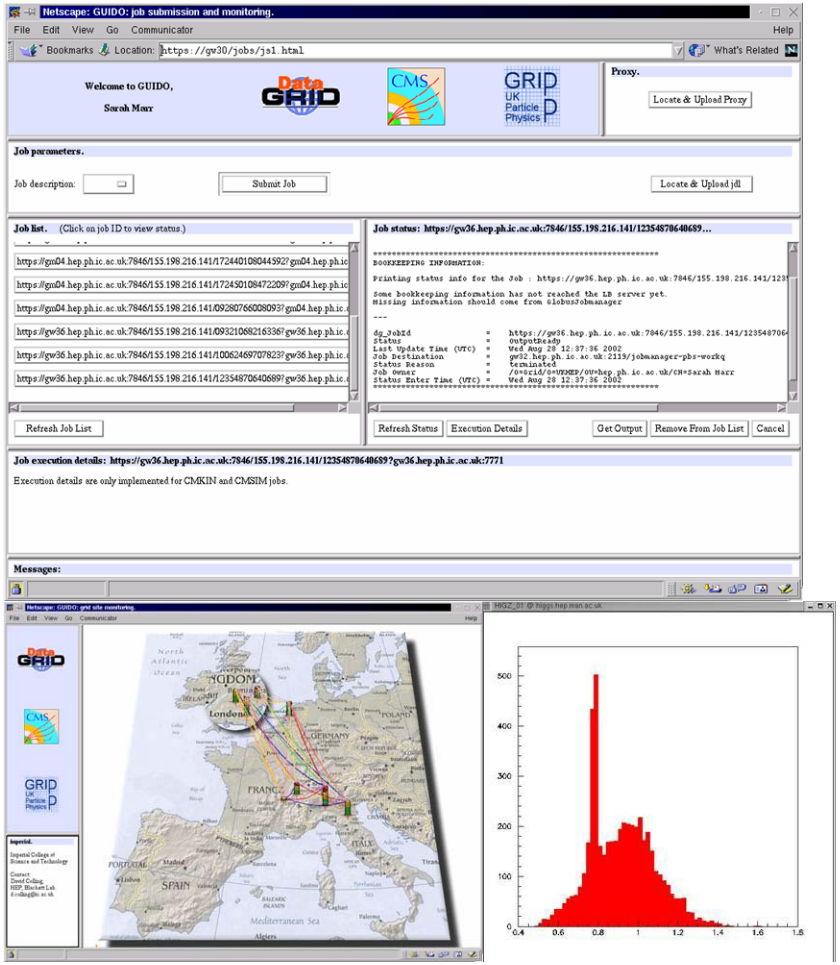

1. Authentication

grid-proxy-init

2. Job submission

edg-job-submit

3. Monitoring and control

edg-job-status
edg-job-cancel
edg-job-get-output

4. Data publication and replication

globus-url-copy, RLS

5. Resource scheduling – use of Mass Storage Systems

JDL, sandboxes, storage elements

From Web to Grid… Who do you trust?

No-one?

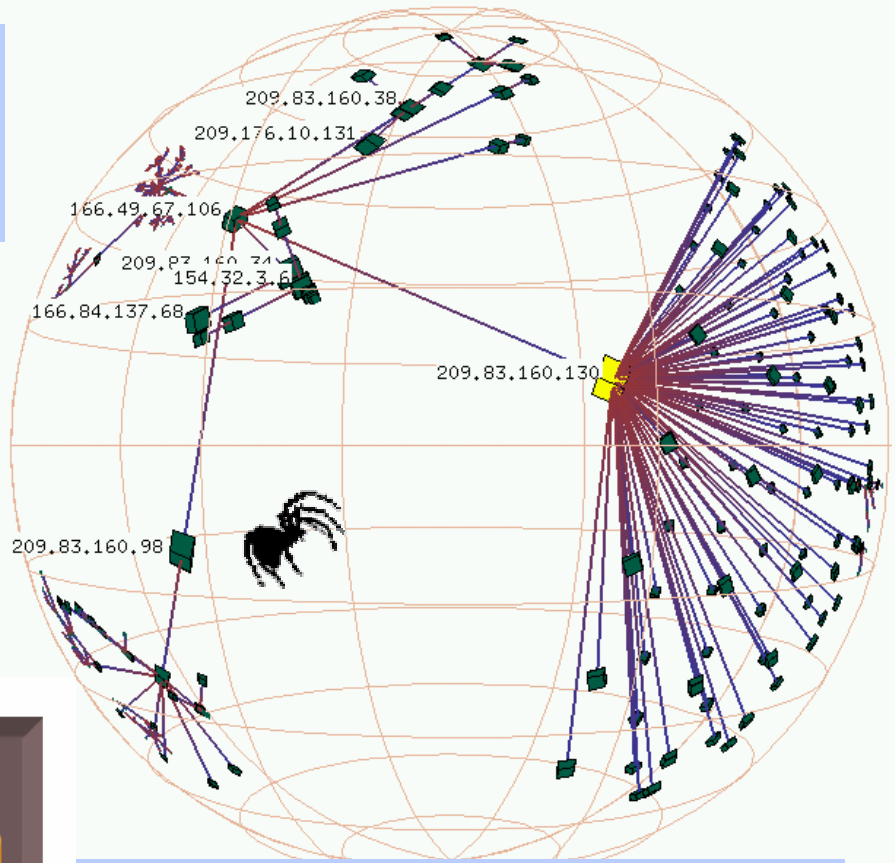

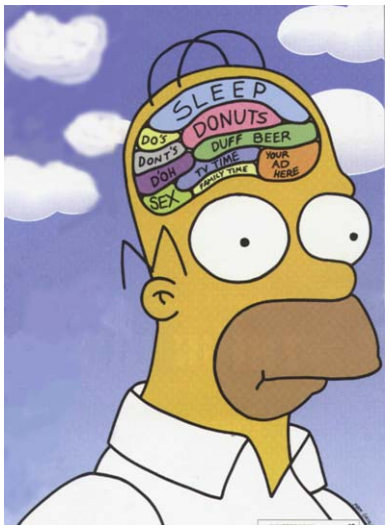

It depends on what you want… (assume its scientific collaboration)

How do I Authorize?

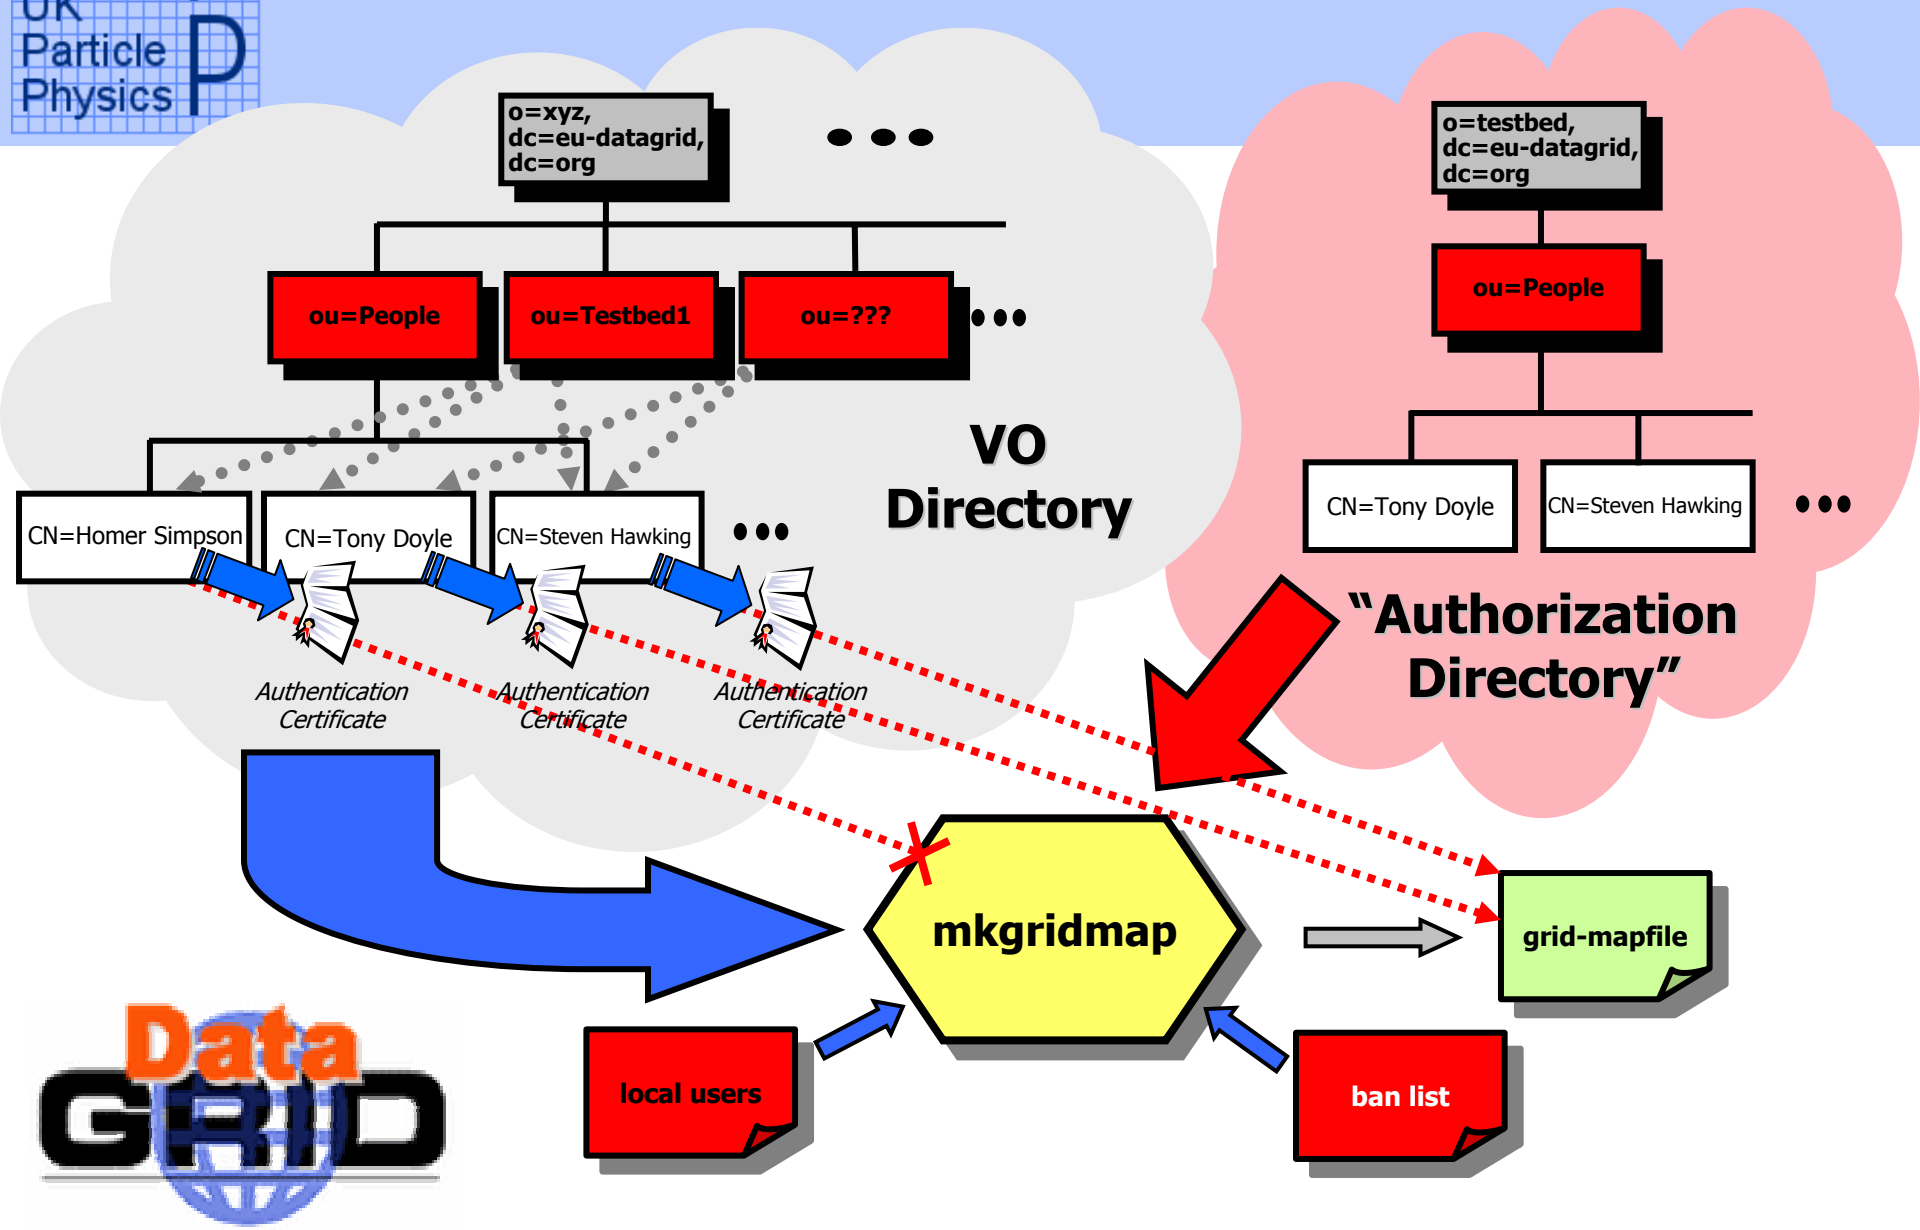

UK Certificate Authority and Virtual Organisation membership

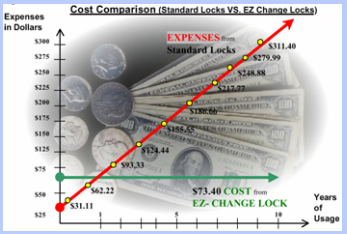

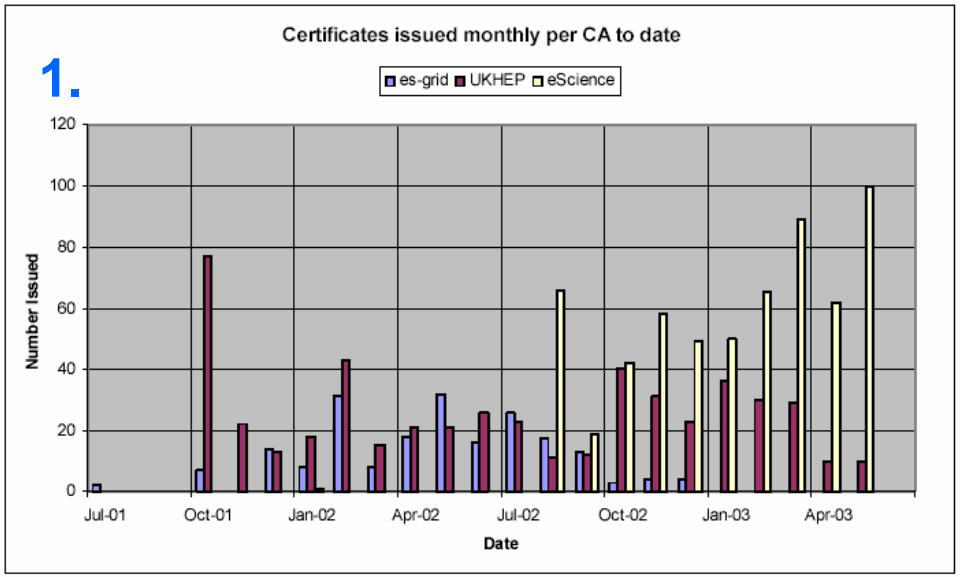

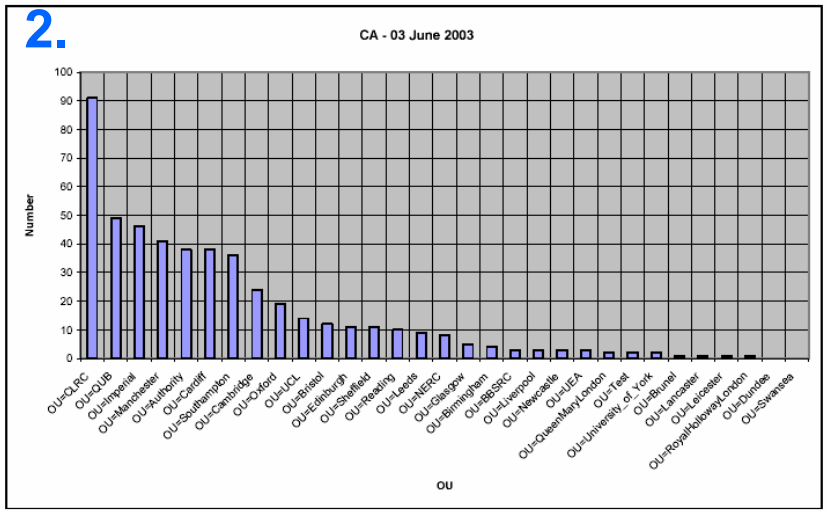

1. **UK e-Science Certificate Authority now used in application testbed**

2. **PP "users" engaged from many institutes**

3. **UK participating in 6 ex 9 EDG Virtual Organisations**

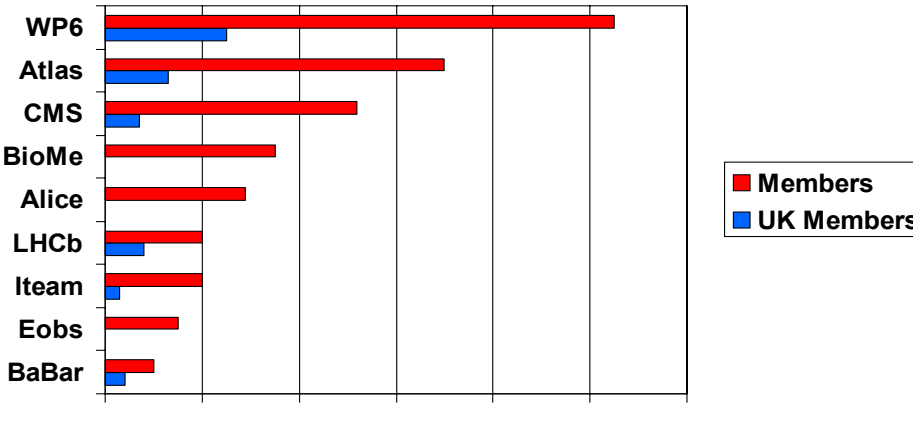

The Challenges Ahead I: Event Selection

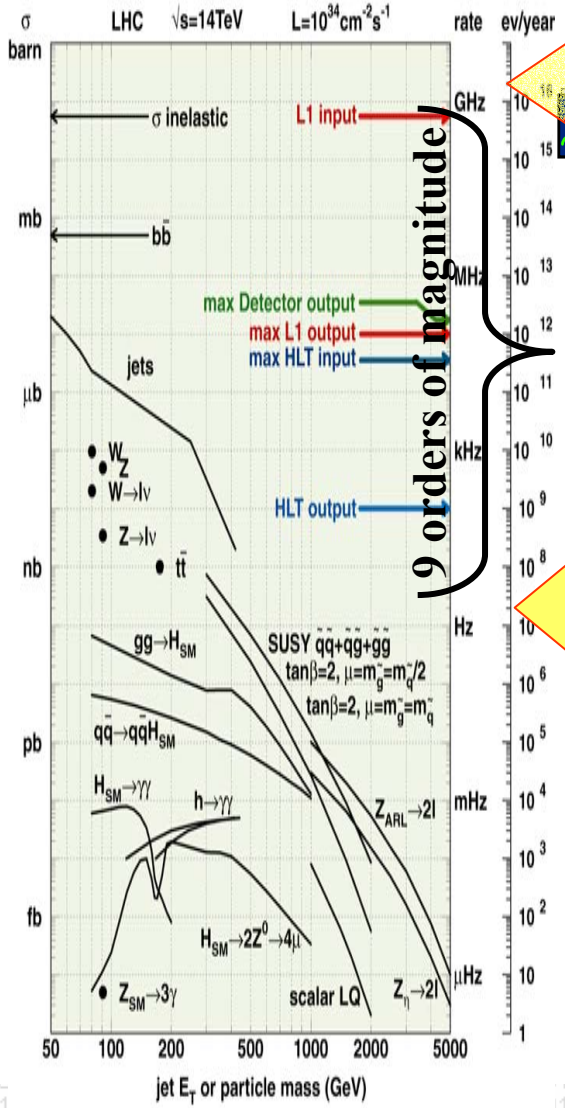

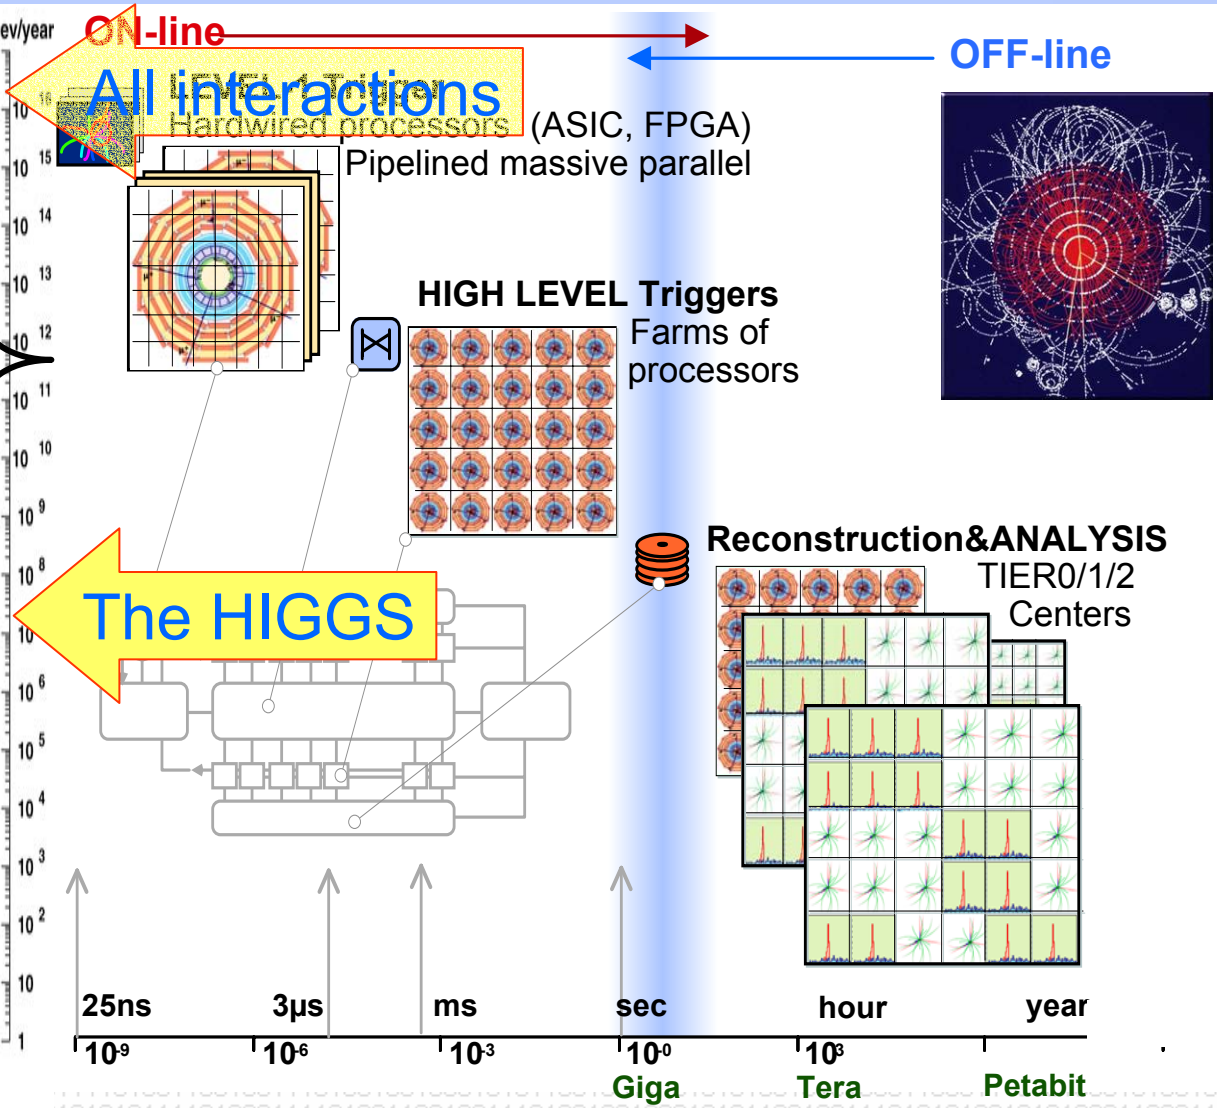

The Challenges Ahead II: Complexity

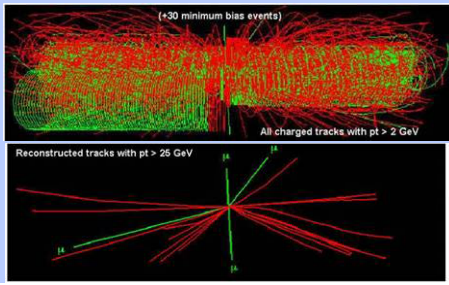

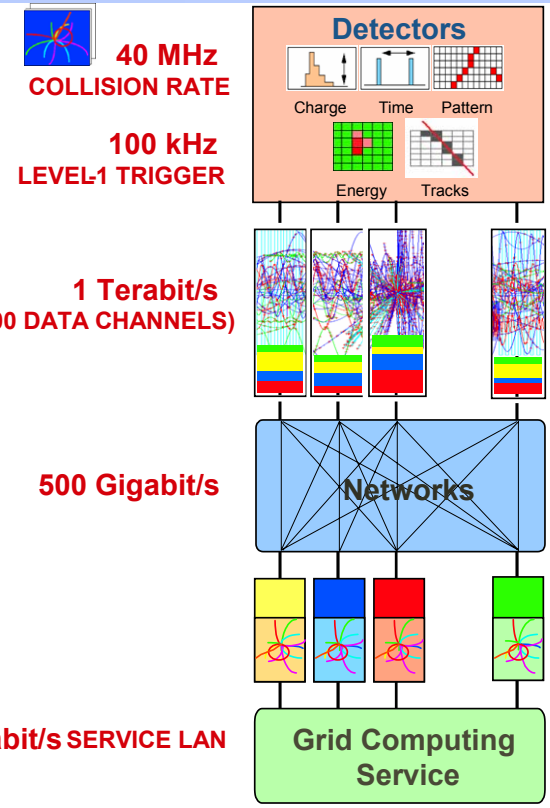

- **Many events**
  - ~10^9 events/experiment/year
  - >~1 MB/event raw data
  - several passes required

→**Worldwide Grid computing requirement (2007):**

~300 TeraIPS

(100,000 of today's fastest processors connected via a Grid)

ATLAS Barrel Inner Detector
$H \to b\bar{b}$

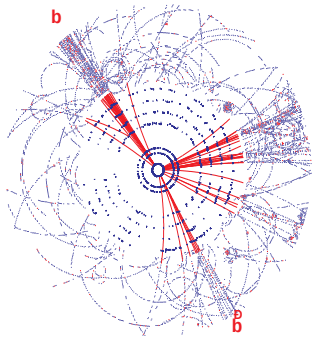

Understand/interpret data via numerically intensive simulations:

- e.g. ATLAS Monte Carlo ($gg \Rightarrow H \Rightarrow b\bar{b}$)

182 sec/3.5 MB event on 1 GHz linux box

The Challenges Ahead III: Experiment Requirements: UK only

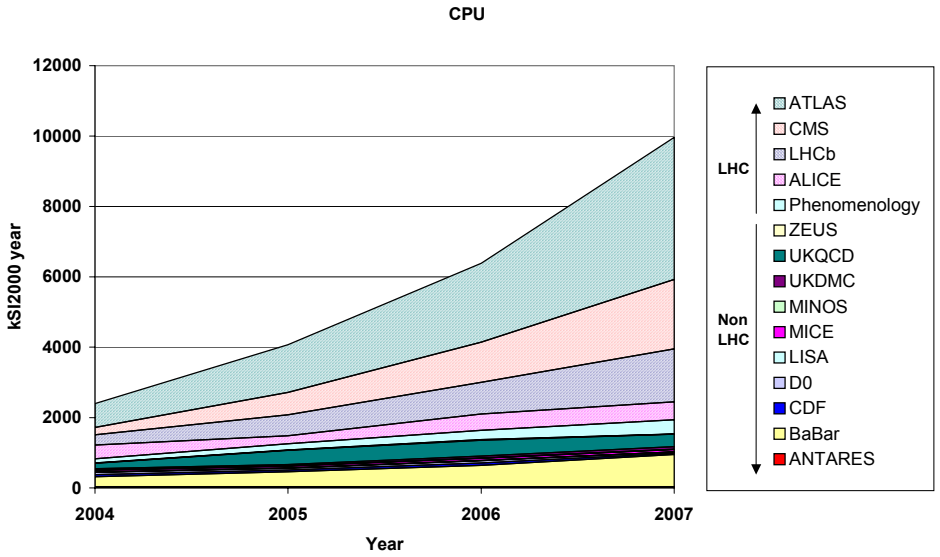

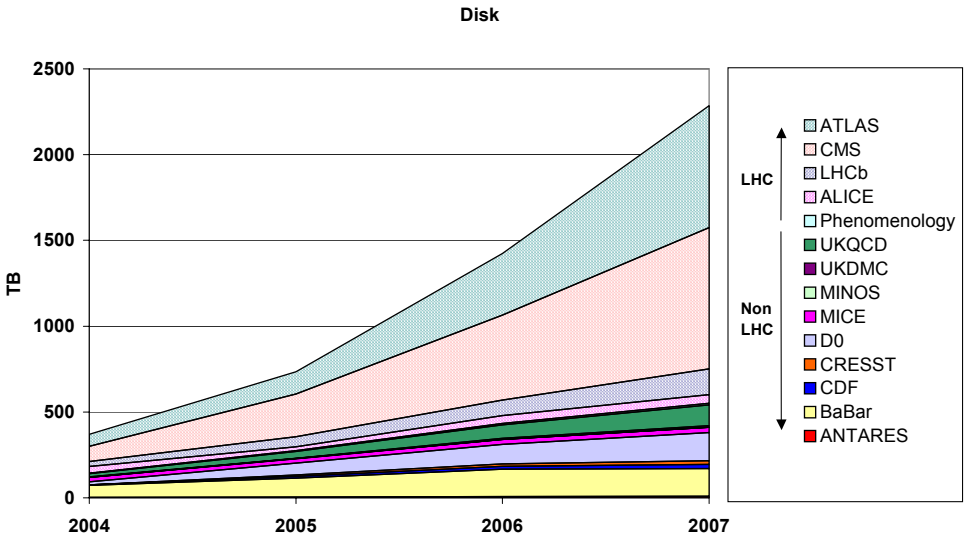

Total Requirement:

Year	2004	2005	2006	2007
CPU [kSI2000]	2395	4066	6380	9965
Disk [TB]	369	735	1424	2285
Tape [TB]	376	752	1542	2623

In International Context - Q2 2004 LCG Resources:

Country	Q2 2004		
	CPU [KSI2000]	Disk space [TB]	Tape space [TB]
Austria	7	1	0
Canada	373	25	30
CERN	700	50	1000
Czech Republic	17	6	5
France	140	27	180
Germany	325	45	77
Italy	945	114	200
Japan	127	34	0
Netherlands	210	10	100
Portugal	2	1	0
Russia	140	29	21
Spain	137	8	46
Switzerland	18	5	20
UK	2042	229	140

The Challenges Ahead IV: Work Group Computing

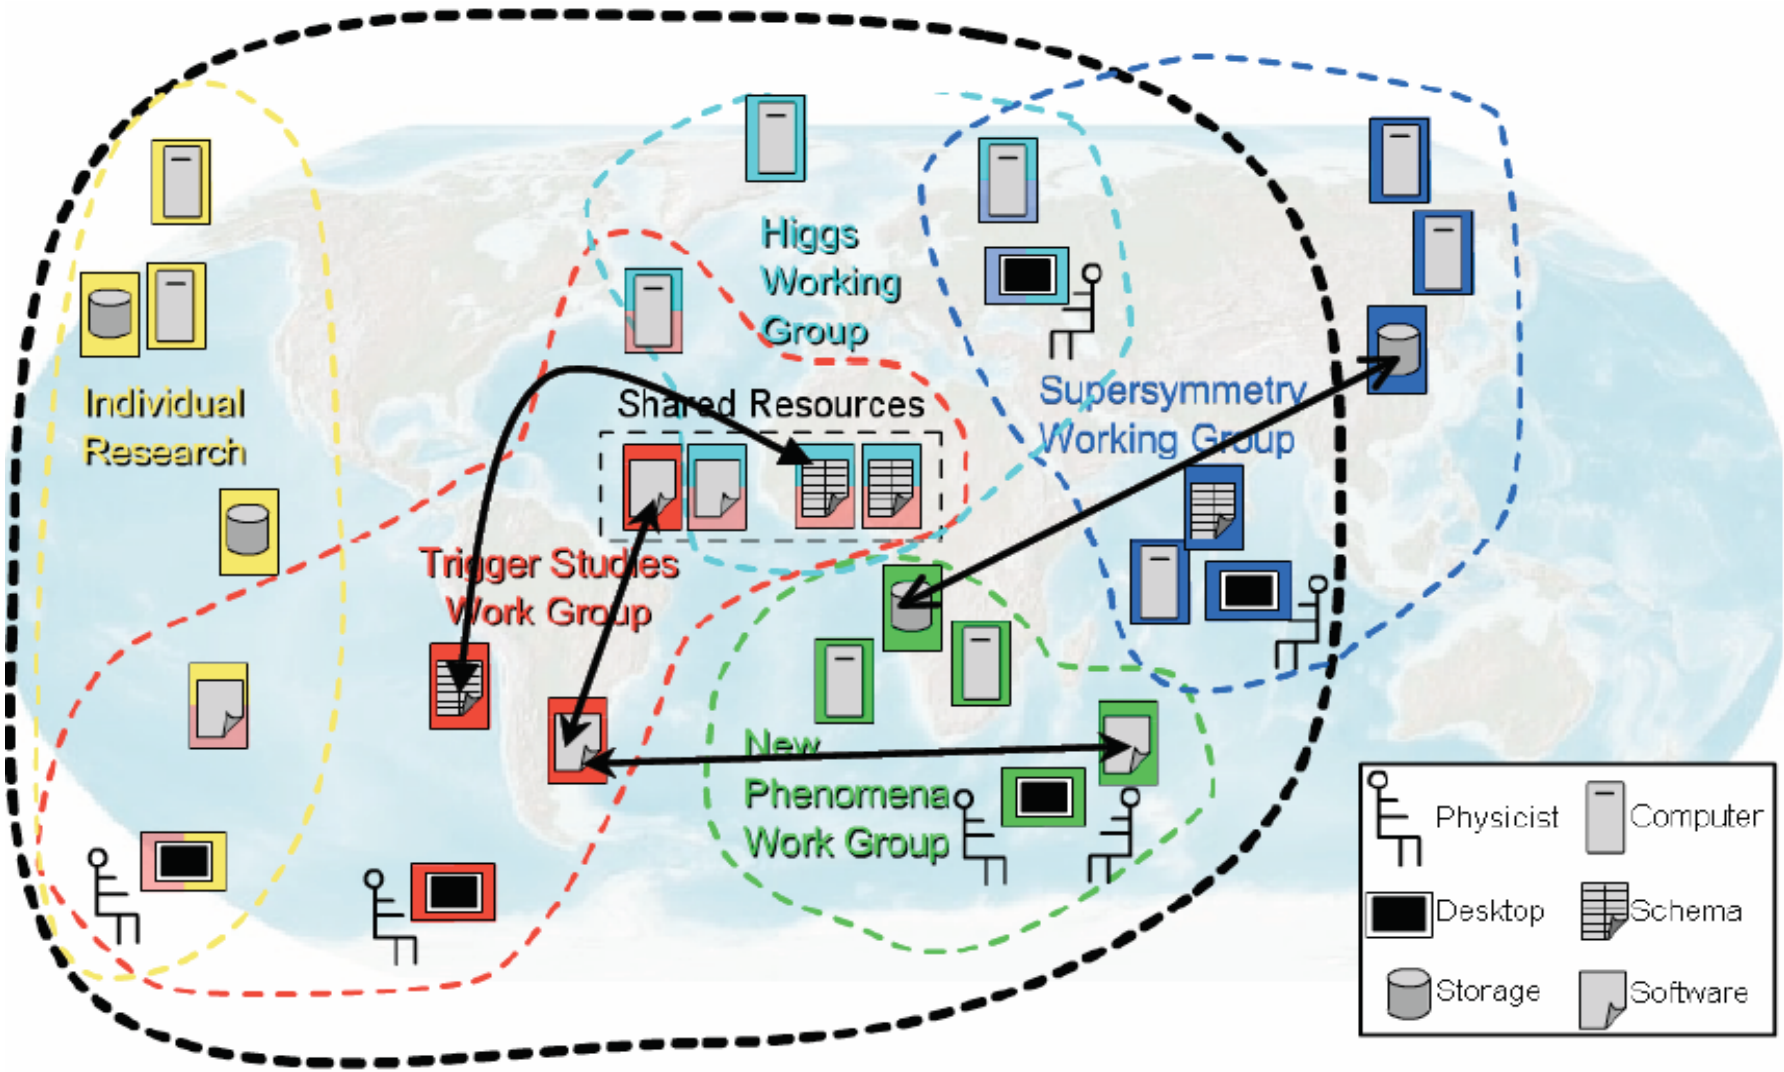

The Challenges Ahead V:
Events.. to Files.. to Events

- **VOMS-enhanced Grid certificates to access databases via metadata**
- **Non-Trivial..**

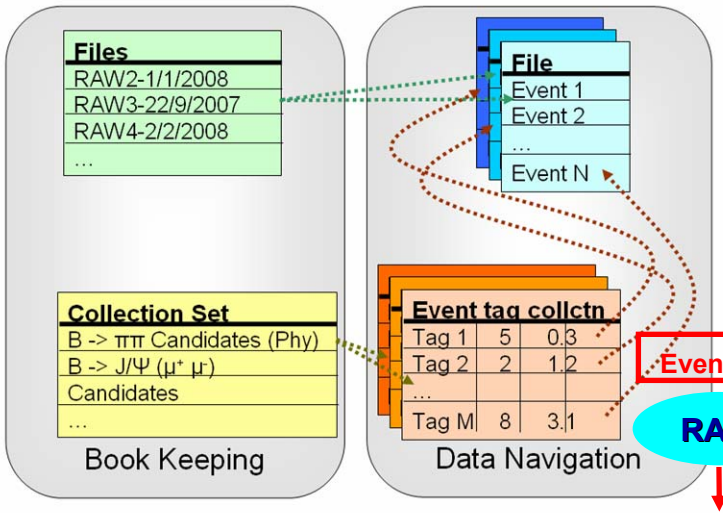

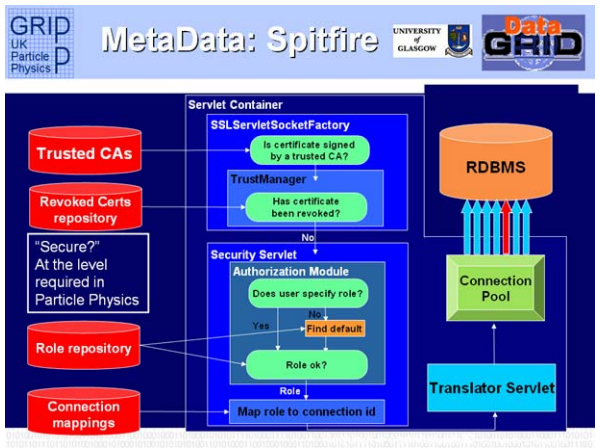

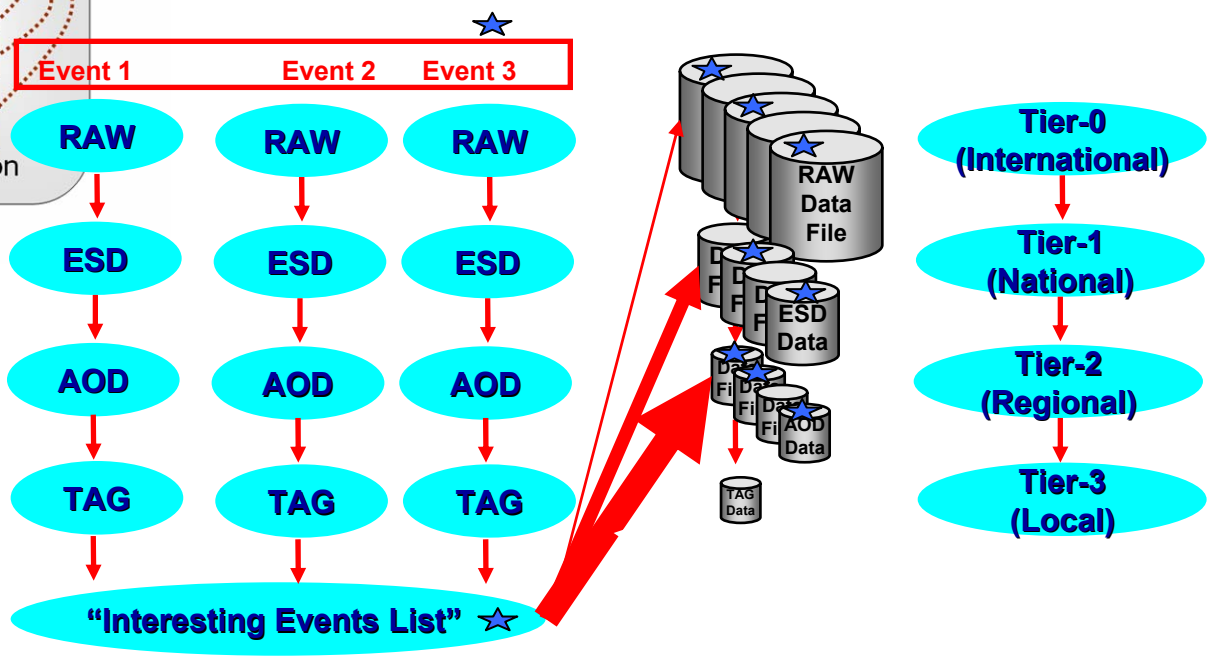

The Challenges Ahead VI: software distribution

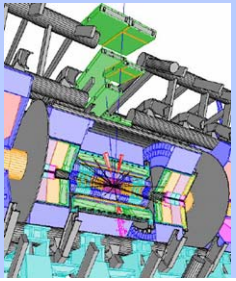

- **ATLAS Data Challenge (DC2) this year to validate world-wide computing model**
- **Packaging, distribution and installation: Scale: one release build takes 10 hours produces 2.5 GB of files**
- **Complexity: 500 packages, Mloc, 100s of developers and 1000s of users**
  - **ATLAS collaboration is widely distributed: 140 institutes, all wanting to use the software**
  - **needs 'push-button' easy installation..**

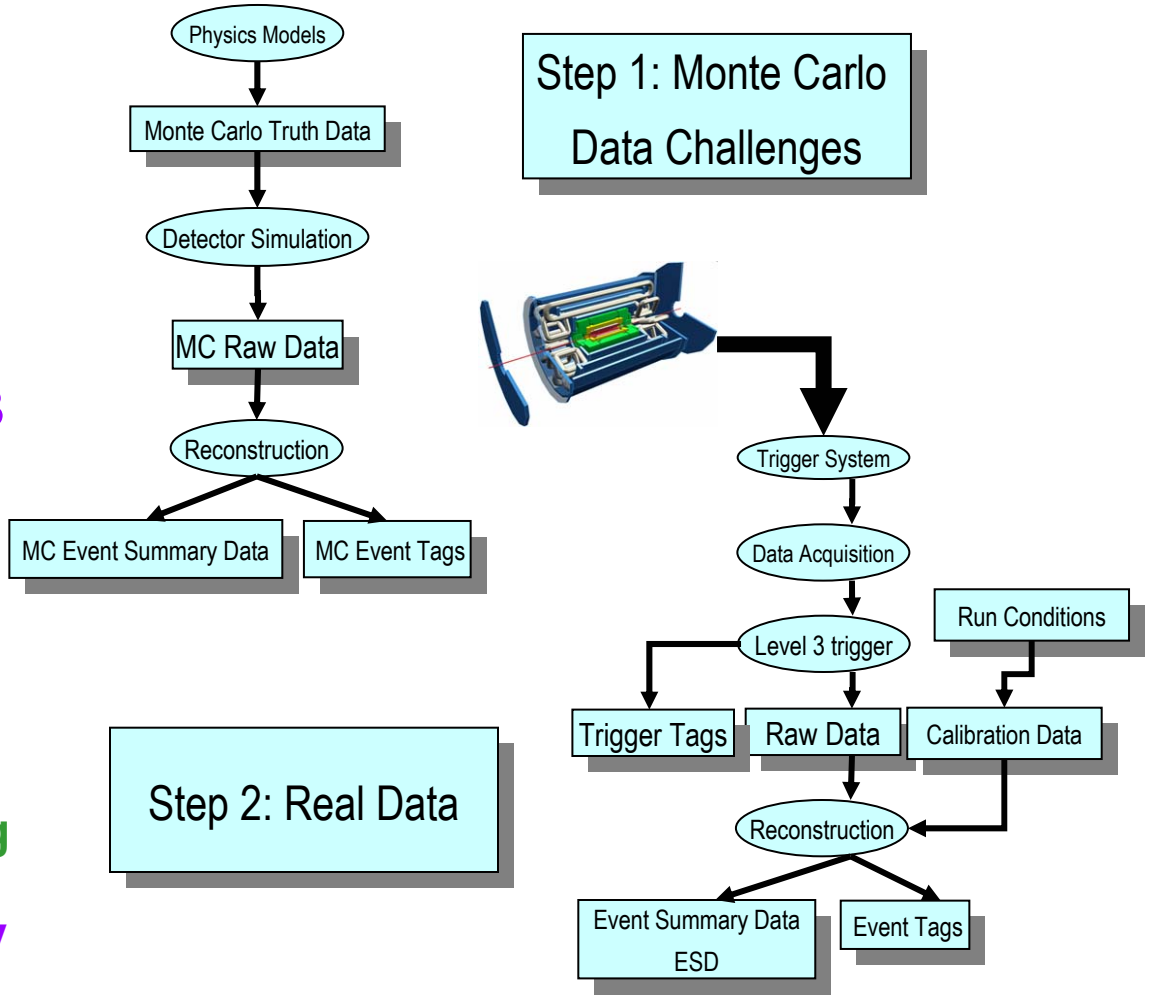

The Challenges Ahead VII: distributed analysis

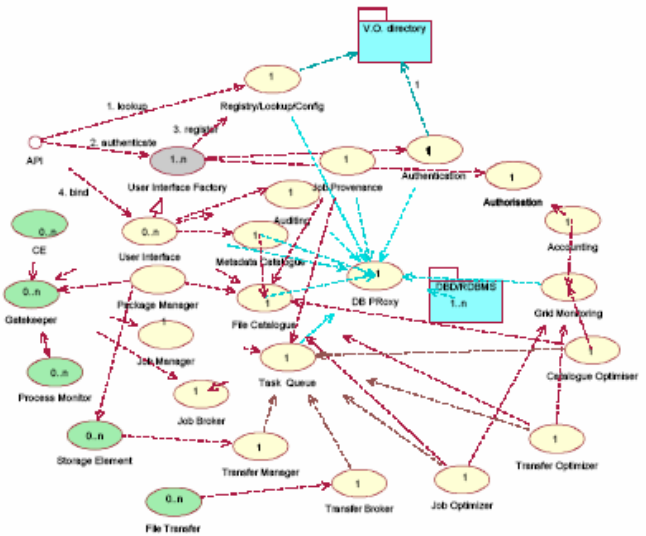

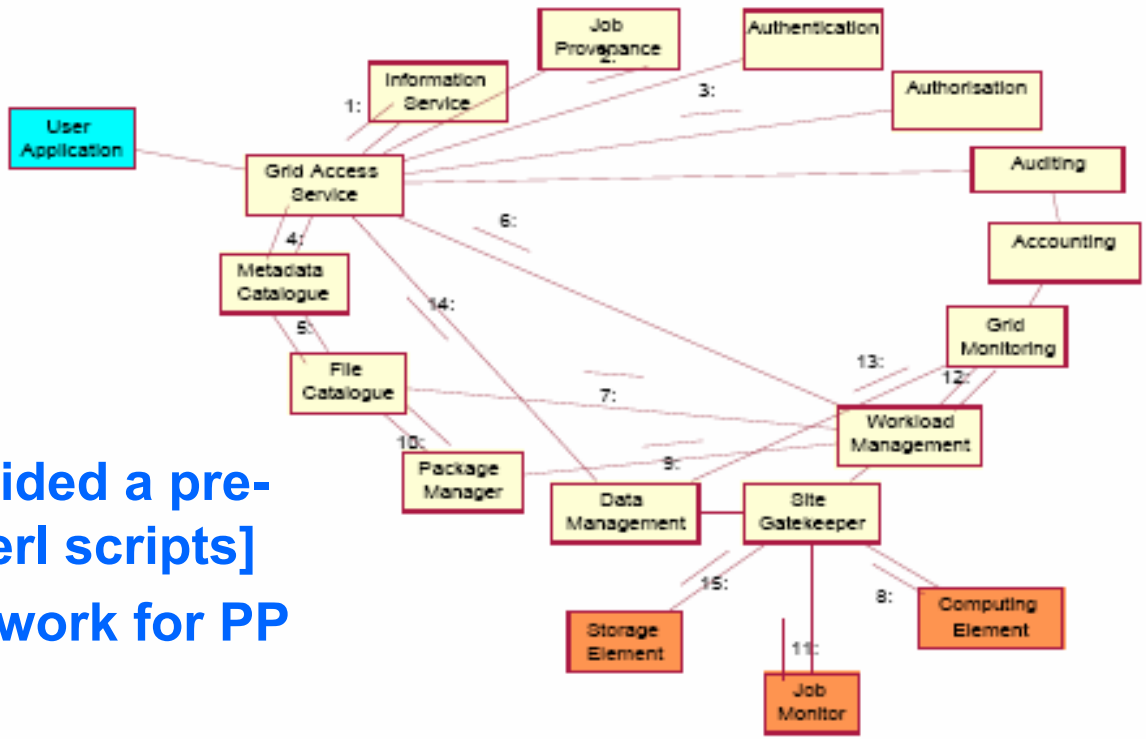

1. AliEn (ALICE Grid) provided a pre-Grid implementation [Perl scripts]
2. ARDA provides a framework for PP application middleware

Philosophy of the Grid?

- "Everything is becoming, nothing is." Plato
- "Common sense is the best distributed commodity in the world. For every (wo)man is convinced (s)he is well supplied with it." Descartes
- "The superfluous is very necessary" Voltaire
- "Only daring speculation can lead us further, and not accumulation of facts." Einstein
- "The real, then, is that which, sooner or later, information and reasoning would finally result in." C. S. Pierce
- "The philosophers have only interpreted the world in various ways; the point is to change it." Marx
- (some of) these may be relevant to your view of "The Grid"…

Historical Perspective

- I wrote in 1990 a program called "WorlDwidEweb", a point and click hypertext editor which ran on the "NeXT" machine. This, together with the first Web server, I released to the High Energy Physics community at first, and to the hypertext and NeXT communities in the summer of 1991.
- Tim Berners-Lee
- The first three years were a phase of persuasion, aided by my colleague and first convert Robert Cailliau, to get the Web adopted…
- We needed seed servers to provide incentive and examples, and all over the world inspired people put up all kinds of things…
- Between the summers of 1991 and 1994, the load on the first Web server ("info.cern.ch") rose steadily by a factor of 10 every year…

Next steps

- **From prototype to production**
  - UK particle physics grid equivalent to 20,000 1GHz personal computers by 2007
  - available for day-to-day use by particle physicists
  - web portal for other e-scientists
- **GridPP will support Enabling Grids for E-science in Europe (EGEE)**
  - to integrate national and international grids, and grids from different scientific disciplines
  - particle physics is a pilot project

Enabling Grids for E-science in Europe

What is The Grid?

The Grid is:

not hype, but surrounded by it

a working prototype running on testbed(s)…

about seamless discovery of PC resources around the world

using evolving standards for interoperation

the basis for particle physics computing in the 21st Century

not (yet) as transparent as end-users want it to be

Acknowledgements



Hayes member Todd Huffman /O=Grid/O=UKHEP/OU=physics.ox.ac.uk/CN=B. Todd Huffman member
CN=G N Patrick member Pete Gronbech /O=Grid/O=UKHEP/OU=physics.ox.ac.uk/CN=Pete Gronbech
member Nick Brook /O=Grid/O=UKHEP/OU=phy.bris.ac.uk/CN=Nick Brook member Marc Kelly /O=Grid/O=UKHEP/OU=phy.bris.ac.uk/CN=Marc Kelly
member Dave Newbold /O=Grid/O=UKHEP/OU=phy.bris.ac.uk/CN=Dave Newbold member Kate Mackay
/O=Grid/O=UKHEP/OU=phy.bris.ac.uk/CN=Catherine Mackay member Girish Patel /O=Grid/O=UKHEP/OU=ph.liv.ac.uk/CN=Girish D. Patel member David
Martin /O=Grid/O=UKHEP/OU=ph.gla.ac.uk/CN=David J. Martin member Peter Faulkner /O=Grid/O=UKHEP/OU=ph.bham.ac.uk/CN=Peter Faulkner
=Grid/O=UKHEP/OU=ph.bham.ac.uk/CN=David Smith
eve Traylen member Ruth Dixon del Tufo /O=Grid/O=U
OU=hepgrid.clrc.ac.uk/CN=Linda Cornwall member /O=

Grid Certificate Pioneers…



/O=Grid/O=UKHEP/OU=ph.qmw.ac.uk/CN=A.J.Martin member Steve Lloyd /O=Grid/O=UKHEP/OU=ph.qmw.ac.uk/CN=S.L.Lloyd admin John Gordon
/O=Grid/O=UKHEP/OU=hepgrid.clrc.ac.uk/CN=John Gordon member

Tony Doyle - University of Glasgow

Final Test at Grid School

Condor

C1) The condor_master daemon launches all other required Condor-G daemons, and restarts them in case of trouble.
a) true b) false

C2) Which one of the following is *not* a Condor-G command-line utility?
a) condor_q
b) condor_history
c) condor_login
d) condor_submit
e) condor_off

C3) The condor_submit_dag command submits a job to which Condor job universe?
a) scheduler
b) globus
c) standard
d) vanilla
e) dependency

C4) Condor-G can currently do which of the following? Check *all* that apply.
a) Monitor your jobs and keep you posted on their progress.
b) Replicate your job's data and update a replica catalog service.
c) Implement your policies for the execution order of your jobs.
d) Add fault tolerance to your jobs.
e) Implement your policies on how your jobs respond to grid and execution failures.
f) Improve the execution speed of individual jobs by overlapping I/O and computation.

C5) When DAGMan notices that one of the jobs it submitted to Condor-G has failed, which of the following best describes what it does next?
a) Immediately stops.
b) Runs as much of the DAG as possible and then outputs a rescue DAG.
c) Immediately outputs a rescue DAG.
d) Runs as much of the DAG as possible and then stops.

C6) Condor-G GlideIn provides which TWO of the following?
a) A mechanism to submit jobs with complex interdependencies.
b) A mechanism to get transparent job migration capabilities.
c) A mechanism to browse available grid computation sites.
d) A mechanism to dynamically schedule your jobs across worksites.
e) A mechanism to stage data at remote worker sites on

EU DataGrid

E1) What are the main applications that DataGrid aims at?
a) business applications
b) data intensive scientific application
c) massively parallel computing problems
d) thermodynamics and modelling

E2) What is the main focus of EDG middleware development?
a) higher level services for workload, data, and information management
b) application level portals
c) basic grid middleware for authentication/file transfer program execution
d) providing a grid operating system

E3) What is *not* required in order to use the EDG testbed?
a) sign up with a VO
b) obtain a proxy certificate with 'grid-proxy-init'
c) obtain a local user account at the resource you want to run the job on
d) obtain a certificate from a recognized CA

E4) What is the main role of the Resource Broker?
a) find out where all replicas are located
b) interface to Storage Systems
c) monitor the status of a fabric
d) matchmaking of job requirements

E5) Your job requires access to a file "input.dat" which is available on some Grid storage (SEs). Which field of the JDL do you use to make the broker aware of this requirement?
a) StdInput
b) InputSandbox
c) InputData
d) Requirements

E6) In order to locate your files on the Grid you need to
a) register the file in R-GMA
b) register the file in the replica catalog
c) register the file in the local application database
d) register the file in the resource broker database

E7) The EDG Replica Manager (EDG Release 2.0) client is responsible for replicating files and active interaction with other services, i.e. the Replica Manager _calls_ the services. One in the list below is not called:
a) Replica Optimization Service
b) GridFTP server
c) Resource Broker
d) Replica Location Service

E8) Which one of the following components is _not_ part of R-GMA?
a) Registry
b) back-end to object-relational database
c) Producer
d) Archiver

E9) In order to filter information within R-GMA you need to
a) set up an archiver that publishes into a new producer
b) modify the R-GMA registry
c) stream data into RLS
d) this cannot be done with R-GMA

E10) LCFG has been chosen by EDG in order to fulfil the following tasks. One of them is *not* correct that you have to identify:
a) configuration of a farm or a single machine
b) act as a central software repository where each member of the DataGrid project can down-load the software via HTTP
c) managing user accounts on a machine
d) local software repository for a LAN

Globus

G1) Which of the following is *not* a true statement about the Replica Location Service (RLS) in Globus GT3.0?
a)The soft state updates sent from Local Replica Catalogs (LRCs) to Replica Location Indexes (RLIs) eventually time out and are removed from the indexes unless refreshed by subsequent updates.
b) The RLS included in the GT3.0 release is not an OGSA service.
c) The contents of the LRCs and RLIs are kept strictly consistent with one another using the immediate mode for soft state updates.
d) Bloom filter compression of LRC updates reduces the network bandwidth and CPU overheads associated with soft state updates.

G2) What type of failures can the GT3.0 Reliable File Transfer Service recover from that cannot be survived by the basic GridFTP transport protocol?
a) failure and recovery of client that submitted transfer request
b) destination host failure and recovery
c) temporary network partition between source and destination
d) removal of file on host machine during transfer operation

G3) Soft state update in MDS-II is used to:
a) increase system scalability
b) decrease latency in access to information
c) increase system robustness
d) simplify the management of MDS.

G4) Grids are distinguished from other distributed computing technology by
a) the number of different types of resources that they use
b) spanning different organizational domains
c) the response time needed for a requested operation
d) the use of Web services to implement them

G5) In GSI, a proxy is distinguished by:
a) the number of bits in its key
b) using a symmetric encryption algorithm
c) having a lifetime
d) being signed by the users private key

G6) To be an OGSI gridservice, one must define which of the following port types? [check all that apply]
a) Gridservice
b) Factory
c) NotificationSource
d) NotificationSink
e) HandleResolver

G7) Every OGSI gridservice is also a webservice.
a) true b) false

Tony Doyle - University of Glasgow

Meeting Current LHC Requirements: Experiment Accounting (pre-Grid)

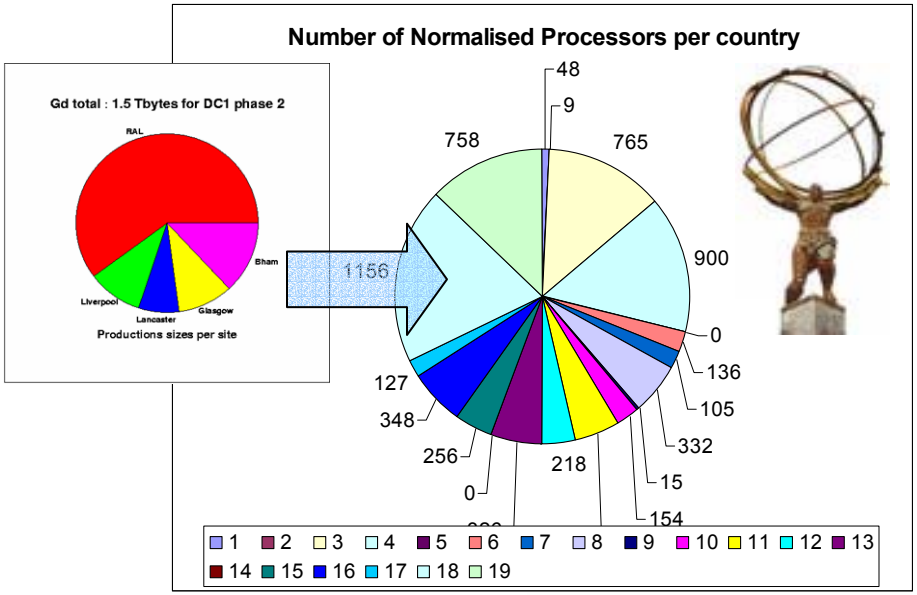

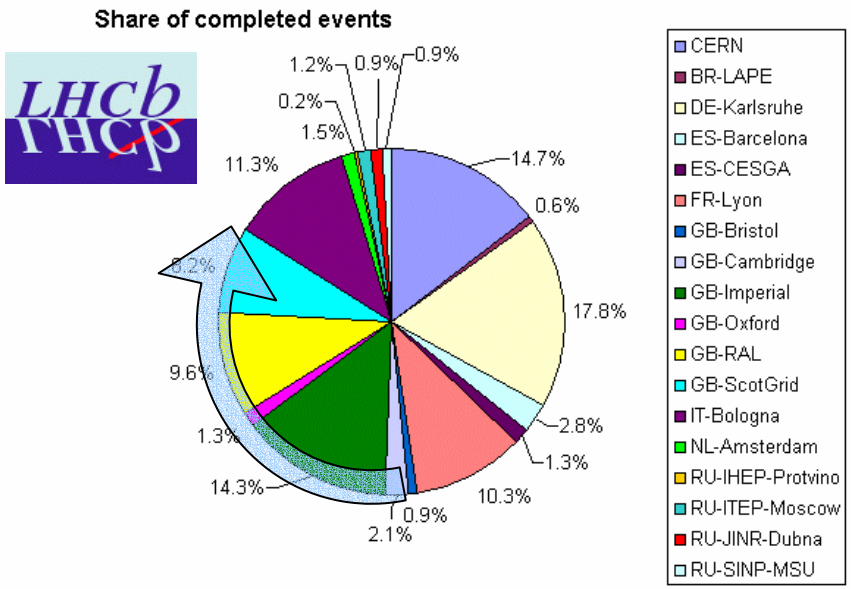

Regional Centre	Simulation	Hits	No Pile Up	2x1033	2x1044	NassID
Bristol/RAL	0.55	0.33	0.04	0.06	0.02	20
Caltech	0.17	0.15	0.00	0.15	0.00	6
CERN	0.89	2.20	1.40	2.66	2.25	300
Fermilab	0.35	0.41	0.00	0.25	0.33	70
ICST&M	0.88	0.59	0.50	0.15	0.12	84
IN2P3	0.20	0.00	0.00	0.00	0.00	1
INFN	1.55	1.18	0.40	0.72	0.71	99
Moscow	0.43	0.14	0.14	0.00	0.00	41
UCSD	0.34	0.30	0.00	0.29	0.30	80
UFL	0.54	0.04	0.00	0.04	0.04	11
USMOP	0.00	0.00	0.00	0.00	0.00	1
Wisconsin	0.07	0.08	0.00	0.06	0.00	12
TOTAL	5.94	5.40	2.47	4.36	3.77	

Experiment-driven project. Priorities determined by Experiments Board.

What is the GridPP1 Project Status?

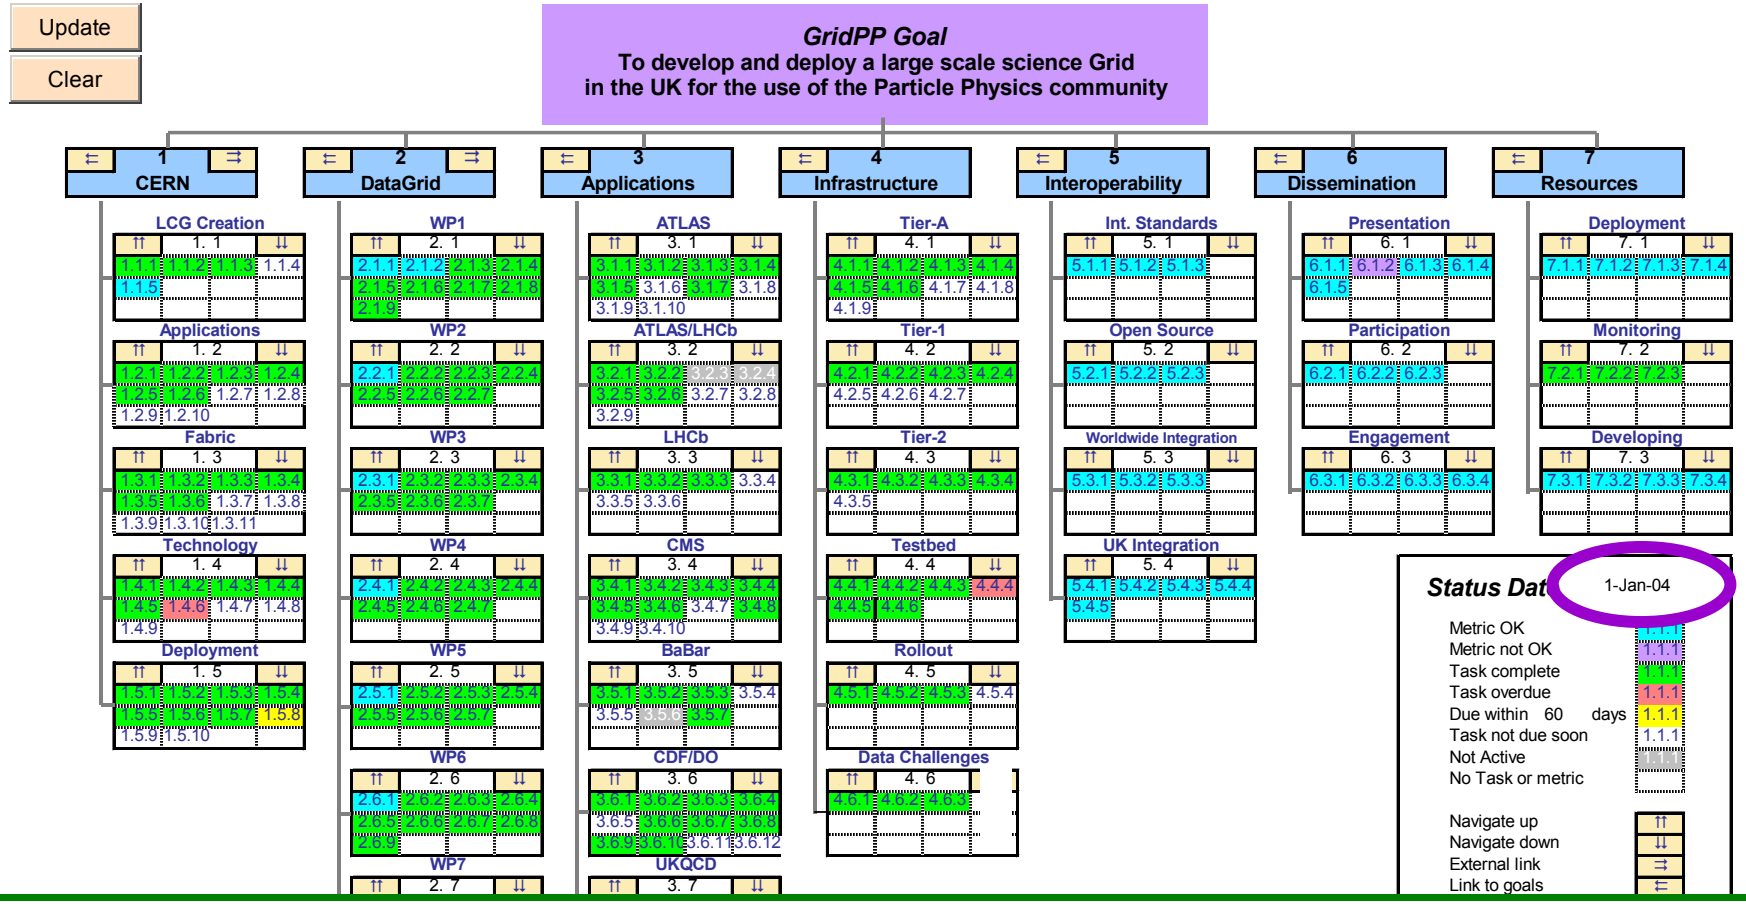

Metric OK	Metric not OK	Tasks Complete	Tasks Overdue	Tasks due in next 60 days	Tasks Deleted	ChangeForms Enacted	Tasks not due
43	1	145	2	1	3	17	42

- 76% of the 190 GridPP1 tasks have been successfully completed

Achievements I

1. Dedicated people actively developing a Grid
2. All with personal certificates
3. Using the largest UK grid testbed (16 sites and more than 100 servers)
4. Deployed within EU-wide programme
5. Linked to Worldwide Grid testbeds

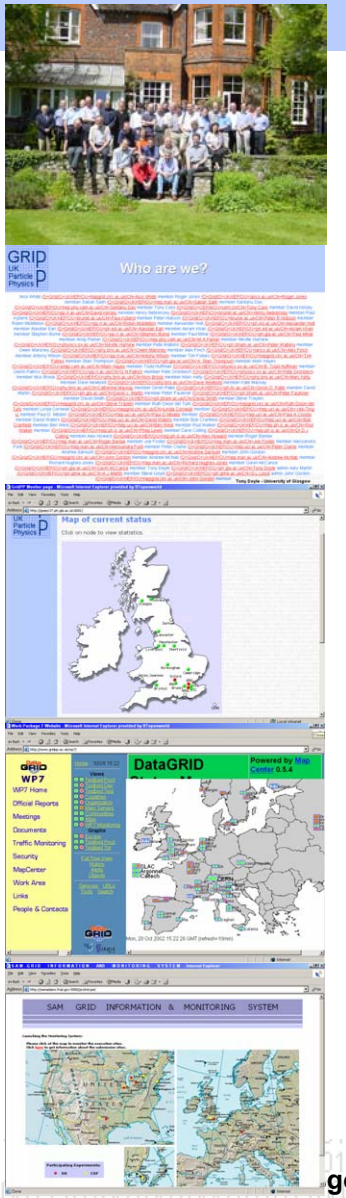

Achievements II

6. Grid Deployment Programme Defined The Basis for LHC Computing
7. Active Tier-1/A Production Centre meeting International Requirements
8. Latent Tier-2 resources being monitored

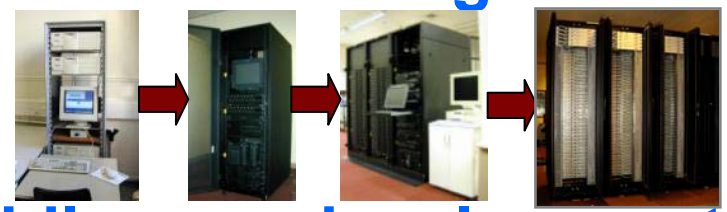

9. Significant middleware development programme
10. First simple applications using the Grid testbed (open approach)