Grid Computing
- a short history of the future

Dr Mark Parsons

Commercial Director, EPCC and NeSC

m.parsons@epcc.ed.ac.uk

BCS Edinburgh, 5th February 2003

National e-Science Centre

epcc

Summary

- a history of the early Grid
- where the Grid is today
- where the Grid is going
- who should worry about the Grid
  - some reasons to be sceptical
- conclusion

a history of the early Grid

National e-Science Centre

epcc

In the beginning

- in 1999 Ian Foster and Carl Kesselman editted a book of papers

  “The Grid – Blueprint for a new computing infrastructure”

- timely, widely read and commented upon

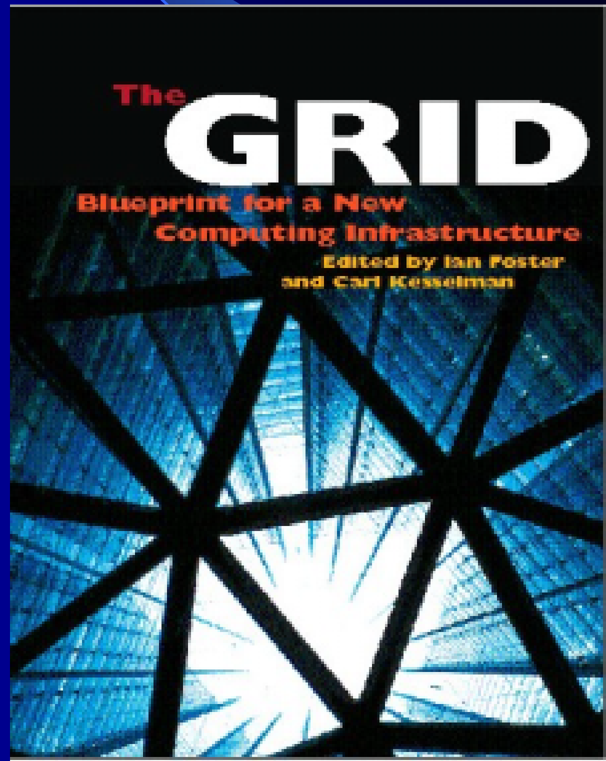

Not really the beginning

- *The Grid* presented its ideas as new but really built on ten+ years of research into meta-computing
- what's the difference?
  – it might work this time
  – metacomputing focussed mainly on joining large distributed systems together
  – the Grid idea is much more hetrogeneous and all encompassing

The vision

- the early Grid vision was analogous to the electric power grid
  - connections provided wherever you are
  - always on, highly reliable
  - pay for what you use
  - shared resource
  - always sufficient capacity (unless you're Californian …)

The vision (cont)

- the vision has been refined over past 3 years and is now accepted as

  … a software infrastructure that enables flexible, secure, co-ordinated resource sharing among dynamic collections of individuals, institutions and resources …

  – Foster, Kesselman & Tuecke (Anatomy of the Grid, 2001)

Why now?

- several forces are driving us towards the Grid
  - the relentless increase in microprocessor performance
  - the availability, reliability and bandwidth of global networking
  - new scientific experiments are producing a data explosion

CPU performance

- in 1990 EPCC owned Europe's fastest computer
  - 800 Megaflops (peak)
- in 2003 EPCC owns Europe's fastest computer
  - 3 Teraflops (sustained)
  - 3,000,000 Megaflops …
- most computers from Toys'R'Us will do more than 1Gflop for less than £900

CPU performance (cont)

- what do we use all this performance for
  – mostly Microsoft Word, Powerpoint etc
- there are enormous untapped computing resources
  – companies like Entropia seek to tap into this
  – some people believe this is the Grid
  – but only works for small class of problems

Networking

- bandwidth has multiplied relentlessly over past 3 years
- JANET from 2000 to 2002:
  - 155Mbps -> 655Mbps -> 1Gbps -> 10Gbps …
- bandwidth not only story
  - reliability, ease of access, European/global connectivity major part of story
- main European research backbone – GEANT – now provides 1-10Gbps connectivity between 32 countries
- the research sector has benefited greatly from over-provisioning of fibre accompanying the dot.com boom

New science

- scientists love to dream
  - they dream of
    - perfect, high-resolution detectors
    - capturing and storing every last nuance of data
    - simulating the data to better understand it
    - processing the data in real time
    - making the data available to any scientist, in any institution, anywhere in the world
- 20+ years in the planning, the silicon revolution is making many of these dreams come true

e-Science

- from biology to particle physics science is now pushing the limits of computing
  – collaborations are now global
  – datasets are shared
  – computer simulation is an accepted tool in the scientists arsenal
- in the UK, scientific discoveries enabled by collaborative computing are now called *e-Science*
  – the Grid technologies enable e-Science

Complexity = £££s

- very large amounts of money are being spent in the UK, US, Europe and Asia-Pacific
  - UK Grid funding £100+ million
  - US Grid funding $500+ million
  - European Commission Grid funding €50+ million
- justified to funding bodies on complexity of problems under study
- interesting to compare with development of HTML and the WWW …

The Original WWW Proposal

To: P.G. Innocenti/ECP, G. Kellner/ECP, D.O. Williams/CN

Cc: R. Brun/CN, K. Gieselmann/ECP, R. Jones/ECP, T. Osborne/CN, P. Palazzi/ECP, N. Pellow/CN, B. Pollermann/CN, E.M. Rimmer/ECP

From: T. Berners-Lee/CN

Date: **12 November 1990**

The attached document describes in more detail a Hypertext project.

- HyperText is a way to link and access information of various kinds as a web of nodes in which the user can browse at will. It provides a single user-interface to large classes of information (reports, notes, data-bases, computer documentation and on-line help). We propose a simple scheme incorporating servers already available at CERN.
- The project has **two phases**: firstly we make use of existing software and hardware as well as implementing simple browsers for the user's workstations, based on an analysis of the requirements for information access needs by experiments. Secondly, we extend the application area by also allowing the users to add new material.
- **Phase one should take 3 months** with the full manpower complement, **phase two a further 3 months, but this phase is more open-ended**, and a review of needs and wishes will be incorporated into it.
- The **manpower required is 4 software engineers and a programmer**, (one of which could be a Fellow). Each person works on a specific part (eg. specific platform support).
- Each person will require a state-of-the-art workstation , but there must be one of each of the supported types. These will cost from 10 to 20k each, **totalling 50k**. In addition, we would like to use commercially available software as much as possible, and foresee an expense of 30k during development for one-user licences, visits to existing installations and consultancy.
- We will assume that the project can rely on some computing support at no cost: development file space on existing development systems, installation and system manager support for daemon software.

T. Berners-Lee, R. Cailliau

UK e-Science funding

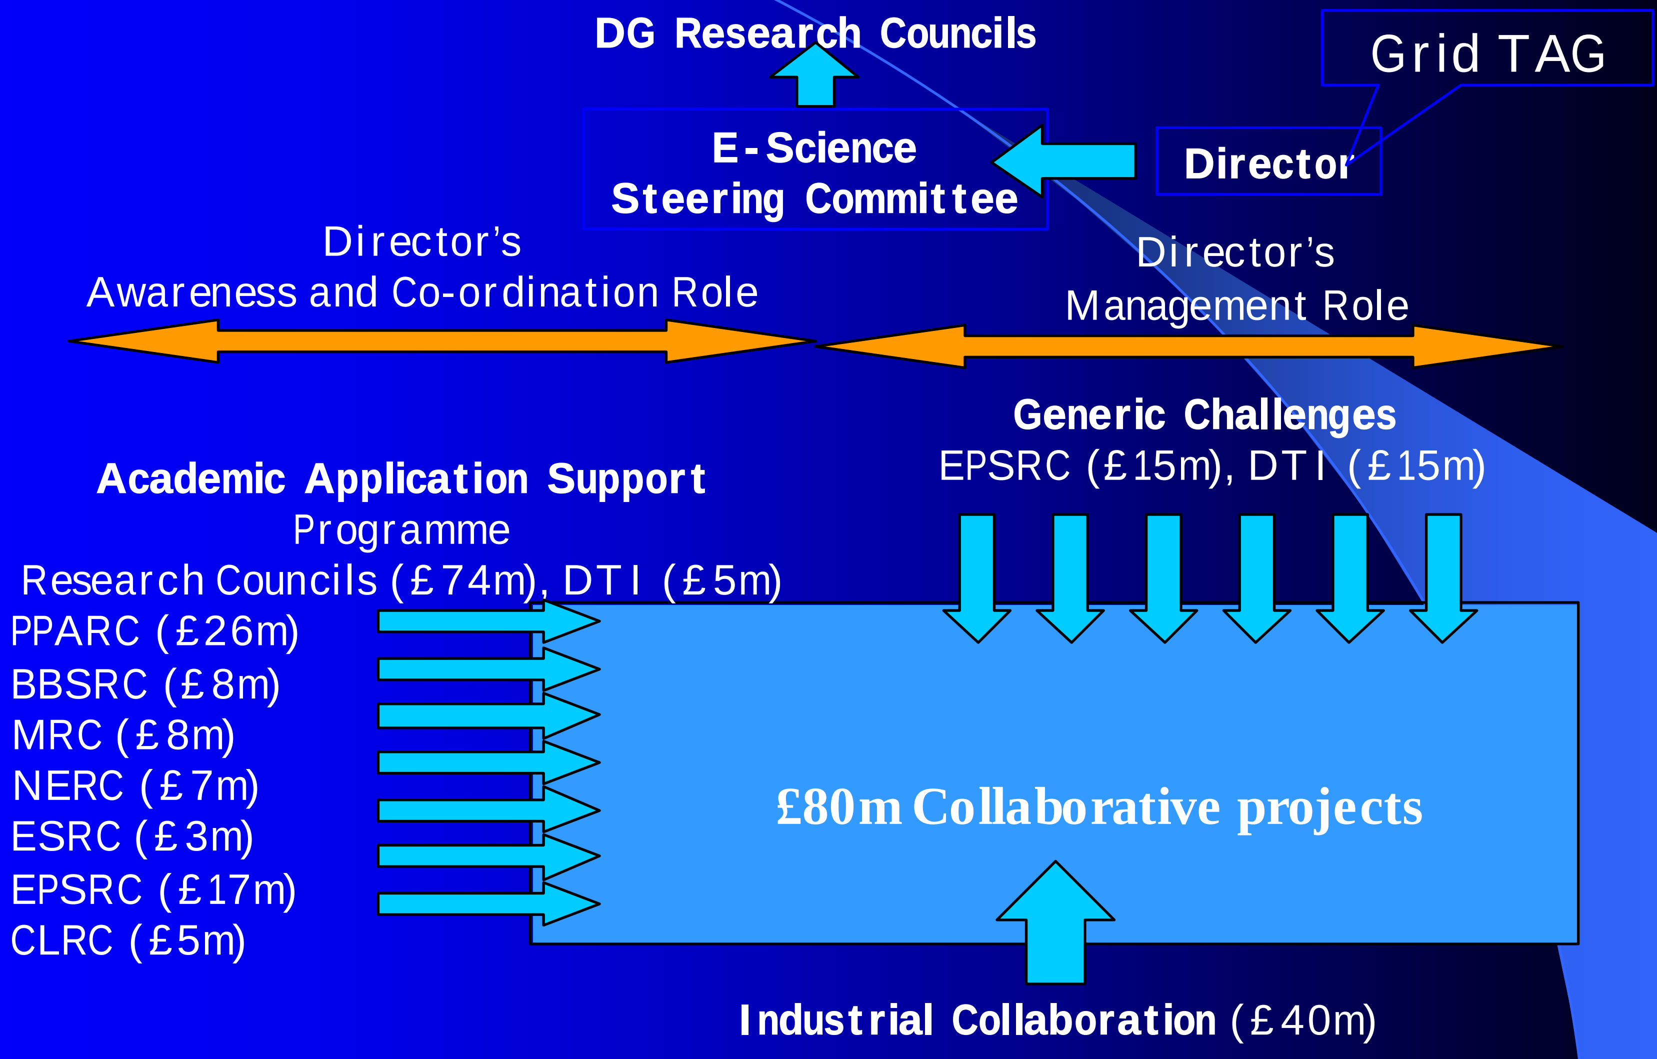

Why the expense?

- why did HTML cost £50,000 and why is the Grid costing 10,000 times more?
  - Tim Berners-Lee has said you couldn't invent the WWW so simply now (CERN only adopted TCP/IP in 1987 …)
  - we work in a vastly more complex software environment
    - during the 1990's computing matured much like physics in the 1900's
    - the domain is now too big for one person to comprehend
  - the scale of the data challenges the Grid will be applied to dwarfs anything attempted before

where the Grid is today

“Big science” drivers

- much of the early development of the Grid is being driven by big science
  - the large hadron collider
  - genomics
  - earth observation
  - astrophysics
  - …
- the challenge is to make the resulting software useful to all – *not just scientists*

The LHC project

- at CERN in Geneva the world's largest machine is under construction
- the LHC will probe the conditions at the start of time
- the fundamental goal is the discovery of the Higgs Boson

The 26Km tunnel housing LHC

LHC (cont)

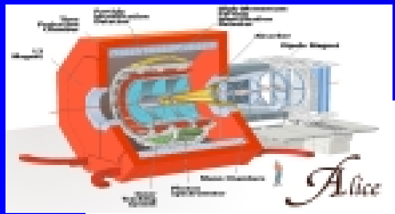
Alice

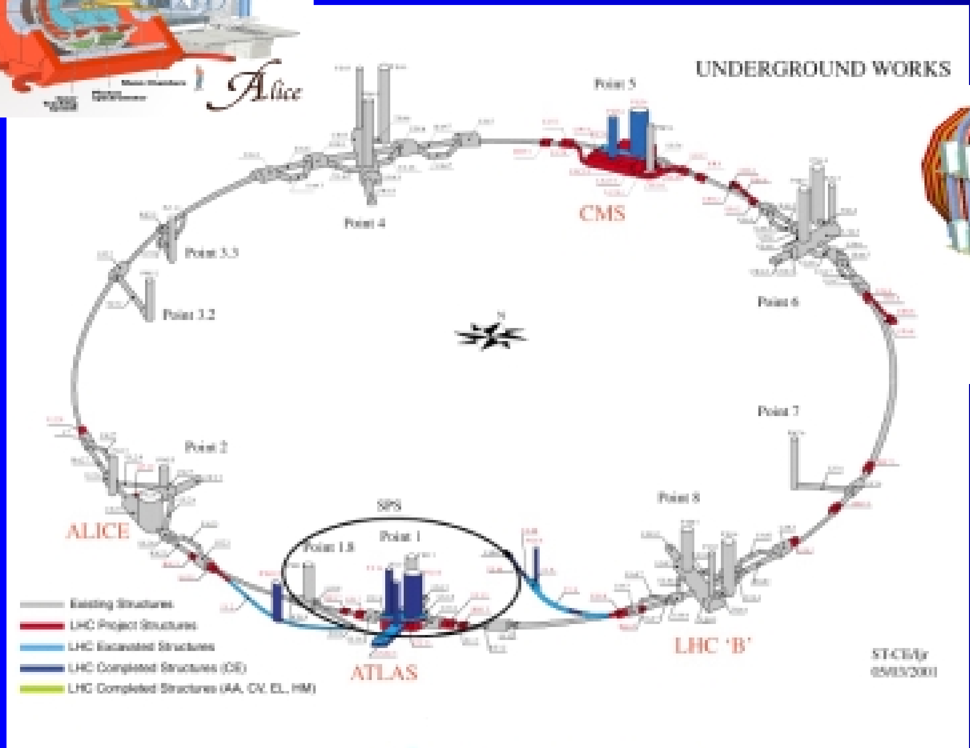
UNDERGROUND WORKS
Point 5
CMS
Point 4
Point 3.3
Point 3.2
Point 6
Point 7
Point 2
Point 8
ALICE
SPS
Point 1.8
Point 1
LHC 'B'
ATLAS
Existing Structures
LHC Project Structures
LHC Excavated Structures
LHC Completed Structures (CE)
LHC Completed Structures (AA, CV, EL, HM)

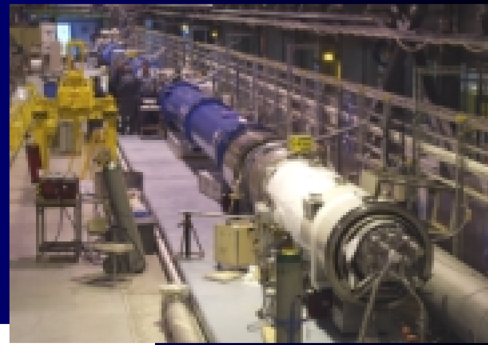

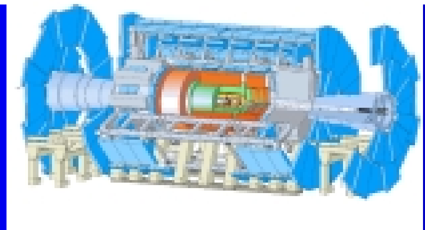

LHC (cont)

- this is a huge computing problem
  - raw recording rate of 0.1 – 1 Gbytes/s
  - accumulating at 5-8 Pbytes/year
  - 10 Petabytes of disk needed per year minimum
  - processing requirements involve 200,000 of today's PCs
- and this is ***PER EXPERIMENT*** …

HEP analysis

- until LHC, for a PhD in HEP you need to learn to do the following:
  - capture events from your detector
  - create Monte Carlo data (10x real)
  - write them RAW to disk or tape
  - process them into Data Summary Tapes (DSTs)
    - with corrections etc
  - process them into miniDSTs
    - specific Physics information pre-processed and stored
  - analyse the data comparing against MC
  - extract Physics measurements

Why is LHC different?

- at LEP a typical analysis:
  - mounted 2-3000 tapes
  - ran for 1-2 months on Cray or IBM or WS cluster
- but
  - data rates were much lower 100Kbytes/s
  - maximum event sample was generally <10million events < 1Tb
- at LHC
  - data samples for 1 year will be at least 1Pb
  - much more complex events – LEP had 50 particles, LHC will have 10,000+

A Higgs?

- the LHC search
  - is the equivalent of looking for 1 person in 100 human populations
  - must sift through 10^8 events per second

LHC summary

- challenges exist at a number of levels
  - vast quantities of data
  - enormous computing requirements
  - highly distributed collaborations
    - 475 institutes worldwide, 6,235 physicists
- the physics is also a challenge
  - but without the computing infrastructure the physics cannot be done …

Earth Observation

- new high resolution satellites are coming on stream eg. ENVISAT
  - €3.5billion 10 instruments
  - 200Mbps download
  - 400Tbytes/year
  - 100 products
  - 10+ European facilities
  - 700 approved users

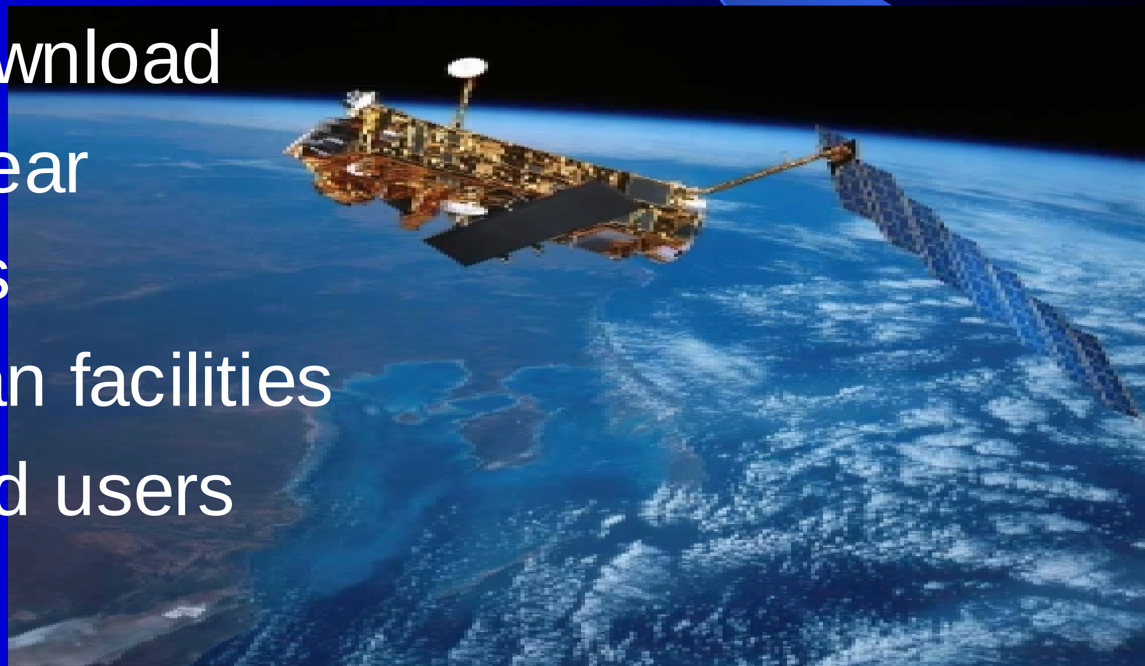

EO Applications

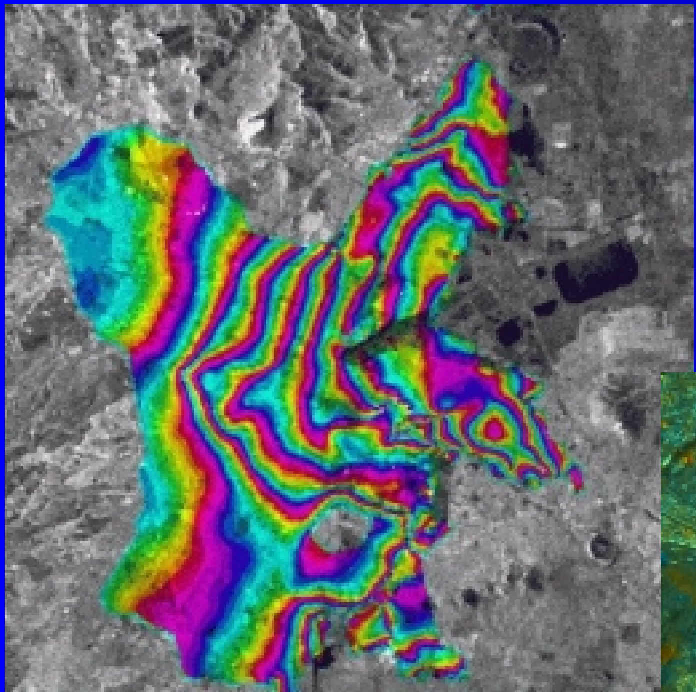

ground deformation prior to volcano

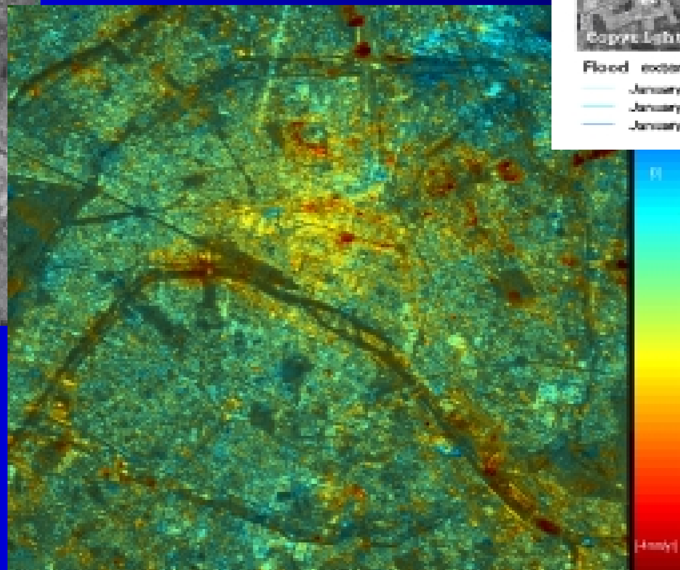

subsidence in central Paris

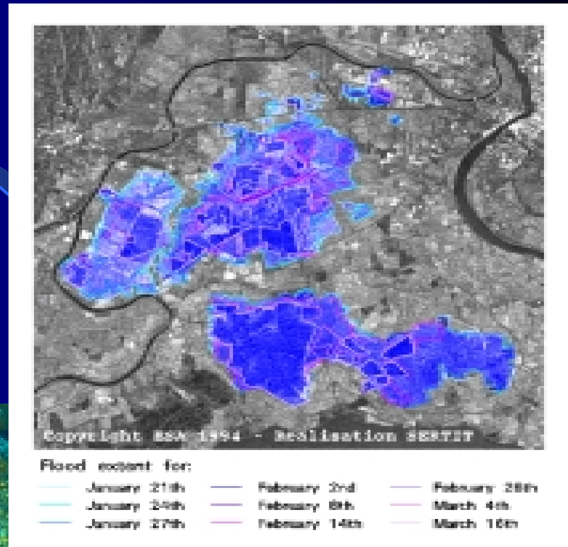

flood prediction

All require Grid computing and storage resources

Genomics

- the mapping of the human genome is
  - creating a data explosion within bio-informatics
  - driving the global bio community towards establishing virtual organisations
  - forcing the creation of new distributed data indexing, searching and matching technologies
- new lab techniques require more and more computing resource
  - CPU
    - biologists are now using HPC services daily
  - data storage
    - new high quality imaging techniques are rapidly emerging

3D Optical Microscopy

- good example of biology Grid project
  - new imaging technique
  - very CPU intensive
  - capture takes 30 mins
  - 1.2Tb data produced
  - processing takes 5+ hours
- Grid project with EPCC cut processing to 15 minutes

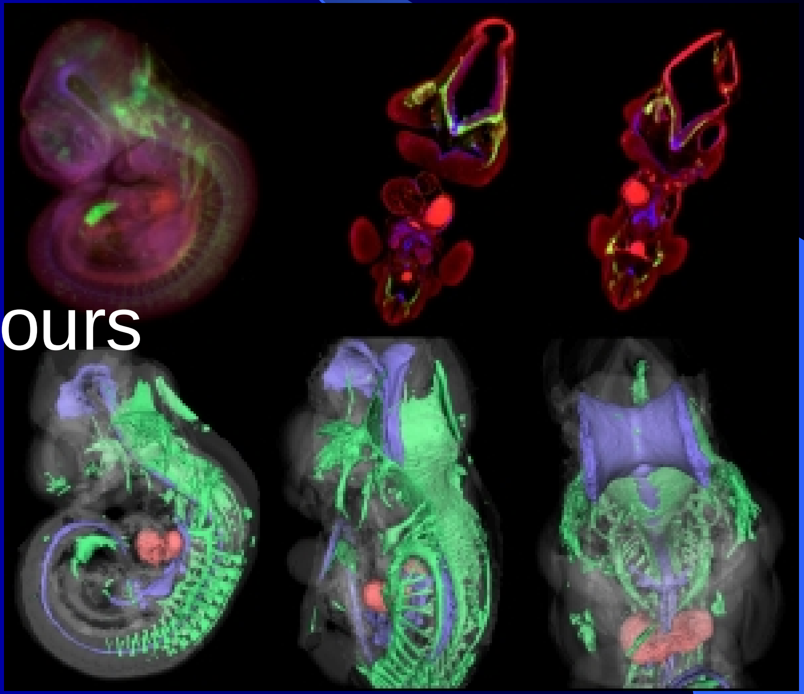

Mouse embryo example

image captured
from specimen

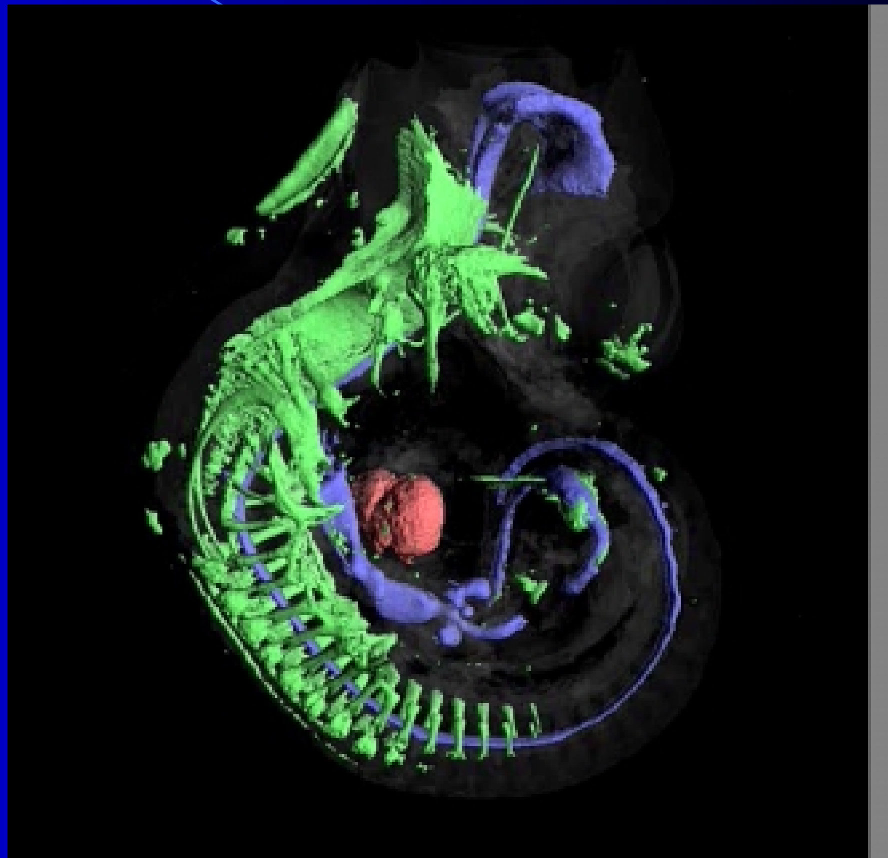

Green - developing nervous tissue
Blue - gut and other organs
Red - ventricles of the heart

How the Grid does this …

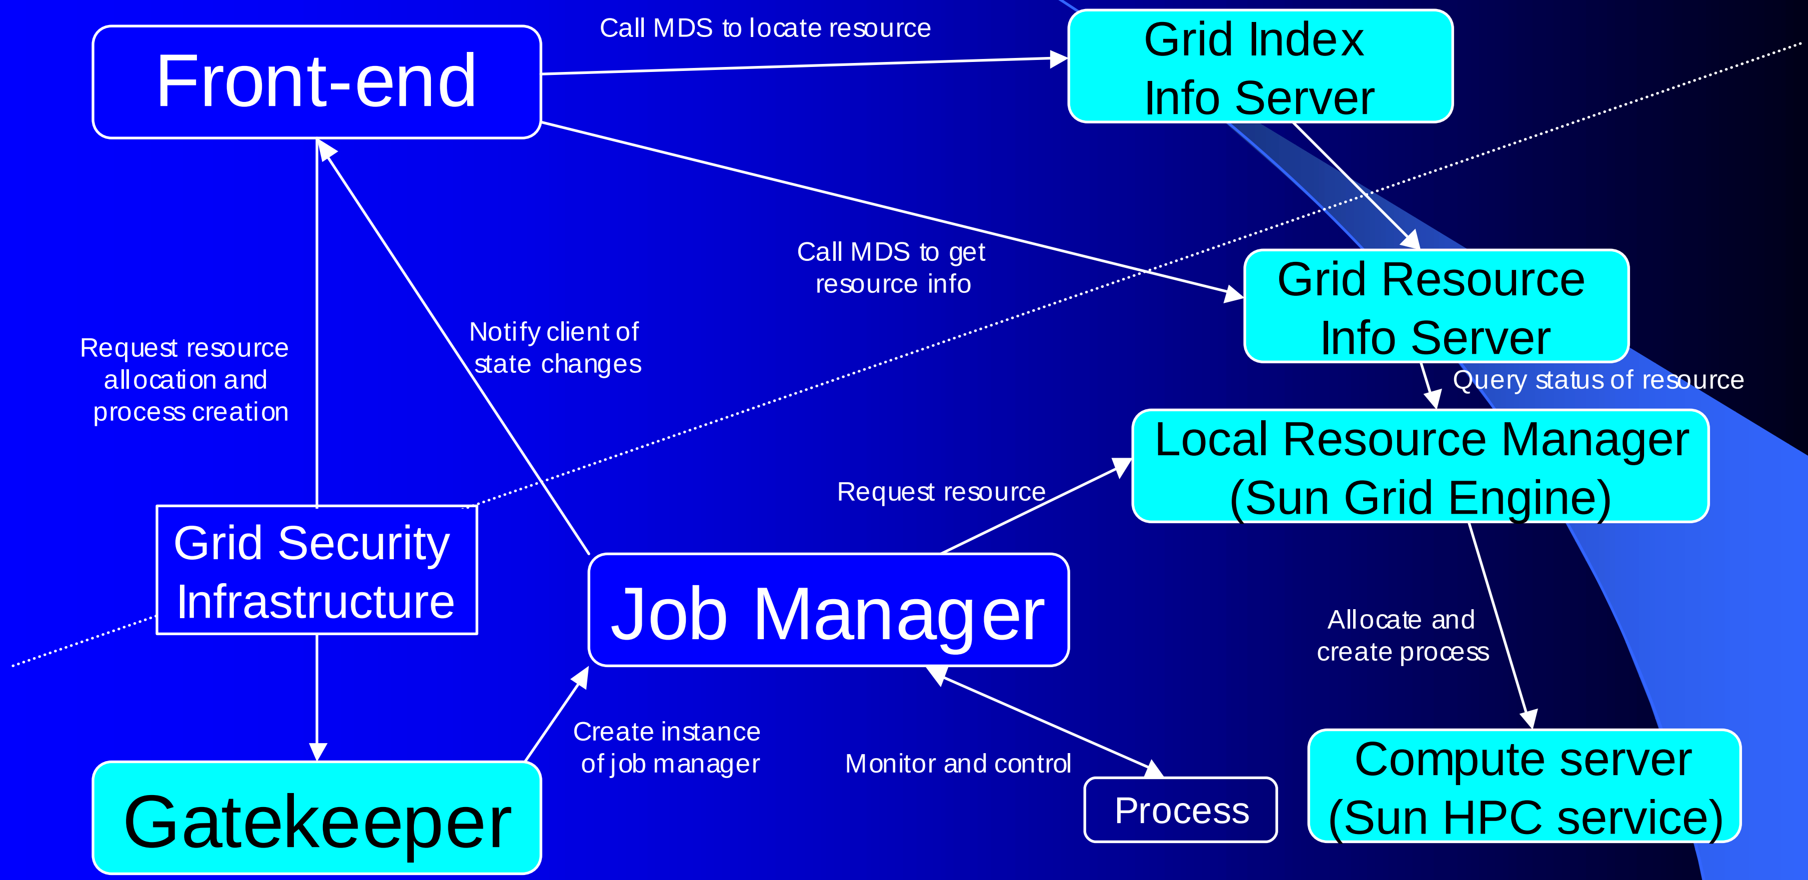

(using Globus 2)

Grid middleware

- Grid software development has been going strongly for 3+ years
- to be part of a Grid you have to run *Grid Middleware* on your computer(s)
- the only game in town is Globus
  - http://www.globus.org
- currently most Grid applications projects use Globus 2.x

A Grid must provide

- Identity & authentication
- Authorization & policy
- Resource discovery
- Resource characterization
- Resource allocation
- (Co-)reservation, workflow
- Distributed algorithms
- Remote data access
- High-speed data transfer
- Performance guarantees
- Monitoring
- Adaptation
- Intrusion detection
- Resource management
- Accounting & payment
- Fault management
- System evolution
- Etc.
- Etc.
- …

Middleware

- there are many different ways to provide such functionality
- Globus 2.x provides a classic, bespoke, sockets based approach to middleware
  - but things are changing rapidly …OGSA is coming
- most current projects use Globus 2.x but are now studying how to migrate

Globus 2.x architecture

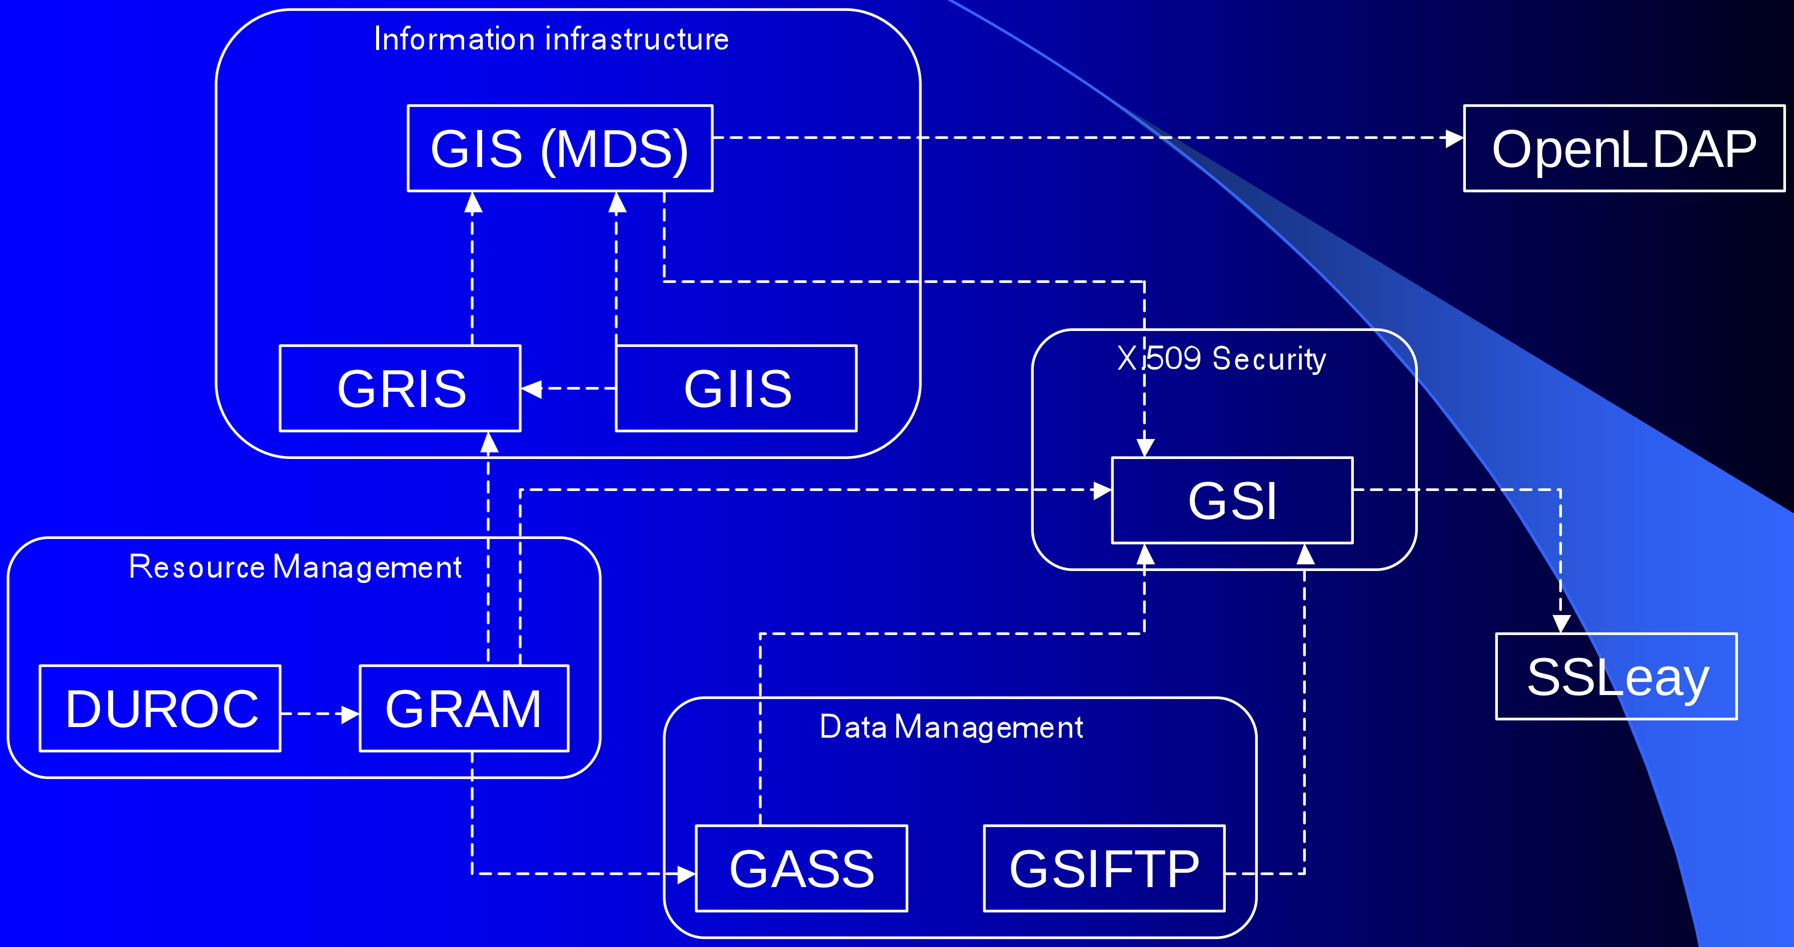

Globus 2.x (cont)

- the Globus middleware relies on X.509 certificates
  - a single sign on using CAs is implemented
- Globus sits above conventional scheduling software like NQS, PBS or LSF
- Globus 2.x takes a very “computing centre” view of the world
- in general using it is a major investment in time and effort

Current projects

- there are a plethora of Grid projects
- range from pure science to industrial and commercial applications
  - main CERN project is €10M EU DataGrid
- in UK DTI/OST funds e-Science Programme
  - see http://www.nesc.ac.uk
- for European Commission funded projects
  - see http://www.gridstart.org

where the Grid is going

what's next

- gradually the deficiencies in the Globus architecture have become apparent
  - it is after all a research project
- during 2002 the *Open Grid Services Architecture* was defined at the Global Grid Forum
  - first reference implementation released on 15th Jan 2003
  - known as Globus Toolkit 3

OGSA

- OGSA reinvents the Grid
  - based on and extends Web Services concept
  - all the Globus 2.x clunkiness is gone
  - strong support from vendors
    - IBM has 200+ developers working with Argonne National Lab on their implementation of OGSA
  - OGSA provides basic services
    - the Grid community must now develop services to sit on top
      - eg. EPCC/NeSC's Data Access & Integration services
  - brings Microsoft via .NET into the game
    - EPCC is working with Microsoft on a .NET implementation of OGSA-DAI

General principles

- one of the main drivers behind OGSA has been the need to
  - better support virtual organisations
  - enable dynamic sharing of resources
    - heterogeneity is *major* requirement
  - standardise on a (small) number of protocols
    - Globus 2.x uses too many – porting is difficult
- OGSA is derived from the Open Grid Services Infrastructure definition (OGSI)
  - see http://www.gridforum.org

General principles (cont)

- OGSA describes a set of implementation and platform independent protocols and services
  - these define technology and infrastructure of the next generation of Grids
- defined in terms of Grid Services
  - closely allied to Web Services
- supports the sharing and coordinated use of diverse resources in dynamic distributed virtual organisations

Web services

- a Web Service is an interface that describes a collection of operations that are network accessible.
- they build on XML based technologies and messaging
- use XML schemas to mark-up and describe services and their operations
- Web Services commonly use the Simple Object Access Protocol (SOAP) as a communication protocol over HTTP
  - when this is not appropriate for performance reasons a different protocol may be used

Grid services

- OGSA pulls together Open Grid Architecture and Web Services to form Grid Services
- Grid Services are dynamic, transient, and have state (i.e. have a finite lifetime) and are defined by a well defined set of interfaces and behaviours
- interfaces and behaviours describing Grid Services can be written in an XML schema based WSDL document
  - other implementations are possible.

OGSA-DAI

- OGSA Data Access and Integration services are the first of many higher level services
  - from EPCC with IBM, Oracle, NeSC and Universities of Manchester and Newcastle
- world first
  - released on same day as Globus Toolkit 3 reference at GlobusWorld, San Diego

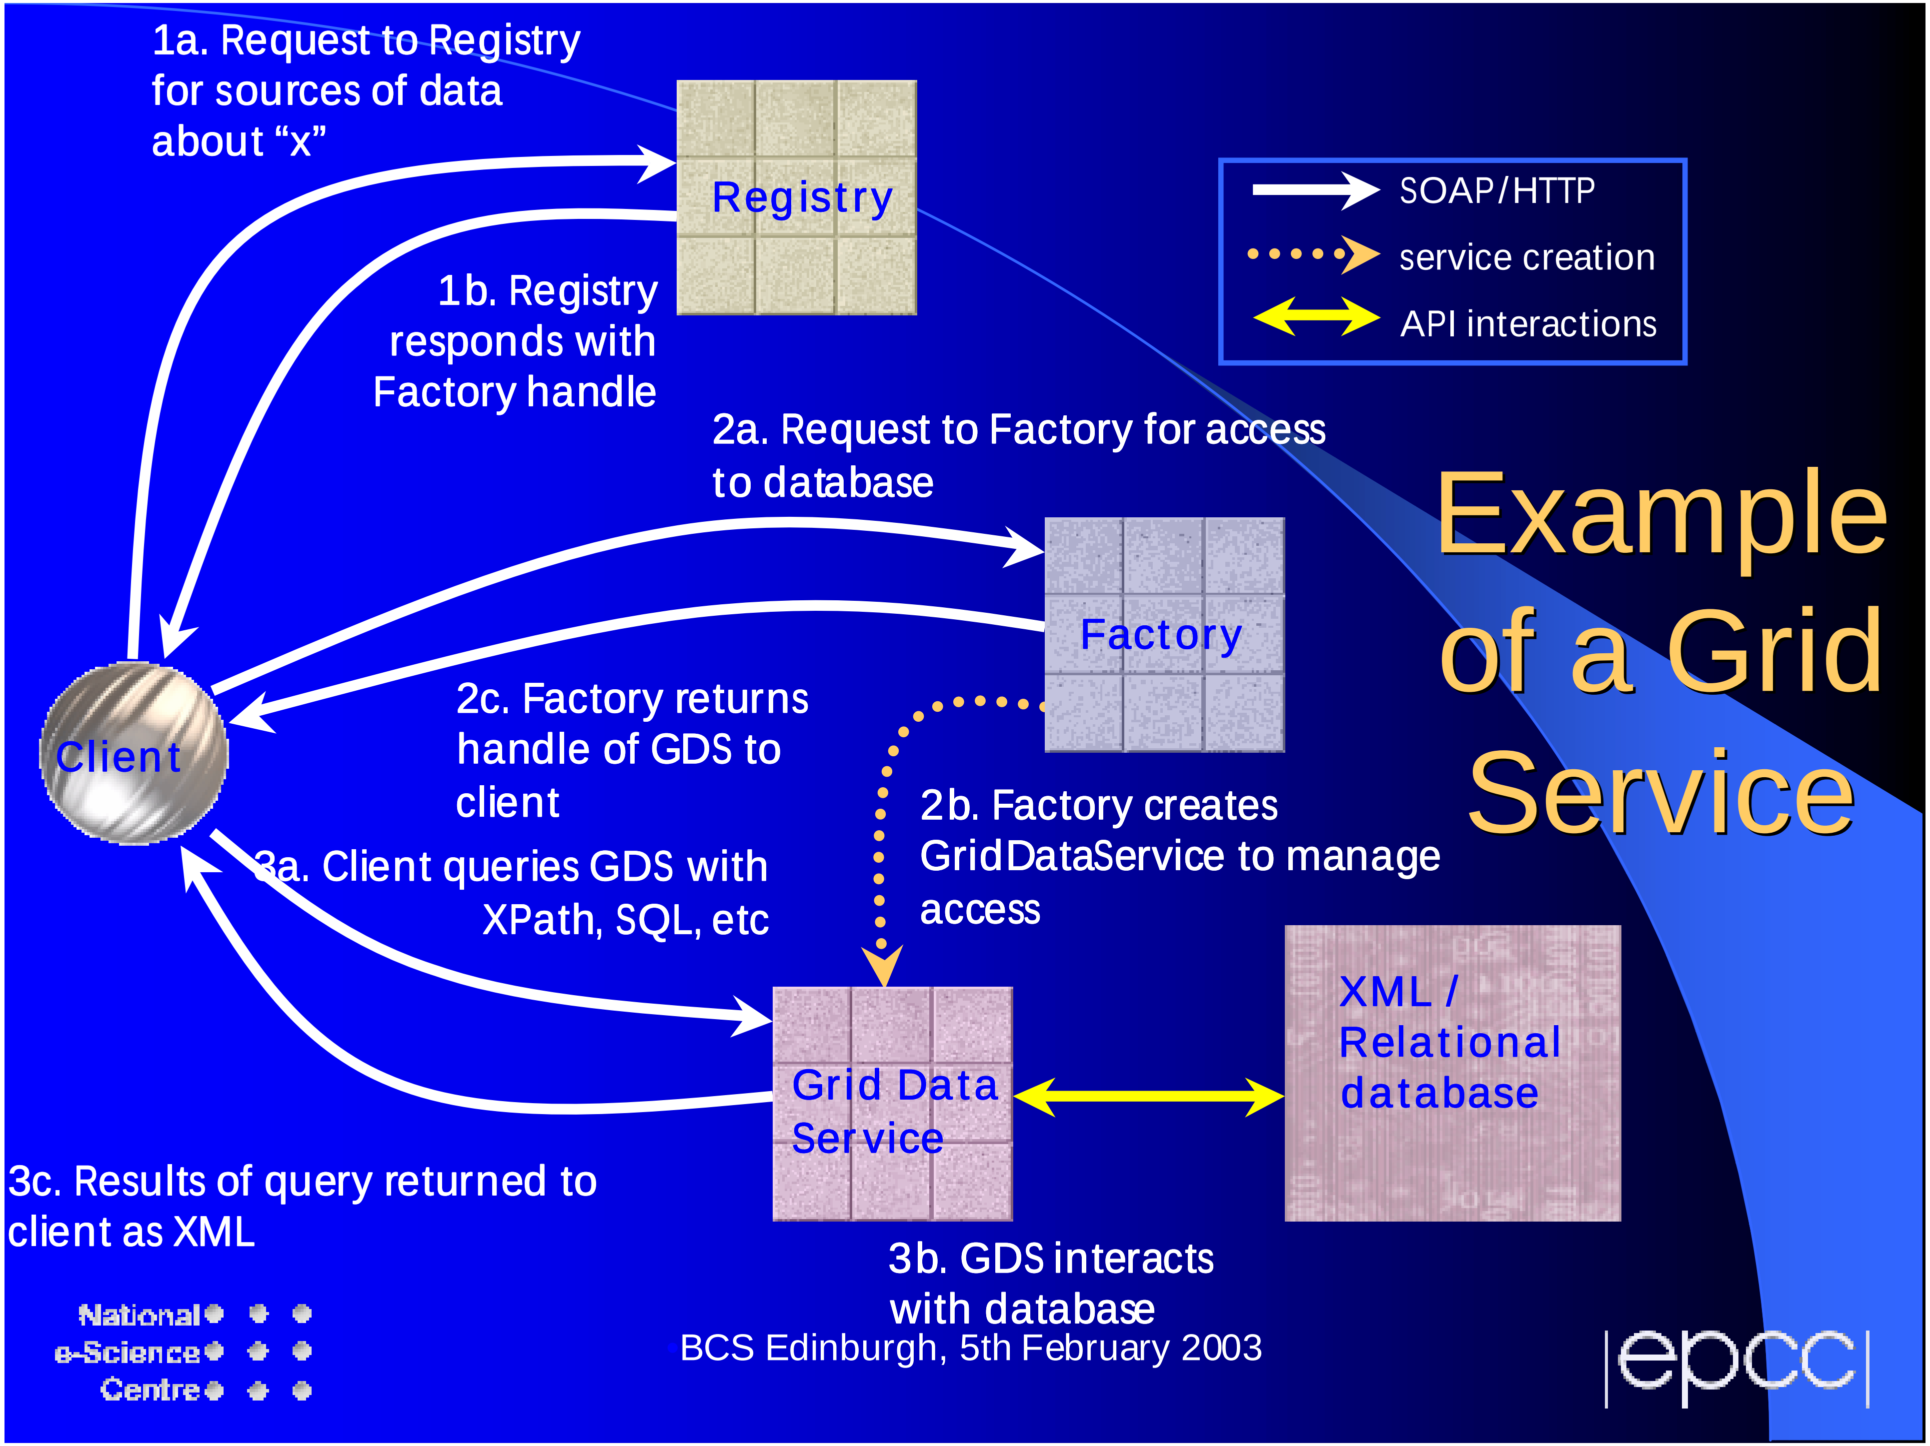
Example of a Grid Service
1a. Request to Registry for sources of data about "x"
Registry
1b. Registry responds with Factory handle
SOAP/HTTP
service creation
API interactions
2a. Request to Factory for access to database
Factory
2c. Factory returns handle of GDS to client
Client
2b. Factory creates GridDataService to manage access
3a. Client queries GDS with XPath, SQL, etc
XML / Relational database
Grid Data Service
3c. Results of query returned to client as XML
3b. GDS interacts with database
National e-Science Centre
epcc

a distributed query

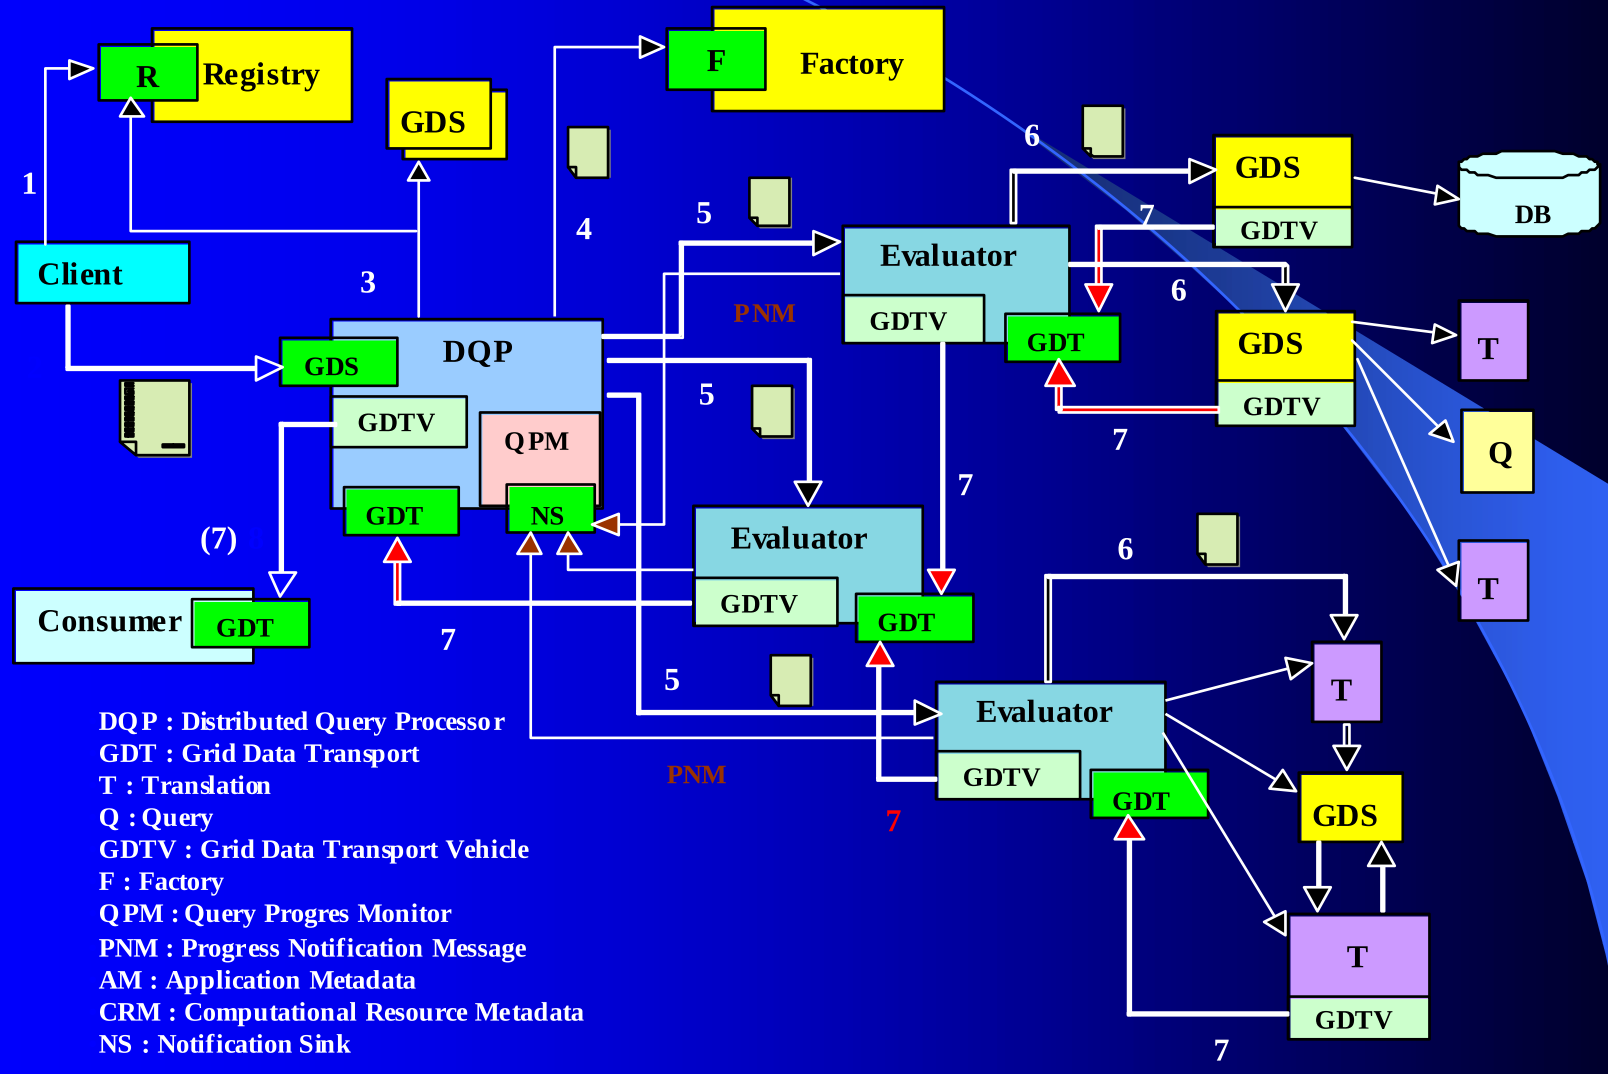

DQP : Distributed Query Processor
GDT : Grid Data Transport
T : Translation
Q : Query
GDTV : Grid Data Transport Vehicle
F : Factory
QPM : Query Progres Monitor
PNM : Progress Notification Message
AM : Application Metadata
CRM : Computational Resource Metadata
NS : Notification Sink

final thoughts on OGSA

- we're just at the start of OGSA
- very few services exist
- most currently funded projects use Globus 2.x and will have to migrate
- timing is very difficult
  - probably will be end 2003 before robust Globus Toolkit 3 exists
- but … this will be *the* Grid underlying technology for next few years

who should worry about the Grid

who should care

- to date Grid development has been led by the scientific community
- the work has been largely irrelevant to industry and commerce
  - other than computer vendors
  - and despite what we write in our proposals …
- only the scientific community should worry about Globus 2.x
  - good functionality for science
  - useful to provide access to distributed resources
    - in particular access to large remote systems

industry and commerce?

- the move to OGSA poses some interesting questions
  - if OGSA truly delivers, many of its services will be very interesting to companies
    - single sign on
    - distributed, secure access to data globally
    - access to and management of heterogenous resources etc etc
- many companies are already using Web Services
  - it won't be a major leap to using Grid Services

industry (cont)

- in the future vendors will package Grid Services
  - eg. IBM may distribute with WebSphere
- Microsoft's .NET strategy is very close to Grid Services
  - .NET is being used for OSGA prototyping in many projects
- in 18 months time industry and commerce should be prepared to evaluate the first vendor quality Grid Services

conclusion

conclusion

- the Grid is a vast area of research
  - only scraped the surface in this talk
- although largely driven by science only now is its true form being understood
  - the link between the Grid and Web Services will transform the opportunities
- the Grid *may* be the first major new technology of the new century
  - we won't know until *you've* tried it …