Enabling Grids for E-science in Europe

What are Grids and e-Science?

Richard Hopkins

National e-Science Centre *Edinburgh, Scotland* *rph@nesc.ac.uk*

EGEE is funded by the European Union under contract IST-2003-508833

Acknowledgements

- This talk is based on a module of the tutorials delivered by the EDG training team and slides from
  - Andrew Grimshaw, University of Virginia
  - Bob Jones, EGEE Technical Director
  - Mark Parsons, EPCC
  - the EDG training team
  - Roberto Barbera, INFN
  - Ian Foster, Argonne National Laboratories
  - Jeffrey Grethe, SDSC
  - The National e-Science Centre
- Prepared by Richard Hopkins

Questions to be answered

- What is e-Science and the role of grids?
- What's in it for working scientists?
- What are grids really like?
- What does EGEE build on?

What is e-Science and the role of Grids?

Our perspective is grids to enable e-science

What is e-Science?

A definition of *e-Science*

- Invention and exploitation of **advanced computational methods** to support **scientific research**

A To generate, curate and analyse research *data*
- From experiments, observations and simulations
- Quality management, preservation and reliable evidence

B To develop and explore *models and simulations*
- Computation and data at extreme scales
- Trustworthy, economic, timely and relevant results

C To enable *dynamic* distributed virtual organisations
- Facilitating collaboration with information and resource sharing and replication
- Security, reliability, accountability, manageability and *agility*

- Grid Infrastructure is what allows **A** and **B** to happen within **C**
- **A** and **B** provide e-science methods as (possibly) grid applications

Scale of Data and Processing Resources

The need for advanced computational methods is to deal with
the scale of available data
the scale of resources potentially available to process it

- Proliferation of high volume data sources
  - Instruments, detectors, sensors, scanners, …
  - Organising their effective use is the challenge
- Enormous quantities of data: Petabytes (10^{15}) per year
  - For an increasing number of communities
  - Analysing the data is the challenge
- Huge quantities of computing: >100 Top/s
  - Moore's law gives us all supercomputers
  - Organising their effective use is the challenge
- Ultra-high-speed networks: >10 Gb/s
  - Global optical networks
  - Bottlenecks: last kilometre & firewalls

The desire to use this to push the boundaries of scientific discovery by computational analysis and simulation – e-Science

The main drivers behind e-Science – A ubiquitous pattern

- The interaction between developments in technology and in applications

Application demand → New technology to meet it

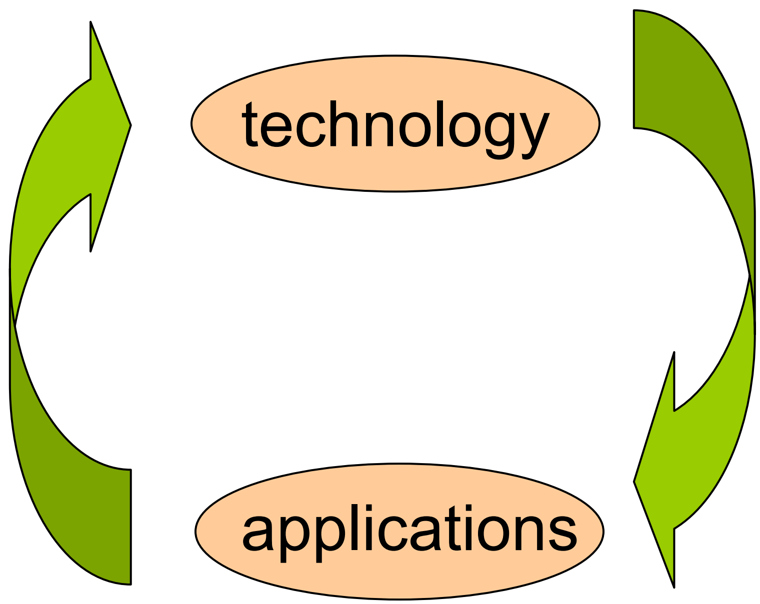

increased technological capability → New applications to use it → increased demands on technology

- A general (un-ending) pattern – e.g. WWW
- For e-Science-
  - Scientific enquiry as the application domain
  - Main technology driver has been Network capability

Exponential Growth

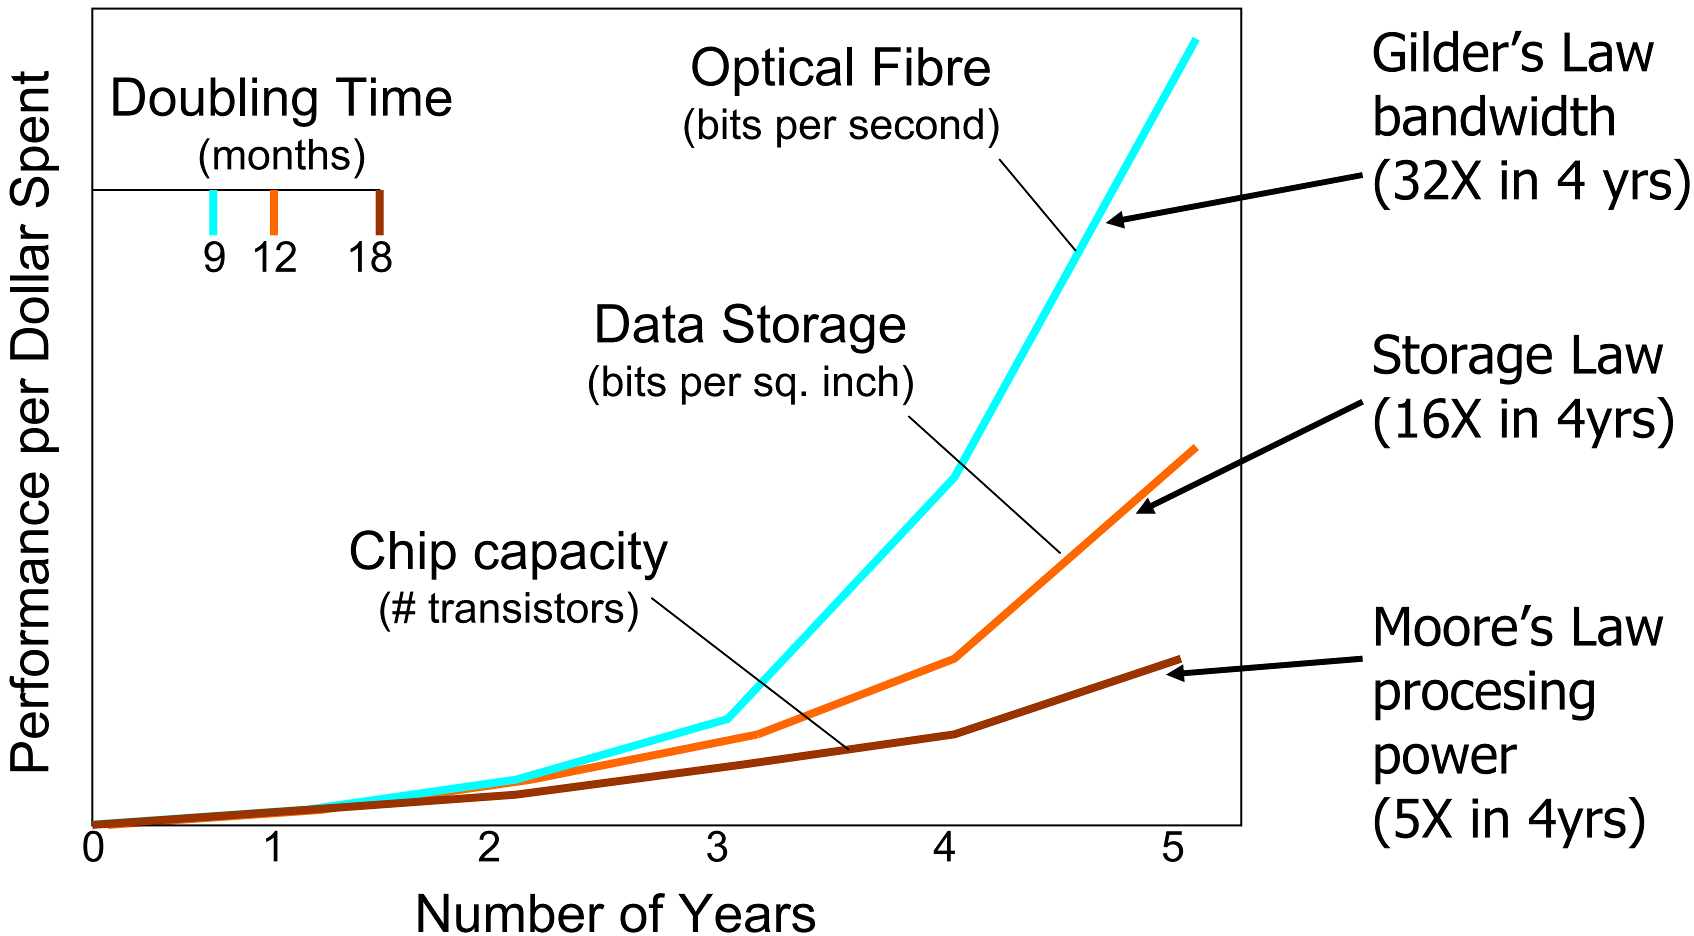

Triumph of Light – *Scientific American*. George Stix, January 2001

High-end example: The LHC Experiments

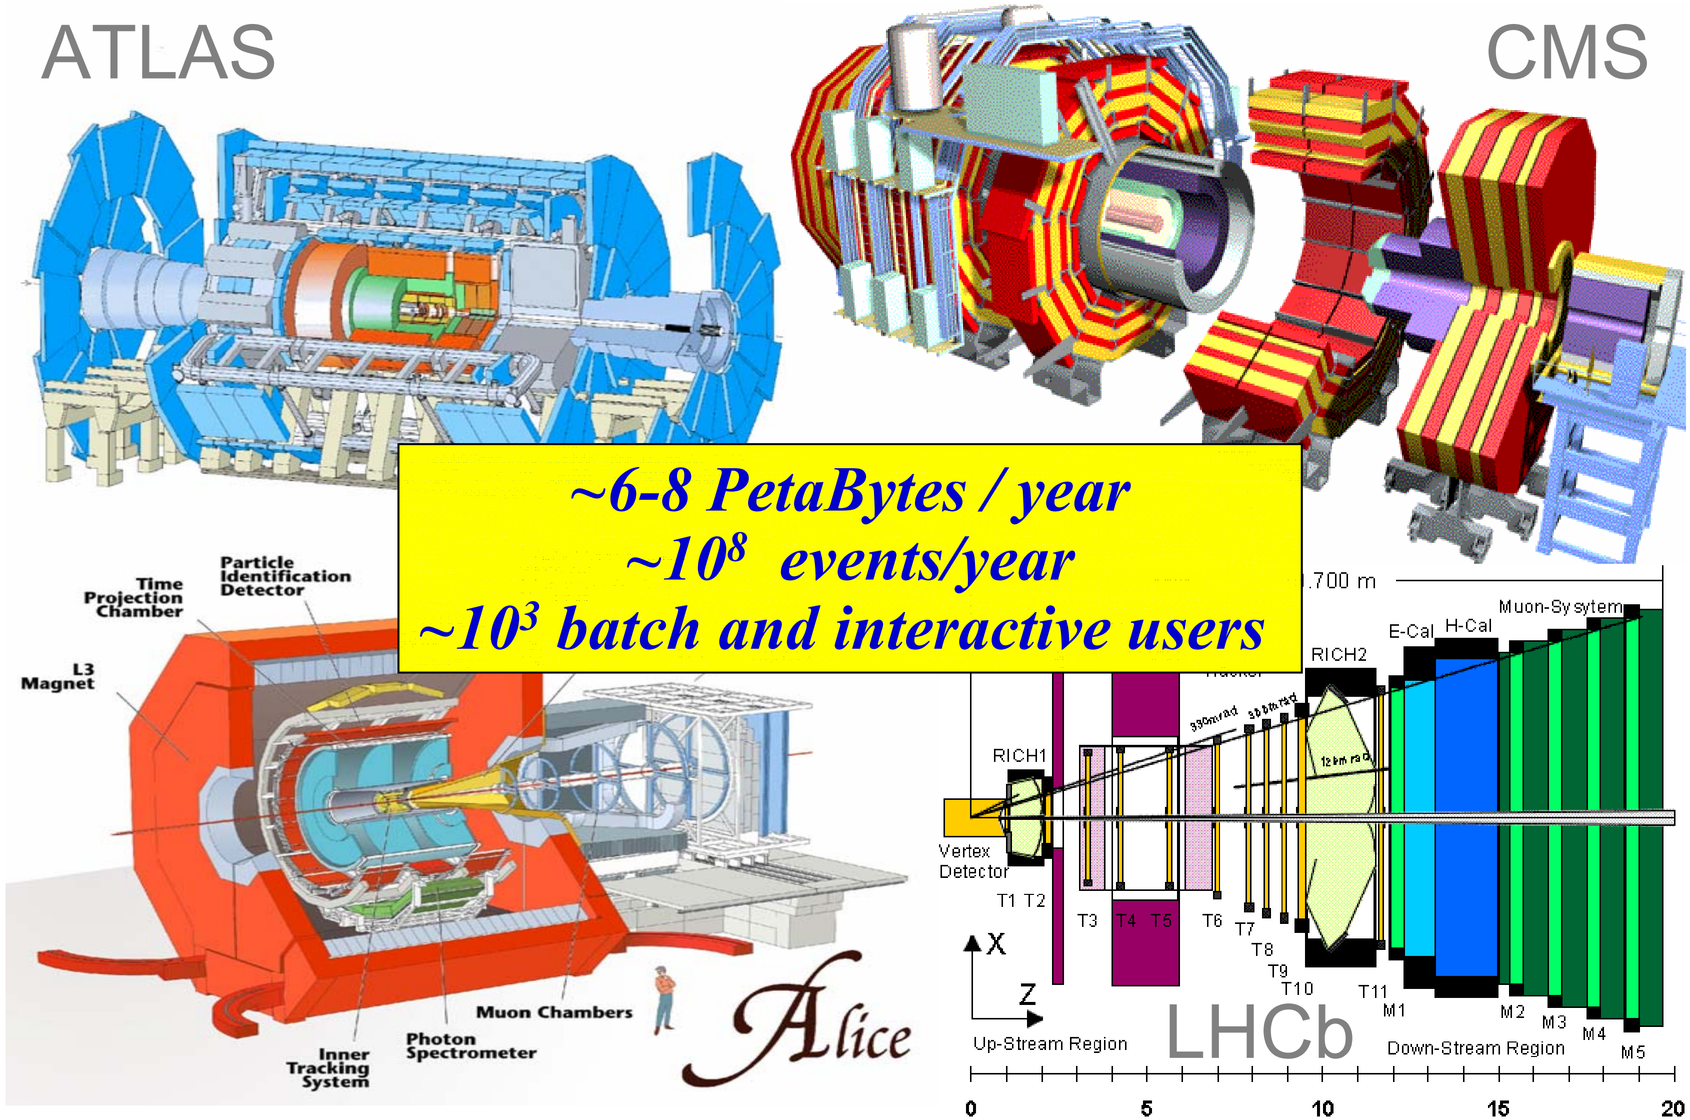

~6-8 PetaBytes / year

~10^8 events/year

~10^3 batch and interactive users

Just a comparison…

6-8 Petabytes
~10.000.000 CD-ROM

Produced each year

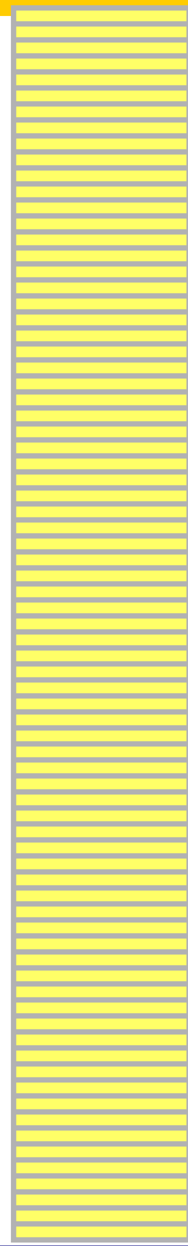

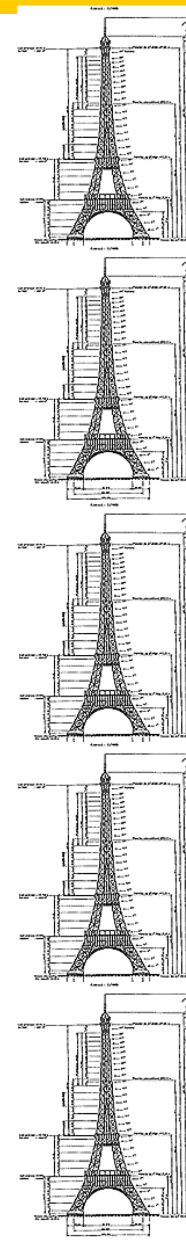

5 times the
Eiffel Tower
~1500 m

Enabling Grids for E-science in Europe

Grids in e-Science

- Grids provide *access* to:
  - Very large data collections
  - Terascale computing resources
  - High performance visualisation
  - Connected by high-bandwidth networks
- Grids support global *collaborations* enabled by the internet
  - increasingly how science is done
  - Necessary to integrate the information to yield understanding
- e-Science is more than Grid Technology

It is what you do with it that counts

Challenges

- Must share data between thousands of scientists with multiple interests
- Must ensure that all data is accessible anywhere, anytime
- Must be scalable and remain reliable for more than a decade
- Must cope with different access policies
- Must ensure data security

Enabling Grids for E-science in Europe

Before Grids

Researchers in many locations need to interact with colleagues, share and access data

FTP, telnet, blood, sweat and tears… and little support for collaboration

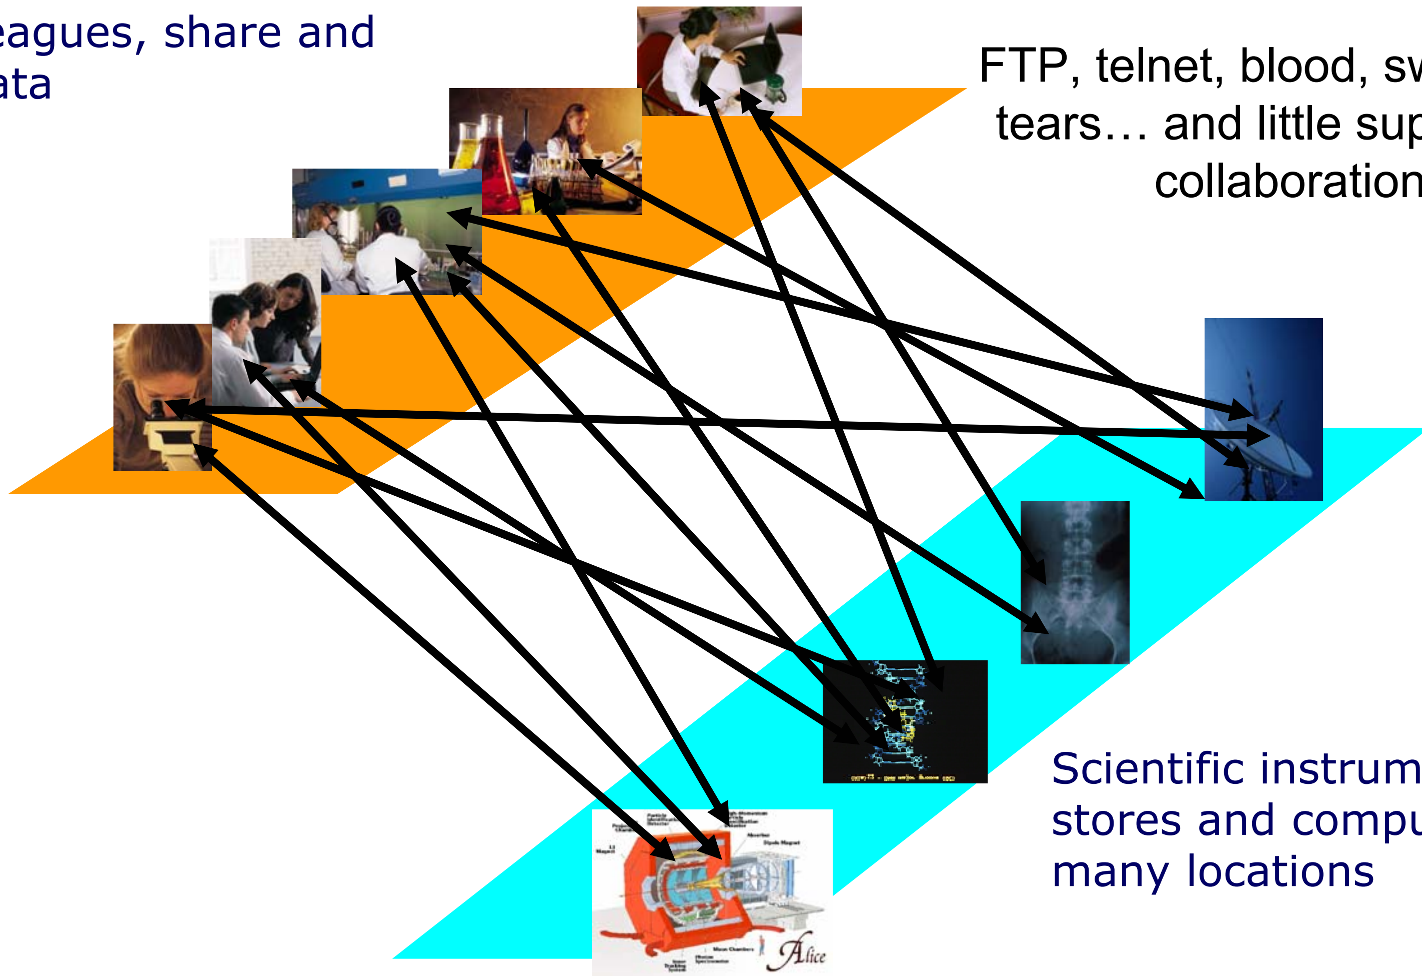

Scientific instruments, data stores and computers in many locations

The Grid Vision

Researchers in many locations need to interact with colleagues, share and access data

The Grid: enables sharing of resources and collaboration

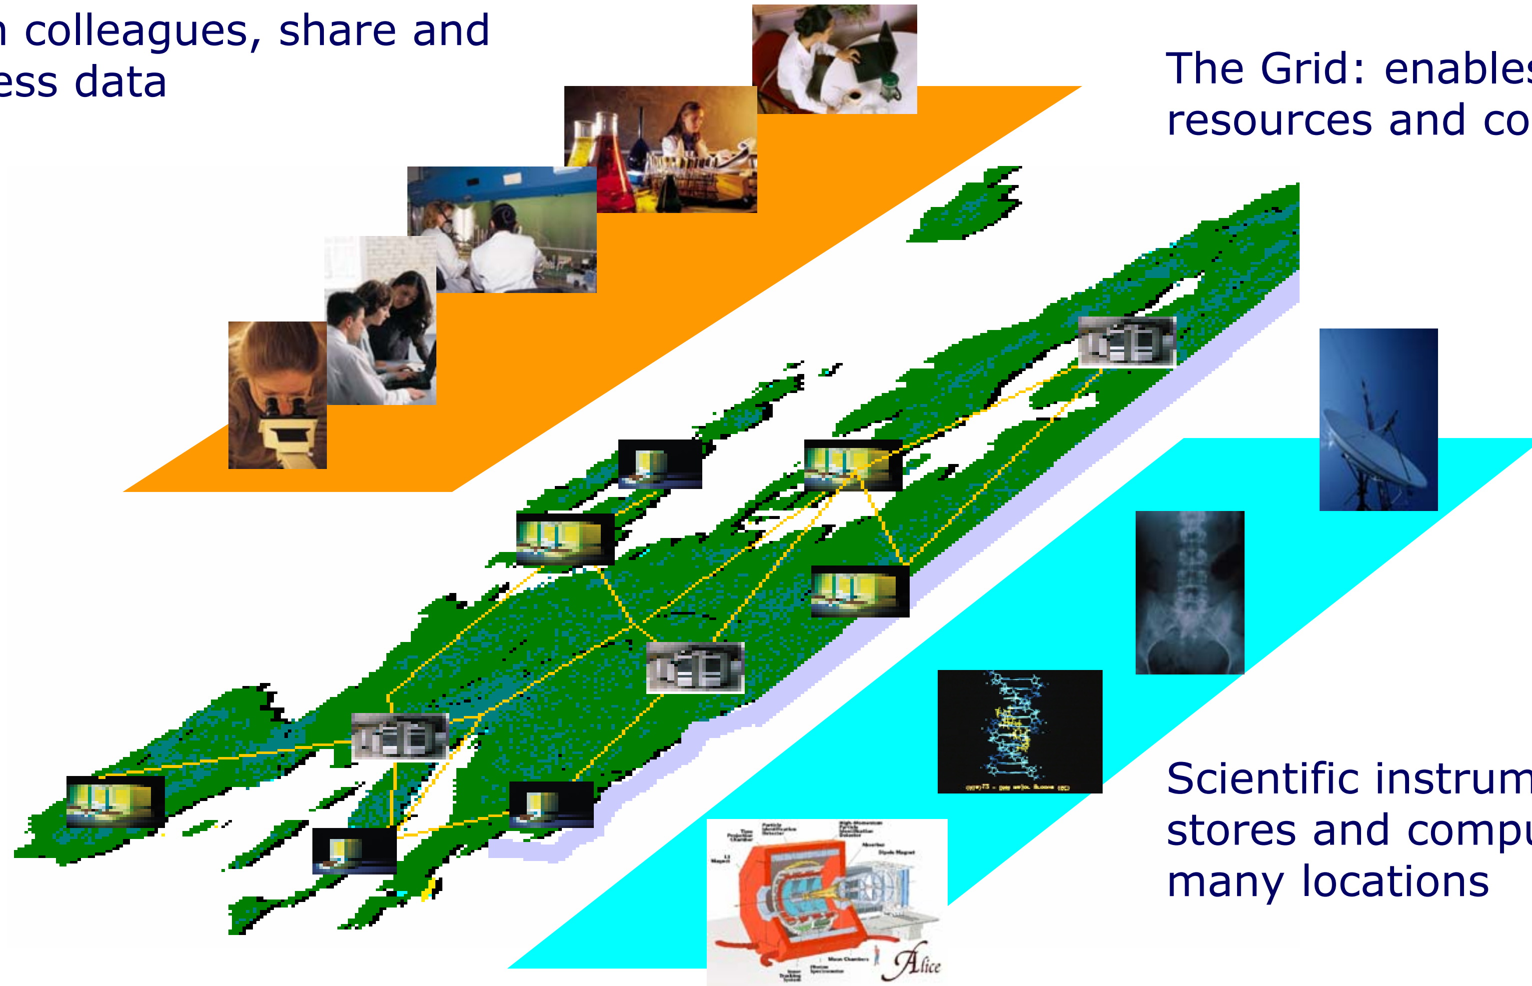

Scientific instruments, data stores and computers in many locations

Slide derived from EDG / LCG tutorials

The Emergence of Global Knowledge Communities

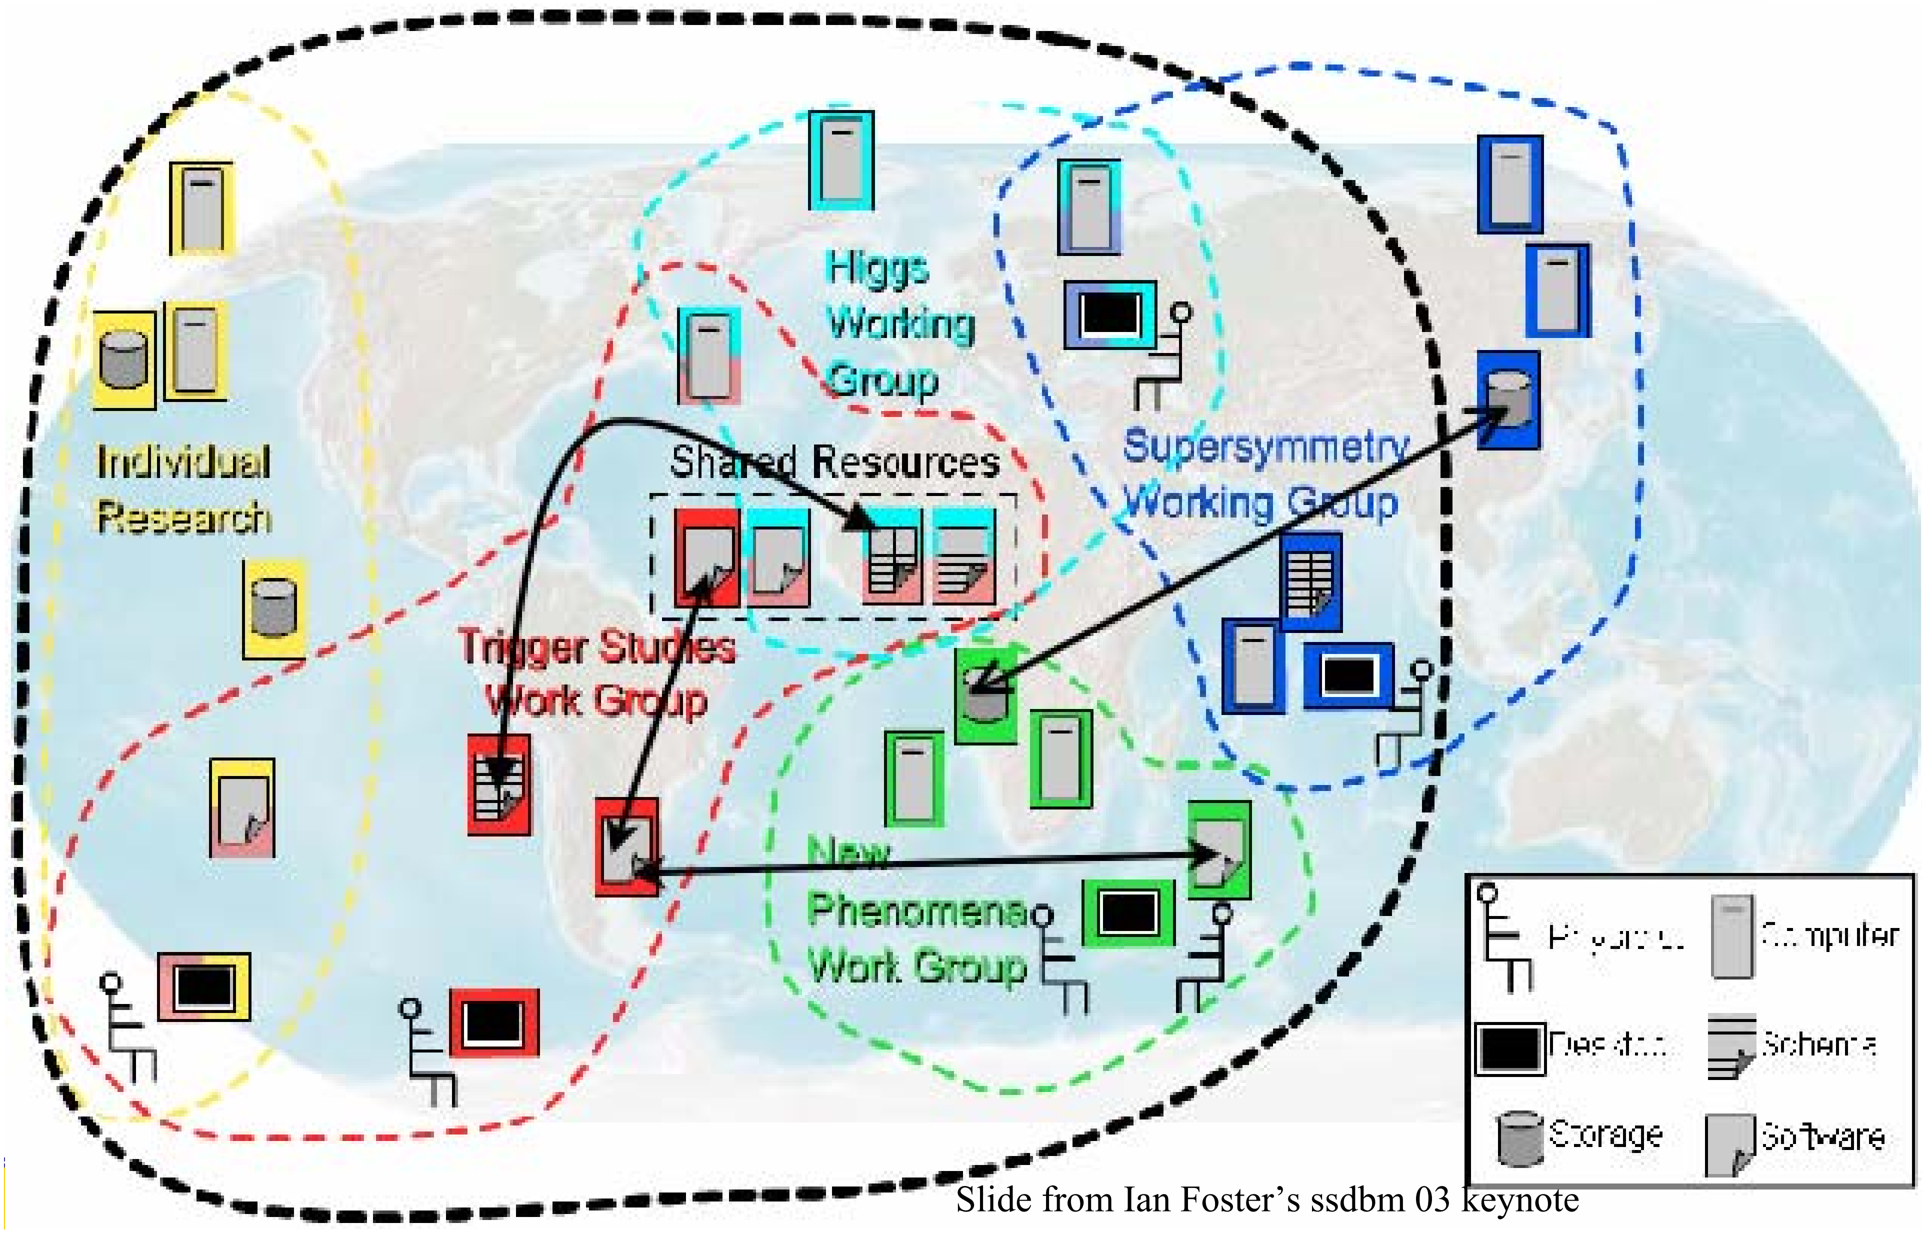

Slide from Ian Foster's ssdbm 03 keynote

eScience Applications

(What's in it for working scientists?)

Grid Applications

- High Energy Physics *(analysising the results from particle collisions)*
- Medical/Healthcare *(imaging, diagnosis and treatment)*
- Bioinformatics *(study of the human genome and proteome to understand genetic diseases)*
- Nanotechnology *(design of new materials from the molecular scale)*
- Engineering *(design optimization, simulation, failure analysis and remote Instrument access and control)*
- Natural Resources and the Environment *(weather forecasting, earth observation, modeling and prediction of complex systems)*

CERN: Data intensive science in a large international facility

- The Large Hadron Collider (LHC)
  - The most powerful instrument ever built to investigate elementary particles physics
- Data Challenge:
  - 10 Petabytes/year of data !!!
  - 20 million CDs each year!
- Simulation, reconstruction, analysis:
  - LHC data handling requires computing power equivalent to ~100,000 of today's fastest PC processors!

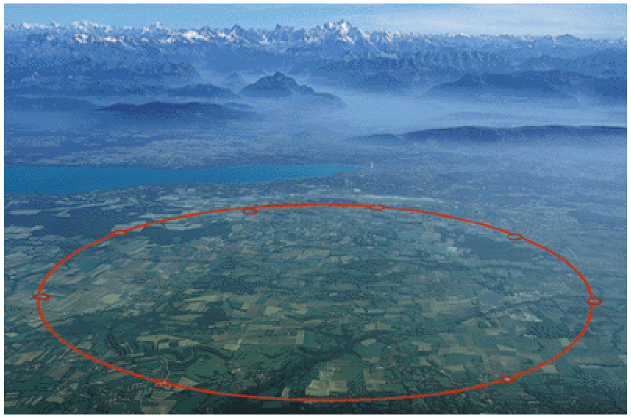

Grid Applications: art

the Thomson flat scanner
developed in 1990

140,000 photo-archives
digitised in
6.000 dots x 8.000 lines
in 5 years (1996-2001)

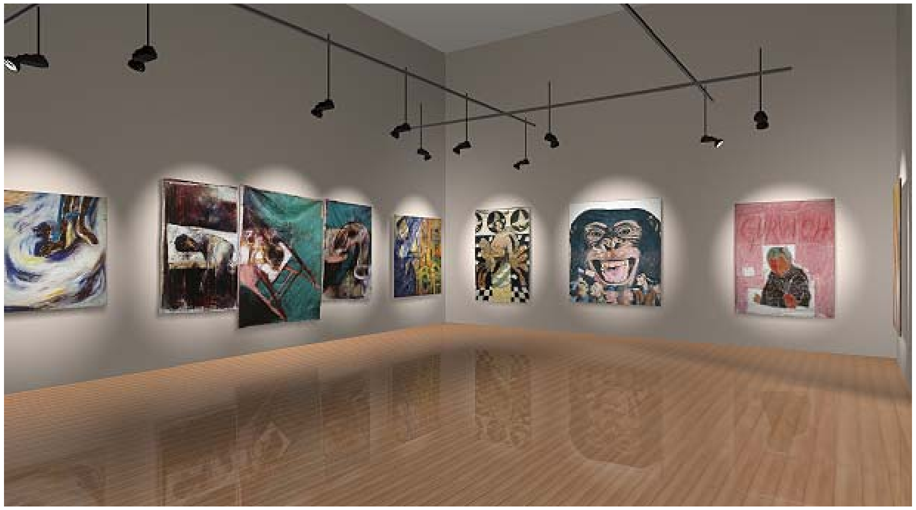

Museo Virtual de Artes El Pais (MUVA)
http://www3.diarioelpais.com/muva/.

Books are being scanned in at
767 MB per page
1/2 Terabyte for Gutenberg Bible

Paintings are being scanned in at
30 GB each
in the EU CRISATEL Project

Grid Applications: Engineering

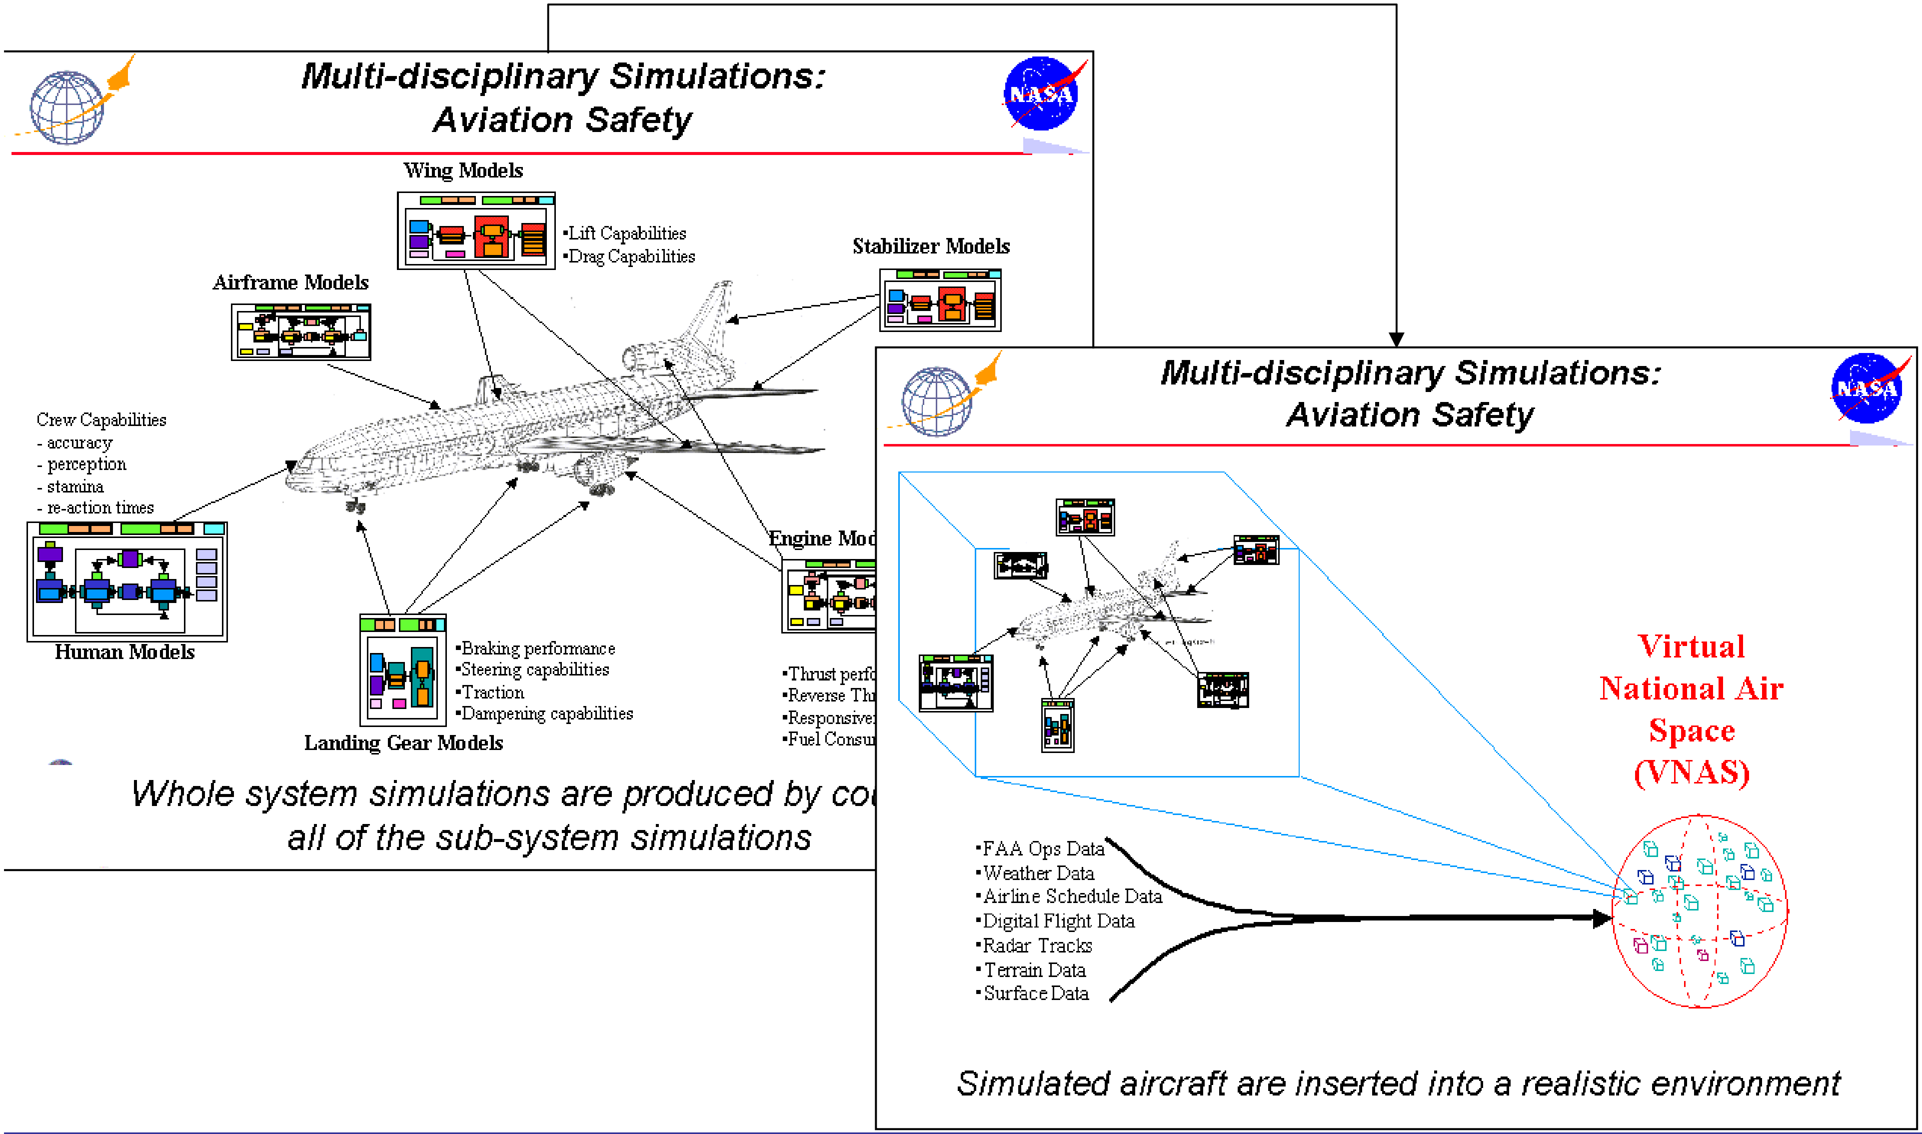

Grid Applications: Earth Observation

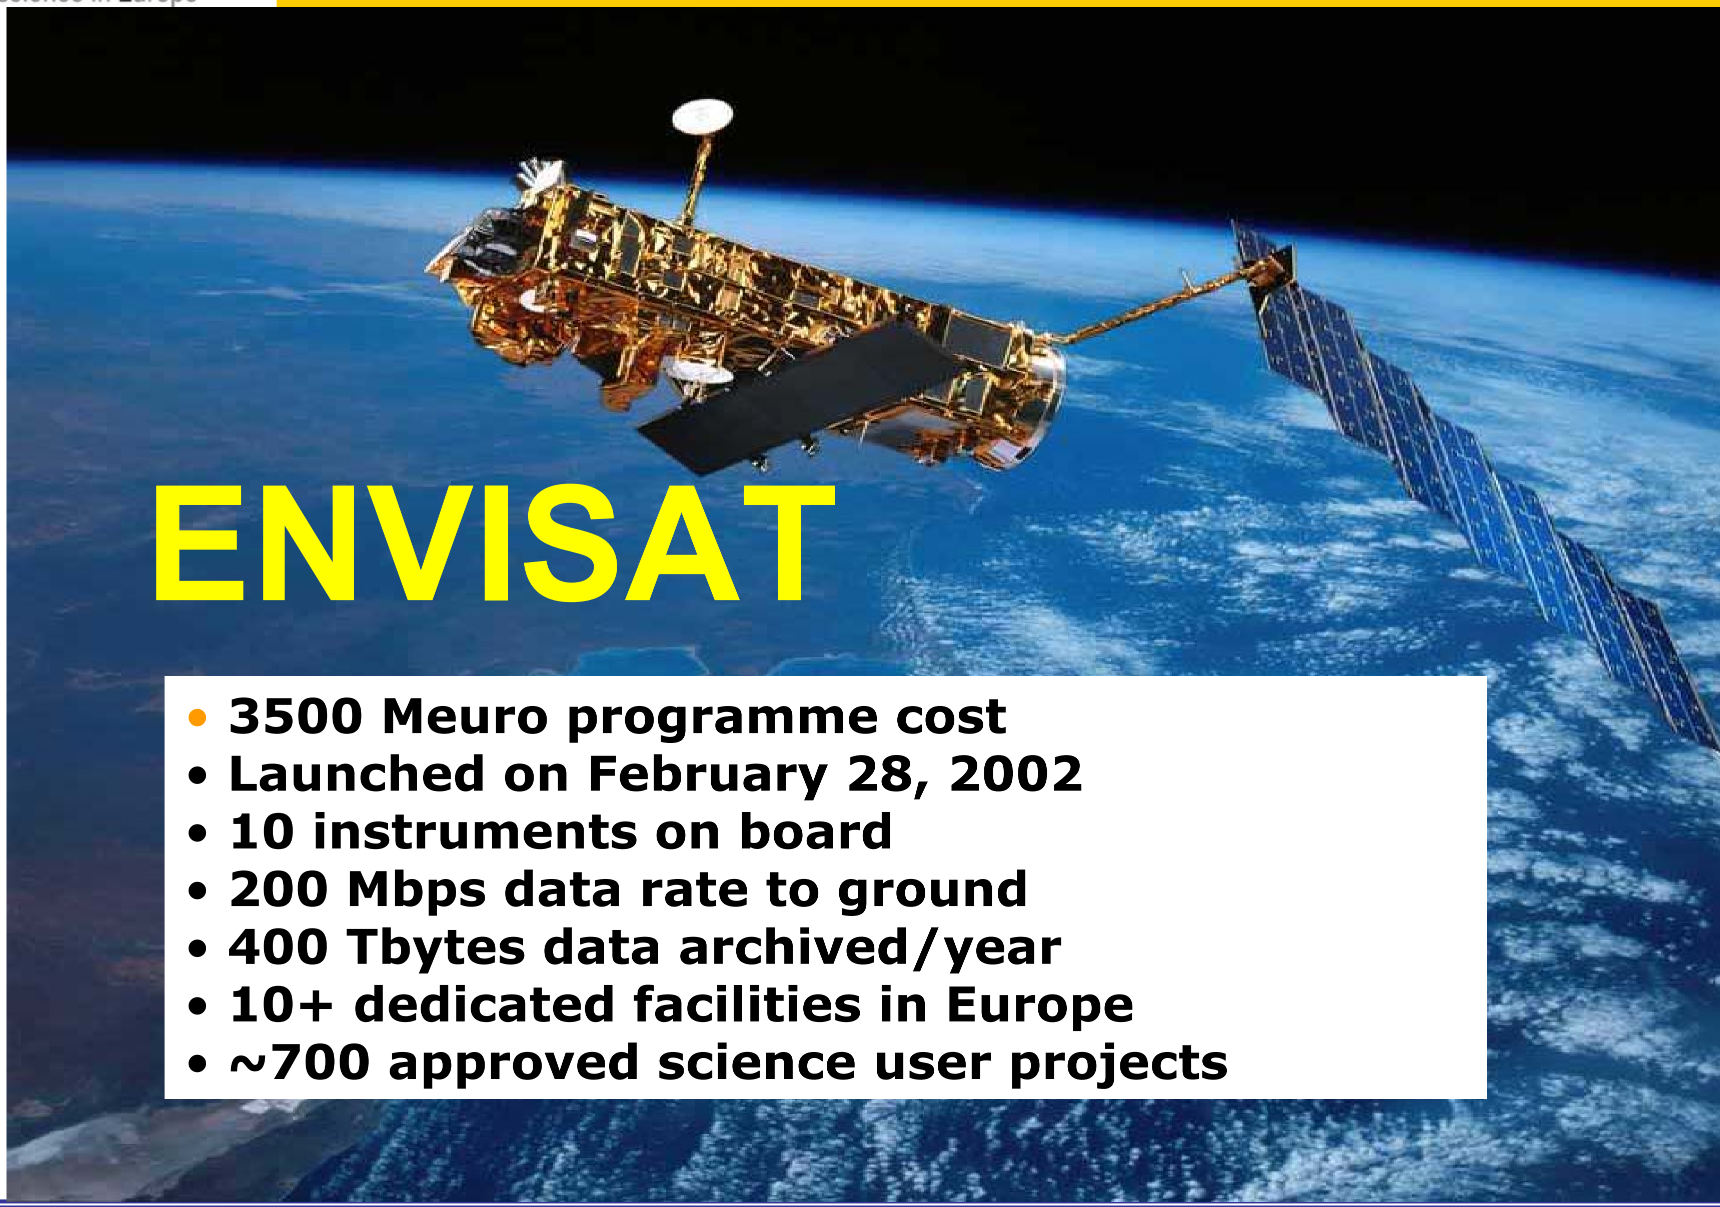

- 3500 Meuro programme cost
- Launched on February 28, 2002
- 10 instruments on board
- 200 Mbps data rate to ground
- 400 Tbytes data archived/year
- 10+ dedicated facilities in Europe
- ~700 approved science user projects

CrossGrid

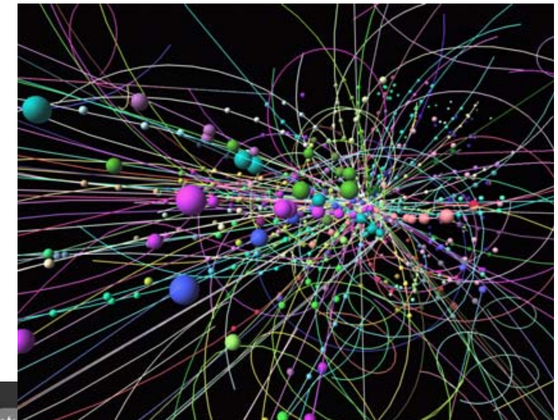

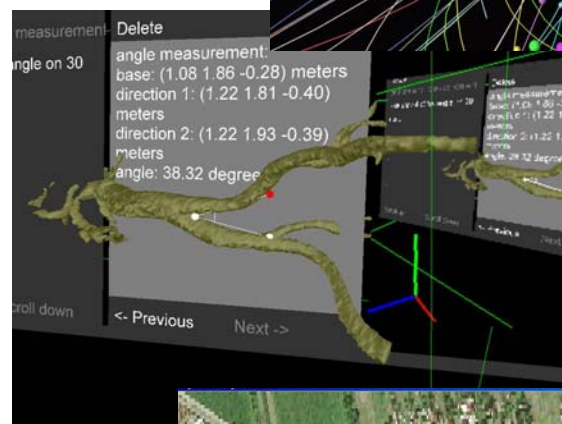

- 1. Interactive biomedical simulation and visualization
- 2. Flooding crisis team support
- 3. HEP distributed data analysis
- 4. Weather forecasting and air pollution modeling

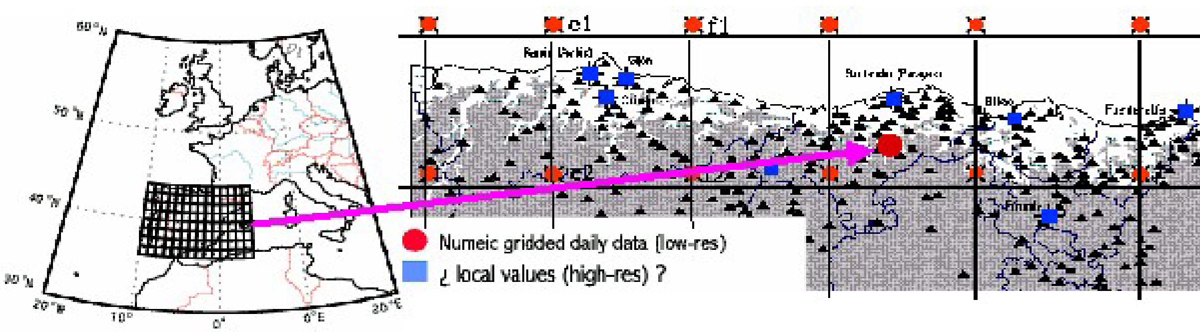

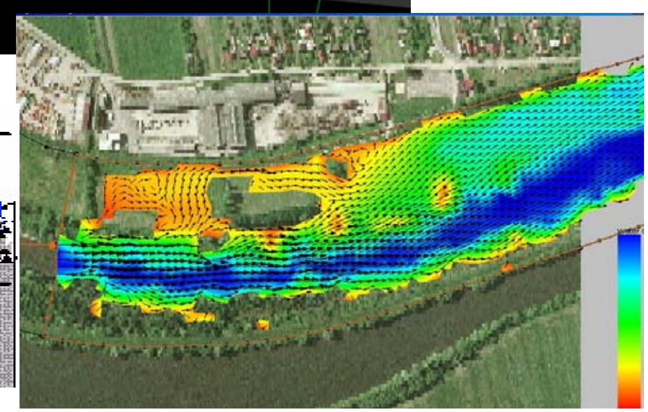

Connecting *People*: Access Grid

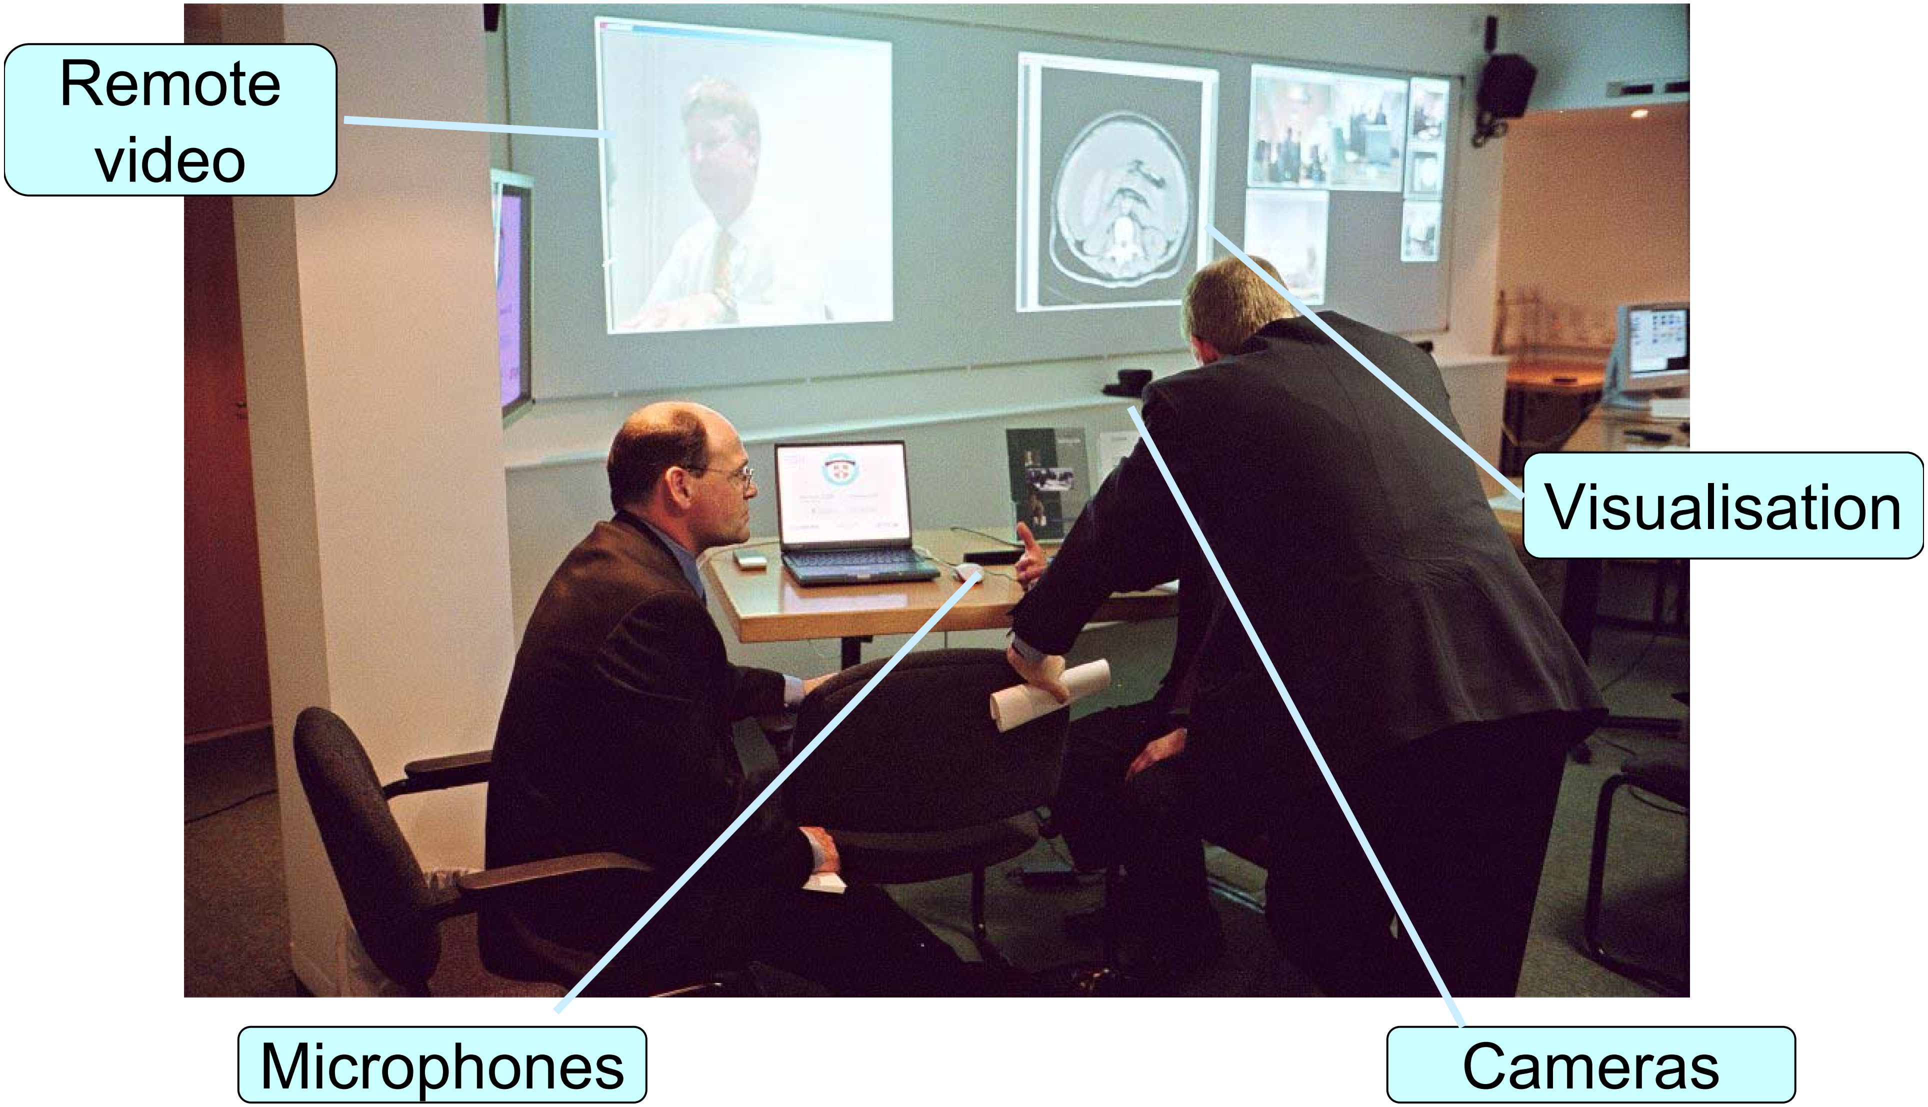

What are grids really like?

Different people give different views

This is just one perspective on the question

Specialised Distributed Computing

Grids come within the general field of

- Parallel / Distributed Computing

A computational task involves coordination of components which occur

- **Simultaneously**

And/or

- **At physically separated locations**

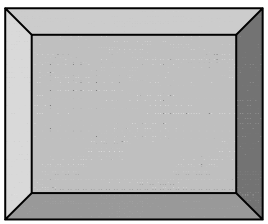

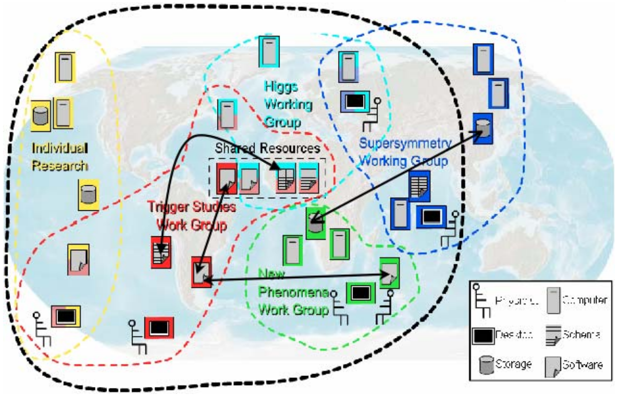

Parallel ⟷ **Distributed**

Multi-processor | **Processor Pool** | **Embedded Systems** | **Flexible Manufacturing** . . . | **GRIDS** | . . . ?

Definitions of Grids

There are a number of different perspectives on the term “Grid” with different definitions

A very inclusive formulation of “Grid Computing” is

- Coordination of computational components
- Of up to international level of geographic separation
- **Crossing organisation boundaries**

Different "Grid" Perspectives

Parallel Computing ⟷ Distributed Computing

↓

GRIDS: Multi-organisational

E-commerce perspective:
• On-demand renting of remote computing resources
• via the Grid – a subset of the web

E-science perspective
• Collaborative use of multiply-owned resources
• on a grid - a persistent distributed environment

- We are taking an E-science Perspective
- Central to that is Virtual Organisations – **VO** s

Grids and Virtual Organisations

- A *Virtual Organisation (VO)* is:
  - People from different institutions working to solve a common goal
  - Sharing distributed processing and data resources
  - Model is most closely that of a Scientific community

- Grid infrastructure enables virtual organisations

  > *"Grid computing is coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations"*
  > (I.Foster)

- A grid supports multiple VOs
  - Membership of a VO determines which grid resources you can use

Enabling Grids for E-science in Europe

What are the characteristics of a Grid?

Numerous Resources

Ownership by Mutually Distrustful Organizations & Individuals

Connected by Heterogeneous, Multi-Level Networks

Different Security Requirements & Policies Required

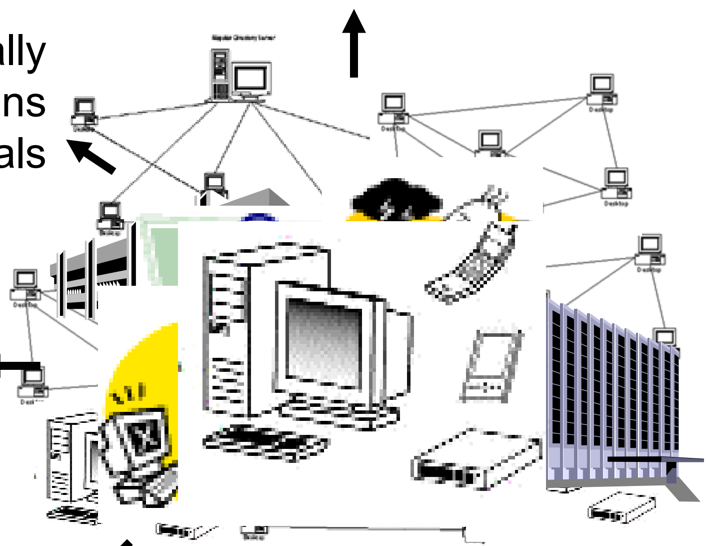

Different Resource Management Policies

Potentially Faulty Resources

Geographically Separated

Resources are Heterogeneous

What are the characteristics of a Grid?

Numerous Resources

Ownership by Mutually Distrustful Organizations & Individuals

Connected by Heterogeneous, Multi-Level Networks

Different Security Requirements & Policies Required

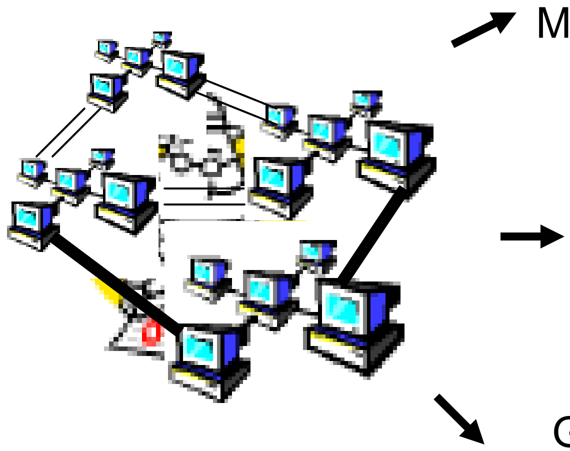

Different Resource Management Policies

Potentially Faulty Resources

Geographically Separated

Resources are Heterogeneous

Requirements on Grid Design

Need a framework for

- Controlled sharing of dynamically changing, replicated resources which accommodates a rapidly evolving diversity of applications and resources.

- “Middleware” is what provides that framework –
- Must meet requirements for
  - Changing physical resources
  - Replicated logical resources
  - Controlled sharing
  - Evolving diversity

Dynamic Physical Resources

A full Grid Architecture accommodates –

- Continually changing physical resources
  - Existence –
    changes in the set of resources that exist
    - Partial failure – there is always something not working
    - Different maintenance schedules
    - Autonomous addition/removal of resources
  - Capacity -
    Changes in what a particular resource can offer
    - Multiple independent users within the grid
    - A resource may also service local non-grid users
      - E.g. offering the grid "spare capacity"

Change occurs autonomously – no central control on this process

Enabling Grids for E-science in Europe

Replication of Logical Resources

A full Grid Framework accommodates –

- Replication of Logical Resources
  - Data – there may be multiple copies of the same data resource
    - to allow efficient multi-user access
  - Processing – the processor pool concept
    - A logical processing resource is a processing capability with particular characteristics (e.g. processing power, software available)
    - A major motivation is to allow polling of processing resources, such that a particular computation can use whatever is available, by expressing its logical processing resource needs
  - Deal with Partial Availability
    - Always some physical resource is unavailable
    - So a logical resource must be accessible via multiple physical resources

Controlled Sharing

A full Grid Framework addresses various access control issues, e.g.

- Different organisations have different access control policies
- Need to provide the end user with a simple interface,
  - E.g. single sign-on
- The high risk of security breaches
- Dynamic VO membership
- Although the scientists may trust each other, the systems administrators are wary of ousiders gaining acccess to their resources

- These issues are dealt with in a subsequent talk

Evolving Diversity

A full Grid Framework accommodates heterogeneity …

At any point in time there are many differences

- Heterogeneous multi-level networks
  - Performance characteristics – latency and bandwidth
  - Protocols used
- Heterogeneous computing resources; e.g. one "resource" might be
  - A single processor work station
  - A multi-processor super-computer
  - Even, a grid!
- Heterogeneous applications
  - Different application areas will use the facilities in very different ways

… and change

- There will be many developments, relatively un-coordinated
- Some applications are long-lived (10 years) and need a stable environment over that period

The Middleware Challenge

- Challenge is to support stable collaboration within evolving diversity
- What is needed is a “middleware” architecture
- To provide
  - Conceptual frameworks
  - Standard Interfaces
- Which
  - Are usable by a wide range of applications
  - Allow effective exploitation of a wide range of resource capabilities –
    - Processing, data and network
- And
  - Are as simple as possible
  - But no simpler
- A funnel between applications and resources

The Middleware Funnel

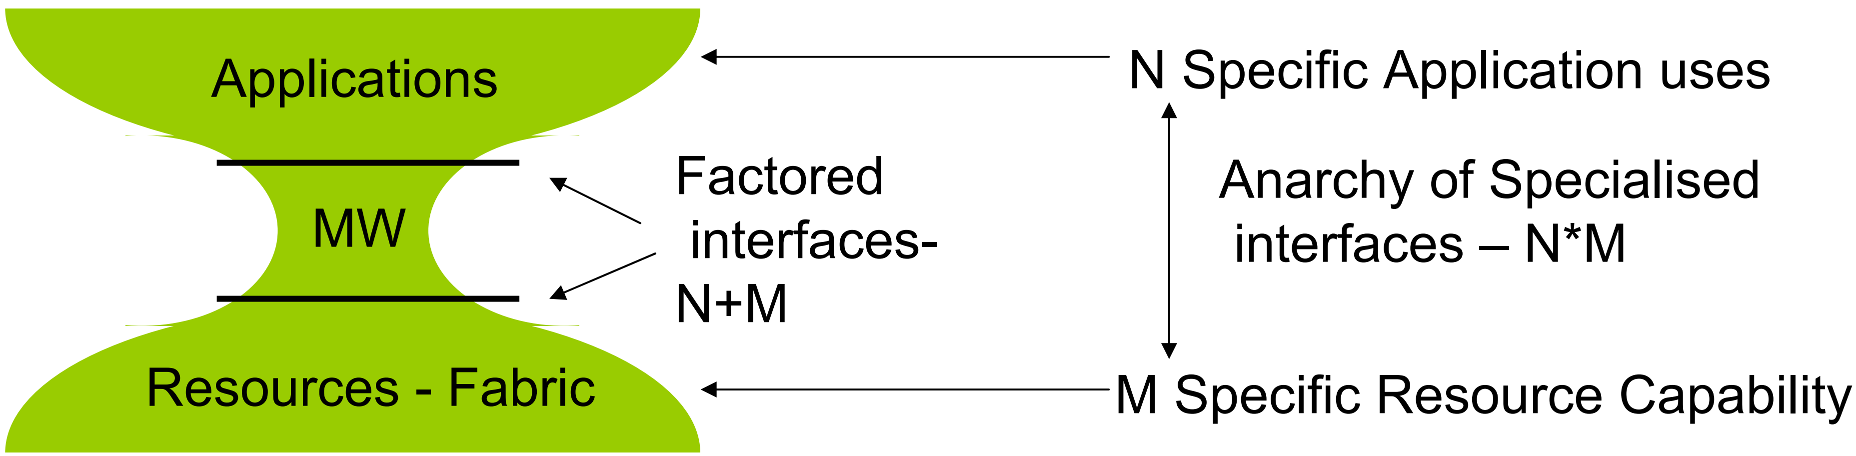

- For a Particular set of Application needs and resource capabilities
- Develop a middleware layer which provides a standard way for the applications to access the resource capabilities – N*M becomes N+M

Middleware Ageing

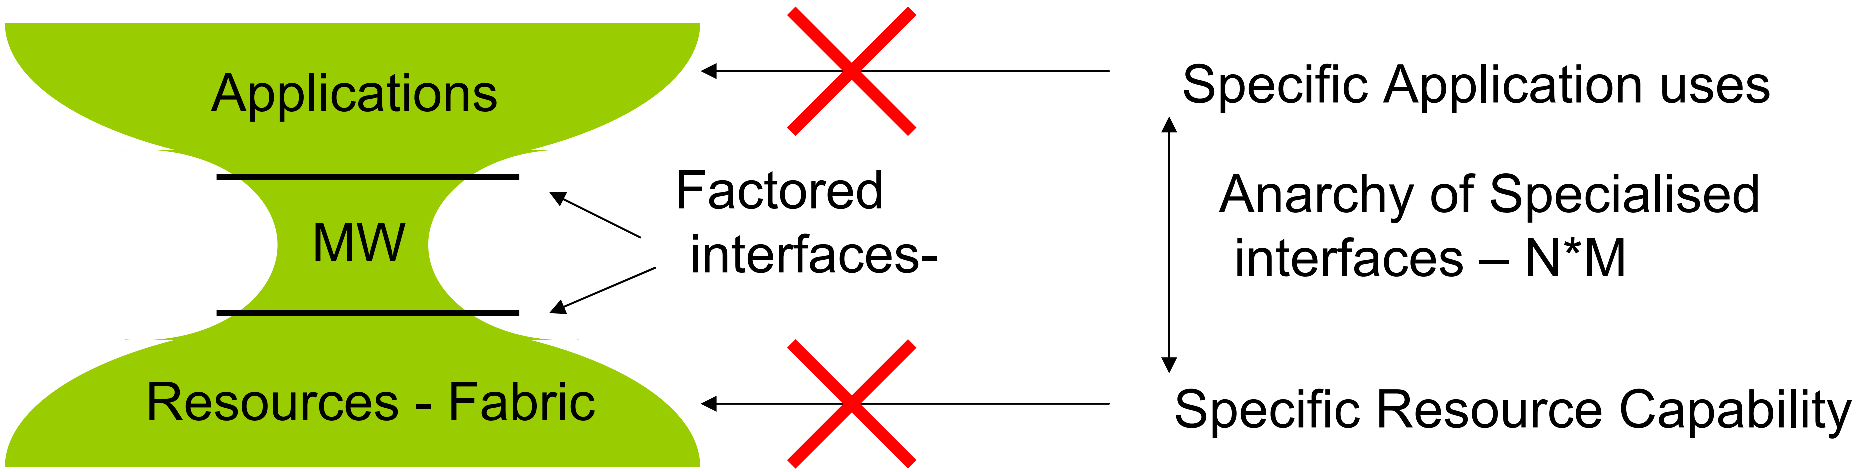

- Eventually there will be application needs and/or matching resource capabilities that can't pass through the funnel effectively
  - In a different domain which we want to integrate
  - or developing over time
- Gives rise to specialised interfaces which start to compromise principles of cooperative use of heterogeneous resources
- A middleware architecture attempts to be general enough to prevent that, but eventually fails …

Middleware Ecology

A way to understand the grid world

- There are a multiplicity of **components**,
- each provides functions for some aspect of grid requirements, e.g.
  - Access control
  - Data replication
  - Matching processing requirements with resources
  - Obtaining information about resource availability
- There are middleware products putting together various components, either
  - Taken from other products – possibly re-engineered
  - Developed specifically for this product
- Eventually any product falls into relative disuse

  but may partially live on by contributing component to a new packaging

A complex situation of many components combined in different ways

EGEE BACKGROUND

- What does EGEE build on?

  Previous Products and Projects

EGEE Middleware

EGEE includes the development of a new middleware product

- Brings together the best of existing components
- Re-engineers them to
  - Fit into a common architectural framework
  - To provide an overall system which is of production quality

- Enables co-operation across existing European grid installations

- Firmly embedded in and evolving out of existing European and other international grid projects

So now some background on that context

eGee Enabling Grids for E-science in Europe

Virtual Data Toolkit

- Grid Middleware components from several projects
  - Packaged and tested together
  - Foundation of EGEE/ LCG

- Globus Toolkit
- Condor
- Chimera
- EDG & LCG tools
- NCSA Tools
- Other Tools

Globus Toolkit

- Grid Security Infrastructure (GSL)
  - X.509 authentication with delegates and single sign-on
- Grid Resource Allocation Mgmt (GRAM)
  - Remote allocation, reservation, monitoring, control of compute resources
- GridFTP protocol (FTP extensions)
  - High-performance data access & transport
- Grid Resource Information Service (GRIS) +
  Monitoring and Discovery Service (MDS)
  - Access to structure & state information
- XIO
  - TCP, UDP, IP multicast, and file I/O
- Others…

Condor

- "Cycle-stealing"
  - Use idle CPU cycles for productive work
- "High Throughput Computing"
  - Using all available compute power over periods of days, weeks,…
  - "Embarrassingly parallel" problems
- Fault tolerance
  - Algorithms must allow for failure
  - Checkpointing and process migration
- DAGMan
  - Workflow specification

Chimera

- Technology for collaborative management of data, programs & computations
- Virtual data system
  - Virtual data catalog
  - Virtual data language
  - Automated data derivation
  - Provenance tracking
- Pegasus
  - AI planning system for Grid workflows

Tools

- NCSA
  - MyProxy
  - GSI OpenSSH
- EDG & LCG
  - Make Gridmap (Authorisation control)
  - Certificate Revocation List Updater
  - GLUE Schema (Monitoring)
- Others
  - VDT System Profiler
  - Configuration software
  - KX509 (X.509 <-> Kerberos)

Previous Major EU GRID projects

European DataGrid (EDG)	www.edg.org
LHC Computing GRID (LCG)	cern.ch/lcg
CrossGRID	www.crossgrid.org
DataTAG	www.datatag.org
GridLab	www.gridlab.org
EUROGRID	www.eurogrid.org

European National Projects:

- INFNGRID,
- UK e-Science Programme,
- NorduGrid

EU DataGrid at a glance

People

500 registered users

12 Virtual Organisations

21 Certificate Authorities

>600 people trained

456 person-years of effort
170 years funded

Application Testbed

~20 regular sites

> 60,000 jobs submitted (since 09/03, release 2.0)

Peak >1000 CPUs

6 Mass Storage Systems

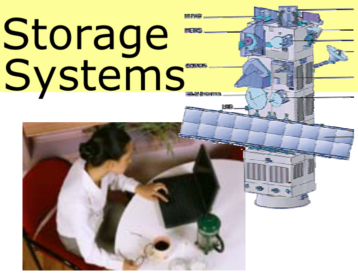

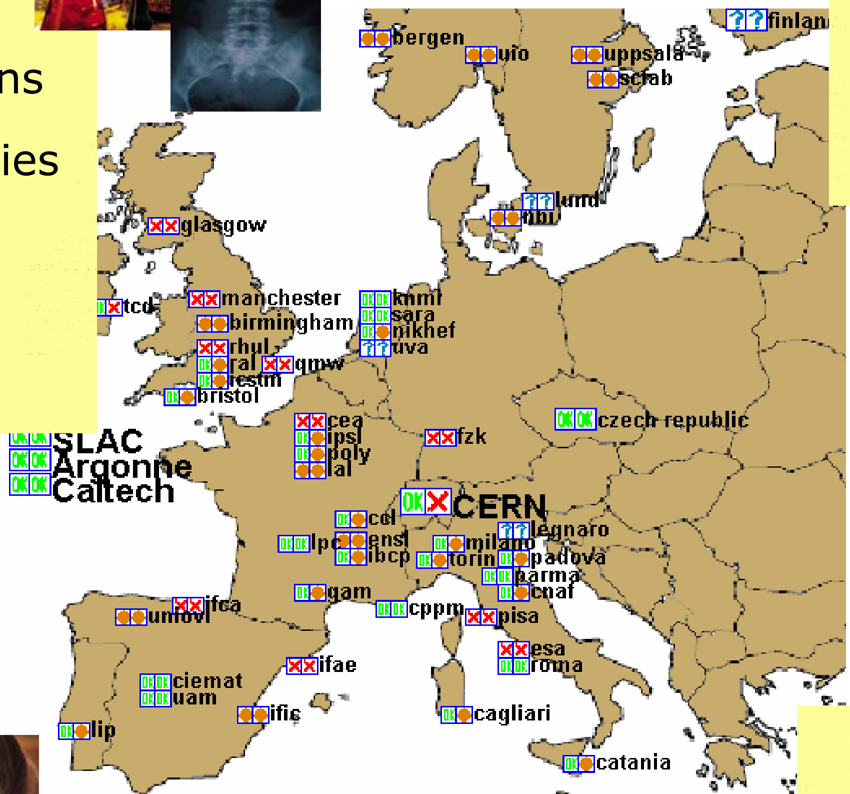

Software

> 65 use cases

7 major software releases (> 60 in total)

> 1,000,000 lines of code

Scientific Applications

5 Earth Obs institutes

10 bio-medical apps

6 HEP experiments

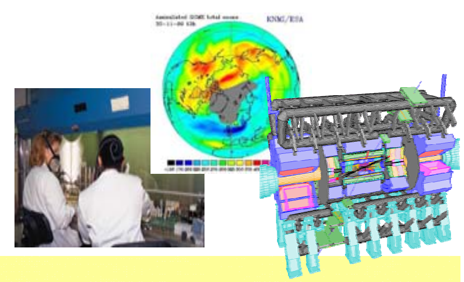

Grid projects

Many Grid development efforts — all over the world

- •NASA Information Power Grid
- •DOE Science Grid
- •NSF National Virtual Observatory
- •NSF GriPhyN
- •DOE Particle Physics Data Grid
- •NSF TeraGrid
- •DOE ASCI Grid
- •DOE Earth Systems Grid
- •DARPA CoABS Grid
- •NEESGrid
- •DOH BIRN
- •NSF iVDGL

- •DataGrid (CERN, ...)
- •EuroGrid (Unicore)
- •DataTag (CERN,…)
- •Astrophysical Virtual Observatory
- •GRIP (Globus/Unicore)
- •GRIA (Industrial applications)
- •GridLab (Cactus Toolkit)
- •CrossGrid (Infrastructure Components)
- •EGSO (Solar Physics)

- •UK – OGSA-DAI, RealityGrid, GeoDise, Comb-e-Chem, DiscoveryNet, DAME, AstroGrid, GridPP, MyGrid, GOLD, eDiamond, Integrative Biology, …
- •Netherlands – VLAM, PolderGrid
- •Germany – UNICORE, Grid proposal
- •France – Grid funding approved
- •Italy – INFN Grid
- •Eire – Grid proposals
- •Switzerland - Network/Grid proposal
- •Hungary – DemoGrid, Grid proposal
- •Norway, Sweden - NorduGrid

Summary


COLABORATION
COLABORATION
Display
FLEXIBILITY
Store data
Internet
Create data
Process data
Standards