Future of Grid Computing

Dr Richard Sinnott

http://csperkins.org/teaching/2004-2005/gc5/

UNIVERSITY
of
GLASGOW

Future of Grid Computing Overview

- Classifications of Grid Computing
  - Future and challenges of different classifications
- Technology and standards gaps
  - OGSA
  - WSRF
  - WS-?
  - Security
- Key application domains for the future
- Funding and industry

Classifications of Grid Computing

- Compute Grids
  - Run multiple jobs with distributed compute and data resources
- Data Grids
  - Grids focused on finding, managing and generally dealing with large distributed data resources
- Complexity Grids
  - Hybrid combination of Data and Compute Grid
- Campus Grids
  - Grid supporting University community computing
- Enterprise Grids
  - Grid supporting a company's enterprise infrastructure
- Semantic Grids
  - Integration of Grid and Semantic Web meta-data and ontology technologies
- Lightweight Grids
  - Grid designed for rapid deployment and minimum maintenance
- Collaboration Grids
  - Grid supporting collaborative tools like the Access Grid, whiteboard and shared application.
- Autonomic Grids
  - Fault tolerant and self-healing Grids
  - Others possible too, e.g. P2P Grids, …

Classifications of Grid Computing

- Compute Grid
- Data Grid
- Complexity Grid
- Campus Grid
- Enterprise Grid
- Semantic Grid
- Peer-to-peer Grid
- Lightweight Grid
- Collaboration Grid
- Autonomic Grid

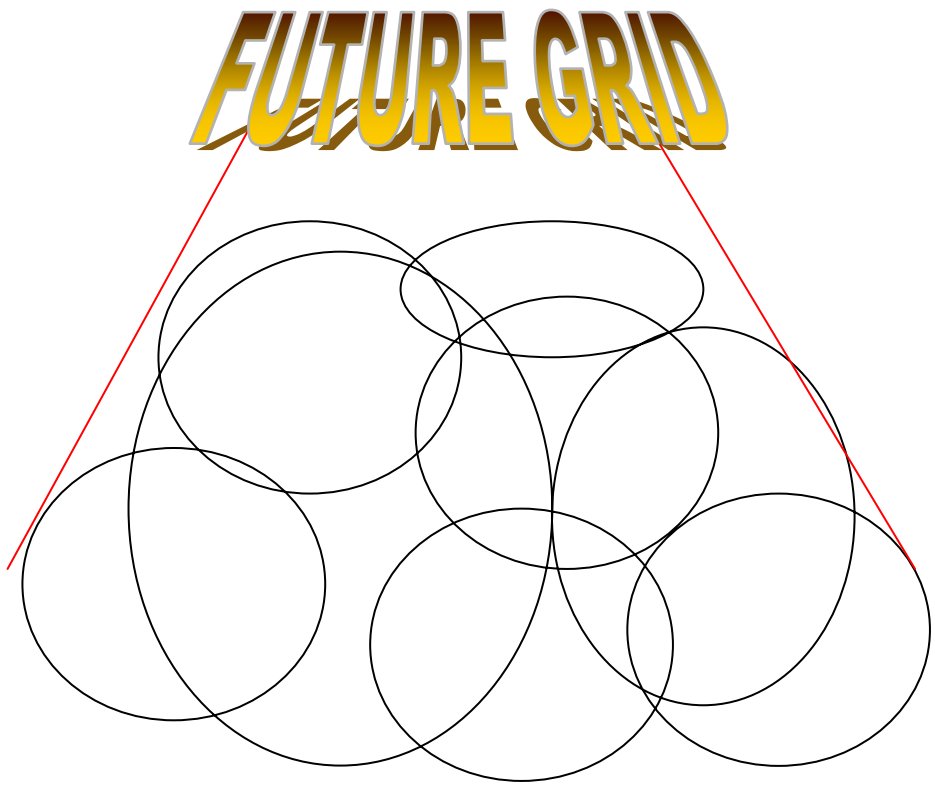

Compute Grid Future

- Compute Grid
  - Run multiple jobs with distributed compute and data resources
    - Aspects of Compute Grids well supported, e.g. Condor
    - Future application domains
      - Physics (monte carlo simulations, statistical physics…)
      - Engineering (simulations, optimising designs, …)
      - Life sciences (high throughput BLASTing, …)
      - …
    - Future of Compute Grids likely to include
      - Robust middleware for
        - » information services
        - » resource brokering services
        - » Meta-job scheduling
        - » fault tolerance, dependability
        - » security
        - » … on heterogeneous infrastructure

Data Grid Future

- Data Grid
  - Grids focused on finding, managing and generally dealing with large data resources
    - Application domains
      - Particle physics
        - » LHC collider
      - Astrophysics
        - » Worldwide virtual observatories
      - Life sciences are likely to form own data grids in the near future
        - » Issue of developing domain
      - **<u>… data is the primary driver for Grids in the future?</u>**

    - Future of Data Grids likely to include
      - robust and stable middleware for accessing and integrating scientific data
      - enhanced services for replication (e.g. faster as per PP domain experiences)
      - advanced and stable meta-data catalogues (per application domain)
      - fast and robust data movement services
      - closer integration with security services
      - data curation and provenance services
        - » more understanding of longevity of data

Complexity Grid Future

- Complexity Grids
  - Hybrid combination of Data and Compute Grid
    - (= Grid vision?)
    - Much to be done to realise these Grids right now
      - Scientific need and agreements (e.g. on naming) coupled with advanced software engineering
    - Note could be argued that can build complexity Grids right now
      - But not without considerable cost
        - » In future should be able to dynamically (and quickly!) establish VOs using both Data and Compute Grid technologies

Campus Grid Future

- To co-ordinate resources and people across campus
  - Likely to involve various faculties
    - At Glasgow it will involve ALL faculties
  - Heterogeneous resources
  - Co-ordination of people critical

- Why?
  - Sharing resources/people is only way forward
    - Cost benefits
    - Resilience of resources, people
  - Being organised is extremely beneficial to the university
  - E-Science = big business
    - Quick and conservative estimate of £13.5M+ monies **TO GLASGOW** from e-Science

Campus Grid at Glasgow

- Consolidation of resources
  - Story started with building around ScotGrid
    - Providing shared Grid resource for wide variety of scientists inside/outside Glasgow
      - HEP, CS, BRC, EEE, …
        - » Target shares established
        - » Non-contributing groups encouraged

- ScotGrid [Disk ~15TB
  CPU ~ 255 1GHz]
- Over 1 million CPU hours completed (June 2004)
- Over 100,000 jobs completed
  - Includes time out for major rebuilds
- Typically running at ~90% usage

Hardware

- **59 IBM X Series 330 dual 1 GHz Pentium III with 2GB memory**
- **2 IBM X Series 340 dual 1 GHz Pentium III with 2GB memory**
- **3 IBM X Series 340 dual 1 GHz Pentium III with 2GB memory and 100 + 1000 Mbit/s ethernet**
- **1TB disk**
- **LTO/Ultrium Tape Library**
- **Cisco ethernet switches**

New..

- **IBM X Series 370 PIII Xeon with 32 x 512 MB RAM**
- **5TB FastT500 disk 70 x 73.4 GB IBM FC Hot-Swap HDD**
- **eDIKT 28 IBM blades dual 2.4 GHz Xeon with 1.5GB memory**
- **eDIKT 6 IBM X Series 335 dual 2.4 GHz Xeon with 1.5GB memory**
- **CDF 10 Dell PowerEdge 2650 2.4 GHz Xeon with 1.5GB memory**
- **CDF 7.5TB Raid disk**

Glasgow e-Science Infrastructure Future Plans

- But not enough…
  - Computer Services second HPC facility (128 processor)
    - (being procured)
  - University SAN (50TB – 25TB mirrored at separate locations across campus)
    - (being procured)
      - ~£850k investment
        - » Expected usage first quarter 2005
  - Access to campus wide resources
    - Physics and astronomy training lab condor pools
    - NeSC training lab condor pool
    - Computer services…???
    - EEE compute clusters and larger SMP machines…???
    - others…???
  - NGS
  - proposals for Scottish Grid Service infrastructure
  - SBRN (and other) equipment funds
  - …

Glasgow e-Science People Plans

- E-Science Strategy/Business Plan
  - Research Computing Director (new position)
  - E-Science lectureship (new position)
  - E-Science applications co-ordinator (moi)
  - Underwriting of existing staff contracts
  - Future running costs

– To be funded through contributions from *university wide* e-Science activities and university Strategic Investment Funds
  - university wide engagement and support of e-Science

Glasgow e-Science Organisational Aspects

- Essential to nurture relationships across university at ***ALL*** levels
  - Social issues almost as important for Grid success as technical issues

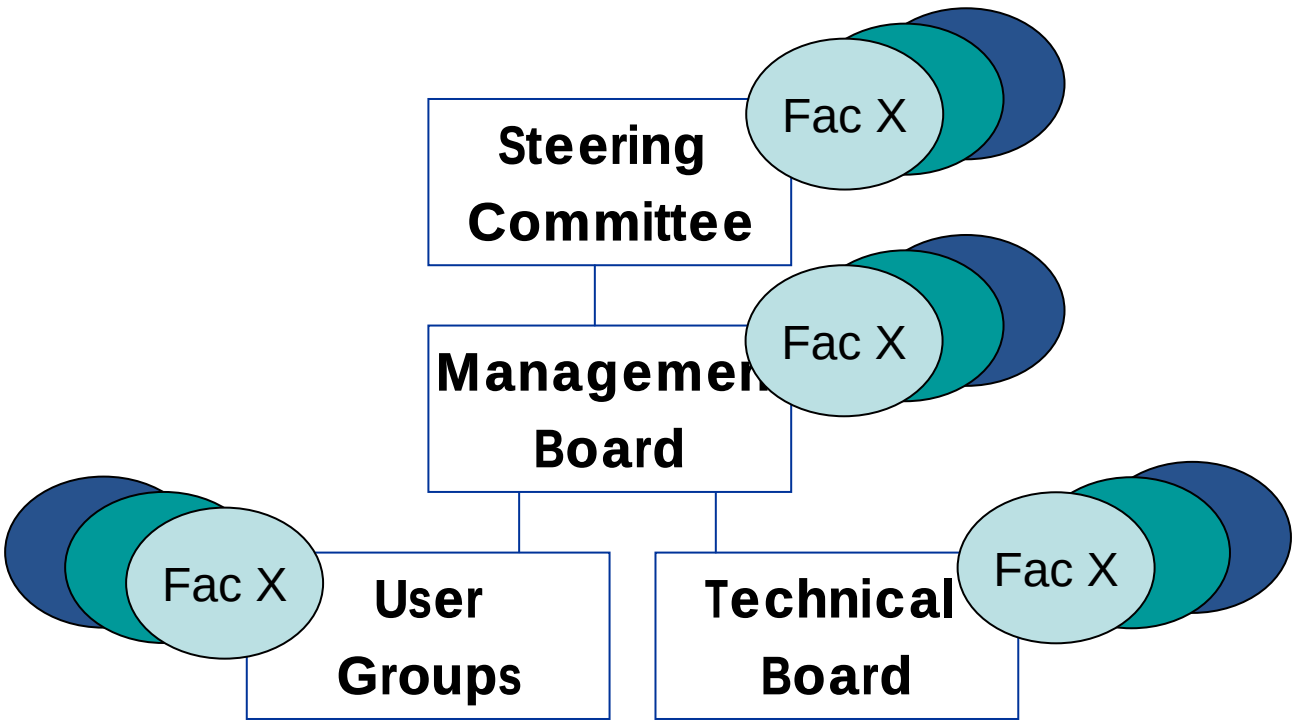

Enterprise Grid Future

- Enterprise Grid
  - Grid supporting a company's enterprise infrastructure

- Major corporations?
  - Need to establish dynamic virtual organisations?
    - More organised (???)
      - Trust of open source Grid solutions?
  - Private Networks and Intranets
  - GSK
    - Have full time staff for keeping up to date copies of ALL bio-data sets
  - SMEs more likely
    - How many need access to HPC facilities?
    - More likely need tailored access to numerous data sets, or mgt of internal operations,
      - e.g. FirstDIG project
  - Maturity of Grid software
    - Needs to evolve into mature "product" before adoption likely
    - Security!!!!!

Semantic Grid Future

- Semantic Grid
  - Integration of Grid and Semantic Web meta-data and ontology technologies
    - extension of current Grid in which information and services are given well-defined meaning, better enabling computers and people to work in cooperation
    - Semantics are key to virtualisation and abstraction in the Grid
    - WHY NEEDED?

The Semantic Grid Problem

- Wish to reuse
  - Data
  - Services
  - Knowledge
  - Software
- Anticipated use
- Unanticipated use

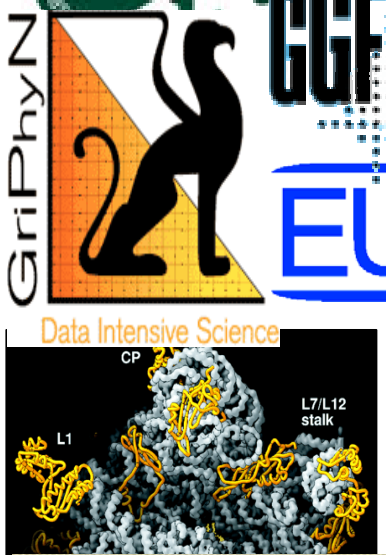

Taken from De Roure/Goble talk

Semantic Grid Future

- Numerous activities on-going in this area
  - Development of ontologies (per application domain)
  - Languages used to capture semantics
  - Prototype toolsets

  - www.semanticgrid.org

Lightweight Grid Future

- Lightweight Grids
  - Simple solutions that work with minimal complexity for installation and usage on a variety of infrastructures/OS
    - (see Chin/Coveney paper)
    - Needed if more heavyweight Grids easier to install/use?
  - Essential for “current” would be Grid adopters
    - No major re-engineering to get a piece of code to run on the Grid
  - EGEE working towards this
    - gLITE

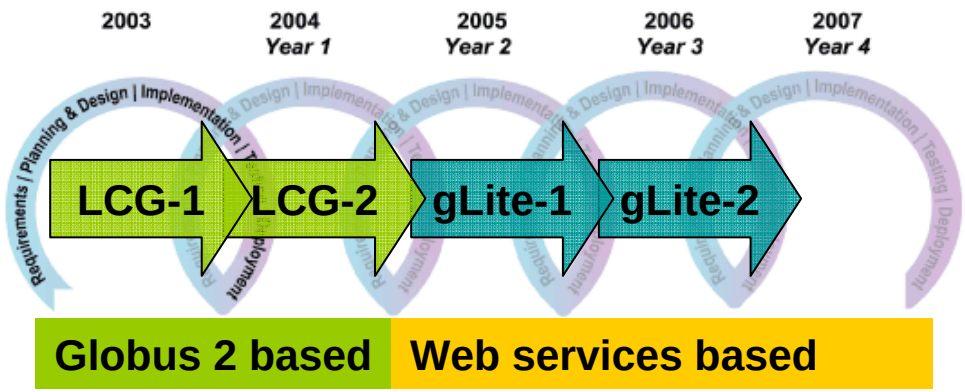

Collaboration Grid Future

- Grid supporting collaborative tools like Access Grid, whiteboards etc
  - Crucial to support e-Science
    - not just about connecting machines/data sources
    - Useful for “virtual” F2F meetings/seminars, but…
- Should be linked to rest of Grid infrastructure
    - also allow scientists to view same simulations, models and decide what to do in real time
      - aka computational steering
      - Trialled in projects like RealityGrid but still not mainstream
      - Security issues of supporting this also

Autonomic Grid Future

- Autonomic Grids
  - Fault tolerant and self-healing Grids
  - Yet another possible Grid of the (distant?) future
    - Once we know how to build fault tolerant, dependable Grids
      - Then look towards automating the processes by which a Grid can manage itself
    - Look towards optimised Grid solutions, e.g. for load balancing as first step towards autonomic computing
    - Profiling of Grids to user demands, e.g. data replication based upon general usage

Future of OGSA Framework

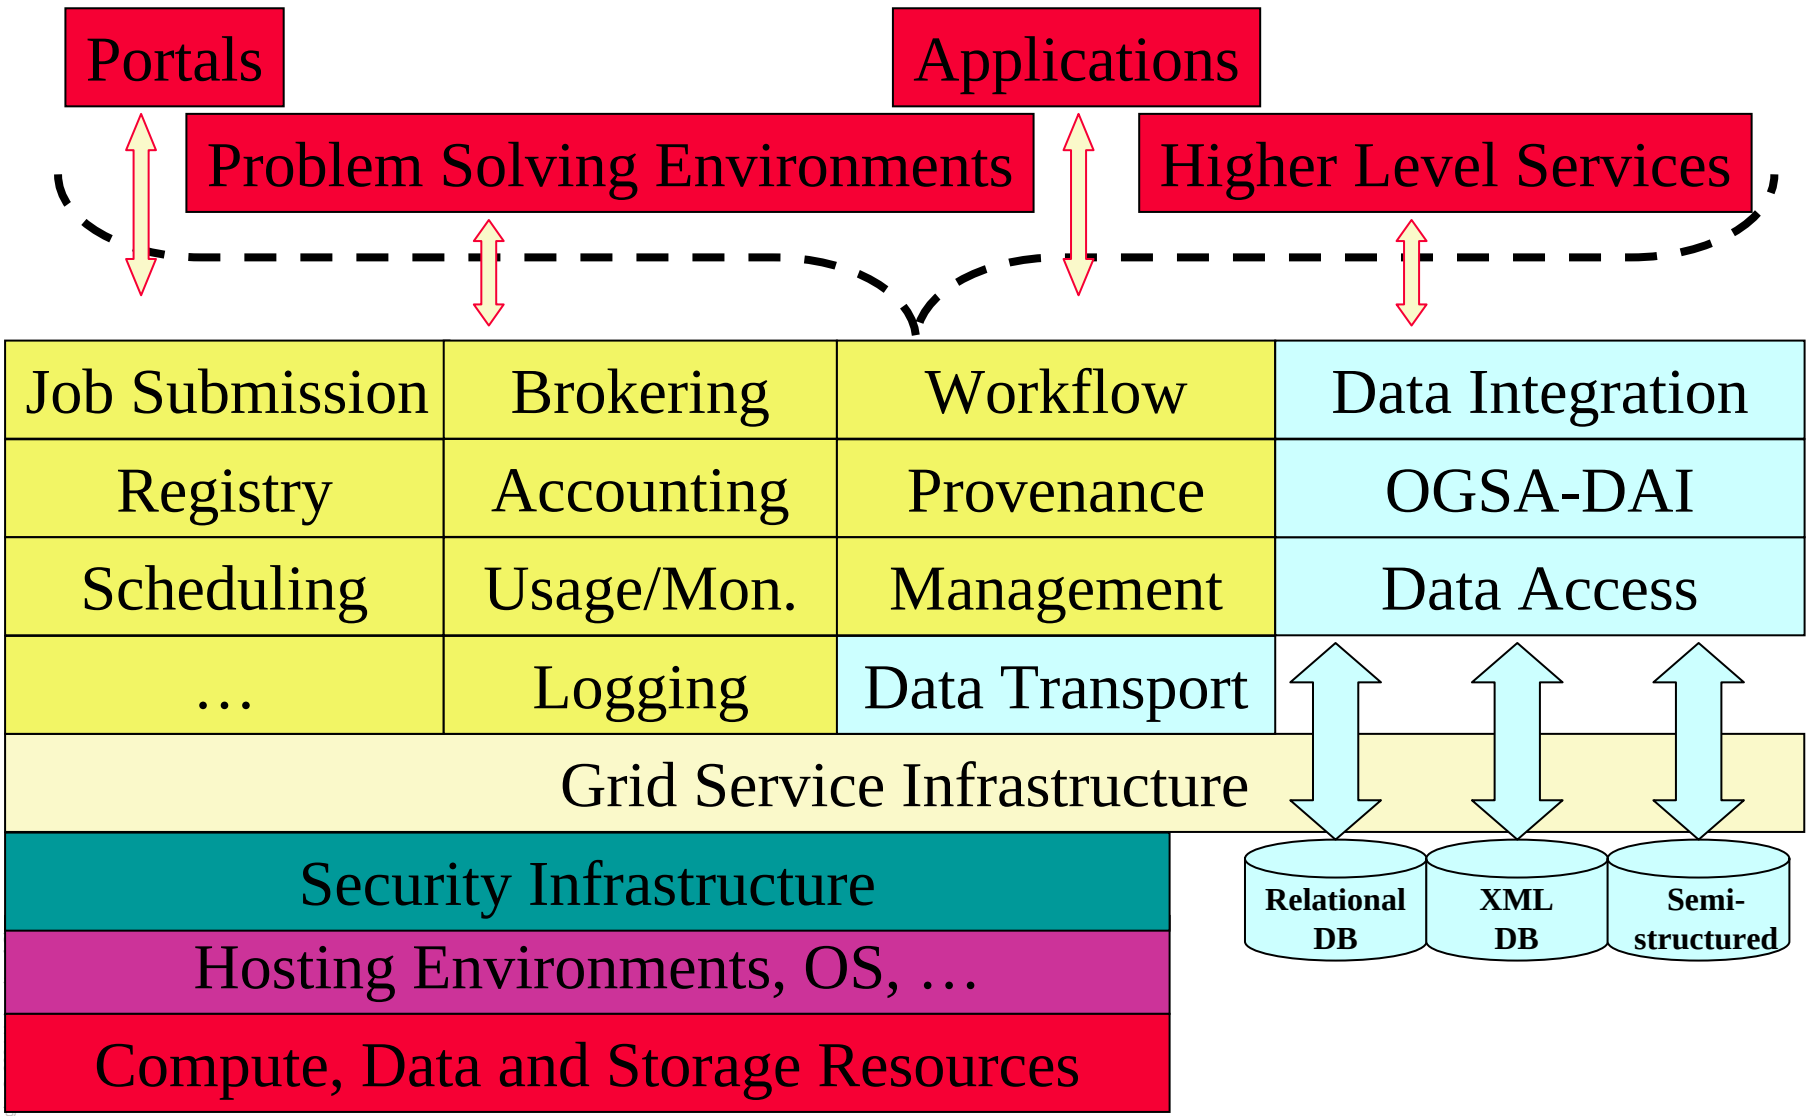

Future of OGSA

- Technologies and standards evolving
  - Move from OGSI to Web Service Resource Framework (WSRF)
    - Preserves OGSI functionality
      - Lifetime, properties, notification, error types, **…**
    - Separates service from resource
      - Service is static and stateless
    - Resource is dynamic and stateful
    - ***Builds on WS-Addressing***
    - Is WS-Interoperability standards compliant

Justification (?) for WSRF

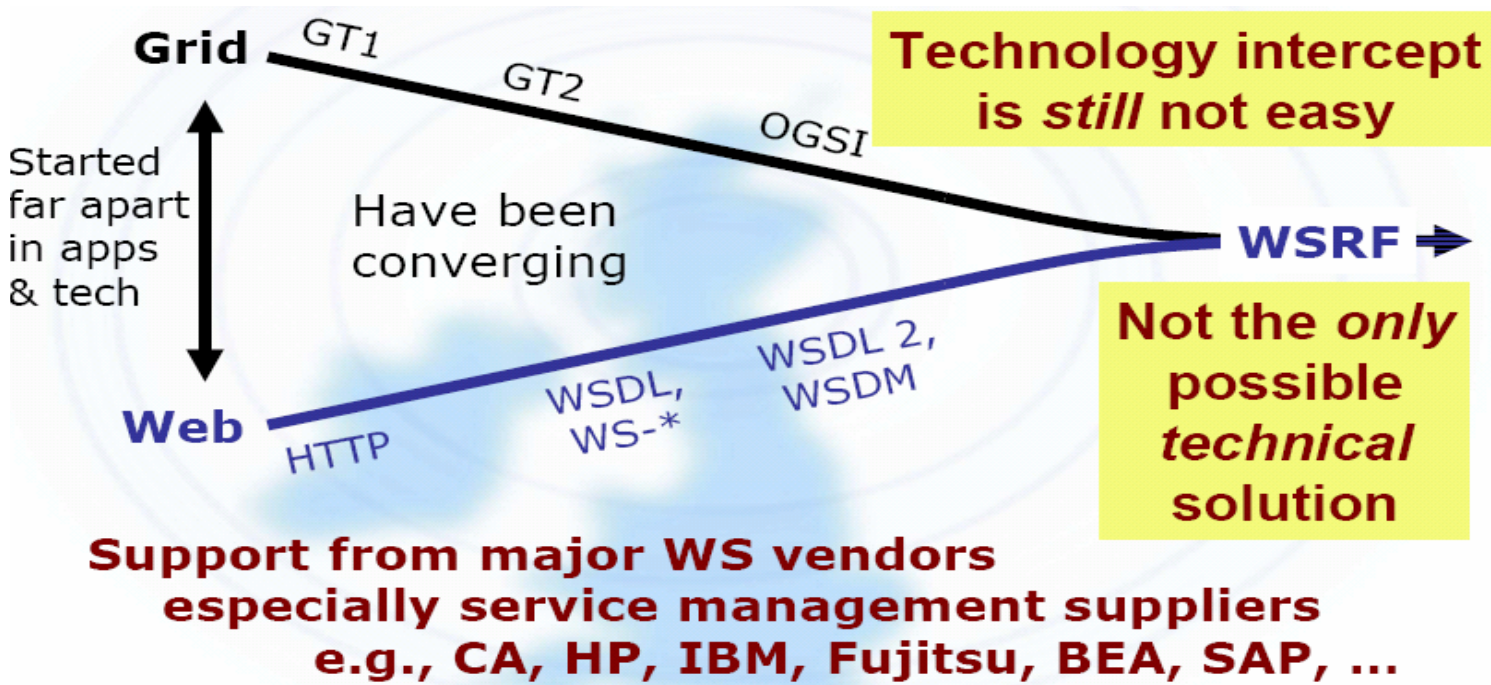

Borrowed from Malcolm Atkinson talk

Timelines for Grid Computing

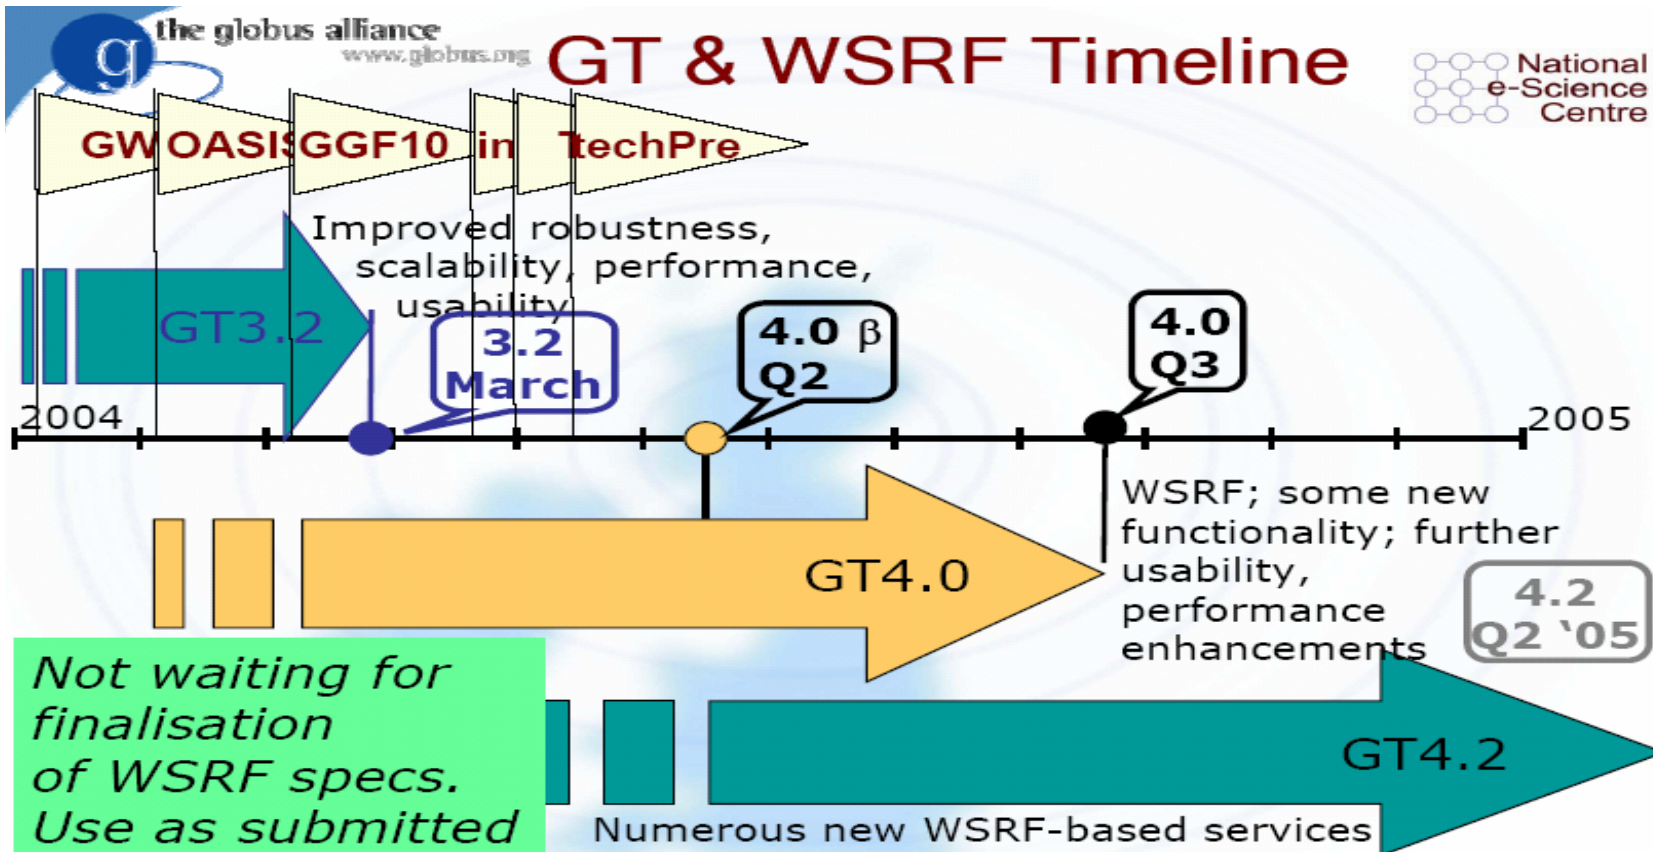

Plus

- EGEE
- OMII
- NGS

Borrowed from Malcolm Atkinson talk

Future of Grid Security

- Security
  - Grid should not weaken existing security infrastructures
  - Scaling up Grid and security to masses needed
  - Overcome issues with PKI
    - take away X.509 certificates from users?
      - » e.g. more widespread usage of MyProxy
  - Need AAA and best practice guidelines in how to establish "Grid security"
    - Need better tools to set up and manage security infrastructures (experiences from Problem set 2)
  - What to do in case of "incidents"?
      - » SC2004 and FBI session!

Future of Grid Security …ctd

- AAA projects (DyVOSE)

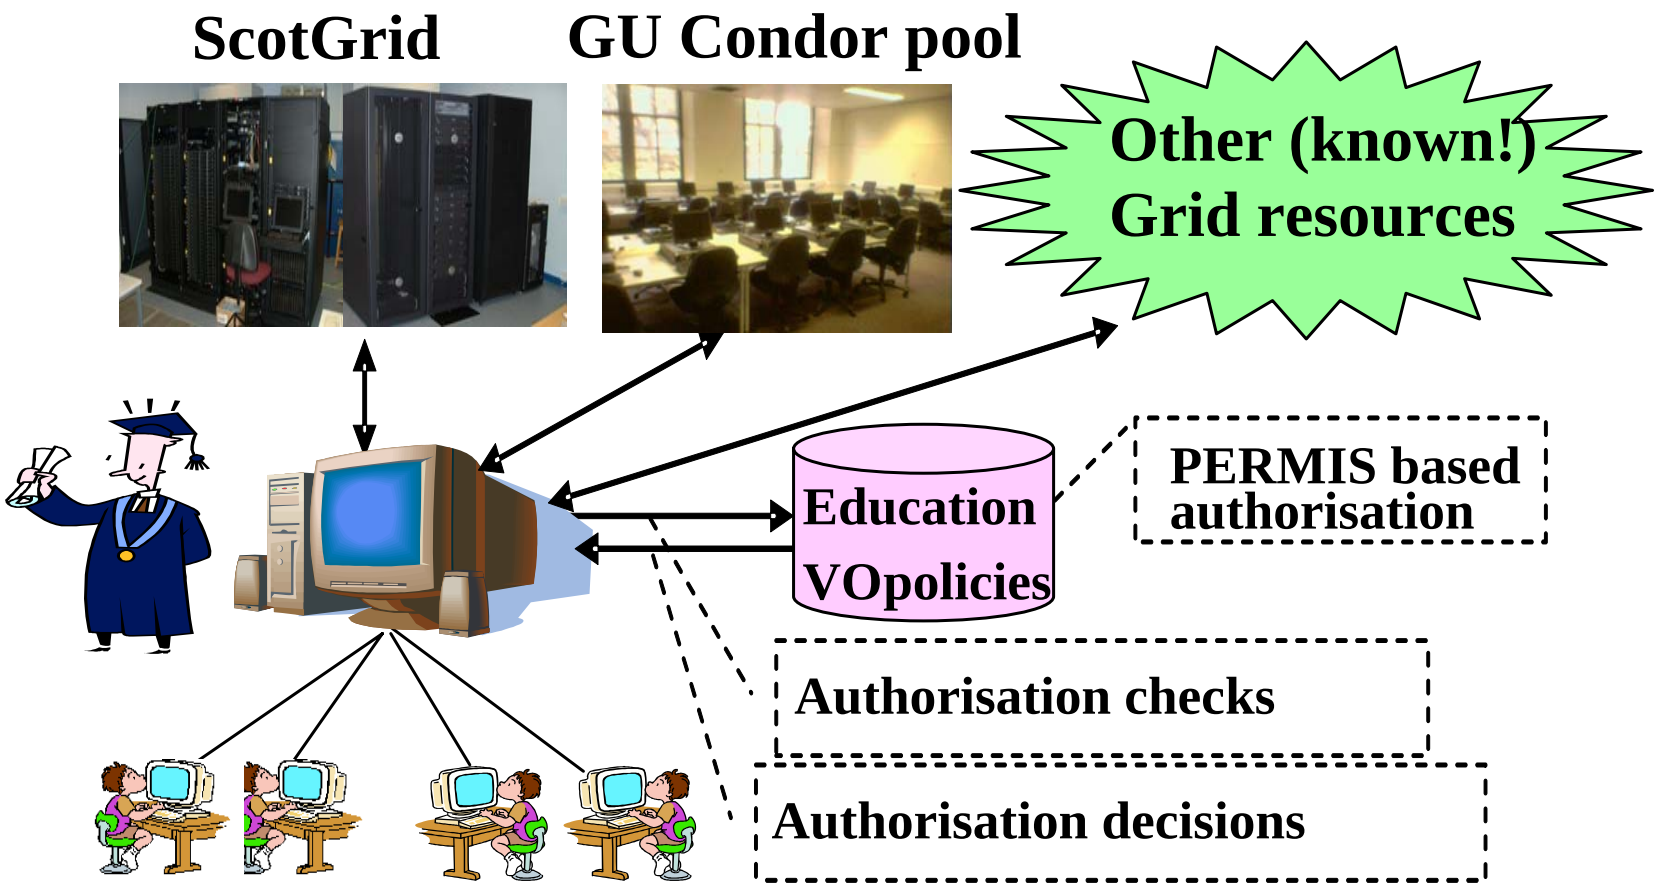

Future of Grid Security …ctd

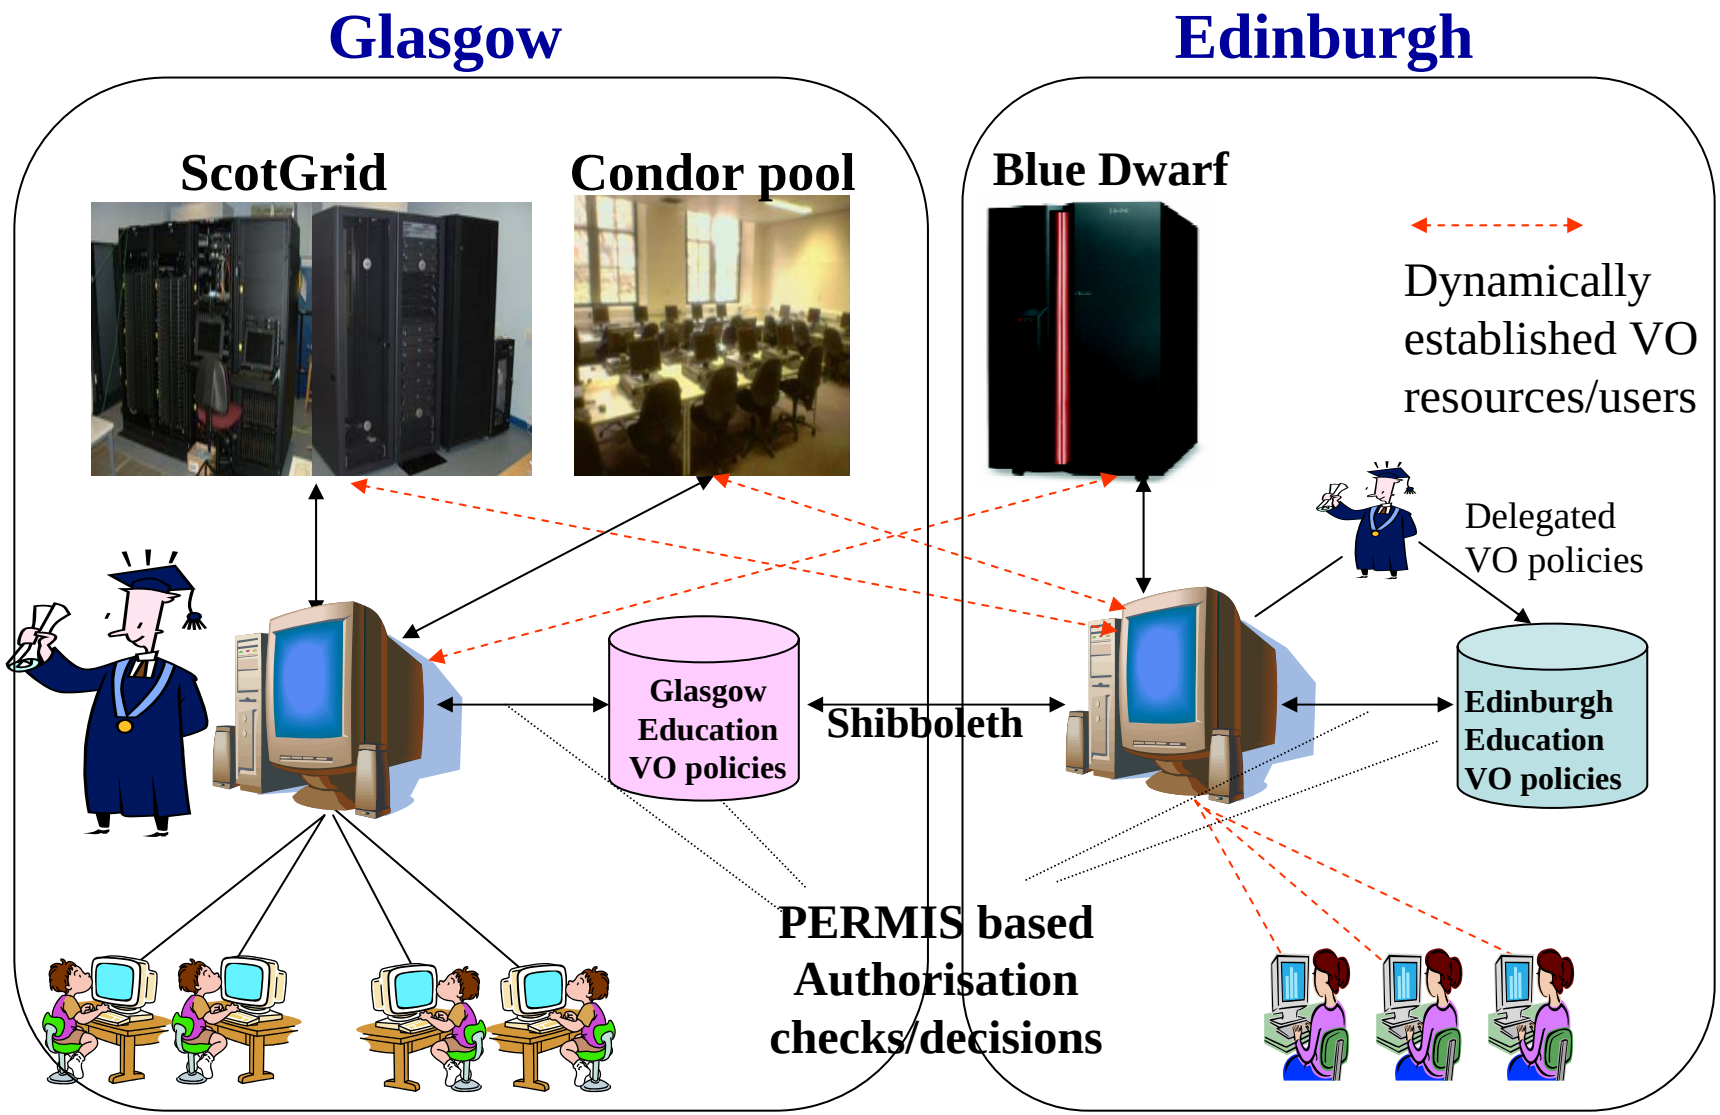

Future of Grid Security …ctd

- Local certificate allocation
  - e.g. to students when matriculating
    - Used for authentication and authorisation
    - Local monitoring and policy enforcement

- Across UK academia that this model together with **Shibboleth** for security aspects of inter-site sharing of resources will be taken up
    - *Shibboleth: A word which was made the criterion by which to distinguish the Ephraimites from the Gileadites. The Ephraimites, not being able to pronounce sh, called the word sibboleth. Hence, the criterion, test, or watchword of a party; a party cry or pet phrase.*
      - Will be seeing more Shibbleth origin, target, WAYF services

Future Key Directions for e-Science

- Physics
  - production level services for particle physics in 2007
  - astrophysics, gravitational waves, ...
- Chemistry
  - Molecular simulations, structure visualisations, ...
- Engineering
  - Mechanical, civil, electronic, aerospace, ...
- Computing science (users and developers)
  - Grid development, usage of Grid, e.g. for information retrieval experiments, ...
- Social sciences
  - Data everywhere, information retrieval, ...
- Environmental sciences
  - Earth modelling, pollution control, ...
- ...
- Life sciences biggest driver for future?
  - Data data everywhere, linkage to clinical, medical, genomic, ...
    - Who sees it, who can change it, ...
    - What does is the significance of this data in comparison to...

Bioinformatics Grid Needs

Workflow / Virtual Organisatio

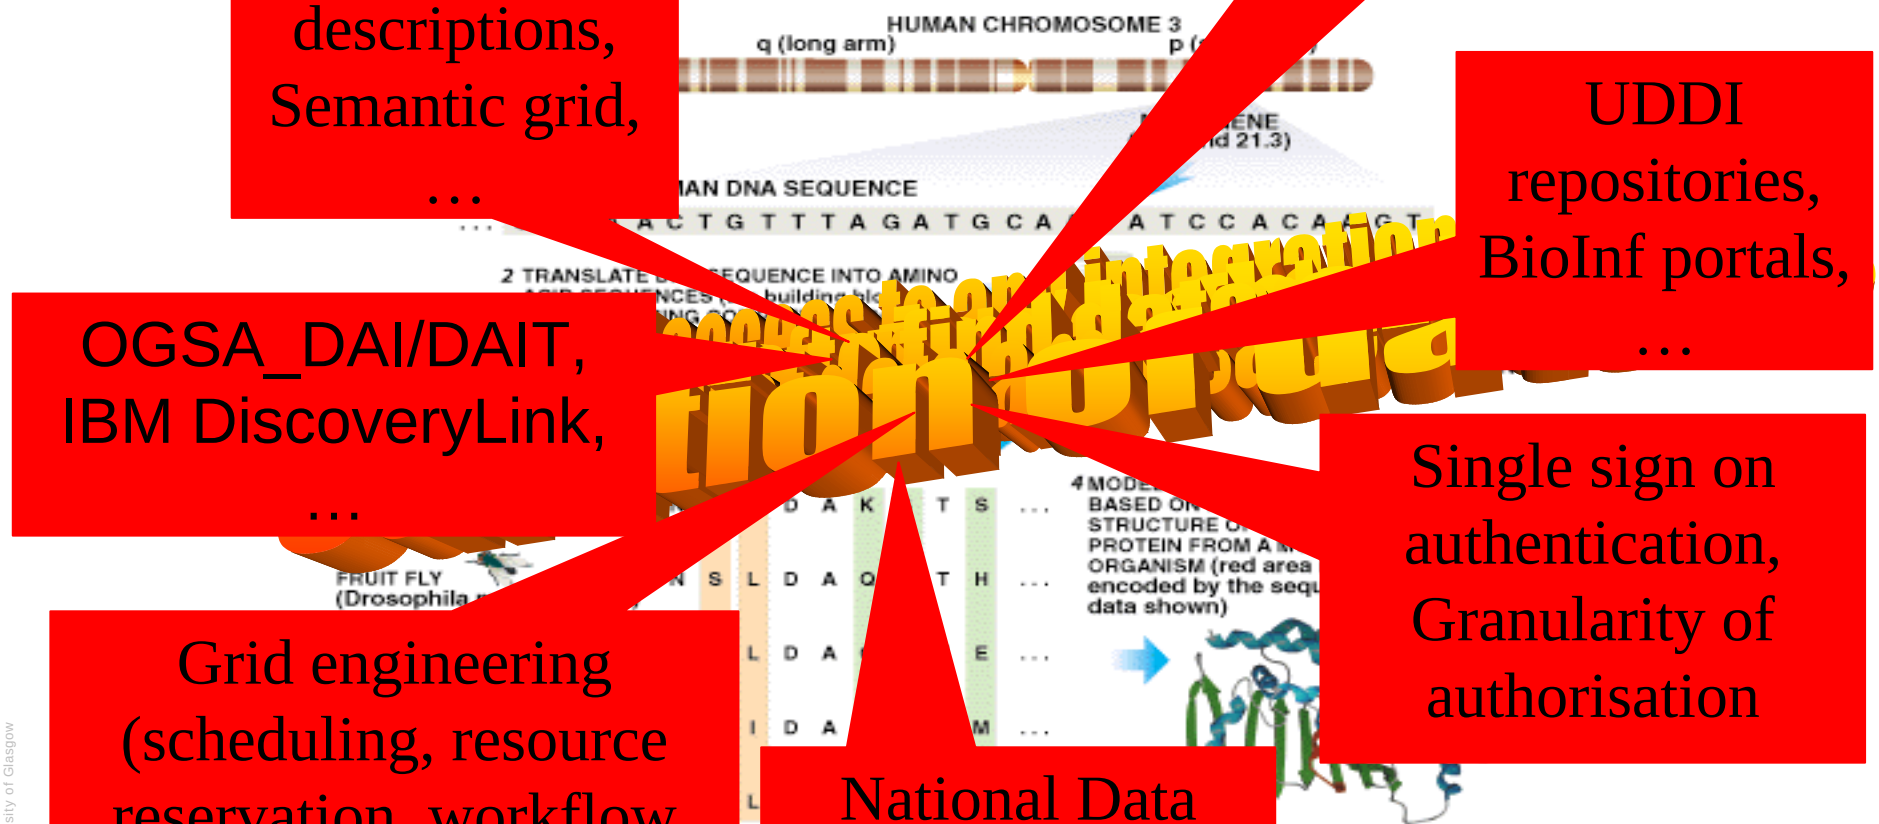

BioInf community, Database schemas, …

WSDL descriptions, Semantic grid, …

UDDI repositories, BioInf portals, …

OGSA_DAI/DAIT, IBM DiscoveryLink, …

Single sign on authentication, Granularity of authorisation

Grid engineering (scheduling, resource reservation, workflow enactment, …)

National Data curation centre

Future Funding for Grids / e-Science

- e-Science/Grid happening and influencing all research / faculties at Glasgow
  - DTI call for £100M out on 29th November
    - Looking towards aligning the following technology themes:
      - design, simulation & modelling
      - pervasive computing
      - nanotechnology
      - imaging
      - smart manufacturing
      - knowledge transfer networks
      - …
  - April 2005 call for £150M pounds
- Government SR allocated £115M for e-Science
  - Spread across ALL funding councils
    - PPARC, EPSRC, MRC, BBSRC, NERC, ESRC, SHEFC, ...
- Plus £150M for each of the three Scottish ITIs
  - TechMedia, Life Sciences, Energy
- Plus European monies
  - Framework programmes, ...

Future of Grid Computing Conclusions

- e-Science/Grid has been big business for university
  - It will in some shape or form continue to be so
- Grid has numerous open research areas that have still to be addressed
  - Is it novel?
  - Perhaps from the applications side...