Building the e-Science Grid in the UK: Middleware, Applications and Tools deployed at Level 2.

Rob Allan[1], David Baker[2], David Boyd[3], Dharmesh Chohan[1], Simon Cox[2], Hakki Eres[2], Ron Fowler[3], Nathalie Furmento[4], Jonathan Giddy[5], Terry Harmer[6], Mark Hayes[7], Nick Hill[3], Jon Hillier[8], Jens Jensen[3], Andy Keane[2], Marko Krznaric[4], William Lee[4], Mark McKeown[9], Alistair Mills[3], Steven Newhouse[4], Stephen Pickles[9], Robin Pinning[9], Andy Richards[1], Asif Saleem[4] and John Watt[10].

[1]CCLRC Daresbury Laboratory
[2]University of Southampton
[3]CCLRC Rutherford Appleton Laboratory
[4]Imperial College, London
[5]University of Wales at Cardiff
[6]The Queen's University of Belfast
[7]University of Cambridge
[8]University of Oxford
[9]University of Manchester
[10]University of Glasgow

Abstract
Over the period September 2002-April 2003 the UK Grid Engineering Task Force and staff at Regional e-Science Centres and CCLRC deployed the Globus Toolkit GT2 at 14 sites and on approximately 80 compute resources to set up the first production-quality e-Science Grid for the UK. This work is proving to be exemplary of what can be achieved using heterogeneous resources on a national scale and is feeding discussions on how to link Grids with multiple virtual organisations across Europe. This paper therefore reports on experiences of this deployment exercise, compares with other efforts worldwide and suggests further work to be done.

Introduction
Members of the UK Grid Engineering Task Force from the Regional e-Science Centres and CCLRC have deployed the Globus Grid middleware and a number of additional tools from their projects to implement the first UK computational e-Science Grid. This so-called "Level 2" Grid now provides a service for a growing number of users of applications pre-installed on Centres' resources and is supporting grant-funded e-Science pilot projects. Over the coming months it will continue to be developed as a robust service platform. The programme of work undertaken to deploy software and applications is described in this report. It was agreed by all members of the UK Grid Engineering Task Force and managed by Rob Allan, Alistair Mills and David Boyd. The workplan was forwarded to the e-Science TAG for approval, and we acknowledge the active participation of Anne Trefethen in project meetings, which enabled us to get very useful guidance in a number of areas.

E-Science Grid deployment phases were defined to be:

Level 0: Resources with Globus GT2 registering with the UK MDS at *ginfo.grid-support.ac.uk* [5];

Level 1: Resources capable of running the Grid Integration Test Script [4];

Level 2: Resources with one or more application packages pre-installed and capable of offering a service with local accounting and tools for simple user management, discovery of applications and description of resources in addition to MDS;

Level 3: GT2 production platform with widely accessible application base, distributed user and resource management, auditing and accounting. Resources signing up to Level 3 will be monitored to establish their availability and service level offered. A GT3 testbed is also included at this level.

Access Grid nodes at the 12 participating sites were used for fortnightly meetings during which progress was discussed, issues raised and new milestones and deliverables agreed. These meetings were documented in the form of minutes and action lists. In the later stages the new Centres of Excellence have started to contribute resources.

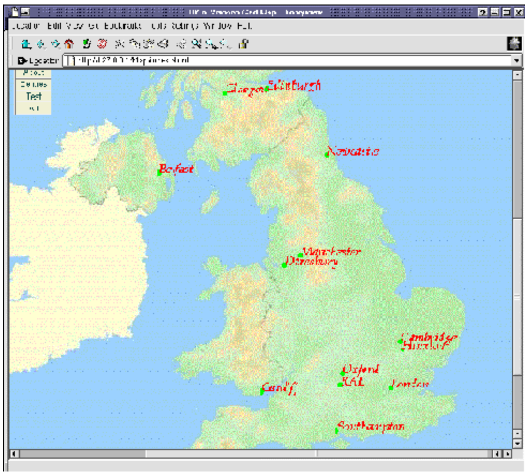

Figure 1: on-line active map of Regional Centres

The agreed milestones and deliverables were also documented in an overall project plan together with required effort levels and this used to check the feasibility of the outcome. Task graphs and Gantt charts together with a task summary were produced automatically using the OpenSched open source software and incorporated into a full project report which included a basic risk analysis.

Key components of the project thus include:

- ETF Coordination: activities are coordinated through regular Access Grid meetings, e-mail and the Web site;
- Resources: the components of the e-Science Grid are the computing and data resources contributed by the UK e-Science Centres linked through the SuperJanet4 backbone to regional networks;
- Middleware: many of the infrastructure services available on this Grid are provided by Globus GT2 software with additional local resource management software such as Condor, PBS, Sun GridEngine, LoadLeveler, LSF;
- Directory Services: a national Grid directory service using MDS links the information servers operated at each site and enables tasks discovery usable resources at any of the e-Science Centres.;
- Security and User Authentication: the Grid operates a security infrastructure based on x.509 certificates issued by the e-Science Certificate Authority at the UK Grid Support Centre at CCLRC;
- Additional tools and applications to be described below.

Work package leaders were responsible for setting the technical agenda of their own work packages and for documenting technical aspects of the work undertaken and the deliverables produced. Proposals were circulated for discussion and agreement. The current working groups are:

- Grid Deployment Management – leader R.J. Allan (CCLRC Daresbury);
- Grid Middleware – leader J. Giddy (Cardiff), previously N. Hill (CCLRC RAL);
- Information Services – leader R.J. Allan (CCLRC DL);
- Security and Firewalls – leader I. Kostadinov, previously J. Hillier (Oxford);
- CA – leader A.B. Mills (CCLRC, RAL);
- Grid User Management – leader S.J. Newhouse (Imperial);
- Applications – leader S. Cox (Southampton);
- Grid Integration Testing – leader D.J. Baker (Southampton);
- Grid Support – leader A.B. Mills (CCLRC RAL);
- Access Grid – leader M. Daw (Manchester);
- Web and Grid Services – leader R.J. Allan (CCLRC Daresbury).

A secure Web site was used for internal document management based on the EDG GridSite software from Andrew McNab (Manchester). The Web site was set up at RAL and maintained by the Grid Support Centre. We also thank the Grid Network Team who were instrumental in the early stages in collecting requirements and for making the university computing services community aware of the need for access through firewalls and for sufficient bandwidth and QoS to run applications and demonstrators.

In the latter stages of the project several participants from the new Centres of Excellence and other e-Science Pilot Projects became involved in deploying and using applications on the Grid infrastructure. International collaborators are also aware of our activities and have provided input in a number of areas. The outcome of our work has clearly created interest at an international level with the integration testing activity, heterogeneous nature of the Grid, its 4-tier MDS hierarchy and diversity of installed applications being cited as leading examples. We are particularly grateful for funding to employ David Wallom, who is participating in setting up a European Grid Support Centre and for his input in the later stages. Ron Fowler also joined the ETF with new funding from DTI and is setting up Grid monitoring tools, which may be of value in the context of a future Grid Operations Centre.

We firmly believe this deployment of the first UK e-Science Grid has been successful and represents one of the first national-scale computational Grids anywhere in the world (as reported at GGF7). Applications running on this Grid have been demonstrated at a number of meetings, including this one, and are listed together with the published papers in a final report on the project [1].

Particular successes have been:

1. Establishment of a coherent collaborating group of people working to common goals (the ETF). We have used Access Grid for regular technical and management meetings;
2. Evaluation, deployment and support for the Globus middleware and management of updates through several releases. This was largely facilitated by the UK Grid Support Centre staff;
3. Issuing of digital certificates and implementation of working solutions to related security issues around users, hosts and firewalls (CA, IP database, dynamic firewall);
4. Implementation of solutions to manage authorisation issues for distributed users and virtual organisations and to do accounting (VOM and RUS);
5. Implementation of several client tools to manage the install base and access to applications on a heterogeneous resource pool (Nimrod, HPCPortal);
6. Implementation of a sophisticated Grid information and operational monitoring system comprising several software components (MDS, GITS, InfoPortal);
7. Implementation on Grid resources of a diverse portfolio of applications with a wide user base;
8. Identifying lessons learned and future requirements.

Issues Addressed

There were a number of technical issues which had to be addressed to complete this task and a number of equally important sociological issues. Among these included:

- Establishment of a coherent collaborating group of people working to common goals;
- Data communication through a limited firewall port range using a "trusted host" database;
- Version management and interoperability testing of Globus releases;
- Integration testing and monitoring of resource status;
- Installation and validation of applications for end users in e-Science pilot projects;
- Issuing of host, service and user certificates from the UK Certificate Authority;
- In-depth support for system administrators and users;
- User management, including multiple conditions of use (currently being addressed by John Duke, JISC);
- Demonstrations of the Grid at Centre open days and pilot project workshops;
- Project management, including regular ETF meetings with 12 sites via Access Grid;
- Deployment of tools to facilitate management of distributed resources, applications and users.

Resources

Collectively, resources at these sites represent a significant pool of heterogeneous computing power. The Level 2 Grid consists of a large number of testbed systems plus the larger (>32 processor) servers shown in Table 1. Other resources are diverse in nature, including a Sony Playstation at University of Manchester and an SGI Onyx visualisation engine at University College London.

Table 1: The larger servers on the Level 2 Grid.

Daresbury	HPCx	1280 Power 4 procs in 40 IBM Regatta p690 nodes, AIX5.1, LoadLeveler
Daresbury	Loki cluster	64x alpha cluster, RedHat 7, PBS
Daresbury	IBM SP	8x IBM Power2 Winterhawk nodes, AIX5.1, LoadLeveler
Manchester Computing	Green	512 proc SGI Origin 3800, Irix, LSF
Manchester Computing	Fermat	128 proc SGI Origin 2000, Irix, LSF
Manchester Computing	Bezier	32 proc SGI Onyx 300, Irix
Imperial College	Condor pool	Large teaching pool
Imperial College	Sun	Large SMP enterprise system
Imperial College	Viking cluster	236 proc Pentium IV
Rutherford Appleton Laboratory	Hrothgar cluster	16x dual AMD Athlon, RedHat Linux, PBS
University College London	Condor pool	Around 500 proc Pentium IV
University of Southampton	Iridis cluster	Around 400 proc Pentium III and Pentium IV, RedHat Linux, EasyMCS

Tools and Middleware

Among the many tools deployed from centre projects were:

- Certificate Authority and Grid Support – OpenCA and Remedy/ ARS used (CCLRC) [8]
- GITS – Grid Integration Test Script (University of Southampton) [4]
- UK hierarchical Monitoring and Discovery Service – Globus MDS provides snapshots of resource status (CCLRC) [5]
- InfoPortal – provides resource-centric and site-centric views of the Grid and related information including an active map (CCLRC) [5]
- GridSite - secure Web access for project management (University of Manchester)
- VOM and RUS portal - user management and accounting (Imperial College) [11]
- “Trusted host” secure ip database - for firewall administrators (University of Oxford)
- ICENI – Grid client framework (Imperial College)
- IeSE and HPCPortal – portal client for moving data and running HPC applications across the Grid (CCLRC) [6]
- Nimrod/G – task farming application management client tool (Cardiff)
- Network monitoring tools (CCLRC) [13]

Several of these tools are described in separate conference papers as referenced above.

Applications

A number of e-Science Grid sites participated in winning entries at SuperComputing 2002. There were 3 prize categories and members of our community were involved in/ linked to all three: (i) Most Geographically Distributed, (ii) Most Heterogeneous, and (iii) Most Innovative Data-Intensive Application. For further details see: http://www.nesc.ac.uk/ news/SC2002_UK_success.html

The ETF Applications WP8 started scoping further requirements in December 2003. An open meeting was held in Southampton on 24/1/2003 and a follow-up meeting in Cambridge in April. At least eight groups have run applications across the multiple sites and provided details of their work: encountering and solving a range of issues. The work reported here would not have been possible without the considerable efforts of a large number of individuals across many e-Science projects and at every e-Science Centre. They have generously given their time to test out the infrastructure, and develop and test applications for deployment for the timely delivery for the launch of the e-Science Grid in April 2003.

The applications listed below have been run both on a regular basis for scientists in the participating projects and also on special occasion for demonstrations and open days at the participating centres. We are continuing to add new applications to this list and to bring in new users as the e-Science Grid becomes a more mature production environment.

1. Monte Carlo - simulations of ionic diffusion through radiation damaged crystal structures. Mark Hayes and Mark Calleja (Cambridge)
2. GENIE - integrated Earth system modelling with ICENI. Steven Newhouse, Murtaza Gulamali and John Darlington (Imperial College), Paul Valdes (Reading), Simon Cox (Southampton) [14]
3. BLAST – for post-genomics studies. John Watt (Glasgow)
4. Nimrod/G - with astrophysical applications. Jon Giddy (Cardiff)
5. DL_POLY - via e-Minerals portal. Rob Allan, Andy Richards and Rik Tyer (Daresbury), Martin Dove and Mark Calleja (Cambridge) [6,7]
6. Grid Enabled Optimisation - vibrations in space application to satellite truss design. Hakki Eres, Simon Cox and Andy Keane (Southampton)
7. RealityGrid – computational steering for chemistry. Stephen Pickles, Robin Pinning (Manchester), Peter Coveney (UCL) [9]
8. R-Matrix – electron-atom interactions for astrophysics. Terry Harmer (Belfast) [12]
9. GITS. David Baker (Southampton) [4]
10. ICENI. Steven Newhouse, Nathalie Furmento and William Lee (Imperial College) [15]

Some of these applications are described in separate conference papers as referenced above.

Work still to be done

By far the most significant outcome of the initial user experiences was to highlight issues of ease of use and reliability of the Grid infrastructure. Users have to apply for a digital certificate, but even when they have this they still need to apply for individual accounts on resources which they wish to access and sign up to diverse usage policies. Implementation of applications across a pool of heterogeneous resources can be difficult and validation mechanisms are required. Discovery of appropriate applications on machines with sufficient capacity (cpu and memory) to run a particular job is also not easy as the MDS system does not produce output in an easily consumable form. Client tools such as Nimrod-G and HPCPortal are being used to address some of these user issues. Reliability issues have also been a concern, for instance upgrading middleware, testing inter-operability of releases and keeping the MDS system running have all required much more manpower than initially envisaged.

A number of areas significant for a fully-functional production Grid environment have hardly been tackled using GT2. Issues include:

- Grid information systems, service registration, discovery and definition of facilities;
- Security and maintenance of the Grid;
- Role-based authorisation;
- Portable parallel job specifications;
- Meta-scheduling, resource reservation and ‘on demand’;
- Linking and interacting with remote data sources;
- Wide-area computational steering;
- Workflow composition and optimisation;
- Distributed user and application management;
- Data management and replication services;
- Grid programming environments, PSEs and user interfaces;
- Auditing, advertising and billing in a Grid-based resource market;
- Semantic and autonomic tools.

There are many general issues surrounding security and deployment of applications and software. For instance, ensuring that all nodes have suitable policies and security patches in place, implementing mechanisms both technical and policy for coping with potential breaches of security and keeping software and patches up to date throughout the Grid.

Given that some projects are now in place to develop services in an OGSA framework (Grid Economies, OGSA-DAI, Workflow Optimisation, etc.), is GT3 a more appropriate vehicle for future development? A production GT2 platform can continue to be supported and forms the core pillars for GT3. This will be discussed further in a mini workshop at the AllHands meeting [10] and in GT3 testbed projects just starting.

We are also now beginning the task of bridging communities with the e-Science Grid, JISC and EGEE starting to share experiences, middleware and tools. A BoF will be held at the AllHands meeting to discuss possible ways forward.

The UK Grid is also growing, with seven new Centres of Excellence joining - UCL and Warwick are already on line. This may mean that the ETF will become too large if all sites are involved, so a different organisational structure may be required. Suggestions were put to the recent Stakeholders' Meeting (London 10/7/03) including:

- A users' forum;
- An administrators' forum;
- Operational management and monitoring group;
- Distinct cross-community working groups, e.g. on Grid Information Services, Security, OGSA.

The participation of a wider community is bringing tools from other e-Science projects into use and also increasing the resource and application base for real scientists.

Finally the e-Science Grid will soon have dedicated resources from the JISC JCSR programme and the ETF anticipates working with the Open Middleware Initiative Institute, OMII, at least to inform on practical experiences, requirements and constraints, e.g. around security.

Acknowledgements

All ETF members and collaborators who contributed to deploying and using the UK e-Science Grid up to Easter 2003.

Many other people who took part in the regular ETF meetings and were instrumental in making the Grid work.

The Centre Directors and e-Science Core Programme Team who gave guidance and resources to the project.

Finally we also thank the AG node operators and members of the Grid Support Centre work worked in the background but were essential to the task in hand.

References

[1] ETF Team *Building the e-Science Grid in the UK* Final report of Level 2 Grid Deployment Project (CCLRC e-Science Centre, May 2003)
[2] http://www.grid-support.ac.uk/etf
[3] http://www.grid-support.ac.uk
[4] D.J. Baker, S.J. Cox, J. Hillier, M. Mckeown and R.F. Fowler *Building the e-Science Grid in the UK: Providing a Software Toolkit to Enable Operational Monitoring and Grid Integration* Proc. AHM 2003 (Nottingham 2-4/9/2003)
[5] R.J. Allan, D. Chohan, X.D. Wang, M. McKeown, J. Colgrave, M. Dovey, M. Baker and S.M. Fisher *Building the e-Science Grid in the UK: Grid Information Services* Proc. AHM 2003 (Nottingham 2-4/9/2003)
[6] A.J. Richards, R.J. Allan, G. Drinkwater and K. Kleese van Dam *Building the e-Science Grid in the UK: An Advanced Grid Programming Environment* Proc. AHM 2003 (Nottingham 2-4/9/2003)
[7] M. Dove et al. *Environment from the Molecular Level: an e-Science Testbed Project* Proc. AHM 2003 (Nottingham 2-4/9/2003)
[8] A.B. Mills et al. *UK Grid Support Centre* AHM 2003 (Nottingham 2-4/9/2003)
[9] J.M. Brooke, P.V. Coveney, J. Harting, S. Jha, S.M. Pickles, R.L. Pinning and A.R. Porter *Computational Steering in RealityGrid* Proc. AHM 2003 (Nottingham 2-4/9/2003)
[10] D. Berry *OGSA: The Platform for delivering UK e-Science Proc.* <AHM 2003 (Nottingham 2-4/9/2003)
[11] A. Saleem, M. Krznaric, S. Newhouse and J. Darlington *ICENI Virtual Organisation Management* Proc. AHM 2003 (Nottingham 2-4/9/2003)
[12] A. Carson et al. *2D R-Matrix Propagation: the Calculation of Atomic Collision Data using Grid Resources* Proc. AHM 2003 (Nottingham 2-4/9/2003)
[13] M. Leese *Network Monitoring* Proc. AHM 2003 (Nottingham 2-4/9/2003)
[14] M. Y. Gulamali, T. M. Lenton, A. Yool, A. R. Price, R. J. Marsh, N. R. Edwards, P. J. Valdes, J. L. Wason, S. J. Cox, M. Krznaric, S. Newhouse, J. Darlington *GENIE: Delivering e-Science to the Environmental Scientist* Proc. of AHM 2003 (Nottingham 2-4/9/2003) http://www.genie.ac.uk/
[15] N. Furmento, W. Lee, S. Newhouse, J. Darlington *Test and Deployment of ICENI, an Integrated Grid Middleware on the UK e-Science Grid* Proc. of AHM 2003 (Nottingham 2-4/9/2003)