

# Study of User Priorities for e-Infrastructure for e-Research (SUPER)

Steven Newhouse  
University of  
Southampton

Jennifer M. Schopf  
National e-Science  
Centre and ANL

Andrew Richards  
CCLRC

Malcolm P. Atkinson  
e-Science Institute

## Abstract

SUPER, a Study of User Priorities for e-infrastructure for Research, was a six-month effort funded by the UK e-Science Core Programme and JISC to inform investment so as to provide a usable, useful, and accessible e-infrastructure for all researchers. It envisages a coherent set of e-infrastructure services that would increase usage by at least a factor of ten by 2010. Through a series of unstructured face-to-face interviews with over 45 participants from 30 different projects we have observed recurring issues relating to collaborative file management and policy, the need for tools to support dynamic virtual organisations, the support of project teams with services and training, operational issues (such as authentication, license management and reliability) and how end-users interact with e-infrastructure services through the available tools. We finish with a list of possible future activities for further prioritisation and adoption.

## 1 Introduction

SUPER, a Study of User Priorities for e-infrastructure for Research, was a six-month effort funded by the UK e-Science Core Programme and JISC to inform investment so as to:

- provide a usable, useful, and accessible e-infrastructure for researchers across a wide variety of disciplines,
- integrate existing and new e-infrastructure facilities and services into a coherent whole, and
- provide the basis to increase the use of the existing e-infrastructures by a factor greater than ten by 2010.

E-infrastructure encompasses the facilities and services that enable more advanced or more effective research through support of access to information, data, computation resources and distributed collaboration. Since multi-disciplinary teams often span national as well as institutional boundaries, support for international collaboration and consistent provisioning across these resources must be considered. Inevitably, multiple organisations will be engaged in constructing and operating e-infrastructure; therefore the recommendations must permit autonomous providers to collaborate in order to provide coherent and consistent facilities.

As a first step in carrying out the SUPER project, we spent several months in late 2006 meeting with approximately 30 groups across the UK who represent a cross section of the research community and were engaged in a

variety of roles and projects. The groups interviewed included current and potential *end-users* of e-infrastructure, generally those conducting research, design, analysis or diagnosis in projects funded by one of the UK Research Councils; *technologists*, or developers who take generic middleware and adapt it for a specific use; *generic tool developers*, those who were building solutions that could be used in many application domains; and *service providers*, including universities and publicly funded research institutes that are setting up hardware and services.

Our meetings lasted from half an hour to half a day and covered a wide variety of topics, concentrating on current use and needs. We considered performing structured interviews based on a standard questionnaire but differences in backgrounds and knowledge of those interviewed made this approach ineffective, so the interviews were not formally structured as such. Instead, we asked what functionality the groups had tried in the past, what their applications needed today from the current Grid infrastructures, and what functionality the groups were considering for near-future plans. Most meetings ended by our inquiring what functionality the group thought was most important but lacking in today's tools or services.

In terms of coverage, we performed face-to-face interviews over during the autumn in Newcastle, Glasgow, Edinburgh, Oxford, Cambridge, UCL, and Reading, covering 45 people from over 30 projects. A third of the projects were funded by EPSRC and a further

third by BBSRC, MRC and JISC. The remaining projects were funded by DTI, EU, Wellcome, AHRC, ESRC, NERC and PPARC, with a significant number of campus activities being funded directly by the local university. A full list can be found in Appendix A.

For the purposes of discussion we have categorised these requirements as “major,” discussed in Sections 2–6, and “minor,” discussed in Section 7. In Section 8 we provide a set of recommendations and identify where they can be supported within existing funding (e.g. OMII-UK, NGS, DCC, JISC e-infrastructure) and where future funding needs to be sought to support these activities. In general, however, we have not identified “best” technical solutions. Such a task will require additional detailed requirements-gathering and analysis that is beyond the scope of this broad study.

More information relating to this activity and input from other sources can be found in the SUPER technical report [1].

## 2 Collaborative File Management and Policy

The early vision of Grids – running applications on remote shared distributed resources ranging from dedicated clusters to shared-use, cycle-stealing desktop machines – is in use in many communities. The availability and accessibility of these resources, a result of the UK e-Science programme, has brought new user communities to larger-scale science than was previously feasible, and has promoted paradigm-shifting science techniques in some domains such as large-scale ensemble and parametric studies.

By far the largest concern of the users we spoke with, centred on the use of large-scale data stored in structured file sets – *how to share data with colleagues* within their project or their wider community. Of particular importance is the management of data stored in files, whether it was software, results or other data, as opposed to data stored in databases. Users are concerned about the data’s long-term storage and curation, about means for accessing these files from the local desktop, and about seamless data transfer to remote resources for further analysis. End-users also want to be able to annotate files with metadata about their contents and provenance, in order to enable searching and reanalysis at a later date – *metadata is the key to being able to share the results*. Many of the groups have explored solutions such as the Storage Resource Broker (SRB); however, SRB is seen as a

heavyweight solution that was hard to deploy and maintain by a project.

Research has experienced a paradigm shift with respect to the changing uses of data as well as the changing use of compute resources. More groups now need to share data more widely under defined access policies, to merge data, and to retain that data longer. There was an identified need for *standard policies and tools for data curation* across RCUK-funded projects. Those policies are either currently not well defined or not well known, but they need to be – for both user roles, expectations on providers and the expected duration of data’s availability.

Any access control structure and its underlying authentication mechanisms must support controlled access by recognised external collaborators and eventually unrestricted access to others. But in general, there is a shift towards much *longer-term storage of data*, some for pragmatic experimental use, and some at the behest of the funding agencies.

Easier *access to data* was also an issue. Many groups now have to manage the file output from computations across multiple locations, including national resources such as the National Grid Service (NGS), as well as campus and desktop resources. Researchers would like to seamlessly store output and access their local files when running an application remotely. This situation leads to requirements for the registration and discovery of files held in different locations.

In order to be able to easily curate and annotate data files, additional tools are needed to *autogenerate metadata* about the data itself and how, where and by what that data was generated – its *provenance*. Once provenance data is collected there will also be a requirement to navigate and analyse the data. It was noted that if users are responsible for the annotation of their data the task generally is left undone. Many groups have begun to create their own metadata after failing to find acceptable community standards. Standards exist for common basic properties, however lower-level and domain specific metadata has yet to be standardised for many communities. The quality of such annotations is often uneven, however, and some communities therefore have started developing common metadata schemas to drive annotation standards.

## 3 Tools to Support Dynamic Virtual Organisations

Although technical aspects dominated many of our discussions, there was a growing

recognition of the importance and difficulty in building an effective multidisciplinary multi-site virtual organisation (VO) within a project. This was also the topic of a workshop that we attended [2]. Established management practices need to be adapted to deal with the challenges of a virtual organisation, where the incentives for individuals may be radically different and not under the control of the project management.

Part of the lack of established practices for a VO is due to the fact that there is currently no clear definition of a VO. Undetermined factors include how dynamic the VO is and if it involves changes in membership of people or in resources. Some VOs are being defined top-down, with centralised enrolment, approval, and resource allocations that may have relationships reflecting off-line legal obligations and strong monitoring of service level agreements. Other VOs have a bottom-up approach, generally reflecting more social network groups with role-based agreements.

While current virtual organisations layered on top of production Grid infrastructures are able to adapt to a dynamic set of resources, it is harder for the virtual organisation itself to be dynamic and to bring in new people or resources that are not already part of a larger collaboration. *Support for dynamic virtual organisations* oriented around people and real organisations is becoming critical for future collaborations and controlled resource sharing (services, data, etc.) but varies considerably between domains, for example, business-to-business, accessing services in a production Grid, or more spontaneous collaborations based around social networking.

Clearly, the virtual organisation model needs to evolve to *more than just a management of roles*. Relationships, obligations, actions and interactions within a distributed team are complex and will change over time. These relationships will also have different levels of enforcement, varying from best effort to legally binding contracts. In addition, services and tools are needed not only to enrol, authorise and end an individual's participation in a virtual organisation, but also to manage the relationships and to register transactions delivered as part of this relationship, in order to provide a non-repudiable record of the interactions.

The tools and services need to reflect the different stakeholders, such as the management of the virtual organisations, service providers, developers and end-users, so as to provide a "backplane" to support collaboration. Most VO

tools are built for system administrators and the interfaces are suited to these advanced users. But the end-users may simply want to know what resources they have access to, without in-depth information about certificate setup procedures or accounting requirements. Currently, there is a *lack of tools oriented towards the end-user* in this space.

For the few tools that exist, usability remains a big issue. Current tools, in part because of a lack of consistency in what is meant by a VO in different settings, often solve problems that aren't what the user is interested in and the tools are often very difficult to understand or use in production settings.

Moreover, the current support for *collaboration within virtual organisations* is extremely limited. A successful virtual organisation needs collaborative tools that extend beyond access control and accounting. Mechanisms such as instant messaging have become widespread in order to support a dispersed project team, but the ability to share other resources such as workflows, services and documents is missing. Audio and video conferencing tools are seen as essential or desirable by many groups, but there is general disappointment in the overheads and ineffectiveness of the currently available tools. Although the need for these tools is being driven by end-users, and although these tools can be installed by end-users, training is still needed for system administrators to ensure that sufficient bandwidth is made available and networks are configured appropriately.

Once a common vocabulary has been established for VOs, different VOs will be able to share resources and interact in an organised manner, as well as share VO management tools. Such activity should not attempt to expect a single VO model; indeed, we encountered as many VO models as projects in our discussions. This is an area where standards, such as XACML, are beginning to gain widespread use. As work progresses in this area, it is important to keep in mind that tools need to be flexible for different communities, we need to respect diversity in how VOs work, and to be careful to not force VOs into one mode of operation that may not be suitable for them.

#### **4 Project Support: Teams, Services, and Training**

A recurring theme in many of our discussions was the importance of an effectively organised team to a project's success. To meet the demands of multiple stakeholders, a balanced,

trained, multiskilled team is critical to the development and deployment of complex software infrastructures.

The *skills and roles* needed by such a team, and therefore its membership, will likely vary over time. For instance, more software development effort might be needed during the early phases of a project, with more effort on deployment and training towards its end. In the early phases of the UK e-Science programme expertise in the emerging middleware infrastructures was brought into an applied science project through collaboration with computer scientists or other specialists. Some of these skills are now being brought into a project through training, collaboration with external groups or national organisations, or new personnel. Nevertheless, several groups commented on how hard it is to find such personnel and bemoaned the lack of available specialists in the necessary fields (e.g., Web services, HPC programming, application performance tuning and Java programming).

Several projects have learned the hard way that it is easy to spend the majority of their time building new middleware and support services, and not getting their end-user science done. Hence, there is a much *broader acceptance of third-party software and services* than two years ago [3]. Even groups that had started out building their own infrastructure are turning to pre-packaged solutions for higher-level services when possible. However, there is a *notable lack of information* about what is truly available and what tools might be better suited to a particular use case. To provide such expertise is, of course, an overwhelming problem – no one practitioner can advise across the many areas of Grid software and services. One promising solution is to use “vendors” such as the Matching Agency [4], offered by the Arts and Humanities e-Science Support Centre, which matches applications with software experts for advice and future work.

Experiences from the surveyed projects also show that no single software provider provides the *end-to-end solution* needed for most successful e-infrastructure deployments. While this is not necessarily a problem from an end-user’s perspective, it means that projects deploying their own infrastructure need the human resources and skills to integrate these components from different providers. In order to avoid these overheads, enabling services are needed that provide access to key resources (e.g., compute, data, file movement). Infrastructure providers, such as the NGS and local campus Grids, play a key role in providing these

low-level services. However, these are unlikely to be of direct use to particular domains or projects. In order to promote uptake into applied communities and projects, additional generic and domain specific services and tools are needed to bridge the gap between the e-infrastructure providers and end-users, such as the Grid-enabled access to data sets provided by MIMAS and EDINA. These higher-level, domain-focused services are easier to integrate into projects and, when hosted by “professional” service providers, provide a solid base for communities to build on. These services enable the software outputs from research projects to be consolidated. As communities build around these services, we would expect them to form the basis of standardisation within relevant forums, and subsequently part of the expected e-infrastructure service environment, as captured within roadmaps such as ESFRI [5] and JISC.

Information relating to the availability of tools and services was not being effectively disseminated to many of the groups we talked with. When groups were aware of the various technology options they felt there was no expertise they could call upon to help them in their decisions. Communication between generic providers and end-users continues to be hampered by the fact that the *groups communicate quite differently*. This situation is also complicated by the fact that different members of the end-user group will want different end tools, a fact that is not always clear to the tool developer. A better understanding of how to manage dynamics between these groups is needed.

At a higher level, today’s e-infrastructure users go well beyond the risk takers who started work in this area and as the end-user community expands, there arises a need for *broad outreach and evangelising* of the benefits provided by adopting the emerging campus, national and international e-infrastructure deployments. Many new projects simply don’t know what is possible or probable; and many new communities don’t know what is easy and available versus hard or even an open research question. There is no baseline of capabilities agreed upon in the broader community, and this situation is significantly hampering additional uptake.

In addition to promoting the adoption of available software and services, *training* needs to be provided for the different stakeholders within a project – the end-users who will using the e-infrastructure services through the tools

and applications, the developers (both of generic and domain specific services) who will be using the deployed e-infrastructure services to build other services and tools, and the deployers (within both a local and national context) who will have the responsibility of managing the required e-infrastructure. Training materials are needed in many forms: formal instructor-led courses, standalone self-help material, worked examples, reference systems, and so forth. Ironically, many infrastructure projects noted the need for training and had funds – but were ignorant of existing training materials (from software providers) or ongoing courses (from organisations such as NGS or NeSC) – and as a consequence there was considerable duplication of activity.

Although training can accelerate skill transfer into a team, it cannot provide experience. Several teams expressed a desire to be able to tap into external expert advice as required, rather than build up such experience internally. Advice on how to best use the deployed software needs to continue beyond high-quality training and self-help courses with access to experts who can provide consultancy on best practices and application related problem solving. In addition, support for working with legacy codes is needed. Advising, in general, may include embedding specialists with applied application groups to ensure an effective two-way dialogue that not only disseminates expertise but also captures detailed requirements to guide future developments (although, admittedly, the scalability of this approach must be considered).

## 5 Operational Issues: Authentication, License Management, and Reliability

Many user communities, while technically capable in their own fields, are not interested in dealing with the complexity of Grid middleware, especially their authentication systems. Rather, they are interested in running applications to support their science.

Authentication remains a concern for most of the groups we spoke with. The adoption by most service providers of X.509 public key certificates as the primary authentication mechanism for accessing services raises many concerns about end-user usability, but not performance. There still exists a strong need for *simple authentication mechanisms at the point of use*, which may include hiding all certificates from end users. Policies for certificate issuing are also becoming a problem as collaborations

extend beyond national boundaries. In general, none of the groups we spoke with using certificates are satisfied with their current solution.

Issues with licenses were much more common than two years ago as a result of the broadening of the Grid user base into more applied areas. In many current deployments, strongly licensed software is used because it is the community standard and because good, open source alternatives are not available and are unlikely to emerge, given the dependence of these communities on verified, quality-assured applications. However, the number of licenses available for a given piece of software is limited, and *currently there isn't a good solution for managing the shifting of licenses* on machines across a given site, let alone within a full VO. For those projects with industrial collaborations, this problem is even more significant because if a non-academic partner uses academically licensed software, both groups can get blacklisted, and thousands of pounds will need to be spent to access needed software.

Given the ability to identify the user and to provide access to a suitable license, the next step is to successfully provision the application on the remote resource. Here, the *need for a consistent environment* was commented on by groups using different resources. Re-establishing one's personal user environment on each of the NGS sites is a simple example of a current impediment to use. Virtualisation technologies are one approach to providing consistent portable environments; however, these are not currently in production use. Moreover, none of the tools in current use have client toolkits that are considered simple or straightforward enough for today's users. Indeed, a common complaint was that, even when the tools were installed, there were ongoing problems with getting the environments to be what was expected. There also was agreement that we need to move away from the remote login approach, but the tools available are simply not stable and reliable enough for hard use.

In general, *reliability was still seen to be lacking* by most users we spoke with. Even with professionally hosted Web services, there seems to be a lack of stability with either constantly changing interfaces or endpoints. As a result, many groups simply do not expect that the services will be up; they just work with whatever they find responsive at run-time. The statement "Always *something* doesn't work" was repeated by several groups. Error messages

are still found to be lacking. The complicated stack of services also doesn't help when application or system debugging is needed. Even the monitoring software may give conflicting results.

## 6 User Interaction

Different communities (and users) use different environments to undertake their work. When offered the use of e-infrastructure services to accelerate their research, users frequently express reluctance to change environments solely for the purposes of accessing these new capabilities. It is therefore critical for increased adoption to understand how the different users (end users, developers, etc.) interact with e-infrastructure services.

In order to enable tool builders to integrate access to e-infrastructure services into existing user environments, there is a *need for stable APIs* to build on. Ideally, the APIs will be based on standards and not local to a single toolkit or implementation. Ultimately, these APIs need to encompass the full breadth of e-infrastructure services that are available, but initially they should focus on job submission and management, service discovery, file movement and management, and data annotation.

These standard APIs can be used as a basis for accessing e-infrastructure services from within nongraphical environments; alternatively, command-line clients can be executed from within existing shell environments. Over the years, many projects have invested considerable effort in developing a suite of scripts to support their interaction with Grid services by using scripting environments (such as Matlab, Python, Jython and Perl) and UNIX shells (such as Bash and Tcsh). Porting these scripts to a new middleware infrastructure, or even to a new release of the current software, represents a considerable risk to a project.

Many groups are transitioning to a *workflow system*, either to better track their intermediate results (see Section 2 for a discussion of data needs) or to keep track of a large number – often tens of thousands – of individual jobs. Most groups are considering an already running system such as Taverna or one of the BPEL implementations, as opposed to building their own from scratch or further extending existing scripts that had been the start of a home grown system.

*Portals have gained considerable acceptance in some communities.* Some projects have successfully exposed common operations involving e-infrastructure services to a

predominantly non-technical (in a computing sense) community, while other projects have found the inflexibility frequently offered by a portal solution a hindrance. The availability of customisable portal toolkits (through the use of portlets) provides an important approach to building solutions for communities where they are needed.

## 7 Other Issues

In addition to the major issues discussed in Sections 4–6, several other issues arose, which we discuss here. We have not included in the report other issues that came up in our discussions but were associated with only a particular community.

The move to the Full Economic Cost (fEC) funding models and the potential loss of investment from the Strategic Research Investment Fund (SRIF) is causing concern among many providers of research computing – especially at the campus level – with the need to *account* for on and off campus use.

The management of file-based collections (as noted in Section 2) is producing requirements relating to *file replication* to provide redundancy and improved access speeds, both for campus and national resources.

End-to-end security (of the underlying resources and the data accessed through them) is frequently achieved by many service providers through the deployment of *firewalls* and by restricting the number of open ports, both of which continue to provide problems to end-users.

Only one group we spoke with had a pressing need for advance reservation and co-scheduling, and that was currently being well addressed with the *ad hoc* systems in place. The need for scheduling across resources with common environments was seen as a lower priority than other pressing data or security needs by most groups.

## 8 Conclusions and Recommendations

As a result of the survey we recommend investment in three broad areas: *software*, *policy* and *support*, with items listed in no particular order. Sustained investment in these areas will provide a set of structured tools, services and environments to support access to e-infrastructure, and a support infrastructure to enable the adoption of e-infrastructures by new user groups.

Software:

- Automatic data annotation and provenance tools to support domain specific schema
- Mechanisms to support controlled and convenient sharing of files between groups
- Reliable documented software base to enable virtual organisations of individuals to gain access to services and resources, with mechanisms to facilitate collaborative research between those individuals
- License management software, across a site and across a VO
- Simplified authentication tools – better user interfaces and focus on ease of use and local (client) deployment
- Integration of e-infrastructure resources into the environments in regular use by applied end-users (e.g., Matlab, Jython, Perl, R, portals)
- Better tools for understanding failures and increasing reliability

#### Policy:

- Development of a best practice document to support research groups in developing their own data curation and file management policy
- Development of common annotation schemes for individual communities to enable consistent metadata labelling within these communities
- Defining a set of recognised roles, actions, and relationships to support the construction of virtual organisations
- Consistent environment guidelines for machine interactions, especially with the NGS, so users see one environment across a Grid

#### Support:

- Better technical consultancy to end-users who wish to move their applications to use e-infrastructure services, developers who wish to use best practice to build e-infrastructure services, and deployers who need to configure e-infrastructure services for maximum performance and reliability –in general, better information with a human behind it (this needs to be a funded, coordinated service with experts–simple lists are worse than nothing, as they are always dated and often misleading)
- Provision of self-help training materials and hands-on tutorials delivered to users throughout the UK through national and local trainers for common tools and software being used within the UK community
- Outreach and education about what is possible to non-traditional communities

We hope that these recommendations will influence the roadmaps and activities of organisations charged with supporting collaborative multidisciplinary science in the UK (e.g. OMII-UK, NGS, DCC) and their funding bodies – the UK research councils and JISC.

## Acknowledgements

We gratefully acknowledge all of the people we talked to; their time and frank discussions are appreciated. In addition we received valuable feedback and comments from Dave Berry (NeSC), Richard Sinnott (University of Glasgow), and David Wallom (University of Oxford).

Steven Newhouse was supported by OMII-UK, EPSRC grant EP/D076617/1. Jennifer Schopf was supported by the Joint Information Systems Committee and the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under contract DE-AC03-76SF00098, Andrew Richards was supported by the JISC NGS grant.. Malcolm Atkinson was supported by the EPSRC Senior Research Fellowship for the e-Science Envoy EP/D079829/1. Facilities were provided by the NeSC EPSRC grant EP/D054370/1.

## References

- [1] “Study of User Priorities for e-Infrastructure for e-Research (SUPER)”, Steven Newhouse, Jennifer M Schopf, Andrew Richards and Malcolm Atkinson, UK eScience Technical Report UKeS-2007-01, April 2007.
- [2] Workshop on Virtual Organisations and Grids: Applications, Patterns and Infrastructures, UK National eScience Centre, 21-22 November, 2006 [http://wiki.esi.ac.uk/Virtual\\_Organisations\\_And\\_Grids](http://wiki.esi.ac.uk/Virtual_Organisations_And_Grids)
- [3] Jennifer M. Schopf and Steven J. Newhouse, “Real Users - Real Requirements,” UK eScience All Hands Meeting 2005, Sept. 2005.
- [4] Arts and Humanities E-Science Support Centre Matching Agency, <http://www.ahessc.ac.uk/matching-agency>
- [5] ESFRI: European Strategy Forum on Research Infrastructures, <http://cordis.europa.eu/esfri/>

## Appendix A: Projects Visited

### **CARMEN (EPSRC)**

[http://bioinf.ncl.ac.uk/carmen/index.php/Main\\_Page](http://bioinf.ncl.ac.uk/carmen/index.php/Main_Page)

Paul Watson, Phil Lord

### **Gold (EPSRC)**

<http://www.goldproject.ac.uk/>

Rob Smith

### **CRISP (DTI)**

<http://www.neresc.ac.uk/projects/index.php>

Rob Smith

### **ComparaGrid (BBSRC)**

<http://www.comparagrid.org/>

Neil Wipat, Dan Swan, Matthew Pocock

### **OGC (JISC)**

<http://www.opengeospatial.org/>

Gobe Hobona

### **NanoCMOS (EPSRC)**

<http://www.nanocmos.ac.uk/>

Asen Asenov, Campbell Millar, Scott Roy,

Richard Sinnott, Antony Stell, Dave Berry

### **GridQTL (BBSRC)**

<http://www.gridqtl.org.uk/>

John Allen, Jos Koetsier

### **DGEMap (EU)**

<http://www.dgemap.org/>

Richard Baldock, Yin Chen

### **EDIKT2 – Chemistry (SFC)**

<http://www.edikt.org/edikt2/>

Paul Madden, Andrew Turner

### **EDIKT2 – Biology (SFC & Wellcome Trust)**

<http://www.edikt.org/edikt2/>

Ilan Davis, Alistair Kerr, Russell Hamilton

### **Edinburgh Centre for Bioinformatics.**

*Medical Research Institute (MRC)*

<http://www.bioinf.mvm.ed.ac.uk/>

Donald Dunbar

### **Integrative Biology (EPSRC)**

<http://www.integrativebiology.ac.uk/>

Blanca Rodriguez

### **NeuroGrid (MRC), GIME – Generic**

**Infrastructure for Medical Imaging (DTI),**

### **CancerGrid (MRC)**

<http://www.neurogrid.ac.uk/>

<http://www.cancergrid.org>

Sharon Lloyd

### **Integrative Biology Virtual Research**

**Environment (JISC)**

<http://www.vre.ox.ac.uk/ibvre/>

Matthew Mascord

### **Earthquake Engineering (EPSRC)**

[\[civil.eng.ox.ac.uk/research/structdyn/presentations/uknees.html\]\(http://www-civil.eng.ox.ac.uk/research/structdyn/presentations/uknees.html\)](http://www-</a></p></div><div data-bbox=)

Tony Blakeborough

### **Oxford Internet Institute (ESRC)**

<http://www.oii.ox.ac.uk/>

Bill Dutton, Marina Jirotko

### **Oxford Campus Grid**

<http://www.oerc.ox.ac.uk/>

David Wallom, Amanda Barnard

### **Reading e-Science Centre (EPSRC)**

<http://www.resc.rdg.ac.uk/>

Jon Blower

### **Centre for Computing in the Humanities**

**(AHRC)**

<http://www.ahessc.ac.uk/>

Stuart Dunn

### **NERC Data Grid (NERC)**

<http://ndg.badc.rl.ac.uk/>

Bryan Lawrence

### **UCL Research Computing**

Clare Gryce, Jeremy Yates

### **Crystal Structure Prediction (EPSRC)**

<http://www.cposs.org.uk/>

Sally Price

### **OMII-BPEL (EPSRC)**

<http://sse.cs.ucl.ac.uk/omii-bpel/>

Wolfgang Emmerich

### **GridPP (PPARC)**

<http://www.gridpp.ac.uk>

Ben Waugh

### **AstroGrid (PPARC)**

<http://www2.astrogrid.org/>

Jeremy Yates

### **RealityGrid (EPSRC)**

<http://www.realitygrid.org/>

Peter Coveney

### **Cambridge Campus Grid (University of**

**Cambridge)**

<http://www.escience.cam.ac.uk/>

Mark Calleja, Mark Hayes

### **GROWL (JISC)**

<http://www.growl.org.uk/>

Peter Brorsson

### **CCPNGrid (BBSRC)**

Wim Vranken

### **AstroGrid (PPARC)**

<http://www2.astrogrid.org/>

Nic Walton