Proceedings of the UK e-Science All Hands Meeting 2005, ©EPSRC Sept 2005, ISBN 1-904425-53-4

eScience usability: the *e*Minerals experience

MT Dove[1,2], TO White[1], RP Bruin[1], MG Tucker[1,3], M Calleja[1,4], E Artacho[1], P Murray-Rust[5], RP Tyer[6], I Todorov[1,6], RJ Allan[6], K Kleese van Dam[6], W Smith[6], C Chapman[7], W Emmerich[7], A Marmier[8], SC Parker[8], GJ Lewis[9], SM Hasan[9], A Thandavan[9], V Alexandrov[9], M Blanchard[10], K Wright[10], CRA Catlow[10], Z Du[11], NH de Leeuw[11], M Alfredsson[12], GD Price[12], J Brodholt[12]

[1] *Department of Earth Sciences, University of Cambridge, Downing Street, Cambridge CB2 3EQ*
[2] *National Institute for Environmental eScience, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA*
[3] *Present address: ISIS Facility, Rutherford Appleton Laboratory, Chilton, Didcot OX11 0QX*
[4] *Present address: Cambridge eScience Centre, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA*
[5] *Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge CB2 1EW*
[6] *Daresbury Laboratory, Daresbury, Warrington, Cheshire WA4 4AD*
[7] *Department of Computer Science, University College London, Gower Street, London WC1E 6BT*
[8] *Department of Chemistry, University of Bath, Bath BA2 7AY*
[9] *Department of Computer Science, The University of Reading, Whiteknights, Reading RG6 6AY*
[10] *Davy Faraday Research Laboratory, Royal Institution, 21 Albemarle Street, London W1S 4BS*
[11] *School of Crystallography, Birkbeck College, Malet Street, London WC1E 7HX*
[12] *Department of Earth Sciences, University College London, Gower Street, London WC1E 6BT*

Abstract

In this paper we consider the *e*Minerals project as a case study in escience usability from the perspective of the support given to the scientist project members. We report both successes and problems, with some solutions for the latter.

1. Introduction

The UK eScience Usability Task Force (UTF) has identified three issues/challenges regarding the usability of escience[†]

1. Maximising escience technologies to support new forms of global communities;
2. Exploitation of escience infrastructure to support knowledge production and expertise in escience;
3. Design, assessment and management in global escience systems.

The *e*Minerals project [1] is a good case study of how scientists and grid specialists have worked together to create an inherently usable grid infrastructure to support simulation scientists. The development of the *e*Minerals project has followed a logical sequence that does not readily map onto the three issues identified by the UTF, but many areas within these issues have been tackled in the work to develop a usable escience infrastructure for the *e*Minerals project.

2. The *e*Minerals project

The *e*Minerals project has as its primary scientific aim the study of environmental processes at a molecular level, using atomistic simulation models. Examples of the types of process that interest the *e*Minerals scientists are nuclear waste encapsulation, adsorption of pollutant atoms (e.g. arsenic) and molecules (e.g. persistent pesticide chlorine-bearing organic molecules) onto mineral surfaces, and weathering. Some of these studies are combinatorial in nature, requiring many similar simulations to be performed as part of a single study.

Consider the study of organic molecules on mineral surfaces, where a given type of molecule may have several tens or hundreds of possible variants. One example, the PCB family (chemical formula $C_{10}Cl_xH_{10-x}$), consisting simply of two benzene rings

[†] The UTF also identified a challenge concerning trust and ethics, but this is not an issue for the work of this paper.

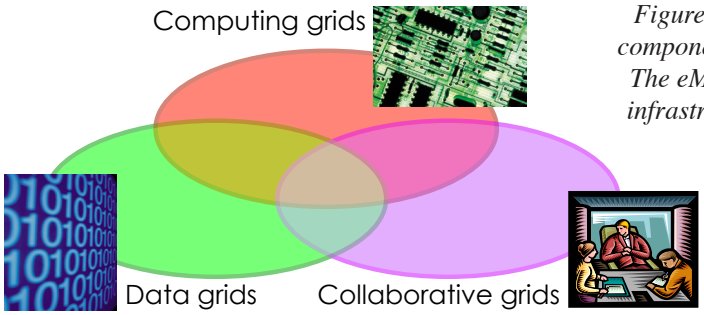

Figure 1: Schematic representation of the three components of the eMinerals grid infrastructure. The eMinerals project has sought to develop an infrastructure that targets all three components.

joined by a bond, with chlorines substituted for some of the hydrogens, has 210 different members (called congeners), differing in the number of Cl atoms (x) and their location within the molecule. It is important to be able to perform calculations on each congener in order to compare the energy of adsorption between molecules with different numbers of chlorine atoms. In addition, each set of 210 calculations needs to be repeated for different adsorption sites on the mineral surface. These and other simulations within the project may each typically require many hours of calculations on a modern PC with 1–2 GByte RAM. A related example is discussed elsewhere in this collection [2].

In addition to the science aims, another aim of the *e*Minerals project has been to develop an escience infrastructure to support the scientists working on molecular-scale computer simulations by developing an integrated compute and data grid infrastructure, coupled with the development of collaborative tools. These three components are illustrated schematically in Figure 1; our aim is to pitch directly for the area of overlap of the compute, data and collaborative grid areas.

The *e*Minerals project team consists of a mixture of simulation scientists, simulation code developers and grid specialists. The close integration of people from the various disciplines within the *e*Minerals team has been one of the keys to ensuring that the resultant escience infrastructure has a high level of usability to the scientists.

3. The *e*Minerals scientists: the users of the escience infrastructure and their requirements

One key to developing a usable infrastructure is understanding the science users, including their needs and their natural ways of working. Our scientists are typical of physical simulation scientists. They have various degrees of computing expertise, but being simulation scientists they naturally have a working-level proficiency in unix. They are capable of editing inherited job-submission scripts, but are not keen to develop new job-submission scripts from scratch. To this extent, they are quite capable of adapting their work to new tools, but even with the high-level of proficiency of the *e*Minerals scientists, usability support has to be built into the heart of any grid infrastructure they are to be expected to use. Like most scientists, their primary motivation is to do the best possible science. Thus, like most good scientists, they are impatient to get results, and will only use tools that actually help drive the science forward. Moreover, due to a number of standard constraints, our team scientists focus more on short-term rather than long-term goals, and any new infrastructure has to be compatible with this goal horizon. Most of our team have not had prior experience of working within a distributed collaborative team (or virtual organisation, VO).

Simulation scientists have high requirements for computing resources. In some cases these can only be met by high-performance resources, but in many cases, as in the example of organic molecules on mineral surfaces described above, these can be met by high-throughput or grid computing resources. Furthermore, the type of problems we are tackling within the *e*Minerals project also require escience methods for management of many jobs (workflow) and many data files.

The primary output from a simulation study

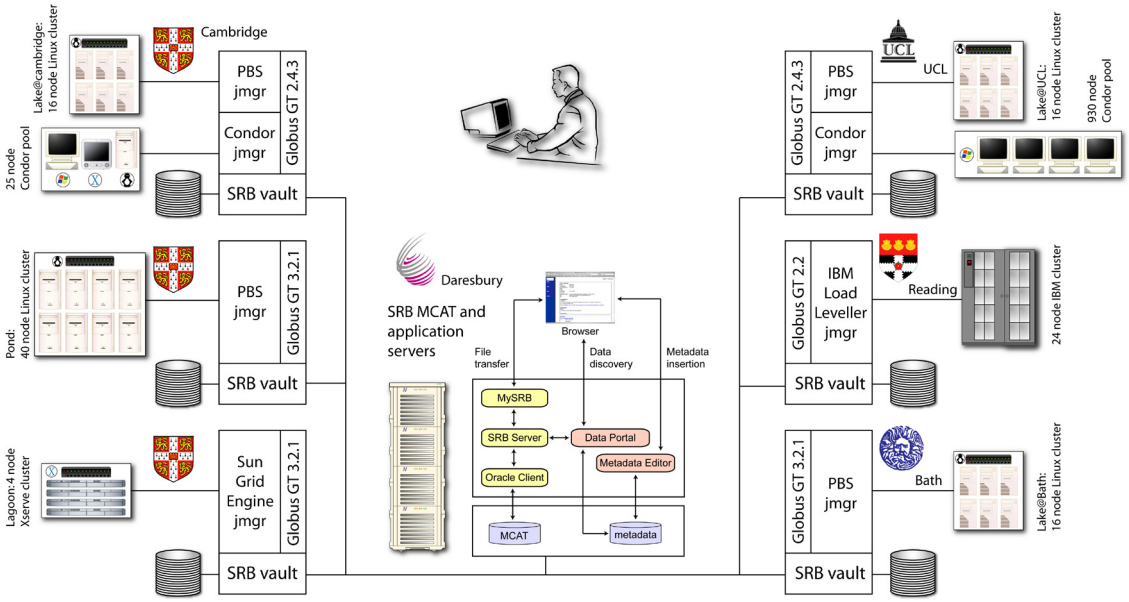

Figure 2: Schematic representation of the eMinerals minigrid, showing the compute resources (Condor pools, Linux and Apple Xserve clusters, and standalone systems), the SRB vaults and SRB MCAT server, and the application server that also runs the eMinerals data portal. The representation also shows the various job managers and the grid middleware employed.

is a set of data files, which typically consist of lists of numbers. Traditionally, scientists manage output files using the standard unix filesystem coupled with separate notes (paper or electronic) to capture information that in principle could be captured as metadata (such as date, simulation details, reasons for running the simulation, relation to other studies etc).

It should be noted that operating within an escience framework may require a change in mindset for the scientist. The traditional approach, which owes as much to constraints of technology as to anything else, involves the user submitting jobs, or batches of jobs, one by one to a specific (and typically limited) set of compute resources, bringing output files back to the user's own computer, and recording the results of one run or batch of runs, before starting the next. The opportunities enabled by grid computing and escience will allow users to have many jobs running across many resources at the same time, with a much higher level of abstraction across the whole process. And yet, within this new abstraction, users may still need to carry out some sets of sequential calculations. The changes in mindset required of the scientists may only partially be helped by the ways that new technologies are implemented.

4. Design of the *e*Minerals minigrid

At the heart of the *e*Minerals project is an integrated compute/data testbed infrastructure called the '*e*Minerals minigrid', Figure 2 [3,4]. The compute resources include Condor pools (one of which is a large campus grid), Linux clusters, an Apple Xserve cluster, and an IBM SP2 parallel computer. Not only is there a heterogeneous mix of hardware and operating systems (by design), the *e*Minerals minigrid also has a mix of job schedulers, including PBS, Sun Grid Engine and IBM Load Leveller. The middleware layer has been built using standard tools, including Condor, Globus (using GT2 functionality) and web services. Authentication is handled by standard GSI security, based on the use of X.509 digital certificates issued either by the UK eScience certificate authority or by the project's own CA.

The data layer is handled using the Storage Resource Broker (SRB), developed by the San Diego Supercomputer Centre [5]. The *e*Minerals minigrid has data vaults associated with each major compute resource, with the central metadata catalogue (MCAT) managed

at Daresbury Laboratory. The users are able to view the entire storage system within the SRB as a single logical file system, with immediate access to files. The SRB is linked to the *e*Minerals DataPortal and metadata editor [6], which allow files to be grouped together as components of studies. The DataPortal enables searching based on the use of topic keywords and metadata, with a key role of enabling effective data archiving and sharing.

As an escience testbed project, we have enforced the rule that resources can only be accessed using client tools and GSI security (i.e. standard ssh logins and ftp file transfers have been prohibited, and gsissh logins are only permitted for code developers and for system administrators needing to perform infrastructure management tasks). These constraints were established as absolutes from the outset, and the task was to design a usable escience infrastructure within these constraints.

Interoperability and design of workflows has been aided by our adoption of the Chemical Markup Language (CML), one of the first applications of XML [7]. We have enabled all our major codes to write CML output.

The final component of the *e*Minerals grid infrastructure is the use of tools to support the scientist and development teams to work as a virtual organisation (VO) [8,9]. This is based on the use of the desktop access grid, supported by a custom-developed multicast application sharing tool, as described below and elsewhere in this collection [10].

We reiterate that our infrastructure was developed through close collaboration between the scientists and grid specialists, which means that lessons learned had an easy feedback into the design process. It also meant that evaluation was immediate, accurate and brutally effective; the ultimate test is that anything that was useful and usable was used, and anything that was less useful was not used. Another test is whether new science is emerging from the escience infrastructure we have set up; results are reported in another paper in this collection [11], in a series of papers in Volume 31(5) of the journal 'Molecular Simulations', and elsewhere [12].

5. Experiences 1: what has proved usable

5.1 Condor and Globus

Early on in the *e*Minerals project, it was noted that the science users adapted to Condor quite easily, much more so than Globus. Our scientists were well able to use the Condor job scripts. We exploited this by using the Condor-G interface to Globus. We have developed a script-based interface to the *e*Minerals minigrid using Condor-G, and incorporating Condor DAGman to enable the use of simple workflows. For use of the *e*Minerals minigrid, a standard workflow is to download input files from the SRB, run a simulation code, and place all output files into the same directory within the SRB. The user provides a small script, an example template of the which is given in Example 1, and this script is then read by a perl program called "my_condor_submit" [3,4]. The user needs to provide relatively little information; the key information required is the resource to be used, and information about the directories and files to be accessed in the SRB (and there is a wildcard option for the surprisingly-common case where the complete list of output file names is not known in advance). This script interface required very little user training.

5.2 The Storage Resource Broker for distributed data management

The SRB proved to be very useful for data management, and again required little user training. The project uses the SRB regularly for management of files by individual users, and also for sharing of data between project members. The SRB also naturally provides a useful archive of users' work processes.

5.3 Resource management

In the initial stages in the use of the *e*Minerals minigrid, we did not have capabilities for resource brokering. Thus users need to target particular resources in their job submission scripts (see Example 1). We have developed a globus-based web resource-monitoring tool, which is illustrated in Figure 3 as accessed

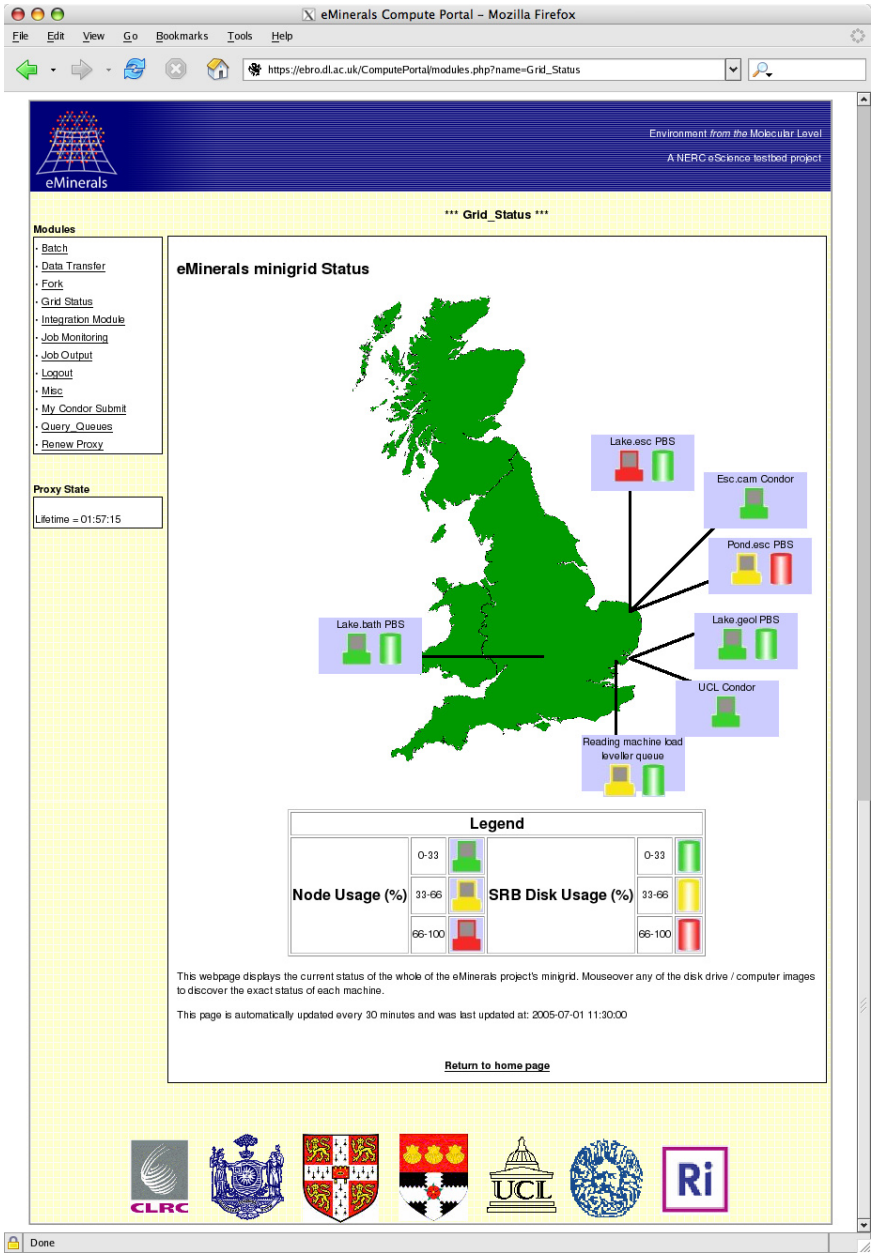

Figure 3. The resource monitoring tool within the eMinerals compute portal. The small icons represent the usage of compute and data storage resources within the eMinerals minigrid.

from the *e*Minerals ComputePortal (see §6.1). We are currently developing a metascheduler for the *e*Minerals minigrid which will make resource selection somewhat easier.

5.4 GSI authentication

GSI authentication proved easy for users to work with. It has been used throughout the minigrid in order to present the users with a single homogeneous security mechanism for compute, data and portal resources. Access to the latter is mediated using the MyProxy credential repository which, in conjunction with the portal interfaces, also allows use of the minigrid when the user does not have access to their normal desktop environment.

5.5 Collaborative tools

To facilitate working as a VO, we have shown that the desktop Access Grid works well [8,9]. It enables team meetings to take place without a lot of forward planning (as would be necessary if using conventional Access Grid suites). To help users, the *e*Minerals team often has informal group meetings, at which the team can discuss usability issues among other things.

We have enhanced our use of the desktop Access Grid with the development of a collaborative multicast-based application

```
Universe            = globus
Globusscheduler     = <minigrid resource>/jobmanager-<jobmanager>
Executable          = <name of executable binary or script>
Notification        = NEVER
#
# Next line is example RSL for a single-processor PBS job
# Modifications are required for other job managers
#
GlobusRSL           = (arguments=none)(job_type=single)(stdin=<filename>)
#
# Next lines are modifications for access to the SRB
#
Sdir                = <some directory in the SRB>
Sget                =     <list of input file names, or * for wildcard>
Sput                = <list of output file names, or * for wildcard>
#
Output              = <standard output file name>
transfer_output     = False
Log                 = <name of log file>
Error               = <name of standard error file>
Queue
```

Example 1: Template script for a job to be submitted to the eMinerals minigrid resources using my_condor_submit (described in the text). This is executed as a condor job submission. Access to the SRB is embedded within the heart of the job submission. The parameters embedded in <...> brackets (coloured red) are fixed by the user, and include file names and parameters associated with required resources.

sharing tool (MAST) [10]. MAST allows geographically distributed participants to share anything from subject specific applications, such as a molecular viewer to generic office programs, such as a PowerPoint presentation. MAST provides a scalable model as it uses multicast for communication within the group. The important point to note here is that the sender only needs to send one copy of the message which is then propagated to all group members. This is in contrast to unicast based tools, such as Virtual Network Computing (VNC).

To complete the discussion on this point, we note that the *e*Minerals VO has made use of two other technologies, namely the wiki concept and instant messaging [8,9]. The *e*Minerals wiki has to be very useful for collating ideas, e.g. for planning team meetings and editing documents.

4.6 Chemical Markup Language

The *e*Minerals project has put effort into enabling many of the simulations codes we use to produce XML output for at least the basic results (it is not practical to work with large atomic configurations in XML). We use the Chemical Markup Language (CML) for this [7,13]. We have found that this has worked very well for a number of reasons. It makes post-processing of output from many runs carried out in a grid infrastructure much easier for the science users. Moreover, using XSLT (eXtensible Stylesheet Language Transformations) we have been able to generate HTML representations of output files with additional transformations of data to SVG (Scalable Vector Graphics) for plotting of data. The important point is that this can be achieved in a platform-neutral way.

6. Experiences 2: what hasn't worked, and steps taken to overcome usability issues

6.1 Client tools

Installation of client tools has often been found to cause difficulties. Tools such as Globus were found to be particularly difficult, sometimes even impossible, for some users, because they are poorly wrapped and installation procedures needed modifying to get them to work. Our initial solution was to set up a number of 'submit machines' that are properly configured and maintained. Users log into these machines in order to submit

their jobs. Although this enables the users to access the minigrid, it is much less desirable than users having direct access to the minigrid from their own desktops.

In order to achieve desktop access to the minigrid, we have been working on two approaches. The first is to develop the *e*Minerals 'ComputePortal' [14], through which users gain access to all the client tools they need without having to have a local installation. The ComputePortal provides a portal interface to the *e*Minerals minigrid using functionality provided by Globus and Condor client tools and libraries. It exposes job submission, job monitoring, data transfer and workflow functionality within an intuitive web interface. In particular, the workflow integrating computation and data management provided by the Condor-G DAGman is also accessible via the portal.

GSI security in this case is handled with a MyProxy server, which users have found to be easy to use. Each of the processes which the user previously used to decide where to run their computation and to setup and submit their job is catered for via the ComputePortal but the only requirement on the user's computer being that they have a modern web browser and a working java installation. The java installation is required to use a MyProxy upload tool so that the user's certificate is never actually sent across the network.

The use of the compute portal also has the benefit of much simplified firewall administration as the only firewall hole required to or from the user's machine is for outgoing https access on port 443 which is generally allowed. This requirement when compared to the complex set of holes required for each of the middleware tools used (Condor, Globus, etc.) is obviously much easier to fulfil.

As the ComputePortal is effectively a central submission resource for all users' jobs it can also cache details of jobs submitted by all users, giving them fairly advanced monitoring capabilities including access to simulation output files etc. This is not always the case in a grid setting. It can also act as an archive of user's work, allowing them to easily check what work was submitted to where and when.

Our second approach is to use the recently developed Web Services interface to Condor [15]. A separate advantage of using Web Services from a usability perspective is that this technology allows multiple independent grid services to be composed into a single unified computational process. This is a particularly important and powerful feature when users want to specify complex workflows that define a sequence of tasks to be performed, which could be batch jobs to be executed but also other types of activities such as data storage operations. An example is a simple workflow to replicate the processes followed in my_condor_submit described above (§5.1), including a web services interaction with the SRB.

As an increasing number of grid technologies provide support for Web Services, these can be composed into user-defined workflows to provide an integrated service using technologies such as the Business Process Execution Language (BPEL). BPEL is a workflow specification language that allows users to specify a set of Web Service invocations to be performed. Though initially intended for business-to-business interaction, BPEL proves to be a powerful tool to orchestrate grid services. A BPEL editor has been in development at UCL in the context of an OMII funded project that aims to make BPEL more accessible to scientists by providing a graphical user interface that hides from the user most of the complications of writing BPEL scripts. Users can use this tool, available as a plug-in for the freely available eclipse development environment, to define their own computational workflows from their desktops and deploy them in our mini-grid environment.

6.2 Firewalls

Firewalls are always a problem, not only for individual users, but also for robust running of the minigrid. We have experienced partner institutions changing firewall policies, leading to resources no longer being available to other partners of the VO. This is an irritation with regard to the sudden loss of computational resources, but unacceptable when it comes

to the sudden loss of access to data within the data grid component of the *e*Minerals minigrid.

The *e*Minerals ComputePortal and web services interface to Condor both help solve some of the problems posed by firewalls, in that much of the communications traffic is directed through a small number of standard open ports.

6.3 Metadata

Although the *e*Minerals minigrid provides tools to manage metadata, namely the DataPortal and Metadata Editor [6], the lack of automatic metadata harvesting has limited the uptake of these tools as the process of metadata entry is still somewhat manual. It would be advantageous to harvest metadata from both the output of the simulation codes and the job submission process itself. The migration towards an integrated portal environment, along with the move towards using CML based data formats, will allow much more automated metadata management strategies to be realised in the near future.

In order to address the lack of automation with respect to metadata capture, a suite of tools known as the RCommands have been developed. These tools have both client side and server side components. The client side components are shell commands that allow metadata to be inserted, modified, searched etc. This is advantageous as the *e*Minerals scientists are comfortable working with shell commands. Moreover, these tools are scriptable. The Rcommands have been integrated into the *e*Minerals my_condor_submit command, allowing for the data output by simulations to automatically be associated with studies in the DataPortal and to have metadata entered automatically as part of the job submission process and in keeping with the my_condor_submit condor style input file.

The server side components are exposed via a web service interface which allows the same functionality to be accessed from within both the DAGMan and BPEL workflows. This separation removes the need to expose the back end database holding the metadata database and also abstracts the end users and client tools from any changes made to the metadata schema.

7. Summary

In this paper we have presented a number of lessons concerning escience usability based on the use of the *e*Minerals project as a case study. By integrating both end-user scientists and grid specialists within the one team, and providing collaborative tools to support the team, it has been possible to develop an integrate compute and data grid infrastructure that is being used in production mode by the simulation scientists. A number are areas of future development have been highlighted.

Acknowledgements

We are grateful to NERC for financial support.

References

1. Dove MT & de Leeuw NH, *Mol Sim* **31**, 297, 2005
2. White TO et al, *All Hands 2005*
3. Calleja M et al, *All Hands 2004*, p 812
4. Calleja M et al, *Mol Sim* **31**, 303, 2005
5. Moore RW & Baru C, in *Grid Computing: making the global infrastructure a reality*, ch 11, 2003
6. Blanshard LJ et al, *All Hands 2004*, p637
7. Murray-Rust P et al, *J Chem Inf Comp Sci*, **39**, 928, 1999
8. Dove MT et al, *All Hands 2004*, p 127
9. Dove MT et al, *Mol Sim* **31**, 329, 2005
10. Hasan SM et al, *All Hands 2005*
11. Alfredsson M et al, *All Hands 2005*
12. Wells SA et al, *All Hands 2004*, p 240
13. Wakelin J et al, *Mol Sim* **31**, 315, 2005
14. Tyer RP et al, *All Hands 2004*, p 660
15. Chapman C et al, *All Hands 2005*