User requirements for UK e-Science grid environments

Bruce Beckles

University of Cambridge Computing Service

Abstract

Following recent analyses of the Globus Toolkit [1] from a UK perspective, e.g. [2], [3], the UK Grid Engineering Task Force (ETF) felt it was necessary that a user requirements exercise [4] be undertaken. This paper presents the background to this exercise, outlines the current proposed plans for this exercise and discusses the need for such an exercise.

1. Introduction

In January 2004, a workshop entitled "Requirements Capture for Collaboration in eScience" [5] was held at the National e-Science Centre at which two papers ([2], [3]) were presented that examined various aspects of the then current versions of the Globus Toolkit pertaining to the Toolkit's suitability for the UK e-Science community. Both these papers found that there were a number of problems with the then current versions of the Toolkit, and both recommended alternatives to the Toolkit be investigated or developed.

Both of these papers were discussed by the UK Grid Engineering Task Force (ETF) – the body at that time officially charged with matters concerning the deployment of computational grids in the UK e-Science community – and a consensus was reached that the ETF should, in collaboration with other relevant bodies, undertake a requirements capture exercise [4] whose aim was to inform the ETF of the end-user requirements for grid technology amongst the UK academic community.

This paper outlines the current plans for this exercise and discusses the need for such an exercise, as well as presenting some useful background material in this area.

2. Background

One of the papers [2] referred to above highlighted particularly strongly the "gap" between the apparent user requirements for computational grid middleware and the capabilities of the then current Globus Toolkit (frequently described as "the leading grid middleware toolkit"). A close reading of other work that has attempted to elucidate the requirements for e-Science in relation to the field of grid technology, such as the e-Science Gap Analysis [6], would also have revealed that the apparent needs of the UK e-Science community were not being addressed by developments in the Globus Toolkit.

Discussions between the author of this paper and members of the development team for the Globus Toolkit have revealed that "end-user requirements" have no direct influence on the Toolkit's development path, and indeed the Globus Alliance [7] does not undertake either requirements capture exercises or user consultation exercises. It is therefore unsurprising that this "gap" continues to show little sign of closing.

The ETF has been attempting to deploy, at a national level, the infrastructure necessary for computational grids which will be of service to the UK e-Science community, in particular, and the wider UK academic community in general. Until recently there has been a very strong preference for using the Globus Toolkit for these activities, so much so that many members of the ETF have privately expressed frustration at the national UK e-Science Core Programme's apparent refusal to consider other grid middleware toolkits. This has meant that it is extremely difficult to construct a credible claim that the grids being built by the ETF are genuinely useful to the stated target audiences (in the sense of meeting their needs).

Although the announcement of the Web Services Resource Framework (WSRF) [8] by the Globus Alliance in late January 2004 precipitated a change in the attitude of the UK e-Science Core Programme which has meant that other grid middleware toolkits are now being evaluated from a technical point of view, it is still the case that, at the time of writing, little has been done to collect and analyse the requirements of the users for whom one or more of these toolkits will eventually be deployed.

Whilst the importance of requirements engineering is still not accepted in all parts of the software development industry, it is generally accepted that, at least in large software development and/or deployment projects, best practice dictates at least the

development of user scenarios, whose purpose is to illustrate the intended use of the eventual goal of the project. The ETF currently does not have such scenarios at its disposal, nor does it have any systematic collection of requirements data which it can draw upon to base its decisions. This suggests that, should the ETF deploy grid infrastructure which does, in fact, meet the needs of the UK e-Science community, it will have done so more by luck than by judgement.

The ETF is not insensible to these problems, hence its agreement to undertake the requirements gathering exercise [4] previously mentioned. Regrettably, the UK e-Science Core Programme has deemed this exercise to fall outside the remit of the ETF (*pers. comm.*), and a similar exercise has yet to be undertaken by any other official body created by the e-Science Core Programme. This means that the ETF is unable to officially devote any resources to this exercise and so, although the proposals presented here have been largely agreed by the ETF, it has not yet been possible to implement them. A discussion of ways of circumventing this problem is given in Section 5.

3. The Need for Requirements Analysis

One of the most remarkable results of the study reported in [2] was just how badly the then current versions of the Globus Toolkit failed to meet the apparent requirements of end-users. Of the eight derived "technical" requirements for computational grids, the versions of the Globus Toolkit examined failed to adequately satisfy any of them (see Section 5 of [2]). In addition, the authors found that they were unable to locate "a clear statement of what [the Globus Toolkit] intended to achieve and how it intended to do this" and that there were no "formal specification or verification documents" for the Toolkit.

An examination of the development team of the Globus Toolkit [9] reveals a large number of distinguished academics and experienced programmers in the field of distributed computing. In addition the project has significant industrial backing and input from large, experienced IT organisations such as IBM and Hewlett-Packard. It would, therefore, be unreasonable to assume that these problems are due to a lack of technical expertise or to poor quality program code. It seems more likely that the problem lies elsewhere.

If one wishes to develop a solution to a problem faced by a particular user community, but no clear statement of the problem exists, then it seems obvious that it will not be possible to develop a solution systematically. In such cases only a combination of luck, intuition and repeated iteration through the space of possible solutions will achieve the desired result. Unfortunately, for any complex problem, the space of possible solutions will be quite large, and even more regrettably, the more complex a problem the less reliable is one's intuition about its solution. This suggests that a clear statement of the problem one is trying to solve is an essential prerequisite to efficiently obtaining a solution.

Consider the problem that faces the ETF – to deploy grid infrastructure that will be useful to the UK e-Science community and the wider UK academic community. The key word in this succinct expression of the ETF's stated aim is "useful". The word "useful" suggests that it is the *utility* of these grids that is important; that these grids are there to serve a particular purpose, and are not "science (or software development) for its own sake".

So how is it possible to evaluate whether or not the deployed grid infrastructure is "useful"? Presumably by assessing the "usefulness" of the deployed grids to the relevant communities of interest. The most obvious assessment would be to ask a representative sample of the communities in question whether these grids were useful to them.

So far, only one study has attempted to do something even remotely similar to this, namely the paper [2] cited at the start of this section. Viewed from this standpoint its results make depressing reading. So how could the ETF do better? Perhaps a clearer statement of the problem would help – how can this be achieved?

Clearly, a more explicit definition of what constituted "usefulness" would make it more apparent what the ETF's design goals should be. There is only one group of people who can definitively state whether or not the grids built by the ETF are "useful" in the sense discussed here – namely, the intended end-users of these grids. So if the intended end-users could state what would be useful to them in the context of computational grids, then it would be possible to know, *in advance*, whether a particular grid, or type of grid, was likely to be useful *before* it was deployed.

This has numerous advantages, some of which are listed below:

- Little time, money and other such scarce resources would be wasted deploying grids for which the end-users have no use.

- The end-users would have grids that would actually be useful to them, and hence they would be both happier and more productive than they are with the current generation of grids.
- A set of objective criteria would be available against which any design or deployment decisions could be adjudicated, instead of such decisions being made on the basis of intuition, or to appease short-term political interests.

The objection most normally raised to this idea can be most simply expressed as follows: "That's all very well, but users do not know what they want, therefore they cannot tell us." There is some truth in the statement "users do not know what they want" (although that statement is unnecessarily and unfairly patronising to users). However this is a problem that has been extensively researched and several highly effective techniques have been developed to address this problem. This is, in fact, one of the central problems of requirements engineering.

A number of investigative techniques and design methodologies, which can be grouped together under the heading "user-centred design", have been developed which effectively address this problem: contextual design [10], interaction design [11], Sommerville's viewpoints and concerns method [12], etc.

These methodologies all have the following characteristics: they do not attempt to ask the cognitively difficult question "what are your requirements for this system?" of the end-user. Instead a more indirect approach is taken to obtaining this information.

For instance, a user might be interviewed in their workplace, and encouraged to demonstrate their normal routine in the problem space under examination to the interviewer, who will be as unobtrusive as possible. The user will be asked to walk the interviewer through how they would approach a particular problem, as though the interviewer were a trainee or "apprentice" who is being shown the ropes.

The interviewer will note the details of the user's approach to this particular task, highlighting apparent problems and inefficiencies, as well as areas where the current approach works well. They will confirm these subjective assessments with the user, either at the time (if this can be done without breaking the flow of the current task), or retrospectively. In this way an accurate and effective picture of the user's current work practice is built up.

Each of the user-centred design methodologies provides a framework for the analysis of this data that allows the designer to determine the areas where improvement is needed, what aspects of current practice need to be retained or supported, and, perhaps most crucially, what exactly is the problem that needs to be solved.

These design methodologies produce a very focussed design in which the requirements of the user are clearly stated and paramount, and as such are particularly well suited to areas, such as grid technology, where the user requirements are particularly poorly understood.

4. Detail of Proposed Plan for Requirements Capture Exercise

With the above considerations in mind, the following plan of action is proposed to address the current difficulties faced by the ETF in determining what is necessary for their deployed grid infrastructure to be useful to their communities of interest. This plan is divided into the following four phases: *review and determination of contacts*, *requirements elicitation*, *requirements analysis* and *validation*. These phases are described below.

4.1 Review and Determination of Contacts

Members of the ETF will be canvassed for their opinions on the issues, areas and people that they feel are relevant to this exercise. Detailed interviews will be conducted with those members of the ETF who are particularly knowledgeable or forthcoming on any of the identified issues.

Through the ETF, the e-Science Centres (both the regional centres and the Centres of Excellence) will be asked to briefly describe (including the general area, e.g. particle physics, and contact details) all the projects which they support or with which they are in contact. For those projects which they support or have particularly well-developed links, they will be asked to identify a contact person who would be willing to discuss that project's requirements.

The list of attendees of the UK e-Science Stakeholders' Meeting on 10 July 2003 (organised by the UK e-Science Core Programme) will be examined for potential contacts. In addition to these activities, each e-Science Centre will be asked to identify likely individuals and/or projects, with which they have links, whose work might benefit from grid technology. Target individuals/projects will probably be those who:

- work with extremely large quantities of data;

- require large amounts of processing power;
- have high data bandwidth or throughput requirements;
- require access to remote computing or data resources which provide otherwise unavailable facilities.

Further contacts will be obtained using a snowball technique [13]: the contacts obtained thus far will be asked whether they know anyone who fits with the above criteria, and any new contacts thus obtained will be contacted and asked the same question.

4.2 Requirements Eliciatation

An analysis of the contacts gathered in the previous phase will be undertaken to determine the key stakeholders and a reasonably representative sample of 'representative' users (the 'representativeness' of a user is determined by their relationship to the underlying 'archetype variables' of the target user community [11]). The individuals thus identified will then be contacted and, if agreeable, will be interviewed in their place of work, using the interview techniques described in [10] or [11]. It is envisaged that five interviews per identified 'user group' will be sufficient, operating on the assumption there will be of the order of ten such groups.

It is hoped that the same interview team (consisting of one or two interviewers) will be able to conduct all these interviews, but this may not be practical depending on the resource and time constraints under which this exercise will operate. The raw data gathered from these interviews will be made available to any interested third parties.

4.3 Requirements Analysis

The data gathered from the interviews conducted in the previous phase will be analysed according to the most appropriate methods ([10] or [11]) to produce a series of user personas and user scenarios that will be used to 'bring to life' the model implicit in this analysis and to make it accessible to developers.

A set of use cases and technical specifications will be produced from these scenarios so that the results of this exercise can be directly utilised by those who see the work of the ETF principally in technical terms. In addition, the data collected and analysed in this exercise should enable a general picture of the requirements for grid technology in an academic context, and an outline design for a general computational grid, to be developed.

4.4 Validation

Following well-established survey techniques [14], a suitably representative sample will be drawn from the contacts gathered in the first phase of this exercise. A questionnaire will be designed which will describe the newly generated requirements, including the requirements originally listed in the questionnaire used in [2] as a control group, and respondents will be asked to rate them in importance. A suitable sampling fraction will be used to ensure statistical representativeness in responses, and there will be two waves of questionnaires, the second wave to chase up non-respondents.

5. A Different Approach

As mentioned in Section 2, the plan outlined in Section 4, though largely agreed by the ETF, has not been carried out because the UK e-Science Core Programme has deemed it to be outside the remit of the ETF. Regrettably, there seems to be no officially sanctioned body into whose remit it does fall that is prepared to undertake it. Therefore an alternative approach, within the limited resources of the author of this paper is being actively considered.

The author of this paper is based in the University of Cambridge, and is responsible for the promotion of e-Science and its related technologies within the University. This aspect of their role means that some of the activities consonant with a requirements gathering exercise of this nature fall within their scope.

The University of Cambridge is one of the leading research universities not only in the UK, but also in the world. This means that it is likely to have academics involved in most areas of research in which the UK academic community is active, and thus it should be possible to find a reasonably representative sample of the UK academic community within the University.

Making use of the author's local contacts, and taking advantage of their location within the University, it should be possible for the author to contact a sufficiently large number of researchers in the University to gather useful data for a requirements gathering exercise similar to that proposed in Section 4. Contact has also been made with academics in the University who have an interest in HCI and/or requirements analysis, who have agreed to provide advice and assistance in such an exercise.

It is hoped that it will still be possible to gather contacts as described in Section 4.1 and

conduct a validation survey (as described in Section 4.4) – as the survey could be conducted electronically, the costs should be minimal. The author continues to actively seek collaborators in related areas, both within the ETF and in UK academic institutions, and it is hoped that, although the original exercise proposed in Section 4 may never be carried out, a similar exercise, albeit on a smaller scale will still produce sufficient information to assist the ETF and other relevant bodies in designing and deploying grids that have a good chance of being actually useful to their target audience.

6. References

[1] The Globus Toolkit: http://www-unix.globus.org/toolkit/

[2] Beckles, B., Brostoff, S., and Ballard, B. A first attempt: initial steps toward determining scientific users' requirements and appropriate security paradigms for computational grids (2004). Proceedings of the Workshop on Requirements Capture for Collaboration in e-Science, Edinburgh, 14-15 January 2004, 17-43.

[3] Chin, J. and Coveney, P. V. Towards tractable toolkits for the Grid: a plea for lightweight, usable middleware (2003). Proceedings of the Workshop on Requirements Capture for Collaboration in e-Science, Edinburgh, 14-15 January 2004, 44-48.

[4] Grid Requirements Interest area of the Grid Engineering Task Force: http://www.grids.ac.uk/ETF/public/Requirements/index.shtml

[5] Requirements Capture for Collaboration in e-Science: http://www.nesc.ac.uk/esi/events/320/

[6] Fox, G. and Walker, D. e-Science Gap Analysis (2003): http://www.nesc.ac.uk/technical_papers/UKeS-2003-01/GapAnalysis30June03.pdf

[7] The Globus Alliance: http://www.globus.org/about/

[8] WS-Resource Framework: http://www.globus.org/wsrf/

[9] The Globus Alliance Team: http://www.globus.org/about/team.html

[10] Beyer, H., and Holtzblatt, K. Contextual Design: Defining Customer-centered Systems. Morgan Kaufmann Publishers, London, 1998.

[11] Cooper, A., and Reimann, R. About Face 2.0: The Essentials of Interaction Design. John Wiley and Sons, New York, 2003.

[12] Sommerville, I. An Integrated Approach to Dependability Requirements Engineering (2003). Proceedings of the 11th Safety-Critical Systems Symposium (Bristol, UK, 2003), Springer, 2003, 3-15.

[13] Oppenheim, A. N. Questionnaire Design, Interviewing and Attitude Measurement. Continuum International Publishing Group – Academic and Professional, 2000.

[14] Rosenthal, R., and Rosnow, R. The Essentials of Behavioural Research (second edition). McGraw Hill Book Company, Singapore, 1991.