# Types of grid users and the Customer-Service Provider relationship: a future picture of grid use

***A paper for the UK e-Science All Hands Meeting, September 2006***

Mark Norman,
University of Oxford

## 1. Abstract

Who will be the grid users of tomorrow? We propose a categorisation of 'future grid' users into the following categories: Service End-User, Power User (with three distinct sub-types), Service Provider and Infrastructure Sysadmin. A further basic type could be argued as Third Party Beneficiary. This paper outlines the possible characteristics of these 'types' of users. For users that have layers of applications or, for example, a portal between them and the grid resource, it is almost certain that heavyweight security solutions, as we have with client digital certificates, are too onerous and unnecessary. It is likely that some users will, however, need client digital certificates, due to the level of control that they may exert on individual grid resources. We also outline a Customer-Service Provider model of grid use. It may be that authentication and authorisation for the SEU 'customers' should be the responsibility of the Service Providers (SPs). This would hint at a more legal framework for delegating authority to enable grid use, but one which could be more secure and easier to administer. Such a model could also simplify the challenges of accounting on grids, leaving much of this onerous task to the Service Providers.

## 2. Introduction

### 2.1. The ESP-GRID project

The Evaluation of Shibboleth and PKI for Grids (ESP-GRID) project's central aim was to achieve a deeper understanding of the potential role that Shibboleth can play in grid authentication, authorisation and security. One of the main outcomes of the project has been that Shibboleth is applicable to some users in some situation and client-based PKI is applicable largely to more technical users in other situations. This gave rise to an examination of the future types of grid users. This arose from a series of brainstorming and consultation sessions with current grid users and developers. The inking within this paper represents some of this output supported by anecdotal evidence and findings in the literature.

### 2.2. When will the grid be really useful?

Before Netscape's browser, Mosaic, was given away free in 1994, the Internet was the domain of the educated and technically knowledgeable. Even within that educated elite, the use of the Internet was dominated by a few research subject areas, possibly arenas in which the development of computing itself had been highly relevant for many years. What changed? The introduction of a graphical interface that was easier to use and was more intuitive did increase the rate of uptake of home computing.

Many interested groups must hope that grid technology must be approaching the metaphoric 'release of the browser' stage some time soon. Whether there will be a surge in take-up, as seen with Internet technologies after 1994, or whether it will be a more steady increase remains to be seen. However, it is the availability and ease of use to the greater community that will make the breakthrough. This paper is focussed mainly upon the educational and research use of grid technology. The engagement of the average citizen with grid technology will take much longer. We believe that the experience of take-up

of the Internet is relevant to the divide between 'researchers experienced in programming or scripting' and the 'rest' of the research community.

### 2.3. Are we talking to the right users?

Anecdotal evidence of researchers refusing to engage and benefit from grid technology suggests that when an application interface is presented that is easy to use, the uptake is strong (e.g. Sinnott, 2006). As the Market for Computational Services Project notes, the inability to use a simple 'service' such as a resource broker in itself leads to a lack of ease of use and little motivation for the end user leading to little or no take-up for real use (Grid Markets, 2003).

Authors have previously noted that the current grid middleware is too intimidating for many users, and have often focussed on the security aspects (e.g. Beckles, 2004a). These aspects are important as they are often the most onerous for the non-computer specialist. In temporary lieu of the work, noted above, to collect requirements from current non-users (Beckles, 2004b), we believe that we should examine the types of users that are emerging within grid computing and consider their generic security profiles as well as their likely access management requirements. This may assist in identifying such users in order to carry out a real world requirements analysis. However, until such an analysis is made, our work is merely a guide to the likely categories of users.

We have, by necessity, very technical users at present. This may distract us from building an accessible grid for future users who may be far less computing-technical.

The following sections of this paper present our view of these users of tomorrow. This is a personal view, based partly on recent experience within the ESP-GRID project and partly on predictions arising from the use of the Internet and the Web.

## 3. Types of grid users

### 3.1. Categories of users

Table 1 presents a summary of the types of grid users that exist, or that will exist in the very near future. Clearly, as with any

Table 1 Grid users of the future

| Type of user | Typical characteristic | Main role |
|---|---|---|
| SEUD | Service End-User (data). Little or no computing expertise. | User of applications served by SPs. Uploads data or runs queries. |
| SEUX | Service End User (executables). Some understanding of code creation. | As SEUD, but runs either executable code or scripts via SPs |
| PUA | Power User Agnostic of grid resource node. High degree of computing expertise. | Develops programs and data but does not care where processing takes place. |
| PUS | Power User requiring Specific grid resource nodes. High degree of computing expertise. | As PUA but may have more platform etc. dependent expertise and some sysadmin expertise. |
| PUDS | Power user Developing a Service. High degree of computing expertise. | As PUA/PUS but developing expertise like SP. |
| SP | Service Provider. High degree of computing expertise. | As PUA/PUS but has expertise in authorisation and possibly identity management. |
| Grid-Sys | Infrastructure sysadmin. High degree of computing expertise. | System administration of grid nodes, possibly with infrastructure delivery and security expertise. |

'categorisation' activity, there will be users who move frequently between the groups, and whom may occupy two or more categories simultaneously. However, we believe that the categories are useful in examining high level requirements, especially those of access control and security.

Note that there are clearly omissions from Table 1. Two notable actors are the Third Party Beneficiary (TPB) and Resource Owner. A TBP could be a person or organisation who/which does not interact directly with the grid but whose personal data are being handled on the grid. Resource Owners clearly have important functions, but they do not necessarily interact with the grid, unless playing one of the seven main roles shown in Table 1 at a particular moment in time. In designing future grids, the requirements of both of these actors would have to be given much thought and would impact upon the likely architecture and security mechanisms of those grids. However, for the purposes of this paper, the general requirements (or expectations) of only the seven main roles are considered in relation to access management and other security needs.

Throughout this paper, the abbreviation SEU is taken to represent SEUD and SEUX where a statement could apply equally to either category.

On the ESP-GRID project wiki at http://wiki.oucs.ox.ac.uk/esp-grid/UserCategoryExampleActivities, we outline some example illustrations of these seven major actors. There was not room, in this brief paper, to describe them here. The majority of today's users come into the PUS and Grid-Sys categories (see Table 3, below).

Note that we have not divided the users into the kinds of grid jobs that result from their activity. This may, however, be another valuable approach. For example, one type of user may run a job that executes (or interacts with) only one grid node (a 'single point' job) whereas another may run a job that is divided, or subsequently splits, into many sub-jobs that interact with many grid nodes. These are useful definitions but are probably applicable to nearly all of the actors in Table 1.

## 3.2. Access management characteristics of these actors

Table 2 describes the access management or security characteristics of the seven user types. The final column of 'Security risk to grid node' tries to capture both the concepts of the threat to the grid resource and the risks (or costs) associated with managing these kinds of users. For example, the threat from an individual user of this type may be fairly low, but the difficulties of managing many users of this type give rise to an associated threat of attackers posing as those users. These are separate concepts, but they have been combined in this case, as it would appear to be appropriate.

The SEUD only ever uses a service, probably presented through some sort of gateway to the grid beyond. Therefore, the security risk to the grid resources from this user should be much lower than the other users. It is assumed that this user cannot interact directly with any grid nodes. Whatever threats that may exist from this type of user, there is the added defence of a restrictive application layer between the user and the grid node.

A similar profile could be expressed for the SEUX. However, a greater risk exists from those users due to executable code being run. Note also that there are similarities between the SEUX and PUA, using a resource broker. The main differences are in computing expertise, the use of a SP and that one is a true end user.

The PUA does not interact directly with any grid node (apart from the resource broker) and therefore should pose a lower

security risk than the other users, apart from the SEU. Nevertheless, code written by or submitted by the PUA will be run on a grid node somewhere and therefore the security risk may be seen as being moderate.

The PUS interacts directly with grid nodes, running code on those nodes. The security risk, from the viewpoint of the resource owners, is therefore much higher from this type of user and, if her identity is not known, it would be a requirement that

Table 2 Access management/security characteristics of the seven user types

| User/actor | Access management/security characteristic | Security risk to grid node |
|---|---|---|
| SEUD | SEUD does not need to be 'known' by a grid access management service (should one exist) as the grid trusts and accounts the SP not the user. SP may need to authenticate, authorise and account for the user as well as possibly taking on 'metering' responsibilities. | Low<br><br>(shielded by gateway/ application). |
| SEUX | The SEUX may have a similar access management characteristic to the SEUD due to the possible greater absolute numbers. The presence of SEUX will probably mandate the automated trace/isolate functionality discussed in section 4.3 on page 10. | Moderate. |
| PUA | The PUA's identity need not be managed by a grid access management service (should one exist) but some sort of mapping to a billing account may be necessary. It could be possible for the identity of the PUA to be concealed behind another entity, as occurs with the SEU. This entity could be a SP providing grid brokering services. Either the SP or the grid access management service is likely to require status (and other) information from an identity manager/provider for authorisation purposes. | Moderate<br><br>(shielded by resource broker). |
| PUS | As for PUA, above in some scenarios. However, in addition, grid node owners may wish to have a direct authentication, authorisation (and accounting) relationship with the PUS. Alternatively, authentication elsewhere may be acceptable if a more transparent assertion of identity is given in order to satisfy the security instincts of grid node owners. | Moderate/high. |
| PUDS | As for PUS but moving into arrangements like SP (see below). May need to begin interacting with and accounting for SEUs in an experimental manner. | High<br><br>(as for SP, see below). |
| SP | A SP may be trusted to provide services only to those authorised to use the grid or the SP may offer services to any end user, and be simply billed by the grid, or by the nodes that it uses. The SP may wish to manage identities and to authenticate SEUs or the SP may be willing to devolve these tasks. The SP probably needs to manage or recognise status (authorisation-related attributes). The SP needs strong/secure assertions of identity/authentication between it and the grid resource nodes. Accounting may be required between the grid resource nodes (or access management service) and the SP and between the SP and the SEU, although this latter requirement may not need to be met using grid middleware.<br><br>As an individual, the SP could use any method (including that of devolved authentication) of access management to his/her machines (to which the SEUs connect or utilise in some way). Moreover, those machines may or may not be considered to be part of the grid. | High<br><br>(impacts security both of grid nodes and of SEUs). |
| Grid-Sys | A Grid-Sys is likely to need to authenticate directly to particular grid resource nodes. However, in theory, it is possible that he may authenticate elsewhere and the node computer may trust that external authentication point (or identity provider). | Moderate<br><br>(High risk but more easily managed). |

it could be traced easily and very quickly, should any 'breach in security' occur. The benefits of a system whereby the user can be traced accurately, when problems occur, should outweigh the benefits (if there are any) of logging explicit identities at each grid node. See Norman (2006) for a further examination of the issues of asserting explicit identity 'up front'.

The PUDS has a similar security profile to the PUS, but is beginning to take on some of the aspects of a SP and therefore could pose a threat to both the grid and to the test SEUs involved in the development. When interacting with the grid, there may therefore be a requirement for the PUDS to be explicitly identified.

As Table 2 indicates, the SP has a complex security profile. The SP (machine) is likely to be trusted by and/or to be explicitly identified to both SEUs and to grid nodes. The SP (user) is also likely to have a profile similar to a PUDS when developing and testing and connecting to grid nodes directly. The SP machines may or may not be considered as part of the grid: these machines may simply be gateways to the grid and not contribute directly to grid computation.

The security profile of the Grid-Sys has been expressed as 'Moderate'. This is due to two opposing influences. Firstly, for each grid node, there will be very few system administrators, almost certainly in single figures. This means that the task of managing these users' authorisation information – and possibly authentication mechanisms – is relatively simple. Secondly and conversely, if an impostor were to be able to bypass the access management system, the risks are very high to the grid node.

### 3.3. An access management scenario

#### *3.3.1. Two major routes of entry to the grid*

Figure 1 outlines a likely scenario illustrating the access management 'behaviour' of these different types of grid users. As we have already established in Table 2, there are a variety of ways in which the access management requirements of each set of users, and of each resource protecting itself from each user, may be fulfilled. The scenario presented in Figure 1 is merely one of many that are possible. Nevertheless, we have depicted the PUA acting in two different ways, as these two ways are likely to be significant.

### 3.4. Proportions of users and our effort in servicing them

The assertions made in Table 3 are opinion only. However, if these are correct, then we need to find a way of engaging with users in the categories that are likely to account for a medium or high proportion of grid use in the future. It is obviously quite difficult to service such users when their current abundance is low and very tempting to over-engage with the current most common type of user.

## 4. The Customer-Service Provider grid relationship

### 4.1. SEUs dominate

It is clear, from our earlier assumptions, that the vast majority of 'users' of the grid, in future, will probably be Service End Users and, individually, these SEUs pose the lowest security threat as their activities are highly controlled by the SP and the service application. If we accept these as basic assumptions, we can see some advantages for the simplicity of a multi-tiered security architecture. As Figure 1 shows, there is trust between the grid and the SP and between the SP and the entity or organisation managing the user's credentials. Furthermore, the SP and the IdP are clear auditable points. We can thus envisage the SP as the true grid user. It is the SP entity that runs jobs on the grid

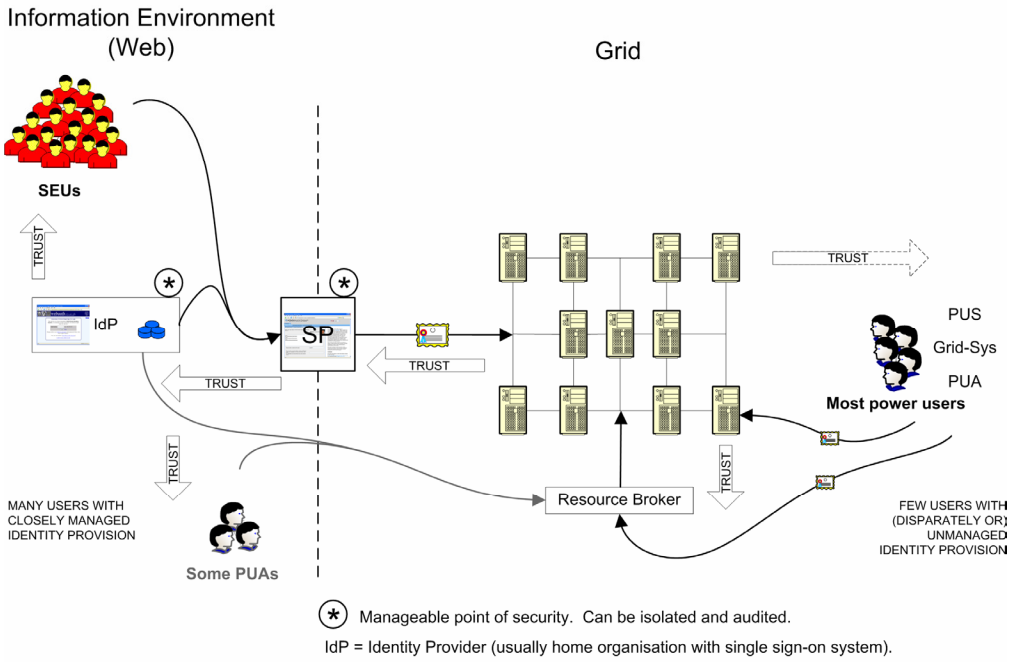


Figure 1 Possible access management behaviour scenario of the different types of grid users. (The PUDS has been omitted as it should contain elements of the PUS and SP).

and, if it were a commercial grid, the owners of the grid nodes could charge the SP for the use of their resource. Thus a clear relationship between the grid and the SP begins to emerge. Where particular authorisation requirements exist – such as "only members of organisation *A* can use this grid at this time" – the grid could mandate that the SP honour that requirement, and the SP could be audited for this. The SP is thus required to take responsibility for authorising users.

Table 3 Likely future proportions of grid users in each category

| Type of user | Proportion of grid users by category | |
|---|---|---|
| | Current | Future |
| SEUD | Low | High |
| SEUX | Low | Medium |
| PUA | Medium | Low/medium |
| PUS | High | Low |
| PUDS | Low | Low |
| SP | Low/medium | Low |
| Grid-Sys | High | Low |

This Customer-Service Provider concept does not need to rely upon any financial requirements. Even in an academic world – but one in which access is restricted to only certain communities – it would be appropriate to run application-based services to high numbers of users in this way.

For ease of use, the vast majority of users will access the power of grids via portals, portlets or similar server-based applications. If we accept that this is true, then we can take the opportunity to tighten up security for all of these users. The portal/server represents a point to which we can – technically or legally – devolve the responsibility for authentication and authorisation. This is a truly synergistic opportunity by:

- improving usability to users who would never benefit from the grid if it meant that they had to perform technical

computing operations to reach that point;

- introducing an 'auditable' point of security to which authentication and authorisation may be securely devolved.

The BRIDGES project built both a data and a compute grid infrastructure accessible by a portal which allowed biomedical researchers to authenticate (using a simple username/password mechanism)[1]. Scientists were then able to upload nucleotide (or protein) sequences and compare them against a variety of local and remote genomic databases. Explorations in rolling out X.509 *user* certificates to the BRIDGES scientists, for identity/authentication purposes, were largely unsuccessful. Instead, solutions utilising X.509 *server* certificates were adopted. Scientists were more comfortable with username/password solutions and to encourage uptake, these requirements were directly addressed. Numerous other challenges were tackled in BRIDGES such as re-engineering of client side tools for simplicity and user friendliness, e.g. to make them "google-like". In short, the scientists wanted a familiar environment in which to work, which shielded them as far as possible from the underlying Grid infrastructure.

Similarly, the Market for Computational Services project (Grid Markets, 2003) asserts that the evolution of the grid is constrained by the fact that users can only use machines where they have accounts. This approach is largely – but not entirely – aimed at Power Users in that the user has to engage at a much more technical level with each grid node. The user experience is greatly simplified in the Grid Markets project by interacting via a central broker. A logical extension to the findings of the BRIDGES and Grid Markets projects means that a lack of usability can mean an absolute lack of take-up, which in turn makes it difficult to survey users regarding their usability comments.

## 4.2. The main threat with the Customer-Service Provider model

The main threat with the Customer-Service Provider model, if implemented efficiently, is likely to be from denial of service (DoS) attacks on SPs. From the grid's or grid node's point of view, the user is the SP. Should any breach in security occur, the normal reaction would be to revoke the SP's privileges, temporarily or permanently. This seems reasonable. However, this means that all users benefiting from the service provided by the SP and the grid will by stymied.

A balance would need to be struck between the risk of this threat and the ability of the SPs to build reasonably safe applications. With such an application as the BLAST technology provided-for by the BRIDGES project, for example, it is difficult to see many threats to the SP other than:

- poor application/API control allowing (for example) SQL insert and update (etc.) statements;
- users submitting jobs incessantly, and thus tying up the databases and the compute cycles;
- submitting a cleverly formulated nucleotide sequence that never resolves and stays busy (as an extreme example).

Clearly, problems will occur, as they do with any multi-user application, but they should be able to be either mitigated-for in advance, or dealt with as they arise.

If a SP provides an application with very poor security then that SP clearly deserves to be suspended until such problems are fixed.

[1] See http://wiki.oucs.ox.ac.uk/esp-grid/NeSC_Shibbolized_Resources

### 4.3. Automated suspension

Rogue, clever, end users (or attackers who are, apparently, end users) will always exist and these need to be quickly identified. A high-level description of the need for automated suspension is discussed in Norman (2006). This could supersede the need for explicit identity assertion 'up front' and may make the C-SP model more secure.

## 5. Shibboleth

Building on the previous sections, we have established how the majority of users of a grid may be 'funnelled' via a server-based application so that requests and jobs may be run on the grid for them. Norman (2006) provides further details as to the case for using Shibboleth with grids and also points out some difficulties. However, the C-SP model would appear to make the use of Shibboleth more attractive.

## 6. Conclusions

The main conclusions of this hypothetical thinking regarding the likely users of future grids are:

- Like the mature web, we predict that most users will require simple, secure, ring-fenced applications to obtain the great benefits of grid technology.
- If such applications are placed in portals (probably using web technology), the security threat profile of this vast majority of users is relatively low (being partly mediated by the application). Thus, heavyweight security solutions will not be needed for the majority of users.
- In such a scenario, Power Users will exist as a small proportion of users. Those users probably merit heavyweight security solutions to be applied to them.
- Where applications are based for the benefit of most users, these provide convenient 'funnels' for such users. Such funnels are suitable for security auditing and therefore are a substantial aid to scalability.
- We have categorised the majority of users as Service End Users (SEUs) who interact directly with Service Providers (SPs). In this Customer-Service Provider model, it is the SPs that interact with the grid directly.
- Grids will be used by many Power Users. We have tentatively named these as PUAs, PUSs, PUDSs, (SPs) and Grid-Sys's. There are more 'actors' in such a system, but we believe that these capture most of the users who interact with the grid directly.
- We described the concept of the SEU-SP interaction as the 'Customer-Service Model'.

## 7. References


Beckles, B. (2004a) Removing digital certificates from the end-user's experience of grid environments. UK eScience All Hands Meeting (2004) http://www.allhands.org.uk/2004/proceedings/papers/250.pdf.

Beckles, B. (2004b) User requirements for UK e-Science grid environments. UK e-Science All Hands Meeting (2004) http://www.allhands.org.uk/2004/proceedings/papers/251.pdf.

Grid Markets (2003). A Market for Computational Services: A Proposal to the e-Science Core Technology Programme. http://www.lesc.ic.ac.uk/markets/Resources/Tag.pdf . Also http://www.sve.man.ac.uk/Research/AtoZ/MCS/RUS/.

Norman, M.D.P. (2006) A case for Shibboleth and grid security: are we paranoid about identity? Proceedings of the 2006 UK e-Science All Hands Meeting

Sinnott, R (2006) Development of Usable Grid Services for the Biomedical Community. Proceedings of *Designing for e-Science: Interrogating new scientific practice for usability, in the lab and beyond*' workshop at the UK National e-Science Centre, January 25-26, 2006.