

1 Introduction

1.1 The Identity Project Background

The Identity Project addressed the current practice and future needs of UK academic institutions in Identity Management ([IdM](#)). The [IdM](#) issues which were investigated included Grid use, Shibboleth installations of varying degree of maturity, collaborative courses and other long term inter-institutional collaborations, internal and shared dynamic virtual organisations, classes of users other than standard staff/student mix, library access schemes, and NHS involvement.

Partners in the project are:

- [Cardiff University](#) (project lead partner)
- [London School of Economics & Political Science](#) (leading the case studies work package)
- [Birkbeck College](#)
- [Goldsmiths College](#)
- [Imperial College London](#)
- [Queen Mary University of London](#)
- [Royal Holloway College](#)
- [School of Oriental & African Studies](#)
- [University College London](#)
- [University of London](#) (associate partner, not funded by the JISC)

Each partner carried out an audit of their [IdM](#) processes, as described in the [Audit final report](#). A large part of this was concerned with issues surrounding membership of an institution. The project also ran a wide-ranging survey, circulated to every HE institution in the UK, which also addressed these issues; this is described in the [Survey final report](#). This report is based on the information gathered through these activities.

The project started on 1 November 2006 and ended on 31 October 2007. The project was funded under the JISC [e-infrastructure programme](#).

Further information about the project generally can be found via the [project web site](#).

This report takes the information gained from the two major investigative work packages, the institutional audits and the survey, to obtain a picture across a range of institutions as to additional IDM practices necessary in an institution with Grid infrastructure. The primary focus of this report is on examining current IDM practice in force at institutions with Grid infrastructure – and what extra burdens are placed upon an institution's IDM to handle the complexity that can exist in an institution with Grid infrastructure.

1.2 Grid

Johnston [1] defines a grid as an environment that provides access and management for the whole range of computing resources needed to solve complex computing and data handling problems. He goes on to say that a grid is a set of services that provide uniform access to a large number of diverse and distributed resources, together with auxiliary services for resource discovery and secure communication based on authenticated, global identity.

In essence, grid computing allows a user at one organisation to access resources at a different organisation. Resources may be specialised services not available at all sites. For example, one site may host a database for the benefit of users at all organisations. Alternatively, the emphasis may be on the coordination of similar resources in order to create a super-service. For example, several compute clusters at different organisations can be harnessed to provide equivalent computing capacity to a vastly more expensive supercomputer.

A grid provides the mechanisms to allow secure remote access to heterogeneous resources in separate administrative domains. The vision of grid computing is that the ability to do this is seamless, that the user does not see the underlying network and security extensions required to make remote resources as simple to access as local resources. By making access to resources simple and scalable grid computing enables a single user to have access to a large number of resources, and a large number of resources provides redundancy and hence high availability of the required capabilities.

Grid identity management typically uses a Public Key Infrastructure (PKI) to provide the “authenticated, global identity”. In a PKI, a user is identified by their possession of a private key known only to them. The private key allows a user to send messages signed with a digital signature which a site can verify as having originated from the user. As part of the verification of such messages, verifiers must have access to the user's certificate. A certificate contains both a globally unique name for the user and a public key, and is signed by a Certificate Authority (CA).

Before signing the public key, the CA takes steps to satisfy itself that the unique name accurately represents the identity of the user. These steps are documented in the CA's Certificate Policy Statement (CPS). Individual sites can refer to the CPS to determine whether they are satisfied with the level of authentication performed by a CA, and hence whether to accept that a certificate signed by the CA represents the user.

2 Identity Management issues

In distributed computing, reliability is a significant problem. If just one of the component resources is unavailable, a distributed computation may fail. Grid computing provides reliability and high availability by providing a user with access to many resources, more than truly needed for a computation. This redundancy allows an unavailable resource to be replaced with an available and equivalent resource elsewhere. To make this work, each user needs access to many resources, some of which they may use rarely or never. Hence there is a high cost-benefit ratio for registering users on each individual system.

The heavyweight procedures typically used for providing access to each resource are not appropriate for grid users. As with any system that is relatively heavyweight compared to the benefits obtained, this tempts users to subvert the mechanisms, by, for example, sharing private keys.

The creation of a global identity for the user is a heavyweight process. The requirement for users to present a physical identity document to the CA makes the process scalable only to a small degree. The UK e-Science CA, with regional representatives at many sites, can handle a few dozen certificate requests a day.

Management of the global identity credentials is very exposed to the user. The user is aware of where the credentials are located, may need to map the credentials from one format to another, and transfers their credentials by physically copying a set of files. The user is responsible for ensuring that the credentials are kept secure, particularly that the private key is encrypted with a secure password and that other users do not have access to the private key file. This is error-prone and can lead to stale credentials existing in some machines.

There is a need to map the user's global identity to a local access management system. Whether this local access management relies on a local identity depends on the technology. If it does, for scalability reasons it generally uses generic local identities, and maps the global identity to a generic local identity as required. Similarly, local activities must be mapped back to the global identifier for accounting purposes.

Typically each user is provided with a single global identity credential. However, an individual may work on several different projects and may have several different roles. For the purposes of access

control and accounting, it may be important for the user to identify under which project or role they are performing a task.

3 Current practice at institutions

From the survey, the most popular grid infrastructure is the Globus Toolkit (in multiple guises).

Globus Toolkit 2 (GT2) was the first distributed computing infrastructure to be associated with the term “grid computing”. It provides secure remote access to three classes of services: 1) initiation and management of program execution; 2) file access and transfer; 3) publication and discovery of resource information. To access these services, each user requires an account on the target resource. Since account usernames and passwords may differ between target resources, the user also has a global identity which is mapped on each target resource to a specific local account. The default mechanism in GT2 for performing this mapping is a simple file with each line containing a global identity and a local username. Once the global identity of a user has been verified, all further actions take place using the respective local user account. This allows the operating system mechanisms to provide access control and accounting.

VDT (Virtual Data Toolkit) is an infrastructure based on GT2. It provides additional services aimed at making installation and use of the grid simpler. A patch to allow pooled accounts was introduced by the European Data Grid (EDG) and subsequently included in VDT. This patch allows a relatively small set of generic user accounts to be created and the mapping from global identity to a user account is deferred until the user actually accesses the target resource. After a period of non-use, the account is reinitialised and returned to the pool of unallocated accounts. The use of pool accounts decreases the administrative burden, since a fixed number of generic accounts can be created instead of adding and deleting accounts as users are added and removed. However, accounting is more difficult as there is no permanent mapping from a local username to a global identity.

Globus Toolkits 3 and 4 (GT3/GT4), while providing the same functionality as GT2, allow additional services to be created using the Web Services model. Using Web Services, it is possible for a resource owner to restrict users to performing tightly defined actions. For example, a Web Service may allow a user to perform a BLAST search, but not to run any other program codes on a resource. Using this restriction, it is possible to remove the requirement for an individual account, since the instantiation of a Web Service can be tracked by the Web Service so the operating system is not needed to map arbitrary operations back to a specific user.

When the UK e-Science program began, and GT2 was deployed at participating sites, external applicants for resource access had to go through each site's account application process, in order to be registered in the site identity management system. In general, this was required in order to satisfy policy requirements or perhaps to reserve a common userid within the organisation. As more users and sites joined the grid, this became unwieldy, both for the users, who had to apply using different processes to many sites, and for the sites, who had to register many users.

The National Grid Service (NGS) is a project to provide a production quality service available to all UK academic users and their collaborators. The NGS core and partner resource providers must accept a single centralised application process to allow users to apply once for access to many resources. (Sites which are unable to do this (including, for example HPCx) can join the NGS as affiliate sites.)

To handle this, the user's global identity is placed in a central LDAP repository. A script runs periodically to read the repository entries and create a mapfile, typically to pool accounts. Note that the LDAP repository is not accessed interactively when a user performs an authenticated action. This is partly due to the scalability and response time of accessing a national-scale LDAP repository, but mainly to avoid making major modifications to the grid infrastructure.

When the National Grid Service was created it was decided to provide a single application process

for users. Details of registered users are placed in an LDAP repository. However, this repository is not interrogated directly by grid resources. The resources periodically download the list of users from the LDAP server and construct a map file from it. The site identity management systems no longer have any information about the external users. Instead the users are given generic accounts, either in a permanently mapped way, or using pool accounts.

The UK e-Science CA requires users to present themselves to a regional (RA) where physical ID (organisation ID card or passport) is matched to their request through the user's knowledge of a 10-digit PIN number. Once satisfied with the identity of the applicant, the RA approves the request by digitally signing the request and forwarding to the CA.

Resources can accept certificates signed by any CA, as long as the CA's Certificate Policy Statement is satisfactory to the resource owner. To provide consistency and simplify the analysis of many CPS documents, a Policy Management Authority (PMA) may group together CA's which satisfy particular criteria. Currently, three PMA's, representing Asia/Pacific, Europe, and the Americas, are grouped under the International Grid Trust Federation.

4 Conclusions

The main identity management challenge for enabling grid computing is to make registration of external users lightweight. To do this, the grid should be adapted to take advantage of the identity checks already in place for organisation identity management systems. The cost of identifying users is high and hence unscalable using the current CA methods. However, organisations put some effort into identifying staff and students at the beginning of their employment or studies.

Combining the grid identity management with the organisation's own checking would provide a scalable method for establishing identity. The recommended method for doing this would be for each organisation to host its own CA and sign certificates for their local users. Running a CA has proved to be a time-consuming operation in its own right. However, by automating the issuance of certificates through the local identity management system, this need not be the case.

Private and public keys can automatically be created for each user when they are entered into the organisation's identity management system. The user can then retrieve a grid proxy certificate through a server (such as a modified MyProxy1? server).

In order to make organisational CA certificates acceptable to other sites, a UK PMA should be created to provide a consortium of organisations that meet a specified standard for issuance of these certificates, and to lobby for the inclusion of their CA certificates into appropriate grids such as the NGS.

References

[1] William E. Johnston, "A different perspective on the question of what is a grid?", Grid Today 1(9), August 2002, <http://www.gridtoday.com/02/0812/100217.html>