Shibboleth-based Access to and Usage of Grid Resources

R. O. Sinnott¹, J. Jiang², J. Watt³, O. Ajayi⁴

National e-Science Centre
University of Glasgow
United Kingdom

¹r.sinnott@nesc.gla.ac.uk

²j.jiang@nesc.gla.ac.uk

³j.watt@nesc.gla.ac.uk

⁴o.ajayi@nesc.gla.ac.uk

Abstract— Security underpins Grids and e-Research. Without a robust, reliable and simple Grid security infrastructure combined with commonly accepted security practices, large portions of the research community and wider industry will not engage. The predominant way in which security is currently addressed in the Grid community is through Public Key Infrastructures (PKI) based upon X.509 certificates to support authentication. Whilst PKIs address user identity issues, authentication does not provide fine grained control over what users are allowed to do on remote resources (*authorization*). In this paper we outline how we have successfully combined Shibboleth and advanced authorization technologies to provide simplified (from the user perspective) but fine grained security for access to and usage of Grid resources. We demonstrate this approach through different security focused e-Science projects being conducted at the National e-Science Centre (NeSC) at the University of Glasgow. We believe that this model will be more widely applicable and encourage the further uptake of e-Science by non-IT specialists in the research communities.

I. INTRODUCTION

One of the critical factors to the success of Grid technologies is ease of use. To encourage wider uptake, the access to large scale computational and data resources such as the National Grid Service (NGS) (www.ngs.ac.uk) needs to be made as simple as possible for the end user. Currently, the end user experience of interacting with such resources typically begins with obtaining an UK e-Science X.509 certificate issued by the UK Certification Authority (CA) at Rutherford Appleton Laboratories (RAL) (www.grid-support.ac.uk/ca). This can often be an arduous process, especially for non-IT experienced researchers, requiring them to follow a detailed recipe for obtaining the certificates and converting them into appropriate formats before they are then able to access the resources. There are also likely to be scalability issues with this approach once the number of certificate holders extends to the wider academic and industrial community. These certificates are used to support Public Key Infrastructures (PKI) where the trusted root of authority is the CA at RAL. Through trusting the process in which certificate requests are processed, subsequently issued and managed by RAL, single sign-on is achieved through recognition of these certificates

by all nodes comprising the NGS. However, the security of PKIs can be compromised in a variety of ways, including the certificate holders not taking appropriate security measures, e.g. writing their private key passwords down in visible places. This issue is further exacerbated since the private key passwords associated with these certificates are necessarily strong. For current and future sporadic users of resources such as the NGS, the likelihood of forgetting complex passwords is increased, and ad-hoc insecure practices such as writing passwords down will occur. A better and more scaleable solution for secure access to large scale Grid infrastructures is thus essential. Shibboleth offers one such possible solution (<http://shibboleth.internet2.edu/>).

The UK academic community is currently in the process of deploying Shibboleth technologies to support local, existing methods of authentication for remote login to resources. Through this model, sites are expected to trust local security infrastructures for example in establishing the identity of users (*authentication*) and their associated privileges (*authorisation*). To support this, the Shibboleth architecture [1] and associated protocols [2] identify several key components that should be supported including IdP, targets and optionally Where Are You From (WAYF) services. Through these components, end users will have single usernames and passwords (which they are more familiar with than PKIs!) which will provide for seamless access to a range of resources at collaborating institutions and service providers. Local security policies at service provider sites can then be used to restrict (authorise) what resources authenticated users are allowed access to.

To support the harmonisation of the Grid and Shibboleth worlds, numerous issues in the current Grid based approach to authentication and Shibboleth need to be addressed. Firstly, the existing approach taken in controlling access to large scale Grid infrastructures via access control lists such as Globus GSI middleware [3] and the use of grid *mapfiles* needs to be either reconciled with Shibboleth components and protocols or alternatively more scalable solutions explored and supported. Secondly, X.509 certificates and the associated Grid security infrastructure should ideally be managed by the system and not left to end users. A third and major hurdle that has to be overcome in utilising Shibboleth technologies in

large scale Grid infrastructures like NGS is trust. Trust underpins the Shibboleth approach to establishing federations. Grid service providers like the NGS need to ensure that local institutions take identity provisioning and associated validation of identity (authentication) processes seriously.

II. BACKGROUND TO SHIBBOLETH AND ADVANCED AUTHORISATION INFRASTRUCTURES

To understand the impact of Shibboleth technologies on Grid security it is first necessary to have an appreciation of the interactions that typically arise with Shibboleth. When a user attempts to access a Shibboleth protected service or Service Provider (SP) more generally, they are typically redirected to a WAYF server that asks the user to pick their home Identity Provider (IdP) from a list of known and trusted sites. The service provider site already has a pre-established trust relationship with each home site, and trusts the home site to authenticate its users properly.

After the user has picked their home site, their browser is redirected to their site's authentication server, e.g. an LDAP repository, and the user is invited to log in. After successful authentication, the home site redirects the user back to the SP and the message carries a digitally signed Security Assertion Markup Language (SAML) [4] authentication assertion message from the home site, asserting that the user has been successfully authenticated (or not!) by a particular means. The actual authentication mechanism used is specific to the IdP.

If the digital signature on the SAML authentication assertion is verified and the user has successfully authenticated themselves at their home site, then the SP has a trusted message providing it with a temporary pseudonym for the user (the handle), the location of the attribute authority at the IdP site and the service provider URL that the user was previously trying to access. The resource site then returns the handle to the IdP's attribute authority in a SAML attribute query message and is returned a signed SAML attribute assertion message. The Shibboleth trust model is that the target site trusts the IdP to manage each user's attributes correctly, in whatever way it wishes. So the returned SAML attribute assertion message, digitally signed by the origin, provides proof to the target that the authenticated user does have these attributes. The attributes in this assertion may then be used to *authorise* the user to access particular areas of the resource site, in principle without the service provider ever being told the user's identity.

How returned attributes are used to make decisions can be done in numerous different ways depending upon the authorisation infrastructure used and the security requirements of the virtual organisation or service provider itself. Numerous authorisation infrastructures exist today [5-9]. The advantages and disadvantages of some of these infrastructures are described in detail in [10-13].

The Shibboleth model of security attributes being required to make an authorisation decision builds on the international X.812 Access Control Framework standard [14] which defines a generic framework upon which numerous authorisation

infrastructures can be supported. In the X.812 model, one of the fundamental scenarios to be supported is where an *initiator* attempts to access a protected target in a remote domain. Two key components support authorised access to the target: a Policy Enforcement Point (PEP) also known as an Access control Enforcement Point (AEF), and a Policy Decision Point (PDP), also known as an Access control Decision Function (ADF). The PEP ensures that all requests to access the target are run through the PDP and the PDP casts the authorisation decision on the request based on a collection of rules (policies).

To make this structure scalable and easily applicable within a Grid environment, a generic API [15] to model the PEP has been proposed and created by the Authorisation Working Group of the Global Grid Forum (GGF) (www.ggf.org). The specification of this API is an enhanced profile of the OASIS SAML v1.1 specification and defines a message exchange between a PEP and PDP consisting of an *AuthorizationDecisionQuery* going from the PEP to the PDP and a returned assertion containing a number of *AuthorizationDecisionStatements*, e.g. boolean decisions stating granted or denied.

Through this API, a generic PEP can be achieved which can be associated with arbitrary Grid services. Thus rather than developers having to explicitly engineer a PEP on a per application basis, the information contained within the deployment descriptor file (.wsdd) when the service is deployed within the container, is used. Authorisation checks on users attempting to invoke "methods" associated with this service are then made using this deployment information, the contents (security policies) of the PDP, e.g. an LDAP repository, together with the DN of the user themselves. Releases of the Globus software since GT3.3 have supported this API.

One authorization infrastructure that has been extended to support this PDP is the Privilege and Role Management Infrastructure Standards Validation (PERMIS) initiative (www.permis.org). PERMIS itself realises a Role Based Access Control (RBAC) authorisation infrastructure offering both a standards-based Java API that allows developers of resource gateways (gatekeepers) to enquire if a particular access to a resource should be allowed, as well as the less tightly coupled SAML AuthZ API. RBAC systems support privilege management infrastructures (PMI).

The relationship between a PMI and authorisation is similar to the relationship between a PKI and authentication. Consequently, there are many similar concepts in the two types of infrastructure. Central to a PMI is the idea of the attribute certificate (AC), which maintains a binding between the user and their privilege attributes. It is similar in notion to the public key certificate in a PKI. The entity that signs a public key certificate is a CA; the entity that signs attribute certificates is called an Attribute Authority (AA). The root of trust of a PKI is often called the root CA, which can delegate this trust to a subordinate CA; the root of trust of a PMI is called the Source of Authority (SoA). The SoA may have subordinate authorities to which it can delegate powers of

authorisation. Certificate Revocation Lists (CRLs), which show a list of certificates that should no longer be accepted as valid, exist in a PKI; Attribute Certificate Revocation Lists (ACRLs) exist in a PMI.

The critical idea in a PMI is that the access rights of a user are not held in an access control list (ACL) but in the privilege attributes of the ACs that are issued to the users. This is the central idea behind RBAC – the privilege attribute will describe one or more of the user’s rights and the target resource will then read a user’s AC to see if they are allowed to perform the action being requested. This de-couples the user’s privileges from their local identity and allows a more dynamic and flexible approach to access control.

Other RBAC based authorisation infrastructures include the Community Authorisation Service (CAS) [5] and the Virtual Organization Membership Service (VOMS) [6]. The central idea behind CAS is that while resource providers can specify a coarse-grained policy, the fine-grained security policy decisions can be delegated to the administrator of the community that is served by CAS. Resource providers grant privileges to the community and establish a trust relationship with the representative of that community. That representative then uses CAS to manage the distribution of privileges within the community. When a user wants to access resources served by CAS, the user issues a request to the CAS server (using their own X509 certificate). If the CAS server decides that the user associated with this certificate has sufficient privileges, then it will issue a proxy credential with an embedded policy giving the user the right to perform the requested actions (assuming that the user has sufficient privilege). The user then uses these CAS credentials to access the resource. The local resource then applies its own local policy to determine the amount of access granted. Currently the primary resource that can be accessed through CAS credentials is gridFTP.

VOMS is a system for managing authorisation data within VOs and has gained widespread acceptance by the HPC-oriented Grid community. VOMS has been developed as part of the European DataGrid project (edg-wp2.web.cern.ch/edg-wp2). The use of VOMS requires a VO administrator to create a separate database for the VO. Each VO user is added to this database and given the appropriate attributes needed to access resources across that VO. This necessarily places a large burden on each VO administrator, since not only must they run their own separate database, they also need to manage it and add all the VO members to it.

Neither CAS nor VOMS is fully aligned with the Shibboleth model of security. VOMS is based upon a centralised server based approach whereas a more scalable, Grid-like model should ideally be based upon a federated model. PERMIS when used for specification and enforcement of local security policies combined with Shibboleth for attribute retrieval offers a model more aligned with the true federated model of the Grid.

The PERMIS RBAC system itself uses XML based policies defining rules, specifying which access control decisions are to be made for given VO resources. These rules include definitions of:

- subjects that can be assigned roles;
- Sources of Authority (SoA), e.g. local managers trusted to assign roles to subjects;
- roles and their hierarchical relationships;
- what roles can be assigned to which subjects by which SoAs;
- target resources and the actions that can be applied to them;
- which roles are allowed to perform which actions on which targets, and the conditions under which access can be granted to roles.

Roles are assigned to subjects by issuing them with X.509 Attribute Certificate(s). A graphical tool called the Privilege Allocator (PA) has been developed to support this process. Once roles are assigned, and policies developed, they are digitally signed by a manager and stored in one or more LDAP repositories.

The PERMIS infrastructure offers very fine grained authorisation capabilities both in terms of policy expression and enforcement. Policy editing tools allow for easy development of the XML based policies. These tools have been developed with HCI considerations included, although we note that the advanced MSc students at the University of Glasgow raised issues with the tools, e.g. the XML that is generated is inconsistent with the tool user interface. For example, the XML has attributes for “subject domain”, whilst the tool has buttons for “where are users from”.

With support for the GGF SAML AuthZ API described previously, PERMIS allows easy linkage between Grid services and authorisation infrastructures.

III. IMPACT AND CHALLENGES OF SHIBBOLETH IN THE CONTEXT OF THE GRID

Shibboleth offers numerous possibilities and potential advantages in the context of the Grid. Single sign-on via authentication at a home site and subsequent acceptance and recognition of the authentication and associated attributes released to remote sites is the most obvious advantage. Thus users need not remember X.509 certificate passwords but require only their own institutional usernames/passwords. Institutions can establish their own trust federations and agree and define their own policies on attribute release, and importantly SPs can decide upon what attributes and attribute values are needed for authorisation decisions.

The uptake and adoption of Shibboleth technologies within a Grid context is not without potential concerns however. Sites need to be sure that collaborating sites have adopted appropriate security policies for authentication. Strength of user passwords and unified institutional account management are needed. Shibboleth is, by its very nature much more static than the true vision of the Grid, where virtual organisations (VOs) can be dynamically established linking disparate computational and data resources at run time. Instead Shibboleth requires agreed sets of attributes that have been negotiated between sites.

Ensuring that an institution in a Shibboleth federation can guarantee the authenticity of a user when accessing a remote resource is crucial to the overall principles upon which Shibboleth and Shibboleth federations are based. In short, institutions in a federation should trust one another. It is the case however, that users at larger institutions may well have numerous usernames and associated passwords that are used to access a variety of services. This is the case at the University of Glasgow for example! To address this a unified institutional user account management system based on nSure technologies (www.novell.com/solutions/nsure) which handles authentication and attributes is being explored in the GLASS project [16]. Through this system, a one to one representation between each user and their corresponding entry in the Human Resource/Registry database – the definitive sources for data will exist. This will support an agreed standard for unique identifiers for each user account with an agreed password policy, e.g. on password strength, and agreement of the definition of department/faculty codes where user accounts should reside.

As well as authentication information, SPs are likely to need further information in order to allow (authorise) access to specific services. In the context of the Grid, membership of the University of Glasgow will not normally be sufficient information for a decision on access to a specific Grid service hosted and managed by a given VO.

The *eduPerson* efforts [17] have identified a core set of attributes that may be of use within an academic environment. The JISC Blueprint for a Production Federation [18] has also explored some potential attributes of relevance to the UK academic community. A small core set of attributes is recommended for IdPs to support that SPs can subsequently use for authorisation decisions. It is essential that interoperability exists between attribute authorities issuing attribute assertions, policy writers defining access policies, and access decision functions that make decisions based on the initiator's attributes and sites target and resource policy. The overlap between Grid technologies (requiring in the first instance attributes for identification) and Shibboleth technologies is required.

The *eduPerson* attributes that have been recognised as providing the necessary core functionality for IdPs and SPs in the UK academic community include:

- *eduPersonScopedAffiliation*: which indicates the user's relationship (e.g., staff, student, etc.) with the institution;
- *eduPersonTargetedID*: is needed when an SP is presented with an anonymous assertion only, as provided by *eduPersonScopedAffiliation*. In this situation it cannot for example provide usage monitoring across sessions. The *eduPersonTargetedID* attribute provides a persistent user pseudonym;
- *eduPersonPrincipalName*: is used where a persistent user identifier, consistent across different services, is needed;
- *eduPersonEntitlement*: enables an institution to assert that a user satisfies an additional set of specific conditions that apply for access to a particular resource. A user may

possess different values of the *eduPersonEntitlement* attribute relevant to different resources.

Each of these attributes can be used to provide the necessary information to SPs to make authorisation decisions. These attributes are versatile and likely to be sufficient for the great majority of applications. Given the fact that Grids form VOs which themselves will have finer grained structuring, it seems sensible that the *eduPersonEntitlement* attribute can be used for this purpose. The *eduPersonEntitlement* attribute can utilise structured XML data representative of large scale Grid infrastructure users and IdPs. This might include the VO they are involved in, the roles that they might have in that VO etc.

It is important to note that these attributes are statically defined and agreed upon between the institutions prior to formulation of VOs or requests to access Grid resources, i.e. they are based upon statically defined PMIs. The JISC DyVOSE project [19] has developed solutions which allow for the dynamic creation and acceptance of attributes. This is more aligned with the dynamic creation of VOs across Grid infrastructures where dynamic delegation of privilege is supported. As the complexity and number of security policies increases, the ability of a given SoA to delegate responsibility to others is necessary. Through extensions to the PERMIS software, DyVOSE now supports dynamic delegation of authority whereby Grid sites can allow an attribute authority controlled by an external SoA to be delegated the ability to assign roles meaningful to a home SoA [20]. Through this, a remote Grid user can hold a role based in their home institution that will allow access to the potentially remote service provider Grid resources.

Perhaps the biggest challenge in moving from static Shibboleth federations with pre-agreed sets of attributes across the federation to more dynamic Grid-like virtual organisations supporting dynamic PMI infrastructures with VO-specific attributes being created and recognised is a semantic one. Remote policies defining rules and regulations in terms of roles, targets and actions on those remote resources requires tool support that can facilitate the discovery, association, merging and promotion or suppression of policies denoting user privileges between sites. In static delegation, the roles at the remote institution would need to be hand written into the policy at the home institution. Dynamic delegation factors away the role assigning powers to subordinate authorities, which may delegate the ability to assign local roles to remote attribute authorities, and vice versa. Thus a Glasgow "Student" role may be assigned to Edinburgh Computing Science users, so they may access the Glasgow resource without the Glasgow SOA knowing about any Edinburgh roles. This trust relationship is agreed beforehand, where it is implicit that the role of Student at Glasgow and Trainee say at Edinburgh are equivalent. Complex delegation allows new intermediate roles with less privilege than their superior role to be defined and assigned to remote attribute authorities. This Delegation Issuing Service to support such dynamic creation and recognition of attribute certificates has been implemented and available for use (www.openpermis.org). This software is currently being used

within the last phase of the DyVOSE project to dynamically link the security infrastructures used for teaching at Glasgow and Edinburgh universities [21,22].

IV. INTEGRATING SHIBBOLETH AND GRID INFRASTRUCTURES

There is much effort to reconcile the Shibboleth and Grid worlds. The GridShib project [23] and the two recently funded JISC projects: ShibGrid [24] and SHEBANGS [25] are exploring use of Shibboleth and Grid. The GridShib project is focusing upon identity federation between the Grid and Shibboleth communities. In real terms the GridShib project is looking towards Grid (GSI) based authentication followed by Shibboleth based retrieval of attributes for making authorisation decisions. It is important to note that the GridShib project does not directly address Shibboleth single sign on to Grid infrastructures. The ShibGrid and SHEBANGS projects are both looking at supporting scenarios where Shibboleth is used for single sign-on and access to the NGS, both with MyProxy [26] at the core. The National e-Science Centre at the University of Glasgow has successfully applied and integrated Shibboleth and Grid technologies in several projects. We provide an overview of some of them here.

We note that in all of the explorations of Shibboleth outlined below we have been part of the UK federation hosted EDINA at the University of Edinburgh (www.sdss.ac.uk).

A. Background to DyVOSE Project

The Dynamic Virtual Organisations for e-Science Education (DyVOSE) project was funded as one of the JISC Core Middleware projects focusing on advanced security infrastructures in the education domain. The basic model being explored in DyVOSE of sites having their own security authorisation policies and associated attributes is very much consistent with Shibboleth where there will likely be several authorities that assert attributes for users. Various domains will then write their own authorization policies based on such attributes.

In teaching the Grid Computing module as part of the advanced MSc in Computing Science at the University of Glasgow a thorough exploration was made of the PERMIS authorisation software for forming static PMIs in a Grid context. In detail, students were initially expected to develop their own security policies for a basic GT3.3 based Grid service which was subsequently used in their main programming assignment. This assignment required that the students were requested to create a policy for a GT3.3 service (*searchSortGridService*) which wrapped a Condor based Java application (this service offered two methods to search (*searchMethod*) and sort (*sortMethod*) a large (5MB) text file (the complete works of Shakespeare). The students themselves were split into groups (*studentteam1*, *studentteam2*) with the authorisation policy to ensure that method *sortMethod* could only be invoked by members of their student group and the lecturing staff, whilst method *searchMethod* could be invoked

by everyone. This set-up was used to illustrate the use of RBAC, where users are allocated privileges based on what role they have been assigned rather than their local user credentials. The students were also requested to secure their service using Globus GSI and also with PERMIS. Performance aspects and benchmarks for the speed of the different systems were recorded by the students and are documented in [10].

The basic Shibboleth scenario currently supported in DyVOSE demonstrates how the Grid based search and sort service can be securely accessed and used via Shibboleth technologies. Specifically it supports scenarios demonstrating how the attributes related to users being members of *studentteam1* (or *studentteam2*) are returned from the IdP at NeSC Glasgow and used to restrict access to the service itself (which has been deployed as a portlet in a GridSphere web portal). In supporting this scenario we have utilised the PERMIS Shibboleth Apache Authorisation Module (SAAM) module [27] which allows use of the PERMIS infrastructure to make authorisation decisions, as opposed to the existing Apache authorisation module (*mod_auth_ldap*). Currently the IdP returns two attributes: the role that the student has (*studentteam1*) and the DN. These attributes are then used and linked through the SAAM module to make authorisation decisions.

To support the Grid aspect of this system we utilised server certificates to overcome the issues in creation of proxy certificates and for submission of jobs via Grid services to the Condor pool at NeSC. Thus client side certificates are not required. Fig 1 shows the GUI and the Shibboleth attributes that have been returned when using the system. Here we retrieve the distinguished name of the user and the attributes indicating which role(s) the person has in this federation. These attributes are dynamically retrieved by the Shibboleth infrastructure and used by the local security infrastructure to allow or deny access to the Grid service portlet.

The screenshot shows a web browser window displaying the GridSphere portal. The page has a header with logos for National e-Science Centre, ESP-GRID, DyVOSE, and Shibboleth/JISC. Below the header, there is a search bar and a section titled 'The Attributes got from Shibboleth AA are:'. This section contains a table of attributes for a user from the University of Glasgow.

Attribute	Value
User from:	UNIVERSITY of GLASGOW
accept-encoding	gzip, deflate
permisRole	studentteam1;studentteam2
connection	keep-alive
Shib-Authentication-Method	urn:oasis:names:tc:SAML:1.0:am:unspecified
Shib-Application-ID	default
cookie	_shibsession_default=d378f4ede98b6af3cf83a832160f402c0
content-length	urn:mace:ac.uk:sdss.ac.uk:provider:identity:labpc-2.nesc.gla.ac.uk
Shib-Origin-Site	en-gb
accept-language	labpc-6.nesc.gla.ac.uk
host	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0; .NET CLR 1.1.4322)
user-agent	cn=richard.sinnott,ou=nesc,o=glasgow
User Distinguished Name	cn=richard.sinnott,ou=nesc,o=glasgow

On the right side of the page, there is a section titled 'ESP-GRID' with text describing the project's goals and a 'Contacts' section with contact information for Dr Mark Norman.

Fig. 1 Shibboleth-Protected GridSphere portal displaying user attributes received from the IdP

From the user perspective, this infrastructure provides a simple model for access to Grid resources. Through a common understanding of the roles needed and basic trust relationships between sites, single sign-on with fine grained authorisation can be achieved.

B. Background to BRIDGES Project

The BRIDGES project (Biomedical Research Informatics Delivered by Grid Enabled Services [28] recently completed at the end of 2005 and involved the universities of Glasgow, Edinburgh with the industrial participation of IBM. BRIDGES was a core project of the UK's e-Science Programme aimed at developing Grid-enabled bioinformatics tools to support biomedical research. The primary source of use cases in BRIDGES was from the Wellcome Trust funded Cardiovascular Functional Genomics Project (CFG) [29] - a large collaborative study into the genetics of hypertension (high blood pressure).

BRIDGES aimed to aid and accelerate such research by applying Grid-based technology. This included data integration tools and support for compute intensive bioinformatics applications such as Basic Local Alignment Search Tool (BLAST). Solutions were developed which provide simplified access to and usage of range of large scale compute resources including all nodes of the UK National Grid Service, the ScotGrid cluster at the University of Glasgow (www.scotgrid.ac.uk), other HPC clusters at Glasgow University and a collection of Condor pools [30].

One of the project requirements was that user authentication should not cause any additional learning or usability overheads for the users. Biology end users range widely in computer literacy and therefore systems providing a single mechanism for users of all abilities should aim at the lowest level of literacy. It was therefore decided to remove digital certificates from the end user environment altogether and replace them with simple username and password authentication at a central project web portal (see Fig. 2). Authentication at Grid sites such as the NGS is instead being carried out by means of a host proxy generated from the Grid server's host credentials. The host's identity is then mapped locally to a project account in the local grid-mapfile on the remote Grid nodes. Thus, all jobs run under the project's identity on the NGS resources, and the logging and monitoring of user activity has to be moved up one level into the domain of the BRIDGES support staff.

We note that whilst we have removed the need for UK e-Science X.509 certificates from the biological end users, we have not omitted security. Rather, we have defined and enforced a much finer grained security model. For example, once a user has logged in to the portal, they have access to the complete set of tools available on the project portal. The finer grain control of what back end resources associated with a tool are accessible for a given user is implemented through the Grid authorisation software PERMIS.

In the original implementation of BRIDGES as depicted in Fig 2, the identity of the user submitting the job was extracted from the portal context, and passed on with the job request. The Grid server sends a lookup request to a dedicated

PERMIS authorisation server maintained by the project team, where secure attribute certificates are used to store information about the roles/privileges a user has.

Three computational security policies (which are enforced) were supported:

- If they are unknown users the job will only be submitted to the local Condor pool (we allow anyone access to the portal, however we restrict what they are allowed to do once there).
- If we recognise the users but they do not have a local ScotGrid account the job will be submitted to the Condor pool and NGS (we currently use all of the NGS nodes and are helping to define the generic datasets and services for the wider life science community on the NGS).
- If we recognise the users and they have an account on ScotGrid then the job will be submitted potentially to the Condor pool, the NGS and to ScotGrid (based on job numbers).

Fig 2. Usage of Server Certificates for Job Submission onto the Grid

The selection of where to submit jobs is based on availability of resources (which is established dynamically). This model of security through portals and server certificates is one way that increased security can be achieved. It does come with certain constraints however on the Grid application developers. We are required for example to keep a track of the users that are submitting jobs (logging of all activity through the portal is recorded and kept). The dangers that might otherwise arise with usage of server certificates for job submission by anonymous end users (from the point of view of the Grid resources the jobs are submitted to), are minimal however. Users that have successfully authenticated themselves at the portal via a username and password are given access to a fixed set of portlets such as the Grid BLAST service. Should a security breach occur and another masquerading user has managed to authenticate at the portal interface, the worst that can occur is that they will be allowed to run many BLAST jobs for example.

This solution is unlikely to be suitable for many Grid researchers who need to compile and tinker with their codes

on the Grid resources. However there are many other researchers (not explicitly Grid-researchers) that require simple, secure access to large scale Grid infrastructures to run known services. Given the number of UK e-Science certificates that have been issued (approx. 3500), it is clear that simpler services tailored to the scientific community with minimal/no Grid learning or overheads are needed to engage with the much larger research communities. For example, there are over 3 million Athens accounts from over 2000 organisations across UK academia. BLAST is one example of such a service. There are likely to be many other such solutions both within the life science as well as other research communities.

The Shibboleth enabled version of the Grid BLAST service did not require users to log in to the project portal. Instead the users were required to log in to their home identity provider and the attributes that were returned were used to enforce subsequent authorisation decisions. We note that it is the case that BRIDGES VO specific attributes could be defined and returned, however provided the Distinguished Name of the user is returned from the identity provider PERMIS is able to make an authorisation decision on the resources that are available to that user.

This model of applying Shibboleth where the user identity is returned and subsequently used to make authorisation decisions, raises issues in the application of Shibboleth. For example, it is typically the case that Shibboleth usage is based on user anonymisation and privacy. For Grid service providers, this model may not be the best approach to encourage wider uptake by the Grid community, e.g. where fine grained/user specific accounting and monitoring for access to and usage of Grid resources is needed.

Fig. 3 Shibboleth enabled large scale BLAST job submission (running 30,000 jobs) across numerous large scale clusters and their monitoring

V. CONCLUSIONS

The access to and usage of Grid infrastructures needs to be made as simple as possible for end users, especially non-IT specialized scientists. Shibboleth provides – from the end user perspective! – a simple way in which these resources can be accessed and used. Through local institutional usernames and passwords access to (authentication) at remote sites within the federation can be supported. Finer grained authorization can be seamlessly provided through the release and acceptance of the necessary pre-agreed security attributes. Combining this with Grid server certificates (which overcome the restriction on users possessing and managing their own user certificates) or approaches where sets of managed certificates are used which are allocated to users when they access the portal (via Shibboleth) provides a model where simple, access and usage is supported. We are currently exploring this latter model and catering for scenarios where the scientists might themselves have their own X.509 certificates, building upon MyProxy based solutions for credential management at the back-end of the portal. Implicit to all of this is usability and making the system security as easy to use as possible for the end user taking heed of previous lessons learned in security technologies [31].

Usability of security infrastructures and usability of Grids more generally is fundamental to the success of e-Science and e-Research efforts. Why should a biologist apply for and take care of their own X.509 certificates when all they really require is to run BLAST on free, national HPC resources.

To support this, a common understanding of the security attributes and their values are needed to be understood by sites involved in a Shibboleth federation. A core set of eduPerson attributes is being explored across UK academia and should be endorsed more widely. This federated model is also based upon trust. Ensuring that sites take all appropriate security measures for authentication and authorization is crucial. Time will tell how this level of trust is upheld (or not!) and the potential ramifications.

It is our intention to explore Shibboleth based access to Grid resources in several other domains where advanced security is essential including the clinical trials domain and electronics domain in several large projects at the National e-Science Centre at the University of Glasgow [32-35].

Acknowledgments

The DyVOSE project was funded by a grant from the Joint Information Systems Committee (JISC) as part of the Core Middleware Technology Programme. The BRIDGES project was funded by a grant from the Department of Trade and Industry. The authors would like to especially thank Prof. David Chadwick and Dr Sassa Otenko at the University of Kent on their inputs on all matters related to the application of PERMIS.

REFERENCES

- [1] Shibboleth Architecture Technical Overview, <http://shibboleth.internet2.edu/docs/draft-mace-shibboleth-tech-overview-latest.pdf>
- [2] Shibboleth Architecture Protocols and Profiles, <http://shibboleth.internet2.edu/docs/draft-mace-shibboleth-arch-protocols-latest.pdf>
- [3] Globus Grid Security Infrastructure (GSI), <http://www.globus.org/toolkit/docs/4.0/security>
- [4] Organization for the Advancement of Structured Information Standards (OASIS), Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V2.0, OASIS Standard, 15 March 2005.
- [5] L. Pearlman, et al., A Community Authorisation Service for Group Collaboration, in Proceedings of the IEEE 3rd International Workshop on Policies for Distributed Systems and Networks, 2002.
- [6] R. Alfieri, et al, Managing Dynamic User Communities in a Grid of Autonomous Resources, CHEP 2003, La Jolla, San Diego, March, 2003.
- [7] D.W. Chadwick, A. Otenko. The PERMIS X.509 Role Based Privilege Management Infrastructure, Future Generation Computer Systems, 936 (2002), Elsevier Science BV, pp. 1–13, Dec. 2002.
- [8] Lepro, R., Cardea: Dynamic Access Control in Distributed Systems, NASA Technical Report NAS-03-020, November 2003
- [9] Johnston, W., Mudumbai, S., Thompson, M. Authorization and Attribute Certificates for Widely Distributed Access Control, IEEE 7th Int. Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises (<http://www-itg.lbl.gov/security/Akenti/>)
- [10] R.O. Sinnott, A.J. Stell, J. Watt, Comparison of Advanced Authorisation Infrastructures for Grid Computing, Proceedings of International Conference on High Performance Computing Systems and Applications, Guelph, Canada, May 2005.
- [11] R.O. Sinnott, A.J. Stell, D.W. Chadwick, O. Otenko, Experiences of Applying Advanced Grid Authorisation Infrastructures, Proceedings of European Grid Conference (EGC), LNCS 3470, pages 265-275, Volume editors: P.M.A. Sloot, A.G. Hoekstra, T. Priol, A. Reinefeld, M. Bubak, Amsterdam, Holland, June 2005.
- [12] D. Chadwick, O. Otenko, A Comparison of the Akenti and PERMIS Authorization Infrastructures, in Ensuring Security in IT Infrastructures, Proceedings of ITI First International Conference on Information and Communications Technology (ICICT 2003) Cairo University, Ed. Mahmoud T El-Hadidi, p5-26, 2003
- [13] A.J. Stell, Grid Security: An Evaluation of Authorisation Infrastructures for Grid Computing, MSc Dissertation, University of Glasgow, 2004.
- [14] ITU-T Recommendation X.812 | ISO/IEC 10181-3:1996, Security Frameworks for open systems: Access control framework.
- [15] Von Welch, Rachana Ananthakrishnan, Frank Siebenlist, David Chadwick, Sam Meder, Laura Pearlman. Use of SAML for OGSi Authorization, Aug 2005.
- [16] Glasgow University Early Adoption of Shibboleth (GLASS) project, www.nesc.ac.uk/hub/projects/glass
- [17] eduPerson Specification, <http://www.educause.edu/eduperson/>
- [18] A. Robiette, T. Morrow, *Blueprint for a JISC Production Federation*, JISC Development Group, Version 1.1: issued 27 May 2005, http://www.jisc.ac.uk/index.cfm?name=middleware_documents
- [19] Dynamic Virtual Organisations for e-Science Education, www.nesc.ac.uk/hub/projects/dyvose
- [20] J. Watt, R.O. Sinnott, A.J. Stell, Dynamic Privilege Management Infrastructures Utilising Secure Attribute Exchange, Proceedings of UK e-Science All Hands Meeting, Nottingham, England, Sept. 2005.
- [21] R.O. Sinnott, J. Watt, O. Ajayi, J. Jiang, J. Koetsier, A Shibboleth-Protected Privilege Management Infrastructure for e-Science Education, 6th IEEE International Symposium on Cluster Computing and the Grid, CCGrid2006, May 2006, Singapore.
- [22] R.O. Sinnott, J. Watt, J. Koetsier, A.J. Stell, DyVOSE Project: Experiences in Applying Privilege Management Infrastructures, UK e-Science All Hands Meeting, Nottingham UK, September 2006.
- [23] GridShib project, <http://grid.ncsa.uiuc.edu/GridShib/>
- [24] ShibGrid, <http://www.nesc.ac.uk/esi/events/622/>
- [25] Shibboleth Enabled Bridge to Access the National Grid Service (SHEBANGS), <http://www.sve.man.ac.uk/Research/AtoZ/SHEBANGS>
- [26] MyProxy Credential Management Service, <http://myproxy.ncsa.uiuc.edu>
- [27] W. Xu, D. Chadwick, A. Otenko, "Development of a Flexible PERMIS Authorisation Module for Shibboleth and Apache Server", 2nd European PKI Workshop, University of Kent, July 2005.
- [28] Biomedical Research Informatics Delivered by Grid Enabled Services (BRIDGES), www.nesc.ac.uk/hub/projects/bridges
- [29] Cardiovascular Functional Genomics project, www.brc.dcs.gla.ac.uk/projects/cfg
- [30] R.O. Sinnott, M. Bayer, Distributed BLAST in a Grid Computing Context, Proceedings of First International Workshop on Distributed Data Mining in Life Science, Konstanz, Germany, September 2005.
- [31] Whitten, A., and Tygar, J. D. Why Johnny can't encrypt: a usability evaluation of PGP 5.0. Proceedings of 9th USENIX security symposium, Washington, 1999.
- [32] R.O. Sinnott, A.J. Stell, O. Ajayi, Development of Grid Frameworks for Clinical Trials and Epidemiological Studies, HealthGrid 2006 conference, Valencia, Spain, June 2006.
- [33] Virtual Organisations for Trials and Epidemiological Studies (VOTES) project funded by Medical Research Council, www.nesc.ac.uk/hu/projects/votes
- [34] Generation Scotland: Scottish Family Health Study exploring Genetics and Healthcare across Scotland, www.nesc.ac.uk/hub/projects/ghi
- [35] Meeting the Design Challenges of nanoCMOS Electronics, EPSRC pilot project to begin October 2006.

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Sinnott, R. O.;Jiang, J.;Watt, O.;Ajayi, O.

Title:

Shibboleth-based access to and usage of grid resources

Date:

2006

Citation:

Sinnott, R. O., Jiang, J., Watt, O., & Ajayi, O. (2006). Shibboleth-based access to and usage of grid resources. In 7th IEEE/ACM International Conference on Grid Computing, Barcelona.

Publication Status:

Published

Persistent Link:

<http://hdl.handle.net/11343/28865>