

# **Final Report for the JISC funded Glasgow Early Adoption of Shibboleth (GLASS) Project**

Date  
14<sup>th</sup> August 2007

## Authors

Prof Richard O. Sinnott <sup>1</sup>	National e-Science Centre	University of Glasgow
Dr John Watt	National e-Science Centre	University of Glasgow
Mr Jipu Jiang	National e-Science Centre	University of Glasgow

## Document History

Dr. John Watt	Draft version 1.0	10 <sup>th</sup> June 2007
Prof. Richard Sinnott	Draft version 1.1	6 <sup>th</sup> August 2007
Mr Jipu Jiang	Draft version 1.2	9 <sup>th</sup> August 2007
Prof. Richard Sinnott	Final version	14 <sup>th</sup> August 2007

---

<sup>1</sup> Contact person.

## Table of Contents

Table of Contents .....	3
Acknowledgements .....	4
Executive Summary .....	5
1. Background .....	6
2. Aims and Objectives.....	7
3. Methodology.....	8
4. Implementation.....	8
5. Outputs and Results.....	9
6. Outcomes.....	11
7. Conclusions.....	11
8. Implications .....	12
9. References .....	12

## **Acknowledgements**

*This project was funded as part of the Joint Information Systems Committee (JISC) Core Middleware Technical Development programme. The project partners at Glasgow would like to thank JISC and the programme managers (James Farnhill, Nicole Harris and Ann Borda) for providing excellent support throughout the course of the project. Thanks are given especially for the extensions to the life time of the project to support the dissemination of project results.*

*Thanks are due to Ian Young of the UK Access Management Federation for his help with registering Shibboleth entities in the UK metadata. Thanks are also given to David Anderson and Peter Mitchell from the University of Glasgow Computing Services for their support in exploitation of the unified account management system rolled out across the university within the context of GLASS. We also acknowledge the inputs from the neuroscientists at the Glasgow Southern General Hospital for their inputs and evaluation of the brain trauma case study explored within GLASS.*

*For the various other case studies used to demonstrate GLASS results, acknowledgements are given to the Department of Trade and Industry for the Biomedical Research Informatics Delivered by Grid Enabled Services (BRIDGES) project; to JISC for the Exploring Shibboleth and Public Key Infrastructures (ESP-Grid) project; to JISC for the Dynamic Virtual Organisations in e-Science Education (DyVOSE) project; to the Medical Research Council (MRC) for the Virtual Organisations for Clinical Trials and Epidemiological Studies (VOTES) project; and to the Biotechnology and Biological Sciences Research Council (BBSRC) for the Grid Enabled Microarray Expression Profile Search (GEMEPS) project.*

## Executive Summary

The overall aim of the GLASS project was to investigate the application of the Shibboleth federated access management software in various applications in current use at the University of Glasgow. The project, in the space of the funded 12 months, has demonstrated that Shibboleth can be applied to a wide variety of applications running in completely independent execution environments, and has provided a solution to the problem of integrating federated access management with production, campus-level account infrastructures.

The core motivation for this project was to explore ways in which a unified user account management system may be utilised by Shibboleth, in particular the Identity Providers, in making authentication assertions about staff and students in the University. In view of the current efforts by UK academia to roll out Shibboleth as an alternative for ATHENS, this work is of utmost importance in ensuring wide uptake of this new technology, but also in assuring systems administrators that this new technology may be seamlessly integrated with their existing systems with minimal, if no, impact on the configuration or schema of their current infrastructure. Consequently, there are gains to be made for the local administrators, who can build a Shibboleth Identity Provider based on their existing systems and not in competition with it, but there are also gains for the federation as a whole as Shibboleth is making security assertions for *real* users, maintained in a production level, highly supported environment from an institution in which they may reliably trust to identify their own users, and to keep this information as up to date as possible. In addition to authentication information, Shibboleth can provide authorisation information about the user for systems to make for example, role-based access control decisions, however this information is often too dynamic for a centralised system to manage. The GLASS project proposed and implemented a cross-campus federated authorisation infrastructure, where the authentication information is asserted by the institution, but attributes for access to systems are maintained by for example individual departments. Through linking of departmental directories utilising a unique identifier assigned by the institution, a user can invoke roles from different departments if access control requires those attributes. This architecture is now currently in use by all NeSC portal-based projects such as GEMEPS, nanoCMOS and VOTES amongst others. This indicates the generality of the solution since these cover domains such as bioinformatics, clinical sciences and electronics.

The use cases for this project fell into two broad categories. The first category related to the integration of Shibboleth with a number of campus resources used by staff and students, in particular the Moodle online virtual learning environment and a custom application for management of student personal records, exam results and course registration called WebSURF. These applications are hosted on completely different execution platforms, and required different plug-ins for session handling, but Shibboleth was shown to provide single sign-on and access based on the Shibboleth asserted identity. Moodle was used to host course material for the students taking the Grid Computing Module of the Advanced MSc in Computing Science at Glasgow. The second use case was the demonstration of Shibboleth authorisation attributes being utilised to gain access to sensitive brain trauma data provided by NHS Scotland under the BrainIT project. The data was accessed through a portal running in GridSphere, with the services available to any particular user being filtered according to the Shibboleth attribute they can present to the service, so a 'nurse' role would see a subset of non-disclosing information, whereas a 'consultant' has a wider set of data fields available to search. This methodology is now standard for portals at NeSC and the use-case has been demonstrated at various workshops and conferences throughout the UK and beyond.

The GLASS project has provided solutions applicable not only to Glasgow, but any institution wishing to adopt Shibboleth as an authentication mechanism that can interface with existing campus services and directories. The project has realised all of the objectives set out in the project proposal, apart from a pair of services which were not available for integration during the project lifetime, namely WebMail and the networked file store. These are, at the time of writing, still not available, but future projects may look into these services in more detail. The results of the project have been disseminated widely through the e-Research community, with current NeSC demonstrations of portal-based technologies utilising the access control models investigated during the project, with the results feeding into several non-Glasgow based projects, including the SEE-GEO Geographical Information Systems project with EDINA.

## 1. Background

It is now widely accepted that requiring users to obtain, manage and present X509 digital certificates for access to Grid resources is a hindrance to wider uptake, due to the considerable technical know-how required to use them correctly [1]. The UK National Grid Service (NGS) [2] demands that its users apply for certificates from the UK e-Science Certification Authority (CA) [3], and that these certificates are converted to the correct format before submitting jobs. This in itself is a challenge for many user communities, however the real risk with inexperienced Public Key Infrastructure (PKI) users comes from complacency with certificate management, as a rogue user could steal the certificate and masquerade as the user the certificate was issued to without the user or the resource being aware until it is too late.

An older, but no less relevant issue with regards to collaboration between institutions is that of resource access. Typically, a collaborator will be set up with an account on a collaborating partner's resource, usually with a different username/password combination than they use at their local institution, and with the access control enforced by the UNIX permissions and file system restrictions. With a large number of collaborators, the scenario arises where a user may have many online identities on each independently administered machine. This model is fast becoming unsupportable in a secure fashion, as multiple passwords tend to be badly managed.

In order to address the security implications of exposing inexperienced users with PKI technologies and multiple online identities, efforts have been made by the UK e-Science community to investigate middleware which allows certificate interactions and the transport of user information to be hidden from the end-user, where information about the user and their privileges can be securely transported between resources on the user's behalf with minimal user interaction. This single sign-on, single identity model is well supported by the Internet2 Shibboleth software [4], which is currently being rolled out across UK academia as the standard federated access management system for educational institutions. Shibboleth allows a user's authentication information at their home institution to be recognised across a large number of trusting sites. The concept of federated authentication is useful as it would be expected that the place that would have the most detailed, most up-to-date and most reliable information about a user's identity would be their home institution. Sites who agree to trust the assertions that an institution makes about the identity of users are said to form a 'trust federation', the configuration of which makes up a large part of the installation of the Shibboleth software. Shibboleth defines several entities which exchange information in the federation, the Identity Provider (IdP) which represents the user's home institution, an optional WAYF (Where Are You From) which allows a user to select their IdP, and a Service Provider (SP) which protects the secure resource.

Shibboleth provides a framework (based on SAML [5]) which allows user information to be transmitted securely between these collaborating sites, yet Shibboleth is not responsible for the user authentication itself. Shibboleth requires external modules or applications to do the authentication, and once successful it can transport this successful authentication assertion to whichever sites require it. Since most institutions already have some form of centralised user account management system in place, it is desirable to have systems like these authenticate users for Shibboleth. Note that Shibboleth is only interested primarily in the binary response of the authentication step, i.e. is this user authenticated or aren't they? It is possible to have completely anonymised use of Shibboleth protected services, where the service provider is never provided with any user-identifying information, only the assertion that this user has authenticated validly at the home institution.

In addition to the federated authentication model which Shibboleth supports, Shibboleth also allows extra information about the user to be transferred to services in the form of SAML attributes. These attributes can range from user identities (common name, surname), to contact info (addresses, telephone numbers), to attributes which are required for systems to make access control decisions (entitlements, privileges, roles). The extra information that home institutions can assert about their users is a matter for institutions and service providers to agree to, hence the attributes that institutions assert tend to be highly heterogeneous. There is also the problem of how to get an institution to agree to release certain user attributes from their central repository, which may involve some changes to their overall schema. To avoid this, the project investigated using composite, departmental attribute authorities as a way of providing access control attributes that would not require production systems to make changes to their directory schemas. In this way, user authentication information is controlled centrally (from the campus account management system), but user authorisation is controlled for example at the departmental level, where the attributes can be more reliably assigned by the relevant departments, where the roles assigned have more meaning.

Once the authentication/authorisation cycle of Shibboleth is successful, the user is returned to the resource they attempted to access. Shibboleth exposes the user information to the web server through internal HTTP headers, which applications may utilise for access control decisions. It is possible for Grid portals to make use of this information to tailor the views that an individual user sees of the portal, or to restrict the set of services that a user can invoke. Grid portal technologies such as GridSphere support access to grid services through web browser technology, keeping complex PKI interactions hidden from the user, and only presenting the services

they may actually invoke through intuitive GUIs. The GLASS project investigated the use of these Shibboleth attributes to enforce access control on several portals in the educational and medical domains which we have Shibboleth-enabled. Another benefit of Shibboleth is that it offers Single Sign-On between federation resources, meaning that the user only needs to input their password once at the beginning of the session, and this login will be automatically picked up by services within the federation.

Shibboleth attributes can allow a resource to enforce its own fine grained authorisation decisions, e.g. based upon Role Based Access Control (RBAC), where access to a resource is constrained not by the identity of a user, but the roles that the user presents to the resource. Currently the Grid adopts an Access Control List (ACL) infrastructure, where the identity of a user (extracted from the subject of their X509 certificate) is mapped to a user or pooled account on the resource – with the permissions of the target account enforcing what the user may or may not do. This works well for a small amount of users, however with a user base numbering in the thousands or tens of thousands, an ACL will very quickly become unmanageable. RBAC is attractive as an access control mechanism as it reinforces the concept that within an organisation there are groups of people who will be requiring the same level of access to a system, for example students at a university, or IT support groups. Instead of having a specific policy for each individual person, like an ACL, the policy can now be expressed in a much shorter way as the mapping of this role to the privileges or abilities on the resource. Now the job of mapping the user identity to the role is done outwith the resource, and as long as the mechanisms for assigning that role are secure, then this can be enough for access. These roles may be as simple as text strings (student, director, sysadmin) or as complex as digitally signed Attribute Certificates (ACs) [6] which allow the attribute itself (as well as its transport mechanism) to be secured well. Further safeguards can be enforced on attributes sent by Shibboleth, for example ‘scoping’ allows attributes to only be released to certain providers, or for attributes to be filtered by name, rejecting attributes that don’t apply to the system. The attributes that Shibboleth sends may also be used by other security technologies such as VOMS [7] and MyProxy [8] to make job submission to the Grid as automatic as possible. Work is ongoing at NeSC in the VPMAN [9] and the EPSRC-funded nanoCMOS [10][11] project investigating protocols which these tools may use to bridge the gap between a user’s home identity and the Grid identity demanded by large scale Grid resources.

The above technologies may also be applied in other realms, in particular, the secure access and integration of databases from a wide variety of scientific disciplines. One area in which security is a primary factor is access to medical data, where strict guidelines are in force over how this data may be accessed and then subsequently processed. Once more, NeSC can draw on the work of its past projects, with the VOTES project [12] providing the core portal functionality that allows us to explore role-based access control of sensitive brain trauma data provided by NHS Scotland amongst numerous other data resources.

## 2. Aims and Objectives

The overall aim of the GLASS project, as described in the project proposal, was to investigate how the integrated directory infrastructure for unified user account management currently being rolled out across the University of Glasgow may be utilised in a Shibboleth environment. The project had two main areas of work, the first being an investigation into the deeper protocols and configuration options of Shibboleth and its dependencies in order to interface with the campus user account system, plus the application of these user attributes in a Grid portal supporting role-based access control to sensitive medical data sets, provided by the BrainIT brain trauma research centre at the Institute of Neurological Sciences at the Southern General Hospital in Glasgow. The second area concentrated on applying Shibboleth access control to various student resources on campus, demonstrating Single Sign-On through Shibboleth to access separate resources, each with its own distinct environment and implementation platform.

In detail (and as outlined in the project proposal) the specific objectives of GLASS were to:

- Integrate Shibboleth with the Novell Nsure [13] account management system already in operation on campus
- Demonstrate Shibboleth-protected access to 4 campus resources (Moodle, WebSURF, NetMail and Filestore)
- Investigate methods of user attribute assignment for role-based access control
- Use Shibboleth to protect a number of Grid portals in a variety of use cases to support role-based access control to sensitive medical data.
- Report on practical experiences in using Shibboleth in established campus environments
- Disseminate configuration details and infrastructure requirements for federated authentication/authorisation.

*These objectives were all met throughout the course of the project and the project results have been widely accepted by the national and international community.*

Some minor changes were necessary to the final objectives, as the NetMail and Filestore systems were not implemented by the University during the lifetime of the project. Some work was done exploring the integration of Shibboleth with the open-source version of NetMail (Hula [14]) however within the time frame of the project it was not possible to directly link this into the main campus email infrastructure.

It should also be noted that during the lifetime of the project, the SDSS Federation migrated to the UK Access Management Federation [15], the impact of which on the project was minimal.

### 3. Methodology

The general methodology used in the GLASS project was that of a cross-campus experiment. This brought with it the advantage that communication with the university partners in the collaboration was possible on a face-to-face basis. This allowed some of the configuration and debugging to be done very quickly. All resources, except for the NSure account management system were hosted and maintained by NeSC Glasgow.

The University allowed a custom connection from our test Shibboleth Identity Provider (IdP) to its backend LDAP database for the NSure system. This connection was read-only and involved no changes to the information stored in the directory. In order for synchronisation with the other databases in the project, the University guaranteed that one specific attribute, the unique identifier (uid), was unique for all users in the system. This attribute may be utilised by Shibboleth to search the extra attribute authorities for user attributes, ensuring that the attributes extracted have been issued to the same person. The uniqueness of the identifier was achieved through direct linkage with the Human Resources department on campus (the definitive source for identification of staff, e.g. those who receive a salary); and with the university registry for student matriculation information.

The University of Glasgow identified 4 key services which would be desirable to implement a Single Sign-On solution. The two major services (and the ones for which Shibboleth access was successfully implemented) were Moodle [16] and WebSURF [17]. Moodle is an online virtual learning environment that deploys easily in Apache allowing course work to be assigned, and discussed online, and for which a simple Shibboleth authentication connector exists. WebSURF is an information management system written by the University of Glasgow which allows students and staff to make changes to their personal records, or to browse exam results. The data for WebSURF was not accessible as they access databases not intended for external use, however the Moodle was fully configurable, and used to host course material for the 2006-2007 session of the Grid Computing Module of the Advanced MSc in Computing Science.

The BrainIT [18] portal was designed to allow a query constructed from several input parameters to be formed into a database query string and submitted to obtain patient results. By enforcing role-based access control on the portal front page using a custom built module, the user is only presented with the parameters that they are allowed to search, meaning that the possible queries that one can construct is tightly constrained by the user's role or function within the organisation. The portal was configured to be accessed through an SSL-encrypted channel (https) which guarantees that no information (address, images) is transmitted unencrypted. This is a key requirement in transmitting medical data.

### 4. Implementation

All Shibboleth-enabled resources were hosted on machines located at the National e-Science Centre. A custom Shibboleth Identity Provider was implemented within the SDSS federation, which was used to assert the authentication of our users. The connection between the IdP and the NSure account system was made using Apache's mod\_authz\_ldap [19] module, which may also be SSL-secured, although we encountered issues with certificates that were unresolvable at the time of writing. The remainder of the machines were built as Shibboleth Service Providers, each with their own distinct environments. The Moodle machine consisted of a standard Shibboleth installation, with a MySQL backend database for the Moodle users. The custom Shibboleth authentication module allows a Shibboleth session to be used as a session login for Moodle, connecting to the appropriate user account if it exists, or creating a new one if the user does not exist. The WebSURF software runs in a JBoss container, which offers no custom Shibboleth connector. However, the SPIE JAAS [20] module allows a JBoss application to utilise a Shibboleth login session and pass identifying information to the application. This module replaces the user as the source of interactions with the application user database, and it uses the user identifier (e.g. name) as a username and creates a secret password. The WebSURF code was provided with the advanced database functionality removed, as these connections required access to non-public databases which lay outside the scope of the project. To accommodate this, we decided to duplicate the specifications of the production WebSURF resource, so the Service Provider platform and environment could mirror the production WebSURF site exactly. The WebSURF code was rewritten, compiled and launched in the container and the SP configured to release the 'fullName' attribute, which is utilised by SPIE JAAS as the user identifier when creating the user session. Both services were implemented with a choice of login types, the normal custom login or through Shibboleth. In the case of the SPIE JAAS module, care must be taken when

presenting the Shibboleth login as a user option. This is because the secret password that the module writes into the application database is never revealed to the user, as this is an interaction within the container. If a user chooses to login manually having previously logged in with Shibboleth, they will find this password is now the login required allowing access and they will not be able to find it out. It would be recommended that it is beneficial to make a permanent choice of one login method or the other, as it is inevitable that problems would occur with a large number of users asserting their preferences inconsistently.

For the use case involving more complicated attribute types, it was shown that the difficulty in altering the schema of the primary identity database of the NSure account system to assert eduPerson [21] style attributes can be removed by allowing the IdP to utilise multiple Attribute Authorities. Two external LDAP servers hosted in different departments were configured to release the eduPersonEntitlement attribute, some users appear in both directories, whilst some appear in only one. The only piece of information in the external servers which is duplicated from the central NSure directory was the guaranteed unique 'uid' attribute. This attribute is used by the Shibboleth IdP to search the directories for the right user. A Service Provider will be provided with the union of these directories' eduPersonEntitlement attributes, enabling an SP to make decisions about user access based on attributes from several departments. The configuration for distributed attribute authorities for IdPs is done entirely on Shibboleth and nothing is done on the Novell system, this being a great motivation for institutions to adopt this approach for asserting user entitlements across campus as it involves no changes to their central directory.

The use-case of accessing secured medical data enforced by role-based access control was implemented on a JSR-168 compliant Grid portlet running in GridSphere. The custom login module for GridSphere was deactivated, and the HTTP headers presented to Tomcat were scanned for a set of attributes. The IdP had been configured to release attributes such as 'common name', 'organization', 'organizationalUnit' and 'jpegPhoto', so we displayed all of these in the user's information window. Note that through the use of attribute scoping, we only released these highly identifying attributes to our own Service Providers. The Grid portal itself shows a choice of input parameters that the user may alter to perform on search on the brain trauma database, such as patient sex, age, medication etc. The filtered view is enforced by a custom authorisation function which restricts the input parameters based on the role contained in the eduPersonEntitlement attribute. These roles were defined as 'nurse', 'consultant' and 'investigator', with the roles expressing a hierarchy of privileges on the system. The 'nurse' role may only edit a small number of non-identifying parameters and the 'consultant' role may edit more sensitive parameters, while also inheriting the parameters of the 'nurse' role. This form of hierarchical privilege management often allows security policies to be expressed more simply, reflecting the real-life organisational infrastructure already in place. The portal creates an SQL query based on the selected input parameters, and this is fed through Globus Toolkit 4 and OGSA-DAI to the backend database which collects the patient details from auxiliary databases hosting subsets of the total patient data. Securing of the auxiliary databases was outside the scope of this project, but forms part of the work for the JISC SEE-GEO [22] project for securing geographical databases.

No serious implementation issues were encountered during the project, and the portal and single sign-on to student resources has been demonstrated at conferences and workshops in the UK and internationally.

## 5. Outputs and Results

The project has produced many different kinds of output. These have included:

- A working medical data access portal which enforces strict role-based access control on its users via Shibboleth attributes
- A campus-wide infrastructure for both authenticating and authorising user access to local systems
- A legacy method of communicating with a production level user account management system with minimal reconfiguration.
- Established liaison with campus IT services for consultation when full Shibboleth adoption is implemented across campus
- Many live demonstrations at national and international conferences and several published papers

The actual publications generated directly or indirectly from the project, e.g. from exploitation of Shibboleth single sign-on using the integrated Glasgow IdP include:

R.O. Sinnott, O. Ajayi, A.J. Stell. *Supporting Grid Based Clinical Trials in Scotland* to appear in Health Informatics Journal, November 2007.

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.

- R.O. Sinnott, J. Watt, O. Ajayi, J. Jiang, *Shibboleth-based Access to and Usage of Grid Resources*, IEEE International Conference on Grid Computing, Barcelona, Spain, September 2006 (18% acceptance rate).
- R.O. Sinnott, O. Ajayi, A.J. Stell, J. Watt, J. Jiang, *Single-Sign on and Authorization for Dynamic Virtual Organizations*, International Conference on Virtual Enterprises, (PRO-VE'06), Helsinki, June 2006.
- R.O. Sinnott, J. Watt, D.W. Chadwick, J. Koetsier, O. Otenko, T.A. Nguyen, *Supporting Decentralized, Security focused Dynamic Virtual Organizations across the Grid*, 2<sup>nd</sup> IEEE International Conference on e-Science and Grid Computing, Amsterdam, December 2006.
- R.O. Sinnott, O. Ajayi, A.J. Stell, *Secure, Reliable and Dynamic Access to Distributed Clinical Data*, Life Science Grid Conference, Yokohama, Japan, October 2006.
- R.O. Sinnott, O. Ajayi, A.J. Stell, J. Watt, J. Jiang, *User Oriented Access to Secure Biomedical Resources through the Grid*, Life Science Grid Conference, Yokohama, Japan, October 2006.
- R.O. Sinnott, O. Ajayi, A.J. Stell, *Security Oriented e-Infrastructures Supporting Neurological Research and Clinical Trials*, 2<sup>nd</sup> International Conference on Availability, Reliability and Security, (ARES'07), Vienna, Austria, April, 2007.
- R.O. Sinnott, O. Ajayi, A.J. Stell, C. Bayliss, J. Watt, J. Jiang, *Supporting Life Science Research through Shibboleth and Community Grid Portals*, HealthGrid 2007 conference, Geneva, April 2007.
- R.O. Sinnott, J. Watt, J. Jiang, *The GLASS Project: Supporting Secure Shibboleth-based Single Sign-On to Campus Resources*, to appear in UK e-Science All Hands Meeting, September 2007, Nottingham.
- R.O. Sinnott, O. Ajayi, J. Jiang, A. J. Stell, J. Watt, *User-oriented Security Supporting Interdisciplinary Life Science Research across the Grid*, International Journal of New Generation Computing, Special Edition on Life Science Grids, editors A. Konagaya, P. Arzberger, T. W. Tan, R. Sinnott, D. Angulo, March 2007.
- R.O. Sinnott, O. Ajayi, A.J. Stell, *Grid-enabled Infrastructure to Support Nationwide Clinical Trials and Studies*, in progress for 4th International Conference on Life Science Grids, Seattle, USA, October 2007.
- R.O. Sinnott, J. Watt, J. Jiang, G. Stewart, A. Stell, D. Martin, T. Doherty, *Federated Authentication and Authorisation for e-Science*, submitted to APAC, Perth, Australia 2007.
- R.O. Sinnott, J. Jiang, J. Watt. *Trust but Verify: Supporting Usable, Security-driven Virtual Organisations*, submitted to e-Science 2007, Bangalore India, December 2007.
- R.O. Sinnott, L. Han, G. Stewart, D. Berry, *Towards a Grid-enabled Simulation Framework for nanoCMOS Electronics*, submitted to e-Science 2007 conference, Bangalore India, December 2007.
- R.O. Sinnott, O. Ajayi, A.J. Stell, *Towards an Optimal Approach for Decentralised Access Control in the e-Health Domain*, submitted to e-Science 2007 conference, Bangalore India, December 2007.
- R.O. Sinnott, J. Jiang, J. Watt. *Secure User and Admin Friendly Virtual Organisations: The Model for Future Grid-based Collaborations*, in progress for 26th International Conference on Parallel and Distributed Computing and Networks, Innsbruck, Austria, February 2008.
- R.O. Sinnott, G. Stewart, C. Millar, *Supporting EEE-Science*, in progress for 26th International Conference on Parallel and Distributed Computing and Networks, Innsbruck, Austria, February 2008.

In addition, the project has presented and demonstrated its results of integrated Shibboleth across a range of scientific disciplines at a range of fora including:

- R.O. Sinnott, *Update on Shibboleth and Grid work at the National e-Science Centre Glasgow*, Open Grid Forum, Washington, US, September 2006.
- R.O. Sinnott, *Grid Challenges and Opportunities*, seminar University of Stirling, October 2006.
- R.O. Sinnott, *Dynamic Virtual Organisations in the Education Domain*, presentation given at e-Science Institute Edinburgh workshop on Virtual Organisations, November 2006.
- R.O. Sinnott, *Single Sign-on for e-Life Science Research*, presentation/demonstration selected to represent UK e-Science, Supercomputing 2006 conference, Tampa Bay, Florida, November 2006.
- R.O. Sinnott, *An e-Voyage showing Shibboleth Single Sign-on in Action*, UCISA Conference, Aston Business School, Birmingham, November 2006.
- R.O. Sinnott, *Life Science Grids*, seminar University of Aberdeen, November 2006.
- R.O. Sinnott, *Usable Grid Security*, seminar Royal Holloway (invited talk), December 2006.
- R.O. Sinnott, *Security-oriented access to microarray data resources*, BBSRC e-Science and Proteomics Workshop, Carden Park, Cheshire, January 2007.
- R.O. Sinnott, *Grid Challenges and Opportunities*, seminar University of Aberdeen, January 2007.
- R.O. Sinnott, *Life Science Grids*, (invited talk) talk given to UK/Swedish Ambassadors at UK-Sweden e-Research event, Stockholm, Sweden February 2007.

R.O. Sinnott, *Grid based e-Health*, talk given to Chief Scientist Office/Scottish Executive, February 2007.

R.O. Sinnott, *Supporting Life Science Research through Shibboleth and Community Grid Portals*, demonstration and presentation given at HealthGrid 2007 conference, Geneva, Switzerland, April 2007.

R.O. Sinnott, *Experiences Developing e-Infrastructures across Biomedical Repositories at NeSC*, presentation given at EU workshop on Towards a European e-Infrastructure for e-Science Digital Repositories, Brussels, Belgium, March 2007.

R.O. Sinnott, *The UK e-Science Environment*, presentation given at UK-Malaysian e-Research workshop organised by UK High Commission, Kuala Lumpur, Malaysia, June 2007.

## 6. Outcomes

The outcomes of the project include the outputs mentioned in the previous section, but the main focus of GLASS lay in the overall knowledge gained in the operation of Shibboleth in a campus environment. Based on the original project proposal, all the outcomes of the three phases of the project have been met on time. The use-case for integration of the NSure system was defined early on, with the necessary configuration being implemented on our test resources. The outcome of this phase directly fed into the next phase, which involved setting up the distributed attribute authority infrastructure, and designing the environments necessary to replicate the University's campus services for the single sign-on demonstration. Finally, this experience was utilised in securing the medical data portal, which involved the removal of the bespoke login portlet and the use of a custom access control function.

Along with design and implementation issues, other aspects of identity management were invoked in order to keep the data used in the GLASS project secure. One of these being the encryption of stored passwords on the LDAP server, which meant that even administrators could not find out the passwords of their users. This method is now being successfully used in the nanoCMOS project.

The main outcome of this project however, lies in the need for the setting up and configuration of these systems to be better understood, and made clearer for administrators. Currently the knowledge required to setup these systems is vast, with Shibboleth alone requiring the accurate configuration of at least 10 raw XML files in order to express policies. The need for simplification of this has fed directly into the OMII-SP project at NeSC, which aims to produce a set of portlets which allow configuration of (amongst other things) Shibboleth attribute acceptance policies for Service Providers. The design of these portlets has arisen directly out of the requirements of the GLASS project.

## 7. Conclusions

The GLASS project has successfully demonstrated that Shibboleth can be integrated with a variety of applications, both custom and proprietary, across a heterogeneous set of resources. The single sign-on capability of Shibboleth means a user may browse between these resources seamlessly by only entering their password once at the beginning of the session, thus minimising the amount of times this information need be transmitted across networks. The end user may invoke a range of identifying information held in attribute stores across campus in order to access specific services, some requiring contextually accurate text strings, others, like grid Services, requiring access through X.509 certificates. All of which is done transparently to the user, meaning the user cannot weaken the security model of PKIs through poor management or lack of technical knowledge. The implementation has also been shown to have little impact on the running of dependent servers, which is a crucial element in persuading institutions to adopt the Shibboleth technology. Its versatility in protecting highly heterogeneous resources has also been demonstrated.

The results of the GLASS project have been demonstrated in various peer reviewed papers and numerous live demonstrations at JISC and at numerous workshops across the country and internationally. The model adopted for distributed attribute authorities is now standard for NeSC based projects, representing a sum of over £20M of project funding, and we have been disseminating information about how other institutions may adopt this scheme. The results of the GLASS project are also directly informing on-going projects such as the JISC funded VPman project, where integration of VOMS, PERMIS and Shibboleth is on-going. We note that we have already integrated VOMS and PERMIS to successfully protect GT4 based services. Work is on-going to provide Shibboleth based access to other services and resources including OMII-UK services and resources such as the NGS. We have also shown how vanilla VOMS can be used with major cluster resources such as ScotGrid. The technology transfer of these results to the NGS is on-going.

The GLASS project is by all standards a great success and has steered the use of Shibboleth at Glasgow in the past year and will continue to do so for the foreseeable future. This is shown through the recent funding success of NeSC Glasgow in two major EU FW7 proposals exploiting Shibboleth and security technologies in the clinical domain worth over 5M€. The first of these AVERT-IT (Advanced Arterial Hypotension Adverse Event prediction through a Novel Bayesian Neural Network) directly extends the work undertaken in the

BrainIT demonstrator project explored within GLASS. The second of these EuroDSD (Investigation of the molecular pathogenesis and pathophysiology of Disorders of Sex Development) builds upon the security driven approach to Grid research needed in the e-Health domain. Both of these projects are looking at exploitation of Shibboleth and fine grained authorisation technologies.

## 8. Implications

Not surprisingly, the implications of the GLASS project are most strongly felt at the University of Glasgow, where the results of the integration experiments may feed directly into the IT services' master plan for student/staff service access for the next 5-10 years, particularly in light of the current efforts in academia to roll out Shibboleth across education institutions as the de facto authentication mechanism. The ease at which Shibboleth may be integrated with current campus level user accounting systems mean that adoption of this technology may be implemented in a matter of days, providing the University can guarantee the uniqueness of the 'uid' attribute for every user across campus, as this is the 'glue' that holds the distributed attribute authority model together. The model proposed by GLASS has already been adopted by the VOTES, VPMAN and nanoCMOS projects, with the latter providing a motivating use-case for other institutions to adopt this approach, particularly those institutions running production, user-accountable Identity Providers, which will normally be resilient to any requests for configuration changes. The model proposed here allows complete departmental flexibility whilst requiring minimal schema manipulation.

NeSC Glasgow is adopting this approach for all its portal based projects, with Shibboleth taking the form of a portal view filter, allowing only certain user roles to invoke certain services. With the abundance of methods for integrating Shibboleth with Grid-style authentication mechanisms based on X.509 (SHEBANGS, ShibGrid, GridShib etc), the final role of Shibboleth in the Grid context remains open. However, this model allows Shibboleth to be used as a reliable, robust first level of security, meaning that any services invoked by the portal will already have had to pass initial authentication and basic attribute-based authorisation before ever being invoked. The problem of marrying a local and 'national' level credential was beyond the scope of this project, however, this is being explored in the VPMAN and nanoCMOS projects, and the model exploited above is complimentary to any 'tier-2' bridging middleware that can be adopted.

The association with NHS and their willingness to provide data and information regarding the implementation of their live databases for searching is a major political win, as the NHS is well-known for the strength of its security requirements. It says a lot for the technology that this group is willing to participate, and will be an important reference point for other disciplines who cite security concerns as major obstacles to collaboration and information sharing.

The success of this project and its output has hardened the notion that wide uptake and popularity of e-Science and the Grid as a whole will be levered by the ease at which these resources may be accessed by non-technical users whose only specialisation is the field of science they currently reside in. Portal technology is fast becoming a key API for any Grid resource, so security models like those investigated by the GLASS project will be an attractive feature for any service providers wishing to hide any complex PKI functionality from non-savvy end users. The output from GLASS will continue to evolve outside the timescale of the project, as it has been so heavily adopted by the various projects at NeSC, and further integration with Grid security middleware will feed into the academic community, providing some compelling use-cases that prove that an all-encompassing solution to federated Grid access is no longer purely theoretical.

## 9. References

- [1] B. Beckles, 'A user-friendly approach to computational grid security', Proceedings of the UK e-Science All Hands Meeting, Nottingham, UK, September 2006.
- [2] The National Grid Service (NGS) <http://www.grid-support.ac.uk>
- [3] J. Jensen, 'The UK e-Science Certification Authority' Proc. Of the 2<sup>nd</sup> UK e-Science All Hands Meeting (EPSRC ISBN 1-904425-11-9), Sep (2003) pp. 336-369
- [4] S. Cantor et al. 'Shibboleth Architecture: Protocols and Profiles' Internet2-MACE (Document ID: internet2-mace-shibboleth-arch-protocols-200509) 10 Sep (2005) <http://shibboleth.internet2.edu/docs/internet2-mace-shibboleth-arch-protocols-latest.pdf>
- [5] .E. Maler et al. 'Assertions and Protocols for the OASIS Security Assertion Markup Language' OASIS (Document ID: oasis-sstc-saml-core-1.1) Sep (2005) <http://www.oasis-open.org/committees/security/>
- [6] D.W.Chadwick, A. Otenko, E. Ball, 'Role-Based Access Control with X.509 Attribute Certificates' IEEE Internet Computing, Mar-Apr (2003) pp. 62-69.
- [7] L. dell'Agnello, R. Alfieri et al. From gridmap-file to VOMS: managing authorization in a Grid environment' Future Generation Computer Systems 21 (Elsevier Science BV), (2005) pp. 549-558

- [8] J. Novotny, S. Tuecke, V. Welch. 'An Online Credential Repository for the Grid: MyProxy' Proc. Of the Tenth International Symposium on High Performance Distributed Computing (HPDC-10), IEEE Computer Society Press, Aug (2001)
- [9] Integrating VOMS and PERMIS for Superior Secure Grid Management (VPMAN) <http://sec.cs.kent.ac.uk/vpman>
- [10] Meeting the design challenges of nanoCMOS electronics (nanoCMOS) <http://www.nanocmos.ac.uk>
- [11] R.O.Sinnott, A.Asenov et al. 'Meeting the Design Challenges of nanoCMOS Electronics: An Introduction to an EPSRC Pilot Project' Proc. Of the Fifth UK e-Science All Hands Meeting (NeSC ISBN 0-9553988-0-0), Sep (2006)
- [12] Virtual Organisations for Trials and Epidemiological Studies (VOTES) <http://www.nesc.ac.uk/hub/projects/votes/>
- [13] NSure, Novell Identity Manager <http://www.novell.com/products/identitymanager/>
- [14] The Hula Project, Novell <http://www.hula-project.org>
- [15] The UK Access Management Federation for Education and Research <http://www.ukfederation.org.uk>
- [16] Moodle – A Free Open Source Course Management System for Online Learning <http://moodle.org>
- [17] WebSURF Student Records Update Service, University of Glasgow (Robert Stewart, Management Information Services) <http://www.websurf.gla.ac.uk>
- [18] BrainIT <http://www.brainit.org>
- [19] mod\_authz\_ldap: X509 Certificates and LDAP <http://authzldap.othello.ch/>
- [20] SPIE JAAS Module Architecture, SPIE Project, University of Oxford, <http://spie.oucs.ox.ac.uk/Wiki.jsp?page=JAASmoduleArch>
- [21] The eduPerson Specification <http://www.educasue.edu/eduperson>
- [22] SEE-GEO project [http://www.jisc.ac.uk/whatwedo/programmes/eresearch\\_grid\\_ogc\\_collision/project\\_see\\_geo.aspx](http://www.jisc.ac.uk/whatwedo/programmes/eresearch_grid_ogc_collision/project_see_geo.aspx)