

Towards Shibboleth-based Security in the e-Infrastructure for Social Sciences

Wei Jie, Michael Daw, Rob Procter, Alex Voss

National Centre for e-Social Science, University of Manchester, UK

wei.jie@manchester.ac.uk

Abstract. The e-Infrastructure for e-Social Sciences project leverages Grid computing technology to provide an integrated platform which enables social science researchers to securely access a variety of e-Science resources. Security underpins the e-Infrastructure and a security framework with authentication and authorization functionality is a core component of the e-Infrastructure for social sciences. To build the security framework, we adopt Shibboleth as the basic authentication and authorization infrastructure and further combine PERMIS advanced authorization technology. As a result, this security framework integrates the advantages of both Shibboleth cross-domain identity federation and PERMIS policy driven role based authorization control, thus presenting a promising security model for secure access to the e-Infrastructure for the social sciences.

Introduction

The UK Economic and Social Research Council (ESRC) has funded a three year project to create an e-Infrastructure for the e-Social Sciences (E-Infrastructure, 2007). This project leverages Grid technology (Foster, 1999) to build an e-Infrastructure on the UK National Grid Service (NGS) and provide integrated access to a variety of resources for social science research, including datasets, tools, services and easy-to-use user environments.

Security underpins the e-Infrastructure for the social sciences as it is one of the foundations of any working Grid infrastructure (e.g., Foster, 1998). Without 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. As an important activity of the e-Infrastructure for Social Sciences project, we will build a security framework enabling secure access to the e-social science resources.

The e-Infrastructure allows resources to be shared between members of virtual organizations (VOs). However, if an organization is to allow its resources to be shared amongst its VO members, it needs to be able to determine who is authorized to access these resources in which ways, and who is not. Access control rests on the two related concepts of authentication (establishing the identity of a person) and authorization (regulating what an identified person is allowed to do, cf. Broadfoot, 2003), which are two fundamental but critical issues in constructing the security framework for the e-Infrastructure.

Authentication will take place before users gain access to the e-Infrastructure. That is, the e-Infrastructure needs to check the identity of users to ensure that they are allowed access to the resources. A challenge that comes with user authentication is that some domain sites may

wish to keep their own authentication systems, so the security framework of the e-Infrastructure needs to reconcile with existing authentication systems and seamlessly make them interoperable.

Authorization requires the e-Infrastructure security framework to implement access control mechanisms which reflect resource owners' privileges, and protect resources by only allowing authorized accesses. Besides, the security infrastructure should provide resource owners flexibility to define their access control policy (be it coarse-grain or fine grain), as well as modify or extend their current policy with ease.

Our security framework leverages the widely accepted Shibboleth (Shibboleth, 2007) technology for the authentication infrastructure and combines this with PERMIS (PrivilEdge and Role Management Infrastructure Standard) (Chadwick, 2003) as authorization technology. In this way, it offers scalable and flexible Grid VO-wide authentication as well as policy-driven, role-based, multi-grained authorization for access to and usage of the e-Infrastructure. In this paper, we focus on the philosophy and principles for constructing the security framework. Firstly, we introduce the background of Shibboleth and PERMIS as well as the motivations for adopting them as the authentication and authorization infrastructure. Then we discuss the issues and challenges of using Shibboleth and PERMIS in the security framework. We conclude the paper with an outline of the direction for future work.

Adopting Shibboleth and PERMIS as the authentication and authorization infrastructure

The UK academic community is currently in the process of strategically deploying Shibboleth technologies as the next generation authentication and authorization infrastructure for accessing web-based resources. Shibboleth is an authentication and authorization architecture (Scavo, 2005) to support inter-institutional sharing of resources (or services) that are subject to access control.

Shibboleth does not carry out authentication itself. Instead, it implements a set of SAML protocols (Saml, 2007) for the secure passing of identity information between institutions and resource providers. It relies on the institution to establish identity (i.e., authentication), and on the resource provider to control access rights (i.e., authorization). In other words, Shibboleth separates the user authentication that is performed by users' respective home institution, and the authorization that is performed by the resource providers to be accessed based on users' attributes that have been passed to it.

The main features and advantages offered by Shibboleth and its implications are summarized in Table I. From the table, we can see that Shibboleth offers several advantages for the security of the e-Infrastructure for social sciences. First, end users will have single usernames and passwords at their own institution which will provide for seamless access to a range of resources at collaborating institutions and services providers. That is, single sign-on via authentication at a home site and subsequent acceptance and recognition of the authentication to remote sites. Second, Shibboleth separates the user authentication that is performed by users' respective home domains from the authorization that is performed by the resource providers to be accessed based on users' attributes. Institutions can thus establish their own trust federations and agree and define their own policies on attribute release, and resource providers can decide upon what attributes and attribute values are needed to make authorization decisions for their resources.

Issues	Features	Advantages	Implications
User Authentication	<ul style="list-style-type: none"> ▪ Shibboleth devolves all responsibility for user authentication to the users' respective institutions, i.e. user authentication only happens at the users' home institutions. 	<ul style="list-style-type: none"> ▪ Remove the need for authenticating users at every instance of remote service provision ▪ Realize SSO authentication across multiple service providers ▪ Make user authentication scalable with the growing size of the virtual organization ▪ Involves less integration work to bring in a new institution or user 	<ul style="list-style-type: none"> ▪ A trust model between the authentication institution and the remote service providers. ▪ User home institution should have a robust and up-to-date authentication system in place. ▪ The need for trust requires the formation of federations which define similar groups who have agreed to a common set of policies.
User Privacy	<ul style="list-style-type: none"> ▪ Users' home institutions only pass information about users' status (attributes) to the service provider rather than their personal identity. ▪ Users' home institutions control what attributes get released to the remote service providers. 	<ul style="list-style-type: none"> ▪ Users' personal identity information is protected ▪ Release the minimum amount of information about users required by the remote service provider for authorization and to balance the needs of the service provider to make valid authorization decisions with the desire to preserve as much of the users' privacy as possible. 	<ul style="list-style-type: none"> ▪ A federation needs to define the users' attributes required to be released. Each institution must adhere to the policy. ▪ Users' home institutions and service provider must negotiate the actual attributes about the user required to make valid authorization decisions.
Access Authorization	<ul style="list-style-type: none"> ▪ Service providers have the freedom to define their own authorization policy and decide upon what attributes are needed for authorization decisions to access services. 	<ul style="list-style-type: none"> ▪ Make authorization management more flexible over the traditional approach of using users and group identifiers. ▪ Enable service providers to implement multi-grained access control and effectively control access based on users' attributes 	<ul style="list-style-type: none"> ▪ Service providers need to adopt an appropriate authorization model which makes authorization decisions based on users' attributes

Table I. Shibboleth Technology Features, Advantages and Implications

Whilst Shibboleth provides a framework for authorization based on attributes, some advanced authorization technologies, in particular, PERMIS can be used to implement the authorization. PERMIS is a project that has built an advanced authorization infrastructure based on attribute-based access control (ABAC). ABAC is a superset of role-based access controls (RBAC) in which access control decisions are made based upon attributes held by the user, and not just upon their organizational roles (as in conventional RBAC). In PERMIS, user attributes are held in X.509 standard attribute certificates (ACs) (Farrell, 2002). An AC is a data structure that binds details about the holder to the attributes that are assigned to them, digitally signed by the issuing attribute authority (AA). Attributes describe a user's rights; target services read the user's AC to see if s/he is 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.

Issues and challenges of Shibboleth in e-Infrastructure for social sciences

Our security framework for social science e-Infrastructure will leverage Shibboleth as the authentication infrastructure and combines it with PERMIS authorization technology to provide scalable and flexible Grid authentication as well as policy-driven, role-based, multi-grained authorization for access to and usage of the resources in the e-Infrastructure. However, the uptake and adoption of Shibboleth and PERMIS within the e-Infrastructure is not without potential concerns. In particular, three critical issues need to be addressed:

- **Integration of Shibboleth with other authentication systems** Currently, many tools and most publicly-curated datasets of social science interest rely on the Athens access management system (Athens, 2007) and community investments in Athens-based security control needs to be preserved. At the same time, the e-Infrastructure will utilize the NGS which relies on the PKI-based Grid Security Infrastructure (Gsi, 2007) built in the Globus Toolkit (Globus, 2007). All these demand a pragmatic strategy for integration with the Shibboleth-based security framework in order to both preserve investment and adapt to the emerging environment.

As such, we need to build gateways between Shibboleth infrastructure and other authentication systems including the above-mentioned Athens and PKI-based Grid Security Infrastructure of the Globus Toolkit. Therefore a Shibboleth-based authentication model needs to be established to interoperate users' identity in different security domains.

We will follow developments in the Eduserv's Shibboleth-Athens gateway (Shibathens, 2007) as well as projects on Shibboleth-enabled NGS including the ShibGrid project (Shibgrid, 2007) and the SHEBANGS project (Shebangs, 2007), with a view to their adoption for the security framework in the e-Infrastructure project. We will also examine related projects such as the GridShib project (Gridshib, 2007) and the ESP-Grid project (Espgrid, 2007) which investigates how Shibboleth offers solutions to the issues of Grid authentication, authorization and security.

- **Using PERMIS for authorization management** When used in conjunction with Shibboleth, PERMIS extends the authorization system used in Shibboleth by introducing hierarchies of roles and conditional decision making. In PERMIS, each site administrator throughout a VO can act as an attribute authority and assign attributes to his/her site members. Thus the allocation of entitlements (or ACs) is distributed throughout the entire VO, and the burden of VO administrators is relieved.

To use PERMIS for authorization, a set of attributes describing the users of the e-Infrastructure needs to be defined, and these user attributes must be negotiated between domain site administrators of a Grid VO. We will refer to the attribute sets used by the UK Access Management Federation for Education and Research (Ukfederation, 2007) when defining the user attributes set. Second, depending on the social science applications running upon the e-Infrastructure, we need to write the authorization policy stating the rules for assigning roles to users and permissions to roles. Lastly, a PERMIS-

based authorization engine should be designed to make access control decisions, and manage user attributes and authorization policies. We will explore the Open PERMIS project (Openpermis, 2007) and utilize its open source package as a foundation in the design and implementation of our authorization engine.

- **Integration of Shibboleth with PERMIS** The goal of the Shibboleth and PERMIS integration is to coordinate the identity federation and attribute assignment function of Shibboleth with the policy-based enforcement function offered by the PERMIS access control infrastructure in order to provide policy-driven attribute-based access control decision making based on user attributes.

There is some existing work addressing the integration of Shibboleth and PERMIS. We will leverage the outcome of certain projects such as the GridShib and GridShibPERMIS (Gridshibpermis, 2007) projects which aim to integrate the combined Globus Toolkit and Shibboleth infrastructure with the PERMIS authorization infrastructure. We will also investigate the BRIDGES (Biomedical Research Informatics Delivered by Grid Enabled Services) project (Bridges, 2007) which leverages Shibboleth as the authentication infrastructure and subsequently uses PERMIS to determine authorization decisions based on user attributes retrieved from Shibboleth.

Conclusions and future work

Security is a fundamental requirement in the e-Infrastructure for the social sciences project. The emergence of Shibboleth and advanced authorization technologies like PERMIS provides a promising solution for the security infrastructure. In future work, we will investigate typical use cases in the social sciences, and design and implement a demonstrator security framework based on the principles discussed in this paper. The security framework is expected to be capable of granting users secure, anywhere, anytime access to a variety of resources in order to facilitate their social science research. As a consequence, Shibboleth would be adopted as an integral component in the strategic approach of the future development of the next generation secure e-Infrastructure for social science education and research.

References

- Athens. (n.d.): Athens for Education, retrieved on May 2007, from <http://www.athens.ac.uk/>.
- Bridges. (n.d.): BRIDGES project, retrieved on May 2007, from <http://www.brc.dcs.gla.ac.uk/projects/bridges/>
- Broadfoot, P. J. and Martin, A. P. (2003). A Critical Survey of Grid Security Requirements and Technologies, Oxford University Computing Laboratory Technical Report, PRG-RR-03-15.
- Chadwick, D.W. and Otenko, A. (2003). The PERMIS X.509 Role Based Privilege Management Infrastructure, *Future Generation Computer Systems*, vol. 19, no. 2, 2003, pp. 277-289.
- E-Infrastructure. (n.d.). e-Infrastructure for Social Science project, retrieved on May 2007, from <http://www.ncess.ac.uk/research/hub/einfrastructure/>.

- Espgrid. (n.d.). ESP-Grid project, retrieved on May 2007, from <http://labserv.nesc.gla.ac.uk/projects/esp-grid/index.html>.
- Farrell, S. and Housley, R. (2002). An Internet Attribute Certificate Profile for Authorization, Internet-draft 2002.
- Foster, I., Kesselman, C., Tsudik, G. and Tuecke, S. (1998). A Security Architecture for Computational Grids, the ACM Conference on Computers and Security, 1998.
- Foster, I. and Kesselman, C. (1999). *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, San Mateo, USA.
- Globus (n.d.). Globus Toolkit, retrieved on May 2007, from <http://www.globus.org>.
- Gridshib.(n.d.). GridShib project, retrieved on May 2007, from <http://gridshib.globus.org>.
- Gridshibpermis (n.d.). GridShibPERMIS project, retrieved on May 2007, from http://www.jisc.ac.uk/uploaded_documents/GRIDShibPermis.pdf
- Gsi (n.d.). Grid Security Infrastructure, retrieved on May 2007, from <http://www.globus.org/toolkit/docs/4.0/security/key-index.html>.
- Openpermis (n.d.). Open PERMIS project, retrieved on May 2007, from <http://www.openpermis.org/>.
- Saml (n.d.). Security Assertion Markup Language (SAML) v1.1. OASIS Standard 200308. OASIS Security Services Technical Committee, retrieved on May 2007, from <http://www.oasisopen.org/specs/index.php#samlv1.1>, 2003.
- Scavo, T. and Cantor, S. (2005). Shibboleth Architecture Technical Overview, Internet 2 document: draft-maceshibboleth-tech-overview-02, available at <http://shibboleth.internet2.edu/docs/draft-mace-shibbolethtech-overview-latest.pdf>.
- Shebangs (n.d.). Shibboleth Enabled Bridge to Access the National Grid Service (SHEBANGS), retrieved on May 2007, from <http://www.sve.man.ac.uk/Research/AtoZ/SHEBANGS>.
- Shibathens (n.d.). Shibboleth-Athens gateway, retrieved on May 2007, from <http://www.athensams.net/news/shibgatewaylaunch.html>.
- Shibboleth (n.d.). Shibboleth project, retrieved on May 2007, from <http://shibboleth.internet2.edu>.
- Shigrid (n.d.). ShiGrid project, retrieved on May 2007, from <http://www.nesc.ac.uk/esi/events/622>.
- Ukfederation (n.d.). UK Access Management Federation for Education and Research, retrieved on May 2007, from <http://www.ukfederation.org.uk/>.