Final Report

iamsect

Inter-institutional authorisation management to support eLearning with reference to clinical teaching

Janet Wheeler, Joint Project Manager

April 2006

Table of Contents

Acknowledgments

This project was funded under the JISC Core Middleware Development Programme. It was a collaboration between:

- IT Services, Durham University;
- Information Systems & Services, Newcastle University;
- Faculty of Medical Sciences Computing, Newcastle University;
- Medev, the Higher Education Academy Subject Centre for Medicine, Dentistry and Veterinary Medicine based at Newcastle University.

The project gratefully acknowledges the support of the parent institutions.

We also wish to acknowledge the following.

- The support given to us by our Advisory Board: Mark Adams, Northumberland, Tyne and Wear Strategic Health Authority; Suzanne Cholerton, Director of Medical Studies, Newcastle University; Trevor Cornwell, NorMAN Management Group; Gareth Davies, JISC RSC Northern; Prof Pali Hungin, Dean of Medicine and Head of the School for Health at Durham University; Lyn Norris and Richard Dunning, Eduserv; John Peacock, City of Sunderland College, Katriona .Watson, Faculty of Medical Sciences, Newcastle University.
- Martin Hatfield, Northumbria University, for his all too brief iinvolvement with the project.
- Nicole Harris and Ann Borda, our programme managers at JISC, and our fellow projects within the Programme – in particular SDSS at Edinburgh and PERSEUS at LSE – for their help and support.

Executive Summary

Background

The Core Middleware initiative arose from the identification of a need for an access management approach which would allow users to access internal and external resources using a single, institutionally controlled identity. Thus problems caused where users are required to maintain multiple passwords for multiple resources in multiple domains would be substantially reduced or eliminated. These factors were instrumental in JISC's decision to begin deployment of Shibboleth technology as the basis for the next generation of access management.

The Web Team at Newcastle University had also identified this need and had already deployed the authentication component of Shibboleth; it was therefore in a position to join forces with Faculty of Medical Sciences Computing to instigate a proposal in response to the JISC call for Core Middleware development projects. The proposal was based on two premises:

- the real and current need to authorise access to institutional resources from medical students shared with Durham University and NHS employees in a controlled manner;
- the realisation that responsibility for staff and students outwith clinical teaching is increasingly being shared across organisations and that there was a consequent need for awareness of the potential of Shibboleth to be raised in order to achieve a critical mass of adopters.

There was therefore value in a project which would attempt to implement Shibboleth to authorise access to resources used in clinical teaching on an inter-institutional basis, and to document and disseminate the process. It should be emphasised that, at the time the proposal was made, there was no Shibboleth-enabled resource or Shibboleth federation extant in the UK, and a only a single implementation of a Shibboleth identity provider; it was therefore difficult accurately to determine the project scope.

Aims and Objectives

The overall aim of the project was to learn and to contribute to knowledge by testing practical approaches, and identifying and documenting the issues and outcomes to guide institutions in their uptake of Shibboleth technology.

The specific objectives were to:

- install Shibboleth, including identification of prerequisites;
- use Shibboleth to authorise access to content across participating institutions;
- scope policy, management and legal issues
- Evaluate the impact on the end user in comparison with ATHENS access management services.
- Document and disseminate all appropriate processes and findings

In the course of the project, the use of Shibboleth to authorise access to content on an inter-institutional basis became the overriding focus at the expense of evaluation of the impact on the end user and, to a lesser extent, the scoping of policy, management and legal issues. The reasons for this were the technical difficulties encountered in the course of Shibboleth-enabled resource implementation, the initial underestimation of the resource required for effective dissemination, and the covering of the deprecated objectives elsewhere in the course of the expansion of the Core Middleware programme.

Methodology

The project methodology was to use a simple staged approach, identifying concomitant policy, management and legal issues and documenting and disseminating the work as it proceeded. The steps were to consist of:

- installation of a Shibboleth Identity Provider at Newcastle and Durham Universities;
- use of Shibboleth for access to electronic journals;
- investigation of authorisation attribute identification, storage and retrieval;
- deployment of Blackboard Learning System as a Shibboleth resource in multiple institutions;
- writing of attribute release policies;
- use of Shibboleth to authorise access to the Newcastle eGuides medical VLE, which is based on the open source Zope web applications platform.

Implementation

Early days

The first phase consisted of the mundane business of setting up the project management structure and web site, initiation of project officer recruitment, identification of other staffing resources, and purchase and setup of a server for the Shibboleth Identity Provider (IdP).

Recruitment of the project officer went well, especially considering the demands of the role, and our choice took up his post at the end of September 04, some 4 months after the start of the project. In contrast, negotiations with the technical author who we had identified to produce the documentation outputs failed at a late stage due to a change in their personal circumstances. This resulted in a lengthy process to identify and subcontract a replacement which did not bear fruit until April 05, and delayed production of some of the project outputs more than we would have wished.

Nevertheless, a guide to the installation of the Pubcookie WebISO (the authentication component of Shibboleth), *Installing Pubcookie on Redhat AS 3.0 and authenticating against Windows Active Directory,* was published on the project web site in August 04. This raised the question of licensing arrangements for the project outputs and it was decided to use the Creative Commons By Attribution licence (creativecommons.org/licenses/by/2.0/).

The road to electronic journal access

Access to electronic journals is the most straightforward use of Shibboleth as in general it does not require any complex access decisions, just the assertion that the accessor is a member of a specific institution. What it does require is

- a Shibboleth IdP at the accessor's institution;
- an electronic journal configured as a Shibboleth resource;
- a Shibboleth federation to which the IdP and resource both belong.

Before the Newcastle University IdP was installed, a decision was taken to use version 2 of the Apache web server rather than the previous institutional practice of using the more mature version 1.3. The installation was achieved in September 04 and was the first institution-level IdP to join the

SDSS development federation. The installation process was documented in *Installing Shibboleth on Redhat AS 3.0*; this and the Pubcookie installation guide were used by Durham University to install their Shibboleth IdP; both guides were subsequently refined as a result of Durham's experiences.

An issue that arose at this point was the requirement of the SDSS Federation for the use of Globalsign X.509 certificates (as recommended by UKERNA). Newcastle have an institutional agreement with Thawte and were consequently using self-signed certificates; Durham were willing to purchase Globalsign certificates but found it difficult to actually do so. It was approximately another 6 months before SDSS were able to accept Thawte certificates and Globalsign were persuaded to sign Durham University's procurement agreement.

Nevertheless, the project gave its first public demonstration of access to an electronic journal via Shibboleth in the course of a short presentation, *Introduction to Shibboleth and the iamsect project,* at the Durham Blackboard Conference in December 04 using the Shibboleth enabled journal *BIOSIS* provided by SDSS. *Practical access to electronic journals using Shibboleth* and *Introduction to Shibboleth Federations* were published in July 05.

Authorisation

In January 05 the project visited the SDSS project in Edinburgh en masse and spent a highly instructive day talking about federated identity assurance, the operation of federations, X.509 certificates, and authorisation attribute selection and storage. As a result of these discussions it was subsequently decided to use the attributes highly recommended by the InCommon federation: givenName, sn, cn, eduPersonScopedAffiliation, eduPersonPrincipalName and eduPersonTargetedID.

The next question was where to get the authorisation data from. It had been initially envisaged that the Windows Active Directory could be used but Newcastle University's Active Directory was found not to contain sufficient information to make access decisions. The investigation of populating it with suitable data was deemed to be impractical as this would have required setting up a clone of the production system which would have to have been bridged off the institutional network. It was intended that Northumbria University, who were one of the original project partners, would investigate the use of the Active Directory and ADAM (Active Directory in Application Mode) as an attribute source but they unfortunately decided to withdraw from the project before this could take place.

As an alternative, Newcastle has been using institutional data feeds; because of the way the institutional data is held, it is necessary to query three separate databases in order to obtain authorisation data. As version 1 of the Shibboleth IdP is only able to accept authorisation attributes from a single source, this situation is far from ideal but should be mitigated with the next version of Shibboleth.

Following the above investigations, the document *Attribute identification and storage* was published and we moved on to consider the authorisation of access to medically restricted content. As far as students were concerned, this proved to be straightforward: medical students were identified from institutional data, the medical restrict attribute set for them and the necessary declaration sent to the federation. Identifying the equivalent staff has proved to be rather more of a headache.

Staff who teach on the medical programme at Newcastle University are not restricted to, or even a superset of, those employed within the Faculty of Medicine, and may come from external institutions or within the NHS. Unfortunately there are no institutional data that identify them as being entitled to access medically restricted material. We have not yet solved this problem; Groups Manager or authentication from multiple sources as being pioneered by Faculty of Medical Sciences Computing may be the answer. In contrast, Durham University do have the requisite information in an institutional database.

Blackboard as a Shibboleth resource

In the original proposal it was intended that the Blackboard development work be done at Newcastle. However, following problems with Newcastle's Blackboard system during 2004, Durham offered the use of their Blackboard development licence to the project. In the final analysis this use of a Linux rather than a Windows-based Blackboard system was most fortuitous but it did result in a delay whilst Durham set up a Blackboard development system.

Combined with the delay in making the Durham IdP operational, this meant that it was October 05 before Durham's Blackboard development server was configured as a Shibboleth Resource Provider

within the SDSS federation. This process was documented and feedback provided to Blackboard detailing some limitations identified in their configuration instructions and out of the box configuration.

Jubilation and bottles of fizzy turned to dismay when difficulty was subsequently found in configuring Blackboard to use the SDSS federation WAYF, which meant that members of the federation other than Durham were excluded from accessing the server whether they were entitled to or not, as they could not access their own IdP in order to authenticate. Enabling use of the WAYF was not achieved until March 06 after contacting the Senior Software Architect at Blackboard. There remains an issue with Durham's firewall which is preventing access from external institutions; it is hoped to solve this soon.

In addition to the above, other issues have been identified that lead the project to believe that the current version of Blackboard is not yet ready for deployment as a production Shibboleth resource. These include: fragility of some links to the Content System; inability to log out without closing the browser (a generic problem with Shibboleth at the current time); loss of Blackboard's "portal direct entry" feature. Durham is, however, actively pursuing fixes as they regard use of Shibboleth with Blackboard as important to their strategy for migration from Athens DA to Shibboleth.

It had been intended to use the work done at Durham to configure a Windows/IIS Blackboard system as used at Newcastle as a Shibboleth resource. This has proved not to be possible as currently Blackboard only support Shibboleth on Apache Tomcat (i.e. for Linux and Unix platforms). A request to Blackboard Technical Support indicated that implementation on a Windows/IIS system is possible but that it would require custom development work on their part which they would expect to be financed by the requesting site.

The first draft of *Installing a Shibboleth Service Provider,* which includes a description of the Blackboard server configuration was published in March 06.

A Zope-based VLE as a Shibboleth resource

The aim of this phase of the project was to Shibboleth-enable eGuides, a Zope-based VLE used in the Faculty of Medical Sciences, and, following a pilot service, to use this system in a production environment. At the time that the project proposal was made, the Zope corporation had indicated that Shibboleth functionality (as contained in their proprietary Zope4Edu content management system) would be made generally available during 2004. Unfortunately this proved not to be the case and so an alternative solution was sought.

A Shibboleth-enabled Zope system was achieved in March 05 by using an Apache web server as a front end to the system and passing requisite information between this and the Zope server. We then discovered that the software functionality used to pass the information between the servers was to be deprecated in future versions of Zope, so that this was not a satisfactory long-term solution. A further problem was that, while this was a solution for a single-server pilot system, it was not suitable for the production eGuides system which utilises multiple load-balanced servers.

The latter consideration amongst others lead to a re-evaluation of the system to be used as a Zope-based resource. A customised version of eGuides configured as a Shibboleth resource was produced for the use of a group of some 100 postgraduate students studying for a Masters in Clinical Education. This was deployed as a production service at the start of the 05/06 academic year, utilising an Apache front end on the Faculty of Medical Sciences Computing (FMSC) server and the test IdP. Problems were encountered in moving the service to the production IdP, which necessitated setting up the front end on a dedicated server.

The users of this system were issued with with standard Newcastle student user identifiers for access and it was hoped that, as some of them are based at Durham, it would be possible for them to test access to the system using their Durham credentials via Durham's IdP. Some 3 months into the academic year, it transpired that a proportion of the system's users had been unable to access it. This was because standard Newcastle identifiers must have their password changed on first use. This cannot be done via a webISO such as Pubcookie but many of the users are clinicians who are hospital-based and never make use of the University Windows services and so were both unaware of this requirement and unable to authenticate.

Back to the drawing board. FMSC have now implemented their own IdP which is integrated with CAS (an alternative webISO). CAS is better integrated with Zope, not requiring the use of the Apache front end, and is able to take authentication details from more than one source – in this case the Newcastle Active Directory is being used as the primary source, with FMSC's LDAP directory as a secondary.

They are also developing a Shibboleth-enabled 'portal' which it is hoped will allow multiple web applications to act as Shibboleth resources.

Documentation of this approach, and of the lessons learned in arriving at it, will form one of the final project deliverables.

This has been possibly the hardest development road of the project, necessitating much research and testing, and the drawing together of many developers in order to co-ordinate the design and implementation of the final system. However, FMSC now believe that they are at a stage where they can easily Shibboleth-enable applications with minimal development work, allowing larger systems to retain their flexibility and ability to support many users without the use of any deprecated software functionality.

Outputs and Results

The systems set up by the project in the course of production of the outputs are described above.

The project has produced the following publications:

[final list].

Other material includes

- A glossary of terms associated with the work.
- A survey of existing Shibboleth-related documentation.
- A 2 page article entitled *Shibboleth and clinical teaching* by members of the project was published in 01 – the newsletter of the HE Academy subject centre for medicine, dentistry and veterinary medicine (print run of 2,500, also available at www.medev.ac.uk). This article has also been distributed to other HEA subject centres. for possible use in their newsletters.

Two major dissemination events have taken place:

- *Shibboleth Technology and the IAMSECT Project,* a general non-technical introduction.
- *Shibboleth – the technicalities*

Presentations have been given at three JISC Programme Meetings and at four conferences:

- Durham Blackboard Users' Conference, 2004;
- Institutional Web Management Workshop, 2005;
- Breaking Boundaries, 2005;
- Durham Blackboard Users' Conference, 2005.

Additional dissemination activities have included:

- a short presentation to members of Hull and York Universities (who share a medical school), followed by a demonstration of Shibboleth in action and a lively question and answer session;
- a presentation to members of Newcastle University outlining the project's implementation work, and explaining how content providers could prepare for Shibboleth;
- a presentation at the Higher Education Academy's Technical Awayday;
- discussions with the Director of the National Knowledge Service, Implementation Directorate, NHS Connecting for Health, with a view to setting up a group to look at single sign-on and authorisation issues affecting the NHS from a wider perspective.

Publications and dissemination materials are available on the project web site (iamsect.ncl.ac.uk) under a Creative Commons licence. The web site itself has received over 17,000 unique visitors during the project lifetime.

Outcomes

The project outputs have eased the task of Shibboleth early adopters, who have been recruited in increasing numbers as the Core Middleware Programme has expanded and succeeded. Particular examples are as follows.

- Newcastle University Library's SAPIR project, which has used the documentation and Shibboleth infrastructure produced by **iamsect** to investigate the use of a Shibboleth-enabled version of the Metalib library portal, and to use Shibboleth in place of Athens to authorise access to electronic journals.
- City of Sunderland College, and Sunderland and Teesside universities, as part of the Shibboleth component of the EPICS (distributed elearning) project, which has been heavily supported by **iamsect**.
- St George's Hospital Medical School have used and helped refine the installation guides in the course of their ADAMS (Authentication & Delivery across Medical courses using Shibboleth) project.

The project has been successful in raising the profile and demonstrating the potential of Shibboleth, especially within the region and including the NHS.

The parent institutions have become firmly convinced of the value of Shibboleth and are wholeheartedly embracing its use.

- Durham University is instigating a project to transfer access management from Athens DA to Shibboleth, and intends to use Blackboard as a Shibboleth-enabled service as soon as that is technically feasible.
- The development done as part of the project should see a number of web-based applications belonging to Newcastle University's Faculty of Medical Sciences Computing Shibboleth-enabled for the forthcoming academic year, with the consequent benefits of intra- and inter-institutional single signon and authorisation.
- Information Systems and Services at Newcastle University are already running a Shibboleth-enabled mailing list manager in production; more services will follow. ISS is also working with SAP in the area of federated identity management.

Finally, the project has sown the seeds of a wider regional collaboration, which can only flourish to the benefit of all concerned.

Implications

iamsect has demonstrated the potential of Shibboleth over and above being a simple replacement for Athens access management. This can only encourage wide adoption during the transition from Athens to Shibboleth during the coming two years.

References

- *Connecting People to Resources*, JISC's plans for access management services in the UK HE and FE community.
- SECURe project (Secure Environment for Certificated Use of Resources) http://www.angel.ac.uk/SECURe/