

Shibboleth and SAML: at last, a viable global standard for resource access management

John Paschoud, London School of Economics & Political Science Library
<j.paschoud@lse.ac.uk>

January 2005

Abstract

The library and publishing communities have used various imperfect solutions for managing access to online resources, but now Shibboleth and an infrastructure based on it look set to become a global public domain standard that has been designed to meet all the requirements. This paper reviews the business issues of access management, briefly describes common mechanisms currently in use, and explains what Shibboleth is and how it works. Britain is one of several countries planning for and investing in Shibboleth for use by its' national academic community, via programme funding of over £7million for technology development projects and infrastructure support, and the history of progress to date with this rapidly moving work is outlined.

Introduction

The academic library and publishing communities have been coping with imperfect solutions to the problem of controlling access, since the earliest days when information resources were available online. The Internet and the Web fulfil the vision of Tim Berners-Lee [T1] in making 'everything available to everyone'; but they do not, yet, include an established, pervasive *and globally scalable* standard for making some things available only to some people - in a way that realistically models the access licences between publishers (of e-journals and other online resources) and the libraries that pay for such privileged access on behalf of their communities of registered users.

A strong candidate to become such a standard is now emerging, from work mainly driven by the Internet2 Programme [2], in the form of the OpenSAML implementation of the OASIS Security Assertion Markup Language (SAML) protocol [3] and Shibboleth [4], a set of Open Source Software (OSS) components which implements SAML and are available for libraries and publishers to use freely themselves, or for third party vendors to implement to provide access management services on a commercial basis.

It remains to be seen, in October 2004, whether Shibboleth and SAML will achieve the global acceptance that they must, to become not merely the most technically sound solution to these problems, but also the de facto standard.

Resource access management: the issues and requirements

The conceptual requirements and business relationships between publishers, libraries and end-users were comprehensively analysed in work led by Clifford Lynch [5] of the Coalition for Networked Information in 1998. The 1999 workshop report by the US-based Digital Library Federation [6] identifies five key properties “for the design ... of systems that enable access for users while respecting the rights and interests of authors and publishers.” These are summarised below:

1. **Simplicity.** The less complex a system of access management, the more readily it can be adopted technologically and organisationally, and the more acceptable it is to all involved in its implementation.
2. **Privacy.** Systems that manage access to the cultural record must protect the privacy of users from detailed tracking and disclosure of use. User privacy must not be compromised.
3. **Good faith.** Agreements on access to scholarly information rely on trust among the parties involved. Users and providers would each prefer to depend, in an access management system that implements these agreements, on reasonable barriers against abuse rather than complex restrictions that inhibit use.
4. **Trusted intermediaries.** Intermediaries play an essential role in providing access to the cultural record as parties trusted by both users and providers and as efficient aggregators of distribution and usage. System design must take the role of intermediaries into account.
5. **Reasonable terms.** Access management systems and license agreements must recognize the distinction between access and use. Overly tight control of access to a resource may impose inappropriate constraints on its use, especially in teaching and research contexts. The most useful system will not limit access to specific user groups known in advance to be interested in a resource, but will be reasonably open to serving unlikely users whose curiosity and research interests may lead them in directions not predicted by those responsible for making the agreements or designing the systems.

Despite a great deal of activity in the field since this work was done, by both academic/public sector players and commercially driven concerns (not least the Microsoft .Net initiative [7]), there has been no significant improvement upon the above set of objectives.

I described, in 1999 [8], the developing circumstances of greater diversity of resources and greater mobility expectations of users that would define even more demanding future requirements for access management technologies. If anything, these were understatements of the problem, and other trends in ways of making library resources accessible such as portalisation [9], have

Shibboleth and SAML: at last, a viable global standard for resource access management

made the need for a better solution and consensus on protocols for access management even more acute.

Current solutions and their deficiencies

For academic librarians, licensing and offering access to the range of online resources that are regarded as necessary to support leading edge research in any discipline, one of the big problems to be tackled is maintaining the balance between adhering to their legal and contractual responsibilities to publishers (to limit access to only those users covered by license terms) and to users (to protect the privacy of personal information registered with them), and to carry out that fundamental function of a library - to offer users the easiest possible path to the information that they need.

A range of methods to limit access is currently in common use, either enforced by resource owners or adopted by libraries:

Common shared 'secret' passwords

When made available to the hundreds or thousands of individuals who may have access rights via even a single library to a resource 'protected' in this way, the 'secret' is unlikely to remain within the authorised community for long. It cannot be changed (and re-communicated to the authorised community) easily or frequently.

Registration of individual users

This is a perfectly reasonable thing to expect of the organisation that has some contractual or business relationship with a user - the library, or the university or college of which the library is part. Most library users will expect to share identifying and other personal information with a library on a similar level of trust as they will have with a bank. (Some users would certainly regard the university or library with which they are affiliated as more trustworthy than their bank, not to misuse such information). The library should also be in a position to know when the eligible status of a user is terminated – usually when a student or staff member leaves the institution.

However, many library users will not expect (or be willing) to give such information to an external resource provider (and, potentially, many different external resource providers) with whom they do not have such a trust relationship. Apart from the legal obligations of a university to protect personal information disclosed to them, there are real and serious reasons for academic researchers wishing to avoid associations between their identity and what they are reading, for example when research fields may be commercially competitive (e.g. pharmacology) or 'politically sensitive' (e.g. experimentation involving animals).

IP address restriction

Assuming that a reasonably static range of IP (network/Internet Protocol) addresses can be reliably identified to correspond with the networked

workstations in a particular place, this method can be adequate (and may be the best) for enforcing resource licence conditions that allow only 'on-campus' users, or any person with physical access to the network within a library building. However, this method of access management (and this type of licence) precludes distance learners and others who expect remote access, and allows all those who can gain building and network access (including, for example, to publicly accessible wireless networks operated by a library), and both of these are reasons to progress from physical location as a licence condition, and IP restriction as a method of enforcing it.

Far too often IP restriction is still used to enforce licences that define access by membership (of the library) rather than location (within it), just because it is relatively simple to implement.

IP address restriction, with authenticated proxy-servers for off-campus users

This is a compromise that libraries, under increasing end-user pressure to offer 'access from anywhere' are adopting. Put simply, it allows a user to be 'virtually' on-campus, by accessing the desired (IP-restricted, as above) resource via an intermediate server, itself within the allowed network address range. The drawbacks are that this method cannot be used where licence conditions actually specify the allowed location of users, and it is technically challenging to configure proxying to some resources, particularly those that have complex or dynamic end-user interfaces; and then to reconfigure it, each time one such resource provider decides to re-engineer such an interface.

A more accessible but less sound variant of this compromise is the use of proxy-servers that don't require authentication of users. Despite some high profile cases (notably that involving JSTOR in December 2002 [10]) in which these have been used for 'pirate' extraction of significant quantities of licensed resources, it's not hard to find these still being operated by libraries.

Athens

For the British academic community, a notable achievement has been the implementation of the Athens [11] system, a national shared service first established in August 2000, which effectively allows publishers and other resource providers to outsource the management of individual usernames and passwords to a central service, and similarly allows universities and colleges to outsource part of their work of administering lists of the resources to which each of their users should have access.

Athens has been made economically possible for the UK F&HE community principally because the element that services academic institutions was (and continues to be) centrally funded by JISC (the Joint Information Systems Committee, itself funded by the plethora of national and regional government bodies that fund UK public-sector post-16 education), and (what has proved to be) a critical mass of access-managed resources are also available to the same institutions from national data service centres (principally EDINA and MIMAS, based at the universities of Edinburgh and Manchester), similarly subsidised by JISC funding. Commercial publishers can license their own

Shibboleth and SAML: at last, a viable global standard for resource access management

access to the Athens service (and use of the proprietary software components involved) for an annual fee, and of course this is an attractive way for many of them of achieving the same level of security as if they were maintaining their own individual access credentials for the 3million-plus registered individual Athens users. Not at all coincidentally, the Athens service is operated by Eduserv, a not-for-profit body which also operates the CHEST service that negotiates national purchasing agreements with suppliers to the education community.

What is Shibboleth?

In brief, Shibboleth is software that implements SAML protocols, separating the functions of *authentication* (undertaken by the library or university, which 'knows' its' community of end-users) and *authorisation* (undertaken by the resource provider, which knows which libraries have licenses for their users to access the resource in question).

However, this requires a non-trivial amount of infrastructure to be established, and the Shibboleth Project led by the Internet2 Middleware Architecture Committee for Education (MACE) is addressing the organisational, as well as technical aspects, and developing a policy framework to cover:

- **Federated Administration.** The Identity Provider (IdP) institution ('home organisation' to the end-user) provides attribute assertions about that user to the Resource Provider (ReP) site. A trust fabric exists between the member institutions and publishers of a Shibboleth Federation, allowing each site to identify the other speaker, and assign a trust level. Identity Provider sites are responsible for authenticating their users, but can use any reliable means to do this.
- **Access Control Based On Attributes.** Access control decisions are made using those assertions. The collection of assertions might include identity, but many situations will not require this (e.g. accessing a resource licensed for use by all active members of the university community, or accessing a resource available to all students in a particular course).
- **Active Management of Privacy.** The Identity Provider and the end-user control what information is released to the Resource Provider. A typical default is merely "member of community". Individuals can manage attribute release via a web-based user interface. Users are no longer at the mercy of whatever privacy policy is adopted by each Resource Provider.
- **A Framework for Multiple, Scaleable Trust and Policy Sets (Federations).** Shibboleth uses Federations to specify a set of parties who have agreed to a common set of policies. This moves the trust framework beyond bi-lateral agreements, while providing flexibility when different situations require different policy sets.

Shibboleth and SAML: at last, a viable global standard for resource access management

- **A Standard (yet extensible) AttributeValue Vocabulary.** Shibboleth has defined a standard set of attributes; the first set is based on the Heduperson [12] object class that includes widely-used person attributes in higher education.

It is important to note that Shibboleth is *not* a method of authentication, as such; it depends upon whatever mechanisms the IdP puts in place. Typically and currently this is likely to be a local 'single sign-on' solution requiring knowledge of a network login username and password, such as the Open Source CAS (Central Authentication Service) [13] originally developed at Yale. This has the benefit of allowing users to use (and, hopefully, remember) the name/password combination that they use most frequently, to gain access to resources that they may only use infrequently. This design aspect of Shibboleth, separating authentication from authorisation, ensures that such passwords are never transmitted (even in encrypted form) outside of the users' own institutional network and the secure (HTTPS) connection between the users' browser and the campus authentication server. If an institution decides to use different and more secure methods of authentication (such as digital certificates carried on 'smartcards' or other storage devices, or biometrics), this is a purely local decision and will not affect access to any resources mediated by Shibboleth.

Underlying the original requirements for Shibboleth was the concept of a need for trusted peer-to-peer access: the individual users at one university accessing resources hosted by another university. This was a natural assumption by most members of the original development group, most of whose roles were in managing campus network infrastructures and resources. In the early stages of wider testing and deployment of Shibboleth, information professionals based in libraries became involved, and realised that Shibboleth could be an even more effective solution to the (much larger-scale) problem of access to published resources, with user communities spanning very large numbers of institutions and national boundaries.

Shibboleth provides a trustworthy way for a large community of organisations, each managing a large community of users and/or many collections of resources, to collaborate in controlled user access to resources with a minimal overhead requirement for exchange and updating of information between organisations, whilst meeting all the business requirements for security of resources and privacy of information about users.

How Shibboleth works

At first sight, the sequence of transactions needed to effect an access decision with Shibboleth seems dauntingly complex and therefore difficult to implement and prone to breakdown. However, it must be remembered that the HTTP (HyperText Transfer Protocol) messages that pass between a web browser and one or more web servers, in the course of viewing a relatively simple page of content, would seem almost as daunting if viewed at a similar level of detail. One of the Shibboleth developers, Michael Gettes of Duke University, keeps in reserve as a presentable explanation of "How Shibboleth

works” the answer, “Magic!”. This may indeed be sufficient for information professionals and policy makers who want merely to use the Web as a medium, rather than tinkering with the way it works.

For others who insist on the full gory details (or just like to tinker), here they are (largely adapted from an explanation and diagrams produced by SWITCH, the Swiss Education & Research Network, which has implemented the first operational national-scale Shibboleth Federation [14]).

It is first necessary to know that (apart from the end-user and her web browser) there are three actors in this exchange:

- The **Resource Provider (ReP)**: This is the host of the resource that the user wants to access, which must make an authorisation decision about whether (and possibly at what level) this user is allowed to access this resource. The ReP operates two Shibboleth software components, the SHIRE (SHibboleth Indexical Reference Establisher), and the SHAR (SHibboleth Attribute Requester).
- The **Identity Provider (IdP)**: This is the institution (library, university, or other ‘membership’ or ‘home’ organisation) with which the user is registered, and which (in most cases) has a contractual relationship with the ReP for access by all or some of its’ members. As well as a local authentication service, based on some database of registered users (typically a Lightweight Directory Access Protocol (LDAP) directory), the IdP operates two Shibboleth software components, the HS (Handle Server) and the AA (Attribute Authority).
- The **‘Where Are You From’ (WAYF) service**: This is a ‘central’ service, operated on behalf of a Shibboleth Federation, which resolves the question of “With which IdP is this user claiming affiliation, and thereby access rights?”. At its’ most basic, a WAYF can ask (via a simple Web form) the user to choose from a list of institutions; but other more automated (and more scaleable) methods of resolving WAYF are under development.



Figure 1: User, Home Organisation (IdP) and Resource Owner (ReP)

Shibboleth and SAML: at last, a viable global standard for resource access management

When a user (Alice) attempts a first access (the most complex case) to a resource in a web browser session, the actors interoperate as follows:




Figure 2: User connects to Resource Owner, and is redirected to WAYF

- (1) Using her browser, Alice connects to the web-based resource.
- (2) Since the web server detects no established session for Alice, the server hands her request over to SHIRE, which redirects Alice's web browser to the WAYF server typically run by the Federation.
- (3) The WAYF server presents Alice a web page from which she selects the name of her home organisation.

Shibboleth and SAML: at last, a viable global standard for resource access management




Figure 3: User selects her Home Organisation and authenticates there

(4) When Alice selects her home organisation, her browser returns the selection to the WAYF.

(5) The WAYF then redirects her web browser to the Handle Server (HS) of her home organization.

(6) From the Handle Server, Alice gets a web login screen of her university, well known to her since she uses the web login already for various web resources offered by her university.

Shibboleth and SAML: at last, a viable global standard for resource access management




Figure 4: Authenticated User is redirected back to Resource Owner

(7) Alice provides her credentials (e.g. username and password) to the Handle Server of her home organisation - normally via an existing and familiar web SSO (Single Sign-On) interface.

(8) Provided the credentials are correct, the Handle Server generates an opaque and digitally signed Handle on behalf of Alice. It gets sent to the resource Alice wants to connect to by another web browser redirection.

Shibboleth and SAML: at last, a viable global standard for resource access management




Figure 5: SHAR fetches User's attributes on behalf of the Resource Owner

(This step is completely invisible for Alice, since it is a server to server communication in the background between the Shibboleth components at the ReP and IdP.)

On the resource side, the Handle received gets passed to the SHAR (Shibboleth Attribute Requestor) component.

(9) The SHAR then sends it via a secure HTTP connection to the Attribute Authority (AA) at the home organisation which generated that Handle.

(10) The Attribute Authority verifies the Handle and its validity internally with the Handle Server. If valid, the AA checks out which attributes it may release to the resource based on the Attribute Release Policy (ARP) of Alice regarding the resource. The AA sends the attributes allowed to release, digitally signed, to the SHAR.




Figure 6: Access Control Manager decides on authorising User's access

Finally, the SHAR passes the attributes received to the Access Control Manager which then, according to its configuration, authorises the access for Alice based on the set of attributes provided. If a Resource requires information about the user for functional purposes (such as personalisation), the Access Control Manager can request and pass these attributes to the Resource.

Once authorised, a session is established and the communication between Alice and the resource within this session does not need any further involvement of Shibboleth components.

What the UK academic community is doing

Shibboleth originated as part of the US Internet2 Middleware Initiative, but there has been a high level of international involvement in its development from a very early stage. In spring 2001 JISC commissioned staff at London School of Economics (LSE), working on the ANGEL Project [15] (which was investigating models for resource access management and single sign-on) to liaise with Shibboleth developers and to try out the installation of Shibboleth software.

The JISC Authentication, Authorisation and Accounting (AAA) Programme [16] was launched in summer 2002, and included a specific focus on Shibboleth. This enabled new projects (including SECURE at LSE [17]) to investigate further whether Shibboleth (and several other possible candidates that seemed promising at the time) could form the basis of a next generation

access management infrastructure within JISC's envisaged Information Environment [18] for academia.

The Internet2 community, comprising predominantly information technologists from subscribing US universities, has been enthusiastic and welcoming to international participation, from a growing number of countries [19]. They initially saw our contribution to their efforts mainly in terms of ensuring 'internationalisation' of the standards and protocols they were developing - that, for example, they were not inadvertently embedding 'US English' terminology for university structures and roles into Shibboleth. However, we also brought some useful perspectives and experience from the resource access problems that are frequently faced by libraries, but often not very apparent to those immersed in the 'harder' domains of network technology. Several other European countries have also been strongly represented in these discussions, including The Netherlands, Spain, Switzerland, Norway and Poland.

One specific element of ongoing international coordination in which we're still involved is extensions to the eduPerson Object Class, to facilitate appropriate levels of institutional, national and global agreement on the common attributes that can be used to describe roles and affiliations of students, academics and other education staff. An initial scoping study for 'UKeduPerson' was commissioned by JISC in early 2004, in parallel with short studies on a number of other issues that would inform subsequent major funding programmes [20].

Starting in spring 2004, 16 new projects were funded by JISC via the Core Middleware Technology Development Programme [21]. These are identifying and filling gaps in Information Environment middleware, progressing the development of tools for GRIDs, and testing prototype services in production environments. Several of the projects are working with Shibboleth, including the PERSEUS Project [22] at LSE, which is integrating Shibboleth with other components to enable sophisticated authorisation management in institutional portal environments.

By November 2004 JISC will have announced a detailed timetable for the creation of a national Core Middleware Infrastructure, support for academic institutions and resource hosts (both the public sector data service providers and UK-based commercial publishers), and plans to achieve the transition from the current Athens service. A Core Middleware Advisory Group, representing institutions, services and other interests, has been convened to provide a high level steer to the process. A total of over £7million JISC-managed funding is being invested over a period of three years, across the Technology Development projects and Core Middleware Infrastructure.

Eduserv has been closely involved in this process, and has been proactive in producing and publicising various initiatives that will make the Athens service increasingly 'Shibboleth-compliant'. For example, the 'AthensDA' ('DA' = Devolved Authentication) version of the service [23], now being adopted or trialled by an increasing number of institutions, has a fundamentally very similar architecture to Shibboleth, but continues (at present) to use proprietary protocols and software components.

The Athens service is at once an enormous advantage and a potential obstacle to the UK academic community in adopting a new infrastructure for access management. HE and FE institutions, and of course commercial information resource providers are all autonomous (to some extent). Their US counterparts are being offered the opportunity to adopt a standard for secure resource sharing and trading that they have never had before, at a cost (of implementing Shibboleth Identity Provider services) that is a relatively insignificant addition to that of the local network and user directory infrastructures they each already support. UK institutions have been buffered from most of the direct costs of this for five years by Athens (funded directly by JISC), and could in effect be asked to take a 'leap of faith' from an established system that works 'well enough', into something new and unknown, with costs that they cannot easily assess in advance, for taking on the responsibilities of devolved authentication.

However, the experience gained (by institutions, JISC and Eduserv) in operating Athens means that the UK community already largely understands many of the organisational and legal aspects of operating something that is very close in both scale and complexity to the Shibboleth concept of a Federation. If the transition from Athens to Shibboleth is planned, financed, managed and supported in the right way, the UK could provide valuable lessons for other countries, and continue to play a world-leading role in the establishment of a new standard in which a global information market can flourish.

References

- 1 Tim Berners-Lee; Weaving the Web; Orion Publishing, 1999.
- 2 Internet2 Programme: <http://internet2.edu/>
- 3 Security Assertion Markup Language (SAML): http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security
- 4 Shibboleth: <http://shibboleth.internet2.edu/>
- 5 Clifford Lynch (ed.); 1998; A White Paper on Authentication and Access Management Issues in Cross-Organizational Use of Networked Information Resources; Coalition for Networked Information; <http://www.cni.org/projects/authentication/authentication-wp.html>
- 6 Digital Library Federation; 1999; Enabling Access in Digital Libraries: A Report on a Workshop on Access Management; <http://www.clir.org/pubs/reports/arms-79/contents.html>
- 7 Microsoft .NET: <http://www.microsoft.com/services/net/default.asp>
- 8 John Paschoud; 1999; All Users are not Created Equal! How to decide Who Gets What from your Hybrid Library, Internet Librarian International 1999 (proceedings); Information Today, Inc. (available from: <https://gate-test.library.lse.ac.uk/dspace/handle/1988/8>).
- 9 Francisco Pinto and Michael Fraser; Access Management, the Key to a Portal: report on the experience of the Subject Portals Project; Ariadne issue 35; <http://www.ariadne.ac.uk/issue35/SPP/intro.html>

10 Dan Carnevale; Security Lapses on Campuses Permit Theft From JSTOR Database; The Chronicle of Higher Education; <http://chronicle.com/free/2002/12/2002121201t.htm>

11 Eduserv Athens Service: <http://www.athensams.net/>

12 eduPerson Object Class: <http://www.educause.edu/eduPersonObjectClass/949>

13 Yale Central Authentication Service: <http://tp.its.yale.edu/tiki/tiki-index.php?page=CentralAuthenticationService>

14 SWITCH AAI Federation: <http://www.switch.ch/aai/>

15 ANGEL Project: <http://www.angel.ac.uk/>

16 JISC Authentication, Authorisation and Accounting (AAA) Programme: http://www.jisc.ac.uk/index.cfm?name=programme_aaa

17 SECURE Project: <http://www.angel.ac.uk/SECURE/>

18 JISC's Information Environment: http://www.jisc.ac.uk/index.cfm?name=ie_home

19 John Paschoud; 2002; At the event: The Internet2 Spring Member meeting; Ariadne Issue 32; <http://www.ariadne.ac.uk/issue32/internet2/>

20 JISC Core Middleware and Shared Services Studies: http://www.jisc.ac.uk/index.cfm?name=prog_midss_studies

21 JISC Core Middleware Programmes: http://www.jisc.ac.uk/index.cfm?name=programme_middleware

22 PERSEUS Project: <http://www.angel.ac.uk/PERSEUS/>

23 Eduserv AthensDA: http://www.athensams.net/development/devolved_authentication/

Publication and Intellectual Property statement

This article has been accepted for publication in the Taylor & Francis journal *New Review of Information Networking*, on or after January 2005. The version of this article contained in this document is an author post-print, published by the author via an institutional repository collection based at: <http://hdl.handle.net/1988/14>

You are free to view, copy, distribute, display or make derivative works of any public materials created by John Paschoud as a member of the LSE Library Projects Team under a Creative Commons License detailed here: <http://creativecommons.org/licenses/by-nc-sa/2.0/>

In summary, this license allows you:

- to copy, distribute, display, and perform the work
- to make derivative works

Under the following conditions:

- Attribution. You must give the original author credit.
- Noncommercial. You may not use this work for commercial purposes.
- Share Alike. If you alter, transform, or build upon this work, you may distribute the resulting work only under a license identical to this one.

For any reuse or distribution, you must make clear to others the license terms of this work.

Any of these conditions can be waived if you get permission from the copyright holder.

Your fair use and other rights are in no way affected by the above.