Ariadne Issue 48 July 2006

ShibboLEAP: Seven Libraries and a LEAP of Faith

Martin Moyle introduces the ShibboLEAP Project, a multi-institution Shibboleth adoption in London, and hopes that later adopters will benefit from its findings.

Introduction

Much of UK Higher and Further Education (HE & FE) has begun to grapple with next-generation access management technology. Many UK developments in this area are underpinned by Shibboleth, which is conceptually simple, but architecturally complex. It is hoped that this article will benefit newcomers to Shibboleth. We offer a brief introduction to Shibboleth technology, in the context of the UK's burgeoning federated access management infrastructure. We go on to describe the ShibboLEAP Project, which saw six University of London institutions implement Shibboleth under the guidance of the London School of Economics and Political Science (LSE). The project's background, aims and core findings are summarised, and the detailed project outputs, including case studies of Shibboleth Identity Provider implementation at each participating institution, are introduced. The project deliverables may be of practical assistance to institutions which decide to implement Shibboleth as a step towards federated access management.

Shibboleth: What Is It?

Shibboleth assists with the secure management of Web resources across domains. The software and its installation have been discussed in this publication before [1], and elsewhere [2] [3], but a brief re-introduction would be timely.

Shibboleth is an output of the Internet2 Shibboleth Project [4], which is developing an architecture and policy framework to support the sharing of access-controlled Web resources. Shibboleth is not an all-in-one solution for federated access management, nor is it a SSO (Single Sign-On) system. Rather, as a component in an organisation's SSO environment, Shibboleth interacts with a local authentication system and a local user database and facilitates the exchange of authorisation and authentication information between organisations and resource providers. The software is standards-based, open-source middleware, consisting of 'Identity Provider' software (for universities and other institutions), and 'Service Provider' software (for any parties, whether internal or external, wishing to provide secure access to a resource). In a Shibboleth environment, an Identity Provider (IdP) passes attribute information about

Web-browsing users to a Service Provider (SP), which then bases authorisation decisions on those attributes. The Identity Provider is responsible for authenticating users, using whatever local technology is appropriate.

This sensible division of labour, whereby the registration and authentication functions are carried out by the Identity Provider, while the Service Provider is responsible for authorisation and accounting, carries with it an element of trust. Service Providers need to be confident that IdPs are keeping accurate records and authenticating robustly. Trust is supported by another important facet of Shibboleth: the federation. A federation is a group of organisations which have agreed to share a common set of policies and rules, forming a 'circle of trust'. Federations allow for scalability in trust and policy arrangements. An organisation can belong to as many federations as is necessary to access all the resources it requires.

Privacy is a key feature of Shibboleth technology. The IdP controls which user attributes are disclosed to the Service Provider after authentication. The information which is released concerns the authenticated user's role, rather than his or her identity - 'member of institution' is the typical default, a level of detail which is very often sufficient for a UK HE user to access a protected resource. Shibboleth can, however, support finer levels of granularity in attribute release - detail such as 'student member of institution on course x' or 'student member of institution y AND member of organisation z' could be made available to a Service Provider in cases where access was so closely-governed as to make it relevant. Meanwhile, role-based access management does not mean that personalisation is unsupported: each user is allocated a persistent resource-specific identifier, which ensures that he or she is 'recognised' on return visits, while anonymity is preserved. Finally, because the release of attribute data would be in vain without a common attribute vocabulary, the Shibboleth Project has begun the process of defining a standard set of attributes, initially based on the eduPerson object class [5].

Shibboleth in the UK

In recent years the JISC, through the Core Middleware initiative [6], has invested significantly in moving the UK towards a federated access management infrastructure founded on Shibboleth technology. At the time of writing, the launch of the new UK Access Management Federation for the UK is imminent [7]. As noted above, Shibboleth permits membership of multiple federations; and indeed, Shibboleth federations are not difficult to establish. However, only the UK Access Management Federation need be joined in order to use Shibboleth authentication to resources which are at present protected by Athens. The current JISC contract for Athens comes to an end in July 2008.

Athens has served the UK splendidly, but Shibboleth technology offers two great advantages over existing access management arrangements. First, Shibboleth permits both internal and external resources to be accessed using a single identity. In practical terms, besides supporting access to third-party resources, it has the potential to

support authentication to internal administrative systems and VLEs, to assist with inter-institutional resource-sharing, such as the sharing of e-learning resources across joint degree programmes, and to facilitate dynamic research collaboration. Shibboleth offers us a glimpse of a future in which users take seamless Single Sign-On to Web resources for granted. The second big advantage of Shibboleth is that it is not only standards-based, but that it is becoming a *de facto* international standard. Federations based on Shibboleth already exist in the US, Switzerland and Finland, and serious Shibboleth projects are under way in several other countries. Take-up of Athens outside the UK has been limited, meaning that willing resource providers have often had to invest in it for the benefit of a relatively small marketplace, and alongside other access management technologies. Consolidation around a global standard will, over time, reduce administrative complexity and costs for suppliers and HE institutions alike.

ShibboLEAP

The ShibboLEAP Project ran from April 2005 to April 2006. It was funded by the JISC under the Core Middleware - Early Adopters Call [8], which was designed to stimulate the building and sharing of experience of the technical, cultural and administrative implications of the transition to Shibboleth technology, for the benefit of the wider JISC community. ShibboLEAP had seven partner institutions, as follows:

- Birkbeck
- Imperial College London
- King's College London
- London School of Economics and Political Science (LSE)
- Royal Holloway
- School of Oriental and African Studies (SOAS)
- UCL (University College London)

All the partners were founding members of the SHERPA-LEAP [9] Consortium, which has been developing institutional eprints repositories for University of London institutions since 2004, with the generous support of the Vice-Chancellor of the University of London. LEAP stands for London Eprints Access Project; hence the somewhat aesthetically displeasing project name for ShibboLEAP ('just a badly-chosen email subject-line that stuck', to quote the Project Manager).

ShibboLEAP had two overriding objectives. The first was to enable a full Shibboleth Identity Provider for all users at each of the seven partner institutions, using their existing directory and other infrastructure services wherever possible, and to document the process for the benefit of later adopters. The second objective was to enable the EPrints software [10], which was then in use by all the partners, as a Shibboleth Service Provider. (The content of the repositories is openly accessible, of course, but access restrictions apply to depositors, editors, and so on.)

The Project was led by the LSE, already capable of operating as a Shibboleth IdP [11] as a result of prior work on the SECURe [12] and PERSEUS [13] projects. The LSE Project Team undertook to help the other six partners through their IdP implementations. The partnership, though a natural grouping because of its existing consortial work, offered for these purposes a happy diversity in terms of size, institutional mission (from large, research-led institutions to continuing education specialists), and academic discipline. Understandably, the size and skills sets of IT support departments varied within the partnership, as did the partners' existing arrangements for identity management and resource access management: these included 'classic' Athens, Athens DA, several different implementations of LDAP (Lightweight Directory Access Protocol), and various data sources used to hold and maintain identity and role attribute information about their user base, together with various methods of populating those data sources. All the partners have strategic goals to work towards SSO. In all, the total 'population' of the partner institutions is estimated at over 150,000 users. We would stop short of describing the partnership as a microcosm of UK HE, but its breadth and overall size certainly made it a potentially interesting testbed.

Main Findings

Technologists thirsty for detail should read no further, but head straight for the project Web site [14], where case study reports from participating institutions are to be found. These include background information on institutional size and mission; descriptions of directory and authentication regimes, and other relevant infrastructural issues; accounts of IdP installation, including any difficulties encountered and copies of relevant configuration files; various expositions of Shibboleth, and other internal dissemination, for different audiences; and thoughts on future plans for Shibboleth integration and deployment. The six case studies are supplemented by an overarching project report, and accompanied by a separate report on the work carried out to enable Eprints2 as a Shibboleth Service Provider.

Technical detail aside, various points of agreement between partners emerged in the course of the project. In the first place, it is clear that the Shibboleth installation process remains complex; arguably too complex. All the partners had frequent recourse to the assistance of the LSE Project Team, whose achievement in having effected a 'solo' implementation of Shibboleth IdP seemed all the more impressive as time wore on. Secondly, while the project funded ample levels of staffing at each institution (a member of IT staff at 0.4 FTE, and a Co-ordinator at 0.1 FTE, the latter often from the Library), it became obvious that Shibboleth implementation requires skills, such as Tomcat expertise, which Unix Systems Administrators do not generally possess. A third obstacle to successful installation was the quality of the Shibboleth documentation, which is notoriously poor. Documentation, formal and informal, is available on the Web, but it lacks clarity, there is often inconsistency between sources, and it is very difficult for a novice installer to work out exactly which piece of documentation is most appropriate to his or her situation. A pre-packaged installer might be too much to ask for, but Shibboleth clearly lacks a single site for high-quality, consolidated documentation. (In the meantime, it is hoped that the efforts of

ShibboLEAP and the other Early Adopters to document their experiences will help to resolve the problem, rather than exacerbate it!)

It also became clear that the demands of Shibboleth on institutional directories can be problematic. Existing directories may not be fit for purpose, especially where *ad hoc* changes applied over time have created data inconsistencies - a commonplace scenario; and some directory products are not supportive of change to schemas. Some ShibboLEAP institutions took the opportunity to replace existing directories, or, at least, to clean them up, prior to taking on the Shibboleth installation. Moreover, directories are usually crucial to the day-to-day running of an institution's technical infrastructure, and engaging with the difficulty of installing the eduPerson object class on such an important service was a risk which few were prepared to take. However, where the required data was already present in the directory, partners were able to use the IdP software to transform the names and values of attributes released by the directory to those required by eduPerson, a simple and robust solution.

Next Steps: The ShibboLEAP Partners

Although the project delivered only pilot installations, and in spite of the installation issues itemised above, the partners all expect to continue to use Shibboleth in some way. The Shibboleth-Athens gateway is a particularly useful tool for connecting Shibboleth IdPs to third-party resources, and an increasing number of such resources are already independently Shibboleth-compliant [15]. All the partners are considering their position with regard to Athens resources, although those who are already using AthensDA are more phlegmatic than those still using 'classic' Athens. JORUM is an interesting case, since it is licensed only to staff from member institutions, and will therefore test Shibboleth attribute release and authentication to a finer level of detail than most other Athens resources. Metalib is also mentioned as a potential candidate for early 'Shibbolising'. One partner is already investigating authentication across institutions for a joint e-learning course, using Shibboleth authentication to Moodle.

All the partners, understandably, envisage taking a resource-by-resource approach to Shibboleth integration, rather than replacing all existing authentication regimes with a 'big bang'. Some concerns were expressed during the project about the resource implications of maintaining a Shibboleth IdP in addition to existing authentication systems, but as the number of Shibboleth-enabled Web resources continues to grow, and as the configuration overheads for such resources are expected rapidly to diminish as experience and expertise bed down in the community, it is likely that an institutional investment in Shibboleth will show increasingly high returns over time.

Next Steps: New Installers

Those planning a Shibboleth IdP implementation might begin by asking a few key questions. What is the institutional directory? Who owns it, how is it updated, and how are changes to it arranged? Does it contain all the information required for resources protected with Shibboleth? Should a new directory solution be considered?

How does the institution currently handle user account management? Are user credentials secure enough for SSO use outside the institution? Is a WebISO (Web Initial Sign-On) solution, such as Pubcookie, already in use? Where will the IdP be installed, on what type of machine, and how will it be connected to the institutional directory? It should be borne in mind that network account technical staff, directory administration technical staff, firewall and security staff, Web staff with Tomcat skills, Library IT staff (and other Library staff with knowledge of relevant external electronic resources), the institutional Athens administrator(s), and all their managers, will need to be involved in the process. From here onwards, unfortunately, the absence of a suite of 'oven-ready' installation models, and the quality of the Shibboleth documentation, mean that things can, occasionally, get a bit difficult. However, the good news is that Shibboleth installation is an increasingly well-trodden path in the UK. The ShibboLEAP outputs, the work of the other Early Adopters, and the JISC's Support Service [16] will all be able to offer some guidance through the complexities of Shibboleth installation and its equally complex documentation.

Conclusion

Shibboleth is being adopted by educational institutions worldwide, including institutions in the UK, in support of federated access management. Shibboleth is not an 'all-in-one' access management solution; but, as a component in an institutional SSO environment, it offers tremendous enabling possibilities. It is architecturally complex, and can be difficult to install, but a number of UK institutions, including the seven ShibboLEAP partners, have already blazed a trail through the installation jungle. This growing body of experience in Shibboleth deployment in the UK will benefit new adopters of this important new technology, which has the potential to help institutions to take major steps forward towards achieving Single Sign-On to secure internal and external Web resources.

Afterword: you say tomato...

The word 'Shibboleth', in this context, has a distinguished history. The Bible (Judges 12, v.1-6) records how the Ephraimites, fleeing Gilead after an unsuccessful attack, were cut off at the river Jordan by Gileadites. Those attempting to cross the river were asked to say 'Shibboleth'. The unfortunate Ephraimites, who pronounced 'sh' as 'si', said 'Sibboleth', and 42,000 of them were duly slaughtered. While this was undeniably a triumph of access management, we are confident that the implications of Shibboleth for the 150,000-plus members of the ShibboLEAP institutions will prove to be rather more uplifting.

Acknowledgements

Thanks to John Paschoud and Simon McLeish of the LSE, both for reading a draft of this article, and for their indefatigable technical assistance to the ShibboLEAP partners during the project.

References

1. McLeish, S. "Installing Shibboleth", Ariadne 43, April 2005 http://www.ariadne.ac.uk/issue43/mcleish/
2. JISC "Connecting People and Resources; Briefing Paper version 2", March 2006 http://www.jisc.ac.uk/index.cfm?name=pub_shibboleth
3. MATU - Middleware Assisted Take-up Service "What is Shibboleth?" http://www.matu.ac.uk/shibboleth_intro.html
4. Internet2 Shibboleth Web site http://shibboleth.internet2.edu/
5. eduPerson Object Class Web site http://www.educause.edu/content.asp?PAGE_ID=949&bhcp=1
6. JISC Core Middleware: Technology Development Programme; JISC Core Middleware Infrastructure Programme http://www.jisc.ac.uk/index.cfm?name=programme_middleware
7. UK Access Management Federation Web Site http://www.ukfederation.org.uk/
8. JISC Early Adopter Programme: http://www.jisc.ac.uk/index.cfm?name=programme_cminfrastructure
9. SHERPA-LEAP Web site http://www.sherpa-leap.ac.uk
10. EPrints Web site http://www.eprints.org/
11. Shibboleth at LSE Web site http://www.angel.ac.uk/ShibbolethAtLSE/
12. SECURe Project Web site http://www.angel.ac.uk/SECURe/
13. PERSEUS Project Web site http://www.angel.ac.uk/PERSEUS/
14. ShibboLEAP Web site http://www.angel.ac.uk/ShibboLEAP/
15. Index of Shibboleth-Enabled Applications and Services http://shibboleth.internet2.edu/seas.html
16. MATU- Middleware Assisted Take-up Service http://www.matu.ac.uk/

Author Details

Martin Moyle
Team Leader, Science; Project Manager, SHERPA-LEAP
University College London

Email: m.moyle@ucl.ac.uk
Web site : http://www.ucl.ac.uk/library