

SHIBBOLETH: THE LAST PIECE OF THE JIGSAW IMPLEMENTING AUTHENTICATION FOR DOCUMENT DELIVERY AT RMIT UNIVERSITY

J. WILSON

*Document Services, University Library, RMIT University
Melbourne, VIC.
julie.wilson@rmit.edu.au*

ABSTRACT

Authenticating users across organizational boundaries presents institutions such as libraries with security and other challenges that prevent seamless access and impede customer service. RMIT University Library was amongst the first libraries in Australia to implement stand-alone interlibrary loan management systems. At RMIT the VDX system was installed to run the Library's Document Delivery service, operating in a shared hosted environment, managed and maintained offsite. Single sign-on was a goal in place for all services and applications at the University but for the VDX service, operating remotely from the institution, authentication ensuring protection of users' privacy was a real issue. After a detailed investigation Shibboleth was identified as the authentication system of choice. This paper examines the Shibboleth implementation from conception and planning to set up, testing and production. Difficulties encountered are outlined as are other Shibboleth implementations and library ILL authentication strategies. It is of note that RMIT University's implementation of the full features of Shibboleth is a first for VDX worldwide.

1. BACKGROUND

Shibboleth is now used by a growing community of businesses, government agencies, research communities and educational institutions world-wide (Morgan 2004:17). Shibboleth's federating software and unique capability to manage identity management for secure access to services and resources is solving multi organisational web authentication problems and enabling a new generation of applications and services. RMIT has applied this web based, open source software system to improve service for library customers, giving instantaneous access to the Library's Document Delivery Service. 'The dream of the vast, authoritative, easy-to-use virtual library' (Misekell 2005: 315) is not quite here, but almost. Shibboleth is a complex piece of software requiring technological infrastructure and experienced system administration. Fortunately these

elements came together providing document delivery requesting anywhere, anytime, anyplace.

2. THE UNIVERSITY

RMIT University is one of Australia's original and leading educational institutions. With its heart in the City of Melbourne, RMIT has an international reputation for excellence in work-relevant education, high quality research and engagement with the needs of industry and the community. With more than seventy thousand students studying at RMIT campuses in Melbourne and regional Victoria, in Vietnam, online, by distance education and at partner institutions around the world, the University is one of the largest in the country. RMIT's increased focus on research is outlined in the 2010 RMIT Strategic Plan 'Designing the Future' (RMIT 2010). It aims to be one of Australia's top research universities and to develop focused areas of excellence in research and scholarship that reflect global engagement with industries and communities. With such a diverse student community requiring service delivery on campus, off campus to distance learners in Australia and offshore, together in a climate of increasing emphasis on research needs, the document delivery service becomes very important in meeting the information needs of students and staff.

3. ORGANIZATION OF DOCUMENT DELIVERY

Document Delivery provides centralised services to all eligible staff and students. At RMIT this includes all staff and students including postgraduates, undergraduates, graduate diploma, TAFE, and Open University Australia students. The service also includes a cross-campus copying service and a postal loans service to remote students. In 2009 ten thousand two hundred and seventy-eight requests were received from RMIT staff and students. Four hundred and fifty-seven items were posted to remote students and nine hundred and seventy-nine cross campus copies delivered. A total of two thousand nine hundred and thirty-eight copies and loans were supplied to other libraries.

RMIT University Library, together with a number of other universities, was amongst the first libraries in Australia to implement a stand-alone interlibrary loan (ILL) management system. OCLC's VDX (Virtual Document eXchange) software was installed in 2001 to support the Library's Document Delivery Service. The software was developed as part of the LIDDAS (Local Interlending and Document Delivery Administration System) project, a joint initiative of the AVCC (Australian Vice Chancellors Committee) and the National Library of Australia. RMIT is a member of CLIC (CAVAL Interlibrary Consortium) comprising six academic libraries using the VDX software which is delivered through a shared hosted system managed and maintained offsite. 'CLIC was established to provide configuration and support to member libraries, to share expertise and to assist libraries in managing a complex automation product' (Jilovsky 2006:1). The system consists of a single database with individual institutional views and a separately branded web interface for each institution. 'OCLC manages the CLIC Hosted Service ie the

hardware, operating systems and networks with the hardware physically located at RMIT '(CAVAL 2009:3).

In 2011 a different business model will be put into place with CLIC to be disbanded following a CAVAL initiated business review. In 2011 CLIC VDX libraries will be moving to a hosted service model directly managed by the vendor and to be hosted offshore.

4. MANUAL CUSTOMER REGISTRATION

From the implementation of VDX in 2001 RMIT opted for a manual registration process for users, deciding against the automatic loading of the University's voluminous student and staff records. As document delivery users at RMIT were assessed to be a small percentage of the total student and staff population it was not considered efficient to load all records into the database. RMIT was also in the process of seeking a new integrated library management system (ILMS), and the work required to revise the patron load process for a new system would have been a low priority through any implementation.

A number of problems quickly became apparent with the manual registration process:

- Registration was not instantaneous; the promised turnaround time for delivery of login details was twenty four hours. Users often reported dissatisfaction with the delay in using the system.
- The online registration form required manual data input. Although simple enough to complete, the tendency for clients to use the default entries in drop down menus (for status and campus) meant document delivery staff had to recheck university directories for correct information.
- Considerable library staff time was spent in managing manual registrations.
- Login details for VDX were different from the university's universal Novell Directory Service (NDS) login. The management of multiple logins was a frustration commented upon by many users.
- User details were not automatically updated or synchronized with any other RMIT system. User information over time became out of date and inaccurate.

An authentication solution was required to improve customer access and reduce workloads for staff.

5. POTENTIAL AUTHENTICATION SOLUTIONS

Authentication solutions employed by other libraries operating International Organisation for Standardisation (ISO) ILL systems vary. Victoria University (Melbourne) and Macquarie University regularly load patron data from their ILMS into the VDX database. Overseas, the California Digital Library which runs a VDX installation for ten University of California campuses also utilises data from an ILMS. The user specifies their home

campus and the system queries the appropriate local system. In most cases the authentication interaction is with the university library system, not the university authentication system. Most University of California campuses use the Innovative Interfaces ILLMS and the Patron Initiated Request (PIR) uses the patron application programming interface (API). In OCUL's (Ontario Council of University Libraries) VDX RACER (Rapid Access to Collections by Electronic Requesting) system users create an account which is stored in the VDX Oracle database and associated with their particular school.

Some innovation has been demonstrated in Australia with regard to authentication for VDX. The University of Southern Queensland (USQ) was the first Australian library to implement a Lightweight Directory Access Protocol (LDAP) solution in 2003 followed by the University of Western Australia (UWA). USQ is a stand alone implementation whilst UWA is part of the Western Australian Group of University Librarians (WAGUL) Consortium. La Trobe University, a member of CLIC, also moved to an LDAP solution several years ago, and in this model users authenticate within the La Trobe network.

In Australia CSIRO document delivery service uses the Relais system. Relais doesn't support LDAP, however it does support direct 'login links'. A script on the library side which works with the institution's LDAP directory allows users to login to the Relais request form.

For university ILL services operating in a hosted environment, such as the VDX system at RMIT, there was at this juncture no method of securely transmitting user data to an offsite server.

6. RMIT'S REQUIREMENTS FOR AUTHENTICATION

The use of a single sign-on , one universal login for all RMIT systems, was identified as early as 2005 as an important service improvement for the Document Delivery Service and one that aligned with the University's planning for all services and applications. Discussions on authentication for VDX commenced in 2006 with the three stakeholders: RMIT Information Technology Services (ITS), CAVAL and the system vendor OCLC (then OCLC Pica). In the first instance an LDAP solution was identified as the favored option and work began on system setup and configuration. Significant security risks were subsequently identified and reported on by RMIT's Security Analyst (Coxhill 2006). The concerns raised focused on the exposure of RMIT's network to the internet and the potential risk of unauthorised access to personal information and to other RMIT corporate systems. In the hypothetical scenario an external (non-RMIT) server on a non-trusted network on the internet would be allowed to execute an LDAP lookup against RMIT's Novell Directory Services eDirectory (a database of all the University's network resources) without authentication or restrictions on what could be looked up or retrieved. RMIT data could be compromised and the LDAP solution was shelved.

It was not until early 2007 that Shibboleth, through a detailed investigation by RMIT ITS, was identified as the favored authentication solution. Shibboleth would provide authenticated access in a secure manner, perfect for the CLIC hosted service set up. This was new territory for OCLC. Shibboleth had been in use in the United States by OCLC since 2003 but the available functionality only allowed for searching remote databases through the web user interface (ZPortal). No patron data was transferred and stored in ZPortal, configuration that was critical to RMIT's proposal. OCLC committed to the development of the product to provide the functionality required by RMIT. The project could proceed.

7. WHAT IS SHIBBOLETH?

Shibboleth is a standards based open source software package for web single sign-on across or within organisational boundaries. It allows organisations to exchange information about their users in a secure, privacy preserving manner. Internet2, supported by the National Science Foundation (US), established its Middleware Initiative in 1999. The Shibboleth Project used an open design and development process with contributions from information technology specialists, corporate partners and interested others world-wide to develop a solution to make secure inter-institutional services possible (Perry 2006:12). The Shibboleth system was the result. Implementations are world-wide and growing. In essence Shibboleth:

- enables single sign-on addressing on and off campus web authentication;
- controls release of user information to authorise actions or customize user's experience;
- protects the privacy of personal information;
- supports multi-organisational federations such as InCommon Federation (US), UK Federation and MAMS (Meta Access Management System) (Australia);

8. HOW SHIBBOLETH WORKS

There are three parts to the Shibboleth system:

1. Identity Provider (IdP) - the software run by an organisation with users wishing to access a restricted service.
2. Service Provider (SP) - the software run by the provider managing the restricted service.
3. Trust Federation - a trust framework (policy and technical) that connects Identity Providers and Services Providers.

The Shibboleth website introduction notes: 'A user authenticates with his or her organisational credentials. The organisation passes the minimal identity information necessary to the service manager to enable an authorization decision. Shibboleth leverages the organisation's identity and access management system so that the individual's relationship with the institution determines access rights to services that are hosted both off and on campus' (Shibboleth 2010). When a user at one institution tries

to access a resource at another, Shibboleth sends the attributes about the user to the remote destination rather than making the user log in to that destination (EDUCAUSE 2010). The remote destination decides whether or not to grant access. Shibboleth preserves the user's identity by only releasing necessary information, information which does not disclose identity.

9. ADVANTAGES OF SHIBBOLETH FOR RMIT

Shibboleth authentication would solve problems identified with the manual registration and login process. The following advantages of authentication were identified:

- Customers are registered instantaneously and can use the system immediately
- Reduced workload for staff – no more manual registration
- Single sign-on; customer uses current RMIT NDS login and password
- Familiar login page – same as for other services
- RMIT decides what personal detail 'attributes' to release
- Password details never leave RMIT network
- Customer details updated on 'live' lookup

10. SHIBBOLETH IMPLEMENTATION

The entire configuration, testing and implementation process will be very familiar to those that have been involved in a software implementation. Here are outlined detailed steps to demonstrate requirements unique to the Shibboleth set-up.

As Shibboleth relies on the release of user attributes to make proper authorisation decisions the first step in the implementation process was a decision on the attributes required for VDX. 'The Shibboleth IdP software plugs into existing institutional identity management and user information services' (Morgan 2004:13) In RMIT's VDX implementation this is the LDAP Directory and the Voyager ILMS, although it is largely the Voyager database which became the definitive source of information for VDX for everyone 'at RMIT. To further explain, the "I" (address) attribute in the NDS is the RMIT building and floor identification; however the more appropriate details for VDX are the user's postal address which comes from Voyager. All attributes used were required to be registered (as part of the trust federation policy and technical framework) - OCLC was delegated the namespace urn:mace:olc.org by (Middleware Architecture Committee for Education (MACE Group) Attributes required for VDX were subsequently defined as: full name; ID - staff /student number; email; postal address - street address, locality, state, postcode, country, telephone; school/department; campus; VDX Category; VDX Group.

Mapping attributes into VDX fields from information in the Voyager database was the most complicated part of the implementation, particularly the VDX Category and VDX Group fields. There are thirty four user categories in RMIT's VDX configuration, each

user category specifying a combination of status (staff, postgraduate, undergraduate, TAFE) and location (remote, offshore or on campus at Swanston, Bundoora West, Bundoora East, Business, Brunswick, Carlton Library sites). In VDX each user category is matched to a user group which specifies database searching permissions. A customer's user category reflects location and status, determining the location for loans to be delivered and dictating the level of service received.

'Remote' status was of particular interest. This was not a status recognized by the University on enrolment but rather a status allocated internally for users who resided interstate or outside Victorian Public Transport Metropolitan Zone 2. Remote customers were eligible for the Postal Loans Service. The mapping is configured to allocate a remote category if the user postcode is outside 3210. Very neat and it worked!

The allocation of the Offshore user category was calculated from the ILMS Patron Statistical Category '135' (Offshore). Patron Statistical Category 144 (International) mapped into an on-site user category reflecting the local residence of this group of users. The full documentation supporting the attribute mapping - 'Shibboleth Attributes', 'Shibboleth Data Connectors' and 'Group/Category Matrix' - is available from the Library

At RMIT configuration and testing involved cooperation between two RMIT Information Technology Services teams – the Unix team who supported Shibboleth and the Application Support team for VDX – OCLC, CAVAL and the Library. Testing involved installation of the configuration on the test CLIC system in conjunction with the Shibboleth test server (Shib dvts). Initial testing was on Shibboleth version 1.3. Shibboleth was upgraded to version 2 in May 2010. It proved quite challenging to have all parties and systems available when required. Users from all student and academic staff groups and locations, including offshore and remote users, were asked to test the system. All results were documented (login time, date, browser, location) and updated regularly over the two month testing period. Several issues were identified: attribute mapping inconsistencies, session creation failures revolving around the domain name system (DNS), firewall problems, metadata file validity and stress on the Voyager database. All issues were resolved relatively quickly.

A Shibboleth tasks document allocating responsibilities and time lines was prepared as the project reached the final stages. The document ensured that tasks, mostly dependant on the successful completion of the previous work, were all completed on time and according to schedule. Web site documentation such as user guides and frequently asked questions (FAQs) were prepared in advance ready for implementation.

11. GO LIVE

The test environment was replicated on the OCLC and RMIT Shibboleth production servers and testing was completed successfully. Notification of the imminent login changes were posted on Library web pages and forwarded in newsletters and other publicity. The look and feel of the new login pages for Document Delivery matched

almost exactly existing Library login web pages. The one difference for users logging in with Shibboleth was an additional web page, a notification page from Shibboleth alerting users that the system was redirecting from Shibboleth to the application, however it is standard practice to notify customers of redirection in this way. It was anticipated that the changeover would cause little confusion amongst existing users and be seamless for new customers. A trouble shooting document 'What might happen?' was prepared for library staff listing anticipated user queries and problems with suggested solutions.

The changeover, well ahead of commencement of the academic year, was timed to catch a new intake of users. The CLIC system was brought down briefly to install configuration for the customer web interface. Shibboleth for VDX went live on January 8th 2010.

12. PROBLEMS ENCOUNTERED IN PRODUCTION

The transition went very smoothly with few queries received or login problems noted. The majority of users took to the new login process with ease. Four months into implementing Shibboleth authentication for VDX there have been few issues identified. Certainly no system issues, and requests placed by staff and students have increased.

Workarounds were required to manage local workflows although the difficulties have arisen from lack of RMIT held user data, not any system failure. Incomplete user data in Voyager meant manual updating of the system for the following groups of users:

- RMIT casual staff without Voyager records;
- Open University Australia students who are not identified with an RMIT school/department;
- International students with overseas addresses (where local address had not been updated on RMIT enrolment records).

13. OTHER SHIBBOLETH IMPLEMENTATIONS AT RMIT

When Shibboleth was first identified as the chosen authentication software for VDX in 2007, there were no other implementations at the University. There are now three others. Morgan (2004:14) noted the use of Shibboleth to solve inter-institutional web access problems for research collaborations and to run university outsourced employee administrative services.

Such uses are now in place at RMIT:

- **Pebblepad** is an ePortfolio tool, currently being trialed at RMIT which enhances opportunities for students to record evidence of formal and informal learning.
- **Trobexis** is the RMIT University online travel management system
- **InfoScholar**, accessed through Queensland University of Technology (QUT) Blackboard, is a generic online tutorial for research students provided as a component of the eGrad School Australia initiative.

It is anticipated that there will be additional Shibboleth services in the future.

14. CONCLUSION

The advantages of Shibboleth authentication for VDX outlined earlier in this paper have been realized. Customers have instantaneous access to lodge and track document delivery requests, staff workload has been considerably reduced and single sign-on has eliminated login confusion. It is of note that RMIT's implementation of the full features of Shibboleth is a first for VDX worldwide.

If the measure of the worth of any development lies in its immediate success then the implementation of Shibboleth speaks for itself. It is difficult to remember what life was like before Shibboleth. Document Delivery staff and users, both academic staff and students, were immediately appreciative of the improved service on offer and with single sign-on now in place there is time to look to other developments.

REFERENCES

- [1] Mikesell, B.L. (2005) 'Anything, Anytime, Anywhere' Journal of Library Administration Vol 41 no 1 p 315-326
- [2] Morgan, R.L. et al (2004) 'Federated Security: The Shibboleth Project' EDUCAUSE Quarterly Vol 27 no 4 p 12 – 17
- [3] 'RMIT 2010: Designing the future' (2010) Retrieved May 15 2010 and located at RMIT University Web site <http://www.rmit.edu.au>
- [4] Jilovsky, C, Pearson, K. & Wilson, J(2006). 'CLIC@CLICL06: a consortial success story' ALIA Biennial Conference Click06 Perth 19-22 September Perth. Retrieved May 15 2010 and located at http://conferences.alia.org.au/alia2006/conference_papers.html
- [5] CAVAL Interlibrary Consortium (CLIC) CLIC Review 2008-2009 (2009) CAVAL Ltd. Internal document. Unpublished p 3
- [6] Coxhill, A (2006) 'Application Security Analysis OCLC Pica VDX Inter-Library Loans Software LDAP Authentication Risks' RMIT University Internal Document. Unpublished 22 Aug
- [7] OCLC ILLiad Web site Viewed May 15 2010 and located at <http://www.oclc.org/illiad/>
Perry, S. (2004)
- [8] Rigda, C (2010) 'Implementing ILLIAD: a Successful Collaboration of the Systems and Interlibrary Loan Departments in an Academic Library' Journal of Interlibrary Loan, Document Delivery & Electronic Reserve Vol 20 no 2 p 77- 82
- [9] Perry, S. 'InCommon; Watch This Space!' EDUCAUSE Review Vol 41 no 5 p 12-13
- [10] Shibboleth Project Web site Viewed 15 May 2010 and located at <http://shibboleth.internet2.edu/>
Shibboleth Information Sheet Overview.pdf <http://shibboleth.internet2.edu/>

Attribute Naming Website <https://spaces.internet2.edu/display/SHIB/AttributeNaming>

[11] EDUCAUSE Web site Viewed 15 May 2010 and located at <http://www.educause.edu/Resources/Browse/Shibboleth/30477>

[12] 'Shibboleth: standards-based single sign-on for the web' Viewed 15 May 2010 and located at <http://nsit.uchicago.edu/internal/insite/2009/january/shibboleth.shtml>