CASE STUDY

February 2009

The UK Access Management Federation
FOR EDUCATION AND RESEARCH

University of Dundee

Interviewed: Andy Swiffin, Senior Network Specialist

BACKGROUND

The origins of the University of Dundee date back to 1881 when University College, Dundee was founded. Today it is one of the UK's leading academic establishments, internationally recognised for its expertise across a range of disciplines. The university has more than doubled in size since 1994 and is currently home to some 3,000 staff and 18,000 students, 20 per cent of whom are distance learners accessing university resources remotely.

TIMELINE

The university took its first steps towards federated services in January 2007 when Andy Swiffin, the university's senior network specialist, recognised the potential of federated access management for Dundee. After studying the options, he initiated a project to implement Shibboleth, carrying out the first trial deployments in June 2007. He installed the first production server in December of the same year and the system went live on 17 February 2008. It was unexpectedly prompted by the launch of Microsoft's DreamSpark initiative, which generated a sudden need for students to provide a UK federation login to qualify for free software.

OBJECTIVES

Dundee's overall aim was to achieve unified authentication for access to external resources. Shibboleth was judged to be the obvious choice, providing single sign-on (SSO) facilities for applications and a common authentication mechanism retaining the same appearance regardless of the resource being accessed. The benefits of the common interface and prevention of the unnecessary release of personal information into applications were the main drivers behind the project.

STARTING OUT

Andy describes his first forays into the federated environment as following very much a trial-and-error approach. 'The first six months represented a huge learning curve,' he says. 'Finding where to start was the real problem – it's a lot easier now. When I began there was only sketchy documentation available – lots of useful documents about the UK federation and deploying an IdP, but not one that pulled it all together with a checklist of things to do; nothing to say "start here".'

In terms of application choices, Andy was aware that both Guanxi and Shibboleth would interoperate with the UK federation, but the recommendation from the UK federation sites was to work with Shibboleth because of the existing body of knowledge.

In the absence of a definitive guide to deploying federated access management, Andy began by gathering together all of the reports that he could find from the Middleware Assisted Take-Up (MATU) Service project, regarded as the pioneer of Shibboleth. He also studied the Internet2 web site (although he says initially the documentation there was not as good as it is now) and tried to fill in the gaps in his knowledge by Googling 'Shibboleth' and collecting and studying pieces of other peoples' configurations. He found the JISC Shibboleth mail list tremendously useful in this regard, as the community there was willing to share configuration files, of which there was a severe deficit in the early days (and to some extent still is now, he points out, as there's no central repository for them).

THE BENEFITS OF LOCAL SUPPORT

Having taken the plunge to adopt Shibboleth and researched the subject as fully as possible, Andy attended his first JISC/UK federation meeting in Birmingham during May 2007. Discovering that federation members gathered only at national events like this, he returned to Scotland to set up the Scottish Federated Access Management Forum, better known today simply as 'McShib', to provide a local support network for Scottish members. He received a huge amount of support from the JISC Scotland North & East Regional Support Centre (RSC) in Edinburgh – it took on all of the administration involved – and the inaugural meeting of McShib members took place in August 2007.

One of the presentations at that first meeting was entitled 'From Zero to a Shib IdP in 30 Minutes', authored by Andy Swiffin and documenting the results of his initial trial deployments.

A major concern for many sites in the early stages was the use of Linux as the main platform for running a Shibboleth IdP – Windows was initially more of a niche choice but that isn't necessarily the case now (more about that later). Although Andy admits that an IT background is an advantage: he also insists that one doesn't have to be a Linux expert to set up an IdP – when he started out, he hadn't used Linux much before and had no real experience of deploying web services like Tomcat, Apache and OpenSSL.

'It isn't that difficult, even without Linux skills,' he says, 'and anyone who is familiar with the command (DOS) prompt in Windows should find it fairly intuitive.'

THE INITIAL DEPLOYMENT

The first IdP deployment at Dundee took place in mid-2007 – not as a production server but to figure out how it could all work together. Andy then attended the pilot of a JISC-funded Netskills 'core essentials' course in July 2007, following which it took him about two weeks working from scratch to develop an IdP that worked against the 'TestShib' service (not the UK federation).

MIGRATION FROM ATHENS

Up to this point, access management had been the domain of the library, which used Athens DA with the default authentication via the university's LDAPS service. Dundee was fortunate in already having a well-established identity management strategy, including a fully populated directory of all staff and students against which Athens DA authenticated. As soon as Andy became involved in access management he made contact with the library IT support personnel. 'We've established a very good working relationship,' he says, 'which has been instrumental in moving us forward into the federated world. All I'm doing is providing the mechanism: they're the people actually making it happen.'

It was initially expected that access to Athens would continue via the Athens-Shibboleth gateway, planned for deployment at the end of February 2008, using Shibboleth as the authenticator to access Athens-based resources. However, the announcement that funding for Athens DA would cease on 1 August 2008 prompted the university to opt to use Shibboleth exclusively rather than subscribe to Athens. Andy believes that this was definitely the right decision for Dundee, as in his experience many of those who did sign up to Athens in the interim are now finding it difficult to make the transition to the UK federation – they have to run two authentication mechanisms, which while not difficult technically presents a user education issue, explaining which logins are required for which resources. This creates unnecessary confusion, whereas

Dundee's users see the same login screen for all resources and the library is set up to take the user seamlessly to the correct authentication mechanism.

When Athens access was discontinued, it was known that there would be a delay before some service providers became available. As a stopgap measure Andy Swiffin and Matthew Phillips from the library's IT support group implemented EZProxy to access these, which proved both straightforward and effective. As a result, they have retained EZProxy and configured it to authenticate through Shibboleth. It was deployed in this way to give users a common experience such that they see the same login pages whether accessing resources through EZProxy or directly through the UK federation.

THE LEARNING CURVE CONTINUES

As one of the first adopters, Andy set about making his experiences with Shibboleth more accessible to others – he describes himself as being 'passionate about making it possible for people to do things'. As mentioned earlier, when he started out, he had no real experience of deploying web services like Tomcat, Apache, Open SSL and hadn't even used Linux much before. While deployment in Linux seemed the obvious thing to do at the time, since then – in order to help out some local colleges – he has also deployed it in Windows. His record for getting a Windows server installed with Apache, Tomcat, Shibboleth and all the pieces of configuration needed to work as an IdP is about ten minutes!

It is probably now easier in Windows XP than in Linux, he says; the advantage lies in taking sections of configuration from his own installation and dropping them in to a new system. However, this only works if the setups are the same – for example, one college using Microsoft Active Directory (AD) in place of eDirectory, Fedora, OpenLDAP or similar proved a lot more tricky, but he learnt a lot from that and now knows how to make AD do the job.

CHALLENGES ENCOUNTERED AND OVERCOME

The main problem that plagued Andy's first attempts at deploying an IdP stemmed from the absence of any clear indication of which versions of applications to use. This caused all manner of conflicts such as particular versions of OpenSSL not working with particular versions of Apache. Trial, error and persistence eventually rectified this situation, but for those now deploying on Windows, Apache for Windows is bundled with the correct version of OpenSSL to be compatible, making for a much more straightforward install. Simply obtain Java, Tomcat and Apache bundled with OpenSSL and it all just works together. With hindsight, Andy says he would have considered deployment in Windows as it seems to be even simpler than in Linux.

TIME AND COSTS INVOLVED

Andy estimates that in total it took a year of his time to get from concept to a production Shibboleth IdP installed and working, although he qualifies this by pointing out that relatively few of today's available resource existed when he embarked on the project, and work on Shibboleth was very much slotted around his other tasks and activities. In terms of server support, Dundee has one well-specified Linux IdP server costing around £7000 – the other server is a virtual machine and the infrastructure providing automatic failover was already in place for other services. Content switching isn't yet deployed, and to date hasn't been required. The only failure took place recently while Andy was on holiday – this was the first and only time that the IdP had stopped working since installation, and simply required Matthew Phillips, in Andy's absence, to stop and restart Tomcat and Apache. Andy explains that the failure was not attributable to Shibboleth but to human error on his part: he had overlooked the requirement to increase the Java cache size from its default value, but this has since been rectified and the fault has not recurred. The setup can therefore be considered very reliable, incurring no expensive ongoing maintenance costs.

HINTS AND TIPS – ANDY SWIFFIN'S ADVICE FOR NEW DEPLOYERS

1. **Deploy straight into the UK federation:** From his experiences in setting up Shibboleth IdPs, Andy Swiffin would now advise anyone to bypass TestShib and instead obtain a certificate from JANET, then deploy direct into the UK federation and test against real suppliers. It seems that some sites struggle to make progress with TestShib as it can be more difficult to get it to work, which may be due to the fact that it is now focused on Shibboleth2 deployments. This often leads users to mistakenly doubt the accuracy of their efforts, whereas deploying straight into the UK federation is actually quite straightforward if one knows which steps to take.

2. **Install EZProxy:** Another good option is to install EZProxy; deployed locally and with Shibboleth authentication in place it will reveal the attributes being released to it, so testing can be conducted in-house. However, this may prove more difficult to set up initially, because effectively one must set up an IdP and a service provider and register both into the UK federation.

3. **Choose whether Linux or Windows is best for you:** Initially, Andy would have suggested working in Linux but nowadays deploying an IdP is perhaps more easily achieved in Windows. The JISC-sponsored Automated Windows Installer further simplifies the installation of an IdP that uses Active Directory.

4. **Get your identities sorted out:** At the earliest opportunity, make sure you have identities to authenticate against. Without them you will have invested a lot of time and effort with few results to show for it!

5. **Get connected:** Join the JISC Shibboleth mailing list – there's no elitism there, everyone is very helpful. Find local mentors – JANET and JISC can help you with this. Consider setting up a local group or meetings – McShib has been instrumental in helping people in Scotland but the situation seems more fragmented south of the border.

6. **Beg, borrow and steal ideas:** In the absence of a central configuration repository, take every opportunity to look at what other people have done and copy what works for your needs.

7. **Be positive:** Remember that it's not as complicated as it first looks. Andy started out with minimal experience of Linux and deployment of web services, but with experience and practice it has become second nature.

WHERE TO NEXT FOR THE UNIVERSITY OF DUNDEE?

Dundee now has a unified authentication setup that leads users to the required resource by a common, familiar screen, with the advantage that the university doesn't have to pay for Athens. 'To date, we've federated everything here that can be federated,' says Andy. 'Externally we go through the UK federation for everything that's accessible, while internally we use Ex Libris MetaLib – and we know that's "Shibbolisable" too. Matthew Phillips in the library is quite keen to bring it into the fold, so that will probably become our next project.'

In the longer term, as other applications present themselves as candidates for federation, Andy says that while he can explain the benefits of federation and exercise persuasion, it will also be dependent on the people involved with each application having the will and the skill sets needed to carry it through.

Our thanks to Andy Swiffin for agreeing to be interviewed for this case study.

APPENDIX TUTORIAL
THREE STEPS TO A WORKING IDP, THE ANDY SWIFFIN WAY

1. Gather the applications for installation

Something that Andy cites as having been a major problem during his first attempts at deployment was the inconsistent behaviour of various versions of software when trying to make them interact with each other. Having established a working combination via a good deal of trial and error testing, the first requirement to deploy a Shibboleth IdP is to assemble the correct versions of all of the components needed. The recommended versions of the main 'building blocks' are as follows:

- **Shibboleth** – the latest supported version.[1]
- **Java 1.5 (JDK 5.0)** – Shibboleth is written in Java
- **Tomcat 5.5.25** – a Java servlet container that executes JavaScript
- **Apache 2.2.6 with SSL support** – a web server to act as a front end for Tomcat
- **OpenSSL 0.9.8g** – to provide Secure Socket Layer functions for encrypted transactions
- **mod_proxy_ajp** – this routes web requests incoming to Apache through into Tomcat.

Online sources of the applications are listed at the end. Note that while Shibboleth 2.1 is the current version of the Shibboleth IdP, 1.3 still seems to be the recommended version for use in the UK federation. However, from the end of March 2009 the UK federation will be in a position to support Shibboleth 2.1 and this will be the supported version.[2]

Although these steps refer to the deployment of a version 1.3 IdP, many of the principles will still hold. Deployment of 2.1 may be simpler as, for instance, authentication modules come bundled with it.

2. Compile the applications

Once the applications are assembled, the next stage is to compile them. While compiling applications may sound complex, it mainly consists of specifying options – the actual compiler is never seen. Normally, only at the first ./config or ./configure stage are any options required.

OpenSSL

First to be compiled is OpenSSL. After extracting the files, the installation process is a matter of opening a command prompt and typing a series of simple commands, in this case './config shared' (so that 'shared' libraries are also installed) followed by 'make' then 'make install' – other applications follow a similar syntax. Most of the time it isn't necessary to understand these instructions, nor the scripts that appear as they are executed.

Apache

Next, the process is repeated with Apache. Here, options need to be specified at the configuration stage but detailed instructions are included on the Apache web site to guide the installer. The most important options are --enable-ssl=shared --with-ssl=/usr/local/ssl; note however that this latter path will be dependent on where OpenSSL is actually installed. The default location has been used in this example and Andy's experience is that accepting the defaults often minimises later problems, although expert installers may of course choose to deviate from this.

Some environment variables for Apache, Java, Shibboleth and Tomcat then need to be set up in a file to run every time on startup. From the web server's bin directory, the startup string is run to start the server, and

[1] Please check the federation technical recommendations for the supported version of the software.

[2] See http://www.ukfederation.org/content/News/2009-02-17-1p3-end-of-life

the site secured by setting up SSL and certificates on port 443 of the server – a useful online guide that helped with this is referenced below. During the development phase Andy attempted to deploy various combinations of Apache versions and OpenSSL versions with varying degrees of success – for instance, the combination of Apache 2.0x and OpenSSL 0.9.8x was found to be problematic. It has since been found useful to use the 'ldd' command (e.g. ldd httpd) to verify exactly ***which*** shared libraries Apache thinks it will be loading and ***where*** it will find them – occasional surprises were encountered with Apache using a version of libcrypt.so that Andy had no recollection of installing! The 'find' command (e.g. find / -name libcrypt.*) is useful for identifying just how many copies of the file there are and where they are located.

Tomcat, Java and mod_proxy_ajp

Installing Tomcat is the next stage – as it's written in Java, this doesn't need compiling but simply placing in the correct folder. Following that, Java is installed in a similar manner. Tomcat can then be started and a web browser directed to port 8080 on the server to check that it works. With Apache and Tomcat both working, they need to be connected with the Apache Tomcat Connector; this is now included with Apache 2.2 and so does not need to be compiled separately. Further configuration entries are required at this point to ensure that Apache sends Java Server Pages (JSPs) and incoming URLs referring to Shibboleth on to Tomcat to deal with. Upon restarting both components, successful configuration is indicated by the test page now appearing on port 80 rather than port 8080.

3. Install the Shibboleth Identity Provider

The Internet2 site includes installation notes for the Shibboleth IdP – the process is a little different from the other applications, using an installer called 'ANT'. All of the default install items are selected and a few Java files copied over to enable ANT to create a zipped Shibboleth-IdP file. Tomcat is then restarted, which unpacks the archive and creates the Shibboleth IdP. In the early days of testing Andy registered his IdP with Testshib to verify the configuration and created certificates by registering with OpenIdP.

However, now that the UK federation is up and running and there are test service providers (SPs) available (e.g. https://target.iay.org.uk/secure/printenv.cgi), he believes it is more practical to configure the IdP for the UK federation from the outset, which avoids having to perform one configuration routine for testing and then repeating it all for production.

Prior to testing, the institution must first join the federation – most will have done this already. A certificate must be generated through the JANET Server Certificate Service and the IdP registered with the federation helpdesk so that it appears in the metadata; step-by-step instructions for all of these processes are available on the UK federation web site. Some further file downloads and configuration are required, e.g. inserting the provider's host name and domain details in various files, before restarting Apache and Tomcat and testing the IdP. For test purposes the simple Apache [AuthUserFile] can still be used to create a dummy user rather than going all the way and authenticating users against the actual campus directory. This enables the IdP to be tested against one of the test SPs to check that the user is successfully authenticated and to view released attributes. While success at this stage proves that the IdP is functional, it isn't practical to register all users in this way, in a file on the system – the intention is to authenticate them from information in the directory or a database.

Tomcat has a mechanism – called Tomcat Realms – for authenticating users either by LDAP from an X500 directory or by using JDBC to access a database. Again, a variety of information about how to do this is available on the Internet, but in essence it involves modifying the Apache configuration to instruct it not to execute the authentication, and instructing Tomcat that it should perform this function instead.

While there are still other small pieces of configuration to carry out and security settings to be finalised, this is the basic setup procedure that Andy uses to establish a Shibboleth IdP.

Andy says that although it sounds complicated, the fact that the install process can be broken down into a number of separately verifiable steps makes it more manageable, and he would encourage anyone to 'have a go' irrespective of prior experience.

Appendix A: Sources of Additional Information

Web components: Apache, Tomcat, OpenSSL and Java

http://tomcat.apache.org/

http://tomcat.apache.org/download-connectors.cgi

http://tomcat.apache.org/connectors-doc/reference/workers.html

http://www.openssl.org/

http://java.sun.com/javase/downloads/index.jsp

Shibboleth

The Internet2 site: https://spaces.internet2.edu/display/SHIB/WebHome

http://www.jiscmail.ac.uk/archives/jisc-shibboleth.html

The UK federation

http://www.ukfederation.org.uk/

http://www.jisc.ac.uk/federation/

http://www.jisc.ac.uk/whatwedo/programmes/programme_cminfrastructure.aspx

Web server guides (take care with version numbers referenced!)

http://www.linux-noob.com/forums/index.php?showtopic=373

http://www.coreservlets.com/Apache-Tomcat-Tutorial/tomcat-5.5.html

http://agiletesting.blogspot.com/2005/10/configuring-apache-2-and-tomcat-55.html

http://www.vanemery.com/Linux/Apache/apache-SSL.html

http://www.devside.net/guides/linux/apache-ssl-deflate

Shibboleth-specific guides

Installing on Novell Suse Linux:

http://www.lrz-muenchen.de/~hommel/shibboleth/shib13c_on_SuSE10.0.html

Setting up Tomcat authentication

https://mams.melcoe.mq.edu.au/zope/mams/pubs/Installation/Tomcat%20Authentication%20for%20Shibboleth%20IdP

https://spaces.internet2.edu/display/SHIB/IdPUserAuthnConfig