

JISC DEVELOPMENT PROGRAMMES

Project Document Cover Sheet

FINAL REPORT

Project

Project Acronym	ASMIMA	Project ID	
Project Title	Adoption of Shibboleth for Multiple Identity Management Applications		
Start Date	April 2005	End Date	March 2006
Lead Institution	Cardiff University		
Project Director	Hugh Beedie		
Project Manager & contact details	Joan Wright wright@cardiff.ac.uk t: +44 (0) 29 2087 4496		
Partner Institutions	n/a		
Project Web URL	http://www.cardiff.ac.uk/insrv/shibboleth/		
Programme Name	JISC Core Middleware Programme (Shibboleth Early Adopters)		
Programme Manager	Nicole Harris		

Document

Document Title	Final Report		
Reporting Period	April 2005 – March 200		
Author(s) & project role	Rhys Smith, IT Officer		
Date	31/03/2006	Filename	ASMIMA_Final_Report.rtf
URL	http://www.cardiff.ac.uk/insrv/shibboleth/ -> “The ASMIMA Project” section		
Access	<input type="checkbox"/> Project and JISC internal	<input type="checkbox"/> General dissemination	

Document History

Version	Date	Comments
V1.0	31 st March 2006	Completed document, sent to JISC

i. Acknowledgements.

The ASMIMA project is a JISC Core Middleware project, funded by JISC. The ASMIMA team would like to thank JISC for supporting the project; and in particular our programme manager, Nicole Harris, for all her help and guidance during the project. The team would also like to thank Richard Annette of MATU and John Bond of EduServ for their all their help and patience!

ii. Executive Summary

The ASMIMA project was a JISC Core Middleware Early Adopter project undertaken by Cardiff University. The main aim of the project was to explore the implementation of Shibboleth as a next generation method of access management to remote resources. The project was split into four major Work Packages: implementing a Shibboleth Identity Provider; using Shibboleth as an alternative to Athens for remotely accessing e-resources; exploring how Shibboleth can help our joint Cardiff University/NHS staff access all resources they are entitled to; and exploring how Shibboleth can be used as an access management technique for an e-science application called BioDiversityWorld.

Of these, the first two Work Packages have progressed very well. We successfully implemented a Shibboleth Identity Provider and have used it to access remote e-resources. This pilot service is intended to become a university-wide production service in time for the next academic session – Sept 2006. There are, however, still a few outstanding issues: a few key resources that are not yet Shibboleth-Athens gateway compliant; and that Cardiff University carefully considering our policies with regards to categories of users and their entitlements. Distinctive contributions from Cardiff University as a result of these sections of the project are that we are well advanced in designing and implementing a fully resilient production implementation of a Shibboleth Identity Provider, and that we well advanced in seriously considering this problem of categorising users and their entitlements in order to fulfil our license obligations.

Some lessons we learned during the implementation of the first two Work Packages cannot be stressed enough to other institutions considering the use of Shibboleth. The first is that implement Shibboleth forces an institution to do several things it probably should already have in place, but likely doesn't – a comprehensive identity management system, proper directory services, and good intra-institution political goodwill. This should include clear policies and guidelines of categories of staff and student and their entitlements. The second lesson is that over the next few years, the landscape of educational Access Management in the UK is going to be very complex – Shibboleth, the Shibboleth-Athens gateway, the Athens-Shibboleth gateway, Classic Athens, and IP Authentication all existing at the same time. Clear documentation and user guides will absolutely be required if users are to make sense of this landscape.

Work Package 3 (exploring interworking with NHS resources using Shibboleth) has been delayed due to difficulties in getting the time of key NHS staff. Progress is being made on a number of fronts, including a meeting in April with a key player. Work is continuing, and a final report for this section of work will be submitted separately at the end of May 2006, with the agreement of our Programme Manager.

Work Package 4 (exploring the use of Shibboleth as an access management method to an e-Science application) has been delayed due to recruitment issues. These issues have all been sorted and work started on February 1st. A final report for this section of work will be submitted separately at the end of May 2006, with the agreement of our Programme Manager.

Cardiff University feels that the ASMIMA project has been a great success, giving us a working Shibboleth implementation ready to turn into a production service in time for the next academic session, and in us learning a lot of lessons that we can, and have, disseminated to the rest of the academic community.

Table of Contents ***Acknowledgements.***

ii. Executive Summary.....	
1. Introduction.....	
2. Related Documents.....	
2.1 ASMIMA related documents	
2.2 Other related documents	
3. Background.....	
3.1 Cardiff University	
3.2 Areas of expertise in INSRV	
3.3 Identity Management and Directory Services	
3.4 Athens	
3.5 Cardiff University and the NHS	
3.6 BioDiversity World	
4. Aims and Objectives.....	
5. Overview of Implementation.....	
5.1 Implementation	
5.2 Federations	
5.3 NHS	
5.4 BioDiversity World	
6. Outputs and Results.....	
6.1 Concrete Internal Outputs	
6.2 Documentation	
6.3 Dissemination Activities	
7. Outcomes.....	
8. Conclusions.....	
9. Implications.....	
Glossary.....	
Appendix 1 – Full Work Package Reports.....	
A1.1 WP1 – Implementation and Testing of Shibboleth IdP	
A1.2 WP2 – Shibboleth-Athens	
Appendix 2 – Test Plan.....	
Appendix 3 – Test Results.....	
Appendix 4 – Full survey of CU's Athens Resources.....	
Appendix 5 – Installation, Configuration, Resiliency and Monitoring guides produced..	

1. Introduction

The Adoption of Shibboleth for Multiple Identity Management Applications (ASMIMA) project is a JISC Core Middleware Programme funded project undertaken by Cardiff University's Directorate of Information Services. The core aim of the project was to implement Shibboleth at Cardiff University for the purpose of using the technology in a number of ways: as a replacement to the Athens service, to help improve joint Cardiff University/NHS staff's computing experience, and to investigate Shibboleth as a method of access management to an e-science application.

This document aims to describe the progress of project, documenting all major milestones and outputs, and more importantly, any lessons learned that we think would be of interest to other institutions considering the use of federated identity management solutions.

The intended audience for this document are twofold: the JISC project management team, as a method of assessing the success of this project; and other institutions similar to Cardiff University that are considering implementing Shibboleth.

This document assumes no previous experience with Shibboleth, but does assume a basic level of familiarity with the concepts of a federated authentication and authorisation model of access management. It does not discuss the technological details of the technology used; instead focussing on general technological and political lessons learned that the project team feel would be usefully disseminated.

2. Related Documents

2.1 ASMIMA related documents

Further documents related to the ASMIMA project that may be of interest are:

- Project website:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> "The ASMIMA Project" section
- The ASMIMA Project Plan:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> "The ASMIMA Project" section
- The ASMIMA Project Work Packages:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> "The ASMIMA Project" section
- Test Plan:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> "The ASMIMA Project" section
- Test Results:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> "The ASMIMA Project" section
- Survey of Gateway-Compliant resources:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> "The ASMIMA Project" section
- Userguide:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> "For Users" section
- Shibboleth installation documents:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> "For Developers" section

- Shibboleth configuration documents:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> “For Developers” section
- Monitoring your Shibboleth solution:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> “For Developers” section
- Designing a resilient solution:
<http://www.cardiff.ac.uk/insrv/shibboleth> -> “For Developers” section

2.2 Other related documents

Other documents related to the content of this report that may be of interest are:

- Internet2 Shibboleth Overview:
<http://shibboleth.internet2.edu/>
- JISC Core Middleware Programme:
http://www.jisc.ac.uk/programme_middleware.html

3. Background

3.1 Cardiff University

Cardiff University (CU) is one of Britain’s major teaching and research universities, and is a member of the Russell Group of premier UK universities. Located in the centre of the capital city of Wales, it has an international reputation for the quality of its work and a mission to be a world class research-led institution and a centre of excellence in its teaching. CU recently merged with the University of Wales College of Medicine (UWCM), resulting in an institution now comprised of approximately 5,000 staff and 22,000 students spread across 28 academic schools and 5 administrative directorates.

CU's Directorate of Information Services (INSRV) is a division of CU that aims to deliver superior computer, library and media services that make a distinctive contribution to CU's research, learning, teaching, community activities and administration functions.

3.2 Areas of expertise in INSRV

INSRV manage many different systems and has a wide range of expertise spread across most platforms, vendors and systems.

On the server side expertise is spread in particular across Unix/Linux systems and tools, and Novell products. On the desktop side, the standard INSRV imaged machines used by the majority of users in CU run Microsoft Windows XP using Novell networking tools and application distribution, whilst a lot of development work in INSRV is done on desktops running various distributions of Linux.

More generally, CU has a newly acquired area of expertise in the area of Identity Management and Directory Services.

3.3 Identity Management and Directory Services

CU is a major user of Identity Management (IDM) systems and is actively engaged in

upgrading its IDM procedures. As part of this, CU is currently in phase 2 of an ongoing project which is implementing an enhanced Identity Management system; this system is expanding to include more features and encompass more internal systems. It ultimately aims to integrate all of its data (held in databases, directories, etc.) to enable access and use to be based upon a person's single identity and administered through a single interface.

CU's directory services are based upon Novell eDirectory, DirXML, and LDAP. Four main trees exist:

- The POBLOGI tree: the “Identity Vault” of the new IDM system;
- The JCCS tree: the main tree used on a day-to-day basis for tasks such as network authentication and workstation management and distributed application management (Novell ZENworks);
- The HARMONICA tree: a tree with a replication of the structure of JCCS tree used by the Citrix service; and
- The FARAWAY tree: a very lightweight replication of the structure of the JCCS tree used by applications for simple authentication and authorisation.

The POBLOGI tree has its own structure, and is a fairly flat tree where users are organised based upon their unique CU identity.

The JCCS tree, HARMONICA tree and FARAWAY tree all have an identical structure organised in a more “usual” way, with multiple levels in the tree. This structure is, however, slightly unusual in that it has multiple bases (CF and CM) for historical reasons – one base for ex-University of Wales Cardiff (o=CF), and one base for ex-University of Wales College of Medicine (o=CM).

3.4 Athens

CU is a major user of the JISC funded Athens Access Management System (AMS), which controls access to web-based subscription services, with over a million CU logins to the Athens service each year. The current Classic Athens service is a centralised repository of user accounts and credentials which is managed at a local level.

This management is currently done through the use of scripts that bulk upload a list of new accounts to be created every night. At the start of a new academic session, these scripts must create thousands of new Athens accounts at a time. Users obtain their Athens usernames and passwords from Library staff. Management of these accounts is a substantial administrative workload.

In an effort to solve this problem, CU was considering switching to AthensDA (a federated access management solution designed by EduServ Athens) before deciding to skip that step and become an early adopter of the next generation of AMS – Shibboleth.

Towards the end of this project, we also received information from JISC that they would stop funding the Classic Athens service as of July 2008. To continue to use the service, institutions will have to pay up to £0.50 per user – this would cost CU approximately £13,500 per year.

3.5 Cardiff University and the NHS

CU has a strong connection with NHS Wales via its University Hospital of Wales site

and elsewhere. It has regular contact with two NHS Trusts: Cardiff & Vale Trust, through the University Hospital of Wales site, and the Velindre Trust, through Velindre Hospital and elsewhere.

The Cardiff & Vale NHS Trust is the largest NHS Trust in Wales and one of the largest in the UK. It provides day to day health services to a population of around 500,000 people living in Cardiff and the Vale of Glamorgan who need hospital treatment, mental health care, care for elderly people and children as well as a growing range of community-based services, including specialist dental services, and new therapies as alternatives to hospital admission.

The Velindre NHS Trust, headquartered in Cardiff, was established in 1994 and has steadily grown to manage a budget of over £130M in 2004/05 and provides a range of specialist services at local, regional and all Wales levels. The Trust comprises of a variety of Divisions who specialise in services as diverse as caring for cancer patients in Velindre Hospital to managing the Digital All Wales Network for NHS Wales from Health Solutions Wales in Brunel House.

Some medic-related members of CU are also members of one of the two trusts and are therefore entitled to access both CU and NHS resources: CU resources usually sitting on the CU network and internet, and NHS resources usually sitting on NHSnet and the internet.

The initial method used which allowed joint CU/NHS personnel to access both types of resources was simply have two computers on their desk, one connected to a CU network outlet and one connected to a NHS network outlet. This is obviously an inelegant solution to the problem.

In an attempt to improve this situation, the All Wales Citrix Service was implemented, using technology from Citrix Systems Inc. This service allows CU members to gain access to the CU network from off-campus, and therefore to resources that are restricted to users on the CU network.

However, a solution for access the other way – i.e. to access NHS resources from outside the NHS network has not yet been implemented; the only way to access these resources currently is to use a computer directly on the NHS network.

NHS Wales' strategy is led by the Welsh Assembly strategy known as Informing Healthcare (IHC). IHC intends to improve healthcare in Wales by making better use of information and technology. A key IHC project is Access To Knowledge (A2K). A2K is a project working to ensure healthcare staff in Wales have easy access to healthcare knowledge and evidence, and have the tools they require to retrieve such evidence. Healthcare staff in hospitals, GP surgeries and NHS Dental practices will be able to electronically access the latest and the best evidence to support them as they care for patients.

3.6 BioDiversity World

Cardiff University hosts the Welsh e-Science Centre (WeSC). WeSC was established as part of the UK national e-Science initiative, and has received additional funding from the Welsh Development Agency, and directly from CU.

Emerging distributed collaborative scientific enterprises require desktop access to very large data collections, very large-scale computing resources, and high performance visualization. WeSC's aim is to develop, implement, and deploy applications to utilise and create e-Science technologies, infrastructure, and services. An essential feature of many of these services is that they will be used collaboratively, and by geographically distributed researchers in fields such as

engineering, physics, earth science, bio-science, and chemistry. In particular, the Centre promotes the collaborative development of large-scale multi-disciplinary applications, and the immersive visualization of large multi-dimensional data sets.

One of WeSC's projects is BioDiversity World (BDW). BDW is a three-year e-Science pilot project funded by the BBSRC to create a GRID-based problem solving environment. This problem solving environment is planned for collaborative exploration and analysis of global biodiversity patterns.

4. Aims and Objectives

CU hoped to achieve several major objectives in the ASMIMA project, some of direct benefit to CU, and some of more general benefit to CU and the wider community.

Of direct benefit to CU, the major objectives we aimed to achieve were:

- A general implementation of Shibboleth in order that CU be on the nationwide cutting edge of federated identity management, in order that our users can make use of emerging technologies;
- An implementation of Shibboleth-Athens as a replacement for the Classic Athens service (skipping the originally intended implementation of Athens DA), in order to ease the administrative burden placed on INSRV in managing both local and Athens accounts for all 27,000 users; also to save CU from having to pay ~£13,500/year come July 2008;
- Improved security, as a Shibboleth-Athens implementation would mean that access to remote resources would be integrated with the existing mechanisms for adding/suspending/deleting standard network accounts;
- Enhanced synergy with our new Identity Management philosophy of a single integrated identity for each individual who is a member of CU;
- A remote access management system more ready for Single Sign On (SSO) – having one identity makes this easier to achieve;
- To explore interworking with NHS via Shibboleth – this would help INSRV and CU set up links to the NHS which could come in useful in future projects;
- To help our joint CU/NHS staff and students gain access to all the resources to which they are entitled;
- To explore the use of Shibboleth as an access management method to the BDW project.

In addition to these objectives of direct benefit to CU, the major objectives that we aimed to achieve for the benefit of the wider community were:

- To gain experience implementing this cutting-edge technology in order to disseminate to other institutions considering implementing Shibboleth;
- To more generally gain an expertise in federated access management technologies so that CU could offer its services to other institutions and organisations;
- To enhance CU's reputation as cutting-edge world-class institution.

5. Overview of Implementation

During the course of the ASMIMA project, we achieved several key milestones and learned many valuable lessons. This section aims to simply give an overview of what happened. For further details, please see Appendix 1 which more fully details each Work Package.

5.1 Implementation

The first step we undertook was to implement a prototype Shibboleth IdP and SP and connect them together, in order to better understand how the technology worked. Once this was achieved, we were then able to properly understand the design requirements for producing a Shibboleth IdP implementation.

Given this basic understanding, we produced a first working version of a Shibboleth IdP, and joined the InQueue federation (the Internet2 testing federation). The experience gained in doing this gave us further experience in the inner-workings of Shibboleth.

Given this now complex understanding of Shibboleth, we were able to produce a design for a fully resilient Shibboleth IdP, which involved using two identical servers with a Citrix NetScaler Layer 4-7 Switch sat in front managing requests to the two servers. This Layer 4-7 switch handles failover and load balancing, whilst an experimental Shibboleth extension created by Georgetown University called hashib (High Availability Shibboleth) managed the task of keeping Shibboleth related information replicated between the two servers.

This resilient design was implemented with only a few hurdles encountered along the way (see Appendices for detail). Georgetown University were interested to hear of our experiences with hashib, as they weren't aware of any other major institutions other than them that were using their extension yet. Our experiences with this extension are all good, however, as it appears to work flawlessly, even though it is technically beta code.

During the implementation of this part of the project, the project team worked closely with the IT Security Team Leader at INSRV, who provided guidance and recommendations for certain aspects of the implementation.

5.2 Federations

Given that we now had a fully resilient Shibboleth-Athens solution, we next went about joining production federations. The first federation we joined was the Athens federation, in order to use the Shibboleth-Athens service.

Technically, joining the Athens federation was straightforward, especially with the help of the Athens Local Authentication Support team. In order to use the Shibboleth-Athens service, we needed to pass a few pieces of information to Athens for each user – a unique identifier, and a name of the Athens Permission Set(s) associated with a user. For the former, we made use of an existing unique identifier for each user from IDMAN, for the latter we created a new attribute in the schema on the FARAWAY tree to hold this information.

Initially we simply manually set the attributes of specific users in order that they could test aspects of the system. Our strategy for production is to have a DirXML rule that populates the attribute driven through our Identity Management System. The initial

rule is to populate users with a “no access” permission set by default and to populate users with a “default access” permission set if they are current staff in the staff database (Compel) or current student in the student records database (SITS).

A working group has been set up within CU to consider all other categories of staff and students, and what resources each are entitled to. This working group comprises senior Information Services staff in consultation with CU’s Registry, Human Resources, Corporate Services, and our NHS Liaison Unit. The list of categories, supplied with this report as a separate spreadsheet, is long and complicated, especially because of our relationship to the NHS. This work is ongoing.

At this point, in order to make the most of our Shibboleth IdP, we made the decision to join the SDSS Development Federation. We registered with the people at EDINA who manage it, and once all the paperwork was completed we became a member.

As we were now a member of three federations (InQueue, the Athens Federation and SDSS), and we had our resilient architecture in place, it was time to thoroughly test our implementation.

To this end, we created a comprehensive test plan (see Appendix 2 for full details). In this plan, we decided to split the tests up logically into systems testing and usage testing. Systems testing comprised of such things as testing of backup and recovery procedures, testing of the resilient architecture and load testing. Usage testing consisted of such things as testing that genuine CU users can log into the Shibboleth service, that people without local accounts cannot log in, that people with the appropriate permissions (Athens Permission Set), and only those people, can access Athens resources to which they are entitled, that the Shibboleth session behaviour is working as expected, etc. Since the main impetus for implementing Shibboleth at CU was to replace the Classic Athens service, a major part of the testing correspondingly focussed on testing the Shibboleth Athens gateway.

While most of the tests were successful, the one area that caused us concern was in testing whether all our current Athens resources worked through the new system – of the 92 Athens resources that CU subscribed to, we discovered that 85 were fully working with the Shibboleth-Athens gateway and that 7 were not. As some of these resources are key Athens resources at Cardiff – resources with high levels of use – we contacted each of the suppliers of these resources, requesting information on their plans to Shibboleth enable their resource. The majority are planning to do this by the end of the summer or to move content to a compliant service. One resource, Westlaw, had no plans at all to Shibboleth enable their service. This was potentially a show stopper for CU: our law school needs access to Westlaw. CU contacted Westlaw to stress this importance, also stressing this issue to JISC and the wider community via the Shibboleth mailing lists. We issued a “call to arms” amongst the community asking any institution considering implementing Shibboleth who subscribe to Westlaw to contact Westlaw and register their interest in a Shibboleth enabled version. Due to this pressure from JISC, CU, and the rest of the community, Westlaw now say that they will endeavour to be compliant by the end of 2006 but do not yet have a firm date. Until then, we have an interim solution in place: we have a generic login to issue to all of the new intake. This solution is not ideal, but hopefully should only have to be done for the next academic session.

Given this successful implementation of a Shibboleth IdP, we next set up comprehensive monitoring of the servers, in order that we could assess the usage of the Shibboleth service. These tools allow us to view the status of our servers and the Shibboleth usage on each server (current and total for all/any federations). This will

allow us to closely monitor the uptake of the Shibboleth service, and understand the usage of the service. Further details of the monitoring are in the Appendices.

Presuming the other six resources keep to their timelines, and this interim solution for the problem with Westlaw is in place in time for the next academic session, we believe that we are ready to move forward to a production Shibboleth service next academic session. An INSRV internal project proposal to do this has been drawn up and the project will run throughout the summer. The key aims of this project are to decide the best strategy for rollout to the whole university, to produce internal usage documentation and to train CU staff in the multiple access methods available.

5.3 NHS

NHS Wales differs from NHS England. Their strategy for access management is being developed by the Informing Health Care programme and, in particular, its Access to Knowledge project.

The IHC programme is developing a single-sign on authentication for the individual health record, which will use Smartcard technology. This single sign-on will invoke a series of further permissions, one of which is the electronic resources in the/ an national e-library.

The technical architects here have said that Shibboleth is as good if not a better strategic option for sign-on access to the NHS e-library per se (but that would need to be confirmed by a strategic assessment). However, in their assessment so far, Shibboleth is unlikely to be sophisticated enough for individual health record single sign on, so they are not progressing it. The A2K project, which has only a further 11 months to run, is using Classic Athens for its e-library portal.

It is CU's view that IHC should give greater consideration to the role of Shibboleth as a component of access management. The University Librarian, Janet Peters, is now a member of the A2K Project Board and has raised in particular the needs of medical students and joint University/NHS staff. The advice from the A2K project is that these issues should be discussed with the IHC Programme Director, Gary Bullock. We have been seeking a meeting with this very busy person for several months and this has now been arranged for 24th April.

We have also had discussions with the CU Dental School. They have a requirement for CU staff who are honorary consultants and specialist registrars to access the Cardiff and Vale Trust Clinical Portal from the CU network. Two solutions are under consideration:

- Implement a Citrix Service similar to the All Wales Citrix service but in reverse. Permission for this is already in place but the costs are expected to be high.
- Implement a secure proxy server to access the NHS network. This is a lower cost solution but would require political agreement.

We are also discussing options with Velindre Trust which is smaller and possibly more flexible. One technically simple option we are exploring is whether permission could be granted for resources that the Velindre Trust has purchased to be added to the Cardiff University Athens resource list. We would then set up the appropriate permission set to restrict access to joint CU/Velindre staff.

By agreement with our JISC programme manager, the final report on this workpackage will be submitted at the end of May.

5.4 Biodiversity World

In the BiodiversityWorld exemplar we are investigating ways in which Shibboleth-based authentication can be incorporated into the enactment of BiodiversityWorld workflows by the Triana system. We have explored the options for achieving this: at present we are creating a new kind of workflow component for Triana which will allow the users to authenticate themselves against an identity provider in the usual https-based manner. The other task that we are planning is to implement authentication mechanisms that restrict BiodiversityWorld resources so that only authenticated users can access the resources and use them when executing workflows.

By agreement with our JISC programme manager, the final report on this workpackage will be submitted at the end of May.

6. Outputs and Results

Several outputs were created as a result of this project. In this section, we catalogue them all.

6.1 Concrete Internal Outputs

Internally to CU, the main outputs of this project were producing a successful, resilient implementation of a Shibboleth IdP, which we are planning on switching to a production service in time for the next academic session – Sept 2006 – allowing us to then gradually phase out the Classic Athens service. Come July 2008, when JISC stop funding the Classic Athens service, this will save CU approximately £13,500/year.

Additionally, we have collaborated with the NHS in an attempt to allow joint CU/NHS users to access all resources to which they are entitled, and we have the beginnings of a political agreement in place to do this. Thus, we are on course to help some of our joint CU/NHS staff's interaction with resources a little easier.

6.2 Documentation

As part of the project, we have gained an in-depth experience in installing and configuring a Shibboleth IdP. We have used this information to produce some documentation giving a step-by-step guide to doing this, which will be made available on the project website shortly.

Additionally, as we are considering how to move this pilot service into production service, we have considered some design aspects that haven't been given much attention by the community. The two main areas we have considered are the levels of resiliency that can be implemented and comparatively how much each would cost; and how to monitor your Shibboleth IdP servers in such a way as to be able to easily and quickly get relevant information. Our experience in these areas has let us produce some presentations and documentation in these areas. These presentations and documents have been made available to the community.

Finally, an output we feel may be of interested to other institutions is our comprehensive test plan for a Shibboleth service, and the Shibboleth-Athens

gateway.

6.3 Dissemination Activities

Part of the project was to disseminate information about Federated Access Management, Shibboleth, why other institutions and organisations would want to adopt the ideas/technologies, and generally inform anyone we can about the future of Access Management in the UK.

To this end, we at CU:

- Presented to WREN (the Welsh Janet User Group), giving them an overview of Shibboleth and how it would relate to our fellow HE and FE institutes;
- Presented to National Library of Wales, giving them an overview of the difference between Classic Athens and Shibboleth Athens, and why they would want to go down the Shibboleth route;
- Presented at the end of a MATU workshop in Blagdon, telling the attendees about how to plan for a production Shibboleth service, including resiliency and monitoring issues and solutions;
- Presented internally in CU to the University Systems Group (USG), Senior Library Staff, and the whole of INSRV;
- Are going to present in June at Gregynog – an annual week-long workshop held for technical and library staff for all HE institutes in Wales, where we plan on giving an overview of Shibboleth and how it would relate to them;
- Are going to present at the Novell BrainStormer '06 event held in Telford in June 2006;
- May be presenting at the Novell TTP (Technology Transfer Partner) conference held in the United States in July;
- Held a meeting with the University of Wales Newport, who are planning on implementing Shibboleth over the summer of 2006 and wanted some advice from someone who has already done it; we have also offered to go help them implement Shibboleth;
- Held discussions with members of the A2K Board and will be meeting with the IHC Programme Director;
- Held discussions with HOWIS staff, giving an overview of Shibboleth and how it could effect NHS trusts across Wales;
- We have agreed to help the JISC Plagiarism Detection service test Shibboleth access to the service when they Shibboleth-enable it;
- Held discussion with ProduceWeb (a HE initiative which supports procurement with the HE sector and research councils), giving them an overview of Shibboleth and the benefits they would gain from setting themselves up as an SP.

7. Outcomes

CU hoped to achieve several major objectives in the ASMIMA project, some of direct benefit to CU, and some of more general benefit to CU and the wider community. Comparing the outcomes against the Aims and Objectives stated in Section 4, CU achieved the vast majority of the objectives.

Of direct benefit to CU, the major objectives we achieved were:

- A general implementation of Shibboleth ready to become a founding member of the new JISC UK Access Management Federation, which is planned to go live in August 2006. This will allow our users to access remote resources through this federation from the moment it is set up;
- An implementation of Shibboleth-Athens as a replacement for the Classic Athens service. The pilot service produced as part of the ASMIMA project is ready to move into full production status to be used institution-wide in time for the next academic session (September 2006);
- Improved security – our Shibboleth-Athens implementation integrates directly with the existing mechanisms for adding/suspending/deleting standard network accounts;
- Synergy with our new Identity Management philosophy of a user having a single integrated identity;
- A remote access management system more ready for Single Sign On (SSO) – having one identity makes this easier to achieve;
- INSRV explored interacting with the NHS – we created some key links and gained valuable experience;
- We are developing plans to enable joint CU/Velindre NHS Trust personnel to gain access to all the resources to which they are entitled;
- We have explored the possibilities of using Shibboleth with the Dental School of CU to allow joint CU/Cardiff&Vale NHS Trust personnel to gain access to NHS resources whilst on the CU network;
- We are exploring the use of Shibboleth as an access management method to the BDW project.

In addition to these objectives of direct benefit to CU, the major objectives that we achieved for the benefit of the wider community were:

- We gained experience implementing Shibboleth, giving us the opportunity to disseminate this information to other institutions considering implementing Shibboleth; this is evidenced by our previous and planned presentations about Shibboleth to the National Library of Wales and Aberystwyth Universities, the University of Wales Newport, the Welsh Janet User Group (WREN), Access To Knowledge Board members, and at the University of Wales annual meeting at Gregynog.
- We more generally gained an expertise in federated access management technologies so that CU could both offer its services to other institutions and organisations and plan its own future technologies around the idea of federated access management. This former is demonstrated by our work with

the University of Wales Newport in helping them plan for Shibboleth-Athens, and with ProcureWeb in alerting them to the possibility of using Shibboleth as an Access Management System for their product; whilst the latter is evidenced by the addition of the idea of allowing federated identity login to the new Modern Working Environment (MWE) tender that CU is in the process of issuing;

- We believe that all of the above has enhanced CU's reputation as cutting-edge, world-class institution.

8. Conclusions

Implementing Shibboleth as an early adopter has been very useful to CU. We have implemented a pilot Shibboleth Identity Provider and are moving to turn this into a production service in time for the next academic session (September 2006). Allowing our users to access resources through Shibboleth or the Shibboleth-Athens gateway rather than the Classic Athens service will considerably ease the administrative burden which is currently placed on INSRV, allowing INSRV to concentrate on delivering our stated aims. Also, Shibboleth will make access remote resources easier for our users, as they will have one less username and password to remember. Easier access will also increase awareness of the range of resources available to them.

The main conclusion that we have drawn from our experiences in this project is that implementing a Shibboleth IdP is technically fairly straightforward (not “easy”) -- *if* you already have the necessary components in place. Shibboleth requires that an organisation do all the things it really ought to have already done, but usually haven't – a comprehensive identity management system, proper directory services, and good intra-institution political good will. This conclusion is very important, as implementing these things is a very time and resource consuming task, but ultimately are very worthwhile things to have.

Even though CU is well advanced with respect to Identity Management and Directory Services, this project has reinforced the need for clear policies and guidelines on the many different categories of staff and students and their entitlements. We will be putting recommendations to the University which may affect the status of some categories. For example, Human Resources has published procedures for the granting of honorary titles, including visiting academics, but honorary title holders are not regarded as members of the University. Given the importance of research collaboration with other institutions, there is a case for reviewing the status of honorary visiting researchers.

The tightening up of these policies may encounter some user resistance. For example, with the exception of Emeritus Professors, retired staff are not members of the University. Prior to the introduction of our Identity Management system, the decision on when to delete the accounts of retired staff was taken by Schools, as retired staff often have a continuing collaborative association with the University. Under the new procedures, their access to Athens resources will be removed when they retire, even if they are permitted to retain a CU network account for a period. Clear explanation to Schools of the legal requirements of our license arrangements will be required.

There will be an ongoing political and administrative workload associated with handling these exceptional cases and the new categories that will continue to

emerge.

It is our recommendation that these issues should be addressed on a national basis. If license agreements were formulated on an FTE+5% basis, they could permit a small number of exceptional cases to have legitimate access to resources. We plan to raise this idea with the NESLI2 national initiative through our contacts with them.

The Shibboleth-Athens gateway works well for us in general, and has just a few last wrinkles to iron out – for example, there are still one or two key resources that are not gateway compliant. Until this is achieved, we will still have to issue Classic Athens accounts to users who could not get the full service they need from Shibboleth-Athens. It will work fine for the majority of users, however.

One major thing any institution in a similar circumstance to us that intends to implement Shibboleth over the next year or two needs to take into account is that during the transition period from Classic Athens to Shibboleth, via the Shibboleth-Athens gateway (and the Athens-Shibboleth gateway) is that producing clear documentation on how users should be accessing resources (via Shibboleth? Shibboleth-Athens? Athens? IP Authentication?) is a must – the landscape during this time of interim solutions is complex and hard enough for an expert to understand, let alone a normal user who just wants to view some restricted content as easily as possible. Documentation that explains this will require skilled writing.

One final thing that we have learnt in this project is that when working with the NHS, one must allow more time and effort than with most other organisations, as generally NHS staff are very busy and hard to get hold of. Also, if one is to work with the NHS, political goodwill and agreements between you and them are a must to even begin to get anything done.

9. Implications

The implications that we think come out of the ASMIMA project are thus:

- Technically, implementing Shibboleth is fairly straightforward, if all the back-end requirements are already in place (see the next point). Given these being in place, what is needed is a timescale of a couple of months, and input from a wide variety of people – your directory services team, your server management team, your user interface team, your security team, and people able to sign legal agreements on behalf of the university;
- In order for an organisation to set themselves as a Shibboleth Identity Provider, that organisation needs to do a lot of things it should already have done, but probably hasn't – comprehensive identity management procedures (including knowing every category of users and their entitlements), and full directory services.
- Organisations need to define clear policies and guidelines on the entitlements of “non-standard” staff and students.
- National renegotiation of license agreements to an FTE+x% basis would reduce the political and administrative effort within institutions.
- Shibboleth is strong in academia, but from our interactions with a wide variety of people and organisations, it is not strong in the commercial world. This is potentially a problem, for example, for ProcureWeb.

- The N-tier issue, which affects BioDiversity world is similar to that faced by Federated Search engines.
- Few commercial portal products are Shibboleth-enabled.

Glossary

A2K – An IHC project working to ensure healthcare staff in Wales have easy access to healthcare knowledge and evidence.

<http://www.wales.nhs.uk/ihc/page.cfm?pid=8788>

Apache – Web server software

ASMIMA (Adoption of Shibboleth for Multiple Identity Management Applications) – The JISC funded project at Cardiff University, part of the Core Middleware Program, specifically as part of the Shibboleth Early Adopters scheme

Athens - Athens is an Access Management system for controlling secure access to web based services

Athens-Shibboleth Gateway – A service that allows Athens users to connect to Shibboleth Service Providers using Athens as an Identity Provider

eDirectory - Novell eDirectory (formerly called Novell Directory Services) is an X.500 compatible directory service software product released in 1993 by Novell for centrally managing access to resources on multiple servers and computers within a given network

EduServ - A not-for-profit IT services group, providing support and solutions for business critical hosting, for e-learning, e-government and e-commerce, and for network identity management. Currently hold the JISC contract to run Athens

FARAWAY Tree – An eDirectory tree used in Cardiff University by web applications that need access to only a small proportion of the attributes held in the full tree

HOWIS – Health of Wales Information Service <http://www.wales.nhs.uk/>

Internet2 – Internet2 is a consortium being led by 207 US universities working in partnership with industry and government to develop and deploy advanced network applications and technologies, accelerating the creation of tomorrow's Internet. Internet2 is recreating the partnership among academia, industry and government that fostered today's Internet in its infancy.

IdP – Shibboleth Identity Provider

IHC - Informing Healthcare is a Welsh Assembly Programme to develop new methods, tools and information technologies to transform health services for the people of Wales. <http://www.wales.nhs.uk/ihc/home.cfm>

JISC – Joint Information Systems Committee

LDAP - A protocol used to access a directory listing, in this case, the FARAWAY tree

Linux (properly GNU/Linux) – An open source operating system

MATU – Middleware Assisted Take-up Service. A pilot service to assist Shibboleth Early Adopters, funded by JISC. <http://www.matu.ac.uk/>

Mod_jk – An interface between Apache and Tomcat

MRTG (Multi Router Traffic Grapher) - Free network traffic monitoring system which displays the results as graphs on the Web

NESLI2 – A UK national initiative for the licensing of electronic journals on behalf of the higher and further education and research communities, 2003-2006.

POBLOGI Tree – An eDirectory tree used by the Cardiff Identity Management System.

SDSS (Shibboleth Development Support Service) – A development Shibboleth federation for managing access to UK academic online resources. Will eventually become SPARTA

Shibboleth – A technology developed by Internet2

Shibboleth-Athens Gateway - A service that allows users with an institutional IdP to connect to existing Athens Service Providers, worked by Athens becoming a Shibboleth Service Provider in its own right

SP – Shibboleth Service Provider

SPARTA – The original name for a UK-wide Shibboleth federation for Higher Education institutes and resources, now just known as the JISC UK Access Management Federation

Tomcat – A Java Servlet container and web server

UK-AC-JCCS Tree – The main Cardiff eDirectory tree used for network authentication, workstation management and distributed application management.

Appendix 1 – Full Work Package Reports

A1.1 WP1 – Implementation and Testing of Shibboleth IdP

A1.1.1 Introduction

This first Work Package was intended to give CU a fully working resilient Shibboleth IdP for use with other applications that can make use of federated identities, and to enable the following three Work Packages.

The bulk of the work was organised into two stages: firstly implementing a simple working prototype version of a Shibboleth IdP as quickly as possible; secondly then designing and implementing a fully resilient version of a Shibboleth IdP ready for production use.

A1.1.2 Aims and Objectives

The main aims and objectives for this Work Package were to:

1. Identify the appropriate platform and basic design for the implementation of the Shibboleth IdP software;
2. Establish contacts within the Shibboleth users' community;
3. Create a full project plan;
4. Create a local website for the project;
5. Implement a simple working prototype of a Shibboleth IdP, tested against at least one Shibboleth SP;
6. Investigate what is necessary to build resilience into a Shibboleth IdP;
7. Undertake a review of Identity Management procedures necessary for maintenance for user attributes;
8. Create a design for a fully resilient implementation of a Shibboleth IdP;
9. Implement this resilient IdP so that it is ready for use with applications that can make use of federated identities;
10. Create a full formal test plan for the Shibboleth IdP;
11. Carry out these tests.

A1.1.3 Implementation

The first step taken was to implement a Shibboleth IdP and Shibboleth SP in a bilateral federation setup, meaning they were setup to only communicate with each other. The IdP and SP were set up on separate virtual machines, each with Debian Linux installed, with the IdP hooked up to a simple database of users. This experience gained setting these up provided a great deal of knowledge about the inner workings and configuration of Shibboleth components.

During this initial implementation, project staff established contacts within the Shibboleth user community, by joining the Internet2 Shibboleth-users mailing list, the JISC-Middleware mailing list, the JISC-Shibboleth mailing list, and later the Internet2 Shibboleth Wiki.

The next step taken was to implement a prototype IdP. We installed the Shibboleth

IdP software on a spare machine, and joined the Internet2 InQueue testing federation. This IdP was linked up to our directory – specifically, the FARAWAY tree, and was implemented using direct Shibboleth apache authentication.

As we were now interacting with our directory and could set up the IdP to release information about our users, we went through a period of reviewing identity management procedures.

At this point, we knew enough about the workings of the Shibboleth IdP to design a resilient solution. We decided that for a production service with a high level of availability, we needed two identical servers set up in a load balanced failover manner, each with a high degree of resiliency. To this end, we purchased two rackmount Dell 1U servers with dual power supplies, dual processors, RAID arrays, etc. Logically in front of these we placed a Citrix NetScaler load balancer. The NetScaler polls each server every 5 seconds to see if the Shibboleth service is still alive (using the built-in Shibboleth status page). If one becomes unavailable, the NetScaler stops passing requests to that particular server. In fact, for further resiliency, we have two Citrix NetScalers set up in a failover manner (if one NetScaler device fails, the other one will take over).

When using a Shibboleth IdP in a load balanced environment, thought has to be given to in-state information. The IdP contains some state information that is maintained in memory - this means that when a client makes an initial request to one IdP server, then the next request, which will likely end up at a different node, will fail because the second server does not have the necessary state information. There are several methods of solving this – see the Internet2 Shibboleth Wiki for further information. We chose to use Georgetown University's beta high availability Shibboleth extension named “HA-Shib”. We were in contact with the developer, Chad La Joie, who gave us some assistance in getting it installed and working, and in return we provided feedback for the beta software. Once configured, HA-Shib, although a beta extension, works perfectly. In essence, HA-Shib maintains a shared memory between Shibboleth IdP servers. When a client makes a request to one IdP server, the state information about that request gets shared in amongst the IdP cluster. This solves the problem described above.

This resilient design was implemented with only a few hurdles encountered along the way. Of these, one of the the largest hurdles we encountered was the discovery that our unusual setup of having multiple bases in our directory hampered a simple hookup between Shibboleth and our directory, as when specifying a connection to an LDAP tree, most software requires that the administrator specify a single base to query. To complicate matter furthers, Shibboleth actually interacts with your directory in two different ways – once for Authentication of a user, usually either through Apache or Tomcat form-based authentication (which attempts to bind to the LDAP tree given the credentials the user has passed to see if their credentials are valid), and once through the Shibboleth Attribute Authority (which queries the LDAP tree for attributes for an authenticated user). Since these two different methods of interaction with the directory are two separate parts of the Shibboleth system, the connection to the directory is actually specified in two different ways, and therefore two different solutions to the multiple base problems were needed.

Solutions were found, after a bit of experimentation. For the Authentication part of the Shibboleth to LDAP interaction, a custom JAAS (Java Authentication and Authorisation Services) connector was written that is able to query both bases, while for the Authorisation part of the Shibboleth to LDAP the use of a FailoverDependency object in the Shibboleth-IdP setup was used (courtesy of some help from MATU),

meaning that the Shibboleth software would first search the CF base, and if no results were returned, it would then try the CM base. These two solutions together allow users from either base to successfully use Shibboleth.

Another technical hurdle that was encountered at this stage was using the very latest version of tomcat (v5.1.15). It transpires that the way tomcat handles authentication was handled changed in a minor way in the new release of tomcat (the developers fixed a bug related to roles). No one had encountered this problem before, so we encountered the problem, figured out the solution and let the community know about this problem through emailing the shibboleth mailing lists, and adding the information to the Internet2 Shibboleth WIKI.

Once this resilient design was implemented and working, we set up comprehensive monitoring of the servers, so that we could monitor their performance and usage. To this end, SNMP was set up on the machines in order to monitor key system information (e.g. CPU usage, Memory usage, network latency, disk usage, etc), and Perl scripts were written that parse the Shibboleth log files in order to monitor Shibboleth usage. These scripts keep track of both current (last 5 minutes) and total usage of the IdP service, for all federations and per federation. MRTG was then setup to create easy to view graphs of all of this information. The scripts and monitoring ideas have been advertised to members of the Shibboleth communities and made available to those who requested them.

The monitoring of Shibboleth usage is going to be rewritten soon, however, by changing the way the statistics are gathered: instead of parsing the Shibboleth log files, a Java filter will be written that directly listens for requests to the IdP and records these requests. This is being done as it is a much more elegant solution and should be more reliable. Once rewritten, the software created to

A1.1.4 Outcomes and Results

The main outcomes of this Work Package were that CU:

1. Identified the appropriate platform and basic design for the implementation of the Shibboleth IdP software;
2. Established contacts within the Shibboleth users' community by joining the Internet2 shibboleth-users mailing lists; joining the JISC middleware and shibboleth mailing lists; joining the Internet2 Shibboleth WIKI;
3. Created a full project plan;
4. Created a local website for the project;
5. Implemented a simple working prototype of a Shibboleth IdP and a Shibboleth SP in separate virtual machines using a bilateral federation setup;
6. Investigated what was necessary to build resilience into a Shibboleth IdP;
7. Undertook a review of Identity Management procedures necessary for maintenance for user attributes;
8. Created a design for a fully resilient implementation of a Shibboleth IdP;
9. Implemented this resilient IdP so that it is ready for use with applications that can make use of federated identities;
10. Created a full formal test plan for the Shibboleth IdP;
11. Carried out these tests.

A1.1.5 Lessons Learned

We at CU have learned several valuable lessons when implementing a Shibboleth IdP, some technical and some political.

Technically, we have learned that if your LDAP implementation is non-standard, you may have less of an easy time in getting Shibboleth to interface with it, and may have to implement custom Java code to get it working.

Another major lesson we have learned is how to set up a resilient solution and monitor it. This information has already been made partly available to the community, and will be made available in full soon.

More interestingly, we have learned that implementing Shibboleth forces an organisation to do some things it probably should already be doing – having fully defined identity management procedures which take account of all staff and student types in one's institution. If Shibboleth is going to hook up to your directory services, then your directory services have to be managed such that they are up to date and accurate, otherwise you risk breaking the terms of service of the federations you are a part of. We at CU cannot stress the importance of these particular lessons enough – fully fledged identity management is not a simple thing to implement and will take a great deal of time and effort to do so.

Another lesson learned is in the resources needed to set up a Shibboleth IdP. At CU, setting up the Shibboleth service has required input from the Directory Services team, the Athens team, the Systems Management team and the User Interface team. As you can see, a wide variety of expertise is needed.

Given no prior knowledge of Shibboleth, it would probably take a full-time employee a couple of months to implement, configure and thoroughly test Shibboleth at an organisation. Given a prior working knowledge of Shibboleth, this figure would likely come down to a matter of weeks.

A1.1.6 Conclusions

This Work Package was completed successfully: a resilient, fully monitored, fully tested implementation of a Shibboleth IdP was completed and is in place. We learned several valuable lessons, and produced some information that we consider will be valuable to the community.

A1.2 WP2 – Shibboleth-Athens

A1.2.1 Introduction

Cardiff University is currently a major user of the Classic Athens service, with over a million Athens logins a year. The duplication of effort required to maintain separate local and Athens accounts is a significant administrative burden.

The size of this operation means that it was already provisionally planned to move to Athens DA in the summer of 2005 but, given current international trends, it was decided that a Shibboleth implementation would seem to be a more logical step.

A1.2.2 Aims and Objectives

The main aims and objectives for this Work Package are to:

1. Reduce the effort currently employed in creating and managing separate local and Athens accounts;
2. Allow access to Shibboleth-Athens resources to Shibboleth enabled users;
3. Allow user access to Beilstein Crossfire (a non-web based resource) through Shibboleth;
4. Documentation and Reports based on testing and evaluating user experience in accessing resources using Shibboleth.

A1.2.3 Implementation

Given the working version of a Shibboleth IdP that resulted from Work Package 1, we joined the Athens Federation in order to use the Shibboleth-Athens federation.

Athens required two specific attributes be passed to them – a unique identifier for a user, and the name of the Athens Permission Set(s) associated with that user.

The first attribute, a unique identifier, was simple for CU to pass across – our new Identity Management system gives each unique individual member of CU an identity number. We had the choice of encrypting the number for security purposes, but since any entities outside CU could not link a plaintext number to the identity of an individual in CU, and passing across the plaintext number would help us in analysing usage of Shibboleth-Athens by user, we took the decision to just pass the number across as-is. We simply mapped this identity number to the attribute that Athens was looking at – namely, an attribute called eduPersonPrincipleID.

The second attribute, the name of the permission set(s) associated with that user, was a little more tricky to set up. Access to Athens resources is managed through the use of these permission sets. Essentially, a permission set is a list of accessible Athens resources. We currently have 5 permission sets:

- “no access”: an empty permission set; for users who aren't allowed to access any resources;
- “default access”: a permission set containing all resources all users are allowed to access;
- three permission sets, with one restricted resource in each (e.g. A medically restricted resource).

To indicate which resources a user is allowed to access, we have to pass the name of one (or more) of these permission sets to Athens. To achieve this we modified the schema of the FARAWAY tree to add a new multi-value attribute to each user. This attribute is the container for the name of the permission set(s) associated with that user.

To populate this attribute we had two options – to set up a system to manually populate the attribute, or have some way of doing it automatically. Obviously, an automatic system would result in an easier administrative task for us, therefore, we went with this latter option.

The strategy is to have a DirXML rule that populates the attribute through our Identity Management System. The initial rule is to populate users with the “no access” permission set by default and to populate users with the “default access” permission set if they are current staff in the staff database (Compel) or a current student in the student records database (SITS). A working group has been set up within CU to consider all other categories of staff and students, and what resources each are entitled to. This work is ongoing.

Initially however, we simply manually set the attributes of specific users in order that they could test aspects of the system.

Given that we were now able to pass across the necessary attributes to Athens, we started to test the Shibboleth-Athens gateway. To do this, we first designed a comprehensive test plan. We split the tests up logically into systems testing and usage testing. Systems testing comprised of such things as testing of backup and recovery procedures, testing of the resilient architecture and load testing. Usage testing consisted of such things as testing that genuine CU users can log into the Shibboleth service, that people without local accounts cannot log in, that people with the appropriate permissions (Athens Permission Set) can access resources to which they are entitled, that the Shibboleth session behaviour is working as expected, etc. See Appendix 2 for full details of the test plan.

We encountered several problems during this testing, a few of which were minor, while a few were major which could potentially effect the takeup of a production Shibboleth-Athens service.

Issue 1 - Non-compliant resources:

We tested all of the 91 web-based Athens resources that CU subscribed to, from both on and off campus. Of these, we discovered that 85 were fully working with the Shibboleth-Athens gateway, and that 6 resources were not gateway compliant. This means that although they work with Classic Athens, they do not work with Shibboleth Athens. These resources were Westlaw, LexisNexis, Dialog, EDINA Digimap, JSTOR and Proquest. As some of these resources are key Athens resources at Cardiff – resources with high levels of use – we contacted each of the suppliers of these resources, requesting information on their plans to Shibboleth enable their resource. Five of these were planning on doing this by the end of the summer or to move content to a compliant service. Digimap have a gateway compliant test system and have successfully tested this with a Cardiff account. They will be releasing this once testing is complete. Finally, one resource – Westlaw – had no plans at all to Shibboleth enable their service.

Westlaw not being gateway-compliant was a big issue – potentially a show stopper for CU. The Law School needs access to Westlaw. CU contacted Westlaw to stress this importance, also stressing this issue to JISC and the wider community via the

Shibboleth mailing lists. We issued a “call to arms” amongst the community asking any institution considering implementing Shibboleth who subscribe to Westlaw to contact Westlaw and register their interest in a Shibboleth enabled version. Due to this pressure from JISC, CU, and the rest of the community, Westlaw now say that they will endeavour to be compliant by the end of 2006 but do not yet have a firm date. So, until this is done, CU needed an interim solution.

One potential solution would be to issue Classic Athens accounts to law staff/students, but INSRV did not wish to do this as it would partly negate the whole reasoning for the Shibboleth project – to get rid of the Classic Athens administration burden. Another potential solution would be to make Westlaw available through our existing Citrix solution. This would allow users to login to Citrix from off campus and access Westlaw as if they were on-campus, using IP authentication to gain access. INSRV was, however, reluctant to do this as we were unsure as to whether the Citrix service could deal with this sudden increase in usage without affecting the performance of the service. We also considered the use of a proxy server such as EZProxy. However, a major CU project, the Modern Working Environment, is addressing the general issue of off-campus access to IP-restricted resources, so we wish to look at this in the light of a general solution. The interim solution that we have decided to follow, was to contact Westlaw, explain this issue, and request a generic login to issue to all of the new intake. This solution is not ideal, but hopefully should only have to be done for this academic session, if Westlaw become gateway compliant by Sept 2007.

Issue 2 - The cookie problem:

The second problem encountered is known as the “cookie problem” - so called because it is a result of a cookie being set on the user's web browser. In essence, the problem is this: when a user signs in using Shibboleth-Athens instead of Classic Athens, their browser remembers this fact (also remembering which institution they are from). If the user then requires to login to an Athens resource using Classic Athens, when they choose the “Athens Login” link on the resource, they will be automatically redirected to their institution's Shibboleth IdP. This is both potentially confusing to the average user, and means a great deal of clicking and following of links if a user wishes to use both Classic and Shibboleth Athens.

We set up a workaround to this problem by setting up a user guide to the Shibboleth service in which we described the idea of “turning on/off” Shibboleth access, and providing a single link for each which set or unset the cookie, as necessary. Thus, while not solving the problem, we set up a method for getting around the problem that should be easily understood by the average user.

Issue 3 - Non-web based resources:

The third major problem encountered was the use of non-web based resources which make use of the Athens service, such as Beilstein Crossfire and EndNote. As with the web-based resources we had issues with, we contacted the suppliers of these services. Beilstein Crossfire informed us that they were presently upgrading their application to be Shibboleth compliant, and desired help in testing their upgrades. As such, we provided them with the details of our IdP service, along with a test login and password. The results of their tests showed their application to successfully be able to use Shibboleth. Endnote, however, informed us that they were aware of Shibboleth and were keeping an eye on its development, but currently were not changing their application to use it. Further pressure, as with Westlaw, is required here from the community.

Issue 4- Speed of Shibboleth-Athens

The fourth major problem encountered, and another possible show-stopper, was an issue as to the speed of the Shibboleth-Athens gateway. Occasionally, it took up to (or over!) 20 seconds to log a user in to a Shibboleth-Athens resource. This would obviously be unacceptable to the majority of our users. We contacted Athens about this issue, and were informed that it was a bug in their implementation. This bug was then fixed, and the speed issues ceased.

Issue 5 - Personalisation:

The fifth major problem encountered was an issue to do with personalisation of resources, and personalisation when migrating from Classic Athens to Shibboleth Athens. Basically, as a user's Shibboleth-Athens account is not linked to their Classic Athens account, users will have to set up any personalisation again on their new account. This is an obvious issue for take-up, but not an insurmountable issue as it will only need to be done once.

However, we have encountered a potential problem. If a user has provided additional information to a resource when logging in to that resource for the first time (e.g. Email address for alerts), and when they log into that resource again for the first time with their Shibboleth credentials and are asked for this additional information, some resources refuse the information the user provides as it is already registered – e.g. Providing the same email address for alerts. This could be an issue for us regarding take-up of a production service, and we will be bring this to the attention of all resources that this occurs on. A simple workaround that was used with Web of Knowledge is to supply the email address in a different format, *mailname@cardiff.ac.uk* instead of *mailname@cf.ac.uk*

We have found one resource, ZETOC, which offers the user the ability to move their existing alerts to a new username but we suspect most resources do not offer this.

There is also the possibility that a Classic Athens user who moves to Shibboleth may still receive “orphaned” alerts set up under Classic Athens. The migration procedure will need to be clearly documented and tested.

Finally, we set up a website for the users, containing user documentation on how to use the pilot service, and how to switch between using Classic Athens and Shibboleth-Athens. Doing this made us realise an important point when it comes to roll-out of the production version of this service – during the transition period from Classic Athens to Shibboleth, via the Shibboleth-Athens gateway (and the Athens-Shibboleth gateway), producing clear documentation on how users should be accessing resources (via Shibboleth? Shibboleth-Athens? Athens? IP Authentication?) is a must – the landscape during this time of interim solutions is complex and hard enough for an expert to understand, let alone a normal user who just wants to view some restricted content as easily as possible. Documentation that explains this in an easy to understand manner is a must.

As this pilot of access to remote resource through the Shibboleth-Athens gateway has been deemed a success, we are moving to turn this into a production service in time for the next academic session (September 2006) and a project team is being formed to manage this.

A1.2.4 Outcomes and Results

The outcomes of this Work Package were:

1. We implemented Shibboleth-Athens, thus giving us a system that will reduce the effort currently employed in creating and managing separate local and Athens accounts; this service is planned to switch from a pilot service to a full production service university-wide ready for the next academic session (Sept 2006);
2. Our implementation allows Shibboleth enabled users to access Shibboleth and Shibboleth-Athens based resources;
3. Thorough testing of all resources by library information specialists is in progress;
4. Our implementation has been tested with allowing access to Beilstein Crossfire (a non-web based resource) through Shibboleth and shown to work;
5. Documentation and Reports based on testing and evaluating user experience in accessing resources using Shibboleth.

A1.2.5 Lessons Learned

We at CU have learned several valuable lessons when implementing access to remote resources through the Shibboleth-Athens gateway.

Technically, we have learned that if you wish to keep within terms of your license agreements with the remote resources you are wishing to access, then you absolutely need to understand what categories of people you have in your organisation and what they have access too; and have an identity management system and directory services in place that can automatically keep track of these users and allow or deny access to Shibboleth resources accordingly in a accurate and fast manner. This kind of system is complex to produce, but can give great benefits to your organisation.

We have also learned that community pressure can help speed up adoption of technologies such as Shibboleth, as shown by the community pressuring Westlaw who have changed their plans from not touching Shibboleth to planning on implementing it once their current major upgrade project is complete.

A1.2.6 Conclusions

This Work Package was completed successfully: our implementation of a Shibboleth IdP has been set up to use the Shibboleth-Athens gateway and users can access Shibboleth-Athens resources through it. It has been thoroughly tested and shown to work with the vast majority of resources, and those which don't are having pressure applied to them to Shibboleth enable themselves.

Several issues were encountered when implementing this Work Package. The majority of them have been resolved, a few remain outstanding however.

CU is planning on switching the current pilot Shibboleth-Athens service into a full production service in time for the next academic session (Sept 2006). We plan on not issuing next year's undergraduate intake Classic Athens accounts, and will only issue Classic Athens accounts only for exceptional circumstances until our Identity Management procedures take into account the results of the working group set up that is aiming to classify all CU users and their privileges.

Appendix 2 – Test Plan

1. Introduction

This document aims to catalogue a comprehensive test plan for the Cardiff University Shibboleth implementation that is a result of the Adoption of Shibboleth for Multiple Identity Management Applications (ASMIMA) project in Cardiff University.

The objective of the testing effort described within is to ensure that the Shibboleth implementation in Cardiff University is ready for full-scale usage within the University: that it is able to perform in a resilient manner when conceivable faults occur within the implementation; and that it able to function adequately when both expected and (reasonably) extreme levels of usage are placed upon it.

The intended audience for this document is the staff in Information Services (INSRV) in Cardiff University, and anyone with an interest in creating a resilient Shibboleth implementation in a large scale organisation.

2. ASMIMA Overview

Cardiff University has secured funding to develop and implement Shibboleth as part of a coordinated series of JISC projects. Shibboleth is a federated authentication mechanism developed as an open standard by the Internet2/Middleware Architecture Committee for Education (MACE). It enables sites that hold user authentication details to securely verify the identity of individuals wishing to access content at another site that requires authentication. Core to the shibboleth design is that fact that user authentication tokens do not need to be replicated or reproduced at the content provider.

Shibboleth uses XML and, more specifically, SAML (Security Assertion Markup Language) for inter-server communications on user authentication and access rights. In addition, PKI (Public Key Infrastructure) is used to securely verify the participants in a Shibboleth federation.

The Shibboleth implementation at Cardiff University is intended principally to provide authentication as an “Identity Provider” to the EduServ Athens content “Service Provider”. EduServ Athens is a JISC awarded contract that provides a single point of access management and authentication services to a wide variety of academic content. Beyond this principle Athens usage, the Shibboleth implementation will be used further afield with other Shibboleth Federations and Services, including SPARTA, the UK Academic Federation currently being set up by JISC. In the medium term, the Shibboleth implementation must integrate with a Cardiff University-wide single sign on and security system.

3. Related Documents

- Project Plan - \\Shrinsrv\INSRV\SHARED\Projects\Current\P05_17 Shibboleth\Project_Plan\ProjectPlan2a.doc
- Project Plan Workpackages - [\\Shrinsrv\INSRV\SHARED\Projects\Current\P05_17 Shibboleth\Project_Plan\Workpackages.doc](\\Shrinsrv\INSRV\SHARED\Projects\Current\P05_17_Shibboleth\Project_Plan\Workpackages.doc)

- IDMAN Docs - \\Shrinsrv\INSRV\SHARED\Projects\Current\P03_02 Identity Management\Project Library\Documentation

4. Shibboleth IdP Implementation Overview

This section intends to give an overview of the physical and logical setup of the Shibboleth implementation that we are going to be testing.

Shibboleth is currently implemented by ASMIMA on two identical servers (idp1.cf.ac.uk and idp2.cf.ac.uk). The outside world accesses the IdP through a CName, idp.cardiff.ac.uk, which is a Layer 4-7 switch whose function is to route Shibboleth traffic to a working IdP server, transparent to the entity accessing idp.cardiff.ac.uk. All traffic will be redirected to the same IdP server, unless the Layer 4-7 switch determines that the server isn't functioning correctly, in which case it switches to the alternative server.

The Shibboleth software running on these servers performs authentication and authorisation tasks by communicating with the eDirectory FARAWAY tree, through LDAP. The FARAWAY tree holds information relating to all users of Cardiff University that the Shibboleth IdP needs to function correctly.

4.1 Hardware

Each of the identical servers is a Dell PowerEdge 1850 1U rackmount server, with dual 3.2GHz HT Intel Xeon processors and 2GB RAM, connected to the Cardiff University network at 100MB Full Duplex. Each server has a redundant power supply, the main supply coming through a UPS and the redundant supply plugged into a wall socket.

4.2 Software

The IdP part of Shibboleth itself is implemented as a Java Servlet, and can be configured to work with many different variations of software and operating system.

In our case, OS the servers are running Red Hat Enterprise Linux AS release 4 (Nahant Update 2, kernel 2.6.9), and the software details are as follows:

- Shibboleth IdP v1.3c – The actual Shibboleth software, written as a Java Servlet;
- Java 1.5.0_05 – The Java Virtual Machine (JVM) that runs the Shibboleth software;
- Tomcat 5.5.12 – The Java servlet container that interfaces the servlet code with the JVM;
- Apache 2.0.54 – Used to control the flow of internet requests to Tomcat;
- mod_jk 1.2.14 – The apache module that handles the communication between Apache and Tomcat;
- openssl 0.9.7a – Handles all of the secure connections and encryption used by the Shibboleth IdP; (the SSL certificates used by the servers are issued by a globally recognised Certification Authority – GlobalSign)
- openldap 2.2.26 – Used by the Shibboleth IdP to communicate with Cardiff University's LDAP service.

4.3 Shibboleth-Athens

When accessing the Shibboleth-Athens gateway, two attributes are passed across from our IdP to the Athens SP. These are the “eduPersonTargetedID” attribute (mapped from the “cardiffidentityno” LDAP attribute), and the “CardiffAthensOptionSet” attribute (mapped from the “cardiffathensoptionset” LDAP attribute).

The two LDAP attributes are sourced from the FARAWAY eDirectory tree. “cardiffidentityno” is the user’s unique identity number assigned by IDMAN; “cardiffathensoptionset” contains a string value, currently either “CUla#default” or “CUla#noaccess”. The first currently gives the user access to all Athens resources; the latter none.

5 Overview of Testing

Testing is to be split into two main areas – “System” testing and “Usage” testing.

The system testing area is intended to test the set up of the resilient architecture that we have designed. The main test areas will be:

- Backup and Recovery performance – testing recovery procedures in case of server failure;
- Extreme load testing – checking what throughput the software can handle through scripted access attempts;
- Resilience testing – how well the resilient architecture copes with simulated failure of one or more core Shibboleth components.

The usage testing area is intended to check that the Shibboleth software is working as required. The main test areas will be:

- Account security – checking only valid users can log in through Shibboleth;
- Athens access – checking that users can access Athens resources via Shibboleth log in; and that users with the “noaccess” permission set assigned cannot access any resources;
- Athens access (all resources) – checking that every single Athens resource works correctly;
- Athens multiple access – checking users can access multiple resources simultaneously;
- Classic/Shibboleth Athens – checking how well classic Athens and Shibboleth Athens coexist;
- Crossfire – seeing if Shibboleth login to Beilstein Crossfire will work;
- Endnote – seeing if Endnote will work with Shibboleth Athens;
- Load testing – checking the software can perform to the stated level of throughput in a real life situation – all people in the biggest computing pool room in Cardiff University attempting to log in simultaneously;
- Logout behaviour – seeing what the effect of logging out of one resource while still connected to another is;

- Monitoring – seeing whether the monitoring tools installed provide accurate statistics;
- Multiple federation access– whether Shibboleth will successfully work with multiple federations;
- Multiple simultaneous federation access – accessing multiple federations simultaneously;
- Personalisation – checking whether users can personalise their browsing experience at resources that allow it through the use of the persistent ID;
- Server session behaviour – checking what happens to current Shibboleth sessions when a server fails;
- Session behaviour – checking Shibboleth is working as intended in that credentials do not survive beyond the browser’s run-time, and upon logout/login to PC;
- Statistics – seeing whether the usage statistics (both on our end and those provided by Athens match the real usage).

5.1 Test Roles

The following personnel will take part in the testing of Shibboleth:

- INSRV Technical Staff (ROS) – “Dev Eng”;
- A set of personnel, list drawn up by INSRV (SAS) with expertise in each Athens resource - “Expert User Group”;
- Large group of users at the start of a training class in large computing pool room – “Tutorial Group”.

5.2 Test Approach

The following phases will occur during the testing of Shibboleth:

1. “Dev Eng” will conduct these tests as soon as the implementation is ready to be tested;
2. “Expert User Group” will be asked to test specific Athens Shibboleth-protected resources and report back their success/failure;
3. “Tutorial Group” will be asked at the start of a class to help with the testing of a new system. They will be instructed by a member of INSRV in exactly what to do, this staff member will get all users to login to Shibboleth at roughly the same time. As a “reward” for helping us in this way, member of this group of users will be able to use Shibboleth from then on – before the rest of the university.

5.3 Test Strategy

#	Testin g	Phas e	Role	Area	Description	Test	Expected Result
1	System	1	Dev Eng	Backup and Recovery	Test recovery procedure	Kill live server in some way (e.g. deleting core	Systems team should restore system from

						system files)	backup within an acceptable amount of time
2	System	1	Dev Eng	Extreme Load Testing	Test the throughput capability of Shibboleth implementation	Use a script to simulate logins at an increasing rate and see what the servers can handle	Should be able to at least handle our stated minimum throughput (10 logins per second)
3	System	1	Dev Eng	Extreme Load Testing	Test the concurrent capability of Shibboleth implementation	Increase hard-coded maximum threads to see what the servers can handle	Should be able to at least support our stated minimum throughput (10 logins per second)
4	System	1	Dev Eng	Resilience	Test resilient hardware	On each server, simulate power failure (pull relevant power cord)	UPS keeps server running
5	System	1	Dev Eng	Resilience	Test resilient hardware	On each server, simulate UPS failure (pull relevant power cord)	Redundant supply not plugged into UPS keeps server running
6	System	1	Dev Eng	Resilience	Test resilient hardware	On each server, simulate disk failure (pull plug from RAID)	RAID compensates server keeps running
7	System	1	Dev Eng	Resilience	Test L4-7 switchover	On each server, simulate software failure (kill vital processes: Apache, Tomcat, etc.)	L4-7 switch automatically diverts requests to the other server
8	Usage	1	Dev Eng	Account Security	Check valid users can login	Attempt to login to Athens resources through shibboleth using a valid user account	Login successful, shibboleth session established
9	Usage	1	Dev Eng	Account Security	Check invalid users cannot login	Attempt to login to Athens resources through Shibboleth using an invalid user account	Login unsuccessful, no shibboleth session established
10	Usage	1	Dev Eng	Account Security	Check valid, but suspended (disabled) accounts cannot login	Attempt to login to Athens resources through Shibboleth using a valid user account that has	Login unsuccessful, no shibboleth session established

						been suspended (disabled)	
11	Usage	1	Dev Eng	Athens access	Check that a given username can log in to Athens protected resources	Log in with a normal user's credentials, attempt to access an Athens protected resource	Able to access the resource
12	Usage	1	Dev Eng	Athens access	Check that a given username with the "noaccess" permission set can log in but not access Athens protected resources	Log in with the credentials of a user with the "noaccess" permission set, attempt to access and Athens protected resource	Unable to access the resource
13	Usage	1	Dev Eng	Athens multiple access	Check that accessing multiple Athens resources work in the same Shibboleth session	Login to Athens through shibboleth, access several resources	Able to access all resources without being asked for credentials after the initial login
14	Usage	1	Dev Eng	Crossfire	Check if Crossfire works with Shibboleth Athens	Load Crossfire, attempt to access Athens-Shibboleth resources	Crossfire should work
15	Usage	1	Dev Eng	Endnote	Check if Endnote works with Shibboleth Athens	Load Endnote, attempt to log in to Athens using Athens-Shibboleth credentials	Endnote should be able to access Athens resources
16	Usage	1	Dev Eng	Logout behaviour	Check behaviour when logging out of one resource when connected to another	Login to two Athens resources concurrently, log out of one, check whether still logged into other	Should stay logged in on other resource
17	Usage	1	Dev Eng	Monitoring	Check accuracy of statistics produced by reporting tools	Use Shibboleth server a specified amount, check logs to confirm, then check reported statistics and compare accuracy	Statistics should be accurate and correct
18	Usage	1	Dev Eng	Multiple Federation Access	Check Shibboleth works correctly with multiple federations	Attempt to access shibboleth protected resources on different federations in the same Shibboleth session	Gain access to Shibboleth protected resources, providing credentials the initial time only session

19	Usage	1	Dev Eng	Multiple Simultaneous Federation Access	Check Shibboleth works correctly when accessing multiple federations simultaneously	Login to shibboleth, then attempt to access resources of two federations (e.g. Athens and SDSS) simultaneously	Gain access to Shibboleth protected resources, providing credentials the initial time only
20	Usage	1	Dev Eng	Personalisation	Check that users can personalise resources using the Shibboleth persistent ID	Login to an Athens resource with personalisation options, apply personal settings, re-login in a different Shibboleth session	Personalisation of the Athens resources should work
21	Usage	1	Dev Eng	Server Session behaviour	Check current session loss	Login to Shibboleth, simulate failure of current idp server, then attempt to access Athens resource	Shibboleth session should have persisted, resource viewable
22	Usage	1	Dev Eng	Server Session behaviour	Check currently logging-in session loss	Login to Shibboleth, during login process (post AuthN and pre AuthZ) simulate failure of current idp server	Shibboleth session should not be properly established, resulting in having to re-provide credentials
23	Usage	1	Dev Eng	Session behaviour	Check correct credential destruction behaviour	Login to Shibboleth, close browser, attempt to access Athens resource	Prompted to re-enter credentials
24	Usage	1	Dev Eng	Session behaviour	Check correct credential destruction behaviour	Login to Shibboleth, logout of computer, re-login, attempt to access Athens resource	Prompted to re-enter credentials
25	Usage	1	Dev Eng	Statistics	Check accuracy of provided statistics	Keep track of Shibboleth usage during earlier tests, compare to reported statistics	Usage statistics should be accurate and correct
26	Usage	2	Expert User Group	Athens access	Check that all Athens resources work with our Shibboleth implementation	Assign each Athens resource to someone (list provided by SS) and check they all work	Able to access all resources
27	Usage	3	Tutorial Group	Load Testing	Test the throughput capacity of our	Have a large group of users simultaneously	All logins should be successful.

					Shibboleth implementation in a real life situation	attempt to sign in to Athens using Shibboleth	
--	--	--	--	--	--	---	--

5.4 Assumptions and dependencies

For this testing to take place, the following must be in place:

- The Shibboleth implementation must be ready and working (including backup procedures, resilient implementation and software);
- Every user who is to take part in the testing (“Dev Eng”, “Expert User Group” and “Tutorial Group” must have the “cardiffidentityno” and “cardiffathensoptionset” attributes set correctly on the FARAWAY tree;
- The “Expert User Group” list must be drawn up and the members of it shown how to use Shibboleth
- The INSRV staff members taking charge of the “Tutorial Group” testing must be given a set procedure to follow outlining the process of the test to take place;

5.5 Scope and limitations of testing

The test strategy defined above, while attempting to test the Shibboleth implementation as thoroughly as possible, do not (and can not) test the implementation with a usage level distribution. However, since our testing should test the implementation at usage levels far above the expected, this is not seen as a problem.

5.6 Test environment

All tests by “Dev Eng” will be conducted on INSRV machines, running the stock INSRV Windows XP Image (820), Linux, and Mac OS X (10.4). All tests will be run in Microsoft Internet Explorer 6, Mozilla Firefox (1.0.x and 1.5 RC x), Opera (7 and 8), and Safari (as appropriate to the platform). Additionally, the appropriate tests will be run from on and off the Cardiff University network (from the home broadband connection of a selection of INSRV staff, in order to check it works on multiple ISPs – including AOL).

The tests run in the computing pool room will be on the stock INSRV Windows XP Image (820), running the web browser of the user’s choice – to make it as realistic as possible.

5.7 Test environment validity analysis

As the tests are going to be run on pretty much every major supported web browser and OS combination any Cardiff University member is likely to use, the test environment is pretty close to future real life use of Shibboleth.

5.8 Outline of system logging capabilities

The following tools will be used to log the testing as it happens, and analyse the results of the testing:

- Built in Shibboleth logging – the log files will show all Shibboleth usage;
- Built in Tomcat logging – the log files will show all Tomcat processes;
- Built in Apache logging – the log files will show all connection attempts;
- Built in mod_jk logging – the log files will show all information relating to Apache – Tomcat interaction;
- Built in linux system logging – the log files will show all information relating the processes happening on the servers;
- MRTG – MRTG will be used for live system monitoring, which includes server information (CPU/Memory/Disk/etc usage) as well as Shibboleth usage through Perl scripts which parse the Shibboleth log files;
- Athens provided statistics – Statistics provided by EduServ that show usage of the Shibboleth – Athens gateway from their end of the chain.

5.9 Personnel pre-training needs

The following training needs to take place before the testing can be completed:

- The members of the “Expert User Group” need to be shown how to Login to Shibboleth and access the particular Athens Resource assigned to them;
- The staff member(s) responsible for leading the “Tutorial Group” needs to be trained in the usage of Shibboleth and what common errors may be received, so they are able to quickly diagnose any errors that may occur during the test.

Appendix 3 – Test Results

#	Description	Test	Expected Result	Actual Result
1	Test recovery procedure	Kill live server in some way (e.g. deleting core system files)	Systems team should restore system from backup within an acceptable amount of time	Delayed. Test put on hold until servers re-racked into final locations
2	Test the throughput capability of Shibboleth implementation	Use a script to simulate logins at an increasing rate and see what the servers can handle	Should be able to at least handle our stated minimum throughput (10 logins per second)	Delayed. Test put on hold until servers re-racked into final locations
3	Test the concurrent capability of Shibboleth implementation	Increase hard-coded maximum threads to see what the servers can handle	Should be able to at least support our stated minimum throughput (10 logins per second)	Delayed. Test put on hold until servers re-racked into final locations
4	Test resilient hardware	On each server, simulate power failure (pull relevant power cord)	UPS keeps server running	Delayed. Test put on hold until servers re-racked into final locations
5	Test resilient hardware	On each server, simulate UPS failure (pull relevant power cord)	Redundant supply not plugged into UPS keeps server running	Delayed. Test put on hold until servers re-racked into final locations
6	Test resilient hardware	On each server, simulate disk failure (pull plug from RAID)	RAID compensates server keeps running	Delayed. Test put on hold until servers re-racked into final locations
7	Test L4-7 switchover	On each server, simulate software failure (kill vital processes: Apache, Tomcat, etc.)	L4-7 switch automatically diverts requests to the other server	Successful. Shib service stopped on each server, Netscaler L4-7 switch detected the service was down within 5 seconds and compensated correctly by only redirecting to the working server.
8	Check valid users can login	Attempt to login to Athens resources	Login successful, shibboleth	Successful. Login attempted with a valid username/password, login granted, access to an Athens resource successful.

		through shibboleth using a valid user account	session established	
9	Check invalid users cannot login	Attempt to login to Athens resources through Shibboleth using an invalid user account	Login unsuccessful, no shibboleth session established	Successful. Login attempted with valid username/password, error page correctly shown.
10	Check valid, but suspended (disabled) accounts cannot login	Attempt to login to Athens resources through Shibboleth using a valid user account that has been suspended (disabled)	Login unsuccessful, no shibboleth session established	Successful. Login attempted with valid username/password, error page correctly shown.
11	Check that a given username can log in to Athens protected resources	Log in with a normal user's credentials, attempt to access an Athens protected resource	Able to access the resource	Successful. Login attempted with a valid username/password, login granted, access to an Athens resource successful.
12	Check that a given username with the "noaccess" permission set can log in but not access Athens protected resources	Log in with the credentials of a user with the "noaccess" permission set, attempt to access and Athens protected resource	Unable to access the resource	Successful. Login attempted with a valid username/password, login granted, athens displays an access not allowed message
13	Check that accessing multiple Athens resources work in the same Shibboleth session	Login to Athens through shibboleth, access several resources	Able to access all resources without being asked for credentials after the initial login	Successful. Login through Shibboleth-Athens, access to random Athens resource successful, subsequent access to different random Athens resource successful.
14	Check if Crossfire works with Shibboleth Athens	Load Crossfire, attempt to access Athens-Shibboleth resources	Crossfire should work	Successful. Registered local username with beta Crossfire service as a DA username, attempted access, access granted, albeit very slowly.
15	Check if Endnote works with Shibboleth	Load Endnote, attempt to log in to Athens using Athens-	Endnote should be able to access Athens	Unsuccessful. EndNote not Shibboleth enabled yet.

	Athens	Shibboleth credentials	resources	
16	Check behaviour when logging out of one resource when connected to another	Login to two Athens resources concurrently, log out of one, check whether still logged into other	Should be logged out of all Athens resources, but stay logged in to Shibboleth	Successful. Logged out of Athens on one resource, attempted to access another Athens resource, prompted for Shib-Athens login.
17	Check accuracy of statistics produced by reporting tools	Use Shibboleth server a specified amount, check logs to confirm, then check reported statistics and compare accuracy	Statistics should be accurate and correct	Partly Successful. Overall recorded statistics correct, monitoring graphs from MRTG however calculate averages and therefore lose some small numbers.
18	Check Shibboleth works correctly with multiple federations	Attempt to access shibboleth protected resources on different federations in the same Shibboleth session	Gain access to Shibboleth protected resources, providing credentials the initial time only	Successful. Logged into Shib-Athens federation, accessed Athens resource, then attempted to access a resource on the SDSS federation. Prompted to choose Cardiff IdP by WAYF, but then automatically logged in as credentials saved in browser session.
19	Check Shibboleth works correctly when accessing multiple federations simultaneously	Login to shibboleth, then attempt to access resources of two federations (e.g. Athens and SDSS) simultaneously	Gain access to Shibboleth protected resources, providing credentials the initial time only	Successful. Logged into Shib-Athens federation, accessed Athens resource, then accessed SDSS resource in another browser tab. Access successfully granted on both.
20	Check that users can personalise resources using the Shibboleth persistent ID	Login to an Athens resource with personalisation options, apply personal settings, re-login in a different Shibboleth session	Personalisation of the Athens resources should work	Delayed. Awaiting results from Information Specialists. See Appendix 4 – Full results for details. Summary: <ul style="list-style-type: none"> • Almost all personalisation works, albeit with a few caveats – see Appendix 1.2 for details
21	Check current session loss	Login to Shibboleth, simulate failure of current idp server, then attempt to access Athens	Shibboleth session should have persisted, resource viewable	Successful. Logged in through Athens, accessed Athens resource, restarted Shibboleth service on both servers, accessed Athens resource, session successfully persisted and resource viewable

		resource		
22	Check currently logging-in session loss	Login to Shibboleth, during login process (post AuthN and pre AuthZ) simulate failure of current idp server	Shibboleth session should not be properly established, resulting in having to re-provide credentials	Successful. Logged in through Athens, simulated failure post AuthN, pre AuthZ. Athens error message displayed as expected.
23	Check correct credential destruction behaviour	Login to Shibboleth, close browser, attempt to access Athens resource	Prompted to re-enter credentials	Successful. Logged in through Athens, accessed Athens resource, restarted browser, attempt to access resource, prompted to enter credentials
24	Check correct credential destruction behaviour	Login to Shibboleth, logout of computer, re-login, attempt to access Athens resource	Prompted to re-enter credentials	Successful. Logged in through Athens, accessed Athens resource, logged out of computer and re-logged in, attempt to access resource, prompted to enter credentials
25	Check accuracy of provided statistics	Keep track of Shibboleth usage during earlier tests, compare to reported statistics	Usage statistics should be accurate and correct	Successful. Compared recorded usage statistics with direct Shibboleth application logfiles. Overall recorded statistics correct.
26	Check that all Athens resources work with our Shibboleth implementation	Assign each Athens resource to someone (list provided by SS) and check they all work	Able to access all resources	See Appendix 4 – Full results for details. Summary: <ul style="list-style-type: none"> • 85/92 resources work correctly • 7/92 have authentication problems currently
27	Test the throughput capacity of our Shibboleth implementation in a real life situation	Have a large group of users simultaneously attempt to sign in to Athens using Shibboleth	All logins should be successful.	Delayed. <i>Test put on hold until servers re-racked into final locations</i>

Appendix 4 – Full survey of CU's Athens Resources

Resource	Login success	Gateway compliant	Notes
ABC-CLIO Serials Databases	Yes	Yes	
Adept Scientific - Adept4Education	Yes	Yes	
BANKSCOPE	Yes	Yes	
BMJ Journals	Yes	Yes	
Blackwell-Synergy.com	Yes	Yes	
Bristol Biomedical Image Archive	Yes	Yes	
British Standards Institution	Yes	Yes	
Butterworths Legal Updater	No	No	To be withdrawn and replaced by Lexis Nexis Butterworths
CHCC Historical Censuses Collection	Yes	Yes	
CHEST Higher Education Site Contacts	Yes	Yes	
CPPE	Yes	Yes	
CSA Illumina	Yes	Yes	
Cambridge Journals Online	Yes	Yes	
Census Dissemination Unit	Yes	Yes	
Census Geography Data Unit (UKBORDERS)	Yes	Yes	
Census Interaction Data Service (CIDS)	Yes	Yes	
Census Learning Resources	?	Yes	Although listed as an Athens protected resource, does not appear to require authentication.
Census Registration Service	Yes	Yes	
Census: Samples of Anonymised Records	Yes	Yes	
CrossFire self-teach modules (MIMAS-XFT)	Yes	Yes	
Dialog Education@Site	No	No	Have no plans to shibbolise Dialog, but are moving their content to DataStar and will be making that compliant instead. They tell us that they have re-started their

Resource	Login success	Gateway compliant	Notes
			discussions on this, and they were due to meet with Athens in December. No date for implementation given yet. We see no difficulty in moving from Dialog to DataStar when the time comes.
Digimap Historic Map Collection	Yes	No	Digimap have a gateway compliant test system and have successfully tested this with a Cardiff account. Resource is still listed as non Gateway compliant as it is not yet in production service
Digimap Ordnance Survey Data Collection	Yes	No	See above
EBSCOhost EJS	Yes	Yes	
EBSCOhost databases	Yes	Yes	
EDINA BIOSIS	Yes	Yes	
EDINA INSPEC	Yes	Yes	
EEBO	Yes	Yes	
ESDS International	Yes	Yes	
EXTENZA	Yes	Yes	
Economic and Social Data Service (ESDS)	Yes	Yes	
Education Media OnLine	Yes	Yes	
Education Media OnLine medical-restrict	Yes	Yes	
Emerald Fulltext	Yes	Yes	
Engineering Village 2	Yes	Yes	
ESDU Data	Yes	Yes	
European Sources Online	Yes	Yes	
FAME	Yes	Yes	
Global Market Information Database(GMID)	Yes	Yes	
Grove Art Online	Yes	Yes	
Grove Music Online	Yes	Yes	

Resource	Login success	Gateway compliant	Notes
HEFCE Extranet	Yes	Yes	
HeinOnline	Yes	Yes	
IHS Technical Indexes Info4Education	Yes	Yes	
IOP's Electronic Journal Service	Yes	Yes	
Ingenta Select	Yes	Yes	
IngentaConnect	Yes	Yes	
JUSTIS Daily Cases	Yes	Yes	
JUSTIS Law Reports (eLR)	Yes	Yes	
JUSTIS Law Reports Digest	Yes	Yes	
JUSTIS Weekly Law	Yes	Yes	
JustCite	Yes	Yes	
Keynote	Yes	Yes	
KnowUK Database (Chadwyck-Healey)	Yes	Yes	
Lawtel	Yes	Yes	
LexisNexis Butterworths	Yes	Yes	
LexisNexis Professional and Executive	No	No	Not planning to Shibbolise LN Executive or LN Professional. Executive is being moved to a new service which will be compliant, while the content of Professional will be migrated to LN Butterworth's, which is already compliant. This is due by August 2006.
Literature Online (Chadwyck-Healey)	Yes	Yes	
MIMAS Landmap	Yes	Yes	
MIMAS Landmap Mediterranean	Yes	Yes	
MIMAS LitLink	Yes	Yes	
Macromedia Athens Student Store	?	Yes	
Mintel Reports	Yes	Yes	
NetLibrary	Yes	Yes	
New Scientist	Yes	Yes	

Resource	Login success	Gateway compliant	Notes
OCLC FirstSearch Service	Yes	Yes	
Ovid Online	Yes	Yes	
Oxford Dictionary of National Biography	Yes	Yes	
Oxford English Dictionary Online	Yes	Yes	
Oxford Journals	Yes	Yes	
Oxford Reference Online	Yes	Yes	
ProQuest	No	No	Have told us that they are investigating Shibboleth, but no date for implementation yet.
RSC Journals Archive	Yes	Yes	
SAGE Online	Yes	Yes	
SCOPUS	Yes	Yes	
ScienceDirect	Yes	Yes	
SilverPlatter Arc2	Yes	Yes	
SwetsWise	Yes	Yes	
TRILT	Yes	Yes	
TVTimes Project 1955-1985	Yes	Yes	
Thomson Gale Databases	Yes	Yes	
UK JSTOR Mirror Service	No	No	Due first quarter 2006. However, experiencing delays due to loss of a key member of staff, but hope to be back up to speed shortly.
WILSONWEB	Yes	Yes	
Web of Knowledge	Yes	Yes	
Westlaw UK	No	No	Westlaw have submitted a business request. No formal statement on timetable but have said that they will endeavour to implement by end 2006. They will make generic accounts available to CU as a

Resource	Login success	Gateway compliant	Notes
			workaround.
Wiley InterScience	Yes	Yes	
ZETOC - BL Electronic Table of Contents	Yes	Yes	Offers the user the ability to move their existing alerts to a new username
internurse.com	Yes	Yes	

Windows Applications

Belstein Crossfire	Yes	No	We have successfully tested access to CrossFire via Shibboleth and the Shib-to-Athens gateway, although currently still listed as Non-Gateway compliant on Athens website as it is not yet in production service
EndNote	No	No	CU has submitted a development request. They say that Shibboleth is something they have on their "radar" and are certainly considering for both EndNote and EndNote Web. Of the databases that are searchable by EndNote only one insists on a username and password: BIOSIS

Appendix 5 – Installation, Configuration, Resiliency and Monitoring guides produced

See the project website at:

<http://www.cardiff.ac.uk/insrv/shibboleth> -> “For Developers” section