A Distributed Authentication Infrastructure for Western Australian Universities

Peter Green* and Alex Reid
* Curtin University of Technology
GPO Box U1987
Perth, Western Australia 6845 AUSTRALIA
E-mail: p.green@curtin.edu.au

The Western Australian Group of University Librarians (WAGUL) had identified the increasing diversity and cost of scholarly resources as having a major impact on the ability of university libraries to deliver those resources. Various avenues for the sharing of these resources were pursued as strategic objectives. One major impediment to this sharing was the need for reliable, automated inter-institutional authentication. WAGUL were successful in 2000 in obtaining a substantial grant from the Commonwealth Development Pool (CDP) to fund a major authentication project to address this strategic objective. The aim of the authentication project was a distributed authentication infrastructure for the five Western Australian Universities. The main components of the solution have now been implemented, and the technical feasibility demonstrated. It remains to populate the directories with live data, which it is expected will have be completed during 2003, and to begin using the distributed authentication infrastructure.

This paper describes the following aspects of the Project:

1. The motivation behind the Project
2. The specifications describing the functionality required of the solution
3. The stages of the solution implementation
4. Technical aspects of the solution
5. Issues of a technical and political nature that had to be overcome
6. Possible uses to which the infrastructure will be put

1 Introduction

Trust is not commonly exchanged amongst competitive entities and thus creation of a distributed authentication infrastructure, where the distribution crosses institutional boundaries, was always going to be challenging, and the challenges would be technical and political. This paper attempts to document the solutions found and describe the challenges still to be overcome.

2 Motivations

The motivations that led to the creation of the WAGUL Libraries Authentication project (WALAP) are amply described in an earlier paper (Green, 2002). It is sufficient to note here that the growing dominance of electronic resources has caused a fundamental shift in how libraries do their business, in particular how they have become de facto gatekeepers rather than guides and curators. Authentication of patrons has become a necessary evil in a profession more famed for advocating open access to information than in restricting its availability. In addition to pressure from publishers, vendors and statutory authorities to control access to resources there are the additional pressures to do more with less and to

ISBN 1-876346-47-7

integrate library authentication with university authentication. These combined pressures provide a powerful motivation for the distributed authentication solution adopted by WALAP.

3 Specifications

The specifications published via a Request for Proposal (WAGUL Authentication Project, 2001) attempted to describe the expected functionality and performance required to meet the project requirements without dictating the solution. This is normal practice when the requirements are understood but the solution is not. This approach allowed vendors to provide one or more solutions in their response as they attempted to balance sometimes conflicting requirements.

The specifications were as follows:

3.1 Interoperability

3.1.1 University systems

Respondents should address the issue of interoperability with existing systems at each University using the information provided in section 6. In particular the Respondent should comment on the limitations or customisation that would be required to attain interoperability. The ability of the proposed system to integrate with existing systems is critical to its success. It should be noted that full interoperability with all nominated systems is not expected to be delivered as an outcome of the Proposal. It is anticipated that additional work, outside of the scope of the RFP, may be required to make some systems interoperable with the Respondent's Proposal. The writing of technical specifications to achieve interoperability is one of the expected deliverables of the Proposal. Respondents should not make any assumptions that existing directory services will be replaced by this Proposal. Respondents should address issues of interoperability with the existing directory services.

3.1.2 Open Standards

It is expected that the proposed solution will be compliant with Open Standards in general and in particular but not limited to:

Authentication Protocols (e.g. Kerberos, Radius)
Directory Services (e.g. X.500, LDAP)
Public Key Infrastructure (e.g. X.509)
Other communications protocols such as TCP/IP

3.2 Continuity

It is envisaged that the proposed solution will form an integral part of many of the Universities' core business activities. As such, uninterrupted access to the system is critical. The Proposal should indicate the mechanisms that would be used to meet the need for continuous access.

3.2.1 Level of Availability

The Respondent should indicate the level of availability that can be expected and provide an indication of how this will be achieved. The level of availability excludes failures to the network or other systems beyond the scope of the proposed solution, but should include scheduled downtime for routine maintenance.

3.2.2 Robustness

The proposed solution should be designed so that the unavailability or failure of one segment of the proposed system should not cause the failure or unavailability of the entire system. For instance the failure or unavailability of the system at one University should not impact upon the functioning at any other University.

ISBN 1-876346-47-7 © 2003 EDUCAUSE IN AUSTRALASIA

3.2.3 Modularity

While the proposed solution is to be designed to allow all five WAGUL partners to work together, the proposed solution should allow for withdrawal by one or more partners, after implementation, without compromising the remainder of the system. The withdrawing partner should continue to enjoy the full functioning of the system at their local level.

3.2.4 Recovery

Respondents should indicate how data would be protected and recovered in the event of hardware or other failure.

3.3 Sustainability

An indication is required as to the ongoing costs of maintaining the system. The actual provision of ongoing maintenance is not expected to form part of the Proposal, but Respondents should include a Proposal to supply ongoing maintenance beyond the first year of operation.

3.3.1 Staffing

An indication is required of the expected number and type of staff required to maintain and administer the system at each University. It is not the expectation of WAGUL that there will be any support staff shared between the Universities.

3.3.2 Licence Costs

An indication is required of ongoing licence or other costs that are not included in the total cost as presented by the Respondent. There is a strong preference for all costs, including license costs imposed by third parties, to be included in the implementation phase. Alternative payment methods could be proposed.

3.3.3 Intellectual Property.

It is a requirement that ownership of any Intellectual Property, such as copyright rights, trade marks, patents and confidential information, created in the process of designing, constructing and implementing the system be assigned to WAGUL and the participating Universities. The Respondent must also ensure that any Intellectual Property created by sub-contractors in the process of designing, constructing and implementing the system will be assigned to WAGUL and the participating Universities. The Respondent will also be responsible for obtaining from its employees and sub-contractors all consents necessary to give WAGUL and the participating Universities unrestricted use of the system, such as moral rights consents.

3.4 Functionality

The functionality that might be delivered by the proposed solution is not limited to the following and Respondents may include information on additional functionality.

3.4.1 Client privacy

The confidentiality and security of data held in the directory service is an important consideration. The Respondent should indicate the methods used to ensure strong protection from unauthorised access. The response in this area should address security in the administrative functions as well as security in machine-to-machine communications.

3.4.2 Partitioning

The WAGUL Universities will generally not want to share access to their own data or to allow administrative access to other partners. The Respondent should indicate how the system could be partitioned or otherwise organised for security and administration to meet this requirement.

3.4.3 Authentication methods

The proposed solution should allow for different authentication methods and not be constrained by any particular authentication system or technology. The authentication of users should not be restricted to any particular physical location. For instance a student visiting another University should be able to be authenticated. The Respondent should indicate any strengths or limitations in this area.

ISBN 1-876346-47-7 © 2003 EDUCAUSE IN AUSTRALASIA

3.4.4 Accounting functionality

The proposed solution should allow for accounting functionality (eg general usage statistics, traffic accounting, charging). The implementation of such functionality is not a requirement of the Proposal but Respondents should indicate the availability of such functionality, detailed information on the capabilities of the accounting functionality and the estimated costs associated with implementing the functionality.

3.4.5 Granularity of authorisation

It is a requirement that the proposed solution should allow for granularity of authorisation. Granularity indicates that authorisation can be determined on individual attributes or combination of attributes. For instance the ability to authorise a student based on a unit of enrolment and enrolment status or a staff member based on substantive position and campus location. Respondents should indicate the how this will be achieved.

3.4.6 Architecture

The Respondent is required to give a clear indication of the logical and physical organisation of their proposed solution. This should indicate the server architecture and the data architecture. The Respondent may wish to propose more than one configuration. It is expected that WAGUL participants will retain control over their own data and the Respondent should indicate how this would be achieved.

3.5 Extensibility

It is envisaged that the proposed solution will have the flexibility to adapt to changing requirements over time. The proposed solution is seen as a starting point for WAGUL and not the delivery of a comprehensive solution.

3.5.1 Scalability

The proposed solution should be scalable beyond the current geographic and size limitations of WAGUL. This may include future partnerships outside the existing WAGUL group or the addition of campuses beyond the current locations. The additional partnerships may not be with WAGUL as a whole but with individual institutions. The proposed solution should be scalable beyond the current volume of data traffic without adverse effect.

3.5.2 Schema must be extensible.

It is a requirement that the directory schema will have a global aspect for data that WAGUL wish to make available for common authentication and a local aspect that will reflect the different local needs of each participant. Objects in the schema should be capable of having repeatable attributes, for example a single student would have multiple entries for an attribute *Unit Code*. In addition the schema should be capable of modification and extension over time at both the global and local level. The Respondent should indicate how these requirements will be met and any implications or costs of future extension or modification.

3.5.3 Directory Structure must be flexible

It is expected that the Directory Structure can hold various types of objects and the relationships between them. These objects would include people, places and things. The Respondent should indicate any limitations in this area or any implications of adding a new type of object to the schema at a later stage.

3.6 Administration

It is a requirement that administration of the directory service can be managed separately by each WAGUL University and then locally devolved in a secure and managed way. The Respondent should indicate any implications or limitations created by this requirement for devolved administration, including limitations such as desktop configuration or bandwidth requirements.

There were a number of proposals that met the specifications and a lengthy process of short listing, respondent presentations and evaluation took place. The proposal from Computer Associates was recommended as the one that best met the specifications and delivered value for money, though there were clearly other respondents capable of delivering a solution. Computer Associates would supply services to design and implement the directory structure as well as supplying the directory software (eTrust Directory).

ISBN 1-876346-47-7 © 2003 EDUCAUSE IN AUSTRALASIA

4 Implementation

The implementation involved stages as follows:

4.1 Phase 1a - Planning & Design

In the beginning much time was spent in getting the design right. This was an aspect of the project that drew heavily on the time and expertise of staff from all five universities, being facilitated by the technical staff from Computer Associates. At the end of this stage the design was well documented (WALAP, 2002a) and the fruit of the time spent at this stage was the minimal change required to the design through the next phases.

4.2 Phase 1b - Prototype build

Once the design had been agreed a prototype was built on machines within Computer Associates' Perth laboratory. This was an essential element of ensuring that the design would work in a test environment. This work included the writing of the web application.

4.3 Phase 2 – Edith Cowan University Production Build

An initial production build took place at one of the universities.

4.4 Phase 3 - Rollout Design

Once the first production build had successfully been concluded, planning was required for the production build at the other four sites. This required arranging time and resources at each site so that the local build could progress without delays, but also so that knowledge transfer could occur with the local technical staff.

4.5 Phase 4 through 7 - Rollout to WAGUL members

The production build at the remaining four universities was conducted sequentially. This allowed for the same Computer Associates technical staff to be used and thus a better quality build and knowledge transfer to occur.

4.6 Phase 8 - Final Integration Testing

Once testing was conducted at each site and the local builds signed off, a final test of the entire system was conducted. A formal series of test cycles was conducted with auditing by the project team to ensure that response times, failover and functionality were delivered as required.

4.7 Phase 9 - Skills transfer and Training

Along with the skills transfer at each site during the roll out, formal product classes in eTrust Directory administration were conducted with a number of technical staff from each university.

From start to finish the implementation took five months and ran within budget expectations. At the end of this time each site had a functional, fully tested directory service with demonstrated distribution but yet to be populated with live data.

5 Technical Solution

The solution adopted by the project had several elements.

5.1 Directory Architecture

To fulfill the key business requirements for the directory solution to be modular and for each university to have their own administration facilities, the directory solution adopted the following distributed directory architecture with each university having their own servers to host their respective directories. This is illustrated in Figure (1).

ISBN 1-876346-47-7

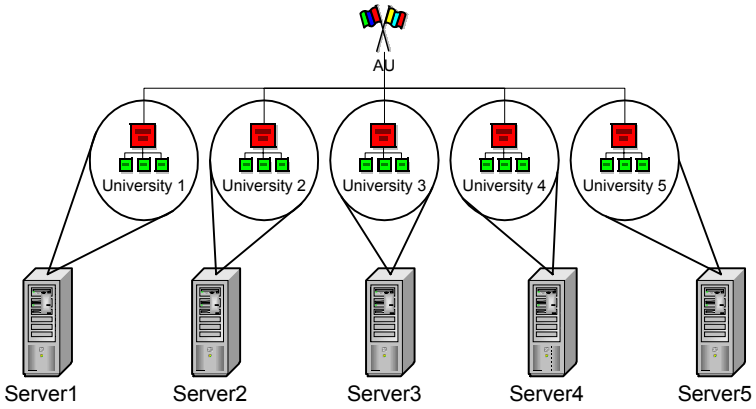

Figure 1

There is a Directory Service Agent (DSA) operating on each machine. These are organized in such a way as to provide for a single Directory Information Tree. Each University controls their part of the Directory Information Tree, delegating administrative responsibility locally using access controls.

If a query is made of the Directory Service, it is passed to the appropriate Directory Service Agents for processing. It should be noted that at no stage does any University have access to the underlying data of any other University. No data is copied from the server that supports one University's part of the Directory Information Tree to another's – this is distribution in its purest form. To achieve redundancy the directories are replicated between DSA's on the production and failover servers.

5.2 Directory Information Tree (DIT)

The DIT chosen by the project is represented in Figure (2).

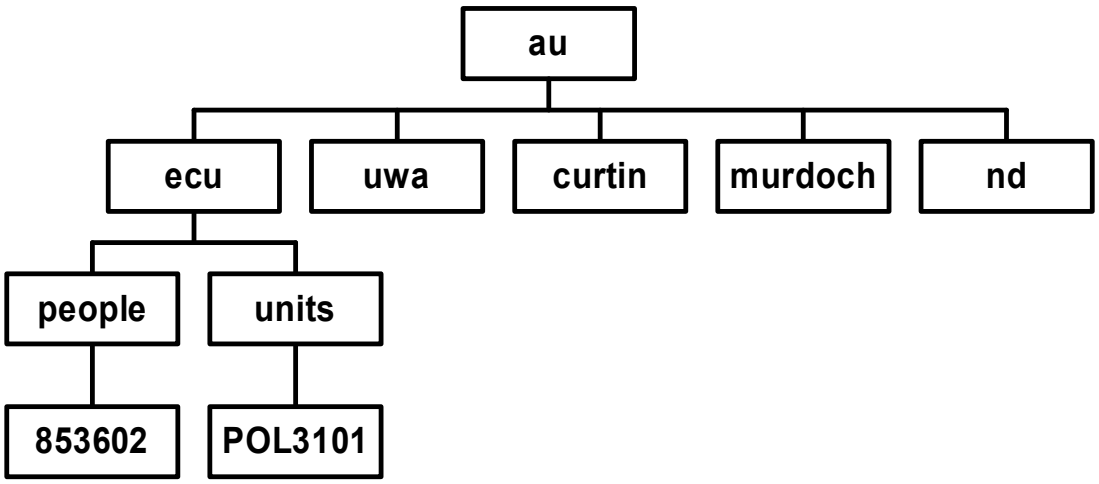

Figure 2

At the top level is the country, at the next the university and the third level comprises objects to hold people or units. The object classes are country (c), organization (o), organization unit (ou) and common name (cn) with the 'name' of the person being an ID number. Thus a person object could have a distinguished name of c=au, o=ecu, ou=people, cn=856302 and a unit object c=au, o=ecu, ou=units, cn=POL3101.

ISBN 1-876346-47-7 © 2003 EDUCAUSE IN AUSTRALASIA

While this is simple to express and takes a few lines to describe, there was considerable discussion and debate about the DIT. The key issue was whether the tree would reflect the organizational structure of each university or not. This is a very common way of designing a tree but there were some reasons for not doing it this way. The first was that the organizational structure of universities is loose, with frequent changes in name and structure, and with persons being in more than one unit and moving between units. The universities have experience at building trees based on organisation structure for directories providing access to resources, for instance file and print services, but were not keen to use that model for a directory proving enterprise level authentication and authorisation. If there had been a need to partition and replicate parts of the tree to different locations or to devolve administration then organizational structure is a good way to build a tree. However partitioning and devolved administration were not intended outcomes and thus did not need to drive the tree design in this case. It is also clear that sites were planning to automate the updating of the directory from information held in primary systems and manual updating would be limited. Taking all of these factors into consideration, a flat tree structure was the logical conclusion.

It was agreed that use of Domain Component (dc) would be the best way to construct the high level nodes. These would reflect the domain names of each university. This use of existing domain names is a standard practice and allows the DIT to be scalable across the education sector in Australia (or elsewhere). It was agreed that a single node for people would be used (ou=people) to hold all members of the university community. This adheres to our desire for a flat structure. There was some brief discussion about having nodes for staff and students, but no advantage was seen in doing this. After some reflection a single node for units (ou=units) was also added. It emerged during discussion on the schema design that being able to point at a unit object might prove useful in the deployment of certain applications. It is also possible that a student of one university might be enrolled in a unit from a different university. The existence of a unit object would allow such a unit to be described using its distinguished name without ambiguity. This would also provide a model for additional nodes that could be created to describe objects other than people. An example could be a node to hold organizational units such as departments or schools. Such objects were seen as being outside of the current expected use of the directory but useful for future developments.

5.3 Schema

The project had looked at some existing schemas, PRIDE (PRIDE, 2001) and eduPerson (eduPerson, 2002) in particular, but decided that Pride was too specific to its own context and that eduPerson didn't add anything to inetOrgPerson (Smith, 2002) for our purposes. Thus WALAP defined two new structural object classes auEduPerson and auEduUnit as described below, while the auxiliary object class eduPerson was included in the WALAP schema. Object class eduPerson 1.5, as defined by the EDUCAUSE/Internet2 eduPerson task force, is an auxiliary object class for campus directories designed to facilitate communication among higher education institutions. It was envisaged that these attributes may be useful in the future, particularly if Shibboleth (Shibboleth, 2003) is ever implemented locally, and so they were included. It is worth noting that Shibboleth is more concerned with attributes and the exchange of information about users in a trust environment than in authentication. While WALAP has included the eduPerson attributes in the directories, these are not essential to the use of Shibboleth and the auEduPerson attributes or others could be used to inform Attribute Release Policies and thus provide the information required by Shibboleth target sites to allow access to resources. There is no overlap between the attributes defined specifically for the project and those from eduPerson. The WALAP solution could be considered a complementary architecture to that proposed by the Shibboleth project.

The auEduPerson object class was created as a sub-class of inetOrgPerson and uses some existing attributes from that class as well as defining specific attributes. Brief descriptions of the attributes for auEduPerson and auEduUnit are found in Figure (3) and Figure (4), but full documentation is available (WALAP, 2002a). An asterisk (*) indicates that an attribute is mandatory.

ISBN 1-876346-47-7 © 2003 EDUCAUSE IN AUSTRALASIA

Attribute Type	Definition	Information
cn *	RFC 2256, X.520	“856302”
sn *	RFC 2256, X.520	Surname, “Smith”
givenName	RFC 2256, X.520	“John”, “Tom”, “Ann” “Mary”
displayName	RFC 2798	“Johnny Smith”
auEduPersonSalutation	auEduPerson	“Ms”, “Mr”, “Dr”, “Prof”
auEduPersonPreferredGivenName	auEduPerson	“Tiger”
auEduPersonPreferredSurname	auEduPerson	“Sugar”
auEduPersonExpiryDate	auEduPerson	“2008-05-22”
userPassword	RFC 2256, X.520	encrypted user password
auEduPersonID *	auEduPerson	“856302”
auEduPersonType	auEduPerson	“student”, “staff”, “others”
auEduPersonSubType	auEduPerson	“undergrad”, “postgrad”
auEduPersonEmailAddress	auEduPerson	“jsmith@ecu.edu.au”
userCertificate	RFC 2256	X.509 certificate for user
auEduPersonLibraryBarCodeNumber	auEduPerson	“72891201”
auEduPersonLibraryPIN	auEduPerson	“8271”
auEduPersonActiveUnit	auEduPerson	“cn=POL3101,ou=units,o=ecu”
member	RFC 2256, X.520	“cn=soccer,ou=sport,o=ecu”

Figure 3

Attribute Type	Definition	Information
cn *	RFC 2256, X.520	“POL3101”
auEduUnitCode *	AuEduUnit	“POL3101”
auEduUnitName	AuEduUnit	“Politics and Government for beginners.”
auEduUnitActiveMember	AuEduUnit	“cn=856302, ou=people, o=ecu”

Figure 4

The project has registered the WALAP schema objects as objects defined under the Australian Academic and Research Network (AARNet).

When a university requires additional attributes to satisfy their local information needs, as is inevitable if the directory is to be used locally, a new object class can be derived from auEduPerson and additional attribute types may be defined. Also, object classes inetOrgPerson, organizationalPerson and person define many more attributes, which can be facilitated for local usage. One of the universities has already gone down this path, defining a local extension to the schema based on auEduPerson.

5.4 Web Application

To facilitate the design and development phases an application was required that would model the intended usage of the distributed infrastructure as shown in Figure (5). To this end a Java/J2EE sample web application was specified and built to demonstrate the functionality of the distributed directory infrastructure for user authentication and authorisation by a web application. This web application performs the simple tasks of taking a user ID and password and authenticating that user regardless of the home institution of the user or the location of the application. Once authenticated, a user can then be authorized to view a certain unit or not depending on their attributes. These actions take place over secured channels (SSL) and the application binds to a single DSA thus demonstrating the distribution of the process between the directories. The web application was also used to test the failover functionality.

ISBN 1-876346-47-7 © 2003 EDUCAUSE IN AUSTRALASIA

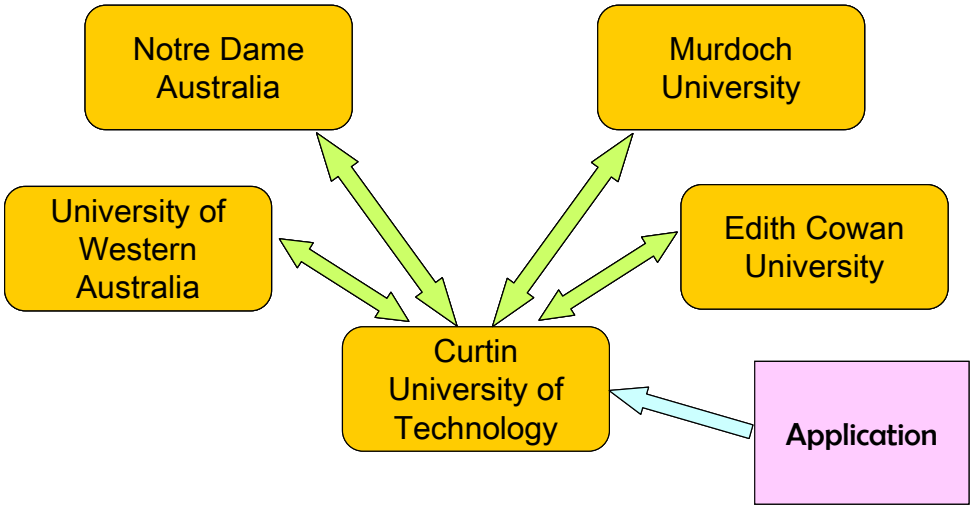

Figure 5

5.5 Hardware

Each of the sites has three servers to house three instances of the directory (with the exception of Notre Dame having two due to their much smaller student base). These are for production, failover and development. Each of the universities had the choice of Intel or Solaris platforms, though all but one chose Intel. Ownership of the servers resides locally.

5.6 Software

The eTrust Directory software provides the directory services required by the project. Each of the five sites has its own directory service, with ownership residing with the university rather than the project. Support and maintenance for three years was included in the purchase of the software. Local ownership was an important aspect of creating confidence in the directory. The vision of the project was that each site would use their directory for multiple purposes. On the one hand they would be part of the distributed directory service, while at the same time providing a level of service to their own constituents that would differ to that provided to the distributed partners. However to do this they need to know that they had complete control over their own service. This is partly achieved through allowing the ownership to reside locally but also through the design, mentioned in more detail below. The ability to control access down to the attribute level is also a critical element.

6 Issues

6.1 Contractual

The largest single delay to the project came between the recommendation of a preferred supplier and the signing of the contract with them. The nature of the project meant that while Edith Cowan University was the 'lead institution' as regards the project (being thus designated as the recipient of the funding from the Commonwealth) and was the signatory on the contract with Computer Associates, the collaborative nature of the project meant that all five universities needed to take some degree of responsibility for the contract. This proved to be an issue that required some time to resolve, given that the contract was not on a fixed cost but on a time and materials basis, thus creating uncertainty about the eventual cost. The eventual solution was an agreement between the five universities, signed by the vice-chancellors, which resolved the relationship, liability and obligation of each party. Once this agreement had been drafted with oversight by legal opinion at each university and signed, the contract with Computer Associates could be signed by ECU and work could begin. However the contract between ECU and Computer Associates didn't cover all of the work

required, particularly phases 4 through 7 (see 4.5 above) where the work was to be conducted at the four other sites. These works required individual agreements but were fixed cost agreements. Fortunately these didn't prove to be difficult and no delay was experienced to the project while they were drawn up and signed.

6.2 Trust

When the auEduPerson schema was first drawn up the workshop group took a broad view as to the sharing of information between the universities. However once the schema had been completed and all parties had time to reflect on the way in which the schema would be used, some concern arose as to the implications of sharing such sensitive data. Two steps were taken to address these concerns.

The first was the drafting of a document (WALAP, 2002b) that detailed the way in which the attributes would be populated and reducing the number of attributes that would be shared in the first instance. To do this much tighter justification was required that an attribute would be essential for authentication or authorisation and not just useful to have available. This reduced the number of attributes to be shared from eighteen to nine, though it was recognized that there was potential for the others to be shared at a future time.

The second step was the drafting of an agreement between the five universities that would cover the expected use of the distributed authentication infrastructure, detailing the actual attributes to be shared and clarifying liability and appropriate use. This Mutual Confidentiality agreement was groundbreaking in terms of formal agreement between all five universities but provides a legal platform for the degree of trust implicit in the use of the distributed authentication infrastructure. At the time of writing the final draft was being circulated.

These two steps, agreeing on the values of a reduced set of attributes and putting in place an agreement, should resolve the legitimate concerns that arose during the project and allow for full production use of the infrastructure.

6.3 Firewall

The deployment of the directory service at Notre Dame University coincided with the deployment of a new VPN infrastructure for the Catholic Education system, including Notre Dame. This had been anticipated but problems with the DSA to DSA communication had not. The firewall through which these communications traveled would rehash the IP address of the remote DSA, which had been expected, but also hash the incoming port number which was unexpected. The DSA to DSA communication, an essential part of the distribution of the directories, required a level of security that was disrupted by these two events. The IP issue could be resolved at the firewall, as expected, but the other remains an outstanding issue at the time of writing, though a solution is in the pipeline. It is understood that elements of the network are unique and the firewall problem experienced at Notre Dame shouldn't be a typical experience of working within a firewall, though it should be acknowledged that this aspect of the project is rather unique, and distribution doesn't typically cross institutional boundaries.

6.4 Automated updating

The largest outstanding issue, at the time of writing, is the automated populating of the directories at each University. For the distributed infrastructure to be of real use, the data must be live and valid. The technical method for achieving an automated update was well understood and covered in the design phase (WALAP, 2002a) and in the training. However the underlying identity infrastructure of each university is different and the updating solution must be built on the data available from staff and student systems. At the time of writing each university is still progressing this work, but expects it to be completed by the end of March 2003.

7 The Future

The project is currently moving towards a production environment. The completion of the automated updating, the signing of the Mutual Confidentiality agreement and the resolution of the firewall issue will see the final hurdles overcome. In the meantime work has begun on enabling key applications to use the distributed directory structure.

These include the following:

- Enabling the Innovative, ExLibris and Horizon library systems to take advantage of the directories at a local and distributed level using LDAP over SSL.
- Enabling LIDDAS, a major WAGUL inter-library loan project, to take advantage of the distributed infrastructure, though this work is waiting on the May 2003 release of VDX v1.3 which contains the LDAP capability required.
- Configuring the online learning application WebCT to authenticate against the directory; however, this work is waiting on the implementation of version 3.8 of WebCT at Curtin University.
- Configuring the Apache web server, using LDAP modules, to use the distributed infrastructure with the aim of preparing tutorials for staff from the five universities.

These works are being conducted by subsets of the project team with the intention of full sharing of knowledge once the objectives have been achieved.

8 Conclusion

The aim of the project was a distributed authentication infrastructure for the five Western Australian Universities. The main components of the solution have been implemented with some final key issues nearing completion. A number of applications are being prepared to take advantage of the infrastructure once it becomes available. To reach this stage, it has taken a strong collaborative effort from many staff at the five universities to overcome technical and political hurdles.

9 References

1. A. Reid (2002). “WALAP - A Distributed Library Directory Service”, *QUESTNet2002,* Gold Coast, Australia, http://questnet.scu.edu.au/uploads/33.pdf .
2. eduPerson (2002). “eduPerson Object Class”, http://www.educause.edu/eduperson/
3. M. Smith (2002). “inetOrgPerson LDAP Object Class”, http://www.ietf.org/rfc/rfc2798.txt
4. P. Green (2002). “Building a shared authentication infrastructure: a matter of trust”, *VALA2002,* Melbourne, Australia, http://www.vala.org.au/vala2002/2002pdf/40Green.pdf .
5. PRIDE (2001). “PRIDE: People and Resources Identification for Distributed Environments”, http://www.ukoln.ac.uk/metadata/pride/
6. Shibboleth (2003). “Shibboleth project”, http://shibboleth.internet2.edu/
7. WALAP (2001). “WALAP Request for Proposal”, http://john.curtin.edu.au/walap/docs/RFP_v1.1.pdf .
8. WALAP (2002a). “Directory Design Document”, http://john.curtin.edu.au/walap/docs/Design_Document_20082002.pdf .
9. WALAP (2002b). “Attribute Value Consensus Agreement”, http://john.curtin.edu.au/walap/docs/Attribute_Value_v1.1.pdf

ISBN 1-876346-47-7 © 2003 EDUCAUSE IN AUSTRALASIA