CAMPUS LINK: Educational Institutes in Synergy

Chris Porter
Department of Computer Information Systems
Faculty of ICT
University of Malta

chris.porter@um.edu.mt

Braden Borg
Department of Computer Information Systems
Faculty of ICT
University of Malta

bradenborg@gmail.com

ABSTRACT

Based on standardized protocols, Federated Identity Management is an emerging IdM approach which allows for identity sharing across institutional domains. Applied in both aeronautical and educational industries amongst many others [3], FIM is a preferred IdM technique since it allows for improved authentication workflows and enhanced security across institutional domains. This is mainly due to the improvement in identity mobility and hence, the reduction of redundant identity information across domains [6, 7]. Trust is fundamental in a federation and this paper will explore the challenges and opportunities of federating identity information across educational institutions in order to provide seamless authentication for students and academics.

This paper will demonstrate how federated identity management can help students and academics within two Maltese educational institutions, MCAST and the University of Malta. It will also show how these can share resources, specifically with the aid of secure and discrete sharing of identities. The key to achieve this is to create a trust relationship between participating entities and clear policies. Only once trust has been established can a federation be created.

A prototype is also presented to demonstrate how such a system could be designed and implemented. Note that this paper focuses on the implementation of federated systems and the way such systems work under the hood. Political issues between institutions, although crucial, are not within the scope of this paper.

Categories and Subject Descriptors

C.2.4 [**Distributed Systems]**: This paper makes full use of the following concepts: *Client/Server* design, *Distributed Applications* and *Distributed Databases*

General Terms

Management, Design, Economics, Security, Human Factors, Theory.

Keywords

Federated Identity Management, Educational Institutions, Electronic Identities, Shibboleth

1. INTRODUCTION

The goal of this project is to show how a circle of trust can work across a number of institutions irrelevant of their backend infrastructures and Identity Management technologies used.

The benefits of such a system are seen from the point of view of both users and institutional administrators.

1.1 Benefits for users:

Users can make use of Single Sign On (SSO) which allows them to authenticate once at their home institution while allowing access to any resource provided by any other educational institutions within the circle of trust. Therefore users would not need to re-register and re-enter their authentication details every time they wish to access a resource outside of their campus, if possible at all. Furthermore users may also access resources from other institutions which are otherwise restricted to registered users of such institutions. Single Sign Out is also provided, allowing the user to sign out from all the institutions he or she has accessed just by signing out from any one institution. Personal details are solely kept by their home institution and no sensitive information is shared across institutions ensuring privacy across domains. Only assertions are shared, with discrete and privacy-focused messages being passed across institutions.

1.2 Benefits for administrators:

Administrators don't need to create and manage identities for users external to the institution. This circle of trust also enhances security efficiency, since if any institution revokes/terminates a specific identity, administrators at other institutions do not need to update their own systems.

2. THE CHALLENGE

The challenge of creating an educational federation is actually that of setting up the circle of trust itself. The educational institutes would need to discuss on the policies they would like to see in place before they enter the federation. For example, an educational institution would not provide services to another institution without knowing that the way it verifies and authenticates its users' identities is secure enough for it to be trusted.

During this research other technologies have been identified which are generally used in such federations. These include

SAML, Public Key Infrastructures (PKI) and digital certificates which help in creating a secure environment when sending requests from one educational institution to another.

The prototype created is made up of several local institutions, including: the University of Malta, MCAST, ITS and ST Martin's Institute, all of which have demo websites and databases that represent them within the prototype. The University of Malta and MCAST act as both Service Providers and Identity Providers. ITS and ST Martin's on the other hand act only as Identity Providers for their students and academics. All four institutes are part of the federation allowing the use of Single Sign On and Single Sign Out for any user wishing to use a resource on the federation. For simplicity's sake the UI of the demo websites was created in a simplistic manner in order show more clearly the complex processes carried out by the underlying architecture. Usability was also a high priority for this project.

Usability testing has been carried out rigorously by means of test cases where a number of people were given access to the system in order to try out the implemented features and provide their feedback.

3. BACKGROUND

3.1 Federated Identity Management (FIM)

From an educational point of view, FIM enables educational institutes to share resources with one another in a secure manner for the benefit of their users' research and educational activities. Also, this helps in creating a user friendly environment for both the users (students and lectures) and administrators who need to setup and maintain the system, by reducing registration and authentication burdens across multiple institutions.

3.2 Roles in a Federation

Vooren [5] identifies three important roles within a FIM, which are the Identity Provider, Service Provider and Identity Selector.

3.2.1 Identity Provider

The IdP has the responsibility of authenticating users and giving them an identity which is trusted within the federation. For example if a user is from UOM then their home identity provider is UOM.

3.2.2 Service Provider

This is the organization which provides the service requested by a user. It is the SP which selects the appropriate access privileges. As an example, assume a user from UOM who wants to access a resource from institution B, then this institution (B) becomes the service provider

3.2.3 Identity Selector

An Identity selector is used when a number of Identity Providers exist which can identify a user. A user can choose an IdP in order to get himself authenticated according to the security requirements the SP puts forth through its security policy for a particular service.

There might be a situation where a user can be authenticated by more than one identity provider therefore a user can choose the most appropriate IdP for that situation.

3.3 Security Assertion Markup Language (SAML)

SAML is managed by OASIS and was created in order to standardize the exchange of security information between online organizations.

The SAML standard is based on XML, and is used "*for describing and exchanging security information* [4]. SAML Assertions are used to exchange information about entities from one party to another. SAML provides "*syntax and rules for requesting, communicating and using these SAML assertions*" [4].

3.3.1 Assertions

SAML assertions contain the subject at hand that is the entity which is being authenticated; a student, a lecturer or even the whole organization. The SAML assertion would be sent to the service provider which in turn makes use of the information appropriately. Assertion are a core component within a federation, through which relevant identity attributes are shared across entities, while respecting privacy and keeping redundancy of identity-related information in check, avoiding potential security breaches [7]. Assertions can vary in structure, depending on what is being shared across institutions, and also depending on the policies implemented, but for the scope of this paper a simplistic view will be kept.

3.4 Public Key Infrastructure

"*A Public Key Infrastructure (PKI) is the key management environment for public key information of a public key cryptographic system*." [1]

A Public Key Infrastructure (PKI) made up of hardware and software components together with policies and people allows data to be transferred over the internet in a secure manner, while ensuring its integrity and confidentiality. This is only possible through the secure creation and distribution of public and private cryptographic key pairs by a trusted authority Non-repudiation and authentication are also major functions of a PKI. The public key can be used by anyone, hence the name, but the private key is known only by the party to whom it was issued, and should be stored in a secure manner. Both keys are created following mathematical algorithms which are related but the private key can't be derived from the public key. Their relation allows the possibility of decrypting data using the private key which had been encrypted using the public key. Therefore if the data is obtained without having the private key then this could not be decrypted.

[1] http://csrc.nist.gov/groups/ST/crypto_apps_infra/pki/pkiresearch.html Retrieved June 3, 2010

3.5 Shibboleth

“The Shibboleth System is a standards based, open source software package for web single sign-on across or within organizational boundaries.” [1]

Shibboleth, an educational-oriented system, uses SAML in order to provide SSO and to exchange attributes between the IdP and SP. Shibboleth provides the users with “*extended privacy functionality*” [1] which basically allows the user to have more control of what attributes are given to different applications across different SPs.

Shibboleth tackles the following 7 issues[2]

1. The need of having a number of passwords for different applications.
2. Managing the number of different accounts for applications.
3. Third party applications can be the cause of some security issues.
4. Privacy of user information.
5. Interoperability issues across institutional boundaries.
6. Institutions can pick their authentication technology.
7. SPs can control access to their own resources

4. DEFINING CAMPUS-LINK

4.1 Problem Definition

Currently with the siloed systems implemented at UOM and MCAST, students and lecturers lack mobility when it comes to access resources from outside their campus. For example if a student from MCAST goes to UOM's library to do some research he or she can't borrow books or access the internet infrastructure since they don't have an identity provided by UOM. Synergy between institutions can be a key to enhance the usage of resources and to improve the students’ and academic’s capabilities.

Copying identity information from one institute to another isn't much of a solution, since this will create additional burden and inefficiencies on the system administrators, as well as security risks and a negative economic impact due to maintenance required on the large number of external user accounts created. What happens if another institution wants to join the community? That would introduce more problems than it would solve, such as the need of keeping all information from all institutions up to date. An institution might not need all information from a user’s profile. Also institutions can have different identity management systems implemented within their backend so combining information and integrating cross-institutional systems could be very cumbersome.

The task at hand involves creating a federated identity prototype which represents different educational institutions which use different technologies for their identity management sub-systems.

The goal of this project is to show how a circle of trust can work across a number of institutions irrelevant of their backend infrastructures and Identity Management technologies used. This project been coded '*Campus-Link*'.

4.2 Scenarios

Three hypothetical scenarios were designed upon which the final prototype was based.

Scenario 1: UOM and MCAST start off the educational federation, both of which are IdPs and SPs for their members. Both UOM and MCAST use SQL Server for their identity system.

Scenario 2: ST Martin's Institute joins the federation. ST Martin's Institute acts only as an IdP for their students. MySQL is used for their backend identity system.

Scenario 3: Later on ITS joins the federation using Active Directory for their backend identity system. ITS is also an IdP, providing no services to its members or any other entities within the federation.

4.3 Functionality

4.3.1 Student and Lecturers

Both students and lecturers can benefit from Single Sign On onto the federation allowing them to use services and download resources from multiple institutions without re-registering at different institutions. Furthermore and in parallel to this these parties can also sign out using Single Sign Out in order to ensure stricter security across sessions within the federation. Global session timeouts are also important whenever users do not log out.

4.3.2 Administrators

Administrators can add IdPs and SPs within the institution’s circle of trust depending on the policies they adopt and on the trust they have in specific external institutions. They can enables and disable this trust through their federation backend systems.

5. IMPLEMENTATION

5.1 Single Sign On

With reference to figure 5, the following steps represent how single sign on is achieved within the prototype:

1. A UOM student signs into the student homepage (of his IdP) using a username and password.
2. Once the user is authenticated the **'CreateIdPCookie'** form is called, which creates a cookie for the user. It stores the user's session id, his or her security clearance, username, the name of the user's Identity Provider and a string which is comma delimited, containing the IDs of the service providers the user has visited. When the cookie is first created this string is empty.

[2] http://shibboleth.internet2.edu/why-shibboleth.html Retrieved January 3, 2010

3. User is then logged in and sent to the home provider homepage
4. The student then decides to consume a resource at another SP, such as MCAST which is password protected and reserved for MCAST students only. MCAST checks if the user has a valid MCAST cookie. If not the MCAST application first asks the student to select his or her identity provider (**WAYF – Where Are You From?**). The student chooses his/her own IdP, in this case the University of Malta.
5. When the student hits the submit button the **'AssertionRequest'** form is called. The **'AssertionRequest'** gathers information from the service provider in order to be sent to the identity provider (such as who issued the request and what type of binding is being used). The binding that is being used is HTTP Redirect which doesn't require the message to be signed since the link would be too long for some browsers.
6. The assertion request is then sent to the appropriate IdP which in this case is UOM. UOM receives the request via the **'SSOService'** form. The code first starts off by checking if the user already has been logged in by checking for a valid cookie. If yes, then it immediately creates an assertion for that user. The assertion would contain information gathered from the cookie such as the username and the user's security clearance. If the user had not been logged in, the application sends the user to the login page before the assertion can be created and sent back to the SP. Before the assertion is sent, the service provider's ID number is added to the cookie.
7. The assertion is then sent to the SP's **'AssertionConsumerRequest'** form.
8. The **'AssertionConsumerRequest'** form receives the assertion, checks if it is valid and starts taking the values stored within it. The user's security clearance is checked in order to know whether the user has the right to access the requested resource. If authorized, the application calls the **'CreateCookie'** form to create a cookie for the user.

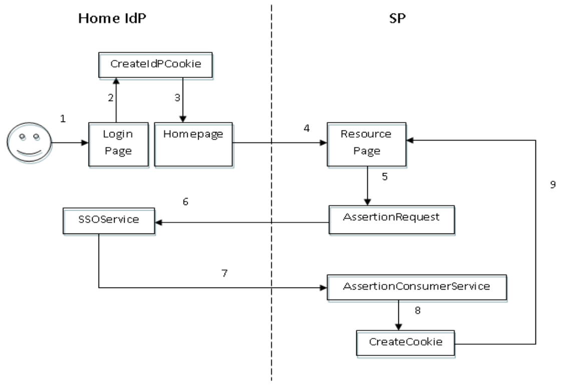

Figure 1. Single Sign On

5.2 Single Sign Out

Single Sign Out allows the user to sign out from his or her identity provider along with all the SPs he or she has also signed into. There are two ways a user can sign out. Either through the user's identity provider or through any service provider he or she has already accessed.

In order to make this possible, within the user's cookie, a string containing the ID's of all the service providers the user has signed into is stored. This allows the Identity Provider to track the user's movement and make Single Sign Out possible.

1. User decides to logout from the IdP
2. When logging out from the IdP, the **'DeleteAllCookies'** form is called. The IdP grabs the first ID from the cookie, removes it from the string and searches for the service provider's logout URL (within the IdP's trusted SP table).
3. Once found, the browser is automatically redirected to the service provider's logout page. That logout page takes note of the user's Identity Provider's name which was stored in the user's cookie when the user first logged in and then removes the user's cookie by calling **'DeleteCookie'**.
4. Once removed, the user is sent back to the Identity Provider. The Identity Provider then again checks the string within the user's cookie to see if the user has logged into any other service providers. If not then the user's identity provider cookie can be deleted.

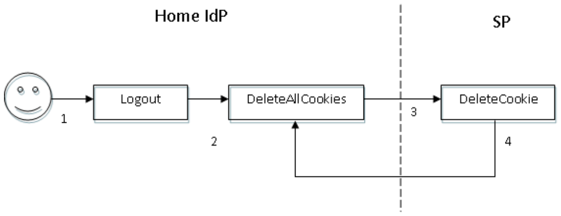

Figure 2. Single Sign Out

6. TESTING AND EVALUATION

6.1 Testing

To test out the system a number of individuals were chosen and assigned to one of the Identity Providers. Each individual was given a username and password in order to access the system. They were then given a set of objectives they needed to achieve along with a number of binary questions for them to answer. This resulted in a number of use cases. All use cases tested where successful for all the features provided by Campus Link.

Usability testing has also been carried out rigorously by means of test cases where a number of people were given access to the system in order to try out the implemented features. Feedback was encouraged, recorded and used for iterative prototype improvement.

6.2 Evaluation

Throughout this study it was found that both students and tutors are willing to access shared resources across educational institutions. This scenario is also applied for international online libraries such as ACM and Springer which allow institutions to act as IdP in order to access restricted resources. Students also agree that this mobility of identities across local educational domains could only enhance their experience.

Of greater importance is the mapping of policies across institutions and the maintenance of trust between the technical and administrative people involved. These have to agree on the minimum authentication mechanisms which can be considered as acceptable so as to maintain a standard level of security across domains. More sensitive resources, sites or physical areas, such as specialized labs, may require a higher level of assurance which needs to be agreed before the federation is created or before new entities join in. Policy documents apart from technical aspects are thus crucial for the success of these kind of identity systems.

7. CONCLUSION

To conclude, it can be said that from the final results and benefits obtained from the prototype, the system was a success and worked as intended. The standardized protocol used, SAML, has proven to be very resourceful and handy to create a federated system regardless of the backend technologies used by the various institutions. The important thing to remember is that trust and proper policy mapping are the keys to create a successful federated system, without which institutions would not be able to work in synergy.

8. REFERENCES

[1] Anonymous. (n.d.). *About Shibboleth*. Retrieved January 3, 2010, from Shibboleth: http://shibboleth.internet2.edu

[2] Anonymous. (n.d.). *PKI*. Retrieved May 11, 2010, from Elock: http://www.elock.com/pki.html

[3] Latamore, B. (2006, May 10). *Boeing Pioneers Federated Identity Management with Partners*. Retrieved April 13, 2010, from ComputerWorld: http://www.computerworld.com/s/article/9000324/Boeing_Pioneers_Federated_Identity_Management_with_Partners

[4] Ragouzis, N., Hughes, J., Philpott, R., Maler, E., Madsen, P., & Scavo, T. (2008). *SAML Technology*. Retrieved December 12, 2009, from OASIS: http://docs.oasis-open.org/security/saml/Post2.0/sstc-saml-tech-overview-2.0.pdf

[5] Vooren, T. v. (2007). Retrieved December 13, 2009, from Everett: http://www.everett.nl/

[6] Watkins, B. (2005, November 02). *Federated Identity Management*. Retrieved April 13, 2010, from TechRepublic: http://articles.techrepublic.com.com/5100-10878_11-5893946.html.

[7] Porter, C. *A ubiquitous system for the provision of secure services in a multi-channel federated environment enabling full identity mobility*, University of Malta, 2008