Federated Identity Infrastructure

for

Andalusian Universities

Victoriano Giralt[1], Carmen López[2], Diego Lopez[3], Luís Meléndez[4], Francisco Sánchez[5]

[1]University of Málaga, SCI, 29071 Málaga, victoriano@uma.es. [2]University of Seville, SI, 41071 Seville, carmen@us.es. [3]RedIRIS, Campus Reina Mercedes, 41071 Seville, diego.lopez@rediris.es, [4]University of Córdoba, SI, 18071 Seville, luism@uco.es, [5]International University of Andalusia, Sede Isla de la Cartuja, 41071 Seville, fco.sanchez@unia.es.

Keywords
identity federation, AAI, Shibboleth, SAML, SCHAC, eduPerson

1. EXECUTIVE SUMMARY

Identity is an strategic piece in any kind of organization, as it is the key to gain access to institutional services and information. Proper management of authentication and authorisation mechanisms is beign transformed into a more and more relevant (and often complicated) task for academic institutions. This fact is brought forward when we take into account the constant change of the user base induced by their open nature. Authentication and authorisation infrastructures (AAI) and, above all, the schemes that allow their integration through federation mechanisms, are a key component of academic and administrative institutional ICT. This is specially relevant for the achievement of a fundamental target in European academic environment, the implementation of the Bologna Process.

1.1.Background

The importance of federated access to services for their strategic goals prompted the Association of Public Universities of Andalusia to set up an working group for specifying the requirements for such federated services and, if possible, set up a demonstrator. RedIRIS, the Spanish National Research and Educational Network, offered its support and, then, chaired the working group.

1.2.Design and Development

The working group has found that technology is mature enough to set up a federation using systems provided by diverse vendors and sources. The state of the art in interoperability standards and the involvement of several group members in work groups defining international standards such as SCHAC or eduPerson, has been key to the success.

1.3.Conclusions

Using the principles defined by the work group, the main deliverables have been both an specification document for the final production federation that will include the 10 Andalusian universities and the Andalusian REN both as Identity Providers and as Service providers, and as many private Service Providers that wish to participate; and a working demonstration federation with the Universities participating in the group, that is in full production with most proposed services, including some experimental ones designed and developed by the group. The federation common elements, have been designed with the main aim of guaranteeing the scalability and continuity of the federation. The work group has also produced several use cases for the federation some of them with corresponding functional proof of concept services.

2. INTRODUCTION

The ICT managers committee of the Association of Public Universities of Andalusia (AUPA) decided that an Identity Federation for the member Universities would be a desired advancement to allow their students, faculty and personnel access shared resources, following the mobility principles of the Bologna process.

Such decision prompted the creation of a technical working group composed of ICT specialists from some of the participating Universities and the Andalusian regional REN (Research and Educational Network) and chaired by RedIRIS (the Spanish National REN). These experts have produced a test bed operational service and a technical specifications document for the final deployment of the Federation.

The main guiding principle for the works has been a strong commitment to standards and the ability to form a multi-technology federation, provided that the state of maturity of Identity Management technologies and the different level of adoption in the participating institutions, has resulted in several products deployed at various sites.

One of the main deliverables required in the technical specification is the development and deployment of a canonical implementation against which products can be tested before being accepted for use in the Federation.

There are services in the European academic space that allow users that belong to participating institutions to use services provided by institutions other than their own, but they use a single underlying technology.

3. OBJECTIVES

The working group set up the following main objectives:

- The federation should serve the needs of any user in the participating institutions, i.e.: students, teachers, researchers, administrative personnel, etc.
- Find a system that will allow the participating institutions to select the technology they prefer for connecting to the federation, based on status of their Identity Management Systems.
- Develop the needed pieces of software required to fill voids detected in other deployments, specially regarding metadata distribution, trust fabric building and diagnostics.
- Deploy a reference implementation for validation and testing purposes.
- Build a series of common use cases that could be used for deploying services at the institutions joining the federation, with a special focus on those departing from the usual browser-server interactions.
- Use as much open source software as possible.
- Use established standard schemas, like eduPerson and SCHAC, for the exchange of personal data, with a high regard for users privacy and consent.

At first, the work group decided that the federation should be based on Shibboleth 1.3 SAML profile, but the changes in the state of the art that have developed during the process, like the appearance of new versions (Shibboleth 2.0) or totally new software (simpleSAMLphp), have made the group to revise that decision and propose the use of full SAML 2.0 as soon as possible, which increases the opportunities of interoperability with commercial and governmental federated services.

4. METHODOLOGY

The project has used collaborative and federated technology from the start. Work has been carried out by technical personnel from the participating institutions.

Each participating institution has installed, at least, one Identity Provider connected to its Identity Management Infrastructure, be it an LDAP directory, a relational database system or any other thing, such that real user identities could be used for the pilot services.

Also, some of the institutions have also deployed federated services for use by the community.

The new software that was needed to implement some of the new core services, has been developed by members of the working group and made available to the community using the RedIRIS forge.

In order to identify needs that should be addressed, the most expert members of the group have pointed out weak points they knew from other federation deployments. Other members of the group have voiced needs and requirements they had in their domains for the federated infrastructure.

Once the requirements and improvements were collected, the working group drafted a document with specifications for the different services.

5. TECHNOLOGY AND DESCRIPTION

The main characteristic of the federation has been the use of disparate technologies in the participating sites. When the project started, the first three universities already had Web Single Sign On system in production and diverse IdM system. This a clear difference from other similar deployments in our area, like the Norwegian FEIDE or the Finnish Haka federations, where all the participating institutions use the same technology.

Eduroam is a very successful federated service deployed all over Europe, and beyond. The group counted with two persons that have been very active in its development and deployment, which has been of great help in identifying issues even before they arose. This service is different from our federation in that it is based in a common credential transport mechanism: a RADIUS hierarchy. And is similar in that user access the services provided by institutions other than their own with credentials issued and validated by their home institutions, and in allowing the later to decide which credentials and validation mechanisms to use.

The federation uses or is deploying several technical and policy elements.

5.1.Federation Metadata Management

Practice has shown that what conforms a federation is the metadata, so the main design principle for the technical foundations of the Andalusian Federation has been the ease of access to, and the validation of, the services metadata.

Metadata allow IdPs and Sps to build the trust links required for the former to assure that the user data are sent to an authorized receiver, and for the later to be able to autheticate the source of said data. Metadata are the back bone of a federeation and should be managed by an entity that al members trust.

Thus the group has produced specifications and reference implementations for:

- Metadata manager with web and web-services interfaces. This is a repository for metadata that can be retrieved both manually, by persons, and automatically, by programs. The retrieved data is properly authenticated using public key criptography mechanisms. A similar approach has been developed independently, and after our specifications were written, by the authors of SimpleSAMLphp. The metadata manager allows for institution designed persons to update the metadata through a management interface, using X.509 certificates with legal binding value in Spain for identification.
- Metadata based WAYF service. Existing Where Are You From service implementations use local configurations for displaying the available IdPs to the users, thus requiring manual intervention for every addition or removal, which leads, often times, to the existence of just one WAYF per federation. Our approach has been to develop a WAYF service that presents the users with information derived from the federation metadata. This information is securely retrieved from the metadata manager and, thus, easily kept up to date. This also allows for the easy deployment of as many WAYF services as the participating institutions see fit.
- Metadata query facilities. The information model for the metadata conforms to the SAML 2.0 specifications and allows for the participating IdPs to easily obtain all the information needed to configure the services to join the federation. This also allows for a Service

Provider wanting to offer its services to the federation to easily obtain the required data with no other intervention.

- Attribute release control. The work group has specified but not yet developed a system for centralized attribute release policy (ARP) management. The ARP system should also allow for locally specified policies at institutional level with web services access to the centralized policies and for user informed consent before sending the data to the SPs. The group has evaluated several existing technologies such as the Australian ShARPE, the works of SWITCH and the controls in SimpleSAMLphp.

5.2.Trust model

TERENA has produced an excellent service, called Server Certificate Service (SCS), for obtaining service and server certificates at a really low cost, it is free for the Universities, the NRENs have absorbed the cost. The procedure for obtaining such certificates is both easy and reasonably secure, which has lead the working group to base the trust model on the use of SCS certificates for the participating entities. Any institution that wants to be part of the federation just needs to designate a representative whose legal personal certificate is sent to the federation operators for access to the metadata manager, obtain an SCS certificate, and upload the information to the service.

- Trust root and trust policy manager. The federation will accept several Certification Authorities (CA) for certificate signing. This trust roots will be distributed off line, outside the federation infrastructure, but there should be mechanisms for obtaining the list of valid root certificates both over normal web queries for persons and over web services for systems. The system will link Internet domains to CA roots, for validation of only certificates pertaining to entities in such domains.
- OCSP. The federation will use the Online Certificate Status Protocol for validating the certificates in the metadata. The service will be offered as an aggregator for those CA that only produce CRLs. This service has not yet been developed.

5.3. Validation service

The federation will provide a service with a minimum of IdP, a SP and a metadata server, with test data, for any one that wants to connect to the federation on either role, to test their set up against a known working mock up. One the service works properly against the validator, it can start the process for joining the production federation.

During the design and test phases, this validation has been done against non production services in some of the participating institutions and, then, against non critical services.

5.4.Identity Provider for Homeless users

The Federation will include a central Identity Provider for those uses that need to access any service but do not belong to any of the participating Universities. Such provider will be managed by the Federation operator. This IdP will have a management interface offered as a SP for the federation to be used for those persons with the required privileges. This service is already operational at the Andalusian Scientific Computing Centre (CICA), while the management interface is still in developing stage.

5.5.Attribute exchange model

The federation has selected the eduPerson and SCHAC schemas for attribute exchange, as both are in wide use inside the academic domain.

5.6.Diagnostics tool

The federation participants will offer a federated service for accessing the logs of the services they provide. The group has already developed a Shibboleth log parsing tool that can be queried for specific users or transactions with a federated access control system, so the service operators can

easily find the relevant information to support the users. This tool is being expanded to be easily adapted to other log formats.

6. DEVELOPMENTS

6.1.Software

The working group has already developed several software pieces to support the requirements of the federation:

- Metadata management tool with metadata query facilities
- Metadata based WAYF service, with extra functionalities (described in other paper in the workshop)
- Log parsing federated tool.

6.2.Documentation

The participating institutions. Describing each of them in detail here is not possible due to space constraints but they will be publicly available at the federation web site (http://confia.aupa.inf/) once the federation enters full public production.

- Andalusian Virtual Campus. The e-Learning systems of Andalusian Public Universities provide support for the students registered in any of them to take a virtual course in any of the others. This system will be Federated both for access control and for exchanging student data using standards that could come out from Bologna process.
- Reciprocal library registration. Federating of the Library Management systems will allow users from any of the participating institutions, physically present at the library of any other one, to get the same level of access to resources, electronic or physical, as in their home organization.
- Federated file swapping repository: Consigna. A tool that has been used as a simple demonstrator for the Federation. The tools allows the exchange of files using a web browser between users of the community and external users. Access to files depend on either client IP address or validated identity of the user against any of the Federation Identity Providers.
- Federated SSH access to systems. The group has developed software for federated and automated provision of accounts for users in need to access Linux based systems over secure shell. The tool allows for either federated retrieval of the users public keys or manual entry.

The group has also produced an animation based on several working services at three different institutions to show how an user identity federation works.

6.3.Dissemination

It is also an aim of the working group to give the project and results as much public dissemination as possible, in order for the wider community to make use of anything useful. The following flash animation (in Spanish) has been used to present the concepts to non technical managers and users inside and outside the participant organizations.

Other dissemination efforts include participating in national and international events presenting the achieved goals.

7. RESULTS

The working group has been able to set up a federation with three production IdPs, one preproduction IdP and three more in the works, at the time of writing this paper, using four different technologies. There are also several production or preproduction services put on line.

This has been done in a really short time, which has reassured that a multitechnlogy federation can be set up with not much effort given the proper level of expertise in the core and with collaborative work for those less proficient.

While we were deploying the test services, some important breakthroughs have occurred, like the release of Shibboleth 2.0 or the appearance of SimpleSAMLphp. This has been of capital importance for the quick set up of some new providers.

Having an easy to deploy, configure and use federated application as the file sharing service has proved a strong driving force for the people doing the deployments as it was a quick a dirty way for testing that their work was successful.

8. BENEFITS FOR THE PARTICIPANTS AND THE COMMUNITY

A federation as the one we are setting up will be of capital importance for the Bologna process, in which all the participant universities are involved by law. The framework allows for persons to seamlessly use services provided by several universities simply enrolling in one of them and using the identity this institution provides.

The new concepts and services we have developed for federation administrators and operators ease the configuration and deployment of new services and identity providers, thanks to simple and convenient metadata and trust management.

9. CONCLUSIONS

The federation framework presented in this paper proposes novel solutions to the problems already detected by practitioners in the evolution of identity federation solution, providing support for a multi-technology environment by means of strong standard commitment and defining a system tightly tied to metadata in what relates to the federation central services. Some specific components for these tasks have been defined and demonstrated and will be deployed soon.

Furthermore, the framework is intended to support services beyond the usual browser-server binomial, and the team is working on use cases related to VoIP and SSH access to computing resources.

10.COPYRIGHT NOTICE

11.REFERENCES

FEIDE (2008). SimpleSAMLphp. Retrieved May 2, 2008 from: http://rnd.feide.no/simplesamlphp

FEIDE (2008). FEIDE identity federation. Retrieved May 2, 2008 from: http://feide.no/content.ap?thisId=1307

CSC (2008). Haka Identity Federation. Retrieved May 2, 2008 from: http://www.csc.fi/english/institutions/haka/index_html

EDUCAUSE (2008). eduPerson. Retrieved May 2, 2008 from: http://www.educause.edu/eduperson/

Internet2(2008). Shibboleth. Retrieved May 2, 2008 from: http://shibboleth.internet2.edu/

SAML 2.0

TERENA (2008). Eduroam. Retrieved May 2, 2008 from: http://www.eduroam.org/

TERENA (2008), SCHAC. Retrieved May 2, 2008 from: http://www.terena.org/activities/tf-emc2/schac.html