UPGRADE is the European Journal for the Informatics Professional, published bimonthly at <http://www.upgrade-cepis.org/>

Publisher
UPGRADE is published on behalf of CEPIS (Council of European Professional Informatics Societies, <http://www.cepis.org/>) by Novática <http://www.ati.es/novatica/>, journal of the Spanish CEPIS society ATI (*Asociación de Técnicos de Informática*, <http://www.ati.es/>)

UPGRADE monographs are also published in Spanish (full version printed; summary, abstracts and some articles online) by Novática

UPGRADE was created in October 2000 by CEPIS and was first published by Novática and INFORMATIK/INFORMATIQUE, bimonthly journal of SVI/FSI (Swiss Federation of Professional Informatics Societies, <http://www.svifsi.ch/>)

UPGRADE is the anchor point for UPENET (UPGRADE European NETwork), the network of CEPIS member societies' publications, that currently includes the following ones:
- InfoRewiew, magazine from the Serbian CEPIS society JISA
- Informatica, journal from the Slovenian CEPIS society SDI
- Informatik-Spektrum, journal published by Springer Verlag on behalf of the CEPIS societies GI, Germany, and SI, Switzerland
- ITNOW, magazine published by Oxford University Press on behalf of the British CEPIS society BCS
- Mondo Digitale, digital journal from the Italian CEPIS society AICA
- Novática, journal from the Spanish CEPIS society ATI
- OCG Journal, journal from the Austrian CEPIS society OCG
- Pliroforiki, journal from the Cyprus CEPIS society CCS
- Tölvumál, journal from the Icelandic CEPIS society ISIP

Editorial correspondence: Llorenç Pagés-Casas <pages@ati.es>
Advertising correspondence: <novatica@ati.es>

UPGRADE Newslist available at
<http://www.upgrade-cepis.org/pages/editinfo.html#newslist>

ISSN 1684-5285

Monograph of next issue (April 2010)

"Information Technology inTourism Industry"

(The full schedule of **UP**GRADE is available at our website)

CEPIS UPGRADE
The European Journal for the Informatics Professional
http://www.upgrade-cepis.org

Vol. XI, issue No. 1, February 2010

* This monograph will be also published in Spanish (full version printed; summary, abstracts, and some articles online) by **Novática**, journal of the Spanish CEPIS society ATI (*Asociación de Técnicos de Informática*) at <http://www.ati.es/novatica/>.

Digital Identity and Identity Management Technologies

Isaac Agudo-Ruiz

There are many technologies for identity management available in the form of open specifications, open source tools and commercial applications. There are currently several standards competing in the field of identity management. In the beginning SAML (Security Assertion Markup Language) was the only viable choice with a high enough acceptance level. Recently, another technology called WS-Federation has also gained some attention from the community. Although this technology is not as mature as SAML, its modular design gives it some advantages over SAML. In this article we mainly focus on WS-Federation and the family of specifications that surround it.

Keywords: Identity Federation, Identity Management, Web services.

Author

Isaac Agudo-Ruiz has an MSc in Mathematics and a PhD in Computer Science from the *Universidad de Málaga*, Spain. In 2002 he joined the CASENET Project, funded by the V Framework Programme of the European Union. In 2004 he received a grant from the Regional Government of Andalusia, Spain, to finish his PhD. Since 2008 he has been involved mainly in the European projects PICOS and SPIKE of the VII Framework Programme, as well as some other national and European projects. <isaac@lcc.uma.es>

1 Introduction

It is hard to find a globally accepted definition of the term *identity* and even harder to precisely define what is understood by *identity management*. A simplistic approach would consist of defining identity management as user accounts management in a software system. This was the general understanding some decades ago but in recent years, with the emergence of the Internet of Services, more complex issues have arisen and the identity of users has become crucial. In early software systems, the identity of users was managed locally by the system administrator and was only valid for that particular application. In the Internet of Services, anyone can become a user of our applications and it is the users' responsibility to "manage" their identities in an appropriate manner.

There are some concepts related with *identity* that can help us understand the scope of *identity management* and its key challenges. Firstly, we have to make clear which *entities* can be attached to an identity. According to the RFC 2828 Internet Security Glossary, the term *entity* refers to "an active element of a system - e.g., an automated process, a subsystem, a person or group of persons that incorporates a specific set of capabilities." Although we mostly think of human beings when referring to entities, we cannot forget that in most cases we interact with computers rather than humans when using the Internet.

Many definitions of the term "identity" can be found in literature. The greatest common denominator of all these definitions is that an identity refers to some set of *claims, qualities or attributes* that make an entity unique and different from all other entities. In other words, it is the individual characteristics by which an entity is recognized or known in a community or in a given context. Consequently, an entity may have several identities depending on the context in which it interacts. For example, a person may be recognized as the CEO in the context of his or her company, but in a different context such as their bank or their house this reference might not be meaningful. Each identity can be referenced through one or more *identifiers* which are no more than special attributes that can be used to uniquely reference an identity.

The question 'who are you' is usually followed by 'what are you allowed to do'. In an environment where each entity may have different identities, the problem of deciding which *privileges* or *access rights* they own is not trivial. While identity is the basis for authentication, privileges and access rights are the basis for authorization. In theory, authentication and authorization can be conceptually separated. In practice, however, authentication and authorization are often combined and implemented in an authentication and authorization infrastructure (AAI), a privilege management infrastructure (PMI), or an alternatively named but conceptually similar infrastructure.

The choice of one identity or another by a given entity determines not only its privileges but also the perception of the rest of the entities in the system. When an entity interacts repeatedly with other entities in the system, some *trust relationships* can be established between them. Those trust relationships do not target other entities directly but rather their visible identities, i.e. the identities they use to interact with the rest of the entities in the system. Entities may behave well when using one particular identity but behave

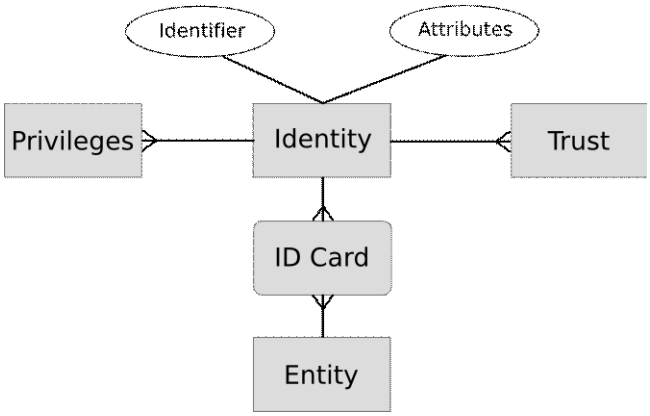

Figure 1: A Simplified Entity-relationship Diagram (ERD) for the Term *Identity.*

badly when using another. The concept of trust is becoming very important in the Internet with the growth of online social communities but is also highly significant in other fields such as sensor networks where the reliability of the network relies on the strength of trust relationships between the nodes of the network.

In order to prove ownership of its identities, an entity makes use of *ID cards*. An ID card is seen as an abstract concept that attests to the legitimacy of an identity and/or its attributes, the same way as we use our passport at passport control.

Figure 1 shows an entity-relationship diagram that relates the terms: *Entity, identity, Trust, Privileges* and *ID Card.* The ID Card can be seen as an associative element in the diagram that relates *entities* with its *identities.*

On the one hand, how identities and privileges are related falls into the field of *privilege management*. On the other hand, how entities and trust are related falls into the field of *trust management*. Thus, the central part of the diagram falls under the umbrella of *identity management*. However it is very difficult to dissociate these three terms: identity management, trust management and privilege management.

2 Interoperability of Identity, or Identity Federation

When people talk about identity, they sometimes underestimate the importance of each of the above mentioned terms and more often than not focus on the notion of the identification card (ID card). In essence, an ID card attests to the legitimacy of an identity and/or its attributes. There are ID cards for all kinds of purposes: passports and ID cards issued by the state, staff ID cards issued by companies, membership and customer cards issued by all kinds of organizations and companies, student cards issued by universities, and so on. While multiple-use ID cards are technically feasible, most ID cards in use today are single-use, meaning that they serve one single purpose or application. There may be many reasons for this fact; one important reason is certainly the fact that an ID card also works as a customer relationship tool (so ID card-issuing organizations are reluctant to share the card with other organizations and potential competitors). The omnipresence of single-use ID cards results in wallets that are filled with all kinds of cards. We know the problem from daily life, and we decide on a case-to-case basis which card to use in a given context.

The situation in the digital world is analogous. We can think of an e-mail address as the most primitive form of identity in the digital world. It consists of one identifier without any attributes. Some people use more than one e-mail account, each of them in a different context as they do with ID cards. They usually have an account for work and another for personal use, but they may have more in order to preserve their privacy, avoid spam or even to be able to access online services that require a particular e-mail account. When we check our mail account we have to first demonstrate the mail server that we are the real owners of the account. For that purpose we typically make use of a combination of a username and a password. We can say that this combination serves as a kind of "ID Card" for our e-mail account to be identified by the mail server. Unfortunately, ordinary e-mail does not provide a means to prove our identity to other users; for that purpose we need to make use of other standards for secure e-mail.

As we mentioned above, physical ID cards are usually not interoperable but nevertheless there is a common understanding of what an ID card looks like. There are eleven workgroups under ISO/IEC JTC1/SC17 working on "Cards and Personal Identification" standards. They have produced a standard that defines the physical characteristics for identity or identification cards, ISO/IEC 7810:2003. Unlike in the physical world, however, in the digital world the form of ID cards has not yet been agreed upon. In fact, there are many ways to implement digital ID cards. An example of a widely used approach that may establish the basis for digital ID cards is digital certificates or public key certificates as defined in the X.509 ITU-T standard [1]. This standard specifies the format of the certificates as well as the algorithms and mechanisms needed for their deployment. Apart from passwords and certificates there are some other mechanisms that we can use to prove our identity. Security grid cards which are used in banking environments, and cryptographic tokens, such as SecurID from RSA, which are widely used in corporate environments are only two examples of many.

One of the main challenges that we have to face in the field of identity management is the interoperability of identities. It is not enough to be able to manage identities within our system; we need to be able to provide mechanisms for the reusability of the identities of our users outside our domain. This requires the establishment of interoperability mechanisms between the different stakeholders of the digital identity business.

When describing the supporting technologies for identity interoperability, or identity federation, that are becoming "de facto" standards, we have to inspect not only the format used to describe the identity or credentials but also

the communication protocols used to transport this information. In fact, the biggest differences between the two approaches that we present in this work focus on the protocols and not on the format of identity. These two technologies, which offer approximately the same functionality, are SAML 2.0 and WS-Federation. SAML 2.0 is a mature OASIS standard specification (2005) and is widely deployed [2]. Most European universities and US Universities are already using SAML with the support of the GÉANT network and the Internet2 consortium respectively. On the other hand, WS-Federation is a novel OASIS Standard built upon the WS-* family of specifications [3]. The latest version of WS-Federation was released in 2009. However, WS-Federation is attracting the attention of the Internet society, mainly because of the success of the Web Services Security Suite (WS-*) [4]. In our opinion WS-Federation has not received enough attention yet whereas SAML has been widely revised in the past. This is the main reason why we pay more attention to WS-Federation in this article. One of the main differences between these two approaches, at least at their beginning, is that SAML was intended to solve the Single Sign-On (SSO) problem in Web environments whereas WS-Federation tries to cover the area of Web services. Both solutions have evolved and have partially converged during the last few years but their origins have left a mark on their development that makes them different in several aspects.

3 Identity Management in the Web Services World

The OASIS consortium is promoting relevant standards for identity management in the area of Web services. In particular, the following four are the most important technical committees (TC) with regard to identity management:

- Web Services Security (WSS)
- Web Services Secure Exchange (WS-SX)
- Web Services Federation (WSFED)
- Identity Metasystem Interoperability (IMI)

In 2006 the OASIS Web Services Security TC released the latest version of WS-Security, initially developed by IBM, Microsoft, and VeriSign. It provides three basics security mechanisms for Web services: sending security tokens as part of a message, message integrity, and message confidentiality. A security token represents a collection of claims, a claim being a statement made about a client, service or some other resource (e.g. name, identity, key, group, privilege, capability, etc.). The specification supports different security tokens as described in the following associated profiles:

- Username Token Profile 1.1
- X.509 Token Profile 1.1
- SAML Token profile 1.1
- Kerberos Token Profile 1.1

This specification works in the application layer, providing end-to-end security, as opposed to point-to-point as provided by TLS/SSL, by incorporating security features in the header of SOAP messages. It provides a standard set of SOAP extensions that can be used when building secure Web services to implement message content integrity and confidentiality. This extension is referred to as "Web Services Security: SOAP Message Security" or simply "WSS: SOAP Message Security". Apart from allowing the use of different security token formats, it also supports multiple trust domains, multiple signature formats, and multiple encryption technologies. This specification cannot be considered as a complete solution for secure Web services but rather as a basic building block that can be used in conjunction with other Web service extensions and higher-level application-specific protocols. Indeed, most of the following specifications rely on this one.

The OASIS Web Services Secure Exchange (WS-SX) TC focuses on the definition of extensions to the previously defined OASIS Web Services Security (WS-Security) specification that enable trusted SOAP multi-message conversation (versus the simpler request-response mechanism) via the establishment of a shared security context, and also the definition of security policies regarding the format of the messages and the kind of tokens included in them. This technical committee is responsible for three specifications:

- WS-Trust.
- WS-SecureConversation.
- WS-SecurityPolicy.

WS-Security relies on the existence of certain trust relationships between participants in communications, i.e. Web Service Providers and requestors. Credentials presented by the requestor have to be trusted by the provider and vice versa. How these trust relationships are established is beyond the scope of WSS and this is what WS-SX TC focuses on, by adding additional primitives that enable the establishing and brokering of these trust relationships between SOAP message exchanges participants.

WS-Trust [5] focus on the definition of a Security Token Service (STS) which issues security tokens in accordance with the WS-Security specification. The specification describes mechanisms for issuing, renewing, and validating security tokens. It establishes the format of messages used to request security tokens and their responses. It also provides mechanisms for key exchange. It makes use of WS-Addressing to describe endpoints.

WS-SecureConversation [6] introduces the context authentication model. This model is based on the use of a new WSS token type called Security Context Token (SCT), which is obtained using a binding of WS-Trust. A security context token implies or contains a shared secret which, while it can be used for signing and/or encrypting messages by itself, is best used to derive other keys for signing and/or encrypting messages within this security context. A security context token can be created by a Security Token Service (STS) defined in WS-Trust by one the participating entities alone or cooperatively via message exchanges between the participants. The mechanisms for distributing SCT are covered in WS-Trust. Security contexts are shared among the communicating parties for the lifetime of a communications session but its lifetime can be extended by renewing it, or reduced by cancelling it.

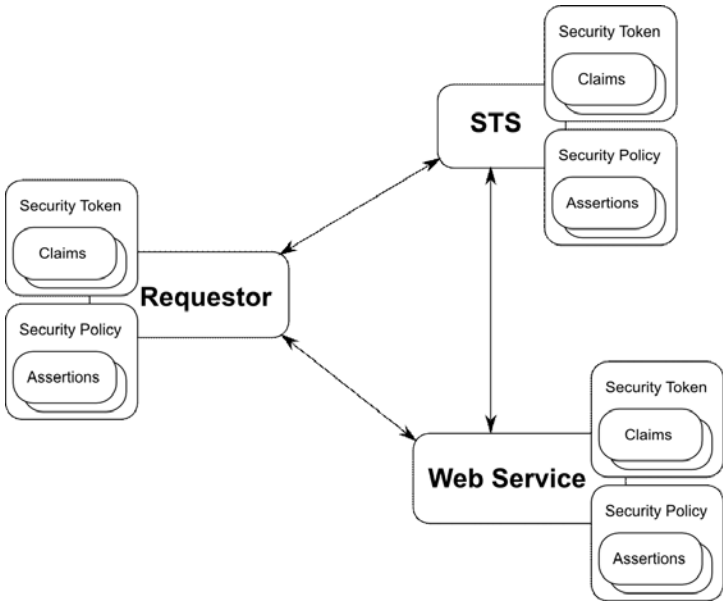

Figure 2: Architecture and Core Elements of WS-Trust.

WS-SecurityPolicy [7] defines a set of security policy assertions that conforms to WS-Policy framework [8], regarding some security features introduced in: WS-Security, WS-Trust and WS-SecureConversation. These policy assertions cover aspects related to which token types, cryptographic algorithms and mechanisms are allowed in a secure exchange of messages. They can also be used for describing security requirements at a more general or transport-independent level. The main purpose of this specification is to define an initial set of assertions which is both flexible and specific enough to ensure proper interoperability of security mechanisms between the participants in the communication. This specification also aims to make policy assertions as simple as possible, so that the policy intersection mechanisms introduced in the WS-Policy framework can provide a narrowed set of policy alternatives which are shared by the two participants who are attempting to communicate.

The typical scenario covered by WS-Federation is one in which resources managed in one realm can be accessed by entities whose identities are managed in other realms. The mechanisms presented in this specification enable authorization decisions to be based on the sharing and interchange of identity, attribute, authentication and authorization assertions between realms.

The federation framework defined in this specification builds on top of the WS-* family of specifications, in particular WS-Security and WS-Trust, providing a rich extensible mechanism for federation. It therefore allows for different types of security tokens, infrastructures, and trust topologies. In order to describe what aspects of the federation framework are required/supported by federation participants, we recommend the use of WS-SecurityPolicy.

All these specifications together allow identities from one realm to be properly recognized in any other realm. However, the problem that the final user faces due to the co-existence of many digital identities is not covered in any of them. In fact, none of them deals directly with digital identities but only with security tokens.

In 2009 The OASIS Identity Metasystem Interoperability (IMI) TC approved the Identity Metasystem Interoperability specification that aims to integrate digital identity into the WS-world using the Information Card Model. In the IMI specification, digital identity is specifically defined as a set of claims made by one party about another party. If we look at the definition of security tokens given in WS-* specifications we see that both terms are very similar. It introduces the term identity selector which allows users to manage their digital identities and use them according to the context of the application. Although information cards are more oriented to Web browsers they can also be used with Web services. This specification also provides an extension to WS-Addressing to describe secure and verifiable identities for endpoints.

An *information card* is a signed XML document representing a digital identity of a subject, i.e. a set of claims. We can consider two kinds of information cards: self-issued or personal information cards which are generated and signed by an individual; and managed information cards which are generated and signed by a third-party Identity Provider. Information cards can be used for both signing-in and signing-up. Self-issued cards would be normally used to sign-up as they do not require a previous trust relationship between the subject and the relying party, whereas signing-in might require a managed card. An identity selector is a piece of software responsible for managing informa-

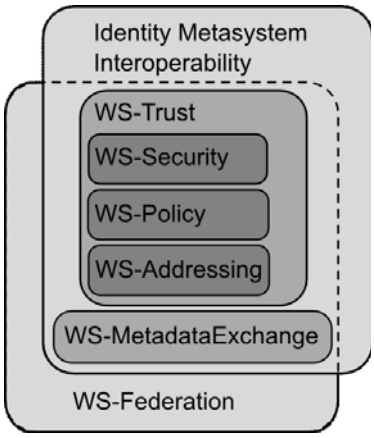

Figure 3: Dependencies of WS-* and Identity Metasystem Interoperability Specifications.

tion cards. It allows users to select the most appropriate card in any given context. It prompts the user with a list of cards that match security policies specified by the application acting as a Relying Party. It is also responsible for retrieving a security token from the associated Identity Provider when the user chooses a card. When retrieving the security token, the subject must authenticate to the STS or Identity Provider. There are four supported authentication mechanisms:

- Username and Password Credential
- Kerberos v5 Credential
- X.509v3 Certificate Credential
- Self-issued Token Credential

The identity card model, also known as CardSpace, was initially promoted by Microsoft but thanks partially to the *promise of open specifications* made by Microsoft (http://www.microsoft.com/interop/osp), some open source implementations of the model have emerged, such as Higgins, Bandit, OpenInfoCard and Pamela. After the approval of the IMI specification as an OASIS standard more and more companies are likely to provide support for it.

Figure 3 shows the modular structure of the family of WS-* specifications. At the lower level we have the WS-Security, WS-Policy and WS-Addressing specifications that are responsible for providing the basic mechanism for the definition of security tokens, associated policies and the addressing mechanism. At the next level we can find the WS-Trust layer which focuses on the Secure Token Service. As we mentioned before, this layer also includes the WS-SecureConversation and WS-SecurityPolicy specifications. At the highest level we can find both the WS-Federation specification and the IMI specification. All those specifications together provide the mechanisms needed to cover all the aspects of identity management in the Web services world.

4 SAML and Related Technologies

The *Security Assertion Markup Language* (SAML) is an XML-based standard for exchanging authentication and authorization data between security domains maintained by the OASIS Security Services Technical Committee (SSTC). There are two main actors in this information exchange: the Identity Provider (Security Token Service in WS-* specifications) and the Service Provider (Relying Party in WS-* specifications). The initial purpose of SAML was to provide a Web browser Cross-Domain Single Sign-On experience, whereas WS-* specifications were originally targeted at providing a security extension for Web services. However, we have already mentioned that WS-Federation also provides mechanisms for passive requestors, i.e. Web browsers. Hence, there is a functionality overlap between both specifications. The SSTC has also published a specification for federation metadata (Metadata for the OASIS SAML V2.0) which has been adopted by WS-Federation in its latest version. We imagine that in the near future we will see more initiatives for the convergence of both specifications

As we can see in Figure 4, the overall idea of SAML is similar to the one underlying WS-Federation passive requestors' profile, although the terminology is different.

On top of the first version of SAML, the Liberty Alliance proposed its Liberty Identity Federation Framework (ID-FF). Liberty Alliance is a large consortium of both companies and non-profit and government organizations which has played an important role in the evolution of SAML. Most of the changes proposed in ID-FF have been incorporated in SAML 2.0. We can say then that SAML is more mature than WS-Federation, but some SAML profiles extending the core functionalities are still under development or have just been approved, e.g. SAML 2.0 Holder-of-Key Assertion Profile Version 1.0 was released on July 2009. Another initiative that supports SAML is Shibboleth, promoted by the Internet2 consortium.

The ID-WSF specification from Liberty Alliance provides mechanisms that allow SAML tokens to be used in Web services. In fact, the Liberty Alliance is currently focusing its efforts on the integration of SAML in the Web

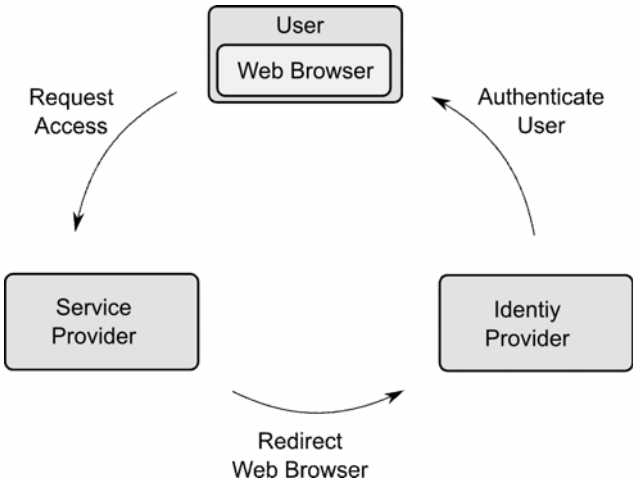

Figure 4: SAML Web Browser SSO Actors.

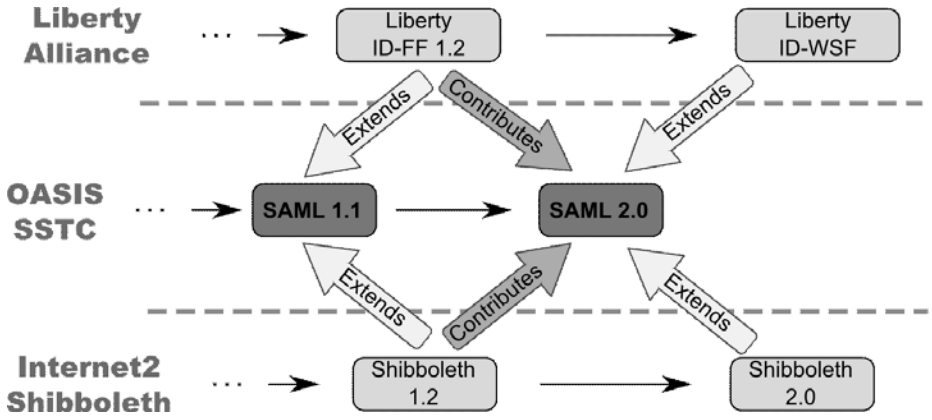

Figure 5: Evolution of SAML and its Related Specifications.

services world, the main reason being that the rest of its proposals have been already adopted by SAML.

Another initiative that has helped SAML evolve is the Shibboleth suite, promoted by the Internet2 consortium. Shibboleth is a set of applications that provides a full identity federation solution on top of SAML protocols. Shibboleth itself is built upon OpenSAML, developed by the same consortium.

The relationships between different versions of SAML, Shibboleth and Liberty are shown in Figure 5. We can see that both Shibboleth and Liberty have fully adopted SAML 2.0, which received feedback from both initiatives. In fact, as members of the SSTC, both Liberty and Internet2 are also contributing to the completion of all SAML 2.0 related specifications.

5 Conclusions

We have seen that both SAML and WS-Federation provide a similar functionality. Indeed, the actors and the abstract information flows are almost the same. There is one entity that attests to the identity of the user and another that trusts this, shall we say, identity statement. In SAML terminology those two entities are called Identity Provider (IdM) and Service Provider (SP) respectively, whereas in WS-* terminology they are called Security Token Service (STS) and Relying Party (RP).

Most commercial identity management solutions support both technologies by letting the two entities, the attesting entity and the trusting entity, speak and understand these two different identity languages. It is up to the system administrator to decide whether to allow the use of both or whether to stick to one of them. If we were to have two different identity federations, each using a different technology, we could have an interoperability problem. A simple solution to help two *a priori* incompatible identity management systems to interoperate would be to place a common entity that switches roles in each of the federations. In one of them it would act as an attesting entity whereas in the other it would act as a trusting entity. This common entity is called an identity bridge and has to be able to speak

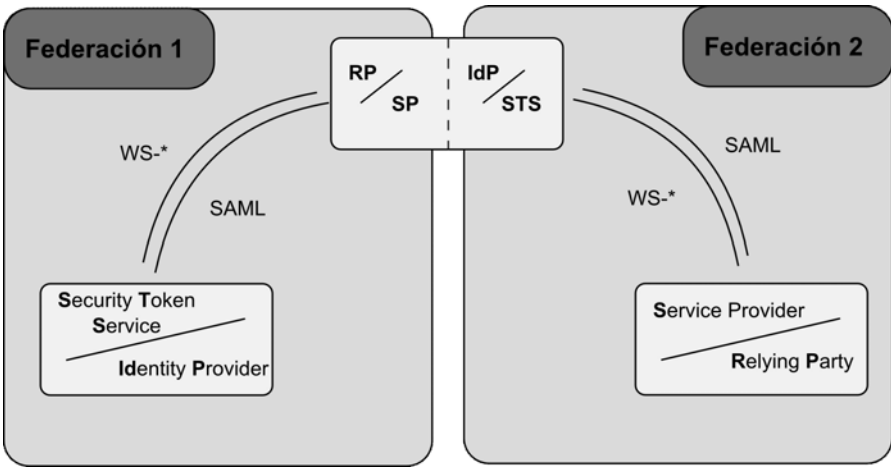

Figure 6: Bridging Federations with Different Technologies.

the languages of each of the federations we wish to interconnect. It will basically act as a translator for them. A simple scenario is shown in Figure 6.

It is worth mentioning that Microsoft, Sun and Novell have made several joint efforts to validate the interoperability of their latest identity solutions with regard to the two specifications reviewed in this article. This underlines the importance of having an open reference specification to enable real identity interoperability.

References

[1] UIT-T Recomendación X.509: Tecnología de la información – Interconexión de sistemas abiertos – El directorio: Marcos para certificados de claves públicas y atributos. (ISO/IEC 9594-8:2001).

[2] "Security Assertion Markup Language (SAML) v2.0" OASIS Security Services TC, March 2005. <http://docs.oasis-open.org/security/saml/v2.0/sstc-saml-approved-errata-2.0.pdf>.

[3] "Web Services Security v1.1" OASIS Web Services Security TC, February 2006. <http://www.oasis-open.org/committees/download.php/16790/wss-v1.1-spec-os-SOAPMessageSecurity.pdf>.

[4] "Web Services Federation Language (WS-Federation) v1.2" OASIS Web Services Federation (WSFED) TC, May 2009. <http://docs.oasis-open.org/wsfed/federation/v1.2/os/ws-federation-1.2-spec-os.pdf>.

[5] "WS-Trust 1.4", OASIS Web Services Secure Exchange (WS-SX) TC, February 2009. <http://docs.oasis-open.org/ws-sx/ws-trust/v1.4/os/ws-trust-1.4-spec-os.pdf>.

[6] "WS-SecureConversation v1.4", OASIS Web Services Secure Exchange (WS-SX) TC, February 2009. <http://docs.oasis-open.org/ws-sx/ws-secureconversation/v1.4/os/ws-secureconversation-1.4-spec-os.pdf>.

[7] "WS-SecurityPolicy 1.3", OASIS Web Services Secure Exchange (WS-SX) TC, February 2009. <http://docs.oasis-open.org/ws-sx/ws-securitypolicy/v1.3/os/ws-securitypolicy-1.3-spec-os.doc>.

[8] W3C Recommendation, "Web Services Policy 1.5 - Framework", 04 September 2007. <http://www.w3.org/TR/2007/REC-ws-policy-20070904/>.