A. Clemm, O. Festor, A. Pras (Eds.): 2005 9th IFIP/IEEE International Symposium on Integrated Network Management (IM 2005) – Managing New Networked Worlds;
IEEE, Nice, France, May 2005

Federated Identity Management: Shortcomings of existing standards

Wolfgang Hommel
Munich Network Management Team
Leibniz Supercomputing Center
Barer Str. 21, D-80333 Munich, Germany
hommel@lrz.de

Helmut Reiser
Munich Network Management Team
Ludwig Maximilian University Munich
Oettingenstr. 67, D-80538 Munich, Germany
helmut.reiser@ifi.lmu.de

Abstract

As the coverage area of conventional *identity & access management* solutions is limited by an organization's boundaries, several approaches for *Federated Identity Management* (FIM), i.e., cross–organizational identity and user data exchange, have emerged. In this paper we demonstrate that even the most important FIM standards have several shortcomings in common which are prejudicial to early FIM adoption in large heterogeneous service infrastructures.

1 Introduction

Identity & Access Management (I&AM) has evolved into one of the most important technical and organizational aspects of service and infrastructure management. Yet, its techniques are limited to the enterprise boundaries by both definition and best practices. Interorganizational cooperations require additional functionality under different technical and legal constraints. In this paper, we derive several shortcomings from a real–world scenario which the three major Federated Identity Management (FIM) standards – SAML, Liberty Alliance and WS–Federation – have in common.

2 FIM scenario: Outsourcing of IT services

The Leibniz Supercomputing Center provides services such as e-mail, web hosting and file storage for staff and students of both Munich universities and several other colleges in the area. The total number of active users exceeds 100,000. Student accounts are presently created and deleted based on large lists of currently enrolled students which the universities send. Staff accounts must be applied for by the departments; each of them obtains a pool of accounts which are managed by a chosen department member. This *delegated administration* saves a lot of work, but introduces security

problems as there is no guarantee that a staff member's account will be deleted when she quits. However, also storing all data at the service provider would raise severe privacy, redundancy and consistency issues.

The idea behind FIM is that access will be granted to everyone whom a trusted business partner vouches for. Each of the partners knows which services its employees are allowed to use. A FIM system makes this information available to the service providers on demand, online and with low delay. Thus, the up–to–dateness of the user data and data quality is much higher than if each user has to maintain her data in multiple places. FIM increases security because there cannot be external accounts forgotten to be deleted, as an identity provider will not vouch for its retired employees any longer. FIM reduces costs and redundancy because organizations do not have to acquire, store and maintain authorization information about all their partners' users anymore. Finally, FIM enhances the protection of privacy, because only data required to use a service has to be transmitted to a business partner.

3 Existing FIM standards and shortcomings

The security assertion markup language (SAML) is standardized in [2]. Data is being exchanged in terms of XML based *assertions*. SAML currently defines *authentication*, *authorization* and *attribute* assertions; furthermore, SAML specifies a request–response–*protocol* which can be used by the service provider to request *assertions* from the identity provider. A *binding* defines how SAML protocol messages are to be transmitted using SOAP over HTTP, and *profiles* determine how SAML can be used by standard web browsers. SAML also provides flexible extension mechanisms; e.g., it can be used with the eXtensible Access Control Markup Language (XACML) to achieve fine–grained access control as it has been done in the area of Grid Computing [5].

The Liberty Alliance is a consortium founded to develop an open standard for FIM [4]. The Liberty architecture is made up of three large building blocks: The *identity federation framework (ID–FF)* defines how data must be exchanged between identity providers and service providers. Although the SAML assertion format is used for the data itself, ID–FF defines protocols, bindings and profiles of its own which are extensions and modifications of their SAML equivalents. The *services interface specifications (ID–SIS)* define two sets of user attributes, the personal and the employee profile, which include basic information for use in B2B and B2C scenarios. The *web services framework (ID–WSF)* specifies SOAP bindings, a discovery service, an authentication service and an interaction service which can be used to ask the user about additional information. Furthermore, Liberty specifies various federated identity trust models; e.g., *circles of trust* which are based on pairwise trust. Alternatively, in the *community trust model* all providers are trustworthy which can prove their membership in the community.

The *web services federation language* [3] is neatly integrated into a series of other *web services specifications* such as WS–Trust and WS–Security. In WS–Federation the user obtains *security tokens* from her identity provider and can pass them to service providers in order to get access to resources. WS–Federation defines a request–

response–protocol which can be used by service providers to acquire security tokens containing attributes actually needed.

We will now demonstrate several shortcomings which are common to all of these standards with respect to the scenario described in the previous section.

Limitation to web services Each of the FIM approaches enables interorganizational web single sign–on. But none of the FIM approaches can be applied to services which are not yet or cannot be fully web enabled, e.g., e-mail and file storage. Although web interfaces exist for both, access through conventional protocols is much more popular and cannot be given up. As there is no support for such legacy protocols, conventional user registration and system provisioning is required.

Persistent data storage Although a service provider can request arbitrary attribute information from an user's identity provider while the service is being used, none of the FIM approaches offers means to notify the service provider about changes in this data later on. SAML assertions also explicitly have a one request property, i.e., the identity provider deletes each assertion after sending it to the service provider. Thus, the service provider cannot contact the identity provider later to request the reassurance that the assertion and the contained data are still valid.

Federation security and privacy control Concerning security, almost only communication security is considered in the standards. Although a public key infrastructure based solution is elegant in theory, it is not trivial to realize: building a common single–purpose PKI for a lot of federation partners, e.g., in supply chain management, would require enormous resources for both setup and maintenance. Furthermore, neither a holistic security view nor methods for the correlation of security related events across organizational boundaries do exist yet in the FIM standards.

Regarding privacy, the users must be able to regulate which information about them is allowed to be sent to which providers. However, there are no concrete definitions of such *attribute release policies* (ARPs) in the specifications yet. Only Shibboleth v1.2 [1], a SAML–based solution, supports simple ARPs, but in a proprietary format.

Syntax and semantics of attributes Each of the FIM approaches supports the exchange of arbitrary attributes, i.e., pairs of keys and values. But it does not help to find a common data scheme which should be used within an identity federation, including the definition of syntax and semantics. Our experience is that in the real world, finding a common data scheme for interorganizational cooperation is far from being as trivial as it might seem to be in theory: Each organization internally uses slightly different terms and wants to stick to them for both political reasons and the costs involved. So far, only the Liberty Alliance attempts to define a common set of user attributes within ID–SIS. But they are likely to meet the same fate as various "standard" LDAP object classes did in I&AM solutions, i.e., require substantial application and federation specific extensions. None of the FIM standards supports the process of finding a

federation–wide data scheme nor do they offer methods to cope with provider specific semantics.

4 Conclusion and future work

In this paper, we have presented a scenario in which FIM would greatly reduce the management complexity and costs, as well as improve the security regarding the management of external users, which can be customers, business partners or temporary accounts. After outlining the three state–of–the–art FIM approaches, namely SAML, Liberty Alliance and WS–Federation, we demonstrated that they share various shortcomings. In the future we will be working on a generic FIM architecture and concepts to mitigate these deficiencies.

Acknowledgment The authors thank the members of the Munich Network Management Team for valuable comments on previous versions of the paper. The MNM Team directed by Prof. Dr. Heinz-Gerd Hegering is a group of researchers of the University of Munich, the Munich University of Technology, and the Leibniz Supercomputing Center of the Bavarian Academy of Sciences. The web server of the MNM Team is located at `http://wwwmnmteam.informatik.uni-muenchen.de`.

References

[1] M. Erdos and S. Cantor. Shibboleth-Architecture v05. `http://shibboleth.internet2.edu/docs/draft-internet2-shibboleth-arch-v05.pdf`, 2002.

[2] J. Hughes and E. Maler. Security Assertion Markup Language v1.1 Technical Overview. `http://www.oasis-open.org/committees/download.php/6837/sstc-saml-tech-overview-1.1-cd.pdf`, 2004.

[3] Ch. Kaler and A. Nadalin. WS-Federation specification. `http://www-106.ibm.com/developerworks/webservices/library/ws-fed/`, 2003.

[4] S. Landau and J. Hughes. A Brief Introduction to Liberty. `http://research.sun.com/techrep/2002/smli_tr-2002-113.pdf`, 2002.

[5] R. Lepro. Cardea: Dynamic Access Control in Distributed Systems. `http://www.nas.nasa.gov/Research/Reports/Techreports/2003/nas-03-020-abstract.html`, 2004.