Future Generation Computer Systems 26 (2010) 1043–1052

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

A conceptual model for attribute aggregation

David W. Chadwick [a,*], George Inman [a], Nate Klingenstein [b]

[a] University of Kent, UK

[b] Internet2 Consortium, USA

ARTICLE INFO

Article history:
Received 11 January 2008
Received in revised form
25 August 2009
Accepted 14 December 2009
Available online 21 December 2009

Keywords:
Authorisation
Identity management
Attribute aggregation
Privacy protection
Trust
SAML2
Liberty Alliance
Level of Assurance

ABSTRACT

This paper describes a conceptual model for attribute aggregation that allows a service provider (SP) to authorise a user's access request based on attributes asserted by multiple identity providers (IdPs), when the user is known by different identities at each of the IdPs. The user only needs to authenticate to one of the IdPs and the SP is given an overall level of assurance (LoA) about the authenticity of the user and his/her attributes. The model employs a new component called a Linking Service (LS), which is a trusted third party under the control of the user, whose purpose is to link together the different IdP accounts that hold a user's attributes, along with their respective LoAs. There are several possible interaction models for communications between the IdPs, the SP, LSs and the user, and each is described. The model is underpinned with a fully specified trust model, which also describes the implications when participants do not fully trust each other as required. Finally, the paper describes how the model has been implemented by mapping onto existing standard protocols based on SAMLv2.

1. Introduction

Many organisations are experimenting with virtual organisations (VOs) and federations. Practical examples abound, such as the Tera-Grid VO [1], the In-Common Federation [2] and the UK Access Management Federation for Education and Research [3]. Microsoft has added identity federation into its latest Vista operating system with Cardspace [4]. Whilst user authentication and authorisation was originally based on globally unique user identities, such as X.500 distinguished names held in X.509 public key certificates, more recently authentication is being based on federated identities [5], and authorisation is being built upon the Role or Attribute Based Access Control models (RBAC/ABAC), for example as exemplified in [6,7]. The typical modus operandi of the latest federated systems is that the user authenticates to an Identity Provider (IdP), and the IdP sends an authentication statement and authorisation attributes to the Service Provider (SP). The SP then grants access based on the user's attributes. Note that only one IdP and one set of user identity attributes are typically involved in this exchange. However, most users have attributes assigned by a number of different authorities or IdPs; for example, the General Medical Council in the UK says who are doctors, organisations say who are their employees and what their roles are, VO managers say who are their VO members, whilst learned societies such as IEEE say who their members are, etc. Unfortunately most VO and federated systems currently suffer from a significant limitation, namely, the lack of a standard approach to aggregating user attributes, asserted by multiple authorities, for the SP to use in its access control decision making. Ad hoc solutions are currently being experimented with, such as Grid-Shib [8], myVocs [9] and D-Grid [7]. Some of these solutions, such as Grid-Shib and D-Grid, are only capable of aggregating attributes from two authorities, namely the VO manager and the user's organisation. myVocs is an alternative solution that places a myVocs IdP–SP server between the real IdP and the real SP. The myVocs server can hold a set of VO specific attributes which it can aggregate with the IdP's attributes. Different IdPs can be involved. But myVocs has severe limitations in its trust model. It requires the SP to trust the myVocs server to both authenticate all users correctly, and to aggregate all users' attributes correctly. The SP has no assurance about the authentic source of any of the user's attributes since myVocs appears to be the authoritative source of all of them. In comparison, in this paper we propose a conceptual model and a standard protocols based solution to the problem of attribute aggregation, in which a user's attributes can be aggregated from any number of IdPs, whilst maintaining user privacy and giving the SP assurance about the authoritative sources of all the attributes.

A couple of use cases might help the reader to envisage why attribute aggregation is needed.

(i) Accessing Electronic Medical Records (EMRs). Only qualified and registered health care professionals can access EMRs. In addition,

* Corresponding author.
E-mail address: d.w.chadwick@kent.ac.uk (D.W. Chadwick).

0167-739X/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2009.12.004

these professionals have to be employed by a local health authority and be currently on duty. The EMR application needs to aggregate attribute assertions from the national professional database, the local health authorities and the duty roster system.
(ii) Online Purchasing with Membership Discount. In order to purchase a mobile phone contract online and obtain a student discount, a user has to prove she is a registered student, has a good credit record, and has a credit card from an issuing bank. The online store needs to aggregate attribute assertions from the user's university and bank and a credit rating bureau.

Before we developed our conceptual model we gathered a set of user requirements for attribute aggregation, primarily from the academic networking community. The user requirements were obtained by widely circulating a structured questionnaire[1] to many email lists. The results were first presented in [10] and are summarised in Section 2. Section 3 defines the conceptual model that satisfies most of the user requirements. Note that it is not possible to simultaneously satisfy all the user requirements since some of them are mutually exclusive, such as the desire to support multi-hop proxying without knowing who the ultimate end-entity is, and the requirement to have attribute assertions digitally signed by their authoritative sources. Section 4 describes the trust model that the conceptual model requires. Section 5 describes the mapping of the conceptual model onto existing standard protocols based on the Security Assertions Markup Language (SAML) version 2 from OASIS [11]. Section 6 concludes and indicates our next steps in this project.

2. User requirements for attribute aggregation

The following requirements were seen to be important for any new multi-source attribute authorisation system by the majority of the questionnaire respondents.

General requirement

1. Attribute aggregation must be usable in a variety of ways: humans via web browsers, applications via APIs, and grid users via grid clients, etc.

Privacy related requirements

2. Privacy protection of user attributes is of high importance, and this should be through the use of technical controls, which are independent of legal means.
3. Service Providers should be able to track users between sessions if required.
4. Service Providers should be able to learn the true identity of users in exceptional circumstances, but only by contacting the user's IdPs.

User consent requirements

5. IdPs and SPs should only be able to communicate with each other to link together the attributes of a user with the user's consent.
6. Service providers should only be able to query multiple IdPs, in order to pull additional attributes for authorisation purposes, with the user's consent.

Protocol requirements

7. The protocols should be able to tunnel through firewalls using existing open ports (i.e. use http/https).
8. The system should use existing standard protocols and only extend them in a standard way if necessary. SAML is the most popular choice for the base protocol.
9. The proxying of information should be supported through multiple hops/proxies.

Trust requirements.

10. The optional ability to sign all assertions should be supported for all message exchanges.
11. The SP should be able to require that all assertions are signed by their authoritative sources.

Usability requirement.

12. It should be easy to use by end-users and require the minimum amount of user interaction.[2]

As we describe the conceptual model below we will show how most of these user requirements have been met.

3. The conceptual model

Before describing the attribute aggregation conceptual model, we initially need to describe the concept of level of assurance (or authentication).

3.1. Level of assurance (LoA)

The level of assurance (or level of authentication) that is provided by an authentication service indicates the amount of reliance that a relying party can have on the identity of the authenticated user. Identity Providers may indicate the LoA when issuing authentication assertions to Service Providers (SPs). NIST in [12] describes four levels of assurance, ranging from 1 to 4, with 4 being the strongest. Asserting an LoA of 4 requires that the user's identity has been verified physically during registration, using official documents such as a passport and birth certificate, and that online authentication is carried out using strong cryptography where the user's key is held in a tamperproof hardware device. Asserting an LoA of 1 does not require the user to have been physically identified during registration, but it does assert that it is the same online user each time (regardless of who this user actually is). An LoA of zero means no assurance whatsoever and indicates that no registration or authentication has been carried out, i.e. the online user could be any member of the public. Consequently level zero is not defined in the NIST guidelines. NIST specifies the levels of assurance based on the combined strengths of the registration and authentication phases and if either is weak, then the asserted LoA must be correspondingly low.

Our model requires that IdPs use the LoA as follows. When each IdP initially registers a user, it will do so at a particular LoA value (which we call the registration LoA). This will be based on the level of identity checking that is performed by the IdP during registration. Self-asserted attributes, where the user states his name, age, address, qualifications, etc., and no checks are performed by the IdP, are given a registration LoA of 1.

The IdP may support a variety of authentication mechanisms which have different strengths (we call this strength the authentication LoA). When a user authenticates to an IdP during a service session, she will be authenticated using a particular authentication mechanism (which has an associated authentication LoA), and allocated a session LoA which equates to this authentication LoA, but

[1] The questionnaire is available at http://sec.cs.kent.ac.uk/shintau/pages/requirements.html.

[2] This last requirement was not part of the user requirements questionnaire, but was mentioned by at least one respondent as additional requirements. In our opinion it should be a "given" for any system that is to gain wide acceptability.