

Extending the WebID Protocol with Access Delegation

Sebastian Tramp¹, Henry Story², Andrei Sambra³, Philipp Frischmuth¹,
Michael Martin¹, and Sören Auer¹

¹ Universität Leipzig, Institut für Informatik, AKSW,
Postfach 100920, D-04009 Leipzig, Germany,
{lastname}@informatik.uni-leipzig.de
<http://aksw.org/FirstnameLastname> (WebID)

² Apache Foundation
henry.story@bbfish.net
<http://bbfish.net/people/henry/card#me> (WebID)

³ CNRS Samovar UMR 5157, Institut Mines-Telecom / Telecom SudParis
andrei.sambra@it-sudparis.eu
<https://my-profile.eu/people/deiu/card#me> (WebID)

Abstract. The WebID protocol enables the global identification and authentication of agents in a distributed manner by combining asymmetric cryptography and Linked Data. In order to decide whether access should be granted or denied to a particular WebID, the authenticating web server may need to retrieve other profiles and linked resources to work out if the requesting agent is member of an authorized group (e.g. friends of the resource owner's friends). If it were required for such resources to be publicly available, this would be a major privacy limitation on a linked Social Network. In this paper we explore different ways in which an agent can act on behalf of a user and we propose an extension to the WebID protocol which allows for delegation of access authorization. This extends the range of application scenarios where WebID authentication can be efficiently deployed while increasing privacy.

1 Introduction

The World Wide Web is a peer to peer communication system designed from the outset to work in a global space, between agents distributed around the world and acting independently, with no center of control, in a space where new agents can at any point join and there is no complete overview of the whole system. These agents need know nothing of each other up to the point of their interaction. These architectural requirements are the reason for the web's declarative functional architecture, with its emphasis on naming and a logic of unalterability of the meaning of names (URIs).

Agents on the web communicate with each other through a limited number of actions: by making requests for resources (GET), by creating resources (POST

or PUT), or even by deleting resources. Creation or deletion of resources usually require authentication of the agent making the request, and so in many cases do requests for information.

A *WebID* is a URI that refers to an agent - person, robot, group or other thing that can act intentionally. The WebID should be a URI which when dereferenced returns a representation whose description uniquely identifies the agent as the controller of a public key [1,12].

The WebID protocol enables the global identification of agents using asymmetric cryptography in a way that fits cleanly with web architecture: namely in such a way that agents can verify each others identity without having had any previous interactions and in such a way as to allow trust to build up in a decentralized manner through a linked web of trust.

It is worth noting that the host serving the WebID profiles controls the identity of every agent whose URI is within that server's namespaces. This service is known as the origin server [3].

The WebID protocol is designed for client authentication⁴. But the origin server does not just respond to requests, the server is also able to make requests. Indeed WebID authentication requires the server to make WebID profile requests to other servers in order to verify the identity of agents⁵.

Things get more interesting in the authorization space. Consider a very natural application of WebID: allowing friends of one's friends access to some resources. This authorization rule will require the web server to fetch each of the resource owner's friends profiles, in order to build up the list of authorized users. But there is a privacy issue involved here: not everyone wants to make all of their social network publicly visible, and some may not want to make any of it publicly visible. Those people may then protect their FOAF profile with access control rules such as one only allowing friends of their friends access to it. How can a server that needs access to these FOAF profiles in order to apply its own access control rules get access to the information? Would the server itself need to be listed as a friend of a friend by each of the users friends? Should the server take on the identity of the user it is fetching resources for? How would it be able to do so? What other solutions are there? These are the questions we will try to answer in this paper.

The rest of the paper is organized in the following way: Section 2 describes preliminary requirements which we had in mind for our solution and clarifies some important terms, Section 3 goes into detail with the WebID specification and adds support for authorization delegation, in Section 4 we describe our

⁴ Server authentication using IETF DANE follows much the same logic, except that the lookup for the identity is not done using the HTTP protocol but DNSSEC [7].

⁵ Note that fetching a WebID profile for WebID authentication should be done anonymously, for fear of authentication deadlocks. For example one can imagine an agent S with profile P_s requesting a resource on server R which requires authentication. S would send R its certificate, thereby requiring R to dereference S 's profile P_s in order to verify the WebID. If P_s itself requires authentication of R and if R sends a certificate containing a WebID with its profile P_r , and if P_r itself requires authentication then it looks like we have a deadlock.

reference scenarios based on two web applications which is followed by Section 5 where we compare our proposal with other protocols, namely CORS, OAuth and Usage Restrictions. Finally, we conclude our work in Section 6 and give directions for future work.

2 Preliminary Requirements

In order to make discussion of the problems easier, we distinguish the following roles in the access delegation process:

1. The *secretary* acts in the name of another agent, the *principal*.
2. The *principal* is the agent who has a secretary that acts on its behalf.

The solution we propose will be based on the following general principles:

Distinguish secretary from principal - Identity should as far as possible be transparent. A secretary should have its own WebID. The motivation for this is: (1) It allows resource guards to permit or deny requests based on this information. (2) Secretary that have many principals do not need to switch their certificate between requests. (3) It makes it possible to describe the relationship between a principal and its secretary using Linked Data.

Easy to use - The one and only place to describe which secretary are allowed to operate for a principal should be the principal's WebID profile. To grant delegated access to a secretary agent, no other actions than adding 1 triple to the WebID profile should be needed. Retracting this grant should involve simply removing it from the WebID profile.

Minimal protocol footprint - By using HTTP and working declaratively by placing statements in documents, we make adoption of the delegation easier and avoid complex protocol developments. We believe that this is a crucial feature of Linked Data in general.

Efficiency - Finally, the proposed solution should scale with growing number of users and connections. In our context this means that an Social Web application should be able to act in the name of thousands of users.

3 Extending WebID for Access Delegation

This section is organized in subsections which build up step by step to the proposed protocol, adding and validating the need at each step for the next one, the final proposed protocol satisfying the previous principles.

3.1 Solution 1: Acting as the user

The simplest solution for any user agent A wishing to act as the user U , is for it to create a public/private key pair, and with the public key thus generated create a certificate C_u with U 's WebID in the Subject Alternative Name position and add the public key to the user U 's profile P_u (see Listing 1). Having done this, the agent A can then connect using the certificate C_u and its associated private key to any service it wishes to, when it wants to work for/as the user U . This has the advantage of requiring no change to the WebID protocol. It does give the agent A all powers of the user U and so this type of privilege should be limited to agents which are either extensions of U such as smartphones are, or ones U can or must fully trust anyway.

```
1 <http://montague.net/romeo/#me> a foaf:Person;
2 foaf:name "Romeo";
3 cert:key [ a cert:RSAPublicKey;
4   rdfs:label "August 2 2012 at 10am CET, in Firefox laptop browser";
5   cert:modulus "cb24...ed85"^^xsd:hexBinary; cert:exponent 65537;
6   ],
7   [ a cert:RSAPublicKey;
8     rdfs:label "August 2, 2012 at 14:30 CET by Montague Family FB Laurence";
9     cert:modulus "2e34...aa24"^^xsd:hexBinary; cert:exponent 65537;
10    ],
11   [ a cert:RSAPublicKey;
12     rdfs:label "August 10, 2012 at 14:30 Rome time with Nokia Cell Phone";
13     cert:modulus "77e5...c342"^^xsd:hexBinary; cert:exponent 65537;
14    ] .
```

Listing 1. Profile allowing three different agents to act as Romeo

Any server requiring authentication of an agent using WebID must implicitly trust the origin server [3] hosting the WebID profile, since the profile hosting server can change any information in the profile.

When the user owns the origin server and there is a one-to-one relation between the two, as it is the case in a personal FreedomBox⁶ (FBx henceforth), then the origin server is just an extension of the user, and in that case trusting the server is the same as trusting the user.

It should be noted that each device can still be distinguished by servers in case of a problem by logging the exact public key that was used to connect to a service. It should be possible then to track down which device's private key had been compromised in case of unusual behavior.

Things become more interesting when an origin server agent A is serving a number of different profiles $P_1 \dots P_n$ each identifying respectively users $U_1 \dots U_n$ and needs to act on behalf of each of these users. This would be the situation for company servers, government agencies, educational institutions, charities, football clubs, etc. This brings up two issues that need to be taken into account, an issue of data perspective and an issue of efficiency.

⁶ <http://www.freedomboxfoundation.org/>

Keeping views distinct First whenever A connects as U_1 to a remote resource R it has to place the data received in a separate graph from the one it stores public non authenticated data in, and to the one it stores a representation to the same resource R seen when connected as the different user U_2 . This is because the resource may return different representations depending on who is connecting to it - one representation for close friends of the owner of R perhaps, and one for more distant ones. This can be illustrated using N3 graphs, by specifying a relation between a view on a resource (using a yet to be settled on vocabulary), and a graph which is the log:semantics of that view (cf. Listing 2). It is very important that the agent receiving information from different users does not merge the information that was destined to different users, or else information leakage will severely reduce the trust other agents have in that server A . This requires carefully keeping identified information that was aimed at different users separate, perhaps by creating different graph stores: one for public - non authenticated - information, and one graph store for each user the server needs to authenticate as.

```

1 # view that the F.B Laurence has as Romeo of Juliet's friends
2 [ a subj:View;
3   subj:of <https://capulet.org/juliet/friends> ;
4   subj:by <https://montague.net/romeo#me> ;
5   subj:using [ :modulus "2e34...aa24"^^xsd:hexBinary; :exponent 65537 ]
6 ] log:semantics {
7   <https://capulet.org/juliet#ms> foaf:knows <https://montague.net/romeo#m>,
8     <https://capulet.org/fb#john>, <http://montague.net/fb#laurence> .
9 } .
10
11 # view that the F.B Laurence has as Lord Capulet of Juliet's friends
12 # notice that the relationship of Juliet knowing Romoe is missing here
13 [ a subj:View;
14   subj:of <https://capulet.org/juliet/friends> ;
15   subj:by <https://montague.net/father#lord> ;
16   subj:using [ :modulus "ec224...532"^^xsd:hexBinary; :exponent 65537 ]
17 ] log:semantics {
18   <https://capulet.org/juliet#ms> foaf:knows
19     <https://capulet.org/fb#john>, <http://montague.net/fb#laurence> .
20 } .

```

Listing 2. Views on Juliet's FOAF profile by Freedom Box Laurence, when acting as Romeo and when acting as Lord Montague

Efficiency Secondly, whenever the Origin Server needs to act as U_1 to a remote server R it will need to open a new TLS connection to that server using the certificate C_{u1} for that user. This could require N parallel connections to the same server S . Opening TLS connections is somewhat expensive, as it requires expensive asymmetric key cryptographic computations. When N becomes large, this inefficiency will be perceived as a serious drawback between two organizations.

Conclusion It is possible for any number of software agents to act AS the user by creating themselves a public key as described above. This does require a very

strong trust relation to exist between the agent acting as the user and the user herself, as the two will be mostly indistinguishable in the linked web of trust. This method also becomes inefficient the greater the number of different users an agent is working for. In addition to that, identification of different secretaries is done using cryptographic keys rather than URIs.

3.2 Solution 2: Origin Server acting on Behalf of a User

In order to reduce the number of open connections to any server down to a minimum, it would be useful if the origin server could identify itself directly when making a request and specify on behalf of which of its users it was acting on per request so as to be able to make multiple requests on the same TLS connection. One way would be for the origin server acting as client to simply use the same public key as the origin server acting as server - that is it could use the same public key as the one used by the TLS server. It could even use the same certificate. This would make clear to any server it was connecting to, that it could rely on this server as much as it could on the identity of any resource, and in particular any WebID profiles served by that server.

Still, such an agent connecting to a remote service on behalf of one of its users would then need to identify which user it was acting on behalf of, or the remote service receiving a request would not know which access control rule to apply to the requesting origin server. For example in our previous example, in the Romeo and Juliet story, we showed how the Capulet family FBx only shows Juliet as knowing Romeo, when Romeo makes the request on the resource, but not when the request is made by Romeo's father, lord Montague. That was an easy access control rule for the Capulet FBx to make when the request was using the WebID protocol, since there it would know via the WebID authentication process exactly which identity the user was making the request as. But if it is the Origin Server connecting to the Capulet FBx, identified as their Montague Origin Server, then what representation should the Capulet FBx return? Should it return the reduced version available to Lord Montague? Or should it return the more complete version available to Romeo and his close friends?

The proposal here is for the Origin Server to add to *each HTTP request* (made over TLS) an `X-On-Behalf-Of` header identifying the user on behalf of which the request is being made with that user's WebID, as shown below:

```
GET /juliet/knows HTTP/1.1
Host: capulet.org
X-On-Behalf-Of: https://montague.net/romeo#me
User-Agent: FreedomBox_FamilyEdition/0.1
Accept: application/rdf+xml,text/turtle,application/xhtml+xml
```

The Capulet FBx Guard having verified that this request does indeed come from the origin server of Romeo's profile `<https://montague.net/romeo>` would then be able to accept that since the requesting server could in any case act as Romeo whenever it chose to do so, this protocol making exchange between the

Montague and Capulet houses more efficient without making it less secure, this be justification enough for it (the Capulet FBx) to serve the representation that Romeo would have received had he connected to the server directly. In Section 5 we will argue that this is very similar to the way CORS deals with access control, namely by adding a header to a request in order to clarify what agent is making the request.

This `X-On-Behalf-Of` header makes it possible for the *secretary* to open only one TLS connection to a remote server and be able to specify for each request on behalf of whom it is making it. The Montague FBx could make a request to the Capulet Freedom Box `X-On-Behalf-Of` Romeo, a few milliseconds later `X-On-Behalf-Of` Lord Montague, and so on. Again as explained in the previous section it is the responsibility of the Montague FBx to keep the returned information separate, and act with discretion whenever it is acting on behalf of one or the other of the users. As such it's role is not unlike the role of Friar Laurence in the play - dedication to people's private matters requires tact and discretion.

Limitations Using the same cryptographic key between the server acting as client and as server risks making the keys available to a wider audience, and so increases the risk of key compromise. It may be useful if it were possible to decouple the identity of the agent acting on behalf of another user - the *secretary* - from the Origin Server.

3.3 Solution 3: Secretary acting on Behalf of a User

The origin server acting as a client on behalf of a user can then be thought of as a keeper of secrets for that user. It should know how to distinguish what remote servers tell it when it is acting on behalf of one user, from what a remote server tells it when it is acting on behalf of another user. The role of the keeper of secrets for a person is known as the secretary role - a prestigious role taken on for example by figures such as the Secretary of State.

If we now identify the secretary that can act on behalf of a user using a WebID, we can generalize the protocol somewhat. That is, when an agent - such as the Montague FBx Laurence with WebID `<http://montague.net/fb#laurence>` - authenticates to a remote server it can use its own WebID. This would allow the *secretary* to have her own public key, and so to reduce the risk of her private key being compromised affecting the server public key.

But how would the remote server know that it can trust that secretary to be acting on behalf of a particular user? It can no longer just compare the TLS keys of the requesting agent to see if it comes from the same Origin Server. We need to make this relation explicit by use of a special RDF relation provisionally called `:secretary`⁷

The remote server can then verify that the identified agent is the secretary of the agent he wishes to act on behalf of (as specified in the `X-On-Behalf-Of` header, by dereferencing that user's profile and verifying that the user specifies

⁷ <http://bblfish.net/work/2012/09/trust#secretary>

the `:secretary` relation there, as it would if Romeo had the FOAF profiles as shown in Listing 3.

```

1 <https://montague.net/romeo#m> a foaf:Person;
2   foaf:name "Romeo";
3   :key [ :modulus "cb24...ed85"^^xsd:hexBinary; :exponent 65537 ];
4   :secretary <https://montague.net/fb#laurence>.

```

Listing 3. Minimal WebID profile including a public key and a secretary relation

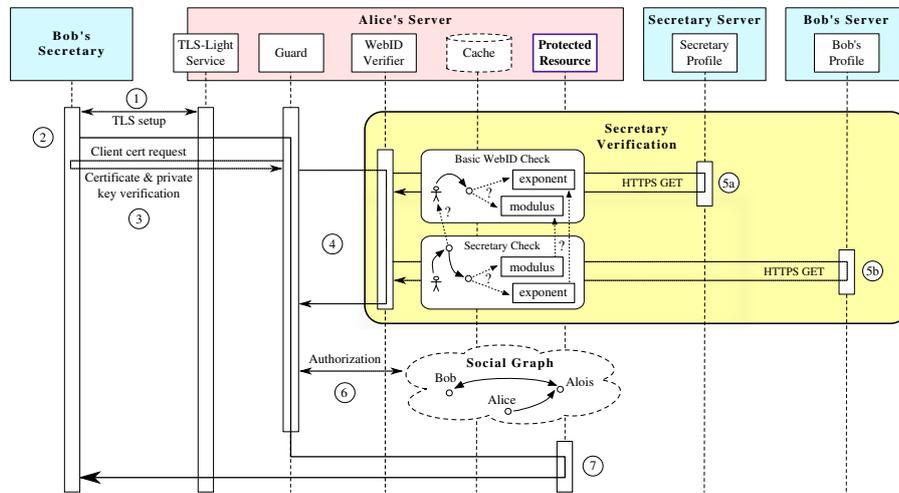


Fig. 1. Extended WebID authentication sequence

The following enumeration describes each authentication step of Figure 1 in detail but concentrates on the context of access delegation: (1) The secretary opens a TLS connection with the server of the protected resource. (2) Once TLS is set up, the HTTP request is sent to the server (e.g. a `HTTP GET`), with an additional `X-On-Behalf-Of` header, which thereby defines the requesting agent as a *secretary* and the referred to agent as the *principal*. (3) The guard intercepts this request, and in turn requests client authentication using TLS session renegotiation. The *secretary* authenticates as itself by sending a certificate containing a WebID referring to it. The TLS-Light service verifies that the *secretary* really is in possession of the private key corresponding to the public key sent in the certificate. This is defined in the TLS protocol [5]. (4) The guard asks the verification agent to verify the secretary WebID which is named in the certificate (5a). This process is exactly as described in the WebID protocol [12]. The guard

also asks the verifier agent to check the secretary claim implied by the `X-On-Behalf-Of` header. (5b) The *principal* agent's relation to the *secretary* is verified by dereferencing the *principal's* WebID profile, and verifying it responds with a true to the SPARQL ASK query `ASK { ?principal :secretary ?secretary . }` where the `?principal` and `?secretary` variables have been bound to the correct URIs. (6) The authentication and verification process having succeeded, the authorization process checking if the *principal* would get access to the requested resource. (7) The resource representation can then be returned or not, depending on the access control rules.

4 Application Scenarios

We describe two different application scenarios where WebID delegation is essential: MyProfile is a WebID identity service application and OntoWiki a semantic data wiki.

MyProfile⁸ is a web service demonstrating how easy it is to both create a WebID profile, and to build up distributed social web applications upon it. Its main purpose is to provide a unified user account through a simple *user profile*. Currently these are tied to the MyProfile project web site, but it has been designed from the ground up to work with distributed Linked Data, making it easy to dissociate the software stack from the MyProfile project domain name, allowing it to be deployed on a machine under the user's control such as a FBx.

It is very important for MyProfile to be able to use WebID access delegation, because a single MyProfile server instance can host multiple users, and must fetch resources for each user asynchronously in order to be able to provide a seamless and rapid user experience. To improve user experience and overall performance, a caching mechanism is used to refresh local copies or "views" of external data. With multiple users coexisting on the same server, the caching mechanism needs to be able to distinguish views of remote resources as seen by different users, as they are served by remote servers depending on their access control and resource filtering policies. As the number of users on MyProfile grows this has to be done efficiently, and so re-using TLS connections where possible is important.

OntoWiki [2] is a web application, which allows publication, exploration as well as manipulation of arbitrary RDF knowledge bases in distributed scenarios. We refer to it as a data wiki, since it adopts the wiki philosophy (ease of editing, tracking of changes, integrated discussions) on the one hand, while focusing on structured information on the other hand. Furthermore OntoWiki is an adaptable application framework, which supports the creation of Linked Data based applications on the web [6]. In addition to the usual features of wikis, OntoWiki provides a sophisticated extension system, such that it can be adapted for a variety of use-cases. Although the wikis usually enable anyone to edit everything, numerous real-world applications require access-control mechanisms. OntoWiki has built-in support for authorization on graph and action level. Furthermore

⁸ <http://myprofile-project.org/>

several authentication protocols can be employed, including amongst others the WebID protocol.

A first use-case for WebID access delegation within OntoWiki arises from the need to *import external data*. Since WebID profiles can contain personal information, access to such data should be restricted with the WebID protocol. Although a user may (or may not in the case of a periodically executed automatic synchronization process) initiate the import procedure manually via the OntoWiki user interface, the actual fetching is done in the background. An OntoWiki instance does not know of any private keys of users of the system. Thus the system is not able to use that information when requesting data. With WebID access delegation though, profiles can be fetched on behalf of the user instead.

Another use-case where access delegation can be employed is within the *Semantic Pingback* protocol [13]. With Semantic Pingback owners of resources can be notified when for example a link to such a resource is created elsewhere on the web. In order to protect the protocol against spam attacks, a Pingback server will fetch the desired resource and check, whether the stated link is indeed contained in the data. The source resource that links to the target resource and thus is fetched by a Pingback server might be access restricted, for example in a scenario where a friending process is initiated [11]. For privacy reasons the owner of the WebID profiles will very likely hide the triples in question (e.g. `foaf:knows`) on anonymous access attempts. With WebID access delegation again, the resources can be fetched by the Pingback server on behalf of the resource owner.

This will require some further changes to the delegation protocol discussed up to now. Specifically the `:secretary` relation currently does not distinguish what kind of responsibilities the *principal* wishes to give to the *secretary*. It is currently assumed that the secretary has full rights. For OntoWiki knowing the Social Network is all that is needed to make access control decisions and full delegation powers may not be needed. The ability to describe more limited secretary relations could be very helpful here.

5 Related Work

OAuth 2.0 [4] is the latest version of the OAuth protocol, which is being presented as an access delegation protocol, enabling users to grant access to third-party services to their personal resources, instead of sharing their passwords with those third-party services. OAuth includes two main parts: obtaining an access token by asking the resource owner (i.e. the user) to grant access, and then using the tokens to access protected resources. Here is a typical example of using OAuth. Ann has a Twitter account and decides to support her favorite project by using Flattr⁹. Since she would rather not create a new account, Flattr offers her the possibility of using her Twitter account to log into Flattr. Once authenticated, Twitter informs her that Flattr would like to access her personal

⁹ A micro payment service: <http://flattr.com>

information. She can now choose if she wants to grant or deny to Flattr access to her Twitter account information.

The advantage here is that Ann only had to use her Twitter credentials to log into Flattr. However, the disadvantage is that Flattr requires an existing trust relationship with Twitter, thus limiting the number of supported services. Even if OAuth provides authentication as a by-product of having the resource owner authorize a third-party client to her resource server, its main focus is on resource authorization rather than on federated identity. In short, OAuth is used to *provide access for external services to local user resources* without disclosing the user's credentials. OAuth is used to allow an agent to *request a user's resources* while WebID delegation is used to allow an agent to *request resources the user could request herself*.

Cross Origin Resource Sharing (CORS) In this paper we proposed to add an `X-On-Behalf-Of` header to each HTTP request made by a *secretary* in order to be able to identify the *principal* she is working for in that particular request. This is quite a major semantic addition to HTTP, but for which there is a widely adopted precedent, namely CORS¹⁰, the Cross Origin Resource Sharing specification at the W3C.

CORS defines two agents, one acting on behalf of the other, and these are: (1) The Browser: it takes the role of the *secretary* as it is acting on behalf of a JavaScript agent - the *principal*. It identifies itself as the browser user. (2) JavaScript code: hosted on the WWW and identified by the Origin Server that hosted it. This JavaScript agent requests resources from the Browser, but it is unable to do it directly. CORS identifies the remote *principal* with an `Origin` header, and uses any HTTP authentication method to identify the *secretary*.

The use cases for CORS are slightly different from the ones for our delegation proposal but the logic is similar. In CORS the browser is using its own identity credentials when connecting to remote servers (e.g. a bank), but it needs to alert the remote service that the request originates from a different agent. The browser also needs to make sure the remote server understands this, especially when authentication is required, to allow the remote server to evaluate whether it trusts the JavaScript agent. In the Section 3 use cases, the remote site trusts the *principal*, but may not be aware of the *secretary*. In our proposal, the identification of the agent in the `X-On-Behalf-Of` header is much more precise. Combined with the ability of the secretary to also identify herself with a WebID, it builds an explicit trust relation between the *principal* and the *secretary*. The vagueness of the `Origin` identity of the *principal* in CORS makes a lot of this close to impossible.

It is a topic of further research to evaluate in which circumstances it is important for the *secretary* to also make sure that the remote server is aware that the connection being made is `X-On-Behalf-Of` the specified *principal*. Otherwise it is possible that the *secretary* herself be made liable for the actions she is making on behalf of her *principal*, especially were she to make PUT, POST or DELETE re-

¹⁰ <http://www.w3.org/TR/cors/>

quests. The CORS protocol determines this by requiring an initial HEAD request to a resource to be made, which will determine whether the server understands CORS. With the help of TLS the server hosting the remote resource could have a WebID profile describing it to be an understander of the `X-On-Behalf-Of` header, making the protocol a lot more efficient and secure once more.

Usage Restrictions such as proposed in [10] can be used to allow clients and servers to specify restrictions on what is allowed to be done with the documents they publish using headers such as `X-UsageRestrictions`. This would for example allow someone posting a picture to a server to specify that this does not entail transfer of ownership.

Similarly, using such an HTTP header, it should be possible for the server serving a protected resource to specify which groups of users a resource representation is intended for. This could remove an important reason for the existence of the `X-On-Behalf-Of` header discussed above. The guard on the server would just need to add the secretary for the principal as one of those agents authorized to access the information. It would know the identity of the secretary for that user by consulting the WebID profile, which it needs to have downloaded anyway.

But how would the guard know that the secretary understood the header and the corresponding restriction vocabulary? After all most servers and clients do not know this vocabulary - it has not even been settled on yet - and so cannot be held responsible for not understanding it. We can reach again here to our WebID profile to settle this: the secretary can specify herself as being an understander of that vocabulary in her profile with a triple such as `:me a trust:VocabularyUnderstander`. The server would check for this statement in the secretary's WebID profile before serving the resource. The advantage of this method would be that it would require less changes to the WebID authentication protocol, could be a lot more flexible but does require the development of a simple usage restriction vocabulary. The secretary relation still remains essential.

6 Conclusion and Future Work

In this paper we justified the need for an agent, such as a Social Web application, to act on behalf of a user in the role of a secretary. The server may need to do things asynchronously in the background for the user (the principal), and may have a number of different users for which this needs to be done efficiently. Our proposal for supporting such a communication schema involves adding a `:secretary` relation from the principal's profile to the agent that acts on its behalf. The secretary then connects to the resource using her own WebID. We then came up with two solutions on how the client can specify that it is acting as a secretary:

- (1) The secretary making the request can add an `X-On-Behalf-Of` header field referring to the WebID of the principal in whose name she acts. This is especially useful if resources return different representations depending on who makes the request, or on behalf of whom a request is made.

(2) The server can specify usage restrictions in the header of the HTTP response, allowing the secretary only to pass the information on to the principal for example. Perhaps this could be done using the WebAccessControl vocabulary¹¹, dgFOAF [9], or ideas coming from the policy creation and management communities [8].

In our future work we will try to close the gap between the policy and access control vocabularies on the one hand and WebID authentication on the other hand. This will require mostly standardization work and seeing which type of solution is the most apt to gain traction in a wider community. As we deploy Social Web servers the need for a such a standard will be felt more and more strongly.

References

1. WebID 1.0: Web Identification and Discovery, 2011. Available at <http://webid.info/spec/>.
2. S. Auer, S. Dietzold, and T. Riechert. OntoWiki - A Tool for Social, Semantic Collaboration. In *Proceedings of the ISWC2006*, volume 4273 of *LNCS*. Springer, 2006.
3. A. Barth. The Web Origin Concept. Technical report, IETF, 2011. <http://tools.ietf.org/html/rfc6454>.
4. Ed. D. Hardt. The OAuth 2.0 Authorization Framework. Technical report, IETF, 2012. <http://tools.ietf.org/html/draft-ietf-oauth-v2-31>.
5. T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol – Version 1.2. Technical report, IETF, 2008.
6. N. Heino, S. Dietzold, M. Martin, and S. Auer. Developing Semantic Web Applications with the OntoWiki Framework. In *Networked Knowledge - Networked Media*, volume 221, pages 61–77. Springer, 2009.
7. P. Hoffman and J. Schlyter. The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA. RFC 6698 (proposed standard), IETF, 2012. <http://tools.ietf.org/html/rfc6698>.
8. L. Kagal, T. Finin, and J. Hendler, editors. *Proceedings of the Semantic Web and Policy Workshop, held in conjunction with the 4th International Semantic Web Conference, 7 November, 2005, Galway Ireland, 2005*.
9. F. Schwagereit, A. Scherp, and S. Staab. Representing Distributed Groups with dgFOAF. In *Proceedings of the (ESWC2010)*, June 2010.
10. O. Seneviratne and L. Kagal. Addressing Data Reuse Issues at the Protocol Level. In *POLICY*, pages 141–144. IEEE Computer Society, 2011. Also published as "Usage Restriction Management for Accountable Data Transfer on the Web" at <http://dig.csail.mit.edu/2011/Papers/IEEE-Policy-httpa/paper.pdf>.
11. H. Story, A. Sambra, and S. Tramp. Friending On The Social Web. In *Federated Social Web Europe 2011*, 2011.
12. H. Story, B. Harbulot, I. Jacobi, and M. Jones. FOAF+SSL: RESTful Authentication for the Social Web. In *Proceedings of SPOT2009*, 2009.
13. S. Tramp, P. Frischmuth, T. Ermilov, and S. Auer. Weaving a Social Data Web with Semantic Pingback. In *Proceedings of the EKAW 2010*, volume 6317 of *LNAI*. Springer, 2010.

¹¹ <http://www.w3.org/wiki/WebAccessControl>