You Are Who You Know

Leveraging webs-of-trust for authentication in identity federations

Bob Hulsebosch, Arnout van Velzen,
Maarten Wegdam & Martijn Oostdijk
InnoValor
Enschede, The Netherlands
e-mail: bob.hulsebosch@innovalor.nl

Remco Poortinga-van Wijnen, Joost van Dijk
SURFnet
Utrecht, The Netherlands
e-mail: remco.poortinga@surfnet.nl

***Abstract*—Digital identity assurance emerges from two aspects: the strength of the authentication solution, or how you identify yourself towards an online service, and quality of the identity proofing and registration process, or how the authentication solution was issued to you. A reliable registration process, however, is often expensive. For example, it may require the establishment of a registration desk, which is not very user friendly as it demands much effort on the part of the user. This paper investigates the feasibility of using webs-of-trust for reliable identity proofing in digital authentication. Webs-of-trust entail communities of people that trust each other, i.e. utilizing social contacts to confirm people's identities. A functional decomposition of an attestation service and protocol for web-of-trust enhanced authentication are provided. A prototype for an attestation service was developed as a proof-of-concept, leveraging LinkedIn as a web-of-trust, and evaluated by users. Finally, characteristics of using web-of-trust for authentication assurance are discussed and a Strengths, Weaknesses, Opportunities and Threats (SWOT) analysis was conducted. Key findings are that while webs-of-trust provide an interesting alternative mechanism for identity proofing that may have merit in use cases where no more efficient registration processes are available, its implementation is complex and mainly challenged by usability.**

Keywords-authentication; web-of-trust; level of assurance; attestation service; identity proofing.

I. INTRODUCTION

Authentication refers to an online process where an entity's identity is verified, typically by providing evidence that it holds a specific digital credential. The strength, or degree or reliability, of the authentication solution is usually expressed in terms of Levels of Assurance (LoA). Two factors are essential in the determination of the LoA [1][2]:

1. The quality of the registration process, i.e., of the identity proofing, registration, and the delivery of credentials that are bound to the registered identity.
2. The strength of the authentication process to establish that a user is who he/she claims to be, which in turn mainly depends upon the strength of the authentication credential.

There is an increasing need for two-factor authentication solutions with cost efficient identity registration. The use of second factor authentication credentials is growing but lack reliable registration processes by which to link a physical person to his/her digital identity information and to his/her authentication credentials during enrolment weaken the overall authentication strength. If this is done poorly, there is little or no assurance that the person using that credential is who he/she claims to be.

Different registration processes and mechanisms apply to identity vetting, proofing, credentialing and linking, and result in different assurance levels. An applicant may appear in person to register or may register remotely. In person registration provides reliable identity proofing, but is expensive (typically from €10 upwards) and not very user friendly (e.g., going to a registration office). Remote registration generally relies on the availability of trusted sources to cross-reference and validate the provided assertions such as name, home address, age, e-mail address, and photo. Remote registration is relatively cheap, but is vulnerable to threats and technically complex. This often leads to weak binding between the user, his authentication credential, and his digital identity. Consequently, the authentication LoA will be low.

An innovative approach to achieve a higher registration LoA, without the cost and overhead of physical registration, is based on the concept of web-of-trust. Using webs-of-trust the authenticity of the binding between an authentication solution and its owner is established via third party user attestations. For instance, if person A claims that user B is using a particular digital identity, it could provide extra confidence for the service provider to allow access to resources that require a certain level of authentication assurance. When Person C also confirms that this digital identity is used by person B, this further increases trust in the digital identity of B. This mechanism can be considered "crowdsourcing of trust". The relations between person A, B, and C, i.e. they share the same social or professional context, could be used to further enhance the level of the authentication assurance.

Particularly in the context of research groups or virtual organizations in which users commonly know each other, such web-of-trust-based authentication LoA enhancement could be executed in an efficient manner. Moreover this approach also promises to capitalize on authentication

 ISBN: 978-1-61208-427-5

functionality provided by social networks such as LinkedIn, Facebook and Google in higher education and research environments. The registration LoA part of the authentication solutions provided by these networks is relatively weak (LoA 1) despite the fact that an increasing number of them are using two-factor authentication (LoA 2 or higher). Web-of-trust based LoA enhancement could help increasing the registration LoA part of these providers and thus could help in increasing the overall LoA.

The objective of this study is to determine the feasibility of using webs-of-trust to enhance the level of the authentication assurance, i.e., having your social connections vouch for your identity. In a way, this implies crowdsourcing assurance for identity verification.

The structure of the paper is as follows. Section II provides some background on webs-of-trust. Section III describes the functional decomposition for an attestation service that enables web-of-trust-based authentication. A protocol for leveraging web-of-trust for authentication and its implementation are described in Section IV. A user evaluation of the prototype that was developed based on this protocol as a proof-of-concept is described in Section V. A Strengths, Weaknesses, Opportunities, and Threats (SWOT) analysis is given in Section VI. Related work is briefly touched upon in Section VII. Finally, Section VIII draws conclusions and provides an outlook for future research.

II. WEB-OF-TRUST

The web-of-trust concept is based on the idea of decentralized trust and social networks. It is used in Pretty Good Privacy (PGP) [3] as an alternative to the centralized trust model that is the basis of a public key infrastructure. In a web-of-trust, each user of the system can choose for himself whom he elects to trust, and who not. Instead of trusting a single entity to validate identities, you validate the identities of the people you know and export this information to a public database. Then, you rely on friends to vouch for the people they know, and those friends to vouch for still more people, and so on until you create a trust chain between any two arbitrary identities. This approach avoids the inherent problems of central authorities, but in practice it is rarely used due to usability issues of tools involved and a lack of user incentive.

A successful web-of-trust should likely be built much like an online social network to obtain the shared experience information for certification, which is a model that hundreds of millions of people all over the world are already comfortable with using. As such, the web-of-trust model can be used to establish the authenticity of the binding between an authentication solution and its owner via third party user attests. Instead of building a whole new web-of-trust, existing trust infrastructures such as PGP, FoaF [4], identity federations, social or professional networks should be readily reused to enhance the registration component of the overall LoA.

LinkedIn is the world's largest business social networking site. One purpose of the site is to allow registered users to maintain a list of contact details of people with whom they have some level of relationship, called Connections. Users may invite anyone (whether a site user or not) to become a connection. LinkedIn provides an interface to obtain basic profile information of users. Information about the connected users in the LinkedIn network of a user can be collected as well. The availability of the information depends on the privacy policy of the connected user. As such, LinkedIn provides sufficient information to determine a reliable set of users that may enhance the level of assurance in someone's identity. The same holds for similar social networks such as Google+, Orkut, and Facebook.

Potentially, the web-of-trust approach combines the best of remote and physical registration practices. There is no need for a physical registration desk as other users in the web-of-trust take over the responsibility to identify users. Confidants in the web-of-trust may use physical presence, phone or email practices for this purpose. However, the attestations from the web-of-trust somehow need to be related to the claimant's digital identity. This needs to be catered for by some kind of attestation service.

III. FUNCTIONAL DECOMPOSITION

Three user-roles can be distinguished in a web-of-trust based authentication scenario:

1. An Asker that wants to use the Attestation Service to enhance assurance of his identity.
2. A Helper that attests for the Asker's identity.
3. A Moderator that wants to have someone's identity (i.e., an Asker) attested.

A functional decomposition results in a number of building blocks that are required to realize web-of-trust based authentication. These are shown in Figure 1.

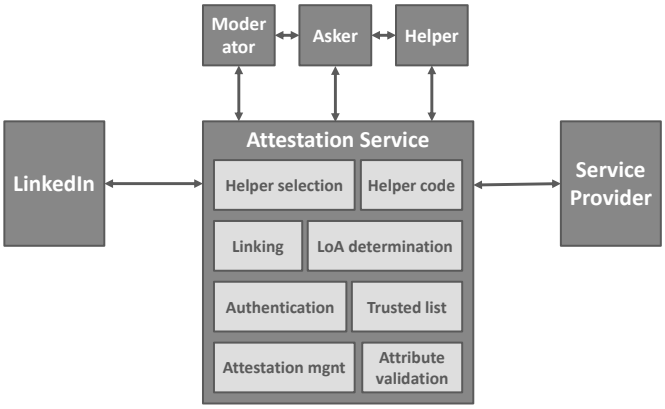

Figure 1. Functional decomposition.

The need for an Attestation Service that facilitates and coordinates the web-of-trust based enhancement of the authentication solution is obvious. Specific functionalities of such an Attestation Service are:

- Authentication of the users (Asker, Helper, Moderator). Authentication could be done in federated manner, via social logon, or locally. Ideally, the Asker as a strong authentication credential with a low LoA due to unreliable registration of the credential to the user's identity.

 ISBN: 978-1-61208-427-5

- Helper selection: who are the best Helpers to attest for the Asker's identity? Candidate Helper selection should be such that it mitigates risks related to herd behavior and fake accounts. Ideally, Helpers come from multiple webs-of-trust and have varying relationships with the Asker (e.g., friend/colleague, recent/longtime, etc.).
- Helper code. The Attestation Service needs to be sure that the selected Helpers are indeed the ones that login to vouch for the Asker's identity. One way to achieve this is by generating a random code that is passed to Helpers that they then have to enter to verify their attests are bona fide.
- Linking of social networks to an Asker or Moderator, i.e., giving the Attestation Service access to the LinkedIn social graph data. This enables the Attestation Service to select meaningful Helpers from the social network. Commonly, users should be able to link their social network accounts to the Attestation Service.
- LoA determination based on Helper attestations. Aspects that could be taken into account are: the number of Helpers, the LoA of Helpers, and the number of invited Helpers that did not vouch. The outcome of the LoA is communicated to the Asker and the service provider.
- Trusted list: establishing a list of trusted Helpers from which Helpers will be primarily selected against the social graph of the Asker. In case of the moderator-scenario, the list consists of the Helpers from the Moderator's social network.
- Attestation management, i.e., keeping track of the attestations given by Helpers, giving feedback to the Moderator or the Asker, asking Helpers to become trusted Helpers.
- Attribute validation could be optional functionality of the Attestation Service. The Attestation Service may ask the web-of-trust to verify self-asserted personal attributes of the user such as a telephone number, age, or address.

IV. PROTOCOL AND IMPLEMENTATON

A. Protocol

The following protocol for web-of-trust enhanced authentication has been implemented in the proof-of-concept:

Step 0: *Building Trust List, Moderation:* A list of trusted potential Helpers may need to be created. A Moderator may make an attestation request for a particular Asker.

Step 1*: Registration of Asker.* Asker registers at the Attestation Service by logging in with his/her federated identity and requests enhancement of authentication. The response of the identity provider contains identity information of Asker. The information at least contains a LoA attribute and value and Asker's federated user identity identifier. Asker is asked to link his/her federated institution account to, e.g., his/her LinkedIn account by logging in with his/her LinkedIn credentials.

Step 2*: Web-of-trust scoping.* The Attestation Service determines who is able to vet for Asker's identity by imposing its trust requirements on the available web-of-trust of Asker. Once the web-of-trust has been determined (in this case LinkedIn) the Attestation Service can start selecting suitable Helpers. Subsequently, Asker is given a vouching code and is asked to contact the Helpers by phone or physically and pass them the code. The use of e-mail is prohibited or deprecated; Asker has to affirm that he/she will adhere to this policy. Asking too many Helpers will burden the Asker as he/she has to contact them.

Step 3*: Passing of vouching code.* Asker calls or meets Helpers and tells them the vouching code. During the phone call or meeting, the Helpers implicitly authenticate the Asker (e.g., via voice or face recognition).

Step 4*: Helper vouching.* The Helper logs in to the Attestation Service with his/her federated identity credentials. The authentication solutions he/she is using must have an equal or higher assurance level than Asker's current level. After successful authentication, the Helper states which Asker he/she wants to vouch for, and the Attestation Service asks the Helper to enter the vouching code. The Attestation Service then validates if the Helper is indeed one of the selected Helpers. If this is the case it asks the Helper to confirm that he/she vouches for Asker's identity. Optionally the Attestation Service may show Asker's personal attributes and asks Helper to validate them. Afterward the Helper logs out. Helper validation can be done in several ways. For instance, the Attestation Service might compare the attributes provided by the identity provider during authentication with those of the selected Helpers from Asker's social network. They should overlap. Another approach is to send the Helper an email with a specific code. The Helper must enter the code together with the vouching code.

Step 5*: LoA determination.* The Attestation Service updates the LoA of Asker based on the number of Helper attestations and their LoA. Mapping web-of-trust-based LoAs to existing frameworks for LoAs like ISO29115 [1] or STORK [2] is not possible; these frameworks do not take web-of-trust mechanisms into account. Consequently, we defined our own web-of-trust-based LoA-framework consists of three levels:

1. WoT LoA1: equal to LoA1 of STORK or ISO29115.
2. WoT LoA2: requires a
   a. minimum of 5 Helpers with LoA1 / WoT LoA1, or
   b. minimum of 3 helpers with LoA2 / WoT LoA2
3. WoT LoA3: requires a
   a. minimum of 8 helpers with LoA2 / WoT LoA2, or
   b. minimum of 5 helpers with LoA3 / LoA4 / WoT LoA3

Also, the number of invited Helpers that did not vouch should be taken into account. These may be considered as 'negative vets'. They have a negative effect on the new LoA. A simple algorithm is to multiply the new WoT LoA with the percentage of positive vets. Note that this is an initial definition of the LoAs, just to get an impression of what it means to step-up to a higher level. The Asker is notified by the Attestation Service about the new LoA, i.e., attestation status.

Step 6*: LoA communication.* Next, Asker can go to a service provider and authenticate himself/herself using his/her federated identity. Multiple solutions are possible for

 ISBN: 978-1-61208-427-5

the communication of the LoA. One possible solution is that the identity provider authenticates Asker at e.g. LoA 1 and communicates this to the service provider. The service provider decides that this is not sufficient and makes a LoA attribute validation request at the Attestation Service. The Attestation Service returns a LoA 2 attribute. This convinces the service provider to allow Asker access to the service. Another solution is that the Attestation Service becomes the (new) identity provider for the Asker, authenticates him/her and communicates the LoA to the service provider. This implies that the Asker must be able to select the Attestation Service as her preferred identity provider.

The different steps are illustrated in Figure 2. The protocol is inspired by the work of Brainard on using vouching by which helpers leverage their strong authentication in order to assist another user, the asker, to perform emergency authentication in case of loss of a second authentication token [5].

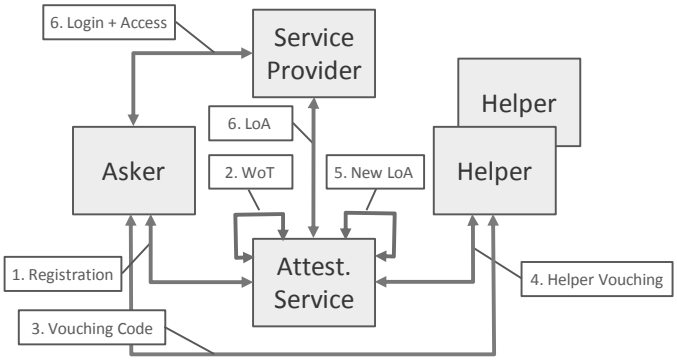

Figure 2. Web-of-trust protocol flow.

B. *Implementation scenario*

A proof-of-concept Attestation Service has been developed. It models a web service for step-up authentication for access to a shared research environment. The Attestation Service allows users to login with a local username and password combination. This can easily by extended to other federated authentication or social login solutions. In case of federated authentication, the attributes that are provided by the identity provider during authentication at the Attestation Service could be used for validation purposes. Furthermore, the Attestation Service offers the user the opportunity to get attested and link the identity provider account to her LinkedIn account.

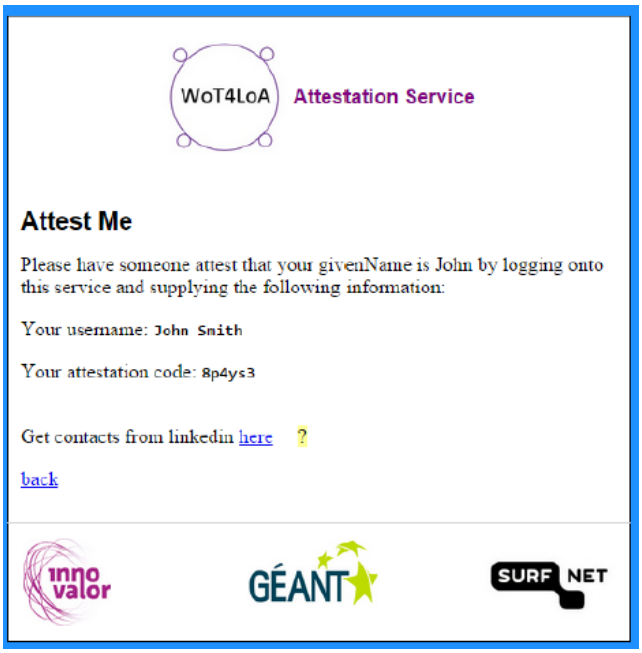

Figure 3. Asker wanting to be attested by Helpers.

The latter provides the Attestation Service the ability to randomly select 5 Helpers from the LinkedIn web-of-trust of the Asker or to select the Helpers from its own list of trusted Helpers. Helpers can put themselves on this trust list by sharing their LinkedIn contacts. Upon submitting an attestation request, the Asker is presented a vouching code that is alphanumeric and consists of five characters, with the instruction to approach the five Helpers (but not by e-mail). Helpers should login to the Attestation Service with their federated account and fill in the vouching code in order to verify the Asker's identity. When all five Helpers have vouched, the hypothetical LoA of the Asker is stepped-up from 1 to 2, granting the Asker access to the concept shared research environment. A Moderator may also request an attestation for an Asker, view the progress of attestations or set his own LinkedIn contacts as the trust list to select Helpers from.

To give an impression of the proof-of-concept several screenshots are shown in Figure 3 and Figure 4.

 ISBN: 978-1-61208-427-5

Figure 4. Helper attesting Asker.

V. EVALUATION

To test the concept of web-of-trust the proof of concept was evaluated by two separate user groups via role-playing scenarios. In addition to specific questions, remarks of the participants as well as non-verbal communication were noted by the observers for evaluation of the prototype.

The outcomes of these two user evaluation tests show that usability is a critical factor for the success of web-of-trust enhanced authentication. Also, the concept is relatively difficult to explain to users. Furthermore, not everyone actively uses social media, e.g., exemplified by not knowing username and password. This could lead to frustration on the part of all three roles (Asker, Helper, and Moderator). The users experienced barriers to contact Helpers and motivating them to provide an attestation. Moreover, it is likely that situations will occur wherein Askers are unable to reach Helpers, e.g., because they do not possess sufficient contact details. Similarly, non-response handling of Helpers could be problematic, since a Helper response cannot be guaranteed. The reliance on others may obstruct or delay authentication and access. So, in order to achieve successful and timely attestation, the whole attestation process should be strictly guided by the Attestation Service. For instance, communication between Helper, Moderator, Asker and Attestation Service could be automated or manually performed through a host of channels, albeit each with their own considerations and trade-offs in terms of responsiveness and 'social pressure'.

According to the evaluation results, there are also trade-offs inherent to Helper selection; not always suitable Helpers were selected. Helper selection is dependent on the information provided by the web-of-trust and the quality of the reasoning algorithm for selecting them. A better selection may be possible if more information is available from the social network used as a source, e.g., the number of likes and comments (cross-) posted on Facebook or the duration of a LinkedIn connection. This information is typically not available to applications outside of the social network itself. Moreover, users may be uncomfortable making the information available to the Attestation Service, as was witnessed by the comments during the prototype evaluation.

Which social network is most appropriate to get attestations from depends on the type of service to be accessed by the Asker. A work-related service would favor the use of a professional social network such as LinkedIn to get attestations from; a leisure or e-commerce type of service might benefit from attestations from the Facebook web-of-trust.

The challenges of automatically selecting the 'right' social network and (then) the 'best' Helpers can be circumvented by restricting the context and work flows for this approach to only Moderator-initiated attestation. The selection of network and Helpers can then conceivably be done by the Moderator, although that does raise the question what additional benefit this approach has if the moderator already has enough information and knowledge to do that selection in the first place (i.e., is attestation really still needed in that situation?). Conceivably removing the social network from the equation altogether and allowing the Moderator to appoint 'delegated Registration Authorities' may work better in those situations.

VI. SWOT ANALYSIS

This section discusses the strengths, weaknesses, opportunities and threats (SWOT) of web-of-trust based authentication approaches, followed by a feasibility analysis to determine whether threats can be mitigated and opportunities leveraged by using the strengths and eliminating the weaknesses.

A. *Strengths*

1) *Cost efficient*

The web-of-trust approach combines the best of remote and physical registration practices. There is no need for an expensive physical registration desk as other users in the web-of-trust take over the identification task, which reduces costs of enrolling strong authentication.

2) *Less intrusive for the user*

Potentially it reduces the intrusiveness for the user as it replaces the cumbersome physical registration overhead by more natural Asker-Helper interactions. Askers, however, may be reluctant to ask a Helper they haven't seen or spoken for quite some time to attest for their identity.

3) *Easy integration in existing federation infrastructures*

The Attestation Service can be easily integrated in an existing identity federation infrastructure. It can leverage the existing federated trust fabric for selecting reliable helpers. The Attestation Service can be positioned as an attribute provider for federated service providers. It can make

 ISBN: 978-1-61208-427-5

assertions about the LoA level of the user. Moreover, contrary to other approaches - such as PGP or FOAF - there is no need for specific client software at the user side.

B. Weaknesses

1) Reliability

ENISA has summarized the possible threats to reputation-based systems. Examples of threats are whitewashing attacks, Sybil attacks, impersonation and reputation theft, bootstrap issues related to newcomers, extortion, denial-of-reputation, ballot stuffing and bad mouthing, collusion, repudiation of data and transaction, recommender dishonesty, privacy threats for voters and reputation owners, social threats such as discrimination or risk of herd behavior, attacking of the underlying infrastructure and the exploitation of features of metrics used by the system to calculate the identity assurance [6]. Most of these threats are also applicable to web-of-trust based authentication. Though the proposed approach does not mitigate all of these threats, their impact is largely influenced by the quality of the Attestation Service's reasoning algorithm. Moreover, using social networks as a web-of-trust for identity attestations makes it more difficult to spoof the system by creating false identities or colluding in groups.

It is relatively easy for an Asker to create multiple LinkedIn, Facebook or Google+ accounts under fake identities and establish via these accounts a web-of-trust of LinkedIn connections or Facebook or Google+ friends (i.e. Sybil attack). This threat is largely mitigated by the fact that the Attestation Service determines the Helpers. Additionally, it can be required for Helpers to have a higher LoA than the Asker; this makes it more difficult to create false Helpers.

False identities can be detected by a relatively poor social ranking. Either they remain disconnected or are connected to a relatively isolated group of 'old friends'. Large-scale analysis of social networks can uncover at least some forms of group collusion. For example, web pages colluding to alter their search engine ranking by linking to one another can be identified and removed if they all have a similar number of links [7]. Alternately, collusion could alter the relative abundance of motifs (small sub-graphs), arousing suspicion if it differs significantly from that of social networks in general [8].

Similarly, herd behavior due to social pressure can be circumvented in a similar manner by selecting Helpers from different webs of trust. Reliable selection functionality may prevent the situation of a group of attackers that collaborate to boost their identity assurance via false attestations.

Services exist that analyze many sources including social networks such as Facebook, Google+, LinkedIn and Twitter to verify cyber identities in real-time. These services are able to detect fake accounts and corresponding identities. An example is Trulioo that offers a service that analyses Facebook profiles and determines whether they're likely to be spoof accounts [9].

However, not all risks can be mitigated completely. Given this weakness, the web-of-trust approach may not be suitable to achieve the highest LoA (i.e., 4), but certainly has the potential to achieve LoA 3.

2) Liability

Another weakness is related to liability. The Attestation Service becomes the authority regarding the authentication LoA of the user. Its owner can, however, not easily be made liable for its LoA claims. The relying service provider has to trust the web-of-trust based LoA claims of the Attestation Service. The fact that both the Attestation Service provider and the relying service provider are in the same federation may help establishing this trust. Additionally a mechanism could be devised that allows service providers to somehow specify trust anchors it 'knows' (e.g., specific persons within institutions) along with their representation in various web-of-trust networks, an approach that fits well if the service providers involved are provided by, or specific to, a virtual organization or collaboration.

3) LoA determination

A web-of-trust based authentication assurance is built from the accumulation of assertions of opinion/judgment by others. It is emergent or generative and is more a matter of judgment than fact. It is an establishment of reputation, as rendered by the attestation service based on a knowable and refutable set of attestations. For example, the trustworthiness that a user's identity is associated to an account is a construct of one or more judgments of other users about this association. Rarely do these sources agree, often because they base their judgment on varying data/experience. There is currently no clear agreement about how to convert the attestations into authentication LoAs. Likely parameters have been determined (number of attestations, the LoA of the helpers, etc.). Evaluation of the model and application in real-life settings has to turn out what suitable parameters are. Inspiration may be obtained from the work of Jøsang [10] and Neisse [11].

In the protocol description, we mentioned that the web-of-trust approach does not fit in the existing LoA frameworks defined by ISO/IEC 29115 and STORK QAA. These frameworks assume there is a central authority that issues the authentication solution and takes care of its binding to a user identity after some form of identity verification. In the web-of-trust based model, the verification role of this central authority becomes less important, i.e., this is done via claims of other users. Adoption of the web-of-trust model in these frameworks is one approach but could take a long time. Another approach is to register our web-of-trust based assurance profiles at the global IANA registry that has been setup for this purpose [12]. The registry is intended to be used as an aid to discovering LoA definitions in protocols that use a LoA concept, including Security Assertion Markup Language (SAML) 2.0 and OpenID Connect. The drawback of a registry approach is that it doesn't provide the registered LoA schemes with any formal status, i.e., it doesn't make them standards that are accepted on a global scale. On the other hand, conforming to standardized frameworks such as ISO/IEC 29115 or STORK QAA provides such a formal status and will make the attestation service more useful in a broader context.

4) Trustworthy exchange of vouching code

The approach implicitly assumes that the Helpers somehow identify and authenticate the Askers via physical

contact or another means that mediates physical communication like a mobile phone call or video session. It doesn't prevent the Asker to send an e-mail to the Helper with the vouching code. This weakness can be mitigated by explicitly asking the Helper to confirm that he had physical or mobile phone contact with the Asker for passing the vouching code. Another option would be to use a customized mobile app that facilitates the exchange of the vouching code to another mobile phone, i.e., the code is only exchanged if the mobile phones are shaken together. This option proves togetherness but excludes the use of remote communication channels such as the mobile phone or a video session. Consequently this narrows down the number of possibilities for exchanging the vouching code in a trustworthy manner.

5) Bootstrapping

Bootstrapping always remains an issue in web-of-trust approaches. The Attestation Service must have sufficient access to social networks or other webs of trust to reliable determine suitable Helpers. Though social networks and interfaces to them are readily available, they need to be made available to the Attestation Service. By making the Attestation Service part of an existing federation and by seducing users to link LinkedIn, Facebook or Google+ accounts to their federated account the bootstrapping problem can be tackled.

6) Usability

Usability is a potential weakness. Particularly in terms of comprehensibility: will the user understand why he/she has to login to the attestation service and pass vouching codes to helpers in order to increase the LoA of their authentication? Users may abort the vouching process because they do not understand why it is needed and consequently may lose confidence in the system. Lack of usability may come at the cost of adoption.

Also, some effort of the Helpers is required. However, Helpers will often have sufficient incentives to attest, e.g., because they need to collaborate with the Asker or want to share something that requires a high LoA. Since the assumption is that the Helpers in some way know and are connected to the Asker via one or more Webs of trust, allowing the Asker to include a reason for vouching in the request may provide further incentive for the Helpers to vouch for the Asker. For instance, the Asker needs to access a Virtual Organization database that is administered by the Helper. These incentives should cater for a reasonably quick enhancement of the user's authentication LoA.

C. Opportunities

1) Useful webs of trust are readily available

Existing webs of trust such as LinkedIn, Facebook, PGP or identity federations are readily available and their exploitation provides sufficient trustworthiness for authentication LoA enhancement purposes.

2) Attribute validation

Many commercial service providers offer discounts for e.g. students or members of a certain community. For these services it is critical to reliably validate the fact if a user is indeed a student or community member, as this is the basis for the discount provided. Other attributes are convenient, but could also be provided by the person directly. As the discounts for students and members are often considerable, these services are highly valuable for users. Attributes such as group membership and age are often used for authorization purposes and must be reliable too.

The Attestation Service can fulfil this need by acting as an attribute validation service. It can ask the Helpers to validate the attributes it has obtained from the Asker's identity provider. Additionally it can ask the Asker to self-assert several attributes (e.g. mobile phone number or gender) and ask the Helpers to validate the assertions. These Helper evaluations will increase the assurance level of the attribute. Similarly to authentication LoAs, this also introduces the need for attribute LoAs. Defining an attribute LoA framework is beyond the scope of this work. An initial attempt is made in the STORK2.0 project [13]. The attribute LoA solution allows the Attestation Service to provide the attributes during authentication, i.e., the service provider is informed about the assurance of the attribute.

Attributes such as student, mobile phone number, e-mail address, group membership and last name are likely to change in time. The reliability of the attestations made by helpers regarding these attributes is time-dependent and has to decrease in time. Consequently, the validation of attributes by helpers should be done one a frequent basis.

The identity providers in existing federations make explicit assertions about the user's identity, e.g., that he/she is a student at the University of Amsterdam. The attestations of other users easily fit into the "claims" architecture of the federated identity infrastructures, and service providers can readily judge the validity of a particular claim based on the authority ascribed to the identity provider in the context of a federated trust framework and the domain. For example, the University of Amsterdam identity provider is arguably definitive regarding the claim that the user is a student, but it is not authoritative for the student's financial status. A project manager is authoritative for the researcher's project membership and a government population register for the age of a student. So, for validation of attributes it is extremely important to know who is authoritative to do so. In a web-of-trust model this can be compensated by using large numbers of attestations: if a large number of helpers attest that a user is of a certain age then this will probably the case. Using large numbers of attestation may also result in large numbers of negative attestations. This may for instance be the case for the validation of membership of a small project team. Only the team members may give positive feedback, whereas the many more other helpers from outside the team may give negative feedback. The context should be taken into account to optimize the validation feedback from the web-of-trust.

D. Threats

1) Loss of privacy

The web-of-trust approach requires intensive linking social network accounts and mining of social network graphs. The Attestation Service potentially obtains insight in the social network of the user and of its connections. Without

 ISBN: 978-1-61208-427-5

proper security measures this may provide a huge privacy threat that will make users reluctant to use the system.
Alternatively, for those concerned that a third party may eventually abuse or be compelled to reveal the social network, decentralized secure computation could produce the aggregate values without a single party having access to the full social network, though such techniques incur substantial computational cost. NodeRank is a decentralized algorithm similar to PageRank that can assign reputations using a social network [14]. Alternately, one can propagate reputation ratings along the social network, where each agent receives information about potential targets through referral chains [15][16]. Cryptographic techniques can further improve decentralized algorithms by allowing precise control over the distribution of information among participants without requiring a trusted intermediary.

E. Summary

Most weakness can be mitigated by the opportunities and threats by strengths. However, two challenges remain to be addressed: usability and liability. The latter can be tackled by integrating the attestation service into the existing trust fabric of the federation (i.e., it becomes a federated service) and possibly by limiting (specific) attestations to a certain context (e.g., membership of a specific organization). The usability challenge strongly depends on how things are presented to the user. This will be the main aspect of the evaluation activity later on in the project.

Looking at web-of-trust LoA enhancement from a business perspective the following question immediately pops into mind: is there a business case for an attestation service? Since there is an increasing need for stronger authentication solutions and physical registration is costly, one would say so. Typically authentication solution service providers could benefit from an attestation service, particularly if standardized frameworks such as ISO/IEC 29115 adopt the approach. An additional value of the attestation service is the opportunity to use it for attribute validation by the web-of-trust. There also is an increasing need for reliable attributes, maybe even more than strong authentication. The sum of all digitally available information about an individual offers enormous potential value [17]. Applications leveraging personal data can boost efficiency, focus research and marketing, and spur the creation of personalized products and services. An important requirement is that the identity attributes are reliable. The attestation service has the ability to meet this requirement

VII. RELATED WORK

The idea of using a web-of-trust is not new and many other reputation systems involve the relationships of participants in the computation of the reputation. Models exist that combine transitive trust (as in certificates or PGP keys) with a reputation rating: If a participant A trusts participant B (with a certain rating) and participant B trusts participant C (with a certain rating), then participant A trusts participant C (with a rating as a function of the other two ratings) [18].

Another way to assign reputation based on social network structure considers each link in the network as an implicit recommendation for a person. Alternatively, weights can be added to the links by allowing users to privately rate their contacts based on characteristics such as trustworthiness. One can then apply a PageRank-like algorithm to assign reputations to individuals [19]. Because PageRank is based on the global structure of the network, it is more difficult to spoof than local network properties, as it is not sufficient to have just anyone recommend a user, but they need to have high reputation themselves.

Brondsema and Schamp have created a system called Konfidi that combines a trust network with the PGP Web-of-Trust [20]. The system implements a metric and mechanism for inferring the trust on the networks formed. The generated network creates trust pathways in between email sender and receiver that can be crawled and using trust mechanisms and metrics, trust values are inferred. This approach has to be extended with LoA-determination functionality to make it suitable for authentication LoA statements.

Calculating trust from social network aggregation is not new [21][22]. These approaches are solely based on the number of claims about a user and do not take into account other trust aspects such as the duration of the connection, presence of the connection in multiple social networks or overlapping attributes like skills and context (e.g. colleague, friend or group membership).

An interesting example is Lenddo [20]. Lenddo is an online platform that utilizes connections, relations and reputation from multiple social media sites such as Facebook to build a credit rating. At Lenddo, everything revolves around the LenddoScore. This number, ranging from 0 to 1,000, is a universal measurement of the user's trustworthiness, with 1,000 being the highest value. Using a proprietary and evolving algorithm, the rating is graphically plotted across categories like Social Data, Trusted Connections, and Financial Performance. This score is what helps the user to obtain approval for loans and services. Lenddo uses social data to ensure that the user is who he says he is. Lenddo also analyzes the user's connections and how strong they are; Lenddo only takes into account the strongest interactions. In many cases this means family, close friends, and coworkers.

These models focus primarily on the calculation of trust and reputation, whereas this work focusses on the translation of crowdsourced trust about an identity into authentication assurance.

VIII. CONCLUSION

Web-of-trust provides an interesting identity proving mechanism that can be used in registration for authentication to attain LoA 2 or 3. Ideally, it should be used in situations where the authentication means has a higher assurance level than its enrollment process (including identity registration and proofing). This could be due to the fact that physical registration was not possible or too expensive. For example, in case of an international collaboration where distance, language or poor electronic communication are barriers to proofing. The main issue of utilizing web-of-trust for

 ISBN: 978-1-61208-427-5

authentication purposes is related to usability, so it is advised to maximize usability in any implementation thereof.

Altogether, this means the applicability of web-of-trust authentication, with regard to necessary level of assurance, alternative registration processes and usability, is use case sensitive. For example, Facebook already has a functionality where users are asked to confirm a photo as their friend, since this concerns an easy extension for a social networking website. The concept of introducing your friends is intuitive, however its implementation for digital authentication is less straightforward.

Future work in the area of web-of-trust for identity management may consist of the following research activities:

- Further optimization of the algorithms and metrics for determining the authentication LoA based on claims from the web(s) of trust.
- Pilot studies to collect user feedback in order to evaluate the approach.
- Further optimization of the algorithms and metrics for determining suitable helper candidates from the web-of-trust. This activity involves complex data mining and analytics.
- Exploration of the use of web-of-trust for other identity-related aspects beyond authentication. Possible aspects are the use of web-of-trust for attribute validation (e.g., is the user indeed a student or older than 18 years?), for authorization purposes, or for linking different user accounts (e.g., the communication of a shared attribute that enables linking).

ACKNOWLEDGMENT

This work is sponsored by the Géant3plus Open Calls program.

REFERENCES

[1] ISO/IEC 29115:2013 Entity authentication assurance framework, available from www.iso.org.

[2] B. Hulsebosch, G. Lenzini, H Eertink, STORK Quality Authenticator Scheme, Deliverable D2.3, March 2009, available from www.eid-stork.eu.

[3] More information about PGP is available at www.pgpi.org.

[4] More information about the Friend of a Friend (FoaF) is available online at www.foaf-project.org.

[5] J. Brainard, A. Juels, R.L. Rivest, M. Szydlo, M. Yung, "Fourth Factor Authentication: Somebody You Know," in ACM CCS, 2006, USA, pp. 168–178, doi:10.1145/1180405.1180427.

[6] E. Carrara, G. Hogben, "Reputation-based Systems: a security analysis," ENISA position paper, October 2007.

[7] M. R. Henzinger, R. Motwani, C. Silverstein, "Challenges in web search engines," Newsletter ACM SIGIR Forum, Vol 36, Issue 2, Fall 2002, pp. 11-22.

[8] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, U. Alon, "Network Motifs: Simple Building Blocks of Complex Networks," Science, vol. 298, pp. 824-827, 2002, doi:10.1126/science.298.5594.824.

[9] More information about Trulioo is available online at www.trulioo.com.

[10] A. Jøsang, R. Ismail, C. Boyd, "A survey of trust and reputation systems for online service provision," Decis. Support Syst., vol. 43, March 2007, pp. 618-644, doi:10.1016/j.dss.2005.05.019.

[11] R. Neisse, "Trust and privacy management support for context-aware service platforms," PhD thesis, University of Twente. CTIT Ph.D. Thesis Series No. 11-216 ISBN 978-90-365-3336-2, 2012.

[12] The LoA Registry, more information available online at http://levelofassurance.org/process.html.

[13] STORK2.0 project, more information is available online at www.eid-stork2.eu.

[14] P. Li, X. Qiu, NodeRank: An Algorithm to Assess State Enumeration Attack Graphs, 8th international conference on Wireless Communications, Networking and Mobile Computing (WiCOM), pp. 1-5, doi: 10.1109/WiCOM.2012.64785852012.

[15] B. Yu, M. P. Singh, "A social mechanism of reputation management in electronic communities," Proc. 4th International Workshop on Cooperative Information Agents IV, Springer-Verlag London, 2000, pp 154-165, ISBN:3-540-67703-8.

[16] G. Zacharia, A. Moukas, P. Maes, "Collaborative reputation mechanisms in electronic marketplaces," Proc. 32nd Hawaii Intl. Conf. on System Sciences (HICSS), 1999, vol. 8, p. 8026, ISBN:0-7695-0001-3.

[17] Boston Consultancy Group, The Value of our Digital Identity, 2012, available online at www.libertyglobal.com/PDF/public-policy/The-Value-of-Our-Digital-Identity.pdf.

[18] F. Kerschbaum, J. Haller, Y. Karabulut, P. Robinson, "PathTrust: A Trust-Based Reputation Service for Virtual Organization Formation," Proc. 4th Int. Conference on Trust Management (iTrust), vol. 3986 LNCS, 2006, pp. 193–205.

[19] L. Page, S. Brin, R. Motwani, T. Winograd, "The pagerank citation ranking: Bringing order to the web," Technical report, Stanford Digital Library Technologies Project, 1998.

[20] D. Brondsema, A. Schamp, "Konfidi: Trust Networks Using PGP and RDF," Proc. Of the WWW'06 Workshop on Models of Trust for the Web (MTW'06), Edinburgh, Scotland, UK, May 22, 2006.

[21] S. Noh, "Calculating trust using aggregation rules in social networks", Proc. 4th international conference on Autonomic and Trusted Computing, Hong Kong, China, pp 361-371, 2007, doi:10.1007/978-3-540-73547-2_38.

[22] H.R. Singh, A. Neelima, L.S. Singh, S.Ib. Singh, "A Model of Computing Trust in Web Based Social Network Using New Aggregation and Concatenation Operators," International Journal of Computer Science and Network, Volume 2, Issue 4, August 2013, IJCSN International Journal of Computer Science and Network, Volume 2, Issue 4, August 2013, ISSN:2277-5420.

[23] Lenddo, more information available online at www.lenddo.com/pages/faq.

 ISBN: 978-1-61208-427-5