31-03-2015

Open Call Deliverable OCU-DS4.1
A Feasibility Study (WoT4LoA)

You Are Who You Know

Leveraging Webs of Trust for Authentication in Identity Federations

Open Call Deliverable OCU-DS4.1

Grant Agreement No.:	605243
Activity:	NA1
Task Item:	10
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	InnoValor
Document Code:	GN3PLUS14-1306-36
Authors:	Bob Hulsebosch, Arnout van Velzen, Martijn Oostdijk & Maarten Wegdam (InnoValor)Remco Poortinga-Van Wijnen & Joost van Dijk (SURFnet)

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7 2007–2013) under Grant Agreement No. 605243 (GN3plus).

Abstract

In the real world groups of people often trust each other's identities and it is common for people to introduce new persons into a group, i.e. these groups form webs-of-trust. Such networks of trusted identities manifest themselves in the digital world as well, most recognizably as social networks. Identities are essential to performing transactions of all kinds in the digital world, making them essential to the fabric of modern society. Consequently, digital identities need to be trustworthy, i.e. there needs to be some kind of assurance that the presented digital identity indeed belongs to the correct user.

Digital identity assurance emerges from two aspects: the strength of the authentication solution, or how you identify yourself towards an online service, and identity registration, or how the authentication means was issued to you. Proofing one's identity is key to the registration process, for example by sending in a copy of one's passport.

This study investigates the feasibility of using webs-of-trust for more reliable identity proofing in digital authentication. To this effect, web-of-trust as an identity validation method was compared to other registration and authentication methods. Also, use cases were identified for web-of-trust-authentication, as well as the need for a specific attestation service that enables people requiring authentication enhancement to have their identities vetted by acquaintances. Furthermore, the functional decomposition of an attestation service was laid out, including its decision space and protocol. A prototype for an attestation service was developed as a proof-of-concept utilizing LinkedIn as a web-of-trust, and evaluated by users. Finally, characteristics of using web of trust for authentication assurance are discussed and a SWOT-analysis was conducted. Conclusions address the impact of web-of-trust-authentication on identity federations and how to calculate its level of assurance, among other questions. 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. For example, there are many options and trade-offs involved in designing an attestation service, which should be clear and easy to users so as not for them to forfeit the identity verification procedure.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

To navigate the digital landscape you need to identify yourself, and within the networked society this has made digital identities crucial to the functioning of much of our social and economic infrastructure. Digital identities are the result of two different processes; authentication, or how you demonstrate who you are, and registration, or how this means of identification was provided to you. Registration itself also consists of two essential processes; proofing that you are who you say you are, and binding and issuing a means of authentication to you. For example, sending in a copy of your passport gives some certainty of your true identity, sending a document with a username and password to your home address reliably communicates an authentication credential to you. The degree of confidence with which a digital identity is true may be captured in a Level of Assurance (LoA, commonly a scale of 1-4). In real life, your friends and acquaintances often introduce you to other people, i.e. they guarantee your identity to some degree. In a community or group of people this implies that the participants trust each other's identities and by extension each other's motives; they form webs-of-trust. Such networks of people manifest in the digital world as well, most recognizably in the form of social networks like Facebook or LinkedIn.

The objective of this study was to determine the feasibility of using webs-of-trust to enhance the level of the authentication assurance. There is an increasing need for stronger authentication solutions that go beyond username and password. The use of second factor authentication credentials is growing but lack solid 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. A solid registration process, however, is expensive as it usually requires the establishment of a registration desk and is not very user friendly, as he/she has to go to the registration desk. The latter requirement can even be practically impossible to meet for remote users.

Key take-away: There is an increasing need for two-factor authentication solutions with cost efficient identity registration.

To use webs-of-trust for identity verification someone who wants access to certain resources (the Asker) or someone who requires the Asker to have access (the Moderator) initiates an attestation process. A set of people who are deemed reliable to vouch for the Asker's identity (the Helpers) are selected from one or more webs-of-trust by some selection algorithm. To this end existing digital webs-of-trust are best utilized, for example through social networking sites. The Helpers are requested to attest and proceed to confirm the Asker's identity (and possibly one or more attributes) through some digital application (the Attestation Service). If enough Helpers have guaranteed the Asker's identity, the level of assurance is high enough to grant the Asker access.

Web-of-trust as an alternative registration solution was compared to other registration processes and authentication solutions to assess where it fits in the authentication landscape and what level of assurance it could enable. It follows that web-of-trust-authentication is suitable to reach at least level 2, as it is potentially as reliable as a copy of an identity-document, or even somewhat stronger. Level 3 registration quality is less easy to achieve as this requires very accurate selection of helpers from the web-of-trust. Moreover, the identity of the helpers must be determined with a reasonable high assurance level as well. The web-of-trust provides identity proofing, but does not arrange provisioning of credentials, so that element of identity assurance is conditional, as is the strength of the credential and the authentication process. Theoretically, web-of-trust for identity validation could thus be combined with any authentication process.

Key take-away: Web of trust provides an alternative or supplementary identity proofing and registration solution that potentially could reach level 3 assurance.

Next three use cases were identified as exemplars of how web-of-trust authentication may be applied. Firstly, a group of collaborating researchers want to grant a new and previously unknown member access to shared resources. Secondly, identities in social networks are commonly self-asserted, hence web-of-trust could give some certainty about the real identity of users. Thirdly, someone may need a higher level of assurance to gain access to certain resources/services with their account, wherein web-of-trust may strengthen the identity validation. Altogether, web-of-trust may have merit in the creation of a new account, to validate an existing identity, or in step-up-authentication.

Key take-away: Valid use cases exist for the deployment of web-of-trust enhanced authentication.

Subsequently, the functional decomposition of a web-of-trust solution was worked out, which resulted in a set of functional and qualitative requirements, as well as the decision space of such a solution. The main insight here was that such an attestation service knows many variations and options that are mainly context-dependent and dictate the ultimate architecture, as well as the process and reliability of web-of-trust authentication. Based on the functional decomposition a recommended protocol was designed and a proof-of-concept developed for implementation in the context of an identity federation for higher education and research. The proof-of-concept entailed a prototype of an attestation service for access to a fictional dashboard upon vetting by LinkedIn contacts.

Key take-away: Specific functionality (i.e. an Attestation Service) is required that facilitates and coordinates the enhancement of the authentication solution via web-of-trust.

Key take-away: The variations and options for the implementation of web-of-trust enhanced authentication functionality (e.g. via an Attestation Service) are numerous and complex.

A web-of-trust enhanced authentication SWOT-analysis was conducted as part of the feasibility study. The main outcomes are that while webs-of-trust comprise an opportunity for efficient and reliable identity proofing, the main challenges are in balancing usability and liability. For example, while attestation is an intensive and complex process it should remain clear and easy to users so they do not reject verification. Also, the resultant level of assurance should be acceptable to prevent unjustified access to resources. This concedes that usability is to be a prominent objective in any web-of-trust implementation. Webs-of-trust could also be leveraged to confirm user information (attributes), which presents a promising opportunity as well.

Key take-away: While web-of-trust is an opportunity for efficient and reliable identity proofing, the main challenges are in usability and liability.

Finally, preliminary user evaluation tests conducted on the prototype confirmed that usability is a critical factor for the success of web of trust enhanced authentication. The concept is relatively difficult to explain to the user. Also, the users experienced barriers for contacting Helpers and for motivating them to provide an attestation. There is a risk that the whole attestation process may take too long when helpers are reluctant to assist. The attestation process should be strictly guided by the Attestation Service in order to achieve successful and timely attestations.

Key take-away: Preliminary user evaluation tests confirm that the complexity of the concept is a major hurdle and indicate that it may be socially uncomfortable and time-intensive.

Overall, web-of-trust provides an interesting identity proving mechanism to be used in registration for authentication to attain LoA 2 or 3. It should be used in combination with authentication means of the desired level of assurance, if and when existing registration processes are impossible (i.e. too expensive or not user friendly). For example, in case of an international collaboration where distance, language or poor electronic communication are barriers to physical proofing. The main issue in all cases lies with usability, so it is advised to maximize usability in any fashion of implementation.

Altogether, this means that 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 constitutes an easy extension for a social networking website. Within the context of federations for higher education and research web-of-trust authentication would have merit if these conditions are met, for example in an international research collaboration. *The concept of introducing your friends is intuitive in real life, its implementation for digital authentication less straightforward*.

1 Introduction

In the real world groups of people often trust each other's identities and it is common for people to introduce new persons into a group, i.e. these groups form webs-of-trust. Such networks of people manifest in the digital world as well, most recognizably in the form of social networks like Facebook or LinkedIn. To navigate the digital landscape you often need to identify yourself, and within the networked society this has made digital identities crucial to the functioning of much of our societal infrastructure. Identities are essential to performing transactions of all kinds in the digital world, making them essential to the fabric of modern society. They need to be trustworthy, i.e. there needs to be some kind of assurance that the presented digital identity indeed belongs to the user.

Digital identity assurance emerges from two aspects: the strength of the authentication solution, or how you identify yourself towards an online service, and identity registration, or how the authentication means was issued to you. Proofing one's identity is key to the registration process, for example conducted by sending in a copy of your passport. This study investigates the feasibility of using webs-of-trust for identity proofing in digital authentication, i.e. having your social connections vouch for your identity. In a way, this implies crowdsourcing identity verification. Before delving into webs-of-trust further, firstly let us look at digital identities in further detail.

1.1 Authentication

Authentication refers to the process where an entity's identity is authenticated, typically by providing evidence that it holds a specific digital identity such as an identifier and the corresponding credentials. Examples of types of credentials are passwords, one-time tokens, digital certificates, and phone numbers (calling/called). The strength 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. The quality of the registration process, i.e. of the identity proofing, registration, and the delivery of credentials which bind an identity to a token (Registration LoA).

2. The strength of the authentication mechanism to establish that a user is who he claims to be, which in turn mainly depends upon the strength of the authentication solution (Authentication solution LoA).

Figure 1: Factors in determining LoA.

Four levels of identity assurance for electronic transactions requiring authentication are commonly used:

Level 1 – Little or no confidence in the asserted identity

Level 2 - Some confidence in the asserted identity

Level 3 - High confidence in the asserted identity

Level 4 - Very high confidence in the asserted identity

These levels have been specified by multiple independent parties – such as NIST[1], STORK[2] and national eID trust frameworks [3] – and have recently been standardised in the ISO/IEC 29115 (Entity Authentication Assurance Framework) specification.

The process by which to link a physical person to his/her digital identity information and to his/her authentication credential is critical to avoid registration fraud. If this is done poorly, there is little or no assurance that the person using that credential to authenticate and access services and information is who he/she claims to be. It could be anyone including impostors that impersonate a claimed identity; it could be multiple people over time, or even subscribers that deny registration. If the linking is weak, even the most complete personal information and the strongest credential will not improve the assurance of identity. These mutual relations between someone's physical identity, digital identity, and authentication credentials can be depicted in a triangle (see Figure 2).

[1] NIST special publication 800-63

[2] The Standardisation Forum. Assurance levels for authentication for electronic government services. A guide for government organisations.

[3] See e.g. www.eherkenning.nl for the Dutch eID trust framework.

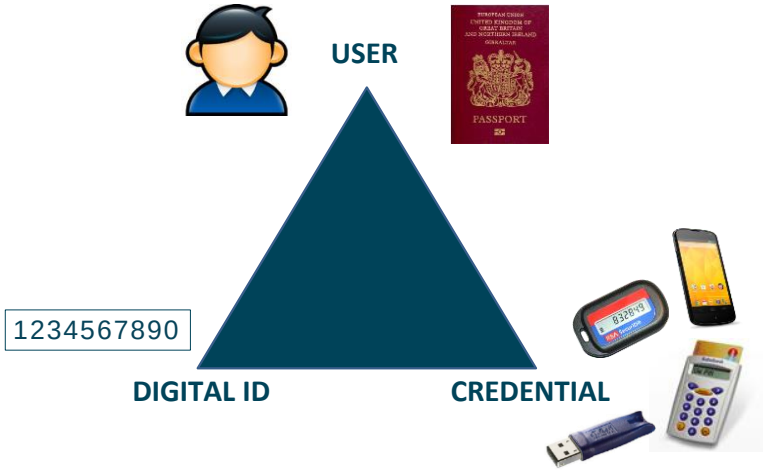

Figure 2: Authentication triangle.

Different registration processes and mechanisms apply to identity vetting, proofing, credentialing and linking, and result in different registration assurance levels. An applicant may appear in person to register, or the applicant may register remotely. In person registration is more secure but 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 more vulnerable to threats and technically complex to achieve. Often this leads to a low authentication LoA (see figure below).

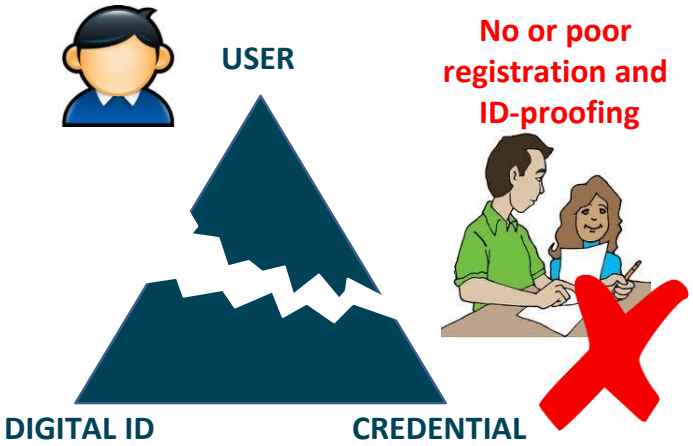

Figure 3: Damaged authentication triangle by poor registration.

An innovative approach to achieve a higher registration LoA, without the cost and overhead of physical registration and complexity of many other remote registration solutions, is based on the concept of web of trust. Web of trust is a concept (e.g. used in Pretty Good Privacy) to establish the authenticity of the binding between an authentication solution (e.g. public key) and its owner via third party user attests. For instance, if person A claims that user B is using a particular authentication solution, it could provide extra confidence for the service provider to allow access to resources that require a high level of authentication assurance. Person C could also claim to know B and his authentication mechanism

thereby even further increasing the trust in the identity of B. It can be considered as a kind of "crowdsourcing of trust". The social or research context of the user could be used to enhance the level of the authentication assurance.

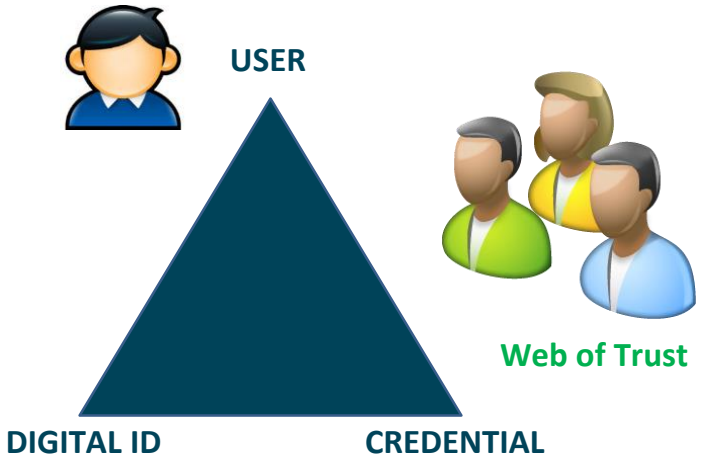

Figure 4: Authentication triangle enhanced by Web of Trust.

Particularly in the context of research groups or virtual organizations in which users know each other, such web of trust based authentication LoA enhancement could be executed in an efficient manner. Moreover this approach also promises to make it easier to use social identities provided by e.g. Facebook and Google in higher education and research environments. The registration LoA part of these popular social identity providers is relatively weak (LoA 1) despite the fact that an increasing number of them are using two-factor authentication (LoA 2). Web of trust based enhanced LoA could help increasing the registration LoA part of these providers and thus could help in increasing the overall LoA.

1.2 Objectives and approach

The objective of the research is to determine the feasibility of using web of trust to enhance the level of the authentication assurance. Specific research questions that will be answered are:

- What are the possible approaches and protocols to increase the LoA associated with identities based on the web of trust model?
- How to define the metric to determine a certain LoA and what factors need to be taken into account in this metric (number of claims, quality of the claims, history of the claims, …)?
    - How many claims are needed to step-up from e.g. LoA 1 to LoA 2?
    - How to guarantee the authenticity of the claims made by others? How to deal with 'negative' claims?

    - How to leverage the web of trust enhanced LoA with virtual organization or research group context where users typically know each other?

- Can a similar approach be used to improve the quality of the attributes related to the user's identity (e.g. group membership, age, gender, student, …)? Such verification of user attributes could be used for authorization purposes.

- What is the impact of this approach on the identity providers (IdPs) that are responsible for user authentication? Can they be made web of trust claims aware and are they able to communicate the enhanced LoA such that service proivders are able to deal with it?

- Does this approach fit in the existing LoA frameworks defined by e.g. ISO, NIST and STORK? Most existing LoA frameworks assume there is a sort of 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. This should be taken into account in the LoA frameworks.

- What will be the highest LoA achievable with this approach?

The proposed approach to address these questions is as follows:

- Determine illustrative use cases that motivate the use of web of trust to enhance authentication assurance.

- To contrast web-of-trust based identity verification to common means of registration and authentication, conduct an inventory of strong authentication and registration solutions including an evaluation of their usability in identity federations for education and research. Evaluation criteria amongst others will be costs, user friendliness, implementation effort for enrolment; scalability, and the level of assurance it offers.

- Perform an analysis of the feasibility of web of trust based LoA enhancement including the description of (e-Research) scenario's, an overview of web of trust based approaches, claims based LoA metrics, architectural approaches, and taking into account specific authentication mechanisms with varying LoAs (username password, two factor).

- Provide protocol description to implement a web-of-trust model to increase the LoA associated with an identity.

- Develop a proof-of-concept for the proposed protocol including integration with a selected strong authentication solution.

- Evaluation of the proof-of-concept and directions for further exploitation of the concept.

1.3 Reading guide

The structure of the deliverable is as follows. Section 2 provides background information about strong authentication and webs-of-trust. Section 3 gives an overview and comparison of different registration and authentication mechanisms. Subsequently, two illustrative use cases are described in section 4. Based on these use cases the functional requirements for web of trust enhanced authentication are derived. Section 5 describes the functional decomposition for an attestation service that enables web-of-trust-authentication, including its decision space and a protocol for leveraging web of trust for authentication enhancement. A SWOT-based feasibility analysis is conducted in section 6. A user evaluation of the prototype that was developed based on this protocol as a proof-of-concept is described in section 7. Finally, section 8 draws conclusions.

2 Web of Trust

2.1 Concept of 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[4]) 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 your friends to vouch for the people they know, and those friends to vouch for still more people, and so on until you can create a trust chain between any two arbitrary identities. This approach avoids the inherent problems of central authorities, but in practice it is barely used due to usability issues of tools involved and the lack of user incentives.

A successful web of trust should likely be built very much like a social networking site, because that is how we obtain the shared experience information for certification, and that is the 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[5], identity federations, social or professional networks should be readily reused to enhance the registration part of the overall LoA.

LinkedIn is the world's largest business social networking hub. Launched in 2003, LinkedIn has millions of users and is implemented in over 200 countries. 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 can invite anyone (whether a site user or not) to become a connection.

LinkedIn provides an interface to obtain basic profile information of users. The fields of the basic profile are amongst others name, headline, location, industry, connections, honours, interests, skills, education, and recommendations. The contact information interface provides fields like telephone number and email address. 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

[4] See PGP website for more information: http://www.pgpi.org/.

[5] Friend of a Friend project (FoaF), see http://www.foaf-project.org/.

that may reliably contribute to enhance the level of assurance in someone's identity. The same holds for other social networks such as Google+ and Facebook.

Particularly in the context of virtually collaborating organizations in which users know each other, such web of trust based LoA enhancement could be executed in an efficient manner. Moreover this approach also makes it easier to use social identities provided by e.g. Facebook and Google. The registration LoA part of these popular social identity providers, however, 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 enhanced LoA could help increase the registration LoA part of these providers and thus could help in increasing the overall LoA.

2.1.1 Reputation systems

The outcome of web of trust identity validation can be considered as a reputation. A reputation system computes and publishes reputation scores for a set of objects (e.g. service providers, services, goods or entities) within a community or domain, based on a collection of opinions that other entities hold about the objects. The opinions are typically passed as ratings to a reputation center which uses a specific reputation algorithm to dynamically compute the reputation scores based on the received ratings.

Entities in a community use reputation scores for decision-making, e.g. whether or not to buy a specific service or good. An object with a high reputation score will normally attract more business than an object with a low reputation score. It is therefore in the interest of objects to have a high reputation score.

Since the collective opinion in a community determines an object's reputation score, reputation systems represent a form of collaborative sanctioning and praising. A low score represents a collaborative sanctioning of an object that the community perceives as having or providing low quality. Similarly, a high score represents a collaborative praising of an object that the community perceives as having or providing high quality. Reputation scores change dynamically as a function of incoming ratings. A high score can quickly be lost if rating entities start providing negative ratings. Similarly, it is possible for an object with a low score to recover and regain a high score.

Reputation systems are related to recommender systems and collaborative filtering, but with the difference that reputation systems produce scores based on explicit ratings from the community, whereas recommender systems use some external set of entities and events (such as the purchase of books, movies, or music) to generate marketing recommendations to users. The role of reputation systems is to facilitate trust (Resnick et al., 2000; Jøsang, Ismail & Boyd, 2007), and often functions by making the reputation more visible.

Reputation systems are often useful in large online communities in which users may frequently have the opportunity to interact with users with whom they have no prior experience or in communities where user generated content is posted like YouTube or Flickr. In such a situation, it is often helpful to base the decision whether or not to interact with that user on the prior experiences of other users.

Reputation systems may also be coupled with an incentive system to reward good behavior and punish bad behavior. For instance, users with high reputation may be granted special privileges, whereas users with low or unestablished reputation may have limited privileges.

A simple reputation system, employed by eBay, is to record a rating (either positive, negative, or neutral) after each pair of users conducts a transaction. A user's reputation comprises the count of positive and negative transactions in that user's history.

More sophisticated algorithms scale an individual entity's contribution to other nodes' reputations by that entity's own reputation. PageRank (Google) is such a system, used for ranking web pages based on the link structure of the web. In PageRank, each web page's contribution to another page is proportional to its own PageRank, and inversely proportional to its number of outlinks.

Reputation systems are also emerging which provide a unified, and in many cases objective, appraisal of the impact to reputation of a particular news item, story, blog or online posting. The systems also utilize complex algorithms to firstly capture the data in question but then rank and score the item as to whether it improves or degrades the reputation of the individual, company or brand in question.

The reputation of a digital identity is closely linked to the reputation of the issuer. Attributes are much more valuable if asserted by a reputed issuer. eBay is one of the systems in which a person can improve its originally very low digital identity reputation based on well rated transactions. Reputations-based systems have a process that will compute trust based on behaviour of claimants and rating of members. This is not necessarily limited to humans, but could be used for identity providers and other roles in federations as well.

Information systems that automatically and systematically gather trust statements of different issuers, accumulate and amalgamate the different subjective opinions and trust values according to the trustworthiness of their issuers in order to compute a resulting estimation of the trustworthiness of a given trustee, which may then serve as basis for decision making (see Fig. 2). This resulting opinion contains (in contrast to the previous trust opinions) not only the opinion of one single individual, but a mixture of opinions of different individuals. To distinguish between these different types of opinions we will use the term trust value for the opinion of one single entity based only on own knowledge and experiences, whereas a reputation value represents a value computed from the opinions of different entities.

In order to get reputation systems to work, empirical facts and circumstances need to be numerically (or symbolically) represented, i.e., the strength of trust relations has to be quantified and measured by an associated trust value. There exists a large number of proposed trust models with different approaches to represent trust values. Three basic types of reputation systems can be distinguished:

- Flat reputation systems (Type A). These systems (e. g., in eBay) are very simple. The reputation values are computed from all available trust opinions of all entities. The opinion of each entity has the same weight, meaning dishonest entities have the same influence on the resulting trust value as honest entities.

- Recursively weighting reputation systems (Type B). These systems try to improve the quality of the computed reputation value by increasing the weight of higher ranked opinions. Reputation values of all entities are therefore computed iteratively: The new reputation values of all entities are computed from the opinions of all other entities weighted by their

reputation values of the last iteration. It is, however, possible that a large group of colluding malicious entities dominates the "public opinion" and manipulates the computed reputation values.

- Personalized reputation system with trust anchor (Type C). These systems (e. g., as proposed by Maurer and Jøsang) aim to resist this kind of attacks. They always start with a "save" set of a priori trusted entities (the so-called trust anchor or trust root), which normally consists of the requester himself. First, only the opinions of the a priori trusted entities are taken into account. Next, also the opinions of entities which have been found to be trustworthy in the previous iteration are taken into account, too. This process is repeated until the opinions of all "reachable" trustworthy entities are included in the reputation value computation. Note that opinions of untrustworthy entities are ignored as long as the opinions of the trust anchor entities are correct. In contrast to the other two reputation systems, this type of system is personalized, because requesters with different trust anchors will in general obtain different reputation values for the same trustee. In the context of virtual organizations, this reputation system seems applicable.

The various reputation system models are shown in Figure 5.

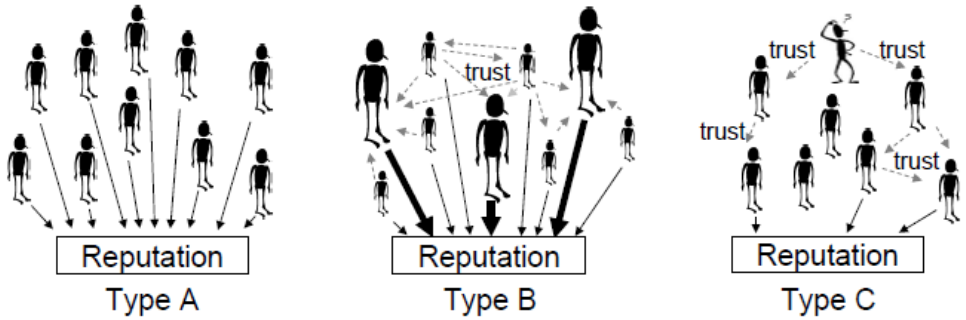

Figure 5: Classification of reputation systems[6].

2.1.2 Use of social networks as web of trust for reputation building

The idea of using a web of trust is not new and many other reputation systems involve the relationships of the participant in the computation of the reputation. Models exist that combine transitive trust (as in e.g. 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).[7]

[6] From Andreas Gutscher, Jessica Heesen, and Oliver Siemoneit, Possibilities and Limitations of Modeling Trust and Reputation, Proceedings of the Fifth International Workshop on Philosophy and Informatics: WSPI-2008; Kaiserslautern, Germany, April 1-2 2008.

[7] Florian Kerschbaum, Jochen Haller, Yücel Karabulut, Philip Robinson: PathTrust: A Trust-Based Reputation Service for Virtual Organization Formation. iTrust 2006: 193-205.

Pretty Good Privacy

Pretty Good Privacy (PGP) is a program used to encrypt and decrypt e-mail over the Internet. It can also be used to send an encrypted digital signature that lets the receiver verify the sender's identity and know that the message was not changed en route. PGP uses a variation of the public key system. In this system, each user has a publicly known encryption key and a private key known only to that user. You encrypt a message you send to someone else using their public key. When they receive it, they decrypt it using their private key.

For sending digital signatures, PGP uses an algorithm that generates a hash from the user's name and other signature information. This hash code is then encrypted with the sender's private key. The receiver uses the sender's public key to decrypt the hash code. If it matches the hash code sent as the digital signature for the message, then the receiver is sure that the message has arrived securely from the stated sender.

Every user in a public key system is vulnerable to mistaking a phony key (certificate) for a real one. Validity is confidence that a public key certificate belongs to its purported owner. Validity is essential in a public key environment where you must constantly establish whether or not a particular certificate is authentic.

In a PGP environment, any PGP user can validate another PGP user's public key certificate. It does that by signing another's key. Multiple PGP can do that and as this process goes on, it establishes a web of trust. However, in such a web of trust a certificate is only valid to another user if the relying party recognizes the validator as a trusted introducer. That is, you trust my opinion that others' keys are valid only if you consider me to be a trusted introducer. Otherwise, my opinion on other keys' validity is moot.

Another way to assign reputation based on social network structure considers each link in the network as an implicit recommendation for the person linked to. 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[8]. 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.

An interesting example is Lenddo[9]. Lenddo is an online platform which 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

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

[9] Lenddo, see https://www.lenddo.com/pages/faq.

Connections, and Financial Performance. This score is what helps the user to obtain approval for loans and services, though currently, only loans and free financial education are offered.

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, best friends, and close coworkers.

2.2 Calculating Levels of Assurance

Various approaches to calculate reputation values exist.[10] The most important ones are:

- Summation and average based: It aggregates the ratings and the overall single reputation score is calculated by summing or averaging. The most well-known summation system is eBay and ratings in this system are represented by numeric rating. The total ratings of a target are added together to represent the target's final reputation score. These models are easy to understand and follow because the reputation is represented by a single number. However, such models only provide a primitive view on the target's performance and the single number representation does not cover the proportion or number of negative/positive ratings of the target.

- Discrete trust models: These models use discrete labels to represent the reputation. By using discrete labels, users can quickly determine a meaning for a reputation measure. Nevertheless, it is not mathematically tractable thus there is no method to determine reputation confidence.

- Bayesian frameworks: Reputation models based on Bayesian frameworks depict the reputation value as probabilities between [0, 1]. These models have been popular for peer-to-peer networks and sensor systems, and they rely on ratings being either positive or negative, and use probability distributions to come up with reputation scores.

Since authentication LoAs are expressed in discrete values, the discrete trust model approach seems the most straightforward approach. For the other two approaches, translation functionality will be required to map a certain reputation value to a LoA value.

[10] For an overview see Neisse, R.: 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, 2011.

3 Authentication solutions

Prior to combining the concept of web of trust with authentication, an overview of existing authentication solutions is required. This section explains the concept of authentication, levels of assurance, and provides overviews of registration LoA and authentication LoA solutions.

3.1 Background

3.1.1 Authentication factors

Authentication systems are frequently described by the authentication factors that they incorporate. The three factors often considered as the cornerstone of authentication are:

1. Something you know (for example, a password);
2. Something you have (for example, an ID badge or a cryptographic key); and
3. Something you are (for example, a fingerprint or other biometric measurement).

Authentication systems that incorporate all three factors are usually considered to be stronger than systems that incorporate only one or two of the factors. Multiple factors raise the threshold for successful attacks. It is increasingly difficult for an attacker to gain control of all required factors when multiple factors are present.

Sometimes a fourth factor is added: Somewhere you are at a certain point in time. This fourth factor basically takes the situational context into account during the authentication process. Context factors of interest are location (based on IP-address, GSM cell ID, or GPS), time, or type of device.

An authentication solution can be implemented in various ways. For instance, by presenting multiple factors to the verifying authentication system or by using a factor to protect a secret that will be presented to the verifying system. For example, a hardware device that holds a cryptographic key might be activated by a password or the hardware device might use a biometric representation to activate the key. This type of device provides two-factor authentication, although the actual authentication protocol between the verifier and the claimant only proves possession of the key.

3.1.2 Quality of the authentication

The quality of the authentication determines the assurance of the identity of the user. The higher the quality, the higher the assurance level of the identity of the user will be. The quality of the authentication depends on many factors. The following aspects need to be considered carefully:

- The robustness of the authentication factors in terms of e.g. tamper resistance, randomness, etc.
- The strength of the binding between the authentication credential and the entity being authenticated.
- The trustworthiness and quality of the information obtained and used for identification.
- The repudiability of an assertion – has a party really been uniquely identified or are there other parties with the same identity as well?
- The liability of the identification – who is responsible for the identification and authentication.

The quality of the authentication is often quantified in terms of Levels of Assurance (LoA).

3.2 LoA

The strength of the entire authentication system is usually expressed in terms of levels of assurance (LoA). Authentication LoA defines the degree of confidence in identifying a user to whom the credential was issued, and the degree of confidence that the entity using the credential is indeed the entity that the credential was issued to. As more diverse resources are being incorporated into federated collaboration environments, service providers may require an assurance level in identifying an entity in an authentication process before an access control decision is made. For resources (data and/or services) with varying levels of sensitivity, the service providers may specify a minimum LoA and require that access is only granted if the LoA satisfies a certain minimum LoA.

LoA frameworks constitute an effort to unify and standardize the perception of assurance into a digital identity for the purpose of sharing digital identities between independent trust domains. Several frameworks for expressing LoAs exist, the most well-known are:

- NIST 800-63 defines four LoAs and describes concrete technical and procedural requirements that apply for each assurance level[11].

[11] William E. Burr, Donna F. Dodson, & W. Timothy Polk, Electronic Authentication Guideline Recommendations of the National Institute of Standards and Technology, Version 1.0.2, April 2006, see http://csrc.nist.gov/publications/nistpubs/800-63/SP800-63V1_0_2.pdf.

- InCommon Identity Assurance Framework[12]. InCommon states that their Bronze and Silver LoAs match with NISTs assurance Level 1 and 2. In order to cover also NISTs assurance Levels 3 and 4, further InCommon profiles are possible, but do not yet exist.

- STORK European eID interoperability framework[13]. The STORK Quality Assurance Framework is largely based on NIST 800-63. Marginal adjustments are made to accommodate European identity authentication and practices. STORK is meant to map national LoA frameworks onto each other in Europe in order to create interoperability.

Recently, ISO standardised four levels of identity assurance in the ISO/IEC 29115 standard (which coincides with ITU-T X.1254)].[14] These four levels of assurance are:

LoA 1 - Little or no confidence in the asserted identity's validity.

LoA 2 - Some confidence in the asserted identity's validity.

LoA 3 - High confidence in the asserted identity's validity.

LoA 4 - Very high confidence in the asserted identity's validity.

The levels are based on the degree of confidence needed in the process used to establish identity and in the proper use of the established credentials. Two factors are essential in the determination of the LoA:

- Registration: the strength of the identity proofing, registration, and the delivery of credentials, which bind an identity to a token. Aspects to take into account are:
    - Quality and robustness of the process (e.g., online, physical presence required, etc.).
    - Reliability of the issuing institution (e.g., government, bank, or social network provider).
- Authentication: the strength of the authentication mechanism to establish that a user is who he claims to be, which in turn depends upon
    - Strength of the authentication solution (passwords vs. smart card).
    - Assertion mechanisms used to communicate the results of a remote authentication to other parties.
    - Strength of cryptography used.

[12] InCommon Identity Assurance Assessment Framework, 9 May 2011, Version 1.1, see http://www.incommon.org/docs/assurance/IAAF_V1.1.pdf, and InCommon Identity Assurance Profiles Bronze and Silver, 9 May 2011, Version 1.1, see http://www.incommon.org/docs/assurance/IAP_V1.1.pdf.
[13] B. Hulsebosch, G. Lenzini, and H. Eertink, STORK Quality Authenticator Scheme, Deliverable D2.3, March 2009, see https://www.eid-stork.eu/index.php?option=com_processes&Itemid=&act=streamDocument&did=577.
[14] ISO/IEC 29115:2013 Entity authentication assurance framework, see http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=45138.

The overall authentication LoA components are shown in Figure 6.

Figure 6: Overall authentication LoA components.

The overall authentication assurance level is determined by the lowest assurance level achieved in any of the areas listed above. Note that the proofing phase LoA also contributes to the assurance of the registered attributes (e.g. name, e-mail address, phone number, student number, etc.). These attributes could be easily be exchanged in existing identity federations for higher education and research.

The paradigm of the LoA approach is that individuals are enrolled and undergo an identity proofing process in which their identity is bound to an authentication token. Thereafter, the individuals are (remotely) authenticated to systems and applications using the token in an authentication protocol. The authentication protocol allows an individual to demonstrate to a verifier that he has or knows the secret token, in a manner that protects the secret from compromise by different kinds of attacks. Higher authentication assurance levels require use of stronger tokens (harder to guess secrets) and better protection of the token from attacks.

In a nutshell:

- LoA 1 requires little or no confidence in the asserted identity. At this level, almost no identity proofing is required (at most an activation link to an email address) and a single authentication factor is sufficient (e.g. a password).

- LoA 2 requires confidence that the asserted identity is accurate. Identity proofing requires the presentation of identifying materials or information (e.g. copy of utility bill or passport). A single authentication factor is sufficient and may include additional requirements compared to LoA 1 factors (e.g. strong password).

- LoA 3 is appropriate for transactions that need high confidence in the accuracy of the asserted identity. At this level, identity proofing procedures require verification of identifying materials and information. LoA 3 requires multiple authentication factors.

- LoA 4 is used for transactions that require a very high level of confidence in the accuracy of the asserted identity. LoA 4 provides the highest practical assurance of authentication. Identity proofing requires physical presence accompanied by a valid photo ID document. Authentication is based on multiple factors. This level is similar to LoA 3 except that only "hard" cryptographic tokens are allowed, cryptographic module validation requirements are strengthened, and subsequent critical data transfers must be authenticated via a key that is bound to the authentication process. The token should be a hardware cryptographic module, which cannot readily be copied and ensures good, two-factor authentication.

Table 1 establishes the generic registration and proofing and technical authentication requirements specific to each LoA. Note that this table is non-exhaustive, for more information we refer to the following two sections on registration practices and authentication solutions.

Table 1: High-level identity registration and proofing and authentication requirements per LoA.

LoA	Registration and proofing	Authentication
1	None or on the basis of user supplied information such as an e-mail address.	Username and password
2	Remote. Use of physical address information to send credentials to, use of shared secrets	Username and strong passwords, knowledge based authentication, two factor authentication
3	Remote. Verification of information provided by applicant including e.g. ID number and account number through record checks either with the applicable agency or institution or through credit bureaus or similar databases, use of physical address information to send credentials to.	Two factor authentication, one-time-passwords via mobile phone or token, non-qualified certificates
4	Physical registration at desk with valid driver's license or passport.	Qualified certificates on trusted hardware.

A LoA can be assigned to all authentication solutions. For service providers this is very convenient, as they do not have to bother anymore about all available authentication solutions. They only have to specify the desired LoA for the service.

LoAs are increasingly becoming a "lingua franca" around the world. Prerequisite for consistent use of LoAs is the definition of a trust framework. A trust framework is a certification program that enables service providers accepting digital identity credentials to trust the identity, security, and privacy policies of the identity and authentication providers who issue the credential and vice versa. Furthermore, the trust framework defines the qualifications necessary to be an assessor for the trust framework; there must be some kind of governance body that is able to assess whether an identity or service provider is in compliance with the policies specified for a certain LoA. E.g., an authentication solution offered by an identity provider must be evaluated and rated according to the specified levels

by the governance body such that a service provider can rely on it without having to know the details of the authentication solution used.

Several attributes provided by the authentication provider will be validated during registration and identification. These attributes for instance include first and last name, e-mail address or virtual organisation membership. A LoA could be assigned to these attributes. In attribute-based access control scenario's, information about the reliability of these attributes could be beneficial for service providers to make their authorisation more reliable. However, in most cases authorisation will be based on other less validated attributes. It will make the registration process too complex if one decides to validate all possible attributes. If specific validated attributes are required, the service provider has to find an attribute service provider that can provide them; this is beyond the scope of this research. Moreover, having attributes with varying LoAs would complicate the use of the already complicated LoA concept even further and might hamper its adoption/acceptation. We therefore solely focus on authentication LoA.

3.3 Registration LoA solutions

The process by which a physical person is linked to their digital identity information and to their authentication credential is critical to deter registration fraud. If this process results in a weak link of the person to either the credential or the digital identity, there can be little or no assurance that the person using that credential to authenticate and access services and information is who they claim to be. If the linking is weak, even the most complete identification process and the strongest credential will not improve the assurance of identity.

The registration process is designed, to a greater or lesser degree depending on the assurance level, to ensure that the registration authority knows the true identity of the applicant. Two processes are important for determining a registration LoA:

1. Identity proofing - the process by which identity related information is validated so as to identify a person with a degree of uniqueness and certitude sufficient for the purposes for which that identity is to be used. This usually involves a minimum set of attributes that reach a high probability that its combination is unique in a given constituency (e.g., the first and last names, date of birth, and place of birth). Identity proofing provides confidence that the user performing an authentication is the legitimate user. Identity proofing will lead to the issuance of an authentication credential to someone. A poorly identity-proofed smart card provides less identity assurance than an adequately identity-proofed password.

2. Identity binding – the process of binding an authentication credential to an identity to which it will be issued. How binding is accomplished and the confidence in the association determines the LoA.

There are two general categories of threats to the registration processes:

- Impersonation of a claimed identity – An applicant claims an incorrect identity, supporting the claim with a specific set of attributes created over time or by presenting false credentials.

- Compromise or malfeasance of the infrastructure – Lack or poor implementation of security measures and policies undermine the reliability of the registration.

This section concentrates on addressing the impersonation threat. Infrastructure threats are addressed by normal computer security controls (e.g., separation of duties, record keeping, independent audits, etc.) and are outside the scope of this document.

Registration fraud can be deterred by making it more difficult to accomplish or by increasing the likelihood of detection. During the registration process methods should be employed to determine that a person with the claimed identity exists, and that the applicant is in fact the person who is entitled to that identity. As the level of assurance increases, the methods employed provide increasing resistance to casual, systematic and insider impersonation.

A number of solutions are known for identity proofing and binding of identity to an authentication solution. This section lists and evaluates them in more detail. The following aspects are taken into account in the evaluation:

- How costly is the solution?
- Is the solution user friendly and convenient?
- What is the implementation effort to roll out the solution; does it scale?
- What level of assurance does it offer? Are there any risks or threats associated to it?

3.3.1 Physical presence with identity credentials

The easiest and safest way to proof a user's identity and bind an authentication token is to require physical presence during registration and enrolment. Usually the user has to show their passport or driving license for identification purposes. During physical presence authentication credentials such as a mobile phone or token can easily and efficiently be bound to the verified identity.

3.3.2 Use of physical address and the postal system

Knowledge of the physical address of the user can be leveraged to improve the binding between the user and their authentication credentials. For instance, the Dutch government makes use of this for the issuance of its national DigiD credentials. The user is asked to enter their social security number ("BSN"), address details and e-mail address on the DigiD website. Subsequently, an activation code is sent to the home address of the user (after verification of the entered information in an authentic government owned registry). The user enters the activation code together with their username to activate the account. Also several banks use the physical address and the postal system to communicate activation or PIN codes.

A prerequisite for this approach is that the credential provider has knowledge of reliable address information of the user. Typically the government and banks are authorities that know their users'

home addresses (there may even be a legal requirement for this). There is a strong reliability on the security of snail mail. The recent fraud in the Netherlands with stolen DigiD activation codes from physical mail boxes illustrates that the postal system cannot be trusted completely. If an authentic registry for address information is not available, the registration authority could ask the user to send a proof of residence (e.g. a recent utility bill) but that is very inconvenient, may violate the user's privacy and requires manual processing.

For higher education, the identity provider (of the university) may have reliable address information.

The use of the physical address is relatively expensive and less user friendly as it breaks the user authentication session; the user may not come back the second time.

3.3.3 Use of e-mail or (mobile) phone

Often when users open an account they are asked to provide their e-mail address. This e-mail address is used to send an activation code or hyperlink that the user must enter or click upon to confirm that they indeed opened the account. Sometimes the mobile phone is used for such purposes as well (e.g. for sending an SMS with an activation code or for calling back the user to validate a callback code that was provided online).

Helpdesk employees use the phone frequently for authentication purposes. They for instance have to provide emergency authentication services for users that have lost their authentication token, forgotten their password or username and call them for help. In such cases helpdesk employees have a range of security assurances to identify the caller prior to providing new credentials. Via the phone they can ask personal questions, identify the voices of callers they know personally, verify caller IDs and so forth.

This form of identity proofing is weak if it does not include human intervention: it only proves that the user has access to a certain e-mail address or phone number. Moreover, it does not prove anything about the identity of the user owning the e-mail address or phone. The lack of solid identity proofing results in a weak binding to the user's authentication solution, i.e. the username and password that he created during the opening of the account. Consequently, the solution is prone to "man-in-the-middle" phishing attacks that try to breach the registration process. The use of human intervention such as is the case for the callback and helpdesk methods provide higher assurance levels but are expensive.

3.3.4 Use of bank account

An identity related to an existing account can be verified via a personal bank account that is not shared with other users. Bank accounts are usually very secure and can only be accessed via strong authentication solutions. Moreover, by law banks are required to know their customers and this implies that many banks require physical presence and proper identification prior to opening a bank account. The combination of strong authentication and identification means that the binding between a user, his personal bank account and his authentication token is very reliable. Other authentication providers can benefit from this.

Paypal for instance makes use of the user's bank account for linking its own user account to a bank account. A user has to use the bank's strong authentication solution to obtain Paypal's charges and link their account with the user's bank account. In a way, Paypal leverages the bank's strong authentication solution to increase its confidence in the user's Paypal identity. The Dutch finance company Alex also makes use of this concept.

There is, however, a slightly unexpected risk here. Several banks, such as the Dutch ING, only require a username/password to login and have an overview of all financial transactions[15]. Fraudsters can use this relatively low level of security and subsequent lack of proper identity verification to spoof transaction details (including Paypal's deposits) and open a Paypal account that is linked to someone else's bank account. This example shows that if one is not careful in taking into account the various LoA aspects things can easily go wrong.

Another risk associated to using a bank account for registration purposes is that the owner of the account may have mandated another user to make financial transactions on his/her behalf. In that case, the mandated user may present himself as the owner of the account during the registration phase of a new authentication solution.

3.3.5 Copy of official identity credentials

As proof of identity a certified copy of a passport or national identity card is sometimes accepted. Such a copy can either be certified or uncertified. The latter offers almost no proof in the sense that it really belongs to the user sending it. For certified copies, the certifier must be able to achieve the following: "having seen the individual and identification document at the same time, I certify this is a true copy and the photograph is a reasonable likeness." Examples of certification authorities are a notary, a consular or embassy official from your consulate or embassy, a police officer, or a qualified accountant or auditor.

Non-certified copies provider a lower level of assurance as they can be taken from stolen or lost identity credentials. A digitized copy, even if it is certified, that falls in the hands of a user with malicious intentions can easily be reused for all kinds of purposes. Binding between a copy and a user remains a weak registration solution. Confidence should be established that the ID document it is still under the control of the entity that it relates to (e.g. it has not been stolen/cloned or being used by a bot). So checks against registers of stolen/revoked official identity credentials are required.

3.3.6 Use of official electronic identity credentials

More and more official identity credentials in the physical world such as a passport or a driving license are equipped with a chip. The chip may contain identifying information such as a social security number, name, address, gender, and date of birth, or even a biometric template of, e.g., a fingerprint. These attributes are signed by the issuing organisation, which in most cases is a government.

[15] See http://www.crimesite.nl/crimesite/159-headlines/20876-bankfraude-fraudeurs-omzeilen-tan-beveiliging-ing.html.

The binding process of these electronic identity credentials to the user is very solid as they require physical presence. The recent electronic enhancement makes them suitable for online authentication processes[16].

When using official ID documents like a passport or a driving license, the authenticity of the document should be checked. This requires trained personnel or the use of online verification services like IDchecker[17]. Furthermore, confidence should be established that the ID document is still under the control of the entity that it relates to (e.g. it has not been stolen/cloned or is being used by a bot). So checks against registers of stolen/revoked official identity credentials are required.

The use of electronic ID credentials (ID card or ePassport) is rather complex but promising for the near future. It requires card readers and a solid enrolment process for binding these credentials with existing accounts. The emergence of NFC enabled mobile phone for payments could provide a boost for this authentication method.

3.3.7 Use of video

In case the user is somehow not able to register in person, video conferencing tools such as Skype could be used. In this case the user identifies him/herself via the videoconference and shows his/her passport or other valid photo-ID to the registrar.

The use of video conferencing tools for identification, however, has several drawbacks: it introduces scheduling overhead and it makes it harder to detect a forged ID. If the user already has an authentication credential, he can be asked during the video-conferencing session, to perform an authentication. This immediately ensures binding between the user and his credential (if not the actual identity itself). If the user does not have an authentication credential he can be asked to provide address information to send the credential to.

3.3.8 Identity verification services

Identity verification services provide a verification-chain framework to identity and service providers, for example in the space of online dating services, while protecting sensitive information. These services typically work as follows:

- Users sign up for a new account on a dating site and are prompted to click through to the site of an identity verifier.

- Users create profiles with details such as their name, age, address, and occupation, etc.

[16] M. Oostdijk, D-J. van Dijk, & M. Wegdam, User–Centric Identity Using ePassports, in Proceedings of SecureComm 2009 Conference, September 14-17, 2009, Athens, Greece; published in LNICST 19, Springer, pp. 296-310, 2009, see http://www.springerlink.com/content/v522837228n3611r/fulltext.pdf.

[17] IDchecker, see www.idchecker.com.

- Verification services electronically check data in public-record databases to verify assertions and prompt users to answer other challenges based on public records.

- If users pass these challenges, they are granted a verified status.

Identity verification services provide value by acting as a mediator in an identity transaction. They create trust by certifying that the user is indeed the person he claims to be, without disclosing sensitive information about the user to the other party.

There are a number of players in this space. Examples are edentiti[18], VeriSign[19], Trufina[20], and Idology[21]. A couple of other services in the space are RapLeaf[22] and iKarma[23]. These services rely on transaction history (RapLeaf) or explicit recommendation and testimonials (iKarma) to evaluate the reliability and trustworthiness of an individual. All of these services provide tight integration at the point of transaction. It is not unrealistic for social network providers such as Facebook, Google+, and LinkedIn to become identity verifiers in the near future as well.

While these companies provide a valuable service, their penetration outside the online dating space seems to be somewhat limited, also depending on country. Some potential issues may be the cause for the low uptake.

One of the main issues is that identity validation services rely on public records. These services typically ask users to provide some personal information, based on which they access public records available for that person. These services then challenge the users to answer questions, based on the information in these public records. If the user answers these questions correctly (i.e., the answers match the information available in public records), the user is considered verified. The availability of suitable public records varies per country, depending amongst others on privacy regulations.

All of the public records are available online for everybody to search and see. E.g., via dedicated search engines such as Intelius or Google access to numerous public records can easily be obtained. If somebody wanted to pretend to be another person, he would certainly check all these public records to provide enough information to answer the challenge questions correctly.

3.3.9 Account linking or federation

Account linking or federation could be used to leverage the user's federated authentication outcome for the issuance of other authentication solutions.

Account federation or linking occurs when a user chooses to unite distinct service accounts and identity provider accounts. Users retain the individual account information with each provider while, simultaneously, establishing a link that allows the exchange of authentication information between

[18] www.edentiti.com
[19] www.verisign.com
[20] www.trufina.com
[21] www.idology.com
[22] www.rapleaf.com
[23] www.ikarma.com

them. For instance, the student's account at the university's identity provider is linked with that of another authentication provider such as Google. In an identity federation, the authentication outcome is reused over a number of (federated) services. Authentication service providers in such a federation can rely on the identity assertions of the identity provider. A good example is the TCS eScience Portal[24] that issues certificates to users based on federated login.

Account linking uses protocols such as SAML, OAuth, or OpenID Connect to create a persistent association between these distinct user accounts. The account link, or name identifier, may be either a unique attribute, such as an email address, or a pseudonym generated by the identity provider to uniquely identify individual users. Pseudonyms can be used if privacy is a concern; they cannot easily be traced back to a user's identity at the partner site.

Optionally, during account linking, additional attributes may be sent with the name identifier. The authentication provider can use these attributes to challenge the user with knowledge-based questions or send activation codes to e-mail addresses or mobile phones via text message. Moreover, by comparing the attributes that are associated to each account (e.g. name, surname, e-mail address, etc.) an enhancement of the assurance of the identity of the user can be achieved, i.e. the matching attributes may proof that both accounts indeed belong to the same user. Care must be taken not to compromise privacy. Similar attributes associated to multiple accounts can be validated against each other thus providing more (or less) confidence in the attributes.

3.3.10 Web of trust

Web of trust is a concept (used in e.g. Pretty Good Privacy) to establish the authenticity of the binding between an authentication solution and its owner via third party user attests. For instance, if person A claims that user B is using a particular authentication solution, it could provide extra confidence for the service provider to allow access to resources with a higher LoA. Person C could also claim to know B and his authentication mechanism thereby even further increasing the trust in the identity of B. In essence, this is a kind of "crowdsourcing of trust". The social or research context of the user could be used to enhance the registration part of the overall LoA.

Particularly in the context of research groups or virtual organizations in which users know each other, such a web of trust based LoA enhancement could be executed in an efficient manner. Moreover this approach also makes it easier to use social identities provided by e.g. Facebook and Google in higher education and research environments. The registration LoA part of these popular social identity providers is relatively weak (LoA 1) despite the fact that an increasing number of them are using two-factor authentication (LoA 2). Web of trust based enhanced LoA could help increasing the registration LoA part of these providers and thus could help in increasing the overall LoA.

The web of trust approach also has its weaknesses. ENISA has summarized the possible threats such as the whitewashing attack, sybil attack, impersonation and reputation theft, bootstrap issues and 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 behaviour, attacking of the underlying infrastructure and the exploitation of features of metrics used by the system to calculate

[24] See https://tcs-escience-portal.terena.org/.

the identity assurance[25]. These threats should be taken into account to evaluate usefulness of the web of trust based solution to enhance the LoA in the context of identity federations.

3.3.11 Evaluation

The identified registration LoA approaches are evaluated in Table 2 below.

Table 2: Evaluation overview of registration LoA approaches.

Registration process	Cost	User friendliness	Enrollment effort	LoA	Explanation
Physical presence with identity credentials	High	Low	High	High	Labor intensive, many procedures; effortful for user, requires trained personnel and/or verification equipment to detect false ID documents; easy reliable binding of identities, though 14% false acceptance rate[26]
Use of physical address and the postal system	Moderate to high	Moderate	Moderate	Low / Moderate	Costly postal process; relatively long time; constructing a process for sending codes by mail; risk of interception or loss
Use of e-mail or (mobile) phone	Low to moderate	High	Low	Low	E-mail is relatively low-cost, text-message may be expensive for large user-groups, callback procedures are time consuming and costly; e-mail and text-message are familiar to user; easily implemented; no identity validation and sensitive to digital fraud
Use of bank account	Low	Moderate	Moderate	Moderate	Small deposit per check; requires some effort of user, some delay; AP has to

[25] Elisabetta Carrara and Giles Hogben, Reputation-based Systems: a security analysis, ENISA position paper, October 2007.
[26] David White, Richard I. Kemp, Rob Jenkins, Michael Matheson, A. Mike Burton, Passport Officers' Errors in Face Matching, PLOS online, August 18, 2014, see http://www.plosone.org/article/info%3Adoi%2F10.1371%2Fjournal.pone.0103510.

					develop payment procedures; bank accounts are well secured; drawback is that bank accounts are sometimes shared by multiple users.
Uncertified copy of official identity credentials	Low	High	Low	Low	Copying is cheap and easy; requires little enrolment effort; no proof that the credentials belong to user
Certified copy of official identity credentials	High	Low	High	Moderate	Certification is costly; inconvenient and takes manual processing; authorities must be involved; requires physical presence; no proof that the credentials belong to user
Use of official electronic identity credentials	High	Moderate	Moderate	Moderate	Requires communication reader (USB reader or NFC); user must have card with him; secure enrolment is needed, could be self-service; susceptible to theft and relay threats
Use of video	Moderate	Moderate	Moderate	Moderate	Conferencing tools are common, but requires scheduling; convenient but intrusive; document falsification is difficult to determine by video
Identity verification services	Moderate	High	Moderate	Low to moderate	Service providers charge; transparent to user; links with provider have to be set up; uses multiple public records
Account linking or federation	Low	Moderate	Moderate	Moderate	Easy to implement; authentication during account linking, privacy risks; additional account linking functionality; only lower or equal assurance

Web of trust	Moderate	High	Moderate	Moderate	Mechanism must be implemented; uses common social relationships for identification; reusing existing webs-of-trust; assurance depends on number of trustworthy nodes that assured an identity, several threats determined by ENISA

3.3.12 Discussion

In general, registration LoA approaches described in this chapter help to increase the overall authentication strength by using an additional communication channel. Using more separate channels leads to stronger authentication levels.

However, higher authentication strength does not necessarily represent the best solution, as security often requires trade-offs in user convenience and/or costs. Particularly the use of regular mail comes at the cost of user convenience as it breaks the user's authentication session. The email and mobile channel provide a more real-time user experience, i.e., activation codes can be entered on the fly, but are less secure.

A generic feature of the registration LoA approaches that are based on remote registration is that almost none of them adequately verify the true identity of the user. Instead they merely validate that, e.g., the banking information provided is capable of routing an electronic deposit to an account or that the mobile phone or e-mail account is capable of receiving challenges. So while validating multiple 'paths' are leading to the same user, it does not necessarily establish the identity of that user. An exception is the use of videoconferencing tools. In this case the user identifies him/herself via the videoconference and shows his/her passport or other valid photo-ID to the registrar. The use of video conferencing tools for identification, however, has several drawbacks: it introduces scheduling overhead and it makes it harder to detect a forged ID.

Other remote registration approaches rely on the availability of trusted sources to cross-reference and validate the provided assertions such as name, home address, age, social security number, and photo. Examples of such sources are the institution's HR system or the government/municipal administration(s). Consultation of the latter source is often restricted by legislation and therefore not available for enhancing authentication; the HR system on the other hand could be used as an alternative source.

An interesting alternative is to combine multiple methods to create a stronger binding with an identity. Possible combinations are:

- Mobile phone and e-mail. The use of the mobile phone is convenient but requires the verification of the phone number somehow (assuming the user's institution doesn't have knowledge of it). This can be done via another channel like e-mail. As the e-mail address is known by the user's institution, it can send a notification to the user saying that a certain mobile phone (number) has been linked to his account.

- Voice recognition via the mobile phone and e-mail. The mobile phone as a user friendly tool to verify the user's identity via voice recognition. The e-mail channel should be used as a second channel for the enrolment of the voice recognition, i.e., for registering the mobile phone number and for recording the voice template.

- Bank account and e-mail. A bank account is highly secured but its use does not prove that it is owned by the user identity that needs to be authenticated. Here an e-mail could be sent as well to verify that the bank account is indeed linked to the user owning the e-mail address.

- Instead of an e-mail address, the physical address can also be used to send activation or other credentials to.

Account linking provides an interesting method to enhance the assurance of the user's identity. The identity attributes associated to the account can be compared. Matching attributes not only provide assurance that both accounts belong to the same user but also help to uniquely identify that user. For instance, the set of attributes that consists of name, surname, date of birth, and place of birth almost uniquely identify a user.

The physical registration approach best verifies the user's identity but this is expensive, time consuming and not user friendly.

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 identification task. User in the web of trust may use physical presence, phone or email practices for this purpose. Somehow, the attestations from the web of trust need to be related to the claimant's digital identity. This needs to be catered for by some kind of federated attestation service that enhances the assurance in the claimant's federated identity with attestations from the web of trust. How this can be done, will be described in deliverable D2.1 of WoT4LoA. Account linking can help for determining the web of trust and the user therein that can vouch for another user's identity.

The table below provides an overview of the LoAs for the different identity proofing and registration methods.

Table 3: LoA overview and corresponding methods for identity registration and proofing.

LoA	Objective	Controls	Method
1	Identity is unique within a context	Self-asserted	Use of email or mobile phone
2	Identity is unique within a context and the entity to which the identity pertains exists objectively	Proof of identity through use of identity information from an authoritative source	Copy of ID-credential Copy of utility bill Account linking / attribute matching

3	Identity is unique within a context and the entity to which the identity pertains exists objectively, identity is verified and is used in other contexts	Proof of identity through use of identity information from an authoritative source + verification	Use of physical address and snail mail Use of video and proof of authentication Use of bank account Use of electronic ID-credentials Identity verification services Web of trust + account linking
4	Identity is unique within a context and the entity to which the identity pertains exists objectively, identity is verified and is used in other contexts	Proof of identity through use of identity information from one or more authoritative sources + verification + entity witnessed in person	Physical presence with ID credential (e.g. passport, ID-card, driving license) and issuance of authentication credential in person.

Combinations of multiple remote methods will increase the LoA till at most level 3. For LoA 4 physical presence is required.

3.4 Authentication LoA solutions

Many inventories of authentication solutions have been published in recent years:

- Kuppinger Cole - Market Overview Strong Authentication –2010;
- Gartner – Market Overview Authentication –2008;
- Gartner – Market Scope for Enterprise Broad-Portfolio Authentication Vendors – 2009;
- Novay –Whitepaper: Authentication solutions state of the art – 2010[27];
- Academic overviews (UvA MSc. thesis C. de Jong – 2008, UvA MSc. thesis D. van den Ende and T. Hendrickx – 2009, RU MSc. thesis Schouwenaar – 2010);
- GLUU overview of authentication solutions – see http://t.co/gl1WLMrfu9.
- J. Bonneau, C. Herley, P.C. van Oorschot, and F. Stajano, The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes, Proceedings of the 2012 IEEE Symposium on Security and Privacy, 2012.

These inventories usually consist of long lists of authentication solutions that are typically classified in four main categories: knowledge-based (e.g. username/password and challenge question), possession-based (e.g. OTP token, TAN list, PKI certificate or EMV CAP reader), out-of-band (e.g. SMS OTP and caller line identification), and biometric solutions (e.g. face and voice recognition).

There are other solutions that do not fit in the four categories. These solutions typically involve new, non-cryptographic authentication techniques like risk based, behavioural biometrics, and geo-location.

[27] See https://maartenwegdam.files.wordpress.com/2011/06/cidsafe-authn-scan-1-0-0.pdf (in Dutch).

We note that so far these solutions have not been integrated in recognized frameworks like the STORK QAA levels and NIST.

In order to evaluate the suitability of the numerous authentication solutions for the education and research community and their identity federations several aspects need to be taken into account:

- How costly is the solution?
- Is the solution user friendly and convenient?
- What is the implementation effort to roll out the solution; does it scale?
- What level of security does it offer? Are there any risks or threats associated to it?
- Is the technology sufficiently mature?
- Is it sufficiently mobile?

Other, more strategic approaches are possible to tackle the authentication challenges. For instance outsourcing two-factor authentication to a trusted third party can help address the challenges (i.e. costs and organisational management overhead) research and education face regarding authentication in the following ways:

- Higher education and research institutions do not need to invest in building and operating a two-factor authentication infrastructure.
- The trusted operator is responsible for reliability and scalability of the service.
- Quick time-to-market.

One solution is two-factor authentication-as-a-service (AaaS). AaaS is an internet-based service that offers an on-demand verification of the user's second factor to service providers. AaaS could be very interesting as a step-up authentication solution for the higher education and research community, and could be offered by federation operators as a federated service. It is up to the federation operator to select and approve the authentication solutions that will be provided via such an AaaS. A challenge here is to make sure that the binding between the user's identity and the selected AaaS authentication token is guaranteed.

When working with authentication tokens the life-cycle of the tokens should be taken into account. Authentication tokens follow a life-cycle which can be described on three levels as depicted in Figure 7. First, on the largest scale, issuers of authentication tokens will, from time to time, come to the conclusion that new authentication technologies need to be introduced. Reasons could be economic (new technologies may be cheaper) or caused by new threats. In general any number of reasons could drive an issuer to switch technologies, and the business model is usually far from simple.

Second, on a medium scale, authentication tokens that are issued to individual users have a limited lifetime in which they can be used. Many tokens have a battery, which naturally limits the token's lifespan. Users will also lose their token or damage a token beyond repair and request a new one.

Security may be a reason to artificially limit the lifespan of authentication technology as well: it is common practice to limit the validity of certificates to a fixed period.

Third, on the finest scale, tokens are used in authentication sessions, starting when the user signs in, and ending when they sign out. Before allowing a user to authenticate with the token, the token may itself want to make sure its rightful user is handling it using a process called card holder verification (CHV) which usually involves verifying that the user knows a PIN code. As some tokens combine pure authentication with other identity related procedures during sessions (such as signing messages or authorizing transactions) this life-cycle involves more than just authentication.

For the evaluation of solutions the life-cycle of authentication tokens and the complexity of it should be taken into account as well.

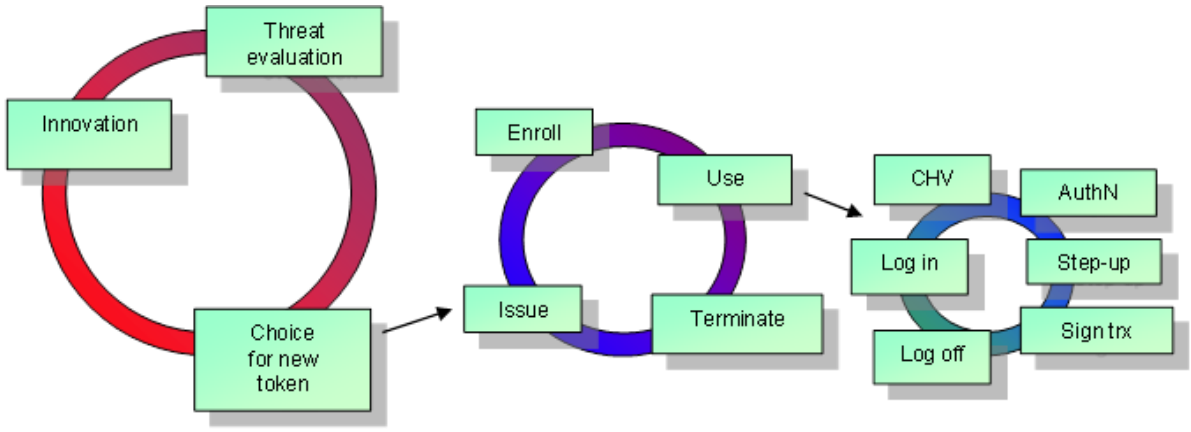

Figure 7: Different life-cycle events of an authentication solution. From left to right: “new technologies cycle”, “enrolment cycle”, “session cycle”.

The quality of the life-cycle of authentication solutions is taken into account in the LoA paradigm. Service providers that need to authenticate the user do not need to worry anymore about various authentication solutions, they just have to determine the required LoA for their services and ask the user authenticate with that LoA.

3.4.1 Classification

Technical authentication solutions (“authentication tokens”) can be classified in many different ways. Typically a taxonomy of tokens is based on characteristics of the underlying technology (see for example the introduction of this paper[28]). A possible (naive, start of a) taxonomy of solutions in given in Figure 8 below.

[28] http://static.usenix.org/event/lisa11/tech/full_papers/Rijswijk.pdf.

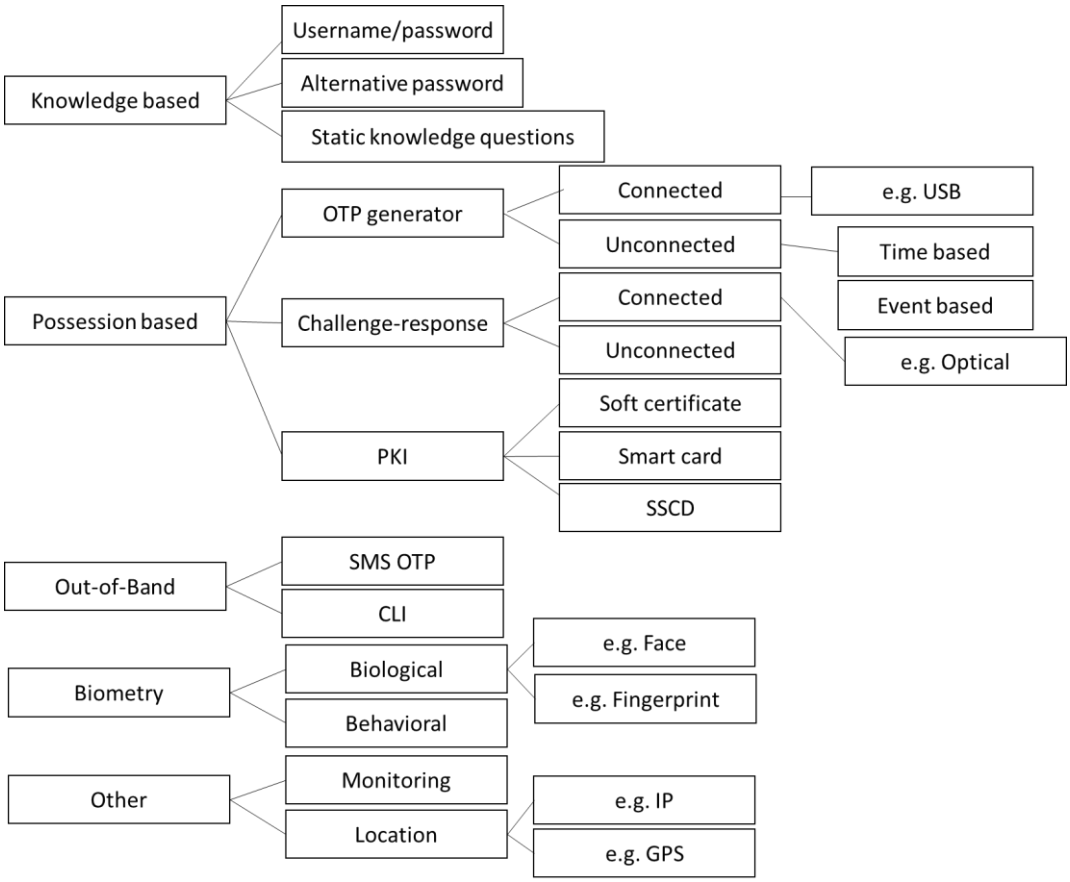

Figure 8: Taxonomy of authentication solutions.

The taxonomy is not complete: there will be new technologies not mentioned here, there will be other ways to connect devices to the web session through which the user is interacting with the service provider, and there are, for example, many more biometric factors in addition to face and fingerprint. The taxonomy also compares apples and oranges, some technologies can be categorized as out-of-band and challenge-response at the same time, risk based technologies do not form a complete authentication solution on their own, but are typically added to take away certain risks associated with weak authentication factors.

Nevertheless, we will have to use some sort of taxonomy like the one above, in order to be able to reason about the different types of tokens, their security, usability, and cost, their benefits and drawbacks, and the possibility to apply these tokens in a web-of-trust setting.

3.4.2 Overview

3.4.2.1 *Knowledge based solutions*

Examples of knowledge-based solutions are:

- Username / password. The user has a public persistent identifier (the username) and corresponding secret password that they memorize. Passwords can usually be chosen by the user, albeit subject to some password policy (i.e. relatively strong passwords can be enforced). The password policy may also contain criteria that force users to regularly pick a new password. During an authentication session the verifying party asks the user to send their username and password.

- Static knowledge questions. In case of lost passwords, or in case of step-up authentication, the verifying party may ask a user a secret question that only the legitimate user is likely to know the answer to. Typical (and therefore dangerous) examples of topics that the user could be questioned about are the user's mother's maiden name, the user's place of birth, and the name of the user's favorite pet. In most cases the user is required to first communicate the answers to the verifying party during enrollment; in other cases the verifying party has authentic information.

- Password replacements. Replacing simple passwords with other types of objects (patterns, images) that can be memorized by users offers some advantages such as better entropy, a primitive form of challenge response (where the user does the "computation" based on a memorized secret) which makes dictionary attacks harder, replay impossible and forcing installed malware to do complex screenscraping to make sense of the challenge or response.

- Gesture based passwords. A particular password replacement could be (hand) gestures. Many smart phones contain sensors that can be used to measure gestures. These include sensors to measure finger movement (tap, swipe, etc.) but also motion sensors that can measure relative acceleration of the device. User studies into gesture-based passwords suggest that not all users find them easy to use (and that his corresponds to the achieved accuracy), see e.g. this early research paper [29].

- Virtual keyboard. By displaying a virtual keyboard on screen on which the user can click on virtual keys (rather than type passwords via the keys of a hardware keyboard) raises the barrier for eavesdropping attackers listening for passwords. This works on desktop/laptop as it would be an alternative path for user input to which malware may not have access, but does not work on tablets, where a soft keyboard is standard and is likely something that malware will have access to.

3.4.2.2 *Solutions based on possession*

Examples of authentication solutions that are based on possession are:

- TAN list. A paper list containing one time passcodes that the user sends to the verifying party to authenticate. A TAN (Transaction Authentication Number) is usually needed to perform a transaction, i.e. not for initial session authentication but for authorization. The passcodes can be numbered (indexed TAN or iTAN) so that the verifying party can ask the user to enter a specific TAN (so that the index number can be tied to the requested transaction). This is susceptible to a MitB if used for authorization of transactions as no transaction information is given to the user.

[29] http://www.ruf.rice.edu/~mobile/publications/liu09mobilehci.pdf

- Matrix card. This printed card contains a grid (e.g., horizontal A - J, vertical 1 - 5). Given a challenge sent by the verifying party (e.g., A5, B7, C9, D8) the user needs to type the corresponding numbers from the grid in response. Each card contains a unique grid, which is tied to a serial number (so that the card can be blacklisted when lost or stolen).

- OTP token. An OTP (One Time Password) token generates fresh (seemingly random) pass codes. During authentication the generates a new pass code (either by having the user press a button, or after some time elapses), which is then entered into the verifying party's web site. The token uses a one-way algorithm *f* to produce codes that are unique for that particular token and are 'fresh' so as to avoid replay attacks.

- OTP token with USB keyboard interface. A USB keyboard OTP token is a variation on the OTP token described above. It saves the user the trouble of having to type in the generated pass code by simulating a connected USB keyboard.

- Challenge response token. An offline challenge response hardware token has a display and a numerical keypad. During an authentication session the verifying party shows a challenge code in the browser, which the user types into the token. The token then shows a response on its display, which the user enters in a form shown in the browser.

- Soft-token or application on PC. Tokens (OTP, Challenge Response, PKI or OOB) do not necessarily have to be implemented in hardware. Software-only tokens take the form of applications running on the user's desktop PC and, while not as secure as their hardware counterparts, can form an extra barrier for an attacker. Credentials used by the soft token can have much more entropy than a password known by the user. A drawback is that the soft token needs to be installed first on the desktop PC before the user can consume the online service, making use on other PCs harder.

- Soft-token or application on smart phone. A token could take the form of an app(lication) on a mobile device. Such an app could implement an OTP token, or a challenge response token, or a connected token with the ability to show transaction details and signature creation. Advantages of this approach are, for the latter case, channel separation (at least, as long as the online service is not consumed on the mobile device) and a relatively rich user interface to show transaction details when compared to dedicated hardware tokens.

- Camera-phone with 2D barcode scanning app. Almost all modern cell phones contain a built-in camera. The verifying party could have the user's browser render a 2D barcode to be read by an app on such a camera phone. The 2D bar code can, in principle, contain transaction details (i.e. QR-TAN). If the information is encrypted using a key shared between the app on the phone and the verifying party's back-end systems, then installed malware on the PC cannot get at it. Camera phones have no problem scanning 2D bar codes (such as QR codes) and these codes can contain in the order of thousands of characters of information.

- SIM Application Toolkit application. A SIM Application Toolkit (SAT) applet is software that is installed on the SIM (at the discretion of the mobile operator who owns and controls the SIM). A SAT applet can receive specially crafted SMS messages over the GSM/UMTS network without the user ever noticing. A SAT applet can interact with the user through a menu style interface. The

user interaction initiated by a SAT applet takes precedence over user interactions by applications running on the phone.

- EMV/CAP reader. A variation on the challenge response token explained above is an EMV/CAP reader. This is a solution where the token itself is a non-personalized smart card reader which uses the EMV chip on a bank's credit or debit card as the source for the secret key. In fact the EMV/CAP protocol mimics the standard EMV protocol (used at a point-of-sale or an ATM for authorizing a payment transaction) in signing a transaction (typically a debit transaction in a web based eBanking transaction).

- Optical EMV/CAP reader. A variation on the EMV/CAP reader described above uses optical sensors which can scan information on the display of the desktop PC on which the user consumes the online banking service. An animated bar code may contain the challenge as well as transaction information, which can be displayed to the user using the display of the token. The single direction communication channel (from back-end to EMV chip) may be secured using keys contained in the EMV chip.

- Connected EMV/CAP reader. A variation on the EMV/CAP reader described above uses a USB connection to the desktop PC on which the user consumes the online banking service. The connection can in principle be used to send transaction details over a secure channel (i.e. it is a form of out-of-band communication) directly to the EMV chip for signing. The display of the reader can show these transaction details to the user and can ask for (informed) consent.

- PKI USB token. A PKI token is a (USB) connected token containing a private key with corresponding public key certificate. The certificate is signed by a trusted CA. The user's desktop PC can read the certificate, but cannot get to the private key. The token can, however, use the private key to sign messages (documents, e-mails, challenges, transaction details) sent to it.

- USB token with secure channel to back-end. The USB connection can be used by a USB token to setup a secure channel to the back-end system. This is similar to the connected EMV/CAP reader described in section ? but can also be applied in non-EMV/CAP setting. The secure channel can be used in a protocol to establish authenticity of the connected token and can also be used, if the token has a dedicated display, to securely send transaction details to the user.

- PKI Smart card. PKI smart cards are PKI tokens with a different form factor and a different type of connection to the user's desktop PC (instead of USB they use either ISO7816 for contact-, or ISO14443 for contactless smartcards). To establish this connection the user needs a general purpose smart card reader, many of which can nowadays be installed driverless[30].

- PKI soft certificate. A soft certificate (a form of soft token) is a PKI token, which stores the private key (as well as the public key certificate) on the user's desktop PC. This causes problems if the user decides to start consuming the online service from another PC (i.e. the credentials will have to be copied to this new PC). Worse yet, installed malware with sufficient privileges may be able to read the private key.

[30] See the CCID device class specification, available from http://www.usb.org/.

- Visual cryptography. In visual cryptography the user has a token containing a transparent area which, when superimposed on an image rendered in the browser (the challenge), yields an OTP for the user to enter. The combination yields a secret, which is only interpretable by the user, i.e. to installed malware the image that is rendered in the browser looks completely random.

3.4.2.3 *Solutions based on out-of-band communications*

Out-of-band solutions use a separate channel to the user (solutions that use the same communication channel for authentication and consumption of the online banking service are called *in-band*). Out-of-band can be either outbound (the IdP initiates the communication) or inbound (the user initiates the communication). Examples are:

- Caller Line Identification. Caller Line Identification (CLI) can be used to identify legitimate users when using the telephony channel (i.e. when using the IVR system for authentication). Although the security of CLI is disputed[31], it forms a factor that prevents attackers to perform fraud on a larger scale.

- SMS OTP (mTAN). SMS OTP or mobile TAN (mTAN) uses short message system (SMS) text messages sent to the user's cell phone to create a second channel (i.e. it is a form of out-of-band authentication) to the user. It basically replaces the paper TAN list by sending a fresh TAN to the user's phone every time they attempt to authenticate. This has the added advantage that the message can also contain transaction details. The additional channel makes it much harder for an attacker to get access to a user's credentials. Attacks on the cryptography in GSM (Nohl et al.[32]) indicate that the security (both confidentiality and integrity) of the SMS channel, at least for the GSM case, should be considered broken. However the strength of this solution stems from the separation of channels and the difficulty for an attacker to control both channels simultaneously.

3.4.2.4 *Solutions based on biometrics*

Biometric solutions fall into two different categories. First, traditionally, biological biometrics use physical characteristics of the user such as fingerprints or iris patterns. Second, behavioral biometrics uses typical behavior of users such as typing speed, GUI interaction, and other context factors. The technology, in general, has many drawbacks. Most notable are: performance (false acceptance/rejection rate), privacy sensitive templates, costly enrollment, non-revocable/renewable, oversight required. Whether and how these drawbacks can be remedied is subject to active research. Examples are:

- Face recognition. Face biometrics use a camera to recognize features of a user's face. If this is based on software running locally on the user's device (PC, smartphone) then the channel between the camera and the verifying party's back-end systems is untrusted. Unless some kind of tamper-resistant camera is used, this kind of biometry is not suitable for remote consumer authentication as an attacker can simply replay a previously recorded sequence of images.

[31] http://www.schneier.com/blog/archives/2006/03/caller_id_spoof.html.
[32] Nohl, K. and Paget, C., GSM – SRSLY?, Chaos Computer Club Congress, December 2009.

Possible countermeasures are to ask the user to say something or to display a moving object on the screen and ask the user to follow the object with his eyes.

- Fingerprint recognition. Fingerprint recognition uses a capacitive or photographic sensor to record the user's fingerprints so that they can be compared to a previously enrolled template. Of the physical biometric factors, fingerprint recognition hardware is probably the cheapest, and can be integrated in hardware tokens to replace PIN as a mechanism to verify ownership. The token itself can then do the matching.

- Vein structure recognition. A relatively new physical biometric factor is vein structure. This uses a photographic sensor to record the 3D infrared vein structure pattern of the user's hand. Vein pattern scanners with USB connection are already commercially available. The security characteristics are comparable to those of face-, and fingerprint recognition.

- Voice recognition. In telephony based authentication (i.e. when users call an Interactive Voice Response (IVR) system) voice recognition may be used as a factor in context based authorization (i.e. determining with some degree of certainty if this is the genuine user, and acting upon this). Having the user read a challenge text makes replay attacks impossible. Obviously, this is not resistant to a man-in-the-middle attack on both PC and telephony channel.

- Keystroke dynamics. Keystroke dynamics is a form of behavioral biometrics. By measuring not *what* is typed, but *how* it is typed (e.g. the time in between key strokes) users can be characterized in a relatively unique manner. Keystroke biometrics may require the installation of client side software, although new web technologies such as HTML5/JS/AJAX may make it possible to record typing times in between typed characters sent to the web server in relative real time. Keystroke dynamics could be used as a factor in risk based authentication described below, e.g. triggering extra authentication steps if the timing pattern appears to indicate that this might not be the genuine user.

- GUI interactivity. GUI behavior, other than key-stroke dynamics, can also be taken into account. For example mouse movements and timing of events characterize the user. Similarly to key stroke dynamics this method could use client side installed software or new web technologies based such as HTML5/JS/AJAX. Just like keystroke dynamics, GUI interactivity can be used as input for risk based authentication described below.

- Hand-written signature recognition. A token could incorporate the possibility to ask a user to show that they can produce a hand-written signature corresponding to the genuine user's signature template stored in the back-end of the verifying party. Some smart phones readily support recording of hand-written signatures[33], proving that the technology has become relatively stable and robust.

[33] See, for instance, Square's application for accepting credit card payments: http://www.youtube.com/watch?v=QSzsFAJAKHI.

3.4.2.5 *Other solution directions*

This section contains solutions (or solution directions),which do not fit in the other categories. These solutions complement the other authentication technologies described (e.g. by detecting attacks, or reducing their impact), but cannot be considered complete independent authentication technologies. Examples are:

- Risk-based authentication. Risk-based authentication and authorization means that contextual factors, such as how the user behaves on a web site (last login, time on site, … but also behavioral biometrics), and transactional factors (the amount, the recipient of the transaction, etc.) are dynamically taken into account when determining which authentication or transaction authorization solutions are acceptable to complete a transaction. A user can, for example, be asked to provide additional step-up authentication before he can continue to consume the online service when contextual factors change (such as the same user logs in from a different IP address, the user engages in a transaction that is out of the ordinary given the user's profile). It is also possible to detect, based on general behaviour of a client computer (identified, e.g., by an IP address), with some degree of certainty that that computer is infected by malware that is actively attempting to use the user's credentials to log in to the service. This can be used as input for risk-based authentication (as well as to warn the user).

- Device identification / characterization and trusted platform module. If the client platform can be identified in a trustworthy way, then the verifying party can more easily detect fraudsters, and use this as input for risk-based authentication. A solution can be to install software on the user's desktop PC, which is configured with some unique persistent identifier. Mobile OSes provide a persistent identifier that apps can use for this. On the browser level cookies and browser characteristics[34] can be used to recognize a returning user.

- IP address based identification / characterization. On the network level IP addresses have been used extensively to recognize returning users, and determine geographical location of the user.

- Physical device authentication. Some tokens and/or devices have characteristics which make cloning / copying of the device itself impossible. This can be achieved using cryptography (think, for example, of active authentication as used in ePassports described above) but also through physical characteristics. An example is magnetic fingerprinting, which is used to distinguish between the original magnetic stripe of a banking card and skimmed clones of that card. Even though the skimmed clones contain identical information, the physical characteristics of the magnetic stripe itself cannot be copied and provide a unique fingerprint of the original card. This solution could, in principle, be used as input for risk-based authentications or other reactive controls.

- Distance bounding protocols. Another instance of physical characteristics that can be measured is the application of so-called distance bounding protocols[35]. These are thought to be usable to

[34] E.g., see http://panopticlick.eff.org/.

[35] See http://www.cl.cam.ac.uk/~sd410/papers/sc_relay.pdf.

detect so-called relay attacks on relatively secure transaction authorization devices (such as EMV banking cards). How and if these can be used for online banking is to be determined.

- Virtualization / compartmentalization. Virtualization and compartmentalization (also known as sandboxing) make it possible to run virtual environments on a single machine This could take the form of distinct virtual machines, one for online and one for regular internet use. Since the virtual machines are separated at the machine level, malware must exploit a remote vulnerability rather than a local vulnerability in order to have access to the online banking browser session.

- Hardened browser. A light-weight form of virtualization is to use a separate browser per service provider. This should preferably be a so-called hardened browser. Browser hardening means removing optional features from the browser such that it becomes less vulnerable to exploits. A logical candidate of a feature to be removed is the plugin mechanism making it much harder for an attacker to run a MitB attack.

- Live CD. An extreme form of virtualization is booting the operating system (including a web browser) from a read-only medium (such as a CD or DVD, or possibly a USB thumb drive) and using such a so-called live CD for specific service providers. As the loaded OS is "fresh" on reboot and cannot be affected in a persistent way, it becomes extremely hard for malware to infect such a system.

- Mutual authentication / site authentication. If the user is given the possibility to recognize authenticity of the verifying party (i.e. the service provider) then phishing attacks can be reduced. This is usually achieved by running a TLS/SSL Web site where the server credentials are signed by a trusted certificate authority. Additional controls can be to display, on the welcome screen, before credentials are entered, some unique information that is shared between bank and user (such as a picture that was picked by the user, or the serial number of the user's authentication token) based on a cookie or known to be valid IP-address.

3.4.3 Evaluation

Solution	Costs	User friendliness	Enrolment effort	LoA	Explanation
Knowledge based					
Username/password, static knowledge questions, password replacements, gesture based passwords, virtual keyboard	Low	Medium/high	Low	Low	Easy, but only remembered if used frequently, also crackable and phishable
Possession based					
TAN list	Low/ medium	Medium	Medium	Medium/ high	Recurring low costs, user has to look it up, physically distributed, sensitive to phishing and MitB
Matrix card	Low	Medium	Medium	Medium	Recurring distribution costs, user needs to look it up, physically distributed, sensitive

					to malware and shoulder surfing
OTP token with USB keyboard interface	High	Medium	Medium	Medium/ high	Relatively high cost, user needs to type it over, physically distributed, sensitive to MitB and phishing
Challenge response token	High	Medium	Medium	Medium/ high	Dedicated hardware costs, user needs to type it over, physically distributed, no transaction details to user.
Soft-token or application on PC	Low	Medium	Medium	Medium	Easily copied software, user friendliness depends on setup, may be sent over the internet, sensitive to malware, see also PKI soft certificate
Soft-token or application on smartphones	Low	Medium	Medium	Medium/ high	Easily copied software, user friendliness depends on setup, may be sent over the internet, sensitive to malware
Camera-phone with 2D barcode scanning	Low	High	Medium	Medium/ high	Easily copied software, easy use and installation, common use as a second factor
SIM Application Toolkit (SAT) applet	High	Medium/high	Low	High	Licensing and distribution costs, works on most phones, PIN code as a second factor, mobile network enrolled, secure operation, weakness at the end-point user
EMV/CAP reader	High	Medium	Medium	Medium/ high	Dedicated hardware tokens, bank card PIN code, physically distributed, susceptible to MitB, possibly with optical or connected EMV/CAP reader
PKI USB token	High	Medium	Medium	Medium/ high	Dedicated hardware costs, suboptimal user experience, physically distributed, best with dedicated display and keyboard, see also PKI smart card
USB token with secure back-end channel	High	Medium	Medium	High	Dedicated hardware costs, user has to type over codes, physically distributed, best with dedicated display and keyboard
Visual cryptography	Mediu m	Medium	Medium	Low	Incurs hardware costs, user has to type over code, visual overlay sometimes imperfect, physically distributed, sensitive to malware and phishing
Out of band communications					
Caller Line Identification	Low	High	Medium	Medium	Common service, works in back-end, requires phone nr., may be forged
SMS OTP	Mediu m/high	Medium	Medium	Medium/ high	Transaction costs, typing codes into webforms, requires phone nr., unfortunate business case for attacker
Based on Biometrics					
Face recognition, fingerprint, vein structure, voice	Mediu m/high	Medium/high	High	Low	May require expensive hardware, could be invasive for users, works best under

recognition, keystroke dynamics, GUI interactivity, handwritten signature recognition					controlled circumstances, may be easily fooled
Other					
Risk-based authentication (context) Device identification/ trusted platform IP address based Physical device authentication Distance bounding protocols Virtualization/compartmentalization Hardened browser Live CD Mutual or site authentication					

3.4.4 Discussion

There is a wide variety of authentication solutions, with different characteristics with respect to cost, user-friendliness, enrolment effort, and level of assurance. The more interesting solutions, given the context of the current project, will be solutions in the low/medium, medium, and medium/high LoA range (i.e., level 2 and 3). Particularly possession or out-of-band approaches are suitable as they intrinsically offer strong authentication but suffer from poor registration practices. These solutions, when combined with a WoT style registration process will have a resulting overall (registration + authentication token) LoA of 2 to 3.

3.5 Summary

Despite large variety and complexity, there are generally two aspects common to all authentication processes that determine the quality of the authentication and the assurance of the user's identity:

1. Registration phase: the process that establishes the identity of an individual and binds an authentication credential to the individual.

2. Authentication strength: the assurance the authentication solution provides regarding the identity of the user.

The authentication process for a specific system may have many variations of these two aspects. For example, a system issuing anonymous credentials, may issue a credential in the registration phase, but not establish identity. Similarly the authentication strength may verify that the credential is being used by its rightful owner, but ignore the identity of the credential-bearer, even if supplied.

The quality of both aspects determines the degree of confidence the authentication process can provide, i.e. the level of assurance. The different levels of assurance require different registration methods (e.g. in-person visits, additional documents, etc.) and different authentication solutions (Knowledge-based, OTP, PKI, etc.). Enhancing authentication assurance is not trivial. It is easy to offer a strong authentication solution; it is less easy to actually proof the true identity of the user and bind it to the authentication solution. Both the strength of authentication solution as well as the identity

registration process that constitute the overall authentication strength must be increased. Furthermore, aspects like cost and user-friendliness must be taken into account.

Different registration processes and mechanisms apply to identity vetting, proofing, credentialing and linking, and result in different registration assurance levels. An applicant may appear in person to register, or the applicant may register remotely. In person registration is more secure but very expensive and not very user friendly. In person registration is required to achieve LoA 4. In case the user is somehow not able to register in person, video conferencing tools such as Skype could be used. In this case the user identifies himself via the video conference and shows his passport or other valid photo-ID to the registrar. The use of video conferencing tools for identification, however, has several drawbacks: it introduces scheduling overhead and it makes it harder to detect a forged ID.

Remote registration is limited to LoA 1 through 3 and generally relies on the availability of out-of-bound channels (e-mail, bank account, phone number) and trusted sources to cross-reference and validate the provided assertions such as name, home address, age, and photo. For most service providers some assurance that they are dealing with someone who has sufficient ties to legitimate systems (e.g., financial, mobile, address, or higher education) is sufficient for trusting the provided user identity. Identity proof-ability can often be a serviceable proxy for this kind of confidence. Financial account numbers, mobile phone numbers, credit card numbers, or student registration numbers are examples of such proofing. Yet most of these solutions do not really identify the actual user, i.e., they merely prove that a valid bank account or phone number exists. The combination of multiple channels or sources increases the overall authentication registration assurance level but may come at the cost of complexity. Linking accounts allows for improved user profiling and identification as it allows for attribute comparison between the accounts. This may come at the cost of privacy. Compared to physical presence, remote registration is relatively cheap but is more vulnerable to threats.

The use of web of trust for the registration part of the overall LoA provides a solution that combines the best of existing remote and physical registration methods for authentication. It takes away the need for a physical registration desk as other users in the web of trust take this task. These users perform the identification of the claimant and assure the binding with their authentication credentials. They must, however, be facilitated in doing this, i.e. some kind of attestation service is required. How the web of trust concept can be used to enhance the authentication registration part is the subject of another deliverable in the WoT4LoA project.

Regarding authentication solutions, there is a wide variety of authentication solutions, with different characteristics with respect to cost, user-friendliness, enrolment effort, and level of assurance. Some solutions are not yet mature enough to be used in a typical WoT setting.

The very low level LoA authentication solutions (e.g. username/password) are not relevant for WoT-enhancement, as the overall LoA could never exceed LoA 1. The very high level LoA authentication solutions (e.g. tamper resistant smart card with qualified signature capabilities) are not relevant too, as they would need LoA 4 (and hence in-person registration) to actually reach an overall LoA of 4.

The more interesting authentication solutions, therefore, given the context of the current project, are solutions in the low/medium, medium, and medium/high LoA range (i.e., level 2 and 3). Particularly the possession or out-of-band approaches are suitable as they intrinsically offer strong authentication but suffer from poor registration practices. These solutions, when combined by a WoT style

registration process will have a resulting overall (registration + authentication token) LoA of 2 to 3. Some of these solutions will use a directed identifier (for instance SMS OTP), which may be known by social connections of the user.

4 Use case scenarios

Web of trust could be applied in registration for a new identity (use case 3), to resolve weak registration in an existing identity (use case 2) or as part of step-up authentication wherein an existing identity as a whole needs a higher level of assurance to attain access rights to resources (use case 1). Although other applications of web-of-trust could be conceived, we consider these three the most common in the realm of digital identities. The following use cases illustrate the use of web of trust for enhancing authentication.

4.1 Use case 1: Collaborating researchers

A group of collaborating researchers from various institutions require access to a highly sensitive database. This occurs, for example, in NRENs such as SURFnet in the Netherlands. Access to the database requires strong authentication. The researchers know each other and their institutions participate in a single identity federation. One of them, Alice, however, does not have a strong authentication solution, i.e. she can only authenticate with an unverified username and password. Consequently she cannot access the database. To solve this issue, the other members assert claims about Alice's identity towards a special Attestation Service. They do this by logging into the Attestation Service and indicate that they want to vet for the user's identities. After successful vetting, Alice's authentication level of assurance is increased by the Attestation Service. During the authentication process of Alice, the service provider can check with the Attestation Service for the authentication level and can decide based on the information obtained whether or not to grant her access to the database.

4.2 Use case 2: Social network as web of trust

Bob has a LinkedIn account. The account is protected with a username and password combined with SMS-authentication. That the account indeed belongs to Bob, however, hasn't been verified by LinkedIn. The consequence is that the overall authentication level of assurance is low. To increase the level of assurance, Bob logs into the Attestation Service with his LinkedIn credentials. This allows the Attestation Service to select several of Bob's connections that it trusts. It asks Bob to contact and request them to vet for his identity. Three connections vouch for Bob's identity and the fact that at least one of the connections already has a higher authentication assurance level means that Bob's level is raised as well by the Attestation Service from level 1 to level 2. Next time Bob logs in with his

LinkedIn account to other services, the Attestation Service asserts that Bob has been authenticated with LoA 2.

4.3 Use case 3: Increasing identity assurance

Eve asks project manager John to become a member of the team. John does not know Eve and wants to know more about her. John asks the Attestation Service to validate Eve's identity. The Attestation Service looks for connections in the social graphs of Eve and Bob that overlap. Eve is asked to contact several overlapping connections and ask them to attest for her identity with the Attestation Service. The Attestation Service aggregates the attestations and informs John about the outcome. Based on this outcome John decides to grant Eve access to project team resources.

5 Functional decomposition

Next the functional layout of a web-of-trust based authentication solution will be explored. A key component would be a digital application since no existing infrastructure for web-of-trust based authentication is available, which is referred to as the Attestation Service. Based on the use cases a number of roles and functionalities can be distinguished. After determining necessary roles and functionalities a decision space was drawn that describes all options to realize a web-of-trust based authentication solution, which results in a decision model. From the decision space a protocol was derived that gives a model implementation of an attestation service. Finally, based on the protocol a proof-of-concept was developed, which is documented in the final section to this chapter.

5.1 Roles

Three user-roles can be distinguished:

- An Asker that wants to use the Attestation Service to enhance his identity assurance.
- A Helper that attests for the Asker's identity.
- A Moderator that wants somebody's identity (i.e. an Asker) to be attested. The moderator role is identified in Use Case 3 from the previous section, i.e. the project manager.

5.2 Functionality

A functional decomposition of the use cases results into the following functionalities that are required to realise web of trust based authentication enhancement. This overview is shown in Figure 9.

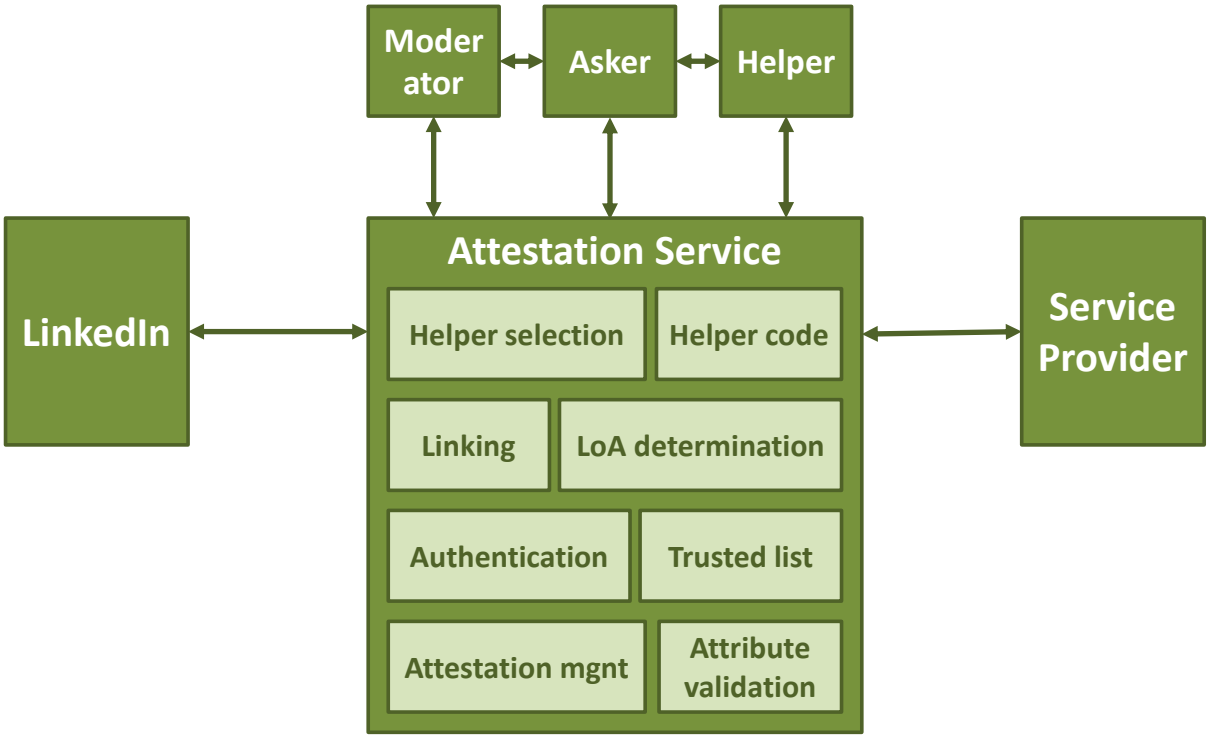

Figure 9: Functional design WoT4LoA Attestation Service.

5.2.1 Attestation Service

The need for an Attestation Service that facilitates and coordinates the enhancement of the authentication solution is obvious. Specific requirements for the Attestation Service are:

- Authentication of the users (Asker, Helper, Moderator). Authentication could be done in federated manner, via social logon, or locally. Ideally, the user has a strong authentication credential with a low LoA due to unreliable registration and binding of the credential to the user's identity. Examples are two-factor authentication solutions like Tiqr or Ubikey.

- Helper selection: what are the best Helpers to attest for the Asker's identity? Candidate helper selection should be such that it mitigates risks related to herd behaviour and fake accounts. Ideally, helpers come from multiple webs of trust (e.g. LinkedIn and Facebook) and have varying relationships with the asker (e.g. friend/colleague, short/long, overlapping skills). If Asker has a PGP key, the PGP web of trust could be utilized as well. In that case the Attestation Service can ask an Asker to provide a PGP key and verify its signatures until it finds a trusted anchor point. In the PGP web of trust a number of anchor points exist. These anchor points are e.g. reputable users that only sign the PGP key of other users when they have physically met or so-called centers of trust whose key is signed most by others. The shorter the path between the Attestation Service's trust anchors and the Helpers, the higher the assurance of the Asker's identity will be.

- Helper code generation to identify Helpers and to associate them to an Asker. The Attestation Service needs to be sure that the selected Helpers indeed are the ones that login to vouch for the user's identity. The Helper code helps to do this.

- Linking of social network to an Asker or Moderator, i.e. giving the Attestation Service to the LinkedIn social graph data. The enables the Attestation Service to select Helpers from the social network web of trust. Users should be able to link their social network accounts to the federated Attestation Service's account.

- Level of Assurance determination based on Helper attestations. Aspects that could be taken into account are: number of Helpers, 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. The Asker shall be notified about an enhancement of the LoA; the service provider shall be informed about the LoA during user login.

- Attestation management, i.e. keeping track of the attestations given by Helpers, informing the Moderator or the Asker, asking Helpers to become trusted Helpers.

- Trusted list: establishing a list of trusted Helpers from which Helpers primarily will be selected. In case of the moderator-scenario, the list consists of the Helpers from the Moderator's social network.

5.2.2 Availability of one or more webs of trust

In order to be able to select reliable helpers for making identity attestations, it must be possible to consult web(s) of trust. Fortunately, these web(s) of trust are available: LinkedIn, Facebook, Orkut, PGP, Google+, … Reusing these webs is preferred above building an own web of trust.

A suitable web of trust must be built very much like a social networking site, because that is the 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, identity federations, social or professional networks should be readily reused to enhance the registration part of the overall LoA.

Pretty Good Privacy (PGP[36]) is based on keys that are signed by others in the PGP web of trust. Signing of keys takes place after physical identification of key holders. The keys are used for encryption purposes and identification purposes.

The Friend of a Friend (FOAF[37]) project created a web of machine-readable pages describing people, the links between them and the things they create and do. FOAF provides a descriptive vocabulary expressed using the Resource Description Framework (RDF) and the Web Ontology Language (OWL).

[36] See PGP website for more information: http://www.pgpi.org/.

[37] Friend of a Friend project (FoaF), see http://www.foaf-project.org/.

The FOAF profiles allow to find, for example, all people living in Europe, or to list all people both you and a friend of yours know.

LinkedIn is the world's largest business social networking hub. Launched in 2003, LinkedIn has millions of users and is implemented in over 200 countries. 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 can invite anyone (whether a site user or not) to become a connection.

LinkedIn provides an interface to obtain basic profile information of users. The fields of the basic profile are amongst others name, headline, location, industry, connections, honours, interests, skills, education, and recommendations. The contact information interface provides fields like telephone number and email address. 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 for the Attestation Service to determine a reliable set of Helpers that can contribute to enhance the Asker's authentication LoA. The same holds for other social networks such as Google+ and Facebook.

Particularly in the context of virtual collaboration organizations in which users know each other, such web of trust based LoA enhancement could be executed in an efficient manner. Moreover this approach also makes it easier to use social identities provided by e.g. Facebook and Google. The registration LoA part of these popular social identity providers is relatively weak (LoA 1) despite the fact that an increasing number of them are using two-factor authentication (LoA 2). Web of trust based enhanced LoA could help increase the registration LoA part of these providers and thus could help in increasing the overall LoA.

5.2.3 Federation infrastructure

There is a need for a federation infrastructure that facilitates login to the Attestation Service, linking of accounts and the communication of the LoA to the service provider.

5.2.4 Attribute validation

Attribute validation could be optional functionality of the Attestation Service. The Attestation Service may ask the web of trust to verify personal attributes of the user such as first name, last name, telephone number, and age. These attributes shall be provided by the user's federated identity provider and the linked accounts (of e.g. LinkedIn or Facebook). Helpers are asked to validate the attributes during their attestation session.

5.3 Decision space

Many implementations of an Attestation Service could be conceived and the desired architecture will depend on its context. Designing an Attestation Service as described above entails a decision making trail encompassing the following questions.

Step 1: Position Attestation Service

The Attestation Service exists somewhere between users and services or content, so the connections to these entities should be mapped. The Attestation Service may be auxiliary to an identity provider or integrated to any degree with an identity provider, i.e. it may be indistinguishable from the identity provider. These choices dictate both the information links and behavior of the Attestation Service, and are true both within the context of an identity federation as well as in other imaginable implementations, e.g. in a reputation system.

Step 2: Authentication and LoA

Firstly, it is decisive when to use the Attestation Service, for example is it part of step-up authentication or primary registration? The following question would be what authentication means it is partnered with. Which begs what LoAs are assigned to those combinations, and how the LoA is calculated and kept.

Step 3: Moderator Role

It should be considered next whether someone should be able to moderate the attestation process. This person (the Moderator) may have the ability to initiate, view, append, edit, delete or overrule attestation processes. All or some attestations could be run by the moderator, or the attestation may exist as a tool that a moderator may choose to utilize for selective cases. The existence of a moderator may undermine the existence of the Attestation Service, foregoing the utility of automation. An example of a moderator is a project manager who wants to validate the identity of a new project member.

Step 4: Trust Network

As explained above, a trust network could be consulted to select Helpers from. This could be a federative network or external network such as a social network. A trust list could be generated real-time from trust networks or kept locally, e.g. based on a periodic copy. Another option is to have people create the trust list, for example Helpers may share their network contacts and details, or a Helper or Moderator may put their own attributes on the trust list. This last action may be prompted by the Moderator. Alternatively, a Moderator may be able to assign a particular trust list to an Asker, for example based on his own social contacts.

Step 5: Initiating an Attestation

An Asker may initiative an attestation process, but a Moderator may also prompt an Asker to get attested. For example, in the case of a project manager who wants a new member of his team to have access rights to certain services or information. Theoretically, the service provider may also initiate an attestation session. This is unlikely to happen in practice as service providers do not want to spend additional effort in authenticating users, i.e. they just want an assurance level that is sufficiently high. Alternatively, the Attestation Service may intelligently start attestation processes, e.g. as part of a periodic check. During initiation certain parameters for the attestation process may be filled in, such as providing attributes of the subject of attestation or setting the trust network(s) to be used.

Step 6: Registration with the Attestation Service

Users (i.e. Askers, Helpers and Moderators) may sign in using their federative identity, preferably through Single Sign On (SSO), or an accepted non-federative account such as a guest account or social login. Alternatively, the Attestation Service may function as an independent identity provider. The registration process with the Attestation Service is directly tied to the selection and verification process. Also, the accounts permitted predict the reliability of the system. Users may be required to sign in again later with an account belonging to a trust network, e.g. a social network that is used to draw Helpers from.

Step 7: Selecting Helpers

Firstly, it should be determined what trust network of an Asker may be utilized, which could be a federative or external network such as a social network or PGP. This trust list may then be compared to the trust network determined at Step 2 and purged further to compile a final list of shared contacts to draw helpers from. One question is how to link the trust networks: based on a unique id, certain attributes such as names, or a combination thereof. An algorithm may be applied to select Helpers from the final list of trusted potential Helpers. Ideally this algorithm is highly intelligent and uses rich data, such as the duration of the relation between Asker and potential Helper, number and date of communications between Asker and potential Helper (Kim et al., 2011), nature of the relationship, etc. In reality however the attributes available are limited, possible variables are the number of selected Helpers, level of assurance, negative vets, trust history between Helpers, time between request and vouching, date first sign-in with Attestation Service, location, family names, employer, etc. Additional attributes may be requested from users during the attestation processes. Another approach could be random selection from either the Asker's own contacts or a list of trusted potential Helpers. Yet another option is to have Askers select Helpers themselves. This approach would greatly simplify the Attestation Service, although it is problematic to link potential Helpers to any trust network, and it is likely to inspire socially desirable behavior.

Step 8: Approaching Helpers

The following decision is how to approach Helpers for verification. An Asker could approach Helpers by in real life, by phone or live-chat, but any low-cues means of communication, especially e-mail, are discouraged, given the inability to confirm identities. Although it would be difficult to dissuade or detect users communicating by e-mail. The Attestation Service could also approach Helpers, for example through verified e-mails or communications channels belonging to the trust networks (e.g. social networks). It is possible to notify (potential) Helpers before and/or after selection and ask for approval. This preserves the anonymity of Helpers, avoids socially desirable behavior because it enables Helpers to reject Askers without them knowing (plausible deniability) and may be more effective since requests are less likely to remain unanswered. It may also be conceivable that Moderators approach Helpers. Also, it is important to consider how to process feedback, i.e. absence or denial of cooperation.

Step 9: Verification process

After the Asker has approached a Helper, this person has to vouch for the Asker by logging into the Attestation Service and confirming the Asker's identity. This may be performed through different means, for example by passing a code, an open hash, a checkbox, a button, etc. During this process the Helper may also be asked to verify attributes provided by the Asker. Also, a Helper may be

presented with the option to give a negative vet, or explicitly decline the Asker access. One consideration then becomes how to process negative vets, whether they are decisive.

Step 10: Communication and notification

Communication between the Attestation Service and the different roles is a matter of concern. Approaching Helpers by the Asker, Moderator and Attestation Service is covered at step 8. Additionally, there is communication between the Moderator and Asker, for example to inform the Asker that attestation is required, and the Attestation Service to the Asker, for example to inform the Asker when attestation has been completed. For each of these interactions channels have to be considered, digitally from the Attestation Service directly, through channels belonging to the trust networks (InMail in LinkedIn), or via text message, mobile app, chat or e-mail, in real life by phone or face-to-face. Next should be decided who is informed of what, i.e. who is entitled to what overview of the attestation process. For example, should Helpers be able to see what other Helpers have attested for an Asker? Should a full history of who vetted for whom be available to all users?

Step 11: Additional functionality

Once a complete model of an Attestation Service has been decided, several additional functionalities could be imagined. For example, continuous authentication by have users vouch for identities and attributes when they login to a service like Captcha. Also, the Attestation Service could be embedded or extended with a reputation system that assigns scores to the actions that take place in the attestations. Another extension could be to implement some open means of communication between users in the Attestation Service, e.g. a forum where users may interact as an antechamber to the Attestation Service. Some legal dependencies should be considered in the design of the Attestation Service, e.g. by users agreeing to a EULA or other waiver at the start of an attestation.

The subsequent steps and the decisions to be made are depicted in Figure 10.

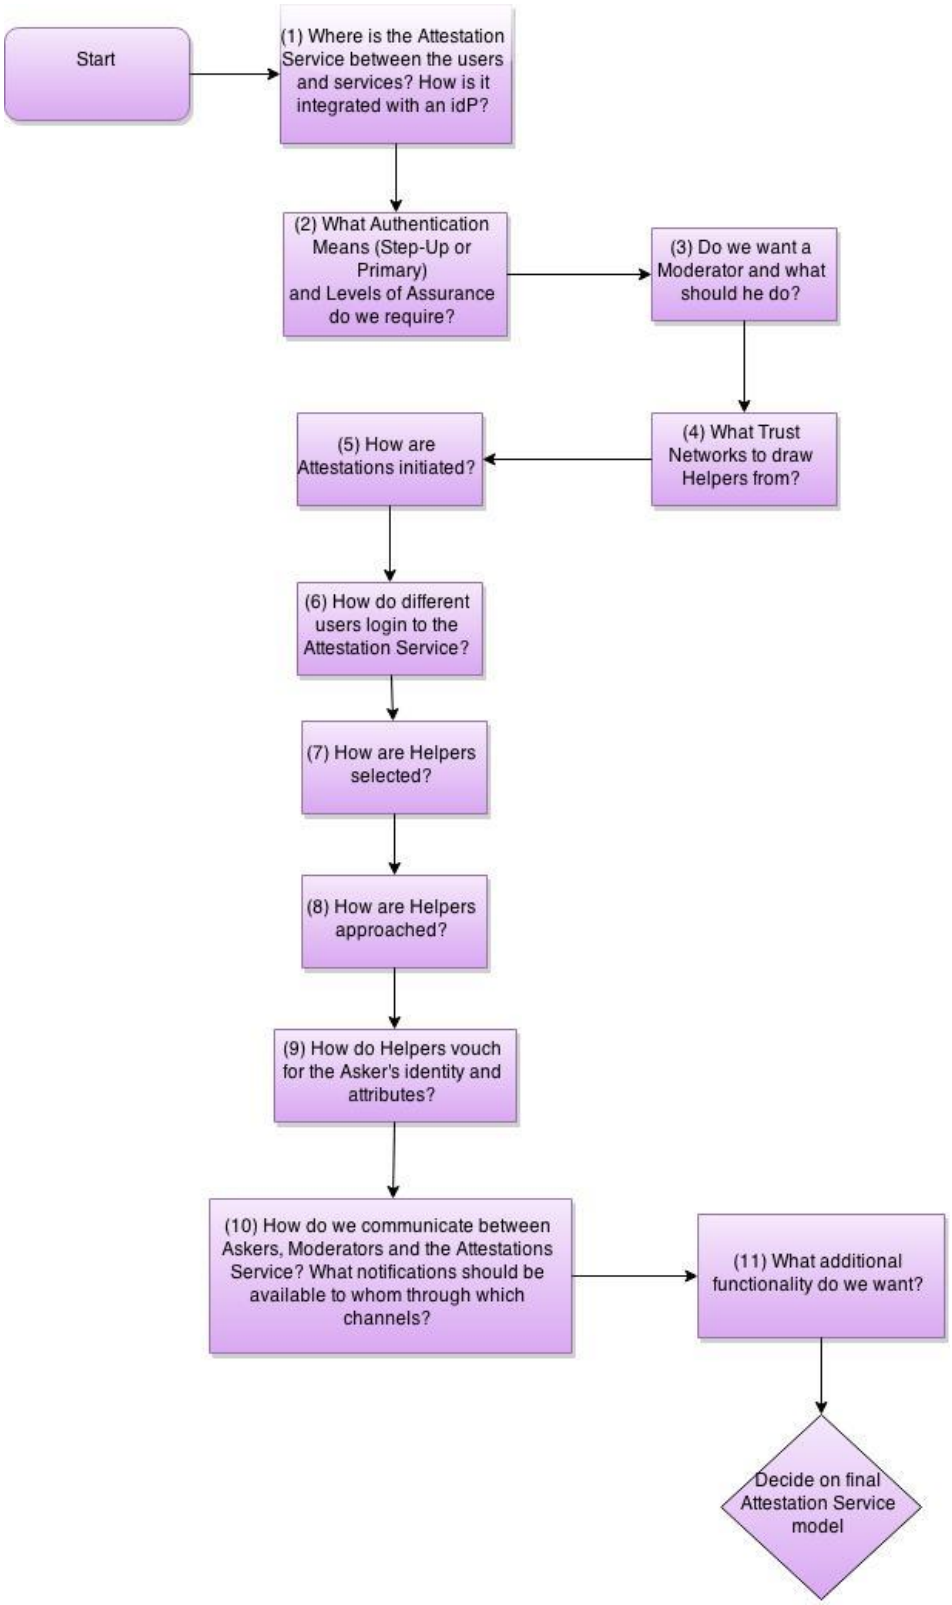

Figure 10 : Decision tree for Attestation Service.

5.4 Protocol description

Based on the analysis of the use cases we propose the following protocol for web of trust enhanced authentication:

Step 0: *Building Trust List, Moderation:* A list of trusted potential Helpers may need to be created, if no such list is available. A Moderator may make an attestation request for a particular Asker.

Step 1*: Registration of Asker.* Asker registers at Attestation Service by logging in with her federated identity and requests for enhancement of authentication. The federated authentication response of the identity provider contains identity information of Asker and is used by the Attestation Service to enhance Asker's authentication assurance. The information at least contains a LoA attribute and value and Asker's federated user identity identifier. Asker is asked to link her federated institution account with e.g. her LinkedIn account by logging in with her LinkedIn credentials. Asker may also be asked to provide her PGP key.

Step 2*: Web of trust scoping.* 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 or PGP) the Attestation Service should know which Helpers and how many are required. Or, in case PGP keys are used, when it should stop with PGP key validation. Asking too many Helpers will burden the Asker as she has to contact them. Subsequently, Asker is given a vouching code and is asked to contact the Helpers by phone or physically and give them the code. The use of e-mail is prohibited or deprecated; Asker has to affirm that she will adhere to this policy.

Step 3*: Passing of vouching code.* Asker calls or meets Helpers and gives them the vouching code. During the phone call or meeting, the Helpers implicitly authenticate the Asker (e.g. via voice or face recognition or by asking questions); this will be used by the Attestation Service to enhance the strength of the authentication of Asker eventually. The passing of vouching codes could be facilitated via a mobile app: by bumping the mobile phones of the Asker and Helper together the vouching code is shared. Such an app not only improves the usability of the vouching process, it also proves that the Asker and Helper have met each other physically.

Step 4*: Helper vouching.* The Helper logs in to the Attestation Service with his federated identity credentials. The authentication solutions he is using must have a higher assurance level than Asker's current level. After successful authentication, 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 vouch for Asker's identity. Optionally the Attestation Service may show Asker's personal attributes and asks Helper to validate them. After that the Helper logs out. If the Helper is not one of the selected Helpers, he will not be able to vouch for Asker's identity. 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. 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 determines the LoA of Asker based on Helper feedback. Given that the scoping of the web of trust is done properly, the number of Helpers and their LoA that actually have given an attestation are key to LoA enhancement. Mapping web of trust based

LoAs to LoAs of existing frameworks like STORK QAA or ISO29115 is not possible; the latter frameworks do not take web of trust based mechanisms into account. Consequently we defined our own web of trust based LoA framework that consists of three levels.

1. WoT LoA1: identical 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 will 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 LoA with the number of positive vets divided by the number of negative 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 via her federated identity provider. The service provider requires LoA 2 authentication. Multiple solutions are possible for the communication of the LoA. A possible solution is that the institutional identity provider authenticates Asker at 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 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 via e.g. Where Are You From (WAYF) functionality.

The different steps are illustrated in Figure 11 below.

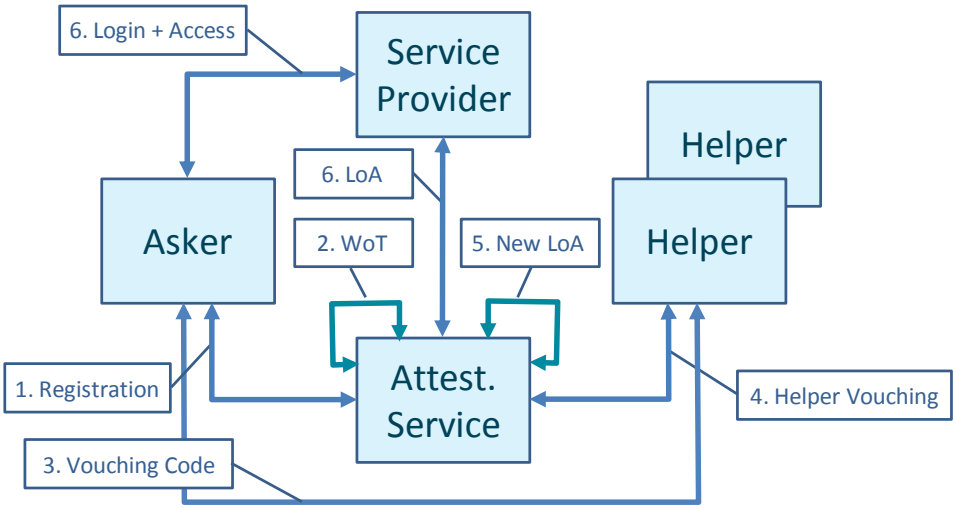

Figure 11: Web of trust protocol flow.

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[38].

5.5 Implementation

A proof-of-concept Attestation Service has been implemented. It currently allows users to login with a local username and password combination. This can easily by extended with other federated authentication or social logon 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. The Attestation Service requests one attribute to be verified. This allows the Attestation Service to randomly select 5 Helpers from the LinkedIn web of trust of the Asker. However, the Attestation Service also keeps its own priority trust list, users can put themselves on the trust list via LinkedIn or share their LinkedIn contacts. The Asker is then presented a vouching code that is alphanumeric and consists of five characters, with the request to approach the five Helpers selected, but not by e-mail. Helpers should login to the Attestation Service with their federated account and fill in the vouching code to verify the Asker's identity. If all five Helpers have vouched, the hypothetical LoA of the Asker is bumped up from 1 to 2. A Moderator may also request an attestation for an Asker, view the progress of attestations and set his own LinkedIn contacts as the trust list to select Helpers from.

Step by step:

[38] Brainard, J.; Juels, A.; Rivest, R.L.; Szydlo, M.; Yung, M.: Fourth Factor Authentication: Somebody You Know, CCS'06, October 30–November 3, 2006, Alexandria,Virginia, USA.

1. Users share their own LinkedIn id or contacts with the trust list; a Moderator may request an attestation process for an Asker
2. Asker is referred to Attestation Service and logs in.
3. The Asker selects ‘get yourself attested’
4. Attestation Service asks Asker for one attribute to have verified.
5. Attestation Service is allowed access to LinkedIn network of Asker.
6. Attestation Service obtains social graph from LinkedIn network of Asker.
7. Attestation Service determines suitable candidate Helpers by contrasting the social graph with its trust list (or randomly selects 5 in the absence of a trust list).
8. Attestation Service generates a helper code and requests Asker to pass the code to the selected Helpers.
9. Asker logs out.
10. Offline: Asker and Helpers exchange helper codes.
11. Helper logs in.
12. Helper enters username and code.
13. The Attestation Service asks Helper to validate attributes. Helper validates attributes.
14. Attestation Service shows identity information of Asker and asks for an attestation: “Is this user indeed Asker?”.
15. Helper attests.
16. Helper logs out.
17. Attestation Service determines the new authentication assurance level of Asker.
18. Asker goes to service provider and is requested to login.
19. Service provider retrieves assurance level from Attestation Service
20. Service provider grants Asker access.

These steps are illustrated in the activity diagram below (Figure 12). The corresponding data model is shown in Figure 13.

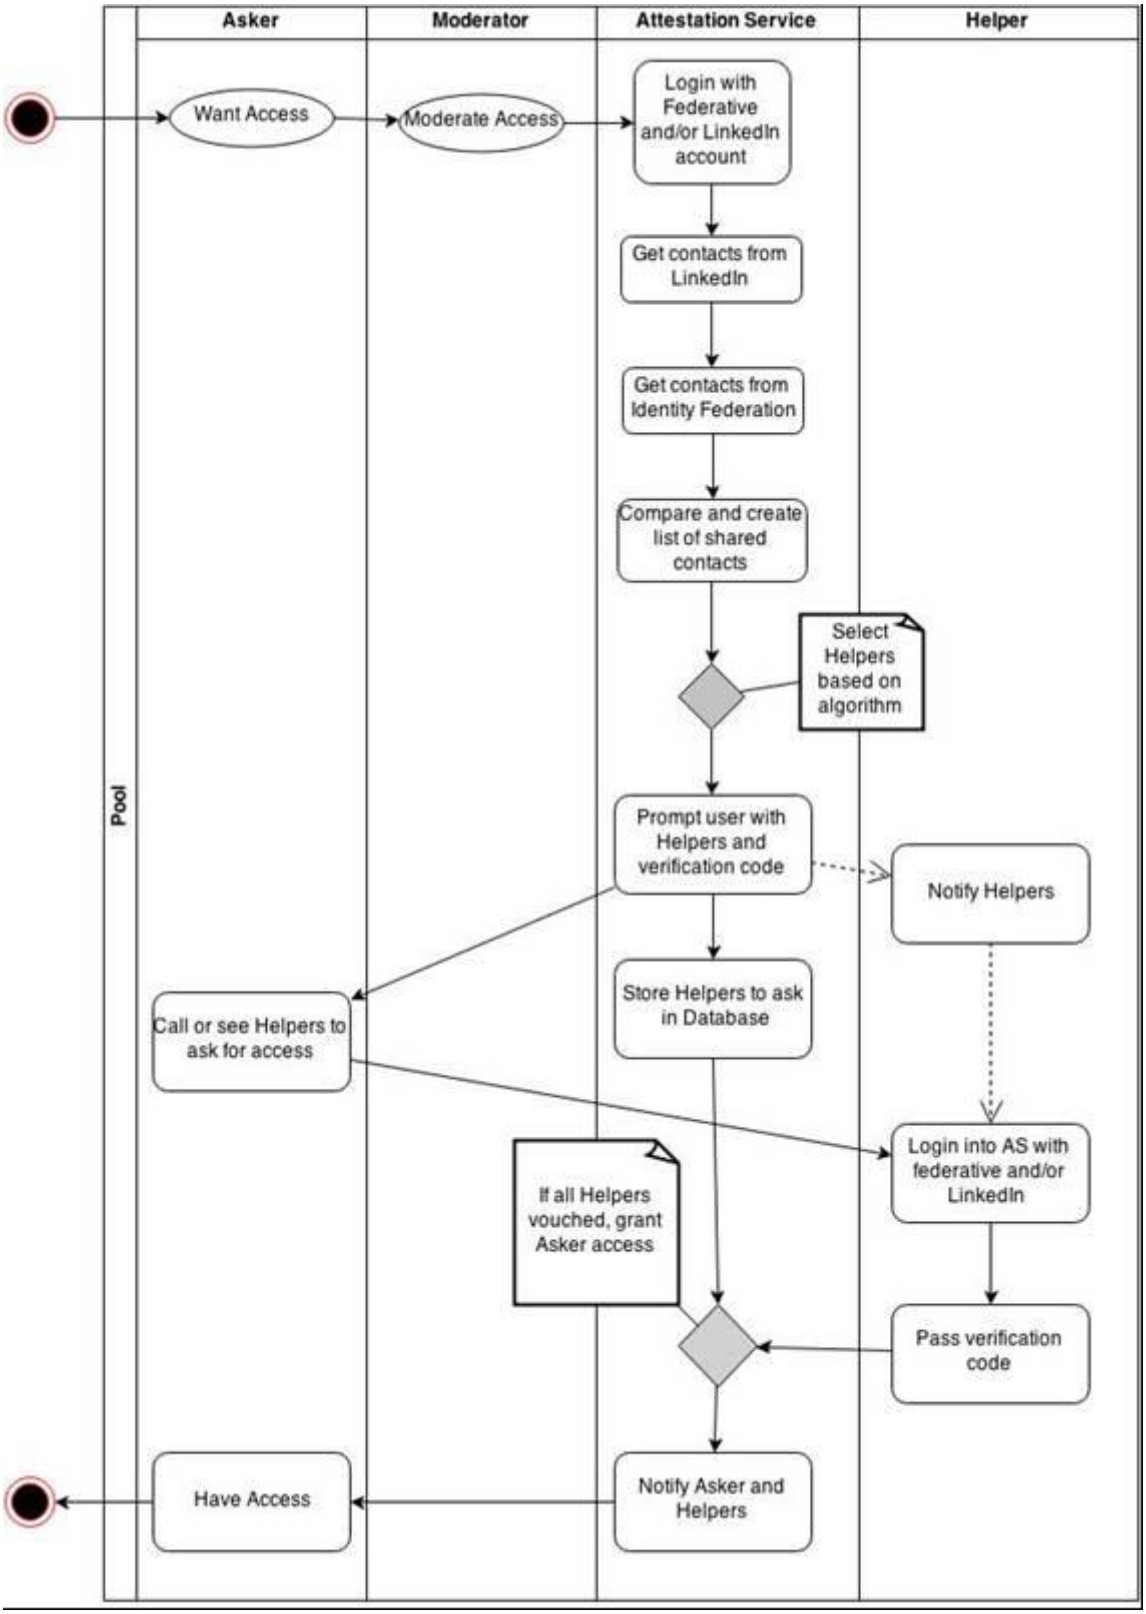

Figure 12: Swimlane activity diagram.

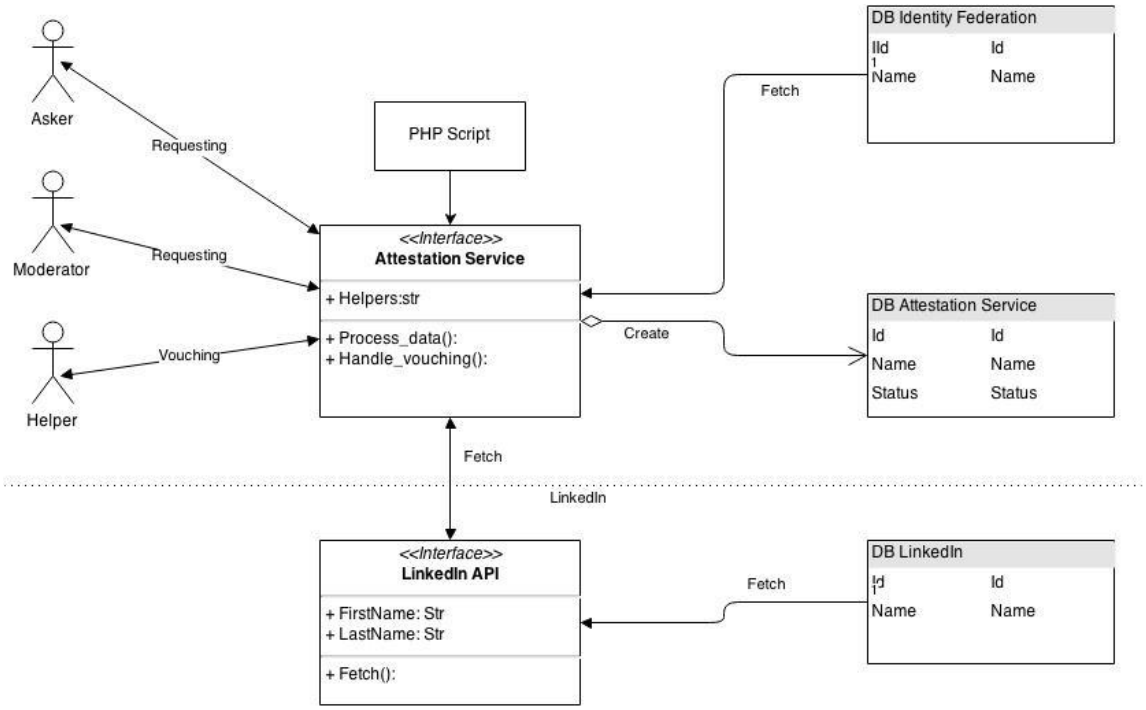

Figure 13 : Data model WoT4LoA.

To give an impression of the proof-of-concept implementation several screen dumps are shown below.

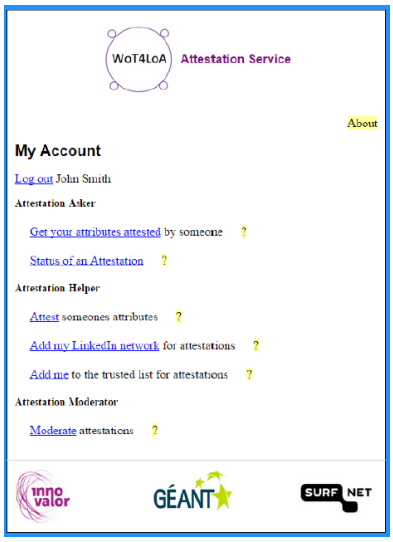

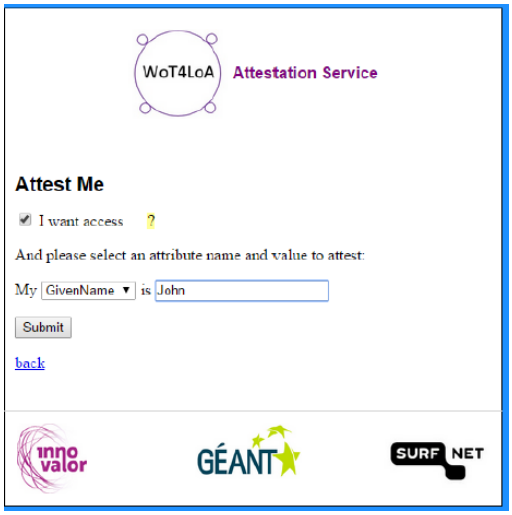

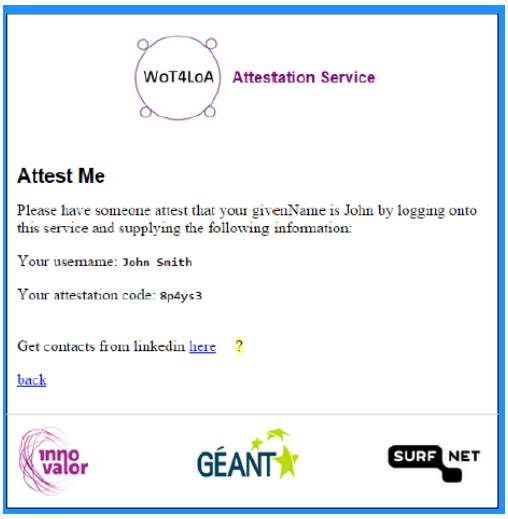
WoT4LoA
Attestation Service
Attest Me
Please have someone attest that your givenName is John by logging onto this service and supplying the following information:
Your username: John Smith
Your attestation code: 8p4ys3
Get contacts from linkedin here ?
back

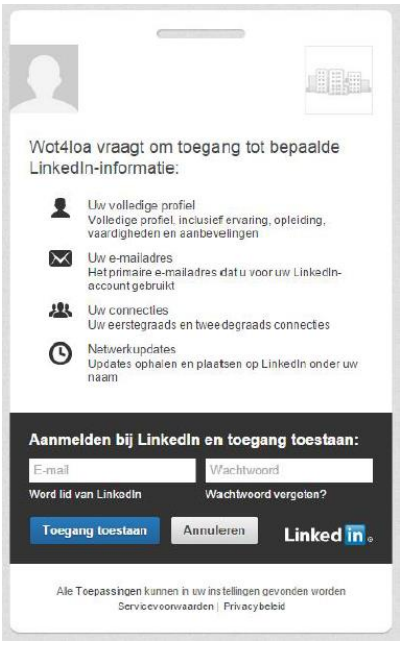
Wot4loa vraagt om toegang tot bepaalde LinkedIn-informatie:
Uw volledige profiel
Volledige profiel, inclusief ervaring, opleiding, vaardigheden en aanbevelingen
Uw e-mailadres
Het primaire e-mailadres dat u voor uw LinkedIn-account gebruikt
Uw connecties
Uw eerstegraads en tweedegraads connecties
Netwerkupdates
Updates ophalen en plaatsen op LinkedIn onder uw naam
Aanmelden bij LinkedIn en toegang toestaan:
E-mail
Wachtwoord
Word lid van LinkedIn
Wachtwoord vergeten?
Toegang toestaan
Annuleren
Linked in
Alle Toepassingen kunnen in uw instellingen gevonden worden
Servicevoorwaarden | Privacybeleid

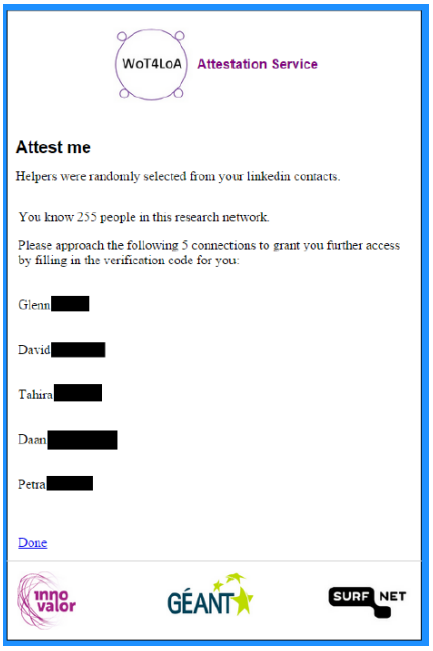
WoT4LoA
Attestation Service
Attest me
Helpers were randomly selected from your linkedin contacts.
You know 255 people in this research network.
Please approach the following 5 connections to grant you further access by filling in the verification code for you:
Glenn [redacted]
David [redacted]
Tahira [redacted]
Daan [redacted]
Petra [redacted]
Done


WoT4LoA
Attestation Service
Your attestation
username: John Smith
listid: 26CjisbsB2
givenName: John
your Level of Assurance is: 1 (you do not have full access)
your verification code is: 8p4ys3
The following people still have to attest: ?
Glenn [redacted]
David [redacted]
Tahira [redacted]
Daan [redacted]
Petra [redacted]
The following people already have attested:
back

Figure 14 : Screenshots Prototype Attestation Service.

5.6 Summary

The functional components and their interrelationships for achieving web of trust based enhancement of the authentication level of assurance of a digital identity have been identified and described. The functional components are derived from a number of use cases.

Important functional building blocks are shown in the figure below.

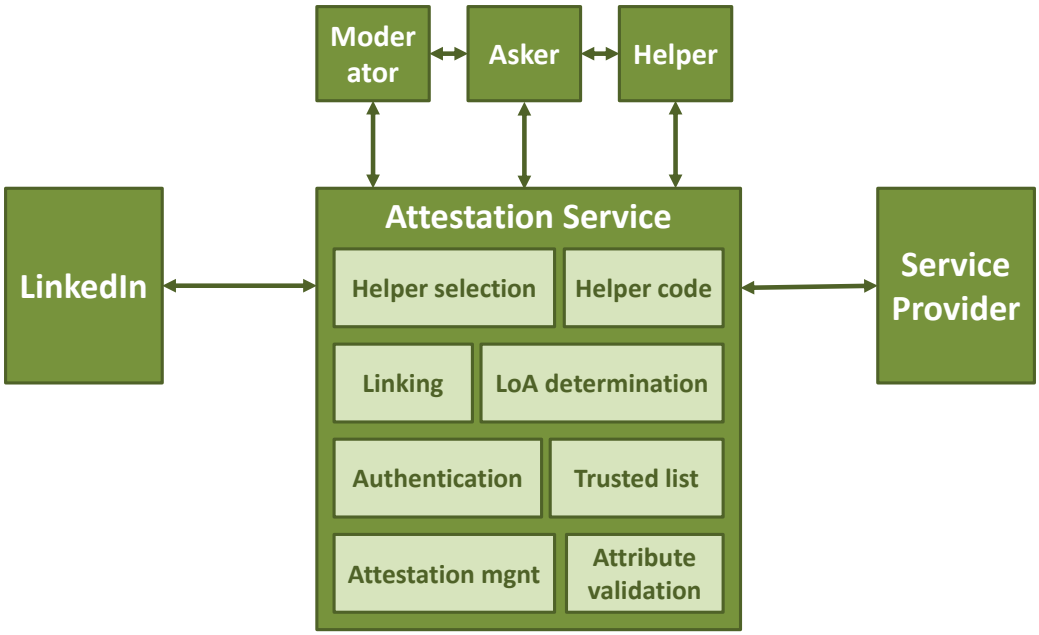

Figure 15 : Functional building blocks.

Candidate helper selection should be such that it mitigates risks related to herd behaviour and fake accounts. Ideally, helpers come from multiple webs of trust (e.g. LinkedIn and Facebook) and have varying relationships with the asker (e.g. friend/colleague, short/long, overlapping skills).

The Attestation Service determines the LoA of an Asker based on Helper feedback. Aspects that should be taken into account are the number of Helpers, LoA of Helpers, and the number of invited Helpers that did not vouch.

A protocol is described for achieving authentication enhancement. A high-level implementation description is provided as well.

There are many different variations within an Attestation Service, depending on its context. Technical feasibility is also by the webs of trust used, e.g. if social media are incorporated, you are bound by their APIs and license agreements.

6 Characteristics and SWOT

The section discusses several characteristics of using web of trust for authentication purposes and conducts a SWOT analysis of the concept.

6.1 Characteristics discussion

6.1.1 Risk assessment

One of the major risks in both scenarios is that the identity of someone with a low LoA authentication solution has been hacked. The Attestation Service has little assurance about the true identity of the user. However, if the identity has been hacked, the hacker will be asked by the Attestation Service to contact helpers from hacked identity's web of trust / social network. It is expected that the helpers will unmask the imposter during contact and refuse to attest for his identity.

6.1.2 Usability of the web of trust model

Ideally, the web of trust approach makes sense in situations that involve a user with a relatively strong authentication solution (LoA 2 or higher) in combination with a poor registration assurance (LoA 1 or lower than the Authentication solution LoA). In these situations the Registration LoA part can be increased via the web of trust approach. This makes the web of trust approach suitable for step-up authentication scenario's that allow the user to add a second factor (token) to his/her first authentication factor (password) without too much registration hassle.

Will the web of trust model also able to deal with account revocation or a change of second factor authentication token? In particular, what to do if a person loses access to the additional factors used to authenticate them such as their phone or computer? Bad guys will try to hijack accounts through account recovery systems, but this poses hard challenges since the recovery systems have to help the real owner who has truly lost access to those other factors. In case of a stolen or lost second factor, the user cannot login anymore at service providers and the Attestation Service will not be consulted. The same holds for revocation of an account: either the federated identity provider of the user's institution will revoke the first factor or the external, non-federated social login provider will revoke

both factors. In case of a change of second factor, the user must link the new second factor to his/her first factor and ask the web of trust again to attest for his identity.

6.1.3 Limitations

Obviously, the web of trust approach can never achieve an overall STORK, NIST, or ISO29115 LoA 4 because that always requires physical presence during registration.

6.1.4 Requirements

A number of requirements can be derived from the use case:

- The need for an attestation service that facilitates and coordinates the enhancement of the authentication solution. Specific requirements for the attestation service are:
    - Determination of the identity of the user;
    - Linking of accounts of users to be able to select helpers from the web of trust;
    - Selection of suitable helpers from the web(s) of trust that can make valuable attestations about the identity of the user;
    - Helper validation prior to vouching – is het helper logging in at the attestation service indeed one of the selected helpers?
    - Collection of attestations from the web of trust;
    - Validation of the attestations;
    - Determination of the authentication strength;
    - Communication of the outcome to the service provider;
    - Optionally: Asking the web of trust to verify other personal attributes of the user such as first name, last name, telephone number, and age.
- The availability of one or more webs of trust that can be exploited by the service to achieve enhancement;
- The need for a federation infrastructure that facilitates the communication of the LoA to the service provider.

6.1.5 Functionality

A dedicated Attestation Service is required that facilitates the process of authentication LoA enhancement. Preferably the Attestation Service is part of the identity federation. Users from outside the federation cannot use the service unless she participates in a federation that is inter-federated with the Attestation Service's federation.

The Attestation Service must be able to select suitable helper candidates from one or more web of trusts that could vouch for a user that is asking for an authentication enhancement. For instance, Helpers of institutions that participate in the same identity federation that the Attestation Service and Asker's institution belong to are preferred. Other candidates are social networks like LinkedIn or Facebook. If Asker has a PGP key, the PGP web of trust could be utilized as well. In that case the Attestation Service can ask Asker to provide her PGP key and verify its signatures until it finds a trusted anchor point. In the PGP web of trust a number of anchor points exist. These anchor points are reputable users that only sign the PGP key of other users when they have physically met or so-called centers of trust whose key is signed most by others. The shorter the path between the Attestation Service's trust anchors and the Helpers, the higher the assurance of the Asker's identity will be. We stress that the Attestation Service reuses existing web of trust structures and does not create its own web of trust (unlike many other reputation or web of trust based systems such as Ebay or AssertID[39]).

6.2 Alternative approaches

6.2.1 PGP

PGP can be considered as an alternative. PGP is based on keys that are signed by others in the PGP web of trust. Signing of keys takes place after physical identification of key holders. The keys are used for encryption purposes and identification purposes. The uptake of PGP is rather limited, despite the fact that it already exists for several years. This is probably due to several reasons:

- It requires client software to work; users have to install software that does the signing of keys and encryption of files/emails.
- The current implementations suffer from poor usability and are unsuitable for the average user.
- Physical identification and presence is required for key signing.
- It is based on transitive trust and in a peer-to-peer setting. This trust model is not always applicable. For instance, a service provider does not want to rely on transitive trust regarding the identity of a customer.

PGP is based on trust anchors. If a trust anchor has signed someone's PGP key, it can be trusted depending on the number of intermediate entities, i.e. chain of trust between the user's PGP key and

[39] Choi, J.N.; Trilli, K.: AssertID – Leveraging Social Networks for Online Identity Verification, http://www.assertid.com.

the trust anchor. In the context of an attestation service this means that the service must have trust anchors as well. Given that the attestation service will facilitate multiple domains and research communities this seems an unlikely and badly scalable model. Nevertheless, PGP provides a web of trust that can be utilised to determine suitable helpers based on signed key chain. Users that have signed another user's key are suitable helper candidates. After all they are supposed to have met physically and know each other.

Brondsema and Schamp[40] have created a system called Konfidi[41] that combines a trust network with the PGP Web-of-Trust (WOT). 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 mechanism and metric, trust values are inferred.

Finally, PGP is not used for authentication purposes. It will have to be extended with LoA-determination functionality to make it suitable for authentication LoA statements.

6.2.2 FOAF

Friend of a Friend is another approach. The Friend of a Friend (FOAF) project created a web of machine-readable pages describing people, the links between them and the things they create and do. FOAF provides a descriptive vocabulary expressed using the Resource Description Framework (RDF) and the Web Ontology Language (OWL). Computers may use these FOAF profiles to find, for example, all people living in Europe, or to list all people both you and a friend of yours know. This is accomplished by defining relationships between people. Each profile has a unique identifier (such as the person's e-mail addresses, a Jabber ID, or a URI of the homepage or weblog of the person), which is used when defining these relationships. Like PGP, FOAF has limited adoption on the web. What FOAF is missing is a trust extension that allows people to create ratings for one another. A proposal for such a trust extension is described[42]. Like PGP, the FOAF infrastructure could be used by the Attestation Service as a web of trust for selecting candidate helpers.

There are numerous approaches to calculate the trustworthiness or reputation of a user. These models often focus on the expression a user's trustworthiness on matters other than personal identification.

6.2.3 Social network aggregation

Social network aggregation is the process of collecting content from multiple social network services, such as MySpace or Facebook, into one unified presentation. The task is often performed by a social network aggregator, which pulls together information into a single location, or helps a user consolidate multiple social networking profiles into one profile. Various aggregation services provide

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

[41] Konfidi, available at: http://konfidi.org/.

[42] Jennifer Golbeck & James Hendler, Accuracy of Metrics for Inferring Trust and Reputation in Semantic Web-Based Social Networks, in International Conference on Knowledge Engineering and knowledge Management (EKAW), Northamptonshire, 2004.

tools or widgets to allow users to consolidate messages, track friends, combine bookmarks, search across multiple social networking sites, read RSS feeds for multiple social networks, see when their name is mentioned on various sites, access their profiles from a single interface, provide "lifestreams", etc.

Social network aggregation platforms allow social-network members to share social-network activities like Twitter, YouTube, Stumbleupon, Digg, Delicious, with other major platforms. Technically, the aggregation is enabled by APIs provided by social networks. For the API to access a user's actions from another platform, the user will have to give permission to the social-aggregation platform, by specifying user-id and password of the social media to be syndicated. Increasingly, social network services are moving away from "walled gardens" to more open architectures. Some sites are working together on a "data portability workgroup", while others are focusing on a single sign-on system such as OpenID to allow users to log on across multiple sites. The OpenSocial initiative aims to bridge the member overlap between various online social network services.

Examples of social network aggregators are FlipBoard[43], Glossi[44], RebelMouse[45], Hootsuite[46], and Flavors.me[47].

Calculating trust from social network aggregation is not new[48], [49], [50]. 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 features like skills and context (e.g. colleague, friend or group membership).

6.3 SWOT Discussion of WoT approaches

A web of trust based LoA approach raises several challenging questions that need to be addressed. The following sections discuss the strengths, weaknesses, opportunities and threats of web of trust based approaches, followed by a feasibility analysis to determine whether threats can be mitigated and opportunities leveraged by using the strengths and eliminating the weaknesses.

[43] See https://flipboard.com/.
[44] See http://slipp.in/.
[45] See https://www.rebelmouse.com/.
[46] See https://hootsuite.com/.
[47] See http://flavors.me/.
[48] Sanguk Noh, Calculating trust using aggregation rules in social networks, in Proceedings of the 4th international conference on Autonomic and Trusted Computing, pages 361-371, 2007.
[49] Sanguk Noh, The Measurable Belief of Trust in Social Networks, volume 202 of CEUR Workshop Proceedings, CEUR-WS.org, 2006.
[50] Heisnam Rohen Singh, Arambam Neelima, Lourembam Suraj Singh, Sarangthem Ibotombi 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.

6.3.1 Strengths

6.3.1.1 *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. So it reduces costs for the authentication service provider.

6.3.1.2 *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 may be reluctant to ask a helper they haven't seen or spoken for quite some time to attest for their identity.

6.3.1.3 *Easy integration in existing federation infrastructures*

The Attestation Service can easily be integrated in the 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 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.

6.3.2 Weaknesses

6.3.2.1 *Security*

One of the challenges is related to several weaknesses the web of trust approach has. ENISA has summarized the possible threats such as whitewashing attack, Sybil attack, 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[51]. The proposed approach does not mitigate all of these threats. Most of them, however, are related to the quality of the Attestation Service's reasoning algorithms that it uses to determine the new LoA.

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.

Registration fraud can be deterred by making it more difficult to accomplish or by increasing the likelihood of detection. 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

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

connections or Facebook or Google+ friends. The requirement for Helpers to have a higher LoA than the Asker makes it more difficult to enhance the LoA however.

False identities would either give themselves away by connecting to their old friends, or remain disconnected, in which case they will have poor social ranking. On the other hand, 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[52]. 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[53].

Similarly, herd behavior due to social pressure can be circumvented in a similar manner by selecting Helpers form 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 leverage big data analytics from 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[54] that offers a service that analyses Facebook profiles and determines whether they're likely to be spoof accounts.

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

6.3.2.2 *Liability*

Another weakness is related to liability. The Attestation Service becomes the authority regarding the authentication LoA of the user. It can, however, not easily be made liable for its LoA claims. The service provider has to trust the web of trust based LoA claims of the Attestation Service. The fact that both parties 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.

6.3.2.3 *LoA determination*

A web of trust based Identity 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

[52] M. R. Henzinger, R. Motwani, and C. Silverstein. Challenges in web search engines. ACM SIGIR Forum, 2002.
[53] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network Motifs: Simple Building Blocks of Complex Networks. Science, 298(5594):824{827, 2002.
[54] Trulioo, see www.trulioo.com.

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 obtain from the work of e.g. Jøsang[55] and Neisse[56].

Assuming that the LoA can be determined: How does the web of trust approach fit in the existing LoA frameworks defined by e.g. 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 web of trust based assurance profiles at the global IANA registry that has been setup for this purpose[57]. 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.

6.3.2.4 *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 e.g. 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.

6.3.2.5 *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

[55] Audun Jøsang, Roslan Ismail, and Colin Boyd. 2007. A survey of trust and reputation systems for online service provision. Decis. Support Syst. 43, 2 (March 2007), 618-644. DOI=10.1016/j.dss.2005.05.019 http://dx.doi.org/10.1016/j.dss.2005.05.019

[56] Neisse, R. 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, see http://ricardo.neisse.name/index.php/publications.

[57] See http://levelofassurance.org/process.html for more information.

seducing users to link LinkedIn, Facebook or Google+ accounts to their federated account the bootstrapping problem can be tackled.

6.3.2.6 *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 – e.g. the need to access a Virtual Organisations database containing sensitive data - may provide further incentive for the Helpers to vouch for the Asker. These incentives should cater for a reasonably quick enhancement of the user's authentication LoA.

6.3.3 Opportunities

6.3.3.1 *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.

6.3.3.2 *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 authorisation 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 (see STORK2.0 D.3.2 – QAA Status Report, 17 April 2013). 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.

Another process flow for validating attributes is possible if the service provider already has a profile of the user, which was accumulated in some other way. Some attribute values now need to be validated against an authoritative source, without the need for revealing additional attributes. In this flow the service provider may ask the Attestation Service to validate a certain attribute. We refer to the Simple Validation Service RFC for more information (see https://wiki.surfnet.nl/display/SvS/RFC%3A+Simple+Validation+Service).

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 attests 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 optimise the validation feedback from the web of trust.

6.3.4 Threats

6.3.4.1 *Loss of privacy*

The WoT 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 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[58] is a decentralized algorithm similar to PageRank that can assign reputations using a social network. Alternately, one can propagate reputation ratings along the social

[58] Pengfei Li; Xiaofeng Qiu, NodeRank: An Algorithm to Assess State Enumeration Attack Graphs, 8th international conference on Wireless Communications, Networking and Mobile Computing (WiCOM), 2012.

network, where each agent receives information about potential targets through referral chains[59], [60]. Cryptographic techniques can further improve decentralized algorithms by allowing precise control over the distribution of information among participants without requiring a trusted intermediary.

6.3.5 Feasibility analysis

Based on the above SWOT-analysis a number of strategic questions should be addressed in order to determine the feasibility of the web of trust approach for authentication LoA enhancement. These questions are:

- Can we use the strengths to leverage the opportunities?
- Can we use the strengths to mitigate the threats?
- Can we eliminate the weaknesses to leverage the opportunities?
- Can we eliminate the weaknesses to mitigate the threats?

The answers to these questions are given in Table 4 below.

Table 4: Feasibility analysis of web of trust for enhancing the authentication LoA.

Feasibility questions		**Opportunities**		**Threats**
		Use social networks	**Attribute Validation**	**Privacy**
Strengths	**Cost efficient for identity providers**	The use of social networks can make the registration process more efficient and trustworthy.	Attribute validation is usually an expensive process. By exploiting the web of trust to do so, it can be made cost efficient as well.	Privacy should be a core business of any identity provider. The money saved on less efficient identity proving could be used for the implementation of privacy enhancing techniques.
	User friendly	Users only have to ask helpers from their social network to give an attestation. It should be made very easy to link	Users benefit from the availability of trustworthy, validated attributes. Attribute validation can easily be integrated and helpers can easily be seduced	Users can usually be tempted to give away their privacy in return for something useful. In this case they don't have to visit a registration desk.

[59] B. Yu and M. P. Singh. A social mechanism of reputation management in electronic communities. In Cooperative Information Agents, pages 154-165, 2000.
[60] G. Zacharia, A. Moukas, and P. Maes. Collaborative reputation mechanisms in electronic marketplaces. In Proc. of the 32nd Hawaii Intl. Conf. on System Sciences (HICSS), 1999.

		social network accounts to the attestation server.	to validate certain attributes.	
	Easy to integrate	Linking of social network account should be very easy. This allows for further leveraging of the social web of trust for attestations.	Readily available attributes from the federated or social network webs of trust can be easily integrated in the attestation server for validation purposes.	Users must be explicitly informed about the consequences. They must give consent.
Weaknesses	**Security**	Combining social network based webs of trust mitigates numerous security weaknesses that are inherent to web of trust. It will result in the selection of helpers that can reliably make attestations about someone's identity.	Combining social network based webs of trust mitigates numerous security weaknesses that are inherent to web of trust. It will result in the selection of helpers that can reliably validate attributes.	Taking away the security issues helps to guarantee privacy.
	Liability	-	-	-
	LoA determination	Leveraging existing webs of trust allows increasing the assurance level of someone's identity.	Validation of attributes results in higher assurance levels of the attributes.	-
	Vouching code exchange	Knowing the profiles of potential helpers allows for trustworthy exchange of vouching codes.	-	-
	Bootstrapping	Tempting the user to link his/her social network account to a federated account helps to boost the web of trust enhanced LoA approach.	-	-
	Usability	-	-	-

Based on the above table, 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 VO). 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 standardised 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[61]. 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.

6.4 Summary

The use of web of trust to enhance the registration part of the overall authentication process seems a feasible approach as it 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 identification task. Users in the web of trust may use physical presence, phone or email practices for this purpose. Somehow, the attestations from the web of trust need to be related to the claimant's digital identity. This needs to be catered for by some kind of federated attestation service that enhances the assurance in the claimant's federated identity with attestations from the web of trust.

There are plenty of webs of trust available that can be used. Social networks provide sufficient information to be able to select suitable candidates that can help to improve the authentication level of assurance of a user.

Technically, there are no obstacles for the implementation of a web of trust based enhancement of the authentication LoA. The facilitating attestation service can easily be integrated in the existing federated identity infrastructure. Moreover, interfaces to social network providers are readily available and standardised without the need for solution-specific client side software.

The main challenges are liability and usability. Liability could be tackled by incorporating the attestation service into the existing trust fabric. Usability will be an important aspect in end-user evaluation tests. The outcomes of these tests are described in the next section.

The ability to easily extend the attestation service's usefulness with attribute validation functionality provides a value adding business opportunity.

[61] Boston Consultancy Group, The Value of our Digital Identity, 2012, see http://www.libertyglobal.com/PDF/public-policy/The-Value-of-Our-Digital-Identity.pdf.

7 User evaluation

7.1 Demonstrations

WoT4LoA has been demonstrated during the TERENA Networking Conference 2014 (19-22 May, Dublin, Ireland) and the GN3plus 2015 symposium (23-26 February, Athens, Greece). Mockups have been shown at the Open Identity Summit 2014 (4-6 November) in Stuttgart, Germany. The demonstrator is available for usage via the GN3plus WoT4LoA website[62].

7.2 Evaluation protocol

7.2.1 Approach

Usability testing can be performed through a survey (with a link to the live WoT4LoA environments, instructions and a questionnaire), by placing the attestation service in a test-environment and gathering feedback, through one-on-one testing with a user by going through a test scenario and asking questions, or a similar approach in a group session. Given the timeframe and complexity of WoT4LoA the group session was preferred, to efficiently gather a rich host of feedback.

Usability testing of the attestation service was done internally at the two project partners in two sets consisting of at least 5 users (5 users at InnoValor, 5 users at SURFnet), wherein all scenarios were covered and followed by a series of questions. In preparation the room was setup and the networks of the facilitators were shared with the trust list of the prototype, to increase the chances of the colleagues participating in the tests to show up as Helpers to approach. After explaining the context of the study and prototype, the participants went through the test scenarios described in the next sections. The questions mentioned in the sections on Test Parameters were used for posing questions during and after going through the scenarios. In addition to the questions, remarks of the participants as well as non-verbal communication were used for evaluation of the prototype.

[62] https://intranet.geant.net/JRA0/WoT4LoA

7.2.2 Test Scenarios

Asker wants *access* and approaches 5 Helpers
Ask someone to play Asker, explain he wants to have access to an NREN collaboration service, and have the Asker take place behind a computer. Start on the URL the Asker wants access to, https://wot4loa.aai.surfnet.nl/home.html, and the Asker will be referred to the Attestation Service, follow the instructions on the screen, which should take you to LinkedIn. Instruct the Asker to assess Helpers to approach and the status of an attestation process (by clicking 'back' and selecting the correct option on the home screen).

Helpers vouch for Asker
The Attestation will have selected two Helpers to approach, these should be in the room, otherwise have the Asker start a new attestation process. Have the Asker leave the room and call the Helpers on their phone. He should then ask the Helpers to login with the Attestation Service and follow instructions on the screen, they should select 'attest' and confirm the attributes of the Asker by entering the verification code passed by the Asker. The Helper is asked to share his LinkedIn contacts with the trust list if attestation was successful, if the Helper does not see it, point this question out. If all two Helpers completed attestation, the last one should note this and tell the Asker outside the room. If not, request the Asker to return to the room and view his progress. If all goes well, the Asker should now be able to access the NREN the collaboration service, ask him to do so.

Moderator wants someone to have access
Now ask a fourth person to take on the role of a Moderator. Explain that the Moderator, e.g. a project manager, wants to make an Attestation Request for a new team member he has not met in person and cannot trust, but who needs access to the NREN collaboration service. Ask the Moderator to go to the attestation service at https://wot4loa.aai.surfnet.nl and attempt to make a custom attestation request using his own LinkedIn contacts for his new team-member John Smith. Then ask the Moderator to view the progress of the attestation. Also, instruct the moderator to share his LinkedIn contacts with the trust list, because he trusts his LinkedIn contacts as trustworthy people to verify just access to NREN services.

7.2.3 Test parameters

Several usability aspects of the prototype have been assessed during evaluation, these test parameters are outlined below.

Test parameter	Example question
Understanding (learning curve)	Did you understand what you were supposed to do?
Social acceptance	Was it awkward to ask/be asked and/or use your LinkedIn contacts?
Routing	Was it clear what to click on next?
Efficiency	How hard was it to gain access?

Feedback	Did the attestation service provide enough information at every step?
Usefulness	How useful do you find WoT4LoA?
Security	Do you feel this contributes to information security?
Use	Would you use it yourself?
Usability	Was the attestation service easy to operate?
Functionality	Did the attestation service function properly?
Design	Was the service well-designed? Consider legibility, language, layout, etc.
Remembrance	Was the process easy to remember?
Temporality	How long or short was the time it took to get attested?
Improvement	What would you change?

7.3 Evaluations

The prototype was evaluated at both InnoValor and SURFnet as follows;

Organization	Date	Time	Nr. of participants	Description
InnoValor	9-2-2015	13:05-14:11	8	After introduction went through Asker, Helper and Moderator roles once with a participant for each role, discussion throughout and questions at the end.
SURFnet	16-2-2015	10:30-12:00	6	After introduction went through Asker, Helper and Moderator roles once with a participant for each role, discussion throughout.

Due to time constraints, the Helper scenario could be simulated only once in both evaluations, despite the protocol stating two Helpers. The facilitators demonstrated the second attestation to show a complete attestation process to participants.

7.4 Analysis

7.4.1 Results

The questions and observations from the two sessions are collated per category below, followed by a discussion of the findings.

User interface and implementation
The participants found the attestation service unclear at times. For example, the multitude of roles and options on a single page was confusing. It was not clear at all times in all roles how to proceed, i.e. the prototype would benefit from better process guidance. For example: help-texts were not viewed, dropdown options were not used, buttons were not noticed, etc.

Moreover, the user interface is considered less than perfect, e.g. the freeform attribute field accepts any input. The Attestation Service was considered cumbersome, e.g. too many fields to enter, the attestation code seemed to serve no real function, etc. The participants also noted that the prototype does not support federated authentication.

Login with LinkedIn for retrieving the contacts also asks permission to post updates. Even though this function is not used by the prototype, users experienced it as off-putting.

Process
There was some misunderstanding in the process as well. The words 'attestation' and 'LoA' were not familiar, the moderator role was not sufficiently explained, certain instructions were not intuitive, and it is unclear why (according to which algorithm) a Helper was selected. The use of an attestation code resulted in confusion as well, as participants thought it necessary to memorize it or write it down, or it proved difficult to communicate a case sensitive code. The process itself was also considered conceptually complicated, it was suggested some people would not understand that only specific helpers are asked to vouch. This misunderstanding can lead to non-Helpers being asked to attest.

Furthermore, the process was considered burdensome for everyone involved; more than one phone call between Helper and Asker occurred and the phone calls were deemed troublesome and awkward.

Web of Trust
Some feedback also directly addressed the web-of-trust solution. Firstly, the reliance on a web of trust (i.e. social media) was considered troublesome, as not everyone knows their LinkedIn account details. Secondly not all contacts from LinkedIn are of equal 'closeness' to people, some may have been added a long time ago and no longer be considered relevant, others may not be suitable to approach as helpers at all or one may simply not have the contact details required to approach them (outside of LinkedIn). While there may be many mutual contacts, there needs to be more intelligence in the Attestation Service to select appropriate and accessible Helpers. It was also unclear how to deal with Helpers that are on holiday or do not answer their phone.

Rich communication channels, such as face-to-face or by phone, are considered more secure but may lead to socially desirable behaviour (it is more difficult to say "no" in person) whereas less rich channels such as e-mail are less secure and may be too easily ignored. One suggestion was to approach Helpers without the Asker knowing; this would limit social pressure, however it does limit the Asker

tracking progress and follow-up, and too little social pressure could take away all incentive for Helpers to participate. People preferred communicating by e-mail (which is sensitive to impersonation) and would do this even when explicitly told not to.

More importantly, it was suggested there is no incentive for Helpers to cooperate, e.g. users would not want to be a Helper too often.

It was stated that attestation takes a lot of time compared to other forms of registration. Also, it may reflect badly on the party insisting on attestation, e.g. “my horrible boss wanted me to get attested”.

Which social network is appropriate to use would vary depending on the context, sometimes this would be LinkedIn, while in other cases Facebook is more appropriate. In other comments, social contacts may not know the requested attributes of an Asker, e.g. year of birth. The users questioned whether validating attributes really constitutes a higher LoA and suggested video or photo assurance might be a more direct way to assure identity. Moreover, access management becomes dependent on third parties (Helpers) that have no real incentive to help. Web-of-trust authentication seems relevant only for high-interest access. Its added value might work best in an international setting where it is more difficult to identify people in more direct ways.

Other

Some actual bugs were also discovered during evaluation: Refreshing the page of the Attestation Service that first presents Helpers to approach, also refreshes this list of Helpers. The Attestation Service matches on username. This means users have to login using their exact case sensitive personal names, as provided by LinkedIn. Also, in one occasion, the Attestation Service demanded two attributes to be verified instead of one, and appended the list of Helpers to approach by the name of “_empty_”. Staying logged into LinkedIn also results in using the wrong LinkedIn accounts in the Attestation Service by accident.

Improvements

Evaluation participants have uttered several expansions and improvements; mostly addressing alternative options for an attestation process. Selecting both reliable and familiar Helpers are likely a contradictory concern, it was suggested to leave this wholly or partly to Askers. Helpers may want a virtual reward, e.g. points, or mutual verification as an incentive to cooperate. Helpers could be enabled to express their certainty about an attribute, e.g. on a scale. The Attestation Service would better approach Helpers as well, typically by e-mail. Also, you could leave the burden to call with the Helper. It should be considered whether an attestation is limited to one service, or could provide access to multiple services to justify the respective user investment. Lastly, no negative option for attribute validation was provided to Helpers (such as: “not born in the year mentioned”).

Some other suggestions concerned the technical implementation of the Attestation Service. The Attestation Service would be greatly improved if options were custom to the role of the user and depicted a to-do list. Clicking through would result in autopilot behaviour; hence explicit attribute entry would be better for example. A phone call script for the Asker would be a welcome addition to clearly explain the purpose and request to the Helper. Askers could determine a verification code themselves, i.e. supply a password, or one could use an image to communicate a code safely.

7.5 Discussion

There was considerable negative feedback on the design of the prototype, and it contained a few development bugs still. This is however considered appropriate for a prototype and participants expressed understanding of the concepts of an attestation service, hence the evaluation was deemed successful.

In feedback on web-of-trust-authentication, several comments addressed the functionality of the concept. For example, not everyone actively uses social media, e.g. not knowing username and password. This could lead to frustration both on the part of all three roles (Asker, Helper, or Moderator). Moreover, it is likely situations occur where Askers are unable to reach Helpers, e.g. because they do not possess contact details. Similarly, non-response handling of Helpers could be problematic, while a Helper response cannot be guaranteed. The reliance on others may obstruct or delay authentication and access. Communicating between Helper, Asker and Service could be automated or manually performed through a host of channels to address some or all of these issues, albeit each with their own considerations and trade-offs such as between responsiveness and 'social pressure'.

There are also trade-offs inherent to Helper selection, more reliable and relevant Helpers are likely sparse or less accessible and are difficult to select automatically based on an algorithm. Although 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, this information is typically not available to applications outside of the social network itself; or if they are releasing this information may make the participants uncomfortable as witnessed by the comments on the prototype requesting permission to post on LinkedIn.

Which social network is appropriate for an attestation depends on the context of the service to be accessed by the Asker. There are many such options for the authentication process, each with different implications for the overall usability and reliability of the whole solution. What choices are most prudent depends on the context of the authentication as well, leading to considerations whether an attestation is applicable only for a single access or multiple services. The need to consider the context when selecting the right social network also makes it difficult to automate (comparable to the difficulties in selecting the right Helpers).

The problems in 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.

The concepts of web-of-trust-authentication are difficult to understand and communicate to users. Furthermore, Helpers may not be familiar with certain attributes of the Asker, e.g. their year of birth, which complicates attribute validation; although one could argue it automatically disqualifies those Helpers for their role in the first place.

In qualifications of web-of-trust-authentication, it was evaluated as socially undesirable, slow and cumbersome, and complex by users. It was even suggested it might reflect badly on services to be accessed or organizations applying web-of-trust-authentication.

7.6 Summary

Key findings from the user evaluation study of the web of trust proof of concept are:

- The prototype was not perfectly developed, but did enable conceptual evaluation of a web-of-trust-authentication solution.
- Web-of-trust-authentication solutions would have a lot of implementation options with different trade-offs and choices would depend on the context of the authentication.
- Web-of-trust-authentication has several inherent complications, mainly pertaining to selection, participation and communication of Helpers.
- While most of said complications could be technically resolved, web-of-trust-authentication remains socially laden, time- and effort intensive, and complex.

8 Conclusions

8.1 Results and key take-aways

There is an increasing need for stronger authentication solutions that go beyond username and password. The use of second factor authentication credentials is growing but lack of solid 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. A solid registration process, however, is expensive as it usually requires the establishment of a registration desk and is not very user friendly, as he/she has to go to the registration desk. The latter requirement can even be impossible to meet for remote users.

***Key take-away**: There is an increasing need for stronger and cost efficient authentication solutions.*

The underlying report describes an investigation on how to leverage webs-of-trust for authentication in identity federations for higher education and research. Identities are established in two process; authentication, or how a user demonstrates his identity, and registration, or how that means was bound to that user. The degree of confidence (level of assurance, commonly a scale of 1-4) in identities is thus determined by the strength of each process. While it is considerably easier to implement a more secure authentication solution, reliably proofing a user's identity and binding the solution is more difficult. Web-of-trust as an alternative registration solution was compared to other registration processes and authentication solutions to assess where it fits in the authentication landscape and what level of assurance it could enable. It follows that web-of-trust-authentication is suitable to reach at least level 2, as it is potentially as reliable as a copy of an identity-document, or even somewhat stronger. Level 3 registration quality is less easy to achieve as this requires very accurate selection of helpers from the web-of-trust. Moreover, the identity of the helpers must be determined with a reasonable high assurance level as well. The web-of-trust provides identity proofing, but does not arrange provisioning of credentials, so that element of identity assurance is conditional, as is the strength of the credential and the authentication process. Theoretically, web-of-trust for identity validation could thus be combined with any authentication process.

***Key take-away**: Web of trust provides an alternative identity proofing and registration solution that potentially could reach level 3 assurance.*

Next three use cases where identified as exemplars of how web-of-trust authentication may be applied. Firstly, a group of collaborating research want to grant a new and previously unknown

member access to shared resources. Secondly, identities in social networks are commonly self-asserted, hence web-of-trust could give some certainty about the real identity of users. Thirdly, someone may need a higher level of assurance to gain access with their account, wherein web-of-trust may strengthen the identity validation. Altogether, web-of-trust may have merit in the creation of a new account, to validate an existing identity, or in step-up-authentication.

***Key take-away**: Valid use cases exist for the deployment of web-of-trust enhanced authentication.*

Subsequently, the functional decomposition of a web-of-trust solution was worked out, which resulted in a set of functional and qualitative requirements, as well as the decision space of such a solution. The main insight here was that such an attestation service knows many variations and options that are mainly context-dependent and dictate the ultimate architecture, as well as the process and reliability of web-of-trust authentication. Based on the functional decomposition a recommended protocol was designed and a proof-of-concept developed for implementation. The proof-of-concept entailed a prototype of an attestation service for access to a fictional dashboard upon vetting by LinkedIn contacts.

***Key take-away**: Functionality (i.e. an Attestation Service) is required that facilitates and coordinates the enhancement of the authentication solution via web-of-trust.*

***Key take-away**: The variations and options for the implementation of such functionality (i.e. an Attestation Service) and the processes it should cater for are numerous and complex.*

A feasibility study resulted in a SWOT-overview of web-of-trust as a registration solution. The main outcomes are that while webs-of-trust comprise an opportunity for efficient and reliable identity proofing, the main challenges are in usability and liability. This concedes that usability is to be a prominent objective in any web-of-trust implementation. Webs-of-trust could also be leveraged to confirm user information (attributes), which presents a promising opportunity as well.

***Key take-away**: While web-of-trust comprises an opportunity for efficient and reliable identity proofing, the main challenges are in usability and liability.*

Finally, preliminary user evaluation tests conducted on the prototype proved indeed that usability is a critical factor for the success of web of trust enhanced authentication. The concept is relatively difficult to explain to the user. Consequently, the user experiences barriers for contacting helpers and for motivating them to give an attestation. There is a risk that the whole attestation process may take too long when helpers a reluctant to help. The attestation process should be strictly guided by the Attestation Service in order to achieve successful and timely attestations.

***Key take-away**: Preliminary user evaluation tests confirm that the complexity of the concept is a major hurdle and indicate that it may even be socially uncomfortable and time-intensive.*

8.2 Answers to research questions

How does the model fit in existing federated identity architectures (i.e. hub and spoke, networked)?

Web-of-trust fits in existing identity federation very well, given that they essentially are federated webs-of-trust. Moreover, it could aid identity federations with new functionality. For example, web-of-trust could be useful to introduce social login in identity federations for higher education and research. Currently, social login is not sufficiently validated, but could reliability in the real identity of users could be enhanced through web-of-trust authentication. Nonetheless, the concept does have some limitations, as outlined above.

How to define the metric to determine a certain LoA and what factors need to be taken into account in this metric (number of claims, quality of the claims, history of the claims, coherency of the web of trust, number of web of trusts consulted, …)?

Calculating the LoA could be very intelligent, however it requires a lot of complexity from the attestation service. Also, this implies the attestation service to be used very frequently for such rich data to be meaningful. Two factors appear to be critical: the desired LoA of the Asker and the LoA of Helpers, i.e. how sure do we need to be of your identity and how sure are we of the identities of the people vouching for you? The Attestation Service determines the LoA of Asker based on Helper feedback. Aspects that should be taken into account are:

- Number of Helpers. A simple algorithm could be:

  New LoA = LoA + LoA*(1 – (1 – H1)*(1 – H2)*(1 – H3)….)

  With H = amount of trust [0..1] for each Helper.

  H depends on:

  - LoA of Helper
  - Coherency of Asker – Helpers web of trust such as
  - Duration relationship between Asker and Helper
  - Overlap between multiple WoTs (e.g. LinkedIn & Facebook)
  - Trust relations between Helpers
  - Number of paths between Helpers and Asker in PGP
  - Path length between Helper and Asker PGP
  - Overlapping skills and endorsements in LinkedIn

- The number of invited Helpers that did not vouch. These may be considered as negative vets. They have a negative effect on the new LoA. A simple algorithm is to multiple the New LoA with the number of positive vets divided by the number of negative vets.

An interesting issue is raised whether Helpers with a lower LoA than the desired LoA of the Asker may reliably vouch. On the one hand, the weakest link principle dictates the lowest LoA is equivalent to the

LoA of the overall solution, so Helpers should always have a higher LoA than the Asker. On the other hand, synergy exists between a given amount of vets, so the whole is greater than the sum of its parts. In other words, the reliability of the web-of-trust determines the LoA affordable to the Asker. What other factors are available for calculating a LoA are determined by the webs-of-trust used. For example, a social network such as LinkedIn is limited in the information publicly available.

Can a similar approach be used to improve the quality of the attributes related to the user's identity (e.g. group membership, age, gender, student, etc)?

Yes, it could be used to verify attributes, with a side note that Helpers should be reasonably aware of attributes (e.g., gender is commonly known, age may be less likely, bank account nr. very unlikely) and requires different algorithms than mere identity for Helper selection and assurance. Also, attributes potentially require more frequent updating (address, phone nr, etc.), yet Helpers are less likely to cooperate to confirm attributes than they are to confirm identity to provide access. It could be suggested that web-of-trust may be even more suitable for attribute validation than for authentication enhancement.

What is the impact of this approach on the IdPs that are responsible for user authentication? Can they be made 'web of trust claims' aware and are they able to communicate the enhanced LoA such that service providers are able to deal with it? What will the impact be for the federation operator?

A web-of-trust authentication solution essentially transforms the federation into a (at least binary) reputation system. Use of reputation by service providers (beyond mere yes / no access rights) requires extra attributes (e.g., at least an attribute that states the LoA), so to some extend the IdP becomes an attribute provider for trust attributes. Another implication of web-of-trust authentication would be the position of an attestation service within the federation network. Where initial authorization systems persist, the attestation service is integrated in IdP. Or the attestation service may replace an institutional IdP and become an external IdP itself. As a registration alternative web-of-trust may make identity proofing easier, thus contributing to a more effective federation. It may even allow new members / institutions to participate that were unable to be reliably proofed before. For example, web-of-trust authentication may be applied in supporting of guest accounts. Moreover, it may be useful in step-up authentication scenarios

Does the approach fit in the existing LoA frameworks defined by e.g. ISO, NIST and STORK?

Most existing LoA frameworks assume there is a sort of 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. This should be taken into account in the LoA frameworks. Adoption of the web of trust model in these frameworks is one approach but could take a long time. Another approach is to define a proprietary web-of-trust based LoA framework and register it at the global IANA registry that has been setup for this purpose[63]. The web of trust based authentication levels of this framework could be mapped to those of other frameworks like STORK or ISO29115 to ensure interoperability.

What will be the highest LoA achievable with this approach?

[63] See http://levelofassurance.org/process.html for more information.

The highest achievable degree of confidence with web-of-trust authentication is STORK/ISO29115 level 3. Note that web-of-trust only refers to the registration process of authentication, hence should be coupled with appropriate issuing of the credential(s) and authentication procedures. While LoA 4 requires physical presence, web-of-trust is as least potentially as or more reliable than physical address, video proofing, bank accounts, and other LoA 3 registration methods. Of course, the actual outcome depends on the implementation, e.g. the algorithm for selecting Helpers should be secure enough.

How to deal with account revocation or a change of second factor authentication token? In particular, what to do if a person loses access to the additional factors used to authenticate them such as their phone or computer? Perpetrators will try to hijack accounts through account recovery systems, but this poses difficult challenges since the recovery systems have to help the real owner who has truly lost access to those other factors.

Web-of-trust authentication is similar to other registration processes regarding these maintenance issues, also it does not imply any type of issuing of credentials, so issues of deprovisioning and authentication fraud are the same as with other forms of registration (for example, self-asserted or over-the-counter). There are a few side notes however, if the authentication solution changes because of updates, misuse or revocation, re-issuing using the same verification process is likely more difficult since Helpers are confused to vouch again for the same user. In that case it is prudent that the reason for reissuing has to be reported to the Helpers.

8.3 Final verdict

Web-of-trust provides an interesting identity proving mechanism to be used in registration for authentication to attain LoA 2 or 3. It should be used in combination with authentication means of the desired level of assurance, if and when no more effective registration process may be implemented. For example, in case of an international collaboration where distance, language or poor electronic communication are barriers to proofing. ***The main issue in all cases lies with 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 constitutes an easy extension for a social networking website. Within the context of federations for higher education and research web-of-trust authentication would have merit if these conditions are met, for example in an international research collaboration. ***The concept of introducing your friends is intuitive, its implementation for digital authentication less straightforward.***

8.4 Future work

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 and pilots 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 extensive data mining and analytics skills.

- 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. For instance, is the user indeed a student or older than 18 years? Getting such attributes from authoritative sources is often not possible or expensive. The web of trust can offer a solution. Attribute validation challenges involve expiration of validations in time, rules for determining that an attribute has been sufficiently validated, and the communication of validation information to service providers such that they can handle it.

    - The use of web of trust for authorization purposes. For instance, helpers from the web of trust grant an asker access to certain resources. Research challenges include governance and administration of access rights in general and revocation in specific.

    - The use of web of trust for identity reconciliation. Reconciling or linking accounts is becoming increasingly popular for single sign-on purposes. Reconciling accounts usually involves the communication of a shared attribute that enables linking. But what if that shared attribute does not exist? The web of trust may be consulted in this case, i.e. to get certainty that both accounts belong to the same user.