Making authentication stronger and more cost efficient with web of trust

Bob Hulsebosch[1], Maarten Wegdam[1], Martijn Oostdijk[1], Joost van Dijk[2] & Remco Poortinga – van Wijnen[2]

[1] InnoValor, PO Box 321, 7500 AH, Enschede, the Netherlands
[2] SURFnet, PO Box 19035, 3501 DA, Utrecht, the Netherlands

e-mail: Bob.Hulsebosch@innovalor.nl, Martijn.Oostdijk@innovalor.nl, Maarten.Wegdam@innovalor.nl, Joost.vanDijk@surfnet.nl, Remco.Poortinga@surfnet.nl

Abstract: Solid registration processes for identity registration including proofing, vetting and binding are essential for strong authentication solutions. Solid typically implies a face-2-face component in the registration process, which is expensive and not user friendly. Alternatives that rely on remote registration often result in weak binding or are overly complex. We propose a web of trust approach in which users can indicate trust in the identity of other users. 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. This paper describes how to achieve web of trust enhanced authentication assurance.

1 Introduction

Service providers traditionally use the familiar username and password combination to authenticate users on their websites. Unfortunately, this approach provides a relatively low level of security for users: passwords can be easy to guess, too short, and difficult to manage. Adding a second factor, e.g., combining what a user knows with something he has, to the authentication process can help to address these issues. Commonly referred to as two-factor authentication, it adds additional security to authentication and raises the level of trust from the service provider to the user.

More and more service providers are beginning to rely on two-factor authentication solutions to stop escalating online fraud, identity theft and to comply with regulations. Many financial organisations such as banks and insurance companies have been using text message- or token-based authentication solutions for transaction verification for years, but recently major websites and businesses not in regulated industries are recognizing the need for stronger online authentication. Not so long ago, Google, Facebook and LinkedIn made two-factor authentication available to all users. The drawback of these two-factor solutions is that their binding to the user's identity is relatively weak. They only ensure with increased reliability that it is the same user, not who the user actually is. Binding an authentication solution to a user whose identity has been verified and registered is not trivial. It often requires physical presence and verification against authentic sources which is cumbersome and expensive.

This paper describes an approach for enhancing the authentication strength by using web of trust in a federated identity ecosystem. The idea is to use the web of trust concept 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 can provide extra confidence for the service provider to allow access to resources that require stronger authentication. Person C can also claim to know B and his authentication mechanism, thereby even further increasing the trust in the identity of B. This approach is a kind of "crowdsourcing of trust" about the identity of the user without requiring a physical registration.

The structure of the paper is as follows. Section 2 provides background information about strong authentication and web of trust. Several illustrative use cases are described in section 3. Based on these use cases the functional requirements for web of trust enhanced authentication are derived. Section 4 describes a protocol for leveraging web of trust for authentication enhancement. Implementation details are provided in section 5. The challenges are discussed in section 6. Finally, section 7 draws conclusions and describes ideas for future work.

2 Background

2.1 Strong authentication and Levels of Assurance

The strength of the entire authentication system is usually expressed in terms of levels of assurance (LoA). The LoA specifies the degree of confidence in identifying a user to whom the credential was issued, i.e. the combination of the strength of the authentication solution used and the quality of the registration process (see Figure 1). The combination of the two – stronger authentication and identity registration – is basically what is needed in order to achieve true strong authentication.

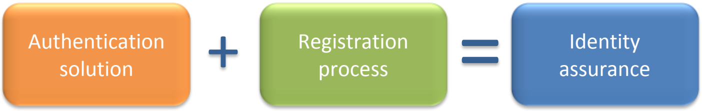

Figure 1: factors that determine the stength of the authentication.

There are several standards for the specification of LoAs. Examples are the ISO/IEC 29115 Entity authentication assurance framework [ISO13] and the STORK Quality Authentication Assurance framework [HLE09]. Both frameworks define four discrete assurance levels varying from almost no assurance in the user identity (LoA 1) to medium (LoA 2), high (LoA 3) and very high assurance (LoA 4).

The LoA paradigm allows service providers to specify assurance levels that correspond to the sensitivity or criticality of the service. Highly sensitive or critical services typically require a higher LoA. This means strong authentication solutions and robust registration processes.

Strong authentication solutions are available and typically consist of two-factor solutions (see e.g. [KUP10] for an overview).

The registration process by which a physical person is linked to his/her digital identity information and to his/her 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 identity, there can be 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 were denied registration. If the linking is weak, even the most complete personal information 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. Specifically, the requirements include measures that:

1. Increase proof in the identity of the user.
2. Increase trust in the binding between the user’s identity and his digital identity.
3. Increase trust in the binding between the user and a second authentication credential.

This authentication triangle of binding is illustrated in Figure 2 below.

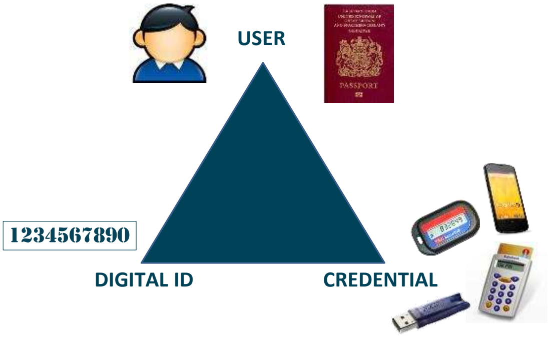

Figure 2: Binding triangle of user ID – digital ID – authentication credentials.

Different registration processes and mechanisms applied to identity vetting, proofing and credentialing result in different registration assurance levels. An applicant may appear in person to register, or the applicant may register remotely.

In-person registration is the most reliable identity proofing process during user registration. It is considered suitable for cases where there is a strong need to be able to determine that a service provider (e.g. a student information system) is dealing with a legitimate user, thus necessitating a stringent identity proofing process during user registration (i.e. a face-to-face process). 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. Other – less attractive and/or appropriate – alternatives (such as use of physical address, email & mobile phone, use of bank account) are discussed in [HUL11]. The STORK and ISO29115 frameworks require physical registration for LoA 4.

Remote registration is limited to levels 1 through 3 and is more vulnerable to threats and technically complex to achieve. Remote registration relies 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. Consultation of the latter source is restricted by legislation and not available for step-up authentication purposes; the HR-system on the other hand could be used as an alternative source. Typically, after a successful validation, a registration activation code is sent to the applicant’s home address. This is cumbersome and expensive.

2.2 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[1]) 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 whom not. Instead of trusting a single entity to validate identities, one validates the identities of the people one knows and exports this information to a public database. Then one relies on friends to vouch for the people they know, and those friends to vouch for still more people, and so on until a trust chain between any two arbitrary identities can be created. 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 must be built very much like a social networking site, because that is how people connect and share information, 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. Existing trust

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

infrastructures such as PGP, identity federation, social or professional networks can be readily used to enhance the registration part of the overall LoA. Particularly in the context of virtual collaboration organizations in which users know each other, 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 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.

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. 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.

3 Use case scenarios

The following use cases illustrate the use of web of trust for enhancing authentication.

3.1 Use case 1

A group of collaborating researchers from various institutions requires access to a highly sensitive database. 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 in to 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 at the Attestation Service for the authentication level and can decide based on the obtained information whether or not to grant her access to the database.

3.2 Use case 2

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, Bob logs in at 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 vet 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 can be raised as well by the Attestation Service. Next time Bob logs in with his LinkedIn account, the Attestation Service asserts that Bob has been authenticated with LoA 2.

3.3 Use case 3

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 asked them to attest for her identity at 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.

3.4 Analysis

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

- The need for an attestation service that facilitates and coordinates the enhancement of the authentication solution. Specific requirements for the attestation service are:
    - Determines identity of user;
    - Links social network accounts of users;
    - Selects suitable candidates from the social network that could attest;
    - Collects and validates attestations from the social network web of trust;
    - Determines the authentication strength;
    - Communicates the outcome to the service provider;
    - Optionally: Asks 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 a web 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.

3.5 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. 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 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. 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 [CTAID]).

4. Protocol description

We propose the following protocol for web of trust enhanced authentication:

***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 implictly 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 stength of the authentication of Asker eventually.

***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 and 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.

***Step 5**: LoA determination.* 1. 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:

  $$\text{New LoA} = \text{LoA} + \text{LoA}*(1 - (1 - H1)*(1 - H2)*(1 - H3)....)$$

  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 or Facebook)
    - Trust relations between Helpers
    - Number of paths between Helpers and Asker in PGP
    - Path length between Helper and Asker PGP[2]
    - 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.

The Attestation Service informs Asker about the new LoA via e-mail.

***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. The 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.

The different steps are illustrated in Figure 3 below.

[2] The ideal scenario in PGP key validation is to have multiple, short paths between the Asker and the anchors the Attestation Service trusts. This provides a strong guarantee that the Asker is indeed who he claims to be. The price, of course, is that it is more difficult to validate keys since the trust anchors must personally sign more keys than if fewer and longer paths are accepted.

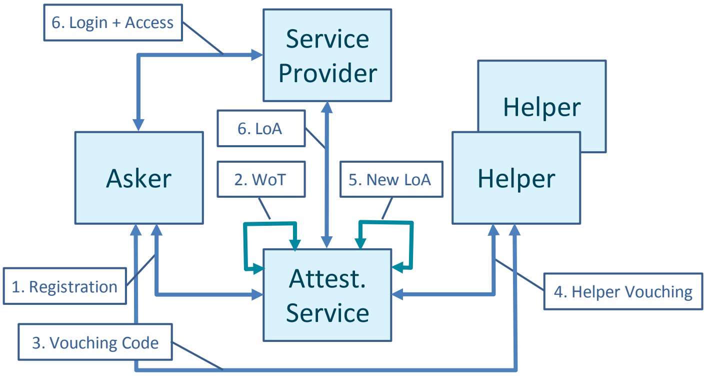

Figure 3: Web of trust protocol flow.

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 [BJR06].

5. Implementation

A proof-of-concept Attestation Service has been implemented. The Attestation Service is part of the test environment of SURFconext[3], the identity federation of higher education and research in the Netherlands. It allows the Asker to login with her federated account. 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 link the identity provider account to her social network account such as LinkedIn. This allows the Attestation Service to select Helpers from the LinkedIn web of trust of the Asker. The vouching code is alphanumeric and consists of five characters. Helpers kan login to the Attestation Service with their federated account. The algorithm for calculating the new LoA is relatively simple for the moment. It takes the number of Helpers into account, their authentication LoA that is provided during login, and the number of Helpers that did not vouch. The communication between the Attestation Service and the service provider for LoA validation is based on a RESTful API[4].

6. Discussion

This web of trust based LoA approach, however, raises several challenging questions that need to be addressed.

[3] SURFconext federation and collaboration infrastructure, see http://www.surf.nl/en/services-and-products/surfconext/index.html.
[4] Representational state transfer (REST), see http://en.wikipedia.org/wiki/Representational_state_transfer.

6.1 Weaknesses of web of trust approaches

One of the challenges is related to a number of weaknesses that are inherent to a web of trust approach. 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 behaviour, attacking of the underlying infrastructure and the exploitation of features of metrics used by the system to calculate the identity assurance [CH07]. Our proposal 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 select candidate Helpers and to determine the new LoA. 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 e.g. multiple LinkedIn accounts under fake identities and establish via these accounts a web of trust of LinkedIn connections. The requirement for Helpers to have a higher LoA than the Asker makes it more difficult to enhance the LoA via this approach. Given the potential weaknesses, the web of trust approach may not be suitable to achieve LoA 4 assurance, but we certainly see the potential to achieve LoA 3.

A potential improvement to traditional web of trust systems would revolve around reducing the validity period of the claims made by other users regarding a specific user account and to allow for automatic prolongation of the trust-based claims associated to the account by subsequent authentication sessions. This would allow for both verification of use of the account and the identity associated to it and user revocation of 'stale' or otherwise undesired credentials. During the refresh process, the user can choose whether to continue to continue or stop endorsing others' accounts; this helps the dynamics of the web by helping to cull out untrusted persons more rapidly.

Further, providing the option for anti-claims, to specifically call out an account as untrusted to others, would significantly mitigate the effect of malicious persons such as spammers gaining access to a web of trust. Allowing for this anti-measure could also form the basis of a sliding trust scale, with trust and anti-trust counting against each other and allowing for unconnected persons to see that a particular account may or may not be trustworthy. Paths connecting persons would be deprecated by paths containing anti-claims; determining whether or not to trust someone with a significant number of anti-claims would be assisted by allowing short comments with them similar to twitter messages (i.e. "this person is a spammer" or "this person is a liar").

6.2 Calculating Levels of Assurance

Various approaches to calculate reputation values exist, see [Ne11] for an overview. 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.

- 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.
- Bayesian frameworks: Reputation models based on Bayesian frameworks depict reputation values as probabilities between [0,1]. These models are popular for peer-to-peer networks and sensor systems, rely on ratings being either positive or negative, and use probability distributions for reputation scores.

Since 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. Calculating trust from social network aggregation is not new [SAN07], [HEI13]. These approaches, however, 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.

Another challege 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.

6.3 Relation to existing LoA frameworks

Closely related to the previous challenge is another one: 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 framework 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[5].

7. Conclusions

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

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

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.

We propose the use of web of trust to enhance the registration part of the overall authentication process. The web of trust approach replaces a physical registration at an authority by outsourcing the actual identity proofing and vetting to a user community. The outcome of the vetting allows for enhancement of the authentication assurance level. Due to various weaknesses of the web of trust model and challenges related to the determination of the actual LoA and due to liability issues, the proposed approach is unlikely to achieve LoA 4 assurance but LoA 3 seems feasible and is suitable for most online services. Future work will involve further optimisation of the algorithms 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.

Acknoledgements

This work is sponsored by the Géant3plus Open Calls program.

References

[BJR06] 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.

[CH07] Carrara, E.; Hogben, G.: Reputation-based Systems: a security analysis, ENISA position paper, October 2007.

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

[HEI13] Heisnam, R.S.; Neelima, A.; Singh, L.S.; Singh, S.I.: A Model of Computing Trust in Web Based Social Network Using New Aggregation and Concatenation Operators, Int. Journal of Computer Science and Network, Volume 2, Issue 4, August 2013.

[HLE09] Hulsebosch, B.; Lenzini, G.; Eertink, H,: 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.

[HUL11] Hulsebosch R. J.: Step-up Authentication-as-a-Service, SURFconext Adoption, 2011.

[ISO13] ISO/IEC 29115:2013 Entity authentication assurance framework, see http://www.iso.org/.

[KUP10] Kuppinger Cole: Market Overview Strong Authentication, 2010, see http://www.kuppingercole.com/report/srmo_stronauth_80310.

[Ne11] 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.

[SAN07] Noh, S.: Calculating trust using aggregation rules in social networks, in Proceedings of the 4th international conference on Autonomic and Trusted Computing, pages 361-371, 2007.