

ENISA Position Paper No. 2

Reputation-based Systems: a security analysis

Editors: Elisabetta Carrara and Giles Hogben, ENISA

October 2007





Reputation-based Systems

ENISA Position Papers represent expert opinion on topics ENISA considers to be important emerging risks or key security components. They are produced as the result of discussion among a group of experts who were selected for their knowledge in the area. The content was collected via wiki, mailing list and telephone conferences and edited by ENISA.

This paper aims to provide a useful introduction to security issues affecting Reputation-based Systems by identifying a number of possible threats and attacks, highlighting the security requirements that should be fulfilled by these systems and providing recommendations for action and best practices to reduce the security risks to users.

Examples are given from a number of providers throughout the paper. These should be taken as examples only and there is no intention to single out a specific provider for criticism or praise. The examples provided are not necessarily those most representative or important, nor is the aim of this paper to conduct any kind of market survey, as there might be other providers which are not mentioned here and nonetheless are equally or more representative of the market.

Audience

This paper is aimed at providers, designers, research and standardisation communities, government policy-makers and businesses.

Reputation-based Systems

EXECUTIVE SUMMARY	3
INTRODUCTION TO REPUTATION-BASED SYSTEMS	5
USE-CASES	8
PRINCIPAL THREATS	12
SECURITY REQUIREMENTS	18
RECOMMENDATIONS	20
CONCLUDING REMARKS	24
CONTRIBUTORS	25
REFERENCES	26

Reputation allows users to form an expectation of behaviour based on the judgements of others, bringing the significant economic and social benefit of being able to trust others who are not known directly. Reputation can encourage good behaviour, as users seek good reputation and benefit from it. It follows that electronic reputation is becoming as valuable an asset as traditional reputation. As new applications embrace reputation-based systems, the value of online reputation will continue to grow – and will become increasingly the target of attacks.

This paper explains the main characteristics of electronic reputation systems and the security-related benefits they can bring when used within applications and services. Four main use-cases are described: online markets, peer-to-peer networks, anti-spam techniques and public key authentication (web-of-trust). From these, the main threats and attacks against reputation systems are derived, as well as the security requirements for system design. This leads to a set of core recommendations for best practices in the use of reputation systems.

Threats and Attacks

Threat Rep. 1 – Whitewashing attack: the attacker resets a poor reputation by rejoining the system with a new identity. Systems that allow for easy change of identity and easy use of new pseudonyms are vulnerable to this attack.

Threat Rep. 2 – Sybil attack (i.e. pseudospoofing): the attacker creates multiple identities (sybils) and exploits them in order to manipulate a reputation score.

Threat Rep. 3 – Impersonation and reputation theft: one entity acquires the identity of another entity (masquerades) and consequently steals her reputation.

Threat Rep. 4 – Bootstrap issues and related threats: the initial reputation value given to a newcomer may lay it open to attacks such as sybils and whitewashing.

Threat Rep. 5 – Extortion: co-ordinated campaigns aimed at blackmail by damaging reputation for malicious motives.

Threat Rep. 6 – Denial-of-reputation: attack designed to damage an entity's reputation (e.g. in

combination with a sybil attack or impersonation) and create an opportunity for blackmail in order to have the reputation cleaned.

Threat Rep. 7 – Ballot stuffing and bad mouthing: reporting of a false reputation score; the attackers (distinct or sybils) collude to give positive/negative feedback, to increase or lower a reputation.

Threat Rep. 8 – Collusion: multiple users conspire (collude) to influence a given reputation.

Threat Rep. 9 – Repudiation of data and repudiation of transaction: an entity can deny that a transaction happened, or the existence of data for which he was responsible.

Threat Rep. 10 – Recommender dishonesty: the voter is not trustworthy in his scoring.

Threat Rep. 11 – Privacy threats for voters and reputation owners: for example, anonymity improves the accuracy of votes.

Threat Rep. 12 – Social threats: Discriminatory behaviour is possible when, for example, in a second-order reputation system, an entity can choose to co-operate only with peers who have a high reputation, so that their recommendations weigh more heavily. Other possible social threats include the risk of herd behaviour and the penalisation of innovative, controversial opinions, and vocal minority effect.

Threat Rep. 13 – Threats to the underlying networks: the reputation system can be attacked by targeting the underlying infrastructure; for example, the reputation information can be manipulated/replayed/disclosed both when stored and when transported, or may be made unavailable by a denial of service attack.

Threat Rep. 14 – Trust topology threats: an attack targets certain links to have maximum effect, for example those entities with the highest reputation.

Threat Rep. 15 – Threats to ratings: there is a whole range of threats to reputation ratings which exploit features of metrics used by the system to calculate the aggregate reputation rating from the single scores.

Recommendations

Rec. Rep. 1 – Perform a threat analysis of the reputation system: before designing or adopting a reputation system, a threat analysis should be performed, and security requirements should be identified.

Rec. Rep. 2 – Develop reputation systems which respect privacy requirements: a more privacy-respecting design of reputation systems is needed without compromising the trust provided by the system.

Rec. Rep. 3 – Provide open descriptions of metrics: the reputation metrics should not be kept secret.

Rec. Rep. 4 – User-interface recommendations:

a. Recommendations for the usability of reputation-based systems: to create a sense of trust, it is important that the user understands the way trust is measured.

b. Differentiation by attribute and individualisation as to how the reputation is presented: a given reputation system should allow a user to customise the reputation according to different attributes (i.e. different aspects/assertions about the reputation subject), and to set a threshold for each of them.

c. A combination of quantitative and qualitative assessments should be offered (instead of just plain score numbers), wherever the application allows.

Rec. Rep. 5 – Promote awareness-raising: users should develop skills in understanding reputation ratings and the trust processes behind them, and developers should be made aware of attacks on reputation-based systems.

Rec. Rep. 6 – Encourage the use of reputation systems in online Social Networking sites: apply reputation techniques within online Social Networks – reputation mechanisms can act as a positive motivator towards good online behaviour.

Rec. Rep. 7 – Encourage research into:

a. Understanding the social components behind reputation systems: there is a need to investigate the social aspects and subtleties that influence reputation models, as several identified threats have social origins.

b. New authentication mechanisms: enhanced authentication mechanisms should be developed as a countermeasure against

attacks such as reputation theft, whitewashing and various automatic attacks.

c. Common solutions to threats against reputation-based systems: despite the variety of use-cases for reputation-based systems, they are often vulnerable to similar threats and attacks – common solutions to defeat these should be investigated.

d. The management of global reputation: how can a user gain control and/or awareness of his overall electronic reputation when it is composed of fragments scattered across the Internet?

e. Anti-phishing tools based on reputation: reputation systems can be used to improve the ability to detect phishing, for example, via toolbars – more research is needed to improve the accuracy of these tools.

f. Use of weightings in the metric: use of appropriate weighting systems, i.e. different weightings to improve the resistance of the metric to attacks (while at the same time maintaining the requirement for transparency in trust-building mechanisms).

Rec. Rep. 8 – Research into and standardisation of portable reputation systems: one possibility is to integrate reputation into authentication transport standards, e.g. OASIS Security Assertion Markup Language (SAML) Authentication Context. As a means of providing evidence for identity claims (e.g. creditworthiness, trustworthiness as a seller), reputation is well suited to be transported in a similar format to, for example, SAML assertions.

Rec. Rep. 9 – The importance of automated reputation systems for e-Government: reputation is informally already an important component of several high-assurance systems such as document issuance and security clearance processes, and automatic ad hoc reputation systems provide scalability and flexibility which are not present in existing systems (such as Public Key Infrastructure (PKI) systems, for example). Policy-makers are therefore encouraged to investigate the possibility of using state-of-the-art reputation-based systems in e-Government systems. Governments should also start investigating the impact of online reputation systems on existing legislation such as the EU privacy directives.

Rec. Rep. 10 – SMEs should embrace the potential of reputation systems: reputation systems can improve competitiveness and customer confidence, and are often easier to understand, implement and manage (than, for example, PKI with certificates).

Introduction to Reputation-based Systems

There is an increasing number of applications which make use of reputation-based systems as risk management mechanisms to facilitate trust.

Reputation¹ is the aggregated opinion that people in general have about someone or something, or how much respect or admiration someone or something receives, based on past behaviour or character [1]. Reputation embeds an expectation about an entity's² behaviour based on information about or observations of its history [2].

Trust and reputation are strongly related: reputation enables trust. Trust typically exists on a personal level, where the personal opinion weighs more than the opinion of others (for example, you can trust someone with a low reputation rating). Reputation expresses the collective opinion, leading to trust or distrust, that emerges as a result of opinions of members of a certain community. Hence, in reputation the 'social context' component is very important in analysing security implications. Reputation is a proxy for acquaintanceship [3]. It is a social-psychological mechanism that is used as the input to a heuristic (rule of thumb/intuitive) decision about trustability, a scalable way to evaluate relationship/transaction risk within an extended community.

There is in fact a significant economic benefit in being able to trust people you do not know personally (i.e. strangers). Opinions we use to build trust in strangers often come from other strangers since online contacts are more numerous, more distributed and more often anonymous than offline contacts. While this makes reputation sound like something unreliable, in reality the majority of people providing reputation feedback, do so honestly [4] [5].

Furthermore, there is evidence to suggest that, under the right circumstances, group decision-making can arrive at surprisingly good quality results³. Reputation in the real world is used to predict the actions of a person or organisation,

hence lowering the risk involved in the trust decision. This same trust assessment and management mechanism is mirrored in the digital world using online reputation systems, which offer aids in making a decision about whether to trust something or someone in a way that is similar to how we interact in society. Therefore, it is more easily understood by users, and triggers higher confidence.

Reputation is typically considered as belonging to soft security. Rasmusson and Jansson [7] suggest the use of social controls to create secure open systems, which do not rely on global authority but on the participants interacting among themselves. While hard security (e.g. typical authentication based on crypto) does not allow attackers to enter but fails once it is bypassed by them, a soft security approach, such as one using reputation, admits entities and allows them to act until they misbehave. Reputation can encourage good behaviour (reputation is valuable, and a user will generally behave well to maintain his good reputation). In addition, the fact that bad behaviour is advertised to the network is a strong discouragement to misbehave. Of course, this approach can also be subject to threats – as we will see shortly.

It is clear therefore that online reputation is a valuable asset, in exactly the same way as a real life reputation. As this study's findings demonstrate, reputation-based systems are emerging as an important risk management mechanism in the cyberspace. As new applications embrace them, the value of online reputation will continue to increase. While this will encourage honest behaviour in order to benefit from a high reputation, at the same time it makes reputation an asset which appeals to dishonest entities, and hence a target for attack. This paper aims to highlight the main threats to and attacks against reputation systems, as well as to identify the security requirements which should be considered in the design and usage of systems where the trust mechanisms make use of reputation.

¹ *Etymology: from Latin reputare, to reckon, think over*

² *In this paper, "entity" refers to a user, a machine, an organisation etc. as owner of the reputation or as voter on the reputation of another entity, depending on the context specified in the text.*

³ *Surowieki [6] describes how experiments have shown that some group decisions are "smarter than the smartest people in the group": under the right conditions, if you ask a large enough group of diverse, independent people to make a prediction or estimate a probability, and then average those estimates, the errors each of them makes in producing an answer will cancel each other out. For example, the sociologist Kate H. Gordon asked two hundred students to rank items by weight, and found that the group's 'estimate' was 94% accurate, which was better than all but five of the individual guesses. However the risk of herd behaviour and the penalisation of innovative, controversial opinions also exist (see chapter 'Principal Threats').*

Characteristics of Reputation-based Systems

There are several reputation-based systems available [8]. The following are some of their distinctive characteristics:

Goals of reputation-based systems

The main goals/reasons for reputation systems to be used in a network of entities interacting with each other are:

- to provide information to help assess whether an entity is trustworthy (trust assessment)
- to encourage entities to behave in a trustworthy manner, i.e. to encourage good behaviour
- to discourage untrustworthy entities from participating in an interaction [9]
- to help seek out new knowledge and resources by querying trusted entities as sources (discovery).

Formation and scope of reputation-based systems

Reputation is an opinion formed on the basis of *aggregated information*. Reputation systems aggregate the experience of the community to offer entities assistance in forming an opinion before interacting with other entities. This aggregated information typically includes the history of the entity (past behaviour) that is reported by:

- direct knowledge: the direct opinion of the assessing entity, when available, for example formed from previous transactions, or from direct observations of certain factors
- indirect knowledge: the opinions that others hold about the given entity.

Reputation is *context-dependent*: for example, someone with a good reputation for selling laptops may have a poor reputation as a car salesman. In other words, a good reputation in one context in general does not assert anything about the reputation of the same entity in another context – although in some cases it may be accepted as an indication of trustworthiness. It follows that reputation needs to be confined to the profile/context to which it is applied (the laptop salesman versus the car salesman). However, similarities between contexts/profiles may be defined to enable sensible re-use between similar contexts.

The fundamental raw material of a reputation is a set of votes or opinions on the attributes of an

entity (e.g. is he/she a trusted seller?, is x the public key of y?, how well the transaction went, how much bandwidth is consumed by that entity etc.). These are then combined using more or less sophisticated algorithms to produce an aggregate reputation score. The ways in which the votes are combined to calculate the aggregate score (the metric algorithm) are typically different according to the application. The chosen metric can be simple (e.g. the mean, concatenation or summation of individual scores) or complex (for example, it may use weightings according to how much the recommender itself is trusted by the community as a whole or by the one using the reputation, or according to the age of the vote). The more complex the metric, the more difficult it may be for a user of the system to understand the actual meaning of the reputation score. As we will see from the identified threats, the metric itself can be a point of failure of the system, since it can be subject to attacks. A simple, immediately understandable metric may not be appropriate to counter attacks that are specific to that application.

The *incentive/punishment scheme* is an important component for a reputation-based system. While incentives and rewards encourage an entity to behave in a trustworthy manner in order to acquire, maintain or increase its reputation and to benefit from it, appropriate punishment (from lowering a person's reputation through to banning him from a community) is needed to discourage misbehaviour including whitewashing of reputation (unfair escape from a bad reputation, see the chapter 'Principal Threats'). To maintain an attack-resilient reputation system, in general the value of the punishment has to be higher than the potential gain from cheating.

Models for assessing reputation

There are three main models for assessing reputation: subjective, objective and hybrid.

- **Subjective reputation systems**⁴ rely mainly on scores provided within a controlled community of users – where the community has well defined purposes, such as selling goods, sharing content, describing the experience of users, sharing knowledge or opinions. The community service provider plays a key role in running the reputation system, and both the provider and the reputation system have a reputation. Users rank 'subjective' reputation scores for other

⁴ Examples of these systems and related community providers include Amazon Reputation Scoring, eBay Reputation Scoring and IMDB (Internet Movie DataBase).

community users, based on their personal interactions and/or the fulfilment of personal aspirations and promises/agreements, which may also vary in time. So, for example, one user may consider an e-mail interesting that others considered spam, and he can even change his mind in the course of the day.

- **Objective reputation systems**⁵ rely on the community service provider or the reputation system to provide factual evidence of the ranking, based on well defined metrics and repeatable criteria. The individual can apply some kind of subjective analysis to this raw data to obtain his personal assessment of reputation (and share it with other users). However the core 'reputation measures' are based on objective evidence, observable by the whole community (e.g. reports, analysis-based scientific metrics and criteria, average bit rate of a video streaming server).
- **Hybrid reputation systems**⁶ are a combination of subjective and objective systems. Usually these systems are based on an objective reputation system, where the factual results are interpreted on the basis of subjective and personal values or motivations.

Main Structure Models of Reputation-based Systems

Reputation systems typically show one of the two following main structure models (a compromise between these two extremes is also possible):

- **Centralised model:** a central authority collects reputation scores (from other entities and using other sources such as its own observation), typically processes them to form an aggregated reputation score for a given entity, and then redistributes this reputation score for use by other entities. Online trading and market communities (see the chapter on 'Use-cases') use this model.
- **Decentralised model:** the entities participating in the community share the reputation information, without the need for a central repository. This model is more suitable for networks that are decentralised by nature, such as peer-to-peer and autonomic systems (see the chapter on 'Use-cases'). It also allows peers to assign different trust values to different sources of reputation scores.

⁵ Examples include Amazon book sales (based on actual sales figures), reports on the punctuality of train operators, the ranking of the performance of companies, universities etc. based on solid criteria such as financial criteria and examination statistics.

⁶ Examples include personal rating and the advertising of books, movies, games etc. – independently from the actual 'objective' sales figures.

Use-cases

Online markets (such as the popular eBay and Amazon) are among the best known examples of applications which make use of reputation-based systems, but are far from being the only ones where such systems can be used. In the following, we describe four use-cases which make use of reputation systems; they were chosen as a means to derive the threats and the security requirements identified in the remainder of this paper. The four selected use-cases are: online markets, peer-to-peer networks, anti-spam techniques and public key authentication (web-of-trust).

Online Markets

The members of an online market are allowed to sell and buy arbitrary items within the community⁷. One of the largest providers is eBay [10] with about 82 million users worldwide at the end of 2006 [11]. The liability for trades within the community is typically delegated to the members involved. After an item has been sold, the seller and buyer have to exchange the item purchased and the payment in a fair way. Most items are physical goods that can be exchanged either directly between the users, or via a transfer service that is offered by many providers for a charge. The first option is common for low-price items such as books or CDs; buyers and sellers exchange money and item directly (by bank transfer and conventional mail).

In most cases of online peer-to-peer market, both the seller and the buyer are exposed to some degree of risk. The risk for the buyer is that the item is not compliant with the sale description, or that the item is not properly delivered despite the payment; the risk for the seller is that his item is sent to the (unknown) buyer but the payment is not good. Many of the exchanges are successful; however, 44,9% of the 207.492 Internet frauds reported to the Crime Complaint Center (I3C) in the US in 2006 were cases of Internet auction fraud [12]. This can hamper the further development of online markets; reputation can help mitigate these risks.

After every transaction, users may give comments and/or scores to each other; these scores are added to the users' reputations, which can then be consulted during other transactions, giving the involved peers a way to assess the risk they are taking and, overall, increasing confidence in the online market⁸.

Peer-to-peer Networks

A peer-to-peer (p2p) network is a network where connected nodes serve as client/server, sharing different types of resources⁹. Peer-to-peer networks are mainly based on two possible architectures:

- a pure distributed p2p network with no mediation by a central entity, such as Bluetooth connections between mobile devices
- a distributed p2p network relying on a centralised (hub-based) infrastructure to 'bootstrap' interactions and/or provide some services (e.g. Instant Messaging (IM) where the list of participants is stored in a central directory), although interactions are peer-to-peer. In this context data storage and processing are local to peers but part of the communication infrastructure is centralised.

Peer-to-peer networks are populated by a set of systems, each potentially managed by a different administrator, implying that different security policies, configurations and settings might be in place. In addition, the 'individual' practices (in terms of security, privacy and personal habits such as what is stored or browsed on the web) might vary. The capabilities of the constituent platforms (laptops, PDAs, mobile phones, servers/desktops etc.) might also differ, limiting the security options that are available and the assumptions that can be made by a remote peer. From the security point of view, each entity belonging to the p2p network cannot be seen as universally trustworthy.

In such open, dynamic environments, reputation-based systems can provide added value in solving a number of security issues, starting from the initial one of granting trust to peers. Other common issues that reputation-based

⁷ Note that these members are largely unknown to each other and sometimes live in different countries with different legislative systems.

⁸ Note that there is a risk that expressed scores and comments are unrealistically positive since users have a disincentive to leave negative feedback, for fear of retaliation (see chapter on 'Principal Threats').

⁹ Peer-to-peer networks include different types of applications, for example file-sharing, PGP/key management systems, Instant Messaging, point-to-point mobile device connections (via e.g. Bluetooth, infrared) and distributed processing. Other applications such as sensor networks, mobile ad hoc network (MANET), ad hoc networks and desktop grids have similarities to p2p networks, as they are by nature dynamic, heterogeneous and distributed.

systems can help solve are pollution¹⁰ and freeloading¹¹ in p2p file-sharing. Traditional reputation relies on explicit feedback; reputation for p2p can also infer information by analysing attributes, such as of the underlying network (subjective and objective reputation). Decentralised reputation often relies on 'referral networks', where each entity receives reputation information from neighbours, weighting the reputation score according to its trust in this neighbour [14]. Hence, p2p exhibits high potential for non-linearity in reputation formation (weightings based on reputation to n degrees). The effectiveness of the reputation system in this use-case depends heavily on the size of the group and the relationships and cultures involved; the automated collection of data input and even the formation of opinions are needed. Current approaches to help overcome the selfish behaviour of nodes use game-theoretical analysis, which looks at strategies, utility and co-operation theory [15].

The p2p reputation system might be influenced by the underlying p2p model, i.e. whether or not this model relies on a central bootstrapping infrastructure. If this is missing, a peer has to formulate an opinion about the other peers; it potentially runs a local instance of a 'reputation system' which might share reputation information that could be the result of interactions and the sharing of information with other peers or might just be locally generated [16]. The central bootstrapping infrastructure, where present, might provide some factual information (such as the time of connection of a peer system on the network, availability, accessibility or the number of shared files) that could be used by peers to bootstrap the reputation of other peers, at least in terms of their commitment to participate fairly in the p2p community. Peers might use their experience gained whilst interacting with other peers to create 'circles of trust', where some degree of reputation and trust is assumed based on shared values and behaviour.

Anti-spam Techniques

The latest e-mail anti-spam (unsolicited mails) techniques¹² leverage reputation to improve the overall filtering efficiency, as reputation-based mechanisms offer a dynamic way to predict whether a sender is trustworthy or the message is spam. Evolving from traditional blacklisting/whitelisting, the trend today is to extend the formulation of reputation to the observation of several parameters (see below), in order to enhance the filtering capability. Fast reaction time to observed misbehaviour is part of the added value of using filtering based on reputation systems. This is particularly valuable in the face of the increasing use of zombies as agents of spam, which are unpredictable and operate for short times; the source of the attack can change continuously, making the attack harder to stop using traditional blacklists.

There are several reputation-based anti-spam solutions available on the market, each looking at a different set of properties over which reputation is calculated [17] (as well as research papers exploiting reputation network analysis for e-mail filtering, such as Golbeck and Hendler [18]). Essentially, these solutions look at the sender's behaviour over time to predict its trustworthiness; the sender is identified [19] by, for example, an IP address, a verified sender domain, the message itself or even Autonomous Systems, hence providing different levels of filtering granularity (e.g. machine vs. domain). Available reputation-based solutions can use reactive information (e.g. participants, often in the context of a community or social network, reporting misbehaviours and complaints; spamtrap data), predictive information (e.g. predicting by observing the behaviour), or a mixture of those.

¹⁰ Pollution in p2p file-sharing refers to downloaded content that may be of poor quality, something other than what was expected, or even harmful, as with malware. To resolve pollution, Walsh and Siner [13] suggest the use of the reputation of objects (such as a file) instead of peer reputation.

¹¹ Freeloading in p2p file-sharing refers to peers not contributing to a fair sharing of resources and free-riding, i.e. consuming content without contributing.

¹² ENISA has conducted surveys on the issue of spam in Europe and on security measures that providers use to mitigate it. See www.enisa.europa.eu/pages/spam/index.htm.

Use-cases

Parameters under observation and collected may include:

- volume of e-mail traffic
- type of traffic (e.g. continuous vs. sporadic)
- compliance with regulations (e.g. the CAN-SPAM in US)
- invalid or malformed e-mail parameters
- response to unsubscribe requests
- time the sender was observed the first time
- feedback from spamtraps
- feedback from user reporting
- content screening.

This aggregated information, collected over time (history) forms the reputation of the sender; a low reputation rating may lead to the message being judged as spam or the IP address or domain being filtered. Reputation mechanisms that observe the sender are typically based on a central entity monitoring some of the parameters listed above; the participants query the central node for the reputation. As an alternative to the central model, a distributed reputation system involves the participants sharing the reputation information with their neighbours, and then calculating the reputation score locally.

Reputation systems can be used in conjunction with mechanisms which try to assess the sender's identity; generally this use-case has poor authentication of subjects. Once the identity (which may be just the sender's domain, such as with SPF [20] and DKIM [21]) has been assessed, reputation systems can be used to predict the sender's behaviour using history and the recommendations of trusted parties. The granularity of the reputation is a factor to be considered: low-grained reputation systems may harm others as well as the attacker – if the bad behaviour of one sender is taken into account for the whole domain or provider, then the other senders associated with that provider will also experience loss of reputation.

Public Key Authentication (Web-of-trust)

One of the most important problems of public key cryptography is the authentication of the public key, i.e. how do you know that a public key that says “This is Alice's key” actually belongs to Alice? Unless the owner of the public/private key pair physically hands her public key over, there is no way of telling just by looking at the public key. To manage this problem, Pretty Good Privacy (PGP) uses the concept of an ‘introducer’ or ‘trusted third party’

whereby an intermediary, trusted by the prospective user of the public key, vouches for the validity of the key ‘belonging to Alice’. In contrast with a Public Key Infrastructure (PKI), which is a system to support the third-party introduction of public keys, via a hierarchy of Certificates Authorities, PGP uses a *web-of-trust*.

It should be stressed that web-of-trust is listed here as a reputation-based system because in this paper a reputation-based system is defined in broad terms, as a system that aggregates subjective (and objective, where possible) opinions as to the validity of an assertion (e.g. Bob is a trusted seller, the IP address that Eve is using is a spammer, or x is Bob's public key). In other words, saying that web-of-trust is a reputation system does not imply that web-of-trust makes any assertions as to the trustworthiness of the person. In web-of-trust, Alice can vouch that the key ‘Pub(Bob)’ is Bob's public key by digitally signing it with her own private key. If Alice's friend Cathy trusts Alice sufficiently as an introducer, Cathy may now also believe that Pub(Bob) belongs to Bob, after verifying Alice's signature. In PGP, Cathy may have ‘complete’ or ‘marginal’ (i.e. partial) trust in Alice as an introducer. PGP also allows each public key to have more than one introducer. As a result, Cathy can then form rules that go something like: “I will accept a public key as valid only if it has been introduced by at least a) one completely trusted introducer or b) three partially trusted introducers”.

PGP allows the formation of such rules and automates the processing of key validity using those rules. As such, PGP uses a reputation system since it aggregates a number of subjective opinions to form trust. If we consider a reputation system as a network of entities that share and evaluate each other's opinions, then we can assume that PGP implements a decentralised reputation system. It has the following characteristics:

- Entities are identified by their e-mail addresses
- Signing a key represents the signer's opinion that the key being signed is valid, i.e. it is really owned by the person with the claimed e-mail address
- Assigning an introducer trust level to a public key asserts trust in that key's owner as an introducer, although this information is private in PGP (this would correspond to second-order reputation)
- The context is the key introducer, with certificate depth and domain constrained parameters.

Other Use-cases for Reputation Systems

There are several other applications where reputation systems are used or proposed as risk management mechanisms [22]; these include the following examples (Baker and Hartrell list others [23]):

- Reputation can aid in quality assessment – for example for content generated in collaborative networks (such as Wikipedia [24] and Slashdot [25]), where there is a need to control the quality of contributions, and in open software communities, which open the code for the analysis of implementers at large and show the reputation of the code writer as an indication of how trustable the piece of code can be considered. Search engines use reputation-based criteria in returning the search results – for example in Google the displayed sequence of the results depends on the number of references to a given page by other reputable pages [26].
- Several applications consider reputation in an effort to solve their resource allocation problems. In collaborative networks such as scientific communities, multiple participants collaborate to solve a single large problem (e.g. Seti@home [27]), and require trust in the chosen participants, in their skills and in their effective willingness and ability to participate. Grid networks, in particular desktop grids (resembling pervasive computing), are turning to (typically decentralised) reputation systems to address the problem of resource allocation with untrusted participants. The issue of ‘sabotage tolerance’ in desktop grids [28], following the master-worker model, sees workers trying to sabotage the computation work to acquire more credits. The typical countermeasure is to duplicate the task among workers and accept the result if the majority of the results from the assigned workers are the same. However, this solution involves wasting computational resources; reputation can be used to select trusted workers whose results will be accepted without the need for redundancy.

Principal Threats

Just as reputation can be attacked in real life, so too can online reputation. The attacker can exploit the community's resources for free or with low cost, for example by increasing his own reputation in an unfair way or by impersonating another peer with a high reputation rating. He can attack another entity directly, with the aim of reducing or even destroying his competitor's reputation. Attacks can also be perpetuated against the community as a whole, damaging its reputation with other communities (e.g. in business competition) [29]. Whatever the reason behind it, the attack must be more advantageous than the gain from building up a good reputation.

During the analysis of the four selected use-cases listed above, a number of threats against reputation systems, and possible related attacks, were discussed and have been categorised below. Listed in brackets below the threat are the use-cases where these attacks are most likely to occur (the absence of a use-case name does not mean that it is immune to that threat/attack). The threats and attacks are often common to all use-cases. So too are the remedies.

Threats Related to Identity and Identity Changes

A number of threats are linked to identity changes. If the system allows for easy identity change, then it is also easy for an attacker to remedy a bad reputation by switching to another identity.

Threat Rep. 1 – Whitewashing Attack

[Use-cases: online markets, p2p networks, anti-spam techniques]

In the whitewashing attack, an attacker rids himself of a bad reputation by rejoining the system with a new identity. Systems that allow for the easy change of identity and easy use of new pseudonyms are vulnerable to this attack (low cost identities or cheap pseudonyms [30]). Indeed Internet culture and protocols have several features, including anonymous interaction and the ability to be untraceable, which can favour identity change. The attack can leverage a sybil attack (see page 13), where the attacker exploits multiple identities, and is also linked to the bootstrap issue (see below). In general, if the reputation rating acquired starts low and an effort is required to build up reputation, then an entity will be encouraged to maintain a more persistent identity. In one

Threats against reputation systems and possible attacks

- Threats related to identity and identity changes
 - Whitewashing attack
 - Sybil attack
 - Impersonation and reputation theft
- Bootstrap issues and related threats
- Extortion
- Denial-of-reputation
- Unfair rating
 - Ballot stuffing and bad mouthing
 - Collusion
 - Repudiation of data and repudiation of transaction
 - Recommender's dishonesty
- Privacy threats
 - Privacy threats for voters and reputation owners
- Social threats
 - Risk of herd behaviour and penalisation of innovative, controversial opinions
 - Vocal Minority Effect
 - Discriminatory behaviour
- Threats via the underlying network
 - Attacks to the underlying networks
 - Trust topology attacks
- Threats to ratings
 - Attacks to link-based ranking algorithms

version of this attack (sometimes called the 'strike and rechange' attack), an entity behaves properly for a time and builds a strong reputation, then suddenly uses this reputation to misbehave. This attack is used, for example, in the online market use-case, when the entity builds up a good reputation by performing a number of low-value transactions well, and then misbehaves in a very high-value one. One possible countermeasure is to apply weightings to the reputation scores according to the size or value of the transaction. However this has the disadvantage that entities are discouraged from engaging in low-value transactions or behaving honestly in them. Alternatively, different reputation scores can be accumulated for different categories of transactions [31].

Principal Threats

Threat Rep. 2 – Sybil Attack

[Use-cases: online markets, p2p networks, anti-spam techniques, web-of-trust]

In the sybil attack¹³, which is also known as pseudospoofing, the attacker creates multiple identities (sybils) and exploits them in order to manipulate a reputation score [32]. For example, multiple identities can be used to provide positive reputation feedback to one designated identity, whose reputation increases in an untruthful way (ballot stuffing, i.e. reputation fraud).

By mounting a sybil attack in the anti-spam use-case, the attacker may sign up for multiple accounts, which can be used to launch spam and then be disposed of, or to improve a given reputation and then spam from there. Alternatively he might mount a whitewashing attack, or leave unfair ratings (ballot stuffing, bad mouthing).

While in centralised environments, a one-to-one correspondence between a real entity and an identity may be ensured by the central authority, the sybil attack appears a difficult issue affecting reputation systems in large-scale decentralised environments such as peer-to-peer networks. The low degree of difficulty in creating sybils is the main vulnerability exploited by this attack, coupled with deficiencies in how the entities leaving feedback, especially those without direct trust by the peer, are treated (e.g. equally). Existing countermeasures against the sybil attack (also used against the whitewashing attack) aim to make the attack unprofitable [33] [34]. To create a new identity (bootstrap), a ‘price’ or ‘entry fee’ [30] could be requested, e.g. a requirement to engage in a time-consuming operation such as resolving a crypto puzzle or downloading a big file. eBay sets an entry fee for sellers, so as to make a change of identity undesirable. To decrease the likelihood of scores sent in bulk, it is possible to ask for reputation feedback as a description of the user’s experience instead of a score, and also to use techniques such as CAPTCHA¹⁴ to establish whether there is a human or an automated script behind the feedback.

Threat Rep. 3 – Impersonation and Reputation Theft

[Use-cases: online markets, p2p networks, anti-spam techniques, web-of-trust]

One entity acquires the identity of another entity (masquerades) and consequently steals his reputation. The aim of the attacker can be twofold: to draw benefit from the good reputation stolen and/or defame the reputation of the victim. In the anti-spam use-case, for example, an attacker can impersonate a legitimate user by spoofing his e-mail address, then he successfully spams and/or makes false recommendations. Entities with a high reputation rating are more likely to be the victims of an impersonation attack. The responsibility to mitigate this problem falls on the underlying system, which needs mechanisms to protect the identity infrastructure. But, while there is a strong policy and enforcement regime in the identity space, this is missing in the reputation space. For example, mitigating reputation theft is harder than mitigating identity theft, because subjective perception is involved [35].

Bootstrap Issues and Related Threats

Threat Rep. 4 – Bootstrap Issues and Related Threats

[Use-cases: online markets, p2p networks, anti-spam techniques, web-of-trust]

The ‘bootstrap issue’ is related to the initial reputation value given to a newcomer who has not yet built any reputation. The choice of the entry value is not trivial. One design option may be to award trust to the newcomer until he misbehaves; however this may encourage bad entities to change identity regularly (opening the way for the sybil attack and the whitewashing attack). Another option is to award no trust to the newcomer who has to work to build up his reputation. This would increase his desire to preserve his identity and reputation; however he would initially be penalised and might thus be discouraged. This is what has been called ‘the social cost of cheap pseudonyms’ [30]. Another proposal is to compromise and let existing peers lend part of their reputation to a newcomer [36]: the recommender puts some of his reputation at

¹³ This attack is named after Sybil, the book describing the case study of a woman suffering from multiple personality disorder.

¹⁴ Completely Automated Turing Test To Tell Computers and Humans Apart (CAPTCHA) refers to a program generating tests that humans can pass but computer programs do not. www.captcha.net/

Principal Threats

risk for the newcomer, is rewarded if the newcomer behaves well and penalised if he misbehaves. Sometimes the bootstrap issue is addressed by having a community already formed with pre-existing trust, where essentially newcomers are not accepted.

Extortion

Threat Rep. 5 – Extortion

[Use-cases: online markets]

The extortion threat is carried out in co-ordinated campaigns aimed at blackmailing an individual by damaging his reputation for profit or general malicious motives. For example, denial-of-reputation (see below) can lock-out the victim who is then blackmailed in order to clean his reputation. In retaliation, a user leaves a negative reputation score to another user, because he received a negative score from the latter. This opens the way for extortion, and in general, it discourages honest users from leaving honest negative scores. A variant of retaliation is engaging with a high-reputation seller and misbehaving, then blackmailing him with the threat of leaving him negative feedback if he does not leave positive feedback as buyer¹⁵. In the online market use-case, traders may choose not to leave feedback (score) after a transaction because they are afraid of retaliation; hence also ‘silence’ should be taken into account by the system [37]. For fear of retaliation, users tend not to express negative feedback. For example, there are reports that highlight how feedback on eBay is unrealistically positive [9] and how there is high correlation between buyer and seller ratings, particularly evident in the case of negative feedback. Gross and Acquisti [38] suggest a solution where the seller gives feedback first and then the buyer is free to give the ‘real’ feedback without fear of retaliation. Also privacy-respecting mechanisms may mitigate the threat of retaliation.

There is a full range of potential damage that this threat can cause: emotional damage to vulnerable individuals (bullying); an increase in criminal activities; damage to reputation; loss of revenue for high-value identities; loss of trust in reputation judgements (as distorted value

judgements reduce trust); and spill-over to physical world events (e.g. the revelation of the private contact data of a third party or the organisation of physical bullying). Possible vulnerabilities that this threat can exploit include the lack of formal management/assurance mechanisms for reputation and the lack of data quality assurance (i.e. anyone can post anything).

Denial-of-reputation

Threat Rep. 6 – Denial-of-reputation

[Use-cases: online markets, anti-spam techniques]

Denial-of-reputation is a concerted campaign to damage the reputation of an entity, in order to isolate the victim – effectively performing a subtle lock-out of the victim, often with the intention to extort. Denial-of-reputation is performed by falsely reporting on the victim’s reputation (i.e. performing bad mouthing by means of collusion or using sybils), or by stealing the victim’s identity and misusing his reputation. Then, the victim is blackmailed (extortion) in order to have the reputation cleaned. Countermeasures to this attack are currently not well developed. Gartner [39] recommends strengthening research into the technology, but also into the associated human behaviour, and to investigate new mechanisms to defeat automated attacks to reputation systems.

Unfair Rating

Threat Rep. 7 – Ballot Stuffing and Bad Mouthing

[Use-cases: online markets, anti-spam techniques]

It is possible for an entity to report a false reputation score [9] [40]. In *ballot stuffing*, a number of users agree (by colluding, or using sybils) to give positive feedback to one entity, to make her quickly gain a good reputation. In *bad mouthing*, the attackers (distinct or sybils) collude to give negative feedback on the victim, to lower or destroy her reputation, or to perpetuate denial-of-reputation. Generally this attack is conditioned in its effectiveness due to statistical reasons. ‘Controlled anonymity’ is used in [40] as a countermeasure.

¹⁵ This activity is called ‘feedback extortion’ on eBay’s help pages and, in a November 2004 survey, 38% of the total respondents stated that they had “received retaliatory feedback within the prior 6 months, had been victimised by feedback extortion, or both” [22].

Principal Threats

Threat Rep. 8 – Collusion

[Use-cases: online markets, p2p networks, anti-spam techniques, web-of-trust]

Collusion means that multiple users conspire (collude) to influence a given reputation rating. Collusion is typically done to perform reputation fraud, for example ballot stuffing or bad mouthing. In the web-of-trust use-case, multiple peers may be colluding to introduce a public key. It depends on the community size whether collusion can manage to influence the resulting reputation. Note that it is also possible to buy the vote of reputation voters, for example by hiring scamming firms which in turn can recruit networks of voters.

Threat Rep. 9 – Repudiation of Data and Repudiation of Transaction

[Use-cases: online markets, p2p networks]

A peer can deny that a transaction happened, or the existence of data for which he was responsible. While it is easier to implement a verification to ascertain the transaction in centralised systems, in decentralised systems the problem is more apparent. Replication of data to multiple peers is one mitigation technique adopted in p2p networks (although it can be made ineffective by the sybil attack or collusion practice). Proof of the transaction should be requested [31].

Threat Rep. 10 – Recommender’s Dishonesty

[Use-cases: online markets, p2p networks, anti-spam techniques, web-of-trust]

A reported reputation is strongly dependent on the trustworthiness of the voter providing reputation feedback. For example, an existing member recommends someone who is not trustworthy and the new member then proceeds to distort other reputation scores. In the web-of-trust use-case, a key introducer may not be trustworthy in his recommendation. This can also be strengthened by collusion action. One suggestion to mitigate this threat is to introduce weightings to a reported reputation score, according to the reputation of the recommender (voter). Another mitigation technique uses only recommenders from a trusted social network.

Privacy Threats

[Use-cases: online markets, p2p networks]

Threat Rep. 11 – Privacy Threats for Voters and Reputation Owners

Giving an honest opinion on a sensitive topic requires that privacy is guaranteed to the voter. The situation is similar to e-voting, where the reputation owner has a strong incentive to try to influence the result, either by threatening the voters or by other means of unfair influence. If the privacy of voters is not guaranteed, there is a risk of distorting the votes due to the voters’ fear, and the risk of different threats to voters from the reputation owners (e.g. extortion or retaliation, mainly when negative feedback is involved).

Similarly, there are threats against the privacy of the reputation owners. There is a conflict between the privacy of the reputation owner and the desire to have a linkable identity for reputation. This privacy threat is present for example in the online market use-case, where the system owner often allows the generation of user profiles including all contexts in which users have been involved [41]. Pseudonyms are often used to enhance privacy; however, pseudonyms can also suffer from linkability. For example, the larger the profile, the greater the possibility of succeeding in linking the reputation to the real holder. Several measures are available however to mitigate linkability [41].

Social Threats

Threat Rep. 12 – Social Threats

Risk of Herd Behaviour and Penalisation of Innovative, Controversial Opinions

[Use-cases: online markets]

Innovative opinions that challenge the status quo are fundamental to the progress of society. However, proposing something new often results in being criticised by the large majority who consider the current situation acceptable. This might result in everyone just reinforcing what the majority thinks and refraining from proposing something that might lead to a bad

Principal Threats

reputation (at least initially) [42] [7]. In effect, this can both penalise creative, independent thought and hinder society's progress. Providing more anonymity to voters may help to mitigate this threat.

Another possible countermeasure is to allow the computation of personalised reputation scores by means of local trust metrics, so that a controversial user might have a low computed global reputation rating (he is proposing something that is currently not accepted by the entire community) but he might have a very high reputation rating in the eyes of a restricted number of users (his trust circle). Personalised reputations allow users to be more free in the expression of their real opinions without fear of being isolated.

Vocal Minority Effect

[Use-cases: online markets]

Votes may only be obtained from those with strong opinions, therefore the reputation may be skewed by the fact that those with moderate opinions do not vote¹⁶. This leads to reputation inaccuracy. Lack of voting should somehow be interpreted, and effort should be devoted by the system to obtaining these missing votes.

Discriminatory Behaviour

[Use-cases: online markets, p2p networks]

An entity can engage in discriminatory behaviour towards others. For example, in second-order reputation systems (i.e. weighing the recommender's trustworthiness), an entity can choose to co-operate only with peers who have a high reputation rating, so that his reputation scores highly because their recommendation is weighted more heavily [31].

Threats via the Underlying Infrastructure

Threat Rep. 13 – Attacks to the Underlying Networks

[Use-cases: online markets, p2p networks, anti-spam techniques, web-of-trust]

The reputation system can be attacked by targeting the underlying infrastructure. Reputation information can be attacked (manipulated/replayed/disclosed) both when stored in the nodes and when transported. It can be interrupted by, for example, a denial of service (DoS) attack or misrouting. An infrastructure which does not implement strong identity management may leave the reputation system open to impersonation threats including sybils (so, for example, anti-spoofing techniques are needed in the anti-spam use-case). The centralised reputation system model suffers because the central entity can become a target and affect the whole system. There is a wide variety of threats that can affect reputation systems at this level, but they are not specific to reputation systems and are therefore not covered in this paper.

Threat Rep. 14 – Trust Topology Attacks

[Use-cases: online markets, p2p networks, anti-spam techniques]

Attacks may try to exploit trust relations among members of the community. The attack may, for example, target links which, if broken (e.g. by DoS), would have maximum effect (such as to weaken or kill the community). The attacker may also investigate which entities have the highest reputation ratings and attack them, since their recommendations have the greatest impact. In peer-to-peer networks, hijacking nodes with a high reputation rating and therefore high throughput may strengthen DoS attacks. One way for the attacker to perform a trust topology attack is to extract the trust mapping by asking for reputation feedback [43]. Seigneur and Gray [17] provide examples of such topology-based attacks against anti-spam systems based on reputation. Revealing trust relationships within a network may also be a threat to the peers' privacy if they preferred some relationships to be kept private.

¹⁶ www.rateitall.com/i-18457-harvard-university.aspx is a good example of a vote skewed by vocal minorities.

Threats to Ratings

Threat Rep. 15 – Threats to Ratings

[Use-cases: online markets, p2p networks, anti-spam techniques]

There is a whole range of threats to reputation ratings and to the metric used by the system to calculate the aggregate reputation rating from the single scores given by the recommenders. These include the following (several of which have already been mentioned above):

- Threats against the secure storage of reputation ratings. In a centralised system, the central depository is the single point of failure. In a decentralised system, the reputation data is replicated in different points of the network. Security measures need to be in place to secure the storage and prevent manipulation (i.e. against attacks to the underlying networks).
- Threats against the secure distribution of ratings (secure transport), including the modification and replay of reputation messages, accidental loss (e.g. via attacks to the underlying networks), confidentiality/privacy of scores.
- Threats against the privacy of voters, e.g. timing attacks.
- Central hubs/peers retaining scores without forwarding them.
- Peers not providing their feedback (solicitation may be needed).
- Linearity of ratings, e.g. the linear combination of ratings may be subject to the ‘strike and recharge’ attack. Weightings should be used.
- Some non-linear algorithms used for the rating calculation can be complex and their result may be unpredictable.

- The reputation scoring itself may be open to diverse attacks, and is in general an important design choice. For example, if it uses a positive-to-negative range (e.g. +1, 0, -1), it gives more weight to past compared with recent behaviour. The choice strongly depends on the specific application and on the specific requirements of the reputation system (robustness to sporadic bad behaviour, robustness to failures, avoiding the risk of overloading a node with a good reputation etc.).

The following are examples of attacks to link-based ranking algorithms, which are well known since they have affected popular applications such as Google Search.

Attacks to Link-based Ranking Algorithms

Some attacks that are possible against search engines rank the returned web pages using link-based algorithms (i.e. according to their interlinking with other high-ranked pages). A *link bomb* (sometimes called Google bomb) [44] is an attack aimed at manipulating the ranking of a returned web page. There have been a number of these attacks reported in the press, particularly against the search engine Google; the motivation behind such attacks may often be political, as well as competitive. Link bombing with commercial motivation may be used by attackers to perform spamdexing (to provide the impression that a web page is popular), which includes *link spam* (taking advantage of the link-based ranking of the search engine) [45]. These types of attacks are effectively attacks to the metrics. The algorithm for the calculation of the page position ranking (i.e. the metric algorithm) needs to take into account the way these manipulation attacks are performed, and to use downplaying components in the algorithms themselves (i.e. giving less weight to certain elements fed into the metric algorithm) [46].

Security Requirements

A number of security requirements for reputation systems are identified below; they include requirements that users expect from services which employ reputation systems, as well as requirements needed if these systems are to remain robust to attacks such as those identified above. These requirements emerged from the four use-cases described above, from which clarifying examples are provided.

- **Availability:** in particular, when the reputation system becomes critical to the functioning of the overall system. For example, centralised systems (e.g. in the p2p network use-case) may be more prone to single-point of failure (the central unit) than decentralised systems. In a decentralised system, if peers whom the user trusts are not available, then the reputation he can obtain from peers might be insufficient.
- **Integrity of reputation information:** the reputation information should be protected from unauthorised manipulation, both in transmission and in storage. This typically translates into security requirements on the underlying networks – e.g. protection of the communication channel, protection at the central reputation repository, and protection at peer-level, for example in the p2p use-case where the reputation information is scattered throughout the network. In the online market use-case, bidding, selling an item and rating another member has generally (as in eBay) to be confirmed by logging in with a pseudonym and a password, then the channel is SSL-protected. Furthermore, the reputation information should be linked to sources (for example in p2p networks, the reputation of a peer is not necessarily known or easily established).
- **Authentication of an entity and access rights** that each entity has in the network/group (to avoid biased contributions to the evaluation of the entity's reputation, for example in p2p networks). Furthermore, identity management mechanisms need to be in place to mitigate the risk of threats related to identity change (sybil attack, whitewashing attack).
- **Privacy/Anonymity/Unlinkability:** privacy should be preserved, as well as anonymity when offered (for example, in the online markets use-case). Privacy should be guaranteed both for the reputation owner and the reputation voter. Anonymity or pseudonymity is often offered on the Internet; person pseudonyms are typically used in the online market use-case. For example, every member of the eBay community is linkable to his pseudonym, and so are all his purchases, sales, ratings and his details. Unfortunately the reputation systems currently in use allow the generation of interest and the behaviour profiles of pseudonyms (such as time and frequency of participation, valuation of and interest in specific items). If the pseudonym becomes linked to a real name, as it typically does for trading partners, the profile becomes related to this real name as well.
- **Accuracy:** the reputation system should be accurate in the calculation of ratings. Accuracy should also consider long-term performance, i.e. the metric may be designed so that temporary minor misbehaviour does not affect the reputation significantly (it could be due, for example, to temporary failure). Other aspects include soliciting (truthful) feedback, but also educating hidden feedback (e.g. lack of vote, observable factors). The ability to distinguish between a newcomer and an entity with a bad reputation should be offered [31]. In the anti-spam use-case, the system should be weighted against false-positives (percentage of legitimate e-mails wrongly identified as spam).
- **Usability/Transparency:** it should be clear to the user how reputation ratings are obtained and what they mean.
- **Fairness:** in particular, in the online market use-case, the bids made are published quickly; this helps members to verify that their own bid is being considered. However, members have to trust the provider that the time of every bid or sale is recorded accurately and that the bids listed were made by the members listed.
- **Accountability:** each peer should be accountable in making reputation assessments. In addition, each member should be accountable for his actions, so that he cannot deny them and is accountable for misbehaviour. If members misbehave, punishments should be considered, both within and outside the system. For example, in the online market use-case, users may complain about unfair behaviour to the provider. In the case of eBay, the dishonest member receives an admonishment and, after multiple admonishments, may even be excluded from the community.

Security Requirements

- **Protection of well-connected entities:** well-connected entities (e.g. those with a high reputation rating) are most likely to be attacked (trust topology attacks). For example in the anti-spam use-case, it is relatively easy to mine the Web to identify the most well-connected e-mail addresses. These should receive a higher level of protection since compromising them puts a high risk on the trust/reputation network.
- **Self-correction:** this might be needed in the case of the overall reputation of each member, since reputation is linked to the subjective opinion of voters. For example, buttons can be used for reporting spam or declaring that mail is not spam [47]. Another aspect of self-correction is the appropriate choice of the period over which reputation is estimated; estimation over a long period allows a strong reputation to be built up, which can make an isolated event (such as a spam mail) negligible. The downside is that an entity will require a longer time to correct a negative reputation.
- **Trustworthiness, including trust in reputation voters:** possible mitigations include making use of existing social networks, and weighting recommendations according to how trustworthy the recommender is (confidence value).
- **Verifiability:** whenever possible, proof should be collected from the interaction that is rated to show that the rating can be verified as correct. In the anti-spam use-case, the messages voted as spam are themselves proof. In the online market use-case, proof is more difficult to collect if the users exchange physical goods. In these situations, a photograph might constitute proof.
- **Security requirements on the underlying networks:** the underlying network should have appropriate security mechanisms in place so that attacks to it do not jeopardise the reputation system, as it is possible to attack the latter by taking advantage of the weakness of the underlying infrastructure (see ‘Threats via the underlying infrastructure’). The security requirements for the underlying networks, however, are not central to the scope of this paper although they represent general prerequisites to the other requirements listed here.
- **Performance efficiency:** the reputation system should have minimal performance impact. For example, in decentralised environments such as p2p networks where the reputation information is scattered throughout the network, there may be an impact on bandwidth and storage.

Recommendations

Recommendations are grouped by target audience.

Recommendations to Providers using Reputation Systems and to Designers of Reputation Systems

Rec. Rep. 1 – Perform a Threat Analysis of the Reputation System

Before designing or adopting a reputation system, a threat analysis should be performed, and the security requirements should be identified. This is security best practice, but it appears not to be common in the design of reputation systems. The threats, and the related attacks as threat vectors, need to be considered in the context of the particular application or use-case, since each of these has specific security requirements. In this paper examples have been identified of security requirements, threats and attacks that should be taken into account in the design and choice of a reputation-based system.

Rec. Rep. 2 – Develop Reputation Systems which Respect Privacy Requirements

Unfortunately the design of current reputation systems allows the generation of user profiles including all contexts in which the user has been involved. Anonymity would increase the accuracy of the reputation system, since it mitigates threats such as extortion and reduces the common fear of retribution for stating a negative (albeit accurate) opinion. A more privacy-respecting design of reputation systems is needed while at the same time preserving the trust provided to the entities by the use of reputations. There are mechanisms that provide both privacy for voters (together with anti-sybil protection) and privacy for reputation owners [48] [49]. They can be implemented by making reputation systems interoperable with privacy-enhancing identity management systems which assist users in choosing pseudonyms (and the reputations associated with them) and informing them about their current privacy [50].

Rec. Rep. 3 – Provide Open Descriptions of Metrics

Reputation metrics should be open rather than closed so that they can be assessed by the greatest possible number of researchers rather than relying on security through obscurity. Threat analysis should be performed to assess whether a metric addresses all the security requirements specific to the application.

Rec. Rep. 4 – User-interface Recommendations

a. Recommendations for the usability of reputation-based systems

To create a sense of trust, it is important that the user understands the system and the way trust is formed and measured within that system and application. In order to meet users' increasing expectations of trustability and reliability, reputation systems should be designed in such a way as to achieve transparency, so as to allow the user to easily understand how reputation is formed (e.g. which factors are taken into account and their weightings), what implications lie behind a given reputation rating, how reputation is verified and how the user can assess the reputation system's own trustworthiness. Searchability is an important aspect of usability, since the user needs to find aspects of reputation relevant to his decision easily.

b. Differentiation by attribute and individualisation as to how the reputation is presented, where possible

A given reputation system should allow a user to customise reputation so as to best accommodate his needs. This would also help transparency. For example, it could be possible to subdivide the reputation rating into attributes (i.e. different aspects or assertions of the reputation object) and allow the user to set an acceptable threshold for each of them (this, of course, is not only a user-interface matter but should also be supported by the system). This would make reputation interoperable with security-token systems like SAML as a means of corroborating assertions (see related recommendation below). Another customisation could allow the user to weight the recommendations he takes into account, according to the confidence he has in them (including his own). The users should be able to switch between local and global trust metrics (or to affect their weighting), for example using a slide-bar (from personalised to unpersonalised predictions). Using local trust metrics, only the opinions of people trusted by the user (and possibly people trusted by people the user trusts) would influence the reputation and the prediction about other people's trustworthiness that the user takes into account [42].

c. Users should be offered qualitative assessment of reputation

Trust is linked to uncertainty, and approaches for supporting trust in computing environments are mainly quantitative. However, there are indications [51] that, when it comes to trust, people prefer to evaluate trust in qualitative terms. Therefore reputation systems should be based on qualitative metrics. If the basis is quantitative, it would be beneficial to translate this into qualitative terms, at the user interface level. However it is also recognised that quantitative indications may offer advantages for a quick evaluation. Hence, using a combination of quantitative and qualitative approaches is recommended, wherever the application allows it.

Rec. Rep. 5 – Promote Awareness-raising

Users should develop skills in understanding reputation ratings and the trust processes behind them, and developers should be made aware of attacks on reputation-based systems. In addition, adding proofs to ratings allows users to have sensible awareness information to help them identify fake ratings.

Rec. Rep. 6 – Encourage the Use of Reputation Systems in Online Social Network Sites

Social Networks (SNs) are one of the most successful technological phenomena of the 21st century. There are a number of threats associated with Social Networking, and the use of reputation systems could be beneficial in addressing threats such as ease of infiltration, squatting, stalking, cyberbullying, SN spam and cross-site scripting. Reputation techniques can help in the assessment of the trustworthiness of the claims that users make about themselves. ENISA's Position Paper on the Security Aspects of Social Networking [52] lists the possible benefits of reputation techniques in SN sites: filtering of malicious or spam comments, filtering comments by quality to increase content quality, increasing the reliability of third party widgets, reporting inappropriate or copyrighted content, reporting profile-squatting or identity theft, reporting of inappropriate behaviour and the posting of high-risk data such as location information. Integrating reputation systems within SNs would reinforce the manual moderation activity of the site owner (which may suffer, for example, because of time and bias limitations), by encouraging the users themselves to take responsibility and build up a more engaged, self-regulated community.

Recommendations to Research and Standardisation Communities

Rec. Rep. 7 – Encourage Research into:

a. Understanding the social components behind reputation systems

There is a need to investigate the social aspects and subtleties that influence reputation models. We have seen how many of the identified threats have social origins, and how often they are difficult to anticipate because they are linked to human behaviour. Research should be encouraged to develop a better modelling of human goodwill, and to look into the possible conflict of individual benefit versus community benefit. Attention should be paid to the factors that can encourage honest, accurate feedback from voters, including reward/punishment mechanisms and privacy guarantees.

b. New authentication mechanisms

Enhanced authentication mechanisms should be developed as a countermeasure against attacks such as reputation theft, whitewashing and various automatic attacks. For example, current mechanisms against automated attacks on reputation-based systems include CAPTCHA, which can distinguish the presence of a human being, but new mechanisms are needed, in particular to mitigate the denial-of-reputation threat [39].

c. Common solutions to threats against reputation-based systems

There are a number of threats that can undermine reputation-based systems and a variety of use-cases where these systems can be used. However, as we have seen, these systems are often vulnerable to similar threats and attacks – common solutions to defeat these should be investigated.

d. The management of global reputation

Users should be aware of and control their overall electronic reputation which is composed of fragments scattered across the Internet. Some applications which are moving in this direction have already appeared [53], but the research community should be encouraged to investigate real automated, scalable mechanisms to allow the user to manage his global reputation.

Recommendations

e. Anti-phishing tools based on reputation

Reputation systems can be used effectively to improve the ability to detect phishing. Their ability to provide advanced and dynamic filtering of spam, as well as of suspected phishing websites, reduces the exposure of users to phishing attacks.

Anti-phishing toolbars¹⁷ are in operation; however more research is needed in order to improve their accuracy, as well as the usability of the results (for example, it is not enough to show a green or red light on the toolbar itself) [54].

f. Use of weightings in the metric

Further research should be undertaken into adding appropriate use of more sophisticated *weighting systems* to the metric, i.e. different weightings to improve the resistance of the metric to attacks (while at the same time maintaining the requirement for transparency in the trust-building mechanisms). Examples include weightings to take into account the degree of trustworthiness of the recommenders (confidentiality value), weightings in favour of the more recent behaviour of the entity and weightings to take into account the length of the entity's history (e.g. to counter whitewashing).

Rec. Rep. 8 – Encourage Research into and Standardisation of Portable Reputation Systems

Attempts to create portable reputation systems have emerged¹⁸, however none of these attempts has gathered a sufficient user-base¹⁹. There is a need for further research and experimentation in this field. In order to fulfil the goal of a truly open architecture, it would be advisable to make the history of received feedback portable, rather than just the reputation, as computed by the third service. In that way, different applications

would be free to use the raw trust data as they prefer and not rely on a reputation as computed by the third service, often using secret algorithms and undisclosed raw data, because this would mean trusting the aggregated reputation provider (eBay or Amazon, for instance). This is not needed in a really open architecture. For example, eBay could make public and portable the list of feedback a specific user has received, rather than just the reputation of that user as computed by eBay. Automatically combining the feedback history from one context with another is a difficult challenge, even more difficult if the respective users have privacy requirements. Ontology languages such as OWL [55] are currently the best means of obtaining interoperability between the semantics of different reputation judgements so that a single aggregate score can be combined from heterogeneous sources. Such languages provide formal semantics of the terms used in reputation scores, including more subtle contextual elements. This allows them to be compared and 'translated' between different contexts.

Standardise Transport Mechanisms for Reputation Data

Reputation and the benefits of a specific reputation should be transportable. Today, reputation is expressed in many different languages and transported in multiple ways including in single-sign-on systems such as Liberty Alliance and OpenID. The initiative should be to encourage specific communities to take one step back, engage with other communities and try to ensure that the various systems already in existence start to work together and develop the means to transport the data from new reputation systems or metrics.

¹⁷ Anti-phishing toolbars (e.g. from Cloudmark, Google, Microsoft, eBay) are used to flag a suspected phishing site to users. They generally look at a number of characteristics of the site and also rely on users' reports. They typically flag the site through a 'traffic light' kind of approach (green is OK, red is a phishing site, with a third colour for an unclassified site – or a variant of this). Some toolbars (such as SpoofGuard) allow the user to set thresholds for the scores of the different characteristics that are measured (i.e. to weight them) [54].

¹⁸ Examples are Rappleaf (www.rappleaf.com/) and iKarma (www.ikarma.com/). Portable reputation means that an entity can build its reputation through interactions, which may be of different natures, in one environment, and use this reputation in other contexts (i.e. the reputation is not confined to a single application). This is opposed to the closed (limited to a given application) reputation systems which are in use today. A federation of reputation providers can offer large population coverage and truly portable reputation, thus mitigating the security implications of a few central reputation providers monopolising the market. This would allow an entity's reputation to be always available, for disparate uses, and would mitigate the bootstrap problem (as it is then unlikely for an entity to begin with an unknown reputation).

¹⁹ One of the reasons is the well-known chicken and egg problem (i.e., these services need to be adopted by many users before more users join them and the services really take off).

Recommendations

Example: Integrate reputation into authentication transport standards – SAML Authentication Context

If we define (as is common practice) [56]:

- Identity as: “a set of claims made by one digital subject about itself or another digital subject” [57]
- Authentication as: provision of evidence for identity claims (the act of establishing or confirming something, or someone, as authentic) [58],

then reputation, as a means of providing evidence for identity claims (e.g. creditworthiness, trustworthiness as a seller) may be classed as an authentication mechanism.

Within the reputation community there is currently a very strong need for an interoperable format for describing reputation, and in practice reputation often falls into the same use-cases described by token transport standards such as SAML [59], since it is just one of a number of means for increasing trust in a claim. Just as SAML transports PKI-based authentication statements, it could also be a vehicle for reputation-based statements about public keys. Similarly, SAML is also able to transport so called ‘attribute statements’ which are essentially assertions with evidence. Thus SAML could easily be a vehicle for reputation information about assertions/attributes. The use of an established standard may help to create interoperability between ‘social trust’, web-of-trust and PKI.

Examples of portable reputation data might be:

- Bob’s reputation score for {delivered the goods on time} is 37% according to a mean score of 2000 votes.
- “Bob’s public key is ASD1ASd1223ASdAHF87” verified by 50 PGP users.

Recommendations to Governments and SMEs

Rec. Rep. 9 – The Importance of Automated Reputation Systems for e-Government

Governments are not making use of automated online reputation, however policy-makers should be encouraged to investigate the possibility of using state-of-the-art reputation-based systems in e-Government systems.

Applications where such investigation may be worthwhile are those related to vetting and proofing-related services, where governments are already using offline reputation as a basis. A

typical example is the issuance of a passport or identity card, which in some countries requires trusted witnesses to vouch for the association of a user’s photograph with the name (an aspect of his reputation). For example in the UK, passport applications must be accompanied by a photograph signed by a witness who is in a respected profession [60]. Governments should examine the possibility of automating and securing these processes, taking advantage of automated electronic reputation systems and the security experience gained in their development. It should also be recognised that automatic ad hoc reputation systems provide potential for a scalability which is not present, for instance, in PKI systems. As an example, a web-of-trust model could be used instead of a European root Certification Authority for certifying authentication mechanisms across European Union (EU) borders. Furthermore, investigation could be initiated to integrate reputation systems into e-participation activities [61], which is a key objective of EU policy.

Governments should also consider the impact of new online reputation systems on existing legislation such as defamation law. Another example is that reputation providers may be based in a specific jurisdiction but have worldwide coverage with their online reputation ratings. Interpretation of and, where necessary, updates in relation to reputation data (including votes), should be added to data protection legislation such as the 95/46 Directive especially because its usage is increasing on Social Networking Sites.

Rec. Rep. 10 – Small and Medium-sized Enterprises (SMEs) should Embrace the Potential of Reputation Systems for their Businesses

Reputation systems have great potential for small businesses:

- to help differentiate or improve their products (e.g. by asking for feedback from the customers)
- to make business-to-business decisions (e.g. which products to buy for internal use, which vendors to trust, which partners to trust for partner agreements and sub-contracting etc.)
- to bootstrap and/or improve business models by subscribing to reputation providers where a high reputation rating would gain the trust of customers.

Reputation-based systems are proving to be easier to understand, implement and manage than, for example, PKI with certificates, as well as being cheaper in many cases than a PKI-based system.

Concluding Remarks

The use of online reputation systems is spreading far beyond the well known cases of e-mail filtering and online markets. The architecture and economics of reputation systems appear to be following a similar path to that of identity systems, which are seeing an evolution from single-application silos (e.g. a single e-mail account), through single-provider identity services, to identity federation and, as a final stage, perhaps user-controlled identity. Reputation services are following a similar trajectory. They began as single-application silos (e.g. eBay reputation) and then, recognising the same problems with lack of portability, they are now evolving into reputation-as-service applications, but tending to have a single-provider model²⁰. However, unlike identity services, this evolution has not yet addressed the problems of aggregating reputation information under a single provider (predominantly related to the necessity to trust that single provider with all of a user's reputation data). This appears to suggest that reputation systems may (and should) move next to a federation model and finally to user-controlled reputation.

Perhaps the most important contribution offered by reputation-based systems is that they reflect how people interact in offline society, hence they are more intuitive to users and trigger higher confidence. Furthermore, they are often easier and cheaper to implement and manage (than, for example, PKI). Due to its dependency on a critical mass, as new applications embrace reputation-based systems, the value of online reputation will continue to increase – and online reputation will be the target of attacks. Despite being used in a variety of applications, reputation-based systems are often vulnerable to similar threats and attacks – common solutions to defeat them can therefore be defined.

²⁰ Examples: www.rateitall.com and www.flickr.com/photos/scoobyfoo/sets/72057594083488682/show/

Contributors

Art Baker, Microsoft
Katrín Borcea-Pfítzmann, TU Dresden
Mark J. Boyd, eBay
Jon Callas, PGP Corp
Elisabetta Carrara, ENISA (Editor)
Marco Casassa Mont, HP Labs
Sorin Chitu, Romanian Ministry of the Interior
Richard Cox, Spamhaus
Stephen Crane, HP Labs
Christian Dietrich, Institute for Internet Security, Germany
Jay Heiser, Gartner
Giles Hogben, ENISA (Editor)
Paolo Massa, IRST, Trento, Italy
Tim Nash, Venture Skills
Farez Rahman, Redkey Digital
James Andrew Reynolds, SECUDE
Christian Rossow, Institute for Internet Security, Germany
Jean-Marc Seigneur, University of Geneva and Venyo
Sandra Steinbrecher, TU Dresden
Denis Trcek, JSI and FAMNIT UP, Slovenia
Rigo Wenning, W3C

Group members participate as individuals. This paper should therefore not be taken as representing the views of any company or other organisation, and does not in any way bind group members when dealing with the issues it covers in other contexts.

References

- 1 *Cambridge Advanced Learner's Dictionary*, Cambridge University Press, Cambridge, 2007
<http://dictionary.cambridge.org/>
- 2 Rahman and Hailes, *Supporting Trust in Virtual Communities*, 2000
<http://ieeexplore.ieee.org/iel5/6709/20043/00926814.pdf?arnumber=926814>
- 3 Stewart, *Social forces and constraint in the attainment of community status*
<http://opensource.mit.edu/papers/stewart1.pdf>
- 4 Report from the ENISA eID Workshop, Paris, May 2007
http://enisa.europa.eu/doc/pdf/Workshop/June2007/Report_eID_Workshop%20Paris.pdf
- 5 Traupman, *Robust Reputations for Peer-to-peer Markets*, May 2007
www.eecs.berkeley.edu/Pubs/TechRpts/2007/EECS-2007-75.pdf
- 6 Surowiecki, *The Wisdom of Crowds: Why the Many Are Smarter Than the Few*
- 7 Rasmusson and Jansson, *Simulated Social Control for Secure Internet Commerce*, 1996
www.sics.se/~lra/nsp96/nsp96.html
- 8 Ruohomaa et al, *Reputation Management Survey*, 2007
<http://cgi.di.uoa.gr/~s3lab/papers/2007/ares2007.pdf>
- 9 Resnick and Zeckhauser, *Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System*, 2001
www.si.umich.edu/~presnick/papers/ebayNBER/index.html
- 10 www.ebay.com
- 11 eBay, *Annual Report 2006*
<http://investor.ebay.com/annuals.cfm>
- 12 *Internet Fraud - Crime Report*, Internet Fraud Complaint Center, Ic3 2006
www.ic3.gov/media/annualreport/2006_IC3Report.pdf
- 13 Walsh and Sirer, *Fighting peer-to-peer spam and decoys with object reputation*
www.cs.cornell.edu/People/egs/credence/credence.pdf
- 14 Aberer and Despotovic, *Managing Trust in a Peer-to-peer Information System*, 2001
<http://lsirpeople.epfl.ch/despotovic/CIKM2001-trust.pdf>
- 15 Cascella and Battiti, *Social Networking and Game Theory to Foster Cooperation*, ENISA eID Workshop, Paris, May 2007
www.enisa.europa.eu/doc/pdf/Workshop/June2007/Papers/reputation/REP_UniversityTrento.pdf
- 16 Cornelli et al, *Implementing a Reputation-Aware Gnutella Servant*, 2002
- 17 Seigneur and Gray, *Default Free Introduction, Rare Self-Introduction Fee, Costly Spoofing: No Profitable Spam*, 2004
www.cs.tcd.ie/publications/tech-reports/reports.05/TCD-CS-2005-03.pdf
- 18 Golbeck and Hendler, *Reputation Network Analysis for E-mail Filtering*, 2004
www.ceas.cc/papers-2004/177.pdf
- 19 Alperovitch et al, *Taxonomy of E-mail Reputation Systems*, June 2007
<http://ieeexplore.ieee.org/iel5/4278983/4278984/04279023.pdf?tp=&isnumber=4278984&arnumber=4279023>
- 20 Sender Policy Framework
www.openspf.org/
- 21 The IETF Domain Keys Identified Mail (dkim) Working Group
<http://ietf.org/html.charters/dkim-charter.html>
- 22 Massa, A Survey of Trust Use and Modeling in Current Real Systems, in *Trust in E-services: Technologies, Practices and Challenges*, 2006
www.gnuband.org/files/papers/survey_of_trust_use_and_modeling_in_current_real_systems_paolo_massa.pdf
- 23 Baker and Hartrell, Information Security and Reputation Systems, in *Information Security Bulletin*, November 2006
- 24 www.wikipedia.org/
- 25 <http://slashdot.org/>
- 26 Page et al, *The PageRank Citation Ranking: Bringing Order to the Web*, 1998, Stanford Digital Library Technologies Project
- 27 <http://setiathome.ssl.berkeley.edu/>

References

- 28 Reputation-based trust management systems and their applicability to grids, CoreGRID Technical Report N. TR-0064, February 2007 www.coregrid.net/mambo/images/stories/TechnicalReports/tr-0064.pdf
- 29 Rahman, *Threat scenarios for a Reputation community*, ENISA eID Workshop, Paris, May 2007 www.enisa.europa.eu/doc/pdf/Workshop/June2007/Papers/reputation/REP_redkeydigital.pdf
- 30 Friedman and Resnick, *The social cost of cheap pseudonyms* www.si.umich.edu/~presnick/papers/identifiers/081199.pdf
- 31 O'Reilly, *Peer -to-Peer: Harnessing the Power of Disruptive Technologies*
- 32 Douceur, *The Sybil Attack*, 2002 <http://research.microsoft.com/sn/farsite/IPTPS2002.pdf>
- 33 Margolin and Levin, *Quantifying Sybil Attacks against Network Applications* <http://prisms.cs.umass.edu/brian/pubs/margolin.tr05-67.pdf>
- 34 Levin et al, *A Survey of Solutions to the Sybil Attack*, 2006 <http://prisms.cs.umass.edu/brian/pubs/levine.sybil.tr.2006.pdf>
- 35 Report from the ENISA eID Workshop, Paris, May 2007 http://enisa.europa.eu/doc/pdf/Workshop/June2007/Report_eID_Workshop%20Paris.pdf
- 36 Garg et al, *Reputation lending for Virtual communities* <http://ieeexplore.ieee.org/iel5/10810/34089/01623817.pdf>
- 37 Dellarocas and Wood, *The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias*, June 2006
- 38 Gross and Acquisti, *Balances of power on eBay: Peers or unequals?*, 2003, Berkeley Workshop on Economics of Peer-to-Peer Systems
- 39 *Gartner's Hype Cycle for Cyberthreats*, 2006, www.gartner.com/
- 40 Dellarocas, *Immunizing Online Reputation Reporting Systems Against Unfair Ratings and Discriminatory Behaviour* <http://ccs.mit.edu/dell/ec00reputation.pdf>
- 41 Steinbrecher, *Privacy-respecting Reputation System for Future Internet Communities*, ENISA eID Workshop, May 2007, www.enisa.europa.eu/doc/pdf/Workshop/June2007/Papers/reputation/REP_UniversityDresden.pdf
- 42 Massa, *Reputation is in the eye of the beholder: on subjectivity and objectivity of trust statements*, ENISA eID Workshop, May 2007 www.enisa.europa.eu/doc/pdf/Workshop/June2007/Papers/position_paper_PaoloMassa.html
- 43 Rahman, *A Framework for Decentralised Trust Reasoning*, PhD Thesis, University College London, 2005.
- 44 http://en.wikipedia.org/wiki/Google_bomb
- 45 <http://en.wikipedia.org/wiki/Spamdexing>
- 46 Kamvar et al, *The EigenTrust Algorithm for Reputation Management in P2P Networks*, in *Proceedings of the Twelfth International World Wide Web Conference, 2003*
- 47 Seigneur et al, *Combating Spam with TEA (Trustworthy E-mail Addresses)*, 2004 <http://dev.hil.unb.ca/Texts/PST/pdf/seigneur.pdf>
- 48 Steinbrecher, *Design options for privacy-respecting reputation systems within centralised internet communities*, in *Proceedings of IFIP Sec 2006, 21st IFIP International Information Security Conference: Security and Privacy in Dynamic Environments*, May 2006
- 49 Voss, *Privacy preserving online reputation systems*, in *International Information Security Workshops*, pages 245-260, Kluwer, 2004
- 50 Mahler and Olsen. *Reputation systems and data protection law*, in *eAdoption and the Knowledge Economy: Issues, Applications, Case Studies*, Amsterdam, 2004

References

- 51 Kovac and Trcek, Methods and models of trust management in distributed systems, *Proc. of the 16th ERK Conference*, IEEE Region 8, 2007, Vol. B, pp. 7-10
- 52 ENISA paper on Security Aspects of Social Networking
www.enisa.europa.eu/doc/pdf/deliverables/enisa_pp_social_networks.pdf
- 53 www.reputationdefender.com/
- 54 Cranor et al, *Phinding Phish: An Evaluation of Anti-Phishing Toolbars*, November 2006
www.cylab.cmu.edu/files/cmucylab06018.pdf
- 55 Web Ontology Language
www.w3.org/2004/OWL/
- 56 http://wiki.enisa.europa.eu/index.php?title=Authentication_Interoperability#Proposal_6._Reputation
- 57 Cameron's Identity Blog
www.identityblog.com/stories/2004/12/09/thelaws.html
- 58 Hogben
www.w3.org/2006/07/privacy-ws/papers/14-hogben-assertion-and-evidence/
- 59 OASIS Security Assertion Markup Language (SAML)
www.oasis-open.org/specs/index.php#samlv2.0
- 60 Note 8, Section 10
www.fco.gov.uk/Files/kfile/passport%20-%20c1notes.pdf
- 61 http://ec.europa.eu/information_society/activities/egovernment_research/eparticipation/index_en.htm

For further information about this Position Paper, contact
Elisabetta Carrara
(elisabetta.carrara@enisa.europa.eu)
or Giles Hogben
(giles.hogben@enisa.europa.eu)

© ENISA – European Network and Information Security Agency, 2007



ENISA - European Network and Information Security Agency
PO Box 1309, 710 01, Heraklion, Crete, Greece
Tel: +30 2810 39 12 80, Fax: +30 2801 39 14 10
www.enisa.europa.eu