Presented at 1st International Conference on Cloud Computing (CloudCom 2009), Beijing
https://link.springer.com/chapter/10.1007%2F978-3-642-10665-1_10

Privacy in a semantic cloud: What's trust got to do with it?

Åsmund Ahlmann Nyre and Martin Gilje Jaatun

SINTEF ICT, NO-7465 Trondheim, Norway
{Asmund.A.Nyre, Martin.G.Jaatun}@sintef.no
http://www.sintef.no/ses

Abstract. The semantic web can benfit from cloud computing as a platform, but for semantic technologies to gain wide adoption, a solution to the privacy challenges of the cloud is necessary. In this paper we present a brief survey on recent work on privacy and trust for the semantic web, and sketch a middleware solution for privacy protection that leverages probabilistic methods for automated trust and privacy management for the semantic web.

1 Introduction

Cloud Computing will be an enabler for the Semantic Web, e.g. by distributing analysis, transformation and querying of data [1]. The Semantic Web as envisioned by Berners-Lee et al. [2] represents a shift from machine readable data towards machine understandable data, allowing machines (e.g. agents) to make intelligent decisions based on the meaning of data on the web.

Similarly, securing the Semantic Web constitutes a shift from current security solutions relying on humans to perform intelligent decisions and assessments, to a semantic security solution where this can be done by automatic and autonomous agents. Providing a basis for such intelligence is assumed to be a highly difficult and complex task [3], but is nevertheless a prerequisite for the anticipated wider adoption of the Semantic Web and semantic technologies.

According to a recent survey on European citizens' perceptions on privacy [4], two-thirds of participants said to be concerned that organisations holding personal information would not handle them appropriately, which is the at the same level as the previous [5]. The survey also showed that four out of five EU citizens feel uneasy about transmitting personal data on the internet due to lack of security, while only one out of five said they used tools for technologies to increase the level of security. This indicates that there is a strong need for better and more reliant privacy control to combat the current threat, and an even stronger one for the future threats.

However, any such privacy enhancing technology is not anticipated to be implemented and deployed in operational environments without providing significant evidence of its correctness and fitness for use. This paper will sketch the first step in creating a privacy middleware for the Semantic Web to be adopted and deployed by the industry.

The remainder of the paper is organised as follows. In section 2 we give a brief overview of existing approaches and solutions to privacy and trust and investigate the current challenges. Next, in section 3 we outline our approach to privacy enforcement through integrated trust and privacy management. Our solution is then discussed in Section 4, before we give our concluding remarks and outline further research in Section 5.

2 Foundations

Although practical security security solutions for the semantic web remain elusive, there is an ample body of relevant security knowledge to draw upon.

2.1 Privacy

The semantic web opens a whole new world of automated data collection and aggregation, surpassing current web searches by far in terms of precision. It is evident that privacy protection will be an absolute necessity for it to be accepted and fully utilised by end users.

Privacy preferences and policies All major web sites with user interaction currently provide privacy policies describing how personal information will be handled. The fact that such policies are not understood (or even read) by users, served as one of the main motivations for the early Privacy Enhancing Technoligies (PETs) [6–8]. The W3C recommendation Platform for Privacy Preferences (P3P) specification [8] utilizes a mark-up language to allow websites to declare their privacy policy in a standardized fashion, which again allow user agents to display the policy in a way users can easily understand. P3P does not provide privacy on its own, but merely helps users make informed decisions about interacting with specific websites. Much of the criticism towards this specification [9] stems from the complexity of the protocol and the failure to adhere to privacy standards and regulations.

With semantically annotated policies, privacy negotiations may be conducted autonomously by agents. Several policy languages have been proposed for both security and privacy policy specification (e.g. [10–12]). By matching users' policies (or preferences) with web services' policies, privacy can be maintained automatically without the need for manual investigation. A review and comparison of current policy languages [13] suggest that policy languages in general are quite expressive, but further work is required especially for improved usage control and minimal information disclosure. Another point being made is the need for user-friendly interfaces and the ability to adapt to changing preferences and requirements.

Privacy through anonymity A way to protect one's privacy is to remain anonymous, e.g. by providing only non-identifiable information. This is common

in the health sector, where e.g. medical status needs to be published in a way so that the patient's identity is not revealed. K-anonymity [14] is one approach to restrict semi-identifiable information such that at least K subjects share any combination. Other approaches include general anonymity, pseudo anonymity and trap-door anonymity. Pseudo anonymity refers to situations where the identity of users are not their real identity (e.g. usernames, subscriber ID, etc) and thereby provide protection from entities that do not know the link between the pseudonym and the real identity. To be anonymous, a user must not identify herself with anything that can link information from different

However, for some types of services, e.g. online social networks (Facebook, LinkedIn, etc.), the benefit is greatly reduced if identity information is not provided. Anonymity is therefore not the answer to all privacy problems.

Privacy regulations Unlike other security mechanisms, privacy is also protected by law. Hence, any privacy policy (and preference) should be according to the privacy legislation of the given country. EU directives on privacy protection [15, 16] place requirements on member states' legislation as to how personal information is stored, handled and shared. The P3P specification has been criticised for its lack of support for such legislation. The architecture proposed in [17] uses the principles from the EU directives as a foundation for its legislation compliance. The architecture is capable of mediating between users, websites and legislation, to ensure that all parties' requirements are satisfied. While most privacy enhancing technologies are focused solely on protecting personal information explicitly given by users, this architecture is determined to protect both active data (controlled by user, e.g. credentials), semi-active data (partly controlled by user, e.g. sensor data) and passive data (uncontrolled by user, e.g. surveillance cameras).

Privacy through usage control Motivated by the shortcomings of current access control mechanisms, Park and Sandhu [18] proposed the generic UCON usage control model. The model is aimed at being generic enough to encompass both traditional access control, Digital Rights Management (DRM) and trust management. As noted by the authors, privacy management (i.e. controlling personal information) may be seen as the reversed version of DRM; where users are placing restrictions on service providers use of information. The basic model is built up of subjects, objects, rights, authorisations, obligations and conditions. Subjects and objects are similar to that of other access control mechanisms, with the distinction that their attributes may be mutable, i.e. they may change due to access requests. Rights are not considered static and the existence of a certain right or privilege is determined by a usage decision function upon requesting to invoke it. Authorisations determine whether the subject is allowed to perform the requested operation on the object. Obligations refer to the mandatory requirements a subject must fulfil before or during usage while conditions describe how environmental or system status may influence usage decisions.

Privacy policy enforcement The current web provides no means of controlling information after it has been published. Anything on the web is visible by all, and is generally hard (or even impossible) to remove. Thus, the best privacy policy would be never to make personal information available to anyone. However, that would also greatly reduce the usefulness of the web, especially the interactive, user-driven services.

Policy enforcement has traditionally (e.g. for access control) been done by a central entity, typically the provider. However, with distributed information and ubiquitous environments, the information provider might be required to enforce restrictions on remote devices. Realising this, Sandhu et al. [19] propose a client-side enforcement strategy based on trusted computing and Privacy Enforcement Implementation (PEI) models.

The approach taken by Lioudakis et al. [17] is to establish a privacy infrastructure similar to that of public keys (PKI). Service providers implement a Discrete Box functioning as a privacy proxy for end-users. Whether to grant requests for personal information is handled by the containing Policy Decision Point and Policy Enforcement Point (PDP/PEP) of the Discrete Box. Policies considered for such a decision include both statutory, service provider and user policies. The idea is that the service provider's privacy proxy guarantees that all applicable policies (regardless of origin) are met whenever access to personal information is granted. To prevent misbehaving privacy proxies, the infrastructure is equipped with a set of Privacy Authorities to supervise service providers adherence to general legislation, user policies and their own specific policies. There are apparent similarities with the Certificate Authority required for the X.509 Certificate infrastructure [20]. Additionally, when applied to the semantic web, each user agent must have its own privacy proxy (Discrete Box), which is a major challenge in terms of scalability.

As stated earlier, the P3P specification [8] offers no enforcement guarantees, and hence the user must determine on its own whether to trust the service provider to adhere to its own policy.

Commercially available privacy management systems (e.g. IBM's Enterprise Privacy Architecture) assume centralized data storage, which leaves them unable to cope with distributed data on the semantic web. The system proposed by Song et al. [21], utilizes social networks as a model to control private data flow within enterprises, not across organizational boundaries.

and Privacy and Identity Management for Europe (PRIME))

2.2 Trust management

Trust, and more specifically trust management, has received considerable attention from security researchers over the past years [22], apparently without being able to make a definite impact on services that are actually deployed on the internet.

The problem with *trust* is that it takes various meanings in various contexts. In a PKI, a certificate is said to be trusted if the link between the owner entity (e.g. user) and the public key is either known in advance, or is confirmed by

a trusted entity. On the current web, the content of a web page is assumed to be trusted if it is provided by a trusted source (as seen by the user). What constitutes a trusted source, is not trivially explained.

Definitions Trust is not easily defined and many definitions exist both within computer science and social sciences [23–25]. Mayer et al. [24] state that organisational studies dealing with trust has been hampered by lack of consensus on contributing factors, trust itself and outcomes of trust. This is supported by a survey of organisational trust [25] and a computer science counterpart [23] where several definitions trust are listed based on different factors and viewpoints. The common factors of these definitions are vulnerability and risk[1], implying that the trustor must be vulnerable to the actions of the trustee and that the inherent risk is recognised and accepted in order to call it trust. Mayer et al. argue that it is the recognition of risk that separates trust from confidence, where the latter does not consciously consider the risk involved. Cooperation is another factor that may be both a contributing factor and an outcome of trust. Trust may result in cooperation and cooperation may result in trust, but they are not dependent on one another. Entities may be forced to cooperate without any trust relation. Similarly, predictability of entities may be a contributing factor of trust, however only if performance is satisfactory. If always performing badly, predictability may lead to decreased trust [24].

We chose to use the definition from [24] where trust is defined as *the willingness of a party to be vulnerable to the actions of another party based on the expectation that the other will perform a particular action important to the trustor, irrespective of the ability to monitor and control that other party.*

Trust models As with trust definitions; several different trust models have been proposed over the years, covering different aspects and views of trust. Many of the models that have been proposed have been targeting a very specific use (e.g. e-commerce) and therefore have sacrificed completeness for simplicity, while others have attempted to specify general and somewhat complex trust models.

Mayer et al. [24] in their attempt to integrate various previous models of trust, focused on a general model. They viewed a trust relation as dependent on the trustor's willingness to trust and the trustworthiness of the trustee (as seen by the trustor). The main factors of trustworthiness were identified as ability, benevolence and integrity. On the trustor's part; disposition to trust and perceived risk were identified as the most influential factors with regards to trust. Furthermore, the outcome of a trust relation (experience) is assumed to influence one or more of the trustworthiness factors and hence the trustworthiness of the trustee.

The work by Marsh [26] was an early attempt to establish a formalism for trust in computer science, and artificial intelligence in particular. The formalism allows agents to compute a trust value based on a set of factors in order to arrive

[1] Deception is a term used by many in information security

at a trust decision automatically. The complexity of the model makes it difficult to use in practise, however as inspiration the model has contributed greatly to advances in research on trust.

Acknowledging that the complexity of several proposed models does not necessarily give better trust assessments, led Conrad et al. [27] to propose a lightweight model for trust propagation. The parameters self confidence, experience, hearsay and prejudice are used to model and assess trust. This computational model also allows agents to compute a trust value to automatically perform trust decisions. The degree of self confidence determines how much influence own experience and hearsay would have on the computed trust value. The more confident; the more dependent on own experience. The prejudice determine the initial value of experience and hearsay, before experience is accumulated.

In the model proposed by Gil and Artz [28] the idea is to arrive at content trust, where the information itself is used for trust calculation. This allows for a whole new range of parameters (such as bias, criticality, appearance, etc) to be used when assessing trusts in resources. The problem of such parameters is that they require user input, which conflicts the assumption of agents conducting the assessment autonomously.

Trust propagation A lot of research has been focused on capturing trust as it is displayed and propagated in social networks, commonly modelled as a weighted digraph where vertices represent entities and edges trust relationships between entities. Golbeck and Hendler [29] describe an algorithm for inferring trust and reputation in social networks when entities are not connected directly by a trust relationship. This is done by computing the weighted distance from the source to the sink. Any distrusted entity is not included in the computation since the trust assessments done by such entities are worthless. Guha et al. [30] introduce the notion of distrust to address the problem of expressing explicit distrust as a contrast to the absence of trust. Absence of trust may come from lack of information to conduct a proper trust assessment, while distrust expresses that a proper assessment have been conducted and that the entity should not be trusted. Furthermore, they argue that distrust could also be propagated and proposes several propagation models in addition to trust transitivity, including co-citation, which is extensively used for web searches.

Huang and Fox [31] claim that not all kinds of trust can be assumed to be transitive. They note that trust based on performance, i.e. an entity performing as expected repeatedly, is not necessarily transitive, while trust based on a belief that the entity will perform as expected often is.

3 Probabilistic privacy policy enforcement

From the discussions above we know that some of the proposed PETs assume that entities will always adhere to and enforce their own policies, either because they are trusted or because there is an infrastructure in place that would not

allow them to misbehave. As a consequence, enforcement is seen as binary, either it is done or it is not.

While assuming that all entities will enforce relevant policies is clearly not a good idea, there are quite some difficulties involved in relying on trusted computing for guarantees.

1. Trusted computing requires an infrastructure (hardware and software) for it to work. Hence, any entity that does not comply with this not allowed to take part.
2. Trusted third parties are needed and are not easily established. Although some have been successfully established for the X.509 Public Key Infrastructure, it is not generally viewed as an unconditional success [32].
3. There may be situations where users do want to communicate with entities not part of the privacy infrastructure, even though this would generally conflict with their privacy requirements. Users would therefore be forced to disable any PET functionality in order to do this.
4. With any such system, there is a critical mass of users/providers that must be attained before users will view investments in such tools beneficial.

Example 1. Consider three websites; one evil, one benign and one somewhere in between (probably owned by Google). All websites provide a privacy policy, possibly very different from one another. Using mere policies, no distinction is made as to the level of trust to be placed in the websites' adherence to their policies, i.e. there is no enforcement. Using trusted computing, only the benign website will be included in the infrastructure, and hence communication with possibly misbehaving websites are impossible (using privacy management).

A user may want to interact with such shady websites despite warnings of misbehaviour and would therefore greatly benefit from a privacy technology that would:

1. Alert the user of the trustworthiness of the website.
2. Record the user's willingness to interact and the willingness to (potentially) be vulnerable to exploit.
3. Provide means to mitigate the risk and calculate criticality and consequence of interaction (e.g. distribution of personal data).
4. Provide anonymity where where appropriate.

We therefore propose a probabilistic approach to policy enforcement, where users are given a probability that their requirements will be respected and polices enforced. Thus when interacting with websites who are known to be less trustworthy, policy adherence is given by a probability metric that the website will actually enforce its own policies. Our enforcement model does not include a privacy or trust model, i.e. it is only occupied with how to handle uncertainty in enforcement and provide a tool for interacting with non-conforming entities while minimising the risks involved. An overview of the resulting middelware solution is sketched in Figure 1, with more details provided in the following sections.

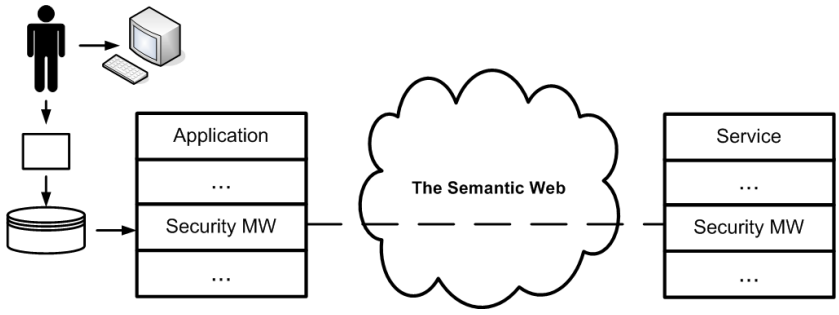

Fig. 1. Conceptual model of a secure semantic web middleware

3.1 Personal Data Recorder

The semantic web offers great opportunities for information aggregation, which is generally difficult to protect oneself from.

Example 2. Consider the situation where a user wanting to stay unidentified has provided his postal code and anonymous e-mail address to a website. Later he also provides age and given name (not the full name) and the anonymous e-mail address. Now, the website is able to combine the data (postal code, age and given name) to identify the anonymous user.

Protecting users from this kind of aggregation requires complete control of what information has been distributed and to whom. In our scheme, this is done by the Personal Data Recorder(PDR), which basically records what data is transmitted to which receivers. Thus in the above example, the second interaction with the website should have been blocked, since it enables the website to reveal the user's identity. The PDR allows the user to view himself through the eyes of the receiving party, and thereby perform aggregation to see whether too much information is provided.

3.2 Personal Data Monitor

The personal data monitor (PDM) is responsible for computing and assessing policies and behaviour, and to update the personal data recorder with inferred knowledge. A problem with the PDR is that it is not capable of handling redistribution of data (receiver forwards the data to other recipients). However, all personal data are assumed accompanied by a privacy policy and obligations. Using the probabilistic privacy enforcement described earlier, the PDM is able to compute the probability that the receiving entity is redistributing information. That is, the PDM will determine the likelihood that the personal information distributed to the receiver will also reach other. This need not be criminal or shady activity either, it is actually quite common in business life. For instance, sending

an e-mail with a business proposition to a specific employee of a company, it is likely that other employees in that company also will receive the e-mail (e.g. his superior). The PDM is in such a case responsible for inferring other recipients and to include such information in the Personal Information Base.

Information that is made publicly available on the Internet, would generally be considered to be available to all. Hence, any interaction later on should consider this information when assessing the kind of information to reveal.

3.3 Trust assessment engine

The Trust Assessment Engine (TAE) is responsible for calculating trust values different entities in order to determine the trustworthiness of other entities. The TAE is thus focused solely on assessing commuicating parties and does not take into account risk willingness, vulnerability and criticality.

3.4 Trust monitor

The trust monitor (TM) is responsible for detecting events that might affect the perceived trustworthiness and the willingness to take risks. The trust monitor is thus responsible for calculating and deciding on what is an acceptable trust level, given the circumstances. Any computed trust value and feedback received from cooperating entities is stored in the trust assessment repository.

3.5 Policy decision point

The Policy Decision Point (PDP) is responsible for the final decision on whether to engage in information exchange and if so; under what conditions. The PDP collects the views of both the TM and the PDM and compares their calculations to the policies and requirements found in the policy repository. The decision is reported back to the TM and PDM to allow recalculation in case the decision alters the calculated trust values or distribution of personal information.

4 Discussion

The practical application of Privacy Enhancing Technologies is limited by the human cognitive capacity – or rather, the lack thereof. However, even on the semantic web, information is ultimately communicated from one human to another, and thus if we want to apply trust to this equation, we have to base ourselves on human notions of trust, which are neither binary nor straght-forward.

In fact, the word “trust” is used to mean many things even in a human context, and is often misunderstood when applied to end-user applications. One example is the Pretty Good Privacy (PGP) program, which allows users to collect public keys of friends ans associates, and subsequently assign a “trust level” to these keys. A common misconception is that this trust level reflects on the degree of certainty that the given key is the correct key for that particular associate;

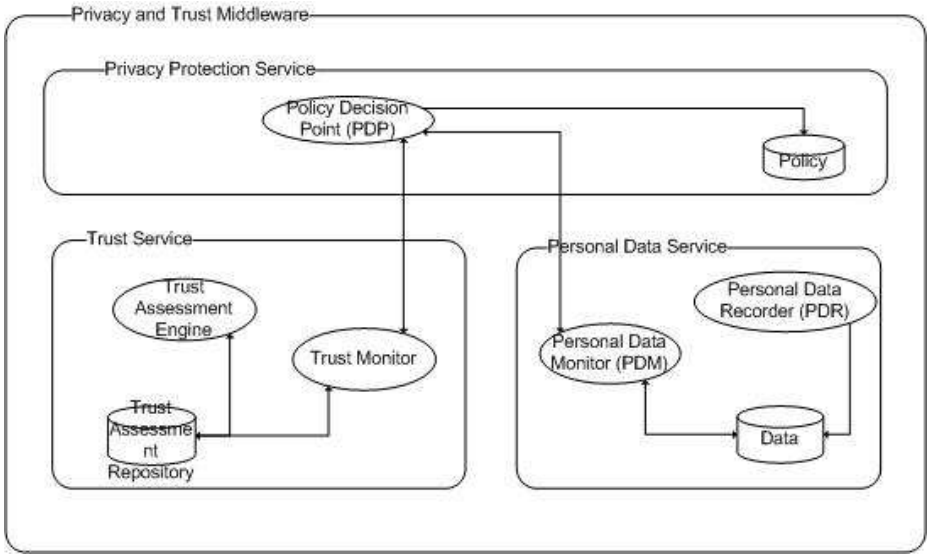

Fig. 2. Middleware architecture for probabilistic privacy management

while in reality it reflects to which degree the user is willing to trust *other* keys signed by that particular associate. These concepts are quite different: While I am confident that I have the right key for my friend Bob on my computer, I also know that Bob is a fool who at regular intervals is sending money to Nigerian princesses, and also freely signs any key that comes his way.

Our proposed middleware relies heavily on the Personal Data Recorder, but it is clear that this element will not be able to cope with passive data collection; in a real-life example, this would require you to e.g. carry a device that could detect all surveillance cameras and record what they capture of your movements. However, since the possibilities for aggregation are so abundant on the semantic web, it is vital that any new PET takes steps to limit unneccesary information spread.

In one way, it may seem that a PET application that introduces anonymous or pseudonymous data would be anathema to the semantic web, since nobody wants innacurate or false data in their system. However, we do not advocate that information that a person *wants* to be disseminated should be anonymized, but rather that the user should be in control of her own information (as in the spirit of European privacy legislation).

One might argue that it would be better if the option to remain anonymous were offered by the various providers, but currently it seems that the providers have no incentive for offering such an option – to the contrary, providers seem to want to go to great lenghts to collect as much information as possible. If we want to be able to communicate with whomever we want, but still want to have

protection against aggregation, it seems the only solution is to lie. Since most humans are bad liers, our privacy middleware will help users to "lie" consistently, allowing them to reap the benefits of the semantic web and cloud computing without sacrificing their privacy.

5 Conclusion and further work

In this paper we have outlined existing approaches to privacy and trust management and the fundamental challenges of the emerging semantic web. We have proposed a new way of handling policy enforcement remotely, based on computing the probability that the recipient will adhere to the established policies. The probability is computed on the basis of trust assertions, user's willingness to trust, and the personal information involved. We believe that such an approach would facilitate a gradual deployment of software since it may prove beneficial to users, regardless of whether other users have adopted it.

We acknowledge that these are early thoughts and that proper justification and simulations should be provided before the benefits of our proposed approach can be rightfully claimed. In particular, a more detailed description of how the Personal Data Recorder and Personal Data Monitor should be designed to meet the goals stated is an important subject for further research. Also, the interface between a trust management system and the personal data service needs to be properly specified to clearly separate the responsibilities of the two, so as to allow for different trust management systems to be utilised. Verification of the approach through simulation or user-testing forms a natural next step.

Acknowledgements

Thanks to Terry Britten and Graham Lyle for inspiring the title of this paper.

References

1. Mika, P., Tummarello, G.: Web semantics in the clouds. IEEE Intelligent Systems **23** (2008) 82–87
2. Berners-Lee, T., Hendler, J., Lassila, O.: The semantic web. Scientific America (2001) 34–43
3. Bussler, C.: Is semantic web technology taking the wrong turn? Internet Computing, IEEE **12** (2008) 75–79
4. : Data protection in the european union - citizens' perceptions. Flash Eurobarometer 225, The Gallup Organization (2008)
5. : Data protection. Special Eurobarometer 192, European Opinion Research Group EEIG (2003)
6. Burkert, H.: Privacy-enhancing technologies: typology, critique, vision. In Agre, P., Rotenberg, M., eds.: Technology and Privacy: The New Landscape. MIT Press, Cambridge, MA, USA (1997) 125–142

7. Goldberg, I., Wagner, D., Brewer, E.: Privacy-enhancing technologies for the internet. In: Proc. of 42nd IEEE Spring COMPCON, IEEE Computer Society Press (1997)
8. Cranor, L., Langheinrich, M., Marchiori, M., Reagle, J.: The platform for privacy preferences 1.0 (p3p1.0) specification. W3C Recommendation (2002)
9. EPIC: Pretty poor privacy: An assessment of p3p and internet privacy. Technical report, Electronic Privacy Information Center (2000)
10. Kagal, L., Finin, T., Joshi, A.: A policy based approach to security for the semantic web. In: Semantic Web - Iswc 2003. Volume 2870 of Lecture Notes in Computer Science. Springer-Verlag Berlin, Berlin (2003) 402–418
11. Bonatti, P., Olmedilla, D.: Driving and monitoring provisional trust negotiation with metapolicies. In: Policies for Distributed Systems and Networks, 2005. Sixth IEEE International Workshop on. (2005) 14–23
12. Damianou, N., Dulay, N., Lupu, E., Sloman, M.: The ponder policy specification language. In: Policies for Distributed Systems and Networks. Volume 1995 of Lecture Notes in Computer Sceince. Springer Verlag (2001) 18–38
13. Duma, C., Herzog, A., Shahmehri, N.: Privacy in the semantic web: What policy languages have to offer. In: Policies for Distributed Systems and Networks, 2007. POLICY '07. Eighth IEEE International Workshop on. (2007) 109–118
14. Sweeney, L.: K-Anonymity: a model for protecting privacy. International Journal on Uncertainty, Fuzziness and Knowledge-based Systems **10** (2002) 557–570
15. EU: Directive 2002/58/ec of the european parliament and of the council concerning the processing of personal data and the protection of privacy in the electronic communications sector. Official Journal of the European Communities (2002)
16. EU: Directive 95/46/ec of the european parliament and of the council of 24 october 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. Official Journal of the European Communities (1995)
17. Lioudakis, G.V., Koutsoloukas, E.A., Dellas, N.L., Tselikas, N., Kapellaki, S., Prezerakos, G.N., Kaklamani, D.I., Venieris, I.S.: A middleware architecture for privacy protection. Computer Networks **51** (2007) 4679–4696
18. Park, J., Sandhu, R.: The UCON_{ABC} usage control model. ACM Transactions on Information Systems Secuity **7** (2004) 128–174
19. Sandhu, R., Zhang, X., Ranganathan, K., Covington, M.J.: Client-side access control enforcement using trusted computing and PEI models. Journal of High Speed Networks **15** (2006) 229–245
20. Housley, R., Polk, W., Ford, W., Solo, D.: RFC 3280: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC Editor (2002)
21. Song, R., Korba, L., Yee, G.: Privacy management system using social networking. In Korba, L., ed.: Systems, Man and Cybernetics, 2007. ISIC. IEEE International Conference on. (2007) 3327–3332
22. Varadharajan, V.: A note on Trust-Enhanced security. Security & Privacy, IEEE **7** (2009) 57–59
23. Artz, D., Gil, Y.: A survey of trust in computer science and the semantic web. Web Semantics: Science, Services and Agents on the World Wide Web **5** (2007) 58–71
24. Mayer, R., Davis, J., Schoorman, F.: An integrative model of organizational trust. Academy of Management Review **2** (1995) 709–734
25. Bigley, G., Pearce, J.: Straining for shared meaning in organization science: Problems of trust and distrust. Academy of Management Review **23** (1998) 405–421

26. Marsh, S.P.: Formalizing Trust as a Computational Concept. PhD thesis, Department of Comuting Science and Mathematics, University of Sterling (1994)
27. Conrad, M., French, T., Huang, W., Maple, C.: A lightweight model of trust propagation in a multi-client network environment: to what extent does experience matter? In: Availability, Reliability and Security, 2006. ARES 2006. The First International Conference on. (2006) 6 pp.
28. Gil, Y., Artz, D.: Towards content trust of web resources. In: WWW '06: Proceedings of the 15th international conference on World Wide Web, New York, NY, USA, ACM (2006) 565–574
29. Golbeck, J., Hendler, J.: Accuracy of metrics for inferring trust and reputation in semantic web-based social networks. In Motta, E., Shadbolt, N., Stutt, A., Gibbins, N., eds.: 14th International Conference on Knowledge Engineering and Knowledge Management. Volume 3257 of Lecture Notes in Computer Science., Northamptonshire, UK, Springer Verlag (2004)
30. Guha, R., Kumar, R., Raghavan, P., Tomkins, A.: Propagation of trust and distrust. In: WWW '04: Proceedings of the 13th international conference on World Wide Web, New York, NY, USA, ACM (2004) 403–412
31. Huang, J., Fox, M.S.: An ontology of trust: formal semantics and transitivity. ACM (2006)
32. Lopez, J., Oppliger, R., Pernul, G.: Why have public key infrastructures failed so far? Internet Research **15** (2005) 544–556