Supporting semi-automated compliance control by a system design based on the concept of separation of concerns

Sebastian Haas[1], Ralph Herkenhöner[2], Denis Royer[3],
Ammar Alkassar[3], Hermann de Meer[2], Günter Müller[1]

[1] Universität Freiburg, Institut für Informatik und Gesellschaft, Abteilung Telematik, haas@iig.uni-freiburg.de

[2] Universität Passau, Lehrstuhl für Rechnernetze und Rechnerkommunikation, ralph.herkenhoener@uni-passau.de

[3] Sirrix AG, Bochum, d.royer@sirrix.com

Abstract: Manual compliance audits of information systems tend to be time consuming. This leads to the problem that actual systems are not audited properly and do not comply to data protection laws or cannot be proven to comply. As a result, personal data of the data subject are potentially threatened with loss and misuse. Automatic compliance control is able to reduce the effort of compliance checks. However, current approaches are facing several drawbacks, e.g. the effort of employing cryptographic hardware on every single subsystem. In this paper a system design is presented that is able to circumvent several drawbacks of existing solutions thereby supporting and going beyond existing mechanisms for automated compliance control.

1 Introduction

With respect to privacy of data subjects, several countries already have established laws and regulations that obligate companies to have extensive measures, which ensure proper and secure handling of *personal data* (PD). However, recent incidents of data loss[1] raise the question why these companies have failed in fulfilling this obligation. The given incidents indicate that weak or ineffective data protection is employed by those companies and that the law enforcement is deficient in detecting the lack of data protection. The latter points show that the current procedures of inspecting the compliance to data protection laws seem to be ineffective and improper. This implication is supported by statistics that cast serious doubt on the effectiveness of governmental inspections.

For example, in 2009, 2.2 governmental inspectors were responsible for 100,000 companies in Germany resulting in an average data protection audit every 39,400 years per company.[2] With regard to the increasing automated processing of personal data, by storing data sets in databases and making these accessible via the

[1] See e.g. http://voices.washingtonpost.com/securityfix/2009/05/hackers_break_into_virginia_he.html (last visit: 25.01.2011) or http://www.wired.com/threatlevel/2009/11/healthnet/ (last visit: 25.01.2011)
[2] http://www.xamit-leistungen.de/downloads/XamitDatenschutzbarometer2009.pdf (in German; last visit: 25.01.2011)

internet, such as in online social networks or in government-driven databases (e.g. central data storages of health cards), this issue poses a significant threat to information privacy.

Privacy can also be considered as an economic problem [Va09]. Thus, assuming that an economically driven company weighs the costs of losing PD (personal data) against the costs of implementing effective measures against data loss. Only an increase in the value and a reduction in the cost of effective data protection can help to improve this situation. Introducing adequate sanctions for not complying with the data protection laws is limited by the effectiveness of law enforcement (i.e. inspection). Hence, increasing the efficiency of compliance audits is a crucial task. Our solution propagates a system design supporting automatic compliance checks, offering companies a tool to maintain and prove their compliance.

Employing the design science paradigm [HMP04] as research framework, this paper presents a schema for data protection supporting automated compliance control that transfers Dijkstra's concept of separation of concerns [Di82] into the domain of data protection (data protection schema and implementation as primary artifacts). The approach has been successfully instantiated as a prototype for mobile health-care services, in which sensitive PD of patients, such as names, addresses, and medical information are processed.

In this paper, the considered scenario is settled in the eHealth domain of home care services. Figure 1a depicts a usual data flow for a home care service. Here, nurses transmit PD `d` to the storage service from which a nursing service is able to retrieve the data. The storage service has the following functions: receiving data with associated policies and storing both embedding data in answer documents and transmitting them to legitimate receivers. As there is no further computation, there is no need for the storage service to access any of the unencrypted PD. The storage service is composed of several complex sub-systems (e.g. load balancers, web servers, databases etc.) or could be a cloud storage database.

The problem addressed in this paper is that current auditing mechanisms for data protection (manual or automated) fail to achieve their goals for complex systems like the described storage service. There is a need for reducing the effort of compliance audits, without decreasing their coverage.

This paper is structured as follows: Following the introduction, Section 2 presents the related work in the domain. Next, Section 3 depicts the developed system for supporting automated compliance control, whose implementation is presented in Section 4. Section 5 discusses the resulting system, including potential attack scenarios and its capabilities. Section 6 summarizes the findings and gives an outlook on future research opportunities.

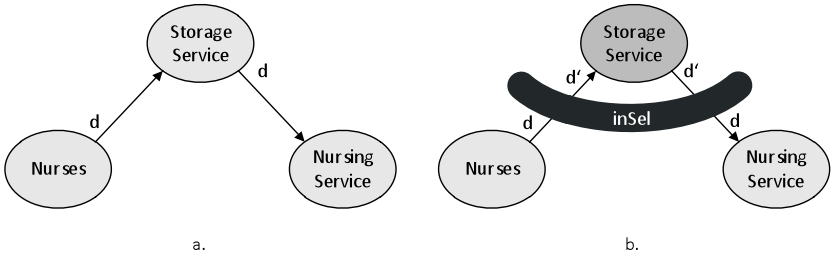

Figure 1: The scenario and the location of the solution inSel.

2 Related Work

By employing automated compliance checks, e.g. by using compliance patterns [Gho007], necessary and verified business process models can be achieved [Lui07], which addresses the issue of compliance on the organizational level. Compliance on the technical level can be achieved by two complementary approaches: by employing *a-priori* mechanisms that enforce the compliance of a system before or at runtime, or by *posterior* controlling that checks compliance after runtime. However, it can be argued that a-priori policy enforcement (e.g., usage control [PS04] and compliance engineering [HJ08]) is too inflexible in some scenarios [Po99, EW07]. Furthermore, posterior controlling is currently under heavy research (e.g. [CC07, EW07, Ac08]). Approaches in this area have in common that log files of a system are used to create a formal model of the usage of PD. Afterwards, the model is used to check, whether the usage was legitimate in relation to a set policy. This approach can be combined with a rollback mechanism for the enforcement of data integrity [Po99].

While these approaches are well suited in theory, there are several drawbacks when applying them to real word systems: a) The reconstructed model is an abstraction of a complex system which might be inaccurate to a certain degree. This is due to, e.g., limited reconstruction capabilities or log inconsistencies. b) A model can only be considered complete, if logs from every component and every layer (i.e. OS, middleware, application, etc.) are used for reconstruction. Otherwise vulnerabilities in the system design or the implementation might lead to unspecified and unlogged behavior of the system. c) Furthermore, in order to ensure authentic logs, secure logging has to be employed on every single subsystem. Any of these drawbacks might render the result of an automated compliance check unusable, as there is no proof that the reconstruction reflects the actual processing of PD.

Applicable measuring points for compliance checks, regarding data protection, are points of enforcement of data access and transfer. Internal data processing is usually regulated by access control mechanisms [PS06] while data transfer is secured

by security gateways [GH06] and the application of layer firewalls [Sc02] and semantic firewalls [As04].

All these approaches are similar in that they enforce security policies that must be compliant with the respective laws. Thus, even with a compliant policy an incompliant transmission might be allowed and will stay undetected if decision or enforcement of the policy fails. If a security system has complete and audit proof logging, they can be extended by compliance checks [Us04]. For retrieving data, these approaches are very similar to our approach. As innovation, our approach assumes that the provider of the sorting service is not a trusted party for storing data. Thus, personal data is protected by using encryption mechanisms against access by the provider of the storage service. Such features are not supported by current approaches.

3 inSel: An approach supporting automated compliance control

The essence of our work is the reduction of the number of components, which must be audited for a meaningful compliance check, and the creation of a trustworthy logging environment. This is achieved applying Dijkstra's concept of separation of concerns to separate security functions from data processing and storage functions. This is done by introducing our solution inSel (**in**formational **sel**f-determination), implementing the security functions in a dedicated security gateway as a barrier for PD (cf. Figure 1b). inSel's purpose is that unencrypted PD text is held off from large parts of the system and access to it will be logged in a secure environment. The main benefit of using this approach is that compliance checks are reduced to checks of the proxy and its secure logs.

3.1 Compliance Control Schema

The schema underlying inSel uses three functions: ***identify*** which separates PD from other data, ***substitute*** which encrypts PD and leaves other data as it is and ***re-substitute*** which recovers the original form of PD. ***substitute*** and ***re-substitute*** are logging access to data. Limitations and applicability are discussed in Section 5.

When correctly applying the schema and adjusting ***identify***, no unencrypted PD is passed to the storage service and every access to PD is logged. A compliance check therefore is reduced to a) a proof of the correct functioning of the proxy as well as b) a check of the respective access logs. Assuming that the correct functioning is correctly described in a model, proof a) is divided into two parts. The first part consists of checking that the schema has been employed this can be done for example by using a trusted computing platform [TCG05] and a certificate of a trusted third party (TTP), showing that the implementation matches the schema. This only has to be done once, as long as the schema and the employed software do not change. Secondly, the correct adjustment of ***identify*** must be proven. This is done manually by a TTP and has to be repeated for every transmission protocol used. b) can be proved by either using manual methods or automatic approaches mentioned

in the related work section. Both approaches can be reduced to a check on the log files of the proxy as the proof of a) states that there is only encrypted PD in the background system.

3.2 Three use cases for compliance

To show the ability for supporting compliance control, three real world use cases were selected.

First, the most general and common case is the *performance reliability*. Purpose of this case is to prove normal operation and, in any case of potential data misuse, to provide evidence on data access and performed operations. To achieve this, there has to be a consistent and audit proof documentation of the data access by subject, object, time and purpose. For reasons of data protection, there has to be access control and consistent and audit proof access documentation.

The second use case is the *inspection of commissioned data processing*. In this case, the data is processed by a third party on behalf of the data controller. To take the German law as an example, the data controller (principal) has to inspect the data processor (agent) for divergence from normal operation. For the reason of data minimization these inspections have to be random (not consistently) but still on a regular basis. A good strategy for triggering the inspection is to do it from both the data controller's and the data processor's side. This minimizes the risk that one single party is able to prevent the inspection of specific operations. Again, a consistent and audit proof documentation of the data access including denied access requests by subject, object, time and purpose is required. Also, there has to be access control and a consistent and audit proof access documentation. For the reason of data minimization, the access should be limited to the documentation of the inspected operation(s).

The third use case is *supporting the right of access*. In Europe the right of access enables data subjects to get information about the processing of their personal data by the data controller. There are specific requirements, to support the right of access in an automated manner [HJ10]. In particular, the legitimate interest has to be checked and the information has to be provided confidentially. Again, a consistent and audit proof documentation of the data access by subject, object, time and purpose is required. Also, there has to be access control and a consistent and audit proof access documentation. For privacy reasons, the access has to be limited to the documentation related to the processing of personal data of the data subject. In particular, personal data of a third person must not be disclosed.

4 Implementation

The schema depicted in Section 3 is implemented by the inSel system, whose architecture is visualized in Figure 2. The system itself is comprised of the following three main components:

- The Secure Hypervisor, representing the basis of the architecture.
- The inSel Core, serving as the security gateway/proxy. Realizing the schema presented in Section 3 and containing the rights management and a user interface which is beyond the scope of this article.
- The Secure Log, offering audit proof log facilities for the substitute and the re-substitute functions.

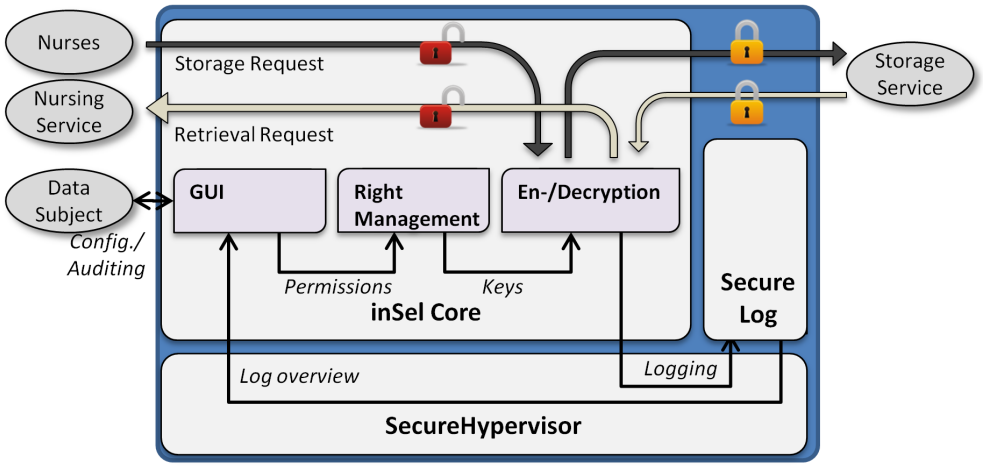

Figure 2: Resulting architecture of inSel.

4.1 The inSel architecture

Operating as a security gateway/proxy, inSel represents the measuring point for compliance checks. This is based on a secure hypervisor, which is the architectural basis of inSel. It comprises of a TURAYA-based [EP10] security kernel and a TPM (Trusted Platform Module) as hardware-based security anchor for performing measurements of the system and the isolated virtual machines (VMs) for the inSel Core and Secure Log. This way, a complete information flow control can be realized. Furthermore, the functionality of the security gateway (inSel Core) and the logging and audit mechanisms (Secure Log) are logically separated.

For every storage and retrieval transaction going through the security gateway/proxy the logging mechanisms stores the following data: identifier of the data subject, identifier of the involved user, type of operation (i.e. storage or retrieval), transferred data types (e.g. name, address), the result of *identify*), a timestamp, and a session identifier (for linking log entries of the same user's session).

For compliance checks, authorized users can send requests to the Secure Audit-Log via the user interface, which only supports pre-defined use cases (e.g. inspection of suspected data misuse or controlling the storage service). This limits the possibility of misusing retrieved logging data and allows the audit of compliance checks itself. Additionally, this can help to provide information to data subjects practicing their right of access by evaluating all transactions for a specific data subject's identifier. For privacy reasons, results of such evaluations should only be provided to the data subject and must not be visible to any other user. This is ensured by the access control mechanism in the compliance interface that will allow requests on information for a given data subject only to the data subject himself.

4.2 Application of the inSel schema

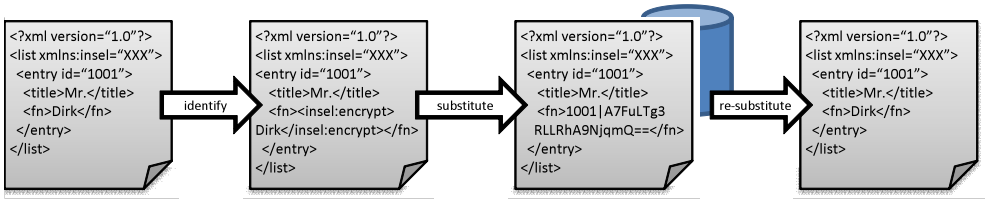

Figure 3: Exemplified XML-Transformations implementing the inSel schema

Before transmitting PD to the storage service, the nurse must authenticate her/himself to inSel. The data is transmitted via HTTP over SSL/TLS in standardized XML documents that are created by the nurse's systems, which are assumed to be trustworthy. Transmitted documents are checked for validity against the RelaxNG schema which is also provided by the nursing service and which is identified via the address of the request. If the document matches the schema, then PD is marked as such by ***identify*** (cf. figure 3). Hence, ***identify*** is implemented as a search on defined XML trees. ***Substitute*** then extracts the data subjects as well as the types of data (e.g. name, address, etc.) from the document according to the schema and encrypts PD and its type with the key of the data subject. After encryption the original PD is replaced by the cipher text and the subjects ID, which is extracted from the XML document. This procedure is repeated for all PD and a log entry is written into the Secure Log, containing data subjects and types. If writing the log is successful the document containing, the encrypted PD is passed to the storage system. These log entries form a protocol showing what data has been collected.

Again, before retrieving data from the storage system, the (employee of) nursing service must authenticate himself to inSel. If the nursing service is not allowed to access the requested URL, the request is discarded and a log entry is written stating the prohibited access. Otherwise inSel retrieves the requested document from the storage service. Relevant data is identified by ***re-substitute*** searching for a predefined pattern ("ID|Base64-string" in the example in figure 3). The matching string is removed and ID and PD are separated. Decrypted PD is inserted where the pattern has been found. After all PD has been processed a log entry is written. The entry contains time, data subject and type of PD accessed as well as an indicator of the success of the decryption. Again, if writing the log is successful the document is passed to the nursing service.

4.3 User Management

In order to manage the users of the inSel system, an integrated rights management component is provided, having the functionality to identify the users on the inSel system (authentication), managing their roles and access control permissions (administration) and to store and manage the encryption keys for the individual users. Configuration of the rights management component will be dedicated to the responsible data controller (in the given scenario to the nursing service). This is done exclusively and for a given instance of inSel in an irreversible manner. Neither the data processor (here the storage service), nor other data controller or the inSel provider himself have the permission to interfere.

5 Discussion

In a productive environment it needs to be assured that the inSel system's availability is guaranteed and that the system itself is capable of scaling to the actual needed demands. While the prototype implementation does not yet allow for scaling and availability mechanisms (e.g. replicating or synchronizing VMs), future extensions need to include appropriate mechanisms for load balancing and cross machine handling of VMs.

The configuration of the function ***identify*** is a vital part of our system. In its current implementation the system takes the RelaxNG schemas (stating which elements are valid) and the categorization of PD (which of the valid elements are PD) from the nursing service, which is assumed to be trustworthy. But there is no guarantee that the type of text elements' content matches its assumed type. Malfunctions or misconfigurations of the nursing services' systems might lead to unwanted disclosure of PD.

The configuration of the rights management is done via the user interface, which is designed under the principles of usable security. It allows data subjects to set access rights to their PD according to their preferences (within legal boundaries). Misconfiguration, leading to unwanted disclosures, is minimized by applying usable security. The rights management, though, is only capable of controlling access

to data stored at the service provider and does not control the usage of PD after it has been transmitted to legitimate data consumers. Usage Control principles could be used at data consumers to counter this problem, which is beyond the scope of this paper.

5.1 Security Discussion

Our approach involves three parties: provider of the storage service, provider of the inSel system, and authorized service users (including data subjects and data consumers). Assuming that there are non-authorized externals, we have four parties to look at with regard to threads against the data subject's privacy. In the following outgoing threads of each of this parties will be discussed.

The provider of the storage service might get non-authorized insights into the stored data; in particular, he might do data mining to increase his knowledge on the data subject beyond the stored encrypted data. Since he stores the data, this is reasonable. In our approach, the personal data that is not required for the storage service is encrypted by the inSel service. This limits the insight into the actual stored data and thus, limits the gain of data mining. The provider of the inSel service might get non-authorized access by sniffing through-going traffic. In our approach, de- and encryption is done inside the inSel Core and incoming and outgoing connections are encrypted via HTTPS. Thus, direct access to the VM is required to gain knowledge of PD. The Trusted Server itself does not allow a direct access to the VMs, as this functionality is centrally deactivated on the system. The only access is via defined interfaces, allowing the remote management of VMs, but not the access to them. Thus, there is no feasible way to read data passing the inSel system.

It might be possible that an authorized user is getting non-authorized access. In our approach, we are using a strict regulation of access control to limit access to only authorized data. Further, the safe use of secure credentials like the new German electronic ID card is supported. Outgoing threads from non-authorized externals are counted by using a very restrictive design of the access points. The inSel service uses proven security mechanisms, such as HTTPS and XML-Encryption with AES, in order to protect and combine it with secure credentials and TPM-based virtual channels. In our approach, we introduced security mechanisms to increase the effort to get non-authorized access to the data outbalance the value of the data. As any other security gateway, our approach is limited to the capability of current cryptographic algorithms and security protocols.

For supporting compliance checks, our approach comes with a secured and predefined interface. Access to the logging data is clearly controlled and regulated. Thus, it is resilient to data mining attacks and non-authorized access. The predefined interface also supports the principle of data avoidance. E.g. for the use case of the inspection of commissioned data processing the access is limited to the documented operations of a specific session ID and can only be done by specific authorized persons. The provided information is: timestamp, ID (linkable pseudonym) of the subject, message type (providing object and purpose of the access), provided

data types (not the data itself), ID (linkable pseudonym) of the data subject. Anyhow, for every new compliance check the interface has to be extended. This may be a drawback for the maintainability and configurability of our approach. But it is reasonable that new compliance checks are not that often introduced, since the underlying regulations and laws do not change that often.

The possibility of existing side channel attacks resulting from the altered architecture or necessary network metadata has not been discussed properly in a systematic way and is topic for future work.

6 Conclusion

In this paper, we presented a schema called inSel that combines the abilities of a security gateway and an automated compliance checking system in a single system. By optimizing the solution for protecting personal data in storage services with multiple accessing parties, the complexity of data protection and compliance controlling could be reduced to a simplified but powerful gateway system. It allows the separation of security functions from data processing and storage functions. Using encryption to protect data from unauthorized usage, inSel operates as a gatekeeper for storing and retrieving data from the data storage. It combines strong access control mechanisms with consistent and audit proof logging. inSel itself is protected from external manipulation by using a TURAYA-based security kernel and a TPM (Trusted Platform Module) as hardware-based security anchor. Further, the access to the logging can be restricted to predefined and use-case related interfaces introducing robust and simplified semi-automated compliance checks. Based on the given scenario, three use cases for compliance control were implemented, proving inSel as a real-world solution.

While our approach clearly supports data subjects and controlling agencies, there are even benefits for companies complying with data protection laws as they are able to prove their compliance to e.g. independent certification authorities and thus obtaining certificates which in turn can be used as a selling argument or to gain customers' trust.

inSel is able to reduce the effort of the compliance check of a complex system by reducing the relevant components that have to be audited. Its applicability has been shown successfully by a prototypic implementation in the scenario of mobile time and service recording of nursing services. In this case, the processing system only stores and prepares personal data for retrieval, allowing for a higher efficiency when conducting compliance checks. For future uses, the system could be adapted and extended, so it can be used in other application domains, such as eMail or many eHealth applications.

Acknowledgements

This work was part of the inSel-project "Informationelle Selbstbestimmung in Dienstnetzen", funded by the German Federal Ministry of Education and Research (BMBF) within the support code 01IS08016(A,B,C). The authors are responsible for the content of this article.

References

[Ac08] Accorsi, R.: Automated Privacy Audits to Complement the Notion of Control for Identity Management. In: Policies and Research in Identity Management. Springer Boston, 261, pp. 39-48, 2008.

[As04] Ronald Ashri, Terry Payne, Darren Marvin, Mike Surridge und Steve Taylor: Towards a Semantic Web Security Infrastructure. AAAI Spring Symposium on Semantic Web Services. Stanford Univ, 2004.

[Ce07] Cederquist, J. G. et al.: Audit-based compliance control. In: International Journal of Information Security. Springer, 2007, vol. 6, pp. 133-151.

[Di82] Dijkstra, Edsger W.: On the role of scientific thought. In: Dijkstra, Edsger W.: Selected writings on Computing: A Personal Perspective. Springer-Verlag New York, Inc., 1982, pp. 60–66.

[EP10] emSCB Project: Towards Trustworthy Systems with Open Standards and Trusted Computing. www.emscb.de. Accessed 01.07.2010.

[EW07] Etalle, S., Winsborough, W. H.: A posteriori compliance control. In: SACMAT '07: Proceedings of the 12th ACM symposium on Access control models and technologies. ACM, 2007, pp. 11-20.

[GH06] Nils Gruschka, Ralph Herkenhöner, Norbert Luttenberger. WS-SecurityPolicy Decision and Enforcement for Web Service Firewalls. In Proceeding IEEE/IST Workshop on Monitoring, Attack Detection and Mitigation, 19–25, Tübingen, Germany, 2006.

[HJ08] Höhn, S. & Jürjens, J. Rubacon: automated support for model-based compliance engineering ICSE 08: Proceedings of the 30th international conference on Software engineering, ACM. New York, NY, USA, 2008, 875/878.

[HJ10] Ralph Herkenhoener, Meiko Jensen, Henrich Poehls and Hermann De Meer. Towards Automated Processing of the Right of Access in Inter-Organizational Web Service Compositions. Proc. of the IEEE 2010 Int'l Workshop on Web Service and Business Process Security (WSBPS 2010), 2010.

[HMP04] Alan R. Hevner, Salvatore T. March und Jinsoo Park. Design Science in Information Systems Research. MIS Quarterly, 28(1), pp. 75–105, 2004.

[Sc02] Scott, David and Sharp, Richard: Abstracting application-level web security. In Proceedings of the 11th international conference on World Wide Web, Honolulu, Hawaii, USA, 2002, pp. 396-407.

[SS94] Sandhu R. & Samarati P.: Access control: Principles and practice. In: IEEE Communications Magazine, 32(9), pp. 40–48, 1994.

[PS04] Park, J. & Sandhu, R.: The $UCON_{ABC}$ usage control model. In: ACM Transactions on Information and System Security. ACM, 2004, vol. 7, pp. 128-174.

[Po99] Povey, D.: Optimistic security: a new access control paradigm. In: NSPW '99: Proceedings of the 1999 workshop on New security paradigms. ACM, 2000, pp. 40-45.

[TCG05] Trusted Computing Group: TPM main specification. Main Specification Version 1.2 rev. 85. Trusted Computing Group, 2005.

[Us04] Andrzej Uszok, Jeffrey M. Bradshaw, Renia Jeffers, Austin Tate, and Jeff Dalton: Applying KAoS Services to Ensure Policy Compliance for Semantic Web Services Workflow Composition and Enactment. ISWC 2004, LNCS 3298, pp. 425–440, 2004. Springer-Verlag Berlin Heidelberg, 2004.

[Va09] Hal R. Varian : Economic Aspects of Personal Privacy. Internet Policy and Economics 2009, Part 4, 101-109