A Framework for Enforcement of Privacy Policies[1]

Ragni Ryvold Arnesen and Jerker Danielsson
Norsk Regnesentral / Norwegian Computing Center
{Ragni.Ryvold.Arnesen, Jerker.Danielsson}@nr.no

Abstract

This paper presents an ongoing work on a framework for enforcement of privacy promises, policies and regulations. The aim is to develop an open framework that can form a basis for discussion of such enforcement, and deployable components that enable integration with legacy systems as well as state-of-art development environments.

Keywords: Privacy framework, policy, enforcement, privacy legislation.

1 Introduction

There are mainly two approaches to the implementation of privacy-enhancing or privacy-assuring technologies and processes. One is to minimize the amount of personally identifiable data through pseudonymisation or anonymisation, or by simply not collecting any data at all. The other approach is to assure that the privacy agreement, e.g. codified in P3P [22], that both data subject and data collector have consented to is enforced. There is no conflict between these approaches. Both are important.

There are circumstances where processing of personal data is useful or necessary. In some cases personal data must be collected due to legislation, or because it is necessary in order to provide some public service. In other cases the collection of personal data may be of benefit to both the data collector and the data subject. An example of such a case is the possibility for the data collector to customize offers to the data subject based on her/his interests, history and current context, e.g. location.

The data subject, i.e. the person whose identity is, or may be, connected to the data, usually has little control over information collected and stored. The notion of privacy when personal data is collected implies some form of trust in the data collecting entity. Systems for mandatory and automated enforcement will contribute to the establishment of such trust, as will a conceived high level of information security. There is in our view a need for an open framework for enforcement of privacy regulations and the privacy promises made by data collectors. Such a framework can form a basis for discussion of technology and processes, and a basis for development and deployment of enforcement functionality.

This paper describes our ongoing work on designing and implementing such a privacy enforcement framework, which comprises functionality necessary for adherence to privacy agreements pertaining to collected data, as well as applicable privacy regulations.

[1] The work presented in this paper is fully funded by the Norwegian Research Council through the research project “Personalized Internet-Based Services and Privacy Protection.”

In addition to privacy enforcement functionality, there is a need for processes that guide the integration of privacy protecting functionality with legacy systems and existing business processes, and that guide the design of new privacy-enabled applications and business processes. We acknowledge this need, but it is not addressed further in this paper.

The remainder of this paper is organised as follows. Chapter 2 gives a brief account of related work in privacy frameworks. Chapter 3 gives an overview of the framework and some of the design principles and rationale behind it, and chapter 4 describes the framework elements in more detail. Finally, our current activities and future plans are described in the conclusion.

2 Related work

There are other ongoing efforts in defining privacy frameworks. The most noticeable being the Privacy Framework [11] developed by the International Security, Trust & Privacy Alliance (ISTPA) and the Enterprise Privacy Architecture (EPA) [15] developed by IBM Research.

The ISTPA Privacy Framework defines a number of services and capabilities that implement the fair information practices, see [18]. A capability is implemented through the invocation of multiple services. The services and capabilities provide functionality that supports both the data subject (e.g. preference definition and validation of preference) and the data collector (e.g. auditing).

EPA is a methodology for introducing privacy awareness, and privacy services and processes into enterprises. It consists of four building blocks: Privacy regulation analysis, management reference model, privacy agreement framework, and technical reference architecture. The privacy regulation analysis identifies and structures applicable regulations in a unified terminology and relates these regulations to the personal data held by the organisation. The management reference model defines processes necessary for a comprehensive privacy management program. The privacy agreements framework is a methodology for privacy enabling business processes. It results in a model of the personal data used in the process, the privacy-relevant players and operations of the process, as well as the rules that govern these operations. Finally, the technical reference architecture is a model of a system for the enforcement of privacy promises. It defines a management system, an audit console and a reference monitor.

3 The privacy framework

The framework presented here is inspired by the life cycle of personal data. That is, collection, various forms of processing (e.g. disclosure to third parties), and finally deletion or depersonalisation. The framework is intended to function as a layer of control between personal data on the one hand, and services accessing and collecting personal data on the other hand.

The framework consists of framework elements (e.g. Access) that together provide the functionality necessary for enforcement of applicable regulations and privacy agreements reached in connection with data collection. Each framework element is composed of components that support the implementation of its functionality (e.g. Reference Monitor).

The framework elements form a basis for discussion of the functionality of a general framework for enforcement of privacy policies. Moreover, the components of the framework elements should be deployable, meaning that they should enable integration with legacy systems and state-of-art development environments. Achieving these two properties of the framework, i.e. basis for discussion and deployable components, is the main objective of our work.

A basic requirement for the framework is that there must be a clear separation of functionality and responsibilities between the framework elements. Further, the framework should be complete in the sense that it should address all functionalities necessary to enforce local privacy policies, legal requirements, and agreements made between data subjects and data collector.

In addition, the framework must have support for traditional security mechanisms, such as authentication of users, and protection of confidentiality and integrity of information. How these mechanisms are integrated into the framework is not addressed in this paper.

4 Framework elements

The framework manages Personal Data Bundles that contain personal data, and the Agreement and access history pertaining to the personal data. IBM Research calls bundling of data and policy the "sticky policy paradigm" [14].

Personal Data Bundles can be introduced into the framework in two different ways: Personal data with pertaining Agreements can be imported from a third party (handled by the Data Import Manager of the Communication element), or personal data can be collected directly from the data subject. From the view of the framework, collected data is assumed to be packaged in Personal Data Bundles. It is further assumed that the Agreement pertaining to the collected personal data is derived from the privacy promise of the data collector and the privacy preference of the data subject. If the privacy promise and the privacy preference are codified in machine-readable formats, such as P3P [22] and APPEL [21], software agents can be used to automatically negotiate and consent to the Agreement on behalf of the parties.

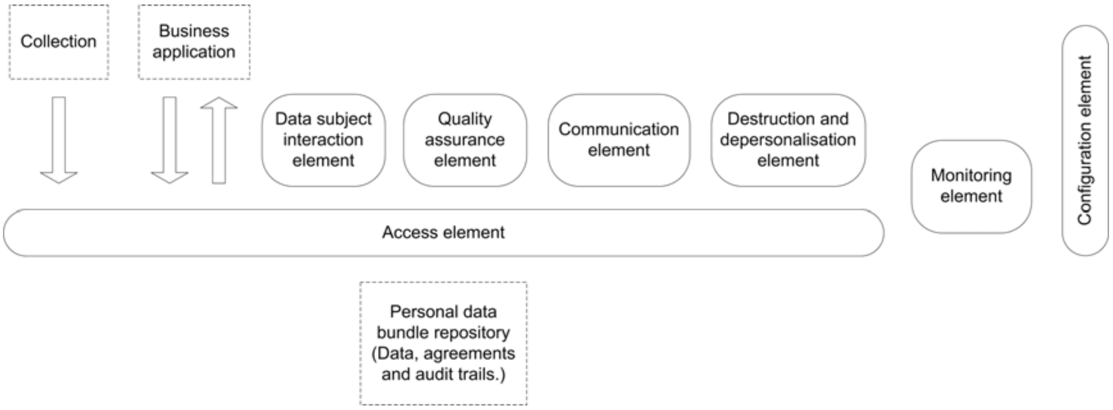
Collection
Business application
Data subject interaction element
Quality assurance element
Communication element
Destruction and depersonalisation element
Monitoring element
Configuration element
Access element
Personal data bundle repository (Data, agreements and audit trails.)

Figure 1 Overview of the framework.

Figure 1 illustrates an overview of the framework. The Access element controls the flow of personal data in its position between the personal data bundle repository, and collection modules, business applications and the other components of the framework. The audit trails that all components generate are analysed by the Monitoring element.

The Configuration element is responsible for assuring that the configuration of the framework elements complies with applicable privacy regulations, that it is consistent with the local privacy policy and that the published privacy promises are consistent with the configuration of the framework.

The following sections describe the Agreement and the Personal Data Bundle concepts, and the framework elements and their components, in more detail.

4.1 Agreement

An Agreement is a set of rules that determine how the personal data the Agreement pertains to can and should be used, and that both the data subject and data collector have consented to. An Agreement is derived from the privacy promise of the data collector and the privacy preference of the data subject.

The data collector's privacy promise forms an important base for all Agreements that the data collector makes. The privacy promise is based on an analysis of the need for personal data to conduct the business processes of the data collector. The data subject's privacy preference defines the Agreements that the data subject is willing to consent to.

The attributes of the Agreement can be divided into two categories: attributes that can be suggested by both the data subject and data collector (they can negotiate the attribute) and those that can only be determined by one party. For example only the data collector can determine the purpose of collection.

Attributes that may be part of an Agreement include:

- Purpose – Why is the data collected? The collected data must only be used for the stated purpose.
- Subject access – Can the data subject access its personal data and the access/usage history of its personal data?
- Disputes – How are disputes solved?
- Remedies – How is a breach of agreement handled?
- Obligations – When performing certain actions, the data processor may be required to take further steps. E.g. if the data is accessed the data subject of the data must be notified.
- Retention – How long will the data be retained? Will it be destructed or depersonalised?
- Disclosure – To which third parties will the collected data be disclosed?

4.2 Personal Data Bundle

A Personal Data Bundle contains personal data and the Agreement regulating how the personal data can and should be used. The access/usage history of the data is also included in the Personal Data Bundle. The Personal Data Bundle may include signatures and the credentials of the data subject and the data collector, to bind the Agreement to the two parties. The credentials of the data subject may also be used in the implementation of subject access (see section 4.6.1).

4.3 Configuration

The Configuration element encompasses functionality for generation of the other framework elements' configuration and functionality for generation of privacy promises. This functionality is automated or semi-automated, in the form of consistency checking, or a combination of both. For example a privacy promise may be generated automatically from the configuration of the framework (see 4.3.3) or it may be constructed more or less manually with the support of consistency checking (see 4.3.4), verifying that the constructed privacy promise is consistent with the configuration.

The automated and/or semi-automated generation of the framework's configuration is based on the local privacy policy and applicable regulations, see Figure 2. The local privacy policy is based on an analysis of the processes of the organisation, their need for collecting and processing personal data, the players involved in the tasks of the processes, etc.

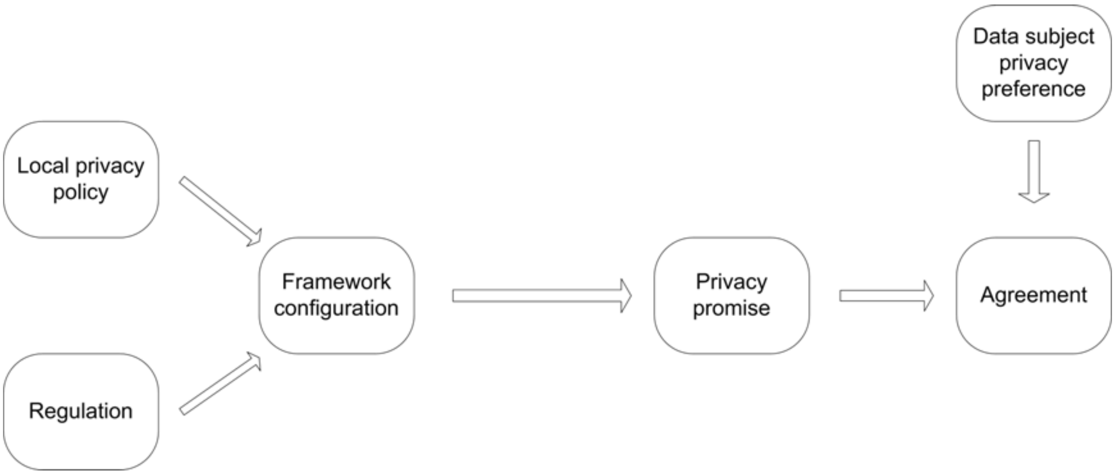

Figure 2 Overview of the configuration process and the making of an Agreement.

In Europe, most national privacy legislations are based on EU directives [5, 6], which in turn are based on the OECD Guidelines [17]. The privacy laws and regulations give rules as to how different types of information can or must be handled, including conditions for, and obligations following from, such handling. To enable automated or semi-automated integration of regulations into the configuration of the framework, the regulations must be codified by a human interpreting them. However, the laws are usually complex and difficult to comprehend, and the process must be repeated whenever the laws change. Hence, this task presents a great challenge. But such an encoding is portable and resource saving, in the sense that once it is defined, it may be used by any system that understands the language the rules are encoded in.

We are currently working on such an encoding of the Norwegian privacy law [16]. This is not an easy task due to the frequently intricate formulation of the paragraphs. Some rules say what you are not allowed to do, some say what you may do, and some say what you must do. There is a multitude of special cases, exceptions to the rules (positive or negative), and exceptions to the exceptions. There are listings of conditions and obligations, but which ones are required often depends on which other conditions are fulfilled and/or which are not.

In addition to the configuration of the framework, the flow of personal data is regulated by the Agreements in the Personal Data Bundles. The Agreements supplement the configuration in the sense that an Agreement must not contradict the configuration.

4.3.1 Configuration Generator

The Configuration Generator provides automated and semi-automated support for the generation of the framework configuration, according to the local privacy policy and codified regulations. The functionality of the Configuration Generator overlaps the functionality of the Legal Compliance Analyser, in the sense that if an ideal Configuration Generator exists there is no need for a Legal Compliance Analyser.

4.3.2 Legal Compliance Analyser

The Legal Compliance Analyser controls that all framework elements are configured in accordance with applicable regulations. The Legal Compliance Analyser is not necessarily an automated tool that analyses everything and outputs an “ok” or identifies were the problems/inconsistencies are. It can also be more of a questionnaire that assures that all relevant checkpoints are gone through so that the configuration complies with regulations and that there are no inconsistencies between regulations and configuration.

4.3.3 Privacy Promise Generator

The Privacy Promise Generator generates a privacy promise based on the framework configuration. The privacy promises define the set of Agreements that the data collector is willing to accept. That is, the privacy promise defines the base Agreement that the data subject may modify through opt-ins and opt-outs defined in the privacy promise.

The functionality of the Privacy Promise Generator overlaps the functionality of the Privacy Promise Analyser. An example of such a generator can be found in [13].

4.3.4 Privacy Promise Analyser

The Privacy Promise Analyser checks that the privacy promises of the organisation and the resulting set of possible Agreements are consistent with the configuration. Moreover, before personal data is imported (by the Data Import Manager, see section 4.8.1), the Privacy Promise Analyser may be used to control that the set of possible Agreements of the imported data is consistent with the local configuration.

4.4 Access

The Access element is responsible for keeping track of all personal data that is held by the organisation and for regulating the access to this data in accordance with the configuration and the Agreements pertaining to the data.

4.4.1 Reference Monitor

The Reference Monitor denies or accepts access to operations on personal data to requestors internal to the organisation, or internal to the domain if the framework is implemented as a security barrier between domains. External requestors or third parties request access through the Disclosure Controller (see 4.8.2).

The Reference Monitor is essentially an access control mechanism, but the type of access control necessary to enforce privacy policies is different from other access control models, such as the well-known Bell LaPadula model. This is mainly because the *purposes* of the data processing, as well as other context information, are important in the privacy case. These differences are discussed in [7], which also presents a formal access control model for the enforcement of privacy policies. Other examples of work on privacy enabled access control are the control service in the ISTPA framework [11] and the privacy policy model presented in [12].

The Reference Monitor bases its decisions on authorisation rules written in a machine-readable policy specification language, e.g. EPAL. EPAL is a formal language for writing “privacy authorization rules that allow or deny actions on data-categories by user-categories for certain purposes under certain conditions while mandating certain obligations.” [2]

Vocabularies need to be built to encompass the specific data- and user-categories, actions, purposes, conditions and obligations pertaining to the system and policies in question. For instance, codifying the Norwegian privacy legislation will require a vocabulary containing the condition “informed consent from the data subject”. Defining useful vocabularies for actions and purposes, and mapping the applicable policies to these, will require analysis of the operations performed by the applications accessing data through the framework. Suitable vocabularies will facilitate efficient and fine-grained access control.

To prevent aggregation or inference of data, the Reference Monitor may also base its decisions on special context conditions like the access history of the data and/or data requestor. In many cases the application does not really need access to the actual personal data; access to the relationships between data may be enough. In such cases, the application should only get access to pseudonymised data. The Reference Monitor may also implement functionality to reduce the accuracy of data, e.g. the granularity of location data, based on the authorisation rules.

4.4.2 Personal Data Broker

The Personal Data Broker acts as a librarian, i.e. handles requests for personal data and locates the requested data. Based on the Reference Monitor's access decisions, it delivers the personal data requested from the different data repositories where personal data is stored. In addition, it assures that documentation is maintained over the personal data held by the organization.

The Personal Data Broker may additionally implement the identity protector concept presented in [10]. The identity protector creates one or several pseudo-domains in which data subjects are known under pseudo-identities. Only the identity protector knows the mapping between identities and pseudo-identities, and the mapping between pseudo-identities, and is thus able to reverse the process and retrieve identities from pseudo-identities.

The Personal Data Broker also triggers the events needed to ensure that any obligations following from the requested type of access are fulfilled.

4.5 Monitoring

The Monitoring element monitors and analyses the audit trails generated by the other elements. Other elements, in particular the Access and Communication elements, implement *proactive* mechanisms whose purposes are to prevent users from doing what they are not allowed to do according to the framework configuration and the Agreements of the accessed data. The monitoring mechanisms, on the other hand, are *reactive* in the sense that they may enable detection of a policy breach and cause some reaction after the breach happened. The proactive mechanisms are the front line of security mechanisms, but the reactive mechanisms are also important, particularly to build and maintain the users' trust in the system.

Monitoring mechanisms are important parts of an internal control system, which is mandated by Norwegian legislation ([16], §14).

4.5.1 Audit Manager

The Audit Manager supports auditing (manual and semi-automated) of the audit trail that the components generate. It implements functionality for searching and reviewing the audit trail.

4.5.2 Privacy Violation Detector

The Privacy Violation Detector continually monitors access to personal data and detects misuse and/or anomaly behaviour. Anomalies can be detected using e.g. data mining methods, see [1] for a survey of such methods.

4.5.3 Remote Privacy Audit Manager

The Remote Privacy Audit Manager provides seal issuing authorities and/or official authorities (e.g. The Data Inspectorate in Norway) with the possibility to remotely monitor and review the site.

4.6 Data Subject Interaction

The Data Subject Interaction element provides access to personal data, access/usage history and Agreements to data subjects. It also provides mechanisms for data subjects to submit complaints, and support for resolving these complaints.

4.6.1 Subject Access Manager

The Subject Access Manager manages data subjects' requests for reviewing and updating their personal data. It may also provide access to the access/usage history of the subjects' personal data. In addition, the data subjects may have the possibility to modify the Agreements that pertain to their personal data.

Subject access might improve the quality of the personal data. If the data subject has access to its personal data he/she might assure that it's correct, especially if there is some sort of incentive for the data subject to do so. Additionally, subject access may provide a powerful tool for detecting agreement violations. The probability of detecting violations increases if data subjects review the access/usage history of their personal data.

Subject access can be managed through electronic means (e.g. the Internet) or through traditional mail delivery. In any case, subject access sets authentication requirements. If the authentication is not strong enough subject access will instead contribute to the impairment of the privacy of the data subjects.

Norwegian legislation gives the data subjects rights to information about the nature of the personal data processing and what information pertaining to the data subject is stored ([16], §18). Data subjects also have a right to demand correction of incorrect or incomplete data, or in some cases also blocking or complete erasure of data (§27).

4.6.2 Dispute Manager

The Dispute Manger offers support in resolving disputes. It offers different mechanisms (e.g. web, email) for submitting complaints to the data collector and/or some other relevant authority. In addition, it may provide support for semi-automatic processing of complaints and compilation of reports on the use of the personal data pertaining to the complaining data subject.

4.7 Quality Assurance

The Quality Assurance element encompasses functionality that aims at upholding the correctness of the stored personal data. Quality assurance is also provided by the data subjects through the Subject Access Manager, but the data collector also has a responsibility and interest in maintaining data quality.

Norwegian legislation demands that the data controller (i.e. the entity storing and processing the personal data) ensures that personal data processed are accurate and up-to-date, and also adequate, relevant and not excessive in relation to the purpose of the processing ([16], §11). Internal control procedures must be implemented to ensure quality of data (§14). If incorrect, incomplete or unauthorised data have been processed, the data controller shall to the extent possible ensure that the error does not have any effect on the data subject, for instance by notifying recipients of disclosed data (§27).

4.7.1 Subject Preference Register Monitor

This component monitors customer preference registers and assures that the framework is compliant with the information in such registers. One example of such a register is the Norwegian reservation register against direct marketing [19]. Here the users may request that their address is removed form address lists used in direct marketing (with a few exceptions), and companies performing such marketing must update their address lists at least every three months.

4.7.2 Validator

The Validator component checks the consistency of incoming data from data subjects or third parties against defined bounds and heuristics. It can also check input data against data collected previously and external sources. The bounds and heuristics should be defined when the data collection or communication is defined.

4.8 Communication

The Communication element provides functionality for importing and exporting personal data in and out of the domain controlled by an instance of the framework. Automatic export and import of data is dependent on standardised exchange protocols. A proposal of such a standard is the Customer Profile Exchange (CPExchange) standard [3], but its adoption has been limited.

4.8.1 Data Import Manager

The Data Import Manager controls the import of personal data and possibly linking and matching of locally controlled personal data with the imported personal data. During import it controls the Agreements of the imported data to verify that the import is allowed. It guarantees the preservation of the Agreements pertaining to the imported data. In addition, it controls that any linking and matching is conducted according to the Agreements pertaining to the data involved in the operation, and that the resulting data is bundled in new Personal Data Bundles with updated Agreements.

Mergers, acquisitions and internationalisation can create a wish to link and match or integrate databases containing personal data. The trend towards one-stop-shop services in the public sector also actualizes the issue of linking and matching. Proper control of linking and matching is essential since separation of data repositories is fundamental to privacy protection.

Norwegian legislation states that data subjects have a right to be notified when data is collected from other parties ([16], §20). Also, if the data controller contacts the data subject or makes decisions regarding the data subject on the basis of personal profiles, the controller must inform the data subject of the sources of the data (§21).

4.8.2 Disclosure Controller

The Disclosure Controller controls the disclosure of personal data to third parties outside the framework's domain. It determines to whom personal data may be passed and under what conditions based on the Agreements of the exported data.

Internationalisation and outsourcing of functions, like Customer Relationship Management (CRM), are two trends that contribute to the transfers of consumer and employee data by businesses. Disclosure is complicated by the fact that different countries or regions have different privacy legislation and some have none. According to Norwegian legislation ([16], §§29-30), personal data may in general only be transferred to countries that ensure an adequate level of protection of the data.

4.9 Destruction and Depersonalisation

The Destruction and Depersonalisation element is responsible for the last step of the life cycle of personal data. After this step the data should no longer be considered personal data, with the exception of pseudonymised data where the depersonalisation can be reverted.

An important principle in Norwegian legislation is that personal data may not be stored longer than necessary for the purpose ([16], §11, §28).

4.9.1 Destruction Controller

The Destruction Controller is responsible for assuring that any commitment to destruction of personal data is fulfilled in time. The data should be destructed in such a way as it is made irretrievable and unreadable.

4.9.2 Depersonalisation Controller

The Depersonalisation Controller is responsible for assuring that any commitment to depersonalisation of personal data is fulfilled in time.

Depersonalisation can be reversible (pseudonymisation) or non-reversible (anonymisation). See [8] for a definition of anonymity and pseudonymity.

How data should be depersonalised is not always straightforward. For example if the record contains name, employer and year of birth it may not be enough to delete or pseudonymize the name field. It may still be possible to identify the person that the record pertains to, especially if one has access to supplementary information that could be linked and matched with the "depersonalised" record. That is, the risk of reidentification depends on the size of the dataset and the entropy of the remaining attributes [9].

5 Conclusion

In this paper we have presented our proposal for an open framework for enforcement of privacy policies. The framework comprises functionality to enforce local privacy policies, privacy legislation and agreements reached between data subject and data collector.

We are currently working on a prototype implementation of some of the components of the Configuration, Access and Monitoring elements in the form of a Java framework and plan to experiment with implementations of other components as well. We are also working on codifying the Norwegian privacy legislation into a machine-readable format. Meanwhile, we will continue to refine and develop the framework, and we are also looking into future possibilities for realising stronger enforcement mechanisms. For example, the technology proposed by the Trusted Computing Platform Alliance (TCPA) [20] may provide possibilities for forcing the recipient of personal data to act in accordance with the agreement bundled with the data.

6 References

[1] Aas K., Huseby R., and Thune, M., *Data Mining - A Survey*. Report No. 942, June, 1999. ISBN 82-539-0426-6

[2] Ashley, P., Hada, S., Karjoth, G., Powers, C., Schunter, M. (ed.), *Enterprise Privacy Authorisation Language (EPAL)*, IBM, 2003. Available via
http://www.zurich.ibm.com/security/enterprise-privacy/epal/

[3] *Customer Profile Exchange (CPExchange) Specification*, version 1.0, October 2000, Available via http://www.cpexchange.org/standard/

[4] Datatilsynet (The Data Inspectorate), *Veiledning i informasjonssikkerhet i kommuner og fylker*, 1999, http://www.datatilsynet.no/dtweb/attachment/783/Kommuneveiledning.pdf

[5] *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 L 281, 23/11/1995, pp. 31-50. Available via http://europa.eu.int/eur-lex/en/index.html

[6] Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications). Official Journal L 201, 31/07/2002, pp. 37-47. Available via http://europa.eu.int/eur-lex/en/index.html

[7] Fisher-Hübner S., Ott, A., *From a Formal Privacy Policy Model to its Implementation*, National Information Systems Security Conference (NISSC 98), 1998. Available at http://www.rsbac.org/niss98.htm

[8] Köhntopp M. and Pfitzmann A., *Anonymity, Unobservability and Pseudonymity – A Proposal for Terminology*, Draft v0.12, Available at http://123.koehntopp.de/marit/pub/anon/Anon_Terminology.pdf

[9] Fisher-Hübner S., *Privacy-Enhancing Technologies*, Karlstad University, Department of Computer Science, PhD course, Slides session 2, Available via http://www.cs.kau.se/~simone/kau-phd-course.htm

[10] Hes, R. and Borking, J. (eds.), *Privacy-enhancing technologies: The path to anonymity. Revised edition.* ISBN: 90-74087-12-4. Registratiekamer, The Hague, August 2000

[11] International Security, Trust & Privacy Alliance (ISTPA), *Privacy Framework*, Version 1.1, ISBN: 0-9721484-1-8, 2002

[12] Karjoth G., Schunter M., *A Privacy Policy Model for Enterprises*, 15th IEE Computer Security Foundations Workshop, June 2002

[13] Karjoth G., Schunter M. and Van Herreweghen E., *Enterprise Privacy Practices vs. Privacy Promises - How to Promise What You Can Keep*, 4th IEEE International Workshop on Policies for Distributed Systems and Networks (Policy '03), Lake Como, Italy, June 2003

[14] Karjoth G., Schunter M. and Waidner M., *Platform for Enterprise Privacy Practices: Privacy-enabled Management of Customer Data*, 2nd Workshop on Privacy Enhancing Technologies, 2002

[15] Karjoth G., Schunter M. and Waidner M., *Privacy-enabled Services for Enterprises*, IBM Research, Zürich Research Laboratory, Switzerland, January 2002

[16] *LOV 2000-04-14 nr 31: Lov om behandling av personopplysninger (personopplysningsloven)*. Available at http://www.lovdata.no/all/hl-20000414-031.html. Norwegian privacy law. An unofficial English translation is available at http://www.datatilsynet.no/lov/loven/poleng.html

[17] *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*. Available at http://www1.oecd.org/publications/e-book/9302011E.PDF

[18] *Privacy Online: A Report to Congress*. Federal Trade Commission, June 1998. Available at http://www.ftc.gov/reports/privacy3/priv-23a.pdf

[19] Reservasjonsregistret i Brønnøysund (The Brønnøysund reservation register), http://www3.brreg.no/oppslag/reservasjon/om_resreg.jsp

[20] The Trusted Computing Platform Alliance (TCPA), http://www.trustedcomputing.org/

[21] W3C, *A P3P Preference Exchange Language 1.0 (APPEL1.0)*, W3C Working Draft, April 2002, Available at http://www.w3.org/TR/P3P-preferences/

[22] W3C, *The Platform for Privacy Preferences 1.0 (P3P1.0) Specification*, W3C Recommendation, April 2002, Available at http://www.w3.org/TR/P3P/