

Defining Consent and Revocation Policies

Ioannis Agrafiotis, Sadie Creese,
Michael Goldsmith, Nick Papanikolaou

International Digital Laboratory
University of Warwick
Coventry, UK

e-mail: {I.Agrafiotis, S.Creese,
M.H.Goldsmith, N.Papanikolaou}@warwick.ac.uk

Marco Casassa Mont, Siani Pearson
Systems Security Lab
HP Labs
Bristol, UK

e-mail: {marco.casassa-mont, siani.pearson}@hp.com

Abstract— In this paper we present the notion of a *consent and revocation policy*, as it has been defined within the context of the EnCoRe project. A consent and revocation policy is different to a privacy policy in that it defines not enterprise practices with regards to personal data, but more specifically, for each item of personal data held by an enterprise, what consent preferences a user may express and to what degree, and in what ways he or she can revoke their personal data. This builds on earlier work on defining the different forms of revocation for personal data, and on formal models of consent and revocation processes.

Keywords - *privacy policies, policy hierarchy, policy refinement, conceptual model*

I. INTRODUCTION

Alan Westin has advanced a view of privacy as control over one's personal information [1]. In this view, an individual has privacy if he or she can exercise control over the use and flow of such data in society. Actual control over personal data shared by individuals today, especially over the Internet, is significantly lacking even though there exist laws and regulations in which such control is enshrined as a basic right. We take the view that control over personal data in our modern society amounts to two main things: (a) having at one's disposal technical means of giving and withdrawing consent for specific uses of personal data by any party interested in that data, and (b) having mechanisms that enforce such consent or non-consent in every instance of data use. In addition, it is desirable to have ways of ensuring that users of (a) are effectively supported in understanding how to effectively *use* such technologies. Implementing suitable mechanisms of *revocation* for personal data poses particular challenges (where revocation is understood most generally as a change of consent, namely, a partial or complete withdrawal of consent), both technological, procedural and legal, which have yet to be fully met.

All modern enterprises rely, in various ways, on collection and processing of personal data about their customers. However, a recent rise in identity theft and other related crimes has made people increasingly aware of the perils of sharing personal data, and enterprises are now required to develop and disclose detailed privacy policies. While privacy policies mostly describe an enterprise's personal data handling practices, there is no commonly accepted way of stating clearly to customers whether (and which) mechanisms for granting and revoking consent for personal data are provided. In other words, what is desirable is a descrip-

tion of exactly which *controls*¹ an enterprise provides to its customers. In this paper we call such a description a *consent and revocation policy*.

As an example, such a policy might state that a customer must provide consent for personal health data to be used for only research purposes, for a fixed time period of 30 days. The policy might also state that consent will be requested for this data to be shared with third parties, although such consent is optional; finally, the policy might specify that the enterprise provides the right for individuals to delete the data from the enterprise's systems and from all third parties with whom such data has been shared.

Consent and revocation policies are intended to be used alongside enterprise security and privacy policies, and capture user preferences in a consistent way. They are intended to bridge enterprise requirements and user requirements, while allowing, on the one hand, enterprises to determine how much control is given to users and, on the other hand, users to determine what happens to their personal data.

There is good reason for consent and revocation policies to be expressed in a machine-readable form, in a manner analogous to website privacy policies [4, 5], for then enforcement can be fully automated; our aim here is not to define the syntax of such policies, rather their structure and general characteristics. The development of enforcement mechanisms for consent and revocation is the object of the EnCoRe project [2, 3].

¹ Note that we use the term *control* in this context to refer to a tool, interface or other mechanism that enables individuals to manipulate the state of personal data held about them.

Consent and revocation appear in many different forms, and the relationship between the two concepts is rather subtle. For instance, they are not symmetric: it is possible for an individual to revoke personal data for which consent has not been given (in Section III we will refer to this as *consentless revocation*). Consent itself is not a simple yes/no (“store this data” or “do not store this data”), but has many subtly different gradations, depending on the type of data it refers to.

In Section II we define the notion of a consent variable, which is an attempt to quantify consent in terms of its constituent parts, while in Section III we review the different forms of revocation, as identified in earlier work [6]. These elements come together in Section IV, where we define consent and revocation policies. The final section describes future work and concludes.

II. FORMS OF CONSENT AND THE NOTION OF A CONSENT VARIABLE

With regard to his or her personal data, there are three principal actions for which an enterprise or other data collector requires the consent of an individual:

- **collection** of data,
- **processing** of data,
- **sharing** of data.

Collection of data refers to the initial process by which data is acquired and stored on the enterprise’s information system. Processing includes any *access* of the data that has been collected and is characterised by a *stated purpose* (e.g. research, marketing, aggregation to derive average customer habits). Data may be shared – internally and externally (e.g. to third parties) so that it can be processed, often elsewhere than the site of data collection.

Thus, consent may be defined as a wish for a datum d to be collected, processed, shared, or any combination of the above. This definition is too coarse, however, for it does not account for subtleties such as these:

- It may be desired to restrict data collection so that it occurs only in one country, so that it is subject to that country’s privacy legislation.
- It may be desired that consent lasts for only a fixed period of time.
- It may be desired to restrict processing of data so that it is used for only certain stated purposes.
- It may be desired that the data is shared with only particular parties, and to banish uses of the data by others.

Thus we claim that consent is parameterised by certain quantities referred to as *consent variables*. As examples we consider: time t for which consent is granted, volume v of data for which consent is granted, the set S of allowed purposes (stated purposes for which consent is granted), the set Π of parties who may access the data.

Thus, consent is fully determined when there are specified:

- the **task** for which consent is given (collection, processing, sharing, or any combination thereof)
- for this task, the **values of the consent variables** of interest.

From this one can extract a mathematical definition of consent and develop a hierarchy of its different forms. Formal models of consent and the attendant logic are addressed in our other work [7].

III. FORMS OF REVOCATION

Revocation corresponds to the withholding or withdrawal of consent. It is manifested in its simplest form as deletion of data, although there are many variations of revocation, as listed below (from [6]):

1. **No Revocation At All:** Personal data remains static, and once it has been disclosed, it is either physically impossible to revoke (how could one ever revoke reputation) or prohibited for various reasons (e.g. law-enforcement, data from police’s DNA database).
2. **Deletion:** Data are completely erased and cannot be retrieved or reconstituted in any way.
3. **Revocation of Permissions to Process Data:** Data subjects withdraw consent that would enable an enterprise to process or analyse their personal data for a specified purpose.
4. **Revocation of Permissions for Third Party Dissemination:** Data subjects withdraw consent that would enable an enterprise to disclose information to a third party.
5. **Cascading Revocation:** This is a variation on any of the above kinds of revocation, whereby the revocation is (recursively) passed on to any party to whom the data has been disclosed. Through this mechanism, data subjects are able to revoke data by only contacting the enterprise that they disclosed their data to originally.
6. **Consentless Revocation:** Personal data for whose storage and dissemination no consent has been explicitly given by the user, but which may need to be revoked. Again, any of the fundamental types of revocation may be invoked. The need to revoke consentless data emerges mainly when a breach in privacy has occurred. In the italics below we describe a characteristic example of consentless revocation.
7. **Delegated Revocation:** This is a kind of revocation which is exercised by a person other than the individual concerned, such as an inheritor or parent/guardian.
8. **Revocation of Identity (Anonymisation):** Data subjects may be happy for personal data to be held for certain purposes so long as it is not linkable back to them personally. Anonymisation may be regarded as a variant of revocation, in that data subjects request a change to data held so that it is no longer personally identifiable.

These forms of revocation should be available as actions that individuals can perform on personal data held about them by an enterprise. Although, it is conceivable that an enterprise might only offer a subset of revocation forms, as they may be limited by technical capability and /or business policy. The EnCoRe project is exploring the various business models which might be adopted.

IV. DEFINING CONSENT AND REVOCATION POLICIES

Having considered the many different forms and variations of consent and revocation, we are in a position to define the notion of a consent and revocation policy, whose purpose is to enable an enterprise to inform customers of:

- a) what consent is required of them for each data field, and
- b) which revocation types are available to them should they wish to exercise control over their personal data.

Thus, a consent and revocation policy is defined over a set of data $D = \{\delta_i\}$ as a set of tuples

$$\{\forall i: (c_i, p_i, d_i, R_i)\}$$

where the c_i, p_i, d_i define together a value of consent (specifying collection, processing and dissemination rights respectively), and R_i is a set consisting of the names of allowed revocation types.

As an example, consider this policy over $D = \{\delta_1, \delta_2\}$:

$$\{ (c(30 \text{ days}), -, d(10 \text{ days}), \{2,3\}), \\ (-, p(10 \text{ days}), -, \{2\}) \}$$

What this policy specifies, assuming the only consent variable of interest is the time for which data is held, is that consent is required to enable the collection of δ_1 for 30 days and its dissemination for 10 days; furthermore, δ_1 can be revoked using revocation types 2 (deletion) and 3 (revocation of permissions to process data) (see Section III). The example policy also states that consent to process δ_2 for 10 days is required, and that deletion may be performed.

Note that the symbol $-$ is used when one of the c_i, p_i, d_i is omitted. Some further formalisation work is needed to make the definitions more rigorous, but we believe that the above presentation is sufficient for our purposes here.

An enterprise is likely to require stipulations on the minimum consent a customer can provide for each data field. In other words, rather than stating in a policy specific consent values that an enterprise needs from a customer, it may be desirable to specify a *range* of values. We regard this as a direction for further investigation.

Automatically enforcing such a policy involves obtaining specific consent from customers through some user interface, and implementing suitable controls corresponding to the revocation types in an enterprise's information system. Such work is being carried out within the EnCoRe project, in particular for case studies involving personal data held by one's employer, and for data held in biobanks.

V. CONCLUSIONS AND FUTURE WORK

We have in this paper explained the significance of consent and revocation preferences for personal data, and introduced the notion of a consent and revocation policy. We have described in detail the different aspects of the notions of consent and revocation respectively, and explained how an enterprise can define policies that make clear what options are available to a customer in terms of controlling his or her personal data.

There is much further work to be done on consent and revocation policies, most notably refining the definition of the R_i , so that the revocation types include parameters analogous to consent variables. In practical terms, we expect to link consent and revocation policies to a real-world access control language, such as XACML, so that enforcement of consent and revocation can be done programmatically. There are also considerations that relate to the internal consistency of such policies, such as preventing conflicts or incompatibilities between the consent requested and the revocation types made available. Further work will also consider the 'usability' of such policies from a data subjects perspective, and how we might make the expression of such policies easy.

Another direction for further investigation is how consent and revocation policies can be mapped into an enterprise's current security and privacy policies, in particular, which aspects can be expressed with existing policy languages, and which require new models and formalisms. We are also hoping to influence standardization bodies. We believe that it is essential to incorporate consent and revocation controls in enterprise information systems that handle personal data, and that the deployment of consent and revocation policies across such systems is a useful means of ensuring that the privacy preferences of individuals will be respected.

REFERENCES

- [1] Alan Westin (1967). *Privacy and Freedom*. New York: Atheneum.
- [2] Marco Casassa Mont, Siani Pearson, Gina Kounga, Yun Shen, and Pete Bramhall (2009). *On the Management of Consent and Revocation in Enterprises: Setting the Context*. Technical Report HPL-2009-49, HP Labs, Bristol.
- [3] EnCoRe. <http://www.encore-project.info>.
- [4] L. Cranor, B. Dobbs, S. Egelman, G. Hogben, J. Humphrey, M. Langheinrich, M. Marchiori, M. Presler-Marshall, J. M. Reagle, M. Schunter, D. A. Stampely, and R. Wenning (2006). *The Platform for Privacy Preferences 1.1 (P3P1.1) Specification*. World Wide Web Consortium Note NOTEP3P11-20061113.
- [5] OASIS eXtensible Access Control Markup Language (XACML). Standard available from <http://www.oasisopen.org/>.
- [6] Ioannis Agrafiotis, Sadie Creese, Michael Goldsmith, and Nikolaos Papanikolaou (2009). *Reaching for Informed Revocation: Shutting Off the Tap on Personal Data*, Proceedings of Fifth International Summer School on Privacy and Identity Management for Life, Nice, France, 7th – 11th September 2009.
- [7] Ioannis Agrafiotis, Sadie Creese, Michael Goldsmith, Nick Papanikolaou (2010). *The Logic of Consent and Revocation*. Submitted.