On the Management of Consent and Revocation in Enterprises: Setting the Context

Marco Casassa Mont, Siani Pearson, Gina Kounga, Yun Shen, Pete Bramhall

HP Laboratories
HPL-2009-49

Keyword(s):

Privacy, Consent, Revocation, Consent and Revocation Management, Privacy Management, Identity Management, EnCoRe Project

Abstract:

The aim of this paper is to set the context for the management of consent and revocation in enterprises, create awareness and so pave the way towards better and improved practices in this area. A number of international laws and regulations mandate (to some degree) that individuals should be enabled to express their consent for the usage of their data and subsequently be allowed to revoke it. Unfortunately the practical implications and management of consent and revocation are not yet fully understood and taken into account, apart from on an ad hoc basis. Key involved aspects are: allowing individuals to retain some control over their personal data; ensuring that consent and revocation can be enforced by data receivers. This paper addresses questions such as the following: What are the key requirements and practical implications of handling consent and revocation, for individuals and organisations (data receivers)? How can we enable people to effectively express their consent when disclosing their personal data and, subsequently, to revoke it? How could we enable organisations to manage and enforce consent and revocation? We focus on an enterprise scenario, as a significant example. We discuss requirements and open issues. We provide a reference model for the management of consent and revocation and illustrate some technologies that could be used to animate it. This is work in progress. Further research will be carried out in the context of the multi-disciplinary, collaborative EnCoRe project.

External Posting Date: March 6, 2009 [Fulltext] Approved for External Publication
Internal Posting Date: March 6, 2009 [Fulltext]

On the Management of Consent and Revocation in Enterprises: Setting the Context

Marco Casassa Mont, Siani Pearson, Gina Kounga, Yun Shen, Pete Bramhall

Hewlett-Packard Labs, Systems Security Lab, Bristol, UK
{marco.casassa-mont, siani.pearson, gina.kounga, yun.shen, pete.bramhall}@hp.com

Abstract. The aim of this paper is to set the context for the management of consent and revocation in enterprises, create awareness and so pave the way towards better and improved practices in this area. A number of international laws and regulations mandate (to some degree) that individuals should be enabled to express their consent for the usage of their data and subsequently be allowed to revoke it. Unfortunately the practical implications and management of consent and revocation are not yet fully understood and taken into account, apart from on an *ad hoc* basis. Key involved aspects are: allowing individuals to retain some control over their personal data; ensuring that consent and revocation can be enforced by data receivers. This paper addresses questions such as the following: What are the key requirements and practical implications of handling consent and revocation, for individuals and organisations (data receivers)? How can we enable people to effectively express their consent when disclosing their personal data and, subsequently, to revoke it? How could we enable organisations to manage and enforce consent and revocation? We focus on an enterprise scenario, as a significant example. We discuss requirements and open issues. We provide a reference model for the management of consent and revocation and illustrate some technologies that could be used to animate it. This is work in progress. Further research will be carried out in the context of the multi-disciplinary, collaborative EnCoRe project.

1 Introduction

It is common practice, for enterprises[1] and other organisations, to collect personal data and other confidential information, in order to provide services to individual consumers and enable business transactions. This data is usually stored, processed, aggregated and shared with third parties. In addition to enterprises, government agencies and e-commerce sites, new web 2.0 social networking sites (e.g. Facebook, LinkedIn, YouTube, etc.), federated and cloud computing services are increasingly collecting and using personal information. This data can range from personal and financial information to personal pictures, videos and descriptions of personal experiences and life matters.

People are raising concerns about the lack of effective control over their data, once it has been released and the fact that this data could be misused or shared with other parties without proper consent. In essence, these concerns involve security and privacy aspects. The main risk is that personally identifiable information (PII) (i.e. personal data) could be used or exposed even if this is not required to achieve an agreed goal or there is no authorization. Individuals, in general, are not aware of how their personal information is actually used, for what purpose and which parties have a copy. An interesting example, highlighting these concerns, is that "88% of individuals want sites to gather their consent" [1] (cf. [2]). Another example is that over 72% of people surveyed do not know that charities do not need permission to sell or rent their names and addresses to other charities [3].

In theory, prior to disclosing their data and giving any consent, individuals should be informed about: the purpose(s) for which it will (or may) be used; the destination of the data; its expected retention time; the amount and the sensitivity of the information exchanged; if and which other parties are involved; whether the data will be shared onward; if the consent to use this data can be revoked. Laws (e.g. EU Data Protection Directive 95/46/EC, HIPAA, COPPA, etc.) and privacy guidelines [4, 5] are already in place and (to some degree) mandate this, but there is little assurance that they are actually enforced. The terms of consent for using personal data for specific purposes, under certain conditions, and any subsequent revocation of this consent will determine how personal data should be handled once it has been disclosed.

In practice, even good-willing organisations struggle to take into account and enforce individuals' preferences and consent statements. Some organisations do not even fully know where customers' and other individuals' data is stored in their IT infrastructures, who can actually access it and how it flows around. The management of revocation of previously stated consent is even more obtuse. Data is usually copied, moved around, in different systems and can be released to other parties. Consent declarations rarely follow this data. In this context, managing revocation could be hard if not impossible.

The concepts of consent and revocation are not new, at least from a legal perspective. However, little progress has been made on understanding the practical implications, in particular within large organisations, including actual requirements and potential technical approaches to satisfy (part of) them. Further complexity comes from changing legal requirements (and the global complications), and hence the difficulty of organisations in being compliant and also knowing whether or not they are compliant. In this paper we take these legal notions into account but it is beyond our scope to provide an analysis of all related laws and regulations.

[1] In this paper, we use the term "enterprise" to refer generically to any commercial or non-state organisation. However, many of the points we make are equally applicable to state organisations.

The key problem is how to effectively enable individuals to express their consent for the usage of their PII data and (subsequently) to revoke this consent if they change their minds. A related important problem is how to enable enterprises to effectively manage and enforce both consent and revocation. Complexity in providing consent and expressing privacy preferences could undermine trust and individuals' willingness to further disclose personal information and engage with web services [6]. *Ad hoc* and reactive approaches to the management and enforcement consent and revocation by organisations also limit their chances to effectively address the problem, especially when this is coupled with the current inadequate audit and tracking approaches for personal/sensitive information. Progress needs to be made in this space: there is an opportunity for research and development of new solutions that take into account current IT and organisational constraints.

The remaining part of this paper is structured as follows. First, we discuss our understanding of consent and revocation. We only briefly touch upon legal aspects: this is a complex area globally and is being looked into within the EnCoRe project (Ensuring Consent and Revocation) [7], a multi-disciplinary, collaborative project that focuses on consent and revocation. We then discuss open issues and problems. We derive a set of basic requirements for organisations to support their management and for enabling individuals to handle consent. We discuss a high level model of what is required (at a conceptual level) to move towards the management of consent and revocation. We illustrate some of our technologies and solutions that could support some of the requirements and animate part of the high-level model. We discuss related work, current results and our next steps, including our involvement in EnCoRe [7].

2 Enterprise Scenario

We consider an enterprise scenario where personal data (of customers and employees) is collected, stored and handled by people within the enterprise, applications and services. This scenario is significant and representative, as we believe that similar needs, requirements and observations (for consent and revocation management) also apply in other related and emerging scenarios, such as cloud computing and federation of web services scenarios.

Personal data can be gathered by different sub-organisations within the enterprise (e.g. customer care department, help-desk and customer-relationship management services, web services to sell products, etc.). A variety of human and automated processes can affect data. Data can be moved around, within the organisation, for processing purposes, e.g. to deal with financial transactions, to handle help-desk requests, for marketing and research purposes. In this scenario, data can also be aggregated and manipulated. Under some circumstances, it can flow outside enterprise boundaries, for example in case some of its services are provided in a context of supply-chain or by outsourced applications.

Individual end-users (data subjects) can interact with the organisation in multiple ways, ranging from potential anonymous interactions (e.g. querying for information online) to the point where personal details need to be disclosed to enable further interactions (e.g. purchase, delivery of goods, etc.). In this context, individuals make decisions about how to proceed, for example in terms of disclosing data, based on a variety of factors, such as reputation of the organisation, perceived level of trustworthiness and degree of risk, perceived benefit of the transaction and clarity of the involved steps.

We specifically consider the case of a good-willing enterprise that is interested in systematically managing consent and revocation. Current identity management solutions and other *ad hoc* security solutions used by enterprises can be used to protect the access and usage of data, but primarily from traditional security perspectives. Information lifecycle management solutions (including the ones dealing with data storage, manipulation and minimisation processes, etc.) generally take into account business or legislative requirements (e.g. on data retention) but they are rarely driven by privacy aspects and/or specifically consent and revocation requirements.

Individual preferences and consent, and their dynamic management are usually not directly factored into these processes, or are so but just in a limited way: when dealing with privacy matters, there is usually reliance on human processes, that are ad-hoc and subject to failures. We are interested in exploring the case where additional solutions and processes could be available to track the flow of data, enforce the terms of consent and revocation when accessing and managing data. In this context, the actual accessed personal data will depend on the context, consent specifications and deployed policies.

Ideally, the enterprise should be aware of where personal information is stored and used within its boundaries and how and where data is flowing, across them. Risk assessment solutions should be available to assess its current compliance situation and make privacy-aware decisions on its practice and processes. Individuals, at the time of disclosing data, should be entitled to provide their consent, freely and on an informed basis, in terms of privacy preferences, related permissions and obligations [10]. Clarity and easiness on how an individual could express this is important. This would include the possibility for the individual to check for the status of their data, actions taken on the data, related consent statements and subsequently partially or complete revoke their consent. Changes to consent specifications and revocations will affect the way people, applications and services (both within the organisation and in third-parties data was disclosed to) access and use related data.

This scenario is far away from what can be achieved today, at least in terms of consent and revocation management. Nevertheless we believe it is important to set a target, analyse the requirements and suggest steps to enable enterprises and other organisations to make progress towards it and improve their practice. In doing this, reality-checks are important. Many ideal solutions could be designed and architected, but their likelihood of being adopted depends on the impact they are going to have on the organisation, the fact that legislations are mandating them, the risks at stake and the eventual mitigation factors they can provide.

3 Analysis of Consent and Revocation

3.1 The Concept of Consent

The concept of consent is not new. It has been discussed in a variety of scientific papers, privacy reports and legislations (e.g. [4, 5]). In the latter, consent is often used as a synonym for "informed consent", i.e. the voluntary agreement of an individual for the specific use of his personal information. Fundamentally, it is a statement that captures the willingness of individuals (data subjects) that their data could be used for specified purposes, under well defined conditions and circumstances. For example, individuals might be happy to allow an organisation to use their financial data for transactional and payment purposes or to allow companies to use their address information for delivery and sending contextual mail. The term "informed" encloses disclosure (the individual must know why his data is collected, who will have access to it, etc) and comprehension (the individual must have an accurate interpretation of what he is asked to consent to) while "consent" encloses voluntariness and agreement [2]. In other words, consent should be *freely* (no coercion) given after the individual has been informed about the implications of his choices and has fully understood these implications. In order to translate a legitimate individual will, consent should be *explicit* and *authentic*, i.e. it should not be given by another entity than a legitimate individual and not be modifiable by other entities. To fulfil the former, the individual should perform an action that proves his consent while for the latter, the source of a consent statement must be authenticated and the integrity of a consent statement must be guaranteed.

Consent can be conditional, e.g. given for specific purposes. It could be given by people on a well defined timeframe or just under well specified circumstances. For example individuals might give consent to use their personal data just for a specific business transaction. Hence there will be no consent for subsequent usages of this data.

Consent is sometimes "confused" with opt-in/opt-out choices. These choices define criteria and conditions under which data can be used, for example for specific purposes (e.g. marketing, research, etc.) or to receive notifications or share data with generic third parties. However, opt-in/opt-out does not usually define aspects of data retention, how personal data should be processed or which parties data could actually be disclosed to. For example, an individual can opt-out from receiving further emails relating to commercial products but it does not necessarily mean that the email address is inaccessible or forces the organisations to remove such information. Consent can be broadly characterised by a wider set of fine-grained privacy preferences and constraints both on personal data itself and how it should be used, such as specific permissions (to disclose data, for specific purposes, to specific entities) and obligations (such as on deletion or minimisation or anonymisation of data items, after a period of time or some events happens).

In other words, individuals' preferences are key aspects that qualify consent, rather than just being a mere case of selection of options, such as for opt-in/opt-out choices.

The specification of consent makes sense when associated with (or referring to) data. It can be fine-grained, i.e. specifically refer to personal data items. It defines constraints and refines organisations' policies on how to deal with data (i.e. access control, privacy and obligation policies etc.). It can be seen as "metadata" associated to personal data.

Enforcement of consent requires that each time data are accessed by the organisation, its constraints (and any additional related policies) specifying how the individuals want their data to be dealt with, must be taken into account. Consent management and enforcement is of little value if it can be easily bypassed by people or applications/services.

Complexity comes from the fact that personal data can be copied into multiple enterprises' systems and moved around. If individuals' constraints and related policies are split from the data, then violations of privacy and their requirements might occur. Therefore, consent and related policies should be bound to data in such a way that they cannot be separated from it, and data should be dealt with as specified by the individual if the needed constraints are to be met by the data processor. This is particularly hard to achieve when data is disclosed by an organisation to third parties (e.g. in a supply-chain). The third party receiving the data should also comply with the individual's consent and associated policies. Hence, the management of consent has strong implications for the receiving organisation (and any involved third party), in terms of setting robust access control and privacy enforcement mechanism, as well as solutions to enforce related obligations and constraints.

Consent declarations are not static. People should be entitled to change their mind, hence consent can change over time and be revoked. For example, a person might have given consent to use her medical data for research purpose, for a predefined period of time. She might then change her mind, and be enabled to revoke that consent, or change it, by qualifying it with different privacy preferences.

3.2 The Concept of Revocation

The concept of revocation of consent is currently not very well understood, in particular in terms of its implications for organisations (data receivers).

In information technology security, revocation designates the process that renders an object (e.g. a public key, a certificate, etc.) invalid. When "revocation" is applied to the broader concept of consent, there are wider and more subtle implications: "revocation" designates the process that permits an individual to invalidate or modify previously given consent, on personal data. This revocation should apply to any copy or instance of this data, within the organisation that initially collected it and in any other third party to which this data was subsequently disclosed. In practice this is hard to achieve.

As with consent, for revocation requests the source of a request must be authenticated. The integrity of a revocation request must be checked. Mechanisms must be in place in organisations to check the current state of consent, and ensure that it has not been revoked. A "non-repudiable" proof that "revocation was executed" should be sent to the requester. Revocation requests could happen at any time and need to be authenticated. Revocation should be enforced in a reasonable time, as requested by the individual.

The management of revocation requests can have consequences on data stored at different places within the organisation. If personal data was copied in various systems of an organisation and also disclosed to third parties, a change of consent (or revocation of consent) must be applied and notified to all the involved entities. Otherwise this would violate individuals' privacy preferences.

It is important to understand that revocation can be fine-grained and be qualified by attributes. It might not just be a matter of "turning off" the entire consent given on a set of personal data but there could be degrees of revocation, affecting specific data items. For example, an individual might give consent to use their personal data (including email and home address) for financial, research and marketing purposes, for a defined timeframe. Later on he or she might revoke their consent to use their email address, for any purposes beyond notifications of financial transaction, for another predefined period of time.

Revocation can be reversible or irreversible. In the above example, the revocation of consent on using the email address could be reversible. The email address is still stored by the organisation, only the way it can be used has changed.

The individual might have qualified their revocation request, by requiring that the deletion of the email address. In this case, revocation is irreversible. Further disclosures of personal data might be then required, along with new consent statements.

Depending on the context, scenario and involved parties there might be significant variability in terms of implications of managing consent and revocation. In the remaining part of this paper we focus on the enterprise scenario, to make some progress in analyzing related open issues, requirements and next steps.

3.3 On the Management of Consent and Revocation in Enterprises

This section aims at briefly discussing some of the implications of dealing with the management of consent and revocation within enterprises, and highlights some of the key open issues.

The lifecycle management of personal data within and across organisations is affected by consent, individuals' preferences and other constraints, along with the enforcement of access control policies and related obligations. Consent and revocation themselves have a lifecycle that need to be managed, impacting how personal data is managed. Figure 1 illustrates some of the different states of data (at different times) along with related consent.

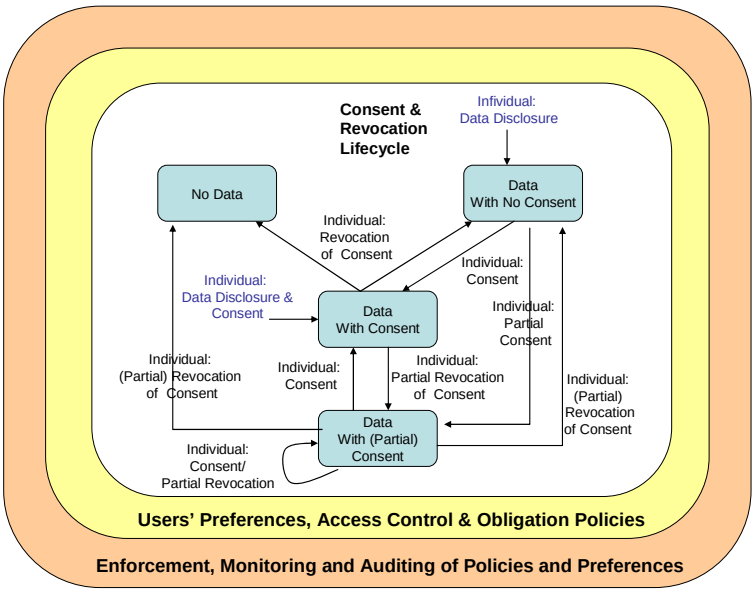

Fig. 1. Consent and Revocation Lifecycle Management

This lifecycle is triggered by individuals giving consent and changing it over time, i.e. revoking or partially revoking it and/or by the management and enforcement of related (obligation) policies. Consent and revocation can be qualified by individual preferences. They have an impact on the enforcement, monitoring and auditing or related privacy-aware access control and obligation policies.

Data could have initially been disclosed by an individual with no expression of consent (so, in theory, this data should not be used but only retained) or have ended up in this situation, after he or she revoked consent. Alternatively, data might

have been deleted ("No Data"), as a result of revoking consent and/or enforcing related individual preferences (inclusive of obligations requiring the deletion of data after a predefined period of time).

However, it is currently hard for individuals, in the first instance, to express and/or change consent: they might not be allowed to do it or it may be just partially enabled. Currently there is a lack of a flexible and systematic approach to the management of consent and revocation, in a way that provides individual control.

Research has been carried out to define privacy-aware access control [9, 23, 24] and privacy obligation policies [10] to enforce aspects of consent, when accessing or manipulation stored personal data.

However, in practice, enforcement of consent and revocation is still a problem: actual control and enforcement points (such as the one required for privacy-aware access control and obligation management systems) may not be deployed within an organisation or might be purely driven by security policies and aspects, etc. Most current approaches still deal with costly human processes that are subject to failure. This is partially due to a lack of effective processes to map consent declarations and changes of consent (e.g. revocation) into actionable policies and automated processes.

It is not uncommon that organisations might not even know where personal data is actually stored, as data might have been copied around or shared. If data has been disclosed to third parties, there is the very complex problem of how to propagate consent and revocation to these third parties and how to have assurance they are going to act on them.

Stickiness of consent to data is a related issue. Consent and revocation statements, along with related individuals' preferences can be seen as "metadata" that must stick to the personal data. This affects the enforcement of related policies and obligations that are contextually affected by this metadata.

However, the data and metadata binding could be broken, when copying data, moving it around or sharing to third parties. This negatively impacts individuals' control and the enforcement of their requests.

A related issue is about the lack of any effective standard to represent and share consent and revocation within organisation and across parties. This undermines the possibility of collaboration between good-willing organisations and requires the definition of limited and *ad hoc* solutions.

How can we help good-willing organisation to improve and check on their practices? On the one hand there is a need for better processes and technological solutions to sustain them. On the other hand, there is a need for risk assurance and for better assurance that all the involved parties are playing as agreed. There is also a need to provide assurance to individuals about current practices and how effectively their consent and revocation requests are handled.

This is a huge and complex research area impacting not only individuals and organisations but also legislation and standardisation.

3.4 Analysis of Threats and Involved Risks

When talking about privacy and specifically in this case about consent and revocation, it is also important to analyse the problem from an economics perspective. There is a balancing act between multiple forces that must be taken into account:

- The actual willingness of organisations to make an effort in this space, driven by customer-awareness, trust and reputation matters
- The involved costs & benefits for organisations and the need to be pragmatic
- Infrastructural and IT aspects of organisations, current investment and operational requirements.

No radical changes will be easily accepted, despite the possible role that laws, legislation, social pressure, etc. can have in influencing decisions. The actual threats and involved risks are the main drivers for organisations, influencing their willingness to accept or reject any management of consent and revocation.

Generic threats of relevance include privacy threats of data exposure, lack of individual control, lack of transparency, lack of legal compliance, unauthorized access to personal information, etc. More specifically, there are the issues of sharing data in a way that the individual would not like, data being collected and/or distributed without the individual's consent, the individual wanting to remove their consent for information to be used a certain way within an organisation but this not being possible, etc.

The actual willingness of organisations to take the involved risks depends on the context. For example strong legislation coupled with penalties and strong pressure coming from customers and civil society organisations might result in enterprises changing their behaviours and being more active on this front. On the other hand, lack of any pressure and legislation might result in little action from an organisational side. Hence, as applies also in general for privacy matters, for the management of consent and revocation the answer to the problem does not just rely on good processes and technological solutions but also on the impact of legislative and social activities.

In this paper we recognise the importance of these factors. In the remaining part of this paper we focus on process and technological aspects but, in the context of the EnCoRe project (in which we are involved) [7], we also aim at addressing involved legislative and social issues.

4 Core Requirements

This section discusses core needs and requirements (both from an individual and enterprise perspective) that, based on our analysis, need to be satisfied in order to manage consent and revocation within organisations.

A list of requirements, from an individual's perspective, follows (in no particular order of priority):

1. Need to provide intuitive ways to explain how personal data is going to be used, in which contexts and for which purposes
2. Need to express consent and revocation in an explicit way, in the context of various policies and processing activities that might apply to the data;
3. Need to check for the current status of personal data, along with the current expression of consent;
4. Need to check for historical changes of consent and expressions of consent.

A list of requirements, from an enterprise/organisation's perspective, follows (in no particular order of priority):

1. Need for an explicit and standard representation of consent and revocation requests along with their associations to personal data;
2. Need for standards to exchange consent and revocation information, within and across organisations;
3. Need to provide individuals with mechanisms to express (fine-grained) consent or revocation on data items, including preferences and any other qualifiers;
4. Need for pragmatic mechanisms to bind consent to data;
5. Need for tracking, across and outside organisations, data that is subject to consent, along with current consent. In other words, knowing where personal data is and which consent expression this data is subject to;
6. Need for ways to express (access control, obligation, etc.) policies that are consent-aware and context-aware. Pragmatic solutions are required, minimising the impact on current enterprise applications, services and processes;
7. Need to take legacy systems into account, inclusive of operational and security (e.g. identity management) systems;
8. Need to provide mechanisms to enforce and audit consent and revocation, in an automated way, involving both access control and lifecycle management of data;
9. Need for data lifecycle management processes that are consent-aware and can react to changes of consent, e.g. revocation;
10. Need for mechanisms to support the lifecycle of consent, based on changes of individuals' preferences and revocations;
11. Need for solutions to assess enterprises' current ability to act on consent and revocation, given the types of handled personal data and related threats/risks;
12. Need for mechanisms to get assurance from business partners (personal data is disclosed to) that they are compliant to consent and revocation requests, along with enterprise policies;
13. Need to provide assurance to individuals and visibility about the effectiveness of an organisation's consent and revocation practices, along with feedback with respect to their data treatment.

These requirements can help to define the high-level model of a potential approach and solution to deal with consent and revocation management within and between enterprises. The next section describes this abstract model, followed by indications of how (aspects of) it might be implemented.

5 A High Level Model of Consent and Revocation Management

This section aims at providing a first, high-level model of consent and revocation management, within and between organisations, taking into account requirements and our analysis. The management of consent and revocation is very complex and it is unlikely that there is going to be a solution that fits all needs, in all potential scenarios. However, we believe it could be valuable to identify the core building blocks and a model that might be used as a reference to build such solution.

We focus on a model for organisations, involving (degrees of) personal data disclosed by individuals, enabling capturing related consent from them, managing and enforcing it, along with revocation – within and outside organisations, based on the flow of this data. Figure 2 provides an overview of this model.

In order to move towards the satisfaction of the requirements identified in Section 4, we envisage that the high-level model should include the following components:

- **Personal Consent & Revocation Assistant:** this component assists and provides user-side capabilities to help ordinary people express their consent (opt-in/opt-out choices, privacy preferences, etc.) and revocation requests, along with explanation of privacy practices provided by the organisations. It can be triggered (e.g. via a plug-in in a web browser) during data disclosure processes, in order to provide suggestions and default values for consent, given the nature of the data. It keeps historical traces of data disclosures, along with revocation requests.;
- **(Virtual) Data Registry**: an organisation, to be able to deal with consent and revocation management, must know, in first instance all the locations where data is stored. This is the goal of the (virtual) data registry, a special repository (or an aggregation of synchornised repositories) that keeps track, for each known individual, where their data has been stored (in operational data repositories) within the organisation, which type of data has been disclosed outside the organisational boundaries and to whom. This component keeps track of expressions of consent and revocations for each managed personal data item. It is a critical component that has to be secured and protected. It also needs to be constantly updated, in order to provide value to the organisation. This can be achieved by a mixture of automation

(provisioning services) and manual intervention by privacy administrators, for example in the context of well known business processes and procedures.

- **Consent and Revocation Provisioning:** this component is in charge of automatically and systematically updating the data registry every time there is a new expression of consent and revocation (either for new, i.e., previously unknown individuals or known ones). It can be seen as a special component, part of consolidated individual record and system provisioning solutions, specialised to deal with the provisioning of consent and revocation. In this context, it is in charge of updating individuals' preferences and constraints that affect the enforcement of access control and obligation policies. An expression of consent or a subsequent revocation of consent will affect how applications and services will access data;
- **Privacy-aware Policy Enforcement**: this component deals with access control on data and its lifecycle management, driven by individual preferences, consent or revocation of consent. It mediates access to data by intercepting applications and services' requests, driven by enterprise policies and individuals' preferences, inclusive of the consent given by individuals (e.g. allowing the usage of personal data for specific purposes, in a given timeframe). It supports privacy-aware lifecycle management of data, driven by obligation policies and individuals' preferences (e.g. deletion or minimisation of data after a predefined period of time). Interception points are required for applications and services when accessing data. In legacy applications this might be a difficult requirement to satisfy;
- **Disclosure and Notification Manager**: this is a component (with a related notification infrastructure) to intercept and keep track of flows of personal data, within and between organisations and to propagate the associated consent information. Applications and services might need to be instrumented with agents that communicate with this component. It interacts with the Data Registry to update data locations and related consent information. Changes of consent and revocations are propagated to known data locations (within and outside the organisation) by means of notifications. Request for reconfiguration might be sent to the consent and revocation provisioning system to reflect any change. These notifications might require (non-repudiable) acknowledgements by the involved parties that are going to be logged and audited. Ideally, the third parties receiving data, consent and notifications should adopt a similar model, to ensure a consistent management of consent and revocation;
- **Audit:** this is an essential component to log and keep track of what happens to data, consent and revocation during operational and administrative activities, inclusive of flows of personal data within and outside the organisation;
- **Risk Assurance**: this is a key offline component that is used by the enterprise's privacy administrators to assess current risks, based on known threats and provide indications of compliance. Based on these threats, known information flows, involved types of data and security aspects (of IT infrastructures, people and applications/services), this component can identify the level of risks, based on known risk heuristics and rules and suggest action items, such as investing in instrumentation of critical applications or forbidding particular flows of information, etc.

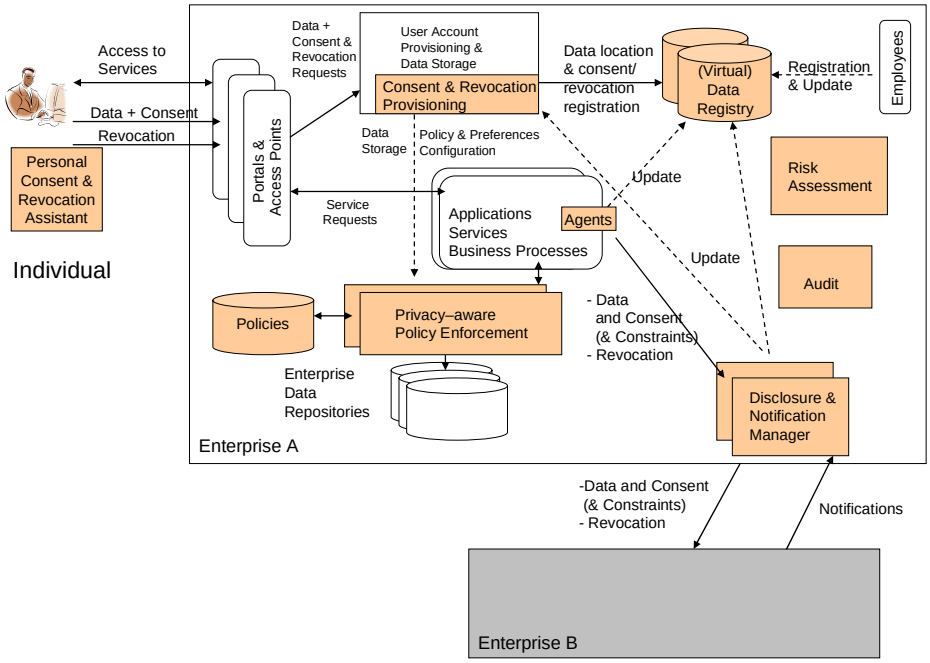

Fig. 2. High-level Model of Consent and Revocation Management

This model is compatible with the PRIME Toolbox [24], with regards to the privacy-aware policy enforcement (inclusive of privacy-aware access control, data handling capabilities [24] and obligation management [10, 11]). However, it aims at explicitly handling consent and revocation, by introducing the concepts of "consent and revocation provisioning" and "data registry", along with the "disclosure and notification manger capability". At the disclosure time, individuals can express their consent. "The Personal Consent and Revocation Assistant" is a critical part of helping people in assessing the situation. This requires further R&D work in HCI, in the directions paved by [27, 33]. Once data is disclosed, the location of where data is stored, along with a copy of the consent statement is stored in the data registry, during the provisioning

phase (consent & revocation provisioning). Accesses to data (by employees, applications and services) are intercepted by the privacy-aware policy enforcement component to ensure the enforcement of security and privacy policies, inclusive of fine grained consent and individuals' preferences. This component handles the lifecycle of data, driven by the enforcement of obligation policies [10]. Should data need to be disclosed to other locations within and outside the organisation, the "disclosure & notification manager" keeps track of it and updates the data registry. Non-repudiable notifications of success or failure (by the involved parties) are logged by the auditing system. Similarly, changes of consent or revocations by individuals are propagated by this component to all the involved parties and the data registry is updated.

This high-level model is an ideal model that makes some strong assumptions. For example, it assumes some "standard" way to deal with account/record provisioning and configuration management, for example by using Identity Management solution suites. This might not be the case for particular organisations that have developed *ad hoc* IT infrastructures or rely on manual administration of accounts and IT systems. Applications and services should also be instrumented, to allow the enforcement of policies (along with expressions of consent and revocation) and interactions with the other model components (e.g. "Data Registry and Disclosure" and "Notification Manager"). This might be possible for new applications and services, by applying a "privacy by design" approach. For legacy applications (more or less secure) *ad hoc* wrappers might be the only possible solution or no changes might be feasible at all.

Another strong assumption is that there is a need to have knowledge of where data is stored and have ways to intercept programmatic or manual flows of data. This again might not be the case, in general, as organisations might not even fully know where customer data or other personal data is stored. However, in absence of this, only limited management of consent and, more importantly, revocation might be possible – hence it is a key requirement knowing where data is.

Finally, to work systematically across organisations, a similar model should be adopted by all the involved parties. In some contexts might be achievable (e.g. in supply chains), in others might require expensive changes and adaptations.

The main point of this model is really to set the context and identify the key aspects that should be considered and get in place in order to ensure an effective management of consent and revocation. We believe this model can be considered as a reference model, to progressively drive the evolution of organisations' infrastructures (handling personal data, consent and revocation). This could be done in multiple stages, starting with the automation of new applications, services and processes, and relying on human processes and auditing for the reaming part of the infrastructures. The risk assurance component can help to assess and determine priorities to focus efforts on.

Some of the components described in this high-level model need further work to be fully designed and implemented. Part of this work might be done in the context of the EnCoRe project [7]. However, there are already existing technologies and solutions that we have developed that can help implement part of the required functionalities of the model. The next section provides an overview.

6 Our Current Approach

In this section we briefly describe how we envisage translating (aspects of) the model presented in Section 5 into something that could be implemented and deployed. Our approach is to map our high-level model into a high-level architecture, refined with current IT solutions used by organisations. In doing this we can exploit the capabilities we already have in this space. Let us consider an organisation that has multi-party interactions, such that there will be data flowing between parties and potentially across legal and organisational boundaries. Existing enterprise capabilities in such a scenario include provisioning systems, identity management capabilities (such as access control, authorisation and authentication), and audit.

We can exploit techniques used within our existing capabilities and R&D prototypes in order to achieve the following:

- Extended provisioning system to keep track of location of personal data and other metadata (inclusive of consent, individuals' preferences, etc) [11]: this system (that can be used to animate part of the functionalities required by the "Consent and Revocation Provisioning" module, Figure 2) is currently part of an identity management suite, in charge of providing wider account/record provisioning capabilities. It needs to be adapted to deal with updates of consent derived from revocations and integrated with a "Disclosure Notification Manager".
- Privacy and consent-aware access control mechanisms [9]: this system can make decisions and enforce privacy-aware access control policies that keep into account individuals' consents and other preferences. It intercepts queries to standards data repositories (e.g. databases, LDAP directories, etc.) and returns sanitized views (if any) on requested data. This component would be part of the "Privacy-aware Policy enforcement" module, Figure 2;
- Obligation management to handle information lifecycle management, driven by individual preferences and organisational policies [10]. It consists of a scalable obligation management system, driven by obligation policies and individuals' preferences (expressed as part of individuals' consents). It manipulates data over time, including data minimisation, deletion and management of notifications to individuals. This system would be part of the "Privacy-aware Policy enforcement" module, Figure 2;
- Compliance and risk assessment support [12, 13]: this system provides capabilities to assess risks within an organisation, based on a model of known threats, questionnaires to extract information by administrators & decision makers, and produce compliance report based on heuristic rules. This system would be part of the "Risk Assessment" module.
- Sticky policies and various degrees of enforcement, including exploitation of trusted computing and cryptography to stick policies to data and ensure that that receives act according to associated policies and constraints, by interacting

with trusted third parties or Trust Authorities[14, 15]. These capabilities can be used to harden the association of policies with data, during notifications and disclosures of data to third parties, to ensure an audit trail [14, 15].

We reckon that R&D efforts need to be carried out on the "Personal Consent & Revocation Assistant", "Data Registry" and the "Disclosure & Notification Manager", to provide effective ways to synchronise the flow of data and notifications with knowledge about data locations and up-to-date consent. Although technologies, automation and solutions can help address aspects the problem, note that this would only form part of the solution. Organisations also need help to assess their practices, processes and behaviours. Risk assessment capabilities [13] may play a role in helping to achieve this.

7 Related Work, Current Status and Next Steps

Existing policy specification, modelling and verification tools include EPAL [16], OASIS XACML [17], W3C P3P [18, 19] and Ponder [20]. However, these do not include explicit modelling of the concepts of consent and revocation. The concept of privacy and the notion of consent has been explored in the context of a taxonomy [21]. Mechanisms for implementation of consent from an HCI point of view have been considered in [22, 24], and should be built upon within our solution.

A technical solution for sticky policies and tracing services is introduced in [14] that leverages Identifier-Based Encryption (IBE) and trusted technologies. This solution requires enforcement for third party tracing and auditing parties. One drawback of this approach is that data is bonded with the policy itself, which makes data heavy-weighted and potentially not compatible to the current information systems. However, its traceability function can be treated as the start point of consent and revocation. An alternative solution, that permits to bind privacy preferences to data and to convey the consent of the individual as well, has been proposed by Pöhls in [26, 27]. The solution relies on a Merkle hash tree [28] whose leaves are the tuples <data item, privacy preferences>. The root of the tree is signed by the individual to indicate his consent for each data item to be handled as specified by the associated preferences. This solution permits individuals to anticipate future aggregation by giving in advance their consent. Revocation of consent is managed with existing solutions [29, 30] used in X.509 public key infrastructure. However, as explained in [26, 27], the solution does not avoid the non-consented use of data.

A unified policy model is discussed in [25], which discusses steps towards privacy management for organisation or across organisations within a federated identity management environment. A coherent entry point was proposed to enable the consumers to easily review and manipulate stored personal information, which may in turn provide improved control over suspension and resumption of individuals' personal data. This method is elegant but its potential application in consent and revocation needs to be further studied.

A Platform for Enterprise Privacy Practices (E-P3P) is discussed in [8]. It separates the enterprise-specific deployment policy from the privacy policy and facilitates the privacy-enabled management and exchange of customer data. The authors also proposed E-P3P language to formalise privacy policies. The authors claimed that the proposed method can possibly be extended to multiple enterprises though such scenario is not discussed. This research effort further led to the Enterprise Privacy Authorization Language (EPAL) [16]. IBM [25] introduced a unified policy model towards privacy management in an organisation or between organizations within a federal identity management environment. The authors proposed a coherent entry point at web sites enabling the consumers to easily review and manipulate stored personal data. This scheme provided improved control over suspension and resumption of individuals' personal data, though limited enterprise scenarios were discussed. Its potential application in consent and revocation needs to be further studied.

A data handling policy is introduced in [23] to define how the personal information should be dealt with at the receiving party. The authors also proposed a DHP language to enable the building of customised policies. This method was developed/implemented in the framework of the PRIME project [24] although the management of revocation is not explicitly addressed. PRIME aimed to develop a working prototype of a privacy-enhancing Identity Management System. Its purpose was to foster market adoption of novel solutions for managing identities, as demonstrated in real-world scenarios. A systematic approach is proposed [31] including an anonymous credential system, an access control system, negotiation functionality, and an automated reasoning system, to serve both individuals and service providers needs in order to implement the EU Directives 95/46/EC and 2002/58/EC in PRIME. Within this architecture, Christer Andersson et al. [32] discussed the socio-psychological factors and HCI aspects that influence the individuals' trust in privacy enhancing identity management, and demonstrated that HCI research, user studies, and socio-psychological research are indispensable steps in privacy system design.

The usability research work that has been done within the first year of PRIME is summarised in [33]; three UI paradigms - role-centred, relationship-centred and townmap-based paradigms - for privacy-enhanced identity management are discussed. Our suggested model and approach is consistent and compliant with PRIME outcomes, but it extends it by explicitly suggesting how to handle consent and revocation, along with their lifecycle. We currently have technologies and working prototypes to animate aspects of the model we illustrated in Section 5. However, we reckon that the management of consent and revocation in enterprise is an important, open issue that requires research and further understanding of the involved problems. This is work in progress: we are going to carry out further R&D work in this space within the EnCoRe project.

8 Conclusions

Despite the existence of various data protection and privacy laws, there are technical and organisational issues as well as different perceptions of the priorities and risks involved in consent and revocation. Organisations have different understandings and approaches on how to deal with consent and revocation: usually they provide human-based, limited solutions. In this paper we have discussed the complexity of this topic. We focused our observations on an enterprise scenario, as a significant example. We analysed these concepts, along with some key open issues and problems. We derived a set of basic requirements of relevance for organisations. We then discussed a high level model of a potential approach to deal with the management of consent and revocation and briefly described technologies and solutions that could be used to implement some of the involved aspects (e.g. privacy-aware access control, obligation management, compliance and risk assessment support, etc.). We illustrated the current status, related work and our next steps: specifically, we are developing our ideas further via our involvement in the EnCoRe project [7].

References

1. Harris, Poll: Business Week, http://www.businessweek.com/ 2000/00_12/b3673010.htm, 2000
2. Friedman, B., Howe, D.,Felten, E.: Informed Consent in the Mozilla Browser: Implementing Value Sensitive Design, Proc. of HICSS '02, 2002
3. Anonymous, Test Your Personal Privacy IQ, Privacy Journal, Dec, 2008
4. Federal Trade Commission, Privacy Online: Fair Information Practices in the Electronic Marketplace: A Federal Trade Commission Report to Congress. Washington DC: FTC, 2000
5. Organisation for Economic Co-operation and Development (OECD), Guidelines governing the protection of privacy and transborder flows of personal data, Paris (1980) and Guidelines for consumer protection for e-commerce www.ftc.gov/opa/1999/9912/oecdguide.htm, 1999
6. Tweney, E., Crane, S.: Trustguide2, An exploration of privacy preferences in an online world, Expanding the Knowledge Economy: Issues, Applications, Case Studies, P. Cunningham and M. Cunningham (eds), IOS Press 2007
7. EnCoRe project, Ensuring Consent and Revocation, Project web site, http://www.encore-project.info, 2008
8. Schunter, M., Waidner, M.: Platform for enterprise privacy practices: Privacy-enabled management of customer data, in Privacy Enhancing Technologies, Second International Workshop, PET 2002, San Francisco, CA, USA, April 14-15, 2002, Revised Papers, vol. 2482 of Lecture Notes in Computer Science, 2003
9. Casassa Mont, M., Thyne, R.: Privacy Policy Enforcement in Enterprises with Identity Management Solutions - 4th International Conference on Privacy, Security and Trust 2006, PST 2006, 2006
10. Casassa Mont, M.: Dealing with Privacy Obligations, Important Aspects and Technical Approaches - 1st International Conference on Trust, Privacy and Security in Digital Business 2004, TrustBus 2004, 2004
11. Casassa Mont, M. Thyne, R.: A Systemic Approach to Automate Privacy Policy Enforcement in Enterprises - 6th Workshop on Privacy Enhancing-Technologies 2006, PET 2006, 28-30 June, Cambridge, United Kingdom,, 2006
12. Elahi, T., Pearson, S.: Privacy Assurance: Bridging the Gap Between Preference and Practice, C. Lambrinoudakis, G. Pernul, A.M. Tjoa (eds.), Proc. TrustBus 2007, LNCS 4657, Springer-Verlag Berlin Heidelberg, 2007, pp. 65-74.
13. Pearson, S., Sander, T., Sharma, R.: A Privacy Management Tool for Global Outsourcing, submitted to TrustBus'09, 2009
14. Casassa Mont, M., Pearson S., Bramhall, P.: Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services, Proc. DEXA 2003, IEEE Computer Society, pp. 377-382, 2003
15. Pearson, S.: Trusted Computing: Strengths, Weaknesses and Further Opportunities for Enhancing Privacy, Trust Management, Proc. iTrust 2005, LNCS 3477, ed: Peter Herrmann, Valérie Issarny, Simon Shiu, 2005
16. IBM, The Enterprise Privacy Authorization Language (EPAL), EPAL specification, v1.2, http://www.zurich.ibm.com/security/enterprise-privacy/epal/, 2004
17. OASIS, eXtensible Access Control Markup Language (XACML), http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml
18. W3C, The Platform for Privacy Preferences, v1.0, http://www.w3.org/TR/P3P/, 2002
19. Cranor, L.: Web Privacy with P3P, O'Reilly & Associates, September 2002. ISBN 0-59600-371-4.
20. Damianou, N., Dulay, N., Lupu, E., Sloman, M.: The Ponder Policy Specification Language, http://wwwdse.doc.ic.ac.uk/research/policies/index.shtml, 2001,
21. Solove, D.J.: A Taxonomy of Privacy, University of Pennyslavania Law Review, vol 154, no 3, January 2006, p. 477. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=667622
22. Patrick, S., Kenny, S.: From Privacy Legislation to Interface Design: Implementing Information Privacy in Human-Computer Interactions, R. Dingledine (ed.), PET 2003, LNCS 2760, pp. 107-124, Springer-Verlag Berlin Heidelberg, 2003
23. Ardagna, C., Vimercati, S., Samarati, P.: Enhancing user privacy through data handling policies, Data and Applications Security, volume 4127, LNCS, pp. 224–236, 2006
24. PRIME, Privacy and Identity Management for Europe, http://www.prime-project.eu, 2008
25. Schunter, M., Waidner, M.: Simplified privacy controls for aggregated services - suspend and resume of personal data, Privacy Enhancing Technologies, 7th International Symposium, pp. 218–232. Springer, 2007
26. Pöhls, H. C.: Verifiable and Revocable Expression of Consent to Processing of Aggregated Personal Data, ICICS 2008, 2008
27. Pöhls, H. C.: Authenticity and Revocation of Web Content using Signed Microformats and PKI Technical Report, number 276-07, University of Hamburg, Germany, February 2007
28. Merkle, R.C.: Secrecy, Authentication, and Public Key Systems, PhD thesis, Stanford, 1979
29. Housley, R., Polk, W., Ford, W., Solo, D.: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, RFC 3280, 2002
30. Myers, M., Ankney, R., Malpani, A., Galperin, S., Adams, C.: X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP. RFC 2560 , 1999
31. Camenisch. J.,, Shelat, A., Sommer, D., Fischer-Hübner, S., Hansen, M., Krasemann, H., Lacoste, G., Leenes, R., Tseng, J.C.: Privacy and identity management for everyone, in Digital Identity Management, pp. 20–27, ACM, 2005
32. Andersson, C., Camenisch, J., Crane, S., Fischer-Hübner, S., Leenes, R., Pearson, S., Pettersson, J.S., Sommer, D.: Trust in PRIME, in Signal Processing and Information Technology, 2005. Proceedings of the Fifth IEEE International Symposium on, pp. 552–559, 2005
33. Pettersson, J. S., Fischer-Hübner, S., Danielsson, N., Nilsson, J., Bergmann, M., Clauss, S., Kriegelstein, T., Krasemann, H.: "Making PRIME usable," in *SOUPS*, ACM, 2005