A Customizable Privacy Assurance System based on Active Feedback

Marco Casassa Mont, Stephen Crane
Trusted Systems Laboratory
HP Laboratories Bristol
HPL-2006-56
April 12, 2006*

privacy, assurance, reputation, feedback, obligation management, trust, identity management

People are often required to disclose Personal Identifying Information (PII) in order to achieve their goals, e.g. when accessing services, obtaining information and goods, etc. Being able to say with absolute certainty that another party can be trusted to properly handle personal data with today's technology is probably unrealistic. Feedback solutions based on reputation mechanisms can address aspects of trust and assurance in relation to how personal data is managed by an enterprise. However they usually rely on subjective feedback which is based on empirical experiences, and typically they do not allow individuals to systematically track and manage their specific experience. In this paper we propose an approach that enables people to monitor the status of their personal data which they have previously shared with an enterprise, service provider or other organisation - under specific conditions previously negotiated - and actively gather information on how adequately the management of these data meets their personal expectations. Ongoing monitoring and notification, and the ability of the client to form a simple record of past interaction, provides the client with greater confidence and assurance in situations where they need to share personal sensitive information with organisations they would otherwise not be able to claim they trust. This feedback process is based on conditions that are specific to the process of sharing PII and provides the client with assurance that an enterprise is a) capable and b) actually fulfilling PII processing preferences that are agreed at the time the data is disclosed, and which ultimately enables the client to form an opinion about the service provided. We present the principles of our approach and architectural components that support a practical implementation. This is work in progress and the research is on-going, carried out in the context of PRIME.

* Internal Accession Date Only
Approved for External Publication

A Customizable Privacy Assurance System based on Active Feedback

Marco Casassa Mont, Stephen Crane

Hewlett- Packard, Filton Road, Stoke Gifford, BRISTOL, BS34, 8QZ UK
{marco.casassa-mont, stephen.crane}@hp.com

Abstract. People are often required to disclose Personal Identifying Information (PII) in order to achieve their goals, e.g. when accessing services, obtaining information and goods, etc. Being able to say with absolute certainty that another party can be trusted to properly handle personal data with today's technology is probably unrealistic. Feedback solutions based on reputation mechanisms can address aspects of trust and assurance in relation to how personal data is managed by an enterprise. However they usually rely on subjective feedback which is based on empirical experiences, and typically they do not allow individuals to systematically track and manage their specific experience. In this paper we propose an approach that enables people to monitor the status of their personal data which they have previously shared with an enterprise, service provider or other organisation - under specific conditions previously negotiated - and actively gather information on how adequately the management of these data meets their personal expectations. Ongoing monitoring and notification, and the ability of the client to form a simple record of past interaction, provides the client with greater confidence and assurance in situations where they need to share personal sensitive information with organisations they would otherwise not be able to claim they trust. This feedback process is based on conditions that are specific to the process of sharing PII and provides the client with assurance that an enterprise is a) capable and b) actually fulfilling PII processing preferences that are agreed at the time the data is disclosed, and which ultimately enables the client to form an opinion about the service provided. We present the principles of our approach and architectural components that support a practical implementation. This is work in progress and the research is ongoing, carried out in the context of PRIME.

1. Introduction

In this paper we describe a novel approach to providing individuals with assurance that their personal information will be only used as they intended at the time of release. In so doing we address the situation where individuals are not permitted to exploit anonymisation technologies if they are to be given access to the services they desire.

We explain the philosophy behind this new approach, describe the key components in our solution architecture and illustrate a typical practical implementation.

1.1 Problems associated with disclosing personal information

People, when interacting with other parties (e.g. enterprises, other people, etc.) might need to disclose Personal Identifying Information (PII) in order to achieve their goals (e.g. access to a service, obtaining specific information, etc.).

Within PRIME[1][2] we have been investigating how PII can be shared during an interactions, between individuals and between individual and organization, in a way that 1) reassures the individual, who is arguably recognised as the owner of the PII, that their information will not be misused or abused, and 2) gives directives to organisations on how to handle PII data, based on individuals' expectations and privacy preferences. The one factor that underpins the ability to share with confidence is trust. In this paper we describe our work to-date in establishing techniques for managing privacy of PII at initial contact and throughout the lifetime of an interaction.

Being able to say with confidence that another party can be trusted to handle personal information with today's technology is probably unrealistic. Unless we can: 1) completely isolate the processing from the operator; 2) rely on the technology and its implementation, we are left having to rely on our faith in the other party. Condition 1) is unrealistic in practice since virtually every practical application is likely to involve some form of human intervention, including access to the information after the 'trusted' processing is complete. Requirement 2) is currently difficult to demonstrate.

Since individuals have difficulty proving 'before the event' that a recipient is trustworthy and will uphold their wishes, the next best approach (as in real life) is to establish an alternative means of privacy enforcement.

In the real world a contract gives an individual a strong indication that the other party intends to carry out an individual's wishes and provides a means to identify deviation from agreed actions, impose sanctions and potentially guide remediation. Of course, the contract is only useful if it is enforceable; it must be agreed by the parties involved and adhered to.

In the context addressed by this paper, a deceitful recipient of PII will most likely always be able to circumvent controls. However, the concept of a contract is useful for a recipient who has every intention of behaving properly, but still wishes to demonstrate so in order to be differentiated from other less scrupulous recipients. The acceptance of this condition simplifies the enforcement challenge. Large corporate organizations, for the most part, have strong reputation brands which they would like to protect, and so take steps to behave honourably and fairly. These are the organizations that are willing, encouraged or even compelled to demonstrate openness and be held

---

[1] PRIME: PRivacy and Identity Management for Europe. European RTD Integrated Project under the FP6/IST Programme. http://www.prime-project.eu.org/

[2] The PRIME project receives research funding from the Community's Sixth Framework Programme and the Swiss Federal Office for Education and Science. This work was supported by the IST (Information Society Technologies) PRIME project; however, it represents the view of the authors only.

accountable for their actions. It is these organisations that, when interacting with their customers, we have in mind when developing our solution.

1.2 Strategies for preserving privacy

Individuals want to be able to release personal information in the confident belief that it will only be used in the way the individual intended, based on their privacy preferences and expectations. Our emphasis is on the individual as the consumer of a service. Organizations that have valued brand and reputation are keen to 'show' individuals that they can be trusted even if they cannot present indisputable facts that support their claim. Of course, even the best-intended organizations make unintentional mistakes. These organizations would welcome solutions that help them keep in check and reaffirm their own trust in their systems.

PII management options fall along a PII Release Spectrum (Fig. 1). At one extreme there is the situation where a user adopts the approach of not releasing any personal identifying information at all. Instead, the user provides the recipient with information that has passed through some form of anonymiser[3].

At the other end of the spectrum is 'unrestricted release' of identifying information. This approach potentially exposes personal information to the greatest level of abuse, but is common practice nowadays for most commerce and services-based interactions.

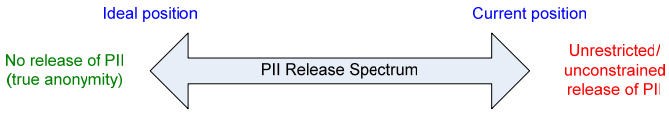

Fig. 1. PII Release Spectrum

The anonymisation approach could be considered the ideal. However, whether the world of commerce is able and willing to adapt existing practices and procedure to the extent that some anonymising techniques demand is still unclear. Furthermore, it is doubtful that complete anonymity is possible in many practical scenarios, e.g. healthcare and travel, where personal information simply must be divulged. Our approach to PII management is to provide the tools both to individuals and organisations that allow PII to be monitored (and to a lesser extent controlled) after it has been released. At its very core, our approach is based on the concepts of feedback and management of historical evidence.

Systems that provide assurance and feedback are not new. In fact feedback mechanisms are now common-place on the Internet, e.g. with service providers such as Amazon (www.amazon.com) and E-bay (www.ebay.com). One criticism of most feedback services is that they are often *subjective* (i.e. they depend on the experiences of other parties – which have to be trusted as well) and difficult to quantify except at the extremes of the feedback, e.g. 0/10 and 10/10. In addition, these systems are usu-

[3] Anomymisation is the de-personalisation of data.

ally general purpose, i.e. they provide feedback on aspects relevant to a wider audience, and not on the specific interests that an individual may have. They require further interpretations in terms of the context surrounding the query. Endorsement schemes like Trust Seal (http://www.trustseal.org/) that provide third-party feedback already exist, but typically validate company credentials rather than the objectives of the individual, and they are not designed to specifically address privacy concerns.

Our aim is to provide an innovative, complementary solution in the privacy space that supports *objective* feedback i.e. feedback tailored to the expectations and needs of individuals, taking into account their own personal experience and attitudes.

2. Addressed Problem

The problem we have addressed in this paper is how to assure individuals that their disclosed personal data are going to be managed according to their expectations and preferences. An important related issue is how to help individuals remember their preferences and expectations so that they can actively check for omissions and violations, and be consistent in controlling what they share.

We aim to address this problem by means of a solution that support the provision of objective feedback, based on clear stated goals and matching evidence and personal experience against these goals. By providing feedback that directly relates to the fulfilment (or not) of specific goals, our solution avoids the need to use an arbitrary aggregated scale.

The challenges that need to be addressed are: 1) identifying assurance constraints that can be imposed by individuals on enterprises and that can subsequently be measured; 2) developing a process that can match these expectations against delivery in both a reactive and proactive manner; 3) ensuring that the feedback service manages personal information in a privacy-aware manner, e.g. it must be able to anonymise feedback information if shared or evaluated by a third-party.

We have developed *client-side* and *enterprise-side* architectures to enable individuals to control their privacy by controlling the release of their Personal Identifying Information (PII). The privacy preserving process essentially employs conditional release mechanism, which permits the explicit expression of conditions and expectations that must be fulfilled by an enterprise to satisfy the terms under which PII is shared. These conditions and expectations – also referred in this paper as *obligations*[8][9] – that represent personal privacy preferences and enable: (1) the client to express, monitor and record the enterprise's fulfilment of its expectations; (2) the enterprise to explicitly understand and implement controls that meet the conditions of release. Privacy obligations dictate the duties and expectations an enterprise has to deal with when handling personal data.

After disclosing personal data to an enterprise or service provider, the client can check the status of data and obtain a concise summary of the behaviour of the enterprise, including evidence derived from past interactions. This is achieved by observing how obligations set for previous releases of PII have been fulfilled.

The proposed system is intended for privacy-aware, collaborative enterprises that recognise the value of handling personal data based on clients' expectations and pri-

vacy laws and the positive impact this can have on their professional reputation and brand.

In our model *preferences* are set by the client and released alongside the personal data to which they relate. Preferences dictate conditions for the release of personal data, such as subsequent deletion of data, periodic notifications of usage, minimisation of attributes, etc. The client also keeps a record of the release in a form that can be automatically analysed on an on-going basis. Clients are provided with tools that enable them to challenge the enterprise and obtain the status of a preference. In addition, preferences may require the enterprise to 'update' the client on the progress and completion of actions relating to their PII, and the status of preferences that are essentially open ended.

3. System architecture

3.1 High-level architecture overview

At the highest level, our architecture involves a simple client-server relationship in which the client releases PII with preferences attached, the details of which are recorded locally. On completion of the transaction relating to the PII, the server responds with a status report which (hopefully) shows that all preferences have been fulfilled. At any time the client can challenge the server with an unsolicited status request. (See Fig. 2.)

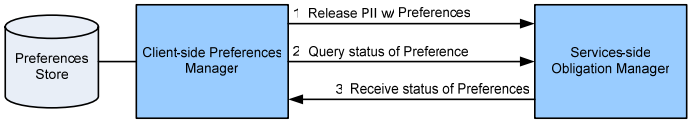

Fig. 2. High-level Client-Server Architecture

3.2 Client side

3.2.1 Overview

As previously explained, clients express their expectations for how the enterprise will handle their personal information through the use of preferences. Preferences define and dictate obligations, each of which requires specific actions that the enterprise is expected to perform. Because obligations dictate specific systems actions and can be rather technical, preferences are used to describe the requirements in human readable terms. Obligations are described in more detail later in this paper and the references. Simple obligations might state "Delete my PII after this transaction" or "Notify me when you share my PII with a third party".

The process in which the client releasing PII to the enterprise therefore consists of four key steps: 1) Setting/selecting preferences and deriving the underlying obligations; 2) Storing a local copy of preferences for future cross-reference; 3) Releasing the PII with the associated preferences and 4) Checking status and recording completion of a required preference. A fifth stage involves the process of recording the outcome for use in influencing future releases of PII to this enterprise, and/or sharing an outcome with other clients. This latter forms the basis of an objective reputation system, described more fully later in this paper. The following figure (Fig. 3) illustrates this process.

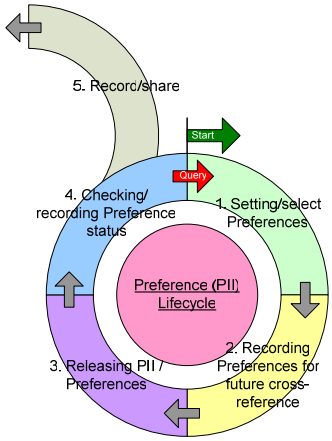

Fig. 3. Process of client releasing PII

3.2.2 Setting preferences

The first action the client does is to decide which obligations the enterprise is to be asked to fulfil. Most users are unlikely to understand the subtlety of this process so, as mentioned above, we provide preference that group together obligations that are likely to be required for a given situation. For example, if the user is shopping online, the PII release requirements may be greater than if they are simply browsing the Web. In practice even preferences could be technically off-putting for most users, so we envisage preferences being provided to the user as a standard set, or by a trusted third party, or even by the enterprise that the user is communicating with. This last category may seem strange in terms of trust, but in practice the set of preferences or obligations available will be limited to that which the enterprise can offer, so it makes sense for the enterprise to state 'up front' what is supported. There is unlikely to be an opportunity for the user to negotiate the detail with the enterprise; they will have to choose to accept what is on offer or decline to proceed. In due course we expect the preference sets to become a de-facto standard.

In an extended model users would specify their preferences, after which the set of obligations would be created by the client system 'on the fly' according to other (typically enterprise) systems parameters.

Having decided on the set of preferences, the next step is to associate these with the PII. Two possibilities exist: 1) the preferences relate to the complete set of PII that is about to be released; 2) each item (or group of items) of PII requires a difference set of preferences. The association must be made locally, and this potential complication simply means that the user must record the preference-PII relationship is more detail. The reason for recording this relationship is that at some time in the future the user will want to check that their preferences have been fulfilled.

The translation from preferences to obligations can be performed by the client or the enterprise. Either way, a standardised way of mapping preferences to obligations is required. It is important to ensure that both parties are working with the same set of underlying obligations. Details of this mapping are not discussed in this paper.

3.2.3 Storing preferences

When preferences are created (or in some cases the resulting obligations are created), a copy of the preferences is placed in the local client-side *store*. The store contains predefined fields that enable the preferences to be processed automatically, e.g. date of creation/fulfilment, evidence of acceptance.

The implication for the user of having to keep a record of all PII released with its associated preferences is that in principle the user could be expected to hold a significant volume of data. In practice there may be duplicate preferences that can be grouped together to reduce the storage overhead. The issue of local storage also raises the need for backup (especially so where a shared client platform is being use), archive and other fundamental security requirements like confidentiality and integrity.

It also seems likely that in practice the set of PII will need to be stored on the client platform along with their preferences, although in principle the PII could be entered 'as required' or sourced from a trusted token with the preferences then assigned at the time of release.

3.2.4 Releasing PII and preferences

Once an appropriate set of preferences has been selected, PII can be released. Various schemes for binding the preferences to the PII can be implemented, but these will most likely only provide evidence and prevent third party interference during transmission. They are unlikely to enhance the trust the client has in the enterprise, although if implemented they could be read by the client as a sign that the enterprise recognises and respects the value that the client has placed in the PII.

3.2.5 Checking status of preferences

The provision of functionality that allows clients to check the status of shared PII is a key requirement. Depending on the nature of preferences associated with PII, different outcomes can be expected. In some cases the client will expect to receive notification that an action has been performed. In others, the client will know when to check for compliance. In the case of a preference that says "Tell me when you share my PII with a third party", the enterprise will be expected to communicate this action to the client. This could take place using either in-band signalling (e.g. a feature of the solution architecture) or out-band signalling (e.g. email).

Checking the status of preferences is likely to be even more onerous for the client than the initial process of setting them. Again, technology can help here by automating the process and making it as transparent as possible. Since the client is aware of the details of the preferences it is also able to challenge the enterprise and record the response. The fact that preferences are recorded in a formalised way using predefined fields means that this process of automation is greatly simplified.

The client therefore need to 'keep an eye' on those preference that are 'active' and be ready to receive notifications from the enterprise. There must also be a way to notify the user when an anomaly arises along, and an agreed set of action that should be performed in order to resolve the situation. The latter point is not discussed further in this paper, other than to say that either the technology could be extended to perform further automated interactions with the enterprise, or the model could be for the matter to be resolved out of band, i.e. manually. In the case where the anomaly is not satisfactorily resolved, other steps may be appropriate, and some of these are discussed later in this paper under reputation.

A typical implementation of this overall release process is shown in Fig. 4. An obligation is set and transferred to the enterprise's *Obligation Management System* (OMS) via standard communication and web service technologies. At the same time a copy of the obligation is transferred to the local *Obligation Store*. The OMS notifies the client as events arise concerning the obligation so that their status can be assessed. At any time the client can solicit the same status information which is assessed in exactly the same way. In this implementation the assessment is influenced by predefined policies that describe how exceptions are to be handled.

A Human Computer Interaction (HCI) component enables the client to interact with the system to 1) understand the details behind a preference status report and 2) determine the overall status of the service provider, possible derived for several interactions over a period of time.

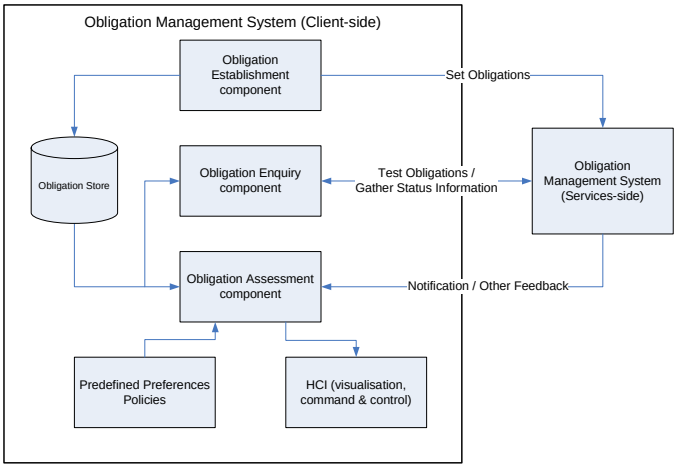

Fig. 4. Client-side Architectural Details

3.2.6 Responding to anomalies - reputation and organisational trustworthiness

Trust that an enterprise will act as expected is built up over time, based in part on past interactions. Evidence that an enterprise is willing to commit to an intended action, possibly in the knowledge that not doing so will incur penalties, is a useful sign of good intentions. Although this paper focus on trust from the viewpoint of privacy, there is a broader application of the concepts discussed in the area of general trust [5][6].

Typically, an individual would either review or present the terms under which the interaction will take place (i.e. a policy or contract). Once accepted, these terms are binding to some degree. As required, the user reviews the interaction and compares outcome against the contract, particularly where the terms specify several points in the process where an assessment can be made (c.f. project milestones). This leads us to the clearly defined process already described, with steps:

- Policy/contract negotiation (between user and organization)
- Fulfilment (by an organization)
- Checking (by an individual)
- Opinion forming (by an individual – essentially retention of evidence to aid trust evaluation during future interactions.)

A third option is where a group of individuals pool their resources and agree to share their experiences based on the status of preferences. This brings forth two advantages: 1) the peer group can help one another to understand the significance of an outstanding preference, and 2) the group start to build a reputation service where they share opinions about the trustworthiness of enterprises they have interacted with.

The idea of a 'home grown' reputation service is potentially particularly a very interesting development. Traditionally reputation services have struggled to provide recommendations that can be evaluated against a common criterion. For example, an individual 'scoring' a server has limited value unless others understand and agree with the underlying scoring rules. Similarly, scoring is influenced by context and personal averseness to risk. We believe our approach has merit because the basis for a trust assessment is clearly defined. Context has a bearing, but this can be captured in the set of preferences pertaining to a given situation, which again can be easily re-interpreted by another individual. Of course, the opportunity to cheat still exists. Either the server or the client could collude in order to affect the shared evaluation, giving a false impression of a server. However, given that the organisation operating the server has entered into the process with stated good intentions, and we have already explained that our approach is not intended to deal with dishonest organisations, this weakness may not be so significant.

The approach we are proposing differs from traditional reputation-based systems and *webs of trust* in that assurance and reputation are based on the fulfilment of the individual's specific expectations, monitored on an on-going basis. We do not rely on how other people interpret how their expectations have been fulfilled. Thus we offer a direct measurement of the trust experienced rather than an indirect one.

In situations where preferences are not fulfilled, and a satisfactory resolution is not apparently possible, publicly publicising the 'reputation' of the enterprise is a useful defence against future breaches of privacy.

3.3 Enterprise-side

Our intention in this section is to provide an indication of how obligations are processed by the enterprise, and how the expectations that clients have attached to their PII could be upheld in practice.

3.3.1 Overview

The enterprise side of our solution aims at: 1) helping enterprises to automate the enforcement of client's expectations and their privacy preferences; 2) providing clients with the required support infrastructure to manage their preferences and obtain fine-grained information about the fulfilment of related privacy obligations. The goal is to provide a solution that supports privacy-aware lifecycle management of stored PII data that is driven by clients' expectations and preferences.

As anticipated in the previous sections, the disclosure of PII data is associated with privacy preferences and constraints that are defined by client. These preferences and constraints are automatically turned into privacy obligations and managed accordingly.

3.3.2 Concept of privacy obligation

Privacy obligations [8][9] are policies that dictate expectations and duties governing how personal data is handled and managed over its lifetime. Privacy obligations include: dealing with data deletion; data transformation (e.g. encryption); sending notifications; executing workflows; etc.

It is important to note that obligation management and enforcement is orthogonal to the management and enforcement of privacy-aware access control policies [3]. For example, deletion of personal data has to happen independently from the fact that this data has ever been accessed.

We define an obligation management model in which privacy obligations are "first class" entities, i.e. they are explicit entities that are modelled and managed. In this model a privacy obligation is an "object" that includes (at least) the following aspects: *Obligation Identifier*; T*argeted Personal Data*; *Triggering Events* (e.g. time-based events); *Actions* (e.g. data deletion, sending notifications). Different categories of privacy obligation need to be managed and enforced by enterprises: *transactional obligations*; *data retention and handling obligations*; *other types of event-driven obligations*. A complementary classification of our managed privacy obligations is based on their activation timeframe and period of validity: *short-term obligations*; *long-term obligations*; *ongoing obligations*.

3.3.3 Obligation management

To deal with the management of privacy obligations, we introduce an *obligation management framework*, based on the following principles:

- *Clients* can explicitly define privacy preferences (e.g. on data deletion, notifications, etc.) on their personal data at the disclosure time (e.g. during a self-

registration process) or at any subsequent time. These preferences are automatically turned into privacy obligations;
- *Enterprise privacy administrators* can further associate other privacy obligations, for example dictated by laws or internal guidelines.

Our obligation management framework handles these obligations by providing the following core functionalities:

- **Scheduling the enforcement of privacy obligations**: it schedules which obligations need to be fulfilled and under which circumstances (events);
- **Enforcing privacy obligations**: it enforces privacy obligations once they are triggered. Enforcement may range from execution of simple actions to complex workflows involving human intervention;
- **Monitoring fulfilment of privacy obligations**: it monitors and audits the enforced obligations, at least for a predefined period of time, to ensure that the desired status of data is not changed and to report anomalies;
- **Supporting overall status and compliance checking of managed obligations**: it consolidates information retrieved during the monitoring and auditing phases to provide a comprehensive checking of the compliance of managed obligations against clients' expectations, privacy laws and other guidelines;
- **Supporting the provision of feedback and notifications to clients**.

Our solution consists of the Obligation Management System (OMS) which provides the obligation management functionalities described in the above obligation management framework. This system is in charge of processing obligations associated to personal data, schedule, enforce and monitor them. Fig. 5 shows the high level architecture of this system.

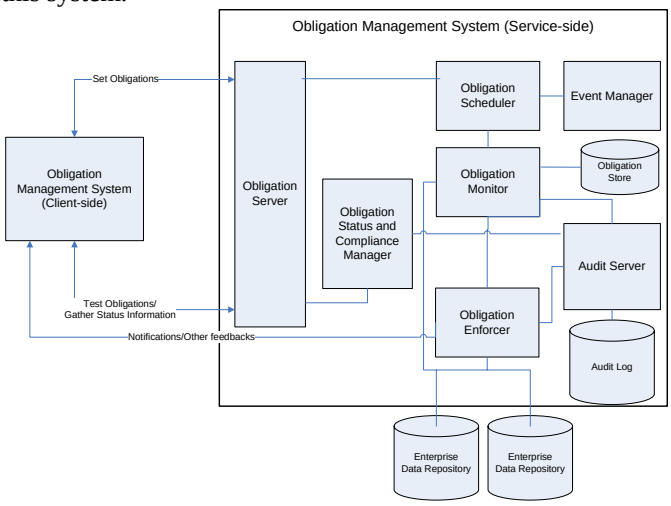

Fig. 5. Enterprise-side Architectural details

The main functional components of the OMS system are:

- **Obligation Server**: this component coordinates the interactions with the external world that involves the setting and management of obligations. This component also provides the capabilities that help clients to gather information about how their personal data is managed;
- **Obligation Scheduler**: this component schedules submitted obligations and listens to relevant events. Should a combination of events trigger an obligation;
- **Obligation Enforcer**: this component enforces obligations. The enforcement consists in executing the obligation Actions;
- **Obligation Monitor**: this component monitors enforced obligations to ensure that the personal data is in the expected status as prescribed by the obligation. This is a key part of the compliance feedback mechanisms that the OMS can provide to the client, along with notifications that can be explicitly requested by clients within their obligations. All these obligation management phases are audited: audit logs are stored in a secure audit system;
- **Obligation Status and Compliance Manager**: this component leverages audit information and internal OMS management data. It aggregates status information of how clients' personal data has been managed, including the successful management of related obligations and any violation. This information is packaged on a client basis and can be directly accessed by the end-client via the Feedback Manager or pushed to it as part of the feedback mechanism.

3.3.4 Enabling the Feedback Mechanism

The Obligation Management System has been designed to have a fine-grained understanding of the privacy preferences and requirements dictated by each client on how their personal data should be managed - in terms of data retention, data deletion, data minimisation and notifications when specific events happens (i.e. accesses to and usage of their data, disclosures, etc.).

The OMS not only is able to enforce these obligations but also actively supports the feedback mechanisms required at the client side. Specifically, the **Obligation Enforcer** and the **Obligation Monitor** support *"operational"* aspects of feedback, i.e. feedback based on the punctual enforcement of privacy obligations or their violations.

The **Obligation Status and Compliance Manager** component instead can provide an *"aggregated"* feedback based on historical events. It collects and manages the aggregation of information on how privacy obligations have been handled over time. It formats this information in a way that it can be processed by users. The client-side will process this aggregated information and compare it against the expected behaviour.

The combination of "punctual" feedback and "aggregated" feedback provides the client-side with information at different level of granularity and relevance on how the enterprise deals with users' expectations and preferences.

4. Related work

Feedback (reputation) systems are not new. Nevertheless, we believe that the consistent challenge that still remains is to ensure that feedback is objective and not subjective. There is likely to always be a human element in the eventual determination, even if the underpinning evidence is indisputable. Examples of solutions that provide feedback based on 'hard evidence' include payment-related reputation and the use of trusted platforms [10].

The work carried by the World Wide Web Consortium's Platform for Privacy Preferences (P3P) [7] that defines a policy language for privacy is relevant to our work. P3P is emerging as an industry standard providing a simple, automated way for users to gain more control over the use of their personal information on the Web sites that they visit. Similarly, AT&T's Privacy Bird[4], which reads privacy policies written in the standard format specified by P3P, is another form of privacy compliance indicator.

We believe that our solution extends the P3P model by creating an *active feedback loop* that enables the individual to have a more active role in understanding how their PII will be used.

5. Current status

The enterprise-side Obligation Management System has already been fully implemented and tested as part of the PRIME project. It has also been integrated, as a proof of concept in a commercial identity management solution [3], in the context of client (and personal data) provisioning and account management. This result provides evidence that it can be integrated with state of the art identity management solutions and that it can interface with the client-side. In terms of client-side component, an early prototype of the *feedback manager* has been build to demonstrate how such system will interact with the client and provide information in a summarised and meaningful way.

The client-side feedback manager and the enterprise-side remote access facilities, including the "Obligation Status and Compliance Manager" component are currently under development. Cryptographic processes to protect and anonymise assurance information when shared with other individuals as part of a wider reputation systems are also currently under development.

6. Open research questions and next steps

As part of the PRIME project we will continue over the next 18 months to research the whole area of feedback-based privacy management. Open questions that remain

[4] For more information about AT&T's Privacy Bird see http://privacybird.com/

research include: How to do it for real? How to leverage current infrastructures? How to increase trust in the statements/evidence provided by the enterprise? How to preserve privacy in shared feedback?

7. Conclusions

We believe that using preferences/obligations to manage personal data disclosed to enterprises is both novel and a pragmatic way to deal with the privacy expectations of business that the average citizen will encounter. We also believe that our work is a pragmatic and effective approach to providing users with assurance about the release of this PII where personalised expectations and past experience are important factors in the successive releases of PII.

In this paper we have explained an approach to providing users with greater assurance in situations where they have no choice but to share their PII. We have chosen to concentrate on the detail of the processes required on the client-side platform rather than the supporting server-side processes. Our approach assumes that the organisation that (in our case) is receiving the personal information is essentially honest and believes that there is merit in demonstrating a respect for the individual's privacy. We have identified some of the problems associated with disclosing PII to unknown parties, and shown that the establishment of agreed preferences is a first step in establishing trust. We introduced the concept of preferences and obligations as a way for an individual to express how their PII should be managed.

About PRIME

PRIME (Privacy and Identity Management for Europe) [1] is the name of a 4-year project, conducted within the EU 6th Framework Programme, which was launched on 1st March, 2004. Its objective is the research and development of solutions to empower individuals in managing their privacy in cyberspace.

PRIME is performing research in the related areas of ontologies, authorisation and trust model, cryptographic mechanisms, secure and privacy-enhancing end-to-end communications, technologies that enable trust in privacy-enhancing IDM solutions, and in assurance through formal evaluations and seals.

PRIME: PRivacy and Identity Management for Europe. European RTD Integrated Project under the FP6/IST Programme. http://www.prime-project.eu.org/

The PRIME project receives research funding from the Community's Sixth Framework Programme and the Swiss Federal Office for Education and Science. This work was supported by the IST (Information Society Technologies) PRIME project; however, it represents the view of the authors only.

References/Bibliography

[1] PRIME Project: Privacy and Identity Management for Europe, European RTD Integrated Project under the FP6/IST Programme, http://www.prime-project.eu.org/, 2005

[2] Marco Casassa Mont; Dealing with Privacy Obligations in Enterprises - HPL-2004-109, 2004

[3] Marco Casassa Mont, Robert Thyne, Kwok Chan, Pete Bramhall; Expanding HP Identity Management Solutions to Enforce Privacy Policies and Obligations for Regulatory Compliance by Enterprises - HPL-2005-110, 2005

[4] Stephen Crane, Marco Casassa Mont, Siani Pearson; On Helping Individuals to Manage Privacy and Trust; HPL-2005-53, 2005

[5] Cofta, Piotr; Crane, Stephen; Towards the Intimate Trust Advisor; First International Conference on Trust Management; May 2003.

[6] Tim Kindberg, Abigail Sellen, and Erik Geelhoed. Security and trust in mobile interactions: A study of user' perceptions and reasoning. Technical Report HPL- 2004-113, HP Laboratories, 2004.

[7] Platform for Privacy Preferences (P3P) Project. http://www.w3.org/P3P/

[8] Casassa Mont, Marco; Dealing with Privacy Obligations: Important Aspects and Technical Approaches; TrustBus 2004; http://www.hpl.hp.com/personal/Marco_Casassa_Mont/Documents/Documents.htm

[9] Casassa Mont, Marco; Dealing with Privacy Obligations in Enterprises; ISSE 2004; http://www.hpl.hp.com/personal/Marco_Casassa_Mont/Documents/Documents.htm

[10] Kinateder, Michael; Pearson, Siani; A Privacy-Enhanced Peer-to-Peer Reputation System; HPL-2004-203