Extending HP Identity Management Solutions to Enforce Privacy Policies and Obligations for Regulatory Compliance by Enterprises♦

M. Casassa Mont, R. Thyne, K. Chan, P. Bramhall
Trusted Systems Laboratory
HP Laboratories Bristol
HPL-2005-110
June 8, 2005*

privacy, IT governance, privacy policy enforcement, privacy-aware access control, privacy obligations, regulatory compliance

This paper describes issues and requirements related to privacy management as an aspect of improved governance in enterprises. It focuses on the privacy enforcement aspect, in particular related to privacy-aware access control and enforcement of privacy obligations: this is still a green field and, at the same time, is a key aspect to be taken into account to ensure compliance both with regulations and an enterprise's IT governance objectives. We introduce our HP Labs work in these areas: core concepts are described along with our policy enforcement models and related technologies. Two prototypes have been built as a proof of concept to: (1) enforce privacy policies on personal data by extending HP Select Access; (2) manage and enforce privacy obligations on personal data, integrated with HP Select Identity. We describe their technical capabilities and our next steps.

* Internal Accession Date Only
♦ HP OVUA 2005, Oporto, Portugal, 10-13 July 2005

Approved for External Publication

© Copyright 2005 Hewlett-Packard Development Company, L.P.

Extending HP Identity Management Solutions to Enforce Privacy Policies and Obligations for Regulatory Compliance by Enterprises

M. Casassa Mont, R. Thyne, K. Chan, P. Bramhall
Hewlett-Packard Labs
Filton Road, Stoke Gifford
Bristol
United Kingdom
{marco.casassa-mont, robert.thyne, kwok.chan, pete.bramhall}@hp.com

Abstract

This paper describes issues and requirements related to privacy management as an aspect of improved governance in enterprises. It focuses on the privacy enforcement aspect, in particular related to privacy-aware access control and enforcement of privacy obligations: this is still a green field and, at the same time, is a key aspect to be taken into account to ensure compliance both with regulations and an enterprise's IT governance objectives. We introduce our HP Labs work in these areas: core concepts are described along with our policy enforcement models and related technologies. Two prototypes have been built as a proof of concept to: (1) enforce privacy policies on personal data by extending HP Select Access; (2) manage and enforce privacy obligations on personal data, integrated with HP Select Identity. We describe their technical capabilities and our next steps.

Keywords

Privacy, IT Governance, Privacy Policy Enforcement, Privacy-aware Access Control, Privacy Obligations, Regulatory Compliance

1. Introduction

Privacy management is important for enterprises and organisations that handle identities and personal data of customers, employees and business partners: it has implications on their compliance with regulations, their reputation and brand [19,20]. Enterprises have been heavily investing in identity management solutions for the last few years: this paper discusses core privacy aspects and requirements that need to be satisfied by these solutions. This includes: authoring, management and enforcement of privacy policies and obligations when provisioning and handling identity information and personal data; auditing and monitoring these policies for compliance.

We focus on the specific problems of enforcing privacy policies and privacy obligations on personal data within enterprises: these areas are still a green field.

Section 2 and 3 describe core privacy management concepts and provides more details about the addressed problems, related issues and core requirements. Section 4 describes related work. Section 5 introduces our work i.e. our privacy policy enforcement models and our technologies.

As a practical demonstration of the feasibility of our work we describe how we leveraged and extended current HP OpenView Identity Management solutions - specifically HP Select Access and HP Select Identity - to respectively enforce privacy policies and privacy obligations. Current results and next steps are illustrated in section 6.

2. Privacy Management

Dealing with privacy is an important aspect of enterprises' regulatory compliance efforts and it is required by law. A lot of work has been done in terms of privacy legislation often driven by local or geographical needs [1,2,3]. Large enterprises that are geographically distributed across different nations might need to comply with different privacy laws.

Privacy policies can be used to represent and describe privacy laws, guidelines and privacy statements. They express rights, permissions and obligations, usually in natural language that needs to be interpreted and understood by people. They need to be enforced and audited.

Most of the technical work currently done in this space focuses on the provision of auditing and reporting solutions to analyse logged events and check them against privacy policies. The enforcement of privacy policies is very important for regulatory compliance. Often privacy policies are hardcoded into applications and services or managed with very vertical, ad-hoc solutions, in specific contexts. This approach is not adaptive to changes and does not scale.

The enforcement of privacy rights, permissions and obligations on confidential and personal data requires the mapping of these concepts (that are most of the time abstract and based on high-level principles) into rules, constraints and access control, the meaning of which must be unambiguous so that it can be deployed and enforced by software solutions. This still requires following best practices and good behaviours. However, automating aspects of the enforcement of privacy policies and reducing the involved costs is important for enterprises.

The (technological) enforcement of privacy permissions and rights requires extended access control and authorization mechanisms on stored personal data that check these privacy permissions against data requestors' rights and intents, data subjects' consent and the stated data purposes [19]. This applies, for example, to enterprise web services or applications that need to access and manipulate personal data for business reasons.

Even more complex is the case of dealing with the enforcement of privacy obligations. Privacy obligations dictate criteria for a privacy-aware lifecycle management of data. They might require the deletion or transformation of confidential data after a predefined (potentially very long) period of time, periodic notifications and request for authorization to data owners or data subjects, fulfilment of opt-in/opt-out choices made by data owners, ongoing compliance with laws' obligations and internal guidelines. The events that trigger the fulfilment of privacy obligations can be completely orthogonal to the ones relevant to privacy permissions. Privacy obligations can have ongoing aspects that need to be monitored and satisfied over a long period of time. All these tasks are challenging for enterprises because of the need for specific IT infrastructures and processes able to manipulate confidential data as dictated by privacy obligations.

3. Addressed Problems, Issues and Requirements

This paper focuses on two core enterprise privacy problems: (1) Privacy policy enforcement on personal data; (2) Privacy obligation management and enforcement.

We address these aspects by analysing and developing a privacy enforcement framework that can be deployed within current enterprise identity management solutions, to leverage current enterprises' investments in this area. In this context, we want to enable privacy management scenarios where data subjects can specify their privacy preferences (that becomes privacy obligations for enterprises), give explicit consent and limitations about the usage of their data and provide them with degrees of control on their personal data. Enterprises must be able to explicitly author, deploy and enforce privacy policies and obligations during accesses, manipulations and transmission of personal data. Enterprises need tools and solutions to achieve this. More details and a list of requirements follow.

3.1 Privacy Policy Enforcement on Personal data

The enforcement of the core privacy principles [1,2,3] on personal data has implications in terms of access control: enterprises must state the purposes for collecting data and data must be accessed only for that reasons. The consent given by data subjects impose limitations on how these data are accessed. Similarly, the limitations on data usage, disclosure and retention dictate conditions and constraints that need to be satisfied before accessing personal data.

Traditional access control systems are necessary but not sufficient to enforce privacy policies on personal data. They are mainly based on "access control lists" and enforcement mechanisms that keep into account the identities of data requestors, their rights and permissions and the types of actions that are allowed/disallowed on the involved resources (data resources). These systems do not keep into account additional aspects relevant to privacy enforcement: the stated data purposes and data subjects' consent - i.e. properties usually associated to collected data - the intent of data requestors and any additional enterprise or customized data subjects' constraints.

It is necessary to build "privacy extensions" of traditional access control systems that can author and enforce privacy policies. To address the above issues and move towards privacy-aware access control systems, it is important to satisfy the following core requirements: (1) Explicit modeling of personal data stored by enterprises; (2) Explicit definition, authoring and lifecycle management of related privacy policies; (3) Explicit deployment and enforcement of privacy policies; (4) Integration with traditional access control and identity management systems; (5) Simplicity of usage of all the involved system; (6) Support for auditing. A more comprehensive analysis and discussion of these aspects can be found in [19].

3.2 Privacy Obligation Management and Enforcement

Privacy obligations on personal data can be defined by people (data subjects), by laws and by enterprises. Privacy obligations dictate responsibilities on how data has to be handled and processed, given specific contexts, for example with respect to disclosure of personal information. Obligations can be expressed in terms of notice requirements, opt-out options, limits on reuse of information and information sharing for marketing purposes. Privacy obligations can dictate very specific requirements. For example privacy obligations can require that personal data must be deleted after a predefined number of years, e.g. 30 years, (long-term commitment) or in a few days if user's consent is not granted (short-term commitment) or their account is closed. Privacy obligations can have "ongoing" and long-term commitments for enterprises or might apply only for a short period of time and be transient. The enforcement of privacy obligations can be independent from access control. For example, the deletion of personal data after 7 years has to happen independently of whether or not these data have ever been accessed. It is important that privacy obligation management solutions address the following core requirements: (1) Explicit modeling and representation of privacy obligations; (2) Association of obligations to data; (3) Being able to timely enforce privacy obligations; (4) Mapping obligations into enforceable actions; (5) Compliance of refined obligations to high-level policies; (6) Tracking the evolutions of obligation policies; (7) Dealing with long-term obligation aspects; (8) Accountability management; (9) Monitoring obligations; (10) User involvement; (11) Complexity and cost of instrumenting applications and services. A comprehensive analysis and discussion of these aspects can be found in [20,21].

4. Related Work

A common approach to enforce privacy policies on personal data consists of hardcoding them within applications and services or building ad hoc solutions. This approach is suitable for very simple and static environments: it shows all its limitations and maintenance costs in case of complex and dynamic organizations that need to adapt to changes. As described in the requirements section, to explicitly address the problem, a model of the relevant personal data is required. Privacy policies dictating how these data must be accessed need to be authored, deployed, enforced and audited. This requires the definition of a comprehensive privacy-aware access control model and systems that implement it.

Relevant work in this direction, for privacy management and enforcement in enterprises is described in [4,5,6,7]. An Enterprise Privacy Architecture is introduced and described in [7]. This approach is further refined and described in the Enterprise Privacy Authorization Language (EPAL) specification [8]. These papers provide general guidelines.

Important related work on actual privacy enforcement on personal data has been done by IBM with their research on Hippocratic databases [9]. The drawback of this approach is that it mainly focuses at the database level, specifically on RDBMS data repository architectures and related data schemas. The enforcement of privacy policies might need to span across a broad variety of data repositories and legacy systems to include LDAP directories, meta and virtual directories, file systems and legacy systems. It might need to incorporate higher-level views and perspectives than just the database-level perspective.

In terms of commercially available solutions, IBM Tivoli Privacy Manager [10, 11] provides mechanisms for defining fine-grained privacy policies and associating them to data. On one hand this solution provides the required privacy enforcement functionalities. On the other hand this approach dictates strong constraints on how applications need to be developed and how personal data has to be stored and administered: it might require some duplications of administrative and enforcement frameworks (it requires the parallel usage of Tivoli Access Manager) and it is vertically-based on other IBM products and solutions.

Other products, such as HP Select Federation [12] and ePok [13], focus on single-sign-on and related privacy aspects: they enforce privacy rules on personal data in federated environment when these data are disclosed by an organization (or an identity provider) to other parties.

Our work specifically addresses the problem of enforcing privacy policies on personal data stored in a broad variety of data repositories and used *within* enterprises. Personal data can be accessed by different types of requestors, including people, applications and services. It includes related aspects of modeling the managed data and authoring privacy policies. Our work aims at not being invasive for applications and services: privacy policies are managed in an explicit way, in conjunction with traditional access control policies and not hardcoded in applications and services. We want to avoid duplications of efforts by providing a single, integrated framework for authoring, administering and enforcing both traditional access control and privacy policies. To demonstrate the feasibility of this, as a significant example, we leveraged and extended HP Select Access [14] to enforce privacy policies on personal data.

In terms of managing and enforcing privacy obligations, relevant work is described in [4,5,6,7,8], in particular the EPAL specification. Their approach to handle privacy obligations is driven by an authorization and access control perspective. However, privacy obligations cannot be managed at their best only from an authorization-based perspective. Privacy obligations can include aspects that are not driven by data accesses, for example the deletion of data.

We believe that modularity and separation of concerns are important aspects. In our approach obligation policies are first-class citizens with their explicit management, as a self-standing component of a more comprehensive policy management framework. Our architecture has high-level commonalities with the architecture described in [4,5,6,7] but in our work we further refine the concept of obligations and their enforcement. We split the enforcement mechanisms in two parts by including a scheduling mechanisms and an enforcement mechanism allowing for workflow automation and human intervention.

Approaches to deal with (privacy) obligations have already been implemented in products, in particular for data retention [15] and in a variety of document management systems. Nevertheless, these approaches are very specific, focused on particular domains and handle simple obligation policies on files and documents, not really on personal data. Our work aims at pushing the barrier even further to create an obligation management framework that can be leveraged in multiple contexts, for different purposes.

A lot of work has been done in representing privacy policies, including obligations such as [16,17]. Relevant work on mechanisms to associate policies to data is described in [4,5,6,7,18]. We can leverage aspects of this work, in particular [18] to provide a stronger association of obligation policies to confidential data.

5. Our Work

This section provides technical details of HP Labs work to enforce privacy policies and privacy obligation on personal data stored within enterprises. To demonstrate the feasibility and deployability of our work in real world identity management solutions, as a significant example we deployed our prototypes within HP Identity Management solutions: we extended HP Select Access to deal with privacy policy enforcement on personal data, integrated with the enforcement of “traditional” access control policies. We have also implemented a prototype of an obligation management system, integrated with HP Select Identity, to represent, schedule, enforce and monitor privacy obligations.

5.1 Privacy Policy Enforcement and Integration with HP Select Access

Our approach to enforce privacy policies is based on a privacy-aware access control model that extends traditional access control models (based on users/groups, users’ credentials and rights, access control lists and related policies) by explicitly dealing with the stated purposes for which data is collected, checking - at the access request time - the intent of requestors against these purposes, dealing with data subjects’ consent and enforcing additional access conditions and constraints on personal data defined by data subjects and/or enterprise administrators [1,2,3] –

see Figure 1.

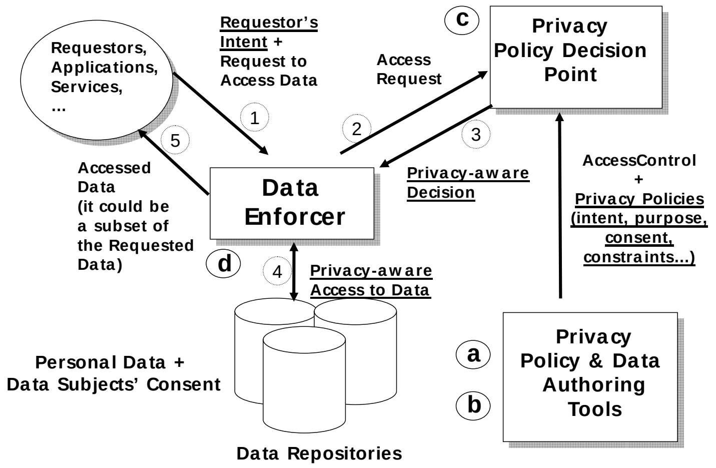

Figure 1: Model of our Privacy-aware Access Control System

The main aspects of this model are:

a. A mechanism for the explicit modelling of personal data that are subject to privacy policies: this mechanism provides a description of data including the type of the data repository (database, LDAP directory, etc.), its location, the schema of these data, types of attributes, etc.;

b. An integrated mechanism for authoring privacy policies along with traditional access control policies: it is a Policy Authoring Point (PAP) to allow privacy administrators to describe and author privacy policy constraints and conditions (including how to check consent and data purpose against requestors' intent and how to deal with data filtering and transformation, etc.) along with more traditional access control policies based on security criteria (such as who can access which resource, given their rights and permissions);

c. An integrated authorization framework for deploying both access control and privacy-based policies and making related access decisions: it is an integrated Policy Decision Point (PDP);

d. A run-time mechanism –referred to as the "Data Enforcer" - for intercepting attempts to access personal data and enforcing decisions based on privacy policies and contextual information, e.g., intent of requestors, their roles and identities, etc. It is a Policy Enforcement Point (PEP). This mechanism is in charge (among other things) of dealing with the transformation of queries to access personal data (e.g. SQL queries) and filtering part of the requested data, if their access is not authorised for privacy reasons.

The data enforcer plays a key role to enforce privacy policies on personal data. At "run-time", attempts to access personal data are intercepted and managed in the following way - Figure 1:

1. A request from a data requestor to access personal data is intercepted by the data

enforcer. Available information about the requestor (credentials, identity, etc.) is collected, along with their intent (that can be explicitly passed as a parameter or could be predefined in the application/service making the request);

2. The data enforcer interacts with the privacy policy decision point by passing information about the request (including the intent) and the requestor;
3. The privacy policy decision point makes a decision, based on available privacy policies and the context (request, requestor's information, etc.). This decision is sent back to the data enforcer. It can be any of the following types:
   - No: access to data is denied;
   - No & conditions: access to data is denied. Some conditions are sent back to the requestors. The satisfaction of these conditions (for example passing the intent or authenticating) could change the outcome of the decision;
   - Yes: access to data is granted;
   - Yes & conditions: access to (part of the) data is allowed, under the satisfaction of the attached conditions. Among other things, these conditions might require data filtering, transformations and manipulations.
4. The data enforcer enforces this decision. In particular, if the decision is "Yes & conditions" the data enforcer might have to manipulate the query (query pre-processing) and/or transform the requested personal data (result prost-processing), before returning the result to the data requestor;
5. Data (or alternatively no data) is returned to the data requestor, based on the enforced decision.

Figure 2 shows a simple example based on this model where an attempt to access personal data is made by an enterprise employee. In this example, the employee's intent (marketing) is consistent with the declared purposes of data (marketing and research). However the employee is trying to access – via an SQL query - more data than she is allowed to. The SQL query is intercepted by the enforcement point (data enforcer) and transformed in a way to include constraints based on data subjects' consent and the filtering of data. In this example privacy is achieved by pre-processing and transforming the query before actually interacting with the database.

To implement the above model we leveraged and extended HP Select Access. HP Select Access [14] is a leading-edge access control product. It provides policy authoring, policy decision and policy enforcement capabilities via the following components:

- Policy Builder: it is a graphical tool to author access control policies (PAP) on resources managed by the system;
- Validator: it is a Policy Decision Point (PDP). It makes access control decisions based on the access control policies (authored with the Policy Builder) and contextual information, such as the identity of a requestor;
- Web Enforcer plug-in: it is a Policy Enforcement Point (PEP) for web resources.

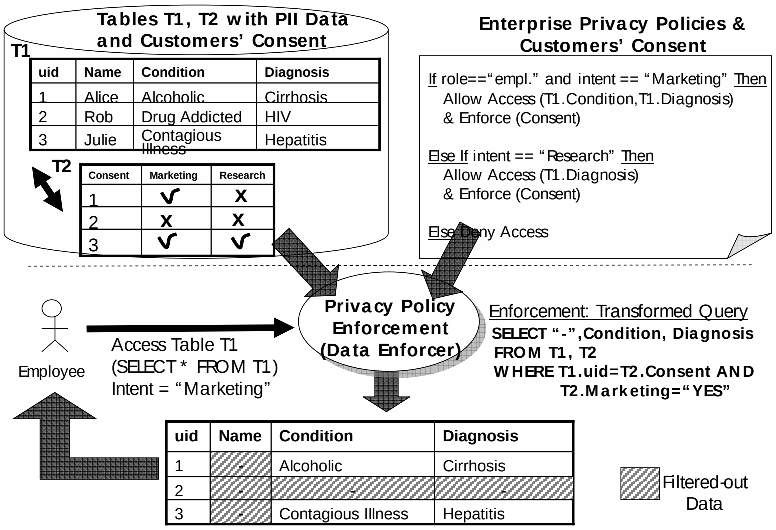
Tables T1, T2 with PII Data and Customers' Consent
T1
T2
Enterprise Privacy Policies & Customers' Consent
Privacy Policy Enforcement (Data Enforcer)
Employee
Access Table T1
(SELECT * FROM T1)
Intent = "Marketing"
Enforcement: Transformed Query
Filtered-out Data

uid	Name	Condition	Diagnosis
1	Alice	Alcoholic	Cirrhosis
2	Rob	Drug Addicted	HIV
3	Julie	Contagious Illness	Hepatitis

Consent	Marketing	Research
1	✓	X
2	X	X
3	✓	✓

```
If role=="empl." and intent == "Marketing" Then
   Allow Access (T1.Condition,T1.Diagnosis)
   & Enforce (Consent)

Else If intent == "Research" Then
   Allow Access (T1.Diagnosis)
   & Enforce (Consent)

Else Deny Access
```

```
SELECT "-",Condition, Diagnosis
FROM T1, T2
WHERE T1.uid=T2.Consent AND
        T2.Marketing="YES"
```

uid	Name	Condition	Diagnosis
1	▨	Alcoholic	Cirrhosis
2	▨	▨	▨
3	▨	Contagious Illness	Hepatitis

Figure 2: Example of Privacy Policy Enforcement

The current commercial version of HP Select Access does not handle data as managed resources: it only deals with traditional access control policies on web resources. New functionalities have been added to HP Select Access in our prototype, to explicitly deal with privacy-aware access control on personal data, as shown in Figure 3.

The following extensions of HP Select Access have been implemented in our prototype:

- **The HP SA Policy Builder has been extended to represent "data resources"** (databases, LDAP directories, virtual-directories, their schemas, etc.) in addition to traditional IT resources (such as web resources);
- **The HP SA Policy Builder has been extended to graphically author privacy policies** on "data resources" in addition to traditional access control policies: a set of additional plug-ins has been implemented, including the ones that check (at the enforcement time) the requestor's intent against the stated data storage purposes, take into account data subjects' consent & data retention policies and describe how the accessed personal data must be filtered, obfuscated or manipulated, etc.;
- **The HP SA Validator has been extended to make privacy-aware decisions**. Plug-ins, correspondent to the ones used in the Policy Builder, have been implemented. This enhanced-version of the Validator can now make "Yes & constraints" decisions as described in our model;
- **A Data Enforcer has been built and added to the framework**: this is a new functionality added to HP Select Access. It is in charge of enforcing privacy decisions made by the Validator. It intercepts incoming calls to data resources, interacts with the Validator, performs fine grained manipulation of data resources and deals with the interpretation and enforcement of additional constraints as defined by the privacy policies.

The data enforcer sits nearby managed data repositories (e.g. databases, LDAP directories, virtual directories, etc.): we envisage that a family of data enforcers (sharing a common logic but differentiated by add-ons dealing with different types of data resources) need to be built, because of the different semantic of different data repositories. The data enforcer currently implemented is a JDBC proxy for RDBMS databases.

The above functionalities address and satisfy the core requirements described in section 3 for privacy enforcement on personal data. More details can be found in [19].

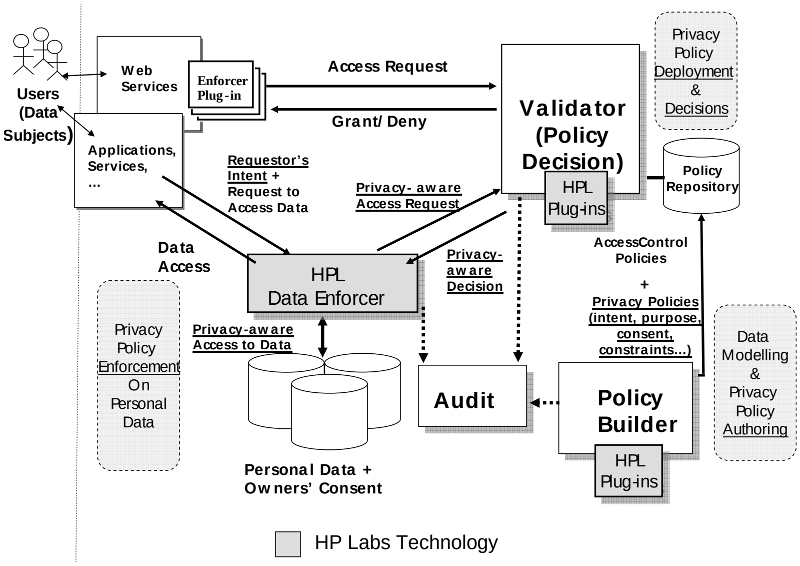
Users (Data Subjects)
Web Services
Enforcer Plug-in
Applications, Services, ...
Access Request
Grant/ Deny
Validator (Policy Decision)
HPL Plug-ins
Privacy Policy Deployment & Decisions
Policy Repository
Requestor's Intent + Request to Access Data
Privacy- aware Access Request
Data Access
HPL Data Enforcer
Privacy- aware Decision
AccessControl Policies
+
Privacy Policies (intent, purpose, consent, constraints...)
Privacy Policy Enforcement On Personal Data
Privacy-aware Access to Data
Audit
Policy Builder
HPL Plug-ins
Data Modelling & Privacy Policy Authoring
Personal Data + Owners' Consent
HP Labs Technology

Figure 3: Extended HP Select Access to deal with Privacy Policy Enforcement

5.2 Privacy Obligation Management and Integration with HP Select Identity

Our technical work focuses on the explicit management and enforcement of privacy obligations for personal data stored by enterprises. In our model, privacy obligations are "first class" entities, i.e. they are explicit entities that are modeled and managed to provide a privacy-aware lifecycle management of personal data: this includes data deletion, data transformation, dealing with notifications, etc. A related obligation management framework is introduced to manage these privacy obligations.

From a technical perspective, a privacy obligation is an "object" (currently expressed in XML) that includes (at least) the following aspects:

- **Obligation Identifier:** it is a unique identifier of the privacy obligation, used by the obligation management framework;
- **Targeted Personal Data**: it contains specific links to the personal data of pertinence, including information on where data are stored and how they can be uniquely retrieved;
- **Triggering Events**: it describes one or more events that are relevant to trigger the

fulfillment of the obligation;

- **Actions**: it describes the set of actions that need to be executed when the obligations has to be enforced;

Different categories of privacy obligation need to be managed and enforced:

- **Transactional obligations**: privacy obligations that need to be immediately enforced, when transactions and interactions involve personal and confidential data. For example, they might require to notify the data subject or create audit logs every time personal data is accessed;
- **Data retention and handling obligations**: these privacy obligations describe criteria for the management and deletion of personal data, usually driven by time-based events. For example, they might require the deletion of data after a predefined period of time (ranging from days to years) or at a specific time agreed with the data subject;
- **Other types of event-driven obligations**: these privacy obligations are triggered by events that relate to contextual and application-relevant information, based on usage of personal data, trust information about the systems dealing with personal data, etc.

A complementary classification of our managed privacy obligations is based on their activation timeframe and period of validity:

- **Short-term obligations**: these obligations have a short period of validity;
- **Long-term obligations**: these obligations might have long term implications in terms of resources needed for their fulfillment (months or years);
- **Ongoing obligations**: these obligations might be short or long termed. They imply an ongoing, periodic, fulfillment of activities related to the management of personal data. For example they might require sending periodic notifications to data subjects about the status of their personal data.

Figure 4 shows the conceptual model underpinning our Obligation Management Framework.

Data subjects can explicitly define privacy obligations on their personal data at the disclosure time (e.g. during a self-registration process) or at any subsequent time. 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**: the system schedules which obligations need to be fulfilled and under which circumstances (events);
- **Enforcing privacy obligations**: the system enforce privacy obligations once they are triggered. The enforcement ranges from the execution of simple actions to complex workflow involving human interventions;
- **Monitoring the fulfilment of privacy obligations**: the system 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.

More details can be found in [20,21]. These functionalities can be accessed by enterprise privacy administrators and potentially also by data subjects, for example to monitor their personal data and check for privacy compliance.

Figure 5 shows the high-level architecture of our obligation management system, based on the model shown in Figure 4.

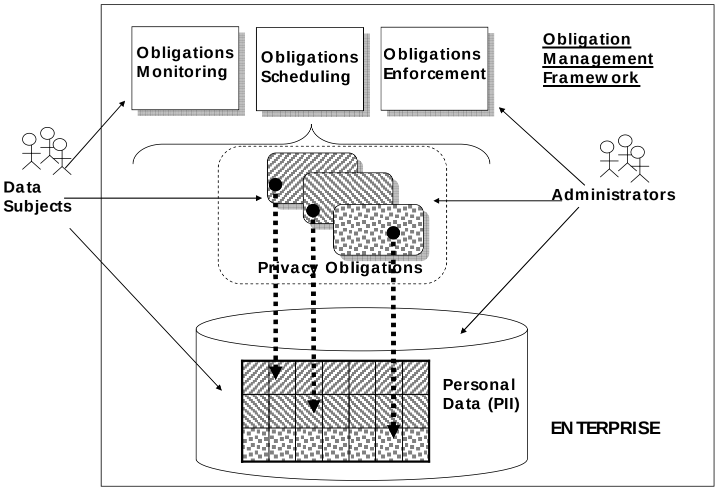

Figure 4: High-level model of our Obligation Management Framework

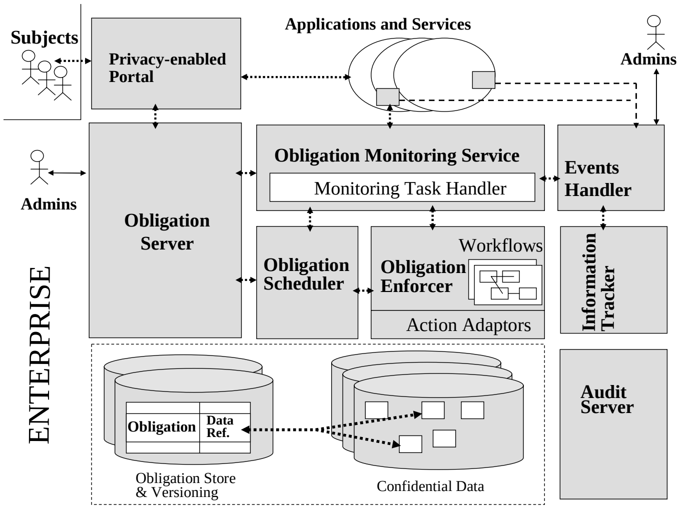

Figure 5: High-level Architecture of our Obligation Management System

Our obligation management system consists of the following modules:

- **Obligation Server**: it deals with the authoring, management and storage of obligations. It explicitly manages the association of privacy obligations to confidential data and their tracking and versioning. It pushes active obligations (i.e. obligations to be fulfilled) to the "obligation scheduler". One or more obligation servers can be deployed (and synchronised), depending on needs;

- **Obligation Store and Versioning**: it stores obligations and their mapping to confidential data. Multiple versions of obligations are also stored in this system;
- **Obligation Scheduler**: it is the component that knows which obligations are active, ongoing obligation deadlines, relevant events and their association to obligations. When events/conditions trigger the fulfilment of one or more obligations, this component activate the correspondent "workflow processes" of the "obligation enforcer" that will deal with the enforcement of the obligation;
- **Obligation Enforcer**: it is a workflow system containing workflow processes describing how to enforce one or more obligations. The enforcement can be automatic and/or could require human intervention, depending on the nature of the obligation;
- **Events Handler**: it is the component in charge of monitoring and detecting relevant events for privacy obligations and sending them to the obligation scheduler. It coordinates its activities with other instrumented components;
- **Obligation Monitoring Service**: it is orthogonal to the scheduling and enforcement components and monitors enforced obligations and the expected status of data;
- **Information tracker**: it is a component that focuses on intercepting events generated by data repositories, databases and file systems containing confidential data and providing this information to the event handler.;
- **Audit Server**.

A working prototype has been implemented in the context of the EU PRIME project [22], as a proof of concept, providing the core functionalities: scheduling, enforcement and monitoring of privacy obligations. At the moment the managed obligations are restricted to handling time-based and access based events. The supported actions include deletion of data and notifications. Short-term, long-term and ongoing obligations are supported. Our work addresses the core issues and requirements described in section 3. More details can be found in [20,21].

We believe that an obligation management system should be considered as an additional component of current enterprises' identity management solutions. These solutions already provide identity management functionalities for identity federation management, user provisioning and account management, access control and privacy management that can be leveraged. Our obligation management system can be integrated with the self-registration, customization and account management capabilities of identity provisioning systems to allow users and administrators to describe and handle privacy preferences and turn them into privacy obligations for the enterprise. In this context our system allows for the explicitly representation and management of privacy obligations, along with the coordination of their overall enforcement and monitoring.

To demonstrate how this can be achieved for real, we integrated our Obligation Management System with HP Select Identity, as shown in Figure 6. HP Select Identity [23] is a state-of-the-art solution to manage digital identities within and between large enterprises. The Select Identity solution automates the process of provisioning, managing and terminating user accounts and access privileges across platforms, applications, and corporate boundaries. Specifically, the key features of the Select Identity system include:

- Centralized Management: provides a single point of control for the management of users and entitlements;

- Provisioning: automates the creation, update, and deletion of accounts and entitlements on information systems across the enterprise;
- Administrative Delegation: enables administrative rights to be distributed among multiple tiers of functional departments, customers, and partners;
- User Self Service: enables end users to initiate access to Services, change passwords, set password hints, and update general identity information through a web browser interface;
- Approval Workflow: automates approval processes required for granting access privileges to users;
- Password & Profile Management: manages and distributes password and user profile information across and between enterprise information systems;
- Audit and Reporting: provides standardized reporting on actions and user account activity.

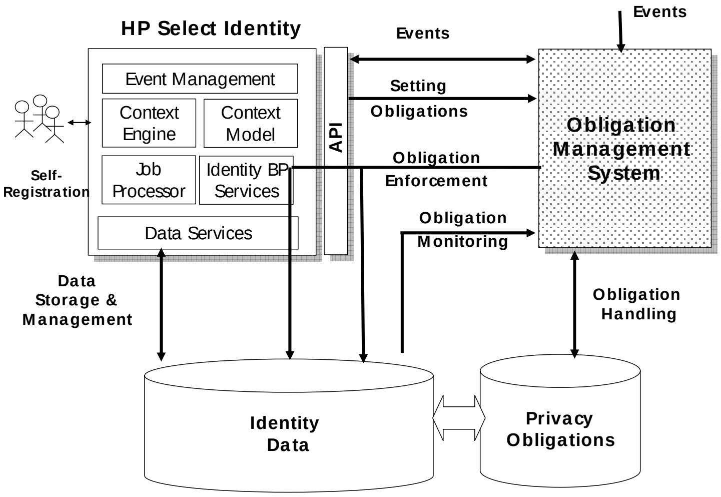

Figure 6: High-level Architecture: Integration of OMS with HP Select Identity

In our integrated prototype we use HP Select Identity self-registration and user provisioning capabilities to specify and capture privacy constraints and preferences on how to handle personal data. These preferences are then processed by our obligation management system that transforms them into privacy obligations. Privacy obligations are then scheduled, enforced and monitored by our system. We leverage the workflow and user/identity management capabilities of HP Select Identity to enforce aspects of privacy obligations. Our system retains control of the supervision of obligations and their monitoring. HP Select Identity enforces obligations constraints, such as deletion of identities, data transformation, etc. At the moment the deletion of personal data (as the effect of enforcing obligations) is achieved by triggering HP Select Identity workflows, whilst the obligation management system handles the notifications to users.

6. Discussion and Next Steps

Our prototypes are proof of concepts. They show the feasibility of our work in addressing core issues and requirements in real contexts: we are refining and extending them for their potential future productisation by HP businesses. A working demonstrator (shown at RSA 2005 and DIDW 2005), based on a healthcare scenario and using both prototypes, has been implemented to show the integration of privacy policy enforcement and obligation enforcement in an identity management context.

At the moment the enforcement of privacy policies in HP Select Access mainly consists in enforcing data subjects' consent, constraints on data purposes and data expirations via data filtering. This has been achieved by intercepting and transforming incoming SQL queries by our data enforcer (query pre-processing). Current performance tests and analysis (done on databases of sizes from 100K to 500K records) are promising. No noticeable loss of performance (i.e. the time spent between sending a query to a RDBMS and retrieving the last returned record) has been registered so far, on common SQL queries. More tests and experiments are in progress on different varieties of SQL queries. We are also planning to: (1) explore the implications of post-processing queries (post-processing of query results) to extend the current set of managed privacy constraints; (2) explore the enforcement of privacy policies on LDAP repositories and virtual directories.

In terms of privacy obligation enforcement, we are currently refining the integration of our obligation management system with HP Select Identity, specifically to leverage as much as possible the provisioning and workflow capabilities of HP Select Identity to enforce obligations' actions. Additional work and research in the space of privacy obligations is going to be done in the context of the EU PRIME project [22]: in particular we plan to work in the area of stickiness of privacy obligations to personal data, management of complex obligation actions, end-to-end graphical management of privacy obligations, compliance feedback and longevity/survivability of the obligation management system.

7. Conclusions

Privacy management is becoming more and more important for enterprises to ensure their compliance to regulation, their governance objectives and address customers' needs and rights.

This paper focuses on privacy policy and obligations enforcement for personal data stored and accessed by enterprises: these aspects are still a green field. We discussed a privacy-aware access control model to enforce privacy constraints (including handling the purpose of data, checking data requestors' intent against data purposes and enforcement of data subjects' consent). We also analysed aspects and concepts related to privacy obligations, considered as "first-class" entities (including data deletion, data transformation, notifications, etc.) and introduced our obligation management framework to schedule, enforce and monitor them.

Working prototypes have been implemented and integrated with state-of-the art identity mangement solutions: specifically we described our work to add privacy policy enforcement to HP Select Access and obligation management and enforcement capabilities to HP Select Identity. These technologies are ready for commercial exploitation. Research and development work continues to refine our technolgies and implement adiditional functionalities.

References

[1] C. Laurant, “Privacy International: Privacy and Human Rights 2003: an International Survey of Privacy Laws and Developments, Electronic Privacy Information Center (EPIC)”, Privacy International. http://www.privacyinternational.org/survey/phr2003/ 2003

[2] Online Privacy Alliance, “Guidelines for Online Privacy Policies”, http://www.privacyalliance.org/, Online Privacy Alliance, 2004

[3] OECD, “OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data.”, http://www1.oecd.org/publications/e-book/9302011E.PDF, 1980

[4] G. Karjoth, M. Schunter “A Privacy Policy Model for Enterprises”, IBM Research, Zurich. 15th IEEE Computer Foundations Workshop, 2002

[5] G. Karjoth, M. Schunter, M. Waidner, “Platform for Enterprise Privacy Practices: Privacy-enabled Management of Customer Data”, 2nd Workshop on Privacy Enhancing Technologies, Lecture Notes in Computer Science, Springer Verlang , 2002

[6] M. Schunter, P. Ashley, “The Platform for Enterprise Privacy Practices”, IBM Zurich Research Laboratory, 2002

[7] G. Karjoth, M. Schunter, M. Waidner, “Privacy-enabled Services for Enterprises”, IBM Zurich Research Laboratory, TrustBus 2002, 2002

[8] IBM, “The Enterprise Privacy Authorization Language (EPAL), EPAL 1.1 specification”, http://www.zurich.ibm.com/security/enterprise-privacy/epal/, IBM, 2004

[9] R. Agrawal, J. Kiernan, R. Srikant, Y. Xu, “Hippocratic Databases”, http://www.almaden.ibm.com/cs/people/srikant/papers/vldb02.pdf, IBM Almaden Research Center, 2002

[10] IBM Tivoli Privacy Manager, “Privacy manager main web page”, http://www-306.ibm.com/software/tivoli/products/privacy-mgr-e-bus/, 2005

[11] IBM Tivoli Privacy Manager, ”online technical documentation”, http://publib.boulder.ibm.com/tividd/td/PrivacyManagerfore-business1.1.html, 2005

[12] HP, “HP Select Federation - Product and Solution Overview”, http://www.managementsoftware.hp.com/products/slctfed/, 2005

[13] ePok, “identity management solution - Trusted Data Exchange Server”, http://www.epokinc.com/, 2005

[14] HP, “HP OpenView SelectAccess - Overview and Features”, http://www.openview.hp.com/products/select, 2005

[15] IBM, “IBM Tivoli Storage Manager for Data Retention”, 2004

[16] C. Bettini, S. Jajodia, X. Sean Wang, D. Wijesekera, “Obligation Monitoring in Policy Management”, 2002

[17] N. Damianou, N. Dulay, E. Lupu, M. Sloman, “The Ponder Policy Specification Language”, 2001

[18] M. Casassa Mont, S. Pearson, P. Bramhall, “Towards Accountable Management of Privacy and Identity Information”, ESORICS 2003, 2003

[19] M. Casassa Mont, R. Thyne, Pete Brmhall, “Privacy Enforcement with HP Select Access for Regulatory Compliance”, HPL-2005-10, 2005

[20] M, Casassa Mont, “Dealing with Privacy Obligations: Important Aspects and Technical Approaches”, TrustBus 2004, 2004

[21] M. Casassa Mont, “Dealing with Privacy Obligations in Enterprises”, ISSE 2004, 2004

[22] PRIME, “Privacy and Identity Management for Europe, European RTD Integrated Project under the FP6/IST Programme”, http://www.prime-project.eu.org/, 2004

[23] HP, “HP OpenView SelectIdentity – Overview and Features”, http://www.openview.hp.com/products/slctid/index.html, 2005