A Systemic Approach to Privacy Enforcement and Policy Compliance Checking in Enterprises

Marco Casassa Mont, Siani Pearson, Robert Thyne
Trusted Systems Laboratory
HP Laboratories Bristol

HPL-2006-44
March 16, 2006*

privacy, privacy policy enforcement, policy compliance, automation, governance, identity management

Privacy management is important for enterprises that handle personal data: they must deal with privacy laws and people's expectations. Currently much is done by means of manual processes, which make them difficult and expensive to comply. Key enterprises' requirements include: automation, simplification, cost reduction and leveraging of current identity management solutions. This paper describes a suite of privacy technologies that have been developed by HP Labs, in an integrated way, to help enterprises to automate the management and enforcement of privacy policies (including privacy obligations) and the process of checking that such policies and legislation are indeed complied with. Working prototypes have been implemented to demonstrate the feasibility of our approach. In particular, as a proof-of-concept, the enforcement of privacy policies and obligations has been integrated with HP identity management solutions. Part of this technology is currently under productisation. Technical details are provided along with a description of our next steps.

* Internal Accession Date Only Approved for External Publication
© Copyright 2006 Hewlett-Packard Development Company, L.P.

A Systemic Approach to Privacy Enforcement and Policy Compliance Checking in Enterprises

[(1)] Marco Casassa Mont, [(1)] Siani Pearson, [(2)] Robert Thyne

[(1)] Hewlett-Packard Labs, Trusted Systems Lab
Bristol, United Kingdom

[(2)] Hewlett-Packard, Software Business Organisation
Toronto, Canada

{marco.casassa-mont, siani.pearson, robert.thyne}@hp.com

Abstract. *Privacy management is important for enterprises that handle personal data: they must deal with privacy laws and people's expectations. Currently much is done by means of manual processes, which make them difficult and expensive to comply. Key enterprises' requirements include: automation, simplification, cost reduction and leveraging of current identity management solutions. This paper describes a suite of privacy technologies that have been developed by HP Labs, in an integrated way, to help enterprises to automate the management and enforcement of privacy policies (including privacy obligations) and the process of checking that such policies and legislation are indeed complied with. Working prototypes have been implemented to demonstrate the feasibility of our approach. In particular, as a proof-of-concept, the enforcement of privacy policies and obligations has been integrated with HP identity management solutions. Part of this technology is currently under productisation. Technical details are provided along with a description of our next steps.*

1 Introduction

Enterprises that handle identities and personal information of data subjects (i.e. customers, employees and business partners) are coming under increasing pressure to improve privacy management, both to satisfy people's expectations and to comply with privacy laws and internal policies. Ultimately, the way they manage privacy aspects has implications for their reputation and brand.

Privacy laws, such as HIPPA, COPPA, EU Data Protection Directives [13] and privacy guidelines, such as OECD [14], dictate that enterprises should clearly state the purposes for which they are collecting personal data and should take into account the consent (or lack of consent) given by data subjects to use their data for these purposes. In addition, personal data should be deleted once its retention is not required anymore. Openness and transparency over how data is processed, manipulated and disclosed to third parties are also key requirements. Data subjects should be notified of changes affecting the management of their personal data and they should retain a degree of control over it. Compliance to all these aspects must be monitored and any violation promptly reported and addressed.

Privacy policies are commonly used to represent and describe these privacy laws and guidelines. They express *rights* of data subjects, *permissions* over usage of personal data and *obligations* to be fulfilled [1,12]. These policies must be understood and refined by enterprises, deployed in their data management processes and IT infrastructures and enforced. They need to be audited and monitored for compliance. Both *operational* and *compliance* aspects must be dealt with. Enterprises that span across different geographical and organisational boundaries might be subject to different privacy laws and related privacy policies: this adds further complexity to the privacy management problem.

Current enterprise practices to privacy management are mainly based on manual processes, good behaviours and common sense. Not only are human processes prone to failure but the scale of the problem highlights the desire for additional technology to be part of the solution. The trend towards complexity and dynamism in system configurations heightens this need for automation to ensure that privacy and security properties are maintained as changes occur, and in addition to check that privacy is delivered as expected.

Enterprises are already investing in identity management solutions to automate the management of digital identities and user profiles. Most of this information is sensitive and must be managed in a privacy-aware way. To be adopted, privacy management solutions must also leverage and be compatible with these identity management solutions.

This paper describes work done by HP Labs to address these issues, in particular how to automate the management and enforcement of privacy policies in enterprises and how to provide support for automatic compliance checking.

2 Addressed Problem

The key problem addressed in this paper is how to automate the management of *operational* and *compliance* aspects of privacy within enterprises. Currently much is done by means of manual processes, which make them difficult and expensive to comply. The introduction of automation still requires following best practice and good behaviour. However, it can help enterprises to reduce involved costs and make the overall process simpler and more effective.

Operational aspects of privacy include ensuring that use of personal data – collected by enterprises – takes into account the stated purposes for which it was collected, consent given by data subjects and other customisable constraints. They also include dealing with privacy obligations that dictate expectations and duties over how to handle data – such as deleting data, notifying users, transforming data, etc. Automating the management of operational aspects includes addressing how to model, deploy and enforce privacy policies and obligations and how to achieve this whilst leveraging existing identity management solutions (specifically, in the context of access control, user provisioning and account management).

Compliance aspects of privacy include ensuring that data is processed and handled consistently with laws, guidelines and data subjects' expectations. It must take into account the run-time behaviour of the enterprise and check for compliance at different levels of abstraction, including internal processes, applications/systems handling personal data, identity management components, systems and platforms running these components and storing personal data. Automating the management of compliance aspects includes addressing how to model all these aspects, how to gather relevant events and information, how to check for compliance and how to provide meaningful reports highlighting compliant aspects and violations.

A systematic and comprehensive approach to privacy is required to address all these aspects. An additional challenge is how to address them in an integrated way to provide full value to enterprises and ensure that current enterprise investments in the identity management area can be leveraged.

3 Our Solution

Our solution consists of three integrated R&D technologies to deal with the automation of privacy management both at the *operational* and *compliance levels*. Specifically it consists of:

- **a Privacy Policy Enforcement System and an Obligation Management System** that address *operational aspects* of privacy. These two systems help enterprises to model, deploy and enforce privacy policies and obligations with respect to managed personal data. As a significant example, we demonstrated the feasibility of integrating these systems with HP identity management solutions to handle privacy aspects;
- **a Policy Compliance Checking System** that addresses *compliance aspects* of privacy. This system helps enterprises to model privacy laws and guidelines, map them at the IT level, analyse related events and generate compliance and violation reports.

These technologies are self-contained and self-deployable: however we have also combined them in an integrated prototype solution to provide a comprehensive, flexible and systemic approach to privacy management. Specifically, the Policy Compliance Checking System, amongst other things, can supervise and report on the availability of the other two privacy management systems and check for inconsistencies within enforced policies, by comparing information coming from different sources. Further details follow.

3.1 Privacy Policy Enforcement System

Fundamentally, privacy policies define the purposes for which data can be accessed, how to take into account data subjects' consent and in addition the actions that need to be fulfilled at data access time, such as filtering out data, blocking access, logging, etc. Access control plays a key role in addressing these aspects [1,3,15].

Our approach to automate management and enforcement of these policies is based on a privacy-aware access control model [1,3,15] that extends traditional access control models (based on users/groups, users' credentials and rights, access control lists and related policies) by: (1) explicitly dealing with the stated purposes for which data is collected; (2) checking – at the access request time – the intent of requestors against these purposes; (3) dealing with data subjects' consent; (4) enforcing additional access conditions and constraints on personal data defined by data subjects and/or enterprise administrators. The main aspects of this model (shown in Figure 1) are:

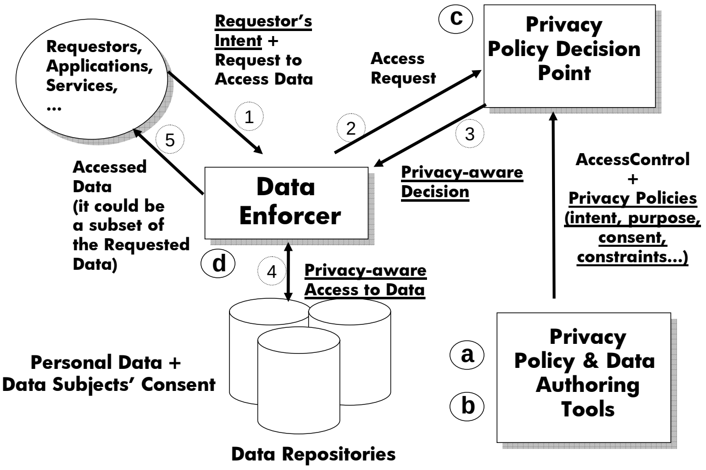

Requestors,
Applications,
Services,
...
Requestor's Intent + Request to Access Data
1
Access Request
2
c
Privacy Policy Decision Point
3
Privacy-aware Decision
5
Accessed Data (it could be a subset of the Requested Data)
Data Enforcer
d
4
Privacy-aware Access to Data
AccessControl + Privacy Policies (intent, purpose, consent, constraints...)
Personal Data + Data Subjects' Consent
Data Repositories
a
b
Privacy Policy & Data Authoring Tools

Fig. 1. Model of our Privacy-aware Access Control System

a) **A mechanism for the explicit modelling of personal data** subject to privacy policies: it provides a description of this data including the type of data repository (database, LDAP directory, etc.), its location, the data schema, type of attributes, etc.;
b) **An integrated mechanism for authoring privacy policies** along with traditional access control policies: this is a *Policy Authoring Point (PAP)* that allows 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 authorisation framework for deploying both *access control* and *privacy-based* policies and making related access decisions**: this 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. This is a Policy Enforcement Point (PEP) in charge (amongst other things) of dealing with the transformation of queries [1] (e.g. SQL queries) to access personal data and filtering part of the requested data, if their access is not authorised for privacy reasons.

The *"data enforcer"* plays a key role in the process of automating the enforcement of privacy policies over personal data. At "run-time", attempts to access personal data are intercepted and managed in the following way – see 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 the requestor's intent (this 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 types of data to be accessed) 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: *Access to data is denied*; *Access to data is fully granted*; *Conditional access to (part of the) data is allowed i.e. under the satisfaction of attached conditions*. Amongst other things, these conditions might require data filtering, data transformations and its manipulation.
4. The data enforcer enforces this decision. In particular, if the decision is a "*Conditional Access*" the data enforcer might have to manipulate the query (query pre-processing) and/or transform the requested personal data (result post-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.

To demonstrate the feasibility of this model, we have deployed it in a commercial identity management solution. We leveraged and extended HP Select Access [4], an HP state-of-the-art system to deal with fine-grained, policy driven, access control management. 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 in order to explicitly deal with privacy-aware access control on personal data, as shown in Figure 2. The following extensions of HP Select Access have been implemented in our prototype [1]:

- **The HP SA Policy Builder (i.e. its Policy Authoring Point) 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 also been extended via plug-ins to graphically author privacy policies** on "data resources" as described in our model;
- **The HP SA Validator (i.e. the Policy Decision Point) has been extended via plug-ins to make privacy-aware decisions**, as described in our model;
- **A data enforcer (i.e. the Policy Enforcement Point) 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 semantics of varying data repositories. The data enforcer currently implemented is a JDBC proxy for RDBMS databases. The implications for legacy applications and services are minimal as our data enforcer is seen as a traditional JDBC driver.

Additional details can be found in [1,3,15].

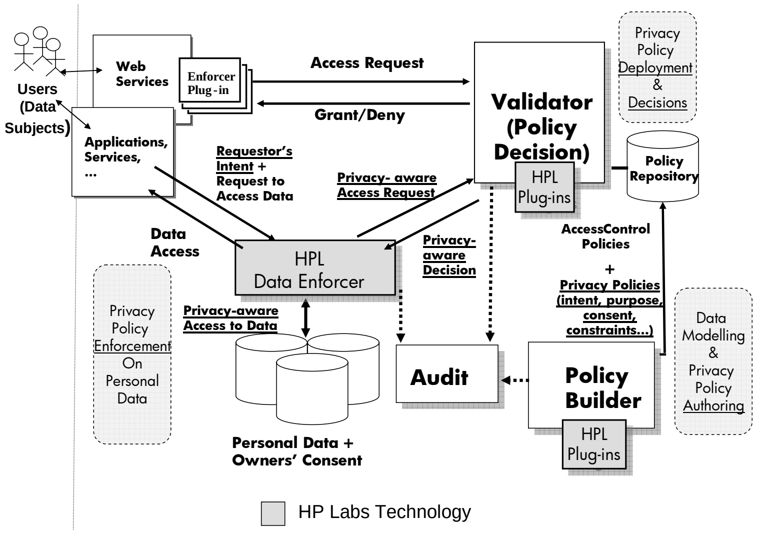

Fig. 2. Extended HP Select Access to deal with Privacy Policy Enforcement

3.2 Privacy Obligation Management System

This work addresses the problem of automating the management and enforcement of privacy obligations for personal data stored by enterprises. Privacy obligations [2,12] dictate expectations and duties on how to handle personal data and deal with its lifecycle management. Privacy obligations include: dealing with data deletion, data transformation (e.g. encryption), sending notifications, executing workflows, etc.

It is important to notice that their management and enforcement is orthogonal to the management and enforcement of privacy-aware access control policies [2]. 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 [2,12], where privacy obligations are "first class" entities, i.e. they are explicit entities that are modelled and managed in order to provide privacy-aware lifecycle management of personal data. In this model, a privacy obligation is an "object" that includes (at least) the following aspects [2,12]: *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 [2] 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* [2].

A related *obligation management framework* is introduced to manage these privacy obligations [2,12], based on the following principles:

- **Data subjects** 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**: the system schedules which obligations need to be fulfilled and under which circumstances (events);
- **Enforcing privacy obligations**: the system 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**: 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 [2,12]. 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 3 shows the high-level architecture of our obligation management system derived from our obligation management framework.

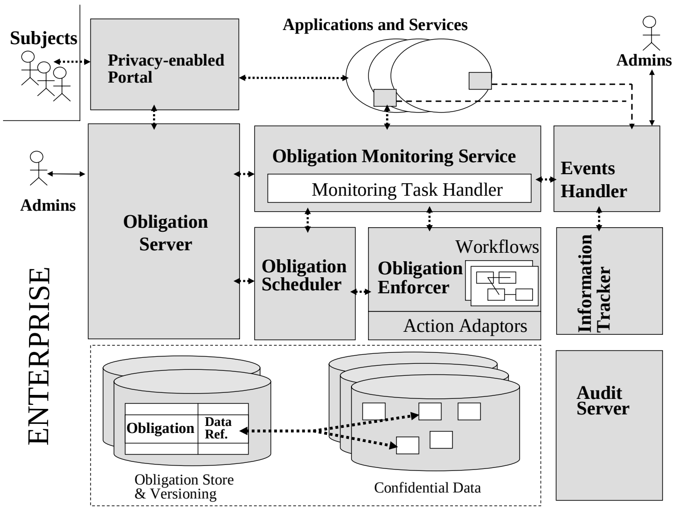

Fig. 3. High-level Architecture of our Obligation Management System

This system consists of the following modules:

- **Obligation Server**: this deals with the authoring, management and storage of obligations. It turns privacy preferences into privacy obligations via a template –based mechanism. It explicitly manages the association of privacy obligations to personal data and their tracking and versioning. Active obligations (i.e. obligations to be fulfilled) are pushed to the “Obligation Scheduler”;
- **Obligation Scheduler**: this 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 activates the corresponding “workflow processes” of the “Obligation Enforcer” that deals with the enforcement of the obligation;
- **Obligation Enforcer:** this 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**: this is 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**: this is orthogonal to the scheduling and enforcement components and monitors enforced obligations and the expected status of data;
- **Information tracker**: this intercepts events generated by data repositories, databases and file systems containing confidential data and provides this information to the event handler;
- **Audit Server**: this logs the various events and messages generated by all the components.

A working prototype has been fully implemented in the context of the EU PRIME project [11], as a proof of concept, providing the specified core functionalities: scheduling, enforcement and monitoring of privacy obligations.

We believe that an obligation management system should be considered as being 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. Specifically, our obligation management system can be integrated with the self-registration, customisation 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 explicit representation and management of privacy obligations, along with coordination of their overall enforcement and monitoring.

To demonstrate how this can be achieved in a practical way, we integrated our Obligation Management System with HP Select Identity [5], as shown in Figure 4.

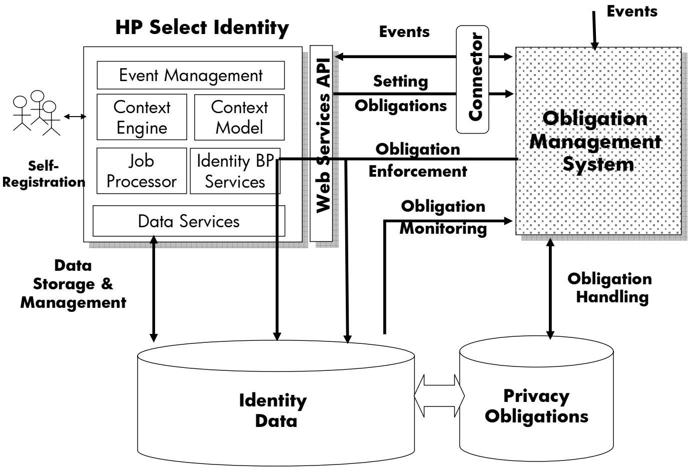

Fig. 4. High-level Architecture: integration of OMS with HP Select Identity

HP Select Identity [5] is a state-of-the-art solution to manage digital identities and personal data within and between large enterprises. The key features of the Select Identity system include: Centralized Management; User Provisioning; Administrative Delegation; User Self Service; Approval Workflow; Password & Profile Management; Audit and Reporting [5].

It automates the processes of provisioning, managing and terminating user accounts and access privileges by keeping all this information consistent and synchronised across provisioned platforms, applications and services – within and between corporate boundaries. Interactions with these third party systems (i.e. data repositories, legacy applications, services, etc.) are achieved via *Connectors*. These third parties can provide feedback to HP Select Identity (via an agent-based mechanism) about changes to their local copies of provisioned data. Changes are communicated to HP Select Identity via its Web Service API.

In our integrated prototype, HP Select Identity and our *obligation management system* interact via an ad-hoc *Connector*. As shown in Figure 4, we use HP Select Identity self-registration and user provisioning capabilities to specify and capture (at the time data is disclosed by users) privacy constraints and preferences about how personal data should be handled. These preferences are then processed by our *Connector* and sent to the obligation management system that will transform 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 privacy obligations.

Specifically, our system retains control of the supervision of obligations and their monitoring. HP Select Identity is leveraged by our system to enforce obligations constraints, such as deletion of identities, data transformation, etc. Currently, deletion of personal data is achieved by triggering HP Select Identity workflows, whilst the obligation management system handles notifications to users.

3.2 Policy Compliance Checking System

This work addresses the problem of automating the assessment of compliance of privacy policies within enterprises; a similar approach applies to best practice guidelines, legislation and risk analysis. Our system verifies whether the data processing system is strong enough to automatically execute the privacy policies reliably: this involves assessment of the deployment of privacy enhancing technologies and the underlying trust, security and IT infrastructure.

We aim to allow enterprises to check the trustworthiness of their system components, as well as those of their business partners to whom they may transfer personal data. For example, a service may be considered trustworthy if it has been accredited by an independent privacy inspector (such as BBBOnLine or TRUSTe), or a platform may be considered trustworthy if it is judged to be in a trusted state and is compliant with standards produced by the Trusted Computing Group.

In order to automate privacy compliance the system assesses the extent to which IT controls (including privacy-enhancing technologies, such as our privacy policy enforcement system and privacy obligation management system) satisfy key privacy principles or goals. To do this the system uses a model that cascades and refines top-level properties down to specific requirements that technologies can analyse, enforce and report on.

An example of technological control influence on a high level goal would be the following: a privacy related goal an enterprise could face is that data is only used for the purposes for which it was collected. This can be satisfied by the sub-goal that the enterprise uses a control that enforces role based access, where roles are associated with processes like marketing or customer support. In addition, the system should check that the control is configured correctly, the control is available, the control has not been subverted and there is proper separation of the duties defined for specific roles. There can be a many-many mapping between the goals and sub-goals: for example, it may be necessary to satisfy a combination of sub-goals in order to satisfy a higher level goal.

The architecture of our prototype system is shown in Figure 5. This system examines distributed system configurations using an agent infrastructure deployed across IT resources, feeds the findings into a reasoning engine and reports the resulting findings in a tree-like structure that can be 'drilled down' to the level of detail required. It uses functional decomposition to model privacy and model-based reasoning to carry out the analysis and generate reports. More specifically, modelling of privacy goals is combined with modelling of organisation resources and the processes around these resources. If desired, semantic web technology can be used to create a common understanding of lower level checks that are carried out.

The system is intended to be used in the following way: first of all, predefined policy sub-trees would be input into our editing tool (shown in Figure 6) by a privacy expert to form a generic privacy model (this only needs doing once, but can be updated subsequently). For each specific system on which the compliance checker is to be run, a privacy officer and/or specialised administrator would tune the constraints and deploy the system. Next, agents would be deployed to resources based on information given in the model, and would gather information over a selected time period.

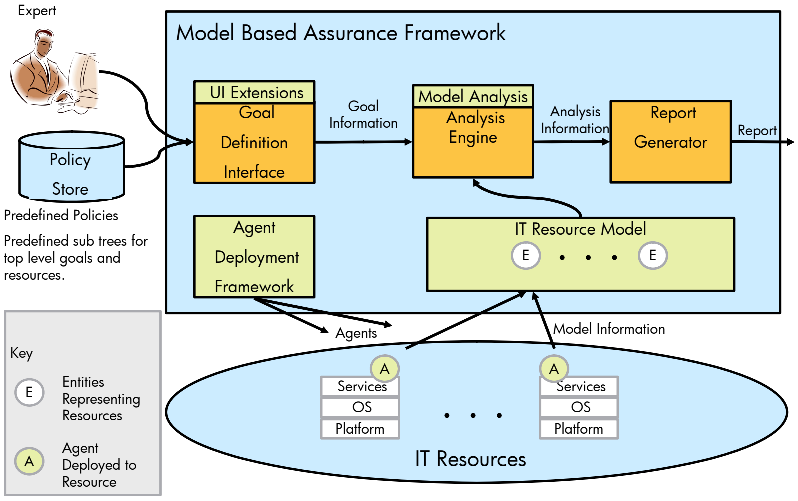

Fig. 5. System Policy Compliance Checking: Architecture

Whenever desired, analysis could be triggered and a corresponding report generated: an example is shown in Figure 7.

Figure 6 shows an example of the tool we developed to enable definition, input and customisation of models that refine and transform privacy policies from high level statements to something that can be executed automatically at a lower level. In this example, the OECD principles [14] for fair information usage were taken as the top layer within the model, there is an intermediate layer of information analysis nodes and a lower layer of technological input. In Figure 6, the model focuses on assessing the deployment of the privacy policy enforcement system described above (SAPE stands for "Select Access Privacy Enforcer"). We also developed other models, including analysis of a range of privacy and security-related IT controls and assurance information.

Figure 7 shows and example of the compliance report generated by our system using the model shown in Figure 6. This report is targeted at company executives, managers and auditors in order to provide information in a transparent way that can highlight areas that are a privacy concern in a dynamic and accountable way, and allow drilling down if desired to obtain further levels of detail.

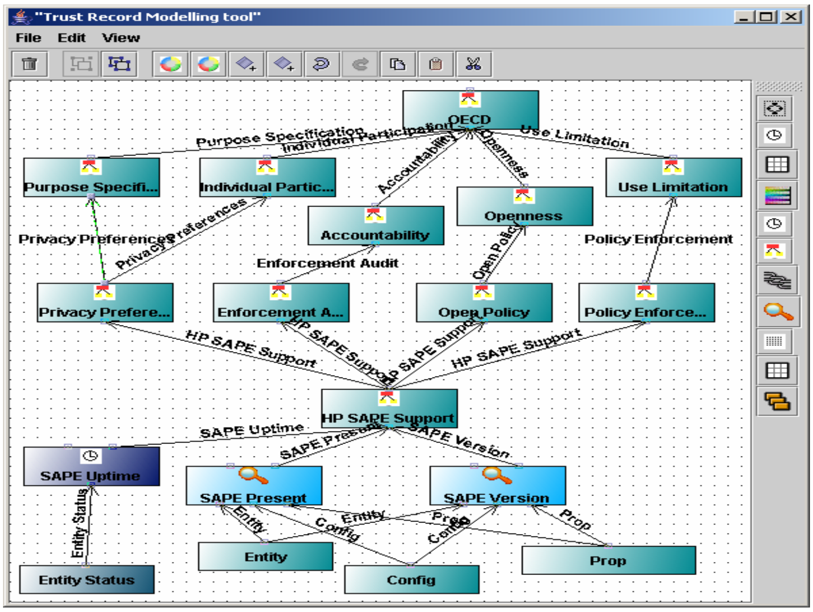

Fig. 6. Example sub-tree within privacy model

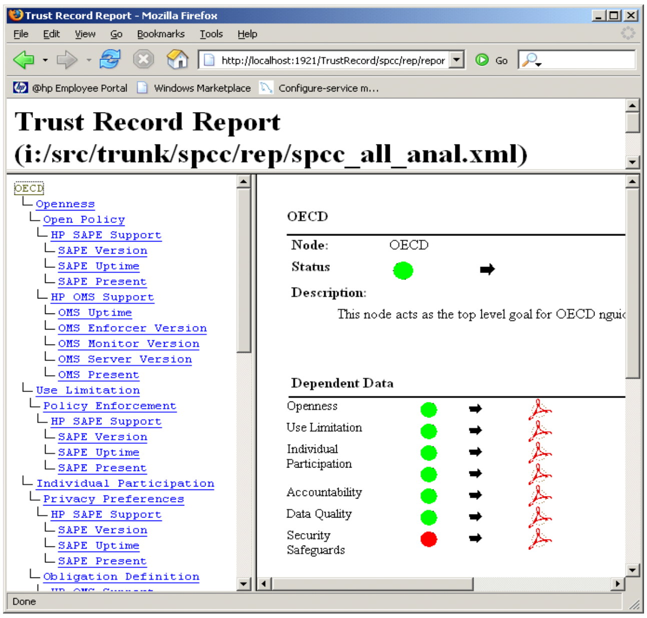

Trust Record Report (i:/src/trunk/spcc/rep/spcc_all_anal.xml)

```
OECD
└ Openness
  └ Open Policy
    └ HP SAPE Support
      └ SAPE Version
      └ SAPE Uptime
      └ SAPE Present
    └ HP OMS Support
      └ OMS Uptime
      └ OMS Enforcer Version
      └ OMS Monitor Version
      └ OMS Server Version
      └ OMS Present
└ Use Limitation
  └ Policy Enforcement
    └ HP SAPE Support
      └ SAPE Version
      └ SAPE Uptime
      └ SAPE Present
└ Individual Participation
  └ Privacy Preferences
    └ HP SAPE Support
      └ SAPE Version
      └ SAPE Uptime
      └ SAPE Present
  └ Obligation Definition
```

OECD

Node: OECD

Status 🟢 ➡

Description:

This node acts as the top level goal for OECD nguic

Dependent Data

Openness	🟢	➡	
Use Limitation	🟢	➡	
Individual Participation	🟢	➡	
	🟢	➡	
Accountability	🟢	➡	
Data Quality	🟢	➡	
Security Safeguards	🔴	➡	

Fig. 7. Example compliance report

4 Related Work

To the best of our knowledge we are not aware of any alternative integrated and comprehensive privacy management solution covering automation of both operational and compliance privacy aspects.

Key related work, in terms of privacy-aware access control, includes IBM Tivoli Privacy Manager [6] and IBM Hippocratic databases [7]. These are very vertical solutions requiring modification of the IT infrastructures in which they are deployed. Their approach might require duplication of effort at the authoring and enforcement time. We believe that our technology is simpler and more integrated.

In terms of privacy obligation management, no other significant work has been done to explicitly handle and enforce privacy obligations as we do. The EPAL [8] language and a related Enterprise Privacy Authorisation architecture do not define obligations in detail and subordinate their enforcement to access control: this is an inadequate approach because some obligations, such as the ones involving deletion of data, are independent of access control aspects.

In terms of policy compliance checking, we are not aware of products/solutions providing this type of model-driven assurance and compliance verifications. Current products and solutions, including Synomos [9] and SenSage [10], "hardcode" their compliance checking process and do not model privacy processes and IT components as we do.

Our approach is modular and addresses both operational and compliance aspects of privacy management relevant to an enterprise: our technologies can be integrated together and with third party solutions to provide a comprehensive way to automate privacy management.

The integration of our technologies with HP identity management solutions demonstrate the fact they can be deployed in real-world middleware solutions of enterprises. In particular, the privacy enforcement technology (integrated with HP select Access) is transparent to applications and services that make attempts to access data (for which privacy policies needs to be enforced) via standards drivers and protocols (e.g. JDBC, LDAP, etc.).

Our technologies have been designed for a general purpose usage and deployment: they can be leveraged, integrated and deployed in other contexts, beyond HP identity management solutions.

5 Current Status and Next Steps

The prototype of the Privacy Policy Enforcement System, integrated with the HP Select Access solution, has been transferred to the HP Software Business Organisation and is currently under productisation.

The Obligation Management System is currently available both as a stand alone prototype – as a key component of the PRIME system [11] – and integrated with HP Select Identity solution. We are currently exploring how to productise this technology.

The Policy Compliance Check System is currently available as a prototype that extends the Model-based Assurance Framework developed in HP Labs, and also as a subpart of the PRIME system [11].

We will further research and refine our work and related technologies. The privacy policy enforcement system can be further extended to include additional privacy constraints and more sophisticated mechanisms to process queries for additional types of data repositories (beyond RDBMS systems), such as LDAP repositories.

The obligation model underpinning the privacy obligation management system needs to be further extended to be scalable and cope with large amounts of personal data. A promising research topic is to explore the management of parametric obligations that apply to a large subset of personal data subject to similar privacy preferences.

The policy compliance checking system also needs to be extended in terms of modelling capabilities and to provide aspects such as data flow and a more complete assessment of the privacy enforcement technologies' ability to deliver compliance.

We will carry on our research and explore how to further exploit them in the context of HP businesses and also in the EU PRIME project [11].

6 Conclusions

Privacy management is an important issue for enterprises and is a core part of their IT Governance initiatives. It involves both operational aspects, related to the enforcement of privacy policies and obligations, and compliance aspects, related to checking for compliance of these policies to expected processes and their deployment within enterprise IT infrastructures. Key requirements for enterprises include introducing more automation, cost reduction and utilising current investments in the space of identity management.

We have described our innovative and systemic approach to address these issues, inclusive of mechanisms to automate the enforcement of privacy policies and checking for their compliance.

Three integrated R&D technologies – providing these functionalities – have been described: a privacy policy enforcement system; a privacy obligation management system; and a policy compliance checking system that, amongst other things, checks for the correct operational behaviour of the other two systems.

Working prototypes have been fully implemented to demonstrate the feasibility of our approach and integrated - as a proof-of-concept - with HP identity management solutions. In particular, the privacy policy enforcement system integrated with HP Select Access is currently under productisation by HP Software Business Organisation.

Additional work and research will be carried on both within HP Labs and in the context of the PRIME project.

7 References

1. Casassa Mont, M., Thyne, R., Bramhall, P.: Privacy Enforcement with HP Select Access for Regulatory Compliance, HP Labs Technical Report, HPL-2005-10, 2005
2. Casassa Mont, M.: Dealing with Privacy Obligations in Enterprises, HPL-2004-109, 2004
3. Casassa Mont, M., Thyne, R., Chan, K., Bramhall:, P. Extending HP Identity Management Solutions to Enforce Privacy Policies and Obligations for Regulatory Compliance by Enterprises - HPL-2005-110, 2005
4. Hewlett-Packard (HP): HP Openview Select Access: Overview and Features - http://www.openview.hp.com/products/select/, 2005
5. Hewlett-Packard (HP): HP OpenView Select Identity: Overview and Features, http://www.openview.hp.com/products/slctid/index.html, 2005
6. IBM Tivoli Privacy Manager: Privacy manager main web page - http://www-306.ibm.com/software/tivoli/products/privacy-mgr-e-bus/, 2005
7. Agrawal, R., Kiernan, J., Srikant, R., Xu, Y.: Hippocratic Databases, http://www.almaden.ibm.com/cs/people/srikant/papers/vldb02.pdf, IBM Almaden Research Center , 2002
8. IBM: The Enterprise Privacy Authorization Language (EPAL), EPAL 1.2 specification. http://www.zurich.ibm.com/security/enterprise-privacy/epal/, IBM, 2004
9. Synomos: Synomos Align 3.0, http://www.synomos.com/, 2005
10. SenSage: SenSage Web site, http://www.sensage.com/, 2005
11. PRIME Project: Privacy and Identity Management for Europe, European RTD Integrated Project under the FP6/IST Programme, http://www.prime-project.eu.org/, 2005
12. Casassa Mont, M.: Dealing with Privacy Obligations: Important Aspects and Technical Approaches, TrustBus 2004, 2004
13. Laurant, C.: Privacy International: Privacy and Human Rights 2004: an International Survey of Privacy Laws and Developments, Electronic Privacy Information Center (EPIC), Privacy International. http://www.privacyinternational.org/survey/phr2004/, 2004
14. OECD, “OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data.”, http://www1.oecd.org/publications/e-book/9302011E.PDF, 1980
15. Casassa Mont, M., Thyne, R., Bramhall, P.: Privacy Enforcement for IT Governance in Enterprises: Doing it for Real, TrustBus 2005, 2005