Dealing with Privacy Obligations: Important Aspects and Technical Approaches

Marco Casassa Mont
Trusted Systems Laboratory
HP Laboratories Bristol
HPL-2004-34
March 8, 2004*

E-mail: marco_casassa-mont@hp.com

obligations, privacy, policies, enforcement, monitoring, stickiness, accountability, identity management

The management and enforcement of privacy obligations is a challenging task: it involves legal, organizational, behavioral and technical aspects. In particular, the management of privacy obligations for identity and confidential data can require ongoing efforts, both in the short and very long term. It can be affected by events. Work has already been done for the management of obligations subordinated to authorization aspects (triggered by interactions and transactional events) and simple long-term obligations for data retention. Dealing with ongoing and long-term aspects of obligations is still a green field and open to research. This area is of particular relevance for enterprises, organizations and government agencies that deal with personal identity information. Privacy and data protection laws already dictate obligations involving ongoing and long-term constraints and duties. This paper explores and analyses the explicit management of privacy obligations for identity information by considering privacy obligations as first-class citizens. We focus on the technical aspects even if we recognize that the problem cannot be solved only by deploying technological solutions. Mechanisms are required to represent, manage, monitor and enforce obligation policies in complex and heterogeneous environments. Policy-driven scheduling mechanisms coupled with secure workflows and auditing techniques can be useful to address aspects of the problem. It is also important to be able to strongly couple these policies to confidential data, track their storage, distribution and deal with relevant events. Our research is work in progress: we illustrate some of our technical work and investigation in this space.

* Internal Accession Date Only
Approved for External Publication

1. Introduction

In the last decade a lot of work has been done in the area of privacy, in particular from a legal and legislative perspective. This includes European Community data protection privacy laws, various US privacy laws (HIPAA, COPPA, GLB, FRC, etc.) and more specific national privacy initiatives. An overview of these initiatives can be found at [1]. Various guidelines are also available on the protection of privacy and flows of personal data, including OECD guidelines [2] that describe concepts such as collection limitation, data quality and purpose specification principles.

Privacy policies are a suitable tool to represent and describe privacy laws, guidelines and privacy statements. Privacy policies, at the very base, express rights, permissions and obligations, usually in natural language that needs to be interpreted and understood by people.

Privacy policies are formulated and stated in a wide variety of contexts including the e-commerce, financial, health care and government sectors. For example, in e-commerce and web sites, privacy policies describe the rights of users about their personal information, the permissions given to service providers and service providers' obligations. These policies let consumers know about web sites' privacy practices: consumers can then decide whether or not these practices are acceptable, when to opt-in or opt-out and who to do business with. Examples of guidelines for formulating online privacy policies can be found at [3].

If on one hand the expression of privacy statements via policies is a significant advancement in communicating privacy rights, permissions and obligations, on the other hand, are quite often difficult to understand, hard to find, they take a long time to read and can change without notice.

Last, but not least, privacy policies might be hard to enforce via IT solutions. The enforcement of privacy rights, permissions and obligations related to 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.

In many cases the full enforcement of privacy policies cannot be achieved only via technological approaches but it still requires that the entities involved in the management of confidential and personal data follows best practices and good behaviours. However, being able to automate aspects of the enforcement of privacy policies and reduce the involved costs is of primary importance and interest for enterprises, web sites, e-commerce and financial organisations that more and more recognise that dealing correctly and honestly with privacy matters can have a beneficial return in terms of branding, trust and business.

Advancements in this direction have already been made when dealing with the (technological) enforcement of privacy permissions. Extended access control and authorization mechanisms have been built to check privacy permissions against users' rights, the purpose of the confidential information (that needs to be accessed) and the declared intents. This is the case, for example, of web transactions and interactions or applications/services within organizations that need to access and manipulate confidential data for business reasons. More details are provided in the related work section.

On the other hand, we argue that the management and enforcement of privacy obligations, as first class citizens, is still a green field and open to research. The events that trigger the fulfilment of privacy obligations can be completely orthogonal to the ones that are relevant for privacy permissions. Privacy obligations can have ongoing aspects that need to be monitored and satisfied. In this paper we analyse some of the related issues and describe possible technical approaches to move towards a more explicit management and enforcement of privacy obligations.

2. Privacy Obligations

It is hard to classify privacy obligations in a manner which is satisfactory for all environments. Different types of privacy obligations have been defined for financial institutions, health-care, enterprises and e-commerce: they have different interpretations, implications and enforcement requirements depending on the context and the legislative framework where they are applied.

The description of responsibilities and commitments dictated by privacy obligations can range from being very abstract to very specific.

Privacy obligations can be very abstract. An example is: “Every financial institution has an affirmative and continuing obligation to respect customer privacy and protect the security and confidentiality of customer information” - Gramm-Leach- Bliley Act (1999).

Other privacy obligations can dictate more refined responsibilities 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.

At the other extreme, privacy obligations can dictate very specific requirements. This is the case where data retention has to be enforced for a long period of time or data is temporarily stored by organisations: 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).

When dealing with privacy obligations, different aspects need to be kept in account:

- The timeframe (period of validity) that applies for obligations: it could be for a short or a long period of time;
- The situations/events that trigger the need to fulfil obligations: it could be triggered by a specific event or be ongoing, for example dictated by law. Events include deadlines, specific transactions/interactions and contextual changes;
- The enforceability of obligations: an obligation can be technically enforceable or its implementation can only happen as the result of guidelines, human behaviours and best practices;
- The target of an obligation and the implications: for example the target can be confidential data, personal profiles, medical or criminal data, etc. In case of long term privacy obligations due to data retention, data has to managed in a particular way, to ensure its survivability and longevity;
- The entities that are responsible for enforcing obligations and criteria specifying their accountability;
- Exception or special cases that applies for obligations.

In general high-level privacy obligations have “ongoing” commitments for organisations, dictated by privacy guidelines. These obligations describe what the acceptable behaviours and best practices are: they usually are abstract and need to be refined, ground and implemented in specific contexts to be enforced.

More refined privacy obligations can still impose commitments over a significant number of years: examples apply for health care, financial and criminal contexts, where data retention laws must be applied. In these cases, the deletion of data due to privacy obligations can happen many years after the date when data has been collected or only when specific (long-term) events happen. On the other hand, obligations might apply only for a short period of time and be transient: for example this is the case of privacy obligations that are valid just for the period of time a user has an account or a profile stored at an e-commerce site.

The topic related to “privacy obligations” is complex and exploring all the possible implications and involved aspects goes far beyond the purpose of this paper. In this paper we specifically focus on enforceable privacy obligations related to personal and confidential data for enterprises and business organisations.

3. Important Issues and Requirements

Important issues need to be considered when dealing with the management and enforcement of privacy obligations. These issues dictate related requirements:

- **Modelling and representation of privacy obligations**: aspects of privacy obligations need to be modelled, including representing which data is affected by the obligation, the events and conditions that trigger the fulfilment of an obligation, actions to be carried on, who is responsible and accountable for their enforcement.

- **Association of obligations to data**: the association of privacy obligations to the targeted confidential data must be strong i.e. not easy to be broken. This aspect is particularly challenging in dynamic environments where confidential data can be processed, moved around or sent to other parties. Breaking the association of data to their associated privacy obligations is, on its own, a violation of these obligations;
- **Mapping obligations into actions**: when possible, actions must be expressed in a way that can be programmatically enforced. Otherwise they should trigger related processes and workflows (involving the human intervention) and clearly state responsibilities;
- **Compliance of refined obligations to high-level policies**: refined privacy obligations are usually an interpretation and adaptation of high-level policies to specific contexts. High-level policies can change and, as a consequence, refined policies need to reflect this. The mapping of high level policies to refined privacy obligations (and the affected data) should be managed explicitly and tools built to spot potential inconsistencies, dependencies and which refined policies needs to be modified;
- **Tracking the evolutions of obligation policies**: as obligation policies can be carried on over long periods of time, they are subject to changes. An important issue is related to the tracking of these changes, for accountability reasons and to deal with the evolution of the contexts and frameworks where these obligations apply. This introduce requirements in terms of dealing with versioning of obligation policies and context tracking;
- **Dealing with long-term obligation aspects**: the fact that obligation policies might require long-term commitments has implications on the longevity and survivability of related processes and the involved data. Events and conditions related to obligations need to be monitored over long period of time. Solutions need to be built in a way that can be easily extended and modified over time. The format of stored data needs to evolve to take into account technological advancements. Openness and flexibility are two important requirements;
- **Accountability management**: as anticipated above, the explicit management of accountability is fundamental to ensure that the enforcement of privacy obligations is carried on with clear responsibilities of the involved parties. Responsibilities should be explicitly defined and communicated. This introduces requirements in terms of auditing, tracking of obligations and their monitoring;
- **Monitoring obligations**: it is important that the fulfilment of obligations is monitored and checked against expected situations and behaviours. Despite all the good intents and enforcement mechanisms, it can always happen that the fulfilment of obligations is omitted. Monitoring mechanisms must be orthogonal to the enforcement mechanisms. Monitoring tasks need to be aware of the set of "active" privacy obligations and access evidence about the enforcement of obligations, such as audit logs. In case of discovery of overdue obligations they should trigger their enforcement and create awareness about the encountered problems;
- **User involvement and awareness**: at the very base, privacy policies and obligations are defined and enforced to preserve user's rights on their personal data. It is important that rights are well understandable by users. Users should also have visibility of which obligations an organisation has with them and potentially monitor their fulfilment. This introduces requirements of transparency about organisational practices, along with the provision of tools that allow users to monitor and directly manage privacy obligations;
- **Complexity and cost of instrumenting applications and services**: last but not least, an important issue is related to the impact that the enforcement and monitoring of obligation policies has on the involved applications and services, both in terms of their instrumentation and costs. As long as possible, a privacy obligation framework should be deployed in a way that requires a minimum impact on applications and services.

Dealing with the management and enforcement of privacy obligations can be reasonably easy when the events that trigger them are well defined and easy to capture. For example, a web transaction between a user and a service provider might require the access or the disclosure of user's confidential data to third parties: obligations might dictate the need for notifying users or requesting their

authorization. Applications and services can be instrumented to intercept these events and trigger the enforcement of relevant obligations.

More complex is, for example, the case of privacy obligations related to ongoing obligations, triggered by the occurrence of specific events and conditions, non-necessarily related to any transaction or interaction. Some of these events might not be so easy to intercept or the software cannot be easily instrumented to deal with them. In addition solutions need to be deployed and kept running for long periods of time to fulfil these obligations.

4. Addressed problems

In this paper we address the problem of dealing with an explicit management of privacy obligations, on an ongoing basis, including long-term privacy obligations. This implies dealing with the explicit monitoring, enforcement, and tracking of privacy obligations.

Related to this we also want to address the problem of dealing with the strong association of privacy obligations to data, enforce accountability and provide more transparency to users.

Work has already been done to deal with some of these issues, in particular related to the representation of privacy policies (and obligations), their enforcement in transactional and interaction-driven contexts and the management of simple long-term aspects of obligations in particular for data retention. In many cases, though, obligation policies are considered as second-class entities the enforcement of which is subordinated to other aspects of privacy policies, such as permissions.

What we believe is missing is a more explicit and comprehensive approach to privacy obligations, where they are considered as first-class citizens and can be managed without their subordination to other aspects such as dealing with privacy permissions and access control/authorization.

5. Technical Details

This section provides technical details about the approaches and solutions under exploration to address the problems stated in section 4.

Figure 1 shows a high-level architecture of a system providing an explicit management of privacy obligations:

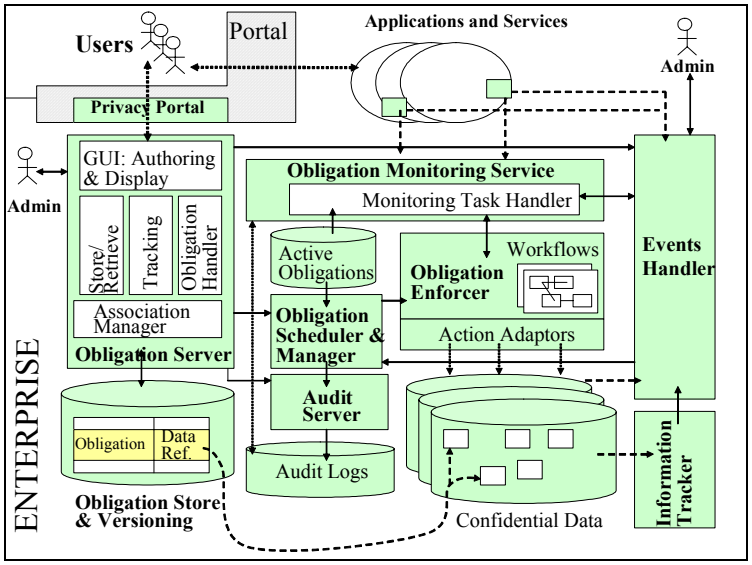

Figure 1: High-Level Architecture

The obligation management system consists of:

- **Obligation Server**: it is the component that deals with the authoring, management and storage of obligations. It allows the management of the association of privacy obligations to confidential data and their tracking and versioning. Administrators and users can access, review and manage privacy obligations of their competence. It pushes active obligations (i.e. valid obligations) to the "obligation scheduler & manager" and relevant events (to monitor) to

the event handler. One or more obligation servers can be deployed (and synchronised), depending on needs;

- **Obligation Store and Versioning**: it is the data repository storing obligations and their mapping to confidential data. More details follow. Multiple versions of obligations are also stored in this system;
- **Obligation Scheduler and Manager**: it is the component that is aware of which obligations are currently active (i.e. stored in the "active obligation" repository), ongoing obligation deadlines, relevant events and their association to obligations. For simplicity this component does not enforce privacy 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**: at its core 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. It is configurable via "action adaptor" plug-ins, specialised in performing specific actions (deletion of data, transformation/obfuscation of data, e-mail notification, etc.);
- **Events Handler**: it is the component in charge of monitoring and detecting relevant events for privacy obligations. These events are defined and pushed by the obligation server. The detection of events can happen via instrumented application/services. They can also be directly generated by users, administrators, the "obligation monitoring service" and the information tracker;
- **Obligation Monitoring Service**: it is the component, orthogonal to the scheduling and enforcement systems that monitors active obligations and if they have been enforced (by analysing and checking for effects of their actions);
- **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. It is aware of the location of confidential data (as described by the obligation policies) and checks for movements and changes happening to this data;
- **Audit Server**: it audits the relevant events and information generated by the overall system components and involved applications/services.

In our model, privacy obligations contain the description of relevant events/conditions, actions, target (i.e. related confidential data) and accountable/responsible entities. A simple XML-based example of privacy obligation is shown in figure 2. The content of this privacy obligation is self-explicative. It is about the deletion of confidential data at a specific point of time. The policy contains a reference to the actions to be enforced (in the example they are two workflow processes for deleting data and notifying relevant entities) and the entities responsible for this obligation.

The privacy obligation policy also contains information about the "targeted data" and specifies the association to this data along with the owner(s). In the example the obligation refers to data stored both in a relational database (accessible via an SQL query) and in a file system. Other mapping mechanisms can be used.

If the system is deployed in a stable and well-controlled environment, managing the association of data to obligations can be handled via a mixture of automation mechanisms and manual intervention (of administrators and users).

Issues arise when the overall environment is dynamic and data can be moved around. In this case, despite all the efforts of handling events and tracking movements, the association of data to obligations policies can be broken or be left in an inconsistent state.

To address this issue we are exploring a variant of the architecture shown in figure 1, where stronger mechanisms are introduces to manage the association of obligations to data. Figure 3 shows these additional components.

Confidential data is obfuscated and strongly associated to privacy obligations by using cryptographic techniques. A key management system is introduced to deal with this task. For example a symmetric key is generated by the key management system and used to obfuscate data. An envelope (e.g. based on PKCS#7) is created: it contains (at least) the hash of the obligation policy along with

the symmetric key. This envelope is encrypted with the public key [3] associated to the key management system.

The triple consisting of ***<obligation policy, encrypted envelope, obfuscated data>*** is stored as a replacement of the original data. The obligation policy must contain a reference to the competent Obligation Server but it could omit the reference to confidential data, as the policy is now directly associated to this data.

In this way, the encrypted confidential data can be moved around and transmitted to other parties without any strict control. The receiving party has to interact with the Obligation Server to decrypt the data: this allows the system to track and audit where the data is, check for relevant obligations and update its obligation store. The basic principles and additional details on how such this approach can be implemented are described in [9].

```
<Obligation Policy>
  <id: 21435 version: 1.0>
  <Obligation trigger descriptor>
     <type: event>
     <subtype: time-based event>
     <parameters>
        <relevant time: DD/MM/YYYY>
     </parameters>
  </Obligation trigger descriptor>
  <Target>
    <Data Owner> refXYZ  </Data Owner>
    <Data Locators>
      <data locator> "select … from …where …" </data locator>
      <data locator> X:\\filesyst1\sub1\sub2\*.doc </data locator>
    </Data Locators>
  </Target>
  <Enforcement>
  <Actions>
    <action> Workflow.delete_data_process  </action>
    <action> Workflow.notify_process </action>
  </Actions>
  </Enforcement>
  <Responsibilities>
      <Responsible> Roles.Obligation_Manager_DataCenter1 <Responsible>
  </Responsibilities>
</Obligation Policy>
```

Figure 2: Simple Example of Obligation Policy

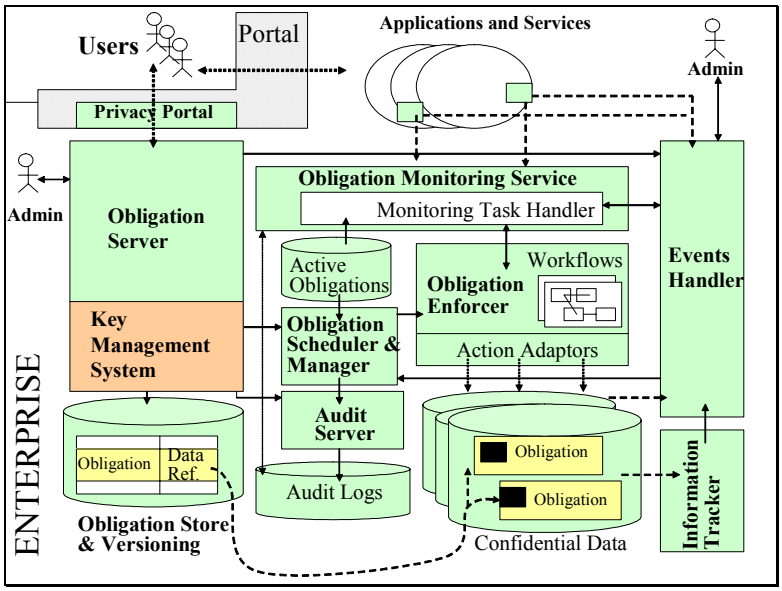

Figure 3: Extended High-Level Architecture

The technical approach described in figure 1 is almost transparent to applications and services that are affected by privacy obligations. The system needs (in most of the cases) only to be aware of relevant events. The second approach, in figure 3, on one hand introduces more control and enforcement of

accountability: on the other hand, applications and services might need to be modified in order to handle the encrypted data and associated process. Data repositories might need to change the way they store information, to accommodate encrypted data (for example databases schemas need to be changed, etc.).

We are currently exploring how a hybrid solution can be used to accommodate different needs and requirements and the overall implications on the underlying environment.

6. Discussion

The system described in this paper centralises the storage of privacy obligations along with their management. It can support the management of versions of privacy obligations over time and enable the tracking of their changes (and related applicability contexts) for auditing and accountability reasons. We are exploring how these aspects can be distributed to avoid potential bottlenecks and central points of failure, without compromising the overall security of the system.

The obligation scheduler coupled with the event handler allows for the management of short and long-term obligations. The monitoring system provides an additional mechanism for spotting enforcement omissions thanks to the fact it can understand the effects of actions dictated by privacy obligations (such as deletion and manipulation of data, notifications, etc.). Information logged by the audit server is used during these monitoring tasks.

Our system explicitly focuses on the management and enforcement of obligations: this does not imply that it has to happen independently by other privacy aspects, such as permissions.

It should be considered as a sub-system of a more comprehensive privacy management framework. Similarly, the representation of obligations is part of the wider task of representing privacy policies.

Even if the system enables automation, when dealing with privacy obligations, it also allows the human intervention in a variety of context. Administrators and users can intervene during the enforcement of obligations, if required (for example to explicitly authorise actions).

Administrators and users can access and manage the privacy obligations of their competence in a monitored and audited way: this increases the transparency of the enterprise's privacy practices and the involvement of the interested parties.

We assume that the enterprise is willing to be compliant with privacy policies and, more specifically, privacy obligations. However the system must be deployed by keeping in mind good security practices, especially for the platforms that will host our system components. As all the described components are critical, they require to be secured accordingly. Additional assurance and accountability can be added by hardening the audit server and involving trusted third parties in the monitoring of the enforcement of obligation policies.

When dealing with long-term obligations it is also important to ensure the reliability, survivability and longevity of the platforms running our system components and the involved data (including the representation of privacy obligation). Work has already been done in this space, including [15,16,17,18,19], and can be leveraged.

7. Related Work

Relevant work in the space of privacy management for enterprises is described in [4,5,6,7]. Enterprise Privacy Architecture is introduced and described in [7], encompassing a policy management system, a privacy enforcement system and an audit console. Paper [6] introduces more architectural details along with an interpretation of the concept of privacy obligations. This concept is framed in the context of privacy rules defined for authorization purposes. This approach is further refined and described in the Enterprise Privacy Authorization Language (EPAL) specification [8].

The above work makes important advancements in exploring and addressing the problem of privacy management in enterprises. Our main comments are on the suggested approach to handle privacy obligations i.e. stress the authorization and access control perspective as the key driver for their representation, management and enforcement.

The above approach and architecture are definitely pragmatic and can be leveraged by current access control mechanisms available within enterprises. However it has still to be fully demonstrated that privacy obligations can be managed at their best from an authorization-based perspective. Privacy

obligations can include aspects that are not really driven by authorization aspects, especially when the set of events that triggers privacy obligations is extended, to include, for example, dealing with the deletion of confidential data at a specific date/event, periodically providing notifications to users about stored confidential data, dealing with ongoing requests dictated by users or laws. We believe that modularity and separation of concerns are important aspects. In particular, the representation, management and enforcement of privacy rights, obligations and permissions should be addressed without imposing any specific or dominant perspective.

In our approach obligation policies are first-class citizens with their explicit management. However the proposed system can be considered as a subsystem of a more comprehensive policy management framework. Even if our architecture has high-level commonalities with the architecture described in [4,5,6,7] we further refine the concept of obligations, we introduce the concept of obligation versioning and tracking. We further split the enforcement mechanisms in two parts by including a scheduling mechanisms and an obligation enforcer where the obligations actions are carried out by flexible workflow processes that allows automation but also people involvement.

Approaches to deal with (privacy) obligations have already been implemented in products, in particular for data retention [10] and in a variety of document management systems. Nevertheless, these approaches are very specific, focused on particular domains and handle simple obligation policies. Our work wants to push 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 [8,11,12]. Work describing the monitoring of obligations in policy management is described in [12]. Relevant work on mechanisms to associate policies to data is described in [4,5,6,7,9,14]. Each mechanism has pros and cons in terms of the implications for existing enterprise applications, services and data repositories. We can leverage aspects of this work, in particular [9] to provide a stronger association of obligation policies to confidential data.

8. Current and Future Work

We are in the process of developing a prototype of the system components described in this paper. Components, when possible, will be implemented as web services and deployed within an enterprise scenario: different types of confidential data and repositories will be considered. Obligation policies will be represented by using an XML format to allow future extensions.

Our work and research is definitely in progress: technical aspects needs to be further refined and investigated especially the ones related to the life-cycle management of privacy obligations and events. The overall implications for the involved enterprise applications and services have to be fully understood. One of the reasons of developing our prototype is to make advancements in these areas by experimenting and refining our concepts.

Tools and mechanisms to address the compliance of refined obligations to high-level policies are also under investigation.

9. Conclusion

The management of privacy obligation is important for enterprises and organisations to preserve their reputation and brand, be compliant with legislation and customers' requirements and increase business opportunities.

In this paper we describe important issues that need to be kept into account when dealing with privacy obligations. In our vision obligation policies (as well as for other privacy aspects, including rights and permissions) need to be considered as first-class citizens in privacy management frameworks.

We introduce a technical approach to deal with the explicit management of privacy obligations, on an ongoing basis, including long-term privacy obligations. We provide a high-level description of system components dealing with the monitoring, enforcement, and tracking of privacy obligations. Related to this we also address the problem of dealing with the strong association of privacy obligations to data, accountability management and users involvement.

Our research and work is in progress. A prototype will soon be developed to test and refine our ideas.

10. 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] OECD - OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data. http://www1.oecd.org/publications/e-book/9302011E.PDF, 2001

[3] Online Privacy Alliance - Guidelines for Online Privacy Policies. http://www.privacyalliance.org/, Online Privacy Alliance, 2004

[4] G. Karjoth, M. Hunter - A Privacy Policy Model for Enterprises, IBM Research, Zurich - 15th IEEE Computer Foundations Workshop, June 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. Shunter, 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] M. Casassa Mont, S. Pearson, P. Bramhall - Towards Accountable Management of Privacy and Identity Information, ESORICS 2003, 2003

[10] IBM - IBM Tivoli Storage Manager for Data Retention, 2004

[11] C. Bettini, S. Jajodia, X. Sean Wang, D. Wijesekera - Obligation Monitoring in Policy Management, 2002

[12] N. Damianou, N. Dulay, E. Lupu, M. Sloman - The Ponder Policy Specication Language , 2001

[13] R Housley, W. Ford, W. Polk, D. Solo - RFC2459: Internet X.509 Public Key Infrastructure Certificate and CRL profile, IETF, 1999

[14] R. Agrawal, J. Kiernan, R. Srikant, Y. Xu - Hippocratic Databases, IBM Almaden Research Center, 2002

[15] R. J. Anderson - The Eternity Service, Proc. PRAGO-CRYPT 96, CTU Publishing House, Prague, 1996

[16] R.J. Ellison, D.A. Fisher, R.C. Linger, H.F. Lipson, T.A. Longstaff, N.R. Mead - Survivability: Protecting your Critical Systems, Proceeding of the International Conference of Requirements Engineering, 1998

[17] J. Kubiatowicz, D. Bibdel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, D. Rhea, H. Weatherspoon, W. Weimer, C. Wells, B. Zao - OceanStore: An Architecture for Global Scale Persistent Storage, University of California, Berkeley, ASPLOS 2000, 2000

[18] P.G. Neumann - Practical Architectures for Survivable Systems and Networks, SRI International, Army research Lab, 1999

[19] J.J. Wylie, M. W. Bigrigg, J. D. Strunk, G. R. Ganger, H. Kiliccote, P.K. Khosia - Survivable Information Storage Systems, IEEE Computer, 2000