



Handling Privacy Obligations and Constraints to Underpin Trust and Assurance

Marco Casassa Mont, Stephen Crane, Siani Pearson
Trusted Systems Laboratory
HP Laboratories Bristol
HPL-2005-54
March 18, 2005*

privacy, trust,
obligations,
assurance, trusted
platforms, prime,
policies

Trust is important to enable interactions on the web, in particular with enterprises. The trust that people have in enterprises can be built, reinforced or modified via a variety of means and tools, including personal experience, analysis of prior history, recommendations, certification and auditing by known authorities. The behaviour of an enterprise and the fact that it performs as predicted and agreed is important to shape its reputation and perception of trustworthiness. In particular, the way enterprises handle privacy has an impact on these aspects. We focus on enterprises that recognise the importance of dealing properly with privacy to increase their reputation and business opportunities. Important problems need to be addressed: how can enterprises provide people with degrees of assurance that they will operate in the way dictated by policies and privacy obligations, according to people's expectations? How can enterprises explicitly manage these policies? How can people check upfront that an enterprise has the right capabilities to handle and process their personal data? How can people have a constant, personalized feedback on the fulfillment of all these aspects? We describe requirements, a model to address the problem and provide technical details. Our work is in progress: initial prototypes have been developed and further work will be done in the context of the EU PRIME project.

Handling Privacy Obligations and Constraints to Underpin Trust and Assurance

Marco Casassa Mont, Stephen Crane and Siani Pearson

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

Abstract. Trust is important to enable interactions on the web, in particular with enterprises. The trust that people have in enterprises can be built, reinforced or modified via a variety of means and tools, including personal experience, analysis of prior history, recommendations, certification and auditing by known authorities. The behaviour of an enterprise and the fact that it performs as predicted and agreed is important to shape its reputation and perception of trustworthiness. In particular, the way enterprises handle privacy has an impact on these aspects. We focus on enterprises that recognise the importance of dealing properly with privacy to increase their reputation and business opportunities. Important problems need to be addressed: how can enterprises provide people with degrees of assurance that they will operate in the way dictated by policies and privacy obligations, according to people's expectations? How can enterprises explicitly manage these policies? How can people check upfront that an enterprise has the right capabilities to handle and process their personal data? How can people have a constant, personalized feedback on the fulfillment of all these aspects? We describe requirements, a model to address the problem and provide technical details. Our work is in progress: initial prototypes have been developed and further work will be done in the context of the EU PRIME¹²³ project.

1 Introduction

Trust is important to enable interactions on the Internet. People quite often have to trust e-commerce sites, service providers, online services and enterprises that they will perform as expected, they will provide the agreed services and goods and that

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

² The PRIME project receives research funding from the Community's Sixth Framework Programme and the Swiss Federal Office for Education and Science. This work was supported by the IST (Information Society Technologies) PRIME project; however, it represents the view of the authors only. The information in this document is provided "as is", and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability

³ Note that the architectures and concepts described in this paper will not necessarily be incorporated into the PRIME architecture, whole or in part.

they will not exploit and misuse personal and confidential information. In this paper we will use the term “enterprise” to generically refer to organizations, service providers, etc.

The trust that people have in enterprises can be built, reinforced or modified via a variety of means and tools, including personal experience, analysis of prior history, recommendations, certification and auditing by known authorities. The behaviour of an enterprise, the fact that it will fulfill agreed tasks in due time and perform as predicted are all important aspects to shape its reputation and perception of trustworthiness. Related to this, the way an enterprise handles privacy aspects has also an important impact on trust. In this paper we explore and focus on how enterprises can handle privacy policies and users’ constraints and expectations in order to underpin trust and assurance.

Privacy has an impact on the way enterprises manage personal and confidential data. Digital identities are one of the most important and valuable assets in the digital society. People disclose (aspects of) their personal data and confidential information to enterprises to engage in business interactions, transactions and obtain the provision of services.

People have expectations about how these data must be managed. Privacy is one of these. Any use of these personal data beyond the agreed purpose and the given consent can provoke major damages to people and the society: this includes spamming, identity frauds, identity theft and violation of fundamental data protection laws and privacy legislation.

In the last decade a great deal of work has been done in the area of privacy, in particular from a legal and legislative perspective. Privacy policies [3] are a suitable tool to represent and describe privacy laws, guidelines and privacy statements, as they allow expression of rights, permissions and 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 with whom to engage in business. If on the one hand the expression of privacy statements via policies is a significant advance in communicating privacy rights, permissions and obligations, on the other hand such statements are quite often difficult to understand: they take a long time to read and can change without notice.

Users might want to express customized requirements and obligations about how their data should be handled and used. Specifically, privacy obligations are a viable tool to dictate to the enterprise ways to process data and interact with users: they might require various things, including the deletion of personal data after a predefined period of time, periodic notifications to users about the status of their data, ways data should be accessed or disclosed to other parties, etc. Users might want to actively check for the compliance of enterprises to these privacy obligations once their data has been disclosed.

In some cases users might also want to get some assurance of the capabilities of an enterprise, even before engaging in any interaction or transaction with this enterprise. This includes obtaining degrees of assurance that the enterprise can actually support specific privacy policies and obligations, that their data will be processed and managed securely, that enterprises’ web services, applications and data repositories are

installed, run and patched according to security standards and good IT practices, that secure and trusted platforms are used.

Checking upfront if an enterprise can satisfy some constraints and policies, imposing privacy obligations on personal data and verifying their fulfillment over time are all important aspects that determine the perception of trust that a user has on an enterprise.

In the last few years, enterprises have started to recognize the importance of these needs and the fact that dealing properly with privacy aspects and addressing people's expectations is a win-win solution: enterprises benefit for this in terms of brand, reputation and business; people can increase their trust in the enterprise's ability to perform in a predictable and law conformant way.

2 Addressed Problem

Privacy and privacy management undoubtedly have an impact on the trustworthiness of enterprises. In this paper we want to explore this area. Important aspects and issues need to be addressed: how can people check upfront that an enterprise has the right capabilities to handle and process their personal data? How can enterprises provide people with degrees of assurance that they will operate in the way dictated by policies and privacy obligations, accordingly to people's expectations? How can enterprises explicitly manage these privacy obligations? How can people have a constant, personalized feedback on the fulfillment of all these aspects?

This paper aims at addressing the above problems. We mainly focus on the enterprise side of the problem: the goal is to build tools and solutions that help enterprises to enforce privacy obligations and allow users to check for enterprises' compliance to these policies and any additional constraints. In particular we are interested in providing solutions to "good-willing" enterprises that are aware of the importance of privacy as a driving factor to underpin trust, reputation and a business enabler. Of course users must have mechanisms to express their requirements and intuitive and usable tools to check for their fulfillment. The solutions deployed within enterprises must support these user-side functionalities. We briefly describe our approach to these aspects and provide references to our work done in this area.

We recognise that the problem cannot be solved only by deploying technologies: behaviors and implementation of correct process are very relevant. However our objective is to build technical solutions that can help enterprises to increase automation and help people to have additional support to make informed decisions about trust.

3 Related Work

Relevant work in the area of privacy and privacy management has been done on the legislative front. 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 regarding 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.

The fact that enterprises run their business and operations in countries and areas subject to these laws is important in providing degrees of assurance to people. In addition, approaches based on seal programs [25], i.e. certification of compliance by third-party authorities can provide further assurance.

Laws and seal programs imply the implementation of processes and behaviours by enterprises that must be continuously monitored: in many cases only general-purpose behaviour and procedures can be checked and audited. The fine-grained effect of their implementation is what eventually has an impact on people and their perception of the trustworthiness of an organization. In addition, these approaches do not take into account specific, fine-grained requirements, needs and constraints dictated by individuals.

The usage of recommendation mechanisms [26,27] - based on people sharing evaluations of enterprises' behaviours - is another well-explored approach for dealing with trust matters. These mechanisms can also be used to evaluate enterprises' compliance to privacy and, as a side effect, have an impact on the perception of the trustworthiness of an organization.

This approach is complementary to the problems we want to address. We want to enable enterprises to proactively control, manage and enforce their privacy policies and provide more assurance to people by involving users in this process.

Most of the traditional recommendation and feedback mechanisms involve people relying on the experiences of a trusted person. In our vision, in addition to this, people should be able to build up their own trust perception of the behaviour of an enterprise by having tools to remember the expectations, policies and privacy obligations that they imposed on enterprises and being able to check for their fulfilment over time.

Relevant work has been done by W3C with their P3P specifications [28] (and related framework) to allow people to describe in more details their privacy expectations and match them against the level of privacy supported by an enterprise. This is important to shape (aspects of) the trust that people might have on the enterprise. However P3P is mainly a "front-end" mechanism, in the context of web services. In its current form it is "passive" i.e. it only checks if people's expectations are matched against promises made by the enterprise. It does not address the problem of allowing users to express fine grained privacy obligations; it does not provide mechanisms to check and prove upfront compliance with fine-grained constraints; it does not provide active feedback on the ongoing execution and fulfillment of privacy obligations and constraints by enterprises. Last but not least, it does not provide enterprises with a full framework for dealing with privacy policies.

The problem of allowing people to check upfront for the compliance of enterprises to privacy policies and other constraints is very hard. In particular, it is hard for enterprises to demonstrate to users, in a fine grained way, that their IT infrastructure is secured and trustworthy and that their data will be processed according to the higher level security standards. This might include certifying that the specific services, ap-

plications, platforms and data repositories that will handle users' personal data are up to the users' expectations. Again, seal programs and certifications do not address these issues at this level of granularity.

We aim at making progress in this space by leveraging work done in trusted platforms and work under specification in the Trusted Computing Group (TCG) [29]. Leveraging trusted platforms and showing proof of their properties and usage in an aggregated way (to involve the relevant systems used by an enterprise) can provide degrees of assurance to users. In this paper we provide more details, and describe preliminary work done in this area.

Relevant work on the management and enforcement within enterprises of privacy policies and obligations dictated by users and laws is described in [4,5,6,7]. An 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 advances 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. to consider the authorization and access control perspective as the key driver for the representation, management and enforcement of obligations. Privacy obligations include aspects that are not really driven by authorization, especially when the set of events that triggers these 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.

Privacy obligations are an explicit tool that can be used by users to describe their privacy requirements. The fulfilment of privacy obligations is fundamental to providing assurance to people and increasing their level of trust. This includes ensuring that privacy obligations are scheduled and enforced in due time, that they are strongly associated to personal data, that any violation is reported and processed and that the expected feedback about their enforcement is given to users.

Part of the work described in this paper focuses on these aspects. In our approach obligation policies are first-class citizens with their explicit management. Compared to related work, we refine the concept of privacy obligations. 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, including providing support to enterprises and user from a trust and assurance perspective.

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 and degrees of assurance to people.

4 Privacy Obligations and Assurance Policies

This section provides more details about privacy obligations and other types of constraints that can be used by people to describe and convey their expectations to enterprises in terms of privacy and assurance. The way they are fulfilled by enterprises affects their reputation and trustworthiness. In particular we make the following distinction between privacy obligations and assurance policies:

- **Privacy obligations:** these are a set of conditions, requirements and expectations that need to be fulfilled by enterprises and organizations. They have operational implications on enterprises i.e. on the way enterprises store, handle, access and disclose personal data. They are usually formulated by users when disclosing personal data: they are associated to these data;
- **Assurance policies:** these are a set of conditions and constraints formulated by people to obtain degrees of assurance from enterprises that their data will be processed according to their expectations, such as compliance to privacy, security and IT standards. These policies are usually formulated by people before engaging in interactions with enterprises. The proofs that enterprises can give about their capability to support these policies is important to reassure people and has an impact on trust.

4.1 Privacy Obligations

It is hard to classify privacy obligations in a manner which is satisfactory for all environments. 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. Abstract privacy obligations can usually be found in laws. More refined privacy obligations can dictate constraints 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. 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) or when their account is closed.

Privacy obligations can have "ongoing" and long-term commitments for organisations or might apply only for a short period of time and be transient.

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

- **The timeframe (period of validity) that applies for obligations;**
- **The situations/events that trigger the need to fulfil obligations;**
- **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:** the target could be confidential data, personal profiles, etc.;
- **The entities that are responsible for enforcing obligations** and criteria specifying their accountability;
- **Exception or special cases that applies for obligations.**

The “privacy obligations” topic 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, systems to enforce and monitor them and provide feedback to users. The fulfillment of obligations has an important impact on the trustworthiness of enterprises.

4.2 Assurance Policies

Assurance policies are constraints and conditions usually expressed by people upfront to their engagement with enterprises. They can require enterprises to provide degrees of proof about their ability to:

- Support the enforcement of predefined privacy policies and obligations with respect to laws and legislation;
- Run their processes, services and data repositories in a secure way;
- Use secure and trusted systems [29], such as trusted computing platforms, to increase the level of security and trust in their operational activities.

Related to the last point, people might want enterprise to provide them with proof that they use trusted computing systems to run critical processes and data storage: this could happen via the issuance of signed statements by enterprises or by allowing users to directly control some of the characteristics of enterprises’ platforms. Again, TCG/trusted computing platform mechanisms could be leveraged to check that these platforms include trusted platform modules (TPMs) [30] certified by trusted authorities/manufacturers and that enterprises’ platforms and their software satisfy predefined integrity constraints.

5 Important Issues and Requirements

Important issues and related requirements need to be considered by enterprises when dealing with the management and enforcement of privacy obligations and support for assurance policies:

a) Issues and requirements for privacy obligations

- **Modelling and representation of privacy obligations:** privacy obligations need to be explicitly modeled. This includes representing which data is affected by an 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;
- **Mapping obligations into enforceable actions:** when possible, actions must be expressed in a way that can be programmatically enforced. Otherwise they should trigger related processes and workflows;
- **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;
- **Monitoring obligations:** it is important that the fulfilment of obligations is monitored and checked against expected situations and behaviours. It can always happen that the fulfilment of obligations is either omitted or violated. Monitoring mechanisms must be orthogonal to the enforcement mechanisms. In case of discovery of overdue obligations they should trigger their enforcement and create awareness about the encountered problems;
- **User involvement:** privacy policies and obligations are defined and enforced to preserve user's rights on their personal data. Users need tools and mechanisms to remember and have visibility of the obligations they imposed to an enterprise and potentially monitor their fulfilment. This introduces requirements of transparency about organisational practices;
- **Accountability management:** the explicit management of accountability is fundamental to underpin trust and assurance in people. This introduces requirements in terms of auditing, tracking of obligations and their monitoring;
- **Complexity and cost of instrumenting applications and services:** to be usable and deployable a privacy obligation framework should be deployed in a way that requires a minimum impact on applications and services.

b) Issues and requirements for assurance policies

- **Modelling and representation of assurance policies:** assurance policies need to be explicitly modelled and their constraints represented in a way that can be programmatically processed. There needs to be a practical way for users to specify such policies and it is also important to be able to model and represent the answers and statements provided by enterprises to people;
- **Checking for compliance of assurance policies:** mechanisms are required to allow enterprises to check the status of their IT infrastructure and processes against constraints and expectations provided by people. These mechanisms must

provide clear indications to people about the degree of compliance with these policies;

- **Compliance verification:** mechanisms need to be provided to people so that if desired their computers can remember their policies and expectations and check, over time, compliance against enterprises’ statements based on the enterprises’ behaviours and also various feedbacks.

Being able to address the above requirements and issues is important for those enterprises that want to be transparent in the way the handle personal data and deal with privacy matters. This ultimately has an impact on the way enterprises are perceived by people and their trustworthiness in performing and executing in an expected and predictable way.

6 Model

This section introduces and describes our model to handle privacy obligations and assurance policies in enterprises. It addresses aspects of the problems stated in section 2 and takes into account the above requirements and issues: it is work in progress. Figure 1 shows the core aspects of this model:

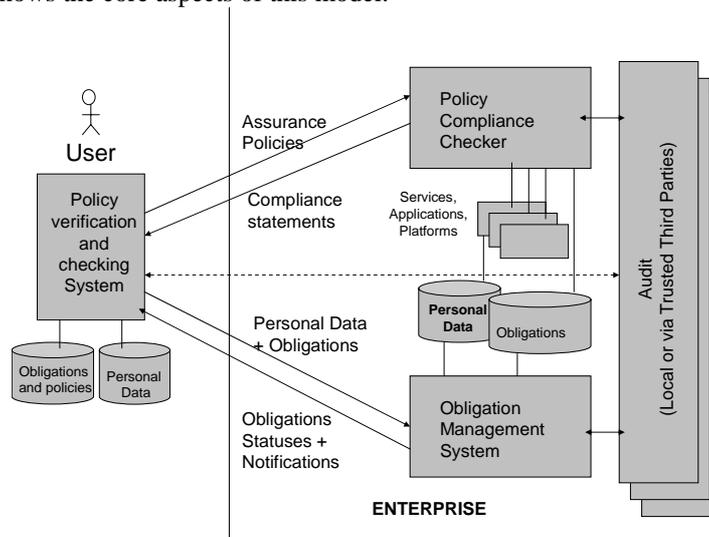


Fig. 1. Model of a system handling privacy obligations and assurance policies

This model supports the following core interactions between users and an enterprise:

- **Users ask an enterprise to demonstrate their support and compliance to a set of policies:** this can be done by users before engaging in any interaction with the enterprise. The “policy compliance checker” module, within the enterprise, issues compliance statements and potentially it supports degrees of verifications

made by users. The outcome is recorded and remembered by the “policy verification and checking system” on the user side for future reference and control;

- **Users disclose their personal data along with their privacy obligations:** user can dictate the set of privacy obligations and constraints they want to be fulfilled on their personal data. These obligations are processed within the enterprise by the “obligation management system”. Periodic feedback and notifications are provided to users, according to their expectations;
- **Users control and verify their expectations and compliance over time:** the “policy verification and checking system”, at the user side, remembers commitments, obligations and promises made by an enterprise. It processes them against evidence and information provided by the enterprise and potential third parties in order to verify their consistency and compliance. This module provides users with intuitive visual clues that help them to make decisions and influence their perceptions of the trustworthiness of an enterprise in executing what has been agreed.

In this model enterprises explicitly check users’ “assurance policies” against their current practices and systems and issues signed statements about their degree of compliance. These statements can be used for future verifications and checks. Enterprises allow users to express customizable obligations when disclosing their data: they manage, enforce and monitor privacy obligations and allow users to be kept in the enforcement loop. Obligations are remembered and periodically checked by users’ side tools.

For example a user might engage for the first time with an enterprise that implements aspects of our model. In addition to other aspects that might be as well supported by the enterprise (such as seals and recommendations by other parties) the user might require the enterprise to assure them about their privacy practices, security and trustworthiness of their IT systems. The user might request the enterprise, by means of assurance policies, to provide them with fine-grained statements about their security systems and business practices and declarations of which privacy policies and obligations they support, specifically to how their data will be handled. The user could go even further by directly checking the trustworthiness of some platforms, via TCG-enabled mechanisms [30], if supported. The user can use their “policy verification and checking system” to verify enterprise statements and promises, remember their expectations and re-check them over time.

If the user is satisfied by these initial statements, they might decide to engage in an interaction or transaction with the enterprise and potentially disclose their personal data. In doing this, the user can specify their privacy obligations, for example in terms of data retention (i.e. deletion after a predefined period of time) and required notification of access, usage and disclosure. Again, the user side “policy verification and checking system” has an important role in remembering these privacy obligations and enabling users to check them over time. For example, if data was supposed to be deleted on the enterprise database at a due time or if the user made the explicit request to receive no further notifications from the enterprise, the user can now actively check if these obligations are fulfilled or violated.

The “policy verification and checking system” is an essential part of our model however its detailed description goes beyond the intent of this paper, as we focus

more on the enterprise mechanisms underpinning policy compliance check and obligation management. A more detailed description of our work on this component and technical aspects can be found in [20, 24].

Last but not least, the way privacy obligations are handled can also be audited by the enterprise, the user and potentially trusted third parties. Auditing is another fundamental aspect to check statements and assertions made by enterprises. It should be at least tamper evident. Users should be able to access audit logs and use them as part of their control and verification activities.

The fact that people are part of an active feedback loop that lasts for the entire duration of their business relationships with an enterprise is very important to allow them to form, review or consolidate their perception on the trustworthiness of an enterprise.

The remaining part of this paper provides more technical details about the systems and solutions underpinning the obligation management system and the policy compliance checker. We envisage these systems as being part of current enterprise middleware, in particular part of identity management solutions. More details of how this could be achieved can be found in [20,21].

7 Technical Details

This section provides technical details on our obligation management system and policy compliance check system. Fig. 2 shows a high-level architecture of an obligation management system. This obligation management system includes the following components:

1. **Obligation Server:** 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 the event handler for their monitoring. One or more obligation servers can be deployed (and synchronised), depending on needs;
2. **Obligation Store and Versioning:** the data repository storing (various versions of) obligations and their mapping to confidential data;
3. **Obligation Scheduler and Manager:** the component that is aware of which obligations are currently active, their ongoing deadlines and relevant events. When events/conditions trigger the fulfilment of one or more obligations, this component activates the correspondent “workflow processes” of the “obligation enforcer” that will deal with the enforcement of the obligation.
4. **Obligation Enforcer:** 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;
5. **Event Handler:** the component in charge of monitoring and detecting relevant events for privacy obligations. These events are defined and pushed by the obliga-

tion 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;

6. **Obligation Monitoring Service:** 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 the involved actions);
7. **Information Tracker:** 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;
8. **Audit Server:** this audits the relevant events and information generated by the overall system components and involved applications/services.

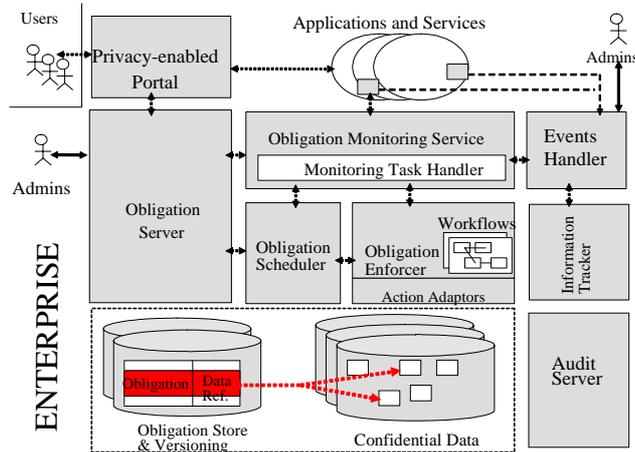


Fig. 2. High-level architecture of an obligation management system.

This system explicitly manages, enforces and monitors privacy obligations. Users can check the status of their data and the enforcement status of privacy obligations. The system also provides the user with the required notifications and feedback, as requested by them. More details are provided in [21,22].

In our model, a privacy obligation describes the relevant events/conditions triggering the obligation, actions to be enforced, the target (i.e. related confidential data) and accountable entities. A simple XML-based example of privacy obligation is described in [21,22]. Details of how to strongly associate privacy obligations to personal data via cryptographic schemas are described in [9,21,22]. It is also important to deal with longevity and survivability aspects of IT solutions when dealing with long-term obligations. Work has already been done in this space, including [15,16,17,18,19], and can be leveraged.

Figure 3 provides technical details about a “policy compliance checker” system.

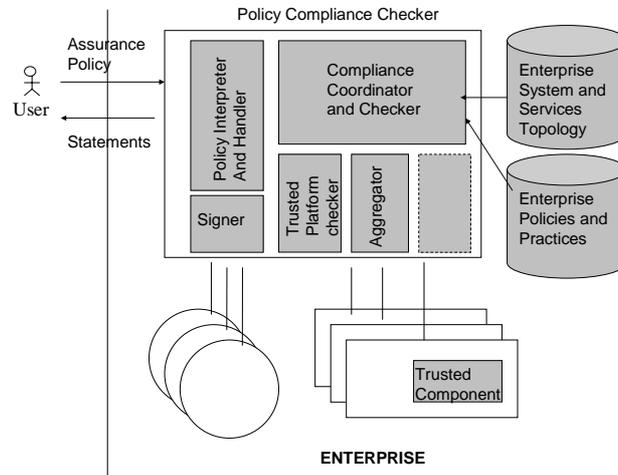


Fig. 3. High-level architecture of a policy compliance checker

This system includes the following components:

1. **Policy interpreter and handler:** the component that interprets an assurance policy and determines if it is well expressed and can be handled by the system. If it can, it is passed to the “compliance coordinator and checker” component;
2. **Compliance coordinator and checker:** this component coordinates the collection of required information to provide support to tests/requirements expressed by a user via an “assurance policy”. It can potentially allow remote users to directly perform some test of platforms, for example trusted platforms. This happens via the usage of the “trusted platform checker” and “aggregator” sub components;
3. **Trusted platform checker:** the component that interacts with and retrieves information about the status of critical platforms running enterprises’ services and applications and that store and handle personal data; it also analyses this information to assess the trustworthiness of those platforms with respect to the current context;
4. **Aggregator:** this aggregates information collected from various enterprise systems, in order to provide a comprehensive result to the user; this may involve analysis to provide an overall trust assessment;
5. **Enterprise system and services topology:** database containing information about the topology of enterprise systems and services; this is used during the checking and aggregation phases;
6. **Enterprise policies and practices:** database containing information about the policies and procedures supported by an enterprises. It is used during the checking and aggregation phases;
7. **Signer:** the component responsible for signing the statements made by the policy compliance checker for integrity and non-repudiation reasons; this could be done via a trusted hardware device such as a TPM [29].

The policy compliance checker component is work in progress. At the moment we are exploring how an enterprise can provide simple assurance statements to users about the usage of trusted platforms [29,30] in an aggregated way, along with active verification by users (using platforms leveraging the same trusted computing technology).

8 Discussion and Current Work

Dealing with assurance policies, generating compliance statements, enforcing and managing obligations, providing feedback and verification mechanisms to users is not a trivial task, especially when the final goal is to underpin trust and assurance in users.

This paper describes preliminary work to address these problems, mainly from an enterprise perspective. Our aim is to provide “good-willing” enterprises with tools and mechanisms that help them to support privacy obligations and allow users to check their expectations in terms of privacy, security and trustworthiness of enterprises’ IT solutions. We consider the case where these enterprises are willing to collaborate with users, make assurance statements and be compliant with privacy policies and, more specifically, privacy obligations. Additional assurance and accountability can be added by hardening the audit server [31,32] and involving trusted third parties in the monitoring of the enforcement of obligations policies.

In the end, what really matters is that users can make informed decisions whether or not they should trust enterprises, in particular from a privacy perspective. We argue that this can be achieved not only by relying on third parties’ recommendations or certifications but also on users’ direct experience and interaction with enterprises. Hence it is important for users to have tools to remember promises, statements and obligations underpinned by enterprises and periodically check them against evidence and feedbacks. More details about how we address these latter aspects can be found in [20].

Progress has been made in designing and refining the architecture of a system based on our model, both at the user and enterprise side. More details can be found in [24]. This work has been done in the context of the EU PRIME project [23].

Working prototypes of the “policy verification and checking system” (user side) and the “obligation management systems” (enterprise side) have already been implemented. More details on the latter prototype can be found in [22].

Progress has also been made in leveraging TCG-based trusted platforms and TPM modules [29,30], in the context of the “policy compliance checker”.

9 Future Work

Our work and research is definitely in progress: technical aspects need to be further refined and investigated, particularly the ones related to the “policy compliance checker” and the full integration of our solution in an enterprise identity management

solution. The overall implications for the involved enterprise applications and services have yet to be fully understood. One of the reasons for developing our prototypes is to make advancements in these areas by experimenting and refining our concepts. Work in this space will be carried on in the context of the EU PRIME project [23].

10 Conclusions

Among other things, privacy and privacy management are important to underpin trust and assurance in enterprises. In this paper we address related problems: how can enterprises provide people with degrees of assurance that they will operate in the way dictated by policies and privacy obligations, accordingly to people's expectations; how can enterprises explicitly manage these obligations; how can people check up-front that an enterprise has the right capabilities to handle and process their personal data; how can people have a constant, personalized feedback on the fulfillment of all these aspects?

We introduce a model and a technical approach to make progress in addressing these problems. Our approach allows people to express, check, remember and verify "assurance policies" against statements and promises made by enterprises. People can define privacy obligations associated to the personal data they disclose to enterprises. On the enterprise side we describe systems for managing and enforcing obligations and for issuing policy compliance statements.

Enterprises adopting this model and related systems can explicitly handle privacy obligations and provide additional assurance to people. People, on the other hand, get more information and detail for making informed decisions about the willingness and trustworthiness of enterprises in maintaining their promises.

Our work is in progress: an initial architecture of our solution and related prototypes have been developed in the context of the EU PRIME project. Future work will be done in this project to refine our concepts and prototypes.

References

1. Laurant, C., 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> (1980)
3. Online Privacy Alliance: Guidelines for Online Privacy Policies. <http://www.privacyalliance.org/>, Online Privacy Alliance (2004)
4. Karjoth, G., Schunter, M.: A Privacy Policy Model for Enterprises. IBM Research, Zurich. 15th IEEE Computer Foundations Workshop (2002)
5. Karjoth, G., Schunter, M., Waidner, M.: Platform for Enterprise Privacy Practices: Privacy-enabled Management of Customer Data. 2nd Workshop on Privacy Enhancing Technologies, Lecture Notes in Computer Science, Springer Verlag (2002)

Marco Casassa Mont, Stephen Crane and Siani Pearson

6. Schunter, M., Ashley, P.: The Platform for Enterprise Privacy Practices. IBM Zurich Research Laboratory (2002)
7. Karjoth, G., Schunter, M., Waidner, M.: 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. Casassa Mont, M., Pearson, S., Bramhall, P.: Towards Accountable Management of Privacy and Identity Information, ESORICS 2003 (2003)
10. IBM: IBM Tivoli Storage Manager for Data Retention (2004)
11. Bettini, C., Jajodia, S., Sean Wang, X., Wijesekera, D.: Obligation Monitoring in Policy Management (2002)
12. Damianou, N., Dulay, N., Lupu, E., Sloman, M.: The Ponder Policy Specification Language (2001)
13. Housley, R., Ford, W., Polk, W., Solo, D.: RFC2459: Internet X.509 Public Key Infrastructure Certificate and CRL profile. IETF (1999)
14. Agrawal, R., Kiernan, J., Srikant, R., Xu, Y.: Hippocratic Databases. IBM Almaden Research Center (2002)
15. Anderson, R. J.: The Eternity Service. Proc. PRAGO-CRYPT 96, CTU Publishing House, Prague (1996)
16. Ellison, R.J., Fisher, D.A., Linger, R.C., Lipson, H.F., Longstaff, T.A., Mead, N.R.: Survivability: Protecting your Critical Systems. Proceeding of the International Conference of Requirements Engineering (1998)
17. Kubiatowicz, J., Bibdel, D., Chen, Y., Czerwinski, S., Eaton, P., Geels D., Gummadi, R., Rhea, D., Weatherspoon, H., Weimer, W., Wells, C., Zao, B.: OceanStore: An Architecture for Global Scale Persistent Storage. University of California, ASPLOS 2000 (2000)
18. Neumann, P.G.: Practical Architectures for Survivable Systems and Networks. SRI International, Army Research Lab (1999)
19. Wylie, J.J., Bigrigg, M. W., Strunk, J. D., Ganger, G. R., Kiliccote, H., Khosia, P.K.: Survivable Information Storage Systems. IEEE Computer (2000)
20. Crane, S., Casassa Mont, M., Pearson, S.: On Helping Individuals to Manage Privacy and Trust. HPL Technical Report, HPL-2005-53 (2005)
21. Casassa Mont, M.: Dealing with Privacy Obligations: Important Aspects and Technical Approaches, TrustBus 2004 (2004)
22. Casassa Mont, M.: Dealing with Privacy Obligations in Enterprises, ISSE 2004 (2004)
23. PRIME: Privacy and Identity Management for Europe, European RTD Integrated Project under the FP6/IST Programme, <http://www.prime-project.eu.org/> (2004)
24. PRIME: PRIME Architecture V0, D14.2.a, editor: Dieter Sommer (IBM Labs), <http://www.prime-project.eu.org/> (2004)
25. Federal Privacy Commissioner, Web Seals: A review of Online Privacy Programs, <http://www.privacy.gov.au/publications/seals.pdf> (2000)
26. Resnik, P., Varian H.R.: Recommender Systems, Communications of ACM, <http://www.acm.org/pubs/cacm/MAR97/resnick.html> (1997)
27. Reputation Research Network: Online papers on reputation and reputation research, <http://databases.si.umich.edu/reputations/index.html> (2004)
28. W3C: The Platform for Privacy Preferences 1.0, <http://www.w3.org/TR/P3P/> (2002)
29. TCG: Trusted Computing Group, <https://www.trustedcomputinggroup.org/home> (2004)
30. Pearson, S. (ed.): Trusted Computing Platforms, Prentice Hall (2002)
31. Baldwin, A., and Shiu S.: Enabling shared audit data, IJIS (to appear) (2004)
32. Baldwin, A.: Enhanced accountability for electronic processes, 2nd international conference on trust management. Lecture notes in computer science, vol. 2995, Springer (2004)