

Privacy and Identity Management for Everyone

Jan Camenisch* abhi shelat* Dieter Sommer*
Simone Fischer-Hübner† Marit Hansen‡ Henry Krasemann‡ Ronald Leenes§
Jimmy Tseng¶

ABSTRACT

The shift from a paper-based to an electronic-based society has dramatically reduced the cost of collecting, storing and processing individuals' personal information. As a result, it is becoming more common for businesses to "profile" individuals in order to present more personalized offers as part of their business strategy. While such profiles can be helpful and improve efficiency, they can also govern opaque decisions about an individual's access to services such as credit or an employment position. In many cases, profiling of personal data is done without the consent of the target individual.

In the past decade, the European Union and its member states have implemented a *legal* framework to provide guidance on processing of personal data with the specific aim to restore the citizens' control over their data. To complement the legal framework, the PRIME (Privacy and Identity Management for Europe) project [13] has implemented a *technical* framework for processing personal data. PRIME's vision is to give individuals sovereignty over their personal data so that:

- ▷ Individuals can limit the information collected about them by using pseudo-identities, certifications and cryptography when performing online transactions,
- ▷ Individuals can negotiate legally-binding "privacy policies" with their service providers which govern how disclosed personal data can be used and which precautions must be taken to safeguard it, and
- ▷ Individuals and service providers can use automated mechanisms to manage their personal data and their obligations towards data which they have collected from other parties.

To accomplish this, the PRIME project has designed and implemented a practical system-level solution which incorporates novel

*IBM Zurich Research Lab, Säumerstrasse 4, 8803 Rüschlikon, Switzerland, email: {jca, abs, dso}@zurich.ibm.com

†Karlstads Universitet, Sweden

‡Unabhängiges Landeszentrum für Datenschutz, Germany

§Universiteit van Tilburg, The Netherlands

¶Erasmus Universiteit Rotterdam, The Netherlands

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DIM'05 November 11, 2005, Fairfax, Virginia, USA.

Copyright 2005 ACM 1-59593-232-1/05/0011 ...\$5.00.

cryptographic protocols, sophisticated security protocols, and artificial intelligence algorithms. This paper describes the architecture of this system. The key features of this architecture have been implemented in a proof-of-concept prototype.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Communication Applications—*privacy*

General Terms

Security, Design, Management

Keywords

Privacy, Identity Management, Protocols, Credentials, Anonymous transactions

1. STATUS QUO PRIVACY PROBLEMS

Many individuals are surprised by the amount of personal data requested when they provision an online service. Why does buying a train ticket require disclosure of a phone number?

In addition to provisioning, many other processes of modern life generate additional data that can be monitored, stored and analysed. Our concern is that the use of these data may be quite different from what people expect and what they want.

A recent survey by Turow, Feldman and Meltzer [19], for instance, reveals that about two-thirds of surveyed Americans do not know that US supermarkets are allowed to sell information about individual purchase decisions to other companies. While this kind of secondary use of personal data requires the consent of the customers in the EU, European shoppers may not be aware that they give this kind of consent when signing up to loyalty programs. In the same survey, the respondents also object to online price discrimination and to most forms of behavioural targeting—at the same time, their online behaviour across websites is monitored, enriched by other (purchased) data, and used for targeted advertisements.¹

Another concern is the prospect of large-scale identity thefts, such as theft of credit card data, social security numbers, and student numbers. In June 2005, CardSystems Solutions, a large credit card payment processor in Tucson, Arizona announced that 40 million credit card numbers may have been stolen by computer hackers. More alarmingly, the scope of the theft was a direct result of

¹Specifically, 87 percent of the respondents objected to the idea that "an online store could charge people different prices for the same products during the same hour." Such price discrimination is legal and can be expected to be used in the future.

the company's illegal practice of retaining transaction records for their unauthorized use:

John M. Perry [CEO of CardSystems] told The New York Times the cardholder data was kept for "research purposes." MasterCard and Visa both require card processors such as the one CardSystems ran in Tucson, Ariz., to expunge that information once it is passed on to the banks. Instead, the Atlanta-based company retained records. "We should not have been doing that," Perry told the newspaper. [18]

If personal data is aggressively analyzed, individuals may only be presented the limited set of choices which are "deemed" appropriate for their personal profile. This in turn may lead to refinements in their profile leading to further restrictions. Another concern relates to how individuals might fear certain types of expression. Fear of creating online traces may, for instance, warrant a father to discourage his daughter from writing a school essay on his employer's business ethics, as this essay could possibly affect his professional career. As a final concern, unbridled data collection and profiling by a government in the name of protecting (national) security may lead to unjust and ultimately unwarranted blacklists.

1.1 Current Solutions Are Inadequate

One common misconception is that people voluntarily give away their personal data. The choice is not a conscious one, but rather due to a lack of alternatives: people have little choice but to fill out the mandatory fields of web forms. One might argue that if market forces prevail, consumers eventually choose providers that cater to their preferences by switching whenever there are incompatibilities between what is offered and what is desired. However, in practice there is a power imbalance. Transaction costs and uncertainties with respect to the risks of personal data abuse prevent people from actually going elsewhere.

Moreover, if a company suffers a security breach and loses its private customer data, the customer, rather than the firm, bears the economic consequences (i.e., identity theft, invasion of privacy). Thus, as long as businesses are able to externalise the costs of security breaches, we can not expect them to adopt better privacy and identity management measures.

On the other hand, there is a well-established legal framework codified in the EU Data Protection Directive 95/46/EC and the E-Communications Privacy Directive 2002/58/EC which protects personal data in Europe. These regulations set the conditions for the processing of personal data, and offer citizens rights on inspection of the data held on them by organisations processing personal data, as well as the right to have errors corrected and data removed from certain databases.

In practice, the complexity of the regulation, incomplete enforcement, the unawareness of people, and sometimes even conscious decisions by businesses and governments not to comply with the rules, render legislation ineffective.

Considering the problems outlined above, the current technological tools are also insufficient from both a functional point of view and from the fact that current tools require too much effort on the part of the user. Moreover, the lack of standardisation and interoperability makes the use of even the simplest privacy technology ineffective.

A NEED FOR CHANGE. Individuals will slowly discover that businesses and governments know a lot more about them and their behaviour than they expected, perhaps because personal data they had consentingly disclosed for one purpose was being used for unauthorized secondary ones. They will notice that their personal infor-

mation is being negligently stored and therefore vulnerable to theft and misuse by perpetrators. Our aim with PRIME is to address this situation from both a customer and firm's perspective.

2. PRIME'S VISION

The goal of the PRIME project is to reconcile privacy and accountability of users' electronic interactions. The PRIME consortium envisions a system in which people can use information services in a reliable and trustworthy way while keeping control over the details of their private life as in the paper-based world. Within the PRIME project, an architecture for the system has been constructed and a first prototype been implemented. In this section, we briefly outline some specific points of PRIME's goal.

USER INFORMED CONSENT AND CONTROL. The user keeps control over which personal data are given to whom and for which purpose and maintains a complete and coherent view of the privacy policy of all their transaction partners.

PRIVACY NEGOTIATION AND DISPUTE HANDLING. When a user discloses personal data, the user can express a privacy policy which states how her personal data should be handled. To make such policies meaningful, a user can *negotiate* with her transaction partners and conclude an agreement that forms contractual provisions on the privacy rights and obligations of the parties involved in the transaction. Such agreements serve as legal contracts that must be fulfilled by the transaction partners.

DATA MINIMIZATION AND IDENTITY MANAGEMENT. Transaction partners only collect personal data that are necessary to perform their part of the transaction. When they no longer need these data for their stated and agreed purpose, they shall delete them. This is in line with the legal principle of data minimization.

The transaction partners' need-to-know data depends on the business process. For example, it might be required that a user prove that she is aged over 18. This could be performed equally well with an ID card, passport, or—in some countries—with a driver's license credential. Besides the fact that the holder of the credential is aged over 18, each of those credentials would reveal additional data not required in the business process if being disclosed to the provider. Depending on the user's preferences, either of the credentials could be the better choice from a privacy standpoint. From a privacy point of view, the best possible approach in this case would be to only prove that the user possesses a credential with an appropriate birth date attribute without revealing any other information.

Furthermore, users may want to release different credentials and personal information to different transaction partners, as the level of trust varies. This "user-controlled identity management" is performed intuitively by everybody in the offline world. People meet face to face and decide in each situation what information to reveal. In the online world, identity management is even more relevant, as information from different contexts can be collected more easily. Use of the same username across multiple transactions can yield comprehensive profile information on the usage, interests or behaviour of the user to the service provider.

SPECTRUM OF ANONYMITY. The PRIME system does not impose full anonymity, but instead supports a broad range of possible transactions. At one end of the spectrum, the parties agree to proceed without the need for any identifying data and the relationship can stay anonymous. Consider, for example, browsing a web page while using an anonymous communication channel.

In low-risk transactions, a non-cryptographic pseudonym, which serves to link the online and offline parts of the transaction, is sufficient. Consider, for example, reserving a seat at the theatre, which

could be accomplished by providing or receiving a random identifier which acts as a proof of the reservation.

Alternatively, cryptographic credentials (and their related protocols) can improve accountability, while, at the same time, keeping the parties of the transaction anonymous. In such systems, identities are only revealed in case of dispute or fraud. Such systems also facilitate the ability for the parties to reveal a small part of their identity, for example to build reputation or to obtain personalized services.

Finally, at the other end of the spectrum, in medium-to-high-risk transactions and law-related transactions, a third-party-issued identity proof such as an identity card, or a witness like a notary might be necessary. Examples of such transactions include taking an exam or buying a high-value asset, like an apartment, a piece of land, etc.

ACCOUNTABILITY. Let us reiterate that properly-designed anonymous transactions can also provide accountability. One way to do so is by requesting the user to verifiably encrypt her identity using the public key of a mutually trusted third-party. Since the encryption is verifiable, the service provider can be assured that the encryption contains the identity vouched for by a certifier. If a dispute arises and certain well-defined conditions are met, the service provider may request that the trusted third-party decrypts the identity of the user.

3. THE PRIME SOLUTION

The PRIME project has designed and implemented a system for identity management system which addresses the goals listed above. The architecture defines the interoperation of several components. Some key components of the architecture handle access control, advanced anonymous credential systems, and automated reasoning. In this short abstract, we focus on the cryptographic mechanisms and briefly sketch the other key components. We begin with a summary of the parties involved.

3.1 The parties involved

USER. A user in the system has certificates, data and policies regarding their data. A user engages in transactions with service providers during which their data may be exchanged.²

SERVICE PROVIDER. A service provider offers services and resources to an interested user by means of *transactions*. A service provider may have certificates and private data, and may also have access control policies over their services and resources. Some service providers, such as shipping companies, do not directly interact with users, but nonetheless handle user data.

CERTIFYING AUTHORITY. A special type of service provider is a certifying authority that issues *certificates*, i.e., digitally-signed statements. By issuing a certificate, a certificate authority vouches for the truthfulness of the statement. (A reputable certificate authority, therefore, has a clear incentive to verify the statements before issuing a certificate asserting them.)

3.2 Cryptographic Tools

We present the key cryptographic mechanisms of the PRIME system with a particular emphasis on credentials and what we can achieve by using them.

SECURE COMMUNICATION. We assume that all communications between a user and a service provider are performed over an encrypted, semi-anonymous channel. The company is authenticated

²Interactions between users are envisioned, but not considered in this abstract.

to the user by means of standard technology (e.g. an X.509 server-side certificate [14]) and the user remains unauthenticated. The TLS protocol provides such a channel.

ANONYMOUS COMMUNICATION. By default, a user will remain anonymous by using an anonymizing network. Anonymization is provided for example by onion routing networks [12], mixnets [8], or crowds [17].

PSEUDONYMS. A pseudonym is the name under which a user is known to one or multiple service providers. Simple pseudonyms are simply random strings which can be generated by the user at any time. Cryptographic pseudonyms such as Idemix pseudonyms [5], are more complicated to generate and may require a certification authority, but allow the owner to cryptographically prove ownership of the pseudonym or issue signatures under the pseudonym.

CREDENTIALS AND PROOFS OF OWNERSHIP OF CREDENTIALS. One of the key building blocks in the PRIME system are cryptographic *credentials*. A *credential* is a piece of data such as a birth date or postal address, or a list of such data items, certified by a third party. A credential is often also called *certificate* or *attribute certificate*. It is important that a credential be bound to its owner by cryptographic means, e.g., by requiring the owner's secret key to use the credential [5]. This is important for accountability of anonymous transactions and the prevention of users sharing their credentials. Binding credentials to hardware is another option to prevent sharing [4].

From a privacy perspective, the use of credentials is preferable to the direct request to the certifying party as it prevents the certifying party from profiling the user because it is unaware of the identity of the user. Credentials can either be realized using traditional attribute certificates, see e.g. [11], where the reference to the user could be the user's real name or a pseudonym, or by so-called *private credentials*, see e.g. [9, 15, 5].

Traditional certificates have the drawback that different uses of the same certificate can be linked to each other. Private certificates do not suffer from this drawback. They allow users to disclose selectively certain personal information and be certain that nothing more than the selected information is disclosed. For instance, a user owning an identity card as a private credential, containing name, address, and birth date, can prove being older than 18, without revealing any of name, address, or birth date, nor making this transaction linkable to any other transaction.

Private credentials can also allow for verifiably encrypting an attribute under a third-party public key to ensure that only the third-party can access the attributes. For instance, consider a credit-card number which has been certified by a bank. By encrypting the attribute, a user could assure a merchant that a payment will be approved without the merchant learning the credit card number. That is, the user provides, in addition to the encryption of the credit card number, a proof that the encryption contains the credit card number attribute of her credit card certificate issued by her bank.

Furthermore, a condition agreed between the party providing a verifiable encryption and the receiver is cryptographically bound to the encryption such that the third party under whose key the encryption has been done will only be able to decrypt if the very same condition is supplied to the decryption process. The condition specifies the circumstances under which the ciphertext may be decrypted. Thus, as mentioned before, private credentials together with encryption of attributes enable transactions to be privacy-protecting and yet accountable.

3.3 System Architecture

Our basic system architecture is explained below. Both user and service provider share essentially the same architecture. Before

describing the architectural components, we present the data model of the architecture and our way of identifying resources; both these concepts apply throughout the architecture.

Resource Referencing Scheme. To be meaningful, an access control system must have a well-defined system for “naming” resources. For this purpose, we use the Uniform Resource Identifier (URI) scheme proposed in RFC 2396 [1] to name every resource in our system.

URIs are general enough to be used as a naming system for categories of data, instances of data, process workflows, specifications, elements of policies, and obligations such as “Delete this data after two weeks.”

Data Model and Ontology. In order to account for interoperability between two parties and the components within a party, a data model or *ontology* is required.

We selected the Resource Description Framework (RDF) as defined in [16] for representing data. RDF allows one to model relations between data as an arbitrary graph, and provides the ability to associate privacy obligations to data without revealing the data itself in the reference. RDF also allows automated reasoning algorithms to determine which credentials satisfy a requirement (e.g. that a driver’s license can be used to prove one’s age) and also compute the set of certificates and attributes which satisfy a requirement while revealing the least information. Moreover, if the RDF has been designed properly, the automated reasoning algorithms can also provide explanations for a particular decision—a feature which is a critical for creating useful graphical interfaces.

At the center of the system architecture is a database which holds certificates and declarations of a party (declarations are uncertified data such as a user’s name or address, which are generated by the user or company). Furthermore, the database contains default policies for the release of information, and logs of previous interactions with other parties. These logs, for example, can contain information on the receivers of particular data attributes of the party and help in deciding on further disclosures to particular parties. A service provider in addition stores the data that it has collected from other parties.

To control access to the database, there is an Access Control component (AC), an Identity Control (IC), and a graphical user interface (GUI) for the overall privacy and identity management task. In addition, service providers will also have an Obligation Management component (OM) which manages all of the privacy obligations the company has assumed regarding the data it has collected from its customers [7, 6].

ACCESS CONTROL. The Access Control component limits access to a party’s resources. A party’s resources include all data in the database and other “external” resources, such as services provided by a company.

The interface to the Access Control component is simple: A request to the Access Control consists of a URI identifying the resource to be accessed, a purpose for the access, and auxiliary information provided by the requester such as certificates or declarations. The reply of the access control can be either of the following: i) A DENY answer with a list of preconditions that the requester has to fulfill in order to access the resources, or ii) a GRANT answer with the requested resource.

Our access control system follows the *attribute-based* paradigm of access control [2, 3] in which access is granted based on the *properties* that the requester has asserted via the auxiliary information. Thus, our system explicitly avoids relying on the identity of the requester to determine how to respond to the request.

As a concrete example, a driver’s license might be used to establish one’s age, and therefore one’s legal right to purchase alcohol.

A driver’s license, however, reveals much more information than the fact that the owner of the license is of legal age. In our system, a credential certifying that the presenter is of legal age would suffice for the transaction.

The attribute-based access control paradigm is particularly valuable in open environments such as the Internet where identity data can often not be meaningfully interpreted. But attributes of a requester, such as their age, the fact that they possess a driver’s license etc. can be interpreted without needing to identify the requester.

Technically, policies are referenced in a way that allows the Access Control component to efficiently gather all policies applying to a particular resource of a request.

IDENTITY CONTROL. The Identity Control component manages all interactive protocols with other parties.

Indeed, the PRIME identity management system uses sophisticated protocols to present certificates in a privacy-preserving manner. Moreover, we envision that negotiations between a customer and a company about privacy policies will eventually also require messages sent back-and-forth. Such coordination is handled by this component.

More specifically, the IC component (a) delegates requests to the AC, (b) handles all credential-related protocols, (c) automatically computes optimal ways to fulfill a request, and (d) manages user input and notification via the graphical user interface.

OBLIGATION MANAGER. The obligation manager maintains all of the obligations which have been accrued by a party through its various transactions. An *obligation* is an event-condition-action (ECA) rule and is generally activated any time that data is stored to the database. The obligation manager triggers a workflow process whenever the conditions defined by the obligation are met. For example, an obligation may be of the form “delete data record 21321 by August 15, 2006.”

4. A SAMPLE TRANSACTION

In practice, a business process governs how a transaction between a user and a service provider might be structured. In order to provide a more detailed example of how data and identity are managed by the PRIME system, we present the following example of a common transaction between a buyer and a seller. We provide a more detailed description of the components involved in the transaction in §3.

Transactions proceed in sub-transactions each consisting of two phases as outlined below and depicted in Figure 1.

NEGOTIATION – PHASE I

1. The user requests information about a product from a service provider.
2. The request is received by the service provider and directed to the AC component. The AC component returns an *offer* which includes a description of the product, a list of requirements in order to buy the product along with corresponding reasons for each of the requirements.

The list of requirements can include the price, a request for the user’s address, billing information, phone number, etc. The list of reasons can explain why certain information, i.e. phone number, has been requested.

The offer also specifies how the data related to this transaction will be treated. In particular, the service provider presents obligations to the user that will be automatically enforced.

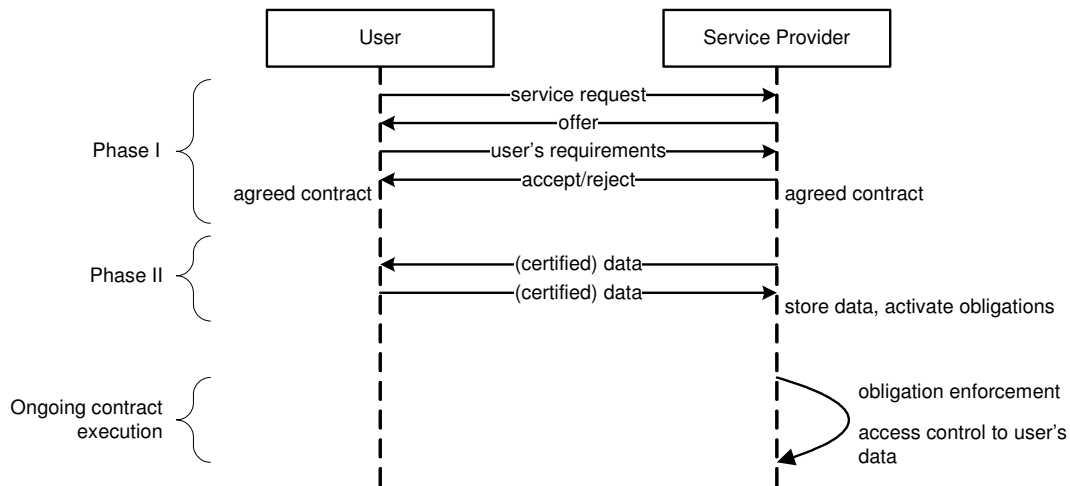


Figure 1: Execution of a transaction

The service provider's AC may also reply with multiple offers for the same product. For example, there can be a standard offer with the retail price, and a special offer with a reduced price which requires that the user provide a loyalty program number.

3. The user's IC component receives the offer and parses it. Each of the requirements are presented to the user's AC in order to determine the counter-requirements for the release of the required information.

The IC may add obligations to the offer, for example, it may add the obligation that the company notify the user whenever the transaction data is transferred to a third party.

The IC presents choices about how the requirements are to be fulfilled to the user via the GUI. For example, a user might have to choose between multiple offers, or choose between various ways to fulfill a requirement (i.e., by using driver's license versus passport). For convenience, the user can configure certain choices to be made automatically (e.g., if possible, use e-coins to fulfill a payment, and if not, use a credit card).

4. The service provider either accepts or rejects the offer.

CONTRACT EXECUTION – PHASE II

After both parties have accepted the same offer, the transaction defines a contract that is executed automatically by the PRIME machinery as follows.

1. Company sends necessary credentials to the user.
2. User's IC uses the received credentials to access user's information via the AC.

The AC responds with the requested data. (If the access policy has changed between the time at which the contract was accepted and the time of this access, the user must decide how to proceed.)

The IC packages the requested data and sends it back to the company. This can involve shows of credentials or sending declarations.

3. The company's IC processes the requested data and determines whether the requested information satisfies the contract. If so, the IC requests the AC to store parts of the user data and to store the related obligations in the OM. The IC also triggers any business processes related to the transaction (e.g., to deposit the e-coins and to ship the good to the user).

4. The OM handles any obligations whose conditions have been triggered. For example, when the company relays the user's address to the shipping company, the OM informs the user that such information has been transferred. Obligations can be completely orthogonal to any services side accesses to the user's data, but can also be related to such accesses as in the example.

The two-phase transaction outlined above is the simplest case of a transaction. In more complicated situations, multiple two-phase sub-transactions might be required. For instance, if a service provider requires a user to prove membership in some group before continuing the interaction, the contract negotiation phase and contract execution phase for the membership sub-transaction can be completed before the main transaction is executed. Note that in such a case, the contract in the first sub-transaction does not yet yield any business obligations such as service provisioning to the service provider.

4.1 Extensions to the System Architecture

To complement the architecture described above, we have implemented some extensions to provide a more robust privacy and identity management toolkit.

Policy Compliance. Why should a user trust the service provider's information processing system? How can one be sure that the system is running the correct software?

The Policy Compliance (PC) component addresses this issue by generating and checking "assurance information" about the trustworthiness of a computer system. The component dynamically generates and verifies declarations or certificates that a computer system is functioning "properly."

In the simplest case, the component can present statically installed certificates such as trust seals or audit certificates which have been issued to a service provider by a third party who has inspected the service provider's platform.

More generally, “trusted hardware modules,” which are tamper-proof devices which house their own protected memory and a simple processor, can be used to generate cryptographic proofs that the service provider’s computer platform is running a specific version of software, that no other rogue processes are executing on the machine, and that no other processes have accessed the customer’s private data, etc. Such proofs do not reveal any business secrets regarding the service provider’s system configuration.

Reasoner. How does a user parse a computer-generated list of requirements, choose which credential satisfies the requirement, and pick a pseudonym to use in the transaction?

To simplify the user’s selection process, the Reasoner component provides automated reasoning over ontologies. An *ontology* is a graph representing which data can satisfy which requirements. One can, for example, encode the fact that “a US driver’s license satisfies the requirement that the owner is older than 16” in an ontology. Our architecture uses OWL [10] as the language for expressing ontologies.

With a well-designed and standardized ontology, the Reasoner can determine which of a set of certificates satisfies a requirement while revealing the least additional information. Additionally, the Reasoner can parse the rather detailed and technical assurance information from a service provider’s PC component, and determine whether this information satisfies some more abstract requirement (e.g., “Level 2” versus “Level 3” security) specified in the user’s policy.

Other Components There are additional software components defined in the software architecture. In particular, an Event Management component provides a framework for handling any kinds of events, for example, events that are generated when user’s data is accessed at the services-side database; such events are needed by the Obligation Management component in order to enforce the ECA rules.

5. CONCLUSION

The identity management system described in this paper serves both user’s and service provider’s needs in order to implement the EU Directives 95/46/EC and 2002/58/EC (whose purposes are to safeguard individuals’ privacy and freedom). To our knowledge, this system is the first one that takes this comprehensive approach to tackle the privacy problem.

Our system includes an anonymous credential system, an access control system based on a novel paradigm, a negotiation functionality, and an automated reasoning system. This machinery performs most of the decision making involved in privacy management and involves the user mainly for making final high-level decisions and for giving consent to data processing. Together, these components give a user the power to easily manage her privacy without being an expert in the field.

At the services side, our access control paradigm puts forth a new approach in services-side identity management. In particular, authorizations are not made on identities as used in today’s scenarios, but rather on properties of the users that typically do not identify them. Thus identities of requesters can just be pseudonyms to which properties are bound which still allows for customer relationship management if the same pseudonym is reused.

Although the system allows a user to act anonymously in many cases, it can at the same time allow the service provider to hold her accountable. That is, law enforcement is sufficiently supported by the ability of third parties to revoke the anonymity of selected transactions in certain situations. Moreover, instead of allowing a single trusted party to revoke the anonymity of a transaction, such power can be distributed among many external parties by employ-

ing standard techniques from threshold cryptography.

Since our system also allows a user to assess the trustworthiness of a service provider, it will be easier for less well-known companies who are willing to fulfill their obligations regarding the handling of customer data to gain trust more quickly than is otherwise possible these days. In particular, privacy seals provide an incentive to service providers to run compliant software and take the enforcement a contract seriously. Trusted hardware will make it much more difficult to tamper with the enforcement mechanisms and thus a company will be able to provide quite convincing evidence of compliance to its users.

For businesses, the selling points of privacy-enhanced services are (a) the development of customers’ trust in the services offered, (b) process improvement as an investment to enable scalability with cost control, (c) and cost reductions from automated privacy handling.

6. ACKNOWLEDGEMENTS

The authors would like to thank the PRIME project and the Swiss government for financial support, and the anonymous referees for their comments. In addition, we would like to thank Gérard Lacoste for his support and useful feedback.

7. REFERENCES

- [1] BERNERS-LEE, T., FIELDING, R., AND MASINTER, L. Uniform Resource Identifiers (URI): Generic Syntax. RFC 2396 (Draft Standard), Aug. 1998. Obsoleted by RFC 3986, updated by RFC 2732.
- [2] BONATTI, P., AND SAMARATI, P. Regulating service access and information release on the web. In *CCS '00: Proceedings of the 7th ACM conference on Computer and communications security* (New York, NY, USA, 2000), ACM Press, pp. 134–143.
- [3] BONATTI, P. A., AND SAMARATI, P. A uniform framework for regulating service access and information release on the web. *J. Comput. Secur.* 10, 3 (2002), 241–271.
- [4] TRAMENISCH, J. Protecting (anonymous) credentials with the trusted computing group’s trusted platform modules v1.2. Tech. rep., IBM Research, Jan. 2005.
- [5] CAMENISCH, J., AND LYSYANSKAYA, A. Efficient non-transferable anonymous multi-show credential system with optional anonymity revocation. In *Advances in Cryptology — EUROCRYPT 2001* (2001), B. Pfitzmann, Ed., vol. 2045 of *LNCS*, Springer Verlag, pp. 93–118.
- [6] CASASSA MONT, M. Dealing with privacy obligations: Important aspects and technical approaches. In *TrustBus 2004* (2004), pp. 120–131.
- [7] CASASSA MONT, M. Dealing with privacy obligations in enterprises. In *ISSE* (2004).
- [8] CHAUM, D. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM* 24, 2 (Feb. 1981), 84–88.
- [9] CHAUM, D. Security without identification: Transaction systems to make big brother obsolete. *Communications of the ACM* 28, 10 (Oct. 1985), 1030–1044.
- [10] DEAN, M., AND SCHREIBER, G. OWL web ontology language reference. W3C Recommendation.
- [11] FARRELL, S., AND HOUSLEY, R. An Internet Attribute Certificate Profile for Authorization. RFC 3281 (Proposed Standard), Apr. 2002.
- [12] GOLDSCHLAG, D. M., REED, M. G., AND SYVERSON, P. F. Onion routing for anonymous and private internet

- connections. *Communications of the ACM* 42, 2 (Feb. 1999), 84–88.
- [13] HANSEN, M., AND KRASEMANN, H. Prime whitepaper. Whitepaper, 18 July 2005. http://www.prime-project.eu/prime/public/press_room/whitepaper/PRIME-Whitepaper-V1.pdf.
- [14] HOUSLEY, R., POLK, W., FORD, W., AND SOLO, D. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 3280 (Proposed Standard), Apr. 2002.
- [15] LYSYANSKAYA, A., RIVEST, R., SAHAI, A., AND WOLF, S. Pseudonym systems. In *Selected Areas in Cryptography* (1999), H. Heys and C. Adams, Eds., vol. 1758 of *LNCN*, Springer Verlag.
- [16] MANOLA, F., AND MILLER, E. RDF primer. W3C Recommendation.
- [17] REITER, M. K., AND RUBIN, A. D. Crowds: anonymity for web transactions. *ACM Trans. Inf. Syst. Secur.* 1, 1 (1998), 66–92.
- [18] SAITA, A. Cardsystems admits stolen data violated policy. SearchSecurity.com, 21 June 2005. http://searchsecurity.techtarget.com/originalContent/0,289142,sid14_gci1099932,00.html.
- [19] TUROW, J., FELDMAN, L., AND MELTZER, K. Open to exploitation: American shoppers online and offline. Tech. rep., Annenberg Public Policy Center, University of Pennsylvania, June 2005. http://www.annenbergpublicpolicycenter.org/04_info_society/Turow_APPC_Report_WEB_FINAL.pdf.