Privacy and Identity Management for Everyone*

Jan Camenisch† abhi shelat† Dieter Sommer†
Simone Fischer-Hübner‡ Marit Hansen§ Henry Krasemann§ Gérard Lacoste¶
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 [14] 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 that 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 cryptographic protocols, sophisticated security protocols, and artificial intelligence algorithms. This paper describes the architecture of this system. Most key features of this architecture have been implemented in a proof-of-concept prototype.

†IBM Research, Zurich Research Lab, Säumerstrasse 4, CH-8803 Rüschlikon, Switzerland, jca,abs,dso@zurich.ibm.com
‡Karlstads Universitet, Sweden
§Unabhängiges Landeszentrum für Datenschutz,Germany
¶Compagnie IBM France, France
‖Universiteit van Tilburg, The Netherlands
**Erasmus Universiteit Rotterdam, The Netherlands

*Part of the work reported in this paper is supported by the European Commission through the IST Project PRIME. The PRIME project receives research funding from the European Community's Sixth Framework Programme and the Swiss Federal Office for Education and Science.

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

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 [20], 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]

[1]Specifically, 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.

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 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. [19]

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 because current tools require too much effort from the user. Moreover, the lack of standardization and interoperability typically renders 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 information 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 service provider'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, a system architecture has been proposed and a first prototype of the architecture has 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. 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. 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 amount of data needed to fulfull a transaction certainly depends on the business process. For example, it might be required that a user prove that she is at least 18 years old. This could be performed equally well with an ID card, passport, or—in some countries—with a driver's license *credential*. As a *side effect* of showing such a credential, however, additional data about the user (which is unnecessary for the transaction) is inadvertantly disclosed to the provider. To satisfy data minimization, a better approach would be to only prove that the user possesses a credential with an appropriate birth date attribute without revealing any other details. In Section 3, we describe how PRIME can achieve this property.

Identity Management. A user may also wish to release different amounts of personal information depending on the trustworthiness of the transaction partner. This "user-controlled identity management" is performed intuitively by everybody in the physical 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 in an automated way. 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. PRIME allows a user to control, record, and manage all the information which has been revealed to the various service providers.

PRIME also addresses services-side identity management, which is principly concerned with authorization. A service provider assigns each of its users a unique identity. Later, when a user provides

this identity, for example, by entering her user name and password, the user is authorized to access the provider's services or resources. PRIME particularly allows for services-side identity management in which users can remain anonymous while simultaneously proving that they have been authorized to use the requested service.

Because the PRIME architecture allows for both user-side and server-side identity management in both anonymous and identified cases, PRIME software can work in today's prevailing identity-based scenarios and at the same time, work in future scenarios where the data minimization principle is applied. This compatibility suggests a good migration path for PRIME-enabled transactions on the Internet.

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.

Accountability. Let us reiterate that properly-designed anonymous transactions can also provide accountability—in other words, a user can be made accountable for misuse of the system or cheating, even though transactions are "anonymous." One way to do so is for the service provider to request from the user a verifiable encryption, in the public key of a mutually trusted third-party, of a credential which contains the address, phone, and identity of the user. Because the encryption is verifiable, the service provider can be assured that the encryption contains a certified identity (vouched for by a separate certification authority). Under normal circumstances, the encrypted identity does not leak any identity information of the user, and anonymity is preserved. If a dispute arises, however, and certain well-defined conditions are met, the service provider may request that the trusted third-party decrypt 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

User. A user in the system has certificates, data and policies regarding their data. Access control policies restrict the access to the user's data for release of the data to transaction partners and privacy policies define how the user wants a potential receiver of her data to handle the 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 that they have received from business partners. A service provider has privacy policies defining how user's data will be handled.

Certification Authority. A special type of service provider is a certifying authority that issues *certificates*, that is, digitally-signed statements. By issuing a certificate, a certificate authority vouches for the truthfullness of the statement. (A reputable certificate authority, therefore, has a clear incentive to verify the statements before issuing a certificate asserting them.) In this paper we use the term credential synonymously with certificate.

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 to the user by means of standard technology (e.g., an X.509 server-side certificate [15]) and the user remains unauthenticated at that point. The TLS protocol provides such a channel.

Anonymous Communication. Traditionally, a user who connects to a service provider's computer system implicitly reveals network information, such as an IP address or MAC address, which can be used to identify the user and link all of the user's transactions to one another. Therefore, special network precautions must be taken for there to be any hope for PRIME to achieve its goals. The solution is for the user to employ an anoymizing network, which is provided for example by onion routing networks [13], mixnets [10], or crowds [18].

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 [6], require cryptographic protocols for their establishment and use, 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, for example, by requiring the owner's secret key to use the credential [6]. This is important for ensuring the accountability of anonymous transactions and to prevent users from sharing their credentials. Binding credentials to hardware is another option to prevent sharing [5].

From a privacy perspective, the use of credentials is preferable to making a direct request to the issuer of the credential because

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

in the later case, the issuer can profile the user by recording *who* makes queries about him. Credentials can either be realized using traditional attribute certificates (see for example [12]), where the reference to the user could be the user's real name or a pseudonym, or by so-called *private credentials* (see for example [10, 16, 6]).

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 using the birth date attribute of the credential, without revealing any of name, address, or birth date, nor making this transaction linkable to any other transaction.

As mentioned in the Accountability section, private credentials allow the user to verifiably encrypt an attribute under a third-party public key. For instance, consider a credit-card number which has been certified by a bank. If the user provides—in addition to the encryption of the credit card number—a proof that the encryption contains a number certified by the bank, the user could assure a merchant that a payment will be approved without the merchant learning the credit card number.

Furthermore, the user could also "cryptographically bind" the encryption to a condition which specifies the circumstances under which the ciphertext can be decrypted. In such a case, the third party would only be able to decrypt the verifiable encryption (and thereby reveal the identity of the user) if the merchant provides a certain value—for example, proof that the user has cheated—to the third party.

Thus, as mentioned before, private credentials together with encryption of attributes enable transactions to be privacy-protecting and yet accountable. See [7] for details about a framework for private credentials and [1] for how to enable access control for the private credential framework.

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 our way of identifying resources, the data model of the architecture, and how we use ontologies; all these concepts apply throughout the architecture.

Resource Referencing Scheme. To be meaningful, an access control system must have a well-defined mechanism for "naming" resources. For this purpose, we use the Uniform Resource Identifier (URI) scheme proposed in RFC 3986 [2] to name every resource in our system.

The use of URIs is well-established practice on the Internet, in particular the URL of a website is a special kind of URI. In addition, URIs are general enough to name data types, services, process workflows, or obligations such as "Delete this data after two weeks." In our context, we can use a URI to refer to a user's personal information without revealing any information about the user or the information—in other words, our system can process "links to user information" in place of the actual information.

Data Model and Ontology. In order to account for interoperability between two parties and the various components of our system, a *data model* and *ontology* are required.

We selected the Resource Description Framework (RDF) of the W3C [17] as the language for describing information about the resources in our system. The RDF language consists of triples of the form *(subject, predicate, object)* which represent the fact that *subject* is related to the *object* by the relationship identified by the *predicate*. If all three values in the triple are URIs, then the statement is well-defined.

In order for RDF statements to be globally understood by different systems, however, there must be common language and a well-defined semantics for *subject*s, *object*s, and *predicate*s. This is especially true, for example, if we want our system to "reason" about RDF statements such as privacy policies and credentials. For example, we would like to automatically determine all of the credentials which can be used to prove—say—one's age, and determine which of them reveals the "least" additional information. In order to support such features, we use the Web Ontology Language (OWL) [11] proposed by the W3C to describe all of the meta-information about *subjects*, *predicates* and *objects*. OWL was expressely "designed for use by applications that need to process the content of information instead of just presenting information to humans" and its expressivity suffices for our needs.

3.3.1 Components

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 party). Furthermore, the database contains default policies for the release of information, the policies for handling of data, 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 [9, 8].

Access Control. The Access Control component limits access to a party's resources and enforces the party's access control policies. 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 mainly consists of a URI identifying the resource to be accessed, a purpose for the access, an operation on the resource (e.g. read, or update, or delete), 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 paradigm of *attribute-based* access control [3, 4] in which access is granted based on the *properties* that the requester has asserted via the auxiliary information. Thus, our system explicitly avoids—when possible from a business process point of view—relying on the identity of the requester to determine how to respond to the request.

As a concrete example, a traditional driver's license credential might be used to establish one's age, and therefore one's legal right to purchase alcohol. Such 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 private credential such as a driver's license certifying that the presenter is of legal age would suffice for the transaction. It would not release more information than that the user is of legal age and possesses a driver's license.

The attribute-based access control paradigm is particularly valuable in open environments such as the Internet. In such environments identity data is difficult to meaningfully interpret as there is no preexisting relationship which would allow that authorizations be granted based on the identity of the requester. 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 (private) 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 specific workflow process defined by the obligation whenever appropriate triggering events occur and 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." A time-based event will trigger the workflow for this obligation.

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 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, such as a phone number, has been requested.

   The offer also specifies how the data related to this transaction will be treated. This is done by expressing the service provider's privacy policy for the data categories being requested. In particular, the service provider presents obligations to the user that will be automatically enforced.

   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 requested 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 the possible 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 otherwise use a credit card). The agreed privacy policy and obligations are presented to the user in an easy-to-understand representation. The user finally has to give their informed consent to the data processing.

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 interactive protocols in which credentials are shown or simple transmissions of uncertified 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 specific parts of the user data under an access control policy that enforces the agreed privacy policy and to store the related obligations in the OM. The OM activates each obligation meaning that they can noe be triggered by appropriate events and conditions. 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 either be completely orthogonal to any services-side accesses to the user's data (e.g., time-driven deletion) or can be related to such accesses as in the example.

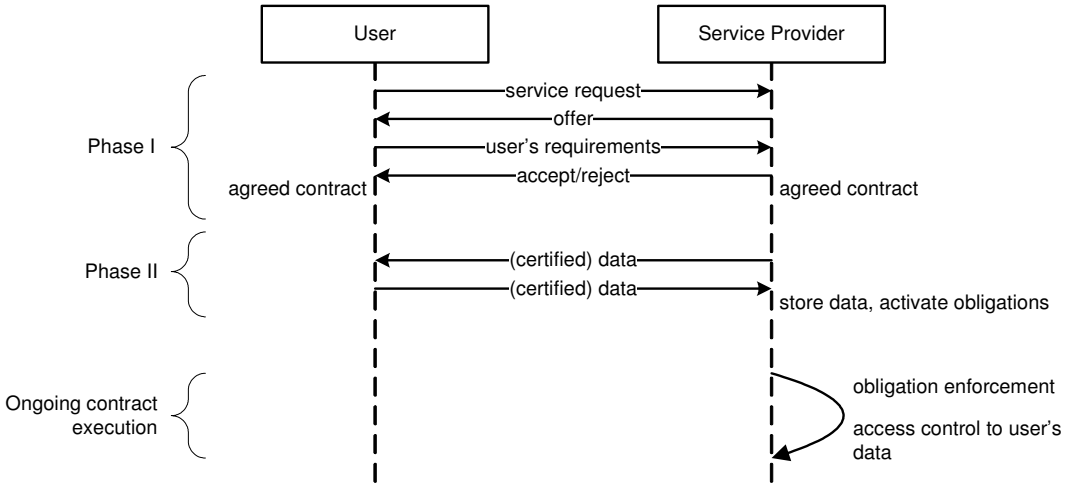

Figure 1: Execution of a transaction

The two-phase transaction outlined above is the standard case case of a transaction that applies for many scenarios. 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, but typically obligations for correctly handling the user's data disclosed in this transaction.

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 are designed not to reveal any business secrets regarding the service provider's system configuration.

Reasoner. How does a user parse a computer-generated list of requirements and choose the set of credentials and perhaps a psuedonym that satisfy them? The Reasoner component addresses this issue by providing automated reasoning over ontologies (see §3).

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 by triggering appropriate workflows.

5. DISCUSSION

PRIME Project. The PRIME project is a European research and development project partially funded by the European Union. The project consists of more than 20 project partners mainly from Europe, comprising universities, public companies, a consumer protection authority, and standardization bodies. As the privacy problem is not only limited to technical aspects, but also driven by legal, economic, and social aspects, the project takes an interdisciplinary approach to the project. Thus there are partners from the technical, legal, economic, and social sciences arenas. The duration of the project is 4 years.

Prototype. The project has developed several prototype implementations of the privacy architecture outlined in this paper. In the first phase of the implementation process, separate prototypes for each of the mechanisms that we employ, among them, attribute access control, private credentials, and automated reasoning, were written. The second round yielded an integrated prototype that integrates the key components of our architecture. This is the first prototype implementation of a comprehensive privacy architecture. The current prototype is still missing some features, such as the negotiation of privacy policies and a complete GUI. The next round of the development process will address these points.

Contribution. The PRIME project tackles the privacy problem in a comprehensive fashion. In particular, our architecture takes a

whole-system approach to the problem and improves on other proposals that are typically limited in scope. For the first time privacy-enhancing technologies from all these different fields have been absorbed into one architecture with the goal of providing better privacy to end users while simplifying privacy management for companies thereby decreasing their expenses.

As a key feature, interactions are to a large extent policy driven. This makes the architecture highly flexible and minimizes the amount of required user involvement. Following the data minimization principle, transactions can be designed in a way that a minimal amount of information can be released by users in order to get access to services. This becomes possible by employing private credential protocols. The advanced protocols allow a user to verifiably encrypt attributes. The user is supported in their decision by automated reasoning and by a simple GUI that makes the software easy to use for average users.

Migration Path. The migration to PRIME technology on the Internet is not a trivial process, and as with all large-scale deployments, requires a well-defined plan . The client software must function correctly with legacy server-side software and guarantee an excellent browsing experience. At the server side, a gradual deployment of the technology provides the best approach for the deployment. This can start with privacy policy negotiation in a first step. As a next step, attribute-based access control and credential systems could be deployed in order to allow for data minimization. This requires that a PKI featuring private credentials be put in place beforehand.

6. 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 as key elements an anonymous credential system, an attribute-based access control system, a policy compliance checking functionality, a negotiation and orchestration 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 certified 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 a single trusted party to revoke the anonymity of a transaction, such power can be distributed among many external parties by employing standard techniques from threshold cryptography. This helps reducing the stong trust assumptions in a single party.

Because our system also allows a user to assess the trustworthiness of a service provider, it will be easier for smaller 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 of contracts 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, and (c) cost reductions from automated privacy handling.

DISCLAIMER The information in this document reflects only the author's views, 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.

7. REFERENCES

[1] BACKES, M., CAMENISCH, J., AND SOMMER, D. Anonymous yet accountable access control. In *Proceedings of the Workshop on Privacy in the Electronic Society 2005* (2005).

[2] BERNERS-LEE, T., FIELDING, R., AND MASINTER, L. Uniform Resource Identifier (URI): Generic Syntax. RFC 3986 (Standard), Jan. 2005.

[3] 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.

[4] 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.

[5] CAMENISCH, J. Protecting (anonymous) credentials with the trusted computing group's trusted platform modules v1.2. Tech. rep., IBM Research, Jan. 2005.

[6] 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.

[7] CAMENISCH, J., SOMMER, D., AND ZIMMERMANN, R. A general certification framework with applications to privacy-enhancing certificate infrastructures. Tech. Rep. RZ 3629, IBM Zurich Research Laboratory, July 2005.

[8] CASASSA MONT, M. Dealing with privacy obligations: Important aspects and technical approaches. In *TrustBus 2004* (2004), pp. 120–131.

[9] CASASSA MONT, M. Dealing with privacy obligations in enterprises. In *ISSE* (2004).

[10] CHAUM, D. Security without identification: Transaction systems to make big brother obsolete. *Communications of the ACM 28*, 10 (Oct. 1985), 1030–1044.

[11] DEAN, M., AND SCHREIBER, G. OWL web ontology language reference. W3C Recommendation.

[12] FARRELL, S., AND HOUSLEY, R. An Internet Attribute Certificate Profile for Authorization. RFC 3281 (Proposed Standard), Apr. 2002.

[13] 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.

[14] HANSEN, M., AND KRASEMANN, H. Prime whitepaper. Whitepaper, 18 July 2005. http://www.prime-project.eu.org/prime/public/press_room/whitepaper/PRIME-Whitepaper-V1.pdf.

[15] 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.

[16] 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 *LNCS*, Springer Verlag.

[17] MANOLA, F., AND MILLER, E. RDF primer. W3C Recommendation.

[18] REITER, M. K., AND RUBIN, A. D. Crowds: anonymity for web transactions. *ACM Trans. Inf. Syst. Secur. 1*, 1 (1998), 66–92.

[19] SAITA, A. Cardsystems admits stolen data violated policy. SearchSecurity.com, 21 June 2005. http://searchsecurity.techtarget.com/originalContent/0,289142,sid14_gci1099932,00.html.

[20] 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.