PRIME white paper

Third and final version,

15 May 2008

Ronald Leenes

Jan Schallaböck
Marit Hansen

welcome

privacy-enhancing identity management

The PRIME project demonstrates the viability of privacy-enhancing identity management. By this we mean identity management solutions that manage the individual's identity online and that also empower the individual to actively protect their own privacy.

The guiding principle in the PRIME project is to put individuals in control of their personal data. The notion of user control has been adopted in many recent user-centric identity management initiatives.

However, most of these initiatives only takes the first steps on the way to a new generation of identity management systems. They do not provide adequate safeguards for personal data and are limited in giving individuals control over their personal data. Effective management of information privacy requires new tools starting with the minimisation of personal data disclosure. Furthermore, users can be empowered with tools that allow them to negotiate privacy policies with service providers. This would require systems that enforce agreed policies by technical means, and keep track of data collection and usage. In addition to user side applications, service providers will be required to put adequate protection mechanisms in place and align business processes to take advantage of these mechanisms.

This white paper describes our vision of privacy-enhancing identity management and how it can be realised in software. It also shows where work remains to be done.

Published PRIME documents

These documents are all available from the project website located at http://www.prime-project.eu/

Excerpt of project "Description of work"	03-2004	Framework Version 2	07-2006
Project presentation	09-2004	Requirements Version 1	05-2005
Overview of existing assurance methods	09-2004	White Paper Version 1	07-2005
Evaluation of early prototypes	12-2004	White Paper Version 2	07-2007
Evaluation of Integrated Prototype	10-2005	Tutorials Version 1	06-2005
Evaluation of initial Application Prototypes	03-2006	General Public Tutorial	03-2006
Evaluation of Integrated Prototype V2	05-2007	Advanced Tutorial	02-2007
HCI guidance and proposals	02-2005	Architecture Version 1	08-2005
Framework Version 1	03-2005	Architecture Version 2	03-2007

table of contents

your notes

online ID reconsidered

motivation

The internet, by design, lacks unified provisions for identifying who communicates with whom; it lacks a well-designed identity infrastructure.[1] Instead, technology designers, enterprises, governments and individuals have over time developed a bricolage of isolated, incompatible, partial solutions to meet their needs in communications and transactions. The overall result of these unguided developments is that enterprises and governments cannot easily identify their communication partners at the individual level. Given the lack of a proper identity infrastructure, individuals often have to disclose more personal data than strictly required. In addition to name and address contact details such as multiple phone numbers (home, work, mobile) and e-mail addresses are requested. The amount and nature of the data disclosed exceeds that usually required of real world transactions, which can often be conducted anonymously – in many cases the service could be provided without any personal data at all. Over the long run, the inadequacy of the identity infrastructure, that takes the above into account, affects individuals' privacy. The availability of abundant personal data to enterprises and governments has a profound impact on the individual's right to be let alone as well as on society at large.

The online world is a complex new environment. Social structures online have to be established within a short time - very much unlike their real world counterparts. At first glance those procedures based on personal contact or paper are transformed into digital procedures for use online. But below the surface, more fundamental differences between the offline and the online world exist, such as the relative permanence of memories and the ease with which experiences can be shared between many of actors across time and space barriers.

We are beginning to understand that these differences are both qualitative (e.g., automated decision making) and quantitative (e.g., more data collected and stored for a longer period) in nature. The speed of developments and potential irreversibility of their effects requires urgent attention on issues such as identity, trust, security, and privacy.

The – sometimes conflicting – interests and issues that have to be reconciled are increasingly well understood. For example for such a conflict is an interest in identifying trading parties on one hand and providing anonymity on the other. The convenience of 'portable' online identities is another example; users do not want to fill in similar forms for each service, yet there is the risk of disclosing more than is required. National security interests – sometimes positioned as overriding civil liberties in public debates – increases the need for proper data protection. And finally, while customer data is an important business asset, they can become a business liability in complying with data protection legislation.[2]

identity management (IDM) ↻
Systems and processes that manage and control who has access to resources, and what each user is entitled to do with those resources, in compliance with the organisations' policies.

---

[1] The Internet has a ID infrastructure often identifying only the endpoint of a communication: IP addresses. These are often unreliable to identify users.

[2] This is particularly so in several States in the US, where security and privacy breaches have to be reported which could subsequently result in liability cases and damage of reputation.

enterprise IDM ↻

Enterprise identity management is understood as idm that primarily serves the enterprise's needs. Control is exercised by the enterprise instead of by the individual (see in contrast user-centric IDM).

user-centric IDM ↻

Different communities agree on rough definitions of the concept, yet a common precise definition is lacking. In this paper it is idm that seeks to place administration and control of identity information directly into the hands of individuals.

federated IDM ↻

A relationship that allows the authentication of an entity verified by one identity authority to be recognised by other identity authorities in the federation [ISO/IEC 2005].

Federated IDM enables single sign-on.

single sign-on ↻

A form of software authentication that enables a user to authenticate once and gain access to the resources of multiple software systems.

Online **identity management** ↻ is in need of reconsideration. The patchwork approach to online identity needs to make way for a more elaborate design that takes into account the various stakes and issues. Indeed the identity management landscape appears to be changing. **Enterprise identity management** ↻ is slowly making way for **user-centric identity management** ↻. Various initiatives, such as the Liberty Alliance project and WS-Federation, aim to pave the way for identity management that 'involves the users in the management of their personal information and how that information is used, rather than to presume that an enterprise or commercial entity holds all the data.' [Liberty Alliance 2006]. Establishing authenticated individual identities within and across organisational boundaries are the primary business drivers behind these initiatives. Their successful adoption depends on improved privacy protection. User control and other elements of privacy protection also gain attention in a broader sense. The '7 laws of identity' [Cameron 2005] initiated by Microsoft's Kim Cameron clearly attracted attention in the identity community.

What these developments show is that industry is adopting the notion of user control over personal data. But so far the interests of the service providers are better served than those of the individuals. In the wake of what is coined Web 2.0, where consumers merge into prosumers, services replace applications, data increasingly drives economic activity, and where generally the landscape becomes more dynamic, this will not do. Individuals will feel a stronger desire for privacy and control over what's known about them. They also require more security, which demands stronger and better authentication and identification, which in turn requires even better privacy protection.

The PRIME project aims to show that seemingly disparate notions such as anonymity and accountability, security and privacy, and informational self determination and enterprise needs can be reconciled. PRIME intends to set the boundaries for the emerging identity management infrastructure with a clear balance of the interests of users, enterprises and society.

document purpose and scope

This white paper provides an overview of the PRIME consortium's vision on privacy-enhancing identity management and illustrates a number of high level requirements incorporated in the PRIME software architecture. The document is intended for the architects of the new identity infrastructures, including information architects, systems designers, corporate decision makers and policy makers.

Numerous identity management white papers have been written by technology solutions vendors from the perspective of the enterprise. These white papers target corporate decision makers' concerns about financial compliance and return on investment.

The PRIME project takes the perspective of the individual and places the individual at the core of user-centric privacy-enhancing identity management. This leads to a different, but not incompatible, set of requirements. The requirements elicited in this document have their roots in the OECD Privacy Guidelines [OECD 1980], the Council of Europe's Convetion No. 108 [CoE 1981], the Fair Information Practices (for instance embedded in the CSA privacy code [CSA 1996]), the EU Data Protection Directive (Directive 95/46 EC) and recent discussions on user-centric identity management. Many of these requirements are discussed in Kim Cameron's '7 Laws of Identity' [Cameron 2005] and the Ontario Information and Privacy Officer's white paper on identity management [Cavoukian 2006].

a scenario

Alice goes shopping

The requirements elicited in the following pages may appear to be ambitious, but software prototypes that demonstrates these features have been developed and evaluated within the PRIME project.

Before looking at them in more detail and describing the PRIME approach to address them, we will first take a walkthrough current practice and the problems it entails in a typical online shopping scenario. The purpose here is to showcase the software architecture required for enabling privacy-enhancing identity management to those organizations interested in deploying these features.

Figure 1 illustrates the exchange of personal data in a typical online shopping scenario today. Alice asks her sister Alicia, whom she dearly trusts, for advice on white wine. On the basis of her sister's recommendation, she orders a box of bottles of Chardonnay at CyberWinery. To this end, Alice has to provide personal data (name, delivery address, and possibly payment data, such as her credit card data). If this is her first and only order, the winery will store only some of these data in their records. But more likely, it will ask Alice to register, arguing that this will make it easier for her to make additional purchases. If she does, Alice will have an account at the winery which not only contains her name and address, but also her purchase history, personal preferences, and likely also her credit card data.

Suppose the winery has outsourced warehousing and delivery to LogisticsProvider, a major logistics company. LogisticsProvider needs to have some of Alice's personal data – name and shipping address – to deliver her order. CyberWinery checks Alice's credit card details at CreditProcessor for credit authorisation. If the order is accepted, CreditProcessor also takes care of processing the payment. Again, Alice's personal data are exchanged between two businesses. CreditProcessor will store transaction details in their records for business and accountancy purposes.

Suppose Alice also takes up Alicia's recommendation to purchase the Ultimate Wine Guide at CyberBooks, *the* online bookstore for Gourmet
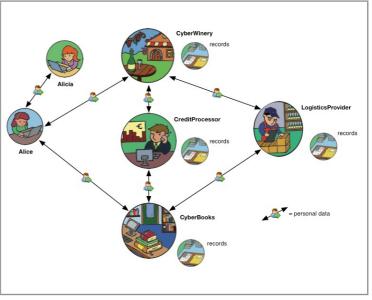

Figure 1: user data exchange

issues involved

web forms: Each service provider uses its own registration and data entry forms. Form fillers can provide some support, but overall the process of registering at different service providers is tedious and redundant.

consent: The user has to provide the data asked for (or fake it) and has no choice but to accept the privacy conditions set by the service provider if she wants to enter into a contract with this provider. Does this 'take it or leave it' approach really reflect the idea of user consent?

excessive data disclosure: More data than required for the transaction itself are often collected for the sake of establishing trust in the user, and for being able to address her for different purposes, such as direct marketing or safety alerts. But is the service provider really ever going to call the customer if they also have her e-mail address?

assurances: The user has to provide data that, when leaked to unauthorised persons, can have serious implications. How can the user, apart from the lock icon in her web browser, be sure the systems and communication are trustworthy? How can she be sure the provided SSL credentials are trustworthy?

transaction history: Users currently do not have reliable digital receipts for online purchases and other (trans)actions. This makes it difficult for them to dispute errors made by service providers and exonerate themselves from blame in the case of ID fraud.

enforcement: Current privacy policies give users limited rights, for instance for opting in/out of certain secondary uses of their personal data. Effective enforcement of these terms is lacking. Users cannot check whether the service provider adheres to the policy agreed upon, nor can they see whether legal obligations are observed. It is even more difficult to verify whether user data has been shared with associates.

user support: The user has to cope with problems mentioned above on her own. She has to devise ways of managing her digital identities, assess privacy policies, assess the trustworthiness of interaction parties and keep track of transaction histories using applications that offer very limited support for these mundane tasks.

proprietary protocols: Data exchange between enterprises is often done using proprietary protocols, although (XML) standardisation is under way. This makes the exchange of data between enterprises, e.g., for single sign-on, cumbersome to set up. This situation becomes more aggravated when non standard data are being used, such as encrypted data and credentials.

books. She again has to register before being able to order, and she basically has to provide the same information she provided to CyberWinery. Consequently, the CyberWinery scenario unfolds again, most likely involving CreditProcessor, and possibly even involving LogisticsProvider as well.

The scenario sketched encompasses many exchanges involving personal data between user and service provider and between service provider and their associates. Many of these data are stored in multiple databases. Some providers can make interesting inferences on the basis of the data they have. CreditProcessor, for instance, knows where Alice does her shopping and the amount she spends, whereas LogisticsProvider even knows what she buys and where. CyberBooks has a much dimmer picture of Alice's shopping habits, they only get to see what they contribute to Alice's collection of cook books.

Overall, this scenario illustrates a number of issues from a user's perspective, especially if she wants to minimise the risk that her data may be abused, for instance for ID fraud, or for profiling and **social sorting** ↻.
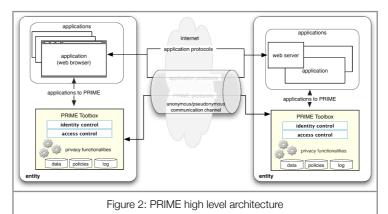
Many of these problems can be addressed by means of novel identity management systems. In this paper we discuss various problems and describe the way the PRIME project aims to resolve them by offering a privacy-enhancing identity infrastructure. We will use Alice's online shopping scenario to unravel the problems and formulate a list of requirements on our way.

social sorting ↻
Classifying and profiling groups of people in order to provide different services, conditions or treatment.

PRIME toolbox ↻
The PRIME toolbox is a collection of concepts and application modules developed in the PRIME project which support or implement user-centric privacy and identity management functionality.

PRIME Middleware ↻
The PRIME Middleware is the software which glues the components together. It provides interfaces for external and legacy applications to use the PRIME architecture.

PRIME enabled shopping

The aim of the PRIME project is to provide privacy-enhancing identity management tools for individuals. PRIME empowers the user by offering them more extensive **(user) control** ✪ over their personal data. The **PRIME toolbox** ↻ offers support for creating, using and keeping track of multiple digital identities and the (certified) attributes associated with them. It allows (certified) attributes to be transferred between entities, such as user and service provider, or between service providers. It also extends the user's control over attributes disclosed to remote entities. The PRIME vision is based on the principle of data minimisation, i.e., disclosing and processing personal data only to the extent necessary. To limit the transfer of personal data for authentication purposes, claims and credentials are used to establish trustworthy relationships. Where necessary, for instance for certifying certain user attributes, PRIME makes use of privacy-enhancing public key infrastructures and trusted third parties. The integrity of claims in PRIME enabled communication is guaranteed by cryptographic techniques. Each party in the interaction makes use of **PRIME Middleware** ↻. The individual users additionally use the **PRIME Console** ↻ to manage their personal data. User applications (such as web browsers) may delegate identity man-

Figure 2: PRIME high level architecture

agement tasks to the PRIME Console and PRIME Middleware. The trustworthiness of the PRIME components should be maximised by technical means (e.g., cryptographic techniques) and non technical means (e.g., certification and assurance).

We will now explore Alice's ventures in the online wine business by going through six phases to illustrate online transactions from before entering the internet to becoming a frequent shopper and beyond.

phase 1: buyer beware

Transactions require a certain level of mutual trust between transaction partners. Each party has to be confident that the other will perform their contractual obligations, will not abuse one's vulnerabilities, and that there are options for redress in case of breaches. In the offline world this confidence stems from factors such as the respect commanded by the brick and mortar that houses (commercial and governmental) institutions and from honourable social institutions such as the legal system that acts as a safety net in case of conflicts.

In the online world tangible signs of trustworthiness are absent to a large extent and therefore we have to rely on other signals for trustworthiness. It is relatively easy to create websites that resemble genuine ones. This method is therefore frequently employed by criminals for all kinds of fraud (including ID theft and phishing attacks). Although people may believe a certain website to be genuine, a reliable level of trustworthiness cannot be established. Assurance that a service provider is genuine and complies to regulations and policies can be provided by third parties (trust assurance), however, some users have difficulties in understanding the scope and value of these trust marks.

In addition, the communication channel needs to be trustworthy because communication can be intercepted, manipulated and suppressed. User provided data, such as credit card data can, when in the hands of the wrong people, have serious implications for the user. Integrity and confidentiality of both communication and data are therefore important requirements for online interactions. The users should substantially be able to trust the entire chain of entities involved in providing a service to be secure against intruders, eavesdroppers, etc - or better: not even need to rely on trust but stay in control. This calls for technical measures, such as, encrypted and properly authenticated communication.

Additionally non-technical measures can increase user confidence that their interactions are safe. Transparency, i.e., providing clearly understandable information to the user on the data processing, is an example. Online processes – like shopping or simply gathering information – are currently rarely transparent and many users do not feel comfortable because of the technology involved in the interaction. Prospective customers often even have to guess or do not understand what the shopping process will look like when engaging in it.[3] Improving the transparency of the processes and making clear why personal data are being collected and what happens with the data makes users feel more comfortable in online interactions and helps to build their trust.

Based on her sister's recommendation, Alice decides that it may be worthwhile looking for wine at the CyberWinery. The store implements a number of measures that reassure Alice of its trustworthiness. CyberWinery's home page shows a trust mark she is familiar with and that she considers trustworthy. The shop also turns out to be PRIME enabled, which means that she knows how the communication will work because it is well documented and she has experience with it. For the shop having a PRIME enabled customer means that it is able to check the validity of certain credentials provided by this customer by means of trusted third parties (see phase 3).

---

[3]   Despite legal requirements (e.g., the e-Commerce Directive 2000/31/EC), many online shops still do not offer clear documentation of the shopping procedures.

PRIME Console ↺
The PRIME Console is the interface to the user's identity management system. It allows users to create partial identities (pseudonyms), associate personal data to these identities, assists the user in understanding privacy policies, decides on the basis of the user's preferences and allows the user to inspect the transaction history.

data minimisation ↺
From the perspective of the individual, this means disclosing as little personal data as possible, from the service provider's point of view it means only collecting personal data that are necessary for the purpose of the interaction or storing data only in anonymised form. See also the (legal) requirement on the next page.

legal requirement: ⊛
justifiable parties
Personal data should only be accessible to entities with a legitimate interest in the data, e.g., by consent of the individual, by legal obligation or for other legitimate purposes. Service providers should implement technical measures to enforce this requirement, especially with respect to the use of personal data by third parties (for secondary uses). This requirement also implies that the user should be able to check the authenticity of the data requester.

phase 2: pre-sales – starting from maximum privacy

Alice's online interest in white wine does not appear to be particularly private or sensitive, when compared with her visits to health insurance websites or medical websites that might reveal information she wants to keep private. However, incorrect and damaging inferences may be drawn from Alice's wine interest when disclosed at the wrong place at the wrong time. Her search for wine during working hours may reflect a drinking habit, even though it just happens to be that she is organising a cocktail party for a colleague. Alice is sensitive about disclosing data that may lead to the wrong conclusions, about leaving online trails about her online transactions and she may even be worried about identity fraud due to recent newspaper reports. She guards her private sphere and wants to remain as unobserved as possible. She adheres to the **data minimisation** ↺ ⊛ principle and starts her online journeys from maximum privacy, choosing to disclose more personal details with her consent and according to the her own preferences.

Alice studies the company's general privacy policy. The shop has implemented the Article 29 working party's recommendation of layered policies [Art. 29 wp 2004]. The shop's home page shows the simple and short outline of the privacy policy and offers a click through to more detailed explanations of the company's policies. The privacy policy states CyberWinery's intentions regarding the protection of personal data. It assures the user that the data obtained by the store during browsing, purchasing and later on for delivery (see phase 3) will be handled as stated in the policy and will only be made available on a need to know basis. Alice is assured for now that the shop meets some basic requirement (see **justifiable parties** ⊛). The policy also states that the shop will allow her to opt out of their direct marketing programme at all times if she cares to join it. It also explains that her IP address is only recorded for statistical purposes, but not for profiling her behaviour.

IP addresses warrant caution because they are in many cases identifying data, albeit not very reliably.[4] They are like breadcrumbs left behind as a trail of the user facilitating linking her behaviour from one site to another. Due to the inadequacy of IP addresses as identifying data, they are sources of false conclusions about internet users. The principle of data minimisation can be applied to IP addresses as well. The shop should refrain from storing them unless there are legitimate reasons to store them. Alice can use an anonymising service, such as TOR or AN.ON, to hide her IP address from the webshop. This would reduce her concerns about leaving IP breadcrumbs. The PRIME Middleware provides interfaces to such anonymising services which makes it easier for the user to use these services.

phase 3: ordering – informed consent and purpose limitation

Autonomy as a central concept implies that individuals should make their own choices and only be bound to contracts they knowingly and voluntarily enter into.[5] As there usually is an asymmetry in both power and information to the detriment of the individual, it is reasonable to protect the individual in their relation with enterprises and governments. To this purpose regulation obliges service providers to state who they are and what their terms and conditions are, and what the effects of contracts they enter into are. This allows individuals to make informed choices and also provides them with information they need if they seek redress in case of contractual breaches, problems, and so forth.[6] The information requirements also apply to the collection and use of personal data because this affects the individual's privacy.

---

requirement: ⊛

user control and consent

In order to maintain the individuals' trust in the information society and guarantee their freedom of choice (autonomy), users should be able to control which personal data are given to whom and for what purpose.

Exercising control requires informed and uncoerced consent for specific uses, which may be revoked at a later date, by the individual.

legal requirement: ⊛

data minimisation

Personal data disclosure should be limited to adequate, relevant and non-excessive data. Implied in this requirement is that data needs to be provided on a need-to-know basis and stored on a need-to-retain basis. This requires the requester to specify the purposes of collection, processing and storing of the data. Data should be deleted at the requester's end as soon as the specified purposes of data collection are met.

When Alice decides to purchase a box of white wines she must disclose some personal data in order to complete the purchase order. To determine which data are reasonable to disclose, she has to dig deeper in the shop's general privacy policy requiring serious effort. She has to consider the information the shop is obliged to provide about the purpose of data collection, the duration the data are kept, etc. On the basis of this information, she may decide that, in her opinion, certain data is excessive and she may decide to proceed, not to proceed, or provide false data. Assessing privacy policies is not easy in current environments. Many general privacy policies state the website's policy in lengthy difficult language that appears to show that the website really has considered all the intricacies of online transactions rather than providing the customer with relevant information. They are generally not written with the average user in mind. Although the statement 'we will share your data with our business partners' in itself is clear, its scope is not. There is often clearly room for improvement.

Consent is understood by many service providers as a necessary requirement for entering into contracts, and for being allowed to collect and use personal data. It is usually implemented, if at all, by means of an 'I agree' button. The user has no choice but to accept the privacy conditions set by the service provider if she wants to enter into a contract.

policies
PRIME replaces the 'take it or leave it' approach to privacy policies by a system of policy negotiation. Both parties can express different kinds of policies relating to authorisations, data handling and preferences. The user is assisted (see **human measure** ✵) by the PRIME Console which helps in setting personal preferences and requirements, in converting preferences from machine readable form to human readable form and vice versa, and in automatically negotiating the user's preferences with the other party. It supports the notion of user roles that allow the user to define policy sets (and their associated personal data) for various frequent uses. The PRIME Console therefore allows the user to delegate reaching a policy agreement to a digital assistant for common interactions and assists the user in more complex interactions.

Alice, for instance, has a preference to reduce the chances of receiving unsolicited email. Therefore she wants to receive order confirmation through a temporary, 'disposable' mail address that retains mail only for one hour. Furthermore, she does not want to receive newsletters, unless the shop offers some kind of incentive after her initial refusal, in which case the PRIME Console has to consult her. She also does not want to have her data distributed to business affiliates.

When the user enters a PRIME enabled website, she can activate the PRIME Console to take over all interactions relating to privacy policies or personal data. User applications may delegate identity management to the PRIME Console which then replaces the traditional webforms by a unified interface to the user's identity management system.

The PRIME Console keeps track of personal data relating to the user, her (negotiated) policies and service customisations, as well as of data disclosure to PRIME enabled services. The Console therefore keeps track of the history of the user's interactions. It can also poll services to provide information about the use of the data (and further disclosure to other parties) by this service provider, as well as the state of policy enforcement because the policies are associated with the data (**sticky policies** ↻). This allows the user to maintain control over her own data and exercise her statutory rights[7] to be informed about the data controller's use of her data in a more effective way.

---

[7] As laid out in for instance the Data Protection Directive 95/46/EC.

requirement: ✵
human measure
The user should be able to understand how she can exercise control over her personal data. Communication should therefore be in plain language using understandable concepts. 'Thingification' should be used for necessary but complex notions, such as roles, rights and obligations (e.g., using business cards to represent data related to a role). Human-machine communication within and between contexts should be unambiguous offering situational normality and predictability. The interface should help to protect the user against identity attacks.

sticky policies ↻
Sticky policies are a way to cryptographically associate policies to encrypted (personal) data. These policies function as a gate keeper to the data. The data is only accessible when the stated policy is honoured.

pseudonyms ↻
A pseudonym is an identifier of a subject other than the subject's civil identity.

person pseudonym:
A substitute or alias for a data subject's civil identity (name).

pseudonyms

Data minimisation is furthermore facilitated by support for pseudonyms. In fact, anonymous, or pseudonymous interactions are the default within PRIME. In many cases a handle to the user (or pseudonym) known by both parties is sufficient for the interaction and for possible follow-up interactions. For instance returning customers can be recognised on the basis of the user's pseudonym, and also tailoring services to her needs and preferences is possible on the basis of a pseudonym. PRIME supports different forms of pseudonyms with different characteristics with respect to linkability between the pseudonyms.

claims

Using **pseudonyms** ↻ instead of civil identities in transactions makes it more difficult to validate **claims** ↻ or attributes.[8] Yet, claims play an important role in minimising data disclosure because often it is not the identity of the user that matters but rather some attribute. For instance, the fact that Alice is over 16 years of age allows her to purchase alcohol, not the fact that she is called Alice. The fact that she can make the warranted claim that payment is assured, such as providing valid, non-revoked, credit card details, should be sufficient reason for CyberWinery to authorise shipment for a box of wine.

Claims in the real world can be certified by third parties. The State, for instance, offer certificates that a certain individual has a certain date of birth and lives at a certain address (passport, ID card, or driver's license). Online certifiers can, by means of cryptographic techniques (security tokens), vouch for certain claims in a secure manner that cannot be tampered with. PRIME offers extensive support for certified claims as well as for the creation of **private credentials** ↻. Private credentials (or certificates) allow for releasing partial information contained in a master certificate, for example, that one is over 18 using the birth date attribute. In addition, it is possible, to provide encryptions of attributes of private certificates in the claim together with a proof that the encryptions actually contain the third-party-endorsed attribute values and not any values put there by the claimant. Alice uses such a private certificate to prove that she is over 18.

What data Alice discloses when ordering her box of white wine depends on her preferences. She may want to reveal her real identity to CyberWinery, but she can also opt for a pseudonym. In the latter case the remainder of the shopping process will be slightly more complex than in the traditional setting where providing name, address and credit card data are sufficient to complete the transaction. If the winery makes use of a delivery service there is no need for them to have her address for the purpose of delivery. Alice can provide CyberWinery with a security token that points to her account with the delivery service. Alternatively, she could send an encrypted token including her address to CyberWinery while only providing the delivery service with the decryption key to her address.

phase 4: after-sales and delivery – retaining control: policy enforcement

Some time after Alice placed her order she is not only curious to know when to expect her purchase, but she is equally eager to know what data CyberWinery actually stored about her. She even had second thoughts about the shop having information about her at all. However, because the PRIME Console created a transaction pseudonym for her, she has trouble remembering which pseudonym was used for the transaction.

This shows two core problems of (data protection in) the online world. The first is that (privacy savvy) netizens will accumulate many digital personae. They use avatars in online games and virtual realities, pseudonyms for other kinds of interactions and finally their civil identity for certain business.

---

relationship pseudonym: A pseudonym that is used in regard to a specific communication partner (e.g., distinct nicknames for different communication partners.

role pseudonym: A pseudonym that is chosen for the use in a specific role (e.g., patient or customer).

role-relationship pseudonym: A pseudonym that is used for a specific combination of a role and communication partner.

transaction pseudonym: A pseudonym that is used for a specific transaction, i.e., for each transaction, a different pseudonym is used.

claims ↻
A claim is a statement made by an entity (the claimant) about another entity (the claim's object) to an entity or set of entities (the claimant's addressee).

private credentials (e.g., Idemix) ↻
Private credentials are secondary credentials that are derived from a certificate issued on a different pseudonym of the same person. Multiple private certificates can be created from a single certificate that are neither linkable to each other nor to the issuance interaction in which the master certificate was obtained. See [Camenisch/Lysyanskaya 2002] for details.

---

[8] If I know your name, I can try to get data about you through all sorts of channels, which is much more difficult if I only know you by transaction pseudonym ghT57897.

Unless there is a way to keep track of what each of these partial identities has done online, privacy protection is difficult in practice. The second problem is the lack of control on information once it has been released. Unlike goods, data cannot be reclaimed without the possibility that a copy is left behind in several possible places. This makes erasing traces hard, unless technology is brought to bear.

PRIME supports the user in staying in control of her partial identities, also after data disclosure. It offers support for managing the (possibly) multiple pseudonyms that make up a partial identity and the revealed (certified) attributes of the user under these pseudonyms. It provides the user with three central means to accomplish this: tracking one's data trail, support for rights enforcement and policy enforcement.

The PRIME Console's DataTrack function maintains a database of the personal data disclosed by the user. It provides a comprehensive overview of what personal data the user has released to whom, under which partial identity (pseudonym), when, and for what purpose (i.e. under what policy). The DataTrack therefore is an essential tool to keep track of one's digital personae.

The DataTrack also assists the user in enforcing her rights under the Data Protection Directive, for instance the right to get information about the data the service provider has about her, the right to correction and erasure. This functionality requires the implementation of PRIME Middleware at the user's side and the server's side. In cases of non-PRIME compliant service providers, the Data-Track will provide the user with hints on how to correctly enforce her rights using legal means.

The most powerful function of the PRIME concept is the technical **enforcement of agreed policies** ⊛ on the service's side when equipped with PRIME enabled Middleware. The machine-readable part of the sticky policies can be processed automatically by the PRIME server Middleware. The system will detect the fulfilment of certain conditions that warrant action on the user's data. For instance, it may detect certain purposes of data collection having been fulfilled, e.g., the order was shipped and hence retaining the shipping address is no longer necessary. In line with the principle of data minimisation it will then be deleted. Or, if the user allows the service provider to store her home address for 6 months for personal offers, the expiry date is attached to the address. The server side PRIME Middleware will then automatically delete the home address at the due date. Ideally, the user's increased control over the data disclosure should lead to the disclosure of less personal data, but better quality data. As a side effect, certainty over policy enforcement may increase the chances of the data being accurately provided instead of being fabricated. This not only is beneficial for the user, but also for the service provider. Automated policy enforcement is also advantageous for service providers because it facilitates compliance with internal policies as well as legal regulations.

phase 5: customer relationship – building the relationship

The quality of the CyberWinery's dry white wine appeals to Alice's taste and she returns to the shop to try out some of their red wines. She becomes a returning customer and before she realises it, she is a frequent customer (being the one with a big house, she hosts many family parties). Alicia's expertise as a wine buff turns out to be limited to white wine, so Alice decides that she may need the shop's recommendations on red and sparkling wines. She might also be interested in getting recommendations based on her previous purchases, similar to recommendations given at Amazon when accessing the site as a frequent customer. Both CyberWinery and Alice may benefit from this. Provided that Alice consented to such a service, CyberWinery could provide it. The PRIME Console facilitates the means to opt-in and opt-out of such a recommendation service at will.

> **requirement:** ⊛
> *policies and policy enforcement*
> Users should be able to express their privacy policies and preferences and negotiate the terms of data disclosure with service providers. The agreed upon policies should be strongly enforced by the identity management systems on both sides of the transaction.

She may do so if she is concerned about the store's ability to build detailed profiles about her, or even combine their data with those of other service providers to create a comprehensive picture of their customers' tastes, budgets and more. Although CyberWinery's recommendations may benefit from such detailed profiles, Alice wants to remain in control.

This desire to benefit from the advice provided by a service provider who is familiar with one's personal history on the one hand, and to remain relatively unknown on the other, leads to identity management issues. PRIME can help to address these. PRIME allows for a reduction of linkability of personal data if the user adopts different kinds of pseudonyms during the interactions. Alice can enter the store and identify herself with a role-relationship pseudonym for browsing and choosing items at CyberWinery that allows the shop to build a 'shopping' history for this pseudonym that is unlinkable to her real identity. Only when she decides to order, she switches to a transaction pseudonym that is only maintained for this specific transaction and is unlinkable to her role-relationship pseudonym. CyberWinery will retain the data associated to Alice's role-relationship pseudonym for further interactions. This does require a certain infrastructure to be in place that allows for a seamless identity switch at Alice's end – items placed in her shopping basket while browsing under her role-relationship pseudonym should be transferred to the real shopping basket she uses when checking out under her transaction pseudonym. The PRIME Middleware allows for this. CyberWinery also has to be trustworthy not to associate the two pseudonyms behind the screens.

There are other concerns during online interactions. What about Eve the notorious eavesdropper? Alice does not have to worry much about people acquiring her personal through interception of her communication because her personal data will be communicated using keys from the service provider and herself unavailable to Eve (public key encryption). Alice will also have some protection against 'man in the middle attacks', such as spoofed websites, because the PRIME Middleware will help her detect whether the site she visits is false, and again her personal data will be communicated using keys from the genuine site and herself.

phase 6: beyond being a connoisseur – Alice's other identities

It appears Alice has found a new hobby. She begins to like good food, good wine and matching company. She also appears to have a good nose and matching taste. She quickly gains a reputation as a connoisseur which also becomes apparent in online communities. In one of them, iConnoisseur, she gains a reputation of being a real expert under her pseudonym Malbecky. iConnoisseur's reputation system is based on the member's rating of the amount and quality of others' contributions. Alice receives 6 out of 10 corks in a whim. When she joins CyberWinery's forum, she learns that the quality of discussion is much lower here and she decides to contribute to improve the forum of her favorite webshop. However, as a newcomer she has trouble being heard. If only she could bring in her reputation.

This anecdote illustrates a common problem in the online world. Netizens build up reputations such as financial creditability, but also valuations and ratings by peers, such as iConnoisseur 'corks' are common. Transferring reputations from one context to the next, without linkability of the underlying partial identities is a feature that will prove valuable in online interactions.

PRIME can handle this kind of reputation transfer because reputations can be transferred into (anonymous) credentials. iConnoisseur can provide Alice with a credential that she can present at CyberWinery's forum. CyberWinery can check the validity of the credential, without being able to establish a link to Alice's pseudonym in the iConnoisseur site.

Now that Alice has become a real connoisseur, she starts thinking about a career shift. She visits many vineyards in Spain, Italy, and France. She notices the steep price differences between CyberWinery and local vineyards and sees a business opportunity. She and her bookkeeping genius of

a sister Alicia set up a small online wine shop which implements the PRIME Middleware to honour their customers' privacy.

Their shop, MerchantSisters, flourishes, but one of their customers, identified as Bob13, plays a trick on them. He (or she) does not pay for a large shipment after a number of successful transactions. The sisters want to claim payment but need a way to address Bob13 who does not respond to mail sent to the email address he provided.

PRIME allows for several new business mechanisms for privacy-enhanced services. The classical approach would be to use a payment system that adopts the first line responsibility for paying the service provider, which is how current services like credit cards deal with the issue. The problem introduced by Bob13 would not have occurred in this situation, or would have been put on the plate of the credit card company.

But with PRIME and its use of credentials and pseudonyms other approaches become feasible. Anonymity and pseudonymity have their limits. As users and service providers should be accountable for their actions when they breach their contractual or legal obligations, also when they are surfing the web. Users can use pseudonyms and credentials to minimise data disclosure as long as there are mechanisms to reveal their civil identity when warranted, and under strict conditions. One of these conditions would be the use of a trusted third party that is contractually bound to reveal the civil identity of the user under certain circumstances (i.e., breach of contract between the MerchantSisters and Bob13 in our case).

Another approach would go even further and have the trusted third party act as a court of arbitration. The contract between the MerchantSisters and Bob13 could contain a clause subjecting both parties to the rulings of this court. In many cases, alternative dispute resolution can work cheaper and faster than regular courts - also effectively lowering the threshold for making sustained claims. Involving the trusted third party as an intermediary preserves Bob13's privacy if the claims of the MerchantSisters prove to be unsubstantiated.

requirement:
multiple identities and accountability

The user should be able to use a range of identifiers with varying degrees of observability and linkability. This means users must have a choice to operate anonymously, pseudonymously or known. Users should also be able to use identities provided by public bodies or enterprises, as well as ones created by themselves, to be able to provide certainty about their identity to other entities and therefore promote accountability when required.

the bigger picture

The preceding pages have illustrated some of the (privacy) issues that individuals and businesses encounter in online interactions and the ways in which PRIME can offer privacy-enhancing solutions to these problems. The scenario introduced a limited application domain, online shopping. The PRIME concepts can also be used in other application domains, and also in other forms of communication. Here are some examples.

The adoption of mobile phones and other mobile communication equipment is enormous. Because the location of these devices can be determined by telecommunications providers, this opens the way to a plethora of Location Based Services (LBS). One of these developments involves pull services. Here, the user initiates a location determination which is then used to provide a location based service, such as pointing out the nearest train station or pharmacy. Push services are also possible. Here the service is activated without the individual's intervention. The location of the device triggers services the user subscribes to. For example a service could inform the user that one of their friends is nearby. These scenario's are likely to involve multiple service providers: the telecom infrastructure provider, content service providers and telecom providers. It may be undesirable for these different providers to have access to the data generated by location based services. For instance, why should the telecom provider, let alone the infrastructure provider, know that Alice is looking for a pharmacy? PRIME technology can be used in LBS provisioning to offer ways to keep these various service providers separate and thereby maintain the unlinkability of the user's personal data. This scenario is the basis of one of the PRIME application prototypes.

Another important area where PRIME concepts can be of service is in citizen government interactions. Current eGovernment services and identity management infrastructures are not exactly ideal from a privacy perspective. Adoption of PRIME technology in eGovernment would open ways for pseudonymous interactions while also allowing identified interaction, when required. This use runs parallel to Alice's shopping scenario. The added bonus is that the government can serve as a credential provider which would leverage privacy-enhancing technology from beyond eGovernment use to private sector use because there is a clear need for certified credentials here as well.

A third area where privacy issues can be tackled by PRIME technology are social networks. Profile sites, self-help discussion forums, and even virtual communities such as Second Life are environments where the users are very open about their interests, attitudes, concerns and behaviour. Though this is not without problems. The mechanisms controlling access to personal data are coarse in most cases. For instance, friends, and friends of friends, can have access to your profile data. It becomes increasingly clear that elaborate schemes are necessary to curb the spread of personal data, for instance by distinguishing types of stakeholders: friends, colleagues, sporting mates, etc. PRIME concepts can help here to define circles of users, decide who gets access to what data, offer encrypted data to be unencrypted only by authorised 'friends', and allows the user to see who had access to what data.

concepts and human computer interaction

The preceding sections have illustrated some of the PRIME concepts[9] and some possible uses. Introducing and adopting privacy-enhancing identity management not only makes online life possibly easier, for instance by enabling portable identities, it also means that individuals and businesses have to adopt different kinds of concepts and modes of operation. Data minimisation also means a change of attitude and culture. But beside this, relatively novel concepts such as roles, use contexts, credentials, and certificates are required. Although most people (implicitly) use the concept of social roles, for instance Alice is Alicia's sister, entrepreneur, tennis player, and possibly also mother, this use of role concepts to delineate access to personal data will be new to them. Yet these kinds of concepts are prerequisites for more elaborate privacy-enhancing identity management systems. Privacy-enhancing identity management is not mature but a field in flux and it is still in the research phase. This means that, although the underlying technical mechanisms are relatively clear, the translation of these to concepts understandable for the normal user are not yet completed. In this respect, the user interface to the identity management system plays an important role because it is the user's instrument and shields the user from the technical intricacies. Much work in this field remains to be done on the level of requirements, the conceptual level, and in designing concrete interfaces. Some approaches in this field are also shown in the PRIME project.

public awareness

Privacy issues abound, and to some extent solutions are also present. Yet the adoption of privacy-enhancing solutions by businesses and individuals has so far lagged behind what may be necessary to bring the Internet to full fruition. This is partly due to a lack of awareness among the general public of the risks involved in the unbounded disclosure of personal data. Reports in the popular press about privacy incidents involving personal data leaks from enterprise and government databases, about profiling and mining an individual's past on profile sites by human resource departments and reports about ID theft surface more frequently. This may slowly increase the public's awareness that to be more careful with their personal data than they think. The PRIME project sees it as one of its tasks to raise public awareness with respect to privacy issues in a more systematic way. White papers such as this one, but also general public tutorials and promotional videos are part of this work package.

economics

Businesses are utilizing data, in particular personal data, and so personal data routinely for daily operations, and as means of customising services, e.g. to employees and customers. Some of these information-processing practices are coming under increasing scrutiny leading to a call for better privacy management in organisations. Some processes may even become impossible to execute because of limitations imposed by privacy regulations and policies. In definitional terms a business process is a structured, measured set of activities designed to produce a specified output for a particular (internal or external) customer or market. The central question that concerns PRIME is how business processes are impacted by personal data, and how they can be reengineered to improve their privacy management. Realizing an adequate level of data protection requires the implementation of a set of organizational/procedural, e.g. segregation of duties and data handling procedures and technical measures. The latter are usually described as 'Privacy Enhancing Technologies' (PETs).

For the implementation of PETs solutions and PRIME in general, a increased level of maturity of the organization is often required. It is highly unlikely that an immature organization will implement PETs,

---

[9]   More detailed (technical) information can be found in the PRIME Architecture V2 and PRIME Framework V2 documents.

let alone that these organizations have any awareness of privacy protection. For privacy in particular we believe that there are two levels: the level where privacy is at best an ad hoc process, with local patches to solve local privacy problems; and the level where privacy is subject to a focused company policy.

The benefits offered by PETs can be quantitative or qualitative. If the application of PET leads to a reduction in costs or increase in revenues (e.g through a bigger market share), then the benefits can be measured and, therefore, are quantitative. Qualitative benefits are tricky to measure and hard to express in monetary terms; however, they can surpass the quantitative benefits. One example is the positive image generated by the application of PETs.

Costs of PETs vary with the selected PETs option. For example if the option is data anonymization the emphasis lies on the one-off investments and less on the structural costs. When data are separated, different domains are created, the data model usually has to be modified, and there is more often a need for customization to implement the PET option. Encryption, for instance, is often cheaper than the application of biometrics with PKI.

reaching out

Finally, in order for privacy-enhancing identity management to be adopted on a large scale not only requires that individuals take notice of the technology. But it also requires service providers to implement the necessary software. Businesses and governments will only do so if they see an advantage for doing this. PRIME investigates and reports on business opportunities, costs and benefits in order to show the viability of adopting privacy-enhancing identity management. It allows businesses and governments, for instance, to comply with data protection legislation more easily. It may also reduce their liability because storing less personal data means less vulnerability to attacks by ID thieves. Not asking for excessive data and offering ways for pseudonymous transactions may also increase the quality of the data they have about their customers.

Another prerequisite for large scale adoption is interoperability. PRIME, or for that matter any identity management system, stands no chance unless it allows interoperability with existing back-end applications and other identity management systems. This calls for standardisation. The PRIME project is therefore actively involved with standardisation bodies, such as W3C and the relevant ISO/IEC Working Groups.

your move

The present white paper has given a glimpse of PRIME project's vision, goals and the way of achieving these. Having read this, it will be clear that much work remains to be done. To achieve the goals set out we need your input, so we kindly invite you to join the discussion and help us bring privacy-enhancing identity management closer to reality.

Please visit us online at http://www.prime-project.eu

references and further reading

- PRIME Architecture V2, https://www.prime-project.eu/prime_products/reports/arch/.
- PRIME Framework V2, 27 July, 2006,
  https://www.prime-project.eu/prime_products/reports/fmwk/.
- [Art. 29 wp 2004] Article 29 Data Protection Working Party, Opinion on More Harmonised
  Information Provisions, November 25, 2004,
  http://www.europa.eu.int/comm/justice_home/fsj/privacy/docs/wpdocs/2004/wp100_en.pdf.
- [CA 2007] Identity Federation: Concepts, Use Cases and Industry Standards, White Paper from
  Computer Associates, Inc., January 2007,
  http://www.ca.com/Files/WhitePapers/identity_federation_white_paper.pdf.
- [Camenisch/Lysyanskaya 2002] Jan Camenisch, Anna Lysyanskaya: A Signature Scheme with
  Efficient Protocols, in: SCN (S. Cimato, C. Galdi, G. Persiano, eds), vol 2576 of Lecture Notes in
  Computer Science, Berlin, Heidelberg: Springer, 2002, pp. 268-289.
- [Cameron 2005] Kim Cameron: The Laws of Identity, May 2005,
  http://www.identityblog.com/stories/2005/05/13/TheLawsOfIdentity.pdf.
- [Cavoukian 2006] Ann Cavoukian: 7 Laws of Identity: The Case for Privacy-Embedded Laws of
  Identity in the Digital Age, Ontario Information and Privacy Commissioner, 2006,
  http://www.ipc.on.ca/images/Resources/up-7laws_whitepaper.pdf.
- [CoE 1981] Council of Europe: Convention for the Protection of Individuals with regard to
  Automatic Processing of Personal Data, Strasbourg, January 28, 1981
- [CSA 1996] The Canadian Standards Association's Privacy Code, http://www.csa.ca/standards/
  privacy/code/
- [IBM/MS 2003] Federation of Identities in a Web Services World, A joint whitepaper from IBM
  Corporation and Microsoft Corporation, July 8, 2003, Version 1.0,
  ftp://www.software.ibm.com/software/developer/library/ws-fedworld.pdf.
- [ISO/IEC 2005] ISO/IEC 1st WD 24760: A framework for identity management, 2005.
- [Liberty 2006] Liberty Alliance Project Whitepaper: Personal Identity, March 23, 2006,
  http://www.projectliberty.org/liberty/content/download/395/2744/file/Personal_Identity.pdf.
- [OECD 1980] OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal
  Data, http://www.oecd.org/document/18/0,2340,en_2649_34255_1815186_1_1_1_1,00.html

appendices

appendix 1 – requirements for Identity Management Systems

At the start of the PRIME project in 2004, the following principles were adopted as guidelines for the design and implementation of privacy-enhancing identity management solutions.

The PRIME design principles

- Design must start from maximum privacy
- Explicit privacy governs system usage
- Privacy rules must be enforced, not just stated
- Privacy enforcement must be trustworthy
- Users need easy and intuitive abstractions of privacy
- Privacy needs an integrated approach
- Privacy must be integrated with applications

The PRIME project continues to adhere to these principles. Yet in the current white paper we have approached requirements for privacy-enhancing identity management from a slightly different angle combining the PRIME principles with requirements brought forward by other initiatives. This has resulted in the following list of requirements.

user control and consent

> In order to maintain the individuals' trust in the information society and guarantee their freedom of choice (autonomy), users should be able to control which personal data are given to whom and for what purpose. Exercising control requires informed and uncoerced consent for specific uses, which may be revoked at a later date, by the individual.

justifiable parties

> Personal data should only be accessible to entities with a legitimate interest in the data, e.g., by consent of the individual, by legal obligation or for other legitimate purposes. Service providers should implement technical measures to enforce this requirement, especially with respect to the use of personal data by third parties (for secondary uses). This requirement also implies that the user should be able to check the authenticity of the data requester.

data minimisation

> Personal data disclosure should be limited to adequate, relevant and non-excessive data. Implied in this requirement is that data needs to be provided on a need-to-know basis and stored on a need-to-retain basis. This requires the requester to specify the purposes of collection, processing and storing of the data. Data should be deleted at the requester's end as son as the specified purposes of data collection are met.

policies and policy enforcement

> Users should be able to express their privacy policies and preferences and negotiate the terms of data disclosure with service providers. The agreed upon policies should be strongly enforced by the identity management systems on both sides of the transaction.

human measure

> The user should be able to understand how she can exercise control over her personal data. Communication should therefore be in plain language using understandable concepts. 'Thingification' should be used for necessary but complex notions, such as roles, rights and obligations (e.g., using business cards to represent data related to a role). Human-machine communication within and between contexts should be unambiguous offering situational normality and predictability. The interface should help to protect the user against identity attacks.

multiple identities and accountability

> The user should be able to use a range of identifiers with varying degrees of observability and linkability. This means users must have a choice to operate anonymously, pseudonymously or known. Users should also be able to use identities provided by public bodies or enterprises, as well as ones created by themselves, to be able to provide certainty about their identity to other entities and therefore promote accountability when required.

appendix 2 – PRIME walkthrough in more detail

In this appendix we provide a somewhat more detailed walk-through of the network interaction in the shopping scenario.[10] We focus on the exchange of claims and credentials. The walk-through starts when Alice orders her box of wine. In the sequence diagram, this is represented as Alice requesting access to a resource.

This triggers the webstore's access control mechanism. This component responds by sending a request for a claim satisfying the condition in the Data Handling Policy (DHP) for the requested resource. In our case the DHP could be that the customer needs to show that she is over 18 years of age. She is offered the choice to provide proof by means of a valid OECD ID document and an encrypted copy of her name and address as appearing on the OECD approved ID document. Alternative she can present a pseudonym used in previous transactions.

access requires:

age > 18 ← OECD_Passport AND

Enc_k(name, address) ← OECD_PhotoID

OR

pseudonym ← CyberWinery

Alice's PRIME Middleware Access control component will in turn respond to this claim request. It will make a release decision whether the requested proof will be provided to the service provider. If the service provider is unknown to Alice's PRIME Console. It may issue a request for the shop to prove that it meets certain requirements, such as complying to certain standards. This optional request is similar to the shop's request for proof of Alice's legal age. The proof will be verified and logged.

If Alice's PRIME Access control module decides that the requested proof can be released, the claim and evidence will be communicated to the service provider. Concretely, the OECD ID passport may be instantiated to a Swiss passport, and the address on an OECD Photo ID may be instantiated to the address as appearing on Alice's Swiss driver's license. These data, together with a DHP' proposed by Alice's system that states the conditions under which the data is disclosed will be sent to the service provider. The DHP' may for instance be that the encrypted name and address may be provided to the shipping service for the purpose of being able to ship the order and the data may be retained for a maximum of three years or whatever is legally obligatory. Alice's PRIME Console will next log which data has been disclosed under which conditions.
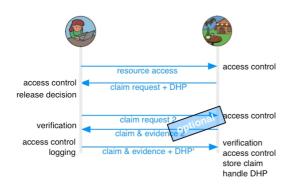

data handling policy'

Enc_k(name, address) to Shipping_Company(Purpose: service),

Max_data_retention(3 years, or legal maximum)

---

[10] For a detailed description of the PRIME system the reader is referred to the PRIME Architecture V2 documents available at https://www.prime-project.eu/prime_products/reports/arch/.

At the service provider's side similar events happen. The data are verified, and if approved access is granted to the service and the claim stored.

In order for Alice's system to make access and data disclosure decisions, it has to make use of various sources as depicted below. It will, for instance, make use of ontologies to decide that Alice's passport is an instance of an OECD passport. It will consult Alice's access control policies to decide what kind of conditions Alice may want to have associated to the data. It will access the credential storage to retrieve the data to be disclosed. And it will store all sorts of logging data in the transaction history.