

Privacy in Identity & Access Management systems

Andreas Pashalidis

Katholieke Universiteit Leuven, Belgium

Chris J. Mitchell

Royal Holloway, University of London, UK

ABSTRACT

This chapter surveys the approaches for addressing privacy in open identity and access management systems that have been taken by a number of current systems. The chapter begins by listing important privacy requirements and discusses how three systems that are being incrementally deployed in the Internet, namely SAML 2.0, CardSpace, and eID, address these requirements. Subsequently, the findings of recent European research projects in the area of privacy for I&AM systems are discussed. Finally, the approach taken to address the identified privacy requirements by ongoing projects is described at a high level. The overall goal of this chapter is to provide the reader with an overview of the diversity of privacy issues and techniques in the context of I&AM.

INTRODUCTION

Identity and Access Management (I&AM) systems support *access control*, namely ensuring that access to certain resources is granted only if the requestor is properly authorized. For example, a company employee that accesses a company VPN (Virtual Private Network) while working from abroad is likely to be granted access by an access control system. Although I&AM systems are closely integrated with access control systems, their main function is to support the system administrators and the end users in performing maintenance procedures, such as managing access credentials, user roles, access rights, rights delegation, auditing, and relationships between organizational units, throughout the lifetime of the system.

Over the last fifty years, many I&AM systems with a wide range of functions have been developed. Such systems are typically composed of a number of modules, each with a specific task. Some I&AM systems are as simple as a database with authorized username/password pairs, while others are complex distributed systems that could include sophisticated policy decision points, interconnection with business process engines, accounting and billing infrastructures, credential negotiation agents, customer relationship management systems, administrative interfaces for the lifetime management of comprehensive user profiles, and provisions for auditing. Many I&AM systems are *closed*, i.e. they are designed for environments where there is a single system provider, such as a company or government organization, that has a very strong relationship with the prospective users.

The focus of this chapter is *open* I&AM systems, i.e. systems that cover multiple organizations. In the context of such systems, users interact with a range of different organizations using one or more credentials. New users may be introduced into the system by multiple parties, or users may be able to independently create new accounts for themselves. In open systems there is clearly a need for interoperability, and thus standardization is probably more important than in closed systems; privacy also plays a central role. Users should, for example, be able to control the degree of dissemination of their personal information to organizations and other users. The particular focus of this chapter is the various degrees of privacy achieved by current open I&AM systems, and what issues need to be addressed in future such systems.

PRIVACY REQUIREMENTS FOR I&AM SYSTEMS

The need for user privacy in open I&AM system arises from the need to reduce the risks of unnecessary or otherwise unwanted disclosure of personal information. In recent years, legislation in Europe, both at EU and at national levels, has become an important driver for the introduction of privacy and transparency enhancing techniques within I&AM systems. This is because many of these laws require businesses to follow the principles of data minimization, data protection, and, in some cases, data retention. The data minimization principle requires that personal data is not disclosed to a transacting partner unless that information is strictly needed in order to carry out the transaction. In order to establish such strict necessity, the purpose of disclosure must be specified for each data item to be disclosed. Data protection and retention require that users have access to, and can update, their personal information when it is stored at an organization, but also that organizations have to keep records in a way that facilitates effective investigation of past transactions. In this context, ‘personal data’ is any data that could potentially lead to the identification of an individual, even if this is only possible in combination with additional information.

The following more concrete requirements arise from the requirement to minimize the personal data that is transferred between parties. We say that a privacy-preserving I&AM system should enable its users to:

- selectively disclose personal data to organizations and other users;
- create multiple identities or pseudonyms;
- attach different pieces of personal information to different identities;
- review data disclosed in the past;
- maintain different identities towards different organizations;
- formulate ‘sticky’ policies that follow personal data and that govern under which conditions the data may be disclosed and used;
- minimize the amount of trust users are required to place in third parties and infrastructural components in general; and
- provide explicit consent for sharing personal information, and enable users to revoke previously given consent.

Of course, achieving all the above in a usable manner, i.e. without placing too great a burden on users and system administrators, is very challenging.

The above requirements can be roughly captured by the following criteria. They can be used to evaluate I&AM systems with respect to their privacy-friendliness.

Trust model

Some I&AM systems are designed so that a remote entity, typically called the ‘Identity Provider’ (IdP), stores and manages the user’s personal information. Users are typically authenticated by an IdP, and are then able to access their own information and forward it to requesting parties. While this has the advantage of mobility – users may use the system from any computer and any location – this model also raises significant privacy issues. This is because the trusted party not only learns the personal data of the user, but will also gain information about the behaviour and relationships of the user, since other parties will refer to the trusted party every time they require user data and or assurances about user authenticity. Moreover, the trusted party must be relied upon not to assert that a user has been authenticated when this has not occurred, and/or to assert false information about the personal attributes of a user (as discussed in (Alrodhan & Mitchell, 2010)).

Of course, the privacy issues arising from the use of a third party IdP can to some extent be mitigated if the user is able to choose which IdP to use. This issue of choice arose starkly in the case of Passport, Microsoft’s initial attempt to solve the identity management problem by making itself the IdP for everyone. As has been widely documented, the notion of trusting Microsoft with large quantities of personal data gave rise to a widespread and violent negative reaction, which clearly took Microsoft by surprise (Kormann & Rubin, 2000). Indeed, this informed Microsoft’s subsequent effort in this space, the CardSpace system, discussed later in this chapter. Of course to be effective, choice requires a rich ecosystem of entities prepared to act as IdPs, and this ecosystem is still at an early stage of evolution. Moreover, even as and when such an ecosystem develops, not all users will be equipped with the means to decide which IdPs they can trust with mission-critical personal data.

An approach in which personal data, such as attributes and certificates, are stored on the user’s own computer and are then disclosed directly to the parties that require it, is likely to be more privacy-friendly. Such an approach, however, is less convenient, since mobility is no longer guaranteed and the users may have to perform a greater number of administrative tasks. It also increases the importance of security management of the platform on which the user information is stored.

Ultimately it all comes down to trust. Users will have to make a trust decision with regard to the handling of personal data, either in terms of the use of a trusted third party (e.g. an IdP) or a personal platform. Sadly, recent history suggests that this is a highly problematic issue, since users are known to make poor trust decisions with regard to the handling of personal data, as the many issues identified with social networking sites have proved (Hogben, 2007).

Multiple unlinkable identities

The notion of a user identity is commonly defined as the set of personal information for that user (attributes, certificates, credentials, and other statements concerning the individual). The user may ‘compose’ one or more personal identities by grouping relevant pieces of personal information. Identities do not need to be consistent; for example, one identity may include the user’s real name and address, whereas a pseudonym (e.g. a nickname) and address (or no address at all) may be included in another. The grouping of data into identities makes it easier for a user to switch between contexts or roles. Note that the literature sometime uses the term ‘virtual identity’ in order to refer to such a composed identity (see, e.g. (Aguar, 2010)).

In closed systems users are typically restricted to a single identity, whereas open systems typically do not have such a restriction. In particular, if the I&AM system interacts with multiple organizations that can each identify users, it may be desirable for users to be able to use different identities with different organizations. Moreover, in order to achieve data minimization, mechanisms should be provided to

prevent collaborating organizations from linking a given user's profile at one organization with the same user's profile at another. While it is relatively easy to let users create and maintain multiple identities for themselves, ensuring that these identities remain unlinkable is not straightforward. In particular, there is always a risk that usage patterns and attribute values leak enough information to link the identities of a given user. The system itself, however, should not prevent privacy-conscious users from maintaining identities that are effectively unlinkable (up to certain inherent limits, discussed in (Pashalidis & Meyer, 2006) and (Pashalidis & Mitchell, 2004)).

Selective disclosure

Selective disclosure requires that it is both possible and simple for a user to disclose only part of an identity to a given requestor. If, for example, the system has registered the user's date of birth, it should be possible to disclose only the user's age or even age group (e.g. 18–25) without having to install a separate identity or undergo a lengthy registration process.

Consent

Personal data can be used for a wide range of purposes. If, for example, a user wishes to buy electronic goods that are shipped via email, then the user must disclose his email address. However, an online shop may wish to employ a user's email address for other purposes, such as research or marketing. A privacy-preserving I&AM system should enable a user to be asked for explicit consent for such secondary uses of personal information. Similar provisions should be implemented regarding data retention times where input is needed from the user. The system should also support the revocation of consent, for example if the user no longer wishes to be contacted by the other party

Privacy respecting sharing of personal information

Sometimes it is necessary for a piece of personal information to be transferred from one organization to another. If, for example, a user orders a book from an online shop, the shop must be able to forward the user's address to the shipping company in order for the book to be delivered. In such situations it should be possible for the user to define a policy that tells the shop for which purposes and to which recipients the data is permitted to be forwarded. The notion of a 'sticky policy' is similar to consent solicitation in that they both enable the user to define acceptable retention periods, purposes, and authorized recipients. The difference, however, is that while consent only applies to the first recipient of personal information, a sticky policy is 'stuck' to the data. This means that the policy is visible and applies to all 'downstream' data processors, i.e. everyone to whom the data is disclosed in the context of a process of the I&AM system. Note that there is an interplay between privacy-respecting sharing of personal information and revocation of consent: if personal data has been already shared, then effective revocation of consent becomes very challenging; effectively revoking consent would, for example, require dynamic updates to sticky policies. To the best of our knowledge, to date no deployed I&AM has tackled this problem.

SOME COMMON SYSTEMS

We next examine some widely discussed protocols used by open I&AM systems. It is important to observe that products offer a variety of user interfaces and varying degrees of usability and functionality. Moreover, as open I&AM systems are developing rapidly both in the technical and legal dimensions, so are the interfaces and functionality of individual products. Hence it is of limited use to evaluate current implementations at a very fine level of detail. It is more valuable to examine the *protocols* that are used

to support I&AM systems. Because these protocols are standardized, they are more likely to be stable than user interfaces and software functionality sets, which change much more frequently.

In particular, we examine the privacy properties of SAML 2.0, CardSpace, and electronic ID. The reader should keep in mind that these systems focus more on identity management rather than access management. Access management infrastructures, typically located in the backend of an organization's infrastructure, are largely orthogonal to the processes that affect the systems below. Nevertheless, much current research, as outlined later in this chapter, is aimed at achieving a tighter integration between identity and access management, for example by enabling policy evaluations to be distributed over multiple domains.

SAML 2.0

SAML, which stands for 'Security Assertions Markup Language', is a set of web services protocols used in web services, and is standardized by OASIS. SAML versions 1.0 (Haller-Baker & Maler, 2002) and 1.1 (Maler, Mishra & Philpott, 2003) were published in 2002 and 2005, respectively. SAML 2.0 (Cantor, Kemp, Philpott & Maler, 2005), specifies protocols enabling organizations to exchange data about users. The typical use case involves a user that is authenticated by an organization called an Identity Provider (IdP), who maintains an account for the user.

The IdP can authenticate users by a variety of methods (Kemp et. al., 2005). The scheme is not restricted to a single IdP, and users may choose their preferred IdP from a list that contains all IdPs that are recognized by the website which they wish to access. The specifications also provide data structures which enable the IdP to send attributes it stores about a user to other websites in a manner that enables the receiving websites to verify the validity of the attributes. It is up to the IdP to provide user interfaces through which users can compose identities and exercise selective disclosure. Selective disclosure can be exercised if the user is given the opportunity to specify policies that tell the IdP which potential recipients are allowed to see which attributes, or by explicitly asking the user to confirm attribute disclosures every time they are about to take place. Both approaches have usability disadvantages.

SAML 2.0 and access management XACML (Extensible Access Control Markup Language), another OASIS standard (Moses, 2005), specifies a format for access control policies as well as formats for messages that can be used by a Policy Enforcement Point (PEP) to request a policy decision from a Policy Decision Point (PDP). The XACML SAML profile (Anderson & Lockhart, 2005) specifies how XACML messages can be sent inside SAML 2.0 messages. This specification involves a close integration of Identity and Access Management technologies. In a typical use case, some of the XACML-enhanced SAML 2.0 messages, typically exchanged between a PEP and a PDP, will carry personal data, such as attributes. This standard enables more elaborate access management because it enables one domain to outsource policy decisions to another domain. However, use of the standard is also likely to increase the risk of privacy breaches because the exchanged messages may contain personal user information (e.g. attributes) and, while this information may be necessary to reach an access control decision, the messages are forwarded across domain boundaries. The particular risk level, however, depends on the details of the deployment.

CardSpace

CardSpace is a software product produced by Microsoft that enables users to manage their identities on their own computer. In Microsoft terminology it is an 'identity metasystem', i.e. a system that aims to accommodate multiple, ideally all, Identity Management systems and offer a unified user experience towards the user. Cardspace offers suitable abstractions for processes such as the creation of identities (i.e. grouping together attributes), authentication of remote websites, and remembering histories of

disclosed personal information. CardSpace has been designed to promote adoption of Identity Management systems by presenting these abstractions to the user in a self-explanatory and easy to use manner, namely in the form of ‘Information Cards’. Each such card can contain a range of different types of personal information, and part or all of the contexts of such a card can be selected for disclosure to a remote website.

CardSpace conforms to the Identity Metasystem Interoperability Standard (Jones & McIntosh, 2009) and supports “U-Prove” anonymous credentials (Brands, 2000). According to (Jones & McIntosh, 2009), ‘Information Cards can be used both at applications hosted on Web sites accessed through Web browser and rich client applications directly employing Web services’. In a typical use case, however, while composition and selection of cards is done at the user’s computer, the personal information itself may be stored at a variety of providers on the Internet. CardSpace can handle different types and formats of credentials, claims and attributes including SAML 2.0 and the recently specified protocol for U-Prove credentials (Paquin, 2010). Even though it is envisaged that Information Cards reside on the user’s computer, most use cases require a wider infrastructure, with IdPs that authenticate users and provide assertions containing personal data. That is, the IdP and potentially other parties such as attribute and storage providers, are likely to be actively involved whenever the user chooses to show a card to a remote website.

As we have discussed, in CardSpace an IdP provides interested parties with statements about the attributes of a user. The system allows the recipient of such statements to be confident that the user with which it is communicating is the rightful holder of such attributes. If the attributes include a unique identifier, then the system thereby provides a means for a party to (indirectly) authenticate a user. That is, in some sense CardSpace combines attribute management with the provision of user authentication services. This property is shared by a number of other identity management systems. However, other identity management systems, such as Liberty¹, deal only with the issue of authentication. By restricting scope in this way, the privacy implications are much reduced, since in Liberty the IdP solely provides statements about whether a user (identified by a pseudonym) has been authenticated.

Finally we observe that CardSpace has the capability to reduce the trust requirement on IdPs not to monitor user activity (as discussed under Trust model above). A CardSpace IdP provides statements about user attributes, but, depending on which cryptographic options are in use, may not be required to know to which party this statement is being provided.

eID

eID, which stands for electronic IDentity, refers to efforts inside the European Union to introduce the electronic equivalent of national identity cards to its citizens. An eID solution typically takes the form of a smart card embedded into a credit-card-sized plastic card. An eID card can be used to authenticate a citizen, and to share information about the citizen that has been verified by the issuer of the card, i.e. the government. One rationale for the introduction of eID is an expected reduction of costs in the public sector resulting from its role in enabling citizens to interact electronically with government services. However, eID applications are not necessarily restricted to government applications; any business could decide to accept eID cards in order to identify or collect information about its customers. In some countries, e.g. in Germany, eID cards are likely to be able to generate so-called ‘qualified signatures’. These signatures can be used to sign legally binding contracts, and, because of special technical protection measures, are expected to have greater legal weight than non-qualified signatures.

A number of countries have introduced eID cards, for example Estonia and Belgium, and other countries are planning such a deployment; indeed, only a minority of EU member states do not have plans to roll out an eID in the future (Naumann & Hogben, 2009). Unfortunately, the eID systems of different countries differ to such an extent that future interoperability may be hindered (Naumann & Hogben,

2009). In order to prevent this from happening, efforts are underway to harmonize the eID landscape and to introduce more stringent privacy measuresⁱⁱ. It is important to keep in mind that many of the differences are due to differing national legal frameworks (Naumann & Hogben 2009, Jentzsch 2010).

Given the diversity of national legal frameworks, it is no surprise that different eID systems have different properties with respect to protecting citizen privacy. Apart from the the German eID system, all deployed systems of which we are aware produce a signature in order to authenticate the citizen. This is a violation of the data minimization principle since, when used for authentication only, the signature reveals more information about the citizen than is strictly necessary (Naumann & Hogben, 2009). The protocol used by the German eID card, called PACE (shown to be secure in Bender, Fischlin & Kügler 2009), circumvents this problem.

The type of data that is stored on different national eID cards, as well as the conditions under which access to this data is granted, also differs greatly from one system to the other. For example, currently the chips used in Belgium, the Netherlands, Portugal and Germany store a picture of a citizen's face, but only the German system restricts access to governmental services, whereas those used in Belgium, the Netherlands and Portugal impose no such access restrictions. Similarly, currently only the Austrian and the German cards support pseudonymous transactions, in the sense that different organizations get to see different identifiers for the citizen/card; other schemes reuse the same identifier for the citizen and/or card across contexts, thereby enabling colluding organizations to breach privacy by linking the transaction histories of any given citizen. It should be mentioned, however, that certain countries, e.g. Belgium, legally prohibit organizations from storing any long-term identifiers that are retrieved from the card, thereby reducing the risk level. For more information on the differences of eID approaches, the reader is referred to (Naumann & Hogben, 2009) and (Modinis-IDM consortium, 2006). A European eID card, called the European Citizen Card (ECC) is currently being specified; this specification supports unlinkable pseudonyms, and the different possibilities to integrate ECC with SAML 2.0 are investigated in (Eichholz, Hühnlein & Schwenk, 2009).

Other systems

SAML 2.0, CardSpace and eID are certainly not the only I&AM systems. We conclude this discussion on existing systems by briefly mentioning some other widely discussed examples of such systems.

The Liberty Alliance Project (usually abbreviated to Liberty), which went public in 2001, is one of the most prominent collaborative efforts aiming at building open standard-based specifications for identity federation systems. The Liberty model is essentially that of an Internet single sign-on (SSO) system. In this scheme, a principal (or a user) can federate its various identities to a single identity issued by an identity provider, so that the user can access services provided by service providers belonging to the same circle of trust by authenticating just once to the identity provider. This relies on a pre-established relationship between the identity provider and every service provider in the circle of trust. As stated above, Liberty does not support the management of personal information, and provide only authentication services.

Shibboleth is an open source federated identity management system that has been developed by the Internet2 consortium. It offers standards-based authentication and authorization systems. Shibboleth mandates identity federation, in which the IdP and the service provider systems consuming user information exchange public key certificates. Unlike in Liberty, the IdP and the service provider do not have to establish long-term shared pseudonyms during the federation process (but they can if they wish). Instead of long-term pseudonyms, the IdP and SP can use short-term random IDs to help preserve user privacy and maintain anonymity. The latest version of Shibboleth, version 2.0, is based on SAML 2.0.

OpenID is an open source identity management system in which IdPs issue their users with 'global' identifiers that can be used to log-in to any service provider. OpenID is somewhat different in nature to

SAML 2.0, CardSpace, Liberty and Shibboleth, and relies on a rather different model. In OpenID, an IdP issues a user with a global identifier (or OpenID) that can be used to log-in to any OpenID-enabled service provider. This identifier is typically a URL, and identifies the IdP that issued it. Obviously, there is no need for pseudonyms in this system, since IdPs and the SPs can refer to a user using the OpenID global identifier. There is no identity federation process in OpenID; however, if a user already holds an service provider-issued identifier, then the service provider may choose to ‘locally’ link this identifier with the user OpenID (i.e. the IdP-issued global identifier). Of course, since everything is based on a global identifier, OpenID does not support any degree of anonymity or pseudonymity, and hence is much less privacy-friendly than the other systems we have discussed.

PREVIOUS RESEARCH PROJECTS

We briefly discuss some of the approaches taken by three recently finished European research projects in the area of I&AM, namely DAIDALOS, SWIFT, and PRIME. The European projects discussed here are all large scale research collaborations involving a significant number, typically between 15 and 40, of partners representing both industry and academia. The lengths of the projects discussed range from 30 months to four years.

DAIDALOSⁱⁱⁱ, a project with nearly 40 partners, involved two consecutive phases that ran from 2003 to 2008, involved a number of mobile phone operators including Deutsche Telekom, France Telecom, Telekom India, and Telefonica (Spain). According to the project website, its overall goal was to ‘design, prototype and validate the necessary infrastructure and components for efficient distribution of services over diverse network technologies’ and to ‘integrate complementary network technologies to provide pervasive and user-centred access to these services’. As a result of the broad scope of the project, it would be unfair to say that its focus was on I&AM. However, a significant part of the project was dedicated to I&AM and the related privacy issues. In this context, DAIDALOS (in its second phase) introduced the concept of virtual identities. A virtual identity is an index of pointers to personal information that may reside at various places in the network. According to the project vision, users should be free to construct as many virtual identities for themselves as they wish, and choose where different items of personal information are stored. The index would be stored at an IdP (perhaps the user’s network operator), and would itself be identified by a random-looking pseudonym.

DAIDALOS virtual identities are cross-layer in nature. This is because it was recognized that lower communication layers trivially enable an adversary to link transactions, even if these transactions are made unlinkable at the application layer. To this end, the adversary simply has to observe the user’s Internet Protocol (IP) address; if the same IP address is used, with high probability the same user is behind the transaction. The project introduced mechanisms that trigger a switch of all identifiers across the stack, namely MAC address, IP address and, if applicable, SIP address, whenever the user switches his VID at the application layer. Moreover, multiple VIDs can be simultaneously active, with the consequence that the user’s device will have an equal number of concurrently active MAC addresses and IP addresses (Aguiar, 2010).

Whether or not the results of the DAIDALOS project with respect to privacy-enhancing I&AM will be taken up by industry and deployed in real mobile networks remains to be seen. Certainly many practical obstacles will have to be overcome, most importantly the introduction of the new infrastructure that enables different operators to interoperate. Moreover, the replication of the entire communication stack whenever the user switches a VID is likely to introduce potentially unacceptable performance degradation.

The SWIFT (Secure Widespread Identities for Federated Telecommunications), was a 30-month project that built on the concept of DAIDALOS virtual identities. The main focus of the project was the integration of virtual identities into the authentication infrastructure of telecommunication operators,

enterprises and ISPs. One important driver was the desire to support flexible charging and billing schemes as well as a form of single sign-on in which the user's ISP acts as an IdP, causing the user to be automatically logged into services on the Internet without further interaction (Azevedo 2008).

One goal of the PRIME (Privacy and Identity Management for Europe) project was to develop a privacy-preserving identity management system. The approach it adopts makes use of cryptographic tools called 'anonymous credentials'. Such credentials enable a level of data minimization that is not possible with conventional public key cryptography. In particular, an anonymous credential enables a user to demonstrate possession of a certified attribute to third parties, while at the same time avoiding the disclosure of any unique identifiers that would enable different demonstrations of the same credential to be linked. The project also developed an architecture that acts as a middleware component between an application and the repository that holds the user's personal information. The architecture combines access control policies that support obligations, negotiation and trust management; for further detail, the reader is referred to (Sommer, Cassasa Mont & Pearson 2008).

While DAIDALOS and SWIFT seem to assume that the user's data will be primarily stored by services (e.g. IdPs) in the network, PRIME's default mode of operation appears to assume that user data is managed on the user's own computer. It should be emphasized that these two approaches are very different, because the former requires third parties to be entrusted with user data while the latter does not. The distinction is, however, a superficial one since, in principle, both modes of operation are possible. Meta-identity systems like Higgins^{iv} and Cardspace (Brands, 2000) make this degree of flexibility explicit.

ONGOING RESEARCH PROJECTS

This section presents currently running European research projects that, amongst other things, aim to improve privacy in open I&AM.

PrimeLife

PrimeLife^v (Privacy and Identity Management throughout Life) is a European project that aims to develop mechanisms that prevent the collection of the massive amounts of personal data that individuals leave behind in their online transactions. The project takes a somewhat holistic approach, looking at the problem not only in specific domains, but in a range of domains and throughout an individual's lifetime. PrimeLife builds on the work done in PRIME^{vi}, but also aims at addressing the requirements listed earlier in this chapter. To this end, a number of mechanisms are being developed within the project. These can be roughly divided into cryptographic primitives, transparency support tools, privacy enhancing technologies, mechanisms for access control, and user interface development.

The cryptographic mechanisms build on previous work on anonymous credentials and related types of cryptosystem. In this area, the project has developed more efficient mechanisms for the encoding of attributes and the revocation of anonymous credentials, as well as protocols that allow users to retrieve information from a server without the server learning which exactly item of information was disclosed. Research topics also include enforcement mechanisms to prevent excessive data sharing in the context of social networks, and a 'trusted wallet' i.e. a software module that can manage sensitive information for multiple security modules.

Transparency support tools are, in the view of the project, tools that enable the user to access data that is stored about the user at third parties, the purposes of data collection, and the risks involved in divulging further information. First results aim to obtain an overview of technologies in this area, but also results on measuring privacy properties such as anonymity and unlinkability have been produced. It should be noted here that the topic of how to obtain and communicate reliable privacy measurements to end users is very challenging and still in its infancy.

Privacy enhancing technologies considered within PrimeLife include mechanisms suitable for the establishment of collaborative groups, the management of trust for privacy-preserving reputation systems, and for querying large collections of personal information without compromising the privacy of the individuals that are represented in the data set. In the area of access control, the project concentrates on how to capture the purpose for which personal data may be requested within access control policies, as well as the confidentiality of the policies themselves, and how users can define access control for data that is stored at external parties. These external parties do not necessarily have to be trusted with the data, but only with encrypted versions of the data. The project also aims to examine how such approaches can lead to better protection of biometric traits.

Finally, the project aims to increase end user awareness about privacy issues and to provide useful controls to users. This is to be achieved by implementing a variety of prototype user interfaces and subsequently conducting user studies in order to gain understanding of what abstractions and metaphors work in practice. For more information, the reader is referred to (Camenisch & Samarati 2009) and (Fischer-Hübner, Wästlund & Raggett 2009), as well as the project website.

TAS3

TAS3 (Trusted Architecture for Securely Shared Services) is a project related to I&AM which aims to develop an architecture that deals with authentication of users and organizations, credential management, the establishment of trust between users and organizations, compliance considerations such as data protection policies, and a seamless integration into established business processes.

One of the main differences between PrimeLife and TAS3 is that, while the former project focuses on improving different privacy-preserving techniques, TAS3 focuses on the specification and development of a concrete architecture that integrates such techniques, while addressing the challenges that arise from this integration. Like PrimeLife, TAS3 also introduces mechanisms for the specification and management of policies that govern access to personal information. In particular, the project specifies a comprehensive authorization infrastructure that takes into account the requirements from different stakeholders: the user's privacy preferences in the form of explicit consent and sticky policies, policies from multiple organizations, and input from a business process engine.

Although the TAS3 architecture is generic and is designed to handle any type of information and personal data, the main scenarios targeted by the project are e-health and employability. In the e-health scenario, sensitive medical data about patients must be made available to doctors, while it must also be ensured that non-authorized persons cannot access a patient's medical data. Moreover, it must be guaranteed that the system can be audited, and hence a trustworthy log file of who accessed which files must be constructed. Emergency situations must also be addressed, where a doctor may need to access a patient's file even if the doctor could not do so in the absence of the emergency.

The employability scenario, on the other hand, focuses on the situation where a user uploads CV data to a special server in order to support a search for a job; instead of the user manually filling out cumbersome forms at every potential employer's site, the system enables potential employers to see the required data from the user's uploaded CV. This scenario highlights the need for the system to be able to handle complex structures in personal data, and to handle complex policies regarding the handling of data with respect to consent, purpose, and forwarding to third parties.

Both the e-health and the employability scenarios provide a motivation for introducing a business process engine that orchestrates the overall information flows and that enables changes to the process to be introduced in a structured manner.

Organizations may not only use different policy formats, but also use different vocabularies when formulating their policies. For example, while one organization might use the term ‘manager’, another may use the term ‘supervisor’ to refer to the same concept. In order to address the resulting semantic interoperability issues, TAS3 is also developing modeling tools that capture the diversity of naming. A dedicated TAS3 component is planned that will translate the affected policies into a common format at runtime. More details of the TAS3 architecture are given in (Kellomäki 2009).

SUMMARY AND CONCLUSIONS

In this chapter we have discussed privacy issues that arise in the context of I&AM, and provided a high level overview of how certain systems that are currently being used or developed address these issues. We have found that privacy protection plays a major role in the current I&AM landscape, and that most protocols are designed with at least some of the privacy requirements in mind. Whilst privacy-protecting protocols and mechanisms are a necessity in order to achieve an overall system that is privacy friendly, their mere presence is not sufficient. Assuming that privacy-protecting protocols are in place, the user interfaces of the system, as well as the degree to which the system enables users to exercise fine-grained control over the dissemination of their personal information, will to a large extent determine the level of privacy that can actually be obtained. That is, the mode of operation imposed by the I&AM infrastructure, including the underlying trust assumptions, determine whether or not it is possible for users to retain their privacy.

Finally, the usage of the system also affects privacy; if only one user is using the system then clearly there cannot be any privacy. That is, whether or not the design and assumptions regarding future usage of the system matches the actual usage when it takes place, is also important. Hence, the issue of privacy protection in the context of I&AM is likely to remain an important and active research area for the foreseeable future.

REFERENCES

Anderson A. & Lockhart, H (2005), SAML2.0 Profile of XACML v2.0, OASIS Standard, 1 February 2005, Retrieved August 23rd, 2010 from <http://docs.oasis-open.org/xacml/2.0>

Aguiar, R.L. (2010), Deliverable DII-122, Updated Daidalos II Global Architecture, (2007). Retrieved August 23rd, 2010 from <http://www.ist-daidalos.org/>

Alrodhan, W.A. & Mitchell, C.J. (2010), Enhancing User Authentication in Claim-Based Identity Management, in: *Proceedings of CTS 2010, the 2010 International Symposium on Collaborative Technologies and Systems* (pp.75-83), 17-21 May 2010. Chicago, Illinois, USA, IEEE, Piscataway, NJ.

Azevedo, R. (2008), SWIFT Deliverable D.402 – Specification of Identity-based Service Architecture with AAA functions, December 2008, Retrieved August 23rd, 2010 from <http://www.ist-swift.org/>

Bender, J. & Fischlin, M. & Kügler, D. (2009). *Security Analysis of the PACE Key Agreement Protocol*. Paper presented at the Twelfth International Conference on Information Security (ISC 2009), Pisa, Italy.

Brands, S. (2000), *Rethinking Public Key Infrastructure and Digital Certificates: Building in Privacy*, MIT Press

Camenisch, J. & Saramati, P. (2009), PrimeLife Deliverable D2.1.1, First Report on Mechanisms, Retrieved August 23rd, 2010 from <http://www.primelife.eu>

Cantor, S. & Kemp, J. & Philpott, R & Maler, E (2005), Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V.2, OASIS Standard, 15 March 2005, Retrieved August 23rd, 2010 from <http://saml.xml.org/saml-specifications>

Chadwick, D.W. & Zhao, G. & Otenko, S. & Laborde, R. & Su, L. & Nguyen, T.A. (2008), PERMIS: A modular authorization infrastructure, in: *Concurrency and Computation: Practice and Experience* (pp. 1341-1357), Volume 20, Number 11.

E. Maler, P. Mishra, R. Philpott (2003), Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V.1.1, OASIS Standard, 2 September 2003, Retrieved August 23rd, 2010 from <http://saml.xml.org/saml-specifications>

Eichholz, J. & Hühnlein, D. & Schwenk, J. *SAMLizing the European Citizen Card*. Paper presented at BIOSIG 2009 Special Interest Group on Biometrics and Electronic Signatures, September 2009, Darmstadt, Germany.

Fischer-Hübner, S. & Wästlund, E. & Raggett D. (2009), PrimeLife Deliverable D4.3.1, UI Prototypes, administration and presentation version 1, Retrieved August 23rd, 2010 from <http://www.primelife.eu>

Hallam-Baker, P. & Maler, E. (2002), Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML), OASIS Standard, 5 November 2002, Retrieved August 23rd, 2010 from <http://saml.xml.org/saml-specifications>

Hogben, G (2007), Security Issues and Recommendations for Online Social Networks, ENISA Position Paper 1, Retrieved August 23rd, 2010 from <http://www.enisa.europa.eu/act/res/other-areas/social-networks/security-issues-and-recommendations-for-online-social-networks>

Jentzsch, N. (2010, May). *Welfare Analysis of Secondary Use of Personal Data*. Paper presented at the Ninth Workshop on the Economics of Information Security (WEIS 2010), Harvard University, USA.

Jones, M.B. & McIntosh, M. (2009), Identity Metasystem Interoperability Version 1.0, OASIS Standard, 1 July 2009, Retrieved August 23rd, 2010 from <http://docs.oasis-open.org/imi/identity/v1.0/identity.html>

Kellomäki, S. (2009), TAS3 Deliverable 2.1 – TAS3 Architecture, Retrieved August 23rd, 2010 from <http://www.tas3.eu>

Kemp, J. & Cantor, S. & Mishra, P. Philpott, R. & Maler, E. (2005), Authentication Context for the

OASIS Security Assertion Markup Language (SAML) V2.0, OASIS Standard, 15 March 2005, Retrieved August 23rd, 2010 from <http://saml.xml.org/saml-specifications>

Kormann, D.P. & Rubin, A.D (2000), Risks of the Passport Single SignOn Protocol, in: *Computer Networks* (pp. 51-58), Volume 33, Issues 1-6

Modinis-IDM consortium (2006). D3.5 Identity Management Initiative Report 1 IIR1. *Deliverable of Modinis Project*. Retrieved August 2nd, 2010, from https://www.cosic.esat.kuleuven.be/modinis-idm/twiki/pub/Main/ProjectDocs/modinis.D3.5_Identity_Management_Initiative_Report_1_IIR1.pdf

Moses, T. (2005), eXtensible Access Control Markup Language (XACML) Version 2.0, OASIS Standard, 1 February 2005, Retrieved August 23rd, 2010 from <http://docs.oasis-open.org/xacml/2.0>

Naumann, I. & Hogben, G. (2009), Privacy Features of European eID Card Specifications, ENISA Position Paper, Retrieved August 23rd, 2010 from <http://www.enisa.europa.eu/act/it/eid/eid-cards-en>

Paquin, C. (2010), U-Prove Technology Integration into the Identity Metasystem V1.0, Retrieved August 23rd, 2010 from <http://connect.microsoft.com/site642/Downloads/DownloadDetails.aspx?DownloadID=26953>

Pashalidis, A & Meyer, B (2006) Linking Anonymous Transactions: The Consistent View Attack, in: *George Danezis, Philippe Golle (editors), Privacy Enhancing Technologies, 6th International Workshop, PET 2006, Cambridge, UK, June 28–30, 2006, Revised Selected Papers* (pp. 384-392) , Springer Verlag, LNCS 4258, pages 384-392, Berlin.

Pashalidis, A & Mitchell, C.J. (2004), Limits to Anonymity when Using Credentials, in: *Security Protocols, 12th International Workshop, Cambridge, U.K., April 26-28. 2004, Revised Selected Papers* (pp. 4-12), Springer Verlag LNCS 3957, Berlin.

Sommer, D. & Cassasa Mont, M. & Pearson, S. (2008), PRIME Deliverable D.14.2d – PRIME Architecture V3, July 2008, Retrieved August 23rd, 2010 from <https://www.prime-project.eu>

ⁱ <http://www.projectliberty.org>

ⁱⁱ See, for example, the STORK project at <https://www.eid-stork.eu/>

ⁱⁱⁱ <http://www.ist-daidalos.org/>

^{iv} <http://www.eclipse.org/higgins/>

^v <http://www.primelife.eu/>

^{vi} <https://www.prime-project.eu/>