Improving Privacy in Identity Management Systems for Health Care Scenarios

Rosa Sánchez-Guerrero, Florina Almenárez, Daniel Díaz-Sánchez, Andrés Marín and Patricia Arias
Dept. Telematic Engineering, Carlos III University of Madrid
Avda. Universidad 30, 28911 Leganés (Madrid), Spain
Email: {rmsguerr, florina, dds, amarin, ariasp}@it.uc3m.es

***Abstract*— Privacy is a very complex and subjective concept with different meaning to different people. The meaning depends on the context. Moreover, privacy is close to the user information and thus, present in any ubiquitous computing scenario. In the context of identity management (IdM), privacy is gaining more importance since IdM systems deal with services that requires sharing attributes belonging to users' identity with different entities across domains. Consequently, privacy is a fundamental aspect to be addressed by IdM to protect the exchange of user attributes between services and identity providers across different networks and security domains in pervasive computing. However, problems such as the effective revocation consent, have not been fully addressed. Furthermore, privacy depends heavily on users and applications requiring some degree of flexibility. This paper analyzes the main current identity models, as well as the privacy support presented by the identity management frameworks. After the main limitations are identified, we propose a delegation protocol for the SAML standard in order to enhance the revocation consent within healthcare scenarios.[1]**

Keywords-Identity management; privacy; user-centric; federation; anonymity; pseudonymity; delegation; health care.

I. INTRODUCTION

Privacy is a very complex and subjective concept, since it has a meaning different for every individual. Privacy is related to the sensitiveness of the people and the context in which the information is used. The digital age we live in imposes scenarios in which users' information is extensively collected and distributed to make it available wherever the users are. Pervasive computing and social immersion have made users active broadcasters of their own life. This fact is extremely worrying, especially for young people, despite they have born into technology (i.e. "digital natives"). They are no conscious about the consequences when it comes to spread their personal information all over the Internet. Nevertheless, the origin of this situation can be aligned to the lack of comprehensive privacy frameworks. Hence, the privacy issues risen when distributing their information are considered by users minimal tradeoffs against the benefit of being connected, since there is no way to manage privacy appropriately. However, individuals want to have control of their information. The improper and unsecured management of such attributes may lead to attacks, frauds, and identity misuse, as identity information can be exploited whenever authentication and authorization based on those identity attributes are required. Malicious parties collect sensitive identity attributes of individuals and use them to impersonate users.

Identity Management (IdM) systems provide frameworks for sharing users identity attributes among different entities, and can be used to keep data under users' control. For that reason, IdM systems are the cornerstone of security systems, because they can be used to keep user confidence while preserve privacy in an appropriate way. Such confidence is preserved when users' attributes are exchanged between service providers (SPs) and identity providers (IdPs) across different networks and security domains. Currently, there are several approaches to identity management being the most popular the federated and user-centric approaches. These approaches provide many services as the popular Single Sign On (SSO) across multiple trusted domains. SSO allows users of one domain to securely access resources of another domain seamlessly, requiring no redundant login processes. Both approaches have benefits and shortcomings, for instance, the federated model has scalability issues which the user-centric model solves, but both of them can be used for a better privacy management.

However, current IdM systems by themselves have several problems that should be solved to handle privacy adequately; for instance, the problem of revoking consent is not covered by any of the aforementioned identity management approaches. Revoking consent is part of the privacy rules [1] in health care. This property consists of revoking access to personal data that had been already shared. The privacy requirements shall depend on users and applications, requiring flexibility to handle different application domains. In this article, we focus on the privacy concept within identity management systems (IdM) in ubiquitous environments. Particularly, we have selected health care scenarios since they are among the most sensitive scenarios. Our privacy aware identity system implements an effective consent revocation with an innovative event management that can be used in other scenarios. For this purpose, we assume that the development of patients care can be broken down into events. These events describe a specific situation and can be related to some participant entities. We propose a delegation protocol, which issues a *sleepyhead credential* containing user's attributes and access privileges that have been granted

[1] This article has been partially funded by the grant CCG10-UC3M/TIC-4992 from the state of Madrid and Carlos III University and by the State of Madrid (CAM), Spain under the contract number S2009/TIC-1650, project E-Madrid

beforehand but that are kept latent. To use these attributes, an activation process is necessary. Our solution proposes events to awake dormant privileges or part of them and incorporates several new features that allow better scalability. For instance, emergency services are the entities which manage trust indirectly by emitting events that will be used to determine which participants can access to user's medical information. The rest of this paper is structured as follows: section II presents the main privacy properties and current identity management approaches, identifying the advantages and drawbacks of each one in terms of privacy. Section III provides a comparative analysis of the privacy support in identity management systems. Moreover, open issues and related works regarding privacy in IdM are described in this section. Then, section IV explains our delegation protocol to enhance privacy in health care scenarios. Section V describes implementation issues. And finally, Section VI summarizes our work and presents the main conclusions and future lines.

II. BACKGROUND

As mentioned before, privacy is complex to handle and needs to cope with different sensitivities that depend, among others, on the context in which the information is used. Privacy comprises several subtopics as anonymity, pseudonymity, unobservability and unlinkability that might have different definitions in the literature [2]. In this paper, we focus on privacy within identity management systems (IdM). Thus, this section presents definition of the main privacy principles, such as anonymity, pseudonymity, unobservability and unlinkability. Likewise, we discuss how these concepts are addressed in identity management systems.

A. Principles of pivacy

- ***Anonymity***: can be defined as the state of being not identifiable within a set of subjects or entities, also called the anonymity set. Another definition provided by the Common Criteria [3], asserts that this property ensures that a user may use a resource or service without disclosing the user identity. Cryptographic techniques, such as encryption, do not guarantees anonymity since an observer could analyze traffic, eavesdrop the sender of the message and follow the message up to the receiver, establishing certain relationships without having access to the unencrypted message. Therefore, IdM systems must provide additional mechanisms, such as opaque identifiers to prevent such inference.
- ***Pseudonymity***: is the use of pseudonyms as identifiers. An advantage of pseudonymity technologies is that accountability for misbehavior can be enforced. Thus, this enables Identity providers, that can link identifiers to real identities, to take appropriate decisions when a user commits a crime or offense in an IdM scenario.
- ***Unlinkability***: ensures that a user may consume multiple resources or services without letting other entities to link these multiple resource or service accesses together. In particular, this property allows users to interact with multiple organizations (SPs or IdPs), each of them able to map a user to a given identity, using different identities. Moreover, IdM systems should provide mechanisms to prevent collaborating organizations from linking a given user profile at one organization with the same user profile at another. While it is relatively easy to let users to create and maintain multiple identities for themselves, ensuring that these identities remain unlikable is not straightforward. In particular, there is always a risk since patterns of usage and attribute values might leak enough information to link the identities of a given user.
- ***Unobservability***: permits a user to access resources or services avoiding other entities, especially third parties, to observe that the resource or service is being used. Regarding identity management, traffic analysis is a well-known example, which tries to violate this principle.

B. Current Identity Models

According to Josang [4] the fundamental privacy protection principle is that exposure of personal information should be minimized. If we transfer this concept to identity management approaches, this means that, the fewer parties involved in the management of the identity information the better. Nevertheless, achieving a good degree of privacy implies observing every of the aforementioned privacy principles. Furthermore, although the property of anonymity is one of the four principles of privacy, IdM systems should support mechanisms to break the anonymity of a user for the purpose of analysis or evidence under certain circumstances (e.g a criminal user, lawful interception). For clarity, we introduce here the main actors in an identity management scenario, that are: 1) the *Principal*, or the End User, who has a particular digital identity and interacts (usually via an user agent) with SPs; 2) the *Service Provider*, which provides services and takes decisions based on the identity information provided by a third party (IdP) about a particular subject, 3) the *Identity Provider* that authenticates users, manage identity information and shares identity information with various SPs upon user request. In this section, we briefly introduce identity management systems with an special focus on privacy. However, other aspects as usability and scalability are also considered when assessing IdM approaches.

1) Federated Identity Model: The identity federation model can be defined as a set of standards, technologies and agreements, that enable SPs to recognize user identities and entitlements from other SPs or IdPs. Thus, this approach is based on groups of SPs and IdPs that have a pre-existing

mutual trust relationship. Consequently, specifications, such as Security Assertion Markup Language (SAML) [5], recommend using Public Key Infrastucture (PKI) [6] for establishing trust relationships. Regarding the terminology of Liberty Alliance [7], the above groups are called members of the circle of trust (see Fig. 1).

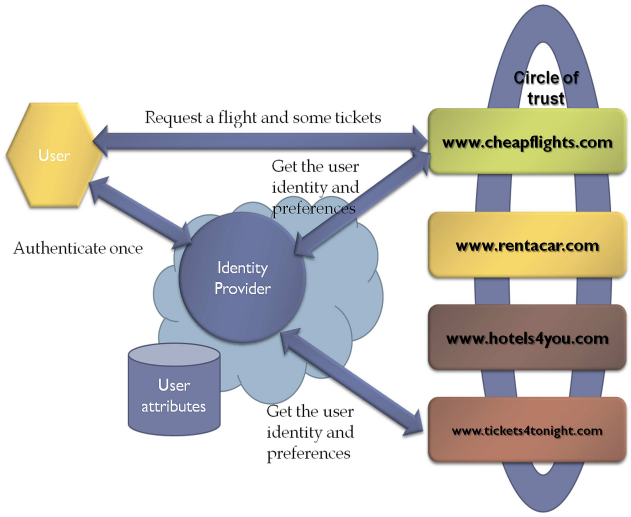

Fig. 1. Federated model scenario. A user, after a successful authentication, can access services from any service provider within the circle of trust. For instance, booking a flight, then renting a car, and finally buying tickets for a show. Note that the IdP stores identity information on behalf of the user.

Identity federation provides several services as Single Sign On across multiple federated domains, which allows users of one domain to securely access resources of another domain seamlessly, requiring no redundant user login processes. Other benefits that federated models bring are user attribute exchange, user account provisioning, entitlement management and personalized service provisioning. Nevertheless, as far as usability and scalability are concerned, this model has several drawbacks. For instance, it creates legal and technical complexity since to be part of the circle of trust, an entity would need to sign a legal agreement. In addition, federated model presents scalability issues when deployed in dynamic open environments due to rigidity and staticity of the agreements between federated organizations. A comparative analysis of the underlying trust mechanisms of the current frameworks for federated identity management can be found in [8]. From a privacy perspective, the federated identity approach has both advantages and disadvantages. Regarding its advantages, it allows users to have multiple identities within a given domain. Similarly, the federated model enables an entity to have different identities or identifiers in different domains. These features make possible, for example, the same identity to have one identifier as patient in a health care domain and another identifier as employee or student in another domain. Moreover, from the SP perspective, the identifier mapping permits different SPs to refer to the same user through different identifiers. Moreover, whereas the IdP needs to know the "real world" identity of a user, this user identity can be anonymous for a specific SP, which provides additional privacy protection. However, it must be noted that users never participate in the trust establishment process so they need to believe that the IdP will behave honestly.

In regard to the drawbacks of this kind of identity model, the main issue is that the privacy protection depends on the privacy policy and the adherence of the IdP or SP to the policy, which can be a threat. For instance, different SPs could be able to match personal information of the same user because of the mapping between identifiers. In order to prevent this problem, identity frameworks such as, SAML and Liberty advise the use of pairwise, directional opaque identifiers.

2) User-centric Identity Model: The user-centric model places the user in the middle of a transaction, thereby this approach gives users total control over their identities as well as control over authentication an attribute exchange processes. In this way, the user is no longer aside of the trust establishment process. However, this does not mean that users should approve every transaction, but that data always flow through the user's identity agent. This approach indeed empowers users and follows better than the federated model the philosophy of minimal disclosure defined by Josang. Moreover, from the usability perspective, the user-centric identity model, solves scalability problems and provides similar services, as SSO, whereas is compatible with the federated model.

In regard to privacy, this model has both advantages and drawbacks. It introduces the concept of meta-idp, which allows users to assert several kinds of claims: user-generated and provider-generated claims. These user electronic identities are typically stored in user's equipment, such as his mobile phone. User-centric identity technologies such as InfoCards [9], allow users to select among their multiple identities through identity selectors to identify itself to a service. Regarding identity selectors, in [10] two types of information cards are specified: *Personal* or *Self-Issued* (claims about the user itself, e.g. phone number, e-mail address, web address); and a *Managed Information Cards*, issued by Identity Providers. The latter can be auditing, non-auditing, or auditing-optional to accommodate the needs of different business models. The identity cards are metaphors of real id cards whereas the identity selector mimics a wallet. However, it is worth mentioning here that, in the case of provider-generated claims, the user must rely on the IdP honesty, as occurred in the federated model (see Fig. 2).

The main disadvantage of user-centric approach is that it requires a complex design in order to avoid privacy and trust issues with authentication and attribute verification. In order to assist the reader in understanding this aspect, we provide the following example. If we consider a real world example in which Bob may show his driver licence to a bartender to prove he is above the legal drinking age, we

can see that Bob is able to use his Id card without the Id card issuer's knowledge. However, if we transfer this example to a user-centric scenario, trust and privacy problems emerge, because no SP is obliged to believe Bob when he asserts that he is old enough to legal buy alcoholic beverages. In this sense, it is necessary that a trusted third party corroborates the above statement by using a *provider-generated* card.

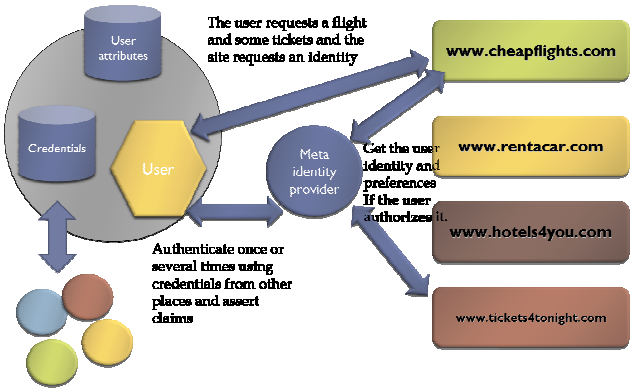

Fig. 2. User centric model. A user can access services from any service provider accepting his/her credentials. For instance, booking a flight, then renting a car and finally buying tickets for a show. Note that the information is provided always by the user.

III. PRIVACY SUPPORT: A COMPARATIVE ANALYSIS

In this section, we compare the privacy support that several federated and user-centric frameworks have, in terms of language support for privacy policies, anonymity, pseudonyms, if they reflect how unlinkability can be achieved, and revoking consent. We focus our comparative analysis on the SSO profile, because this characteristic is common to the two identity management technologies described below. In this section, some proposals that seek to address various privacy issues will be presented, as well as some open privacy challenges in IdM systems.

SAML is a federated specification, which supports two types of identifiers to refer to users: transient or one-time identifiers and persistent identifiers. On the one hand, transient identifiers ensure that a user anonymously accesses a service during SSO process, since these identifiers are created for use during a session and they are destroyed at the end. Thus, correlation between identifiers is avoided. On the other hand, the persistent identifiers provide a persistent federation and remain active until they are explicitly deleted. The permanent federation implies an account linkage process, which relates two accounts associated to a user in different SPs. Note that it is recommended to use different pseudonyms for each SP, in order to avoid different SPs belonging to the same federation to infer user behaviour.

SAML supports partial anonymity in the sense that the IdP itself is able to know which user corresponds to each identity. Indeed, SAML does not provide a solution from preventing IdPs from tracking user's visits to SP. Regarding privacy policies, this technology allows to obtain a principal's consent or describe specific attributes to satisfy requirements to preserve privacy within a health care community, through the XSPA-SAML profile [11]. Nevertheless, SAML standard states that privacy must be considered, but concrete decisions are left to the implementers.

In regard to the Liberty Alliance federated model, it defines an Identity Governance Framework (IGF) [12], which enables the creation of policies or contracts between an Attribute Provider (AP) and a SP. Therefore, IGF includes two XML syntaxes: Attribute Requirement Markup Language (CARML) and Attribute Authority Policy Markup Language (AAPML). Moreover, IGF defines basic privacy constraints such as usage, storage, propagation and display of identity data. Thus, an attribute provider creates statements to access and use protected attributes. At the same time, a SP may specify whether the requested attributes will be discarded after usage. Furthermore, the SP could request to modify the data or forward it to another SP. However, in [13], Liberty proposes a multi-level policy approach, which does not consider any specification or rules for storing user preferences in a manner that would facilitate the SPs to match the privacy policy levels in the attribute request with the levels in user's preferences. As SAML, Liberty offers long-term and one-time pseudonyms. Correspondingly, it must be noted that this specification only allows a user to have one long-term pseudonym per SP to prevent user tracking across different transactions. This is a big limitation. In addition, it does not protect against SPs cooperating to share user pseudonyms in order to track users behaviour. In order to overcome these problems, a set of rules and recommendations are proposed in [14].

In the case of InfoCards, it includes authenticated anonymity and pseudonymity, as well as the ability to express privacy policies of SPs or Relying Parties (RPs). This user-centric framework is characterized by defining a message flow that eliminates direct communication between the IdP and the SP. Moreover, InfoCards allows the identity selector to encrypt the SP identity to prevent the IdP from learning the SP identity when it receives a request for a token. Note that, this identity selector applies user-centric principles in collecting user consent. Both features together are necessary to ensure that an IdP cannot learn which SPs visits a given principal. The SAML Enhanced Client Proxy profile (ECP) is similar, but currently it only has the first characteristic. However, some IdPs may require knowledge of the RPs identity before issuing a requested token, or even if the IdP cannot learn the visited SPs, user profiling is possible by colluding parties.

Regarding OpenID, privacy considerations are not addressed in the main specification and SSO can be performed between previously unknown parties without any configuration. Thus, there is no trust model; the protocol

operates in accordance with the *trust-all-comers* philosophy. Although for some services requiring no verification this model may be sufficient, this mechanism is too simple and unsafe for many other applications, leading to privacy breach. Nevertheless, an OpenID extensions called PAPE (Provider Authentication Policy Extension) [15], provides the means for a RP to request previously agreed upon authentication policies being applied by the OpenID Provider and for an OpenID Provider to inform an RP what policies will be used. Therefore, the decision to trust can be based in the knowledge of the authentication mechanism employed. Hence, with this user-centric framework, RPs must decide for themselves which providers are trustworthy, being able to enforce policies to the OpenID Providers response.

In Table I we summarize the main privacy features of each identity framework. To conclude, all the analyzed technologies typically handle privacy by means of pseudonyms which can be transient or permanent. The only exception is OpenID, which follows the trust-and-accept-all-comers principle and privacy is not addressed. Moreover, it must be noted that, InfoCards and SAML ECP profile address better the principle of minimal disclosure. However, the problem of revoking consent is covered by none of the above IdM technologies. Thus, if personal data has been already shared, the effective revocation of consent implies an important challenge to address. For instance, it requires dynamic updates to sticky policies. Other proposals, as PKI-based solutions, attempt to solve the problem by issuing tokens that expire after some time. However, these approaches have problems since the token duration is very short so, when another entity needs to use this credential, it has to ask the user for permission again because the time has expired. On the contrary, if the token duration is longer than necessary, user's sensitive information may be exposed to entities which should not have access to that information. As far as the attribute exchange between different trust domains is concerned, current specifications focus on trust relationship between SPs and IdPs assuming that trust between users and providers are implicit. However, privacy policies that allow users to understand privacy implications (in terms of attribute exchange or delegation between different security domains) and to give their consent are poorly defined, complex to implement or out of the scope.

Current identity frameworks support partial anonymity, since authorities, as the IdP, provides obfuscated identifiers. In [16] a ring signature and a SAML extension are proposed, thereby a user can sign a message on behalf of a group and the IdP can verify the signature and confirm that the user belongs to a specific group without revealing user's identity. However, this proposal has several privacy problems. For instance, the IdP should provide a list of candidates who have similar access rights to the users in order to create the ring, which affects group members' privacy. In addition, there is no intervention possible, so if a member of the group commits a crime, the IdP cannot determine the identity of the perpetrator. The approach in [17] includes homomorphic encryption techniques to enable the IdP or SP to know the result of the aggregated information from a user group, without knowing users individually. This mechanism is used for maintaining privacy of user's opinions in a reputation protocol. However, the proposed distributed architecture has scalability issues which could be overcome by means of caching mechanisms.

TABLE I. SUMMARY OF PRIVACY FEATURES IN IDENTITY MANAGEMENT

IdM Technology	Anonymity and Pseudonymity	Unlinkability and unobservabity	Privacy Languages	Revoking consent
Federated model (SAML/ID-FF)	Partial anonymity (IdP knows user identity). No solution from preventing IdPs from tracking is provided.	Transient and permanent identifiers. Different pseudonyms for each SP recommended. Confidentiality of transaction recommended. Cryptographic mechanisms do not prevent from traffic analysis attacks.	The XSPA-SAML profile enables to obtain user's consent and describe attributes to preserve privacy in health care. An identity governance framework is defined.	Not addressed
User-centric Model (InfoCards)	Included in the specification	Message flow eliminates direct communication IdP-SP. Identity selector may encrypt SP identity to prevent the IdP from learning.	Allow to express privacy policies of RPs.	Not addressed
Hybrid Model (OpenID)	Not addressed	Not addressed	Not addressed	Not addressed

IV. SLEEPYHEAD CREDENTIAL-BASED DELEGATION

As we have mentioned before, the effective consent revocation is an open privacy challenge in identity management systems. Consider a health care scenario as depicted in Fig. 3. An emergency service should attend Alice because she has suffered an accident. On the one hand, in this scenario attribute exchange and delegation process cannot be completely user-centric, since in cases of serious accidents the user cannot be able to give her consent. On the other hand, federated models raise privacy concerns since medical records may be available to every entity within the circle of trust, even if there is no emergency.

Effective consent revocation is difficult due to complexity of the management (of entities, credentials, and privacy rules), scalability, and control of long-lived delegation chains. In this section, we describe our proposal that consists of a delegation protocol, which issues a *sleepyhead credential* (SC) to overcome the above limitations. The *sleepyhead credential* contains user's attribute identifiers (i.e. her medical history), as well as access privileges, which have been granted to beforehand but they are latent. Thus, to use the aforementioned attributes or privileges, an activation process is necessary. Particularly, within health care scenarios, we model patients' life cycle as event-driven. Events are fired by trusted entities when specific circumstances are met, and routed to required entities. We propose using these events to awake the dormant privileges or part of them. Moreover, in order to prevent unauthorized access, we require some entities (like the IdP1 in Fig. 3) to use a **Privacy Engine**, responsible for analyzing events and activating the strictly needed attributes and privileges for each event.

A. Hypotheses

In the following, we describe the assumptions on which our sleepyhead credential-based delegation protocol has been built. We assume the existence of an event engine, which uses the SIP-Specific Event Notify [18] specification to send events to entities (by means of broadcast or to registered entities). Also note that, both SPs and IdPs can take the role of subscribers and notifiers, either subscribing to different events or notifying them. As regards emergency services, they are responsible for notifying events to the subscribed entities. In order to clarify this last aspect, consider that some parts of the patient's medical history reside in different IdPs and depending on the required treatment, it is necessary to consult several parts of the medical record, thereby an IdP can act as both client (subscriber) and server (notifier). In addition, communication between any two parties is confidential, and messages exchanged between them are transported over HTTPS.

On the other hand, note that it is necessary to take into account security considerations regarding SIP `SUBSCRIBE` and `NOTIFY` messages, given the high sensitivity of health care data considered in the proposal. Therefore, both subscription and notification messages must be authenticated and authorized, for instance to prevent the participating entities from subscribing multiple times or redirecting the subscription of their neighbor either intentionally or accidentally. In this sense, SIP can use different security mechanisms such as HTTP Digest or TLS. We recommend TLS for secure and encrypted SIP communications. Besides, all users utilize transient identifiers in order to preserve their anonymity while enabling the IdPs the accountability enforcement in case of user's misbehavior, according to the main principles of privacy specified in II. Essentially, the hypotheses concerning secure network communication may be satisfied by the existence of a PKI.

Finally, we assume an underlying trust relationship based on PKI for entities belonging to different domains.

B. Description of the Health Care Scenario

In our health care scenario, Alice can authenticate through a credential to the hospital that stores her medical history (IdP1). Alice suffers an accident. The emergency service (SP1) requests access to Alice's medical records in order to send them to an ambulance company (SP2), in another trusted domain, which needs to access to the patient's medical records to provide her the appropriate treatment. Thus, as events happen, they are notified to the involved parties, such as the medical record service (IdP1) and the ambulance (SP2) which treats Alice. So the IdP1 may know which ambulance should be allowed to access to medical histories.

Furthermore, each event describes a purpose, which enables to filter the access to certain parts of medical history according to a policy. Thus, in this example the following events could be distinguished:

- *Event 1*: There is an accident. SP1 notifies this event and calls all ambulance services close to the area.
- *Event 2*: An ambulance from SP2 arrives on the scene and requests access to Alice's medical history. It must give a description of the severity of problem to allow IdP1 to give access to certain parts of Alice's medical records or her full history. To illustrate this, consider that Alice has broken her femur, losses her consciousness and needs surgery. In this case, access to the whole medical record could be provided. However, if the problem is minor, as a sprained ankle, SP2 is allowed to access only to trauma and drug allergies sections of the history.
- *Event 3*: Although not depicted in the Figure 3, another possible event would be fired if Alice is taken to hospital (SP3). The hospital diagnoses her with trauma during the triage and determines that Alice requires an operation. Therefore, a doctor belonging to (SP3) could read Alice's records.

It must be noted that, events may be fired by authorized entities, like the emergency service or a hospital urgency service. Likewise, events happen asynchronously and the

duration of each event lasts from the beginning of the event itself (t_1) until another event arrives (t_2) whose circumstances and context have changed; and it may contain new requested attributes or privileges. Thus, certain attributes or privileges previously granted will be deactivated and new components of the *sleepyhead credential* will be activated.

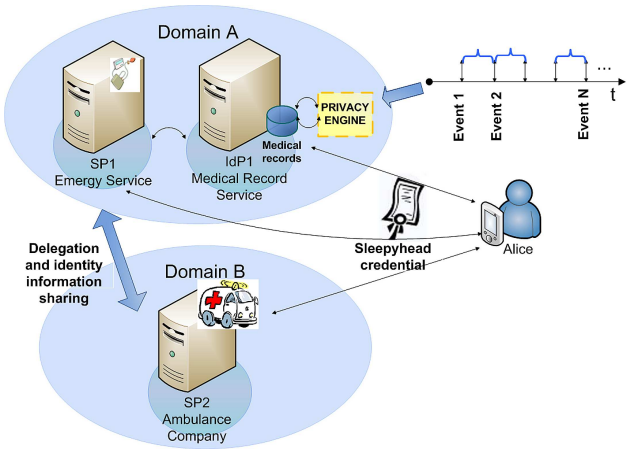

Fig. 3. Health care event-based scenario accross different domains.

C. *SAML-compliant Sleepyhead Credential*

The SC has been defined as a new SAML assertion, according to the SAML proposal for delegation information defined in [19]. Thus, the *sleepyhead credential* is created with the following tuple:

$$\mathbf{SC} = \{AttP_1, AttP_2, \ldots, AttP_n\} \quad (1)$$

Where each component $AttP_i$ represents attributes and access privileges, which have been granted to beforehand but they are latent. In Fig. 4, we can see the structure of a SAML-compliant SC, which includes the following elements for delegation restriction: `EventFilter`, which defines filters that will be used by the *Privacy Engine* to analyze the received events and decide whether any attribute(s) may be activated; the `TrustedEventSources`, which contains entity names whose events active the credential; and the `EntityMedicalRepository`, which specifies the location and distribution of attributes and medical records.

The IdPs will be the entities responsible for storing and managing the sleepyhead credentials, since as we have mentioned before, in cases of serious accident, the user may not be able to provide his credentials.

D. *Privacy Engine*

It is responsible for activating the latent attributes and privileges (following the principle of minimal disclosure) depending on the different event filters and the defined privacy policies. To this end, it analyzes the different elements which compose each event (i.e. issuer, situation, degree of severity), as well as their purposes (i.e. health care treatment, operation, emergency treatment) and applies the corresponding privacy policy. This policy includes the set of consent directives and other privacy conditions (i.e. object filtering, user, role, and purpose) that constrain enforcement.

On the other hand, the *Privacy Engine* includes an audit service for events, attribute activation, and access control decisions. It monitors how user data is being used without compromising user's identity. To accomplish this, the fields that are logged must be able to show the auditor what information about the user is being accessed without divulging the actual information. Note that, this audit service itself will not physically prevent privacy breaches from occurring but it can act as a deterrent and allow individuals and regulatory bodies to monitor how data is being shared in order to prevent from *linking* and *traffic analysis attacks*.

```
<complexType name="DelegationRestrictionType">
<complexContent>
<extension base="saml:ConditionAbstractType">
<sequence>
<element ref="del:EventFilter" maxOccurs="unbounded"/>
<element ref="del:TrustedEventSources" maxOccurs="unbounded"/>
<element ref="del:EntityMedicalRepository" maxOccurs="unbounded"/>
</sequence>
</extension>
</complexContent>
</complexType>
```

Fig. 4. Sleepyhead Credential Assertion.

V. IMPLEMENTATION ISSUES

In order to evaluate our proposal, an identity management infrastructure has been deployed using a C library called Lasso [20], which implements the full SAML 2.0/ID-FF stack. The IdP, called Authentic [21], has been developed from it. This library uses OpenSSL as underlying cryptographic library and Apache2 as the web server. Regarding SPs, we have also used ZXID [22]. Furthermore, with the aim to simulate the system of medical events through the *SIP-Notify-Event* specification, we have deployed a Sailfin Application Server [23] and implemented a set of modules that handle the associated logic to subscribe or register events, as well as send appropriate notifications to each of the participating entities. Exchanged messages contain an `Event` header that indicates the event type to which the entity is subscribed. As for the `Expire` header, it specifies subscription duration. Finally, event descriptions are sent through XML messages embedded in SIP requests.

Based on the described infrastructure we have introduced the modifications proposed in section IV. For this purpose, we are developing the *Privacy Engine* module, including the functionality to receive a structure, which represents the event filter and a hash table, which contains the event sources. Thus, this building block is in charge of checking that each of the sources that caused the event is in its *Dynamic Trust List* (DTL) [8]; and making findings relating to conditions or restrictions in the event filter in order to

determine what attributes or privileges can be activated. Therefore, the IdP has been modified to use the new privacy functionality. On the other hand, we have extended the Lasso library, defining a new structure that represents the new SAML assertion, as well its different fields and associated attributes. Such assertion is exchanged through SAML messages. In addition, it must be noted that the SP and IdP have been extended by implementing the SAML-based delegation protocol. Thus, we are currently working in order to integrate the new software components with the SIP-based event system to offer a really enhanced privacy experience and apply audit services for events.

VI. CONCLUSIONS AND FUTURE WORK

We have reviewed and analyzed the main identity models and current frameworks to preserve privacy in identity management systems, identifying its main drawbacks. Current approaches assume previously established trust between users and providers, support partial anonymity, define privacy policies poorly, or are complex to implement. Another important challenge in the context of privacy in IdM system is the effective consent revocation. However, privacy is a very complex and subjective concept, and depending on the users and applications of the IdM system, privacy requirements may vary. Thus, we have proposed a delegation protocol based on SAML, to overcome the challenge of effective revocation consent within health care scenarios.

Our solution proposes using events to awake dormant privileges or part of them and it incorporates new features that allow better scalability, since the emergency services are the entities which manage indirectly trust. Moreover, we have described some implementation details that we are currently facing. It must be noted that, the usage of the system also affects privacy and should be present in users consents. Besides, the auditing processes should verify that the design and assumptions regarding future usage matches its actual usage.

We shall test this last issue on real health care scenarios in order to demonstrate how the privacy is managed by the system actors. Likewise, how a sleepyhead credential is revoked itself, ending with the permissions must be analyzed. We shall take into account different privacy requirements for identity attributes, including biometric and health care data. Further research could be done to contemplate how to keep user's privacy during the exchange and sharing of attributes in different trust domains, also considering usability.

REFERENCES

[1] Health Insurance Portability and Accountability Act (HIPAA).: Summary of HIPAA Privacy Rule. Available at http://www.hhs.gov/ocr/privacy/hipaa/understanding/summary/privacysummary.pdf, September, 2011

[2] A.Pfitzmann and M Hansen, "Anonymity, unobservability and pseudonymity: A proposal for terminology". In Hannens Federrath editor, Designing Privacy Enhancing Technologies (PET'00), volume 2009 of LNCS, pages 1-9. Springer-Verlag, 2001.

[3] The Common Criteria Project Sponsoring Organizations, "Common Criteria for Information Technology Security Evolution" -part 2, version 2.1 August 1999.

[4] J. Audun, M. Alzomai and S. Suriadi, "Usability and Privacy in Identity Management Architectures", Proceeding ACSW '07 Proceedings of the fifth Australasian symposium on ACSW frontiers Volume 68, 2007.

[5] Ragouzis, N., Hughes, J., Philpott, R., Maler, E., Madsen, P., Scavo, T.(eds).: Security Assertion Markup Language (SAML) V.2.0 Technical Overview. OASIS Comittee Draft 02. March, 2008.

[6] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, W. Polk. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280, IETF Network Working Group, May 2008.

[7] LA.: Liberty ID-FF Protocols and Schema Specification. Available at http://www.projectliberty.org, September 2011.

[8] P. Arias, F. Almenárez, A. Marín and D. Díaz, "Enabling SAML for Dynamic Identity Federation Management", Wireless and Mobile Networking IFIP Advances in Information and Communication Technology, 2009, Volume 308/2009, pages 173-184.

[9] Information Cards.: Information Cards Foundation, 2009. Available at http://informationcard net/, September 2011.

[10] A. Nanda and M B. Jones (eds.):Identity Selector Interoperability ProfileV1.5. July 2008.

[11] XSPA Profile.: Cross-Enterprise Security and Privacy Authorization XSPA) Profile of Security Assertion Markup Language (SAML) for Healthcare Version 1.0. Available at http://docs.oasisopen.org/security/xspa/v1.0/saml-xspa-1.0-cs01.html, May 2011.

[12] Liberty Alliance Project, "An Overview of the Id Governance Frameworl". Available at http://projectliberty.org/liberty/content/download/3500/23156/file/overview-id-governance-framework-v1.0.pdf, May 2011.

[13] Liberty Alliance Project, "Privacy preference expression languages". White report. Available at http://www.projectliberty.org, May 2011.

[14] B. Pfitzmann, IBM Zurich Research Lab, Switzerland, "Privacy in enterprise identity federation policies for Liberty 2 single sign on". March 2004.

[15] D. Recordon, M. Jones, J. Bufu, J. Daugherty and N. Sakimura, "OpenID Provider Authentication Policy Extension 1.0 ". Available at http://www. openid.net, September 2011.

[16] Y.Yang and J.Yang, "Towards Unconditional Anonymity: Privacy Enforcement Model in Web Services", 2008 IEEE Congress on Services Part II, 2008.

[17] F. Gómez, J. Girao and G. Martínez, "TRIMS, a privacy-aware trust and reputation model for identity management systems", Computer Networks, Volume 54, Pages 2899-2912, November 2010.

[18] A. B. Roach. Session Initiation Protocol (SIP)-Specific Event Notification. IETF Network Working Group. RFC 3265. June, 2002.

[19] S. Cantor (ed.). "SAML V2.0 Condition for Delegation Restriction". Committee Draft 01, 10 March 2009. Available at http://docs.oasis open. org/security/saml/Post2.0/sstc-saml-delegation-cd-01.pdf, September 2011.

[20] Lasso, Liberty Alliance Single Sign-On. Availabe at http://lasso. entrouvert.org/, September 2011.

[21] Authentic: Liberty-compliant Identity Provider. Available at http:// authentic.labs.libre-entreprise.org/, September 2011.

[22] SymLabs.: ZXID: Open SAML implementation in C. Available at http://www zxid.org, September 2011.

[23] Project Sailfin: Open source Java application server project. Available at http://sailfin.java.net., September 2011.