Universidad
Carlos III de Madrid

EPS POLYTECHNIC SCHOOL

Department of Telematic Engineering

Ph.D. Dissertation

Contributions to the Privacy Provisioning for Federated Identity Management Platforms

Author: Rosa María Sánchez Guerrero

Co-Advisor: Dr. Florina Almenarez Mendoza

Co-Advisor: Dr. Daniel Díaz Sánchez

April 10, 2017

Contributions to the Privacy Provisioning for Federated Identity Management Platforms

By

Rosa María Sánchez Guerrero

Directed By

Dr. Florina Almenarez Mendoza

Dr. Daniel Díaz Sánchez

A Dissertation Submitted to the Department of Telematic Engineering and the Committee on Graduate Studies in Partial Fulfillment of the Requirements for the Degree of DOCTOR OF PHILOSOPHY

Approved by the Supervisory Committee:

General Chair

Chair

Secretary

Grade

Leganés, __ _____ _____

A mis padres, Pascual y Olimpia

A mi hermana, Olimpia

A mi "muñequete", Javier

A Sergio

. . . Cuando la gratitud es absoluta las palabras sobran.

"Dream small dreams. If you make them too big, you get overwhelmed and you don't do anything. If you make small goals and accomplish them, it gives you the confidence to go on to higher goals."

John H. Johnson

Acknowledgments

A lo largo de estos años de investigación he contado con la ayuda de muchas personas que me han acompañado y cuyas aportaciones tanto a nivel profesional como personal que han hecho posible esta tesis doctoral. Todas ellas se merecen mi más sincero agradecimiento, que quiero expresar en estas líneas.

En primer lugar, me gustaría agradecer especialmente a mis tutores, Florina y Dani su alto grado de implicación tanto en el plano personal como profesional. Gracias por la confianza depositada en mí cuando empecé a trabajar con vosotros en 2008, por haberme transmitido esa pasión por el mundo de la seguridad y la gestión de identidad (a veces muy poco valoradas), por vuestros inestimables consejos, revisiones, etc. y por haberme animado y apoyado para continuar con la tesis cuando me marché a la empresa. También quiero agradecer especialmente a Andrés por escuchar y contribuir a las diferentes propuestas e ideas que finalmente dieron como resultado esta tesis doctoral. Qué fácil te resultaba ayudarme "a pegar el corte" de aquellas líneas que tanto tiempo me habían llevado cuando los artículos excedían el límite permitido por una revista o congreso. Gracias también por eso. A Patri, por todo lo que he aprendido contigo acerca de federaciones, reputación y riesgo; y sobre todo, por haber sido una de las mejores compañeras de trabajo que he tenido hasta el momento y por ser una excelente amiga.

También me gustaría dar las gracias a los compañeros del departamento de Ingeniería Telemática. A Carlos Delgado, por darme la oportunidad de trabajar en el grupo GAST. No puedo olvidar a mis compañeros y amigos con los que he compartido despacho e incontables horas de trabajo, charlas y cafés. La lista es muy extensa y aunque no mencione nombres, vosotros sabéis que formáis parte de esto. Gracias por los buenos y malos

momentos, por aguantarme y por escucharme.

A Nicolás Jaremek, por su implementación para la gestión de historiales clínicos en situaciones de emergencia en Android.

A mis responsables en Telefónica, Azucena y Víctor, por brindarme la oportunidad de aplicar mis conocimientos de seguridad y gestión de identidad en entornos reales. A mis compañeros de trabajo Paco, Ángeles y Mauricio por escuchar y valorar mis ideas.

Dicen que los finales nunca fueron buenos, pero éste no puede ser mejor. Todo esto nunca hubiera sido posible sin el apoyo y cariño incondicional de mi familia, mis padres, mi hermana, mi sobrino Javier, mis amigos, sin el amor y el estímulo de Sergio. Esto es también vuestro premio.

Abstract

Identity information, personal data and user's profiles are key assets for organizations and companies by becoming the use of identity management (IdM) infrastructures a prerequisite for most companies, since IdM systems allow them to perform their business transactions by sharing information and customizing services for several purposes in more efficient and effective ways.

Due to the importance of the identity management paradigm, a lot of work has been done so far resulting in a set of standards and specifications. According to them, under the umbrella of the IdM paradigm a person's digital identity can be shared, linked and reused across different domains by allowing users simple session management, etc. In this way, users' information is widely collected and distributed to offer new added value services and to enhance availability. Whereas these new services have a positive impact on users' life, they also bring privacy problems.

To manage users' personal data, while protecting their privacy, IdM systems are the ideal target where to deploy privacy solutions, since they handle users' attribute exchange.

Nevertheless, current IdM models and specifications do not sufficiently address comprehensive privacy mechanisms or guidelines, which enable users to better control over the use, divulging and revocation of their online identities. These are essential aspects, specially in sensitive environments where incorrect and unsecured management of user's data may lead to attacks, privacy breaches, identity misuse or frauds.

Nowadays there are several approaches to IdM that have benefits and shortcomings, from the privacy perspective.

In this thesis, the main goal is contributing to the privacy provisioning for federated identity management platforms. And for this purpose, we propose a generic architecture that extends current federation IdM systems. We have mainly focused our contributions on health care environments, given their particularly sensitive nature. The two main pillars of the proposed architecture, are the introduction of a selective privacy-enhanced user profile management model and flexibility in revocation consent by incorporating an event-based hybrid IdM approach, which enables to replace time constraints and explicit revocation by activating and deactivating authorization rights according to events. The combination of both models enables to deal with both online and offline scenarios, as well as to empower the user role, by letting her to bring together identity information from different sources.

Regarding user's consent revocation, we propose an implicit revocation consent mechanism based on events, that empowers a new concept, the *sleepyhead* credentials, which is issued only once and would be used any time. Moreover, we integrate this concept in IdM systems supporting a delegation protocol and we contribute with the definition of mathematical model to determine event arrivals to the IdM system and how they are managed to the corresponding entities, as well as its integration with the most widely deployed specification, i.e., Security Assertion Markup Language (SAML).

In regard to user profile management, we define a privacy-awareness user profile management model to provide efficient selective information disclosure. With this contribution a service provider would be able to accesses the specific personal information without being able to inspect any other details and keeping user control of her data by controlling who can access. The structure that we consider for the user profile storage is based on extensions of Merkle trees allowing for hash combining that would minimize the need of individual verification of elements along a path. An algorithm for sorting the tree as we envision frequently accessed attributes to be closer to the root (minimizing the access' time) is also provided.

Formal validation of the above mentioned ideas has been carried out through simulations and the development of prototypes. Besides, dissemination activities were performed in projects, journals and conferences.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Digital identity can change balance sheets and change our future. But solutions need to put the customer front and center. Only when there is trust that organisations are handling information responsibly and providing sufficient individual benefit, will data be shared in a sustainable way.

Liberty Global, 2012

Contents

1.1 Motivation

The advent of the digital era and the growth of the Internet have improved the experience of organizations and users when storing, analyzing and exchanging personal data. According to [1] digital identity management is a critical component and a key pillar to reduce

complexity and enhance user experience in economic, governmental and social activities. Modern federated IdM infrastructures enable users to share attributes among different participants, typically between Identity Providers (IdPs), Service Providers (SPs) and users. Thus, SPs can link a person's electronic identity stored across multiple distinct domains and offer personalized services according to the user preferences.

In this way, identity information, user's profile and sensitive personal data have become very valuable assets for companies and organizations. On the one hand, they allow customization of services in a cost effective and faster manner. On the other hand, interactions and transactions can be more effective. So, in a digital context, users' information is extensively collected and distributed to provide new added value services and to improve availability. Whereas these new services have a positive impact on users' life, they also bring privacy problems.

In this thesis, we understand privacy as the right that every person has to control access to her own personal information.

For instance, in sensitive scenarios, such as health care environments, in the last few years many initiatives have developed new ways to manage, share and organize medical information. In such a way users can manage their data easily keeping, storing, sharing or disposing, for example, their clinical diagnosis. Apart from these functionalities, emergency access to sensitive information like health records is sometimes needed by carers otherwise unrelated to the normal care of a patient. Such accesses can only be consented in a general way, since the specific providers involved will not usually be known in advance. Furthermore, due to its volume, information is stored locally and/or fetched from third parties using cross domain services. However, misuse and unauthorized access to such information may violate user's privacy, cause fraud or even crimes that may harm people's health.

For these reasons, an adequate federated identity management system is necessary, since the success of these kind of systems is directly related to users' confidence in the system to manage their personal data, while preserving their privacy. For this purpose, **IdM systems are the preferred target where to deploy privacy solutions**, since they mediate in every users' information exchange.

In this sense, nowadays we are experiencing a duality in our lives. There is our physical

identity and there is our digital counterpart [2], which brings us to an important moment in history for privacy on the Internet. Modern cloud-based platforms and social networks are intensifying human interactions online whereas they are handling huge amounts of our personal data. Moreover, a growing number of "digital natives" are actively broadcasting parts of their lives through social networks. However, average users rarely have privacy concerns, and feel comfortable accepting default privacy settings enforced by these services. Most ignore the risks as minimal tradeoffs in comparison to the benefits of accessing services. This is specially worrying in youngsters, unconscious about the consequences of spreading personal information over the Internet.

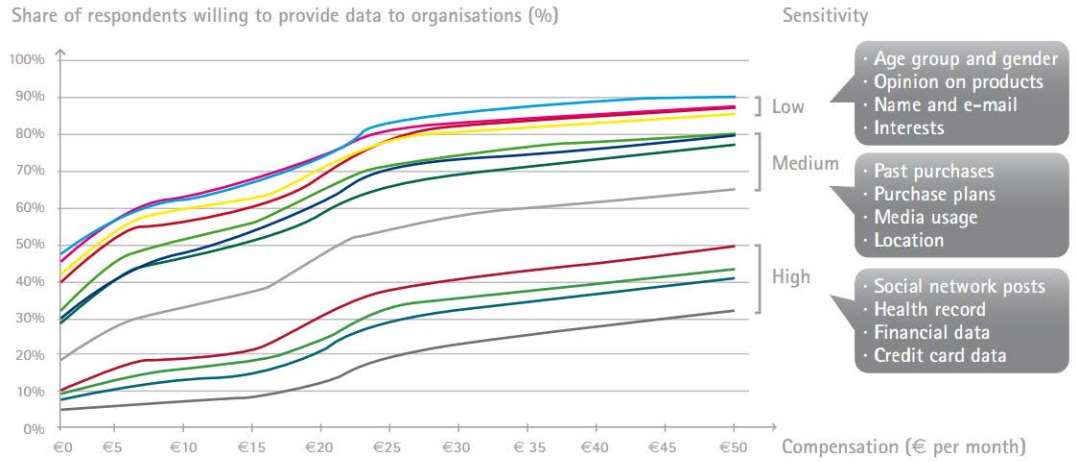

Figure 1.1: Sensitivities of data types for consumers (©[2]).

In the above mentioned, we cannot underestimate the value placed on our personal data and need to appreciate why privacy is an essential part of our lives both on and offline. The study realized in [2] and represented in Figure 1.1 illustrated that while low-sensitivity data (e.g., age or gender) was shared by more than 40% of survey consumers even without receiving a direct benefit, most participants were unwilling to part with highly sensitive information (like a health record) even for large rewards. Remarkably, social network posts belong to the category of highly sensitive personal data.

To the individual, privacy is more than keeping sensitive information hidden. The identities we have developed online are complex and becoming increasingly important to us. Sometimes, we look for complete anonymity and at others, we change our digital personalities depending on who we are interacting with. For many, full disclosure of our digital selves would be unthinkable.

Summarizing, individuals want to have control of their information since the improper and unsecured management of their information 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. Nevertheless, the complexity of some privacy agreements exposed by services providers and the increasing number of participants overwhelm users. This issue can be appreciated in Figure 1.2, which shows the relationship between expressed concern and privacy valuation. This "weak link" is not unprecedented, arising that the disconnect between individuals' concern and behavior applies to the broader array of data-sharing scenarios.

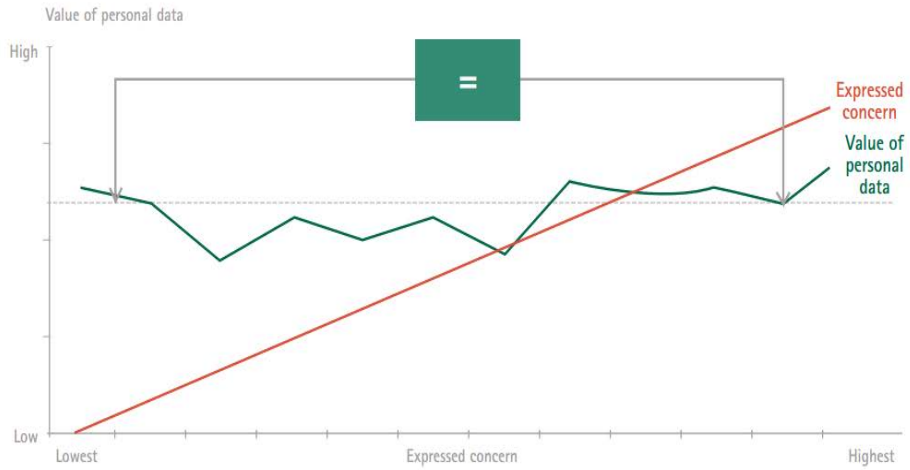

Figure 1.2: Expressed level of concern vs. valuation of personal data (©[2]).

The motivation behind this problem is the lack of comprehensive privacy frameworks. Despite IdM systems are the ideal site to deploy privacy mechanisms, current IdM specifications are not ready to deal with divulging and revocation of users' online identities.

It is essential to realize mechanisms for fine-grained control of sensitive data, storage and maintenance, which enable controlled sharing across different stakeholders, while data protection against unauthorized use and minimal disclosure according to user's consent preferences is provided.

Moreover, user's consent revocation fits with the privacy view as control over the use and flow of one's personal information [3]. Revoking consent allows users grant or withdraw consent of specific actions over data to certain individuals. So, this mechanism is useful

to enforce user's role in the task of preserving her privacy.

With these requirements in mind and aiming at contributing to improve the identity landscape, and more specifically the above privacy aspects, we provide some **contributions to the privacy provisioning for federated identity management platforms**.

1.2 Objectives

The overall objective of our research is oriented to **extend the federated identity management infrastructures with an adequate supply of the users' privacy at different levels while required security services are covered**. As Harriet Tubman said: ''*Every great dream begins with a dreamer. Always remember, you have within you the strength, the patience, and the passion to reach for the stars to change the world".*

Therefore, to achieve our objective, we will pursue the following specific goals:

- **Realizing a study of the state of the art on current techniques used to protect the user's privacy in sensitive and open environments**, among which include anonymity, semi-anonymity, pseudonyms, etc., identifying their shortcomings and limitations.
- **Analyzing the research challenges in privacy and security**, paying special attention to data protection models that could be applied in sensitive and open environments.
- **Analyzing security and privacy risks** in the different open and distributed environments to detect potential threats, such as identity theft, information leakage, manipulation of information, etc.
- **Designing a privacy and security model** that covers the requirements identified in the previous analysis.
- **Integrating the proposed privacy and security model within an identity management infrastructure** to optimize the ability of the system to use data and minimize risks to user's privacy and data integration in a flexible and distributed environment. This model will offer users greater awareness of the use of their digital identity online by introducing monitoring systems in order to enable users to

balance security, privacy and usability depending on their needs. Therefore, our contribution will allow traceability of data and user interaction in the system to support audits. In addition, delegation, **flexible and time-independent revocation consent mechanisms** will be included, as well as **representation schemes and management of user credentials and attributes**, to provide efficient **selective information disclosure** in dynamic environments and emergency scenarios.

Furthermore, as it is also an important part in the development of a doctoral thesis, we also aim to achieve the following goals:

- Evaluation and validation of the contributions. Such validation will be mainly performed on health care scenarios. Other test scenarios, such as cloud computing environments will be also considered. This will also allow us to verify that the thesis contributions have practical applicability in the above scenarios. For instance, the simulations with real data from hospital emergency services are considered relevant to be environments where extremely sensitive information is handled.
- Dissemination of the results through publication, collaboration in research projects and participation in conferences.
- Identification of new lines of research that can be derived from this work.
- Completion of the writing and public defense of the thesis dissertation.

1.3 Development Plan

This section describes the work methodology to carry out the elaboration of the doctoral thesis, including some general and key points for its development:

1. Gather the bibliography related to identity management in order to study and analyze the existing gaps in regard to privacy-enhancing techniques.
2. Design an architecture with the necessary elements to permit effective consent revocation and selective disclosure of users' identities.
3. Design an event-based model, which allows to substitute time constraints and explicit revocation by activating and deactivating authorization rights according to events. Our approach is to **integrate this concept in IdM systems**, which can be an

interesting alternative for scenarios where revocation of consent and user privacy are critical.

4. Design and develop a flexible, efficient and standards-based solution to guarantee selective identity information disclosure and preserve user's privacy. We propose **a privacy-aware profile management approach that empowers the user role, allowing her to amalgamate several service providers as well as user-generated claims into a single credential**.

5. Perform evaluation and validation tests of the designed models in order to demonstrate the benefits of the proposal and its feasibility.

6. Obtain the main conclusions from the performed research work and identify new research lines to be followed.

7. Write and publish papers with the partial results that are obtained during the different phases of the research.

1.4 Interest of the Research

In regard to publication and dissemination, the content of this thesis was developed as a research line in several national R&D projects, called "España Virtual"[1], CONSEQUENCE[2], EMRISCO (EMergency Response In Smart COmmunities) and INRISCO (INcident MonitoRing In Smart COmmunities)[3]. The España Virtual and CONSEQUENCE projects included specific working packages for "Security and Identity Management", whilst the EMRISCO and INRISCO research projects incorporated the Work Package 3 for "Security and Privacy"; where our ideas on security and privacy were contributed.

Moreover, dissemination was also accomplished through publication of scientific papers. The main papers that support the interest of the research presented in this thesis are detailed below. For each contribution, we briefly explain the kind and date of publication (i.e., whether conference or journal) and its contents, showing which part of the dissertation they support. It is to note that all the journal papers correspond to journals indexed in the JCR. We also reference other complementary works we have published that, though

[1] http://www.espanavirtual.org/
[2] http://consequence.it.uc3m.es/
[3] http://www.inrisco.org/

they do not deal with core aspects of this thesis, are derived from the ideas presented here (e.g., application scenarios, use-cases, etc.) and integrated in those works. The criterion for ordering the results is their relevance to the dissertation, so more relevant papers are listed first.

Main Contributions to Journals:

1. Title: *Enhancing Privacy and Dynamic Federation in IdM for Consumer Cloud Computing.*
   Authors: R. Sánchez-Guerrero, F. Almenárez, P. Arias, D. Díaz-Sánchez and A. Marín.
   Journal: IEEE Transactions on Consumer Electronics ISSN: 0098-3063. Printed version in Vol. 58, Iss. 1, pp. 95 - 103, February 2012. Impact Factor as of 2012: 1.087. Q2, category: Electrical and Electronic Engineering [4].

   This work presents an identity management architecture based on privacy and reputation extensions compliance with the SAMLv2 standard. The document, extended from our conference paper in [5], analyzes the main identity management and privacy challenges to be tackled in consumer cloud computing. Likewise, the proposal provides modules that enable users to access cloud services and share digital content without disclosing their real identity and to have enhanced awareness over their identity through audit and monitoring services.

2. Title: *An Event Driven Hybrid Identity Management Approach to Privacy Enhanced e-Health.*
   Authors: R. Sánchez-Guerrero, F. Almenárez, D. Díaz-Sánchez, A. Marín, P. Arias and F. Sanvido.
   Journal: Sensors, 12(5) pp. 6129-54, May 2012. Impact Factor as of 2012: 1.953. Q1, category: Instruments & Instrumentation [6].

   In this paper we analyze the main current identity models to preserve privacy in identity management systems, when these are applied to sensitive scenarios such as e-health, parental control, etc. We focused on the issue of effective consent revocation, which is not supported by any of the models, presenting an event-based mechanism enabling a new concept, the *sleepyhead* credentials, which enables to achieve a more flexible revocation consent model by activating and deactivating

authorization rights according to events. The contributions of the paper include a hybrid model supporting delegation compliance with the SAMLv2 standard, a mathematical model describing the event-based model, and an evaluation of the overhead introduced by the system.

3. Title: *Collaborative eHealth meets Security: Privacy-Enhancing Patient Profile Management.*
   Authors: Rosa Sánchez-Guerrero, Florina Almenárez, Daniel Díaz-Sánchez, Patricia Arias and Andrés Marín.
   Journal: Accepted for publication (January 2017) in the IEEE Journal of Biomedical and Health Informatics (J-BHI). Impact Factor as of 2016: 2.093. Q1, category: Health Information Management.

   In this article we propose a privacy-awareness profile management approach based on a novel Adaptive Extended Merkle structure and compliant with EHR standards, that promotes the user role to realize a selective disclosure of her information to the different health stakeholders. Thus, the contribution empowers users to combine several healthcare providers as well as user-generated claims into a single credential, while avoiding the creation of bogus patient's EHR profiles. To achieve this, the work defines and evaluates a light structure that can be saved on handled devices, as well as adapting efficiently to changes over time and enriching compositions of the patient's medical history, thanks to the designed algorithm.

Main Contributions to International Conferences:

1. Title: *Improving Privacy in Identity Management Systems for Healthcare Scenarios.*
   Authors: Rosa Sánchez-Guerrero, Florina Almenárez, Daniel Díaz-Sánchez, Andrés Marín and Patricia Arias Cabarcos.
   Conference: International Symposium on Ubiquitous Computing and Ambient Intelligence (UCAmI'11). Rivera Maya, Mexico, December 05-09, 2011 [7].

   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 SAML-based delegation protocol with the aim to improve the revocation consent within healthcare scenarios.

2. Title: *A model for dimensioning a secure event-driven health care system.*

Authors: Rosa Sánchez-Guerrero, Florina Almenárez, Daniel Díaz-Sánchez, Patricia Arias Cabarcos and Andrés Marín.
Conference: Wireless and Mobile Networking Conference (WMNC 2012). Bratislava, Slovakia, September 19-21, 2012 [8].

This works builds on the flexible event-based user consent revocation mechanism defined in [6] for health care scenarios. It extends the network dimensioning for the event-driven hybrid model estimating the SIP-Event Notify message overhead generated during the system operation, that means, to calculate the introduced overhead by the subscription and notification event messages exchanged between the entity system in order to activate/deactivate the necessary attributes and privileges, depending on the different health care events that arrive to the system. For this purpose, we used a set of statistical data collected by the HES (Hospital Episode Statistics) online service[4] and considered two main simulation scenarios: a large hospital and a small-medium hospital.

3. Title: *An Identity Aware WiMax Personalization for Pervasive Computing Services.*
Authors: Rosa Sánchez-Guerrero, Daniel Díaz-Sánchez, Florina Almenárez, Andrés Marín, Patricia Arias Cabarcos and Davide Proserpio.
Conference: International Symposium on Ubiquitous Computing and Ambient Intelligence (UCAmI'11). Rivera Maya, Mexico, December 05-09, 2011 [9].

   This article proposes the introduction of identity management with privacy extensions in WiMAX, as a previous step to the definition of identity aware WiMax personalization of pervasive computing services. Personalized content can be achieved through a SAML-based IdP, which includes services for user profile management and device profile selection. The user, device and subscription profile are associated under a WiMAX session linked to a specific user identity.

4. Title: *Introducing Infocards in NGN to enable user-centric identity management.*
Authors: D. Proserpio, F. Sanvido, P. Arias, R. Sánchez, D. Díaz-Sánchez, A. Marín and F. Almenares.
Conference: IEEE Global Communications Conference (GLOBECOM 2010). Miami, Florida, USA, December 06-10, 2010 [10]

[4] http://www.hscic.gov.uk/

This work seeks to improve the navigation experience and security in multiservice and multiprovider environments the user must be empowered to control how her attributes are shared and disclosed between different domains. Thus, a manner to combine the gains of a SAML federation between service and identity providers with the easiness for the final user of the Inforcard System using the well known architectural schema of IP Multimedia Subsystem is explained.

Related Contributions to Journals and International Conferences:

In addition, the following publications complement the core ideas in the above mentioned papers by the definition and the integration of application scenarios:

1. Title: *Media Gateway: bringing privacy to Private Multimedia Clouds connections.*
   Authors: D. Díaz-Sánchez, F. Almenárez, A. Marín, R. Sánchez-Guerrero and P. Arias.
   Journal: Telecommunication Systems, Vol. 55, Issue 2, pp 315-330, February 2014. Impact Factor as of 2014: 0.705 [11].

   This article describes a solution that enables limited devices to access contents located in private clouds, with the cooperation of network providers. It includes a comprehensive and efficient solution for managing content among federated home environments. As part of the purpose of empowering the user role as well as to improve user experience, we placed significant efforts on interoperability and privacy protection when it comes to accessing cloud resources from other networks.

2. Title: *FamTV: An architecture for Presence-Aware Personalized Television.*
   Authors: P. Arias-Cabarcos, R. Sánchez, F. Almenares, D. Díaz-Sánchez, A. Marín.
   Journal: IEEE Transactions on Consumer Electronics. ISSN: 0098-3063. Printed version in Vol.57, no.1, pp.6-13, February 2011. Impact Factor as of 2011: 0.941. Q2, category: Electrical and Electronic Engineering. [12].

   In this article, extended from our conference paper in [13], presents a presence-aware personalized architecture, which enables to combine the advantages of content-filtering and presence-aware technologies for personalization. It provides a security and privacy layer to establish privacy-enabled configurations and rules. These elements are desired, since TVs are usually located in a common place at home and most of the time there are more than one viewer. In this way, whether a user is

watching TV, her personal widgets, such as social network comments, are configured to become visible in the screen. When another user enters into the room, her presence is recognized and the elements considered private are automatically hidden. This work received the **Chester W. Sall Award** for the 2nd place best paper in the IEEE Transactions on Consumer Electronics 2011.

3. Title: *A H.264 SVC Distributed Content Protection System with Flexible Key Stream Generation.*
   Authors: Daniel Díaz-Sánchez, Rosa Sánchez-Guerrero, Andrés Marín López, Florina Almenares and Patricia Arias.
   Conference: IEEE International Conference on Consumer Electronics (ICCE-Berlin 2012). Berlin, Germany, September 03-05, 2012 [14].

   This article describes a distributed system for content encoding and protection that generates a flexible key stream that simplifies the receiver. The proposed key management scheme is based on Merkle hash trees generalized by a Markov's state chain, based on a directed acyclic graph, which allows to provide the key streams as a graph in which every node represents a video quality and depends on the previous lower qualities using a Markov's state chain. In such a way, a receiver does not need to decode the entire graph, by inferring the probability of selecting one path or another.

4. Title: *Family Personalization Service.*
   Authors: D. Díaz-Sánchez, R. Sánchez-Guerrero, P. Arias-Cabarcos, I. Bernavé and F. Almenares.
   Conference: IEEE International Conference on Consumer Electronics (ICCE-Berlin 2011). Berlin, Germany, September 06-08, 2011 [15].

   This paper, builds on the ideas in [12], and describes a personalization system which empowers to automatically configure devices surrounding users. The system deals with privacy-based filtering and group preference modeling.

1.5 Organization of the Thesis

With the aim to achieve the goals outlined in the above sections, the organization of this dissertation is as follows:

Chapter 2 presents the state-of-the-art on technologies and latest research related to the thesis. It objectively reviews the different existing identity management models; provides an overview of techniques for privacy-awareness profiles management, a background of current specifications for managing of personal information focused on e-health standards; and summarizes related work being carried out by individual researchers, international research projects and organizations involved in standardization.

Chapter 3 proposes a generic infrastructure to solve the privacy limitations of current identity management models. Based on this high-level infrastructure description, **Chapters 4**, **5**, **6** and **7** go deeper into the main components and extensions of the architecture.

More specifically, in **Chapter 4**, we give a comparative analysis of the privacy support in identity management systems emphasizing the importance of the revocation consent property. Besides, we propose an event-based mechanism enabling implicit revocation of user's attributes and rights, and formalize a mathematical model to represent the event-driven system behavior. **Chapter 5** is dedicated to the validation of the ideas presented in the previous chapter.

Next, **Chapter 6** describes our approach for selective privacy-enhanced user profile management, which enables to construct enriched compositions of user's profile. An algorithm to sort the proposed structure based on patterns of access compliance with standards for personal information management is also described. Moreover, in **Chapter 7** we explain the evaluation carried out to validate the selective privacy-enhanced user profile management proposal.

Finally, **Chapter 8** summarizes the results and discussions presented in this thesis. Furthermore, since the need for further work and exploration is necessary in any useful research, we also describe the future lines that can be followed from the ideas presented here.

Apart from the the aforementioned chapters, we have included **Appendix A**, which contains a glossary with all the acronyms used in the document.

Chapter 2

State of the Art

Contents

2.1 Identity Management Models

According to Jøsang et al. [16] 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 main privacy principles: anonymity, pseudonymity, unlinkability, unobservability, selective disclosure and revoking consent. A detailed definition of the above privacy principles will be provided in section 2.2.2. Furthermore, although the property of anonymity is one of the main principles of privacy, it would be desirable that IdM systems support mechanisms to break the anonymity of a user for the purpose of analysis or evidence under certain circumstances (e.g., a malicious user, lawful interception).

For clarity, we introduce here the main actors in an identity management scenario. (i) the Service Provider (SP)[1], which provides services and takes decisions about a particular subject based on the identity information provided by (ii) the Identity Provider[2], that authenticates users, manages identity information and shares identity information with various SPs upon user request. (iii) The Principal, or the End User, who has a particular digital identity and interacts (usually via an user agent) with both SPs and IdPs.

Identity models can be categorized into the following three styles: federated identity, user-centric identity and hybrid identity models.

In the federated identity model, user or identity data are distributed across various identity providers, which have a trust relationship among each other. Such trust relationships are based on agreements between the SPs and IdPs belonging to the federation and they are usually established on organizational level, whereas enforcement is carried out on technical level.

[1]The term Relying Parties is also frequently used to refer to SPs
[2]The term Asserting Parties is also used to refer to IdPs

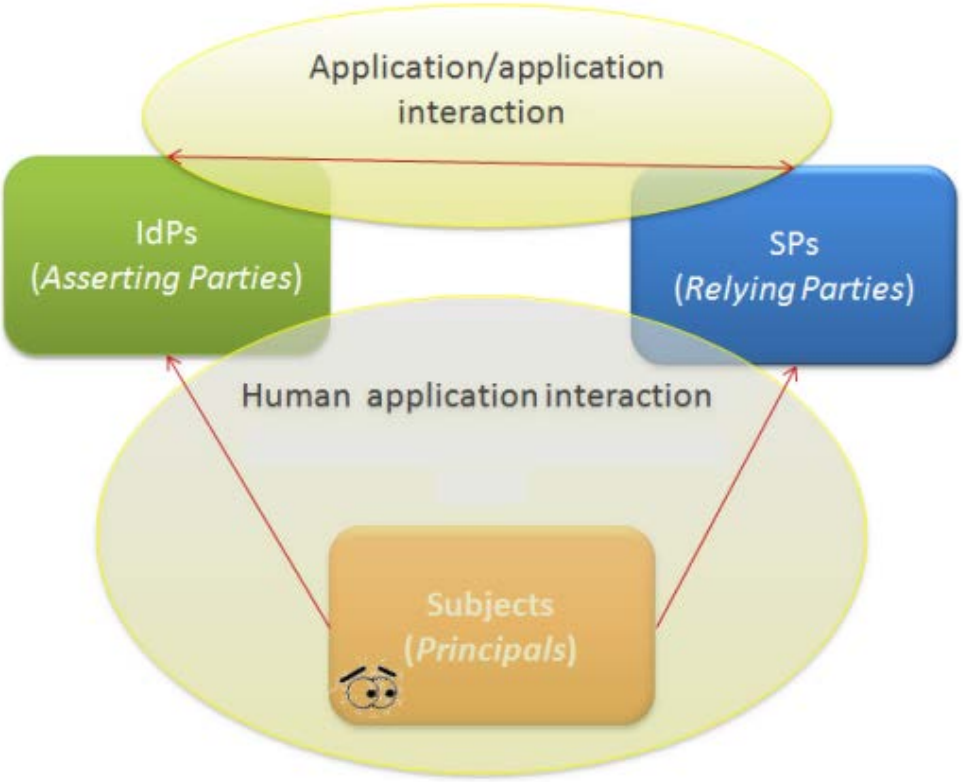

Figure 2.1: Federated identity model: involved entities and interactions between them.

Thus, the federation identity model enables users of one domain to securely access resources of another domain seamlessly, without the need for redundant user login processes. But also it provides means for cross-domain user account provisioning, cross-domain entitlement management and cross-domain user attribute exchange. However, the user does not actively participate in the above processes, since all interactions between the applications for authentication and attribute exchange, are carried out between SPs and IdPs, as can be seen in Figure 2.1.

In its turn, as shown in Figure 2.2, the user-centric identity model situates the user in the middle of a transaction. So, the user is no longer aside of the trust establishment, authentication and attribute exchange processes. Nevertheless, user's attribute exchange process cannot be completely user-centric, since in some scenarios the user is not always online to grant her consent. On the other hand, federated identity models raise privacy concerns because user's identity information may be available to every entity belonging to

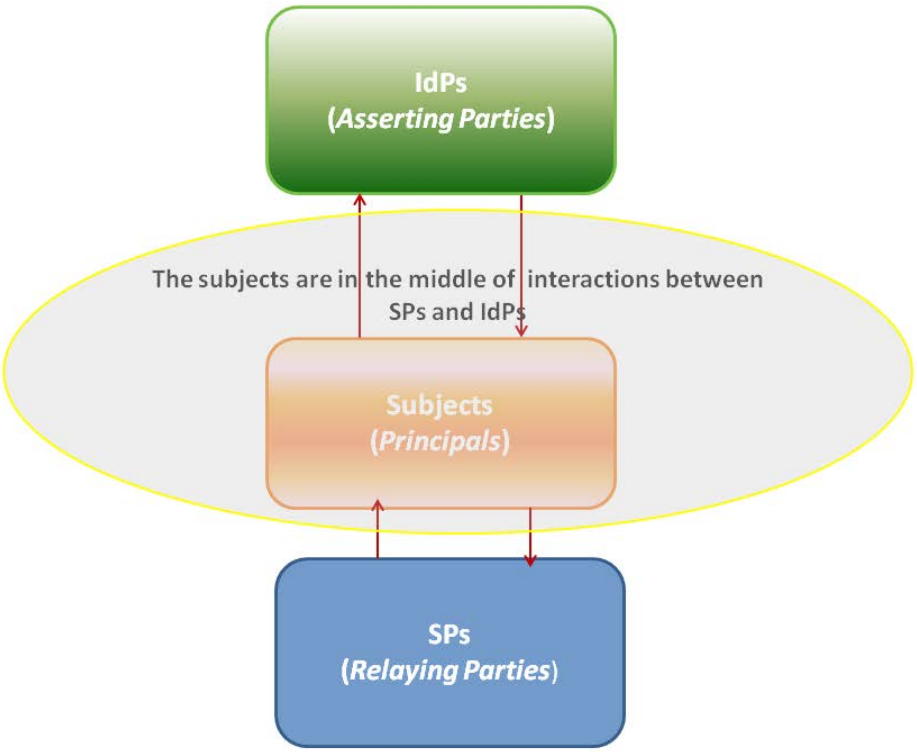

Figure 2.2: Interactions and involved entities for the user-centric identity model. Note that, the user is placed in the middle of transactions between SPs and IdPs.

the federation.

Therefore, a combination of both models is necessary. In this thesis, we will use the term hybrid identity model to refer to the blend of the federated and user-centric approaches.

2.1.1 Federated Identity Models

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. Federated models bring user attribute exchange, user account provisioning, entitlement management and personalized service provisioning. The main existing federation protocols and specifications are detailed below.

Security Assertion Markup Language (SAML)

Security Assertion Markup Language (SAML) [17] is an standard developed by the Security Services Technical Committee (SSTC) of the standards organization OASIS (the Organization for the Advancement of Structured Information Standards). SAML v2.0 [18] is announced as OASIS standard in March 2005 and it represents the convergence of Liberty Identity Federation Framework (ID-FF) [19] and other identity management initiatives, as Internet2 Shibboleth [20] or Web Services Security (WSS) OASIS [21]. However, despite SAML was standardized in 2005, corrections and new specifications (rectification of August 2007 "Approved Errata for SAML V2.0" draft technical review in March 2008, etc.) have emerged.

SAML defines an XML-based framework to allow the exchange of security assertions (about authentication, authorization decision, and attributes) between entities. The aim is to establish open standards to easily conduct online transactions while protecting the privacy and security of identity information. These standards enable identity federation and management through features such as account linkage, and profiles, especially for the simple session management. The SAML standard defines four key elements:

- **Assertions** [22], which define security claims of an entity within a system. These statements can be of three types: authentication assertions, attribute and authorization decisions. Authentication assertions indicate that the user has been authenticated by an application, whereas attribute assertions are more specific and include information about the user attributes (e.g., Alice belongs to a certain group with more privileges). Finally, assertions authorization decision, define something that the subject is entitled to do (e.g., Alice has permission to buy certain product). These assertions and other standard components are based in XML, which allows to be used them in other contexts. In order to better illustrate the format of a SAML Assertion, an XML fragment containing an example assertion with authentication and attribute statements is provided in Figure 2.3:
- **Protocols** [22], which define how assertions are requested.
- **Bindings** [23], which define the lower-level communication or messaging protocols (such as HTTP-GET, HTTP-POST, Simple Object Access Protocol (SOAP), Reverse SOAP (PAOS), etc.) that the SAML protocols can be transported over.

```
<saml:AuthnStatement AuthnInstant="2009-06-17T18:45:10.738Z">
   <saml:AuthnContext>
      <saml:AuthnContextClassRef>urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified
      </saml:AuthnContextClassRef>
   </saml:AuthnContext>
</saml:AuthnStatement>

<saml:AttributeStatement>

   <saml:Attribute Name="portal_id">
      <saml:AttributeValue xsi:type="xs:anyType">060D00000000SHZ
      </saml:AttributeValue>
   </saml:Attribute>

   <saml:Attribute Name="organization_id">
      <saml:AttributeValue xsi:type="xs:anyType">00DD0000000F7L5
      </saml:AttributeValue>
   </saml:Attribute>

   <saml:Attribute Name="ssostartpage"
      NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-format:unspecified">

      <saml:AttributeValue xsi:type="xs:anyType">
         http://www.salesforce.com/security/saml/saml20-gen.jsp
      </saml:AttributeValue>
   </saml:Attribute>

   <saml:Attribute Name="logouturl"
      NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-format:uri">

      <saml:AttributeValue xsi:type="xs:string">
         http://www.salesforce.com/security/del_auth/SsoLogoutPage.html
      </saml:AttributeValue>
   </saml:Attribute>
</saml:AttributeStatement>
```

Figure 2.3: Example of assertion with authentication and attribute statements.

- **Profiles** [18], which are combinations of SAML protocols and bindings, together with the structure of assertions to cover specific use-cases. For each combination of use case and binding, some variations of the same profile can be obtained. Eight different associations and specific transport protocols for a use case (Web SSO profile, Attribute profile, etc.) are contemplated by the standard.

The combination of the aforementioned building-block components enables a set of use-cases to be supported. The most relevant use case for which SAML is applied is multi-domain web Single Sign-on (SSO), which allows a user to authenticate at a single site and gain access to other sites without the need for re-authentication, i.e., reusing the same identifier, act of authentication, and login session across multiple sites. The Single LogOut (SLO) use-case enables near-simultaneous logout of active sessions associated with a principal through SOAP, HTTP Redirect, HTTP POST or HTTP Artifact bindings. Moreover, the logout can be directly initiated by the user, or initiated by an IdP or SP because of a session timeout, administrator command, etc.

Another important profile is the SAML Enhanced Client or Proxy (ECP) [24], which is a single sign-on authentication profile that specifies a client application capable of directly determining and contacting the user's IdP, without getting redirected by the SP. It is par-

ticularly useful for client-side and server-side applications with a fixed set of services. The ECP profile focuses on applications with enhanced functionality, for instance, supporting more sophisticated protocols and bindings (e.g., SOAP and PAOS). This profile also allows to empower the user or principal role, as well as to improve user's privacy by minimizing direct interactions between SPs and IdPs.

Apart from the aforementioned SSO use cases, SAML covers a huge range of use-cases, namely:

- **Federation via Persistent Pseudonym Identifiers**: In this use-case, an IdP federates the user's local identity principal with the principal's identity at the SP by means of a persistent SAML name identifier.
- **Federation via Transient Pseudonym Identifiers**: A temporary identifier is used to federate between the IdP and the SP for the life of the user's web SSO session.
- **Federation via Identity Attributes**: In this use-case, attributes of the principal, as defined by the IdP, are used to link to the account used at the SP.
- **Federation via Out-of-Band Account Linking**: The establishment of federated identities for users and the association of those identities to local user identities can be performed without the use of SAML protocols and assertions.
- **Federation Termination**: This use case consists of the termination of an existing federation.

We describe below the first two use cases, as they are the most relevant here.

In the Federation via Persistent Pseudonym Identifiers scenario sketched out in Figure 2.4, the processing is as follows:

1. John attempts to access a resource on the SP (*cars.example.co.uk*). John does not have any security context on this site and is unknown to it. The resource that the user attempted to access is saved as `RelayState` information.
2. The SP sends to the IdP (*airline.example.com*) a SAML `<AuthnRequest>` message requesting that the IdP provide an assertion using a persistent name identifier for John. As the service provider desires the IdP have the flexibility to generate a new

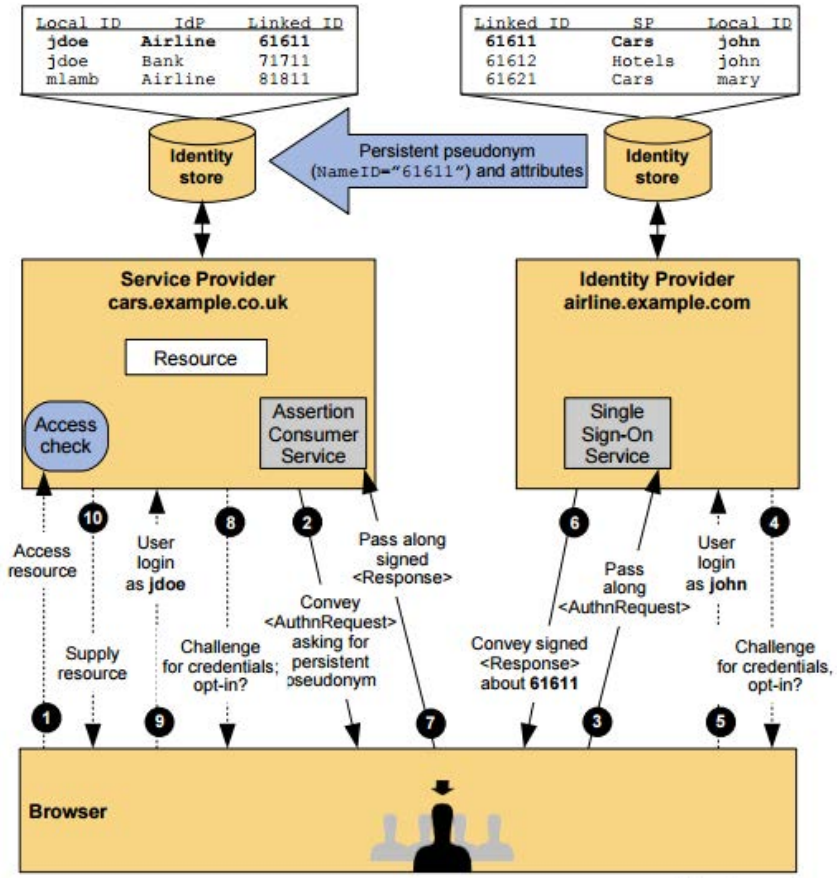

Figure 2.4: Identity Federation with Persistent Pseudonym Identifiers example (©[18]).

identifier for the user should one not already exist, the SP sets the `AllowCreate` attribute on the `NameIDPolicy` element to "true".

3. In Steps 3 and 4 the user authenticates as *john* and a local security context is created for the user at the IdP.

4. In Step 5 the IdP looks up user *john* in its identity repository. Then it verifies the `AllowCreate` attribute and creates a persistent name identifier (61611) to be used for the session at the SP. Then the IdP builds a signed assertion where the subject uses a transient name identifier format. The name *john* is not contained anywhere in the assertion.

5. The SP validates the digital signature and assertion on the SAML Response in Steps 6 and 7. Moreover, the supplied name identifier is then used to determine if a previous federation has been established. Whether a previous federation has been established

then go to step 9. If no federation exists for the persistent identifier in the assertion, then the SP needs to determine the local identity to which it should be assigned. The user will be challenged to provide local credentials at the SP. Optionally the user might first be asked whether he would like to federate the two accounts.

6. A local session is created for user *jdoe* and an access check is then made to establish whether the user *jdoe* has the correct authorization to access the desired resource at the SP.

7. Finally, if the access check is correct, the desired resource is returned to the user agent (e.g., a web browser.).

Regarding, the Federation via Transient Pseudonym Identifiers use-case, the difference from the previous scenario is that, the IdP creates a transient name identifier (294723) for John to be used for the session at the service provider (See Figure 2.5). Then, the IdP builds a signed assertion where the subject uses a transient name identifier format. Thus, the Federation via Transient Pseudonym Identifiers use-case avoids having to manage user identifiers and passwords at the SP and enables users partial anonymity, since the IdP knows which user corresponds to each identity.

Summarizing, in regard to privacy, 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 SPs. 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 Cross-Enterprise Security and Privacy Authorization (XSPA)-SAML profile [25]. Nevertheless, SAML standard states that privacy must be considered, but concrete decisions are left to the implementors.

Liberty Alliance

The Liberty Alliance initiative began in September 2001, after the union of a consortium of companies, suppliers and institutions with a common interest: to provide standards for federated identity management. The federation model of Liberty, provides guarantees of privacy and security as well as an open and standard single registration mechanism. The development of the set of standards and recommendations has been divided into three

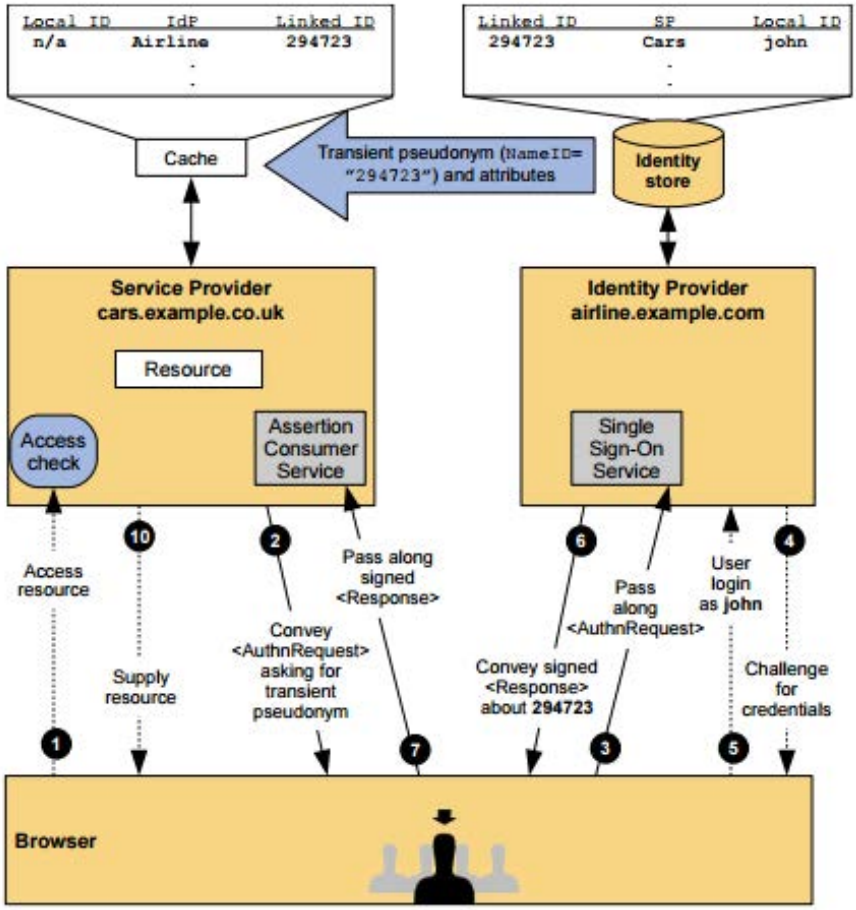

Figure 2.5: Identity Federation with Transient Pseudonym Identifiers example (©[18]).

phases:

- **Phase 1- Liberty Identity Federation Framework:** Proposes the use of federated network identity for troubleshooting network identity. Identity Federation Framework was the first specification of Liberty and defines the roles of the participating entities in a federated identity environment, as well as the required protocols to perform tasks related to identity federation, single sign on, use of pseudonyms global and single logout. The concept of "Circle of Trust" (CoT) is a key pillar, which is defined by Liberty as a federation of SPs and IdPs that have business relationships based on Liberty architecture and operational agreements and with whom users can transact business in a secure and apparently seamless environment.

  The specifications that make up the ID-FF module, reflect the description architecture) [19], abstract protocols, XML Schema Definitions (XSDs) [26], the recom-

mended deployment lines, specification of mandatory and optional requirements, etc.

- **Phase 2- Identity Web Services Framework (ID-WSF):** It is based on ID-FF to provide a framework for web services based on federated identity. It provides the necessary mechanisms to share attributes based on permissions, discovery service profiles and user interaction. The specification includes support of SAML and allows to implement web services in a standard environment ensuring end-to-end security.

- **Phase 3- Identity Services Interface Specifications (ID-SIS):** ID-SIS, meanwhile, uses both ID-FF and ID-WSF to provide network services. ID-SIS specifics web services interfaces that support high-level particular use cases, such as profiles, geographical location, presence, etc.

Apart from the core frameworks defined above, during the last years of the project Liberty Alliance also released Identity Governance Framework (IGF) [27] dealing with identity governance and privacy issues. In February 2007, the Liberty Alliance started to work on the IGF, releasing the first version publicly in July 2007.

The Identity Governance Framework defines a set of standards to help enterprises easily determine and control how identity related information is used, stored, and propagated in appropriate and secure ways. IGF 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 (ARML) 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 their usage. Furthermore, the SP could request to modify the data or forward it to another SP. However, in [28], 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.

WS-Federation

The Web Service Federation Language or WS-Federation [29] is an OASIS standard that forms part of the larger Web Services Security framework (WS-*). More specifically, WS-Federation describes how to use WS-Trust [30], WS-Security [21] and WS-Policy [31] all together in order to provide federation between security domains. This enables high value scenarios where authorized access to resources managed in one realm can be provided to security principals whose identities and attributes are managed in other realms.

Moreover, WS-Federation introduces mechanisms to manage and broker trust relationships in a heterogeneous and federated environment. This includes support for federated identities, attributes and pseudonyms. "Federation" refers to the concept that two or more security domains agree to interact with each other, specifically letting users of the other security domain accessing services in the own security domain. For instance, two companies that have a collaboration agreement may decide that employees from the other company may invoke specific web services. These scenarios with access across security boundaries are called "federated environments" or "federations". Each security domain has its own security token service, and each service inside these domains may have individual security policies. WS-Federation uses the WS-Security, WS-SecurityPolicy and WS-Trust specifications to specify scenarios to allow requesters from the one domain to obtain security tokens in the other domain, thus subsequently getting access to the services in the other domain. WS-Federation contemplates the following security services:

- **Privacy services**: Requests made to service providers for security tokens or authorization decisions may include information that is subject to personal or organizational privacy requirements. WS-Federation defines extensions to the Request Security Token and Response syntax defined in WS-Trust for a requestor to express its privacy requirements and likewise for a Security Token Service (STS) to indicate to the requestor the mechanisms actually used for issuing the token. This includes extensions for indicating parameters that an STS must use if it issues a token, as well as identifying individual sensitive claims in a security token for which the values should be protected by encryption. WS-Federation also defines a model for how privacy statements can be obtained using mechanisms defined in HTTP, HTTPS, WS-Policy or WS-MetadataExchange [32].

- **Pseudonym service**: may store tokens associated with a pseudonym to simplify retrieving a pseudonym and associated tokens in a single security token request. It may provide distinct pseudonyms for different scopes, that is, for access to different resource providers. WS-Federation describes how a pseudonym service that is combined with a STS may map pseudonyms to issued tokens. This includes describing how the mapping may be automatically performed based on the target service for the token. It also defines extensions to the WS-Trust Request Security Token and Response messages syntax for requestors to manually specify how pseudonyms should be mapped.

- **Authorization services**: may be implemented as a special type of Security Token Service which provides decision brokering services for participants in a federation. While the internal processing of an authorization enforcement is implementation specific, interoperability between services in a federation requires a common model for interacting with authorization services. WS-Federation defines an authorization model that meets these requirements. The protocol also defines two extensions for rich authorization capabilities.

- **Authentication types services**: to facilitate interoperability, WS-Federation has identified and defined a set of Universal Resource Identifiers (URIs) for specifying the common authentication types and assurance levels that can be used in Request Security Token and Response messages.

- **Attribute services**: WS-Federation defines a model for either party to access attribute services based upon the security token service concept and reliant on the token issuance protocol defined in WS-Trust.

Figure 2.6 illustrates two realms with associated attribute/pseudonym services and some of the possible interactions. It must be noted that, it is assumed that there is a trust relationship between the realms. Hence a requestor has knowledge of the policies of a resource (including its STS) and can obtain its identity token from its IP/STS (1a). It also interacts with the resource's Identity Provider/Security Service Token (IP/STS) (Step 2) in order to obtain an access token for the resource. In this example the resource IP/STS has registered a pseudonym with the requestor's pseudonym service (Step 3) possibly for service-driven mappings. In Step 4, the requestor accesses the resource using the pseudonym token, which enables to the resource to gain additional information in Step 5

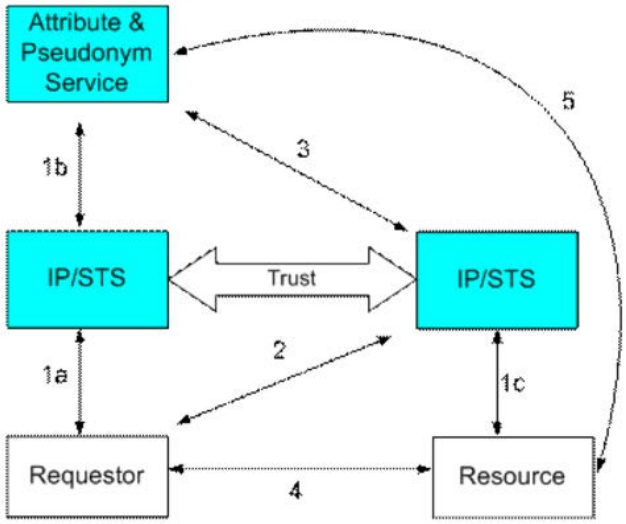

Figure 2.6: General Model for Pseudonyms and Attributes services (©[29]).

from the requestor's attribute service if authorized based on its identity token (previously in Step 1c). Note that trust relationships will require to exist in order for the resource or its IP/STS to access the requestor's attribute or pseudonym service. In subsequent interactions, the requestor's IP/STS may automatically obtain pseudonym credentials for the resource (See Step 1b) whether they are available. In such cases, Steps 2 and 3 are omitted. It worth be mentioned that, in Step 1a a service-consumable identity is returned when the mapping occurs at the IP/STS. In this way, pseudonym services could be integrated with identity providers and security token services. Similarly, a pseudonym service could be integrated with an attribute service as a specialized form of attribute.

Discussion

The federated approach is based on groups of SPs and IdPs that have a pre-existing mutual trust relationship. Consequently, specifications, such as SAML, recommend using Public Key Infrastructure (PKI) [33] for establishing trust relationships. Regarding the terminology of Liberty Alliance and WS-Federation, the above groups are called members of the *circle of trust* and security realm, respectively (see Fig. 2.7).

Nevertheless, as far as usability and scalability are concerned, this model has several draw-

backs. For instance, it adds further 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 [34].

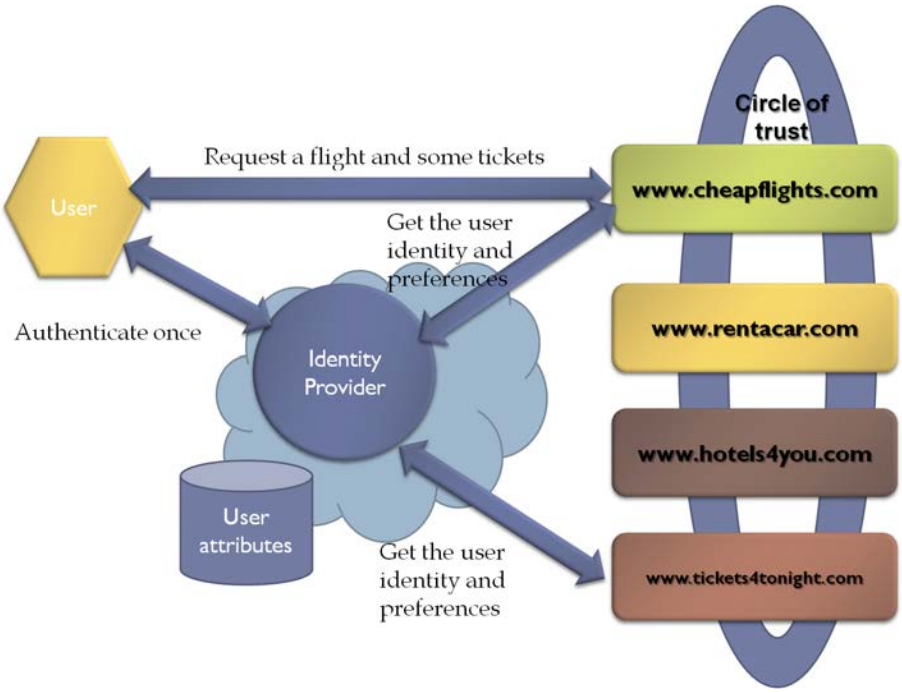

Figure 2.7: Federated model scenario. A user, after a successful authentication, can access services from any service provider within the circle of trust. For instance, booking a flight, then renting a car, and finally buying tickets for a show. Note that the IdP stores identity information on behalf of the user.

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 a single identity to have different identifiers in different domains, e.g., as patient in a health care domain, as employee or student in another domain, etc. Moreover, from the SP perspective, the identifier mapping permits different SPs to refer to the same user through different identifiers. 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 opaque process so they need to believe that the IdP will behave honestly.

The main drawback of this kind of federated identity models 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, SAML, Liberty and WS-Federation, advise the use of pairwise, directional opaque identifiers.

2.1.2 User Centric Identity Models

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 and 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 Jøsang. 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. The main existing user-centric protocols and specifications are explained below.

OpenID Connect

OpenID Connect 1.0 [35] is the current version of OpenID. It is a simple (JSON[3]/REST[4])-based identity protocol built on top of the OAuth2.0 [36] protocol and JWT (JSON Web Token) [37] family of protocols. OAuth is an open standard for authorization, which gives users the ability to grant third-party access to their resources without sharing their passwords. It also provides a way to grant limited access (in scope, duration, etc.). OAuth allows users to share their private resources (e.g., photos, videos, contact lists, bank accounts) stored on one site with another site without having to hand out their credentials,

[3] JavaScript Object Notation (JSON)
[4] Representational State Transfer (REST)

typically username and password.

The concept of OAuth is based on the metaphor of a valet key of car, since it only gives third parties a controlled (limited) access to the car [38]. OAuth mimics the valet key metaphor by providing sites with just enough information to accomplish what the user has requested, but not allowing third-party sites access to any other user information. OAuth defines two different types of Tokens (Request and Access) and covers both online and offline scenarios (i.e. the user is not present). Thus, it only allows users to hand out to third parties tokens (instead of credentials) to their data hosted by a given service provider. The tokens could be granting a printing service access to photos without sharing username and password.

To achieve the aspects and scenarios above described, the framework provides and specifies the following roles, concepts and definitions:

- **The Service Provider or Resource Owner**: It is the term used to describe the website or web-service where the restricted resources are located. For instance, it can be a photo sharing site where users keep albums, an online bank service, a microblogging site, or any other service where user's private information is kept.
- **The User**: The users have personal data they do not want to make public on the Service Provider, but they do want to share it with another site. In OAuth, the protocol stops without manual interaction with the user at least once to receive permission to grant access.
- **The Consumer or Client**: This can be a website, a desktop program, a mobile device, a set-top box, or anything else connected to the web trying to access the User's resources. The Consumer is the one getting permission to access resources and the Consumer is where the useful part of OAuth happens.
- **The Protected Resources**: They are data controlled by the Service Provider (e.g., photos, documents, contacts, activities or any URL with a need for access restrictions.), which the Consumer can access through authentication.
- **The Consumer Key**: It is a value used by the Consumer to identify itself to the Service Provider.
- **The Consumer Secret**: It is a secret used by the Consumer to establish ownership

of the Consumer Key.

- **The Request Token**: It is a value used by the Consumer to obtain authorization from the User, and exchanged for an Access Token.
- **The Access Token**: It is a value used by the Consumer to gain access to the Protected Resources on behalf of the User, instead of using the User's Service Provider credentials.
- **The Token Secret**: It is a secret used by the Consumer to establish ownership of a given Token.

So, the OAuth 2.0 framework enables a Relying Party (RP) to obtain profile information about the end user, but does not provide any means for the RP to obtain information about the authentication of the end user.

In OpenID Connect, in addition to obtaining profile information about the end-user, RPs can obtain assurances about the end user's identity from an OpenID Provider (OP), which itself authenticates the user. Thus, OpenID Connect enables clients to verify the identity of the user based on the authentication performed by an authorization server. Moreover, the specification suite is extensible, allowing participants to use optional features such as encryption of identity data, discovery of OpenID Providers, and session management, when it makes sense for them.

OpenID Connect involves interactions between four entities:

- **The OpenID Provider**: It provides methods to authenticate an end user and generates assertions regarding the authentication process and the attributes of the end user.
- **The Relying Party**: It offers protected on-line services and consumes the identity assertion generated by the OP in order to decide whether or not to grant access to the end user.
- **The End User**: She accesses on-line services of the RP.
- **The User Agent**: It is typically a web browser, that is employed by an end user to transmit requests to, and receive responses from, web servers.

In order to allow a RP to verify the identity of an end user, OpenID Connect adds a new

type of token to OAuth 2.0, namely the *id_token*. This complements the access token and code, which are already part of OAuth 2.0. These three types of token are all issued by an OP with the following functions:

- ***Code***: It is an opaque value, which is bound to an identifier and a URL of the RP. Its main purpose in OpenID Connect is to serve as a mechanism to RP to retrieve other OP cards. With the aim to help minimize threats arising from its possible exposure, this code has a limited validity period and is typically set to expire shortly after issue to the RP [36].
- ***Access Token***: It is a credential used to authorize access to protected resources stored at a third party (e.g., the OP). Its value is an opaque string representing an authorization issued to the RP. It encodes the right for the RP to access data held by a specified third party with a specific scope and duration, granted by the end user and enforced by the RP and the OP.
- ***Id_Token***: It contains claims about the authentication of an end user by an OP together with any other claims requested by the RP. Claims that can be inserted into such a token include: the identity of the OP that issued it, the user's unique identifier at this OP, the identity of the intended recipient, the time at which it was issued, and its expiry time. It takes the form of a JSON Web Token [37] and is digitally signed by the OP.

Both an *access token* [39] and *id_token* [30] can be verified by invoking the web API of the issuing OP.

In this way, the RP generates an authorization request on behalf of the end user and sends it to the OP through a user agent. The OP provides ways to authenticate the end user and asks the end user to enable the RP to access the user attributes. It also produces an authorization response, which includes of two type tokens: *access tokens* and *id_tokens*, where the latter contain claims about user authentication. Thus, the RP can use a received access token to access end user's attributes using the OP-provided API, and after receiving an *id_token* the RP learns about the user authentication, as shown in Figure 2.8.

Furthermore, OpenID Connect supports four authentication flows [41], i.e. ways in which the system can operate, namely *Authorization Code Flow*, *Hybrid Server-side Flow* (or *Hybrid Flow*), *Client-side Flow* (or *Implicit Flow*), and *Pure Server- side Flow*.

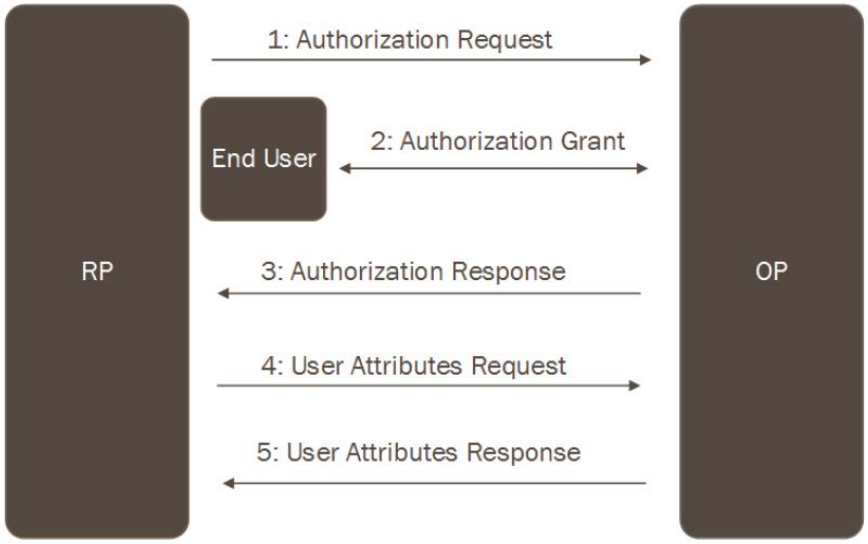

Figure 2.8: OpenID Connect Protocol Overview (©[40]).

Summarizing, OpenID Connect is mainly an authentication protocol and allows attribute exchange with the user's consent.

Information Cards

Information Cards (aka infoCard) are metaphors of real id cards whereas the identity selector mimics a wallet. We use them with a new kind of digital wallet called a selector, which enables users to manage in an easy, visual way their different electronic identities. The Information Cards technology is an open, neutral industry standard for safer digital identity supported by a non-profit organization - the Information Card Foundation (IFC) [42]- composed by companies and individuals working together to evolve the internet identity ecosystem. The foundation is currently organized in a series of active Working Groups that deal with issues such as standardization, implementation, interoperability with deployed identity technologies, and best practices, etc. Moreover, they published several white papers and specifications on Information Card technology and practice, being [43] and [44] the core documents that define the identity formats and protocol flows.

InfoCards are like business cards which let users decide what information will be disclosed during an interaction, keeping personal data under control and they also let Relying Parties to get the information they need directly from the users. Regarding identity selec-

tors, in [43] 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 also *Managed Information Cards*, issued by IdPs. The latter can be auditing, non-auditing, or auditing-optional to accommodate the needs of different business models. The identity cards are the digital version of the cards we carry in our wallet today. Windows CardSpace [45], Higgins project [46] or Open Source Identity Selector [47] are Identity Selector implementations used nowadays. Furthermore, *InfoCards* support several data formats and authentication methods such as XML, SAML, and OpenID. *InfoCard*-based identity management systems typically use Web Services Security protocols (WS-*) and SOAP. WS-Trust [30] is the protocol used to obtain and exchange security tokens. Moreover, the integrity of the tokens is preserved using an XML-Signature as part of the WS-Security [21] protocol.

U-Prove

U-Prove [48] [49] is an advanced cryptographic software designed for electronic transactions and communications to overcome a long-standing dilemma between identity assurance and privacy [1], [50]. The technology is part of Microsoft's drive to promote an open identity and access model for individuals, businesses and governments, based upon the principles of the identity metasystem [51]. The dilemma is addressed by enabling minimal disclosure of identity information in electronic transactions and communications. To ensure minimum disclosure the U-Prove Agent software acts as an intermediary between websites. Furthermore, U-Prove defines two classes of tokens: *claim tokens*, which can encode arbitrary claims, and *ID tokens* that specify globally unique, secure, and privacy-protecting persistent identifiers for users.

A U-Prove token (UPT) is a cryptographically protected container of claim information that is issued for a Prover (the client) by an Issuer (the Claim Provider), and presented to a Verifier (the Relying Party). Each UPT corresponds to a private key needed to present the token, and contains an Issuer's signature that vouches for its origin and integrity. A UPT is conceptually similar to a X.509 certificate or SAML assertion, with two major differences:

- A UPT is generated jointly by the Prover and the Issuer in an interactive issuance protocol. It contains no correlation handles identifiable by the Issuer outside the

certified claim values. In particular, its public key and Issuer's signature have been randomized by the Prover in the issuance protocol; as such, these values are never seen by the Issuer. Consequently, the Prover cannot be tracked on the basis of these values when using the UPT, even if the Issuer and the Verifiers collude (even if they are the same entity).

- When presenting a UPT, the Prover can hide any subset of the encoded claims, without invalidating the Issuer's signature generated on all the claims. In particular, the Prover can hide all the claims or disclose all of them (like presenting a signed SAML assertion or a X.509 certificate).

Figure 2.9 sketches out the protocol operation when a user configures her privacy options to combine several claims. A Prover gets multiple UPTs certifying the same set of claims in one instance of the issuance protocol (multiple UPTs are obtained to preserve unlinkability between Verifiers). To present any subset of the certified claims to a Verifier (immediately after issuance for on-demand tokens, or a later time for long-lived tokens), the Prover creates a presentation proof for a selected UPT by applying the corresponding private key to a cryptographic challenge. It should be noted that blending multiple claims only takes place if user customizes her privacy options.

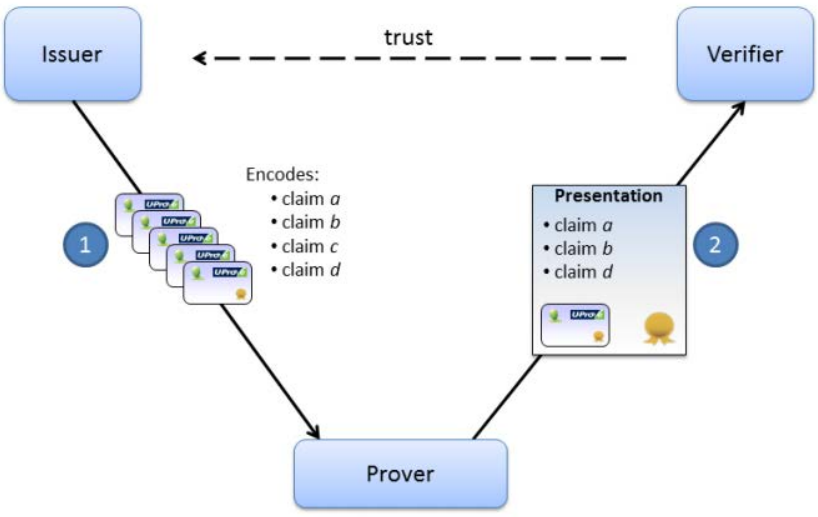

Figure 2.9: U-Prove token data flow (©[52]).

In summary, U-Prove aims to allow users to share data in a manner that protect their

privacy, since they can now choose to share or otherwise. The specification includes a mechanism that separates the retrieval of information from trusted third parties from the release of this information to the destination site. This implies that the organization issuing the information is prevented from tracking where or when information is used. The destination site is similarly prevented from linking users to their activities.

Discussion

In regard to privacy, the user-centric model has both advantages and drawbacks. It introduces the concept of *meta-idp*, which allows users to assert several kind of claims: *user-generated* and *provider-generated* claims. These user electronic identities are typically stored in user's equipment, such as her mobile phone. User-centric identity technologies, such as InfoCards, allow users to select among their multiple identities through identity selectors to identify herself to a service. 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 2.10).

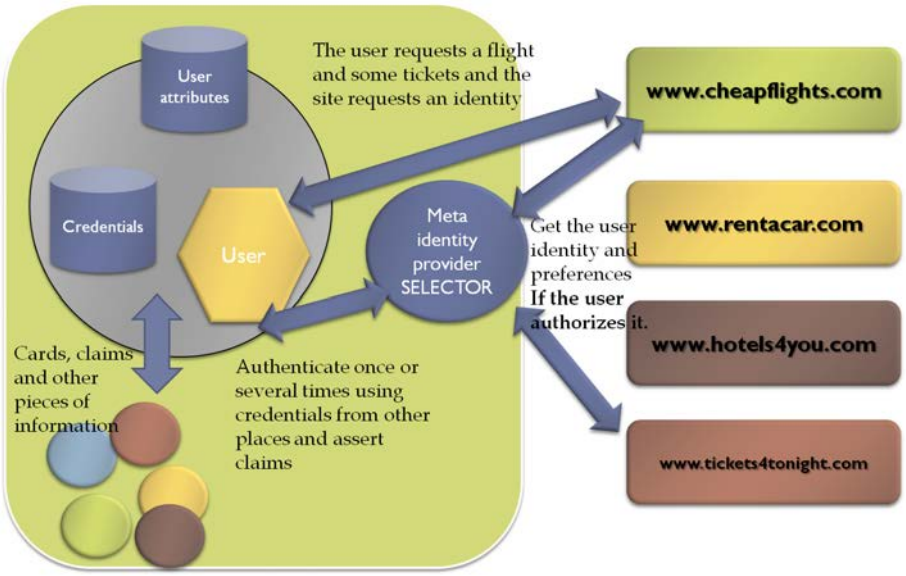

Figure 2.10: User centric model. A user can access services from any service provider accepting his/her credentials. For instance, booking a flight, then renting a car and finally buying tickets for a show. Note that the information is provided always by the user.

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 license 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.

IdM Framework	**Main Purpose**	**Format data and types of tokens**	**Security and Privacy considerations**	**Implementations**
OpenID Connect	Focuses on usability and simplicity aspects for identity management. It is simple enough to integrate with client and basic apps.	*Access*, *Request* and *id_tokens* for JSON and REST. The adoption of JSON Web Token and JSON Web Signatures facilitates attribute exchange and processing on mobile applications.	A huge proportion its security of is based on the utilization of Transport Layer Security (TLS). It is built on the JWT family of protocols. Any credential management mechanisms is defined. Use of the *Authorization Code* and *Access Tokens* in a transient manner recommended.	Google, Microsoft, Deutche Telekom, PayPal, etc. offers this protocol.
InfoCards	Focuses on maximizing the ease of use and the individuals' control over their identities.	*Self-issued* and *managed* information cards. XML, SOAP, SAML, OpenID and WS-* are supported.	XML-Signature, WS-Security. Direct communication between SPs and IdPs is eliminated and identity selectors are defined.	Windows Card Space, Higgings, Open Source Identity Selectors.
U-Prove	Focuses on security and privacy aspects of identity management	Provides *claim tokens* and *ID tokens*, which include persistent identifiers. U-Prove tokens are jointly created by the user's computing device and a credential authority.	Privacy by Design: Minimal disclosure, user control and selective disclosure. Claims are computed as short zero-knowledge proof for the device-protection attribute. Trusted terminal required	U-Prove profiles have been implemented in a variety of technologies; the WS-Trust [52] was used in CardSpace. Open-source U-Prove C# Crypto SDK and JavaScript SDK are provided.

Table 2.1: Outline of features of current user-centric IdM specifications

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.

In Table 2.1 we summarize the main features of the user-centric IdM specifications reviewed in this section. With the aim to show the differences between OpenID Connect, Information Cards and U-Prove, aspects related to its main purpose, format data and security tokens supported, security and privacy considerations, as well as available implementations, are highlighted. A complete comparative analysis of the privacy aspects that cover the federated and user-centric IdM technologies described in this section can be found in Chapter 4.

2.2 Privacy Overview

2.2.1 Basic Concepts and Definitions

Privacy is defined by Windley [53] as the *"protection of the attributes, preferences, and traits associated with an identity from being disseminated beyond the subject's needs in any particular transaction"*. However, privacy is a complex and subjective concept with different meanings to different people, that depend on the context in which it is used.

In this thesis we focus on the privacy involving the capability to control what information an individual discloses or withholds about herself, including determining who has access to such information, and for what purposes the information may or may not be used. With the advances in technology, users are increasingly subject to privacy threats. For instance, users may become concerned if they discover that visited websites collect, store, and possibly share personally identifiable information about them.

According to the U.S. Office of Management and Budget (OMB)[5] [54], and also as adopted by National Institute of Standards and Technology (NIST) [55], Personally Identifiable Information (PII) is defined as *"any information about an individual maintained by an agency, including (1) any information that can be used to distinguish or trace an individual's identity, such as name, social security number, date and place of birth, mother's maiden name, or biometric records; and (2) any other information that is linked or linkable to an individual, such as medical, educational, financial, and employment information"*. In

[5] http://www.whitehouse.gov/omb

a digital context, the term PII refers to the information that can be used to uniquely identify a person.

In this sense, it is important to distinguish the three types of user's attributes that may revealed: direct identifiers, quasi-identifiers (QIDs), and sensitive attributes. Direct identifiers are attributes that can explicitly re-identify individuals, such as name, mailing address, phone number, social security number, other national IDs, and email address. On the other hand, quasi-identifiers are attributes which in combination can lead to identity disclosure (e.g., gender, date of birth, zip code, diagnosis codes, etc.) [56]. Last, sensitive attributes are those that users are not willing to be associated with (e.g., psychiatric diseases, human immunodeficiency virus (HIV), cancer, etc.).

Based on the above types of attributes, the following classes of privacy threats can be found:

- **Identity disclosure or re-identification**: It occurs when an attacker can associate a user with her records in a published dataset [57].
- **Membership disclosure**: This threat occurs when an attacker can infer with high probability that an individual's record is contained in the published data [58]. For instance, consider a dataset which contains information on only HIV-positive patients. The fact that a patient's record is contained in the dataset allows inferring that the patient is HIV-positive, and thus poses a threat to privacy. It should be noted that membership disclosure may occur even when the data are protected from identity disclosure. Some scenarios where protection against membership disclosure is required are analyzed in [58] [59].
- **Attribute disclosure or sensitive information disclosure**: It occurs when an individual is associated with information about their sensitive attributes [60]. This information can be, for instance, the individual's value for the sensitive attribute.

Regarding a taxonomy of privacy from a perspective of law, Solove defined in [61] sixteen different types of privacy violation. The author classifies the identified privacy violations in four categories, which are:

- **Information Collection**: surveillance and interrogation.
- **Information Processing**: aggregation, identification, insecurity, secondary use and

exclusion.

- **Information dissemination**: breach of confidentiality, disclosure, exposure, increased accessibility, blackmail,appropriation and distortion.
- **Invasion**: intrusion and decisional interference.

Technological measures to protect against privacy violations focus mostly on preventing the unintended leakage of information; while other types of violations fall out of the scope of technological systems and legal measures are needed in order to prevent them. Technical systems can better protect against the following particular privacy threats:

- **Surveillance**: considering adversary capable of monitoring electronic transactions, privacy-enhancing technologies aim to reduce the risk of surveillance by concealing information about the content and circumstances of electronic transactions from adversary. When users are able to keep transaction contents confidential and to act anonymously, they protect themselves against surveillance threats.
- **Interrogation**: the technical property that protects a user from being forced to disclose information is called plausible deniability. Systems that provide plausible deniability make it impossible for adversary to prove that the user is concealing information.
- **Aggregation** the property that prevents the aggregation of information as related to each other or to a particular subject is unlinkability.
- **Identification**: Identification is connecting data to individuals. Anonymity, unlinkability and confidentiality properties prevent this connection to be revealed.

2.2.2 Principles and rules of privacy

In the context of sensitive environments, like health care scenarios, user privacy is a fundamental right that is being challenged as user records are digitized. However, users should not be required to sacrifice their right to privacy with the aim to obtain services, such as healthcare. In this sense, privacy principles and rules as anonymity, pseudonymity, unobservability, unlinkability, selective disclosure and revoking consent have to be guaranteed. These concepts might have different definitions in the literature [62]. The works

presented in [63] [64] [65] include additional privacy properties such as undetectability, role interchability or undeniability/non-repudiability.

However, as it is shown in [62], not all the aforementioned privacy properties can be achieved at the same time. For instance, complete anonymity is incompatible with non-repudiability or accountability, which are desirable properties in healthcare scenarios. Thus, in this section we will focus on the definition of properties from a privacy standpoint that have been addressed to be analyzed later in Chapters 4 and 6.

- ***Anonymity*** can be defined as the state of being not identifiable within a set of subjects or entities. According to [66], this property ensures that a user may use a resource or service without disclosing the user identity. The relevant threats are: disclosure of identity or leakage of information leading to disclosure of identity, often described as "usage profiling". Encryption does not guarantee anonymity, since an observer can still analyze traffic, eavesdrop the sender and follow the message up to the receiver, establishing certain relationships; therefore, IdM systems must provide additional mechanisms, such as opaque identifiers to prevent such inference. Moreover, it is necessary that IdM systems support partial anonymity or semi-anonymity mechanisms (the IdP would know which user corresponds to each identity) for the purpose of evidence or analysis under certain circumstances (e.g., lawful interception).

- ***Pseudonymity*** is the use of pseudonyms as identifiers of subjects (or sets of subjects when we generalize it a bit). The subject which the pseudonym refers to is the holder of the pseudonym. An advantage of pseudonymity is that accountability for misbehavior can be enforced. Thus, this enables service providers, that can link identifiers to real identities to make appropriate decisions when a user commits an attack. In addition to the anonymity services, pseudonymity provides methods for authorization without identification (at all or directly to the resource or service provider).

- ***Unlinkability*** ensures that a user may consume multiple resources or services without letting other entities to link these multiple resource or service accesses together [63]. In particular, this allows users to interact with multiple organizations, each of them able to map a user to a given identity, using different identities. IdM systems should provide mechanisms to prevent collaborating organizations from link-

ing a given user profile at one organization with the same user profile at another.

- ***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. Moreover, this property is closely related to anonymity, since in terms of item of interest (IOI), unobservability means anonymity of the subject(s) involved in the IOI even against the other subject(s) involved in that IOI. As we had anonymity sets of subjects with respect to anonymity, we have unobservability sets of subjects with respect to unobservability. Sender unobservability then means that it is not noticeable whether any sender within the unobservability set sends a message. Recipient unobservability then means that it is not noticeable whether any recipient within the unobservability set receives a message. Relationship unobservability then means that it is not noticeable whether anything is sent out of a set of could-be senders to a set of could-be recipients. In other words, it is not noticeable whether, within the relationship unobservability set of all possible sender-recipient-pairs, a message is exchanged in any relationship.

- ***Selective Disclosure*** is the act of disclosing personal information in one online transaction setting and not disclosing it in another. Thus, users may create various online personalities and share and receive specific information with SPs or IdPs for each personality. For instance, users can have a private personality where they disclose private data such as the number of their private mobile phone, and a work personality where they disclose work-related data such as the office location and their meeting schedule.

- ***Revoking Consent*** enables users give or invalidate consent of specific actions over data to certain individuals. So, user's consent revocation adjusts the privacy view as control over the use and flow of one's personal information. IdM systems should provide this mechanism in order to enforce properly user's role in the task of preserving her privacy.

2.2.3 Current Privacy-Preserving Models

In last years, privacy-preserving techniques has seen quick advancement due to rapid increase in storing and maintaining personal data about individuals. The personal data

can be misused, for a variety of purposes. Maintaining the privacy for high dimensional database has become major aspect. With the aim to enhance these concerns, various models based on anonymization mechanisms for preserving privacy of different types of data like medical data, transportation system data, etc. have been proposed. In this section, we review the main existing privacy-preservation techniques that deal with the above privacy concerns.

k-Anonymity Model

Several algorithms with the aim to guarantee k-anonymity have been proposed in literature [57] [67] [68]. The anonymization problem is usually modeled as an optimization problem, where a given information-loss cost metric has to be minimized. The solution search space consists of all possible transformations of data (e.g., generalization, suppression) that lead to a k-anonymous dataset. Optimal solution is the one that achieves the least information loss, and thus ensures maximum utility in the released data.

k-Anonymity model aims at hiding every individual in the released dataset among at least k-1 others. The size of parameter k defines the level of the desired privacy. Sweeney [57] [67] used generalization and suppression to transform quasi identifiers in a way that at least k records have identical QID values. Thus, any of them may correspond to the same individual-target, for any target of the (at least) k-sized group. The main concepts are detailed below:

Definition 2.1. An *Equivalence Class* is a set of records in an anonymized database that have identical values in the QID attributes.

Definition 2.2.: A table T is considered ***k*-anonymous** with reference to a quasi identifier set of attributes $QID = \{Q_1, Q_2, \ldots, Q_d\}$, if the size of every *Equivalence Class* in T is at least k.

The value of parameter k defines the level of privacy in this guarantee [69]. Satisfying k-anonymity offers protection against identity disclosure, because it limits the probability of linking an individual to their record, based on quasi-identifier attributes, to $1/k$. So, a lager k results to larger *Equivalence Classes* and restricts the probability of a privacy breach. However, this also restricts utility in the released data, as generalizations hide more information that would be useful for instance, in data mining scenarios. Figure 2.11 shows

a 2-anonymity example in which generalization and suppression techniques are applied for quasi-identifiers attributes (including age, gender and zip-code).

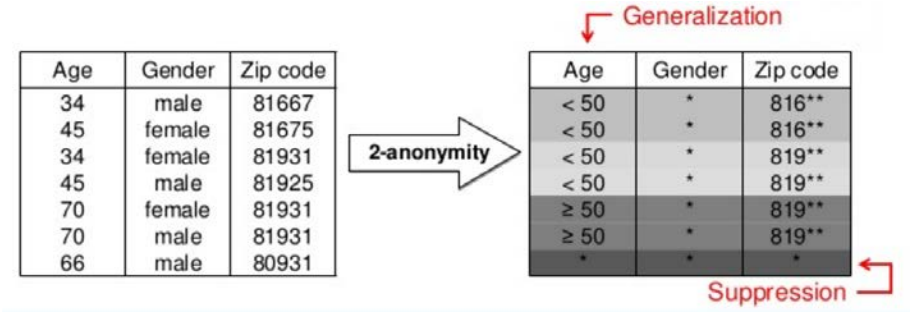

Figure 2.11: 2-Anonymity example for quasi-identifiers attributes.

The work presented in [68] provides an approximate algorithm that minimizes the number of suppressed values. The authors describes an approximation bound of $O(k \log k)$.

Regarding approaches that limit the search space by considering only global recoding, several proposals can be found in the literature [70] [71] [58] [72]. In [70] an algorithm for single-dimensional global recoding is proposed. The k-anonymization algorithm presented in [71] explains a dynamic programming approach called *Incognito*. It tries to find optimal solutions for any metric by considering all possible generalizations, but only for global, full-domain recoding. Full-domain means that all values in a dimension must be mapped to the same level of hierarchy. For instance, in the country→continent→world hierarchy, if France is mapped to Europe, then Canada must be mapped to America, even if the generalization of Canada is not necessary to guarantee anonymity.

Nevertheless, the computational cost of *Incognito* grows exponentially, so it cannot be used for more than 20 dimensions. A different approach is considered in [58], where the authors aims at using natural domain generalization hierarchies to reduce information loss.

Finally, the approach presented in [72] proposes complete k-anonymity. However, complete k-anonymity, may harm data utility unnecessarily because it is extremely difficult for attackers to know for instance, all the diagnoses information in a patient record [56].

L-Diversity Model

k-Anonymity is proven vulnerable to different attacks, especially when the attacker has access to background knowledge [60]. In 2006 Machanavajjhala et al. [60] proposed a new

extended model for preserving privacy, called l-diversity. The authors classifies the vulnerabilities of k-anonymity in two different attack models: background-knowledge attack and homogeneity attack:

- **Background Knowledge Attack**: In this kind of attack, the adversary can use an association between one or more quasi-identifier attributes with the sensitive attribute or public knowledge of the target to eliminate possible values of the sensitive attribute. For instance, if a young individual's QID can be linked to an equivalence class, where all values of the sensitive attribute "disease" are either arthritis, Alzheimer syndrome or flu, it can be inferred that the target's sensitive info is probably "flu", since the first values are highly unlikely to occur to a young person.
- **Homogeneity Attack**: In this attack, all the values for a sensitive attribute within an equivalence class are the same. Therefore, even though the data is k-anonymous, the value of the sensitive attribute for any record in that group of size k can be predicted with 100% accuracy.

Consequently, while k-anonymity is effective in preventing identification of a record, it may not always be effective in preventing inference of the sensitive values of the attributes of that record. A k-anonymous table is safe from record linkage, but vulnerable to attribute linkage. Therefore, l-diversity was proposed as an extension of k-anonymity which not only maintains the minimum equivalence class size of k, but also focuses on maintaining the diversity of the sensitive attributes in every class. The l-diversity model for privacy is defined as follows:

Definition 2.3. A group of records that belong in the same *Equivalence Class* q^* is ***l*-diverse**, if it contains at least l "well-represented" values for the sensitive attribute S. A table T is considered ***l*-diverse** if every *Equivalence Class* q^* in T is l-diverse.

The l-diversity principle ensures the existence of l "well-represented" values in every block of records (equivalence class), without further clarifications on what exactly "well-represented" means.

The simplest interpretation of "well represented" is distinct and leads to distinct l-diversity [60], which requires each anonymized group to contain at least l distinct sensitive attribute values. Another interpretation leads to recursive (c, l)-diversity, which requires each group in T to contain a large number of distinct sensitive attribute values, none of

which appears "too" often. Other principles that guard against value disclosure by limiting the number of distinct sensitive attribute values in an anonymized group are (α, k)-anonymity [73] and p-sensitive-k-anonymity [74].

However, these privacy principles still allow attackers to infer that an individual is likely to have a certain sensitive attribute value when that value appears much more frequently than other values in the group [75].

t-Closeness Model

The main goal of the t-Closeness model [76] [77] is to limite the distance between the probability distribution of the sensitive attribute values in an anonymized group and that of sensitive attribute values in the entire dataset. Thus, this model requires the distribution of a sensitive attribute in any equivalence class to be close to the distribution of the attribute in the overall table. The authors demonstrated in [76] two new types of attack that l-diversity and similar models fail to prevent:

- **Skewness Attack**: Satisfying l-diversity (see Definition 2.3) does not prevent sensitive attribute disclosure when the distribution of an attribute in the overall table is skewed. For instance, authors consider a medical table with the results of HIV examinations, where 98% of the total population were found negative. An equivalence class with equal number of positive and negative results of the examination is assumed. Moreover, it satisfies distinct 2-diversity and a recursive (c, 2)-diversity can be imposed. However, individuals linked to this equivalence class face a severe privacy risk, as they are found 50% possible of being positive to HIV, compared to 2% of the total population.

- **Similarity Attack**: In this attack, an adversary can learn important information when the values of the sensitive attribute in an equivalence class are distinct but semantically similar. Consider an l-diverse equivalence class and the sensitive attribute values that appear in records of this class are bronchitis, pneumonia, flu and lung cancer. An adversary can infer that the target has some lung related disease.

In order to prevent skewness attack, Li et al. [76] proposed a privacy model, called t-Closeness. The main concepts of this model are described below:

Definition 2.4. The t-Closeness principle is satisfied when an equivalence class is said to have ***t*-closeness**. In other words, the distance between the distribution of a sensitive attribute in the equivalence class and the distribution of the attribute in the whole table is no more than a threshold t. A table is said to satisfy t-closeness if all its equivalence classes have ***t*-closeness**.

However, this model has several limitations and weaknesses. It is not suitable for preventing attribute linkage on numerical sensitive attributes. In addition, enforcing t-closeness would severely affect the data utility as it requires the distribution of sensitive values to be the same in all equivalence classes. This would significantly damage the correlation between the QID and sensitive attributes.

To reduce this information loss, another version of the t-Closeness model is presented in [77].

Definition 2.5 The (n,t)-Closeness principle establishes that an equivalence class E_1 is said to satisfy **(*n*,*t*)-closeness** if there exists a set E_2 of records that is a natural superset of E_1. Thus, the distance between the two distributions of the sensitive attribute in E_1 and E_2 is no more than a threshold t. Furthermore, a table is said to satisfy **(*n*,*t*)-closeness** if all of its equivalence classes satisfy **(*n*,*t*)-closeness**.

It should be noted that, this enhanced model achieves a better balance between privacy and utility compared to the stricter t-closeness.

Other Privacy Models

Other techniques to preserve privacy proposed in the literature are δ-presence [58], (c, k)-safety [78] and differential privacy techniques [79] [80].

Firstly, the δ-presence model [58] intents to prevents an adversary from learning whether an individual owns a record in the microdata. The authors proposed to bound the probability of inferring the presence of any potential target record within a specified range $\delta = (\delta_{min},\delta_{max})$ with the aim to prevent linkage. The main principle of this model is detailed below.

Definition 2.6. Given an external public table E and a private table T, where $T \subseteq E$, a generalized table T^* satisfies $(\delta_{min},\delta_{max})$-presence, if $\delta_{min} \leq P(t \in T|T*) \leq \delta_{max}$ for all

$t \in E$.

Where $P(t \in T|T*)$ is the probability that record t is included in original table T, given its anonymous version T^*.

δ-presence can indirectly prevent record and attribute linkages because if the attacker has at most $\delta\%$ of confidence that the target victim's record is present in the released table, then the probability of a successful linkage to her record and sensitive attribute is at most $\delta\%$. Though δ-presence is a relatively "safe" privacy model, it assumes that the data publisher has access to the same external table E as the attacker does. This may not be a practical assumption.

Regarding the (c, k)-safety model proposed in [78], it guarantees that, whether an adversary may have already known at most k pieces of implicational knowledge, they will not be able to infer any user's sensitive information with confidence higher than c. In this work the authors study the adversary's worst-case background knowledge and provide a polynomial-time algorithm that estimates the amount of disclosure of sensitive information in the worst case scenario.

Other approaches as [81] introduce more realistic classes of attackers. These adversaries may have some external knowledge and a new characteristic called *stubbornness* indicating the strength of their prior knowledge.

Finally, in the last decade several models based on differential privacy have been proposed [79] [80]. Intuitively the notion of differential-privacy says that any possible outcome of an analysis should be almost equally likely, independent of whether any individual is included or removed from the data set. In this way, the data of any specific individual can never seriously affect the result of the analysis. A more detailed explanation of this model is found in [79].

Techniques used to achieve differential privacy include the addition of *Laplacian noise* [79] [80] to each of the k outputs that is calibrated to the sensitivity of the analysis. Nevertheless, the tolerance for the addition or removal of any one tuple in the dataset is very restrictive for the production of a private release. Thus, the distortion of the information due to the noise that must be added to ensure such a strict model is severe.

Discussion

The privacy models reviewed in this section attempt to produce anonymous releases, while maintaining the maximum possible data utility, by minimizing a given information loss metric. Detailed definitions of different metrics of data utility can be found in [70] [82]. Consequently, they are prone to attacks can result to identity or attribute linkages. t-Closeness approach has less information loss than l-diversity and k-anonymity models but these techniques still leads to extensive information loss. Differential privacy uses *Laplacian Noise* addition and is less susceptible to minimality attacks. It provides strong privacy guarantees, while leaving little room for data utility maintenance.

On the other hand, privacy approaches that apply to complex data sharing scenarios while enabling selective information disclosure and effective revocation consent need to be addressed. For instance, consider the case of multiple healthcare providers and several identity medical records services that wish to obtain a richer view of patients' Electronic Health Records (EHRs). Healthcare providers and identity medical records services may contribute different parts of patients' EHR, whereas each user or healthcare service provider involved in the above scenario must have access only during the strictly required time to the minimum necessary information. This scenario presents several interesting challenges.

Firstly, as will be discussed in Chapters 3 and 4, more effective and time-independent revocation mechanism are necessary to reach a balance between usability, privacy and security. Secondly, as will be explained in Chapters 3 and 6, data contributed by different providers need to be integrated in an efficient and privacy-preserving structure. In the next section, we take a closer look at the existing structures for privacy-awareness user profiles management.

2.3 Structures for privacy-awareness profiles management

The use of authenticated dictionary structures (ADTs) to construct user' credentials offers desirable properties to preserve user's privacy in IdM systems, since ADTs enable to prove the presence of an attribute without requiring to reveal any other attributes of the structure. Furthermore, ADTs enable to combine and group user's attributes from

different information sources while preserving user's privacy. ADTs are data structures that support both update queries and tamper-evident membership queries. For instance, a tamper-evident membership query is of the form "does element e belong to S?".

The answer to such a query consists in providing a tamper-evident proof showing that element e participated to the construction of the root value of data set S. So, the problem we address involves three parties: a trusted source, an untrusted directory, and a user. The *source* defines a finite set S of elements that evolves over time through insertions and deletions of elements. The *directory* maintains a copy of set S. It receives time-stamped updates from the source together with update authentication information, such as signed statements about the update and the current elements of the set.

The *user* performs membership queries on the set S, but instead of contacting the source directly, it queries the directory. The directory provides the user with a yes/no answer to the query together with query authentication information, which yields a proof of the answer assembled by combining statements signed by the source. The user then verifies the proof by relying solely on its trust in the source and the availability of public information about the source that allows to check the source's signature. The data structure used by the directory to maintain set S, together with the protocol for queries and updates is called an ***authenticated dictionary***.

In order to better understand how ADT structures work, firstly we will review definitions of hash functions and their main properties. Then, we will describe several existing privacy-preserving techniques based on structures formed from hash functions or attribute encryption.

2.3.1 Basic Concepts and Definitions

The term "hash function" is due to computer science, and refers to a function which compresses an arbitrary length input bit string to an output bit string of fixed finite length [83]. These functions are primarily used to speed up the process of finding stored data. However the term "hash function" has since been adopted by cryptographers, and desirable properties of cryptographic hash functions, such as "collision resistance", have been identified [84]. The main concepts are detailed below:

Definition 2.7. A function $f : D \to R$ is preimage resistant if for a given $y \in R$ it is

computationally infeasible to find an $x \in D$ such that $f(x) = y$.

Definition 2.8. A function $f : D \to R$ is second preimage resistant if for a given $x \in D$ it is computationally infeasible to find $x' \in D$ with $x' \neq x$ such that $f(x') = f(x)$.

Definition 2.9. A function $f : D \to R$ is collision resistant if it is computationally infeasible to find $x, x' \in D$ with $x' \neq x$ such that $f(x) = f(x')$.

These properties imply that a malicious adversary cannot replace or modify the input data without changing its digest. Thus, if two strings have the same digest, one can be very confident that they are identical. There are some well known upper bounds on the computation required to exhibit a preimage, second preimage or collision [85].

One of the first authors to define a hash function was Merkle [86], who defined a hash function f to be any function with the following properties:

1. The function f can be applied to an input of any size.
2. The output of f is a bit string of fixed length.
3. The output $f(x)$ is computationally easy to calculate for any x.
4. The function f is preimage resistant (see Definition 2.7).
5. The function f is second preimage resistant (see Definition 2.8).

This definition is also sometimes known as a *weak hash function* [85]. Other authors define hash functions without preimage resistance or second preimage resistance [87], or with additional properties such as collision resistance [88], or such that the input should include a key [89]. Many more authors use the term hash function without specifying the precise properties they require. We will refer to the output of a hash function as the "output", the "hash value", or simply the "hash".

2.3.2 Hash trees

Balanced Trees

The balanced tree data structures were developed in the 1960s and 1970s and provide a guaranteed worst-case running time that is proportional to $log\ N$ for search, insert and delete operations, where N is the number of nodes in the tree prior to the operation.

These algorithms are based on modifying the elementary binary search tree data structure to guarantee that the length of every path to an external node is proportional to *log N*. Examples of such algorithms are 2-3 trees, 2-3-4 trees, AVL trees, red-black trees and B trees.

A 2-3 tree is a tree data structure, where every node with children (namely internal node) has either two children (2-nodes) or three children (3-nodes), whereas a 2-3-4 tree is a self-balancing data structure, where every internal node has either two, three, or four (4-nodes) child nodes. Regarding AVL trees, they are self-balancing binary search trees. In an AVL tree, the heights of the two child subtrees of any node differ by at most one; if at any time they differ by more than one, rebalancing is done to restore this property.

A red-black tree in its turn, it is a binary search tree with one extra attribute for each node: the colour, which is either red or black. Finally, B trees are self-balancing tree data structures and generalizations of binary search trees in that a node can have more than two children.

In [90], Guibas and Sedgewick showed that all of these algorithms can be implemented with red-black trees, which are binary search trees with one extra bit of storage per node: its color, which can be either red or black. By constraining the way nodes can be colored on any path from the root to a leaf, red-black trees ensure that no such path is more than twice as long as any other, so that the tree is approximately balanced. In other words, each link in a binary search tree is assigned a color (red or black) that can be used to control the balance, and that this framework can simplify the implementation of the various algorithms.

Furthermore, each node of the tree now contains the fields *color*, *key*, *left*, *right*, and *p*. If a child or the parent of a node does not exist, the corresponding pointer field of the node contains the value *NIL*. We shall regard these *NIL*'s as being pointers to external nodes (leaves) of the binary search tree and the normal, key-bearing nodes as being internal nodes of the tree.

[90] also describes a way to maintain a correspondence between red-black trees and 2-3-4 trees, by interpreting red links as internal links in 3-nodes and 4-nodes. Since red links can lean either way in 3-nodes (and, for some implementations in 4-nodes), the correspondence is not necessarily 1-1 (See Figure 2.12). The basic operations that balanced-tree algorithms

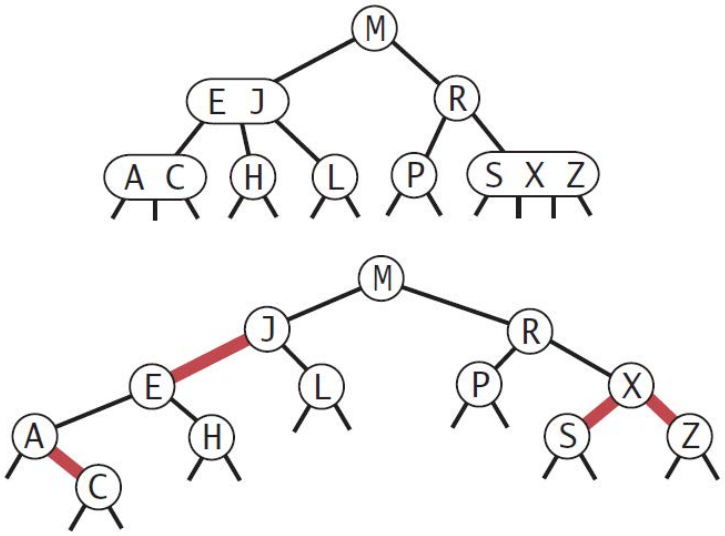

Figure 2.12: Example of Red-black representation of a 2-3-4 tree.

use to maintain balance under insertion and deletion are known as rotations. In the context of red-black trees, these operations are easily understood as the transformations needed to transform a 3-node whose red link leans to the left to a 3-node whose red link leans to the right and viceversa.

In red-black trees, also a simple operation known as a color flip is used. In terms of 2-3-4 trees, a color flip is the essential operation: it corresponds to splitting a 4-node and passing the middle node up to the parent. A color flip obviously does not change the number of black links on any path from the root to the bottom, but it may introduce two consecutive red links higher in the tree, which must be corrected. Figure 2.13 illustrates a color flip operation example to split a 4-node.

In regard to the security applications of these structures, multiple works that explore certificate revocation and the publication of data collections on the Internet can be found in the literature [91] [92] [93] [94]. [91] propose a binary red-black hash tree for certificate revocation meanwhile [92] and [93] use the same kind of tree structure to guarantee secure sharing of information and authorized access in collaborative cloud scenarios.

On the other hand, in [94] Naor and Nissim used techniques from incremental cryptography in order to dynamize hash trees to support the insertion and deletion of elements. In their scheme, the source and the directory maintain identically implemented 2-3 trees. Each

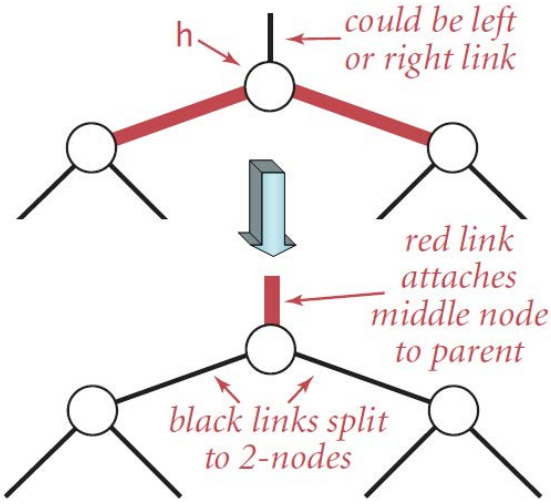

Figure 2.13: Flipping color example to split a 4-node.

leaf of such a 2-3 tree T stores an element of set S, and each internal node stores a one-way hash of its children's values. Hence, the source-to-directory communication is reduced to $O(1)$ items, since the source sends insert and remove instructions to the directory, together with a signed message consisting of a timestamp and the hash value of the root of T. A directory responds to a membership query for an element x as follows: if x is in S, then the directory supplies the path of T from the leaf storing x to the root, together with all the siblings of the nodes on this path; else (x is not in S), the directory supplies the leaf-to-root paths from two consecutive leaves storing y and z such that $y < x < z$, together with all siblings of the nodes on these paths. By tracing these paths, the user can recompute the hash values of their nodes by recalculating the hash value for the root, which is then compared against the signed hash value of the root for authentication.

There are nevertheless some drawbacks of this approach. Dynamic 2-3 trees are not trivial to program correctly, as it is. In addition, since nodes in a 2-3 tree can have two or three children, one must take special care in the structuring of the query authentication information sent by the directory to the user. Namely, all sibling nodes returned must be classified as being left children, middle children (if they exist), or right children. Recomputing the hash value at the root requires that a user be able to match the computation done at the source as regards a particular leaf-to-root path.

Eventually, as it will be shown in chapter 6, in the case of healthcare scenarios, patient's medical history or records do not have to follow a strict binary, ternary or quaternary

structure. Therefore, it is required to have a structure that allows to bring together user's information in a more flexible manner.

Merkle Trees

The well-known Merkle's tree [95] was the first authenticated dictionary structure. A Merkle tree is a binary tree, where leaf nodes are labeled by the hashed values of the elements of a set, S, and internal nodes are labeled by the hashed values of concatenated labels of their children. The root value is then the label of the root node, and the proof that element e belongs to S consists of the labels of all sibling nodes on the path from the leaf node representing e to the root node. In other words, a Merkle's tree is a hash tree T for a set S stores the elements of S at the leaves of T and a hash value $f(v)$ at each node v, defined as follows:

- If v is a leaf, $f(v) = h(x)$, where x is the element stored at x and h is a collision-resistant cryptographic hash function.
- Else (v is an internal node), $f(v) = h(f(u), f(w))$, where u and w are the left and right child of v, respectively.

The authenticated dictionary for S consists of the hash tree T plus the signature of a statement consisting of a timestamp and the value $f(r)$ stored at the root r of T. An element x is proven to belong to S by reporting the sequence of values stored at the siblings of the nodes on the path from the node storing x to the root. Each of these values must be identified as being stored at a left or a right child node, so that the user can correctly recompute the root's hash value and compare it to the current signed value.

It is important that all this order and connectivity information be presented to the user, for without it the user would have great difficulty recomputing the hash value for the root. This hash tree scheme can be extended to validate that an item x is not in S by keeping the leaves of T sorted and returning the leaf-to-root paths. Then, hash values associated, for two elements y and z such that y and z are stored at consecutive leaves of T and $y < x < z$, or y is undefined and z is the left-most leaf or z is undefined and y is the right-most leaf.

Figure 2.14 shows a perfect binary Merkle tree with height 2, because it is the extension of a perfect binary hash tree to a Merkle tree. In addition, we will refer to the height of a Merkle tree as the height of the underlying hash tree.

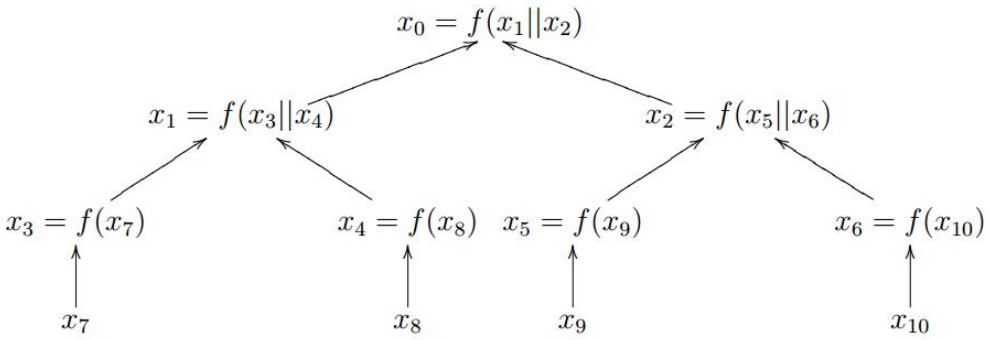

Figure 2.14: Example of a Merkle Tree.

The main goals of Merkle's trees are the following: 1) to make one-time signature schemes feasible and 2) to provide an efficient key management, that reduces the amount of public keys and their size.

Several enhancements to the original Merkle tree can be found in the literature. The improvements proposed in [96] and [97] basically consist of: 1) to use pseudo random number generators with a seed value to generate the private keys of the one-time signatures and 2) to use many smaller Merkle trees instead of one big tree. Furthermore, several approaches that use Merkle tree-based structures to provide security or privacy mechanisms in different application scenarios, such as cloud computing, smart grid, wireless networks and reputation systems can be found in the state of the art [98] [99] [100] [101] [102] [103].

In [98] the authors propose to integrate a Merkle tree with the homomorphic-authenticator-based technique while in [99] the authors suggest authentication scheme that considers the smart meters with computation-constrained resources. The work presented in [100] combines a dynamic reputation mechanism based on subject logic with the multi-level security technology in order to provide a privacy-aware secure hybrid wireless protocol. [101], in its turn, considers a reputation system which relies on Merkle trees, blind signatures, non-interactive zero-knowledge proofs and signed blocks of data to achieve anonymity preserv-

ing, minimize the workload on the trackers and to fairly distribute the record maintenance tasks to the service providers.

Eventually, Merkle trees can be generalized by a structure called "hash DAG", based on a Directed Acyclic Graph [104], thus allowing to extend the original Merkle tree (a binary tree) to an *m-ary* Merkle tree. Hence, our thesis contribution is based on an extended and unbalanced Merkle tree, because this structure enables richer clustering and to node verification using a single signature. Thus, we offer potentially a large number of verifiable attribute combinations by means of a single verification tree that empowers the user to realize a **selective disclosure** of her information to the different entities.

2.3.3 Skip Lists

Skip lists are another kind of ADT structure introduced by W. Pugh as an alternative data structure to search trees [105]. The main idea is to add pointers to a simple linked list in order to skip a large part of the list when searching for a particular element. While each element in a simple linked list points only to its immediate successor, elements in a skip list can point to several successors.

Thus, a skip list stores a set S of elements in a series of linked lists $S_0, S_1, S_2, \ldots, S_t$. The base list, S_0, stores all the elements of S in order, as well as sentinels associated with the special elements $-\infty$ and $+\infty$. Each successive list S_i, for $i \geq 1$, stores a sample of the elements from S_{i-1}. The method used to define the sample from one level to the next determines the kind of skip list being maintained. The default method is simply to choose each element of S_{i-1} at random with probability $1/2$ to be in the list S_i. But one can also define a deterministic skip list [106], which uses simple rules to guarantee that between any two elements in S_i there are at least 1 and at most 3 elements of S_{i-1}. In either case, the sentinel elements $-\infty$ and $+\infty$ are always included in the next level up, and the top level, t, is maintained to be $O(\log n)$, where n is the number of list elements.

In both the deterministic and the randomized versions, the top level is guaranteed to contain only the sentinels. We therefore distinguish the node of the top list S_t storing $-\infty$ as the start node S. A node of S_{i-1} storing an element that does not exist in S_i is said to be a plateau node. A node that is not a plateau node is said to be a tower node. Thus, between any two tower nodes of the same list, there are some plateau nodes. In

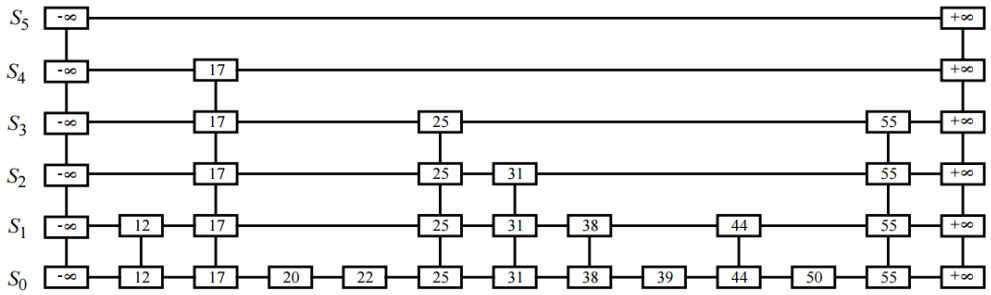

Figure 2.15: Example of a skip list.

deterministic skip lists, the number of plateau nodes between two tower nodes is at least one and at most three. In randomized skip lists, the expected number of plateau nodes between two tower nodes is one (See Figure 2.15).

Moreover, a skip list data structure supports the following operations: find, which allows to determine whether elements x is in S; insert and remove, which enable to both insert x into S and remove x from S, respectively.

2.3.4 Identity and Attribute-Based Encryption Approaches

Identity-Based Encryption (IBE) [107] enables for a sender to encrypt a message to an identity without access to a public key certificate. The IBE systems view identities either as a string of characters or a set of descriptive attributes [108].

Shamir [107] first proposed the concept of Identity-Based Encryption. However, it was not until much later that Boneh and Franklin [109] presented the first Identity-Based Encryption scheme that was both practical and secure. Their solution made novel use of groups for which there was an efficiently computable bilinear map. Canetti et al. [110] proposed the first construction for IBE that was probably secure outside the random oracle model. To prove security they described a slightly weaker model of security known as the Selective-ID model, in which the adversary declares which identity he will attack before the global public parameters are generated. Boneh and Boyen [111] give two schemes with improved efficiency and prove security in the Selective-ID model without random oracles.

Concerning attribute-based encryption (ABE) models, in these approaches each user has a set of attributes and access policies that are defined to determine that the users with certain attributes are authorized to access the shared data. ABE cryptosystems [108] crowd in two categories: ciphertext-policy ABE (CP-ABE) [112] systems and key-policy ABE (KP-ABE)) [113] systems. In the first, the users' secret keys are associated with sets of attributes, and a sender generates a ciphertext with an access policy specifying the attributes that the decryptors must have. Regarding KP-ABE solutions, the users' secret keys are labeled with access policies and the sender stipulates a set of attributes; only the users whose access policies match the attribute set can decrypt.

It is worth to be mentioned that, in an attribute-based encryption system ciphertexts are not necessarily encrypted to one particular user as in traditional public key cryptography. Instead both users' private keys and ciphertexts will be associated with a set of attributes or a policy over attributes. A user is able to decrypt a ciphertext if there is a "match" between his private key and the ciphertext. In their original system Sahai and Waters [108] presented a threshold ABE system in which ciphertexts were labeled with a set of attributes S and a user's private key was associated with both a threshold parameter k and another set of attributes S'. In order for a user to decrypt a ciphertext at least k attributes must overlap between the ciphertext and his private keys.

The primary drawback of the Sahai-Waters [108] threshold ABE system is that the threshold semantics are not very expressive and therefore are limiting for designing more general systems. Goyal et al. [113] introduced the idea of a more general key-policy attribute-based encryption system. In their construction a ciphertext is associated with a set of attributes and a user's key can be associated with any monotonic tree access structure. The approach of Goyal et al. can be viewed as an extension of the Sahai and Waters techniques where instead of embedding a secret sharing scheme in the private key, the authority embeds a more general secret sharing scheme for monotonic access trees.

This approach solves granularity problems of the approach presented in [108] and it works well to manage more or less permanent user' profile attributes that do not vary over time. But in those cases in which the attributes of the user's profile are very changing, it could present scalability problems.

2.4 Current standards for personal information managing in sensitive scenarios

Sharing of Electronic Health Records among healthcare providers has experienced a noteworthy growth in the last years, since it enables physicians to remotely monitor patients' health and enables individuals to manage their own health data more easily. The EHR is a longitudinal collection of electronic healthcare information about individual patients as well as populations. A typical EHR system consists of several subsystems, such as appointments and scheduling; admission, discharge, and transfer (ADT); prescription order entry; dietary planning; routine clinical notes; lab and radiology orders; picture archiving, and smart card sign-on.

However, these scenarios face significant challenges regarding security and privacy of the extremely sensitive information contained in EHRs. Several government bodies, non-government organizations, public hospitals and private clinics have joined forces to implement numerous security measures ranging from regulatory control to security process and technologies. For instance, the HIPAA (*Health Insurance Portability and Accountability Act*) [114] is a well-known federal law, which protects health information and ensures that patients have access to their own medical records, while giving new responsibilities to those in charge of protecting this information.

Nowadays, there are several EHR standards as Health Level Seven International (HL7) [115], OpenEHR [116] and International Organization for Standardization (ISO) EN 13606 [117] compliant with the HIPAA and that we have studied, in order to carry out the validations of our thesis contributions. Particularly, **we have selected health care scenarios, since they are among the most sensitive environments**.

Current e-health standards are based on a dual model architecture, which defines two conceptual levels: reference model and archetype model. The reference model defines the set of entities that form generic building blocks of the electronic healthcare record. The archetypes define clinical concepts in the form of structured and constrained combinations of the entities contained in the reference model, so clinical knowledge is defined at this level. Both OpenEHR and ISO EN 13606 use this modeling architecture, which has also influenced HL7 CDA.

On the other hand, in order to facilitate the interoperability and provide integration capabilities in the exchange of such EHRs, initiatives as Integrating Healthcare Enterprise (IHE) Profiles [118] have emerged, but holistic implementations of IHE based e-Health infrastructures to share EHRs are currently rare. IHE describes real world use cases. In addition, it tries to solve security and privacy issues by a modular approach, which defines Integration Profiles and use cases to address several issues, such as communication of claims about the identity of an authenticated users, exchange of medical or care data or recording of the patient privacy consents.

Nevertheless, the current e-health standards do not deal with some aspects of privacy, such as selective identity information disclosure or fine-grained control of sensitive data that enable open sharing across different healthcare stakeholders in a secure way.

In chapter 6 will describe and detail our proposal to address the above privacy issues. For that thesis contribution, we have selected OpenEHR, because it offers an open and extensible framework, as well as archetypes for many clinical terms widely used in hospitals and summary EHR systems in multiple countries. Here we summarize three well-known EHR specifications in order to provide a brief background that contextualizes the thesis.

2.4.1 OpenEHR

The OpenEHR Foundation [119] is a not-for-profit company, which was established in 2000 in the UK. Members of OpenEHR work closely with standardization bodies, including ISO TC215, CEN TC/251, HL7 and national standards bodies. It publishes e-heath related specifications, open source software, domain models (archetypes) and educational material around a platform architecture. Its specifications include: the OpenEHR Reference Model (RM), consisting of the primary information models (IMs), the archetype model (AM) which includes the Archetype Definition Language (ADL) and Archetype Object Model (AOM), and the OpenEHR service model (SM), which defines interfaces to major software services in a health information environment. Release 1.0.2 [116] of its information models are widely used in the industry. The Archetype Definition Language (ADL) v1.4 [120] and its associated Archetype Object Model (AOM) specification [121] are ISO standards. OpenEHR is a founding member of the Clinical Information Modelling Initiative (CIMI)[6],

[6] http://informatics.mayo.edu/CIMI/index.php/Main_Page

and OpenEHR's ADL 2 draft standard[7] is the language of CIMI models. Since 2012, industry interest in OpenEHR's open platform health computing approach has grown substantially. There are now around ten "Industry partner" (financially supporting) vendor companies (primarily from Europe). These are responsible for OpenEHR systems running in a growing number of hospitals and health authorities worldwide. Furthermore, there are active semantic health modeling programmes or projects based on OpenEHR archetypes in Australia, Brazil MoH, Moscow City, Norway MoH, Slovenia MoH, and United Kingdom NHS.

On the other hand, OpenEHR provides standards for: clinical (EHR) and demographic data (the OpenEHR Information Models [116]); clinical (EHR) and demographic content models, and connection points to terminology through the OpenEHR *archetypes* and *templates*; guidelines, portable queries and key services and APIs including REST services generated from archetypes.

The OpenHER Architecture

Strategically, the OpenEHR approach enables a platform-based e-health software market, in which vendors and developers solutions interface via standardized information models, content models, terminologies and service interfaces. This gives procurement stakeholders new choices, enabling them to avoid product and vendor lock-in; retain ownership of the data for secondary use, as well as let their clinical experts be directly involved in solution development, via archetype authoring. Thus, it also allows application developers to concentrate on their applications, and simply plug in to a reliable back-end.

The OpenEHR technical approach is "multi-level modeling within a service-oriented software architecture", in which models built by domain experts are in their own layer (See Figure 2.16).

This allows domain experts to be directly involved in defining the semantics of clinical information systems, and it also makes using terminology much easier. OpenEHR offers an international repository of these models, known as "archetypes" and its archetype specification is an ISO standard (ISO 13606-2). These are now being used by several national governments as described above. OpenEHR also defines specifications for clinical

[7] http://www.openehr.org/releases/AM/latest/docs/ADL2/ADL2.html

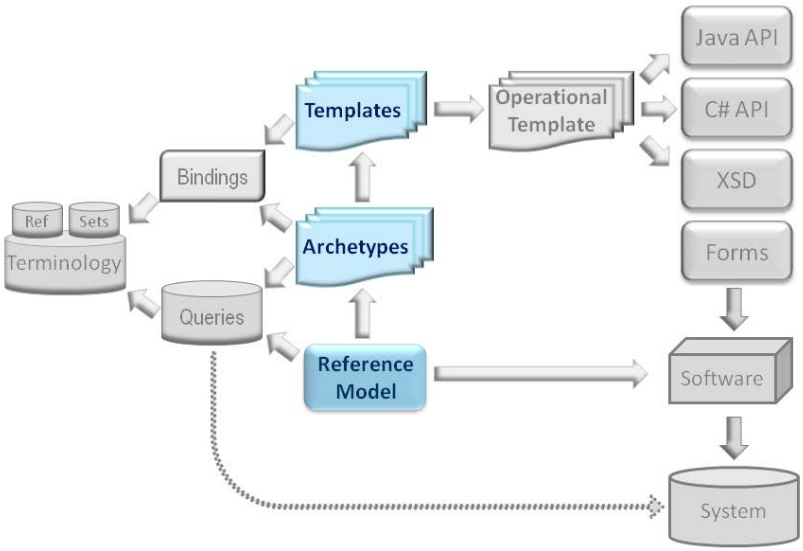

Figure 2.16: OpenEHR Technical Architecture: The OpenEHR technical approach is multi-level modeling within a service-oriented software architecture, in which models built by domain experts are in their own layer (©[122]).

information models, EHR Extracts, demographics, data types and various kinds of service interfaces. These have been used in hospitals and summary EHR systems in various countries. It has also included a leading edge Guideline Definition Language (GDL), originally developed by Cambio in Sweden, as a specification.

A second dimension via which the OpenEHR modeling approach can be viewed is single-source modeling. Via this approach, archetypes and templates are definitive models of semantics, without commitment to specific messaging or document standards. Instead, these concrete expressions are now generated artifacts, i.e., document and message schemes are no longer manually modeled. Once single-source modeling is established, other outputs including user interface (UI) forms and software source code. This means that a model for 'microbiology result' only needs to be done once to enable reports, UI forms, documents and other message formats to be generated.

Archetypes and Templates

The archetypes are key elements of the OpenEHR methodology. They are detailed and domain-specific definitions of clinical concepts in the form of structured and constrained

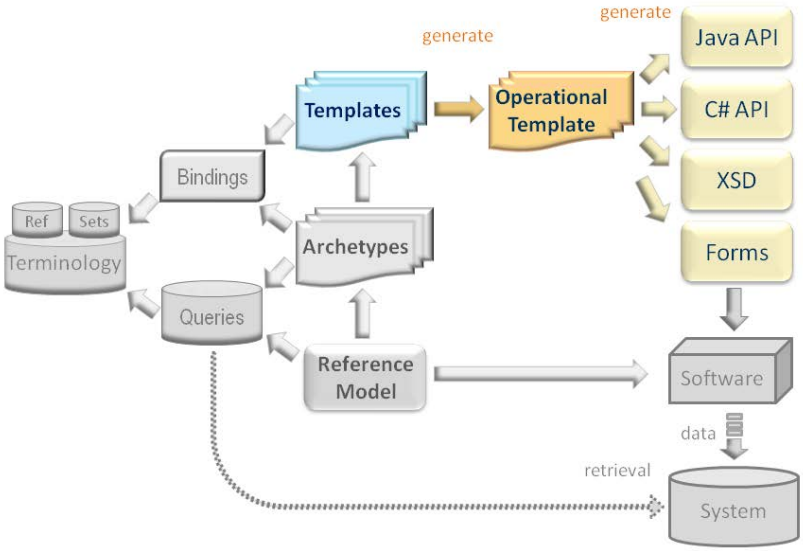

Figure 2.17: OpenEHR Technical Architecture: The OpenEHR single-source modeling approach (©[122]).

combinations of the entities of the reference model. Archetypes refer to clinical concepts and represent healthcare and application specific concepts such as blood pressure, examination of the chest, heart rate, etc. In other words, they are reusable, structured models of clinical information concepts that appear in EHRs, such as “test result” , “physical examination” and “medication order”, and are expressed in terms of constraints on the reference model. All data in OpenEHR EHRs are instances of reference model entities, configured by archetypes. Archetypes also act as mediators between data and terminology. They are external to the software, and due to their adaptability, health information systems can be computed at a semantic level and enable functions such as decision support and research query.

Clinical concerns and the technical design of data storage can be separated through archetypes and the two-level modeling. The first level involves the technical concerns, information structure and data types using the Reference Model, while the second level handles clinical domains and semantics. Key attributes of health records are managed by the reference model and do not need to be addressed by each archetype. The archetypes represent discrete specifications of clinical information, which are as inclusive as possible;

they offer great advantages since data can be specified in an understandable manner to health professionals and IT staff.

Regarding, the templates, they are locally defined models of screen forms, and ring together a selection of archetypes, terminologies, language and other details relevant to the particular local use of archetypes. For example, concepts such as "referral" and "prescription" are modeled as templates, which in turn use archetypes for more fine-grained concepts. OpenEHR and the ISO EN 13606 [117] communities specify them using the ADL [120]. This language provides an abstract syntax, which can be used to express archetypes for any reference model in a standard way. In addition, an archetype can include other archetypes and can be used in combination to form templates. Moreover, archetypes are envisaged as a clinical guide for clinicians.

Security and Privacy Considerations

The OpenEHR specifications and core component implementations do not explicitly define many concrete security and privacy mechanisms. OpenEHR supports some of the key requirements in a flexible enough way that deployments with substantially different requirements. These include EHR/demographic separation and an EHR-wide access control object. At the level of versioned objects, commit audits are mandatory and digital signatures and hashes are available. The main security policy principles supported by OpenEHR are detailed below:

- **Non-repudiation**: OpenEHR supports digital signing of communications in order to guarantee non-repudiation of information exchanged between health care systems.
- **Access control and access lists**: An actor, named "gate-keeper", is responsible for controlling access to the EHR access control configurations. The gate-keeper is established at the EHR creation time as being one of the identities known in the EHR (e.g., the patient, a parent, legal guardian or another responsible person.). Moreover, it is in charge of determining who can change the access control list.
- **Privacy**: The content of the health record is separated from identifying demographic information. Moreover, patients can label compositions in the EHR as having one of a number of levels of privacy.
- **Audit trailing**: The EHR status, access control objects and other changes made to

the EHR, including content objects are audit-trailed with user identity, time-stamp, reason, optionally digital signature and relevant version information. If the updater is the patient, a pseudo-identifier can be used.

- **Non removable**: As health record information cannot be deleted, the logical deletion implemented in version control is accomplished by labeling the data in such a way as to make it appear deleted.

Summarizing, there are some key benefits to OpenEHR's approach. Firstly, an EHR system does not need to know a priori about any of the clinical data it will process, such as vital signs, diagnoses or orders. Models for those things are developed separately. Models for data sets and forms are also developed separately, and UI form components are now generated from these definitions. This enables a new generation of EHR systems that routinely adapts to new requirements. Another benefit is portable queries and decision support logic, since queries in OpenEHR are based on content models, not physical database schemas. Coupled with EHR service interface APIs, these are enabling a new class of decision support tools. Thus, with the Guideline Definition Language[8], it is finally possible to express clinical logic that is truly agnostic to clinical domains, natural languages and reference terminologies.

However, with respect to security and privacy, OpenEHR imposes only a minimal security policy profile, which could be regarded as necessary, but generally not sufficient for a deployed system and other aspects would still need to be implemented in layers whose semantics are not defined in OpenEHR. For instance, mechanisms that limit the time during which given health professionals can see the patient record as well as privacy settings to define access control behaviour to patient's EHRs should be implemented. In this sense, there are various efforts in progress, including the CEN EN13606 [117] part 4 work, the ISO PMAC (*Privilege Management and Access Control*) [123] work being done in TC/215[9] based on the generic security standard ISO/IEC 17799 [124]. Nevertheless, no large-scale shared EHR deployments exist and security solutions to date are still developmental.

[8] http://www.openehr.org/releases/CDS/latest/GDL.html
[9] http://www.iso.org/iso/iso_technical_committee?commid=54960

2.4.2 The ISO/EN 13606 Standard

The CEN/ISO EN 13606 [117] is a European norm from the European Committee for Standardization (CEN) also approved as an international ISO standard. It is designed to achieve semantic interoperability in the electronic health record communication. The main objective of this standard is to define the way that EHRs are exchanged, but it specifies neither the internal architecture of an EHR system nor the way data are stored. Hence, the ISO 13606 standard aims at defining a rigorous and stable information architecture for communicating part or all of the EHR of a patient between EHR systems, or between EHR systems and a centralized EHR data repository. It may also be used for EHR communication between an EHR system or repository and clinical applications or middleware components (such as decision support components) that need to access or provide EHR data, or as the representation of EHR data within a distributed (federated) record system.

ISO 13606 follows a dual model architecture, which defines a clear separation between *information* (Reference Model) and *knowledge* (Archetype Model) [125]. The former is structured through a Reference Model that contains the basic entities for representing any information of the EHR. The latter is based on archetypes, which are formal definitions of clinical concepts (i.e., discharge report, glucose measurement, family history, etc.) in the form of structured and constrained combinations of the entities of a Reference Model.

Furthermore, it provides a semantic meaning to a Reference Model structure and specifies how health data should be aggregated and the context information that must accompany every piece of data in order to meet ethical and legal requirements. The interaction of the Reference Model to store data and the Archetype Model to semantically describe those data structures; provides an unseen capability of evolution to the information systems.

Typically, a Reference Model contains sets of primitive types, classes that define the building blocks of EHRs and set of auxiliary classes, that describe the context information to be attached to an EHR annotation. It may contain classes to describe demographic data and to communicate EHR fragments. Figure 2.18 (extracted from ISO/EN13606-1) shows a simplified scheme of the basic generic components that support information and the relationships between those components.

In regard to the archetypes, they are structured and constrained combinations of entities of a Reference Model that represent a particular clinical concept, such as a blood pressure

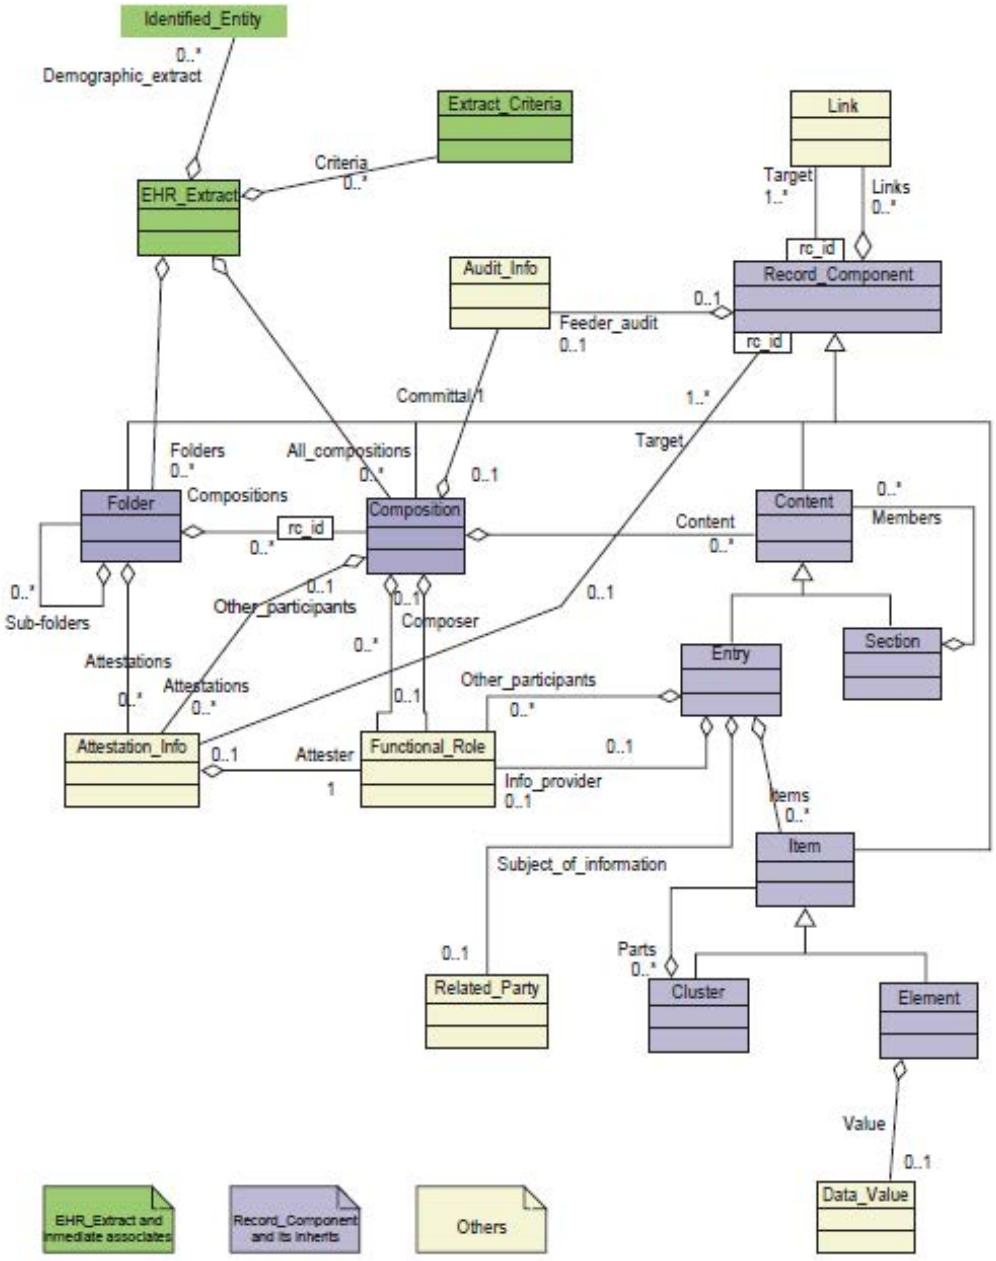

Figure 2.18: ISO/EN13606 Reference Model (simplified scheme from ISO/EN13606-1©[117]).

measurement or a laboratory analysis result. The archetypes can be also defined by further constraining other archetype, called parent archetype, in order to obtain a more adequate or fine grained representation of the clinical concept or to limit the value range of an attribute (see Figure 2.19). Since this part of the ISO/EN13606 norm leverages appropriate parts of the OpenEHR model for defining archetypes, OpenEHR and ISO/EN13606 share the basis of the archetype model [126].

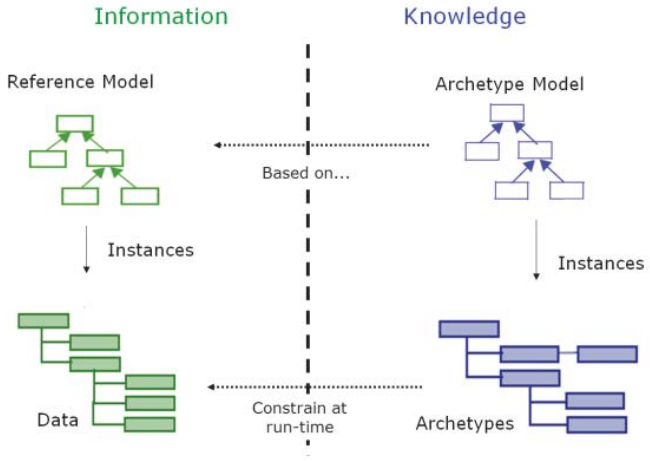

Figure 2.19: Relationship between information (instances of Reference Model) and knowledge (instances of Archetype Model)(©[126]).

Although the main feature of ISO/EN 13606 is the dual model, it is also important to define other aspects in order to achieve interoperable exchange of EHR, such as nomenclature issues, security issues and interfacing for querying:

- ***Reference Archetypes and Term lists:*** This part establishes a normative set of coded terms, each one defining a controlled vocabulary for a Reference Model attribute contained in ISO/EN 13606-1. It includes different groups of terms such as terms related to the subject of an Entry (`SUBJECT_CATEGORY`), the category of information of any `ELEMENT` or `CLUSTER` (`ITEM_CATEGORY`), the status of a particular version of a record component (`VERSION_STATUS`), etc.

- ***Security*:** The ISO/EN 13606 specification describes a methodology for specifying the privileges necessary to access EHR data and some other general security requirements that should apply to EHR communications. It also defines both general and

specific access policies able to deny or grant access to medical records to identified parties or specific functional roles.

- ***Interface Specification:*** This part describes a set of interfaces to request access to the information and resolve the request. Three specific interfaces are defined to request a specific EHR extract, one or more archetypes and to inquire a specific EHR audit log extract.

Finally, another remarkable feature of ISO/EN 13606 is the alignment it presents to other relevant standards such as HL7, OpenEHR, IHE [118], etc.

2.4.3 HL7

Health Level Seven International is a not-for-profit, ANSI-accredited standards developing organization which was founded in 1987. It is dedicated to providing a comprehensive framework and related standards for the exchange, integration, sharing, and retrieval of electronic health information that supports clinical practice and the management, delivery and evaluation of health services.

"Level Seven" refers to the highest level of the ISO communications model for Open Systems Interconnection (OSI) the application level. The application level addresses definition of the data to be exchanged, the timing of the interchange, and the communication of certain errors to the application. The seventh level supports such functions as security checks, participant identification, availability checks, exchange mechanism negotiations and, most importantly, data exchange structuring [115].

The Reference Information Model (RIM) is the cornerstone of the HL7 Version 3 development process. An object model created as part of the Version 3 methodology, the RIM is a large pictorial representation of the clinical data (domains) and identifies the life cycle of events that a message or groups of related messages will carry. It is a shared model between all the domains and as such is the model from which all domains create their messages. Explicitly representing the connections that exist between the information carried in the fields of HL7 messages, the RIM is essential to increasing precision and reducing implementation costs [115]. The RIM consists of templates, vocabulary, and XML standards.

Concerning HL7 templates, they are data structures based on the HL7 RIM which express the data content needed in a specific clinical or administrative context. They are prescribed patterns by which multiple observation result segments may be combined to describe selected, gross observations, etc. The observation result segment is primarily used to carry key clinical observation/results reporting information within report messages, which must be transmitted back to the requesting system. Some observations may be quite simple, such as the blood pressure concept in healthcare, which involves a set of expected observations (i.e., systolic, diastolic, patient position, method, etc.) Other more elaborate diagnostic procedures may involve hundreds of related pieces of information, including anatomy, orientation, sequences of measurements, etc.

Templates provide a means of coupling the multiple observation result segments needed to send the observation with separately encapsulated rules for combining/validating them for the particular observation. Based on user need and preference, the template offers the user the advantage of defining the collection of observation result segments needed and the corresponding set of validation rules once, and once, defined, the structure can be used again and again. Since they are based on a specific user's needs/requirements, templates can be "plug and play" at a given user site [115].

2.5 Related Work

Today, there are several approaches to identity management being the most popular the federated and user-centric approaches. As will be shown in Chapter 4, 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. Current IdM systems are not ready to deal with some aspects of privacy, which are especially critical in sensitive environments, such as healthcare scenarios. Specifically, user's consent revocation, is not covered by any of the aforementioned identity management approaches. This fits with the privacy view as control over the use and flow of one's personal information [3]. Revoking consent allows users grant or withdraw consent of specific actions over data to certain individuals. So this mechanism is useful to enforce user's role in the task of preserving her privacy.

This property is part of the privacy rules described by the HIPAA for health, OECD (*Or-*

ganization for Economic Cooperation and Developments) principles and GLBA (*Gramm-Leach-Bliley Act*) for financial institutions. This is not in the European Union, where a person's consent cannot be given prospectively and where consent must be fully informed [127]. The E-Privacy Directive [128] addresses particular concerns in the use of electronic communications to deal with personal data; however it alone does not provide an adequate revocation solution, conferring as it does only limited rights on individuals to prevent types of processing by withdrawing consent to such processing [129].

Furthermore, a flexible, efficient and standard-based privacy-awareness profile management solution to guarantee selective identity information disclosure and preserve user's privacy is indispensable. This section reviews both the related work carried out by researchers (individually or in the framework of a research project), as well as the standardization initiatives that are related or may contribute in any aspect to provide privacy-enhancing tools that empower the user role and permit secure information sharing across different stakeholders, while data protection against unauthorized use and minimal disclosure are provided.

2.5.1 Previous Work

Privacy and security are crucial issues are becoming important concerns for researchers when managing personal information in very extremely sensitive environments, such as e-health applications. So, this aspect must be taken into account when exchanging and sharing data between different service providers and third parties, preventing personal or sensitive data for being misused.

However, for instance, in the case of healthcare organizations, because of the complex nature of data access for diverse purposes, often give broader access privileges and adopt "Break the Glass" (BTG) policy to facilitate timely and effective care. Rostad and Edsburg [130], for example, report that 99% of doctors were given overriding privileges while only 52% required overriding rights on regular basis, the security mechanisms of health information systems were overridden to access 54% of patients' records. Another common pitfall of BTG policy is that such broad-based privileges could be misused by employees.

Nevertheless, in this sense, there is still scarce work on effective revoking consent mechanisms within sensitive environments.

The approach presented in [131] is close to our work. The authors propose an activity-oriented access control model to protect the confidentiality of health information, which consists of three levels: user level, activity level and privilege level. Thus, this proposal is based on user activity to authorize access privileges and defines two revocation mechanisms called single-step and multi-step revocations. However, possible activities in the hospital and policies need to be defined in advance.

Bhatti and Grandison [132] proposed a privacy management architecture (PRIMA) model that leverages artifacts such as audit logs arising from the actual clinical workflow to infer and construct new privacy protection rules. In particular, PRIMA implements a policy refinement module that periodically examines the access logs and identifies new policy rules using sophisticated data-mining techniques. These audit logs could, as well, be used by privacy officials to determine privacy violations, which in itself is a complex process and often requires merging data from disparate sources [133]. Unfortunately such data merging may cause potential disclosure of patients' sensitive information to the investigators against the patients' consent.

In [134], a privacy-aware role-based access control (P-RBAC) is presented, in order to express privacy policies. These policies are seamlessly integrated with access control. In this same way, [135] presents the notion of consent and revocation policies to express user's preferences, within the context of the EU FP7 EnCoRe project [136]. This work has been extended in [137] and [138] by proposing a conceptual model for privacy policies that can be integrated with XACML (eXtensible Access Control Markup Languages). They also consider revocation of personal data as well as previously granted privileges. Nevertheless, these works propose traditional revocation mechanisms such as temporal information being used to predefine policies. They do not take into account dynamic scenarios or information systems designed for emergency situations.

On the other hand, as far as privacy-enhancing tools for user profile management and selective disclosure of identity are concerned, we found several research initiatives in the field. Many authors have suggested security and privacy as key issues to address in eHealth [139] [140] [141], but these issues as a whole have not yet been covered extensively for application scenarios. The focus is normally on security related issues in general wireless sensor networks.

Nowadays, several approaches to provide privacy-preserving techniques can be found in

the literature [108] [113] [112] [142] [143] [144] [145] [146] [147]. Firstly, in attribute-based (ABE) encryption models each user has a collection of attributes and access policies are defined to determine that the users with certain attributes are authorized to access the shared data. ABE cryptosystems [108] can be classified as ciphertext-policy ABE (CP-ABE) [112] systems and key-policy ABE (KP-ABE) [113] systems. In the CP-ABE systems, the users' secret keys are associated with sets of attributes, and a sender generates a ciphertext with an access policy specifying the attributes that the decryptors must have. In regard to KP-ABE approaches, the users' secret keys are labeled with access policies and the sender specifies a set of attributes; only the users whose access policies match the attribute set can decrypt. In [148] authors suggest a multi-authority CP-ABE scheme to empower to the patient to associate an expressive access tree structure and on-demand attribute revocation.

However, these ABE schemes require a priori access policies, which are not always available in EHRs because the policies to access health records are sometimes determined after key generation. [149] addresses this issue by considering a dynamic ABE paradigm, which provides a delegation mechanism that allows users to redefine the access policy and delegate a secret key without making the redefined access policy more restrictive. Nevertheless, how to construct fully secure hierarchical identity-based encryption systems in prime-order bilinear groups under simple assumptions remains as a challenging problem [150].

Secondly, cloud-based approaches as [142] and [143] propose privacy-aware schemes based on query authentication to enable data confidentiality, the query result integrity of sensitive data, secure storage and secure computation auditing. The work presented in [144] integrates a PRF-based key management for unlinkability, a search and access pattern hiding scheme based on redundancy for privacy-preserving data storage. This approach also combines ABE-controlled threshold signing with role-based encryption to provide access control and auditability. A signature algorithm that allows for controlled changes to the signed data is proposed in [151]. This work studies techniques that cryptographically link the integrity of the original and modified datasets for practical types of modifications such as redaction, pseudonymization and data deidentification.

Thirdly, when it comes to information disclosure, spatio-temporal cloaking and ADT-based approaches enable to preserve user's privacy. Spatial cloaking or perturbation allows to hide the participant location inside a cloaked region using spatial transformations,

generalization, or a set of dummy locations in order to achieve location privacy [152]. In [153] authors propose a privacy-preserving emergency call scheme called PEC, enabling patients in life-threatening emergencies to fast and transmit emergency data to the nearby helpers via mobile healthcare social networks. Once an emergency happens, the patient's mobile device runs the PEC to collect the emergency data including emergency location, patient health record and patient physiological condition. Then, the PEC generates an emergency call with the emergency data inside and epidemically disseminates it to every user in the patient's neighborhood. Moreover, the PEC has been designed to withstand multiple types of attacks, such as identity theft attack, forgery attack, and collusion attack. However, this kind of works do not address privacy issues related to management of user profiles.

Other approaches closer to our work are some identity frameworks like U-Prove [48] allow selective disclosure of claims and pre-signed tokens that could be used when the entity responsible for issuing medical records is offline. Furthermore, there are other proposals such as identity agents [154], or veryIDX [155]. [154] proposes user-controlled identity agents, which allow defining in advance disclosure policies, monitoring credential usage, storing credentials based on a minimal disclosure scheme. The credentials are constructed using Merkle trees, but the details about how patient's attributes are built or can be shared by means of EHR standards are not provided. VeryIDX enables multi-factor identity credential verification, by using a cryptographic commitment and an aggregated zero-knowledge proof of knowledge (ZKPK).

Despite some current works [145],[156], [146] propose the use of authenticated dictionaries or opportunistic computing mechanisms to provide selective information disclosure, none of these works deals with neither building of the ADT structure based on EHR standards nor combining subtrees that allows claims from different sources to be in a single credential in order to make easier the tasks of management of patient attributes, profiles and preferences. In this context, typical research directions are related to development of more efficient and effective ADT-based structures and algorithms, in terms of storage overhead, times of signature generation and verification or query and update times.

[145] proposes a signature scheme on the structure of the tree as defined by tree traversals (pre-order, post-order, in-order), that improves protection against information leakages. [156] describes a secure and privacy-preserving opportunistic computing framework,

called SPOC, for m-Healthcare emergency. The authors introduce an attribute-based access control and a privacy-preserving scalar product computation technique that allows a medical user to decide who can participate in the opportunistic computing to assist in processing his overwhelming personal health information data. Eventually, in [146], the authors propose a multiway extension of the authenticated version of the skip-list data structure and study the authentication cost that is associated with this model when authentication is performed through hierarchical cryptographic hashing. However, due to the heterogeneity of data types in healthcare scenarios, this kind of structure requires a complex implementation.

2.5.2 International Projects

On the other hand, there are a number of key research projects primarily funded by the European Commission - that are involved (or have been involved) in privacy and identity management related topics. We gather here the most relevant ones. More specifically, within the Seventh Research Framework Programme of the European Union from 2010 to 2016, several new projects related to privacy and identity management started.

ABC4Trust [157] project has defined a common, unified architecture for privacy-attribute based credentials systems to enable comparing their respective features and combining them on common platforms. Its main contribution is the specification of the data artifacts exchanged between the implicated entities (i.e. issuer, user, verifier, revocation authority, etc.), in such a way that the underlying differences of concrete privacy-attribute based credentials implementations are abstracted away through the definition of formats that can convey information independently from the mechanism-specific cryptographic data.

It also defines all technology-agnostic components and corresponding APIs a system needs to implement in order to perform the corresponding operations, perform the selection of applicable credentials for a given policy or to trigger the mechanism-specific generation of the cryptographic evidence.

STORK 2.0 (Secure idenTity acrOss boRders linKed) [158] and FutureID [159] projects aim at building an identity management infrastructure for Europe in support of a single market of online services. ENDORSE [160] is concerned with providing a Legal Technical Framework for Privacy Preserving Data Management. The EnCoRe [136] project aims to

offer the data subject consent and revocation controls through which an individual could manage the flow of her personal data. The overall vision of this project is "make given consent as reliable and easy as turning on a tap and revoking that consents reliable and easy as turning it off again" [161].

PICOS (Privacy and Identity Management for Community Services) [162] investigates and develops a state-of-the-art platform for providing trust, privacy and identity management in mobile communities. PrimeLife [163], works on privacy-enhancing technologies that can enable citizens to execute their legal rights to control personal information in on-line transactions. Thus, the project is advancing the state-of-the-art in the areas of human computer interfaces, configurable policy languages, web service federations, infrastructures and privacy enhancing cryptography. For this purpose, PrimeLife works with the relevant open source communities and standardization bodies.

The SWIFT [164] project (Secure Widespread Identities for Federated Telecommunications) leverages identity technology as a key to integrate service and transport infrastructures for the benefit of users and the providers. It focuses on extending identity functions and federation to the network while addressing usability and privacy concerns. The research within SWIFT includes a gap analysis to identify challenges in existing identity frameworks, a requirements list to address these gaps and a generic architecture based on the requirements.

Summarizing, ABC4Trust, STORK 2.0, FutureID, PICOS and PrimeLife are more concerned with privacy enhancing tools whilst ENDORSE and EnCoRe are concentrated on providing privacy preserving and user's consent mechanisms. In regard to SWIFT, it focuses on improving federation functions.

Finally, thought all the projects embrace security and privacy considerations to some extent, none of them provides a comprehensive solution to the specific case of revoking consent when the user is unconscious and cannot give her express consent, as we develop in this thesis. In conclusion, current proposals do not provide a general solution that address how to provide a hybrid and time-independent identity management model, which includes delegation consent mechanisms. Finally, we aim to address more dimensions of privacy, such as a richer patient profile management while enabling to bring together several sources of EHRs to be part of a unique credential, that are not considered in the presented approaches.

2.5.3 Standards Developing Organizations and Related Bodies

Several standardization developing organizations and related bodies are working on identity management topics that conform fundamental pieces to provide enhancing privacy tools for sensitive and dynamic environments, such as health care scenarios. In the following, we name this organizations and explain their work and how it relates to the vision of privacy preserving techniques in identity management platforms.

Organization for the Advancement of Structured Information Standards (OASIS)

OASIS leads several efforts in the standardization of federation standards. As previously documented in this chapter, SAML, WS-Federation, and Identity Interoperability Mestasystem for Information Cards, are federation frameworks standardized by OASIS. In addition, apart form these mature identity standards, OASIS created other Technical Comitees (TCs) that are also related to identity management. The most relevant groups that are addressing issues of privacy are:

- OASIS Cross-Enterprise Security and Privacy Authorization TC.

  This group, created in 2008, focuses on the development of healthcare profiles of existing OASIS standards, such as SAML and XACML, used to exchange interoperable security and privacy attributes within and between organizations. Moreover, this TC concentrates on specifying healthcare profiles of existing OASIS standards to support reliable, auditable methods of confirming personal identity, official authorization status, and role attributes. This work aligns with security specifications being developed within the U.S. Healthcare Information Technology Standards Panel (HITSP). In their main document, "*Cross-Enterprise Security and Privacy Authorization (XSPA) Profile of SAML v2.0 for Healthcare Version 2.0*", they define a set of attributes for using SAML assertions that can be used to capture requests for exchange of healthcare information.

- OASIS Privacy Management Reference Model (PMRM) TC.

  The PMRM group, formed in 2013, works to provide a standards-based framework that help business process engineers, IT analysts, architects, and developers imple-

ment privacy and security policies in their operations. The most relevant technical work produced by the committee so far is "*Privacy Management Reference Model and Methodology (PMRM) Version 2.0*". It provides a guideline for developing operational solutions to privacy issues. It also serves as an analytical tool for assessing the completeness of proposed solutions and as the basis for establishing categories and groupings of privacy management controls.

Kantara Initiative

Kantara Initiative was announced on 2009, by leaders of several foundations and associations working on various aspects of digital identity. It is intended to be a robust and well-funded focal point for collaboration between members of the identity community. Kantara Initiative is bridging the enterprise, mobile, government and Web communities to provide the industry with a clear path for moving interoperable identity systems forward, advancing adoption and meeting marketplace. The organization is structured into working groups that deal with different aspects of identity management. The groups that are more related to privacy and focused on sensitive environments, such as healthcare scenarios, are the following:

- *Consumer Identity Work Group.* The purpose of this group is to foster the development of a consumer-friendly, privacy-protecting, high assurance "identity layer" for the internet that enables consumers to fully exploit the potential of the internet without fear of identity theft. They are working on proposing technical and policy solutions that address current threats to privacy and identity, and socializes these solutions with appropriate parties to help foster their implementation. Specifically, this group will describe how emerging identity technologies, protocols, frameworks, laws and regulations, etc., can be leveraged to: (a) enable businesses to know, with high confidence, the identities of individual consumers with whom it engages in high-value online transactions, without jeopardizing the privacy of the consumer' Personally Identifiable Information; and (b) enable individual consumers to prevent others from impersonating them in high-value, online transactions.

- *Healthcare Identity Assurance Work Group.*

  This group works on designing, implementing and testing reference applications for

secure access to health information. Two use cases are proposed that would be developed and supported as part of the work group. One is for consumers to be able to access their health records with a standardized login system, and secondly, a way for healthcare workers to access secure health information. The goal of this activity is to engage the broadest community participation to facilitate the adoption of the reference implementations and specifications by the healthcare industry, worldwide.

- *Privacy and Public Policy Work Group.*

  This group is intended to ensure that the Kantara Initiative contributes to better privacy outcomes for users, data custodians and other stakeholders, by defining privacy-related principles and good practice applicable to a broad range of prevalent technology platforms.

- *Consent and Information Sharing Work Group*

  The goal of this working group is to identify and document the use cases and scenarios that illustrate the various sub-sets of user driven information, the benefits therein, and to specify the policy and technology enablers that should be put in place to enable this information to flow.

  Project VRM and other related parties wish to build a framework around which a new type of personal information can be enabled to flow, and in doing so improve the relationship between demand and supply. The contention is that when individuals are forced to sign organization-centric privacy policies or terms of use then this places limitations on the information that will be shared. If such constraints were removed, and capabilities built on the side of the individual, then new, rich information will flow.

More specifically, the following initiatives are focused on improving security and privacy measures in health care systems:

OpenID Foundation

The OpenID Foundation promotes and protects the OpenID community and technologies. It is a non-profit international standardization organization of individuals and companies committed to allowing and encouraging OpenID technologies. The foundation was formed

in June 2007 and serves as a public trust organization representing the open community of developers, vendors, and users. This entails managing intellectual property and brand marks as well as fostering viral growth and global participation in the proliferation of OpenID.

The most important group that is working on privacy for eHealth is the *HEART Working Group*. It intends to harmonize and develop a set of privacy and security specifications that enable an individual to control the authorization of access to RESTful health-related data sharing APIs, and to facilitate the development of interoperable implementations of these specifications by others. Recently, in February 2016, the OpenID Foundation members have approved of the following specifications as OpenID Implementer's Drafts: *"Health Relationship Trust Profile for OAuth 2.0"* [165], *"Health Relationship Trust Profile for OpenID Connect 1.0"* [166] and *"Health Relationship Trust Profile for User Managed Access 1.0"* [167]. The main goals of these specifications are to increase baseline security, provide greater interoperability, and structure deployments in a manner specifically applicable to (but not limited to) the healthcare domain.

International Organization for Standardization TC 215 Health Informatics

TC 215 works on the standardization in the field of information for health and Health Information and Communications Technology to achieve compatibility and interoperability between independent systems. Besides, it aims at ensuring compatibility of data for comparative statistical purposes and reducing duplication of efforts.

Today it has members from 31 countries and 162 ISO standards have been published under the direct responsibility of ISO/TC 215. This technical committee is divided in eight working groups, being the most relevant in our topic area the *ISO/TC 215/WG 4: Security, Safety and Privacy*. Its main goals are establishing guidelines for security management in healthcare and defining standards for technical and management measures in order to: 1) enhance the confidentiality, availability and integrity of health information; 2) prevent health information systems from adversely affecting patient's security and privacy; and 3) ensure the accountability of users of health information systems.

Chapter 3

Architecture Proposal

Historically, privacy was almost implicit, because it was hard to find and gather information. But in the digital world, whether it's digital cameras or satellites or just what you click on, we need to have more explicit rules - not just for governments but for private companies.

Bill Gates, 2013

Contents

3.1 Chapter Overview

This chapter provides a global view of the architecture proposed to address the privacy challenges presented in Chapter 1. The description starts in Section 3.2 with a brief introduction to general architectural model that is common to federated identity management infrastructures as references to introduce our extensions. Based on this basic model, a requirement analysis is performed. Next, Section 3.3 introduces the contributions made in this thesis to extend the functionality of the basic architecture satisfying the stated requirements. The new components and the extended functionalities are explained and accompanied by a flowchart that illustrates how each architectural component interacts with the rest. Finally, Section 3.4 ends with the main conclusions extracted from the architecture proposal.

3.2 Design Principles

3.2.1 Requirements Analysis

With the aim to define the IdM architecture, we first analyze the current architecture of Identity Management systems, since it will be the starting point to add the new functionalities. In this respect, there is not a single IdM architecture definition. The only thing that can be found in the literature are standards, specifications, as well as guidelines and best practices for implementation. But every organizations deploying IdM solutions define their own architectures. In the specific case of SAML, which is the most important and complete framework, there are two big implementations: Shibboleth[1] in education environments and deployments in industrial/governmental environments. By extracting the common features of them, we have elaborated a component diagram for a generic identity federation architecture, depicted in Figure 3.1. Note that, we use generic names for the components, however, these modules may be denominated differently across implementations despite their functionality is the same. As it can be seen in the picture, the architectural components in the providers are:

[1] http://www.internet2.edu/products-services/trust-identity/shibboleth/

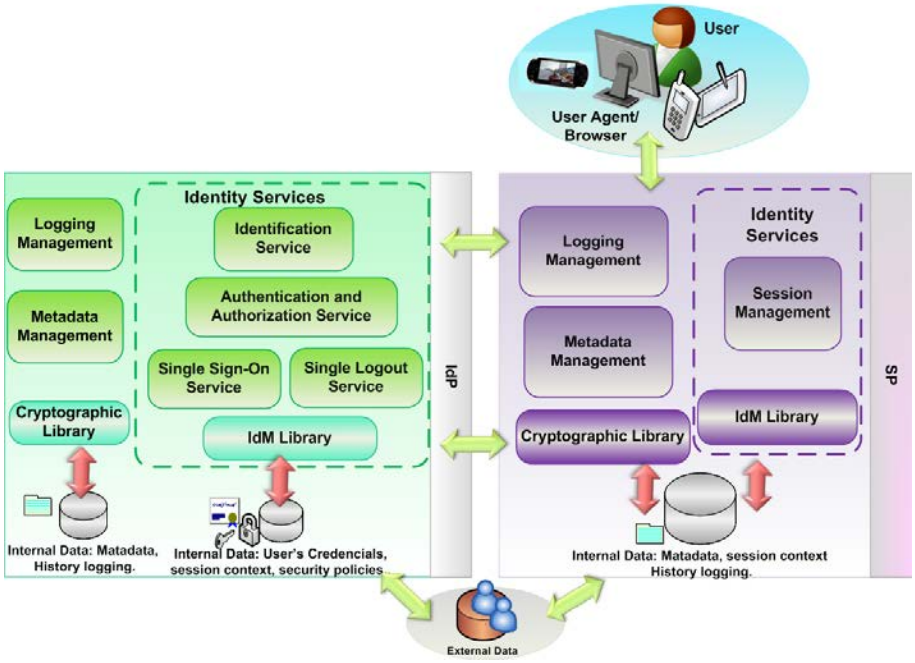

Figure 3.1: Generic Architecture for Federated Identity Management. Common features of current IdM implementations.

- **Identity Services module**: This module embraces services provided by the identity framework for identification, authentication, authorization, Single-Sign On, Single LogOut and session management. Both providers, the SP and IdP, implement this module in order to enable the communication and exchange of user identity data between them.

- **Metadata Management module**: This module is responsible for maintaining a circle of trust, which contains information related to certificate lists, static list of trusted entities, information about profiles supported by an entity, etc. The information contained in the metadata can be associated with either a single entity or a group of entities and they are intended to facilitate the deployment of IdM systems.

- **IdM Library module**: It provides the core functionality of the SP and IdP and offers the necessary support for the main elements defined by current IdM specifications (e.g., assertions, protocols, bindings, profiles, etc.). For this purpose, it interacts with the *Metadata Management* module to store the metadata of providers with either the IdP or the SP has established a trust relationship ; and the *Session*

Management module to handle users' identifiers and session data. These data are typically stored in local SP repositories. In the case of the IdP, it can be observed that the *IdM Library* module also relies on *Identification* and *Authentication Services* to authenticate users who access to the identity-based services through local or external identity repositories.

- **Cryptographic Library module**: This module consists of a set of cryptographic libraries, which offer several features: certificate generation for providers and communication with the *IdM Library* module for the tasks related to encryption/decryption, signing/validation of messages exchanged between the IdP and the SPs. It also contributes to the metadata generation for these entities.

- **Logging Management module**: This module is usually implemented by providers with the aim to monitor user and service activities. The registries are used by the identity services, but an interface may be also provided for auditors (external parties).

- **Data repositories**: They contain information used by the rest of the components, i.e., user's credentials, logs, session data, messages, metadata documents, security policies, etc. These data can be stored either locally by providers or in external repositories.

Finally, in Figure 3.1, can also be observed that the user interaction with the IdPs and SPs is performed through a user agent, which is typically a web browser.

Functional Requirements

To complement the architecture described above, our main goal in this thesis is to define and validate (through proof-of-concepts and simulations results) some extensions to provide a more robust privacy and identity management toolkit. As discussed in the introductory Chapter in section 1.1, in the Chapter 2 in section 2.5 and it will be analyzed in more detail in Chapters 4 and 6, the following aspects need to be tackled:

1. **Current identity architectures lack of a suitable user revocation consent mechanism, which encompasses scenarios in which the user cannot always be aware to grant or revoke her consent expressly**. IdM systems are the ideal target to deploy the privacy mechanisms, since they handle and orchestrate every

users' attribute exchange. However, current IdM specifications are not ready to cope with some aspects of privacy. Particularly, the user consent revocation has not yet been adequately addressed.

It is required that IdM systems protect user's privacy and allow authorized entities (including humans) to access users's information conveniently in order to avoid attacks, frauds or identity misuse. Furthermore, there are very sensitive environments, such as health care scenarios, where the user is not always be aware to give her consent (e.g., the user suffers an accident and loses consciousness.), so it is necessary an identity management model which includes consent delegation mechanisms taking into account the different events that happen in the system.

The proposed mechanism will offer users adequate interfaces and components to configure and modify her security and privacy policies according to her privacy preferences, levels of criticality of the event, etc.

2. **Need for a time-independent revocation mechanism**. A traditional way to mitigate revocation challenges is to limit the lifetime of security tokens by reducing the time-to-live to a magnitude of seconds or to minimize minutes of the vulnerability window in cases of compromising the token. The downside to this is that the systems' usability probably will be reduced since the user will have to re-authenticate to obtain a new valid security token. On the opposite side, users will obtain a better experience when token expiration is set to hours, days or months, while the risk of compromising information and identity theft increases. In order to address this problem, the proposed architecture will incorporate components to provide a flexible event-based user consent-revocation mechanism, which enables to substitute time constraints and explicit revocation by managing authorization rights in accordance with events.

3. **Inclusion of new roles**. As explained above, in those environments that handle sensitive data, it is required to contemplate scenarios in which user is not online to grant her consent. For instance, a user could delegate another user to make a payment to buy an object when she was offline, could not be located and the price of the object was below a certain threshold.

   Besides the roles defined in the current IdM specifications, it is necessary to consider

new roles in order to offer revocation services based on events and support such interactions. Thus, we identify the following roles, which can be adopted by different entities (user, SP or IdP):

- ***Delegator***: It is the role adopted by an entity when it gives its permissions to access to the attributes or services provided by others.
- ***Delegatee***: It is role played by an entity when it receives permissions to access to the attributes or services from others by a *Delegator*.

In this way, we achieve that the different actors of IdM scenario have access to the user's attributes that are allowed the time strictly necessary, since the revocation of attributes and privileges will take place implicitly. It is worth noting that entities can act as both roles of *Delegator* or *Delegatee* in each time.

4. **Current IdM specifications do not sufficiently support users regulation the release and use of their own identity information**. Nowadays, users are expected to deal responsibly with the privacy agreements exposed by the participants in a digital transaction, but the complexity of some agreements and the increasing number of participants overwhelm users. In most cases, users accept Service Level Agreements (SLAs) without reading service conditions or being aware of how they their data will be shared, distributed or used.

   So, it is necessary to include a component which allows to offer a richer management and vision of user's profiles. It also will empower the user's role enabling her to combine multiple sources of identity in a single credential to present to the providers and to disclose her personal information selectively, while minimizing direct interactions between SPs and IdPs. Furthermore, we will address the four **fundamental privacy principles** reviewed in section 2.2: anonymity, pseudonymity, unlinkability, unobservability, selective disclosure and revoking consent.

5. **Need for an enhanced awareness over users' online identity use**. The proposed IdM architecture with privacy extensions should minimize privacy risks and offer users greater awareness of the use of their digital identities by introducing monitoring mechanisms that enable users to balance security, privacy and usability according to their needs. In addition, the architecture proposal should allow user interaction in the system to support auditing. To accomplish this, the solution will

contemplate a module, which will introduce monitoring tools and an audit service focused on data sharing.

To cover the above requirements, we present a solution for federated IdM infrastructures based on the inclusion of a *Privacy Engine*, which includes several modules and sub-modules and has distinct characteristics depending on the different roles of the entity where it is located. The detailed explanation of the proposed architecture is developed in the section 3.3.

Architectural Strategies

This section describes the design decisions and strategies that affect the overall organization of the architecture and its higher-level structures.

- **A Layered Approach**. The proposed architecture is defined by following a layered approach, where all *Privacy Engine* and identity management related functionalities are grouped together in an *IdM Layer*. It provides simple interfaces towards the application and services layer, thereby abstracting the internal design and structure. So, the focus of our privacy-awareness IdM architecture is to define the *IdM Layer*, paying special attention to the *Privacy Engine*, and its interfaces to the upper layers (e.g., Application). With this respect, it does not analyze the internals of the other layers, but it only concentrates on defining the logic and interfaces necessary for those layers to use the functionality of the *Privacy Engine*.

  Equally important in the architecture is the specification of the data artifacts exchanged between the implicated entities, in such a way that the underlying differences of concrete *Privacy Engine* components will be abstracted away through the definition of formats that can convey information independently from the mechanism-specific cryptographic data.

  Therefore, we will define all technology-agnostic components of the *Privacy Engine*. That is, the issuance and presentation, disclosure or revocation of user's attributes and credentials are interactive processes, potentially involving multiple exchanges of messages. Chapters 4 and 6 define the models, flow operations, etc. and specify the data exchanged during the issuance, presentation, revocation of user's credential and attributes; and the generation and disclosure of enhanced user's profiles. There are several existing protocols, in which these privacy-awareness credentials, attributes

and profiles can be embedded, such as SAML, or new ones could be defined in the future.

- **Building Privacy-enabled applications**. The implementation and deployment of proposed architecture will be embedded into example applications showing how to integrate the proposed extensions and components into different use cases for health care, social networks, cloud computing or consumer electronics scenarios. In this way, application developers could integrate the *Privacy Engine* modules in their applications, without having to know how its layers are internally structured.

Finally, it is important to note that we contribute on a specific part of the architecture[2], but without losing the global perspective.

3.3 Architecture Description

The proposed IdM infrastructure incorporates the functionality to allow Identity Providers, Service Providers, and enhanced clients to share common knowledge. The enhanced client is a software element for non-HTTP uses cases, which enables to minimize direct interactions between SPs and IdPs, and provides full control to users over their identities, thereby improving mainly privacy. The proposed architecture for the elements of the IdM system is represented in Figure 3.2. Such image shows the logic blocks, in a layered model, as well as the relationship between them.

At the top of the architecture, we have the *Application and Services Layer*. It contains applications and services offered by providers (SPs or IdPs). Note that, this layer is also located on the the enhanced clients, containing client applications. Next, in the underlying level we can see the *IdM Layer*, which offers the basic functionality of each role defined in the IdM specifications. In addition, such basic functionality is extended by adding the *Privacy Engine* component.

Finally, the *Data Layer* accommodates the information required for the upper layers to operate, that is: security and privacy policies, privacy preferences, history logs, user's data

[2]For reasons of simplicity, we do not include the architectural components that cover layers for establishment of dynamic trust relationships or location, since these topics are out of the scope of the research presented in this thesis. Moreover, the existence of an *Event Engine*, which follows a notification model including well know workflows triggered by well known trusted entities is assumed and the *Privacy Engine* relies on it.

and personal information (e.g., structure data, documents, etc.).

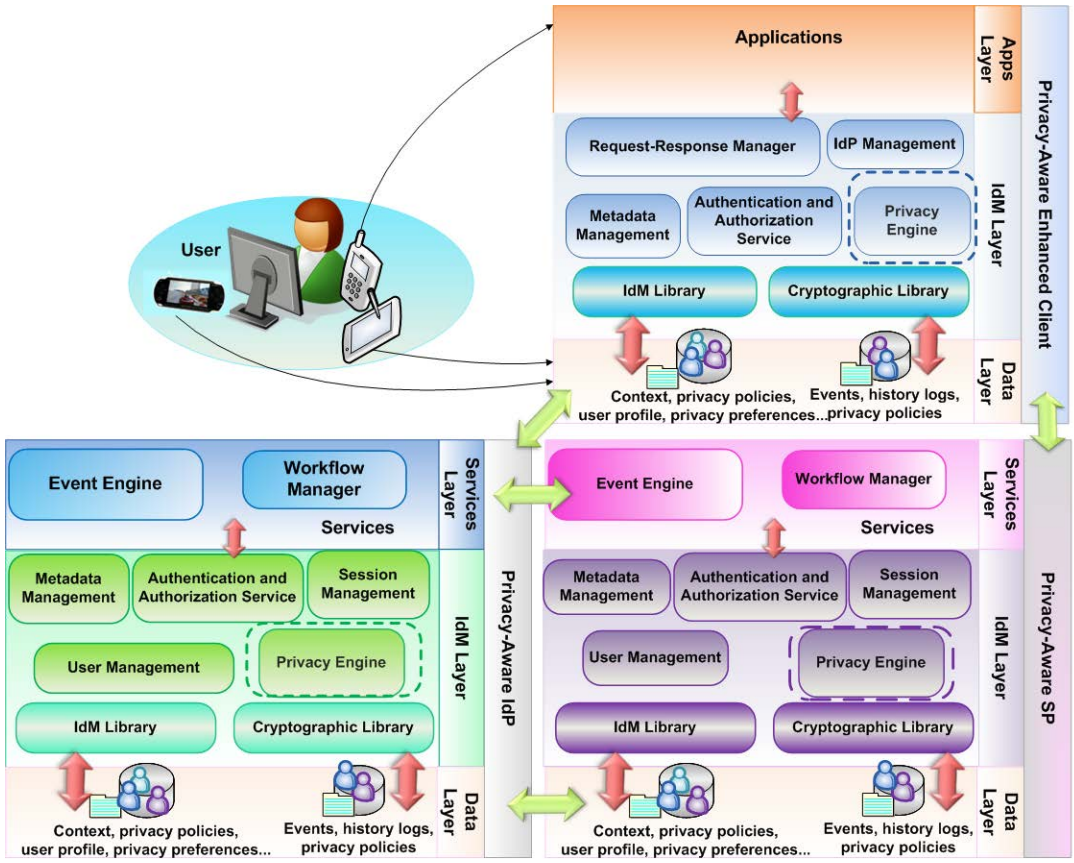

Figure 3.2: Enhanced-Privacy IdM architecture. This shows the building blocks in the different roles in an identity management system, such as providers (SP, IdP), users and enhanced-clients.

3.3.1 Architecture Overview for Enhanced-Privacy IdM

Application and Services Layer

As will be detailed in the Use Case 2 in section 3.3.3, in very complex or heterogeneous scenarios, applications typically rely on mechanisms for event management and workflow management tools, in order to handle and carry out efficiently the different requests that arrive at the system and involve the interaction of multiple individuals and resources (e.g., Alice has an accident and as a result, she loses consciousness, a citizen notifies the emergency services that a person has suffered an accident, the emergency service sends an

ambulance, etc.).

So, with the aim to manage internal and external events that happen in the system, the different tasks that are triggered involving access or exchange of sensitive information; the *Privacy Engine* will rely on the following components to implement the logic for revocation consent event-driven:

- **Event Engine**: This component is in charge of handling and collecting relevant events received and sent to the *Workflow Manager*. Essentially, it communicates with the sub-modules *Event Broker* to orchestrates the communications and operations required by *Workflow Manager*, to broadcast the event sent by producers to the consumers that are interested in those types of events; and the *Event Processing* to analyze the events associated with each transaction. The *Event Engine* supports a producer/consumer model: components need to register as producers or notifiers, to share events with the *Event Engine* by declaring who they are and the types of events that they are going to produce. Likewise, event consumers or subscribers need to register with the *Event Engine* by declaring who they are and which types of events they wish to receive. In the proposed architecture, the *Privacy Engine* components called *Privacy-Aware User Profile Handler* is a notifier; *Audit Service* and the *Action Monitoring* components are consumers (See section 3.3.2). It should be noted that SPs and IdPs can play the role of notifiers and subscribers, in some circumstances, either notifying different events or consuming them. Furthermore, each entity can be subscribed to multiple types of events and each event type can be attended by several notifiers.

- **Workflow Manager** is responsible for the execution of workflows. It instantiates workflows from workflow definitions, and decides which activities of the workflow have to be executed next. To control and monitor workflow execution and to handle failure situations, the *Workflow Manager* maintains so-called workflow control data. For instance, these workflow control data describe the actual execution stage of a workflow and its activities, or record the execution chronology. In particular, workflow control data cannot be manipulated by applications or users. During workflow execution, this module invokes applications if they have been assigned to the activity to be executed next. For activities which have to be executed by users, a worklist handler maintains the worklists stating which activities are to be executed by which

staff members, and propagates this information to the respective user interfaces.

IdM Layer

The logic modules that make up the *IdM Layer* contribute to manage sessions, issue and processing of requests and responses of authentication and authorization, etc. In this sense, the proposed architecture supports a flexible privacy-aware user profile management and an event-driven revocation mechanism, which enable user to have more control over her identities and her disclosed information in different transactions. On the other hand, this layer has cryptographic modules based on an underlying PKI for secure communications and exchange of needed metadata in the dialogue between the SP and the IdP about the user. Also note that, we can think of metadata lists being equivalent to trust lists, since a provider considers trustworthy the entities whose metadata is stored in its repository. The core modules of the IdM layer, which are common both to providers and enhanced clients, are detailed as follow:

- **Metadata Management** Both SPs and IdPs implement a configuration component, over which the services rely. This component is in charge of accessing local data stores to determine if a provider involved in a current identity-related transaction is trusted. This decision is made basically by consulting the local data stores to check if the entity is contained in a list of trusted entities.

- **Authentication and Authorization Service** functionality depends on its location. So, it receives and processes the authentication request messages from either the SP or the enhanced client, regarding to the enhanced client or IdP, respectively. In the SP, it issues such authentication and authorization request. The modules in each entity interact to verify the user requesting a service is really who he claims. For this purpose, it supports multiple authentication mechanisms including PKI, username/password, etc. In regard to the authorization process, the security assertions and the attributes exchanged convey authentication decisions, profiles and attributes to services providers allowing them to decide what services or resources the user can access. For that, this module issues (IdP) or verifies (SP), and manages authentication assertions and attribute statements. The aim is to facilitate authentication and user management to users and services improving user experience while reducing complexity and management costs.

- **Privacy Engine** is responsible for managing user identifiers (e.g., pseudonymous), handling the revelation and revoking of enriched user's profiles and attributes; and auditing how user data is being accessed without compromising user's identity. To accomplish this, the *Privacy Engine* includes an audit component for events, attribute activation, and access control decisions. The fields that are logged by monitoring tools and verified by audit tools show the auditor what information about the user is being accessed without divulging the actual information. In this way, this module provides multiple and partial identities, which allows users to access services and share digital content without necessarily revealing their name and true identity to everyone. The use of different pseudonyms and privacy-awareness profiles enables to support differing ranges of identification and authentication strengths. Finally, note that, the audit component 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*.

These modules are supported by basic libraries such as IdM and cryptographic, which implement IdM functionalities and cryptographic algorithms and protocols, respectively. In the user's side, these libraries implement the minimal functionality, taking into account limited devices. Thus, the enhanced clients incorporate "lite" library versions.

Regarding the provider's side, we can also find other two additional modules whose functionality is specific to service provision:

- **Session Management** is responsible for managing user identifiers, as well as the session data of those users accessing SPs or IdPs services. The SP together with IdP determine when user's session is active. The SP creates session identifiers for every user once user has already been authenticated and registered in the service. Such session identifiers are linked to users' profile. The *Session Management* module also may check several user profiles and select the most appropriate content for a specific service (e.g., video on demand). Thus, the *Session Management* module communicates with other modules handling authentication and attribute exchange, and with the *User Management* and the *Privacy Engine* modules to request the user's profile, related to the different services, and matches it with the user's profile policy, the device profile and any other enforced IdP policies.

- **User Management** is in charge of dealing with credential storage, management users' profiles according to their preferences and policy enforcement. Regarding credential management, the user can store her credentials (e.g., username/passwords, digital certificates, etc.) that are required by applications that can be accessed for instance from TV services, social networks, payment or health care services, user-centric private clouds, which offer a unified perspective on the user's activities, etc. Moreover, this module also communicates with the *Privacy Engine* to build and search selective privacy-enhanced user's profiles and attributes. On the other hand, the *User Management* module interacts with the enhanced client in order to determinate which IdP is appropriate according to the service requested and user's preferences. This last aspect, allows to the proposed architecture to act on behalf of the user and perform authentication in the different applications providing a seamless, personalized and improved user experience.

Besides, the enhanced client incorporates other two additional modules whose functionalities are:

- **IdP Management** is responsible for determining the most adequate identity provider depending on the requested service, the user's preferences related to privacy and security (e.g., type of credential) and the context. For instance, in some contexts the user may not want to reveal any personal information, whereas in other contexts she may wish for partial or full disclosure of identity. To achieve these tasks, this module collaborates with the *Request-Response Manager*, the *Metadata Management* (to configure trust relationships with the IdPs) and the *Privacy Engine* (See the Use Case 1 explained in section 3.3.3).

- **Request-Response Manager** receives authentication, authorization or attribute requests from the applications. It is the interface between *Authentication and Authorization Service* and applications. These requests can be originated by authentication request statements from the SP. After carrying out the processing and verification of requests, this module issues or redirects responses to the applications for being resent to the SP. Note that, the *Request-Response Manager* is able to use a reverse SOAP (PAOS) binding [23] to manage the requests and responses of authentication such as is specified.

3.3.2 Privacy Engine: Components and Relationships

Once an overview of the proposed architecture and its layers has been provided, this section goes into the details of each individual modules that constitute the *Privacy Engine* as well as their contribution to meet the requirements identified in section 3.2.1 and the relationship between them.

Privacy is preserved thanks to the enhanced clients and the *Privacy Engine*, which has been incorporated in each entity. As explained at the beginning of section 3.3.1 and it will be shown through the use cases described in section 3.3.3, the enhanced client allows to give users more control over their personal information, identities, as well as control over authentication and attribute exchange processes eliminating the direct communication between the SP and the IdP.

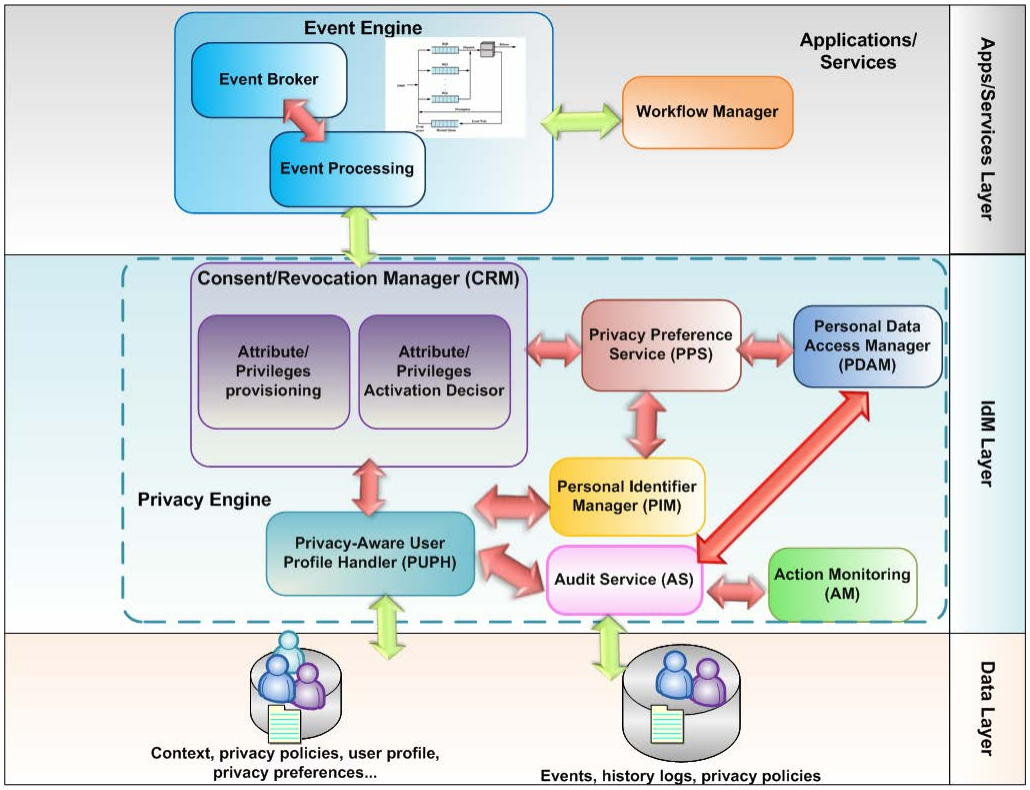

Figure 3.3: Detailed view of the Privacy Engine components.

The *Privacy Engine* component carries out an appropriate management of user identifiers, profiles, revocation of attributes/permissions according to user's preferences and events sent by the *Event Engine*, as well as to monitor how user data is being accessed by

SPs or IdPs without compromising user's identity. Figure 3.3 sketches out the different *Privacy Engine* components, whereas Figure 3.4 illustrates the different functionalities of the *Privacy Engine* depending on the entity in which it is located.

The IdP and the enhanced client have in common the following modules, which improve the active role of the user and also consider scenarios in which the user may be offline:

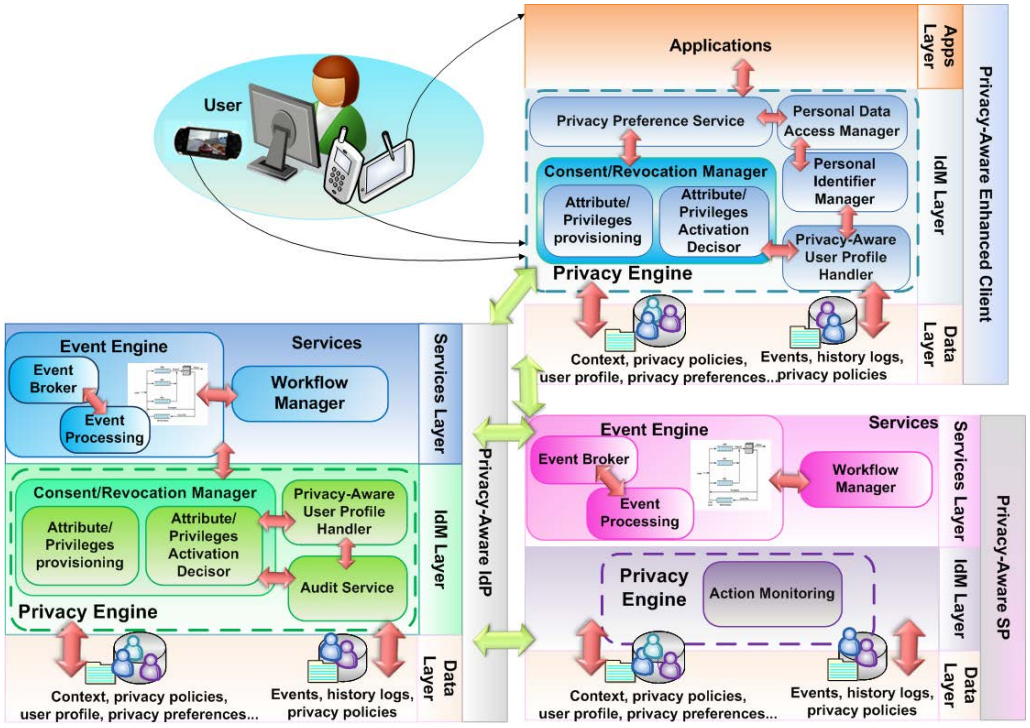

Figure 3.4: Detailed view of the Privacy Engine components from the perspective of the entities. The Privacy Engine has different functionalities depending on the different roles of the entity where it is placed; for instance, management of user identifiers, profiles and privacy preferences or audit and monitoring functions.

- **Consent/Revocation Manager** is responsible for handling and revoking attributes and privileges (following the principle of *minimal disclosure*) depending on the different event filters and the defined privacy policies. This component interacts with various *Privacy Engine* components in order to cover the requirements (1), (2) and (3) identified in section 3.2.1, including: the *Privacy Preferences Service*, the *Privacy-Aware Profile Handler* and the *Audit Service* (See the use cases illustrated in section 3.3.3). To this end, this component has been divided into two sub-modules to separate the responsibilities of the following main functions:

(a) **Attributes and Privileges Management**. The *Attribute/Privileges Provisioning* sub-module enables to the user (*Delegator*) to establish different terms or conditions for access to her personal information (e.g., different parts of her medical history) when specific events happen. Likewise, this component is responsible for provisioning and publishing the required credentials, attributes/permissions in order to the consumers subscribed to those events can receive relevant updates.

(b) **Access revoking**. The *Attribute/Privileges Activation Decisor* sub-module analyzes the different elements which compose each event sent by the *Event Engine* (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 to decide whether the *Delegatee* can access to the requested attributes. The policy includes the set of consent directives and other privacy conditions (i.e. object filtering, user, role, purpose) that constrain enforcement. In addition, this component updates the corresponding privileges and actives new ones, if necessary, depending on the information forwarded by the *Event Engine*.

- **Privacy-Aware Profile Handler** is in charge of managing and storing user profiles to fulfill the requirement (4) specified in section 3.2.1. As it will be detailed in the Use Case 3 in section 3.3.3, this module facilitates that users can instruct the applications on how different identity attributes should be coalesced when a given SP requires information. Thus, in some scenarios where the user interacts with several service providers that need to access to user data stored in different identity repositories, those profiles would be eventually merged for each SP with other profiles from SPs, user devices or IdPs by either a client tool or a trusted identity provider (if the user is offline). Then, the user's profile, including only the strictly required information, would be sent to each service provider. To accomplish this, this component employs a flexible data model by allowing the combination of certified and self issued attributes from different sources with different pseudonyms.

On the other hand, the module incorporated in the enhanced client has the following components, that contribute to achieve the requisites related to enhanced awareness over users' identity use and disclosure (See section 3.2.1). To facilitate the understanding of the

interactions between the *Privacy Engine* components in the client site with the rest of the architecture elements, it is recommended to look at the use cases described in section 3.3.3:

- **Privacy Preferences Service** provides an interface which communicates with the *Consent/Revocation Manager* for configuring user's preferences about the handling of personal data and specifying options for the use and release of sensitive information. The set of variables that define the customizing privacy preferences include for instance, the user identifier, information type, the requestor, the requested operation (i.e. query, create, modify), as well as a set of pre-defined policies by the IdP.

  On the other hand, this module also includes the tracking preference expression (DNT) feature that allows user to express their personal preferences regarding cross-site tracking to each service or application. Whereas the Do-Not-Track approaches that recently have been incorporated in some commercial browsers, only enable or disable tracking characteristic and this is applied to all services accessed by user without including any preference set for configuration. Even if we use anonymous profiles and private browsing options we are tracked [168]. So, we believe that, it is necessary to maintain a trade-off between degree of tracking and user's privacy to obtain an adequate personalization degree in the different services.

  Nowadays, we can find a lot of services, for example, personalized catch up TV services or location prediction applications, which require access to certain attributes related to habits, preferences and user needs to properly adapt to user behavior, to predict future patterns in his preferences and to offer a really customized user experience. Therefore, the improvement of our proposal over current solutions comes from offering the user the option of selecting and detailing which attributes may be traced depending on the user's trust placed in the service, the sensitivity of a specific attribute, desired personalization degree, etc.

- **Personal Data Access Manager** allows user to check the accuracy of her personal information and visualize how his data is being used by both the SP and the IdP. For the latter purpose, this module receives notifications from the *Audit Service* located in the IdP; thereby allowing the user to obtain automatically updated information in a seamless and dynamic manner.

- **Personal Identifier Manager** is responsible for managing different kind of identi-

fiers such as pseudonyms (transient or permanent identifiers), social networks identifiers, etc., in a flexible and personalized manner in order to enable user to choose between multiple identities when interacting with services. For this, the module interacts with the *IdP Management* component to obtain and associate user identifier in each IdP and conveys this information to the *Privacy-Aware Profile Handler*. Note that this module uses different pseudonyms for each SP in order to avoid different SPs belonging to the same federation to infer user behavior.

In regard to other tasks of the *Privacy Engine* in SPs and IdPs, as it will be shown in the use cases 2 and 3 in section 3.3.3, they are related to the functions of auditing (i.e. IdP *Audit Service*) and monitoring (i.e. SP *Action Monitoring*) of how each SP accesses user data without compromising user's identity. So auditing and fraud detection at SPs could be tackled. The IdP *Privacy Engine* module includes an *Audit Service* that focuses on data sharing, which captures any transaction or event where user data is requested, shared, created or modified from a service provider, including information such as the sender, receiver, target identity, as well as identifying the user attributes accessed and the purpose for which they were accessed. It must be noted that, the actual values of the attributes involved in each event are not logged in order to ensure that events are recorded in a consistent manner amongst all the SPs using the *Action Monitoring* module. The *Action Monitoring* module uses an XML-based event structure defined by us to log events to the *Audit Service*, which includes the following elements:

- **UserID** specifies an opaque identifier or pseudonym. It refers to the principal whose personal information is accessed.

- **SPName** specifies the entity name, which is accessing to user data. It identifies the service provider and is a unique identifier of each SP (EntityID) contained in its metadata.

- **AttributeName** is a compound field that contains the attribute names accessed by the SP. In this case, attribute names must be consistent across the federation.

- **Scope** indicates the scope in which user data are being used, as well as how many SPs are sharing or exchanging a specific user's attribute.

- **Purpose** specifies the purpose usage for attribute requested by a SP.

- **AccessTime** indicates the instant time in which an attribute was accessed by a provider. It is a timestamp of the event.
- **UserDelegateeID** specifies an opaque identifier or pseudonym. This field is optional, because it is only logged in delegation cases and refers to principal in whose behalf on user data is being accessed.

Summarizing, since the *Consent/Revocation Manager* and the *Privacy-Aware User Profile Handler* components include the functionality that constitutes the main contribution of the thesis, a separate chapter is dedicated to define each of these components. Thus, Chapter 4 and Chapter 5, develop and show validation results of an event driven hybrid IdM approach to be implemented as part of the architecture by providing a more flexible revocation model. Then, Chapter 6 and Chapter 7 explain the mathematical model to represent selective privacy-enhanced user profiles and the validation carried out, respectively.

Finally, in order to complete the general picture of the architecture, next section explains the general behavior of the architectural components through several use cases and operation flowcharts.

3.3.3 Use Cases and High-Level Interactions

In order to show the behavioral part of the architecture, this section illustrates how the different components interact with each other through several use cases, referencing which parts of the proposed architecture take part in the process. The aim is to provide just a conceptual understanding so low level details are not yet given, but they will be addressed in subsequent chapters.

- **Use Case 1: Privacy-enabled configuration**: This use case is about a user that configures her security and privacy preferences for access to her medical records. The *Privacy Engine* module in the user's site configures various internal components to ensure the fulfillment of these preferences.

  Figure 3.5 shows the involved steps and the affected *Privacy Engine* components.

  The following steps and components are involved:

  1. Alice provides her specification of related privacy preferences, via the interface offered by the *Privacy Preferences Service* component. This specification in-

cludes that a high critical level-based policy that would enable access to the participant entities in particular event to full patient's medical history (e.g., if Alice is unconscious). A low level-based policy would grant access to the parts of medical history that Alice specifies according her privacy preferences is also defined. Moreover, she selects as trusted IdP her city hospital and specifies the credentials to access to a social network where she will share her daily activity and food intake progress.

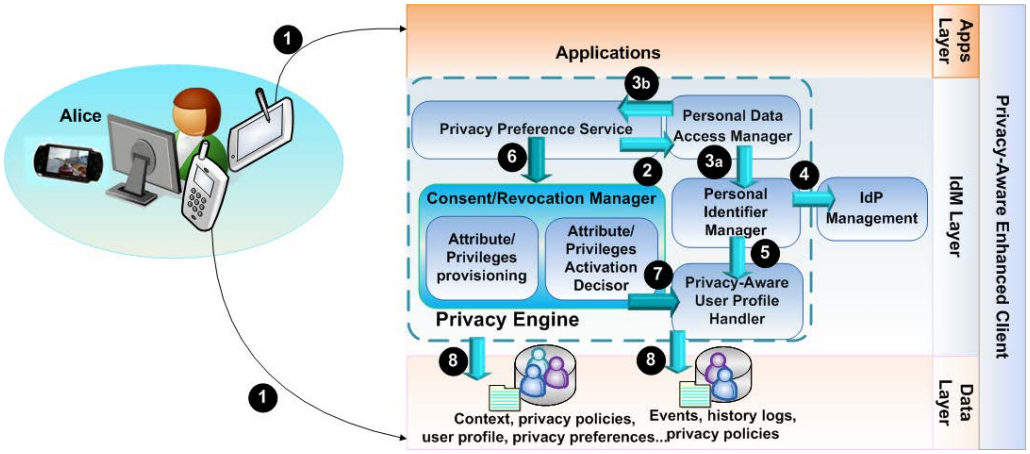

Figure 3.5: Privacy-enabled configuration use case: involved components and high-level interactions.

2. The *Privacy Preferences Service* sends the information provided by Alice in Step 1 to the *Personal Data Access Manager* in order to verify its accuracy.

3. If the verification result is successful, the *Personal Data Access Manager* forwards the information related to the chosen credentials and IdP to the *Personal Identifier Manager* (Step 3a). Otherwise, it returns an error response to the *Privacy Preferences Service* in the Step 3b.

4. In the Steps 4 and 5 the *Personal Identifier Manager* interacts with the *IdP Management* and the *Privacy-Aware Profile Handler* components, respectively, to request the storage of the associated privacy preferences, along with the additional metadata (references to personal data, selected providers etc.).

5. The *Privacy Preferences Service* also communicates with the *Consent/Revocation Manager* in Step 6 to manage the corresponding attributes and privileges according to privacy preferences established in Step 1.

6. The *Consent/Revocation Manager* asks for storing the information associated to the user's credential and attributes, as well as user's consent and privacy preferences to the *Privacy-Aware Profile Handler* in Step 7.

7. Finally, the *Privacy-Aware Profile Handler* calculates an enriched user profile and stores this information through the *Data Layer* in the Step 8.

- **Use Case 2: A user delegates access to others to certain parts of her personal data when she is not able to give her express consent**:

  This use case is about a user, Alice, who suffers an accident, that triggers several events. The emergency service (SP_1) requests access to Alice's medical records in order to send them to an ambulance company (SP_2), that needs access to her medical records (managed by IdP_1) to provide Alice the appropriate treatment.

  The enhanced-privacy IdM architecture, deployed within SP_1, SP_2 and IdP_1, handles the overall process of disclosing data and grant/revocation access to user's attributes, consistent with the Alice's privacy preferences and the credentials and attributes configured in the Use Case 1.

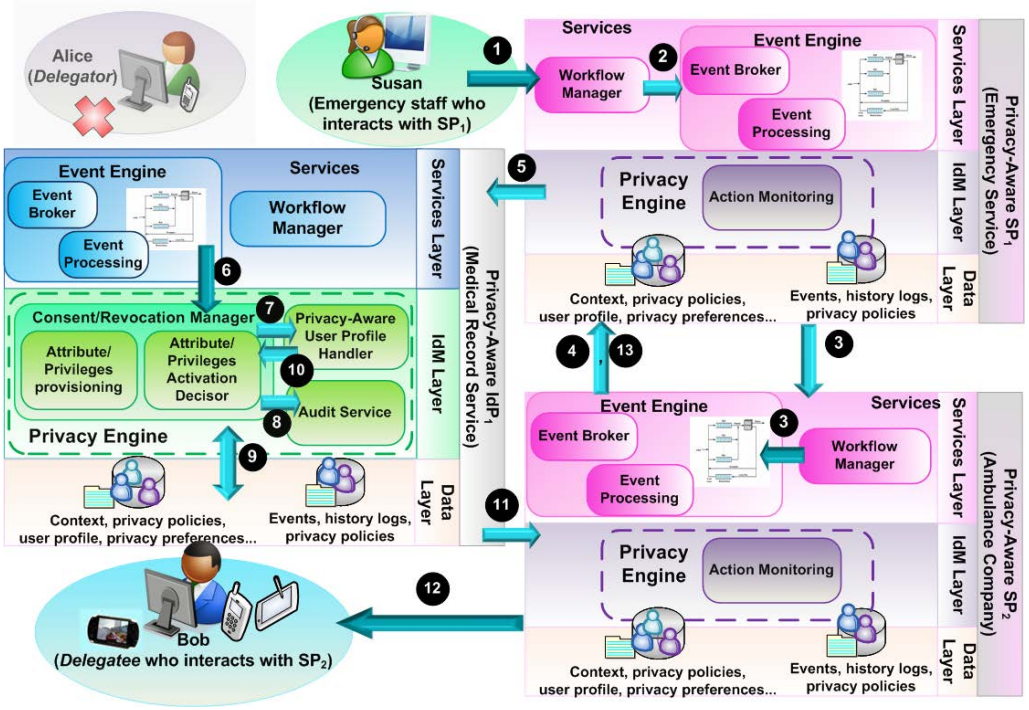

Figure 3.6: Flow of interactions, involved entities and *Privacy Engine* components for a use case in which the user is unconscious.

Figure 3.6 illustrates the involved steps and the affected *Privacy Engine* components along with the participant entities and their roles.

The following interactions and components are involved:

1. The emergency service (SP_1) receives a call which informs that Alice has suffered an accident. An employee of emergency service, Susan, completes a task form introducing the required information (e.g., the address where the ambulance company must go, etc.) by means of the worklist handler offered by the *Workflow Manager*.

2. In Step 2, the *Workflow Manager* conveys the task data as an event type to the *Event Engine* in order to it can be analyzed and sent it to the consumers subscribed to that event.

3. The *Event Engine* deployed within the SP_1 processes the received event and broadcasts it to the ambulance company (SP_2) in Step 3.

4. When the SP_2 arrives at the scene of the accident (in Step 4), it identifies the patient and notifies her identity to the SP_1 (See the interaction between the *Event Engines* of SP_1 and SP_2). At the same time, the SP_2 requests access to Alice's medical history in order to give her a calming to stabilize her.

5. To disclose the strictly parts of Alice's medical records, in Step 5 the SP_1 interacts with the medical record service (IdP_1) to check the credential containing attributes and permissions configured by Alice in the Use Case 1.

6. The occurred events are sent by the *Event Engine* deployed within IdP_1 to the *Consent Revocation Manager* in Step 6.

7. In Step 7, the *Consent Revocation Manager* examines the event received and requests to the *Privacy-Aware User Profile Handler* rigorously necessary attributes to be sent them to the SP_2.

8. The *Consent Revocation Manager* also communicates with the *Audit Service*, in Step 8, to record the request for access to parts of Alice's medical history requested by the SP_1 (e.g., Alice's `UserID`, Bob's `UserDelegateeID`, the attribute names to represent the part of Alice's medical history related to therapeutic precautions that includes allergies, the scope and purpose to indicate that these

data will be used by SP_1 and SP_2, etc.).

9. The *Privacy-Aware User Profile Handler* and the *Audit Service* interact with the *Data Layer* to obtain the requested data without revealing information related to non requested attributes, and to record the audit logs, respectively in Step 9.

10. In Step 10, the *Privacy-Aware User Profile Handler* sends back to the *Consent Revocation Manager* the user's attributes recovered in the previous step.

11. The IdP_1 forward a response to the SP_2, which includes the information about Alice's medical history in Step 11 and it is returned to Bob in Step 12.

12. The *Event Engine* of SP_2 notifies to the SP_1 that Alice is being taken to hospital A (SP_3, not represented in Figure 3.6) in Step 13.

When Alice arrives at the hospital A (not depited in Figure 3.6), SP_3 sends an event to the emergency service indicating that Alice has come to the hospital. Next, hospital staff requests access to her medical records and the SP_1 performs the following tasks. It communicates with the IdP_1 to request revoking access to Bob to Alice's medical records. If the response of IdP_1 is successful, SP_1 grants access to Alice's medical history to hospital A staff. In this case, similar iterations between the SP_3, the IdP_1 and their involved *Privacy Engine* modules to those reflected in steps 5, 6, 7, 8, 9, 10 and 11 will be repeated.

- **Use Case 3: Personal data disclosed to a third party**:

  This use case is about a user, Alice, who wants to watch a video her friends have shared with her through a social network (SP_1) she is member of. This video content is hosted and served through a by a third party (SP_2).

  Figure 3.7 displays the involved steps and the affected *Privacy Engine* components within the providers and the enhanced client. Note the active user role through the privacy-aware enhanced client, which allows to minimize direct interactions between IdPs and SPs.

  When Alice accesses the application through the social network she is warned the SP_2 requires the following attributes from her profile: feeds, friends, adultness and email (See Step 1). Alice, who does not want to share with the SP_2 more than the

necessary attributes, let the SP_1 to send the feeds and friends to the SP_2. However, Alice wants to keep her preferred email account (used in the social network) unknown to the SP_2, so she instructs the social network to get it from her mobile device (hereinafter IdP_1). In regard to adultness, SP_2 would need to get that attribute from an authority that vouches for it. For that reason, Alice instruct the SP_1 to get it from IdP_2. As it has been shown in the Use Case 1, the proposed IdM architecture empowers Alice to configure personalized privacy policies specifying where are those attributes taken from. The interactions between the involved components of the *Privacy Engine*, are represented in gray in Figure 3.7.

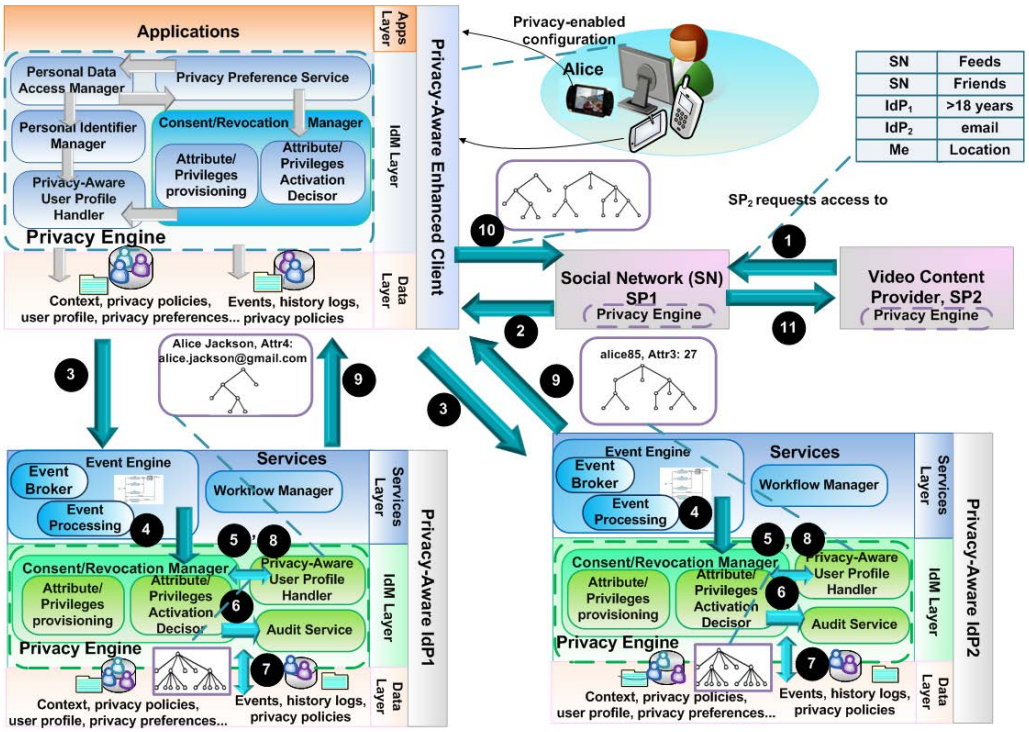

Figure 3.7: Flow of interactions, involved entities and *Privacy Engine* components for a personal data disclosed to a third party.

The following steps to disclose Alice's personal data to SP_2 are:

1. Hence, upon SP_2 request, in Step 2 the SP_1 demands from IdP_1 and IdP_2 the corresponding attributes (e.g., adultness, email, etc.) by means of the enhanced client.

2. The enhanced client forwards this request to IdP_1 and IdP_2 in order to let SP_2 to verify the aforementioned attributes in Step 3.

3. In steps 4, 5, 6, 7, 8 and 9 the construction of the user's profile from various sources of information (IdP_1 and IdP_2) when the user is aware is outlined. The interactions carried out between the *Event Engine*, the *Privacy-Aware User Profile Handler*, the *Consent-Revocation Manager* components and the *Data Layer*, are analogous to those described in the Use Case 2.

   Thus, IdP_1 and IdP_2 give back a sub-tree to the enhanced client in Step 9, which includes a verification path and a root node signed by either IdP_1 or IdP_2, respectively. Moreover, it is worth mentioning here that, the returned sub-tree may be either simple (enabling to verify a single attribute) or complex (allowing to check multiple user's attributes simultaneously).

4. The enhanced client comprises a single tree that contains the user's profile from the information received from the identity providers and sends it to the social network in Step 10.

5. The social network transmits the enriched user's profile to the SP_2 in Step 11.

Although it not is shown in Figure 3.7, other sensitive attributes, such as user's location, have to be directly requested to user devices by the SP_1.

Finally, after presenting this general view of the architecture, we will develop the functionality and low level details of the main *Privacy Engine* modules in the next chapters.

3.4 Conclusions

It is clear that current IdM architectures are limited to provide appropriate tools for effective user revocation consent, which encompass situations where the user is aware to grant or revoke her consent expressly without compromising her privacy. Furthermore, issues related to how users can regulate the use and disclosure of own identity information are not enough addressed by existing IdM specifications.

In this chapter we proposed an extended architecture to fill these gaps. The architecture is composed of a set of logical modules that separate and encapsulate the functionalities required to achieve on the one hand a flexible event-based revocation mechanism and on the other hand, an efficient privacy-enhanced user profile management approach, that

guarantees selective identity information disclosure. The pillars of the architecture are the *Consent/Revocation Manager* and the *Privacy-Aware User Profile Handler* modules of the *Privacy Engine* component, which constitute the main contribution of the thesis. The mathematical models implemented by each part are thus developed in the following chapters.

In conclusion, the extension of the architecture satisfies the intended goals, since it makes possible to cope with both online and offline scenarios, while preserving user's privacy and allowing her to better control over her online identities.

Chapter 4

Revocation Consent Proposal: An Event Driven Hybrid IdM Approach

Here in your mind you have complete privacy. Here there's no difference between what is and what could be.

Chuck Palahniuk, 2009

Contents

4.1 Chapter Overview

This chapter describes the proposed event-based mechanism empowering a new concept, the *sleepyhead* credentials, which allows to substitute time constraints and explicit revocation by activating and deactivating authorization rights according to events. Our approach is to integrate this concept in IdM systems in a blend of the federated and user-centric model supporting delegation, which can be an interesting alternative for scenarios where revocation of consent and user privacy are critical. Moreover, in the case of federated models, they bring up privacy concerns because medical records may be available to every entity within the *circle of trust*, even whether there is no emergency. Section 4.2 highlights the importance of revoking consent and provides a comparative analysis of the privacy support in identity management systems. Section 4.3 explains contributions made in this thesis to enhance privacy in health care scenarios. A mathematical model, which describes the event-driven system behavior, is also illustrated. Finally, section 4.4 summarizes the presented work and presents the main conclusions.

4.2 Understanding the Problem of Revocation Consent

Credential-based authorization offers interesting advantages for ubiquitous scenarios involving limited devices such as sensors, or personal mobile equipment: it offers a computational cost more reduced than its competitors for issuing, storing, and verification; and it naturally supports rights delegation. The main drawback is the revocation of rights. Revocation requires handling potentially large revocation lists, or using protocols to check the revocation status. Moreover, current identity management technologies are not ready to cope with user consent revocation properly.

This is a relevant issue in regard to privacy-enhancing mechanisms, especially in sensitive scenarios when sensitive data and profiles are shared. In a health care scenario, the system must protect user's privacy and allow authorized entities (including humans) to access medical records conveniently. Furthermore, in these scenarios the user is not always be aware to give her consent, so it is necessary a hybrid identity management model which includes consent delegation mechanisms. Besides, privileges permitting access to user attributes should be revoked in an effective way.

To achieve these goals, we will have to provide solutions that allow to perform some required tasks as discussed in the remainder of this section.

4.2.1 The need for an appropriate revocation in current IdM frameworks

Nowadays, there are several federated and user-centric identity management frameworks, (already detailed in Chapter 2), but they have not addressed this privacy rule. The privacy support of the current federated and user-centric identity management technologies is analyzed below. 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 behavior.

Regarding Liberty, as SAML, it 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 behavior. In order to overcome these problems, a set of rules and recommendations are proposed in [169]. WS-Federation, in its turn, contemplates privacy and pseudonym services by defining extensions to the WS-Trust Request Security Token and Response messages in order to specify how privacy statements can be obtained using mechanisms defined in HTTP, HTTPS, WS-Policy, etc., as well as how pseudonyms should be mapped.

In the case of the user-centric technologies, U-Prove and InfoCards handle privacy through pseudonyms and privacy languages to express privacy policies. In addition, U-Prove and InfoCards deal with unlikability and unobservability issues by defining a message flow that

eliminates direct communication between the IdP and the SP. U-Prove enables selective disclosure of user's claims through a pre-signed token that can be used even if the IdP is offline. InfoCards, in its turn, 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 visit a given principal. The SAML Enhanced Client Proxy profile is similar, but currently it only has unlikability. 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.

On the other hand, U-Prove enables the use of services with minimum disclosure of personal information. OpenID Connect identifies a set of personal attributes that can be exchanged between identity providers and the applications that use them, and includes an approval step so that users can consent (or deny) the sharing of this information.

Furthermore, this specification provides an extension called PAPE (Provider Authentication Policy Extension) [170], 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 Provider's response.

Table 4.1 summarizes the main privacy features grouped by IdM models. The technologies that have been analyzed handle privacy by means of pseudonyms which can be transient or permanent. Moreover, it must be noted that, InfoCards, U-Prove and SAML ECP profile address better the principle of minimal disclosure. Current identity frameworks support partial anonymity, since authorities, as the IdP, provides obfuscated identifiers. We want to stress the importance of privacy policies, since they are the basic means that allow users to understand privacy implications in terms of attribute exchange or delegation between different security domains. Though privacy policies are critical for users to give their consent, they are often poorly defined, complex to implement, or simply out of scope of the specifications.

IdM Technology	**Anonymity and Pseudonimity**	**Unlinkability and unobservability**	**Privacy Languages**	**Selective Disclosure**	**Revoking consent**
Federated Model (SAML/ID-FF, WS-Federation)	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**	**Not addressed**
User-centric Model (InfoCards and U-Prove)	Included in the specifications	Message flow eliminates direct communication IdP-SP. Identity selector may encrypt SP identity to prevent the IdP from learning.	Allows to express privacy policies of RPs.	U-Prove enables minimum information disclose if user has personalized her privacy options previously	**Not addressed**
Hybrid Model (OpenID Connect)	Partial anonymity (IdP/OpenID Provider has knowledge of user identity). Recommendations to avoid IdP/OpenID Provider from tracking are not provided.	Opaque persistent or transient identifiers. Use of encrypted channels is recommended. Cryptographic mechanisms do not prevent from traffic analysis attacks.	Not addressed	**Not addressed**	Allows a client to invalidate its tokens if the end-user changes her identity or uninstalls the respective application.

Table 4.1: Summary of privacy properties in Identity Management

4.2.2 The need for a time independent revocation system

The problem of revoking consent is not covered by the aforementioned IdM technologies. In this sense, OpenID Connect is the only specification that addresses partially this issue through time-based and not flexible enough revocation mechanisms, which require an express request from user. Thus, if personal data have been already shared, the effective revocation of consent implies an important challenge to address. For instance, it requires dynamic updates to sticky policies. Consider that Bob, a doctor, has been assigned to Alice for pre-diagnose. He should be authorized to access Alice's medical records (i.e. blood test) but after that evaluation, Bob privileges should be revoked. If the revocation is based on time, as PKI-based solutions, Alice should wait some time until Bob privileges expire to be sure he is no longer able to access her records. In this case, the time window after Bob finishes the pre-diagnosis until the privileges are revoked is a window of opportunity Bob has to access Alice records without explicit permission. This fact difficults the accounting, since after the first legal access, further (illegal) accesses to records will be related to the initial authorization.

To overcome this problem in time based revocation, the validity time of a given privilege set can be set to a very short period of time, so the opportunity window is reduced. 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 during that time.

As a result of this previous discussion, our motivation is to provide a flexible event-based user consent-revocation mechanism. So, in the previous scenario, after Bob finishes the pre-diagnosis, a new event is issued (i.e. needs surgery) and Bob privileges are automatically revoked.

4.3 Towards a Hybrid IdM Event-Driven Consent Revocation Approach

Attribute exchange and delegation process cannot be completely user-centric, since in cases of critical 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. Moreover, to meet the need of a time-independent revocation mechanism, we envision an event-based approach. Here we aim to fill these gaps by designing a hybrid model, which allows users to configure and track access to their medical records while the identity providers are the entities in charge of storing and managing users' *sleepyhead credential*. For the design and validation of the hybrid IdM event-aware model, will follow these steps:

1. **Establishing a set of assumptions during the design of the hybrid IdM event-aware**: It includes the way in which the delegation protocol, along with the event engine, allows to revoke user consent.

2. **Defining a time-independent revoking consent mechanism**: The delegation protocol will include the issuing of a *sleepyhead credential* containing user's attribute identifiers (i.e. her medical history), as well as access privileges, that have been granted beforehand but that are kept latent. To use these attributes, an activation process will be necessary. Moreover, in order to prevent unauthorized access, we require some entities to use a **Consent/Revocation Manager** component of the *Privacy Engine* module (see Chapter 3), responsible for analyzing events and activating the needed attributes and privileges for each event.

3. **Formalization of a mathematical model for the event-based proposal**: Within health care scenarios, patients life cycle can be modeled as event-driven [171] [172]. The proposed model, based on Markov chains with priorities, will enable to trusted entities to fire events when specific circumstances are met, and routed to required entities. These events will enable to awake the dormant privileges of the *sleepyhead credential* or part of them.

4. **Validation of the feasibility of the contribution ideas related to *sleepyhead*-credential delegation protocol and the involved *Privacy Engine* components by developing a proof-of-concept**: Such proof-of-concept will allow to face some important challenges posed by the proposal, namely integration with identity and SIP-based event frameworks, as well as modification of the SAML standard.

5. **Evaluation through simulation scenarios**: We will estimate the overhead of activating/deactivating attributes and privileges, as subscription and notification event messages exchanged, depending on the different health care events, since it is

the main extra part that has been added to the system. For this purpose, we will consider and studied two main simulation scenarios: 1) **A general case**; and 2) **a real case**, which includes a large hospital and a small-medium hospital.

4.3.1 Hypotheses

This section describes the assumptions on which our *Sleepyhead Credential*-based delegation protocol has been built. The *Sleepyhead Credential* is a credential with attributes and privileges granted beforehand, which remain latent until certain circumstances happen.

We assume the existence of an event engine, which follows a notification model based on the SIP-Specific Event Notify [173] specification to send events to entities (by means of broadcast or unicast to registered entities). We assume that the entities persisting the medical records act as IdPs and those requesting access to medical records acts as SPs. SPs, as hospitals, emergency services and even individuals, as doctors, can issue events that will be routed to appropriate medical record holders (IdPs) in order to unblock medical records.

We assume that events follow well know workflows usually triggered by well known trusted entities as emergency services. Moreover, events are achieved allowing rogue entities to be traced if they interfere in the process. It should be noted that SPs and IdPs can take the role of subscribers and notifiers, in some circumstances, either subscribing to different events or notifying them. 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). Besides, each entity can be subscribed to multiple types of events, as well as each event type can be attended by several notifiers.

Patients life cycle modeled as event-driven, as well as Poisson distributed arrival rate of health care events and service rate (i.e. patient arrivals at an emergency service) are widely adopted and well-known by existing research work in the literature, such as [174] [175] [176]. Therefore, we suppose that the arrival process of the events to our system conforms to Poisson distribution with parameter λ and the processing time of these events conforms to exponential distribution, then from the queueing networks result [177], the outgoing process of `NOTIFY` messages is also Poisson distribution (see Figure 4.7).

Concerning security, communication confidentiality should be granted specially for sensitive environments like health care. For that reason we assume the use of HTTPS with mutual authentication to handle message exchange. We assume as well that HTTPS certificates have been correctly issued and distributed. Furthermore, 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 IdPs accountability enforcement in case of user's misbehavior, according to the main principles of privacy specified in Chapter 2. A Public Key Infrastructure can be easily used to support secure communication channels. Finally, we assume an underlying trust relationship based on PKI for entities belonging to different domains.

4.3.2 Implicit Event-based Revocation through Delegation

The main novelty of the event-based user consent revocation model proposed here is the inclusion of a unique credential, named *sleepyhead* credential, which is an event-driven delegation mechanism in which a user delegates to one or more participants the use or access to certain personal information when certain conditions are met. Therefore, the permissions are pre-granted but they are dormant and will not be activated unless established circumstances happen. In this way, we achieve that the different actors of IdM scenario have access to the user's attributes that are allowed the time strictly necessary.

With the introduction of a delegation solution based on events, the revocation of attributes and privileges will take place implicitly. Here we aim to adapt the existing knowledge on delegation protocols to adapt and define a solution that is applicable to provide a more flexible revocation mechanism in IdM environments.

Delegation is a process of an identified entity, called a delegator, giving some of the delegator's privileges to another identified entity, called a delegatee. The delegatee receives

the privileges to act on behalf of the delegator at a service provider [178]. The notion of delegation is widely used as an effective access control method. For instance, many grid systems adopt delegation frameworks to enhance efficiency and scalability [179]. Digital rights management systems provide another example delegation services are offered to consumers, so that they can delegate their access rights to a protected piece of media to a number of devices.

Delegation Models and Sleepyhead Credential Definition

As a previous knowledge base for defining the delegation protocol in the IdM context, we first reviewed the most notable delegation models.

Currently, there are two main models for delegation: the direct delegation model and the indirect delegation approach. In the first case, a delegator directly delegates a set of her privileges or attributes to a delegatee, which uses the delegated identity information to carry out specific tasks or retrieve information. In the second model, a delegator indirectly delegates set of her attributes or privileges to a delegatee through one or more entities. In this latter model there will be more than one delegation step before the delegatee can use the delegated data. Figures 4.1 and 4.2 sketch out these two delegation models.

In an IdM context, service providers trust the identity provider to manage user identities and authenticate users. In such an IdM environment, IdP can, in addition, act as the delegation authority that manages delegations. The delegator assigns delegations at IdP. The delegatee is to perform the delegated tasks at the specified service provider. The SP can obtain delegation assertions from the IdP. Why do we need a delegation authority? In the considered sensitive health care scenario, the user is not always conscious to provide her consent expressly, so the IdPs are the entities responsible for handling, tracking and storing users' *sleepyhead* credentials. The alternative of using the identity providers as trusted delegation authorities has been also proposed in other works, such as [178] [180], but our proposal avoids IdPs to implement and update complex revocation lists, since the revocation of the user's attributes and privileges is being produced based on a series of events that disable them through a single credential, that we have called that credential, the *Sleepyhead Credential*.

Moreover, our approach supports both delegation models: direct and indirect delegation.

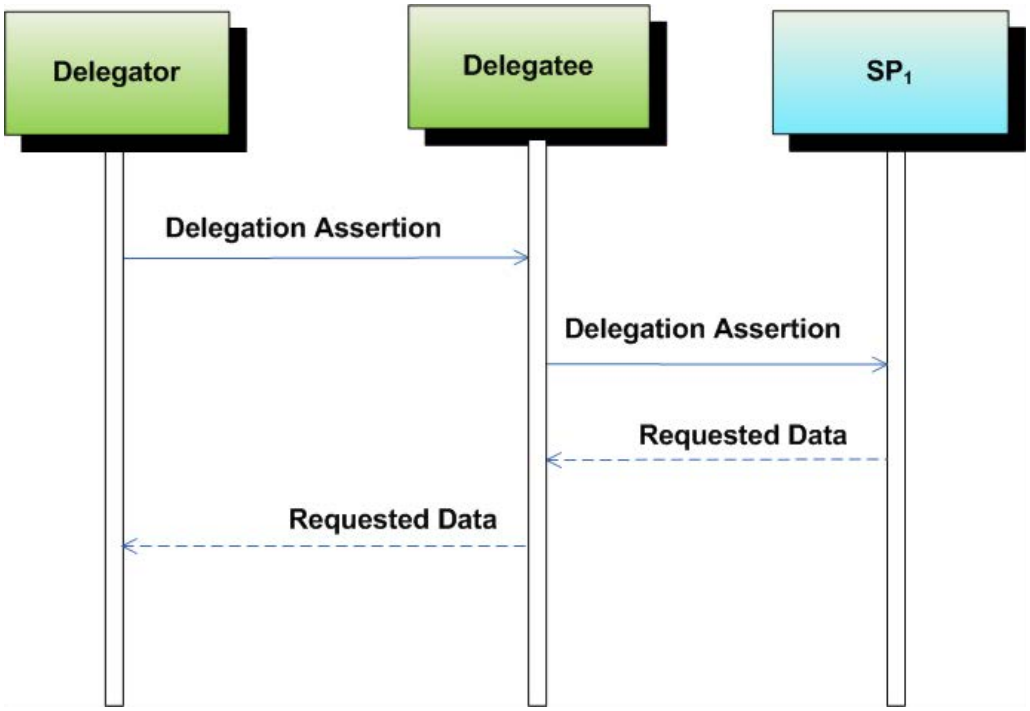

Figure 4.1: Direct delegation model.

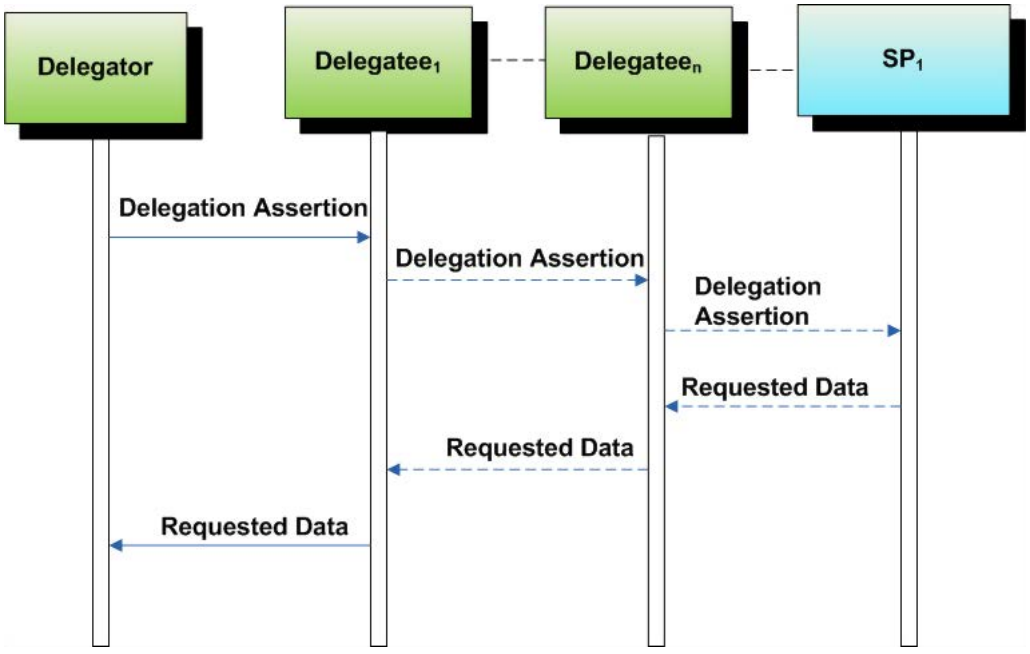

Figure 4.2: Indirect delegation model.

It is worth be mentioned here that the delegation assertion must prove specifically the delegator's consent to the delegation. The delegator may also want to impose certain

conditions on the delegation (e.g., whether or not the delegation assertion is delegatable to another entity depending on some circumstances or events are met, the type of information that can be retrieved, etc.); these conditions must be established in the delegation assertion.

The *Sleepyhead Credential* (SC) is created with the following tuple:

$$\mathbf{SC} = \{AttP_1, AttP_2, \ldots, AttP_n\} \tag{4.1}$$

Where each component $AttP_i$ represents attributes and access privileges, which have been granted to any entity beforehand but that remain latent until activated. A *sleepyhead* credential is composed of fields, including the following elements for delegation restriction: `EventFilter`, defines filters that will be used by the *Consent/Revocation Manager* to analyze the received events and decide whether any attribute(s) may be activated; `TrustedEventSources`, contains entity names whose events activate the credential; `EntityMedicalRepository`, specifies the location and distribution of attributes and medical records.

Protocol Overview

In order to work with the *sleepyhead* credential, first it is necessary to create, registry and publish it within the IdM system. In Figure 4.3, we show a sample of this process:

1. The user (*Delegator*) requests the generation and registration of her *sleepyhead* credential (SC) to IdP_1 through an enhanced client (ECP_1) install on her mobile in Steps 1 and 2.

2. The user authenticates against the IdP_1 to registry the SC in Steps 3, 4 and 5[1].

3. Whether the authentication process is successful, the ECP_1 provides user the *Privacy Preferences Service* in Step 6.

4. In step 7 the *Delegator* establishes her privacy preferences (e.g., different terms or conditions for access to the different parts of her medical history, whether or not an

[1]The specific authentication mechanisms used are out of the scope of the research presented in this thesis.

attribute or privilege is re-delegatable to another entity, the type of information that can be retrieved, etc.) by means of the ECP_1.

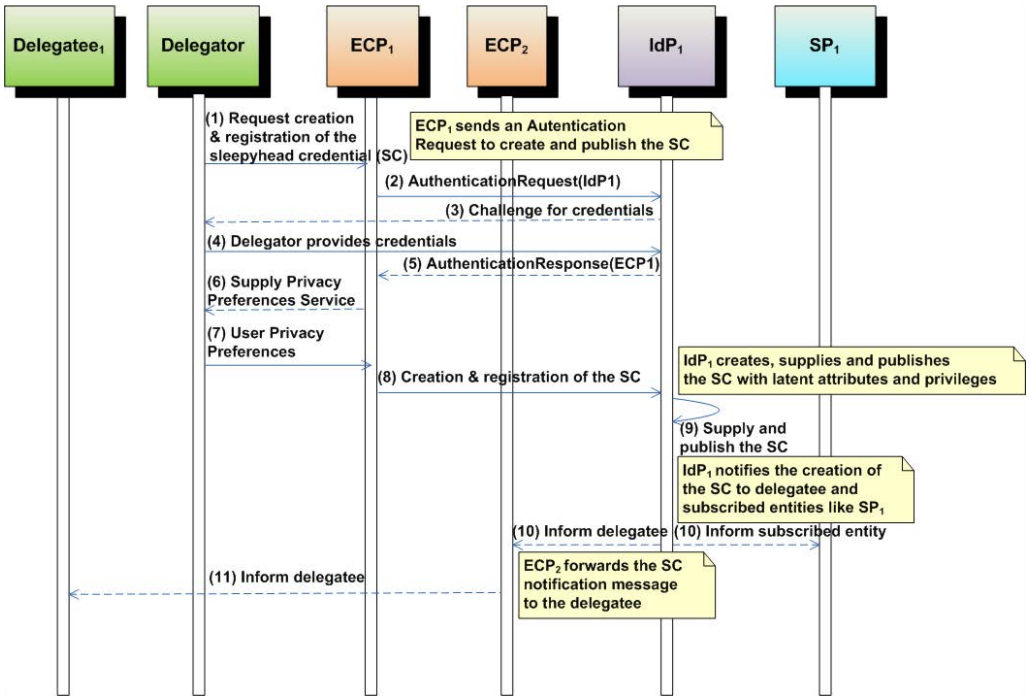

Figure 4.3: Sample *sleepyhead* credential generation sequence diagram.

5. Next, the ECP_1 forwards the above user configuration privacy information to the IdP_1 in order to create the SC in Step 8.

6. In Step 9, the IdP_1 creates, supplies and publishes the *sleepyhead* credential with dormant attributes and privileges; and notifies its creation to the participant entities (SP_1, ECP_2, $Delegatee_1$, etc.) in Steps 10 an 11 in order to physicians and emergency personnel can use the SC.

The protocol messages involved to operate with the *sleepyhead* credential in the hybrid IdM approach proposed, conceptually depicted in Figure 4.4, are:

- `DelegationRequest`. This message is used to ask for delegation, it contains the following fields:
  - `Message_ID`: Message identification number.
  - `Issuer_ID`: The entity (SP or IdP) asking for user's data contained in the *sleepyhead* credential.

- `Subject_ID`: The subject of the delegation or *Delegatee*, i.e., the entity (user, SP or IdP) whose authorization to access to certain attributes/privileges of the *sleepyhead* credential is being evaluated.
- `Delegation Restrictions`: Restriction conditions for the delegation.
- `EventFilter`: The set of filters which allow to determinate which attributes or privileges must be activated/deactivated.
- `TrustedEventSources`: The list of entities (SPs or IdPs) whose events enable to activate the *sleepyhead* credential.
- `EntityMedicalRepository`: The location of attributes and medical records that make up the *sleepyhead* credential.

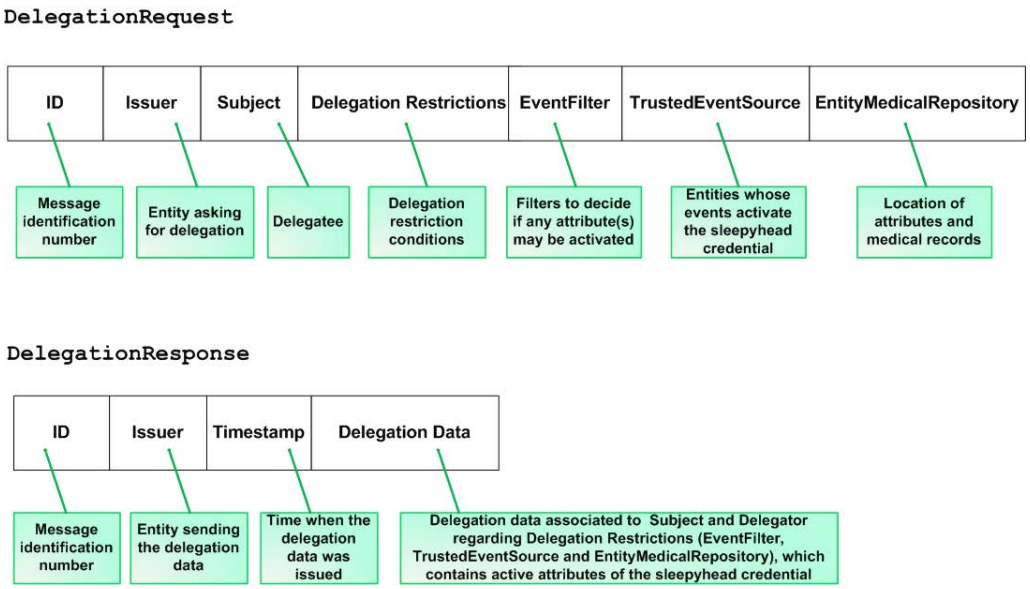

Figure 4.4: Sleepyhead Credential-based Delegation Protocol Messages.

- `DelegationResponse`. This message is used to convey *sleepyhead* credential data in reply to a `DelegationRequest`, it contains the following fields:
  - `Message_ID`: Message identification number (the same as in the request).
  - `Issuer_ID`: : The entity (IdP) sending a delegator's *sleepyhead* credential data.
  - `Timestamp`: Time when the delgation message was issued.
  - `Delegation data`: Associated to the entity identified by `Subject_ID`, in regard

to a set of delegations restrictions comprised of the event filters, trusted event sources and entity medical repositories. The Delegation data contains active attributes of the *sleepyhead* credential of a specific *Delegator*.

Accordingly, the above messages are used to access to user's data through the *sleepyhead* credential. Figure 4.5 illustrates a sample of the protocol sequence diagram for a particular transaction example:

1. If an entity requests access to a service or resource which requires some attributes or privileges (e.g., $AttP_1$), then a `DelegationRequest` message is constructed and sent to the trusted delegation authority (IdP_1) in Steps 1.

2. In Steps 2, 3 and 4 whether the user ($Delegatee_1$) does not have an authentication context established, the IdP_1 challenges her for authentication. A valid security context is created when the user provides valid credentials.

3. When IdP_1 receives a message asking about attributes or privileges of a specific subject (i.e., a `DelegationRequest` message), the IdP_1 must check the delegation restriction conditions (the event filters and entity names whose events activate the credential like SP_1) to determine if the *Delegatee* can access to the requested information (See Steps 4 and 11). Then, it must construct and send a response with the delegation data back to the requester (i.e., a `DelegationResponse` message).

4. As can be seen in Figure 4.5, the event denoted by $Event_1$ allows to awake AttP_1, while $Event_2$ deactivates these attributes and privileges when the $Delegatee_1$ interacts with SP_1. So, the SP_1 provides access to the requested resources or services during the $Event_1$ and it refuses access to AttP_1 when this event ends in Steps 9 and 15, respectively.

SAML-compliant Sleepyhead Credential

Since SAML is the best known identity management framework, we detail how to implement the implicit event-based revocation through delegation protocol over it. For this purpose, SAML offers extension mechanisms that can be used. Including the *sleepyhead* credential-based delegation support to SAML implies modifications to assertions. So, the *sleepyhead* credential has been defined as a new SAML assertion according to the SAML proposal for delegation information defined in [181]. The SAML assertion is defined as

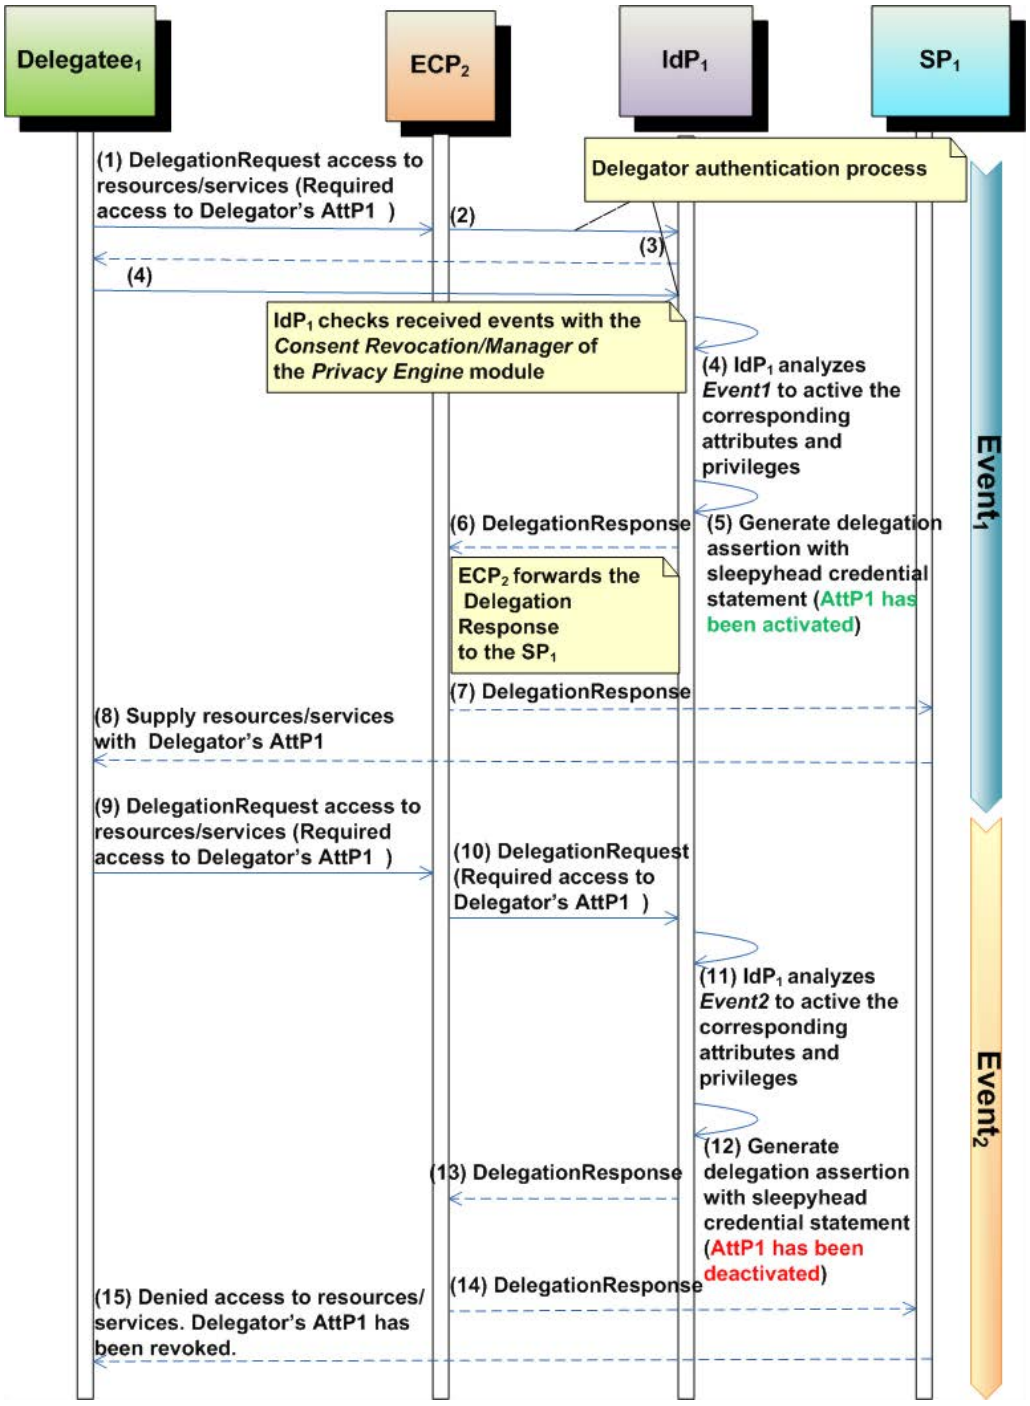

Figure 4.5: Sample use of the *sleepyhead* credential sequence diagram.

follows:

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

Besides, the structure of the `Sleepyhead Credential Assertion` has an initial part or header, whose content is the same that is defined in the standard assertions. This common section includes the assertion identifier, the names of the issuer and the subject, and information about the instant in which the assertion was issued. The XML tags are `<Assertion ID>`, `<Issuer>` , `<Subject>` and `<IssueInstant>`, respectively. And the content for this tags will be the value of `Message_ID`, `Issuer`, `Subject_ID` and `TimeStamp` defined in the `DelegationResponse` primitive.

This `Sleepyhead Credential Assertion` is exchanged using the SAML "Authentication Request Protocol". Therefore, the SAML requests and responses are exchanged using the bindings defined in the specification and they are compliant with the rules defined for extending the schema.

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

We have implemented a proof-of-concept prototype, which is explained in the Chapter 5 and provides further technical details on the implementation issues.

Furthermore, it is worth noting that the choice of implementing the *sleepyhead* credential as SAML assertions is an advantage in the sense that most of the IdM protocols are able to convey SAML tokens. Therefore, messages can be used in other applications. When no SAML bearing mechanism is available, a translation service, to extract the assertion

contents and translate into another token format can be used.

4.3.3 Health Care Application Scenario

In order to show the benefits of our approach, in this section, we describe a potential healh care application scenario. Let us start with some naming conventions for our health care scenario. We will use the term IdP for entities archiving medical records, and with the term SP we refer to consumers of the medical records: hospitals, ambulances and even individuals (doctors). Concerning event management, we assume that events follow well know workflows usually triggered by well known trusted entities as emergency services. Furthermore, events are accomplished enabling rogue entities to be traced if they interfere in the process. It must be noted that SPs and IdPs can act as both subscribers and notifiers, in some circumstances, either subscribing to different events or notifying them. Alice can authenticate by means of a credential to the hospital that persists her medical history (i.e. IdP_1). In case of an accident, IdPs will provide access to SPs to Alice's medical record (or part of it), according to the events. SPs should demonstrate to IdPs, by means of any sort of authentication, that they are elegible (trusted entity as a hospital) to access medical records.

In our scenario, Alice suffers an accident, that triggers several events. The emergency service or 911 (SP_1) requests access to Alice's medical records in order to send them to an ambulance company (SP_2), in another trusted domain, which needs access to her medical records to provide Alice the appropriate treatment. Thus, as events happen, they are notified to the involved parties, such as the medical record service (IdP_1) and the ambulance (SP_2) which treats Alice. So the IdP_1 may know which ambulance should be allowed to access to medical histories.

Furthermore, every medical record access request should be related to an event and bound to a purpose, which enables the IdP to filter the access to certain parts of medical history according to a policy. Thus, in this scenario the following events could be distinguished:

- *Event 1*: There is an accident. SP_1 notifies this event and calls all ambulance services close to the area.
- *Event 2*: SP_2, that is subscribed to SP_1 events, arrives on the scene and requests access to Alice's medical history. It must give a description of the severity of the

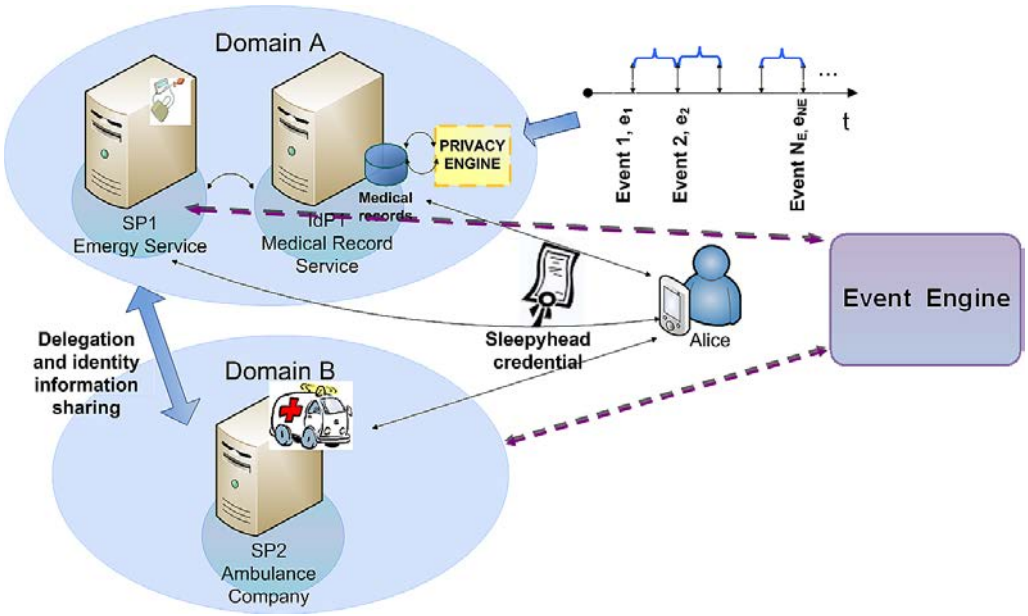

Figure 4.6: Health care event-based scenario across different domains.

problem (event) to allow IdP_1 to give access to certain parts of Alice's medical records or her full history to treat her (purpose). To illustrate this, consider that Alice has broken her femur, has lost 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, SP_2 is allowed to access only to trauma and drug allergies sections of the history.

- *Event 3*: Although not depicted in the Figure 4.6, another possible event would be fired if Alice is taken to hospital (SP_3). The hospital diagnoses her with trauma during the triage and determines that Alice requires an operation. Therefore, a doctor belonging to the hospital (SP_3), 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.

4.3.4 Mathematical Formalization of the Event-based Model

The purpose of this section is to mathematically describe how our event-driven system operates. In this sense, Markov's chains provide support for problems involving decision on uncertainties through a continuous period of time. Specifically, Markov models consider the patients in a discrete state of health, and the events may represent the transition from one state to another. Moreover, these approaches enable to model repetitive events and time dependence of probabilities [182].

So, we assume that events arrive to the system according to a homogeneous Poisson process with rate λ and to be consistent with an exponential distribution. A summary of the definitions and parameters that are used in this section is shown in Table 4.2. Equation 4.2 defines the set of entities of the system (SPs or IdPs), equations 4.3 and 4.4 denote the notifiers and subscribers of the system, respectively. Finally, equation 4.5 describes the set of events that can be triggered by the system:

$$ES = \{es_1, es_2, \ldots, es_{N_{ES}} | ES|\} \tag{4.2}$$

$$N = \{n_1, n_2, \ldots n_{N_N} | N|\} \tag{4.3}$$

$$S = \{s_1, s_2, \ldots s_{N_S} | S|\} \tag{4.4}$$

$$E = \{e_1, e_2, \ldots e_{N_E} | E|\} \tag{4.5}$$

Where,

$$\begin{aligned} &N_N, N_S \leq N_{ES} \\ &N_N + N_S \leq 2N_{ES} \\ &N_N^t = p_N^t N_N \leq N_N \\ &N_S^t = p_S^t N_S \leq N_S \end{aligned} \tag{4.6}$$

Furthermore, being $\{N^1, \ldots, N^t, N^u, N^v, \ldots, N^{N_E}\}$ a subset of N, i.e, notifier subsets of events e_1, e_t, e_u, e_v, etc., then

Parameter	Definition
N_{ES}	Total number of entities in the system
N_E	Number of possible event types (matches the Markov's chain states)
N_N	Number of notifier entities in the system
N_S	Number of subscriber entities in the system,
N_N^t	Number of notifiers in the system delivering events of type e_t
N_S^t	Number of entities subscribed to events of type e_t
p_N^t	Percentage of notifiers in the system delivering events of type e_t
p_S^t	Percentage of entities subscribed to events of type e_t
M^t	Message size to be transferred, considering the overhead introduced by the protocol, when an e_t event is delivered
λ^t	Rate of e_t event arrival
$\lambda^{t,k}$	Rate of e_t event notification rate for notifier n_k
λ_j^t	Rate of arriving events of type e_t to entity e_j
$\lambda_j^{t,k}$	Rate of notified events of type e_t by the n_k notifier that arrives to the entity e_j
$P_{en,j}^t$	Percentage of notified events of type e_t by the n_k notifier that arrives to the entity e_j
μ	Service time for notification of events of type e_t
c	Number of servers or notifiers attending notification of e_t
$\rho = \lambda/(c\mu)$	Congestion of the system with parameters λ, μ and c
K	Maximum number of notification messages that can be buffered by the queue
P_N	Probability of having n notification messages in the system
P_0	Probability of having 0 notification messages in the system
L_q	Notification message queue size
L	Average notification messages in the system
W_q	Average waiting time of notification messages in the queue

Table 4.2: Definition of the parameters for the event model

$$\left|\bigcup_{t=1}^{N_E} N^t\right| = \sum_{t=1}^{N_E} \left|N^t\right| - \sum_{t,u} \left|N^t \bigcap N^u\right| + \sum_{t,u,v} \left|N^t \bigcap N^u \bigcap N^v\right| + \ldots + (-1)^{N_E-1} \left|\bigcap_{t=1}^{N_E} N^t\right| \tag{4.7}$$

$$N_N = \sum_{t=1}^{N_E} \left|N^t\right| - \sum_{t=1}^{N_E} \left|\bigcup N^t\right| \tag{4.8}$$

Similarly with subscribers, it can be established that, if $\{S^1,\ldots,S^{N_E}\}$ are subsets of S, then

$$N_S = \sum_{t=1}^{N_E} \left|S^t\right| - \sum_{t=1}^{N_E} \left|\bigcup S^t\right| \tag{4.9}$$

Also, we define M^t as the message size to be transferred (considering the overhead introduced by the protocol) when a event of type e_t is delivered. Moreover, λ^t and $\lambda^{t,k}$ are the e_t event arrival rate and e_t event notification rate for notifier n_k, respectively. λ_j^t, $\lambda_j^{t,k}$ and $P_{en,j}^t$ are the rate of events of type e_t that arrives to entity e_j; and the rate and percentage of events of type e_t notified by the n_k notifier that arrive to the entity e_j, respectively. Thus, we define the rates as:

$$\lambda_j^{t,k} = P_{en,j}^t \lambda^{t,k} \tag{4.10}$$

$$\lambda_j^t = \sum_{k=1}^{N_{N^t}} \lambda_j^{t,k} \tag{4.11}$$

$$\lambda^t = \sum_{k=1}^{N_{N^t}} \lambda^{t,k} \tag{4.12}$$

$$\lambda = \sum_{t=1}^{N_{N_E}} \lambda^t \tag{4.13}$$

Hence, an entity can be subscribed to several notifiers. Once an e_t event happens, the corresponding `NOTIFY` messages are scheduled to be sent to all the entities which are subscribed to this type of event. Figure 4.7 illustrates the process when an e_t event is received, N_S^t `NOTIFY` messages are generated. Thus, $\lambda^t N_S^t$ is the arrival rate for messages notifying e_t events and there are N_N^t notifiers or servers. In addition, as explained before, the service time for an e_t event `NOTIFY` message is assumed to be exponentially distributed with mean $1/\mu$. Therefore, if we consider a queueing system, that has c servers (being $c = N_N^t$) with K finite capacity, Poisson distributed incoming rates and exponential distributed service rates, this queueing system can be denoted by $M/M/c/K$ [183].

It worth be noted that, when arriving health care events are placed in different queues, each of which has a different service priority. We propose a priority discipline for different categories of events and then a first-in-first-out discipline for each category. For instance,

urgent events will have a higher priority than non-urgent events, since an emergency department should treat patients with life-threatening injuries before others. Moreover, knowing the average frequency of events per period it is possible to derive the probability of a certain number of events to arrive to the system for a given period. This is derived using Poisson's probability distribution:

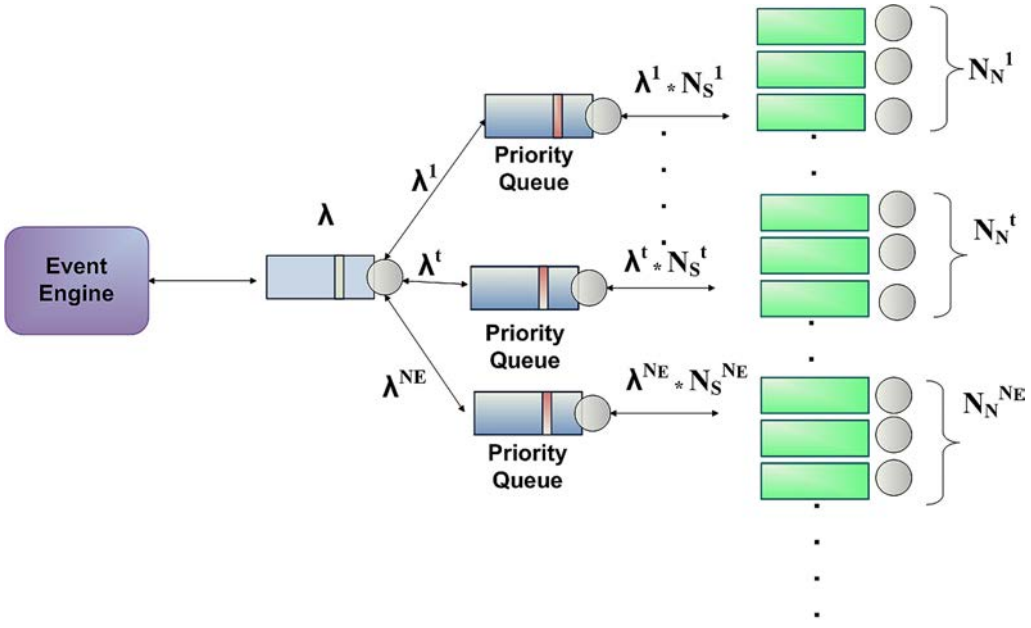

Figure 4.7: Event queueing system.

$$P_n = \begin{cases} \frac{\lambda^n}{n!\mu^n} P_0 & \text{if} \quad l \leq n < c \\ \frac{\lambda^n}{c^{n\ c} c!\mu^n} P_0 & \text{if} \quad c \leq n < K \end{cases} \tag{4.14}$$

$$P_0 = \begin{cases} \left(\sum_{n=0}^{c=1} \frac{r^n}{n!} + \frac{r^c}{c!} \frac{1-\rho^{K\ c+1}}{1-\rho} \right)^{-1} & \text{if} \quad \rho \neq 1 \\ \left(\sum_{n=0}^{c=1} \frac{r^n}{n!} + \frac{r^c}{c!}(K-c+1) \right)^{-1} & \text{if} \quad \rho = 1 \end{cases} \tag{4.15}$$

Besides, we have to calculate the average length of each *Notify queue* in order to estimate the overhead `SIP-Event-Notify` messages:

$$L_q = \frac{P_0 r^c \rho}{c!(1-\rho)^2}[1 - \rho^{k-c+1} - (1-\rho)(K-c+1)\rho^{K-c}] \tag{4.16}$$

$$L = L_q + r(1-P_K), \quad W = \frac{L}{\lambda(1-P_k)}, \quad W_q = \frac{L}{\lambda(1-P_k)} - \frac{1}{\mu} \tag{4.17}$$

For 4.14, 4.15, 4.16 and 4.17 we define $\lambda = \lambda^t N_S^t$ and $c = N_N^t$. Eventually, according to [184] and using 4.17, the average SIP-Event-Notify messages to subscribe to m resources and receive the corresponding notifications (irrespectively of the number of resources that the entity subscribes to) can be defined as:

$$\text{Average\_SIP-Event-Notify\_messages} = 6 + 2L \tag{4.18}$$

4.4 Conclusions

In this chapter we have reviewed and analyzed the main identity models and current frameworks to preserve privacy in identity management systems, identifying its main lacks and drawbacks. Specifically, in this contribution we have addressed the relevant issue of effective consent revocation, since it is not covered by none of the analyzed IdM technologies and it is a must requirement in sensitive environments, such as health care scenarios.

Current approaches are focused on temporal information-based revocation mechanisms being used to predefine policies. Thus, we have proposed a hybrid IdM event-driven model, which includes a *sleepyhead credential*-based delegation protocol compliant with the SAMLv2 standard to provide a more flexible user consent-revocation mechanism within health care scenarios. The main advantages are that this credential is issued only once and would be used any time. Time-based credentials have to be periodically re-issued, for short windows of time to minimize unauthorized accesses, as required.

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, in this chapter we have presented a mathematical model based on Markov's chains and theory queues to determine and study different health care event arrivals to the system and how they are handled and notified to the corresponding subscribed entities. Finally, implementation issues are presented in Chapter 5 together with the validation tests concerning the event engine.

Chapter 5

Event Driven Hybrid IdM Approach Validation

Contents

5.1 Chapter Overview

With the aim to evaluate our proposal, on the one hand we have carried out a proof-of-concept focused on validating the delivery process (over SAML) of security data and information related to the different events that are happening in the system. The *sleepy-head*-based delegation protocol operation, the *Privacy Engine* and the interactions among different entities of our system have been examined, as well, for correctness. In this way, we can prove the feasibility of our proposal and make easier the determination of deployment

requirements. On the other hand, we considered necessary to perform some simulations to estimate the SIP-Event-Notify message overhead generated during the system operation. This is very important in our approach, because it is the main extra part that has been added to the system. The performance of credential issuing is not relevant with regard to time-based credential issuing systems.

As we do not have, so far, a sufficiently rich ecosystem or a large number of entities to conduct a real experiment, we have simulated an event engine through Matlab/Simulink tool [185], including various event sequences with different arrival rates and service times in a random way.

5.2 Event Driven Hybrid IdM Proposal Validation

5.2.1 Implementation Details

We have deployed our own identity management infrastructure using Lasso [186], a C library which implements the full SAML 2.0/ID-FF stack. The IdPs of the systems, are based on Authentic [187], a software that has been developed from Lasso. This library uses OpenSSL as the underlying cryptographic library and Apache2 as the web server. Regarding SPs, we have used ZXID [188]. Furthermore, with the aim to simulate the system of medical events through the *SIP-Notify-Event* specification, we have deployed a Sailfin Application Server [189] 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. Moreover, we have deployed a Registrar Server, to ensure participants are authenticated and registered before exchanging any subscription or notification messages. Once registered, entities might exchange messages containing 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.

The identity management architecture used for the experiments is depicted in Figure 5.1. We used that architecture to introduce the modifications proposed in section 4.3. It can be seen the different interactions between the entities (an IdP and two SPs) through the exchange of SIP and SAML messages. Firstly, the SIP clients are registered in the

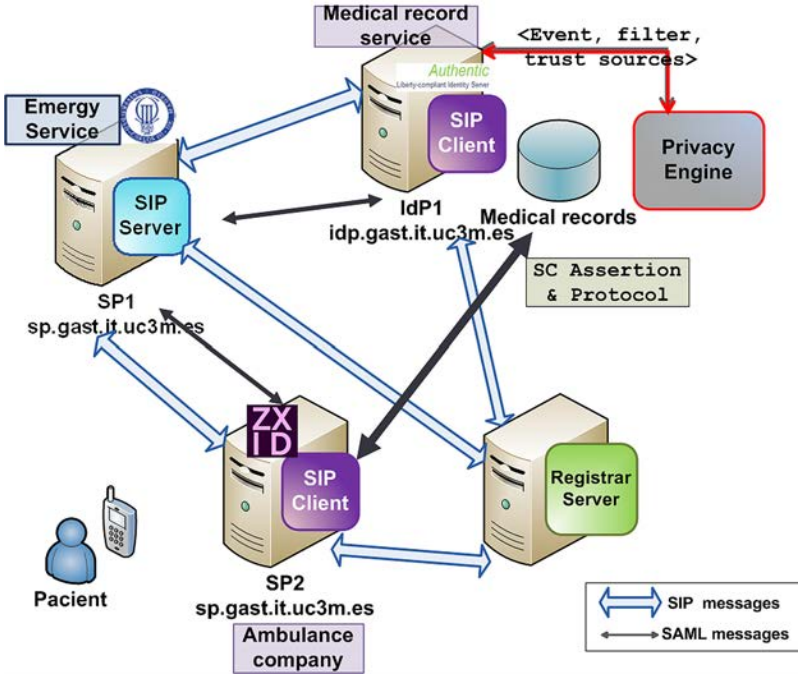

Figure 5.1: Test architecture for the hybrid IdM event-driven proposal.

Registrar Server by sending REGISTER messages. Then, the SIP clients subscribe to different events by means of SUBSCRIBE Requests. The SIP Server notifies events to the subscribed entities through NOTIFY Responses. Once events are received, they are analyzed by the *Consent/Revocation Manager* component of the *Privacy Engine* module and the *sleepyhead* credentials are exchanged through SAML requests and responses.

For this purpose, we developed a preliminary version of the *Consent/Revocation Manager* component, that we are currently refining and integrating with the rest of *Privacy Engine* components (See section 5.2.2), able to receive a data structure, which represents the event filter, and a hash table, which contains the event sources. This building block is in charge of checking that event notifiers are in its *Dynamic Trust List* (DTL) [34]; and accomplishes the decision making process that would determinate under which conditions or restrictions can be the attributes of the medical record disclosed, or the privileges activated. Since the IdP is the entity handling the medical records as well as the credentials, the *Privacy Engine* operation is closely related to the IdP. Thus, the *Privacy Engine* in our solution is an addition to the original IdP.

Besides, we have extended the Lasso library defining a new structure that represents

the *sleepyhead credential* SAML assertion, as well as its different fields and associated attributes. Such assertion is exchanged through SAML messages. To do this, we have modified the metadata exchanged by the IdP and the SPs in order to include the URL or the endpoints to which the *sleepyhead credential* messages would be sent or from they would be received, i.e. the location of the `Sleepyhead Credential Consumer Service`. 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 to apply audit services for events.

5.2.2 Proposal Adoption and Lessons Learned

The proof-of-concept prototype and the architectural design on which the prototype is based, were contributed to the national R&D project "España Virtual".

España Virtual is a CENIT9 project funded by CDTI (Center for the Development of Industrial Technology), Spanish Government. The main goal of the project, that lasted between 2008 and 2012, was to establish a bridge between geography and the Internet through the definition of an architecture, protocols and standards of Internet geography, with a special focus on processing data, 3D visualization, virtual worlds and interaction between users. An specific working package for "Security and Identity Management", where our ideas were developed [190] was included to give a flexible and secure support to the envisioned rich ecosystem of services.

Once we tested the architecture depicted in 5.1, we defined and developed new use-cases to be integrated in the scenarios of the project. More specifically, besides the SPs and IdPs located in the domain of University Carlos III and in one of the participant companies, we offered the possibility of introducing external third parties (e.g., Facebook, Twitter) as IdPs. We successfully tested the introduced privacy extensions with a dynamic federation with these 3rd parties, and the subsequent delegation of user authentication to them.

Note that, some interoperability issues have been addressed; therefore, we have modified the source code of the Authentic to make the IdP-compliant with SAMLv2/ID-FF. The ECP has been deployed in mobile and embedded devices. This has also been integrated with deployed providers. Both providers and ECP are being extended by implementing

the SAML-based *Privacy Engine*. In the ECP, the interfaces to configure user's privacy preferences and check personal data have been developed as an integrated application in Java for limited devices. We are also testing the interaction of these modules with the audit and monitoring services. On the other hand, we have tested guarantee of user's privacy by providing explicit user consent to allow third parties access to certain parts of user profile.

More details on the code, structures used, modification points, interfaces, configuration, threat analysis and so on can be found in the project deliverables [190] [191]. On the other hand, apart from the validation in the context of the above mentioned project, another type of validation was made through the design, implementation and publication of derived use-cases and architectures based on the ideas presented here. Initially, the architecture, main related concepts and implementation issues were outlined in [7] and [6]. [6] and [8] also describe different use cases and simulation scenarios to evaluate our thesis contribution and show the main evaluation results.

Furthermore, the new applications and proposals that reuse the privacy-preserving mechanisms presented in this thesis are the following:

- In the context of **Cloud Computing**, we combined privacy improvements with a dynamic federation architecture layers-based [4] to enable the global scalability and usability that are required for the successful implantation of Cloud technologies, while preserving user's privacy. Thus, our approach empowers users to access cloud services and share digital content without necessarily revealing their true identity to everyone, thanks to the use of multiple identities, for instance depending on the context. Besides, it provides a framework that enables to keep to a trace-off between user's privacy and degree of tracking to obtain an adequate personalization degree in the different services. Moreover, the solution includes a module enables users to have enhanced awareness over their online identity use by introducing monitoring tools and an audit service focused on data sharing through the *Personal Cloud*. This proposal build on the privacy and security model for identity management architectures presented in this thesis.

- In the context of **smart and mobile television**, we proposed FamTV[1] [12].

[1]This work received the Chester W. Sall Award for the 2nd place best paper in the IEEE Transactions 2011 (http://www.icce.org/index.php/awards2013)

FamTV is an architecture focused on offering presence-aware personalized TV by using the benefits of content-filtering and presence detection technologies. But it includes a security and privacy layer where privacy-enabled configurations and rules will be established. This is essential having into account that TVs are usually located in a common place at home and most of the time there are more than one viewer. Thus, if a viewer is watching TV and her personal widgets, such as social network comments, are configured to also appear in the screen. When another user comes into the room, her presence is detected and the elements considered private are automatically hidden.

- In the context of **mobile Internet access and next generation networks**, we use the IdM architecture proposed in this thesis to complement existing WiMAX security solutions [192] [9], while enabling secure personalization of services, as well as improvement of QoS management and user experience and privacy.

The main lessons learned through this phase, as well as the remarkable conclusions and limitations found can be summarized in:

The proof-of-concept has served to face some important challenges posed by the proposal, namely integration with identity and SIP-based event frameworks, as well as modification of the SAML standard. The integration between SP and IdP implementations, even if supporting the same set of IdM specifications, is not a straightforward task. This fact confirms that the introduction of automation features is important. The designed privacy extensions is easy to integrate with current open source identity management toolkits. In our tests, firstly we modified directly the source code of the used toolkits (Authentic and Zxid) to include the new functionality. Next, the key parts of the proposed modules were programmed as external APIs to be used from other SP/IdP implementation. The implemented prototype proves that the core principles of the architecture are workable. Thus, the *sleepyhead*-credential delegation protocol and the *Consent/Revocation Manager* are realizable and can be included in SAML-based frameworks. In the implementation, tests were focused on delivery process (over SAML) of event information and security data.

Further validation is carried out by means of simulation, as it is described in the next section.

5.3 Validation through the Event Engine Simulations

To carry out an adequate dimensioning of the designed system and provide a robust and scalable solution, it is necessary to estimate the `SIP-Notify-Event` messages overhead introduced in order to determinate the required number of notifier entities to attend and notify the event messages according to the subscribed entities and the rate of event arrival without congesting the system. For this purpose, we used the MATLAB mathematical tool [185]. An event engine has been simulated by generating generic health care events that arrive to an $M/M/c/K$ notification queueing system and they are served as described in section 4.3.4.

For the experiments, we used a set of statistical data collected by the HES (*Hospital Episode Statistics*) online service, which contains data related to 18,51 million accident and emergency attendances from April 2013 to March 2014 [193] at major A&E departments, single specialty A&E departments, walk-in centres and minor injury units in England. Furthermore, these data set reflect detailed information provided by 205 English hospitals about emergency attendances and average times depending on the hour of arrival, day, month, etc. Note that, these data are offered at health care provider level. We will differentiate between several scenarios according to the frequency of event arrival, that means, the number of events related to emergency attendances per minute.

Firstly, we will distinguish a general case in which the required number of notifiers and subscribers will be analyzed depending on the mean frequency of health care events conforming to hour and day of event arrival. It must be noted that, the aforementioned average frequency event was calculated in a weighted manner for all health providers (SPs or IdPs) listed in the data set. Afterwards, we will characterize the analysis of the introduced overhead contemplating two specific cases: 1) a *small-medium hospital* and 2) a *large hospital*. After examining the data set, with the term of *small-medium hospital* we will refer to a health provider that received an average of 8.391 emergencies p.a (the Herefordshire Community NHS Trust hospital) and we will use the term of *large hospital* for a hospital with 237.701 emergencies p.a (the Heart of England NHS Foundation Trust hospital).

5.3.1 Simulation of A General Case

We first analyze the proposed event model supposing a general case. For this analysis we use the parameters $1/\mu = 1.25$ events per minute, $K = 100, 200$ and observe the behavior of the event engine in terms of overhead with varying number of notifiers (c) from 15 to 200 and subscribers (10 and 20). In Figures 5.2 and 5.3 we present several plots of Equation 4.18 depending on rates and time of event arrivals.

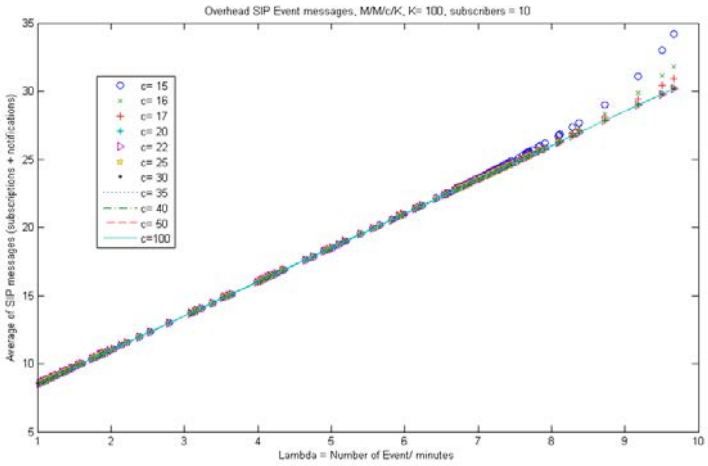

(a) Overhead messages according to notification arrival rates, variable number of notifiers, K=100 and subscribers=10.

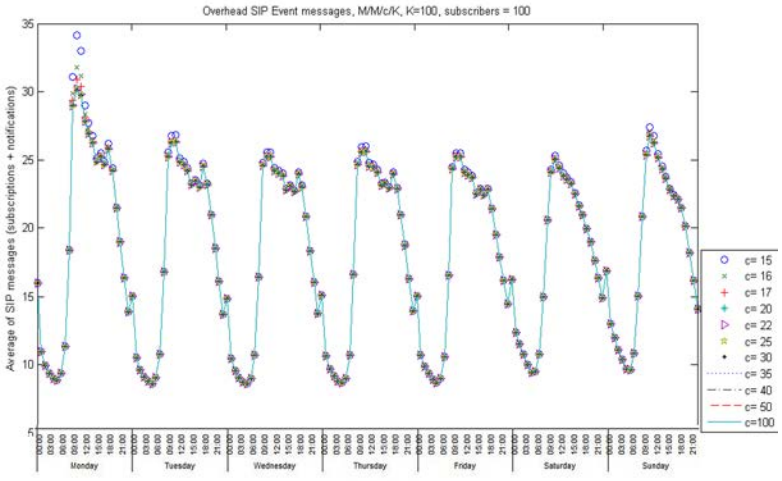

(b) Overhead messages according to notification arrival rates per day and hour, variable number of notifiers, K=100 and subscribers=10.

Figure 5.2: Overhead SIP-Event Notify messages by varying the notification arrival rates and the number of notifiers for a general case. Parameter values: K=100 and subscribers=10.

Subfigures 5.2(a) and 5.3(a) show that the SIP-based event engine introduces an assumable message overhead, which grows linearly with the arrival frequency of events and the number of subscribed entities to the system. As it can be observed, increasing the K value does not have a significant impact on the overhead, but it must be carefully selected to avoid loss of messages without congesting the system. Moreover, we study the system behavior according the hour and day of event arrival (see Subfigures 5.2(b) and 5.3(b)).

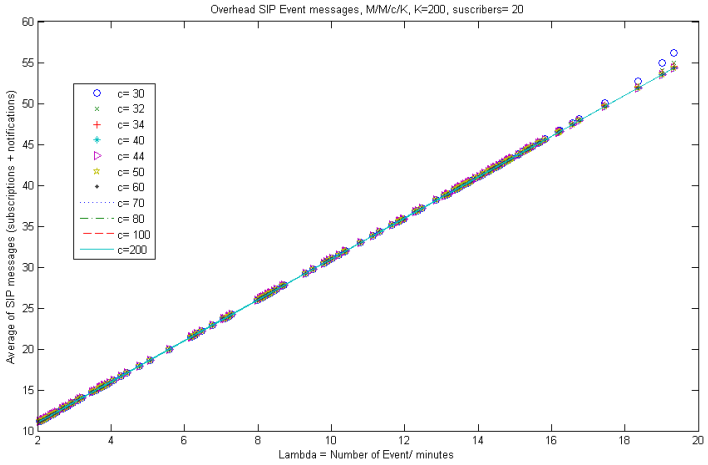

(a) Overhead messages according to notification arrival rates, variable number of notifiers, K=200 and subscribers=20.

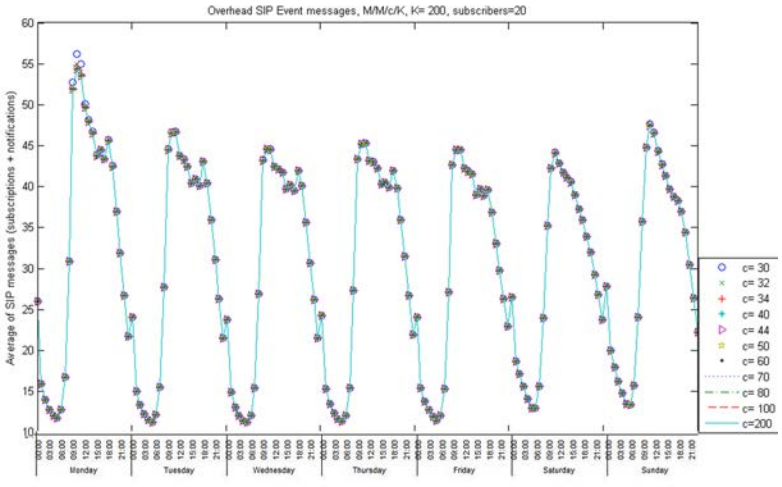

(b) Overhead messages according to notification arrival rates per day and hour, variable number of notifiers, K=200 and subscribers=20.

Figure 5.3: Overhead SIP-Event Notify messages by varying the notification arrival rates and the number of notifiers for a general case. Parameter values: K=200 and subscribers=20.

In these cases, as we can see, the overhead introduced by the event engine is quite similar

for every day of the week, showing a slight increase on Monday from 9:00 a.m to 12 p.m. It is due to the increment of emergency cases. Besides, we can appreciate that generally the highest and lowest rates of arrival of events coincide with the slots from 9:00 a.m. to 12:00 noon and from 12:00 midnight to 6:00 a.m, corresponding to the periods of highest and lowest message overhead respectively.

5.3.2 Simulation of A Real Case

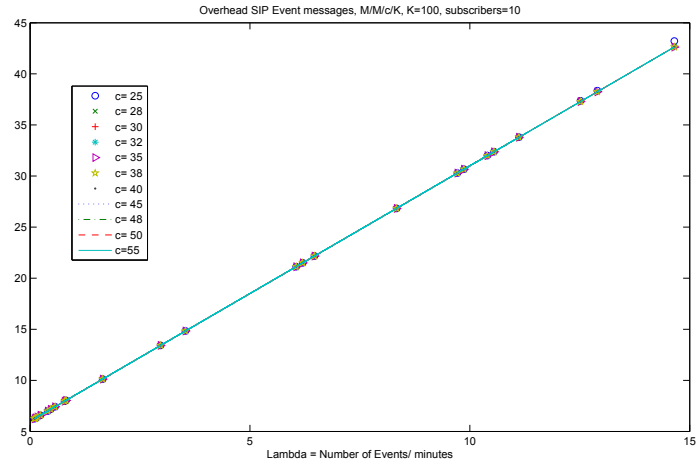

(a) Overhead messages according to notification arrival rates, variable number of notifiers, K=100 and subscribers=10.

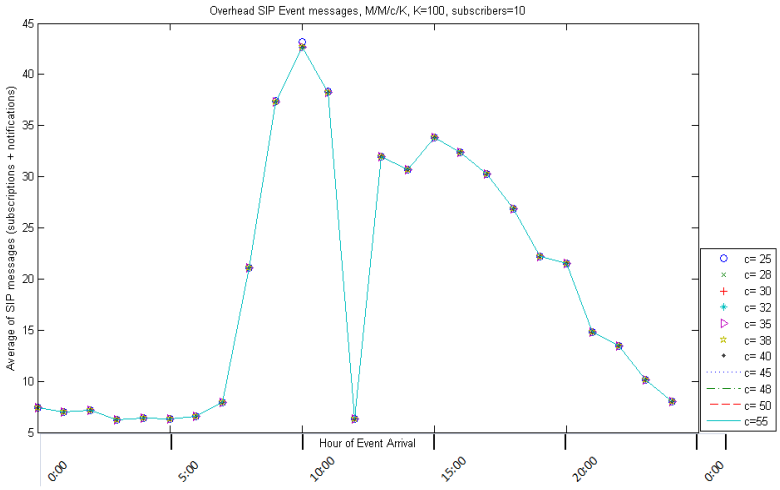

(b) Overhead messages according to notification arrival rates per hour, variable number of notifiers, K=100 and subscribers=10.

Figure 5.4: Overhead SIP-Event Notify messages by varying the notification arrival rates and the number of notifiers for a *small-medium hospital* case.

Similarly, whether we particularize the study for the cases of *small-medium* and *large hospital*, we can observe that in the same way as in the results shown in Figures 5.2 and 5.3, the introduced overhead depending on the rate of arrival of events varies linearly (See Subfigures 5.4(a) and 5.5(a)).

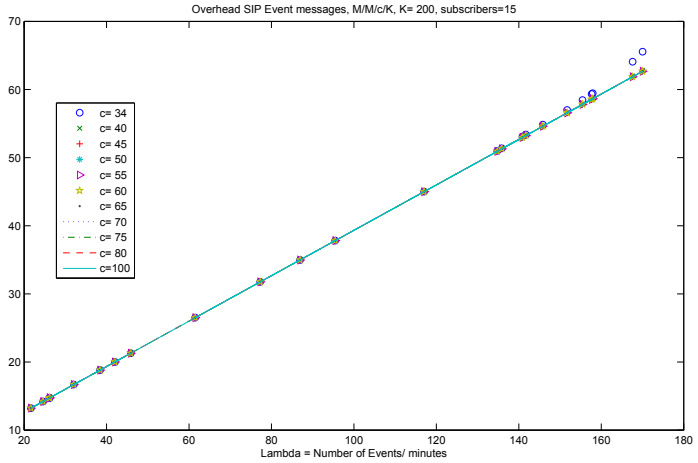

(a) Overhead messages according to notification arrival rates, variable number of notifiers, K=200 and subscribers=15.

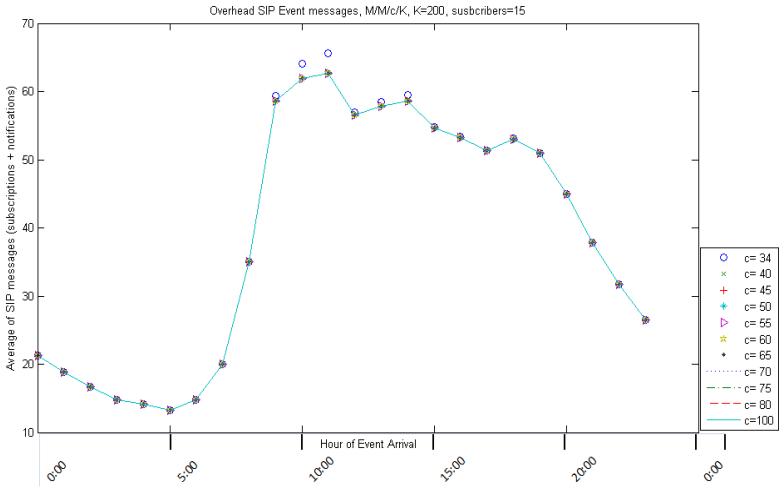

(b) Overhead messages according to notification arrival rates per hour, variable number of notifiers, K=200 and subscribers=15.

Figure 5.5: Overhead SIP-Event Notify messages by varying the notification arrival rates and the number of notifiers for a *large hospital* case.

For the simulation results presented in Figure 5.4 we use the parameters $1/\mu = 0.8$ events per minute, $K = 100$, number of subscribers $= 10$ and distinct number of notifiers (from 25 to 55). It should be noted that, periods of highest and lowest overhead are similar to

the obtained for the general case, except for the decrease occurred at 12 noon, which corresponds to a low health care event rate. In the case of the *large* hospital (See Figure 5.5), we analyze an scenario composed by 15 subscribed entities, a variable number of notifiers (from 34 to 100), with $1/\mu = 3.5$ events per minute and $K = 200$.

5.4 Conclusions

We have covered the validation of the architecture proposed through the implementation of a proof-of-concept and testing prototypes showing that the main ideas of the revocation consent proposal are feasible.

On the other hand, we have observed in the performed simulations, that the message overhead starts to grow exponentially when the event arrival rate doubles the notifiers. This is seen for 34 notifiers (see Figure 5.5); so the parameter that represents congestion of the system is approaching to its saturation value ($\rho \approx 1$). Therefore, it is necessary to achieve a trade-off between the number of entities subscribed to different type of events, the number of notifiers that serve them and the maximum size of the queue for each event type, in order to avoid loss of messages. Note that, a quantitative validation showing the advantages achieved by our revocation consent event-based proposal in compare to time-dependent solutions has not been provided, due to the scarce deployment of consent management infrastructure that avoids to obtain mean values to perform a fair comparative.

To conclude, simulations results show that the event engine introduces an assumable message overhead, since if we consider the the large amount of information exchanged in health care systems, the size of the messages to manage events for implicit revocation is small. However, important parameters such as the number of notifiers and the maximum size of the notification queue, must be controlled in order to avoid loss of messages without saturating the system. In addition, it must be noted that, the usage of the system also affects privacy and should be present in users consents. The auditing processes should verify that the design and assumptions regarding future usage match its actual usage.

Finally, as a future step, it would be interesting to test the architecture in a real world health care environment to validate not only if the architecture is feasible, well-designed and meets the specified requirements, but also if it is useful in real life scenarios.

Chapter 6

Selective Privacy-Enhanced User Profile Management Proposal

The closing of a door can bring blessed privacy and comfort - the opening, terror. Conversely, the closing of a door can be a sad and final thing - the opening a wonderfully joyous moment.

Andy Rooney, 2011

Contents

6.1 Chapter Overview

In different situations of our real life we are who decide with whom we share our secrets and reveal our personal information. For instance, in distinct areas of our either work or personal lives, we just share sensitive information with people from our closest circles of trust and we want to be informed of the third parties involved. It is necessary to transfer these situations to the digital world, by allowing the concept of "man-in the loop" [194] in order to enable users to have complete control over their digital identities and their disclosure.

The purpose of this contribution is to provide precisely a user's profile management mechanism that reveals identity information as selectively as a human would do. This aspect is specially critical in sensitive scenarios, such as health care. For this reason, we will focus on validating and evaluating our proposal in these environments.

Current standards and specifications to share electronic health records [115] [119] [118] are not ready to cope with some aspects of privacy. Specifically, it is necessary to develop techniques for the storage, maintenance, and fine-grained control of sensitive data that permit controlled sharing across different healthcare stakeholders, following minimal disclosure whereas data protection against unauthorized use and minimal disclosure according to patient's consent preferences is provided.

With these premises as foundation, this chapter focuses on studying and defining a flexible privacy-awareness approach for the management of patient's EHR profiles, which is implemented by the *Privacy-Aware User Profile Handler* component of the architecture for privacy provisioning. It is based on a generalized, adaptive and unbalanced Merkle structure. As discussed in Chapter 2 and it will be discussed in greater depth in this chapter, these kind of structures enable to combine user's identity information in a richer and more flexible manner, since user's profiles do not need to follow strict binary, ternary or quaternary structures.

The contribution enables to bring together various patient identity sources to be part of a single credential, while avoiding the creation of bogus patient's EHR profiles. In this sense, solutions based on basic structures need healthcare providers must either see all of the claims or trust the providers of all information, which it is not ideal from a security and privacy point of view. Thus, with our proposal a healthcare service would be able to

accesses the specific personal information without being able to inspect any other details and keeping user control of her data by controlling who can access.

To accomplish this, our contribution relies on the EHR Information Model established by openEHR [119], because it offers an open and extensible framework, as well as archetypes for many clinical terms widely used in hospitals and summary EHR systems in multiple countries. In addition, these structures are publicly available, which facilitates the implementation and adoption of our proposal.

This contribution pursues to bring a data structure that can be saved on constrained devices, as well as adapting efficiently to changes over time thanks to the proposed algorithm. This privacy-aware user profile management contribution re-uses some of the technological mechanisms explained in Chapter 4, such as a hybrid identity management architecture to provide interoperability, a consistent user experience and to control the information exchange in both online and offline scenarios.

The rest of the chapter is organized as follows. Section 6.2 illustrates a use-case motivating the work and highlights the advantages of our solution. Section 6.3 explains a mathematical model to describe the selective privacy-enhanced patient profile management behavior and how the offered privacy mechanisms that can be integrated some use cases are also provided. Next, section 6.4 remarks some security and privacy considerations. Finally, section 6.5 gives the principal conclusions.

6.2 Motivation

6.2.1 Use case for management of EHR profiles

In order to show the benefits of our approach, in this section, we describe a potential use-case that can be realized by applying our proposal. Alice is a diabetic patient who also has hypertension and kidney problems. She finds difficult to manage her condition effectively. Let us assume within a given domain, such as the State of California, we have several healthcare communities in San Francisco and Los Angeles. As Alice travels frequently, has received healthcare in each of these communities. She is undergoing kidney surgery in the hospital in LA (hospital A) next month. The attending physician, Bob, will need to use her hospital information system to query across multiple domains for healthcare

information about this patient (e.g., chronic conditions, critical diseases, past surgical, family history, laboratory results, blood glucose, blood pressure, etc.).

On the other hand, in one of her visits to San Francisco, Alice was admitted to hospital B. Her doctor, Robert, advised Alice to subscribe to a diabetes management program offered by hospital B. As a part of the program, Alice wears a hospital-provided device that continuously monitors her activity level and calories burned, and installs software on her mobile phone. The software processes data it receives from the monitor along with contextual information such as Alice's location. Alice decides to join a social network for diabetics, whose privacy settings enable her to share information with the group (e.g., her daily activity and food intake progress) and to allow complete access to her personal health information to her family members. Once a week, Alice records her weight, blood glucose and blood pressure, using devices that send the measurements wirelessly to her mobile phone. Due to her participation in the management program, Alice's insurance company offers to reduce her premium if she shows significant improvement in controlling her diabetes. In this dynamic scenario, different parts of Alice's medical history can be distinguished and merged as EHR profiles to construct an M-ary Merkle tree according to the openEHR Information Model specification [116] (see Fig. 6.1):

- ***Basic Information***: Patient ID, social security number (SSN), weight, blood glucose, blood preassure and blood group.
- ***Patient Preferences***: Alice has the choice to remain anonymous in the group for diabetics.
- ***Patient Consents***: Alice's husband can access to her complete personal health information (PHR). To demonstrate progress, Alice must provide the insurance company access to certain parts of her health data. She instructs her PHR to provide aggregate information of her activity, diet and physiological parameters.
- ***Therapeutic Precautions***: it considers allergies (e.g., penicillin) and alerts.
- ***Lifestyle***: it includes exercise and food intake progress.
- ***Care Plan***: combinations of goals, targets, monitoring, education concerning the diabetes management plan.
- ***Laboratory Results***: for instance, blood tests.

```
<eee:EHR xmlns:v1="http://schemas.openehr.org/v1">
 <v1:value>example-ehr-id</v1:value><eee:time_created><v1:value>2015-09-
08T19:05:46.29+02:00</v1:value></eee:time_created>
 ........
<all-compositions>
 <data xmlns="http://schemas.openehr.org/v1" xsi:type="COMPOSITION"
   archetype_node_id="openEHR-EHR-COMPOSITION.encounter.v1">
   <name><value>Basic Information</value></name>
   <archetype_id>at001<archetype_id><content>
    <SECTION>
     <name>PatientID</name><archetype_id>at000<archetype_id>
     <meaning>SOAP</meaning>
     .........
    </SECTION>
    <SECTION>
      <name><value>Blood preasure</value></name>
       <items xsi:type="ELEMENT" archetype_node_id="at0002"><name>
         <value>Episode identifier</value></name> <value xsi:type="DV_TEXT">
         <value>2c4a06c2-e3bd-4cd3-a6bb-1fd83df66107</value></value></items>
          <health_care_facility>Hospital B</health_care_facility>
         <content xsi:type="OBSERVATION" archetype_node_id="openEHR-EHR-
       OBSERVATION.blood_pressure.v1">
         <data xsi:type="ITEM_LIST" archetype_node_id="at0003">
          <name><value>data</value></name><items archetype_node_id="at0004">
          <name><value>systolic</value></name><value xsi:type="DV_QUANTITY">
<magnitude>190</magnitude><units>mm[Hg]</units></value</items>
            <items archetype_node_id="at0005"><name><value>diastolic</value></name>
             <value xsi:type="DV_QUANTITY"><magnitude>105</magnitude>
           <units>mm[Hg]</units></value></items>
       </data>
       ..............
    </SECTION>
    <COMPOSITION>
      <name><value>Care Plan</value></name>
      <archetype_id>at0007<archetype_id><content>
     <SECTION>
      <name>Diabetes Management Plan</name><archetype_id>at0007<archetype_id>
      <meaning>SOAP</meaning>
      .........
     </SECTION>
   <COMPOSITION>
     <name><value>Patient Consents</value></name>
      <archetype_id>at000<archetype_id><content>
     <SECTION>
      <items id="text">Informed Consent Details</items><items id="description">Additional
details about the specifics of informed consent.</items></items>......
.........
    </SECTION>
  </content></COMPOSITIOM>
 </all_compositions>
```

Figure 6.1: This figure shows a XML fragment of a patient's EHRs, which can be represented with a tree structure

- ***Prescriptions***: medication orders related to Alice's chronic conditions.
- ***Family History***: Alice's father died of myocardial infarction at 62.
- ***Physical Examinations***: observations appointment, admission and discharge at hospitals A and B.

Due to the significant variation of Alice's diurnal blood-glucose levels, activities related to diabetes management plan, her data and location will be accessed more often.

6.2.2 Advantages of the proposed Adaptive Extended Merkle (AEM) tree-based management

The aim of this section is to justify the choice of the proposed structure, an AEM tree, rather than others, such as binary Merkle or red-black trees or skip lists reviewed in Chapter 2. To select the adequate structure, our work has taken into account on the one hand, the indispensable properties in the context of privacy in healthcare and identity management environments, given the extremely sensitive nature of the information handled and on the other hand, performance needs of these scenarios. Regarding privacy features, the proposed AEM tree provides the following advantages:

- **A richer view of the EHRs by assembling different parts of medical records as profile groups and user's preferences**. Patient's medical history or records do not have to follow a strict binary, ternary or quaternary structure. It is necessary to have a structure that enables to group information in a more flexible manner. Concerning skip lists, their construction needs to order elements and form the first list using the ordered elements. Subsequent lists are built on top of the list by selecting randomly some of the elements from the list immediately below. This will be repeated until there is only one element. However, EHRs may contain attributes and data types of different nature, making it difficult to find a valid sorting criteria for all elements that will become part of the list. Let us consider a patient's EHR example, in which the user has several contacts and critical diseases. To address the treatment of a specific disease, the intervention from 1 to N departments of different hospitals (e.g., surgery, chemotherapy) may be necessary. Each department may implement from 0 to N treatments and each treatment, has a date and may have from 1 to N participants (e.g., doctor, nurse, surgeon, etc.). To achieve a

more flexible structure storage, we study and define a profile management based on a M-ary, adaptive and unbalanced Merkle tree. The tree is distributed suitably according to the frequency user's attributes are accessed. The set of attributes with similar access frequency may have semantic relationships, which will allow to build different profiles to be part of a patient's medical history. To this end, our proposal provides an algorithm for sorting the tree based on patterns of access according to the EHR Information Model [116]. Thus, the attributes frequently required will be placed closer to the root, whilst clinical data whose relevancy to the clinical care of the patient fades in time (e.g., most measurements made on the patients or in pathology) will be located in the lower levels.

- **Combining several sources of EHRs to be part of a single credential**. We use the use case described above to illustrate the potential benefits of combining multiple sources of identity and selective information revelation in collaborative health care environments. In the scenario, solutions based on basic structures would require the healthcare provider must either see all of the claims or trust the providers of all information. This solution is not ideal from a security and privacy point of view. Hence, our approach includes an optional branch to some internal nodes of the full tree and it enables that healthcare providers (hospitals A and B) do not have access to all information about Alice. Healthcare providers are only responsible for claims related to their subject area. Furthermore, the used hash minimizes the need of individual verification of elements along a path and, instead, it would suffice with a root's hash check and the user only has to keep track of one credential. This also enables multiple attributes verification through a single verification tree without revealing information related to non requested attributes.

- **Adaptive performance**. Considering the large information handled and the variability of data is much smaller than in the case of social networks and cloud computing scenarios, it is desirable to have an agile storage structure on read operations. In healthcare environments, response times of insertion or modification operations can be penalized in favor of applying more robust security and privacy mechanisms to protect sensitive information in accordance with the regulatory and legislative frameworks. Although the use of the Merkle Trees makes more difficult to add or update attributes without recomputing parts of the tree as well as changing the root

itself, our work provides an algorithm to improve this aspect, by sorting the tree as we envision frequently accessed attributes to be closer to the root.

6.3 Improving Privacy in e-Health: An Adaptive Extended Merkle Tree-based Management

Before explaining the mathematical model, it is necessary to briefly review the important concepts of the architecture proposed in Chapter 3 in order to identify stakeholders that manage EHRs.

6.3.1 Architecture

We considered our IdM architecture with the following actors: 1) **Service Providers**, which provide services and consume the identity data coalesced by the healthcare providers from several sources. For instance, this role is played by the insurance company in the use case presented in section 6.2.1; 2) ***Identity Providers***, which are entities issuing medical records (e.g., hospital A and hospital B); and 3) ***Users*** (e.g., Alice) with a particular digital identity who interact with SPs and IdPs (see Fig. 6.2).

The information sent to the healthcare providers may contain pieces of data stored in several identity providers and user devices. User's devices would act as an Identity Meta-system , meta-IdP[195], in order to provide interoperability, a consistent user experience and control of the information exchange. In this way, the role of the patient (Alice) is empowered letting her to participate in the process. So, the trust and attribute disclose processes are no longer opaque as in other identity models. The patient is given the ability to configure interactions with healthcare providers and third parties (e.g., the insurance company and the social network), by detailing which attributes may the healthcare providers take from her profile and which ones can be taken from an identity provider.

It worth be noted that, the meta-IdP can be also instantiated in the health care provider to cope with scenarios in which the patient is not online to accept the transaction. More technical details about the IdM architecture has been discussed in Chapter 3.

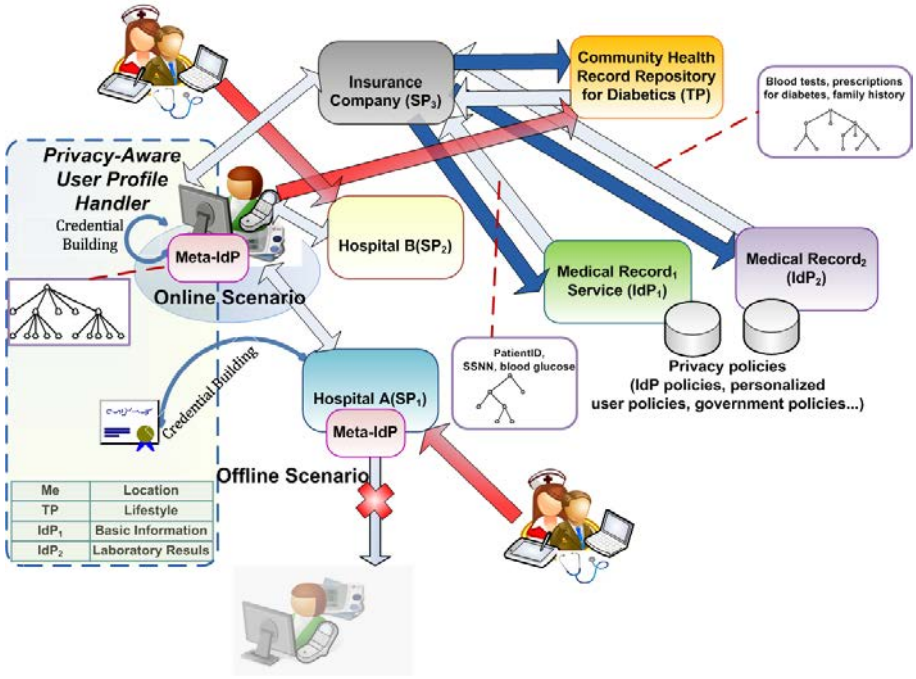

Figure 6.2: This figure illustrates the IdM Architecture. Note the meta-IdP may be instantiated in either the user device or the healthcare provider

6.3.2 Mathematical Formalization and Definition

The purpose of this section is to describe how to handle patient's EHR profiles through a novel AEM tree to convey patient claims to other entities.

An AEM tree is essentially an *M-ary* and unbalanced tree, i.e., each node may have up to M maximum children. Thus, each node holds the hash of the concatenated values of its children nodes. Leaf nodes hold the identity attributes as well as other information as (e.g., node tag, semantic annotation, attribute value, attribute nature and type as *self-issued* - non verifiable - or *provider issued* - verifiable). Node's children influence the node hash and so does the node with its parent until the root node, so a large number of separate data can be tied to a single hash value (root node). In this way, given an attribute and its hash tree, if hashes related to the attribute are consistent until the root and the signature of the root node is valid, it is possible to verify that any of the leaf nodes of the tree are authentic without revealing any further data. Thus, selective patient's attribute disclosure and verification is achieved. To help the consumer to distinguish among different sources, the meta-IdP labels nodes by appending a bit to the end of the hash, true if the attribute

is *provider-issued* (i.e., age, nationality), or false if the node attribute is *self-issued*.

AEM trees are constructed according to the following. A template node, named N, contains several attributes Att_n (of any nature) and children $N_1, \ldots, N_n$. Some node types may contain a Hash value H_N from a summary obtained from its attributes and related to its children. Moreover, they contain a node identifier N_{Id}:

$$N \leftarrow ([H_N], N_{Id}, Att_1, \ldots, Att_v), [\{N_1, \ldots, N_n\}] \tag{6.1}$$

There are several node types: leaf nodes, named LN, contain attributes but no children;

$$LN \leftarrow (N_{Id}, Att_1, \ldots, Att_u), \{\} \tag{6.2}$$

profile nodes, named PN, contain attributes and descendants that should be kept together since they constitute a profile or set of interrelated claims;

$$PN \leftarrow (H_N, N_{Id}, Att_1, \ldots, Att_v), \{N_{p1}, \ldots, N_{pn}\} \tag{6.3}$$

inner nodes, named IN, are structural nodes containing no identity attributes but the necessary hash values to build a verification path from any leaf node to the root (that will be signed by the provider);

$$IN \leftarrow (H_N, N_{Id}, Att_1, \ldots, Att_n), \{N_{i1}, \ldots, N_{in}\} \tag{6.4}$$

the root node, named RN, contains several attributes including an identifier, a time stamp TS (generated during signature) and a signature Sig over the hash value related to its children. Its children contain, as well, a hash value related to their children, until a leaf node. In this way, a provider can certify all the data placing one signature in the root node over the hash value allowing the tree to be lopped by the meta-IdP removing branches without affecting the hash whenever hash values until the claim to be proven are known. Moreover, the root node has a set of special children nodes that contain pseudo identifiers $(P_{id1}, \ldots, P_{idn})$ that are randomly generated when the data structure is signed. Thus a signed tree can be used several times enabling unlinkability. Besides, due to selective disclosure properties of the AEM tree, a degree of unobservability is offered, since the meta-

IdP allows the patient to handle health resources while keeping clear of other entities (e.g., the social network) have access to more information than it is necessary.

$$RN \leftarrow (Sig, TS, H_N, N_{Id}, Att_1, \ldots, Att_n), \{N_1, \ldots, N_n\}, \{P_{id1}, \ldots, P_{idn}\} \tag{6.5}$$

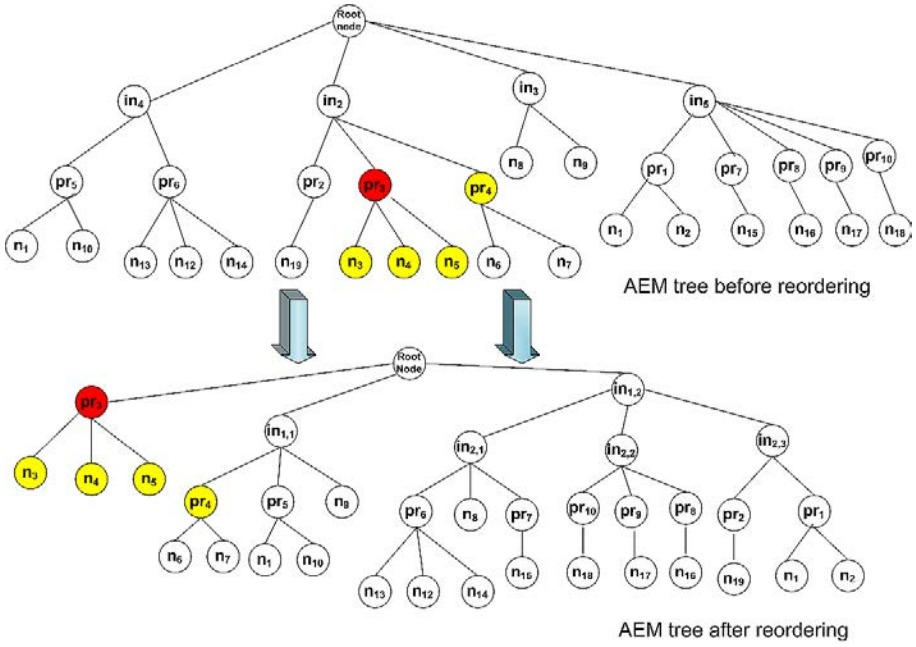

Figure 6.3: Privacy Framework based on extended Merkle trees for EHR profile management.

Figure 6.3 illustrates Alice's medical history according to the use case described in section 6.2.1. We also summarize the information contained in the profile nodes in Table 6.1.

Node	Type of profile node
pr_1	Physical examinations
pr_2	Allergies
pr_3	Patient's basic data
pr_4	Patient's basic data
pr_5	Self-issued lifestyle attributes
pr_6	Physiological parameters
pr_7	Patient Preferences
pr_8	Events or conditions in Alice's family members
pr_9	Laboratory measurements
pr_{10}	Information related to patient consent

Table 6.1: Summary of Profile Nodes

It must be noted that, nodes n_8 and n_9 store PHR data based on compliance with Alice's activity and diet. Leaf nodes n_8 and n_9 contain Alice's location information and her blood group, respectively. Finally, the rest of LN are children of the aforementioned PN.

Furthermore, our AEM tree annotates the query frequencies of its attributes, as a criteria for subsequent optimizations. In this way, the most frequent attributes (e.g., PHR data related to diabetes, Alice's basic information and her location) are placed in the upper levels of the AEM tree, making more efficient and faster the verification process. Hence, the AEM tree can be dynamically optimized according to node frequencies given some structural constraints (tree depth and node children from 0 to M). So, we denote by m the maximum degree of a node, i.e, the maximum number of branches that emanate from each node, the parameter h represents the AEM tree height. In addition, we use the term $L\%$ to represent the contribution percentage variance of the access frequency of the nodes remaining to be placed in the tree. Equations 6.6 define the set of possible AEM tree nodes (including leaf, profile and inner nodes), equation 6.7 their query frequencies, and equation 6.8 denotes the maximum number of AEM tree nodes and the maximum number of leaves, respectively:

$$N_{AEMTree} = \{N_1, N_2, \dots N_k\} \tag{6.6}$$

$$F = \{f_1, f_2, \dots, f_k\} \tag{6.7}$$

$$k = \sum_{i=1}^{h} m^i, \quad M = m^h \tag{6.8}$$

Where,

$$f_1 < f_2 \dots f_{k-1} < f_k, \quad k \geq M \tag{6.9}$$

As mentioned before, a sorting algorithm can be triggered to improve searches. The algorithm works as follows. Step 1: we order the set of AEM tree nodes containing data (LN and PN) by query frequencies in ascending order. Step 2: we take the p nodes, named P, that contribute ($Cvar_i$) to the $L\%$ of the frequency variance (var) of the remaining

nodes (see Equations 6.10, 6.11 and 6.12).

$$P = \{p_1, p_2, \ldots p_p\} \tag{6.10}$$

$$var = \sum_{i=1}^{RM} (f_i - \bar{f})^2 / RM, \quad std = \sqrt{var} \tag{6.11}$$

where RM are the nodes pending to be placed.

$$Cvar_i = (f_i - \bar{f})^2 / RM \tag{6.12}$$

Step 3: the algorithm iterates over the nodes in P. Until P is empty, we take the first node in P, p_i and check:

$$m^h - (m-1) > k - 1 \tag{6.13}$$

The equation 6.13 evaluates if p_i can be placed in this level (since it reduces the maximum number of leafs) leaving room for the rest of the nodes $(k-1)$. If so, we place the node p_i, remove p_i from P and go back to step 3. Otherwise, m new internal nodes are added to the tree and go the next level. Finally, once nodes in P have been placed we move to step 1 where the variance and dispersity for the remaining nodes are recalculated and the following p nodes contributing the $L\%$ of the variance are chosen. This process is repeated until the number of nodes to place in the AEM tree is equal to zero.

It must be noted that, the sorting algorithm is also applied each time a node is inserted or updated in the AEM tree. As the new node does not have historical of access frequency, it will be placed at the "most disadvantaged" positions of the AEM tree, as happens in the real life situations when someone starts at the bottom and works her way up. If this new node is frequently consulted, it will prove itself and its position will improve. As regards the update attributes, whether a node is very frequently accessed when the proposed algorithm is applied, it will be in a good position. Otherwise, it will be located at the lower levels.

Figure 6.3 illustrates an example of how our distribution algorithm works for an AEM tree with parameters $m = 3$ and $h = 3$. Nodes with the highest frequencies are shown in red color. Medium frequencies are have been drawn in yellow, while the rest of the nodes are represented in white. The search algorithm looks for nodes upside-down and left to right, so after running the distribution algorithm, pr_3 and its children (n_3, n_4 and n_5) are put in a higher level and further to the left. The node called pr_4 is also relocated in a position that favors its search and verification when these attributes are shared with hospital A, B and the social network. It must be noted that, the *PN* pr_{10} and pr_7 will be checked before disclosing attributes held by pr_3 and pr_4. The rest of the nodes are placed according the same criteria position respecting the parameters m and h.

6.4 Security and Privacy Considerations

The unlinkability and "partial anonymity" of the proposal stems from the corresponding IdPs services. Users claims, asserted by the IdPs, are only exposed according to the privacy rules and informed consents. When a restricted view is required, opaque parts of the EHR are incorporated and verifiable thanks to the hashes and the opaque and transient identifiers provided by the IdPs. Using Merkle Hash Trees to enforce privacy has been already explored in other works like [101] [102] [103]. Our searching and sorting algorithms may introduce information leakage suitable for a differential analysis. Let us consider a well informed attacker who performs selected searches to initiate new sorting of the AEM tree: measuring the sorting time, the attacker can perform estimations and even models of the attributes and relationships of parts of the EHR beyond her authorization. To prevent such privacy breaches, we propose to introduce random delays in the sorting algorithm. Besides the number of executions of the sorting algorithm should be limited within a given period of time.

6.5 Conclusions

In this chapter we have defined the privacy-enhanced user profile management model that is implemented by the *Privacy-Aware User Profile Handler* component of the proposed architecture to the privacy provisioning for federated IdM. The model is based on a novel

Adaptive Extended Merkle structure, which allows user to have more control over their identities and profiles, by letting them to bring together identity sources comprised of different medical and community health records repositories with identity information stored in their personal devices.

Moreover, we have provided and evaluated the algorithm that allows to build enriched compositions of the patient's medical history and to sort the tree based on patterns of access compliance with open EHRs standards, which empowers to promote their implementation in health information systems. Implementation issues and simulations results concerning the selective privacy-enhanced user profile management proposal are presented in Chapter 7.

Chapter 7

Selective Privacy-Enhanced User Profile Management Validation

Contents

7.1 Chapter Overview

This chapter is dedicated to cover the validation of the ideas presented in Chapter 6. Firstly, section 7.2 shows the simulation results concerning the proposed structure to achieve a selective privacy-enhanced user profile management. Likewise, the integration of the *Privacy-Aware User Profile Handler* with the rest of the *Privacy Engine* components has been tested and it is described in section 7.3.

Finally, section 7.4 concludes by remarking the main results derived from all the the validation tests performed.

7.2 Evaluation of the Sorting Algorithm

Currently there is no standardized set of EHRs on which to test or develop or validate algorithms or protocols. Besides, access to full EHRs is quite limited. To solve these problems, previous research, such as [196] [197] [197] demonstrated the success of synthetic data application. Moreover, there is great value in using synthetic data over sanitized data to ensure user's privacy. Sanitized data, particularly anonymization, is commonly misconceived to ensure confidentiality of private or sensitive data [198]. A disadvantage of using anonymized data is exemplified by the ability to cross-correlate background knowledge with other databases to re-identify individual data records [199].

According to the McGraw-Hill Dictionary of Scientific and Technical Terms [200], synthetic data are *"any production data applicable to a given situation that are not obtained by direct measurement"*.

Thus, synthetic data are generated to meet specific needs or certain conditions that may not be found in the original, real data. This is useful when designing any type of system because the synthetic data are used as a simulation or as a theoretical value, situation, etc. In such a way, researchers may generate synthetic data to aid in creating a baseline for their studies and testing. For instance, intrusion detection software is tested using synthetic data [196].

For the above reasons, we have created synthetic data in order to validate our proposal. These synthetic data include quantitative (e.g., age, weight, blood glucose, blood pressure, etc.) and qualitative data (e.g., demographics and socioeconomic data, lifestyle choices, prescriptions, patient's preferences and consents, etc.) to evaluate different profiles. Furthermore, with the aim evaluate the behavior of the sorting algorithm described in Section 6.3.2, we have developed a prototype in Java and conducted preliminary experiments on the performance of our data structure by generating sets of random requests with different probability distributions. Thus, we have performed two different tests for different tree structures:

- ***Biased tree structure***: In this case, a few nodes have a high probability to be requested. Examples of healthcare data that can be represented as biased tree structures, they are patients' medical records with chronic diseases, since particularly lifestyle, current problems and medications which are continually required by query-

ing. This kind of attributes would be located in the upper level of the AEM tree. Regarding, the attributes consist of clinical data whose relevancy fades fairly quickly, including most measurements made on the patients or in pathology. Attributes in this category are thus potentially very numerous over the patient's lifetime, but of decreasing relevance to the clinical care of the patient in time; it therefore makes sense to place them in lower levels.

- ***Uniform tree structure***: In which each node has the same probability to be requested. These kind of trees can be used to manage a young and healthy person's profiles. This person will go to the doctor for her annual reviews. Therefore, the attributes related to basic information and lab results, such as blood tests will be consulted with similar frequencies.

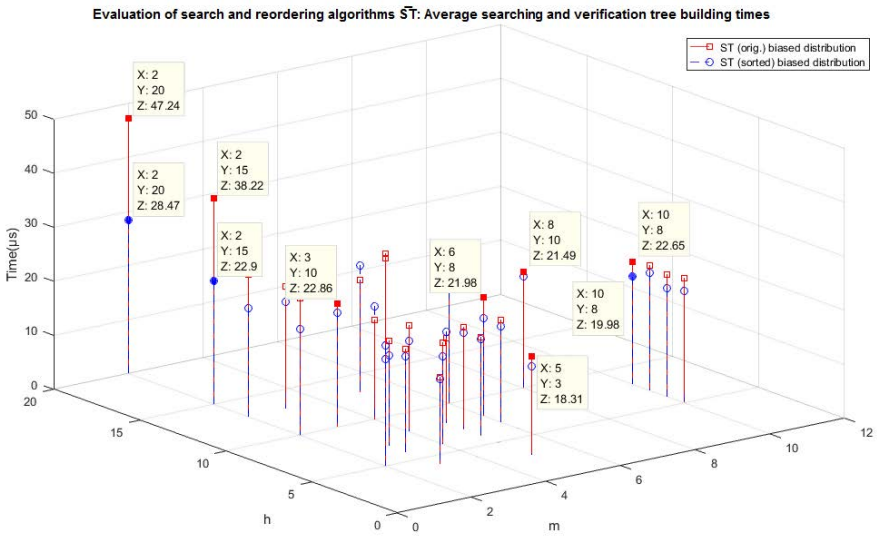

Figure 7.1: The average searching and verification tree building times ($\bar{ST}$) globally for biased structures.

Likewise, in these experiments we have studied the behavior of the frequency-based adaptive distribution algorithm for different AEM tree sizes by modifying both their height and maximum number of children permitted per node. For each operation, the average search verification tree building times, $\bar{ST}$, was computed over $400,000$ trials. The experiment was conducted using a machine equipped with an Intel CORE i7 2760QM with 8G of memory running at 4GHz. Cryptographic hashing was performed using the standard Java implementation of the SHA-256 algorithm.

We have summarized the evaluation results in Figures 7.1, 7.2 and 7.3 through 3D graphics that show the obtained times for searching and verification tree building. These times are represented by the Z axis, when different sizes of m-ary trees are used. To reflect the changes of the trees, the X axis, represents the number of maximum children that each node may have (denoted by m) and the Y axis pictures the different heights of the trees (parameter h) used for the experiments.

For the results depicted in Figures 7.1 and 7.2, we have used biased structures, whilst Figure 7.3 presents findings for uniform structures. Figures 7.1, 7.2 and 7.3 show in red the searching and verification tree building times for the different m-ary trees (the changes of m and h are represented in the axes X and Y, respectively) without applying the proposed sorting algorithm. The findings when the proposed sorting algorithm is executed are shown in Figure 7.1 in blue and in Figures 7.2 and 7.3 in green.

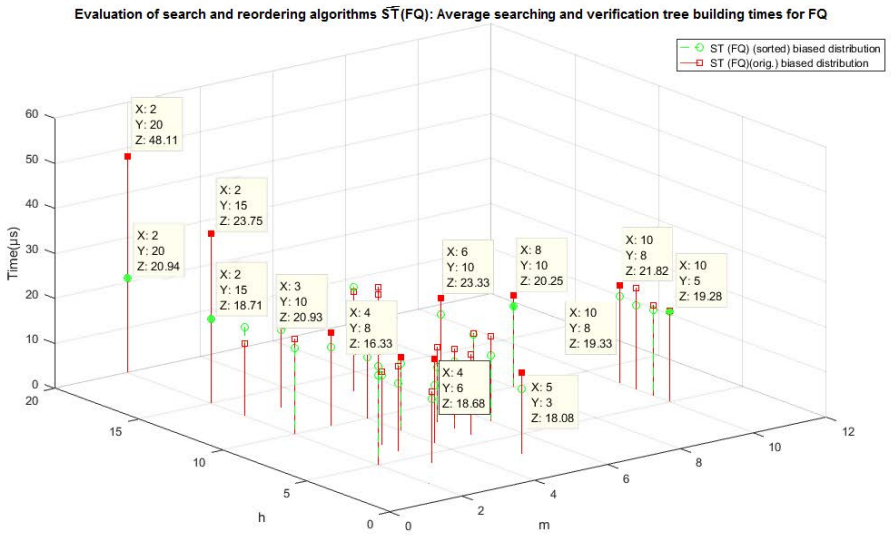

Figure 7.2: The average search and verification tree building times ($S\bar{T}$) for the frequent queries (FQ - those that constitute the 50% of the queries).

We have taken as reference binary trees ($m = 2$ and $h = 5$, 10, 15 or 20) and m-ary trees (e.g., $m = 4$, 5 or 6 and $h = 6$ or 8) and evaluated the average search time ($S\bar{T}$) and the average verification tree length ($V\bar{T}L$) for every node and for the set of nodes that are most frequently queried ($S\bar{T}$ (FQ) and $V\bar{T}L$ (FQ)). Note that, the $S\bar{T}$ includes searching and verification tree building times.

In Figure 7.1, we can appreciate that our algorithm reduces the total searching and ver-

ification tree building times over a 40.08 % and 39.73 % for $m = 2$ and $h = 20$ and 15, respectively. Moreover, the proposed algorithm also decreases the verification path length. For the binary tree cases, the average verification path length before reordering are equal to 7.00 ($m = 2$, $h = 5$) and 13.05 ($m = 2$, $h = 15$), whereas the value of this parameter is reduced to 4.50 and 5.51 after running the algorithm, respectively.

Regarding the outcome of uniform query distribution test (see Figure 7.3), the average searching and verification tree building times are slightly enhanced especially when m decreases and h increases (see the value of the Z axis for instance when $m = 2$, $h = 13$, 15 or 20).

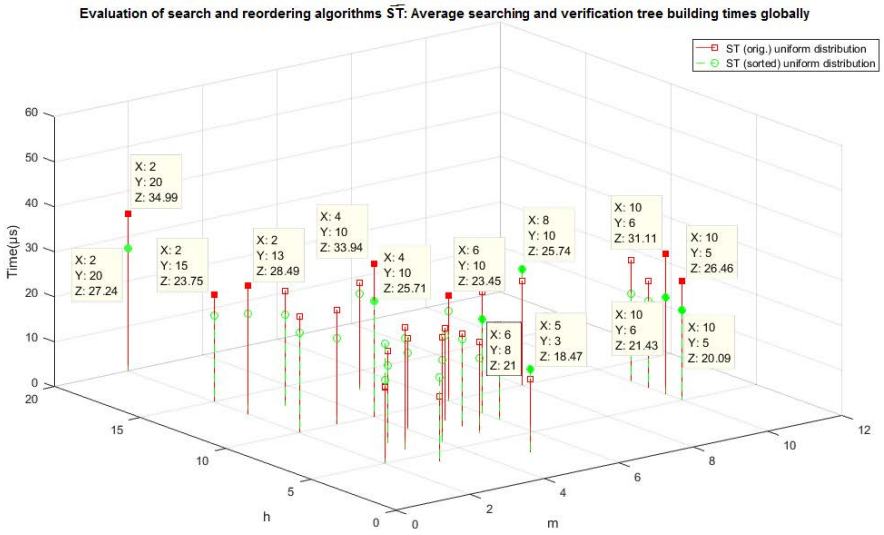

Figure 7.3: The average search time and verification tree building times ($\bar{ST}$) globally for uniform structures.

It is must be noted that, the time spent by the distribution algorithm (OT) is not significant when compared to the improvement over the total search time (Total ST). Although they are not shown in Figures 7.1, 7.2 and 7.3, for instance, for $m = 2$, $h = 15$ and $m = 2$, $h = 20$, the OT are 1670.745 μs and 2711.752 μs, while the total search are 38,71s and 43.57s, respectively. For m-ary trees, the OT are 1531.469 μs and 1463.632 μs when $m = 6$, $h = 8$ and $m = 5$, $h = 6$; respectively. In other words, the time spent by the distribution algorithm is about the 0.04% with respect to the total search time, whereas the reduction of the total searching and verification tree building times after applying the proposed algorithm reaches in some cases for biased structure around the 40%.

On the other hand, our proposal would be dependent of the user and her/his device. It is an application that would allow users to manage their own sensitive records in a flexible and efficient manner. In general, attributes or tags enable faster processing, especially when they are used in a tree and the tree shape varies depending on the use thereof. Furthermore, if they are combined with a tree, the search is predictable and bounded, what could contribute to make attractive our contribution to be used in cloud-based healthcare systems.

In addition, the proposed AEM tree can help to create for instance, a tree for a child in which her basic information (i.e., allergies, blood group, etc) is frequently accessed and has security measures different from other attributes that require, for instance, parental approval.

Finally, the AEM tree can store references to other health records, by governing access to them, but not necessarily containing them. This type of indirections may be necessary in order to increase flexibility storage and it is not incompatible with security, since the AEM tree and its metadata as access information can be used to obtain heavy health data protected and stored by a healthcare provider. Many test results, such as computed tomography (CT), positron emission tomographyŰcomputed tomography (PET-CT), etc. have their own protection and software management in order to be later interpreted by doctors. Besides, these kind of tests take up much space size of a DVD for low resolution results.

7.3 Implementation Issues

With to aim to test the integration of the *Privacy-Aware User Profile Handler* with the other components of the *Privacy Engine* (the *Privacy Preferences Service*, the *Personal Identifier Manager*, the *Consent/Revocation Manager*, the *Audit Service*, etc.), we have developed an Android application which allows the user to configure and edit her security and privacy policies flexibly in order to medical staff, emergency services, etc. can access to various parts of the patient's medical history.

The prototype application works with medical records, which consists of a Java class that contains an AEM tree to represent the different sections and subsections of a medical

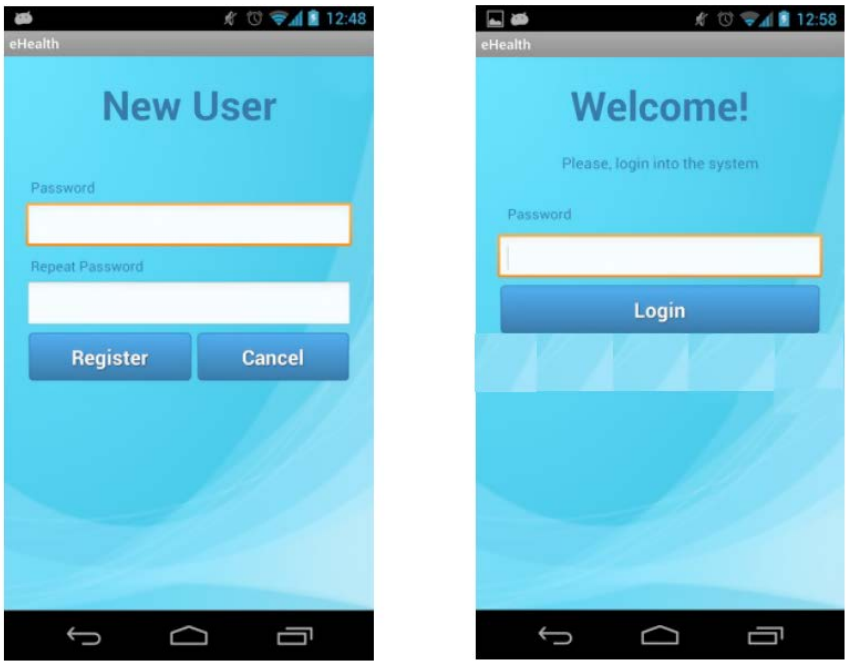

Figure 7.4: User Login and Register Pages

history. For this purpose, the type nodes described in section 6.3.2 have been used.

After installing the developed application, the first time it displays to user a register page. Once the user is registered, when the application is run again, a login page will be as can be seen in Figure 7.4.

When the user is authenticated successfully, the application enables her to perform the following tasks:

- **Consulting the different sections that make up her medical history** through the option named *See Record*.

  When the user selects this option, the application displays a column of buttons, as shown in Figure 7.5.

  Each of these buttons corresponds to a child node of the AEM tree. By clicking on any of these buttons are two possible cases as follows. If the selected node has child nodes, then the application updates the screen showing the corresponding buttons to these children. If one of these nodes in turn has child nodes, the same behavior is repeated. Otherwise, i.e., the pressed node has no child nodes, the application displays a pop-up to inform the user that it has been reached the end of the branch

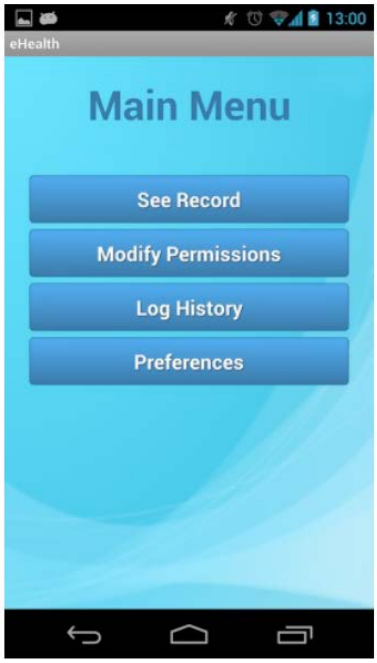

Figure 7.5: View of the main menu of the prototype application.

to show. In both cases, the application updates the text called "Level", located at the top of the screen, with the number of the level of the tree depth, being "Level 1" the initial level.

- **Creating new policies, editing existing policies and creating new policies from other existing** by means of the choice called *Modify Permissions*. As can be observed in Figure 7.6, the application enables user to select a set of nodes of the AEM tree and to define a specific privacy policy to be applied to the chosen nodes through the button "Create Policy". In this example, the following security levels and entities can be configured to access to certain parts of user's medical history: "Emergency Normal", "Emergency Severe", "Emergency Critical", "Family Doctor", "Ambulance", "Radiologist", "Genetic Research" and "Dentist".

  Once the user has selected the security and privacy policies and has pressed the "Allow Access" button (located at the bottom of the screen) an XML-based file, which contains the information of the security and privacy policies is created. Then, this file is encrypted and it can be stored in either the user's device or a trusted IdP.

- **Knowing which entities access her personal** data by choosing the *Log History*

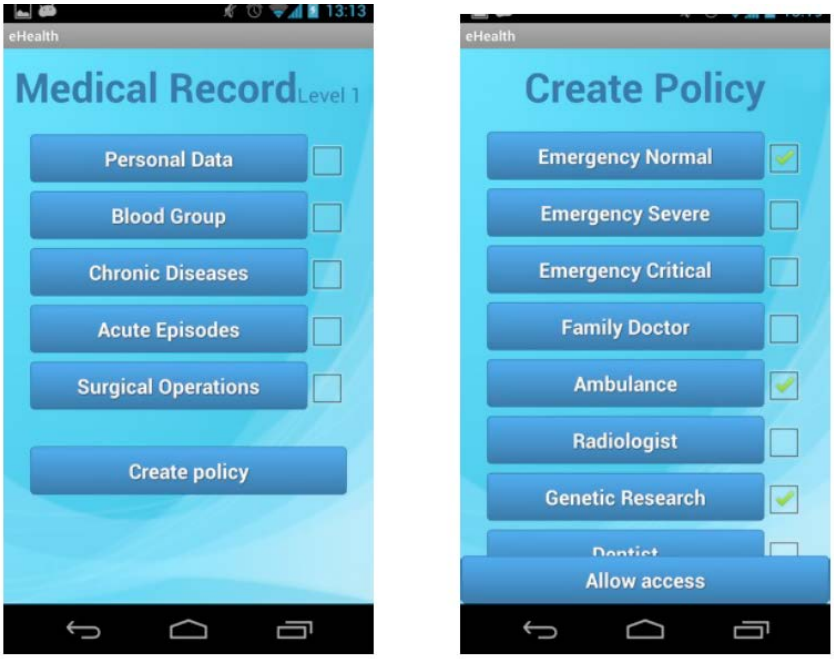

Figure 7.6: View of a patient's medical record and options to modify permissions.

option. For this purpose, the application interacts with the *Audit Service* and the *Action Montoring* components described in Chapter 3.

- **Configuring her privacy preferences** (*Preferences*). This option has allowed to test the correct integration between the rest of the Privacy Engine components explained in Chapter 3 (i.e. the *Privacy Preferences Service*, the *Personal Identifier Manager* and the *Consent/Revocation Manager*) with the *Privacy-Aware User Profile Handler* module.

7.4 Conclusions

In this chapter we have presented the work carried out to validate the ideas presented in regard to the *Privacy-Aware User Profile Handler* module. Validation has been performed through the simulation experiments, implementation and testing of prototypes, as well as through the dissemination and publication in journals and conference, as well as through contribution to R&D projects.

After all these steps we conclude that the selective privacy-enhanced user profile man-

agement proposal is feasible and the evaluation results for different m-ary trees showed that the time spent by the distribution algorithm is not significant when compared to the improvement over the total search time for both biased and uniform structures. The proposed sorting algorithm reduced the total searching and verification tree building times until a 40.08 % when parameter values were $m = 2$ and $h = 20$ in the case of biased structures. In addition, suitable results were also obtained for other m-ary trees by reducing both the searching times and the verification path length. In regard to the findings for the uniform query distribution experiments, the average searching and verification tree building times were lightly enhanced, especially when m decreased and h increased.

Moreover, the main novelty of our proposal is that it is based on open standards to manage user's profiles and it enables to combine different identity information sources to be part of a unique credential.

But also, we identify other aspects still need to be covered:

- The architecture can be further tested through its deployment in real world scenarios.
- Further research is needed to test the integration of our profile management solution using real data and including images and genetic information, as well as more complex information models and archetypes defined by the openEHR Foundation.
- Discussing on the relationship between the actors in implementing privacy protection, particularly the accountability and liabilities of the government, the private sector and the patient should be addressed.
- In the validation of the selective privacy-enhanced user profile management proposal it is required to study and limit the number of executions of the sorting, since the proposed sorting algorithm may introduce information leakage suitable for a differential analysis.

Chapter 8

Conclusions and Future Lines

I don't try to describe the future. I try to prevent it.

Ray Bradbury, 2005

Contents

8.1 Main Contributions

8.1.1 Technical Contributions

The main technical contributions of this thesis towards the fulfilling of the goals presented in the introductory chapter in order to contribute to the privacy-enhanced provisioning for federated IdM platforms are detailed below:

1. **Study of identity management models and gap analysis.**

Identity information, user's profiles and sensitive personal data recently have become key assets for companies and organizations. On the one hand, communications and transactions can be more efficient. On the other hand, they enable service personalization in a cheaper and more agile way.

Current IdM infrastructures allow identity federation and management through features such as account linkage, and profiles, as well as simple session management. In such federation scenarios take part the following main roles: Service Providers, that are entities consuming identity data issued by a trusted third party, Identity Providers that are entities asserting information about a subject, and Users that are the subjects of the assertions. In addition, enhanced clients can be introduced in order to take part of a federation. In this sense, the user acts as an intermediary between providers. The enhanced client also helps to empower the user's role respect to the privacy, because user makes decisions over data control. Moreover, IdM systems mediate in every users' attribute exchange, so they are the preferred target where to deploy privacy solutions.

Due to the importance of the identity management paradigm, the industry and research community have produced a number of standards and specifications representing the fundamental building blocks to accomplish identity federation. However, identity management is still a relatively new technology and so a number of gaps and challenges remain open and need to be filled and solved in order to achieve maturity and wide-scale deployment. With the aim to determine these gaps, Chapter 2 reviewed the state-of-the-art on IdM technologies.

There are several approaches to IdM being the most popular the federated and user-centric models. This summary is the basis for the gap analysis, since both approaches have benefits and shortcomings, for instance, the federated model has scalability issues which the user-centric model solves, and user centric identity requires constant user intervention.

Regarding privacy, the federated and user centric models also present advantages and disadvantages, as discussed in more detail in Chapter 4.

In this dissertation, we center on the privacy problem. The actor Marlon Brandon said: "*Privacy is not something that I'm merely entitled to, it's an absolute prerequi-*

site". According to that, privacy is more than keeping sensitive information hidden. The identities we have developed online are complex and becoming increasingly important to us. Sometimes, we seek complete anonymity and other times, we modify our digital personalities depending on who we are interacting with. For many, full revealing of our digital selves would be unimaginable.

However, current IdM specifications lack of comprehensive privacy frameworks, which allow users to have more control over the use, grant/revocation privileges and disclosure of their online identities. These are key aspects, specially in sensitive environments where improper and unsecured management of user's information may lead to attacks, identity misuse, frauds or privacy breaches.

We decided to tackle this problem and define new enhanced techniques to achieve a suitable supply of the users' privacy at different levels while necessary security services are encompassed. The proposal is based on the introduction of a selective privacy-enhanced user profile management model and flexibility in revocation consent by including an event-based hybrid IdM approach, which enables to substitute time constraints and explicit revocation by activating and deactivating authorization rights according to events and it is an interesting alternative for scenarios where revocation of consent and user privacy are critical. The combination of both models allows to cope with both online and offline scenarios, as well as to empower the user role, by letting users to combine the sources of identity contained in different identity repositories with identity information stored in their personal devices.

2. **Study of privacy aspects in identity management and sensitive environments.**

   In order to acquire the necessary background to design a privacy and security model to provide users greater awareness of the use, disclosure and revocation of their digital identity online, we started by studying how privacy aspects are taken into account in IdM nowadays. Hence, Chapter 2, provides an overview of privacy principles; and summarizes related work being carried out by individual researchers, international research projects and organizations involved in standardization. Then, after this objective review, Chapter 4 gives a comparative analysis of the privacy support in current IdM specifications by emphasizing the relevance of the revocation consent aspect.

The technologies analyzed manage privacy through pseudonyms which can be either transient or permanent. InfoCards, U-Prove and SAML ECP profile address better the principle of minimal disclosure. Current identity frameworks support partial anonymity, since authorities, as the IdP, provides obfuscated identifiers, but do not offer appropriate mechanisms for user's consent revocation. Thus, if personal data have been already shared, the effective revocation of consent implies an important challenge to address.

Furthermore, another aspect not addressed by existing IdM specifications, it is how to combine several claims from distinct identity sources into a unique credential to facilitate management of user's profiles and preferences and to give selective disclosure of identity.

3. **Study of structures for privacy awareness.**

In the state-of-the-art revision in Chapter 2, we have also summarized the main privacy-preserving techniques. As documented in the mentioned chapter, several approaches have been explored. Firstly, identity and attribute-based encryption approaches, in which each user has a defined set of attributes and access policies to decide that the users with determined attributes have privileges to access the shared data. Nevertheless, these approaches require a priori access policies, which are not always ready in real IdM platforms, since the policies to access user's data are sometimes determined after key generation. Secondly, other cloud-based approaches propose privacy-aware schemes focused on providing secure storage, computation auditing, data confidentiality and the query result integrity of sensitive data. Thirdly, spatio-temporal cloaking and ADT-based approaches enable to preserve user's privacy by allowing users to disclose their personal information in a selective manner through perturbations, spatial transformations, generalization or labels of all sibling nodes on the path from the leaf node representing the requested data to the root node.

In order to select the suitable structure, in this study we considered on the one hand, the essential properties in the context of privacy in IdM and sensitive ecosystems, such as healthcare scenarios and on the other hand, performance needs of these environments. The following desirable features were identified:

- It is required to have a structure that allows to assemble user's identity information in a richer and more manageable way, since user's profiles do not need to follow strict binary, ternary or quaternary structures. Regarding skip lists, it is complicated to establish a valid criteria to sort all elements of the list when user's profiles include heterogeneous attributes and data types of different nature. On the other hand, solutions based on basic structures would require service providers must either see all of the claims or trust the providers of all information, which is not ideal from a security and privacy perspective.

- It is desirable to verify multiple attributes *self-issued* and *IdP-issued* by means of a single credential without divulging information related to non requested attributes.

- In sensitive scenarios like healthcare, it is needed an agile storage structure on read operations. The performance for insertion or modification operations can be punished in favor of using more robust security and privacy techniques to protect sensitive information according to regulatory and legislative frameworks.

4. **Design of an extended IdM infrastructure with privacy enabling mechanisms.**

   Once the privacy gaps were identified, Chapter 3 contributes with the design of the required architecture to fulfill our goals. It is clear that current IdM architectures are limited to provide appropriate tools for user revocation consent, which encompass situations where the user is aware to grant or revoke her consent expressly without compromising her privacy. Moreover, aspects related to how users can regulate the use and disclosure of own identity information are not addressed sufficiently by current IdM specifications. Thus, based on the general architectural model that is common to IdM systems, we introduce our extensions to extend its functionality. The architecture is composed of a set of logical modules that separate and encapsulate the functionalities required to achieve an adequate privacy provisioning. The pillars of the architecture are the *Consent/Revocation Manager* and the *Privacy-Aware User Profile Handler* modules of the *Privacy Engine* component, which constitute the main contribution of the thesis. The mathematical models implemented by each part are later developed in the subsequent chapters.

In summary, the extension of the architecture meets the contemplated objectives, since it enables to cope with both online and offline scenarios, while preserving user's privacy and empowering her to better control over her online identities. In order to validate the feasibility of the architecture, we have developed proof-of-concepts and prototypes, which are presented in Chapter 5 and Chapter 7.

5. **Proposal of an event-driven hybrid IdM approach for revocation consent.**

   Having a flexible and effective consent revocation model is essential, but is not covered by none of the analyzed IdM technologies. With this premise as a foundation, we developed the hybrid IdM event-driven model that is implemented by the *Consent-Revocation Manager* component of the enhanced-privacy IdM architecture proposed in Chapter 4. The model includes the definition of a *sleepyhead* credential-based delegation protocol, which allows implicit revocation by delegating user'attributes and rights according to events happen in the IdM system. Specifically, we decided to focus on health care scenarios the implementation and development of the main validations of the proposed model, since they are among the most sensitive scenarios, but it can be also used in other scenarios. For this purpose, we assumed that the development of patients care can be divided into events. These events describe a particular circumstance and can be related to some participant entities.

   The main benefits are that this credential is issued only once and would be used any time; while time-based credentials have to be periodically re-issued, for short windows of time in order to minimize unauthorized accesses, as required.

   Basically, this contribution proposes using events to awake dormant privileges or part of them and it incorporates new characteristics that enable better scalability, since the emergency services are the entities which handle indirectly trust.

6. **Formalization of a mathematical model to an event-driven revocation consent.**

   The proposed event-driven revocation consent model that is to be included by participants in identity management scenarios is explained in Chapter 4.

   A mathematical model based on Markov's chains and theory queues is described to determine and study different health care event arrivals to the system and how they

are managed and conveyed to the corresponding consumer entities.

For this purpose, we assumed the existence of an event engine, which follows a notification model based on the SIP-Specific Event Notify specification to send events to entities. Furthermore, we considered that the entities persisting the medical records act as IdPs and those requesting access to medical records take the role of SPs. SPs, as hospitals, emergency services and even individuals, as doctors, can issue events that will be routed to appropriate medical record holders (IdPs) in order to unblock medical records.

In addition, we considered that events follow well know workflows usually triggered by well known trusted entities as emergency services. SPs and IdPs can act as subscribers and notifiers, in some circumstances, either subscribing to different events or notifying them. Besides, each entity can be subscribed to multiple types of events, as well as each event type can be attended by several notifiers.

Finally, the proposed model considers the events arrive to the system according to a homogeneous Poisson process with rate λ and to be consistent with an exponential distribution. Markov's chains were also included to provide support for problems involving decision on uncertainties through a continuous period of time. Thus, Markov models consider the patients in a discrete state of health, and the events may represent the transition from one state to another.

7. **Implementation and validation of the proposed event-driven hybrid IdM approach.**

In Chapter 5, we presented the result tests of the event-driven hybrid IdM approach. Validation was conducted on the one hand, with the aim to measure the SIP-Event-Notify message performance and overhead generated during the system operation, and on the other hand, to demonstrate the delivery process of security data and information related to the distinct events that take place in the system. In regard to the performance of credential issuing, it is not relevant with respect to time-based credential issuing systems.

Concerning the assessment made through Matlab simulations to calculate the introduced overhead by the subscription and notification event messages exchanged, an event engine was simulated by creating general health care events that arrive to

an $M{=}M{=}c{=}K$ notification queueing system and they were served as described in section 4.3.4.

For the experiments, we utilized a collection of statistical data gathered by the *Hospital Episode Statistics* online service. Firstly, we discriminated a general case where the required number of producers and consumers of events was be examined in accordance with the mean frequency of health care events for each hour and day of event arrival. Next, we considered two specific cases: a *small-medium hospital* and a *large hospital*.

Both simulations results showed that:

- the event engine introduces an assumable message overhead, because whether we take into account the huge volume of information exchanged in health care systems, the size of the messages to handle events for implicit revocation is small.
- the number of notifiers and the maximum size of the notification queue are important parameters to be taken into account, which must be controlled in order to keep away from losing messages without going down the system.

Further work could be done specifically in regard to this part of the thesis, as it will be pointed out in section 8.3.

In turn, the integration of event-based revocation model within the proposed architecture was validated through the implementation of a prototype and its deployment in the context of a national R&D project. Also, different modules of the architecture were used as part of other works published in [6] [4] [7], that illustrate use-cases to improve user's privacy in health care and cloud computing environments.

8. **Formalization of a privacy-aware user profile management model.**

Once the study of structures for privacy awareness was carried out, the main techniques were analyzed and the requirements were identified, Chapter 6 contributes with the design and formalization the privacy-enhanced user profile management approach. The adoption of the structure based on a novel Adaptive Extended Merkle is justified and reasoned in sections 2.3, 2.5 and 6.2.2. The general goals of the model are to empower the user role, by letting users to combine the sources of identity

contained in several identity repositories with identity information stored in their personal devices and to provide an application that would allow users to manage their own sensitive attributes in a selective, flexible and efficient manner.

Thus, we worked on the definition and development of a sorting algorithm for the structure based on Merkle tree extensions taking into account the acquired background and that, attributes or tags enable faster processing, especially when they are used in a tree and the tree shape varies depending on the use thereof. Valid examples are the online tags and free text-based searches. Furthermore, if they are combined with a tree, the search while not optimal, since it varies depending on the type of tree; it is predictable and bounded.

The proposed sorting algorithm is also applied each time a node is inserted or modified in the AEM tree. In the case of a new node is inserted, it is located at the "least benefit" positions of the AEM tree, since it does not have a historical of access frequency. Whether this new node is frequently accessed, it will prove itself and its location will enhance. Regarding update operations, if a node is very frequently consulted when the proposed algorithm is executed, it will be in a great position. Otherwise, it will be placed at the lower levels.

Nevertheless, during the definition of the privacy-aware user profile management model we faced the problem of how to apply our approach to a specific structure that combines attributes of different nature and types in order to validate and adjust the model, as well as to provide a solution compliant with existing standards to handle user's profiles, which leads to the following contribution.

9. **Study of specifications to manage personal information in sensitive scenarios.**

   We reviewed the main current e-health specifications in Chapter 2 in order to obtain the necessary knowledge to implement and integrate the privacy-aware user profile management proposal within extremely sensitive scenarios according to EHR standard formats.

   E-health standards nowadays are based on a dual model architecture, which specifies two conceptual levels: reference model and archetype model. The reference model describes the set of entities that form the general building blocks of the electronic

healthcare record. In regard to the archetypes, they characterize clinical concepts in the form of structured and constrained combinations of the entities contained in the reference model, so clinical knowledge is defined at this level. Both OpenEHR and ISO EN 13606 use this modeling architecture, which has also influenced HL7 CDA.

On the other hand, in order to facilitate the interoperability and provide integration capabilities in the exchange of such EHRs, initiatives as Integrating Healthcare Enterprise Profiles have emerged, but holistic implementations of IHE based e-health infrastructures to share EHRs are currently rare. IHE describes real world use cases. Besides, it tries to deal with security and privacy issues by a modular scheme, which includes integration profiles and use cases to address several issues like recording of the patient privacy consents, exchange formats for medical or care data and announcement the identity of an authenticated user.

However, current e-health standards do not cope with some privacy aspects, such as fine-grained control of sensitive data or selective identity divulging that allow controlled sharing across different healthcare stakeholders in a secure way.

We have chosen the OpenEHR standard to validate our user profile management scheme, because it provides an open development of archetypes, templates, etc. to represent health data. Due to the open nature of OpenEHR, these structures are publicly available to be used, which empowers to promote their implementation in health information systems.

10. **Implementation and validation the privacy-aware user profile management model.**

    Finally, we validated the selective privacy-enhanced user profile management contribution in Chapter 7. Section 7.2 includes two simulation results for different kind of tree structures: uniform, in which each node has the same probability to be accessed, and biased, in which a few nodes have a high probability to be requested.

    To this end, we worked with synthetic data to evaluate different profiles. The use of synthetic data and their validity to evaluate the proposal was justified and reasoned in Chapter 7. These synthetic data included quantitative health data and qualitative data. The evaluation results showed that the proposed sorting algorithm reduced the total searching and verification tree building times until a 40.08 % when parameter

values were $m = 2$ and $h = 20$ in the case of biased structures. Moreover, good results were also obtained for other m-ary trees by decreasing both the searching times and the verification path length. Concerning the outcome of uniform query distribution experiment, the average searching and verification tree building times were lightly improved, especially when m decreased and h increased. Moreover, the integration of the *Privacy-Aware User Profile Handler* with the rest of the *Privacy Engine* components was tested.

Summarizing, the time spent by the distribution algorithm is not significant when compared to the enhancement over the total search time. The proposed sorting algorithm introduces important benefits and reduction of the total searching and verification tree building times for biased structures. Locating attributes more frequently requested in the upper levels of the AEM tree, enables to obtain efficiently and securely this user's information. For instance, patients' medical records with chronic diseases, since particularly lifestyle, current problems or medications which are continually required by querying. Furthermore, the algorithm allows to build enriched compositions of the patient's medical history and to sort the tree based on patterns of access compliance with open EHRs standards. But also, we distinguish other aspects still need to be addressed, as it will be described in section 8.3.

8.1.2 Other Contributions

Apart form the technical contributions described above, other kind of activities were also carried out in the context of this dissertation that convey an added value to its realization. More specifically:

- **Dissemination.**

  The dissemination tasks consisted on the publication of papers and contribution to conferences and journals where the main ideas of the thesis were subject to peer review, evaluation and discussion. The main publications are detailed in Chapter 1, section 1.4.

  On the other hand, we collaborated on security forums and mailing lists related to development tools, such as Zxid, Lasso, Authentic and OpenSSL. Furthermore, part of the contents in this thesis were developed as a research line in several

national R&D projects, called "España Virtual"[1], CONSEQUENCE[2], EMRISCO (EMergency Response In Smart COmmunities) and INRISCO (INcident MonitoRing In Smart COmmunities)[3]. The España Virtual and CONSEQUENCE projects included specific working packages for "Security and Identity Management", whilst the EMRISCO and INRISCO research projects incorporated the Work Package 3 for "Security and Privacy"; where our ideas on security and privacy were contributed.

- **Identification of new research lines.**

  In section 8.3, considering the limitations and the points of improvement of our ideas, we identify a set of open issues for future research works.

8.2 Conclusions

The research on the current IdM frameworks and privacy techniques that best suit the purposes of the thesis resulted in the next conclusions:

- Identity management technologies facilitate handling of identity processes and policies among the collaborating entities. They also enable secure resource sharing among these entities and mediate in every users' information exchange. So, IdM systems are the preferred target where to deploy privacy solutions.

- Privacy and the way organizations treat consumers' data is becoming an increasingly important competitive factor. The explosive innovation that occurred as a result of the arrival of the digital age and technological progress have all changed the rules of the game. Hence, the uses of personal data are relatively advanced- rivaling, and in some cases surpassing, the Internet sector. Partnering and sharing data with external parties, such as producers, consumers, retailers and other social related parties, is not uncommon. For these reasons:

  - Companies and organizations need balanced solutions protect privacy, generate value and create competitive advantage.

  - Consumers want more control of privacy, yet at the same time, want more con-

[1] http://www.espanavirtual.org/
[2] http://consequence.it.uc3m.es/
[3] http://www.inrisco.org/

venience in the way they perform transactions and interact with organizations, but the reality is quite different. Few consumers know how their data are used and fewer can control them.

- The literature review showed the wide variety of specifications and open development tools to implement and deploy IdM-based solutions, as well as the variety of privacy preserving techniques. The preservation of user privacy in federated IdM is a very important factor, but as discussed in Chapter 2 and Chapter 4, this factor does not have catered sufficiently in current frameworks or models. More specifically, on the one hand, there is still scarce work on effective revoking consent mechanisms within current IdM technologies. In this sense, there are currently two main options to revoke attributes or privileges in an IdM context: to set a very short period of time for the validity time of a given privilege; or to set the token expiration to hours, days or months. In the first case, the systems' usability probably is reduced since the user have to re-authenticate to obtain a new valid security token. In the second case, security problems may arise.

  On the other hand, none of the analyzed works deal with neither building of the enriched structures to represent user profiles based on standards nor combining identity information from different sources to be in a single credential.

- Not all data is considered equal when it comes to privacy. Research studies like [2] show that consumers consider some types to be much more sensitive than others and accordingly, they are less willing to share it. For instance, financial data, health records, and social network posts are also considered very private; while location data and past purchases less so. This led to focus on sensitive environments, mainly healthcare scenarios, to perform the proposals and validations of this thesis.

- We have also learned about the standardizations efforts by several government bodies and non-government organizations to implement security measures to share and exchange electronic health records. More specifically, emergency access is a necessary part of access control and will be necessary under emergency conditions, although these may be very different from those used in normal operational circumstances. These procedures are written instructions and operational practices for gaining access to necessary EHR during an emergency. Beforehand, organizations must decide what circumstances would warrant emergency access to EHR information.

However, heterogeneous nature of environments that manage EHRs, security and privacy concerns make difficult the maintenance and deployment of these systems.

- The architecture proposed in this thesis tries to solve the above problems and enables the provision of appropriate tools to manage user's privacy by including more flexible implicit consent revocation and privacy-aware profile management. The new delegation-based revocation mechanisms allow to crowd out explicit revocation and time constraints by activating and deactivating authorization rights in consonance with events. The new user profile management mechanism allows to guarantee selective identity information disclosure in a seamless and scalable manner. The combination of the *Consent Revocation Manager* and the *Privacy-Aware User Profile Handler* components also avoids the incessant mediation of the user and deals with both online and offline scenarios.

- The proposed event driven hybrid IdM model for user's revocation consent provides a quantitative formalization to determine the different event arrivals to the system and how they are managed and broadcasted to the required entities; and introduces flexibility. The main features of the model are:

  - Revocation is implicit and is achieved through the delegation of attributes and privileges that are pre-granted in advance and turned on/off according to the *sleepyhead* credential-based delegation protocol and the events happened within the IdM infrastructure.

  - The difference with related works focused on time-based credentials is that the *sleepyhead* credential is issued only once and would be used any time, without requiring to be periodically re-issued.

Furthermore, we have developed a proof-of-concept in order to demonstrate the feasibility of the *sleepyhead-credential*-based delegation protocol and its integration with the rest the *Privacy Engine* components. Such proof-of-concept served to face some important challenges posed by the proposal, particularly integration with identity and SIP-based event frameworks, as well as changes and adaptations of the SAML standard.

Finally, we performed simulations to estimate the introduced overhead by the subscription and notification event messages exchanged between the entity of the pro-

posed IdM architecture in order to awake/slumber the necessary attributes and privileges. The simulations showed that the message overhead started to grow exponentially when the event arrival rate doubles the notifiers. So, the parameter that represents congestion of the system is approaching to its saturation value ($\rho \approx 1$). However, as we will see in section 8.3, the validation of the model can be extended.

- The proposed privacy-enhanced user profile management model is based on a novel adaptive extended Merkle tree to give users more control over their online identities, as well as to amalgamate identity information collected in their personal devices with sources of identity contained in different repositories in secure and efficient manner. In this sense, the proposed model would be dependent of the user and her/his device.

  Moreover, the AEM tree can store references to other identity data, by governing access to them, but not necessarily holding them. This type of indirections may be necessary in order to increase flexibility storage and it is not incompatible with security, since the AEM tree and its metadata as access information can be used to obtain heavy data protected and stored by the providers. For instance, many test results, such as CT, PET-CT, etc. have their own protection and software management in order to be later interpreted by doctors. Besides, these kind of tests take up much space - size of a DVD for low resolution results.

  On the other hand, the proposed AEM tree can help to create for instance, a tree for a child in which her basic information (i.e., allergies, blood group, etc) is frequently accessed and has security measures different from other attributes that require, for instance, parental approval. Likewise, the AEM tree can be used in another application scenario, which contemplates experiments in patients with complex diseases but requiring notification or confirmation to cross data like sex, race or religion, etc. by allowing audited access.

In conclusion, we have developed an initial approach towards the privacy provisioning for federated IdM platforms. Privacy is increasingly becoming an area of competition for organizations, which can differentiate themselves by providing the right privacy controls and privacy-by-default product design. With easier-to-use privacy protection features, and privacy education, organizations might be able to significantly boost consumers' willingness to share data.

Control and convenience are important aims. They are often conflicting aims, too. Balancing them will not be easy, but it will be critical. We believe indeed that individuals are more willing to take risks when they feel more in charge. Yet while this might be a contributing factor in their decision process, we would argue that choices and control simply help individuals adapt their sharing to their specific preferences.

However, our proposal is a set of preliminary ideas, partially validated and, some of them, prototyped. As a preliminary work, many limitations and weaknesses exist that open the room for further research and improvement. With the aim to identify these lacks and open issues, the following section presents the main future research lines that can be derived from here.

8.3 Future Research Lines

In this thesis we have contributed to evolve federated identity management platforms towards more flexible, privacy-enhancing and secure models. Despite having addressed all the objectives set at the beginning of this thesis, the work carried out during the process has also opened many interesting paths to explore that can be considered as new objectives to focus on. We recognize several areas where the architecture and components presented here can be improved or extended, all of them explained below:

1. **Study and definition of a semantics of events standard format for consent revocation in identity management.**

   Event-based publish/subscribe revocation systems offer a convenient abstraction for data producer and consumers, as most of the complexity related to addressing and routing is encapsulated within the network infrastructure; and avoids exposing sensitive information to other entities for longer than strictly necessary.

   The event-based revocation model proposed in this thesis is based on the existence of an event engine, which follows a notification model to send events to entities (by means of broadcast or unicast to registered entities). The procedures to extract the information of events consist on analyzing the information available for their category in order to activate or deactivate the required attributes or privileges in the *sleepyhead* credential. We developed our event driven hybrid IdM model assuming

that this information was available.

In this sense, ontologies have the potential to become richer as the stakeholders contribute new knowledge. However, it is almost impossible for a single ontology to cover all the domain of complex and sensitive ecosystems like healthcare, since medical services cover many areas such as patient care, clinical and administrative decisions, assisting devices, patient diagnostics.

As a first step, a domain ontological model for organizing knowledge and information of events in the heterogeneous domain of embedded devices and emergency systems could be explored. This work could provide a conceptualization for knowledge in emergency management and the representation of device, SPs and IdPs entities along with their attributes that can enable information and knowledge interoperability among other systems. Different reasoning tools such as Protégé, Racer, Pellet or FACT++ could be used to perform intelligent reasoning on the Web Ontology Language (OWL) ontologies.

In a next step, other specifications and techniques to provide a more generic solution to define a semantics of events standard format to improve the revocation of attributes in the context of identity management should be studied.

2. **Further analysis of the event driven hybrid IdM model for user consent revocation.**

   Further analysis is needed in the topic of self revocation of *sleepyhead* credentials. We also plan to take into account different privacy requirements for identity attributes, including biometric and health care data. Further research is also needed in preserving user privacy during the exchange and sharing of attributes in different trust domains, also considering usability of the system.

3. **Further analysis of the privacy-aware user profile management model.**

   Developing an implementable model for user profile management is an art involving many separate design problems and choices. We have defined the main features for the privacy-aware user profile management model to be applicable in sensitive environments, such as healthcare scenarios. We started a simulation testbed to prove the benefits of the model, showing the total searching and verification tree building

times; and the verification path length reached by applying the proposed sorting algorithm. But further work is still required.

Among others, the following aspects could be addressed:

- Study and definition of mechanisms to determine and adjust the optimal size of the structure dynamically. Initially, this value would be preconfigured on user's devices taking into account mainly their memory constraints, user privacy preferences, the number of confidence IdPs with which the user interacts more usually, default number of attributes more frequently consulted, etc. As a future line, mechanisms that allow to adapt the values of m and h dynamically, could be explored. For instance, if it is detected that the search times and verification path lengths exceed a certain threshold, it should be taken into consideration parameters such as, the number of more frequently accessed nodes, their variance and dispersity, etc. to better adapt the structure size.

- Security analysis and enhancement. The proposed sorting algorithms may introduce information leakage suitable for a differential analysis. Let us consider a well informed attacker who performs selected searches to initiate new sorting of the AEM tree: measuring the sorting time, the attacker can perform estimations and even models of the attributes and relationships of parts of the EHR beyond her authorization. To prevent such privacy breaches, random delays could be introduced in the sorting algorithm. In this issue a new research line is opened in which, the number of executions of the sorting algorithm should be studied and limited within a specific time period. This, and other possible attacks should be simulated and investigated in order to improve the privacy-aware user profile management model to take into account high security and privacy considerations.

4. **Further validation of the revocation consent proposal to contemplate more event categories and define new priority levels.**

   Regarding the priority queues of the event driven hybrid IdM approach, the developed prototype and simulations carried out to demonstrate the workability of the presented ideas only classify the simulated health care events into two categories: "Urgent Events" and "Non-urgent Events". Thus, the system will serve first the

events assigned with higher priority. Further research could be done to consider new event categories and define new priority levels. Hence, currently, we are working on improving the event engine, by defining a Simulink model that enables to represent our multi-server notification queueing system in a more realistic manner. Basically, this model consists of the following components: Time-based Generators, Priority Queues, Event-based random Generators and Servers acting as notifiers.

5. **Further validation of the privacy-aware user profile management model with real data.**

   To validate the privacy-aware user profile management model, we created synthetic data, which included quantitative and qualitative data in order to evaluate different EHR profiles and focused on studying in the AEM structure by considering issues related to its size and measuring the performance of the proposed sorting algorithm. Further research is needed to test our user profile management model in additional mobile and cloud-based e-health scenarios, using real data like images and genetic information considered attached to the tree basic information, as well as more complex information models and archetypes.

   This work line could be initiated by establishing collaborations with the Spanish Health system[4], since a Digital Clinical History of the National Health system[5] has been developed recently. Its main goal is to enable users digital access to a set of personal data by guaranteeing the security requirements established to protect such information. Currently users can consult all primary care reports and only reports of specialized care that have been generated in hospitals with digital support. The rest of hospitals and clinical reports of specialized care will be gradually incorporated.

6. **Additional validation and evaluation in real scenarios.**

   A quantitative validation in order to show the advantages obtained by the proposed event-based revocation model with respect to time-dependent revocation solutions has not been performed, because of the scarce deployment of consent management systems that avoids to get mean values (e.g., mean time of credential validity) to do a fair comparative. Although metrics are not defined yet in this area and metrics

[4] https://www.msssi.gob.es/organizacion/sns/libroSNS.htm

[5] http://www.madrid.org/cs/Satellite?cid=1142675996204&language=es&pagename=PortalSalud%2FPage%2FPTSA_pintarContenidoFinal&vest=1142675932772

used for public key certificate revocation [201] do not apply directly to our proposal, we could use metrics as credential issuing overhead, revocation list management cost, etc. These metrics can be calculated without taking into account qualitative advantages, such as credentials being expired when these would be required, for example. So, as part of the validation process, it would be necessary the implementation and deployment of a full prototype that follows the privacy-enhanced IdM architecture in real healthcare and cloud computing scenarios. In addition, metrics related to delegation chain length are considered as future line.

On the other hand, it must be noted that, the usage of the system also affects privacy and should be present in users consents. The auditing processes should verify that the design and assumptions regarding future usage match its actual usage. For future work, we want to test this last issue on real health care scenarios in order to demonstrate how the privacy is managed by the system actors.

7. **Integration of the the privacy-aware user profile management model with other specifications to handle personal information in sensitive scenarios.**

   With the aim to apply the proposed privacy-aware user profile management model to a specific structure compliant with existing standards to manage user's profiles, we carried out our validations with OpenEHR. This specification supports the creation, storage, maintenance, and querying of complete EHRs, as well as the development of an open and semantic-connected platform for eHealth systems.

   Further work is required to integrate the privacy-aware user profile management model contribution in other specifications. As a first step, it would be interesting to start the integration with the ISO/EN 13606, since it is a European norm from the CEN and it is being used at different public and private projects and deployments. In addition, another remarkable feature of ISO/EN 13606 is the alignment it presents to other relevant standards such as OpenEHR or HL7.

8. **Integration of the privacy extensions for identity management in different specifications.**

   The main goal in this thesis was to define an infrastructure generic enough to be applicable to any federation specification. Thus, we studied the principal documents of the distinct identity management technologies as a basis for the specification.

Although, whenever particularization was required to go deeper on the description of the model, we based on the SAML specifications. The developed prototypes were implemented over SAMLv2/ID-FF. So, further work is necessary to integrate and implement the proposal in the rest of the specifications.

Appendices

Appendix A

List of Acronyms

For the purposes of the present document, the following abbreviations apply:

AAPML Attribute Authority Policy Markup Language

ABE Attribute-Based Encryption

ADL Archetype Definition Language

ADT Authenticated Dictionary Structure

AOM Archetype Object Model

AP Attribute Provider

API Application Programming Interface

ARML Attribute Requirement Markup Language

CENIT Consorcio Estratégico Nacional de Investigación Técnica

CIMI Clinical Information Modelling Initiative

CoT Circle of Trust

ECP Enhanced Client Proxy

EHR Electronic Health Record

ETSI European Telecommunications Standards Institute

FIM Federated Identity Management

GDL Guideline Definition Language

HIPAA Health Insurance Portability and Accountability Act

HITSP Healthcare Information Technology Standards Panel

HL7 Health Level Seven International

HTTP HyperText Transfer Protocol

IBE Identity-Based Encryption

IdM Identity Management

IdP Identity Provider

IETF Internet Engineering Task Force

IGF Identity Governance Framework

ISO International Organization for Standardization

ITU-T ITU Telecommunication Standardization Sector

JCR Journal Citation Report

JSON JavaScript Object Notation

LDAP Lightweight Directory Access Protocol

Lasso Liberty Alliance Single Sign-On

OASIS Organization for the Advancement of Structured Information Standards

OECD Organization for Economic Cooperation and Developments

PII Personally Identifiable Information

PKI Public Key Infrastructure

PMAC Privilege Management and Access Control

QID Quasi-IDentifier Attribute

REST Representational State Transfer

RIM Reference Information Model

RM Reference Model

RP Relying Party

SAML Security Assertion Markup Language

SDO Standards Developing Organization

SLA Service Level Agreement

SLO Single LogOut

SOAP Simple Object Access Protocol

SP Service Provider

SSL Secure Sockets Layer

SSO Single Sign-On

SSTC Security Services Technical Committee

STS Security Token Service

TLS Transport Layer Security

UI User Interface

UPT U-Prove Token

URI Uniform Resource Identifier

WSS Web Services Security

XACML eXtensible Access Control Markup Language

XML eXtensible Markup Language

XRI eXtensible Resource Identifier

XSD XML Schema Definition

ZKPK Zero-Knowledge Proof of Knowledge

Bibliography

[1] OECD. The role of digital identity management in the internet economy. June 2009.

[2] The value of our digital identity. Online, 2012. `http://www.libertyglobal.com/PDF/public-policy/The-Value-of-Our-Digital-Identity.pdf`.

[3] William G. Ryan. Privacy and freedom: Alan f. westin atheneum publishers. *Business Horizons*, 10(4):106–106, 1967.

[4] Rosa Sánchez Guerrero, Florina Almenares, Patricia Arias, Daniel Díaz Sánchez, and Andrés Marín. Enhancing privacy and dynamic federation in idm for consumer cloud computing. *IEEE Trans. Consumer Electronics*, 58(1):95–103, 2012.

[5] Rosa Sánchez Guerrero, Patricia Arias, Florina Almenares, and Daniel Díaz Sánchez. Trust-aware federated idm in consumer cloud computing. *In IEEE International Conference on Consumer Electronics (ICCE)*, pages 53–54, 2012.

[6] Rosa Sánchez-Guerrero, Florina Almenárez, Daniel Díaz-Sánchez, Andrés Marín, Patricia Arias, and Fabio Sanvido. An event driven hybrid identity management approach to privacy enhanced e-health. *Sensors*, 12(5):6129–6154, 2012.

[7] Rosa Sánchez-Guerrero, Florina Almenárez, Daniel Díaz-Sánchez, Andrés Marín, and Patricia Arias. Improving privacy in identity management systems for healthcare scenarios. In *Proceedings of 5th International Symposium on Ubiquitous Computing and Ambient Intelligence*, UCAmI'11, 2011.

[8] Rosa Sánchez Guerrero, Florina Almenárez Mendoza, Daniel Díaz Sánchez, Patricia Arias, and Andrés Marín. A model for dimensioning a secure event-driven health

care system. In *5th Joint IFIP Wireless and Mobile Networking Conference, WMNC 2012, Bratislava, Slovakia, September 19-21, 2012*, pages 30–37, 2012.

[9] Rosa Sánchez-Guerrero, Daniel Díaz-Sánchez, Florina Almenárez, Andrés Marín, Patricia Arias Cabarcos, and Davide Proserpio. An identity aware wimax personalization for pervasive computing services. In *Proceedings of 5th International Symposium on Ubiquitous Computing and Ambient Intelligence*, UCAmI'11, 2011.

[10] D. Proserpio, F. Sanvido, P. A. Cabarcos, R. Sanchez Guerrero, F. Almenarez-Mendoza, D. Diaz-Sanchez, and A. Marin-Lopez. Introducing infocards in ngn to enable user-centric identity management. In *Global Telecommunications Conference (GLOBECOM 2010), 2010 IEEE*, pages 1–5, Dec 2010.

[11] Daniel Díaz-Sánchez, Florina Almenarez, Andrés Marín, Rosa Sánchez-Guerrero, and Patricia Arias. Media gateway: bringing privacy to private multimedia cloud connections. *Telecommunication Systems*, 55(2):315–330, 2014.

[12] Patricia Arias Cabarcos, Rosa Sánchez Guerrero, Florina Almenárez Mendoza, Daniel Díaz Sánchez, and Andrés Marín López. Famtv: An architecture for presence-aware personalized television. *IEEE Trans. Consumer Electronics*, 57(1):6–13, 2011.

[13] Patricia Arias, Rosa Sánchez Guerrero, Florina Almenares, and Daniel Díaz Sánchez. Presence-aware personalized television. *In IEEE International Conference on Consumer Electronics (ICCE)*, pages 769–770, 2011.

[14] D. Díaz-Sánchez, R. Sánchez Guerrero, A. Marín López, F. Almenares, and P. Arias. A h.264 svc distributed content protection system with flexible key stream generation. In *Consumer Electronics - Berlin (ICCE-Berlin), 2012 IEEE International Conference on*, pages 66–70, Sept 2012.

[15] D. Díaz-Sánchez, R. Sánchez, P. Arias, I. Bernabé Sánchez, F. Almenares, and A. Marin. Family personalization service. In *2011 IEEE International Conference on Consumer Electronics -Berlin (ICCE-Berlin)*, pages 145–149, Sept 2011.

[16] Audun Jøsang, Muhammed Al Zomai, and Suriadi Suriadi. Usability and privacy in identity management architectures. In *Proceedings of the fifth Australasian symposium on ACSW frontiers - Volume 68*, ACSW '07, pages 143–152, Darlinghurst, Australia, Australia, 2007. Australian Computer Society, Inc. ISBN 1-920-68285-X.

[17] F. Hirsch, R. Philpott, and E. Maler. Security and Privacy Considerations for the OASIS Security Assertion Markup Language (SAML) V2.0. Technical report, OASIS Standard, March 2005.

[18] J. Hughes and E. Maler. Security Assertion Markup Language (SAML) V2. 0 Technical Overview. Technical report, OASIS SSTC Working Draft, sstc-samltech-overview-2.0-draft-08, 2008.

[19] S. Cantor, J. Hodges, J. Kemp, and P. Thompson. Liberty ID-FF Architecture Overview, Liberty Alliance Project, draft-liberty-idff-arch-overview-1.2. Online, February 2017.

[20] Internet2 Shibboleth. Online, March 2017. `http://shibboleth.net/`.

[21] A. Nadalin, C. Kalin, P. Hallam-Baker, and R. Monzillo. Web Services Security: SOAP Message Security 1.0. Online, 2004.

[22] S. Cantor, J. Kemp, R. Philpott, and E. Maler. Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML)v2.0. Technical report, OASIS Standard, March 2005.

[23] S. Cantor, F. Hirsch, J. Kemp, R. Philpott, and E. Maler. Bindings for the OASIS Security Assertion Markup Language (SAML) v2.0. Technical report, OASIS Standard, March 2005.

[24] Scott Cantor. SAML V2.0 Enhanced Client or Proxy Profile Version 2.0. Technical report, OASIS Committee Specification 01, August 2013.

[25] Cross-Enterprise Security and Privacy Authorization (XSPA) Profile of Security Assertion Markup Language (SAML) for Healthcare Version 1.0. Online, February 2017.

[26] S. Cantor, J. Kemp, R. Aarts, J. Beatty, C. Cahill, X. Serret, and G. Whitehead. Liberty ID-FF Protocols and Schema Specification, Liberty Alliance Project, draft-liberty-idff-protocols-schema-1.2. Online, May 2005.

[27] Prateek Mishra. Liberty IGF Privacy Constraints Specification, Liberty Alliance Project, Version 1.0. Online, 2009.

[28] Liberty Alliance Project, Privacy preference expression languages. White report. Online, February 2017.

[29] M. Goodner and A. Nadalin. Web Services Federation Language (WS-Federation) Version 1.2. OASIS Standard. Online, May 2009.

[30] A. Nadalin, M. Goodner, M. Gudgin, A. Barbir, and H. Granqvist. WS-Trust 1.3. OASIS Standard. Online, 2009.

[31] Web Services Policy 1.5 - Framework. Technical report. Technical report, World Wide Web Consortium (W3C), 2007.

[32] WS-MetadataExchange.: Web Services Metadata Exchange. Online, August 2006.

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

[34] Patricia Arias Cabarcos, Florina Almenárez Mendoza, Andrés Marín-López, and Daniel Díaz-Sánchez. Enabling saml for dynamic identity federation management. In Jozef Wozniak, Jerzy Konorski, Ryszard Katulski, and Andrzej Pach, editors, *Wireless and Mobile Networking*, volume 308 of *IFIP Advances in Information and Communication Technology*, pages 173–184. Springer Boston, 2009. ISBN 978-3-642-03840-2. 10.1007/978-3-642-03841-9_16.

[35] Health Relationship Trust Profile for OpenID Connect 1.0. Online, March 2017. `http://openid.net/specs/openid-heart-openid-connect-2015-12-07.html`.

[36] D. Hardt. The OAuth 2.0 Authorization Protocol, IETF Network Working Group, draft-ietf-oauth-v2-31, July 2016.

[37] Sakimura N. Bradley J. Jones, M. JSON Web Token (JWT). Online, May 2015. `https://tools.ietf.org/html/rfc7519`.

[38] P.J Connolly. OAuth is the "hottest thing" in identity management. *eWeek*, 27(9): 12–13, May 2010.

[39] Google Inc.: Google OAuth 2.0 Client-side. Online, March 2017. `https://developers.google.com/identity/protocols/OAuth2UserAgent`.

[40] Wanpeng Li and Chris J. Mitchell. Analysing the security of google's implementation of openid connect. In *Proceedings of the 13th International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment - Volume 9721*, DIMVA 2016, pages 357–376, New York, NY, USA, 2016. Springer-Verlag New York, Inc. ISBN 978-3-319-40666-4.

[41] Google Inc.: Google OpenID Connect 1.0. Online, March 2017. `https://developers.google.com/accounts/docs/OpenIDConnect`.

[42] Information Cards.: Information Cards Foundation. Online, June 2016. `http://informationcard.net/`.

[43] Arun Nanda and Michael B. Jones. Identity Selector Interoperability Profile V1.5, July 2008.

[44] M.B. Jones and M. McIntosh. Identity Metasystem.: Interoperability Version 1.0. Technical report, OASIS Standard, 2009.

[45] Vittorio Bertocci, Garrett Serack, and Caleb Baker. *Understanding Windows Cardspace: An Introduction to the Concepts and Challenges of Digital Identities*. Addison-Wesley Professional, first edition, 2007. ISBN 9780321496843.

[46] Higgins Project.: Higgins project website. Online, 2009. `http://www.eclipse.org/higgins/`.

[47] Open Source Identity Systems.: Open Source Identity Systems Wiki. Online, June 2016. `http://osis.idcommons.net/`.

[48] Christian Paquin. U-Prove technology overview v1.1. revision 2. Online, April 2013. `https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/U-Prove20Technology20Overview20V1.120Revision202.pdf`.

[49] Christian Paquin and Greg Zaverucha. U-Prove Cryptographic Specification V1.1 (Revision 3), December 2013.

[50] Microsoft Connect. Online, March 2016. `https://connect.microsoft.com/`.

[51] Kim Cameron. The Laws of Identity. Online, November 2005. `http://www.identityblog.com/stories/2005/05/13/TheLawsOfIdentity.pdf`.

[52] Christian Paquin. U-Prove ws-trust profile v1.0. Online, February 2011. https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/U-Prove20WS-Trust20Profile20V1.0.pdf.

[53] Phillip Windley. *Digital Identity*. O'Reilly Media, Inc., 2005. ISBN 0596008783.

[54] US Government Accountability Office (GAO). Privacy: Alternatives Exist for Enhancing Protection of Personally Identifiable Information. Online, 2008. http://www.gao.gov/new.items/d08536.pdf.

[55] Tim Grance Erika McCallister and Karen Scarfone. Guide to Protecting the Confidentiality of Personally Identifiable Information (PII). Technical report, NIST Special Publication 800-122. Recommendations of NIST, 2010.

[56] Grigorios Loukides, Joshua C. Denny, and Bradley Malin. The disclosure of diagnosis codes can breach research participants' privacy. *JAMIA*, 17(3):322–327, 2010.

[57] Latanya Sweeney. K-anonymity: A model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):557–570, October 2002.

[58] M. Ercan Nergiz and Chris Clifton. Thoughts on k-anonymization. *Data Knowl. Eng.*, 63(3):622–645, December 2007.

[59] Mehmet Ercan Nergiz and Christopher W. Clifton. d-presence without complete world knowledge. *IEEE Trans. Knowl. Data Eng.*, 22(6):868–883, 2010.

[60] Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkitasubramaniam. L-diversity: Privacy beyond k-anonymity. *ACM Trans. Knowl. Discov. Data*, 1(1), March 2007.

[61] Daniel J. Solove. A Taxonomy of Privacy. *University of Pennsylvania Law Review*, 154(3):477–560, Januar 2006.

[62] Andreas Pfitzmann and Marit Köhntopp. Anonymity, unobservability, and pseudeonymity; a proposal for terminology. In *International workshop on Designing privacy enhancing technologies*, pages 1–9. Springer-Verlag New York, Inc., 2001. ISBN 3-540-41724-9.

[63] Marit Hansen Andreas Pfitzmann. A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management, August 2010.

[64] Meilof Veeningen, Benne De Weger, and Nicola Zannone. Modeling identity-related properties and their privacy strength. In *Proceedings of the 7th International conference on Formal aspects of security and trust*, FAST'10, pages 126–140, Berlin, Heidelberg, 2011. Springer-Verlag. ISBN 978-3-642-19750-5.

[65] Ken Mano, Yoshinobu Kawabe, Hideki Sakurada, and Yasuyuki Tsukada. Role interchange for anonymity and privacy of voting. *J. Log. and Comput.*, 20(6):1251–1288, December 2010.

[66] Common criteria for information technology security evolution, part 2, version 2.1, August 1999.

[67] Latanya Sweeney. Achieving k-anonymity privacy protection using generalization and suppression. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):571–588, October 2002.

[68] Adam Meyerson and Ryan Williams. On the complexity of optimal k-anonymity. In *Proceedings of the Twenty-third ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '04, pages 223–228, New York, NY, USA, 2004. ACM. ISBN 158113858X.

[69] Bradley Malin, Grigorios Loukides, Kathleen Benitez, and Ellen Wright Clayton. Identifiability in biobanks: models, measures, and mitigation strategies. *Human Genetics*, 130(3):383, 2011.

[70] Roberto J. Bayardo and Rakesh Agrawal. Data privacy through optimal k-anonymization. In *Proceedings of the 21st International Conference on Data Engineering*, ICDE '05, pages 217–228, Washington, DC, USA, 2005. IEEE Computer Society. ISBN 0-7695-2285-8.

[71] Kristen LeFevre, David J. DeWitt, and Raghu Ramakrishnan. Incognito: Efficient full-domain k-anonymity. In *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, SIGMOD '05, pages 49–60, New York, NY, USA, 2005. ACM. ISBN 1-59593-060-4.

[72] Yeye He and Jeffrey F. Naughton. Anonymization of set-valued data via top-down, local generalization. *Proc. VLDB Endow.*, 2(1):934–945, August 2009.

[73] Raymond Chi-Wing Wong, Jiuyong Li, Ada Wai-Chee Fu, and Ke Wang. (α, k)-anonymity: An enhanced k-anonymity model for privacy preserving data publishing. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '06, pages 754–759, New York, NY, USA, 2006. ACM. ISBN 1-59593-339-5.

[74] T. M. Truta and B. Vinay. Privacy protection: p-sensitive k-anonymity property. In *22nd International Conference on Data Engineering Workshops (ICDEW'06)*, pages 94–94, 2006.

[75] Aris Gkoulalas-Divanis, Grigorios Loukides, and Jimeng Sun. Publishing data from electronic health records while preserving privacy: A survey of algorithms. *Journal of Biomedical Informatics*, 50:4 – 19, 2014. Special Issue on Informatics Methods in Medical Privacy.

[76] N. Li, T. Li, and S. Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. In *2007 IEEE 23rd International Conference on Data Engineering*, pages 106–115, April 2007.

[77] N. Li, T. Li, and S. Venkatasubramanian. Closeness: A new privacy measure for data publishing. *IEEE Transactions on Knowledge and Data Engineering*, 22(7): 943–956, July 2010.

[78] D. J. Martin, D. Kifer, A. Machanavajjhala, J. Gehrke, and J. Y. Halpern. Worst-case background knowledge for privacy-preserving data publishing. In *2007 IEEE 23rd International Conference on Data Engineering*, pages 126–135, April 2007.

[79] Cynthia Dwork. *Differential Privacy*, pages 1–12. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006. ISBN 978-3-540-35908-1.

[80] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Conference on Theory of Cryptography*, TCC'06, pages 265–284, Berlin, Heidelberg, 2006. Springer-Verlag. ISBN 3-540-32731-2, 978-3-540-32731-8.

[81] Ashwin Machanavajjhala, Johannes Gehrke, and Michaela Götz. Data publishing against realistic adversaries. *Proc. VLDB Endow.*, 2(1):790–801, August 2009.

[82] Vijay S. Iyengar. Transforming data to satisfy privacy constraints. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '02, pages 279–288, New York, NY, USA, 2002. ACM. ISBN 1-58113-567-X.

[83] H. C. A. van Tilborg. Encyclopedia of Cryptography and Security. Technical report, Springer-Verlag New York, Inc., 2005.

[84] Thomas Page. *The application of hash chains and hash structures to cryptography.* PhD thesis, Egham, Surrey TW20 0EX, England, August 2009.

[85] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone. Handbook of Applied Cryptography. CRC, October 1996.

[86] R. C. Merkle. Secrecy. Authentication, and Public Key Systems. PhD thesis, Stanford University, 1979.

[87] William Robert Speirs, II. *Dynamic Cryptographic Hash Functions.* PhD thesis, West Lafayette, IN, USA, 2007. AAI3278689.

[88] Wenbo Mao. *Modern cryptography: theory and practice.* Prentice Hall Professional Technical Reference, 2003.

[89] Douglas R. Stinson. *Modern cryptography: theory and practice.* Chapman & Hall/CRC, 2002.

[90] Leo J. Guibas and Robert Sedgewick. A dichromatic framework for balanced trees. *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, 0:8–21, 1978.

[91] Hiroaki Kikuchi, Kensuke Abe, and Shohachiro Nakanishi. Online certification status verification with a red-black hash tree. *IPSJ Digital Courier*, 2:513–523, 2006.

[92] Charalampos Papamanthou, Roberto Tamassia, and Nikos Triandopoulos. *Pairing-Based Cryptography - Pairing 2010: 4th International Conference, Yamanaka Hot Spring, Japan, December 2010. Proceedings*, chapter Optimal Authenticated Data

Structures with Multilinear Forms, pages 246–264. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. ISBN 978-3-642-17455-1.

[93] Sheng-Cheng Yeh, Ming-Yang Su, Hui-Hui Chen, and Chun-Yuen Lin. An efficient and secure approach for a cloud collaborative editing. *Journal of Network and Computer Applications*, 36(6):1632 – 1641, 2013.

[94] Moni Naor and Kobbi Nissim. Certificate revocation and certificate update. In *Proc. 7th USENIX Security Symposium*, pages 217–228, Berkeley, 1998.

[95] Ralph C. Merkle. A certified digital signature. In *Proceedings of the 9th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '89, pages 218–238, London, UK, 1990. Springer-Verlag. ISBN 3-540-97317-6.

[96] Johannes Buchmann, Luis García, Erik Dahmen, Martin Döring, and Elena Klintsevich. CMSS - An improved merkle signature scheme. In Rana Barua and Tanja Lange, editors, *Progress in Cryptology - INDOCRYPT*, volume 4329 of *Lecture Notes in Computer Science*, pages 349–363. Springer, 2006. ISBN 978-3-540-49767-7.

[97] Johannes Buchmann, Erik Dahmen, Elena Klintsevich, Katsuyuki Okeya, and Camille Vuillaume. Merkle signatures with virtually unlimited signature capacity. In *Applied Cryptography and Network Security*, volume 4521 of *Lecture Notes in Computer Science*, pages 31–45. Springer, 2007. ISBN 978-3-540-72737-8.

[98] Cong Wang, Kui Ren, Wenjing Lou, and Jin Li. Toward publicly auditable secure cloud data storage services. *IEEE network*, 24(4):19–24, 2010.

[99] H. Li, R. Lu, L. Zhou, B. Yang, and X. Shen. An efficient merkle-tree-based authentication scheme for smart grid. *IEEE Systems Journal*, 8(2):655–663, June 2014.

[100] Hui Lin, Jianfeng Ma, Jia Hu, and Kai Yang. Pa-shwmp: a privacy-aware secure hybrid wireless mesh protocol for ieee 802.11s wireless mesh networks. *EURASIP Journal on Wireless Communications and Networking*, 2012(1):1–16, 2012.

[101] Rémi Bazin, Alexander Schaub, Omar Hasan, and Lionel Brunie. A decentralized anonymity-preserving reputation system with constant-time score retrieval. *IACR Cryptology ePrint Archive*, 2016:416, 2016.

[102] Elli Androulaki, Ghassan O. Karame, Marc Roeschlin, Tobias Scherer, and Srdjan Capkun. *Evaluating User Privacy in Bitcoin*, pages 34–51. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. ISBN 978-3-642-39884-1.

[103] Przemyslaw Kubiak, Miroslaw Kutylowski, and Jun Shao. How to construct state registries-matching undeniability with public security. In *Intelligent Information and Database Systems, Second International Conference, ACIIDS, Hue City, Vietnam, March 24-26, 2010. Proceedings, Part I*, pages 64–73, 2010.

[104] Thomas Page. The application of hash chains and hash structures to cryptography, technical report, August 2009.

[105] William Pugh. Skip lists: a probabilistic alternative to balanced trees. *Commun. ACM*, 33(6):668–676, June 1990.

[106] J. Ian Munro, Thomas Papadakis, and Robert Sedgewick. Deterministic skip lists. In *Proceedings of the Third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '92, pages 367–375, Philadelphia, PA, USA, 1992. Society for Industrial and Applied Mathematics. ISBN 0-89791-466-X.

[107] Adi Shamir. Identity-based cryptosystems and signature schemes. In *Proceedings of CRYPTO 84 on Advances in Cryptology*, pages 47–53, New York, NY, USA, 1985. Springer-Verlag New York, Inc. ISBN 0-387-15658-5.

[108] Amit Sahai and Brent Waters. *Advances in Cryptology – EUROCRYPT 2005: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005. Proceedings*, chapter Fuzzy Identity-Based Encryption, pages 457–473. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005. ISBN 978-3-540-32055-5.

[109] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the weil pairing. In *Proceedings of the 21st Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '01, pages 213–229, London, UK, UK, 2001. Springer-Verlag. ISBN 3-540-42456-3.

[110] Ran Canetti, Shai Halevi, and Jonathan Katz. *A Forward-Secure Public-Key Encryption Scheme*, pages 255–271. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003. ISBN 978-3-540-39200-2.

[111] Dan Boneh and Xavier Boyen. Efficient selective-id secure identity based encryption without random oracles. In *Proceedings of Eurocrypt 2004, volume 3027 of LNCS*, pages 223–238. Springer-Verlag, 2004.

[112] J. Bethencourt, A. Sahai, and B. Waters. Ciphertext-policy attribute-based encryption. In *In IEEE Symposium on Security and Privacy*, page 321, 2007.

[113] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *Proceedings of the 13th ACM Conference on Computer and Communications Security*, CCS '06, pages 89–98, New York, NY, USA, 2006. ACM. ISBN 1-59593-518-5.

[114] Patrick Gunn, Allen Fremont, Melissa Bottrell, Lisa Shugarman, Jolene Galegher, and Tora Bikson. The health insurance portability and accountability act privacy rule: a practical guide for researchers. *Medical Care*, 42(4):321–327, 2004.

[115] Health Level Seven International.: Health Level Seven (HL7). Online, June 2016. `http://www.hl7.org`.

[116] R. Chen. EHR Information Model. Release 1.0.2. Technical report, openEHR, August 2008. `http://openehr.org/releases/1.0.2/architecture/rm/ehr_im.pdf`.

[117] ISOEN-13606. Online, September 2016. `http://www.iso.org/iso/home.htm`.

[118] "IHE IT infrastructure (ITI) technical framework volume 1 integration profiles". revision 12.0, September 2016.

[119] OpenEHR.: An open domain-driven platform for developing flexible e-health systems. Online, March 2017. `http://www.openehr.org/`.

[120] R. Chen. ADL.: Archetype Definition Language 1.4 specification. Technical report, openEHR, December 2008. `http://www.openehr.org/releases/AM/latest/docs/ADL1.4/ADL1.4.html`.

[121] C. Ma, R. Chen, T. Cook, and S. Heard. AOM.: The openEHR Archetype Model Archetype Object Model. Technical report, openEHR, November 2008. `http://openehr.org/releases/1.0.2/architecture/am/aom.pdf`.

[122] Koray Atalag, Thomas Beale, Rong Chen, Tomaz Gornik, Sam Heard, and Ian McNicoll. openehr - a semantically-enabled health computing platform, March 2017.

[123] ISO 22600-1:2014 Health informatics – Privilege management and access control – Part 1: Overview and policy management. Online, June 2016. `http://www.iso.org/iso/home/store/catalogue_tc/catalogue_detail.htm?csnumber=62653`.

[124] ISO/IEC 17799:2000 Information technology – Code of practice for information security management. Online, June 2016. `http://www.iso.org/iso/home/store/catalogue_ics/catalogue_detail_ics.htm?csnumber=33441`.

[125] ISO/EN13606 Electronic Healthcare Record (EHR) Communication. Parts 1: Reference Model, Part 2: Archetype Model, Part 3: Reference Archetypes and Term lists, Part 4: Security and Part 5: Interface Specification. Online, January 2017. `http://www.iso.org/iso/catalogue_detail.htm?csnumber=40784`.

[126] T. Beale. OpenEHR Architecture Overview. Technical report, OpenEHR Foundation, January 2017. `http://www.openehr.org/releases/1.0.1/architecture/overview.pdf`.

[127] Joe Looby. Discovering europe: How to navigate europe's privacy protections, December 2010. National Law Journal.

[128] EU Directive 2002/58 on Privacy and Electronic Communications (E-Privacy Directive), December 2009. amended by Directive 2009/136.

[129] Liam Curren and Jane Kaye. Revoking consent: A "blind spot" in data protection law? *Computer Law Security Review*, 26(3):273–283, 2010.

[130] Lillian Røstad and Ole Edsberg. A study of access control requirements for healthcare systems based on audit trails from access logs. In *In Proceedings of the 2006 Annual Computer Security Applications Conference, Miami Beach*, 2006.

[131] Le Xun Hung, Sungyoung Lee, Young-Koo Lee, and Heejo Lee. Activity-based access control model to hospital information. In *Embedded and Real-Time Computing Systems and Applications, 2007. RTCSA 2007. 13th IEEE International Conference on*, pages 488 –496, aug. 2007.

[132] Rafae Bhatti and Tyrone Grandison. Towards improved privacy policy coverage in healthcare using policy refinement. In *Secure Data Management, 4th VLDB Workshop, SDM 2007, Vienna, Austria, September 23-24, 2007, Proceedings*, pages 158–173, 2007.

[133] A Ferreira, R Cruz-Correia, L Antunes, P Farinha, E Oliveira-Palhares, D W. Chadwick, and A Costa-Pereira. How to break access control in a controlled manner. In *Proceedings of the 19th IEEE Symposium on Computer-Based Medical Systems*, CBMS '06, pages 847–854, Washington, DC, USA, 2006. IEEE Computer Society. ISBN 0-7695-2517-1.

[134] Qun Ni, Alberto Trombetta, Elisa Bertino, and Jorge Lobo. Privacy-aware role based access control. In *Proceedings of the 12th ACM symposium on Access control models and technologies*, SACMAT '07, pages 41–50, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-745-2.

[135] Edgar A. Whitley. Informational privacy, consent and the "control" of personal data. *Inf. Secur. Tech. Rep.*, 14:154–159, August 2009.

[136] Encore (*Ensuring Consent & Revocation*), european community's seventh framework programme [fp7/2007-2013]. Online, January 2017.

[137] I. Agrafiotis, S. Creese, M. Goldsmith, N. Papanikolaou, M Casassa Mont, and S. Pearson. Defining consent and revocation policies. In *Proceedings of 2010 IFIP/PrimeLife Summer School*, 2010.

[138] Casassa Mont, Marco, Pearson, Siani, Creese, Sadie, Goldsmith, Michael, Papanikolaou, and Nick. A conceptual model for privacy policies with consent and revocation requirements. In Simone Fischer-Hübner, Penny Duquenoy, Marit Hansen, Ronald Leenes, and Ge Zhang, editors, *Privacy and Identity Management for Life*, volume 352 of *IFIP Advances in Information and Communication Technology*, pages 258–270. Springer Boston, 2011. ISBN 978-3-642-20768-6. 10.1007/978-3-642-20769-3_21.

[139] F. Kargl, E. Lawrence, M. Fischer, and Yen Yang Lim. Security, privacy and legal issues in pervasive ehealth monitoring systems. In *Mobile Business, 2008. ICMB '08. 7th International Conference on*, pages 296–304, July 2008.

[140] Moshaddique Ameen, Jingwei Liu, and Kyungsup Kwak. Security and privacy issues in wireless sensor networks for healthcare applications. *J. Med. Syst.*, 36(1):93–101, February 2012.

[141] Jun Zhou, Zhenfu Cao, Xiaolei Dong, Xiaodong Lin, and A.V. Vasilakos. Securing m-healthcare social networks: challenges, countermeasures and future directions. *Wireless Communications, IEEE*, 20(4):12–21, August 2013.

[142] Lifei Wei, Haojin Zhu, Zhenfu Cao, Xiaolei Dong, Weiwei Jia, Yunlu Chen, and Athanasios V. Vasilakos. Security and privacy for storage and computation in cloud computing. *Information Sciences*, 258:371 – 386, 2014.

[143] Min Yoon Miyoung Jang and Jae-Woo Chang. A new query integrity verification method with cluster-based data transformation in cloud computing environment. *International Journal of Smart Home*, 9(4):225 – 238, 2015.

[144] Member Sherman S. M. Chow Yue Tong, Jinyuan Sun and Pan Li. Cloud-assisted mobile-access of health data with privacy and auditability. *IEEE Journal of Biomedical and Health In- formatics (JBHI)*, 18(2):419–429, 2014.

[145] Ashish Kundu and Elisa Bertino. Structural signatures for tree data structures. *Proc. VLDB Endow.*, pages 138–150, August 2008.

[146] Shin Dongwan, R. Lopes, and W. Claycomb. Authenticated dictionary-based attribute sharing in federated identity management. In *6th International Conference on Information Technology: New Generations (ITNG '09)*, pages 504 –509, april 2009.

[147] Layla Pournajaf, Li Xiong, Daniel Garcia-Ulloa, , and Vaidy Sunderam. Survey on privacy in mobile crowd sensing task management. Technical report, Technical Report TR-2014-002, Department of Mathematics and Computer Science, Emory University, 2014.

[148] Huiling Qian, Jiguo Li, Yichen Zhang, and Jinguang Han. Privacy-preserving personal health record using multi-authority attribute-based encryption with revocation. *International Journal of Information Security*, 14(6):487–497, 2015.

[149] Bo Qin, Hua Deng, Qianhong Wu, Josep Domingo-Ferrer, David Naccache, and Yunya Zhou. Flexible attribute-based encryption applicable to secure e-healthcare records. *International Journal of Information Security*, 14(6):499–511, 2015.

[150] Waters B. Lewko, A. Why proving hibe systems secure is difficult. *EUROCRYPT'14, LNCS*, 8441:58–76, 2014.

[151] Stuart Haber, Yasuo Hatano, Yoshinori Honda, William Horne, Kunihiko Miyazaki, Tomas Sander, Satoru Tezoku, and Danfeng Yao. Efficient signature schemes supporting redaction, pseudonymization, and data deidentification. In *Proceedings of the 2008 ACM Symposium on Information, Computer and Communications Security*, ASIACCS '08, pages 353–362, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-979-1.

[152] Gabriel Ghinita. Privacy for Location-based Services Synthesis. Technical report, Lectures on Information Security, Privacy, and Trust, University of Massachusetts, Boston, April 2013.

[153] Xiaohui Liang, Rongxing Lu, Le Chen, Xiaodong Lin, and Xuemin Shen. Pec: A privacy-preserving emergency call scheme for mobile healthcare social networks. *Communications and Networks, Journal of*, 13(2):102–112, April 2011.

[154] D. Mashima, D. Bauer, M. Ahamad, and D.M. Blough. User-centric identity management architecture using credential-holding identity agents. *Digital Identity and Access Management: Technologies and Frameworks, IGI Global*, December 2012.

[155] F. Paci, R. Ferrini, A. Musci, K. Steuer, and E. Bertino. An interoperable approach to multifactor identity verification. *IEEE Computer*, 42(5):50–57, 2009.

[156] Rongxing Lu, Xiaodong Lin, and Xuemin Shen. Spoc: A secure and privacy-preserving opportunistic computing framework for mobile-healthcare emergency. *Parallel and Distributed Systems, IEEE Transactions on*, 24(3):614–624, March 2013.

[157] ABC4Trust EU Project (2015). Attribute based Credentials for Trust. Online, December 2016. `https://abc4trust.eu/index.php/home/fact-sheet`.

[158] STORK 2.0 Project (2012). Secure idenTity acrOss boRders linKed 2.0. Online, June 2016. `https://www.eid-stork2.eu/`.

[159] FutureID Project (2012). Shaping the Future of Electronic Identity. Online, June 2016. http://www.futureid.eu/.

[160] Endorse Project (2013). Legal Technical Framework for Privacy Preserving Data Management. Online, December 2016. http://ict-endorse.eu/.

[161] EnCoRe Project (2013). Ensuring Consent and Revocation. Online, December 2016. http://www.hpl.hp.com/breweb/encoreproject/index.html.

[162] PICOS Project (2011). Privacy and Identity Management for Community Services. Online, December 2016. http://www.picos-project.eu/.

[163] Project (2011). Privacy and Identity Management in Europe for Life. Online, December 2016. http://www.primelife.eu/.

[164] SWIFT Project (2010). Secure Widespread Identities for Federated Telecommunications. Online, December 2016. http://www.ist-swift.org/.

[165] Health Relationship Trust Profile for OAuth 2.0. Online, February 2016. http://openid.net/specs/openid-heart-oauth2-1_0-ID1.html/.

[166] Health Relationship Trust Profile for OpenID Connect 1.0. Online, February 2016. http://openid.net/specs/openid-heart-openid-connect-1_0-ID1.html/.

[167] Health Relationship Trust Profile for User Managed Access 1.0. Online, February 2016. http://openid.net/specs/openid-heart-uma-1_0-ID1.html/.

[168] Silvia Puglisi, David Rebollo-Monedero, and Jordi Forné. On web user tracking of browsing patterns for personalised advertising. *International Journal of Parallel, Emergent and Distributed Systems*, pages 1–20, 2017.

[169] Birgit Pitzman. Privacy in enterprise identity federation: Policies for Liberty 2 single sign on, March 2004.

[170] D. Recordon, M. Jones, J. Bufu, J. Daugherty, and N. Sakimura. OpenID Provider Authentication Policy Extension 1.0. Online, February 2017. http://www.openid.net.

[171] Alain Mouttham, Liam Peyton, Ben Eze, and Abdulmotaleb Saddik. Event-driven data integration for personal health monitoring. *Journal of Emerging Technologies in Web Intelligence*, 1(2), 2009.

[172] A. Baarah, A. Mouttham, and L. Peyton. Improving cardiac patient flow based on complex event processing. In *Applied Electrical Engineering and Computing Technologies (AEECT), 2011 IEEE Jordan Conference on*, pages 1 –6, dec. 2011.

[173] A. B. Roach. Session Initiation Protocol (SIP)-Specific Event Notification. RFC 3265 (Proposed Standard), June 2002. Updated by RFCs 5367, 5727, 6446.

[174] Melissa L. McCarthy, Scott L. Zeger, Ru Ding, Dominik Aronsky, Nathan R. Hoot, and Gabor D. Kelen. The challenge of predicting demand for emergency department services. *Academic Emergency Medicine*, 15(4):337–346, 2008.

[175] Niels K. Rathlev, John Chessare, Jonathan Olshaker, Dan Obendorfer, Supriya D. Mehta, Todd Rothenhaus, Steven Crespo, Brendan Magauran, Kathy Davidson, Richard Shemin, Keith Lewis, James M. Becker, Linda Fisher, Linda Guy, Abbott Cooper, and Eugene Litvak. Time series analysis of variables associated with daily mean emergency department length of stay. *Annals of Emergency Medicine*, 49(3): 265 – 271, 2007.

[176] Michael J Schull, Alex Kiss, and John-Paul Szalai. The effect of low-complexity patients on emergency department waiting times. *Annals of Emergency Medicine*, 49(3):257–264, 264.e1, 2007.

[177] *Modeling and Analysis of Computer Communications Networks*. Khanna Publishers, New Delhi, 1984.

[178] Roel Peeters, Koen Simoens, Danny De Cock, and Bart Preneel. Cross-context delegation through identity federation. In *BIOSIG 2008 - Proceedings of the Special Interest Group on Biometrics and Electronic Signatures, 11.-12. September 2008 in Darmstadt, Germany*, pages 79–92, 2008.

[179] M. Ahsant, J. Basney, and O. Mulmo. Grid delegation protocol. technical report. Technical report, YCS-2004-380, University of York, Department of Computer Science, July 2004.

[180] Hidehito Gomi, Makoto Hatakeyama, Shigeru Hosono, and Satoru Fujita. A delegation framework for federated identity management. In *Proceedings of the 2005 Workshop on Digital Identity Management*, DIM '05, pages 94–103, New York, NY, USA, 2005. ACM. ISBN 1-59593-232-1.

[181] Scott Cantor. SAML V2.0 Condition for Delegation Restriction. Committee Draft 01. Online, March 2009.

[182] R C Sato and D M Zouain. Markov models in health care. *Einstein Sao Paulo*, 8: 376–379, 2010.

[183] Donald Gross, John F. Shortle, James M. Thompson, and Carl M. Harris. Wiley-Blackwell (an imprint of John Wiley & Sons Ltd), 4th edition edition, 2008.

[184] Bo Zhao and Chao Liu. Efficient sip-specific event notification. In *Networking, International Conference on Systems and International Conference on Mobile Communications and Learning Technologies, 2006. ICN/ICONS/MCL 2006. International Conference on*, page 1, april 2006.

[185] MathWorks, Using MATLAB Manual. Online, December 2016. `http://www.mathworks.com/`.

[186] Lasso: Liberty Alliance Single Sign-On. Online, June 2016. `http://lasso.entrouvert.org/`.

[187] Authentic: Liberty-compliant Identity Provider. Online, June 2016. `http://authentic.labs.libre-entreprise.org/`.

[188] SymLabs.: ZXID: Open SAML implementation in C. Online, June 2016. `http://www.zxid.org/`.

[189] Project Sailfin: Open source Java application server project. Online, June 2016. `http://sailfin.java.net/`.

[190] E.6.1.5 Informe de avance en la investigación en Seguridad e Identidad. Entregable del Proyecto CENIT España Virtual, 2010.

[191] Informe de avance en el activo experimental "Georreferenciación semántica del conocimiento". Entregable del Proyecto CENIT España Virtual, 2010.

[192] Rosa Sánchez Guerrero, Daniel Díaz Sánchez, Florina Almenárez Mendoza, Patricia Arias, Davide Proserpio, and Andrés Marín. Introducing identity management in wimax to enable secure and personalized services. In *4th Joint IFIP Wireless and Mobile Networking Conference, WMNC 2011, Toulouse, France, 26-28 October, 2011*, pages 1–5, 2011.

[193] Accident and Emergency Hospital Episode Statistics (HES). Online, December 2016. http://www.hscic.gov.uk/searchcatalogue?productid=17200&q=title%3a%22accident+and+emergency+attendances%22&topics=0%2fHospital+care&sort=Relevance&size=10&page=1#top.

[194] Roberto Pinto, Tobias Mettler, and Marco Taisch. Managing supplier delivery reliability risk under limited information: Foundations for a human-in-the-loop DSS. *Decision Support Systems*, 54(2):1076–1084, 2013.

[195] Microsoft's Vision for an Identity Metasystem. Technical report, Whitepaper, 2005.

[196] Emilie Lundin Barse, Håkan Kvarnström, and Erland Jonsson. Synthesizing test data for fraud detection systems. In *19th Annual Computer Security Applications Conference (ACSAC 2003), 8-12 December 2003, Las Vegas, NV, USA*, pages 384–394, 2003.

[197] Linda Moniz, Anna L. Buczak, Lang Hung, Steven Babin, Michael Dorko, and Joseph Lombardo. Construction and Validation of Synthetic Electronic Medical Records. *Online Journal of Public Health Informatics*, 1, 2009.

[198] Bhume Bhumiratana and Matt Bishop. Privacy aware data sharing: balancing the usability and privacy of datasets. In *PETRA*, ACM International Conference Proceeding Series. ACM, 2009. ISBN 978-1-60558-409-6.

[199] Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy*, SP '08, pages 111–125, Washington, DC, USA, 2008. IEEE Computer Society. ISBN 978-0-7695-3168-7.

[200] Editors O. Mcgraw-Hill. *The McGraw-Hill Dictionary of Scientific and Technical Terms, Seventh Edition (Mcgraw Hill Dictionary of Scientific and Technical Terms)*. McGraw-Hill Professional, 7 edition, October 2009. ISBN 0071608990.

[201] Girma Enideg Nigusse. Master's thesis, evaluating public key certificate revocation schemes: Towards conceptually versatile revocation scheme. Online, August 2007. http://es.scribd.com/doc/29894385/GirmaNigusse.