THEME ARTICLE

Towards a cross-context identity management framework in e-health

Mina Deng, Danny De Cock and Bart Preneel
IBBT-COSIC, Katholieke Universiteit Leuven, Leuven-Heverlee, Belgium

Abstract

Purpose – Modern e-health systems incorporate different healthcare providers in one system and provide an electronic platform to share medical information efficiently. In cross-context communications between healthcare providers, the same information can be interpreted as different types or values, so that one patient will be issued different identifiers by different healthcare providers. This paper aims to provide a solution to ensure interoperability so that multiple healthcare providers will be able to collaborate in one e-health system.

Design/methodology/approach – This paper primarily focuses on how different healthcare providers, instead of the patients, are able to interact and share information on a common e-health platform.

Findings – In the course of the work, it was found that previous e-health solutions mainly have a limited view of patient information, where a user-centric approach for identity management is usually restricted to a single healthcare provider. Interoperability in an e-health system becomes more problematic when more actors collaborate, and hence linkability from one context to another should not be straightforward. However, some form of linkability, such as the possibility to follow up a patient's medical treatment, is desirable in the e-health sector, even when it needs to cross different contexts. Therefore, the authors have designed an identity management mechanism to ensure semantic interoperability when data is exchanged among different authorized healthcare providers.

Research limitations/implications – The paper points out that the next generation of e-health will move towards federated e-health and will require user-centricity and transparency properties so that patients are able to specify and verify the disclosure of their medical information.

Originality/value – This paper proposes a new service for cross-context identity management in e-health systems, improving interoperability between agencies when context-specific information is transferred from one healthcare provider to another. How the proposed cross-context identity management service can be integrated in an e-health system is explained with a use case scenario.

Keywords Health services, Medical information systems, Data security, Internet

Paper type Research paper

Introduction

In recent years, the technological evolution of e-health systems has drawn increasing attention from both industry and research communities worldwide. The goals of e-health systems are fourfold:

This work was supported by the Concerted Research Action (GOA) AMBioRICS 2005/11 of the Flemish Government and by the IAP Programme P6/26 BCRYPT of the Belgian State (Belgian Science Policy). The authors were supported by the Interdisciplinary Institute for BroadBand Technology, Belgium.

-
- (1) Primarily, to provide ubiquitous access to lifelong clinical records of a patient to all relevant stakeholders (including the patient), anytime, anywhere, on any device.
 - (2) To integrate and enrich clinical, medical and operational knowledge in order to support lifelong health guidance of citizens within a community, region and country.
 - (3) To streamline the workflow into shared clinical and operational pathways in order to enable disease management and optimally support the clinical process.
 - (4) By combining the three preceding goals, to facilitate inter-professional collaboration, while guaranteeing the privacy of the patient.

The major technical challenges facing e-health services are facilitating efficiency, information retrieval and availability, and cross-context interoperability. The rapid aging of populations combined with pressure on budgets for healthcare delivery and technological advances are the driving forces behind these challenges. Hence, in the realm of e-health, security and privacy issues have a deep impact. In this paper, “privacy” means to protect the private information of certain entities. Security techniques, such as access control structures, are adopted in e-health systems to ensure that only involved parties have access to sensitive data. The circle of trusted parties should not be extended by moving from a paper-world to an e-health administration. A patient expects a trust relation with medics; however, as in the past with a doctor’s secretary, trust in the digital world might lead to the need to trust a system administrator too. The adoption of user-centric federated identity management systems can help to keep the circle of trust as small as possible.

In traditional identity management (IDM) systems identity information is hosted and managed by an identity provider using, for example, directory services (LDAP, 2006). This has various advantages from the service provider’s point of view, such as being cost effective and easily scalable. The disadvantage is that by applying such an approach, the user loses some control over personal identity information. Recently, this has resulted in the emerging trend for user-centric identity management (Josang and Pope, 2005). In user-centric IDM, the user (e.g. the patient in an e-health application) is put in the centre of interest and is given control over personal information. In particular, this means that the user can influence or even specify which information should be forwarded or revealed to a particular service provider. This has the obvious advantage of better protecting the privacy of each individual user. However, responsibility for storing and updating data then lies with the user.

Existing e-health solutions mainly take a limited view of patient information, utilising a healthcare provider-centric approach for identity management, mostly limited to a single healthcare provider. Interoperability has become an important issue, especially as more and more healthcare providers collaborate in an e-health system. If we refer to each healthcare provider, such as a hospital, general practitioner (GP) or clinical research lab, as a “context”, it is common practice for each entity, such as a patient, to have a unique identifier within a particular context. When context-specific information is transferred from one context to another, the same information can be expressed by means of different types or values. Typically, healthcare providers use different terms for the same entity – strictly relying on dictionaries may be very misleading. Besides, all the information exchanged among healthcare providers in the

e-health system needs to be uniquely identified. Therefore, there is a need to have an identity management mechanism in which a mapping of context-specific identifiers and a conversion of context-specific information occurs when data are exchanged among different contexts. Note that the above problem is only a cross-context issue when identifiers should not be shared directly among contexts. Linking information from one context to another should not be straightforward, hence the need for a privacy-friendly but interoperable IDM system. Previous work has mostly emphasised the e-health IDM issues from a provider-centric viewpoint. However, existing work on e-health systems reveals an unsatisfactory provision for the interoperability problem in cross-context IDM systems. Therefore, a new framework is proposed, where each context may deploy a user-centric IDM system, with multiple IDM systems collaborating and information transfer from one IDM system to another.

In this paper, we move towards this framework by defining a new model for identity management with an identifier conversion mechanism to ensure information interoperability in e-health. Note that we aim to provide an electronic platform to interconnect multiple healthcare providers and enable them to update and share patients' health information efficiently. In other words, the focus is on how different healthcare providers, instead of the patients, are able to collaborate on a common e-health platform efficiently. Though all e-health applications bear the ultimate goal to improve the health situation of the public, it is out of the scope of this paper to explore this further. The users of the proposed e-health platform are healthcare providers, such as GPs, hospital physicians, etc. A patient, being the holder of health information, can be identified either by the global identifier (e.g. the national registry number or national social security number), or by local identifiers issued by each particular healthcare provider the patient has visited. From the access control point of view, the patients can only read their own health records, and it brings the benefit of increased transparency. However, write access by the patients to any health records is prohibited. This is because the patients are not trustworthy, and this technical enforcement is necessary in order to prevent patients from altering health records to their own benefit.

This paper builds on the results of the IDem project (IDem, 2008) and the EHIP project (EHIP, 2008), and will explain how the proposed cross-context IDM service that was developed within the IDem project can be integrated in the framework of EHIP. The objective of the EHIP project, successfully completed in Flanders, Belgium, was to create a clinical data sharing infrastructure among multiple healthcare providers. The information sharing platform is designed so that information, such as patient e-health records, is always available and accessible at the time and place it is required, and only to authorised actors. Several key players in e-health, including leading sector companies, several university research groups and large hospitals, contributed to ensure that the project outcome was valid within a genuine context. In addition, a lifetime view was projected, which will be instrumental in guiding the transition in healthcare systems from provider-centric towards patient-centric.

As the legislation discussed in the following section requires, a patient's national identification number is not to be transferred cross-context. Therefore, we provide a privacy friendly solution where each healthcare provider assigns a unique context-specific identifier for the patient and the patient's medical record respectively. When information, such as a patient's medical data and patient

identifiers, is transferred cross-context, context-specific information is converted by our identity management mechanism for interoperability. **Towards a cross-context identity**

The rest of the paper outlines the EHIP platform architecture with functional components, followed by an explanation of the proposed cross-context identity management service in e-health, and a definition of the reversible algorithm to issue and convert context-specific identifiers. Next, how the proposed cross-context IDM service can be integrated in an e-health system is illustrated, with the EHIP platform as a case study. In particular, the motivating scenario, the system model and players, the attack model and assumptions are defined. Services of each entity, the command flow for a service request, and the protocol with the proposed scenario are discussed accordingly. As the final step, related work is introduced, and the future outlook for e-health systems together with a conclusion is provided.

425

Legislation background

In 2006 the United States Department of Health and Human Service (2006) issued the Insurance Portability and Accountability Act. This healthcare regulation demands the protection of patients' data shared from the original source of collection. Since 2005 the processing and movement of personal data has been legally regulated by the European Union (EU) through Directive 95/46/EC (EU, 1995). A citizen's right to privacy is also recognised in Article 8 (EC, 1987) of the European Convention for the Protection of Human Rights and Fundamental Freedoms. The debate surrounding the usage of single national identification numbers has longstanding historical roots. EU countries have sought to regulate their national number(s) in a variety of ways. Article 8.7 of Directive 95/46/EC (EU, 1995) provides that "Member States shall determine the conditions under which a national identification number or any other identifier of general application may be processed", indicating that governments should carefully consider how they allow national numbers to be used. Regardless of how national identification numbers are regulated in each respective state, they constitute "personal data" by nature within the meaning of Directive 95/46/EC. Articles 16 and 17 of the Directive impose on the controller a general confidentiality and security obligation, including the obligation for the controller to take all reasonable measures "to prevent all other unlawful forms of processing" (EU, 1995). Regardless of the possible perception that this might lead to massive data aggregation and profiling by the government, on the value of which we give no judgement, it is manifestly clear that the national number is not intended for use outside the governmental context.

The EHIP architecture

Classic community healthcare systems utilise an e-portal functionality to provide a many-to-one connection between many GPs and one hospital, based on propriety solutions. The disadvantages of this approach are twofold. First, it is impossible to interconnect different entities, such as hospitals. Second, one GP needs multiple portals to access patient data in different hospitals. As an improvement, a centralised infrastructure is enabled in community healthcare. Note that this does not need to imply that data physically resides in a central data store. The advantage of this approach is that users gain a consolidated overview of the clinical data of the patient, to which clinical research institutes and healthcare providers can interconnect.

The EHIP infrastructure aims to promote community healthcare and international standards on different forums. It provides a horizontal infrastructure for e-health applications compatible with international standards of Cross-Enterprise Document Sharing (XDS, 2007) by Integrating the Healthcare Enterprise (IHE, 2007) and technologies such as Web 2.0 to be used in web-based portals (AJAX, 2006; Portlet, 2008), which are interoperable within the Belgian e-health digital platform BeHealth (2006) infrastructure and hospital IT systems, with respect to security and privacy. It also provides a vertical application-based prototype for hospitals and GPs to share a patient's electronic health record (EHR), such as medical summary, clinical results and patient discharge letters.

Figure 1 shows the functional components of the EHIP platform. Based on a service-oriented architecture where each subsystem exposes its functionality through a service interface, it uses a central document registry to contain the metadata of all available documents, and distributed document stores where medical documents are stored in local repositories of the corresponding healthcare providers. The EHIP platform also contains a gateway to support healthcare providers with limited resources, such as small practices that cannot afford a repository. Further, the platform provides internet-enabled access to the resources through a web portal, which facilitates actions such as accessing the platform after-hours. Documents in the platform share a common content model as Clinical Document Architecture (HL7/CDA, 2007), so that all parties, despite their heterogeneous internal systems, gain easy access. The architecture employs federated security, in which security is embedded in middleware. Federated policy enforcement at hospitals and GPs' surgeries with a central policy management are deployed for access control, that is, authentication and

Figure 1.
The EHIP infrastructure

authorisation, in compliance with the EU data protection directive. According to XDS, the EHR security is covered by the following integration profiles (XDS, 2007): the Audit Trail and Node Authentication Profile to provide audit trail and the Cross Enterprise User Authentication Profile to provide a federated identity management framework based on SAML that enables single sign-on functionality across multiple enterprises. Wuyts *et al.* (2008) provide a security-enhanced software architecture with an analysis of potential security risks threatening XDS-compliant systems, such as the EHIP infrastructure.

Cross-context identity management in e-health

In this section, we look at identity management from an e-health prospective. In general, there are two types of identifiers in the EHIP system: a global identifier of a patient (e.g. the national identification number) and context-specific identifiers. The context-specific identifiers in the system are used to locally identify a patient and the patient's medical record within a specific context (e.g. a healthcare provider).

As mentioned in previous sections, all healthcare providers may have heterogeneous internal systems, and each healthcare provider typically issues its own unique context-specific identifier to each patient as well as to the patient's medical record, that will be stored in the local repository of the corresponding healthcare provider. This means the one patient will be issued different identifiers from different healthcare providers, and similarly the patient's medical records stored with different healthcare providers will be assigned different document identifiers. According to the legal restrictions explained, for privacy protection it is not advised to share the patient's global identifier among contexts.

We now attempt to expand the notion of e-health identity management to multiple contexts interacting and communicating with each other. One complication that occurs is that administrations need to exchange information coming from different contexts. For example, one healthcare provider tries to query the medical record concerning a patient from another healthcare provider, so that context-specific information is exchanged from one context to another. Further, the personal information exchanged needs to be uniquely identified, but the same identifier should not be shared among contexts. Whenever information is exchanged between different contexts a mapping and conversion of identifiers is required. In order to exchange information between contexts, identifier mapping and conversion is performed by a trusted party that is available for each context (Modinis, 2006). Since linkability of information from one context to another is desirable but not yet feasible, a manageable system for information interoperability is required. Therefore, our goal is to provide a cross-context IDM system, compatible with all the internal systems employed by the entities in the e-health platform, to translate and convert context-specific information and identifiers used and exchanged between the concerned entities.

Figure 2 depicts a cross-context information exchange in the e-health application context. An administrative organisation of a healthcare provider can be separated as front and back offices. The front office is connected with portals and local repositories. It directly interacts with its users, while the back office provides services for system support, such as identity management, authentication, authorisation, information sharing and auditing. The identity provider from the back office issues context-specific identifiers to its patients. Each healthcare provider is responsible for the issue and use

Figure 2.
Cross-context information
exchange in federated
e-health

of context-specific identifiers within its context. Accordingly, one healthcare provider cannot prevent another healthcare provider from issuing context-specific identifiers for its patients within a particular context.

When healthcare providers communicate, information can be exchanged through a mediating service provided by the e-health platform. The mediating service, that is to say a trusted party available for each context, is responsible for mapping and converting context-specific information, such as identifiers, exchanged between the communicating parties. This paper does not focus on how information is exchanged exactly, since it depends on semantic models, and application- and communication-specific scenarios. Instead, the contexts and entities involved in this communication are explored. In later sections, we explain how context-specific information can be converted and exchanged between contexts within a real-life scenario.

In Figure 3 the abstract structure of cross-context IDM applied in a communication is presented as an interconnected solar system. In e-health, node A, B, E, F and H

Figure 3.
An abstract structure of a
cross-context IDM system
with context-specific
information conversion

denote back offices. They can be the central server of a hospital or a gateway connecting various portals of different departments of a hospital. In a cross-context communication between a service provider and a service requester, the conversion of context-specific identifiers is performed by a trusted agent as the mediator, denoted by node C, D and G.

Cross-context IDM in the EHIP platform: a case study

In this section, we introduce how context-specific identifiers can be computed from global identifiers, and the reverse operation. Next, we explain how context-specific identifiers can be converted and exchanged across-context in the EHIP platform in a real-life scenario.

Reversible identifiers

Figure 4 shows the generation of a patient's context-specific identifier from his or her global identifier. Anon is defined as a deterministic algorithm to issue a context-specific identifier. The context-specific reference of any length is the input to a pseudo-random function with a secret key, such as a hash function, and results in a fixed-length message digest as a prefix. Then the prefix is concatenated to the global identifier and encrypted using a symmetric encryption algorithm with a second secret key, such as Advanced Encryption Standard (AES). The result is the context-specific identifier. Note that the secret keys for encryption and pseudo-random function should be different.

Figure 5 shows the reverse process to convert the patient's global identifier from the context-specific identifier, defined as Deanon. The context-specific identifier is the input to a symmetric decryption algorithm, controlled by the secret key. The result is the prefix concatenated with the global identifier. The prefix is easily removed allowing the global identifier to be recovered.

Figure 4.
Algorithm to issue a
context-specific identifier

Figure 5.
Algorithm to convert a
context-specific identifier

System model

Consider the following scenario: suppose in the EHIP network, a generic hospital H1 intends to query a medical record of patient A from a psychiatric hospital H2, through the EHIP registry R.

The system model we consider consists of the following players:

- (1) Hospital H1 is a generic hospital context, whose back office contains a file repository FR1, a patient ID provider PIP1, a document ID provider DIP1, and a document anonymiser DA1.
- (2) Hospital H2 is a psychiatric hospital context, whose back office contains a file repository FR2, a patient ID provider PIP2, a document ID provider DIP2, and a document anonymiser DA2.
- (3) Patient A is a patient who requests healthcare services from a healthcare provider. Let GID denote A's global identifier (e.g. the patient's national number). Let PID_j denote A's context-specific identifier (a pseudonym) assigned by the healthcare provider H_j . Let DID_j denote the pseudonym for A's medical record $DocA_j$ stored in the context H_j .
- (4) EHIP Registry R is a central registry that maintains a link between a patient's global ID and the locations of each healthcare provider that stores the patient's medical records.

Attack model and assumptions

The objective of an attacker Eve is to obtain private information about a particular patient. In the proposed scenario, Eve may either try to obtain the patient's global ID or the patient's sensitive medical information from the psychiatric hospital H2.

In order to do so, Eve has several options. First, she tries to request the patient's global identifier from the identity providers of each healthcare provider or the central registry. Then, Eve tries to request the sensitive medical data directly from the hospital H2. After which, Eve tries to steal the secret keys of any identity providers or the document anonymiser in the system in order to access the sensitive data. In addition, Eve tries to break into the system. Finally, she could try to eavesdrop to obtain the desired content. An attacker can be categorised as either internal to the e-health network, or external. An internal attacker can be either an authorised or unauthorised recipient of the e-health system services. We assume all attackers external to the e-health network are unauthenticated and unauthorised entities to the system.

The following assumptions hold for the proposed system. All entities that have been authenticated and authorised by the system are assumed trustworthy. The system is not protected against malicious entities that are able to authenticate themselves and who are authorised to use the system's services. We assume that all security-enhancing functionalities employed in the system are robust and well-deployed. All secret keys of the entities in the system are stored physically secure. The communication takes place through a secure communication channel.

Proposed approach

Services provided by each entity

In the hospital context H_j , let KP_j denote the secret key of the patient ID provider PIP_j , and let KD_j denote the secret key of the document ID provider DIP_j . The ID providers PIP_j and DIP_j are able to provide two services: $IDIssue$ and the $IDConvert$. Let $KDoc_j$ denote the secret key of the document anonymiser DA_j , and DA_j is able to provide two services: $DocAnon$ and $DocDeanon$. Both H_j 's file repository FR_j and the EHIP central registry R can provide the service $Query$.

Each service is described as follows:

- *IDIssue*. A service to issue context-specific identifiers from global identifiers.
- *IDConvert*. A service to convert context-specific identifiers to global identifiers.
- *DocAnon*. A service to pseudonymise a document by encrypting part of a document that contains sensitive medical information.
- *DocDeanon*. A service to convert a pseudonymised document back to the non-pseudonymised version by decryption.
- *Query*. A database query service with the input of some attributes and the output of some other attributes from the database.

Proposed approach to request a service

As shown in Figure 6, before a service is delivered to a service requester from a service provider, the service requester needs to be authenticated and authorised each time. The information flow to request a service contains the following steps:

Figure 6.
Check service request
commands flow

- (1) A service requester sends a service request to a service provider.
- (2) The service provider forwards the request to its security facade to check the requester.
- (3) The security facade checks the requester’s authenticity and authorisation.
- (4) If the checks are passed, the security facade informs the service provider to deliver the required service. Otherwise, the service delivery is denied.
- (5) The service provider delivers the service to the service requester.

Protocol of the proposed scenario

Figure 7 presents the protocol of a scenario in which the generic hospital H1 queries a medical record of patient A from the psychiatric hospital H2 through the registry R in the EHIP system. Note that R serves as the mediator for the cross-context communication between H1 and H2. The registry R of the e-health platform interacts with the ID providers of the two contexts (hospitals H1 and H2) as well as issues and converts context-specific identifiers.

As shown in Figure 8, information is transferred cross-context in the following steps:

- (1) H1 queries R with the context-specific patient ID PidA1 of a patient A.
- (2) R requests the IDConvert service with PidA1 from H1’s patient ID provider PIP1 in order to obtain the global ID GidA of A.
- (3) After the authentication and authorisation check of R by H1, PIP1 performs IDConvert on PidA1 with a secret key Kp1, and then delivers the resulting global identifier GidA to R.

Figure 7.
The protocol of the
cross-context query of a
medical record scenario in
the EHIP infrastructure

Figure 8.
Command flow in the
cross-context query
scenario in the EHIP
infrastructure

- (4) R queries its database to retrieve the corresponding location of the hospital where A's medical record is stored, which is assumed at H2.
- (5) R requests the IDIssue service from H2's document ID provider DIP2 with GidA.
- (6) After the authentication and authorisation check, DIP2 provides R with the context-specific document ID of A's medical record stored in H2.
- (7) R sends DidA2 and H2's location information to H1.
- (8) H1 sends a query to H2 with DidA2.
- (9) H2 queries its file repository FR2 and retrieves A's medical record DocA2.
- (10) H2 requests the DocAnon service from its document anonymiser DA2 for an anonymised version of DocA2.
- (11) After the authentication and authorisation check, DA2 performs the document pseudonymisation service DocAnon on DocA, resulting in the pseudonymised medical record AnonDocA2.
- (12) H2 delivers the requested medical record AnonDocA2 to H1 through a secure channel.

Security analysis

In this section, we present a security analysis of our proposed solution, in correspondence to the attack models. We will show, under the predefined assumptions, why those potential attacks cannot be performed successfully in our system. Firstly, it is important to assume that all entities that are authenticated and authorised to obtain a certain service from a service provider are trustworthy, and vice versa. We also assume that any entity who fails to pass a service provider's authentication or authorisation check should not obtain the corresponding service.

As explained above, the system is not protected against malicious entities that are checked as authenticated and authorised. Now we consider the attacker Eve as an unauthenticated or unauthorised entity, and examine the five misuse cases Eve may perform. In the first and the second misuse cases, Eve tries to request the patient's global identifier with the IDConvert service from any of the healthcare provider's identity providers or from the central registry. Alternatively, she tries to request the sensitive medical data directly from the hospital H2. However, neither case is possible because in order to receive the service, Eve has to pass the authentication and authorisation check by the service providers. The third attack Eve may perform is to attempt to steal the secret keys of any identity providers or the document anonymiser in the system so that she could perform the IDConvert service or the DocDeanon service to obtain the desired information. This is unfeasible for Eve, since we assume all secret keys of the entities in the system are stored securely. In the next misuse case, Eve may try to hack the system and breakdown the security. This option is not feasible either, because all the security-enhancing functionalities employed in the system are assumed to be robust and well-deployed in the proposed system. If Eve fails to perform all the above attacks, Eve can still try to eavesdrop on the communication content. However, presumably the communication is taking place through a secure communication channel with client-side and server-side authentication.

We can conclude from the above analysis that the security of the proposed system depends on the security of the secret key, of the communication channel and of the underlying system security infrastructure, such as the security of the authentication and authorisation mechanisms.

Related work

Some of the popular user-centric identity management systems developed over the past years include Shibboleth (Shibboleth, 2001; Scavo and Cantor, 2005), Liberty Alliance (Liberty Alliance Project, 2006, 2008), CardSpace (CardSpace, 2007), and Idemix (Idemix, 2008; Camenisch and Herreweghen, 2002). In the literature, there are some identity management schemes proposed for e-health that utilise a user-centric approach. Peyton *et al.* (2007) used a simple e-prescription scenario to analyse the business and technical issues to be addressed in a Liberty Alliance-based federated identity management framework for e-health. They looked at the potential impact of privacy compliance on three existing components of the framework – Discovery Service, Identity Mapping Service and Interaction Service. They proposed a fourth component, Audit Service, to address potential privacy breaches in Liberty Alliance. Au and Croll (2008) recently proposed a new framework for a consumer-centric identity management for distributed e-health. In their framework, the healthcare consumer maintains a pool of pseudonym identifiers in their personal secure device for use in different healthcare services, perhaps in the form of a smart card. Without revealing consumer identity, health record data from different medical databases distributed in various points of clinical service can be collected and linked together on demand. In particular, pseudonym identifiers are cryptographically generated by a trustee, and the binding of an identifier to the identity key or another identifier is certified by a Key Binding Certificate issued by the trustee. Hence, security of the interactions among different entities in the architecture is guaranteed by certification and cryptographic technologies.

Some results have been published on privacy protection and secondary use issues of EHR. Iacono (2007) discussed the importance of protecting the privacy of patient data kept in an EHR in cases where it leaves the control- and protection-sphere of the healthcare realm for secondary uses such as clinical or epidemiological research projects, healthcare research, and assessment of treatment quality or economic assessments. The work focused on multi-centric studies, where various data sources are linked together using Grid technologies. It introduced a pseudonymisation system that enables multi-centric universal pseudonymisation, meaning that a patient's identity will result in the same pseudonym, regardless from which participating study centre the patient data is collected. Pommerening and Reng (2004) addressed the issue of secondary uses of EHR, such as health economy and healthcare research, or disease specific clinical or epidemiological research. For these uses in general, the patient identity must be anonymised or pseudonymised. Their work described possible model architectures, developed for medical research networks, but useful in broader contexts.

In Europe there have been several research projects on cross-border identity management. The concept of context-specific identifiers was introduced in the Modinis Study on Identity Management (Modinis, 2006). Modinis-IDM (2006) was an EU-funded identity management project that focused on interoperability aspects of identity management systems used in the EU member states. Its aim was to build on

expertise and initiatives in the EU member states to progress towards a coherent approach in electronic identity management in e-government in the EU. The study addressed interoperability issues in cross-context IDM in e-government, without ignoring differences in legal and cultural practices within the EU framework for data protection. GUIDE (2006) was also an EC-funded project, with the aim to create an architecture that will enable open and interoperable e-government electronic identity services in the EU. Its objective concerned interoperability across national systems and structures within broader transnational policy, legislative and socioeconomic boundaries. The PRIME (2008) project looked at the applicability issues of using the federated identity management system Idemix open source initiative and digital credentials in detail. The main contribution of this European research project was a broader understanding of the dependencies between the different components in such a system. These dependencies are reflected by both identity management architecture and an integrated prototype. FIDIS (2008) is an EU-sponsored Network of Excellence targeting various aspects of digital identity and privacy. FIDIS's areas of interest include new forms of ID cards, use of identifiers in information systems, technologies used for citizen's identification and profiling. Research projects in Belgium, such as Identity Management for eGovernment (IDEM, 2008), focus on the identity management aspects that are relevant in a heterogeneous e-government context, and compare the different governments in Flanders, Brussels and Wallonia that have to interoperate with the federal services.

There are some governmental or industrial partners in Belgium in the related fields of IDM or e-health. Crossroads Bank for Social Security (CBSS, 2007) is active in the field of IDM of e-government in the social sector. This organisation provides technical solutions to function as a mediator for cross-context communications among different sectors, and has proposed an algorithm to issue one-way only context-specific identifiers. Custodix (2008) is a company active in the e-health sector. Generally, Custodix is a trusted third party that provides security solutions based on privacy-enhancing techniques at international level. The services place special emphasis on anonymisation and pseudonymisation.

Practical implications and future trends

User-centricity: patient specification

E-health systems are evolving towards user-centricity, where the patients are able to control the granularity of healthcare information disclosed to the third party, and to specify the content of the health information and to which healthcare provider the information can be disclosed, the purpose of processing the information, etc.

Transparency: patient verification

Transparency will be emphasised more in e-health systems such that the patients should be able to access and query the logs of their health records in order to verify if their records were access according to the refined rules.

Federated e-health with interoperability

Another trend is federated e-health, including federation of identity information and healthcare information. Traditional e-health systems are typically integrated within each healthcare provider; however, interconnection between different healthcare

providers is lacking. In this case, the patients function as the only link, carrying their medical files from one healthcare provider to another. Application Programming Interface (API) is available to specify how an individual e-health system works in each healthcare provider, and there is no interaction between different APIs.

The next generation e-health systems are able to offer collaboration between healthcare providers by using a data bus interconnecting different service providers, while each API will not only be available but will also be used for the exchange of health information. In this setting, each healthcare provider has its own database, and the database in one healthcare provider is accessible from the outside by another healthcare provider. As previously explained, information with the same meaning can be interpreted as different types or values by different healthcare providers. Therefore, interoperability is required to ensure information federation in e-health systems. The healthcare information stored in the local database is transferred and translated by the data bus and can be shared between two or more healthcare providers. Interoperability can be realised by implementing and using the APIs specified within each provider.

Conclusions

In this paper, we have presented a new service for cross-context identity management in the e-health application domain, aiming to improve interoperability when context-specific information is transferred between contexts. Previous e-health IDM solutions have mostly had a limited view of patient information, where a user-centric approach for identity management usually was restricted to a single healthcare provider. Interoperability becomes more problematic in an e-health system when more actors collaborate, such as hospitals, GPs, clinical research labs, pharmacists, etc. In such systems, it is common for a patient to be issued different unique context-specific identifiers from different healthcare providers. In cross-context communications, the same information can be expressed by means of different types or values. Since identifiers are not shared directly among contexts, linkability from one context to another should not be straightforward. However, other forms of linkability, such as the possibility to follow-up a patient's medical treatment, is desirable in the e-health sector, even when it needs to cross different contexts. Therefore, in the e-health context, we have designed an identity management mechanism in which a mapping and conversion of context-specific identifiers or information occurs when data is exchanged among different authorised healthcare providers. Further, we have proposed an algorithm for issuing and converting context-specific identifiers, based on cryptographic techniques. As an illustration of the concept, we have presented our research activities on the IDM aspect of the EHIP project, with a real-life use case scenario to explain how the proposed cross-context IDM service can be integrated in the EHIP e-health platform.

As the next step of our research, we are looking for solutions on how to extend the framework to provide federated authentication and federated authorisation mechanisms. Further work is needed to define methods for federation and management of authorisations.

References

- Au, R. and Croll, P. (2008), "Consumer-centric and privacy-preserving identity management for distributed e-health systems", *Proceedings of 41st Hawaii International Conference on Systems Science (HICSS-41 2008)*, IEEE Computer Society, p. 234.

- AJAX (2006), "Ajax at the open directory project", available at: <http://dmoz.org/Computers/Programming/Languages/JavaScript/AJAX>
- BeHealth (2006), "BeHealth", available at: www.behealth.be/
- Camenisch, J. and Herreweghen, E.V. (2002), "Design and implementation of the idemix anonymous credential system", Technical report, IBM Research Division.
- CardSpace (2007), "Windows CardSpace", available at: <http://netfx3.com/content/WindowsCardspaceHome.aspx>
- CBSS (2007), "Crossroads bank for social security ", available at: www.ksz-bcss.fgov.be/En/CBSS.htm
- Custodix (2008), "Custodix home", available at: www.custodix.com/
- EC (1987), *European Convention on Human Rights*, Martinus Nijhof Publishers, Leiden.
- EHIP (2008), "eHealth information platforms", available at: www.ibbt.be/en/project/e-hip
- EU (1995), "Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data", *Official Journal of the European Communities*, Vol. 281, pp. 31-50.
- FIDIS (2008), "FIDIS home: future of identity in the information society", available at: www.fidis.net/
- GUIDE (2006), "Creating a European identity management architecture for eGovernment", available at: <http://istrg.som.surrey.ac.uk/projects/guide/overview.html>
- HL7/CDA (2007), "HL7/CDA release 2.0, clinical document architecture", available at: www.hl7.org/library/standardsfinon1.htm
- Iacono, L.L. (2007), "Multi-centric universal pseudonymisation for secondary use of the HER", *Studies in Health Technology and Informatics*, Vol. 126, pp. 239-47.
- IDEM (2008), "Identity management for eGovernment", available at: <https://projects.ibbt.be/idem/>
- Idemix (2008), "Idemix: pseudonymity for e-transactions", available at: www.zurich.ibm.com/security/idemix/idemix
- IHE (2007), "Integrating the Healthcare Enterprise (IHE) overview", available at: www.ihe.net/
- Josang, A. and Pope, S. (2005), "User centric identity management", in Clark, AL. (Ed.), *Proceedings of AusCERT 2005, May, Brisbane, Australia*.
- LDAP (2006), "LDAP directory technology review – an investigation into the use of directories for developing a management application which can manage both data and network services, specifically for the web hosting industry", available at: <http://ldap.artzweb.net/>
- Liberty Alliance Project (2006), "Liberty alliance project whitepaper: personal identity", 23 March, available at: www.projectliberty.org/liberty/content/download/395/2744/file/Personal_Identity.pdf
- Liberty Alliance Project (2008), "Liberty technology glossary working draft", available at: <http://xml.coverpages.org/draft-liberty-tech-glossary-08.pdf>
- Modinis (2006), "Modinis study on identity management in eGovernment – the conceptual framework version 1.1", available at: www.cosic.esat.kuleuven.be/modinis-idm/twiki/bin/view.cgi/Main/ConceptualFramework
- Modinis-IDM (2006), "Modinis study on identity management in eGovernment – home", available at: www.cosic.esat.kuleuven.be/modinis-idm/twiki/bin/view.cgi/Main/WebHome

- Peyton, L., Hu, J., Doshi, C. and Seguin, P. (2007), "Addressing privacy in a federated identity management network for eHealth", *Eighth World Congress on the Management of eBusiness.WCMeB 2007*, IEEE Computer Society, Toronto, p. 12.
- Pommerening, K. and Reng, M. (2004), "Secondary use of the EHR via pseudonymisation", in Bos, L., Laxminarayan, S. and Marsh, A. (Eds), *Medical Care Compunetics 1*, IOS Press, Amsterdam, pp. 441-6.
- Portlet (2008), "JSR-286, the Java Portlet API 2.0", available at: www.jcp.org/en/jsr/detail?id=268
- PRIME (2008), "Privacy and identity management for Europe", available at: www.prime-project.eu/
- Scavo, T. and Cantor, S. (2005), "Shibboleth architecture. Technical overview. Working draft 02", Technical report, Internet2/MACE
- Shibboleth (2001), "Shibboleth overview and requirements", available at: <http://shibboleth.internet2.edu/docs/draft-internet2-shibboleth-requirements-01.html>
- United States Department of Health and Human Service (2006), "HIPAA administrative simplification: enforcement; final rule", *Federal Register/Rules and Regulations, European Convention on Human Rights*, Vol. 71 No. 32.
- Wuyts, K., Scandariato, R., Claeys, G. and Joosen, W. (2008), "Hardening XDS-based architectures", *Third International Conference on Availability, Reliability and Security 2008, ARES08, Barcelona*, pp. 18-25.
- XDS (2007), "Integrating the healthcare enterprise: IT infrastructure technical framework, revision 4.0", available at: www.ihe.net/Technical_Framework/

Appendix. Definitions of terms

Access control

Access control is the protection of resources with technical, regulatory and organisational measures against access or use by unauthorised entities.

Authentication

Authentication is the corroboration of a claimed set of attributes or facts with a specified, or understood, level of confidence.

Authentication may be used during any IDM process. Authentication serves to demonstrate the integrity (i.e. equivalence to a corresponding reality) and origin (i.e. the source) of what is being pretended (the claimed information). Authentication can be unilateral or mutual. Unilateral authentication provides assurance of the identity of only one entity, while mutual authentication provides assurance of the identities of both entities.

Authorisation

Authorisation refers to the permission of an authenticated entity to perform a defined action or to use a defined service or resource; the process of determining, by evaluation of applicable permissions, whether an authenticated entity is allowed to have access to a particular resource.

Usually, authorisation is in the context of authentication. Permission is granted or denied based on the result of data or entity authentication, and on the allowed activities, as defined within the system. Once an entity is authenticated, it may be authorised to perform different types of access, each of which is referred to as a role.

Confidentiality

Confidentiality refers to the state of keeping the content of information secret from all entities but those authorised to have access to it.

Context

A context is a sphere of activity, a geographic region, a communication platform, and an application, a logical or physical domain.

Practically, a context is only relevant in an interaction. Cross-context refers to activities over two or more contexts.

Digital identity

A digital identity is a partial identity in an electronic form.

A digital identity can be created on the fly when a particular identity transaction is desired. A digital identity is by definition a subset of the identity, and can in effect be considered a manifestation of an entity's presence in an electronic IDM system (i.e. it is the subset of attributes belonging to an entity that is accessible through a specific IDM system).

Federated identity

A federated identity is a credential of an entity that links an entity's partial entity from one context to a partial entity from another context.

Identification

Identification is the process of using claimed or observed attributes of an entity to deduce who the entity is.

The term identification is also referred to as entity authentication. The identification of an entity within a certain context enables another entity to distinguish between the entities it interacts with.

Identifier

An identifier is an attribute or a set of attributes of an entity that uniquely identifies the entity within a certain context.

An entity may have multiple distinct identifiers referring to it. Identifiers uniquely identify an entity, while characteristics do not need to. However, it should be noted that identifiers can consist of a combination of attributes, whereas characteristics are always one single attribute.

Identity management (IDM)

Identity management is the managing of partial identities of entities, that is, definition, designation and administration of identity attributes as well as choice of the partial identity to be (re-)used in a specific context.

Privacy

Privacy is the right of an entity – in this context usually a natural person – to decide when and on what terms its attributes should be revealed.

In an IDM context, privacy is mostly used as a synonym of informational privacy, that is, the interest of a natural person to control or at least significantly influence the handling of data about themselves, also taking into account the nature of the applicable attributes and the entity in charge of data management.

Pseudonym

A pseudonym is an arbitrary identifier of an identifiable entity, by which a certain action can be linked to this specific entity.

The entity that may be identified by the pseudonym is the holder of the pseudonym. A pseudonym is typically a fictitious name that can refer to an entity without using any of the entity's identifiers. As identifiers, pseudonyms are context-bound, and one pseudonym is not necessarily valid across multiple identity management systems. An entity is pseudonymous if it

relies on a pseudonym as identifier. The procedure by which all person-related data within a data record is replaced by one pseudonym is pseudonymisation.

Trust

Trust is a quality of a relationship between two or more entities, in which an entity assumes that another entity in the relationship will behave in a fashion agreed beforehand, and in which the first entity is willing to act on this assumption.

Whether or not to trust depends on a natural person's decision. It is possible, but not necessary, that several entities trust each other mutually in a certain context. Trust decisions of legal persons depend on the decisions made by the legal person's responsible natural persons. Trust may be limited to one or more specific functions, and may depend on the fulfilment of one or more requirements.

Trusted third party (TTP)

A trusted third party is an entity trusted by multiple other entities within a specific context and which is alien to their internal relationship.

About the authors

Mina Deng received her MSc in electrical engineering from the Katholieke Universiteit (KU) Leuven, Belgium, in 2004. She is currently a PhD student at the Computer Security and Industrial Cryptography (COSIC) research lab in the Department of Electrical Engineering at the KU Leuven. She also works as a security and privacy researcher for the Interdisciplinary Institute for BroadBand Technology (IBBT) Belgium. She has been involved in several European research projects, such as SPEED-Signal Processing in the Encrypted Domain (2007 to present), the European Networks of Excellence FIDIS – Future of Identity in the Information Society (2004 to present) and ECRYPT – European Network of Excellence for Cryptology (2004 to 2008) by the European Commission-Framework Program 6 (EC-FP6). Ms Deng's research interests include cryptography, security and privacy, identity management, and electronic service applications such as DRM, e-commerce, e-health, etc. She received the Best Student Paper award of the IEEE International Symposium on Electronic Commerce and Security 2008 for the paper Attacks on two buyer-seller watermarking protocols and an improvement for revocable anonymity, the Outstanding Achievement Award for Best Student Research Paper from the New Zealand State Services Commission in 2008, and the Barco/VIK-Prize for the Belgian Flemish best engineering thesis in 2003. Mina Deng is the corresponding author and can be contacted at: mina.deng@esat.kuleuven.be

Danny De Cock received a Master's degree in Computer Science from the KU Leuven (Belgium) in 1996. After his studies, he started working as a fulltime researcher at the KU's ESAT/COSIC research group, which specialises in computer security and industrial cryptography applications. One of his main interests is the specification and follow-up of the Belgian eID card's functionality and the corresponding services. Related to this he is conducting research on the analysis of identity management systems involving strong authentication tokens and their validation mechanisms, on electronic voting mechanisms, on the analysis of smart card and memory card applications and their protection mechanisms, and on the security analysis of wireless chip cards, electronic banking and payment systems, and biometric authentication systems.

Bart Preneel received a Master's degree in electrical engineering and a Doctorate in Applied Sciences (Cryptology) from the KU Leuven (Belgium) in 1987 and 1993 respectively. He is currently a Professor at the KU Leuven. He was visiting professor at five universities in Europe and was a research fellow at the University of California at Berkeley. He has authored and co-authored more than 200 reviewed scientific publications and is the inventor of two patents. His main research interests are cryptography and information security. Professor Preneel is

president of the IACR (International Association for Cryptologic Research) and of L-SEC vzw (Leuven Security Excellence Consortium), an association of 60 companies and research institutions in the area of e-security. He is a member of the Editorial Board of the *Journal of Cryptology*, the *IEEE Transactions on Forensics and Information Security*, and the *International Journal of Information and Computer Security*. He has participated in more than 20 research projects sponsored by the European Commission, for four of these as project manager. He has been programme chair of ten international conferences (including Eurocrypt 2000, SAC 2005 and ISC 2006) and he has been invited speaker at more than 30 conferences. In 2003 he received the European Information Security Award in the area of academic research, and he received an honorary Certified Information Security Manager (CISM) designation from the Information Systems Audit and Control Association (ISACA).