Identity in Federated Electronic Healthcare

Mina Deng*, Riccardo Scandariato[†], Danny De Cock*, Bart Preneel* and Wouter Joosen[†]

*IBBT-COSIC, Katholieke Universiteit Leuven, Belgium

[†]IBBT-DistriNet, Katholieke Universiteit Leuven, Belgium

Abstract—In electronic healthcare several research and standardization activities are emerging that promote federation. In this scenario, the medical information present at different healthcare providers, such as hospitals, general practitioners, test laboratories, etc., are shared for an improved quality of experience from the patient perspective. However, sharing of medical data on a large scale exposes the patient to several privacy-related threats, such as massive data aggregation or profiling. Therefore, the selection of a privacy-preserving identification scheme is a primary requirement in federated e-health. This paper presents an identity management infrastructure that minimizes the above-mentioned threats.

Index Terms—Identity management, privacy, e-health.

I. INTRODUCTION

Traditional electronic health solutions were mainly concerned with a limited view on the patient information, taking a provider-centric viewpoint, and mostly limited to a single provider. A paradigm shift is taking place in the e-health domain, which is evolving from provider-centric towards patient-centric healthcare. One important requirement in order to improve the quality of experience of the patient is the continuous and transparent availability of medical information, independently from the location where the information has been actually stored. Although a patient will typically visit different healthcare providers over time, and hence the medical information will be dispersed over several locations, the medical record of a patient should be available anytime and anywhere, in a location-independent way. To this aim, healthcare providers, such as hospitals, general practitioners, research laboratories, etc., are federating to share their medical data.

Medical data is of sensitive nature, and therefore several laws and regulations mandate to protect the privacy of the patient [1]. In particular, the federation scenario presents a specific privacy threat. Indeed, this domain makes intensive use of identity information. For instance, in order to retrieve all the necessary data relevant for the ‘treatment’ of a patient, there must be a mechanism to cross-reference medical documents across healthcare providers. That is, it should be possible to search and retrieve documents from several locations on the basis of the patient identity. Naturally, access to such documents is restricted by authorization rules, which, yet again, make an intensive use of identity information about both the healthcare professionals and the patients. Examples clarifying the role of identity in the authorization process are provided later on in this paper.

From a functional perspective, the simplest solution would be to use of global identifiers across the different providers, or

‘contexts’ from this point on. However, this is not a feasible strategy for two reasons. First, healthcare providers require to maintain control over the process of issuing identifiers. This is mainly due to legacy constraints. Second, if medical data sources would use global identifiers, the risk of massive data aggregation and profiling would be much higher. An attacker that got to know the content of two medical databases could be able to correlate the data quite easily.

To accommodate these conflicting forces, namely the need of cross referencing documents and the avoidance of global identifiers, some solutions have been proposed that employ a mediating component. Local identifiers are used within each context and the mediator provides translation services from one context to another. However, if the mediator maintains the translation information on board, such as in the form of a lookup table, it becomes a likely target for attackers. An attacker could steal that information and use it to perform the correlation mentioned above. State-of-the-art solutions in the e-health domain are vulnerable to such attack scenario.

Because existing work reveals an unsatisfactory provision for the interoperability problem in cross-context identity management, we propose a new service to manage identifiers in e-health systems. Specifically, this paper proposes a cryptographic algorithm to be used in issuing context-specific, hence local, identifiers. Local identifiers are derived from a unique global identifier in a reversible way. The algorithm is meant to be used by the identity providers located at each healthcare provider. Further, for cross-context interoperability, a stateless mediation service is presented. The mediation service leverages the reversibility property of local identifiers and does not maintain any cross-referencing information on board. Further, the entity that functions as the mediator is not fixed and may vary.

The rest of the paper is organized as follows. The relation between identity and authorization in federated e-health is discussed in Section II. How to manage identifiers in e-health and an infrastructure for context-specific identifiers translation is proposed in Section III. Cross-context identification and authorization in an e-health system is illustrated in Section IV. Related work is introduced in Section V and Section VI provides a conclusion.

II. AUTHORIZATION IN FEDERATED E-HEALTH

It is well known that identification plays a key role in supporting authorization. From the study of typical authorization rules we realized that such role is even more fundamental in the federated e-health domain. In the EHIP research project

we have developed the security architecture of a multi-party sharing platform. The platform is a communication infrastructure that allows many healthcare providers to collaborate by sharing the medical information they produce. In collaboration with clinical partners, we have elicited and analyzed the low level policy rules used in a real hospital setting. Consequently, we have extracted the authorization rule types that are relevant in the federated case.

Roles have been adopted in the past as the cornerstone technique to manage permissions in e-health, e.g., in the context of the UK National Health Service [2]. In fact, we observed that role is less central than expected in deciding whether an access request to medical information should be granted or not. Rather, we discovered that existing relationships between patients and physicians, besides other context-dependant parameters, such as time and location, are of primary importance in the authorization process. Hence, establishing identity of involved parties is often a primary pre-requisite to authorization. In the remaining of this section we illustrate some typical policy rule types and highlight the identity-related information that is important for the decision process.

A. Authorization in federated e-health

This section describes some generic authorization rules, each requiring the establishment of the identity of a specific patient in order to be enforced. Identity is typically used to verify the presence of a certain relationship between the patient and the physician requesting access to the patient data. Each rule type is described according to the same template: first we give a general description of the rule type, then we provide one example of a possible instantiation, and finally we provide a detailed explanation of the rule with particular focus on the role played by identity.

1) Patient-physician treatment relationship

Rule: Physicians who treat a patient, either as supervisor or executing physician, are granted access to patient data related to that treatment.

Example: A screening center has access to the mammographic pictures of the radiology center to perform a reading, because the screening center is implicitly treating the patient.

This policy provides an example of the treatment relationship, which is the relation between a patient and the physicians that are dealing with the patient during a treatment process. This relationship can be explicit or implicit. In an explicit relationship the treating physician is explicitly assigned, for instance by name, to the patient. Note that there is a clear relationship, as seen by both the physician and the patient. The implicit relationship is illustrated by the example, where the radiologist from the screening center is implicitly assigned to the patient by performing his function and can be considered as part of the treating process of the patient. Note that there is no direct relationship between the patient and the radiologist.

The policy will grant access to the patient data if a relationship exists, and will deny access if no relationship has been established. To decide whether or not a relationship exists,

the identity of both the requester, such as a physician, and the patient must be established. Note that in a cross-context access request, identities are expressed in the ‘vocabulary’ of the requester, i.e., using identifiers that are local to the requester’s context, which may not be meaningful to the authorization service of the context where the requested data belongs to.

2) Patient-department relationship

Rule: A physician is granted view access to the patient’s data, if the patient resides or resided less than two weeks ago on a department to which the physician is assigned to.

Example: When a patient is transferred between hospitals, the physician of the hospital where the patient resided less than two weeks ago, can also access relevant data of the patient from the other hospital.

For this policy, the patient history has to be taken into account. The transfer of the patient between departments, or more in general, between healthcare institutions, needs to be tracked. The time the patient has spent in the hospital has to be considered as well. This policy is clearly related to the treatment relationship case. However, in this case, physicians no longer holding a current treatment relationship, can still access the patient’s data.

3) Physician-department relationship

Rule: A specific physician can view patient data that originated within one of the departments the physician is assigned to.

Example: A physician can remotely access data of the patient via a web portal if the data was created by the physician’s department.

This policy is enforced by establishing the physician’s affiliation. The example described above is rather narrow. This could be extended to data within the same discipline, spread over several healthcare institutions, instead of just within one department. Obviously, this rule requires that the patient-department relationship is verified, as in the previous case.

4) GP-patient relationship

Rule: A general practitioner (GP) retains the access to the medical reports concerning the patient as long as she remains registered as the patient’s GP.

Example: A GP can always access medical reports of all of her patients.

A GP needs specialized rules, in contrast with other healthcare providers, because a GP does not belong to a healthcare institution. Therefore, the GP will not be granted access on the basis of a treatment relationship or because she belongs to a certain department. Rather, access decisions are only based on the long-lasting relationship with the patient.

5) Identity in obligations

Rule: A physician can overrule an access denial, provided that a detailed reason is specified. The system is obliged to log the identity, the reason, the access time, and the accessed resources.

Example: Before a surgical operation, an anesthetist does not automatically get access to the information pertaining the allergies of a patient, because at that time the patient is not yet admitted, so the anesthetist is not a treating physician. An

anesthetist can overrule the denial in order to better prepare for the operation. Overruled access is logged.

It is a strong requirement from the regulatory perspective to establish the identity of the physician that overruled the decision of the authorization service, and the identity of the the patient for which such overruling took place. Therefore, policies exist describing what and how to log and they all require that the individual's identity is traced for auditing and possible legal reasons.

B. Identity and authorization

An interesting result of this study is that role-based access control does not suffice in the federated e-health scenario. This section has identified several cases where verifying identity, rather than role-related credentials, is a pre-requisite to the enforcement of cross-context e-health authorization rules. Further, in real world scenarios there are many, often complex, exceptions to the baseline rules described above, such as the following one: “no access to application X except for personnel of unit 500, for department PNE, LOG, PSY, unless they are assistants in training or if they have user-id ABC or XYZ.” This shows that identifiers play a key role in these cases.

In summary, the policies described above have illustrated that establishing identifiers is necessary to enforce authorization rules, which involve:

- current and historical treatment relationships: identities are used to evaluate the access rights of the physician on a need-to-know basis;
- visit history of the patient: identities are used to verify the relationship with a department, a discipline, and so on;
- long-lasting relationships: such as contractual relationships between patients and the GPs;
- exceptions: identities are directly referenced in the rules;
- auditing: identification is required by policy.

III. MANAGING IDENTIFIERS IN FEDERATED E-HEALTH

In this section, we propose an algorithm to issue and convert context-specific local identifiers to global identifiers, and vice versa. The algorithm is leveraged to build a privacy-friendly, cross-context identification infrastructure.

A. Reversible local identifiers

In general, there are two types of identifiers in an e-health system: a patient's *global identifier*, such as national identification number, and *context-specific local identifiers*, used to locally identify a user within a specific healthcare provider. Each healthcare provider may have heterogeneous internal systems, and is responsible to issue context-specific identifiers for its patients. In other words, the same patient will be issued with different local identifiers by different healthcare providers. According to legislation restrictions, sharing global identifiers directly across contexts may lead to massive data aggregation or profiling from government or corporations.

Fig. 1. Algorithm to issue/recover a context-specific identifier.

Figure 1 depicts a deterministic algorithm to issue a context-specific identifier from a global identifier. The algorithm's public input are, namely a global identifier and a context reference string. The private input are two symmetric secret keys, one for a pseudo-random function, and the other for a symmetric encryption function. The algorithm provides a fixed length context-specific identifier as output. In particular, the context-specific reference of variable length is the input of the pseudo-random function, such as HMAC-SHA-256, and this results a 256-bit message digest as a context-prefix. Then the prefix is concatenated with the global identifier of fixed length, and they are encrypted using a symmetric encryption algorithm, such as AES-CBC mode. The final result is the context-specific identifier. Note that the secret keys for the encryption and the pseudo-random functions may be different. For interoperability, the process to issue identifiers is reversible (see the upward arrows in the picture).

B. Interoperability infrastructure

In an e-health system with multiple healthcare providers collaborating, interacting and communicating with each other, one complication occurs when administrations need to exchange context-specific information between different contexts. For instance, a healthcare provider tries to query a patient's medical record from another healthcare provider. Further, the exchanged information needs to be uniquely identified. Recall that the same global identifier should not be shared directly between contexts for privacy reasons. Since interoperability from one context to another is desirable but not yet feasible, a service for information interoperability is necessary.

Whenever information is exchanged cross-context, an identifiers mapping and conversion is required. As investigated in previous work, the translation must be performed by a trusted third party which is available for all the communicating contexts [3]. Accordingly, the goal of the infrastructure we propose is to include a new service managing identifiers in e-health, which is compatible with all internal systems of healthcare providers, and that translates context-specific information exchanged between different healthcare providers.

PDP_A	H_A 's policy decision point
PDP_B	H_B 's policy decision point
IDP_A	H_A 's identity provider
IDP_B	H_B 's identity provider
$DocID_A$	Doctor D 's context-specific local ID in H_A
$DocID_B$	Doctor D 's context-specific local ID in H_B
PID_A	Patient P 's context-specific local ID in H_A
PID_B	Patient P 's context-specific local ID in H_B
GID_D	Doctor D 's global ID
GID_P	Patient P 's global ID

TABLE I
NOTATIONS AND ABBREVIATIONS

Figure 2 presents a cross-context communication between two healthcare providers in an e-health system. The functional components in each healthcare provider are: a file repository to store medical documents connecting to system portals, an identity provider that offers identity management services, such as issuing and converting local identifiers, and a policy decision point (PDP) as part of the security service to interpret access control rules for authentication and authorization. Each healthcare provider is responsible to manage and issue local identifiers for its users within its context. Accordingly, a healthcare provider cannot prevent other healthcare providers from issuing local identifiers in a particular context. When healthcare providers communicate, information is exchanged through a mediator, which is a trusted party accessible for both healthcare providers. The mediator translates context-specific information exchanged between the communicating parties. Now we focus on the building blocks of the entities involved in a communication. However, how information is exchanged exactly depends on applications. In Section IV, we will provide a scenario as an example to explain how context-specific information can be converted and exchanged among healthcare providers through a mediator.

IV. CROSS-CONTEXT IDENTIFICATION AND AUTHORIZATION IN E-HEALTH

A. System model

Assume that a patient P has received medical treatments from a generic hospital H_A and a psychiatric hospital H_B . Consider the scenario that a doctor D , at a hospital H_A , requests the patient P 's medical records from the two hospitals H_A and H_B . To preserve patient's privacy, the system ensures that patient's medical records can only be retrieved legitimately by authorized parties, such as a doctor with a given consent. Accordingly, access control rules are implemented by the policy decision point PDP at each hospital. The mediator M , for the communication between H_A and H_B , interacts with the hospitals' identity providers IDP_A and IDP_B to translate context-specific identifiers. Notations and abbreviations are depicted in Table 1.

B. Proposed protocol

As shown in Figure 2, information is transferred among different healthcare providers according to the following steps:

- 1) In order to retrieve medical records of a patient P , a doctor D logs in at a terminal in the hospital H_A using his user name and password.
- 2) The identity provider IDP_A of the hospital H_A provides the doctor a token, containing the doctor's local ID $DocID_A$ and the patient's local ID PID_A .
- 3) The doctor sends this token to the hospital's repository.
- 4) The repository sends the token to H_A 's security server for authentication and authorization.
- 5) According to H_A 's access control policy decision point PDP_A , the request can be either permitted or denied.
- 6) If the doctor's request is permitted, H_A 's file repository sends the patient's medical record to the doctor.
- 7) As requested by the doctor, H_A 's repository queries H_B 's repository with the doctor's local ID $DocID_A$ and the patient's local ID PID_A .
- 8) H_B 's repository sends the request to H_B 's security server for authentication and authorization.
- 9) H_B 's access control policy decision point PDP_B requests the mediator M for the translation of the local IDs $DocID_A$ and PID_A .
- 10) M sends $DocID_A$ and PID_A to H_A 's identity provider IDP_A for conversion. (see Section. III-A)
- 11) After H_A 's security server authenticates M , IDP_A converts the doctor's and patient's local IDs $DocID_A$ and PID_A to their global IDs GID_D and GID_P , and transfers the global IDs back to M .
- 12) M then sends the global IDs GID_D and GID_P to H_B 's identity provider IDP_B , to request the local IDs from H_B . (see Section. III-A)
- 13) After H_B 's security server authenticates M , the identity provider IDP_B issues and sends the doctor's and patient's local IDs $DocID_B$ and PID_B back to M .
- 14) M replies H_B 's policy decision point PDP_B with the doctor's and patient's local IDs $DocID_B$ and PID_B .
- 15) Then PDP_B specifies the access control rule based on the relation between the doctor's identifier $DocID_B$ and the patient's identifier PID_B . Accordingly, the security server permits or denies H_A 's request.
- 16) If the doctor's request from H_A is permitted, H_B retrieves the patient's medical record from its repository, and transfers the medical record to H_A 's repository.
- 17) Finally, H_A 's repository replies the doctor with the retrieved patient's medical record from H_B .

V. RELATED WORK

Over the past years, various popular user-centric identity management systems have been developed, such as Liberty Alliance [4], Shibboleth [5], CardSpace [6], and Idemix [7]. The Liberty-like federated identity management systems mainly utilize a trusted central service provider as a trusted third party, to maintain a look-up table of users' identifiers of different contexts. When information is transferred from one context to another, the central service provider uses the directory table to facilitate single-sign-on. The drawback of this approach is that the directory table can easily be the target of attackers;

Fig. 2. The protocol of the cross-context query of medical records in an e-health system.

once the security of the directory is compromised, the whole system's security is compromised. Our solution provides three improvements. First, instead of consulting a directory table in a trusted central service provider, we use a mediator to offer mediation services for each communication between contexts. Second, instead of one central service provider, the entity that functions as a mediator may vary. Third, the mediator doesn't maintain a directory table containing all the user's identifiers but in each communication, the mediator will interact with the ID provider of each communicating party for identifier translation. Furthermore, the identifiers issuing and converting processes by ID providers are controlled by cryptographic functions secured by secret keys. Hence, security of the interactions between different contexts in the architecture is guaranteed by cryptographic functions and security of the cryptographic keys.

In the literature, some identity management schemes a user-centric approach have been proposed for e-health. Peyton et al. [8] use a simple ePrescription scenario to analyze the business and technical issues in a Liberty Alliance federated IDM framework. They discuss the potential impact of privacy compliance on three existing components of the framework, namely, Discovery Service, Identity Mapping Service and Interaction Service; and propose a fourth component Audit Service to address potential privacy breaches in Liberty Alliance. Au and Croll [9] recently proposed a consumer-centric IDM framework for distributed e-Health. The healthcare consumer maintains a pool of pseudonym identifiers in a personal secure device, such as a smart card. Without revealing consumer identity, health record data from different distributed medical databases can be collected and linked together on demand. In particular, pseudonym identifiers are cryptographic keys, that are 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.

VI. CONCLUSION

Two conflicting forces are present in the federated healthcare scenario: inter-operability between healthcare providers sharing medical document must coexist with the privacy

requirements protecting the patients. State-of-the-art solutions provide inter-operability by means of a mediator component that maintains a look-up table storing all local identifiers across contexts. In this architecture, privacy is potentially at stake because of the data aggregation threat. An attacker can get to (illegitimately) own the information that is used by the mediator in order to map references across contexts. In this circumstance, the attacker is in a privileged position to correlate patients information on a large scale.

This paper improved the above scheme by introducing an algorithm to issue reversible local identifiers that does not require any look-up information to be maintained by the mediator component. As a consequence, the overall solution reduces the sensitivity, privacy-wise, of the mediator component. Further, instead of having one, fixed central service provider, the entity that functions as a mediator may vary. Further, this paper investigated, by means of a working example, the interplay between the proposed privacy-friendly identity scheme and the authorization mechanisms that are typically in place in a federated healthcare scenario.

REFERENCES

- [1] "HIPAA administrative simplification: Enforcement; final rule. United States Department of Health & Human Service." *Federal Register / Rules and Regulations*, vol. 71, no. 32, 2006.
- [2] D. Eysers, J. Bacon, and K. Moody, "OASIS role-based access control for electronic health records," *IEEE Proceedings Software*, vol. 153, no. 1, pp. 16–23, 2005.
- [3] Modinis, "Modinis study on identity management in e-government – the conceptual framework."
- [4] Liberty, "Liberty alliance project whitepaper: Personal identity," 2006.
- [5] T. Scavo and S. Cantor, "Shibboleth architecture, technical overview," Internet2/MACE, Tech. Rep., 2005.
- [6] CardSpace, "Windows cardspace," 2007.
- [7] J. Camenisch and E. V. Herreweghen, "Design and implementation of the idemix anonymous credential system," IBM Research Division, Tech. Rep., 2002.
- [8] L. Peyton, J. Hu, C. Doshi, and P. Seguin, "Addressing privacy in a federated identity management network for ehealth," in *World Congress on the Management of eBusiness (WCMeb)*, 2007.
- [9] R. Au and P. Croll, "Consumer-centric and privacy-preserving identity management for distributed e-health systems," in *Hawaii International Conference on Systems Science (HICSS)*, 2008.