A PATIENT CENTRIC AND PRIVACY ENHANCING PROVABLE SECURE ARCHITECTURE ON IHE-XDS

Slamanig D[1], Stingl C[1]

Abstract

In recent years many countries have installed eHealth initiatives and working groups in order to develop strategies to harmonize the exchange of health related information using open environments. A central aspect of eHealth is the electronic healthcare record (EHR) which integrates all relevant medical information of a person and represents a lifelong documentation of the medical history. Thereby the most promising approach is the so called virtual EHR, where documents remain in their respective information systems and a central index holds references to these documents. As proposed in Austria, this document management can be realized by means of the IHE-XDS specification. According to the Austrian strategy on IHE-XDS, we introduce a provable secure architecture which enables a patient to moderate her EHR and to authorize other persons who are directly involved in the treatment of the patient. Since eHealth portals can be accessed via the Internet, additional security and privacy issues arise that have to be considered carefully. Besides traditional security properties, we are also focusing on a less-considered threat, that we call the disclosure attack. This means that a person can be asked or even forced to open her EHR, e.g. at a job interview. In this architecture we use a concept of identity management based on pseudonyms that successively realizes a cryptographic strong and provable secure authorization concept. Moreover it prevents disclosure attacks and helps to reduce the amount of utilizable metadata and thus improves the protection of patients' privacy.

1. Introduction

In recent years the information and communication technology has improved the quality of health care systems by means of an integrated and highly qualitative provision of medical services. In many countries eHealth initiatives and working groups have been installed in order to develop strategies to harmonize the exchange of health related information using open environments (e.g. the eHealth initiative in Austria). A central aspect of eHealth is called the electronic healthcare record (EHR) which integrates all relevant medical information of a person and represents a lifelong documentation of the medical history. In general, there are two types of EHR architectures. One type is called the virtual EHR, where the documents remain in their respective information systems, e.g. hospital information systems (HIS), and the system provides an index which holds references to these documents. According to IHE-XDS the latter system is referred to as document-registry and so called document-repositories are linked with the information systems (see *Figure 1*). The second architecture is called the central EHR, where all documents are collected in a central repository. Weighting out the pros and cons of these two approaches it is highly probable that there will

[1] Department of Medical Information Technology, Carinthia University of Applied Sciences, Klagenfurt, Austria

be far more virtual EHR solutions in the future because information systems are widespread in healthcare and it is unlikely that they will be replaced by one centralized information system. Considering eHealth, one can recognize an alteration to a patient centric approach of health care, where the patients herself should also be able to moderate her EHR. This is in general known as the informational self determination which means that patients can decide for themselves whom to divulge and disseminate their medical data to. For this reason and because of the nearby ubiquitous availability of the Internet, the moderation of EHRs can be implemented by means of web based applications, so called eHealth portals. By virtue of their sensitive character it is crucial that medical data can only be accessed by the patient herself and persons who are directly involved in the treatment of the patient. Since eHealth portals are available on the Internet additional privacy issues arise that have to be considered carefully. *Figure 1* illustrates an eHealth portal which can be accessed across the Internet in order to manage a virtual EHR.

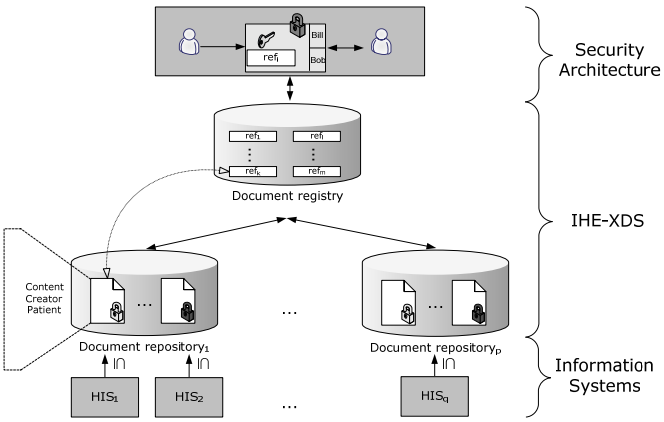

Figure 1: Security Architecture on IHE-XDS

Today most Internet based services support the confidentiality and integrity of transmitted data by means of transmission encryption (SSL/TLS). But these protocols do not protect the privacy of (personal related) data, since an internal attacker (insider, administrator) or a foreign entity which penetrates the portal (hacker), can easily acquire plaintext data even if transmission encryption is used. Therefore, documents need additionally to be encrypted at the application layer in addition to transmission encryption (see *Figure 1*) in order to prevent unauthorized access to documents at the eHealth portal respectively. This end-to-end security at the application layer is in accordance to the existing Austrian laws. Since an eHealth portal constitutes a central access point, we want to point out that there exist additional attacks which can expose potentially compromising data of a person. It must be emphasized that even if content encryption at the application layer is used, the sole observation of the metadata of an eHealth portal, e.g. access logs, operations conducted in the system, can violate the privacy of users, and thus, there emerges the need for methods that improve the users` privacy [5]. Especially when considering the metadata of a system possible attacks and privacy aspects emerge which are rarely discussed in the literature so far. Besides, since the EHR of a person can be accessed time and location independent a person can be asked or even forced to present it, e.g. at a job interview. Although the enforced disclosure would be illegal the disclosure attack is highly realistic in our opinion. Our architecture uses a concept of identity management based on pseudonyms [1] that helps to reduce utilizable metadata and thus enormously improves the patients` privacy.

1. 1. Related Work

Demuynck et al. address privacy issues in EHRs and present protocols based on digital credentials [1,5] which enable trusted doctors to anonymously access all EHRs in a central system and to revoke their anonymity in case of abuse or fraud [2]. Thus for any observer it is impossible to track down which patient is treated by which doctor. However, the main drawback of their protocols is that a patient is unable to selectively disclose data to different doctors (a doctor has access to all or none of the patients documents). Mathe et al. [4] provide a modelling environment capable of representing a functional patient portal, whereas they try to overcome security and privacy threats by automatic analysis of their models. However, they mainly focus on the service itself and do not address privacy issues that these services do entail. In general authorization in eHealth portals is realized by means of access control lists (ACL), role based access control (RBAC) as well as identity based access control (IBAC) (cf. [3]). We denote this level of security as application security which does only guarantee a certain level of overall security. This means that the security massively depends on the implementation of the application (cf. [7] use privacy policies which are implemented at the application layer). Thus, an attacker who is able to manipulate the application level mechanisms can gain full access to all documents, their content in absence of content encryption and is furthermore able to determine the set of qualified users for the documents.

2. Architecture

As illustrated in *Figure 1* we assume that end-to-end encryption at the application layer (content encryption) is applied to all documents of the document repositories. The shares, each containing the respective content encryption key of an encrypted document for a specific person, are stored in the eHealth portal. Thereby all cryptographic operations are performed by client computers. Thus the content of a document in plaintext is solely accessible at the client computer and for an authorized user respectively. It can be assumed that the content of an EHR is hierarchically structured. Besides privacy issues one main aspect of the presented concepts is their straight forward and high-performance implementation. In contrast to the above mentioned application security, the security of the presented architecture does not depend on the application level of the portal. We denote this as provable security, whereas the term is understood in that way, that the only way to defeat the system is to break the underlying atomic primitives (e.g. a public key cryptosystem). In our concept this level of security can be obtained by realizing access rights for documents as so called shares. A share is given from a grantor *Bill* to a grantee *Bob*, contains a reference to an encrypted document in a respective document repository, a cryptographic key for this document and is encrypted with the grantees public key (see *Figure 2*). We will now successively improve the amount of privacy provided by the system by introducing pseudonymization and identity management respectively.

2. 1. Pseudonymization

A serious drawback of the approach illustrated in *Figure 2* (a) is the linkability of shares and respective users (*Bill* and *Bob*). Hence any observer (an insider, etc.) is capable of determining e.g. the number of shares assigned to a specific user. In order to prevent these attacks we use pseudonymization methods in a first step. Pseudonymization of person related data (U_i, ref_i) is the process of replacing the person identifier ("*Bill*" and "*Bob*") U_i by the value $P_{Ui} = E_k(U_i)$, where E_k is an encryption function with a corresponding secret key k. Since k is kept secret it is practically impossible to invert E_k without the knowledge of k and thus compute U_i given the value P_{Ui}. However, an instance which is in possession of k can compute $D_k(P_{Ui}) = U_i$ using the corresponding decryption

function. Hence (P_{Ui}, ref_i) can not be linked to U_i anymore. In *Figure 2* P_{Bob} represents a pseudonym of *Bob* which is unlinkable to him.

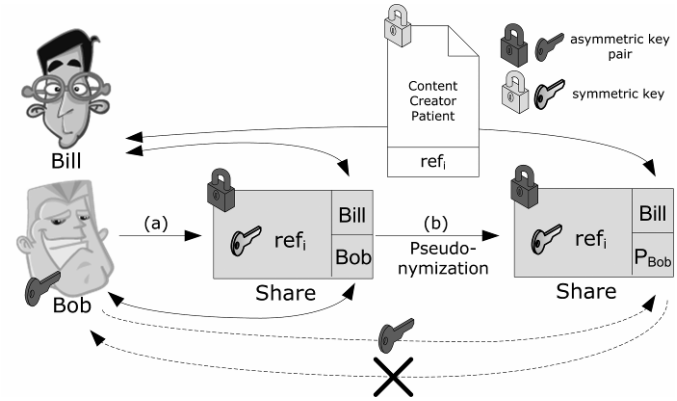

Figure 2: Bill creates a share for Bob.

A second possibility to realize pseudonyms and what we prefer is to realize pseudonymization in an eHealth portal by letting every user U_i randomly choose a second identifier P_{Ui}, i.e. a pseudonym. This pseudonym is used by her to identify her shares (see *Figure 2 (b)*). In order to prevent the linkage between a user and a pseudonym the pseudonym is solely stored in an encrypted fashion, $E_{kUi}(P_{Ui})$, in the user repository. The unlinkability holds, since P_{Ui} is independently chosen from U_i and furthermore $E_{kUi}(P_{Ui})$ can only be inverted by U_i, who holds the corresponding key k_{Ui} (which can be derived from a appropriately chosen password or pass phrase defined by U_i).

2. 2. Disclosure Attack

Since the EHR of a person can be accessed time and location independent, besides attacks from eavesdroppers or insiders, a person can be asked or even forced to present it, e.g. at a job interview or an insurance contract conclusion. The enforced disclosure is of course illegal, but nevertheless it can not be avoided basically by means of legal regulations and/or (cryptographic) standard techniques. There are several groups of people that can be highly interested in the actual state of health of a specific person, e.g. (potential) employers. Besides the enforced disclosure an employer could convince a job candidate to open her EHR using plausible arguments (e.g. the candidate needs to be in a really good physical condition in order not to harm her health and may "voluntarily" prove this by opening her EHR.). Additionally the health awareness in our society could open the possibility that healthy people disclose their EHR voluntarily in order to take advantage at the job hunt. This would discriminate people who are not willing to present their EHRs. Thus there exists the need for a mechanism in order to hide highly compromising information (e.g. a cured burn out depression) from people who don't need to know that information at all. A second important property besides hiding the information is their plausible deniability. If a user is in an awkward predicament to disclose his EHR she is able to deny the existence of compromising documents and nobody is able to prove that black is white.

2. 3. Identity Management

As stated above, even though pseudonymization reduces the information an attacker can obtain enormously, there still exists the much more precarious disclosure attack which can lead to the disclosure of the complete EHR of a person. As countermeasure we propose the use of identity management. In this context identity management can be described by means of dividing the identity of a person into sub-identities $I = \{I_{pub}, I_1, ..., I_k\}$ whereas each of these is represented by a user chosen

pseudonym, e.g. P_{Ik} , and all identities are unlinkable. A user can assign a subset of her EHR to each of these sub-identities, i.e. create shares for this sub-identity. Thereby, these subsets do not need to be disjoint (see *Figure 3*) and additionally these sub-identities open the opportunity to clearly structure the content of the EHR.

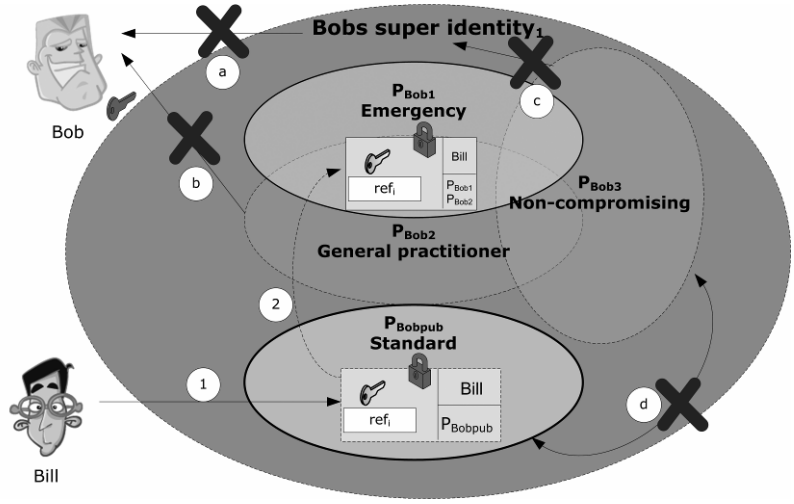

Figure 3: Identity management

Typical identities could be the so called standard identity, the emergency identity, identities for several medical fields and a non-compromising one. The standard identity I_{pub} is used by other users to grant shares and solely represents a temporary storage of shares which are immediately distributed to other sub-identities. The emergency identity should on the one hand include all emergency relevant information but on the other hand must not contain any compromising data (enforced disclosure!). Shares to documents of this identity are given to all emergency doctors of the system. Every identity for a medical field contains all relevant data which are necessary for respective consultations. The non-compromising identity does not include any compromising data and can be presented in case of an enforced disclosure. All identities $I_1, \ldots, I_k$ of a user are pseudonymized. Once again it must be mentioned that only the user herself has a complete picture of all her identities, is able to present the content of one identity and additionally can plausibly deny the existence of any other of her identities. This means, that even if she opens one of her identities it is impossible to obtain any information about the existence of other identities of her. In order to gain a complete picture of her identities this identity concept offers the possibility to optionally create so called super-identities (see *Figure 3*), whereas these are indistinguishable from sub-identities in the system. The purpose of a super-identity is to include sub-identities in order to provide a convenient access to these using only one authentication. It must be emphasized that the authentication information for every identity needs to be completely different. Hence it must not be possible to derive one from another.

2. 4. Implementation

For the implementation of our architecture we used the Microsoft Internet Information Server and the Oracle DBMS and the .NET Framework. In our first prototype we focused on mobile devices running a Microsoft Windows Mobile 5.0 or 6.0 and the .NET Compact Framework 2.0. The motivation for the use of mobile clients resulted from a research project in the field of personal health monitoring where low cost and mobility are fundamental requirements.

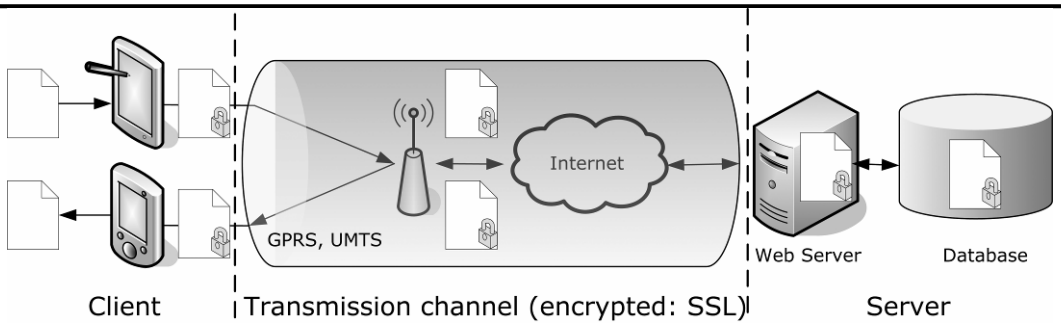

Figure 4: Schematic illustration of our first prototype

However, the architecture is not limited to the use of mobile devices and can easily be adapted to standard clients. The data exchange was realized by means of Web Services. On the server platform as well as on the client platform we used the C# programming language. As cryptographic framework we used the cryptographic standard components of .NET, including the Advanced Encryption Standard (AES), the RSA algorithm and suitable Pseudorandom Number Generators.

3. Conclusion

In this paper we introduced a security architecture on IHE-XDS using identity management and pseudonymization in order to improve the privacy of patients. Furthermore we present the so called disclosure attack which could be mounted by any person due to the existence of a time and location independent accessible EHR. In order to prevent these attacks one can apply the presented methods which enables a patient to open a predefined set of non-compromising data if she is enforced to present her EHR. Furthermore the concept enables her to plausibly deny the existence of additional identities. Besides pseudonymization and identity management there are several mechanisms like anonymous authentication, obfuscation techniques [6] and the sharing of identities between users (e.g. relatives) which are not addressed in this paper but considered to be highly relevant in context of medical data. Within a research project in the field of personal health monitoring the proposed security architecture was implemented and currently we are evaluating the application concerning the usability and the performance of the protocols.

4. References

[1] D. CHAUM. Security without identification: transaction systems to make big brother obsolete. Commun. ACM 28.

[2] L. DEMUYNCK, B. DE DECKER. Privacy-Preserving Electronic Health Records. Proc. of CMS 2005, pp. 150-159.

[3] A. Ferreira, R. Cruz-Correia, L. Antunes, et al. How to Break Access Control in a Controlled Manner. In Proc. of CBMS. IEEE Computer Society, Washington, DC, pp. 847-854, 2006.

[4] J.L. MATHE, S. DUNCAVAGE, J. WERNER, et al. Towards the security and privacy analysis of patient portals. SIGBED Rev. 4, 2 (Apr. 2007), pp. 5-9, 2007.

[5] D. SLAMANIG AND C. STINGL. Privacy Aspects of eHealth. In Proc. ARES 2008. pp. 1226-1233, IEEE Computer Society, Barcelona, Spain, 2008

[6] D. SLAMANIG, C. STINGL, G. LACKNER, AND U. PAYER: Schutz der Privatsphäre in einem webbasierten Multiuser-System. In Proc. of DACH-Security 2007, pp. 98-110, IT-Verlag 2007.

[7] G. YEE, L. KORBA, AND R. SONG. Ensuring Privacy for E-Health Services. In Proc. of ARES'06. IEEE Computer Society, pp. 321-328, 2006.