Privacy in Commercial Medical Storage Systems

Mehmet Tahir Sandıkkaya[1,2], Bart De Decker[3], and Vincent Naessens[4]

[1] Istanbul Technical University,
İTÜ Bilgisayar Mühendisliği Bölümü, TR 34469, İstanbul, Turkey
sandikkaya@itu.edu.tr
[2] Katholieke Hogeschool Sint-Lieven,
Gebroeders De Smetstraat 1, 9000, Gent, Belgium
[3] Katholieke Universiteit Leuven,
Celestijnlaan 200A, 3001, Heverlee, Belgium
bart.dedecker@cs.kuleuven.be
[4] Katholieke Hogeschool Sint-Lieven,
Gebroeders De Smetstraat 1, 9000, Gent, Belgium
vincent@msec.be

Abstract. Today, people grow older than some decades ago. This inevitably leads to an increasing number of commercial players in the healthcare domain. Privacy is a major concern in many eHealth application, especially when sensitive personal data is stored in databases. This paper presents a secure, fair and privacy-preserving solution to enforce the patient's privacy preferences on his or her personal medical records. The proposed cryptographic tools and protocols are thoroughly explained. Moreover, a prototype implementation validates the concept. Finally, it is shown that a convenient, modular and generic system based on lightweight cryptographic primitives can be realized as proposed.

1 Introduction

As the present day population is growing older than ever before [16], the need for better healthcare services is drastically increasing [16]. Though companies are seeking a role in the scene, traditions, strict liabilities and governmental healthcare institutions (GHI) constrain commerical initiatives. Furthermore, they are discouraged due to strict security and privacy regulations [13] and legislation on personal medical information (PMI). Given these considerations, this paper presents a generic and applicable PMI storage approach for commercial healthcare providers (CHPs).

Since the ambient medical assistance and monitoring systems [11] [12] cover a wide range of challenges, our solution is exemplified in this domain. In such systems, patients are continuously monitored while data is gathered by an autonomous device. The device can be seen as a sink that receives data from several sensors in the house, or as a mobile device that receives input from sensors attached to the patient's body. The gathered data is then sent to the CHP headquarters for further analysis. PMI records can only be released with the patient's consent or by jurisdiction. These permissions can be specified in a privacy preference language.

M. Szomszor and P. Kostkova (Eds.): E-Health 2010, LNICST 69, pp. 247–258, 2011.

The PMI is analyzed by software and/or by data nurses at the headquarters. A report and a summary of monitored values are generated and stored to support the medical professionals' diagnosis. In most scenarios, CHPs mediate between the patients and the medical professionals.

According to current regulations, PMI records may only be viewed by the GHIs and by the patient's physician [13]. Thus, the above described situation does not comply to the regulation if the PMI records are kept unencrypted in the commercial databases. On the other hand, if the patients' privacy preferences were cryptographically enforced in commercial databases, CHPs would be able to store and manage the PMI. This approach has two clear advantages. First, it offers opportunities for commercial eHealth organizations, enabling outsourcing PMI storage and management. Second, it centralizes data storage, easing adaptation of security and privacy mechanisms on the PMI.

When the commercial players are invited to the scene, issues of accessibility of emergency physicians, patients' and physcians' privacy concerns, and statistical analysis should be considered in detail.

Emergency physicians must have easy access to the PMI records in the CHPs' databases without the patient's prior consent. For instance, the first emergency physician reaching the location of a car accident must be able to receive the medical history of the injured people by using proper identification, even if such PMI records would be inaccessible under normal circumstances.

Furthermore, the patients' and the physicians' privacy concerns should be considered by CHPs. In the proposed solution, independent from CHPs' behaviour, privacy is preserved by two measures. First, database access is fairly monitored. Appropriate logging prevents repudiation and makes entities accountable. Second, the identities and relationships of any set of users is kept secret. Nevertheless, it is possible to disclose the hidden identities to a legal authority in case of a dispute.

Statistical analysis has two sides. On one hand, the PMI is strictly private as its name states. It seems reasonable to let people decide whether to share their medical records with researchers. Moreover, statistical data may ease violating privacy. On the other hand, people are not allowed to keep their PMI secret if such behaviour may cause public risks. Therefore, statisticians (researchers), as trustworthy officers of GHIs, are legally allowed to view patients' medical records. Based on these principles, statistical analysis and research access is rendered necessary for two reasons in this study. First, epidemics can be foreseen and preventive measures can be carried out. These studies end up with better public healthcare policies. Second, patients' non-identifiable PMI can be analyzed to detect critical health conditions. Such critical conditions, as in quarantine incidents, may lead to identity disclosure by GHIs.

The rest of the paper is organized as follows: in section 2 the primitive building blocks will be described and assumptions will be explained. Section 3 presents the proposed solution. Evaluation and related work are explained in sections 4 and 5. Finally, this paper ends with general conclusions.

2 Primitives and Assumptions

2.1 Cryptographic Primitives

Apart from the symmetric cryptography and PKI infrastructure, some more advanced cryptographic primitives are used in the paper. They are explained in this section. A *pseudonym* defines a random number that corresponds to a real identity. In the proposed scheme, such a random number is included in a certificate as a credential, together with properties linked to the pseudonym. The *pseudonym certificate* is signed by the entity that issues the pseudonym to the identity.

A *cryptographic hash chain* is an iterative application of a cryptographic hash function to an initial message. A master key, which is a symmetric key of n bits, is used as the initial message in a hash chain. By selecting a hash algorithm that produces an output of n bits, a series of keys can be generated. To illustrate this principle: assume $hash(m) = h$ means that the hash of message m is h and the z^{th} power of the $hash(\bullet)$ function implies z successive repetitions. Thus, $hash^2(m)$ defines $hash(hash(m))$. So, to calculate the z^{th} secret key K_z, based on the master key M, one can use the equation $K_z = hash^z(M)$ [8] [15].

This approach has a very useful advantage for the proposed scheme. If someone knows (1) any of the intermediate keys, (2) the hash algorithm used, (3) the index of the intermediate key (say, z^{th}), he or she can calculate all the following keys. However, the desired feature in the proposed scheme is its reverse. When an intermediate key is known, all previous keys must be computable. To achieve this, upper limits are introduced and the indexes are subtracted from the upper limit, expressed formally in a formula: if L indicates a limit, key R_k will indicate K_{L-k}. The limit should be relatively large to avoid reaching $R_L = K_{L-L} = M$. On the contrary, when the limit increases, the computation time increases too. However, some of the intermediate hashes can be cached as in rainbow tables [7]. The number of intermediate caches is a trade-off between computation time and storage.

Similar to cryptographic hash chain, an *indexed symmetric key generator* can be defined with a few modifications. The exclusive or (XOR) of the master key M and an index t is fed to the hash function to produce a symmetric key O_t, i.e. $O_t = hash(M \oplus t)$. Remark that the key is not dependent on any other indexed key and the index is not necessarily a counter. A timestamp can function as well as any ordered index.

2.2 Assumptions

Some assumptions were made through the text for readability. Cryptographic hash functions, random number generators, symmetric cryptosystems and asymmetric cryptosystems are assumed to be flawless in the proposed scheme. All entities have valid proofs of their identities and their qualifications. The communication channels are secure and authenticated.

Patients use a certified trusted device with communication, encryption and key management capabilities. Although such devices can manage many patient

profiles and many healthcare providers based on their settings, devices will be limited to one user and one provider for the sake of clarity in the rest of the text. Note that the protocols show only the main success scenarios.

Moreover, the security and privacy of the PMI kept in the GHIs or statistic institutions are intentionally left out of scope, as we assume CHPs as the application domain.

3 Proposed Scheme

3.1 Privacy Requirements

PMI stored in the database must be unidentifiable. Otherwise, identifiable PMI can be sold to third parties (e.g. to an insurance company) or rumours and speculations about a patient can be spread.

PMI records of the same patient must be able to be linked for certain purposes. If the PMI of a unique patient is not complete, physicians cannot see the full picture during diagnosis and treatment. Moreover, researchers and statisticians often study unidentifiable PMI to point out correlations between patient behaviour, diseases and epidemics.

Physicians and patient-physician relationships must remain hidden. The scheme must consider not only the patients' but also the physicians' privacy concerns. Besides, the relationship between a patient and a specific physician can disclose sensitive information. For instance, if a patient is seeing an oncologist more often than before, his or her disease can be easily guessed.

3.2 Access Control Requirements

Break-the-glass procedures. Emergency physicians rely on the medical history during an emergency action. If the PMI cannot be accessed, emergency physicians can make erroneous decisions.

PMI must be accessible for statistical purposes. Statisticians must be able to access PMI records without knowing the PMI records' owners' identities. In this type of access, PMI records can be read in bulk by the statisticians.

PMI access must be administered and monitored fairly. The data holder that keeps the PMI must not allow uncertified entities to access the PMI. A non-repudiable, fair and independent logging mechanism must record the user's actions and requests.

Patient-physician relationships must be handled easily. As the relationships between the physicians and the patients are subject to change over time the scheme must offer an efficient and dynamic mechanism to manage relationships.

3.3 Setting

In the proposed solution, the patients and physicians are pseudonymous. The PMI which is kept in the commercial databases is encrypted with two different symmetric keys and these keys are inaccessible to the CHPs. Thus, the CHPs can only see the obfuscated PMI content. However, physicians, emergency physicians and statisticians must have access to the PMI. In the scheme, two trusted parties manage key distribution They reveal the required keys to privileged entities.

One symmetric key is for fulfilling privacy requirements related to the patient-physician relationships. This key and its index are updated at each change of the relationships. Even after the key is updated, the current key allows to compute older keys. It allows the physicians to decrypt previous PMI records.

The other symmetric key is used to control the PMI access. This key's index is the current timestamp. These symmetric keys cannot be computed with respect to any other timestamp. Thus, any entity must retrieve the required key from the monitoring entity to access PMI.

3.4 Entities

In the prototype implementation, each entity is assigned exactly one role. It allows to clearly separate responsibilities and increases modularity. This separation is at the design level: multiple responsibilities can be taken by a trusted party or a single responsibility can be shared among a set of entities in a real-world application:

- *Human actors* are (P) the patient, (D) the physician, (E) the emergency physician, and (S) the statistician.
- *Trusted parties* are (I) the identity provider, which generates pseudonyms, (R) the registrar, which keeps track of the patient physician relationships, and (O) the observer, which controls access to PMI.

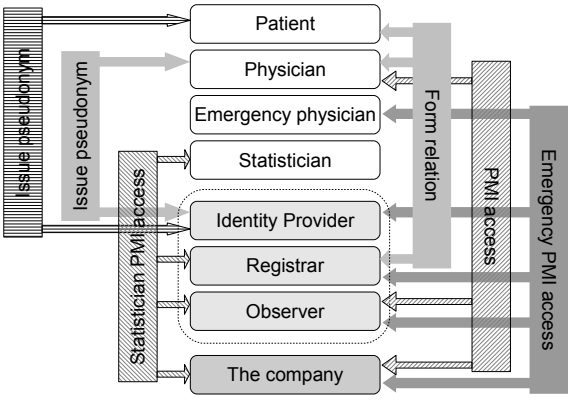

Fig. 1. Protocols and entities involved in the protocols

Table 1. Symbols and notation used in the protocols

$\mathcal{X}$	Regular certificate of X	$\mathbb{X}$	Pseudonym certificate of X
$\delta^{\mathcal{X}}$	Private key of certificate $\mathcal{X}$	$\epsilon^{\mathcal{X}}$	Public key of certificate $\mathcal{X}$
MR^X	Master relation key of X	MO^X	Master observation key of X
R_k^X	k^{th} relation key of X	O_t^X	Observation key of X at time t
$\mathbb{X}_s$	Subject of certificate $\mathbb{X}$	Pr^X	Proof of X's qualification
$sign_\delta(m)$	Signature of m with key δ	$enc_K(m)$	Encryption of m with key K
Rel_Y^X	Proof of relationship between patient X and physician Y		
$Data_{k,t}^X$	PMI record of X with relation index k and observation index t		
$Exp_{k,t}^X$	Explanation for the corresponding PMI record		
$X \xrightarrow{\mathbb{X},\mathcal{Y}} Y$	A channel that is mutually authenticated with certificates $\mathbb{X}$ and $\mathcal{Y}$		

– The *data holder* is the company (C), which keeps patients' PMI. Each record in the database consists of at least the owner's pseudonym, a relation index, an observation index, an explanation and the data itself. The latter is encrypted.

3.5 Protocols

The Physician Receives a Pseudonym Credential. D proves his or her medical qualifications to I and requests a pseudonym. I creates a pseudonym certificate and the related qualification for D.

The protocol is shown in Table 2. $\Pr^{\mathcal{D}}$ can be a signed message from an authority. For instance, $\Pr^{\mathcal{D}}$ can be $sign_{\delta^{GHI}}$('D has a valid medical diploma'). Similarly $\mathbb{D}$ defines the pseudonym certificate of D. It is signed by I to prove there exists a binding between the real identity and the pseudonym.

The Patient Receives a Pseudonym Credential. Upon P's request, I creates a pseudonym and sends a certificate that includes the pseudonym to P. P creates master keys and then sends them to both R and O.

Table 2. The physician receives a pseudonym credential

1. $D \xrightarrow{\mathcal{D},\mathcal{I}} I\ \epsilon^{\mathcal{D}},\ sign_{\delta^{\mathcal{D}}}(Pr^{\mathcal{D}})$
2. $I \xrightarrow{\mathcal{I},\mathcal{D}} D\ \mathbb{D},\ Pr^{\mathbb{D}}$

Table 3. The patient receives a pseudonym credential

1. $P \xrightarrow{\mathcal{P},\mathcal{I}} I\ \epsilon^{\mathcal{P}},\ sign_{\delta^{\mathcal{P}}}(P)$
2. $I \xrightarrow{\mathcal{I},\mathcal{P}} P\ \mathbb{P}$
3. $P \xrightarrow{\mathbb{P},\mathcal{R}} R\ MR^{\mathbb{P}}$
4. $P \xrightarrow{\mathbb{P},\mathcal{O}} O\ MO^{\mathbb{P}}$

Table 4. Forming a patient-physician relationship

1. $P \xrightarrow{\mathbb{P},\mathbb{D}} D\ k,\ \mathrm{R}_k^{\mathbb{P}},\ \mathrm{Rel}_{\mathbb{D}}^{\mathbb{P}}$
2. $P \xrightarrow{\mathbb{P},\mathcal{R}} R\ k$

Table 5. The patient sends data

1. $P \xrightarrow{\mathbb{P},\mathcal{C}} C\ \ enc_{R_k^{\mathbb{P}}}(enc_{O_t^{\mathbb{P}}}(Data_{k,t}^{\mathbb{P}})),\ Exp_{k,t}^{\mathbb{P}}$
2. $P \xrightarrow{\mathbb{P},\mathcal{O}} O\ t$

The protocol is shown step by step in Table 3. In the table, $\mathbb{P}$ defines the pseudonym certificate of P. $MR^{\mathbb{P}}$ defines the master relation key and $MO^{\mathbb{P}}$ defines the master observation key.

$MR^{\mathbb{P}}$ is used as the master key of a cryptographic hash chain by R. Thus, older relation keys can be extracted from the last relation key. $MO^{\mathbb{P}}$ is used as the master key of an indexed symmetric key generator by O. Thus, an observation key cannot be extracted from another observation key or timestamp.

Creating a Patient-Physician Relationship. Since we assume that P and D have met previously and they have decided to form a relationship, the protocol is initiated by P. P sends the proof of their relationship right after the authentication. He or she updates his or her relation key at each run of this protocol and shares it and its index with D. The index of the relation key is also submitted to R.

The protocol is shown in Table 4 step by step. $sign_{\delta^{\mathbb{P}}}$('$\mathbb{P}_s$ has a relationship with $\mathbb{D}_s$ till April 21, 2011.') is an example of how $Rel_{\mathbb{D}}^{\mathbb{P}}$ can be instantiated. Then, P calculates the $R_k^{\mathbb{P}}$ based on his or her master relation key and sends the current relation key together with its index. The latest index is sent to R as a countermeasure against malicious physicians to prevent them to ask R for future relation keys.

The proposed scheme supports multiple physicians related with the same patient. In that case, the relation key must be updated for each newly formed relationship to achieve forward secrecy[1]. If P forms a new relationship when he or she already has relationships with other physicians, a newly created relation key must be distributed to the other physicians. There are two options to distribute the key. (1) P himself distributes the updated key to each related physician. (2) The relation key is not distributed at once. When a physician realizes that the relation key is updated (when he or she cannot decrypt the fresh PMI records anymore), he or she requests the key from R.

The Patient Submits Data to the Company Database. The data gathered by P's sink is submitted to the company database in regular intervals as shown in Table 5. Since every access to PMI is definitely logged by O, PMI records do

[1] In this context, forward secrecy means that former physicians of a patient cannot access newly created PMI.

Table 6. The physician queries the PMI

1. $D \xrightarrow{\mathbb{D},\mathcal{C}} C$ $\mathbb{P}_s$, 'query', $Pr^{\mathbb{D}}$, $Rel^{\mathbb{P}}_{\mathbb{D}}$
2. $C \xrightarrow{\mathcal{C},\mathbb{D}} D$ $\forall k \forall t\, Exp^{\mathbb{P}}_{k,t}$, k, t

Table 7. The physician accesses the PMI

1. $D \xrightarrow{\mathbb{D},\mathcal{O}} O$	$\mathbb{P}_s$, k, t, [access], $Pr^{\mathbb{D}}$, $Rel^{\mathbb{P}}_{\mathbb{D}}$
2. $\mathcal{O}$	Observer logs $\mathbb{D}$, t, [access]
3. $O \xrightarrow{\mathcal{O},\mathbb{D}} D$	$O^{\mathbb{P}}_t$
4. $D \xrightarrow{\mathbb{D},\mathcal{C}} C$	$\mathbb{P}_s$, k, t, [request], $Pr^{\mathbb{D}}$, $Rel^{\mathbb{P}}_{\mathbb{D}}$
5. $C \xrightarrow{\mathcal{C},\mathbb{D}} D$	[response]

not need to be signed, yet they are encrypted twice. The symmetric keys used for encryption are $R^{\mathbb{P}}_k$ that indicates P's k^{th} relation key and $O^{\mathbb{P}}_t$ that indicates P's t^{th} observation key. It should be noted that, any record that belongs to P can be uniquely identified with k and t. The value t is the current timestamp and used as an index to calculate the current observation key. The explanation is kept in plain text to ease querying the PMI records.

PMI Query by the Physician. The query protocol is straightforward. D directly requests explanations from C. C responds with a set of explanations. The protocol is shown in Table 6. As D receives the explanation and the indexes during this protocol, he or she can access any relevant data afterwards.

PMI Access by the Physician. D can both read and write to the database. The read action can be exemplified by receiving previous records and the write action can be a submission of a consultation after D diagnoses P in person.

During the protocol, D sends the required proofs to O and requests permission. If the proofs are valid, O releases the appropriate observation key after logging the access. O keeps track of the index of the observation keys. Therefore it does not send the corresponding observation key if D is asking for a future or non-existing key. However, t is accepted and recorded by O where the access type is 'write' and t is the current timestamp.

In Table 7, [access] indicates the access type, i.e. 'write', [request] indicates any appropriate request, i.e. $write(enc_{R^{\mathbb{P}}_k}(enc_{O^{\mathbb{P}}_t}(Data^{\mathbb{P}}_{k,t})), Exp^{\mathbb{P}}_{k,t})$, and [response] indicates any appropriate response, i.e. boolean true.

Ending a Patient-Physician Relationship. The necessity of a relationship ending protocol depends on the aforementioned key distribution options in the "forming a patient-physician relationship" protocol shown in 3.5. (1) If the relation keys are distributed by P, a protocol will not be necessary, as the former physician will never receive the updated relation key. (2) In the second option, the relation keys are distributed by R. If this option is chosen, P sends the former physician's pseudonym to R and, accordingly, R manages a revocation list

and denies relation key requests of the former physician even if the proof of the relationship he or she holds is still valid.

It should be noted, the former physician can still have access to P's former PMI records[2] after the relation key is updated. If the former physician keeps the required keys, he or she will be able to query and retrieve P's former PMI records from the company database while the proof of the relationship he or she holds is still valid. This access is licit as former records have been seen before. Still, it can be easily avoided by managing another revocation list at the company side. In that case, P must send the former physician's pseudonym also to C.

Emergency PMI Access by the Emergency Physician. Likewise a related physician, an emergency physician needs read and write access to the PMI. During an emergency, an emergency physician could learn only the real identity of a patient. Accordingly, E needs to obtain P's pseudonym to access his or her PMI stored in the company database. The protocol is straightforward: E obtains the pseudonym, queries the company database, receives P's last relation key from R and required observation keys from O, and finally, accesses the PMI.

Statistical PMI Access by the Statistician. Statistical access protocol consists of the following steps: S queries the company database for all patients, and receives patients' last relation keys from R and required observation keys from O. Finally, he or she accesses the PMI.

4 Evaluation

Evaluation of the Requirements. The PMI is stored in the company database. From the company's point of view, a pseudonym is assigned to each PMI record and the PMI's content is inaccessible. Hence, the company cannot link stored PMI to an identity. Yet, the medical history of a unique patient can be linked. Authorized entities can access the PMI. Patients and privileged physicians can retrieve pseudonyms from the identity provider. Their relationships are also based on their pseudonyms. Emergency physicians can also access any patient's PMI with a valid proof of their duty. PMI is also accessible for statistical purposes. Statisticians can access any patient's PMI with a valid proof of their duty. Logs by the observer allow to keep track of PMI access. Finally, the relationships can be updated according to the aforementioned protocols.

Evaluation of the Security and Privacy properties. According to the data access requirements of the entities and the keys used for encrypting the PMI, some attacks can be designed. Under previously mentioned assumptions, paths to the PMI of a specific patient are shown in Figure 2. Four different paths can be determined to access any PMI record in the company database. The first one is asking the patient himself, which hardly is a threat against patient's privacy. Two other possibilities are corrupting the physician or the emergency physician.

[2] The records that the former physician has created or the records that were created before the former physician's treatment.

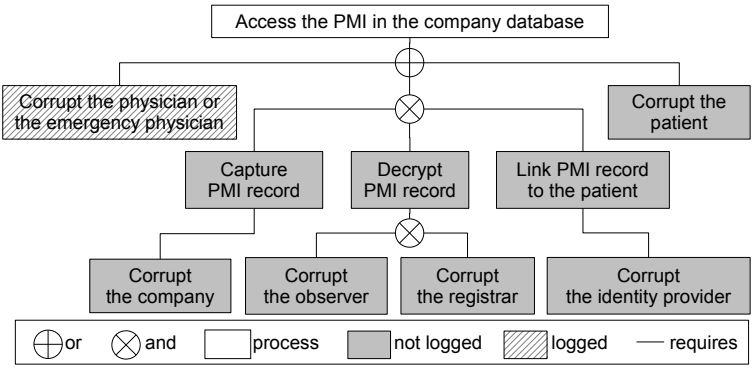

Fig. 2. Possible paths to identify a PMI record

However, even if these parties can be corrupted, the proposed scheme will still log their actions and the corrupted parties can be detected. The last option requires the corruption of several entities to retrieve the encrypted PMI, to decrypt it and to link the data to the patient. In this case, the logging is not possible as the entity that does the logging is also corrupted. Nevertheless, it is unlikely to form a coalition with four independent entities.

The proposed scheme also hides the relationships between the patients and the physicians. The relationship of two identities can only be seen if the identity provider and one of the other entities collaborate.

Besides, denial-of-service (DoS) attacks can be easily identified within the scheme while the secure communication channel assumption holds. When a party refuses to cooperate, other parties can distinguish the malicious party. For instance, if the company refuses to record the PMI appropriately, the observer's logs expose such a misbehaviour. If the observer refuses to provide the required keys, human actors suddenly notice the non-functional entity.

5 Related Work

Efforts through rule-based, service-oriented or ontology-based smart homes are on the scene for a while [17]. Besides, ambient assistance living systems [11] [12] tend to join these efforts. Both of these areas focus on building healthcare at home but they do not thoroughly study security and privacy problems. In addition, some work has already been done for secure and privacy-friendly medical data collection [14] [10]. However, not offering all of the features enlisted here.

Privacy preserving access control systems have been studied around databases for several decades [2]. Yet, this area is not explored especially for medical databases, on which [9] conducts a survey. For Hippocratic databases, which are appropriate for privacy-preserving secure medical data storage, [1] can be seen.

Digital pseudonyms and anonymous credentials, which emerged with [5] and developed seriously in time [3] [4], are other related research areas.

A similar approach has been shown in [6], which will be compared with the proposed scheme to emphasize the differences. First, in contrast to the proposed scheme, the PMI records are not encrypted in the former study. In [6], the PMI records are kept in plain text with an index only known to the physicians with whom the patient has a relationship. Therefore, to query the database and fetch the record, the obfuscated index, which points the actual record, must be known. On the contrary, in the proposed scheme, the PMI record is encrypted. Second, the former study introduces anonymous credentials, however the anonymous credential system is used as a pseudonymous system. Moreover, each access to the PMI records requires a heavy 'credential show' method which typically runs up to a few seconds. In contrast, the proposed scheme directly uses pseudonyms and symmetric cryptography and it is not noted that an access takes more than a few milliseconds in the prototype implementation. In general, cryptologic building blocks used in the proposed scheme are lightweight. Unfortunately, a quantitative comparison could not be carried out as the former system has not been implemented. Third, the former system requires a new credential when a relationship is ended, which enforces credential revocation. In the proposed scheme, such a credential or pseudonym update is not necessary. Fourth, the former system dependent to RSA, as it uses some of its distinct features. The proposed scheme is prepared to be generic and independent from any particular cryptosystem. Fifth, in the former system, access control and logging are done by different entities at different levels which renders liabilities fuzzy and makes corruption detection difficult. In comparison with the former, the proposed scheme clearly distinguishes the responsibilities of partakers. If there is a conflict, the entities' logs clearly point at the step where the problem has occurred.

6 Conclusions

This paper presented a solution to store commercial medical data. The proposed scheme preserves patients' and physicians' privacy. The entities' responsibilities in the scheme are derived from the requirements and their clear cut separation makes the architecture generic. The users and their PMI have been kept pseudonymous while break-the-glass procedures are possible. Non-identifying but linkable PMI records allow statistical data mining. The dynamic nature of patient-physician relationship management makes the scheme more applicable. The observer, which has the responsibility to collect evidence, keeps the audits of the actions that take place in the system. Therefore, possible conflicts can be resolved by jurisdiction. Finally, the concept is proven by a prototype implementation. As a result, our scheme has offered a convenient and generic solution to the commercial medical data storage problem. Our work is currently under inspection by a Belgian commercial eHealth enterprise.

References

1. Agrawal, R., Kiernan, J., Srikant, R., Xu, Y.: Hippocratic databases. In: VLDB, pp. 143–154. Morgan Kaufmann, San Francisco (2002)
2. Bertino, E., Sandhu, R.S.: Database security-concepts, approaches, and challenges. IEEE Trans. Dependable Sec. Comput. 2(1), 2–19 (2005)
3. Brands, S., Légaré, F.: Digital identity management based on digital credentials. In: Schubert, S.E., Reusch, B., Jesse, N. (eds.) GI Jahrestagung. LNI, vol. 19, pp. 120–126. GI (2002)
4. Camenisch, J., Van Herreweghen, E.: Design and implementation of the *demix* anonymous credential system. In: Atluri, V. (ed.) ACM Conference on Computer and Communications Security, pp. 21–30. ACM, New York (2002)
5. Chaum, D.: Untraceable electronic mail, return addresses, and digital pseudonyms. Commun. ACM 24(2), 84–88 (1981)
6. Demuynck, L., De Decker, B.: Privacy-preserving electronic health records. In: Dittmann, J., Katzenbeisser, S., Uhl, A. (eds.) CMS 2005. LNCS, vol. 3677, pp. 150–159. Springer, Heidelberg (2005)
7. Hellman, M.: A cryptanalytic time-memory trade-off. IEEE Transactions on Information Theory 26(4), 401–406 (1980)
8. Lamport, L.: Password authentification with insecure communication. Commun. ACM 24(11), 770–772 (1981)
9. Lin, C.-C., Duann, J.-R., Liu, C.-T., Chen, H.-S., Su, J.-L., Chen, J.-H.: A unified multimedia database system to support telemedicine. IEEE Transactions on Information Technology in Biomedicine 2(3), 183–192 (1998)
10. Maglogiannis, I., Kazatzopoulos, L.: Enabling location privacy and medical data encryption in patient telemonitoring systems. IEEE Trans. Inf. Technol. Biomed. (2009)
11. University of Illinois at Urbana-Champaign. I-living the assisted living project (August 2009), `http://lion.cs.uiuc.edu/assistedliving/`
12. University of Virginia. Smart in-home monitoring system (June 2009), `http://marc.med.virginia.edu/projects_smarthomemonitor.html`
13. The European Parliament and the Council of the European Union. Directive 95/46/ec of the european parliament and of the council of October 24, 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 L(281), 31 (1995)
14. Schartner, P., Schaffer, M.: Efficient privacy-enhancing techniques for medical databases. In: Fred, A.L.N., Filipe, J., Gamboa, H. (eds.) BIOSTEC (Selected Papers). CCIS, vol. 25, pp. 467–478. Springer, Heidelberg (2008)
15. Schneier, B.: Applied Cryptography. Wiley, New York (1996)
16. Steg, H., Strese, H., Loroff, C., Hull, J., Schmidt, S.: Europe is facing a demographic challenge ambient assisted living offers solutions
17. Wu, C.-L., Liao, C.-F., Fu, L.-C.: Service-oriented smart-home architecture based on osgi and mobile-agent technology. IEEE Transactions on Systems, Man, and Cybernetics, Part C 37(2), 193–205 (2007)