Requirements for Integrating End-to-End Security into Large-Scale EHR Systems*

Agnes Gawlik, Lennart Köster, Hiva Mahmoodi and Marcel Winandy

Horst Görtz Institute for IT-Security, Ruhr-University Bochum, Germany
{agnes.gawlik, lennart.koester, hiva.mahmoodi, marcel.winandy}@trust.rub.de

Abstract

Electronic Health Records (EHR) are becoming a growing trend in the healthcare industry. Especially when applied across healthcare organizations, EHRs provide benefits such as financial incentives and a more complete view of a patient's history. However, they also face security issues regarding the confidentiality and privacy of the patients' data, especially when the EHRs are stored at third-party providers or in the cloud. In general, confidentiality can be ensured by using cryptographic mechanisms or access control. Unfortunately, both techniques diminish the usability of the EHR if they are applied straightforwardly. Privacy and confidentiality have to be ensured in a way that does not restrict usability as it reduces the benefits of the EHR. This paper presents experiences from a requirements analysis we made during ongoing projects. We summarize the requirements for integrating end-to-end confidentiality into large-scale EHR systems in a usable fashion. In particular, we show (i) which data granularity is useful to be encrypted without interfering with access control, (ii) requirements for an authorization mechanism to access encrypted data, (iii) a privacy classification of typical metadata in EHRs, and (iv) interoperability issues that must be solved to allow for secure and usable EHR implementations.

1 Introduction

The tendency towards adoption of Health Information Technology (HIT) and particularly Electronic Health Records (EHR) has increased in most countries in recent years. The potential benefits of EHR have lead to substantial interest on the part of policy makers to speed up adoption and use across the globe [19]. Patients' medical data are stored in Electronic Health Records, which enable different care providers to retrieve all required information about the same patient [32]. An EHR is a patient-centered, longitudinal, comprehensive and prospective container of a patient's medical data [11], aiming at increasing the quality and efficiency of integrated healthcare [14]. Moreover, using EHRs can help to reduce the number of unnecessary examinations and reduce costs as physicians are being provided with a complete history of the patient's treatments.

Electronic health records may be maintained in a centralized environment at one healthcare provider or be decentralized and spread across different sites [27]. This requires interoperability of all systems communicating within the EHR network [5]. While decentralized storage avoids privacy risks from aggregating huge amount of medical information from several patients at a single place, it also requires each site to maintain the local infrastructure, providing access from other sites, and to ensure that only authorized access is possible. This could overburden healthcare providers, especially smaller organizations without their own IT department. In contrast, storing data at a central place, operated by third party providers, avoids costs of

*Proceedings of the Amsterdam Privacy Conference (APC 2012), 1st International Workshop on Engineering EHR Solutions (WEES), 2012.

maintaining the infrastructure at healthcare institutes. We expect this to become well established among healthcare organizations, and in the future even healthcare providers might opt to store medical data in the cloud. Therefore, interoperable EHR systems are assumed to provide financial incentives [33]. However, the adoption of standardized, interoperable EHRs faces a number of problems. While there are legal, financial, organizational and technical challenges and barriers [4], EHR architectures often lack appropriate security and privacy mechanisms that are practical for its users at the same time.

On the one hand, EHRs need to ensure the confidentiality of the patients' data and adhere to local privacy laws. In countries with strict privacy laws, patients have to authorize every access to their EHR by any physician [2]. In addition, the security and privacy of the medical data become even more important when the medical data is stored in the cloud [21]. The cloud provider is not part of the patient-physician confidentiality. They can be expected to be semi-honest, i.e., honest but curious, giving rise to insider threats [6]. As the cloud provider cannot be trusted, data has to be protected at all times during transport from the sender to the receiver, ensuring end-to-end security.

On the other hand, security mechanisms integrated into EHR systems must be usable in practice and should not complicate existing workflows. For instance, authorization of access by the patient imposes a demanding set of requirements onto the EHR. Authorization is a problem especially in cases in which the patient is in no mental or physical condition to authorize a physician by complex technical means. As privacy laws have to be adhered to and can not be ignored in order to increase usability, their effects on the usability of the EHR have to be identified and taken into consideration when designing measures to protect the EHR.

Security and confidentiality of data can be achieved by utilizing data encryption to protect the data itself and access control schemes to limit access to data. Both approaches however have drawbacks. In case of security solutions depending solely on access control the insiders can easily obtain data by circumventing the access control. The administrators have physical as well as OS-level access to the EHR servers, making all application-level protections such as access control ineffective. A solution depending on encryption requires key management, though we may not know the receiver of the data when it is encrypted. This is a common scenario for EHRs as it is not predictable who might become interested in a document in the future. Hence, standard public key cryptography is insufficient to satisfy this property in a usable manner.

In this paper we summarize our experiences gathered in the requirements analysis during several projects [9, 24] aiming at the development of a usable security architecture for EHR systems. We identified four new requirements which we will lay out. These requirements previously received only little attention and we argue that further research is needed in the areas of usable encryption, privacy of metadata, diversity of usage and interoperable as well as standardized security. Differences in security among organizations in the EHR network may weaken overall security, we argue to standardize security and give examples of research efforts in this direction.

2 Related work

There are a number of works focusing on privacy and security of EHRs. However, only few consider end-to-end security.

Boonstra and Broekhuis [4] identify general barriers of adopting EHRs and give a high-level view of the resulting problem areas. The barriers they identified are of financial, technical, psychological, social, legal and organizational nature as well as problems introduced by the new workflow of the acquired system. While we focus on end-to-end security of EHRs, their work

describes other problems of adoption and, as such, is complementary to ours.

Haas et al. [13] state four privacy requirements for EHRs: (i) binding privacy policies formulated by the patient, (ii) proof of the enforcement of these policies, (iii) anonymity of users, and (iv) a trust boundary which only includes parties involved in the treatment of the patient and certification authorities. While we agree with the notion of a trust boundary encompassing only parties involved in the treatment, recent headlines have shown certification authorities not to be trustworthy in general [22].

Kahn et al. [20] identify technical and security requirements for EHR systems. Their security requirements, however, focus on compliance to the Health Insurance Portability and Accountability Act (HIPAA), and the technical requirements focus on the semantic interoperability of data exchange.

Löhr et al. [23] propose a security architecture in distributed e-health scenarios that includes the platform security of the end-user devices. In their solution, though, the EHR server is trusted and protected from other servers in a centralized e-health infrastructure. Here, we focus on end-to-end security where we do not trust the EHR server.

Generally two main techniques are used to secure health records: encryption and access control, being combined most of the time. However, pure access control approaches [3, 29] do not provide the needed protection as we can not trust the EHR server where the access control is enforced.

Van 't Noordende [31] discusses the architecture of the Dutch Electronic Patient Dossier (EPD) and describes its security weaknesses and risks. The author argues that due to the decentralized storage of patient medical records in the information systems of the care providers, the encryption of medical documents is less relevant and a central access control provides sufficient confidentiality. This argument is true only if (i) the IT administrators of the information systems of healthcare organizations are fully trusted, (ii) the storage of medical records is not outsourced to (potentially untrusted) third parties and (iii) unauthorized physical access to the file servers may never happen. It is obvious that these constraints are very hard to satisfy. Therefore we believe that the data must be regardless of the storage location kept encrypted. The only location where patients' medical data may be in plain-text available is the end system of a health professional.

3 Model

We foresee the storage of electronic health records to move from the individual storage systems of healthcare providers to a (logically) centralized EHR system, operated by a (partially trusted) third party, eventually even a (fully untrusted) cloud provider. We consider the EHR provider to be semi-honest meaning he may try to get a look at the data and look for opportunities to benefit from the information. As a result, data stored at the EHR provider has to be secured at all times in a way which does not leak information to the EHR provider. The considered usage scenario is presented in Figure 1.

It is of essential importance to note that in our model the whole EHRs are not necessarily stored at one physical location and may be distributed across different providers. However as it does not affect the security requirements we represent the EHR storage servers as one single logical component in the cloud. Also in a decentralized approach, some healthcare providers may wish to outsource storage of medical data to a cloud provider. Therefore, there are no fundamental differences in the security requirements of both approaches and the represented requirements apply also to decentralized systems, in which storage of medical data at third party providers is preferred to establishment and maintenance of an own IT infrastructure.

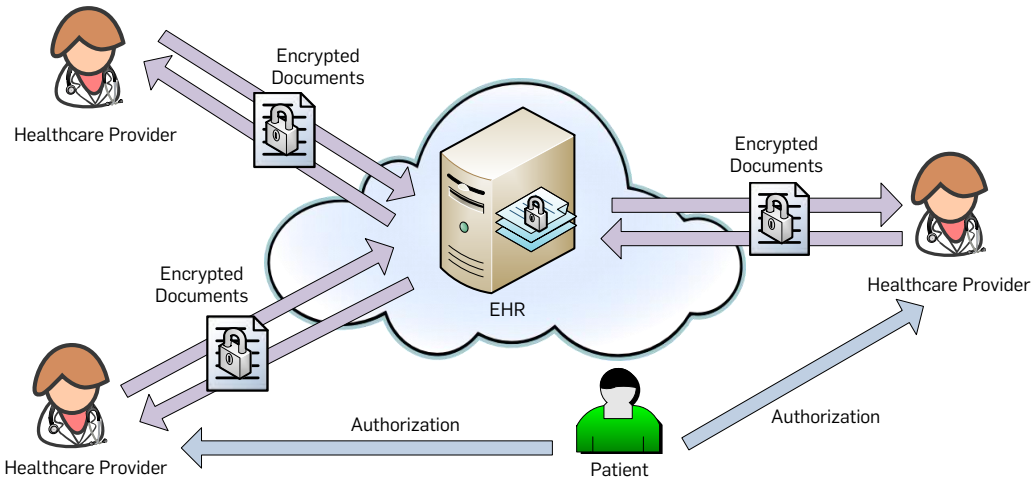

Figure 1: EHR storage at a third party site

The requirements we outline in this paper are grounded in the following concrete use case scenario. Patients authorize a healthcare provider to access their EHR for a limited time through an authorization mechanism which is described in section 6 shortly. The healthcare provider is then able to encrypt and upload a number of medical documents to the patient's EHR. Similarly another healthcare provider can, after authorization by the patient, retrieve those documents in the patient's EHR and decrypt them. An access control component is available that verifies permissions of the health professionals. A healthcare provider is able to run an arbitrary number of storage/retrieval transactions on an EHR within the validity time of the authorization granted by the EHR's owner. Moreover, the patients need not be always physically present in order to authorize access to their EHRs. The authorization can also take place remotely, enabling the healthcare providers to work on patients' EHR data without patients' physical presence at healthcare organizations.

4 General Requirements

In regard to EHR systems, a number of issues should be considered to ensure confidentiality of patients' electronic health data. In this context, patients should be able to define specific access privileges to their personal health data. Exemplarily, a patient's electronic health record should be only accessible for the healthcare personnel in charge. Furthermore, a patient must not be identified on the basis of electronic health data. As corollary, the following general security requirements arise:

- *Confidentiality of patients' health data and patients' identity:* All stored medical and personal data of a patient must be kept confidential within an EHR system to avoid violation of the patient's (data) privacy.

- *Authorized access to patients' EHR data:* Every patient has to authorize health professionals before they are able to access the patient's EHR data. In addition, health professionals

(or users in general) must be authenticated by the EHR system before they can access the patient's EHR.

- *EHR data transfer to third parties must be confirmed by the patient:* The transfer of EHR data to third parties, e.g. for secondary use, must be explicitly confirmed by the patient to guarantee data privacy. Patient's consent is mandated also by the law in some countries, e.g. § 4 BDSG[1] in Germany implies this requirement.

- *Data storage must be protected:* EHR data that is transferred or stored at sites other than a health professional's system must be sufficiently protected to avoid unauthorized access and data manipulation. At this point, achieving end-to-end confidentiality and integrity is necessary to protect the data both when it is transferred to and stored at the EHR provider's system.

Most of these security requirements can be satisfied with cryptographic mechanisms, but their integration into the EHR infrastructure has to meet usability requirements of health professionals and patients. The focus lies on the integration of security aspects into the actual workflows. Confidentiality can be achieved by applying encryption, and cryptographic signatures can ensure the integrity of electronic health data. Independent from the used algorithms, it is important that the patient is in charge of the key material that encrypts or decrypts the patient's electronic health data, whereas the health professional is in charge of signature keys to sign the data when they are uploaded or changed. Health professionals have to be authorized by the patient before they are able to read and update the patient's health data. However, this authorization may hold for a group of health professionals when they belong to the same healthcare organization (e.g., department of a hospital).

In addition, EHR solutions should be diversified by providing a number of interchangeable methods to realize an objective. It's a matter of fact that recent hardware and software technology has to be applied in new secure EHR solutions. However over-challenging people with technology and obliging them to learn a lot to operate an EHR solution causes dissatisfaction, leading the solution to failure. Easy-to-use alternative techniques should be available, if one technique fails to satisfy the requirements of a particular person or situation.

All mentioned requirements have to be considered within EHR systems to ensure protection and confidentiality of patients' medical data.

5 Cryptographic Protection

The sensitive nature of health data requires EHRs to be stored in encrypted form. Due to requirements such as end-to-end security and privacy, straightforward encryption methods cannot be employed, rather we need a secure key management mechanism including secure generation, storage, delivery and retrieval of keys. This section unveils a number of hidden problems and common mistakes in adoption of cryptographic techniques to protect health data in large-scale EHR systems.

5.1 Undirected Communication

When a health professional encrypts a data object, the decrypting receiver, any legitimate system user, may be unknown. Thus, straightforward solutions like public key encryption cannot

[1] http://www.gesetze-im-internet.de/bdsg_1990/___4.html

be employed easily. End-to-end undirected communication requires sophisticated key management mechanisms which should disseminate the key to authorized health professionals securely. Privacy, efficiency and usability also have to be taken into consideration when disseminating the decryption key. If only the patient is in possession of the encryption/decryption key, as adopted in German electronic Health Card [7], the patients' physical presence is required at every practice or hospital where their data is supposed to be encrypted or decrypted. This is not feasible or practicable in many situations. Since only the patients have the keys, access to their data is lost if the keys are lost. Storing a backup key at an untrusted third party must be strictly avoided, as it endangers the confidentiality of the patient's data [34]. A third party acting as a part of the key management solution must not be able to obtain the decryption key. Availability of the decryption key in plain-text at any system node except the decrypting receiver end must be prevented by all means.

We point out the problem of usable key management originating from the undirected nature of communications in large-scale EHR systems. A comprehensive solution should take security and privacy as much into consideration as usability. Unusable security solutions would render users in attempting to avoid or circumvent the protection mechanism (e.g., giving access keys to all people if it takes too long to authenticate users). Essentially, the security mechanisms must be integrated into existing workflows without disturbing their operation or increasing the time needed to execute.

5.2 Granularity of Data Objects and Encryption

It is necessary to choose the granularity at which data objects are encrypted in a reasonable way. The data encrypted as a unit must be small enough to avoid data redundancy and big enough to contain information meaningful to health professionals. To determine the encryption granularity of data objects one should consider both medical and technical aspects. A data object is no more decomposable if it would lose its medical meaningfulness and completeness if decomposed further. Therefore, specification of the smallest meaningful unit is a difficult task with regard to medical aspects and technical performance. The level at which encryption is performed may vary in different systems. At best the granularity of encryption should be specified based on technical standards. Considering the health communications standard Health Level 7 (HL7) [15] and its Clinical Document Architecture (CDA), one can think of a CDA document as the smallest data object that is medically meaningful. As the standard changes, the granularity of encryption may also need to change accordingly.

5.3 Relationship to Access Control

Encryption and access control create synergies for the benefit of system security. Some approaches [3, 29] find access control alone to be adequate. However, neither access control nor encryption alone can guarantee the security of EHRs by themselves. Systems such as [25, 1] use both access control and cryptographic techniques. Without access control, encrypted EHRs may be accessed by all system users regardless of their privileges. Future decryption of these EHRs could be feasible. Access control per se does not provide protection against insiders. Regardless of centralized or decentralized storage of EHR data an insider can circumvent access control and gain access to medical documents easily. If the hardware is physically accessible, e.g. as a result of hardware theft, access control can be circumvented easily similar to any other application-level protection (e.g. audit trails). Decryption of stolen data may become feasible in the future. Secure outsourcing of EHRs to the cloud cannot be realized by simply relying on access control, since the cloud provider has full control over the outsourced data.

Encryption and access control should be regarded as two independent security components in the EHR system. Data objects have to be encrypted independent from associated access control schemes. In addition, access control should not care what ever it is protecting. Neither access control nor encryption can count on the other to act correctly. Access privileges to an EHR and its contents can be set at levels of broader scope than a single document. As a result, encrypted objects inside an EHR can have the same access permissions, but no two objects with distinct access rights should be allowed to be encrypted with the same key. Once one has the decryption key to decrypt one of the two data objects it is possible to decrypt the other one as well if the access control is bypassed. Encryption should be the last protection layer in the chain of protection mechanisms due to it requiring much effort to be broken.

6 Access Authorization

As already mentioned, access to an EHR must be authorized by its owner. Though the EHR data are kept in encrypted form on the EHR storage system, the health professional who wants to gain access to a patient's EHR needs to be able to decrypt the data and eventually encrypt new data to add it to the EHR. As we consider the EHR provider as untrusted, a typical access control and authorization mechanism implemented on the central system will not suffice. Hence, giving access authorization to a health professional is essentially the same step as giving the health professional the ability to encrypt/decrypt the data. Therefore, an access authorization mechanism should be chosen that technically authorizes the usage of corresponding encryption/decryption keys.

6.1 Requirements of Authorization Mechanism

To enable patients to actively confirm their consent with transmitting their medical data to an EHR system, and to allow other health professionals to access the data later, we need an authorization mechanism that essentially authorizes the ability to encrypt and decrypt EHR data. An important requirement is that only the patient should be able to give this authorization. Hence, a conceivable solution would be to establish one or more *authorization secrets* that are kept by the patients. Whenever the patients want to allow a health professional to access their EHR, they give one of these authorization secrets to the health professional.

Essentially, this means that authorization and encryption are conjoined operations. A key management scheme may be a priori established based on the authorization secrets in order to meet the requirements and constraints outlined in section 5.1. A possible realization could be achieved by having the cryptographic keys for encryption/decryption on a smartcard that only the patient possesses; when the patient wants to authorize access to the EHR, the smartcard is connected to a terminal at the healthcare organization and the patient confirms the authorization by entering the corresponding PIN of the smartcard into the terminal.

6.2 Usability Requirements

However, while there are technical requirements on the authorization mechanism which could be solved as depicted above, there are also requirements on the usability and practicality of the mechanism in the normal workflow of health professionals.

Hence, the following additional requirements have to be met by the authorization scheme:

- *Support of various authorization procedures:* A patient has to be physically and mentally present in order to supply the necessary authorization information, e.g. a smartcard PIN,

required to access the patient's associated electronic health data. In real situations, a patient can also be represented by a representative, e.g. a relative or a confidant. Moreover, in case of an emergency that a patient may be unable to authorize, a representative should be able to act as a surrogate for the patient.

- *Simplicity besides modernity:* Authorization should be possible by simple methods as well as modern, eventually more complex, technical means. The authorization secret should be transferable by telephone, SMS, e-mail, smartphone and so on. On the one hand, older or disabled patients could be over-challenged with too complex technical means (e.g., an app on a smartphone). On the other hand, there is an upward trend towards modern devices among younger people. Hence, special care must be taken to enable authorization for elder patients with impairments, bedridden or physically handicapped patients and at the same time allow other people to use modern authorization means.
- *Time-limited EHR data access:* Medical diagnoses are often inserted or changed into the patient's EHR after the patient's medical attendance. Since a patient's EHR cannot be accessed by a health professional without a corresponding authorization secret, it has to enable an authorized health professional to access the patient's EHR for a specified time frame. Hence, medical diagnoses can be read or updated even after a medical attendance. The configured time frame must not be exceeded by the health professional.

7 Privacy and Metadata

Metadata contain information about the contents of a document and are often pseudonymized in order to protect the patient and the document creator. Metadata can be used as an attack vector by means of analysis of statistical databases, e.g. information about zip code, birthday and gender allows one to identify pseudonymized users [8]. Patient identities have been deduced from cancer type, zip code and time of diagnosis as well [10, 28].

If the anonymously stored data does not involve any personal data about the patient, it may still be possible to identify the author of a text by using stylometrics [26]. This can be done by comparing data from the EHR with medical publications or PhD thesis from various authors and thus identifying the physicians who accessed the medical documents of a specific patient. One can draw some conclusions about the patient by analysis of the data extracted this way; e.g. the specialities of the accessing health professionals. This kind of attack may be prevented by using predefined text modules.

Stingl and Slaming [30] pointed out different attacks on patient health records including the statistical analysis of metadata which can be mostly launched by internal attackers. Even if metadata is encrypted the attacker can draw different conclusions from the data. With pseudonymization and/or a reduced set of metadata this attack can be partly prevented at the expense of lesser interoperability.

7.1 Interpretation and Evaluation

Metadata are interpreted and evaluated by health professionals in order to search the bulk of medical documents within an EHR and select some for retrieval. Metadata have to contain meaningful meta information to be used as search criteria for document filtering and have to be stored in a way which makes information extraction possible, e.g. in plain-text. Plain-text achieves a maximum degree of searchability which is the main objective of adopting metadata. As long as no security critical data can be obtained from plain-text metadata this solution is

the fastest and the finest. The healthcare professionals can search within the metadata after they are authorized to access a patient's EHR and their access permissions are controlled by the access control component of the EHR system. It is necessary to use pseudonyms and avoid identity revealing information within the plain-text metadata attributes. Latter aspect is not easy to achieve because medical and technical background information are needed to provably decide which metadata attributes can be used without jeopardizing security.

A big problem we face when de-identifying the metadata is that not all metadata attributes are pseudonymizable. For example we cannot pseudonymize the time of document creation. This is the point where pseudonymization may fail as a general solution. Another approach to metadata protection is encryption, which was adopted in [16] for instance. While encrypted metadata are much easier to protect, their searchability degrades drastically. Available technical solutions for computing on encrypted data such as Fully Homomorphic Encryption [12] have not matured yet and hence are not efficiently applicable to e-health metadata. Therefore encrypting metadata causes some usability concerns. Encrypted metadata should be decrypted to be assessable. Moreover a health professional may need to decrypt the whole metadata of a patient's EHR for a single access.

7.2 Metadata in E-Health Standards

The main focus of e-health standards is interoperability; security has not been a priority to date. Hence, metadata attributes defined by standards such as IHE [17] and HL7 [15] have an informative nature and do not consider security issues primarily. One may find metadata attributes which reveal confidential information either directly or indirectly in these standards. A big challenge is how to secure risky metadata attributes without sacrificing interoperability and informativeness. We have to find a trade-off between security and privacy on one hand and interoperability and informativeness of standard-compliant metadata on the other hand. An ideal solution would be to consider security during the development phase of standards by excluding privacy sensitive attributes. Till then we need to adopt the restrictive solution of using a reduced set of metadata attributes. We classify IHE metadata attributes into three categories: privacy-critical, potentially privacy-critical and privacy-non-critical. For maximum privacy we use only attributes in the privacy-non-critical category.

8 Interoperability Issues

Interoperability is a major success factor of an EHR and can be achieved by establishing an automatically executed protocol for the exchange of information between organizations. Such a protocol can either be agreed upon between two healthcare providers by themselves or established communication standards such as HL7 [15] can be used. Development of custom solutions will however result in fragmentation of the application landscape and increased development costs as interfaces will have to be designed for each point-to-point connection [33]. On the other hand, standards provide pre-defined messages and responses for data exchange and can cut development costs. In addition, implementation frameworks, such as IHE [17] describe an infrastructure for information exchange and may utilize standards as a means for communication.

Similarly to customized communication, each organization may employ its own security measures, resulting in varying security across organizations within an EHR network. Successful attacks on organizations with a weaker notion of security might act as an entry point to the entire EHR network. Also, if organizations do not have an interoperable authentication process,

each organization may require a proof of identity compliant with its individual security policies. Such a scenario would severely hinder interoperability. In order to mitigate the risk of a diversified security environment, organizations should not only standardize their communication but their security policies and measures as well. Applying the principles of openness and auditing to healthcare interoperability solutions and EHR security might further the understanding of possible attack vectors and holes in current security measures and implementations. A well audited and scrutinized framework can at least increase security by design and help identify design and architectural flaws.

Modularization of security relevant components in healthcare information systems may not only lead to increased security but also facilitate exchange of security components if a security primitive is broken or an implementation found to be faulty. An additional side effect of standardization is the overall increase in security for all organizations as organizations will have to audit their security measures and policies if they want to join the EHR network. Additional costs for security might prove to be a barrier for smaller organizations if the financial incentive provided by the EHR is eaten up by the increased cost of security.

With regard to IHE there have been few research efforts to secure the framework. In 2007 Wozak et al. [35] developed a security infrastructure aimed at ensuring end-to-end security for the exchange of information via the IHE Cross-Enterprise Document Sharing (XDS) profile. In 2008 Wuyts et al. [36] published an improved XDS security architecture. IHE itself published the Document Encryption (DEN) [18] profile in late 2011. DEN provides means to securely exchange documents between healthcare providers using XDS. At this point in time, the profile is still only up for trial implementation and not yet part of the standard. It shows however that the need for increased security measures has been recognized and efforts to harden the security of the framework are underway.

9 Diversity in Usage

Healthcare professionals are reluctant to use EHR systems because they lack computer skills and need to spend a lot of time to learn how the complex system effectively and efficiently works [4]. Not all patients are computer-literate as well. Moreover, a human needs mental readiness to operate a device, remember a pass-phrase, etc., which may not be taken for granted if a patient is sick. An EHR is bound to fail if it would require the patient to purchase several pieces of hardware, understand their function, remember multiple passwords and PINs or present information such as cryptographic keys.

A successful EHR solution should enable sophisticated as well as simple usage methods, so that if a technique is difficult to use or gets obsolete over time, the system still remains usable. Technology overuse without provision of alternative usage methods diminishes usability of the EHR's services. Since EHRs are as such a technology-oriented trend, the diversity of usage is an often overlooked requirement in many existing EHR systems.

10 Conclusion

Ensuring the security of EHRs is of major importance and faces technical as well as usability problems. In particular, we consider scenarios where we only trust the healthcare providers and not the EHR systems that are eventually stored in the cloud. We laid out several new usability requirements we identified during our work in ongoing projects. In addition, we presented arguments to implement end-to-end security and pay more attention to the topic of metadata.

We also gave insight into the problems we faced when considering the granularity of encryption and diversity of techniques to be used. We finally argue that future EHR systems need to be interoperable and standardized, but more importantly, their security has to be ensured on a cross-organizational level realizing end-to-end confidentiality of the medical data.

Acknowledgement

This work has been partially funded by the German federal state North Rhine-Westphalia and supported by the European Regional Development Fund under the project eBPG [9].

References

[1] J. A. Akinyele, M. W. Pagano, M. D. Green, C. U. Lehmann, Z. N. Peterson, and A. D. Rubin. Securing electronic medical records using attribute-based encryption on mobile devices. In *Proceedings of the 1st ACM workshop on Security and privacy in smartphones and mobile devices*, SPSM '11, pages 75–86. ACM, 2011.

[2] B. Bergh, N. Bach, A. Brandner, and O. Heinze. EHR access rights and the role of the patient. In O. Dössel, W. C. Schlegel, and R. Magjarevic, editors, *World Congress on Medical Physics and Biomedical Engineering, September 7 - 12, 2009, Munich, Germany*, volume 25/12 of *IFMBE Proceedings*, pages 316–319. Springer Berlin Heidelberg, 2009.

[3] O. Boehm and R. Kuhlisch. eCR security architecture v1.2 services and interfaces. 2008.

[4] A. Boonstra and M. Broekhuis. Barriers to the acceptance of electronic medical records by physicians from systematic review to taxonomy and interventions. *BMC Health Services Research*, 10(1):231+, Aug. 2010.

[5] D. J. Brailer. Interoperability: the key to the future health care system. *Health Affairs (The Policy J. of the Health Sphere)*, Vol. 10:19–21, Jan. 2005.

[6] Cloud Security Alliance (CSA). Top threats to cloud computing, version 1.0., Mar. 2010. `http://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf`.

[7] S. Duennebeil, A. Sunyaev, C. Mauro, J. M. Leimeister, and H. Krcmar. Integration of patient health portals into the german healthcare telematics infrastructure. In *Proceedings of the Fifteenth Americas Conference on Information Systems*, 2009.

[8] C. Dwork. A firm foundation for private data analysis. *Commun. ACM*, 54:86–95, 2011.

[9] eBusiness Plattform Gesundheitswesen. `http://www.ebpg-nrw.de/`.

[10] K. El Emam, E. Jonker, L. Arbuckle, and B. Malin. A systematic review of re-identification attacks on health data. *PLoS ONE*, 6(12):e28071, Dec. 2011.

[11] S. Garde, P. Knaup, E. Hovenga, and S. Heard. Towards semantic interoperability for electronic health records. *Methods Inf Med*, 46(3):332–343, 2007.

[12] C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC '09, pages 169–178. ACM, 2009.

[13] S. Haas, S. Wohlgemuth, I. Echizen, N. Sonehara, and G. Müller. Aspects of privacy for electronic health records. *I. J. Medical Informatics*, 80(2):26–, 2011.

[14] K. Hayrinen, K. Saranto, and P. Nykanen. Definition, structure, content, use and impacts of electronic health records: A review of the research literature. *International Journal of Medical Informatics*, 77(5):291–304, May 2008.

[15] Health Level 7. `http://www.hl7.org/`.

[16] J. Heurix, M. Karlinger, and T. Neubauer. Pseudonymization with Metadata Encryption for Privacy-Preserving Searchable Documents. *45th Hawaii International Conference on System Sciences*, pages 3011–3020, Jan. 2012.

[17] Integrating the Healthcare Enterprise. `http://www.ihe.net/`.

[18] Integrating the Healthcare Enterprise. Document encryption (den). `http://www.ihe.net/Technical_Framework/upload/IHE_ITI_Suppl_DEN_Rev1-1_TI_2011-08-19.pdf`.

[19] A. K. Jha, D. Doolan, D. Grandt, T. Scott, and D. W. Bates. The use of health information technology in seven nations. *International Journal of Medical Informatics*, 77(12):848–854, 2008.

[20] J. S. Kahn, V. Aulakh, and A. Bosworth. What it takes: characteristics of the ideal personal health record. *Health Affairs*, 28(2):369–376, 2009.

[21] C. Klein. Cloudy confidentiality: Clinical and legal implications of cloud computing in health care. *J Am Acad Psychiatry Law*, 39(4):571–578, 2011.

[22] N. Leavitt. Internet security under attack: The undermining of digital certificates. *IEEE Computer*, 44(12):17–20, 2011.

[23] H. Löhr, A.-R. Sadeghi, and M. Winandy. Securing the e-health cloud. In *IHI 2010: Proceedings of the 1st ACM International Health Informatics Symposium*, pages 220–229. ACM, 2010.

[24] MediTrust project: Secure client systems for healthcare IT. `http://www.rubtrust-meditrust.de`.

[25] S. Narayan, M. Gagné, and R. Safavi-Naini. Privacy preserving EHR system using attribute-based infrastructure. In *Proceedings of the 2010 ACM workshop on Cloud computing security workshop*, CCSW '10, pages 47–52. ACM, 2010.

[26] A. Narayanan, H. Paskov, N. Z. Gong, J. Bethencourt, E. Stefanov, E. C. R. Shin, and D. Song. On the feasibility of internet-scale author identification. In *IEEE Symposium on Security and Privacy (IEEE S&P 2012)*, pages 300–314. IEEE Computer Society, 2012.

[27] C. M. O'Keefe, P. Greenfield, and A. Goodchild. A decentralised approach to electronic consent and health information access control. *Journal of Research and Practice in Information Technology*, 37(2), 2005.

[28] Southern Illinoisan vs. The Illinois Department of Public Health, Supreme Court of the State of Illinois, docket no. 98712, 2006.

[29] R. Steele and K. Min. Role-based access to portable personal health records. In *Proceedings of the 2009 International Conference on Management and Service Science*, pages 1–4. IEEE, 2009.

[30] C. Stingl and D. Slamanig. Privacy-enhancing methods for e-health applications: how to prevent statistical analyses and attacks. *Int. J. Business Intelligence and Data Mining*, 3:236–254, December 2008.

[31] G. van't Noordende. Security in the dutch electronic patient record system. In *Proceedings of the second annual workshop on Security and privacy in medical and home-care systems*, SPIMACS '10, pages 21–32, New York, NY, USA, 2010. ACM.

[32] C. Waegemann. The five levels of electronic health records. *MD Comput*, 13(3):199–203, 1996.

[33] J. Walker, E. Pan, D. Johnston, J. Adler-Milstein, D. W. Bates, and B. Middleton. The Value Of Health Care Information Exchange And Interoperability. *Health Affairs (The Policy J. of the Health Sphere)*, Vol. 10:10–18, Jan. 2005.

[34] M. Winandy. A note on the security in the card management system of the German e-health card. In *Electronic Healthcare, Third International Conference, eHealth 2010*, volume 69 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering (LNICST)*, pages 196–203. Springer, 2012.

[35] F. Wozak, E. Ammenwerth, R. B. R. Mair, R. Penz, T. Schabetsberger, and R. Vogl. A security infrastructure for shared electronic health records - role based access control as IHE XDS extension towards end-to-end security. In *Medinfo 2007: Proceedings of the 12th World Congress on Health (Medical) Informatics; Building Sustainable Health Systems. Amsterdam: IOS Press*, pages 1922 – 1926, 2007.

[36] K. Wuyts, R. Scandariato, G. Claeys, and W. Joosen. Hardening XDS-based architectures. In *Proceedings of the 2008 Third International Conference on Availability, Reliability and Security*, ARES '08, pages 18–25. IEEE Computer Society, 2008.