Rights management technologies: A good choice for securing electronic healthrecords?

Citation for published version (APA):

Petkovic, M., Katzenbeisser, S., & Kursawe, K. (2007). Rights management technologies: A good choice for securing electronic healthrecords? In N. Pohlmann, H. Reimer, & W. Schneider (Eds.), *Highlights of the Information Security Solutions Europe Conference on Securing Electronic Business Processes (ISSE/SECURE 2007) 25-27 September 2007, Warsaw, Poland* (pp. 178-187). Vieweg. https://doi.org/10.1007/978-3-8348-9418-2_19

DOI:

[10.1007/978-3-8348-9418-2_19](https://doi.org/10.1007/978-3-8348-9418-2_19)

Document status and date:

Published: 01/01/2007

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Rights Management Technologies: A Good Choice for Securing Electronic Health Records?

Milan Petković · Stefan Katzenbeisser · Klaus Kursawe

Information and System Security Department

Philips Research

{milan.petkovic | stefan.katzenbeisser | klaus.kursawe}@philips.com

Abstract

Advances in healthcare IT bring new concerns with respect to privacy and security. Security critical patient data no longer resides on mainframes physically isolated within an organization, where physical security measures can be taken to defend the data and the system. Modern solutions are heading towards open, interconnected environments where storage outsourcing and operations on untrusted servers happen frequently. In order to allow secure sharing of health records between different healthcare providers, Rights Management Techniques facilitating a data-centric protection model can be employed: data is cryptographically protected and allowed to be outsourced or even freely float on the network. Rather than relying on different networks to provide confidentiality, integrity and authenticity, data is protected at the end points of the communication. In this paper we compare Enterprise/Digital Rights Management with traditional security techniques and discuss how Rights Management can be applied to secure Electronic Health Records.

1 Introduction

Advances in information and communication technologies are expected to bring large benefits in the healthcare domain: the introduction of interoperable Electronic Health Record (EHR) systems can reduce the cost of the healthcare system and enhance the overall quality of treatments, whereas Remote Patient Management (RPM) services will limit the time a patient stays in hospital. Nevertheless, to date EHRs and RPMs are being used on a rather small scale. Besides problems with regard to the integration of different systems and logistic issues, concerns about information security and privacy are primary reasons for the lack of deployed systems. For example, EHR systems are facing strict security and privacy regulations (such as EU Directive 95/46 or HIPAA in the US) to which they have to comply.

Modern healthcare communication architectures tend to be open, interconnected environments: sensitive patient records no longer reside on mainframes physically isolated within a healthcare provider, where physical security measures can be taken to defend the data and the system. Patient files are rather kept in an environment where data is outsourced to or processed on partially untrusted servers in order to allow de-centralized access for family doctors, medical specialists and even non-medical care providers. The currently employed server-centric protection model, which locks the data in a database server and uses a traditional access control model to permit access to data, cannot efficiently deal with the requirements of the new healthcare infrastructures.

In order to allow sharing of records among different healthcare providers or with external parties, end-to-end security techniques facilitating data-centric protection can be employed: data is cryptographically protected and allowed to be outsourced or even freely float on the network. Rather than relying on different networks to provide confidentiality, integrity and authenticity, data is protected at the end points of the communication. However, this is not straightforward to achieve. A system may consist of a number of devices with different constraints and limitations which require specific lightweight cryptographic techniques that have to be interoperable. Furthermore, it may be necessary to perform operations while data is on transit at a specific node of the network. Finally, in emergency cases, access to electronic patient records must be instantaneously available, irrespective of the employed protection model. This opens a number of questions including key management, trust management and secure auditing.

In this work, we compare Enterprise/Digital Rights Management with traditional security techniques and discuss how Rights Management can be applied to secure Electronic Health Records. Instead of securing the flow of information by imposing security requirements at each link/component of the system, Rights Management technologies provide protection by enforcing the use of data according to granted rights or licenses. Using a set of additional technologies, such as encryption, user authentication, logging, and decentralized trust management, it enables protecting critical information and maintains control over its distribution and access. Furthermore, those techniques allow addressing conflicting privacy requirements: e.g., making data available in life-threatening emergency situations, while keeping it confidential to unauthorized people.

2 EHR Security and Privacy Requirements

The Healthcare Information and Management Systems Society¹ [HIMSS03] defines EHRs as follows:

“The Electronic Health Record (EHR) is a secure, real-time, point-of-care, patient centric information resource *for clinicians*.

- The EHR aids clinicians decision making by providing access to patient health record information *where and when they need* it and by incorporating evidence-based decision support.
- The EHR *automates and streamlines the clinician's workflow*, closing loops in communication and response that result in delays or gaps in care.
- The EHR also *supports the collection of data for uses other than direct clinical care*, such as billing, quality management, outcomes reporting, resource planning, and public health disease surveillance and reporting.”

Digital health records are used at different levels in healthcare. First, they are used in hospitals at the departmental level (e.g., at a radiology department) where medical data on examinations or treatments are stored in Electronic Medical Records (EMR). To improve sharing of information within institutions, different departmental systems are increasingly integrated into Hospital Information Systems (HIS) such that digital patient records can span over different departments. Such records, which also include administrative data, are called Electronic Health Records (EHR).

The next step after implementation of EMR and HIS systems is the integration of information systems of larger institutions that have multiple sites, or systems of co-operating care-providers in a region. This is an important first step towards the creation of a real EHR, as it increases the availability of patient

¹ HIMSS is the “healthcare industry’s membership organization exclusively focused on providing global leadership for the optimal use of healthcare information technology (IT) and management systems for the betterment of healthcare”.

records and allows sharing data among different healthcare providers and organizations. Sometimes this intermediate step is referred to as Continuity of Care Records (CCR), because it support care chains (also known as care pathways) within a region.

Finally, the term EHR is often used in relation with national Electronic Health Record infrastructures. These systems are currently under development in several countries of the European Union and the US [Char06]. A national EHR infrastructure connects all health records of a patient stored in various hospital EHR and/or departmental EMR systems deployed by different healthcare providers. The EHR is meant to provide an integrated, holistic, patient-centric and life-long view on the contents of all these systems².

The growth of EHRs from EMRs over HIS and CCRs requires the use of technologies that allow the integration of a plethora of systems into a coherent framework. For this reason, most EHR systems use web service standards like SOAP, ebXML or WSDL. In the context of health records, HL7 provides additional technology that builds on the strengths of the web-service suite of standards [HL7].

From the security point of view, the earliest departmental and hospital centric solutions were based on the “walled fortress” paradigm (firewalls, physical isolation, etc.). As these systems are opening to regional or national EHR systems, this paradigm no longer suffices. Therefore, early versions of EHR systems (that do not yet offer full functionality) utilized a selection of security mechanisms designed to complement web services, such as Transport Level Security, Security Assertion Markup Language, intra-enterprise Role-Based-Access Control, cross-enterprise user authentication, and XML digital signatures.

Around 2010-2015, many countries expect their national EHR infrastructure to be fully functional, connecting all healthcare providers in one distributed system. According to their plans, the national systems will at that time provide functionality for patient-controlled access to EHR systems, where patient policies (such as exceptions or sealed envelopes that hide certain critical information), as well as care provider policies may be integrated with default policies in order to govern access to electronic patient records.

In general, we can state the following central requirements for the security and privacy of EHRs, stemming from legal requirements:

- **Data Integrity**

The general requirement of the HIPAA Security Rule can be stated as “Covered entities must ensure the confidentiality, integrity, and availability of all protected health information (PHI) the covered entity creates, receives, maintains, or transmits.” HIPAA refers to data integrity as to the condition that PHI has not been altered or destroyed in an unauthorized manner. This includes prevention of authorized individuals making unauthorized changes to PHI as well as unauthorized people altering PHI.

- **Data Confidentiality**

A further crucial aspect is data confidentiality: healthcare data should be effectively protected against improper disclosure when it is stored, transmitted and processed. Moreover, the data should be available only to authorized parties. As with data integrity, data confidentiality is also a HIPAA requirement as part of the HIPAA *Security* Rule. Since confidentiality is closely related to privacy, it is also the focus of the HIPAA *Privacy* Rule. This rule sets standards for how PHI should be controlled. Covered entities are required to describe in general terms how health data

² Sometimes a national EHR infrastructure is also called a “virtual HER”. In this highly distributed system different pieces of health records are maintained by a community of different healthcare providers (such as hospitals, pharmacies, general practices, and laboratories).

will be protected; in addition, they have to specify the patients' rights to obtain information on and enforce control over the confidentiality of their data. This includes receiving an account of how information has been used, requesting limits on access to and additional protections for particularly sensitive information, and requesting confidential communications relating to that information. HIPAA does not require consent from patients to use or disclose health data for routine healthcare operations (such as treatment and payment) or to fulfil legal requirements (such as protection of public health). However, a patient's consent is needed for the use of the data for other purposes such as research, marketing and fundraising.

- **Data Availability**

The availability of information resources in the healthcare environment is another aspect of security which is of utmost importance. It refers to the assurance that up-to-date information is available when needed at a required level of performance and at the appropriate place. Data confidentiality, while being an essential requirement of any healthcare system, is contradictory to data availability. A clear example of this conflict is a medical emergency situation, where data confidentiality is much less important than data availability.

- **User Awareness and Control**

Recent trends in the healthcare sector towards personal, user-focused healthcare also demand more patient involvement. Patients are taking a more active role, obtaining disease information, discussing diagnoses with doctors, tracking symptoms and managing their illnesses. This also has some backing in the legislation such as HIPAA, according to which patients have, among others, the rights to request additional restrictions on the PHI, or the right to amend and inspect the stored data.

3 Traditional Solutions

Currently, most healthcare related IT-systems are "islands", i.e., all data resides within one administrative domain. This domain is not or hardly accessible from the outside, and the set of users operating on the data is reasonably small and static. In those settings, data is usually protected by a combination of access control and system security.

3.1 Access Control

An access control system matches the data, the accessing party, and the data policy to determine whether or not access to the data should be granted. To this end, the user first needs to be *authenticated*, i.e., the users' identity (or alternatively a special property such as membership in a special group or the role of the party) is established. Secondly, the system evaluates the data policy to determine if access should be granted; this can be by means of a simple table, or by a complex access control language like EPAL or XACML.

Special challenges in a medical environment are that access to the data is very context dependent, and roles of medical personnel may change quickly – an expert on a certain disease can rapidly turn from visitor to acting doctor. While under normal circumstances medical personnel not involved in the treatment of a patient should not be able to access arbitrary data (as sadly demonstrated recently, this may be a substantial security factor), a doctor also should never be blocked from data access in an emergency. Other context information that applies in a highly dynamic and heterogeneous environment such as in healthcare are the security settings and the location of the user (e.g., a doctor may not see certain information if working from home).

3.2 System Security

To provide meaningful protection, access control mechanisms need to be flanked with a number of security measures. Data needs to be protected both during transport and storage; the operating system needs to ensure that access control cannot easily be circumvented; users must be reliably authenticated; and finally, communicating devices must be able to assess the trustworthiness of platforms they communicate with.

While solutions to all those issues exist, it is hard to create a secure system in a practical environment, and many subtle problems remain. While encryption and authentication methods can be considered reasonably safe (assuming the devices get access to cryptographic keys), platform security is still a hard to solve problem.

User authentication – especially in a hospital environment – must be efficient and secure; in practice, it has often turned out that inefficient solutions are circumvented altogether, for example by the party with the highest authority logging in the morning and everybody using that account.

Platform security is an inherently hard problem as long as standard operating systems are used; given the number of known attacks on such operating systems, maintaining a secure configuration requires substantial expertise. As long as the system is not connected to the outside world and centrally managed, a reasonable secure configuration can be maintained. If, however, the system is opened up to allow communication with remote systems, new approaches such as Trusted Computing are needed to solve some of the issues. Security mechanisms like remote attestation of software allow detection of remote security properties, and also allow to give some evidence that a remote platform will enforce agreeable policies.

3.3 Web Service Security Technologies

Once large heterogeneous systems are interconnected, the interoperability between the systems starts playing a major role. One option for an interoperable communication standard is the use of emerging web service technologies. They bring together several standards to cover security aspects. Here, we briefly survey the most important standards, mainly meant to improve XML security. This is because XML is considered the universal format for structured documents on the Internet and is playing an increasingly important role in the exchange of data.

Several tools have been developed to improve the security of XML files, which basically fall into two groups. The first improves XML document security itself by using encryption and digital signatures within a document. The second provides this functionality outside the XML document.

In the first category we find standards such as XML Signatures and XML Encryption; in the second category we find, among others, XKMS and SAML. These standards are briefly described in this section.

XML Signatures

The XML Signature [XMLSi] specification provides a very flexible digital signature mechanism. A requirement of this specification is that signatures should apply either to a part or to a complete XML document. This is very relevant in the healthcare domain when a single XML document may have a long history, in which the different components are authored at different times by different parties.

XML Encryption

XML Encryption is generally performed using a combination of public key cryptography and symmetric key cryptography. Typically, a symmetric key is used to encrypt the content, and is then encrypted

using public key cryptography. Both the encrypted content and encrypted symmetric key are then sent to the recipient. The recipient may obtain the original content by first decrypting the symmetric key with his private key and then the content with the obtained symmetric key. In this way, end-to-end security for applications that require secure exchange of structured data is provided [XMLEn].

XKMS

XML Key Management Specification (XKMS) [Ford01] is a standard that provides an interface between an XML application and a Public Key Infrastructure. XKMS greatly simplifies the deployment of Public Key Infrastructures by transferring complex processing tasks from the client application to a Trust Service. It is designed to help the distribution of public keys to enable signature verification and encryption for recipients. It makes it possible to revoke or update information associated with the key pair if the circumstances change.

SAML

The Security Assertion Markup Language (SAML) standard [SAML] defines a framework for exchanging security information between online business partners. More precisely, SAML defines a common XML framework for exchanging security assertions between entities. An assertion is a package of verified security information concerning subject's (entity or human) authentication status, or access rights.

4 Rights Management Technologies

Digital Rights Management (DRM) is mainly visible in the domain of entertainment and copyright enforcement. To fulfil the needs of content providers, a number of DRM systems have been developed, such as Microsoft Windows Media DRM, IBM's Electronic Music Management System (EMMS), Sony's Open MagicGate or Thomson's SmartRight. Early DRM systems have been device-based, binding obtained content to one device on which it could be consumed. Current research is focussed on overcoming usage barriers as well as making DRM personalized and interoperable. Thanks to the pervasiveness of DRM clients (such as Microsoft Windows DRM or the Apple iPod) as well as ongoing standardization efforts, Rights Management Technologies are becoming commodity, present in many professional and consumer devices. Consequently, this opens a door for applications of DRM in other domains such as e-Business, e-Government, and in healthcare.

Rights Management Technologies or Enterprise Rights Management (ERM) are increasingly used to protect business documents in order to counter the threat of unauthorized access and distribution of corporate data. ERM has a long-term potential for automating compliance with new regulations, rules and policies, applying among others to the financial and healthcare domain (e.g. HIPAA, the Graham-Leach-Bliley Act for preserving confidentiality of consumer financial information, and the Sarbanes-Oxley Act for integrity of accounting information). Examples of ERM systems are Microsoft's Rights Management Services (RMS) platform, Authentica and SealedMedia.

In contrast to DRM, ERM applications encapsulate sensitive data objects (rather than entertainment content) and protect, control, and monitor their use and dissemination. Data confidentiality and integrity are fundamental requirements inherently supported by ERM systems. Furthermore, ERM provides means for tracing use of the data (e.g. who, when and where was the data accessed, how the data was used, etc.)³. Finally, by putting the user of the ERM system into the role of the data owner, ERM gives the data owner control of data dissemination paths.

³ Note that a DRM system is a distributed system which involves, in addition to servers at the site of the data owner, also clients deployed at the user. Therefore, it can provide a complete log on the data use.

Figure 1. Creating and viewing protected information.

Microsoft Rights Management Service is an example of an Enterprise Rights Management system. Its architecture is briefly sketched in Figure 1. The system enables protection of sensitive information from unauthorized use by allowing the data owner to define usage rights and conditions. In particular, the data owner, who has a client licensor certificate (CLS), protects the data by encrypting it within a protected data container. He also defines a publishing license containing a set of usage rights and conditions for that data container. Technically, a RMS-enabled application encrypts a data file with a symmetric key; the symmetric key is then in turn encrypted with the public key of the author's Windows RMS server and inserted into the publishing license which is bound to the data container. Once this process is finished, the author can start distributing the protected data. When a recipient receives protected data, he needs an RMS-enabled application or browser to access it. This application sends a request for a usage license to the author's RMS server. Next, the Windows RMS licensing server authenticates the recipient and checks if the recipient is authorized (RMS checks the publishing license to see if the recipient requesting a usage license has been granted rights). If the recipient is authorized, the RMS server creates a usage license (including the required decryption key) for the recipient and sends it to the recipient's RMS-enabled application. The application checks the validity of required chains of certificates and revocation lists, decrypts the content and enforces the rights that have been granted.

In the domain of healthcare, some pilots have already been set up to control distribution and usage of Electronic Health Records with existing ERM architectures. The aim is that healthcare providers can securely share confidential patient files with business associates and patients in accordance with HIPAA using the protection of the underlying ERM technology. The ERM framework enforces policies governing access to sensitive information, but also ensures protection if information is distributed beyond organizational boundaries.

5 Rights Management Technologies in the Healthcare Domain

With the increasing role of IT in the healthcare industry as well as more networked applications, medical IT systems can less and less be considered as "islands"; rather, they become a large, heterogeneous network of systems with different security requirements, guarantees, and access policies. Thus, the classical "walled fortress" approach is no longer applicable; data management solutions must take into

account that data moves between different domains. This naturally leads to a data-centric protection model as already deployed in DRM and ERM systems.

Table 1: Comparison of classical security technologies and ERM technology.

Enterprise Rights Management Systems	Classic protection technologies
<i>Usage control:</i> Data access/usage is controlled by the object owner wherever the data is physically located.	<i>Access Control:</i> Data access is controlled by the owner of the system where the data resides. If data is disseminated, the object owner loses control.
<i>Homogenous end-to-end protection</i> due to a single licensing authority.	<i>Heterogeneous end-to-end protection</i> due to a plethora of methods, such as link protection (TLS, IPSec), database security (DAC, RBAC), authentication (PKI) and storage protection.
<i>Inherently distributed architecture.</i> Distributed security architecture (object-centric protection), supporting offline operations with the same security levels.	<i>Predominantly centralized architecture.</i>
<i>Requires trusted ERM clients,</i> lower dependence on overall system integrity.	Security heavily depends on operating system security.

Table 1 compares ERM technology to data protection technologies as currently employed in small-scale EHR systems. The main advantage of ERM lies in its implicit access and usage control. In traditional access control approaches, the system owner controls access to data. When information leaves the system, it is up to the receiving party to set up new access control policies. On the other hand, ERM technologies make sure that data is accompanied with a license that defines who can access it, wherever the data is physically located. Therefore, the data owner always controls the data access, even on remote systems. In addition to access control, ERM provides means for usage control, i.e., the data can be used only as specified by the data owner (for example printing a document could be forbidden). Consequently, ERM provides homogenous end-to-end protection that is data-centric, while classical technologies need different security techniques and complex security architectures to reach the same goal. Furthermore, the data-centric protection approach increases data availability and allows off-line access and use of data which are important functional requirements in the healthcare domain: the proper elements of a patient data could be accessible anywhere and anytime only by entitled (authorized) persons, as well as the patient himself. Moreover, ERM enforces that access to off-line data is also granted based on predefined rights. With respect to audit control, their homogeneous security architecture allows for an easy supplement that can support auditing in the full system of distributed devices.

However, deployment of ERM/DRM technology in the healthcare domain is not straightforward. We can identify, among others, the following special requirements for healthcare ERM systems:

- Many parties from different domains and with different rights may be involved in accessing and modifying the documents. It is thus implausible to implement central management. Furthermore, there is a large uncertainty in who will eventually need to access a data object. While normal ERM protected documents are usually meant to be read by a mostly fixed set of parties in the first place, almost every employee in the healthcare system may potentially need to see every piece of data; nevertheless, most of those employees never need – or should get – access to the data.
- There is no clearly defined data ownership; for example, the patient must not be able to self-prescribe drugs, while the doctor in charge has no right to change the patients' home address. Consequently, even the definition of access rules/policies can come from different parties.
- Data access rights are extremely context dependent. For example, a doctor that normally should not see any data may need full access in an emergency, or the doctor may need to see normally blocked data because she suspects a particular disease. In many cases, this context cannot be

determined automatically, but only verified by a human after the incident; this requires careful auditing and some automated verification procedures.

- Access control must be extremely efficient to fit into the standard processes in medical care; given the short time doctors currently have to spend with the patients, it is unacceptable to build a system that significantly slows down the doctor.
- Even small fragments of patient records (such as the outcome of an HIV test) may be critical; as opposed to multimedia DRM, it is not even acceptable that partial information leaks.
- Roles can change quickly; doctors routinely call in external experts for advice, which then need access to the patient data. Furthermore, these experts may be at home working on their personal PCs, which hamper efficient transfer of access rights.
- In some cases, hospitals need to work with medical data for research purposes; thus, it should be possible to export anonymised data for such applications.
- Data needs to be protected while being used in a highly distributed way by different systems with complex and maybe legacy architectures, some of which may not have a trustworthy data management system.
- Medical data is rather side channel prone, i.e., a data item may reveal substantial information about its context. For example, the fact that someone takes an HIV test demonstrates that he considers himself at risk, and the fact that a teenage girl had repeated visits to her gynecologist may be revealing her pregnancy. Defining rules that cover the side information without disrupting normal healthcare is not straightforward.

6 Conclusions

Due to their construction, Enterprise Rights Management technologies are already able to satisfy many of the privacy and security requirements related to Electronic Health Records, such as data confidentiality and data integrity. ERM technology goes beyond the classical security technologies and provides a data-centric protection allowing the data owner to stay in control of data access and usage independently of the way in which data is distributed and where it is physically located. By protecting the data itself, ERM allows for highly distributed models and the off-line usage of data. In turn, this increases data availability which is one of the most important requirements in healthcare. Finally, by giving the patient the role of data owner, the ERM technology could provide means for patients to have more control on data sharing (e.g., allowing the patients to share their records with family and friends), as well as give them more awareness on how their data is used.

However, neither ERM nor DRM technologies can be, as such, straightforwardly applied in the healthcare domain. The new domain brings some additional security requirements. For example, a new requirement is to be able to cope with Role-Based Access Control, which is usually applied in the healthcare domain. Although a simple notion of roles, groups and domains exists in ERM, it is still challenging to provide for the required dynamics and context awareness. Therefore, ERM must be enhanced with techniques that address a number of open issues in order to be successfully applied in the healthcare domain.

References

- [Char06] R. Charette, Dying for Data, IEEE Spectrum, October 2006, pp. 16-21
- [Ford01] W. Ford, P. Hallam-Baker, B. Fox, B. Dillaway, B. LaMacchia, J. Epstein, J. Lapp, XML Key Management Specification (XKMS), 2001, W3C <http://www.w3.org/TR/xkms>
- [HIMSS03] Healthcare Information and Management Systems Society (HIMSS), EHR Definition, Attributes and Essential Requirements; 2003; <http://www.himss.org/content/files/EHRAttributes.pdf>
- [HL7] Health Level Seven (HL7), <http://www.hl7.org>
- [SAML] Security Assertion Markup Language, Version 2.0, OASIS Security Service TC, <http://www.oasis-open.org/specs/index.php#saml2.0>
- [XMLEn] XML Encryption, <http://www.w3.org/Encryption/2001>
- [XMLSi] XML Signatures, <http://www.w3.org/Signature>