The Active Bundle Scheme for Protecting Electronic Medical Records

Raed M. Salih
Leszek Lilien
Department of Computer Science
Western Michigan University
Kalamazoo, MI 49008
{raedmahdi.salih, leszek.lilien}@wmich.edu

Abstract: Adoption of the electronic medical records (EMRs) or electronic health records (EHRs) by healthcare providers will improve the quality of the American healthcare and reduce the annual bill. However, it will also increase privacy threats due to easier dissemination of EMRs/EHRs than "paper" medical records. Current privacy protection solutions for patient EMRs/EHRs have two main limitations: (1) they require an extensive exchange of messages between computer systems of healthcare providers; and (2) they depend only on data encryption.

In this position paper, we propose a solution that provides protection for the patients' EMRs/EHRs disseminated among different authorized healthcare provider systems. This is achieved through the use of the construct named *active bundles* (ABs). ABs keep EMRs/EHRs as *sensitive data*, include *metadata* containing privacy policies, and encompass a *virtual machine* that enforces privacy policies.

Keywords: Active bundle, Electronic Health Record (EHR), Electronic Medical Record (EMR), Patient's confidentiality, Patient's privacy, Privacy policy.

INTRODUCTION

Exchange of patient data among healthcare provider systems is a factor necessary for improving the quality of healthcare. It is being facilitated by the expansion of information technology applications. Examples of *healthcare providers* include: physicians, surgeons, dentists, pharmacists, nurses, etc., as well as the organizations employing them.

Medical records are an account of patients' long-term health history, which contains medical history, prognoses, laboratory and radiology tests, treatment descriptions, etc. Medical records have been passing through technological transformation from a physical folder form to a digital form since the end of the last century.

Electronic Medical Records (EMRs) are different from *Electronic Health Records (EHRs)* in terms of interoperability of medical information (Garets & Davis, 2006); EMRs are legal records of patients that are created in a single healthcare provider facility like a hospital or a clinic, while EHRs are the summaries of EMRs collected from more than one healthcare provider that a patient has visited. EMRs/EHRs have similar goals: (1) improving safety, quality, and efficiency of patient care; (2) reducing the cost of healthcare provider delivery; and (3) enriching the health-services research and public health monitoring (Hall, 2010). To save space, we will use the term "EMR" while in the vast majority of cases we mean "EMR/EHR."

A patient has the right to read his EMR, but he has no right to modify part(s) of information in his EMR, especially records that are related to the healthcare provider, for example, a physician's diagnosis and treatment. On the other hand, a physician has the right to read these records and no right to modify personal patient information like address and contact information.

The *ownership* of an EMR, like the ownership of any property, should represent the state or fact of exclusive legal rights and control over the property. We define a *legal guardian for an EMR* as a person or institution who has the legal authority (and the corresponding duty) to care for the patient's EMR in terms of issuing permissions for creating, reading, or modifying the EMR.

The ownership of EMRs is a central issue in healthcare because medical information has a commercial value; for example, some companies are making a profit by selling physician's prescribing routine to pharmaceutical companies (Benaloh, Chase, Horvitz & Lauter, 2009). According to American Medical Association (AMA), healthcare providers own the medical information that they are collecting (Hall, 2010) while federal and state privacy law and the Health Insurance Portability and Accountability Act (HIPAA) regulations guarantee the protection of a patient's privacy through the control of sharing the patients' healthcare information among healthcare providers. In addition to being accessed by healthcare providers, EMRs are also reviewed by many other parties, such as health insurance companies and federal or state governments. As the result, neither the ownership nor the guardianship gives the patient complete rights over the patient's EMR.

Among all these challenges along with what future government plans introduce, patients' privacy and confidentiality are considered important topics by many researchers. We define *privacy* as "the right of an entity, acting in its own behalf, to determine the degree to which it will interact with its environment, including the degree to which the entity is willing to share information about itself with others" (Shirey, 2000). *Patient privacy* refers to the patient's right to keep and control their health information. Patients determine when, how and which portions of their health information are disclosed or disseminated (Patient confidentiality, Encyclopedia of surgery, 2011). Patient confidentiality is the right of a patient to keep his healthcare or medical information private and confidential unless the patient gives permission to disclose it to another party. We believe that patient privacy or patient confidentiality gains in importance because there is a risk that patients' personal or private information (such as the social security number, or home address) might be disclosed during interactions among healthcare provider systems or other systems authorized to access it.

The goal of this position paper is arguing that an active bundle scheme can protect privacy of patients by protecting their EMRs.

BASIC CONCEPTS OF THE ACTIVE BUNDLE SCHEME

An active bundle or AB (Ben Othmane & Lilien, 2009) is a software construct, which bundles together the following three components: (1) *sensitive data*, which can contain a patient's EMR, so it is protected from privacy violations (2) *metadata*, which contain information describing sensitive data and prescribing its use; they include a privacy policy for the sensitive data (which control the access to sensitive data or their portions), as well as the rules for AB dissemination; and (3) Virtual Machine (VM), which controls and manages how its AB behaves, thus making the AB active; the essential task of the VM is enforcement of the privacy policy specified by metadata (Ben Othmane, 2010; Ben Othmane & Lilien, 2009).

In more detail, there are four protection mechanisms that the VM provides on behalf of its AB:

1) Enforcing privacy policies: Allows a host to access all or a part of the sensitive data according to the privacy policy specified in metadata.
2) Integrity checks: Verify integrity of data, metadata, and the VM of an AB. When a check fails, the VM can evaporate the bad data, or apoptosize (self-destruct) its AB (Ben Othmane, 2010).
3) Apoptosis: Destroys irretrievably the *entire* AB (including its data, metadata, and VM) in cases when: (1) a visited host's trust level is lower than the required trust threshold specified by the privacy policy; or (2) the integrity check fails.
4) Evaporation: Destroys irretrievably a *part* of the sensitive data that the visited host is not authorized to access. Typically, evaporation destroys all AB data with the "required trust level" *above* the trust level of the visited host (no metadata or VM code is affected); only less sensitive data (that a host with a lower trust level may access) remain.

Delivering sensitive data to a destination host involves two processing steps. First, during AB *creation* an active bundle encapsulates its EMR (sensitive data), metadata, and the VM. A created AB becomes ready to be sent to a destination (carrying its EMR). Second, when an AB reaches its destination (a visited host), the AB's *enabling* process starts; Enabling includes enforcing privacy policies (by the AB's VM), followed by obtaining the visited

host's trust level, then performing (by the VM) the appropriate subset (as specified by the AB's privacy policy) of the apoptosis, integrity checking, evaporation, and data disclosure data activities.

The current implementation of the AB scheme (Ben Othmane, 2010) uses Trusted Third Parties (TTPs) for maintaining (and providing to ABs) the trust levels of visited hosts.

LITERATURE REVIEW

Due to space limitation, we present related work very concisely. Bhattacharya *et al*. (2006) proposes a middleware architecture called (Privacy Broker) to enforce the legal privacy requirements. It uses: (1) a unique key to encrypt and decrypt data that it retrieves from a database and (2) capability certificates that verify and evaluate all users' requests and enforce policies to access patient's data.

Benaloh *et al*. (2009) propose an encryption system, called Patient Controlled Encryption (PCE), to protect patient privacy. Using the PCE, the patient controls and shares his EMR with other authorized entities through generating and distributing a set of subkeys.

Akinyele *et al*. (2010) suggest using a self-protecting EMR inside and outside of the hospital environment. The solution uses attribute-based encryption and a cloud system. The solution uses a set of policies and encryption/decryption keys that allow a patient: (1) to read his Personal Health Record (PHR); and (2) to read, write and manage his EMR through his mobile device that interfaces with a cloud system such as Google Health.

PROBLEM STATEMENT AND RESEARCH HYPOTHESIS

Protecting patient privacy is considered the main problem in healthcare information technology and healthcare informatics. The current solutions for protecting patient EMRs have two main limitations: (1) they require not only an extensive exchange of messages between caregivers to protect data, but also exchange of numerous control messages among caregivers' systems; and (2) they depend only on encryption (in which data decryption keys must be provided to specific caregivers).

In contrast, the AB scheme eliminates both limitations. First, it does not require so many control messages between the AB and the visited host in order to deliver an EMR from a source to a destination. Second, the AB scheme protects privacy of data not only by encryption but also by enforcement of privacy policies. Additionally, our approach does not need to distribute decryption keys (which are typically used in other privacy solutions) for all authorized healthcare providers; instead AB relies on TTP that provides the decryption keys.

Multiple *versions* of an EMR, with different owners (or guardians), can exist for a single patient. Multiple versions are generated when the EMR creator does not know that another EMR exists, or when the EMR creator is unable to receive a copy of an existing EMR. Multiple versions can also stem from a single "initial" EMR in the process of its dissemination and updating.

An EMR updated by multiple entities becomes *multi-owned*, that is, its different portions (even its individual records) are owned by different entities. This happens when an EMR is disseminated among guardians, and a receiving guardian updates its own EMR with information from the received EMR; information used for the update was created (and is owned) by the owner of the received EMR. Multi-owned EMRs can overlap (two EMRs *overlap* when both include at least one record owned by the same owner.)

There are a few issues related to multiple versions and multi-ownership of EMRs. First, the vagueness of EMR ownership rights and limitations of federal and state privacy regulations give strong reason for EMR owners to isolate and avoid sharing their EMRs (Hall, 2010; Jha *et al.*, 2009). Second, an owner lacks tools to protect patient's privacy rights for arbitrary EMR fragments (down to the single record level) owned by him. Third, given a set of

owners of a patient's EMR, a rational attacker will facilitate his attack by requesting the EMR from the owner with the most lax privacy policies (of course, EMRs obtained from different owners might differ).

AB provides solutions for the above issues as well. First, the EMR owners using ABs do not have to isolate EMRs or avoid sharing them because ABs prevent unauthorized EMR accesses by visited hosts with insufficient trust levels. Second, ABs provide tools protecting owners' (or guardians') privacy rights for arbitrary EMR fragments (down to the single record level).

Third, AB protects a patient's EMR even if the EMR is owned by different owners with privacy policies of differing strength. Even if an attacker tries to exploit the laxness of some owners' privacy policies, he still faces the AB's VM that enforces the required level of privacy via AB's privacy policies.

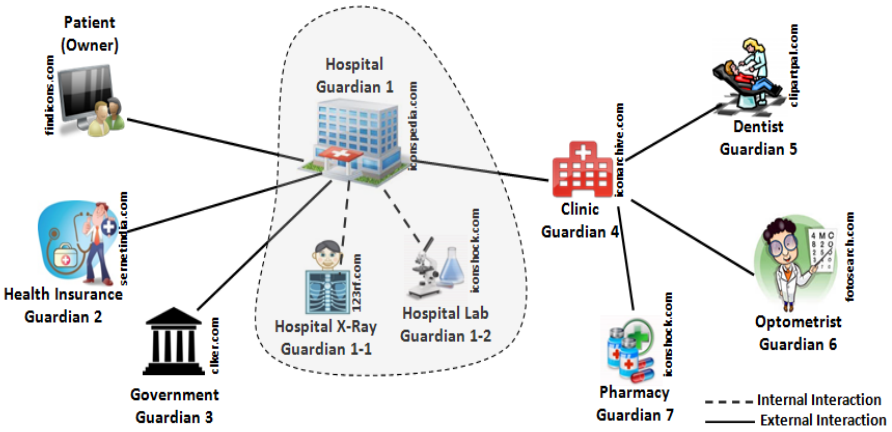

Figure 1: EMR dissemination.

Figure 1 illustrates an EMR dissemination. The hospital represents the main guardian for a patient's EMR. The hospital might send a copy of the patient's EMR to other guardians. For example, a clinic (Guardian 4) receives from the hospital and keeps a copy of a patient's EMR before, during, or after his visit. The patient's EMR is updated by multiple guardians, which increases the risk of its disclosure to unauthorized parties.

Our hypothesis is that the active bundle (AB) scheme can protect a patient's EMR during its entire lifetime; in particular, the AB scheme can prevent the EMR privacy problems discussed above.

THE PROPOSED SOLUTION: USING AB FOR PROTECTION OF PATIENTS' EMRS

Figure 3 shows a basic scenario for a patient visiting a clinic for the first time while his most up-to-date EMR is located in the general hospital that he usually visits.

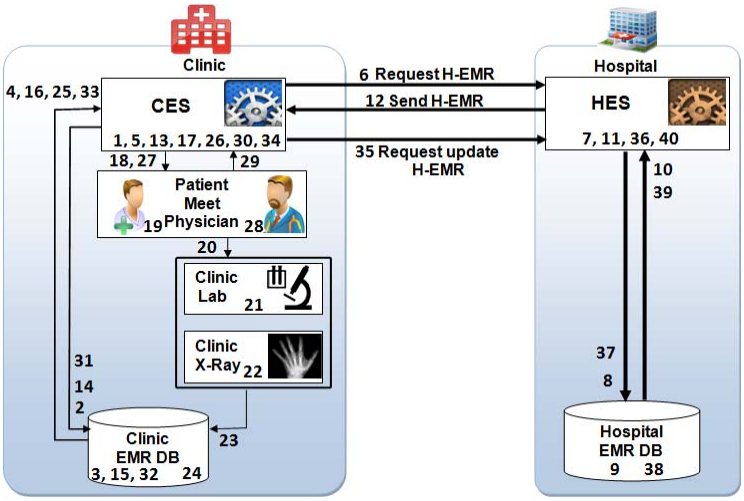

Figure 2: A patient visits a new clinic.

The scenario starts at the clinic EMR system (*CES*) with the local request (Step 1) for the patient's EMR from the CES's database. Because this is a new patient, the database fails to find C-EMR, that is, the patient's EMR for the clinic (Steps 2, 3 & 4). After realizing that there is no record for the patient (Step 5), the CES sends a request for a patient's H-EMR, that is, his EMR from his hospital EMR system (*HES*), (Steps 6). The HES receives the CES's request and sends a local request to the HES's database. After finding the patient's H-EMR (Steps 8, 9, & 10), the HES sends back the patient's H-EMR to the CES (Step 10). The CES receives the patient's H-EMR (Step 12) and uses it to create the patient's C-EMR in the CES's database (Steps 13, 14, 15 & 16). Then, the CES sends the C-EMR to the clinic's physician (Steps 17 & 18). The C-EMR is available for the physician's review (Step 19). The physician meets the patient in her office in the clinic to diagnose him. The physician requests tests from the clinic's laboratories or the X-Ray department (Step 20). New records (requests for tests) are created and sent immediately to the CES database (Steps 21 & 22), which updates the patient's C-EMR (Step 24).Then, the tests results are posted to CES database for updating (Step 23). The patient's C-EMR is updated (Step 24) and sent (Step 25) to the CES, which processes it (Step26). The physician is notified when the patient's test results are available (Step 26). She can now review the test results (Steps 27 & 28), and make her decision about the patient's case; for example, she finds that the patient must be seen by a surgeon in his hospital. The physician gives his instructions to his nurse who adds them to the patient's C-EMR (Steps 29, 30). The nurse schedules an appointment for the patient with a surgeon in the hospital and includes that in patient C-EMR. The CES updates the patient's C-EMR (including the surgeon's appointment information) in the CES database (Steps 31, 32, 33, 34). The CES sends the up-to-date patient's C-EMR to the HES (Step 35), which uses it to update the patient's H-EMR (Steps 36,37,38,39 & 40). The scenario for Figure 2 ends when the HES receives a notification that the H-EMR update is completed (Step 40).

The patients' EMR in the above scenario can be protected by using an AB encapsulating the patient's EMR. Both the CES and HES can send and receive the AB containing the patient's EMRs (its H-EMR and C-EMR versions). Figure 3 shows high-level structure of the AB scheme that provides the lifetime protection of the patient's EMR, from its creation until its destruction.

In detail, an AB is created and enabled as follows. An AB created in the HES packages a copy of patient's H-EMR as sensitive data, the associated privacy policy as metadata, and an executable code as VM (each component of the AB can be encrypted). The decryption keys for the AB are sent to a TTP for secure management. The AB is now ready for dissemination. In particular, it is ready for being sent from HES to CES when a request from the CES arrives (as shown in Step 6 in Figure 2). Once the AB is received by the CES (on the visited host), it can be enabled at the CES, as shown in Figure 3. Enabling starts with the AB's VM enforcing privacy policies to evaluate the CES credentials, and deciding which (if any) EMR portion from the received H-EMR may be disclosed. The VM uses a TTP to obtain the trust level of the CES. The TTP provides the decryption keys to the AB's VM only if CES has an adequate trust level for accessing at least data portion of the EMR.

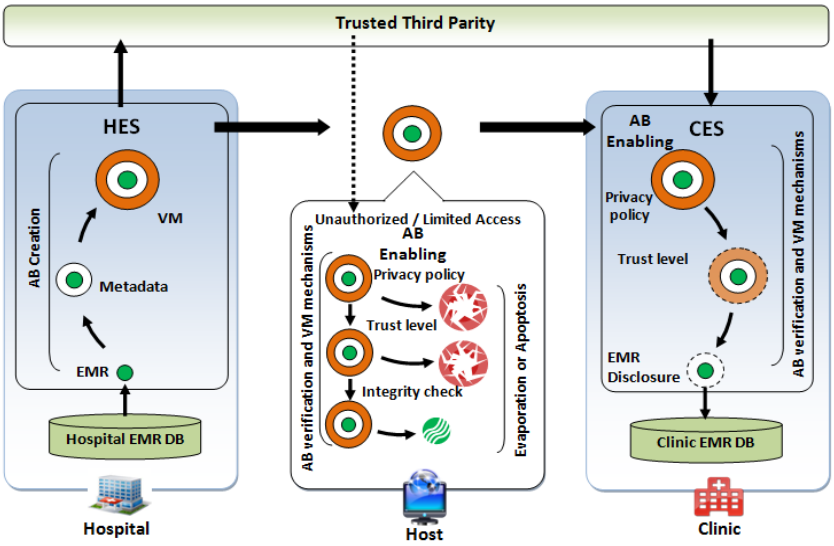

Figure 3: The AB protection lifecycle for a patient's EMR.

During AB enabling (Ben Othmane, 2010; Ben Othmane & Lilien, 2009), the CES is allowed to access only the data made accessible to it by the privacy policy-according to the following verifications steps: (1) *Evaluate visited host's trust level*: A TTP certifies the CES's trust level. The CES can access all or part of the patient EMR (sensitive data) only if its trust level is not lower than the threshold trust level specified by the privacy policy; and (2) *Run AB's integrity check*: The integrity check, specified in the metadata and VM, verifies AB's integrity. It compares the computed AB's hash value with the expected AB's hash value provided within metadata. Depending on the results of the above verification steps, VM performs one or more of the following activities on the AB under verification and on its EMR (Ben Othmane, 2010; Ben Othmane & Lilien, 2009) as shown in Figure 3: (1) apoptosis, (2) evaporation, or (3) data disclosure.

CONCLUSIONS, WORK STATUS AND FUTURE WORK

We propose a solution that provides protection for the patients' EMR during entire EMR lifetime, including its dissemination among different healthcare provider locations. We argue that this can be achieved through the use of the active bundle (AB) scheme.

Presently, we are working on validating the proposed approach via a simulation. We are also investigating whether our solution can fit into the legacy EMR-processing software. Finally, we will investigate building software that will allow each patient's control over her EMR.

There are many other issues left for longer-term research, including the following ones: (1) developing an AB scheme that does not rely on TTPs (Ben Othmane, 2010); in particular, developing an Agent-Based Active Bundle Scheme (ABABS) to protect privacy and confidentiality for both patient and healthcare provider; (2) improving the current AB scheme to include an automatic trust negotiation that prevents information leaks during credential exchanges; (3) making the intelligent agents supporting ABs able to interact with other intelligent agents (incl. other ABs) and smart healthcare environments (Salih, 2011); (4) investigating the use of the AB mechanism to protect patients' privacy in public and private cloud computing for healthcare; and (5) investigating enhancements of the self-protection level in the AB scheme.

REFERENCES

Akinyele, J. A., Lehmann, C. U., Green, M. D., Pagano, M.W., Peterson, Z. N. J., & Rubin, A. D. (2010). Self-Protecting Electronic Medical Records Using Attribute-Based Encryption. *Cryptology ePrint Archive: Report 2010/565*, 2010. Retrieved on Jun 25, 2011, from http://eprint.iacr.org/2010/565.pdf.

Ben Othmane, L. & Lilien L. (2009, August). Protecting Privacy in Sensitive Data Dissemination with Active Bundles. *Proceeding of the 7th Annual Conference on Privacy, Security and Trust, 2009 World Congress on Privacy, Trust and Management of e-Business, Saint John, New Brunswick, Canada*, pp. 202-213, doi: 10.1109/CONGRESS.2009.30.

Ben Othmane, L. (2010). *Active Bundles for Protecting Confidentiality of Sensitive Data throughout Their Lifecycle*. (Doctoral Dissertation). Department of Computer Science, Western Michigan University. Kalamazoo, Michigan, USA.

Benaloh, J., Chase M., Horvitz E., & Lauter K. (2009, November). Patient Controlled Encryption: Ensuring Privacy of Electronic Medical Records. *Proceedings of the 2009 ACM Workshop on Cloud Computing Security, New York, USA*, pp.103-114, doi:10.1145/1655008.1655024.

Bhattacharya, J., Gupta, S.K., & Agrawal, B. (2006, January). Protecting Privacy of Health Information through Privacy Broker. *Proceedings of the 39th Annual Hawaii International Conference (HICSS), Hawaii, USA, vol. 9*, pp. 89b, doi: 10.1109/HICSS.2006.402.

Garets, D. & Davis, M. (2006). *Electronic Medical Records vs. Electronic Health Records: Yes, There Is a Difference*. White Paper, A HIMSS Analytics TM.

Hall M. A. (2010). Property, Privacy, and the Pursuit of Interconnected Electronic Medical Records. *Iowa Law Review, vol. 95(2)*, pp. 631-663. Retrieved on July 13, 2011, from http://www.uiowa.edu/~ilr/issues/ILR_95-2_Hall.pdf.

Jha, A. K., DesRoches, C. M., Campbell, E. G., Donelan, K., Rao, S. R., Ferris, T. G., Shields, A., Rosenbaum, S., & Blumenthal, D. (2009). Use of Electronic Health Records in U.S. Hospitals. *New England Journal of Medicine, vol. 360(16)*, pp.1628-1638, doi: 10.1056/NEJMsa0900592. Retrieved on July 13, 2011, from http://www.nejm.org/doi/full/10.1056/NEJMsa0900592.

Patient confidentiality. Encyclopedia of surgery (2011). Retrieved on July 13, 2011, from http://www.surgeryencyclopedia.com/Pa-St/Patient-Confidentiality.html.

Salih, R.M., Ben Othmane, L., & Lilien, L. (2011, May). Privacy Protection in Pervasive Healthcare Monitoring Systems with Active Bundles. *Proceeding of the Ninth IEEE International Symposium on Parallel and Distributed Processing with Application Workshops (ISPAW 2011), Busan, Korea*, pp. 311 - 315, doi: 10.1109/ISPAW.2011.60.

Shirey, R. (2000). RFC 2828 - Internet Security Glossary. Retrieved on July 13, 2011, from http://www.faqs.org/rfcs/rfc2828.html.