Insider Threats and Access Control in e-Health Systems

Suhair Alshehri, Sumita Mishra, and Rajendra K. Raj
B. Thomas Golisano College of Computing & Information Sciences
Rochester Institute of Technology
Rochester, New York 14623, USA
Suhair.Alshehri@mail.rit.edu, Sumita.Mishra@rit.edu, and Rajendra.K.Raj@rit.edu

***Abstract*—A technological issue in improving healthcare while lowering costs is the rapid and reliable sharing of patient and other medical information across e-Health systems. Such sharing must be done securely and meet the privacy guarantees imposed by applicable laws and regulations; for instance, the Health Insurance Portability and Accountability Act (HIPAA) in the United States requires the use of appropriate access control mechanisms to protect healthcare information. Despite such legal requirements, currently implemented access control models in e-Health systems are often inadequate as demonstrated by the large and increasing numbers of successful attacks on healthcare systems. In particular, current access control models do not provide sufficient protection for healthcare systems from attacks by insiders, i.e., authorized personnel. This paper examines insider threats and presents an approach to protect healthcare information from unauthorized or improper use, disclosure, alteration, and destruction by healthcare personnel. Using a holistic approach toward modeling access control, a threat model for access control in healthcare systems is first designed and constructed, and then used to assess the effectiveness of common access control models that are used in or have been proposed for healthcare systems including Role-Based Access Control and Attribute-Based Access Control, as well as a new access control model, BiLayer Access Control, which was proposed to provide the advantages of both Role-Based and Attribute-Based models without their disadvantages.**

I. Introduction

Information sharing has become crucial in modern healthcare systems; for example, the United States Health Information Technology for Economic and Clinical Health Act (HITECH) of 2009 encourages healthcare providers to share information for improving healthcare quality and lowering costs [1]. These benefits of sharing information need to be balanced with security and privacy concerns, especially when personally identifiable healthcare information is involved. The Health Insurance Portability and Accountability Act (HIPAA) in the U.S. specifies strict requirements for the protection of such identifiable health information [2], with *access control* as a major requirement.

The increase in reported incidents of successful insider attacks shows that currently deployed access control mechanisms are inadequate in protecting against such attacks. The Privacy Rights Clearinghouse tracks data breaches [3] that typically have compromised data elements such as social security numbers, account numbers, and driving license numbers that can be exploited by rogue insiders. These reported breaches are categorized by organization types such as healthcare, educational institutions, government and military, businesses (retail, financial, and other) and non-profits [3]. The healthcare domain is particularly sensitive and it receives more than its share of attacks, especially from insiders. In a PricewaterhouseCoopers survey of more than 600 healthcare providers, insurers, pharmaceuticals, and life sciences professionals, 40% reported an improper use of protected health information by internal parties [4]. A recent report [5] stated that U.S. and German companies experience the most expensive data breach incidents while Brazil and India had the least costly data breaches. Several other reported healthcare information breaches by insiders [6]–[9] have cost healthcare organizations in penalties between $50,000 for one-time violations to $1.5 million for repeat violations across all HIPAA violation categories [10].

The ease and frequency of such inappropriate accesses compels an examination of traditional and current approaches used for access control in healthcare systems, and particularly how these approaches handle threats and attacks, especially from *insiders*. The impact of insider attacks can be significantly worse than that of outsider attacks [11]. One major reason is that insiders already have authorized credentials that allow them some level of access within an organization, thus leading to easier opportunities to cause damage. One approach of assessing how current access control mechanisms mitigate insider threats is evaluating these methods against an insider threat model.

Given the lack of a formal threat model designed specifically for access control, this paper introduces an insider threat model based on a holistic approach toward modeling access control in healthcare systems. Threat modeling is a process for understanding and analyzing the security of a system by following a systematic approach to identify potential security threats to the system [12]. The paper also provides an assessment of how the developed insider threat model performs against the two major access control models, Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC), as well as a new access control model, BiLayer Access Control model (BLAC) [13], that was proposed to utilize the best features of RBAC and ABAC.

The rest of this paper is organized as follows. Section II summarizes the background in threat model development and access control. Section III develops and constructs a threat model for access control in healthcare systems. In Section IV, the constructed threat model is used to assess the effectiveness of access control approaches for mitigating insider threats.

978-1-4799-6644-8/14/$31.00 ©2014 IEEE

Final remarks are presented in Section V.

II. Background

This section provides the background concepts in threat modeling and then examines relevant access control models.

A. Threat Modeling Methodologies

Threat modeling permits the application of a structured approach to security and helps to address top threats that may have the biggest potential to impact a system. Several threat modeling methodologies have been developed. Microsoft Threat Modeling Methodology [14], OWASP Application Threat Modeling [15] start with identifying assets that could attract attackers, understanding the target application by creating use-cases to understand how the application could be used, identifying entry points to determine how attackers could interact with the application, and analyzing data flow diagrams (DFDs) to demonstrate how data travels through the different paths in the application. Next, potential threats are identified using a threat categorization methodology such as Microsoft STRIDE model [16], or the Application Security Frame (ASF) [17].

Other methodologies such as the Process for Attack Simulation and Threat Analysis (PASTA) [18], and Trike [19] differ from Microsoft and OWASP threat modeling. The former identifies business objectives and security and compliance requirements, and the latter takes risks into perspective. In this work, the Microsoft Threat Modeling Methodology [14] was adopted for constructing the insider threat model, because it was the most suitable method for access control mechanisms.

B. Access Control

Among the different access control models, RBAC is extensively used at present in e-Health and other enterprise systems. ABAC, which was proposed to address some of RBAC's shortcomings, has its own shortcomings and newer approaches are being investigated [20]. These approaches are briefly discussed below.

RBAC. RBAC [21] has been widely deployed in healthcare systems. It regulates access to objects based on subjects' roles or their job functions, i.e., permissions to perform certain operations are assigned to roles, and subjects are assigned to those roles. For example, roles can be *Primary Care* and *Cardiology*. Permissions can be *Read a record* and *Modify a record*. RBAC's benefits include its simplicity in terms of access administration and user permission review [20].

RBAC roles typically are not sufficiently granular to restrict data access to only the *right* users. For example, consider a role that is defined as *Cardiology* and is associated with a set of permissions. Any user holding this role would be allowed to perform the operations associated with the role *Cardiology*. Thus, if a patient record is authorized to be accessed by the role *Cardiology*, all cardiologists within a healthcare organization would be able to access the record (for legitimate or illegitimate reasons), although not all those cardiologists may be involved in that patient's care. This lack of sufficient granularity in RBAC and its extensions may lead to improper access in violation of privacy, e.g., the HIPAA Privacy Rule [22].

ABAC. ABAC [23] is an access control approach proposed for providing fine-grained control. ABAC uses attributes, which are characteristics that are associated with each subject, object or environment, to define access requests and policies. For example, attributes within access requests are compared against attributes stated within the policies to determine whether to allow or deny these requests. The use of attributes permits the support for fine-grained access control. For instance, in the example discussed above, access to the patient record can be restricted to the specific cardiologists involved in that patient's care as these doctors' attributes can be included in the ABAC policy. Despite these benefits, ABAC complicates the process of making access decisions due to the large number of rules needed to be evaluated. For n attributes, ABAC may require up to 2^n possible policies [20]. Also, management of privileges, user revocation, and permission review for a particular user are difficult to perform as a large set of rules must be executed [13].

Other newer approaches. Given both RBAC's and ABAC's limitations, the development of an access control model that uses attributes and policies while maintaining the advantages of RBAC has been called for [20]. Three possible approaches were identified by Kuhn [20], but all three approaches have drawbacks [13].

Alshehri and Raj [13] proposed the BiLayer Access Control (BLAC) model that uses attributes and policies while preserving the advantages of RBAC. BLAC uses the concept of *pseudorole*, which is a set of values of static subject attributes. It is not a real role, as roles in RBAC are defined independently from subject attributes. BLAC associates subjects with pseudoroles and objects with policies that specify whether access requests by subjects are accepted or rejected. BLAC uses a two-step evaluation procedure. When an access request is made, the policy associated with the requested object is first checked to see whether the pseudorole of the requester satisfies the PseudoRole function within that policy (first step). If the requester holds the satisfied pseudorole, rules within the policy are further evaluated to check if the access request fulfills the specified values of subject, object, operation, and environment attributes to grant or deny the request (second step). This two-step access control allows BLAC essentially to utilize the approach of RBAC in the first step and ABAC, as needed, in the second step.

III. A Threat Model for Access Control

The main objective for designing and constructing a threat model in this paper is to help improve the security of healthcare systems from the perspective of access control. The focus of concern in this threat model being developed here is the protection of *patient healthcare data* from unauthorized or improper use and disclosure (*confidentiality*), and unauthorized or improper modification and destruction (*integrity*) by healthcare providers. Patient healthcare data, as used here, primarily refers to electronic protected health information (e-PHI), as described in the HIPAA Security Rule [24], that is created, received, maintained or transmitted in an electronic form by healthcare providers.

The *use* of healthcare data is defined by HIPAA as the sharing, employment, application, utilization, examination, or analysis of protected healthcare information within an organization that maintains such information [25]. HIPAA specifically defines the *disclosure* of healthcare data as "the release, transfer, provision of, access to, or divulging in any other manner of information outside the entity holding the information" [25]. The ability to carry out operations over e-PHI, including the use, disclosure, modification and destruction, is denoted as *access* [26].

Access to healthcare data is classified as authorized and unauthorized based on a set of access policies defined by healthcare organizations or healthcare laws and regulations. Authorized access in turn can be, however, classified as either proper or improper. More formal definitions of these terms are provided below, and Figure 1 illustrates the relationship among these access types.

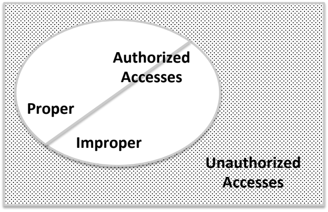

Fig. 1. The classification of access in healthcare systems

- **Authorized access:** healthcare providers have access rights to healthcare data according to the set of security policies enforced by a healthcare application.
- **Unauthorized access:** healthcare providers have no access rights to the data, but have deliberately circumvented the application to gain access.
- **Improper access:** healthcare providers have access rights to the data granted to them by the application, but have used their access to perform operations they are not truly entitled to.

As our focus is on insider attackers, the main adversaries are the authorized users, i.e., insiders, who have some level of authority to access data depending on their identity attributes, and the security policies developed by their healthcare organization.

In the following paragraphs, the steps to construct the threat model are presented.

1. Identifying the security objective. This step permits the model builder to focus on the process of constructing the threat model. The main security goal of this threat model is to minimize unauthorized and improper use, disclosure, modification, and destruction of patient healthcare data by insiders, based on a set of access policies defined by healthcare organizations and healthcare laws and regulations.

2. Creating the application overview. This step permits the model builder to understand the main functionalities and subjects of the target application. Identified here are the application architecture including the application key components, main usage scenarios, roles of subjects, and how the application components interact with each other and with external entities, i.e., healthcare providers. Based on the purpose of the threat model, the identification of these items is tied to the access control.

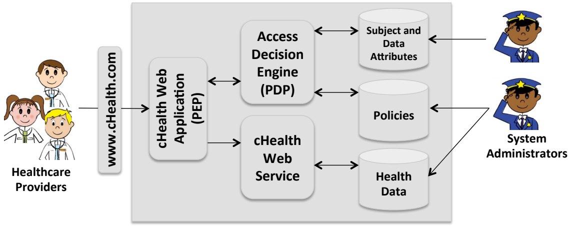

Fig. 2. An architecture of an application implementing a generic access control model

The overall architecture of a general healthcare application implementing a generic access control model is illustrated in Figure 2. To understand the target healthcare application fully, we develop a *use case* from the healthcare domain to describe the main usage scenarios and roles of subjects.

In a medical center with two hospital affiliates, hospital A and hospital B, multiple healthcare providers use a healthcare application called "cHealth" to manage patients' healthcare data in both hospitals. The roles of healthcare providers can be physicians, nurses, and administrative and billing staff, in addition to application administrators to maintain the application and access polices. Each healthcare provider is defined by a set of attributes, for example *name*, identification number (*ID*), *gender*, the field of the healthcare *provider*, their *department*, and their office *location*. These attributes are stored in a database.

Healthcare data is stored in another database in a file-based form that conforms to XML specification. Healthcare data is defined by attributes, for example, patient name, patient MRN (Medical Record Number), patient DOB (Date of Birth), and the ID of the physician responsible for treating the patient. These attributes are also stored in a database. User and data attributes are typically provided and managed by trusted entities, however, managing attributes is out of scope for this paper scope. Healthcare data is organized into a hierarchical data structure: (1) demographical, (2) clinical, and (3) billing, to provide fine-grained access control.

Typical usage scenarios are identified below to describe cHealth characteristics.

- Physicians and nurses create, read, and modify the demographical and clinical sections for patients who are under their responsibility in normal situations, with the exception of psychotherapy notes.
- Physicians and nurses create, read, and modify the demographical and clinical sections for non-patients in emergency situations, with the exception of psychotherapy notes.
- Healthcare providers do not delete data in any section.

- Administrative staff create, read, and modify data within the demographical section when they are on duty.
- Billing staff create, read, and modify data within billing section and read data within demographical section when they are on duty.
- Application administrators delete data after a predefined time of creating them.
- Healthcare providers generate access policies for the newly created data.
- Application administrators modify access policies for the created data.

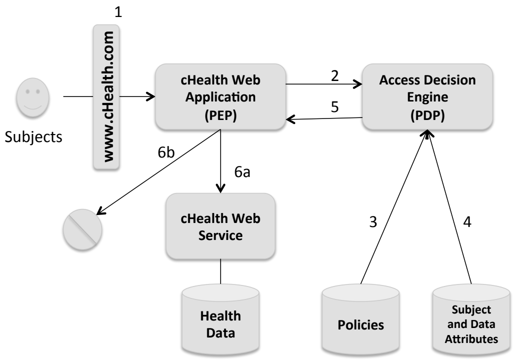
1
www.cHealth.com
Subjects
cHealth Web Application (PEP)
2
5
Access Decision Engine (PDP)
6b
6a
cHealth Web Service
3
4
Health Data
Policies
Subject and Data Attributes

Fig. 3. The interactions among the components of access models

Healthcare providers are entitled to access patient healthcare data through their web browsers based on their identity attributes, and according to their organizations' policies. The interactions among the components of the cHealth application and the healthcare providers are controlled by the access control model to grant or deny access requests. These interactions are described in Figure 3, and listed below.

1) Subject logs in and requests data through cHealth Web Application.
2) cHealth Web Application creates a web request and sends it to Access Decision Engine.
3) Access Decision Engine retrieves relevant policies.
4) Access Decision Engine retrieves attribute values related to subject, data, or environment.
5) Access Decision Engine makes access decision based on the access control model, and sends access decision to cHealth Web Application.
6) cHealth Web Application enforces authorization decision: if accept, cHealth Web Application permits subject to access and perform requested operation over requested data via cHealth Web Service (6a). If deny, cHealth Web Application rejects subject's access request (6b).

3. Decomposing the healthcare application. This step helps the model builder understand the target application in detail and how the internal components interact with one another and with external entities. The data flows and entry and exit points within the healthcare application are identified.

Figure 4 shows a high-level data flow diagram (DFD) between the application components. The purpose of the DFD is to understand how data is processed within the internal components. The rectangles denote external entities, and circles represent functions performed on data, or performed on other functions based on data. The two parallel lines and curved and directional arrows indicate databases and data movement. The curved and dashed arrows represent trust boundaries that refer to changes in access control levels as data flows through the application.

Entry and exit points refer to the interfaces that external entities use to interact with the application whether to send requests or to process data, or respond to requests or send data. In the healthcare application, the login page that subjects use to log in to the cHealth Application before requesting data access, is considered an entry point. It is denoted as the first step in the interaction process based on the access control model illustrated in Figure 3. The cHealth main page is an entry and exit point for all successfully logged-in subjects to carry out one or more of the usage scenarios identified earlier. As the goal of the desired threat model is to identify threats posed by insiders, the cHealth main page is the only point considered as it is controlled by the access control model in order for the subjects to perform operations over data.

4. Identifying the threats. This step permits the model builder to identify the threats that may compromise our security objective. Generating an attack tree is a method of representing threats against an application in a graphical or outline form [27]. An attack tree consists of a root node and child nodes, where the root node denotes a threat, and child nodes represent various methods to realize that threat. An outline for the created attack tree for the concerned security objective identified for the healthcare application is shown in Figure 5.

The construction of the threat model results in the effective identification of a set of threats and alternative approaches used to launch these threats that are relevant to our security objective. In the next section, access control models are briefly assessed against the identified threats to test their efficacy in mitigating the risk of insider threats.

IV. Assessing Access Control Approaches for Identified Threats

This section assesses how the three access control approaches, RBAC, ABAC and BLAC, perform against the identified threats by insiders (i.e., healthcare providers), which is the focus of this paper. Two fundamental types of threats exist: (1) unauthorized access of information, and (2) improper access of information. Access again refers to the set of operations—the use, disclosure, alteration, and destruction of data—that healthcare providers may perform unauthorizedly or improperly over healthcare information.

A. *Threat 1: Gaining Unauthorized Access*

Healthcare providers, or attackers, can gain unauthorized access to healthcare information via several methods. Note that the same methods can also be carried out by outsiders, i.e., unauthorized users who have no access to data but try to gain such access by illegitimate means; however, our focus in this paper is on attacks launched by insiders.

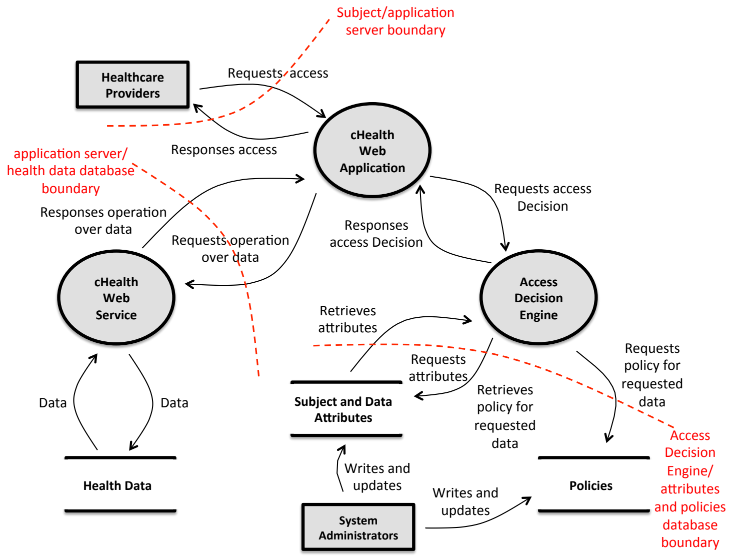
Subject/application server boundary
Healthcare Providers
Requests access
Responses access
cHealth Web Application
application server/ health data database boundary
Responses operation over data
Requests operation over data
Requests access Decision
Responses access Decision
cHealth Web Service
Access Decision Engine
Retrieves attributes
Requests attributes
Requests policy for requested data
Retrieves policy for requested data
Data
Data
Subject and Data Attributes
Health Data
Writes and updates
Writes and updates
Policies
System Administrators
Access Decision Engine/ attributes and policies database boundary

Fig. 4. The data flow diagram

Insiders may obtain credentials from legitimate healthcare providers authorized to access the target healthcare information in several ways including: (1) asking for and obtaining credentials from authorized users, (2) using authorized users' unattended logged-in machines, (3) stealing or illegally obtaining credentials from authorized users, (4) stealing devices that contain the credentials of authorized users, and (5) stealing devices or storage containing the protected heath information.

In these cases, the insiders are able to break the *authentication* scheme being used. That is, the application maps these insiders (attackers) to the identity attributes associated with the authorized healthcare providers. If the attributes associated with authorized providers, along with the attributes associated with the object, operation, and environment, satisfy the policy of the target healthcare information being attacked, the attackers (insiders) would be able to access the target information. In other words, the strength of the access control model to guard against unauthorized access depends on the robustness of the authentication scheme being used. Due to their use of attributes, ABAC and BLAC are able to prevent attacks as attackers can only spoof subjects, but not attributes. In RBAC, however, the insider attacker is likely to have access to a larger set of healthcare information due to RBAC's lack of granularity. ABAC and BLAC use attributes and policies, which must all be satisfied for access, thus reducing the healthcare information that can be threatened by the attacker.

B. *Threat 2: Gaining Improper Access*

Authorized healthcare providers may be able to perform improper operations over healthcare information using their own credentials. Such improper access is possible as most healthcare applications that implement RBAC are typically regulated using the *roles* of healthcare providers. That is, once a set of healthcare providers are assigned to a role, all

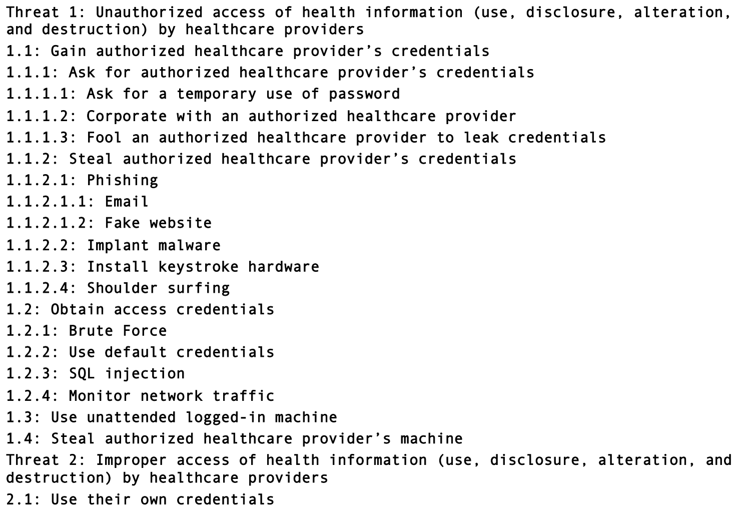

```
Threat 1: Unauthorized access of health information (use, disclosure, alteration,
and destruction) by healthcare providers
1.1: Gain authorized healthcare provider's credentials
1.1.1: Ask for authorized healthcare provider's credentials
1.1.1.1: Ask for a temporary use of password
1.1.1.2: Corporate with an authorized healthcare provider
1.1.1.3: Fool an authorized healthcare provider to leak credentials
1.1.2: Steal authorized healthcare provider's credentials
1.1.2.1: Phishing
1.1.2.1.1: Email
1.1.2.1.2: Fake website
1.1.2.2: Implant malware
1.1.2.3: Install keystroke hardware
1.1.2.4: Shoulder surfing
1.2: Obtain access credentials
1.2.1: Brute Force
1.2.2: Use default credentials
1.2.3: SQL injection
1.2.4: Monitor network traffic
1.3: Use unattended logged-in machine
1.4: Steal authorized healthcare provider's machine
Threat 2: Improper access of health information (use, disclosure, alteration, and
destruction) by healthcare providers
2.1: Use their own credentials
```

Fig. 5. The generated attack tree

providers assigned to this role will be assigned to the same permission set. Such an assignment does not take into account the providers' involvement in patient treatment, as required by the HIPAA Privacy Rule [22]. In such an RBAC setting, healthcare providers may still gain improper access. Auditing mechanisms only detect improper access after the event; the goal must be prevention, not subsequent detection.

Due to fine granularity and flexibility of both ABAC and BLAC, healthcare information accessible by providers is further restricted by attributes and a set of fine-grained access control policies. For example, when using BLAC, it is possible to restrict healthcare providers to only access health information of their own patients. Access by emergency department providers can also be limited to access requests within relevant locations. Attributes and fine-grained access policies in BLAC can thus significantly decrease improper access and reduce the amount of exposed healthcare data compared to RBAC.

Our analysis shows the fine-grained features of both ABAC and BLAC enable them to mitigate insider threats better than RBAC. Compared to ABAC, BLAC has reduced complexity of access control evaluation, user revocation and user permission review [13]. Among these three schemes, BLAC is shown to be the most effective access control approach for mitigating insider threats in e-Health systems.

V. Concluding Remarks

Laws and regulations governing healthcare data privacy, such as HIPAA in the U. S. require access control mechanisms to ensure the privacy and security of shared healthcare information. The large numbers of data breaches involving patient health information caused by insider attacks in healthcare shows that current access control models are inadequate.

Access control mechanisms can often mitigate unauthorized access by external users, i.e., outsiders, but it is more challenging to mitigate insider threats as they already have some authority to access data in the system. This paper designed and constructed a threat model for access control to address the security objective of minimizing unauthorized and improper use, disclosure, modification, and destruction of patient health information by insiders. This model was then used to evaluate the effectiveness of current access control models. The analysis indicated that ABAC and BLAC mitigate insider threats better than RBAC. However, as BLAC has lower complexity than ABAC, it can be more efficient.

Acknowledgment

S. Alshehri was supported by a Ph.D. fellowship from the Ministry of Higher Education, Saudi Arabia. This material is also based upon work partly supported by the U.S. National Science Foundation under Award DUE-1303269. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

[1] US Department of Health & Human Services - HITECH, "HITECH Act Enforcement Interim Final Rule," 2014, http://www.hhs.gov/ocr/privacy/hipaa/administrative/enforcementrule/hitechenforcementifr.html.

[2] US Department of Health & Human Services - Privacy Rule, "Standards for Privacy of Individually Identifiable Health Information; Final Rule," 2002, http://www.hhs.gov/ocr/privacy/hipaa/administrative/privacyrule/privrulepd.pdf.

[3] Privacy Rights Clearinghouse, "Chronology of Data Breaches: Security Breaches 2005—Present," August 2014, http://www.privacyrights.org/data-breach.

[4] PwC's Health Research Institute, "Old data learns new tricks: Managing patient security and privacy on a new data-sharing playground," September 2011.

[5] QNext Blog, "Data Breach Stats: Medical/Healthcare," November 2013, http://qnext.com/blog/data-breach-stats-medicalhealthcare.

[6] CBC News, "Doctor probed for improper health record access," Dec 2011, http://www.cbc.ca/news/canada/edmonton/doctor-probed-for-improper-health-record-access-1.1045719.

[7] Erin McCann, "EHR vendor to report HIPAA breach," *Government Health IT*, Mar 2013.

[8] K. Roney, "Titus Regional Medical Center Nurse Fired Over HIPAA Violation," *Beckers Hospital Review*, Jan 2012.

[9] J. Vijayan, "Three fired for accessing records of tucson shooting victims," *Computerworld*, Jan 2011.

[10] Marc Winger, "HIPAA Increases Financial Penalties For Repeat Violations To Address Increasing Healthcare Data Breaches," *Zephyr Networks*, Feb 2013.

[11] E. Bertino and G. Ghinita, "Towards mechanisms for detection and prevention of data exfiltration by insiders: keynote talk paper," in *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security (ASIACCS '11)*, Hong Kong, 2011.

[12] J. Ingalsbe, L. Kunimatsu, T. Baeten, and N. Mead, "Threat modeling: Diving into the deep end," *Software, IEEE*, vol. 25, no. 1, 2008.

[13] S. Alshehri and R. Raj, "Secure Access Control for Health Information Sharing Systems," in *IEEE International Conference on Healthcare Informatics (ICHI 2013)*, Philadelphia, 2013.

[14] Microsoft Corporation, "Threat Modeling," 2013, http://msdn.microsoft.com/en-us/library/ff648644.aspx.

[15] The Open Web Application Security Project (OWASP), "Application Threat Modeling," 2013, https://www.owasp.org/index.php/OWASP:About.

[16] Microsoft Corporation, "Threats and Countermeasures," 2013, http://msdn.microsoft.com/en-us/library/ff648641.aspx.

[17] Microsoft Corporation, "Web Application Security Frame," 2013, http://msdn.microsoft.com/en-us/library/ff649461.aspx.

[18] Marco Morana and Tony UcedaVelez, "Threat modeling of banking malware-based attacks using the P.A.S.T.A. framework," 2011, https://www.owasp.org/index.php/AppSecEU2011.

[19] Octotrike, "Trike Threat Model," 2013, http://www.octotrike.org/.

[20] D. R. Kuhn, E. J. Coyne, and T. R. Weil, "Adding attributes to role-based access control," *Computer*, vol. 43, no. 6, pp. 79–81, Jun. 2010.

[21] R. Sandhu, E. Coyne, H. Feinstein, and C. Youman, "Role-based access control models," *Computer*, vol. 29, no. 2, Feb 1996.

[22] US States Department of Health & Human Services, "The HIPAA Privacy Rule," 2002, http://www.hhs.gov/ocr/privacy/hipaa/administrative.

[23] V. C. Hu, D. Ferraiolo, R. Kuhn, A. R. Friedman, A. J. Lang, M. M. Cogdell, A. Schnitzer, K. Sandlin, R. Miller, and K. Scarfone, "Guide to Attribute Based Access Control (ABAC) Definition and Considerations," January 2014.

[24] US Department of Health & Human Services, "The HIPAA Security Rule," 2003, http://www.hhs.gov/ocr/privacy/hipaa/administrative.

[25] US Government Printing Office, "Code of Federal Regulations. Title 45 - Part 160 - General Administrative Requirements, Subpart A, Sec. 160.103," 2007, http://www.gpo.gov/fdsys/pkg/CFR-2007-title45-vol1/xml/CFR-2007-title45-vol1.xml.

[26] US Government Printing Office, "Code of Federal Regulations. Title 45 - Part 164 - Security and Privacy, Subpart C, Sec. 164.304," 2007, http://www.gpo.gov/fdsys/pkg/CFR-2007-title45-vol1/xml/CFR-2007-title45-vol1.xml.

[27] B. Schneier, "Attack Trees," *Dr. Dobb's Journal of Software Tools*, vol. 24, no. 12, Dec 1999.