5-1-2013

Insider threat mitigation and access control in healthcare systems

Suhair Alshehri

Sumita Mishra

Rajendra Raj

Insider Threat Mitigation and Access Control in Healthcare Systems

Suhair Alshehri, Sumita Mishra and Rajendra K. Raj

Department of Computer Science
B. Thomas Golisano College of Computing & Information Sciences
Rochester Institute of Technology
Rochester, New York 14623, USA
sxa3788@rit.edu, sxm1145@rit.edu, rkr@cs.rit.edu

Technical Report
May 2013

Abstract

Rapid and reliable information sharing of patient healthcare information has become critical for achieving better care with lower costs. However, such healthcare information sharing requires to be done securely with privacy guarantees, as required by law. Among its other requirements, the Health Insurance Portability and Accountability Act (HIPAA) requires the use of appropriate access control mechanisms to protect healthcare information. Despite these legal requirements, currently implemented access control models in the healthcare domain are typically 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 healthcare personnel. This paper examines how healthcare information can be protected from unauthorized or improper use, disclosure, alteration, and destruction by healthcare providers. Using a holistic approach toward modeling access control, the authors construct a threat model for access control in healthcare systems. The constructed model is then used to assess the effectiveness of current access control mechanisms such as Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC), as well as the BiLayer Access Control (BLAC) model, which was proposed as a flexible, higher-performance replacement for both RBAC and ABAC.

1 Introduction

Information sharing has become crucial in modern healthcare computing environments; for example, the 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. 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) specifies strict requirements for the protection of such identifiable health information. A major requirement specified by HIPAA is access control.

An *access control* model is typically designed to protect data from unauthorized or improper use and disclosure (*confidentiality*), and unauthorized or improper modification and destruction (*integrity*). Such protection can be achieved by ensuring that decisions for access requests by subjects (users) to protected objects (data) to perform certain operations are regulated by a set of access control policies. The access decision making process may involve the *role* played by subjects in an organization or a set of identifiable *attributes* associated with each subject.

The increases in reported incidents of successful insider attacks on healthcare systems in recent years shows that currently deployed access control mechanisms are inadequate in protecting against such attacks. For example, 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 [1]. The University of Iowa Hospitals and Clinics reported that 5 employees inappropriately accessed the electronic health records of 13 University of Iowa football players [2]. Three other breaches led to the disclosure of nearly 1.1 million records at the Utah Department of Health, Emory Healthcare, and South Carolina's Department of Health and Human Services [3]. Several other healthcare information breaches by insiders have been reported [4–8]. Such data breaches cost healthcare organizations in penalties between $50,000 for one-time violations to $1.5 million for repeat violations across all HIPAA violation categories [9].

These inappropriate accesses require a re-examination of the traditional and current approaches used for access control in healthcare systems and how these approaches handle threats and attacks, especially from insiders. The impact of insider attacks compared to outsider attacks, has been studied [10,11]. Some threat models have also been defined [12,13]. This paper uses earlier work in these areas to construct a threat model for access control in the context of a healthcare system. It also provides an initial assessment of the Role-Based Access Control (RBAC), Attribute-Based Access Control (ABAC) and a proposed access control model, BLAC [14], that was developed to combine the best of RBAC and ABAC, as discussed in Section 2.1.

The rest of this paper is organized as follows. Section 2 summarizes the background in access control and threat model development. Section 3 describes the construction of a threat model for access control in healthcare systems. In Section 4, the constructed threat model is used to assess the effectiveness of the current access control approaches, RBAC and ABAC in addition to the BLAC model, on the mitigation of insider threats. Final remarks are presented in Section 5.

2 Background

This section discusses the relevant work underlying this paper and sets the stage for presenting our work. We first present the background concepts in access control and introduce the proposed BLAC model for healthcare, followed by related work in threat modeling.

2.1 Access Control Mechanisms

Of the different access control mechanisms in healthcare systems, RBAC is commonly used in current systems and ABAC has been proposed as a replacement. As both these approaches have concerns, the BLAC model was proposed to address the deficiencies of RBAC and ABAC. These three mechanisms are discussed below.

RBAC. Access control based on RBAC [15–17] has been widely deployed in healthcare. RBAC 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," "Cardiology," and "Neurology." Operations can be "Read a record," "Add a record," and "Modify a record." RBAC's benefits include its simplicity in terms of access administration and user permission review [18]. However, due to RBAC's lack of granularity, insider threat is a common problem faced by organizations that implement RBAC. Roles, as typically defined in RBAC, are not sufficiently granular to restrict data access to only the "right" healthcare providers. For example, consider a role that is defined as "Cardiology" and is associated with a set of permissions. Any healthcare provider holding this role would be allowed to perform the permissions 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 purposes), although not all those cardiologists may be involved in that patient's care. Thus, the lack of sufficient granularity in RBAC, and its extensions [19–23], may lead to improper access in violation of the HIPAA Privacy Rule [24].

ABAC. Another access control approach proposed to provide fine-grained control is ABAC [25,26]. Attributes are sets of labels or properties that are associated with each subject (user), object or environment. ABAC uses attributes to define access requests and policies, for example, subject (user) attributes within access requests are compared against attributes stated within the ABAC access policies to determine whether to allow or deny these requests. The use of such attributes permits ABAC to support fine-grained access control. For instance, in the example discussed above, access to the patient record can be restricted to specific cardiologists involved in that patient's care as these doctor's attributes can be included in the ABAC policy. Despite these benefits, ABAC and its variants [27–31] still have several limitations. 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 rules. Also, management of privileges, user permissions, and permission review for a particular user are difficult to perform as a large set of rules must be executed [14].

RBAC Standard Revision. Given these limitations of both RBAC and ABAC, NIST [32] called for the development of an access control model that includes the use of attributes while maintaining the advantages of RBAC. Three possible mechanisms have been identified: (1) *dynamic roles*, in which attributes are used to dynamically assign roles to subjects, (2) *attribute-centric*, in which roles are defined as another attribute of subjects, and (3) *role-centric*, in which attributes are used to constrain the permissions assigned to roles. Alshehri and Raj [14] show that all these three mechanisms have drawbacks, and proposed the BiLayer Access Control (BLAC) model that uses attributes and policies while preserving RBAC's advantages.

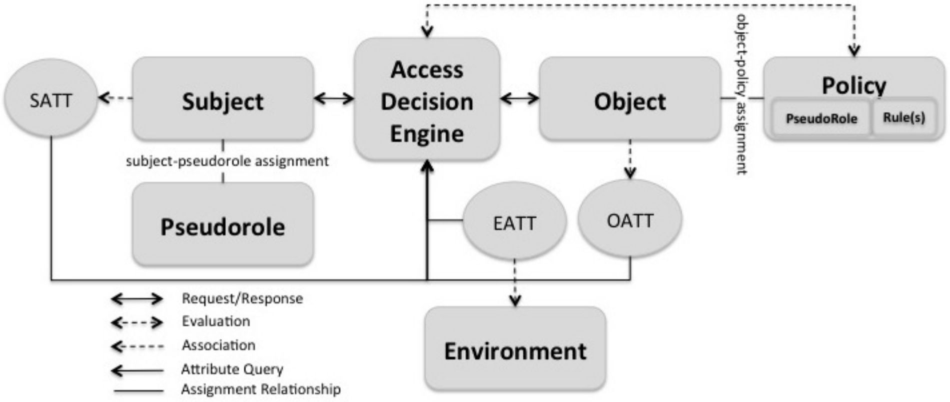

Fig. 1. The BLAC Model

The BLAC Model. The BLAC model [14] is designed to overcome the limitations by RBAC and ABAC discussed above, and to address the call by NIST for the development of an access control model that supports the use of attributes and policies while maintaining the advantages of RBAC. In the BLAC model, attributes are classified as *static* and *dynamic* based on how frequently their values change over time. Static attributes are used to generate a *pseudorole*, which is a manufactured construct consisting of a set of static attributes of subjects that resembles the concept of roles in RBAC. Static and dynamic attributes are used in policies to constrain pseudoroles.

In the BLAC model, subjects are associated with pseudoroles and objects are associated with policies to specify whether access requests by subjects are accepted or rejected. A policy, as defined in the BLAC model, consists of two elements: (1) a Boolean function called PseudoRole, and (2) a set of one or more rules. Each rule consists of four sub-elements defined as Boolean functions to specify the range of values that must be satisfied for the subject, object, action, and environment attributes. The BLAC model is illustrated in Figure 1.

The evaluation of access requests, as depicted in Figure 2, illustrates the two-step process used in the BLAC model:

1. 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 layer).
2. 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, action, and environment attributes to grant or deny the request (second layer).

This two-step access control allows the BLAC model essentially to utilize the approach of RBAC in the first step and ABAC, when needed, in the second step. The BLAC model thus integrates RBAC and ABAC to leverage the strengths of each while not suffering their disadvantages.

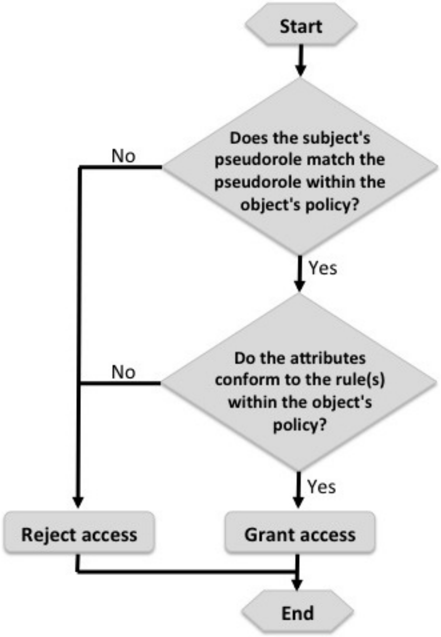
Start
Does the subject's pseudorole match the pseudorole within the object's policy?
No
Yes
Do the attributes conform to the rule(s) within the object's policy?
No
Yes
Reject access
Grant access
End

Fig. 2. The Two-step Access Control in the BLAC Model

2.2 Threat Modeling Methodologies

The other related work is in the area of threat models. Several threat modeling methodologies exist, including Microsoft Threat Modeling Methodology [33], Microsoft Threat Modeling for Web Applications [34], OWASP Application Threat Modeling [35], Process for Attack Simulation and Threat Analysis (PASTA) [36], and Trike [37].

Microsoft and OWASP methodologies start with identifying assets and understanding the target application by creating use-cases, identifying entry points, and analyzing data flow diagrams (DFDs). Next, potential threats are identified using a threat categorization methodology such as Microsoft STRIDE model, or the Application Security Frame (ASF). Finally, identified threats are ranked based on the security risks they pose. Risks can be determined using a simple High, Medium, or Low scale, or Microsoft DREAD threat-risk ranking model. PASTA and Trike differ from Microsoft and OWASP threat modeling by identifying business objectives and security and compliance requirements in the former, and taking risks into perspective in the latter. In this work, various techniques from these methodologies are adopted to meet the need of the target system and the identified security objective.

3 Constructing a Threat Model for Access Control

The main objective for constructing the threat model is to help improve the security of healthcare systems from the perspective of access control. The focus of concern in this threat model 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 [38], 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 [39]. 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" [39]. The ability to carry out operations over e-PHI, including the use, disclosure, modification and destruction, is denoted as *access* [40].

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 3 illustrates the relationship among these access types.

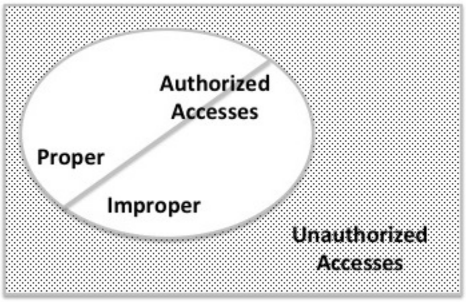

Fig. 3. Classifying Access in Healthcare Systems

- **Authorized access:** healthcare providers have access rights to healthcare data according to the set of security policies enforced by a healthcare system.
- **Unauthorized access:** healthcare providers have no access rights to the data, but have deliberately circumvented the system to gain access.
- **Improper access:** healthcare providers have access rights to the data granted to them by the system, 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 generate 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 system overview. This step permits the model builder to understand the main functionalities and subjects of the target system. Identified here are the system architecture including the system key components, main usage scenarios, roles of subjects, and how the system 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.

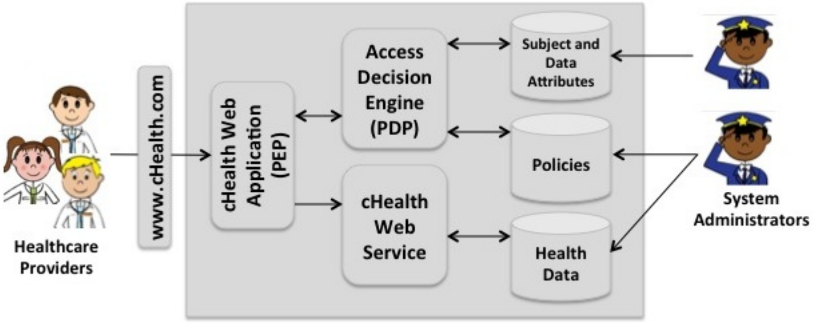

Fig. 4. An Architecture Overview

The overall architecture of a general healthcare system implementing a generic access control model is illustrated in Figure 4. To understand the target healthcare system fully, we develop a real world *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 system called "cHealth" for managing patients' healthcare data in both hospitals. The roles of healthcare providers can be physicians, nurses, and administrative and billing staff, in addition to system administrators to maintain the system 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.
- System administrators delete data after a predefined time of creating them.
- Healthcare providers generate access policies for the newly created data.
- System administrators modify access policies for the created data.

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 system and the healthcare providers are controlled by the access control model to grant or deny access requests. These interactions are described in Figure 5, 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 system. This step helps the model builder understand the target system 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 system are identified.

Figure 6 shows a high-level data flow diagram (DFD) between the system 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 system.

Entry and exit points refer to the interfaces that external entities use to interact with the system whether to send requests or to process data, or respond to requests or send data. In the healthcare system, 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 5. 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.

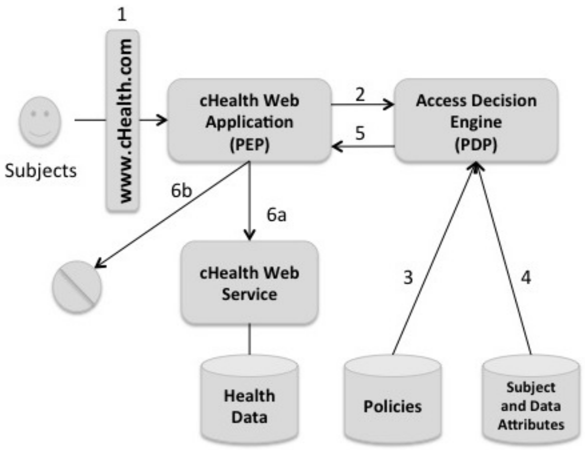
1
www.cHealth.com
Subjects
cHealth Web Application (PEP)
2
5
Access Decision Engine (PDP)
6b
6a
cHealth Web Service
Health Data
3
4
Policies
Subject and Data Attributes

Fig. 5. The Interactions Among the General Access Control Model's Components

4. Identifying the threats. This step permits the model builder to identify relevant threats that may compromise our security objective indicated earlier. Generating an attack tree is a method of representing threats against a system in a graphical or outline form [41]. 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 system is shown in Figure 7.

The construction of the threat model results in effectively identifying 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 assessed against the identified threats to test their efficacy in mitigating the risk of insider threats.

4 Preliminary Assessment

In this section, we assess how access control models, namely RBAC, ABAC, and BLAC, perform against identified threats by insiders (i.e., some healthcare providers themselves), 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 over healthcare information.

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.

Insiders may obtain credentials from legitimate healthcare providers who are authorized to access the target healthcare information in several ways including:

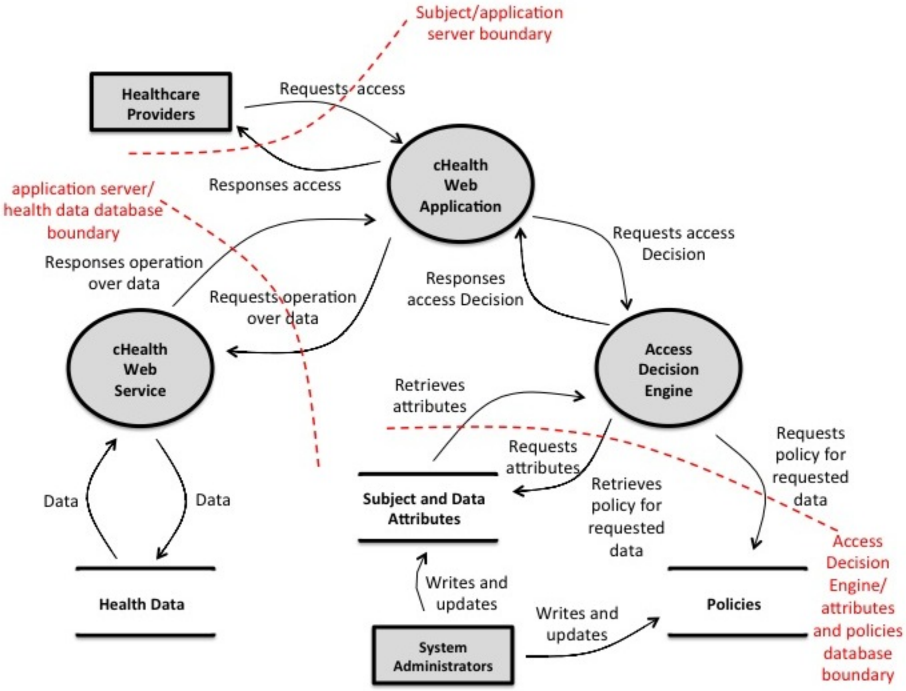
Subject/application server boundary
Healthcare Providers
Requests access
Responses access
cHealth Web Application
application server/ health data database boundary
Responses operation over data
Requests operation over data
Requests access Decision
Responses access Decision
cHealth Web Service
Access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

Fig. 6. Data Flow Diagram

- By asking for and obtaining credentials from authorized users,
- Using authorized users' unattended logged-in machines, or
- Stealing or illegally obtaining credentials from authorized users,
- Stealing devices that contain the credentials of authorized users, and
- Stealing devices or storage that contains the target protected heath information.

In these cases, the insiders are able to break the *authentication* scheme being used. That is, the system 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, action, 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 its utilization of object, environment and action attributes in its policies, the BLAC model will be able to prevent attacks as the attacker can only spoof subjects, but not these other attributes. Similar variants of ABAC will also exhibit this property. However, in RBAC, the insider attacker is likely to have access to a larger subset of healthcare information due to RBAC's lack of granularity. The BLAC model and similar variants of ABAC utilize attributes and policies, which must all be satisfied to grant access, thus reducing the subset of healthcare information that can be threatened by the attacker.

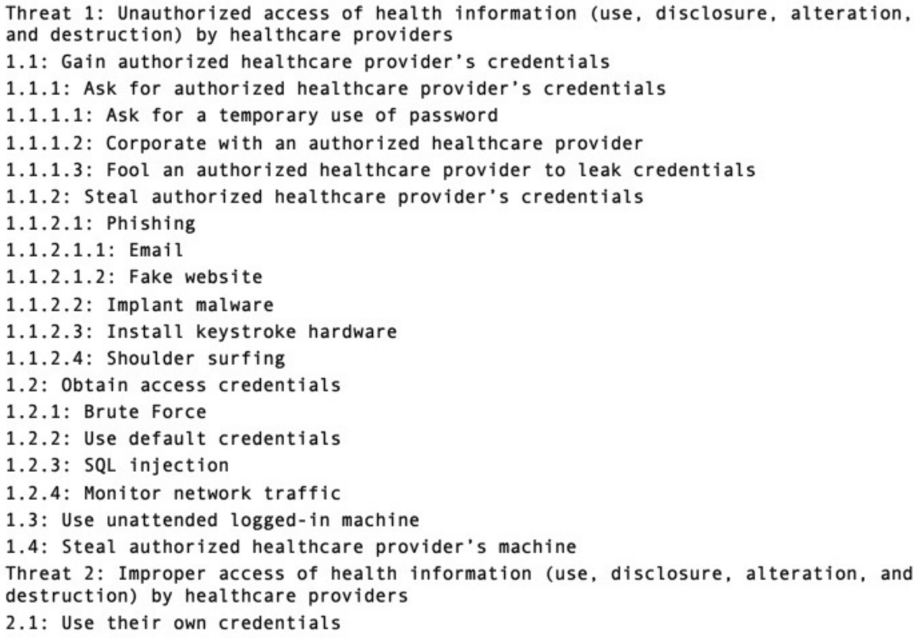

```
Threat 1: Unauthorized access of health information (use, disclosure, alteration,
and destruction) by healthcare providers
1.1: Gain authorized healthcare provider's credentials
1.1.1: Ask for authorized healthcare provider's credentials
1.1.1.1: Ask for a temporary use of password
1.1.1.2: Corporate 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. 7. The Generated Attack Tree

Gaining Improper Access. Authorized healthcare providers may be able to perform improper operations over healthcare information using their own credentials. Such improper access may be possible as most healthcare systems that implement RBAC are typically regulated using *the role of healthcare providers*. That is, once a set of healthcare providers are assigned to a role, all 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 the treatment of each patient, as required by the HIPAA Privacy Rule [24]. In such an RBAC setting, it is possible for healthcare providers to gain improper access. Even the use of auditing mechanisms that may log such improper access is not sufficient as improper access would have already occurred; the goal here must be the prevention, not the subsequent detection.

Due to the fine granularity and high flexibility of the BLAC model and variants of ABAC described in section 2 and detailed in [14], the set of healthcare information that providers can access is further constrained by various attributes and a set of fine-grained access control policies. For example, when using the BLAC model, it is feasible to specify that healthcare providers can only access health information of patients that these providers have direct treatment relationships with. Also, access by emergency department providers can be limited to access requests within the hospital locations. Thus, the BLAC model would significantly decrease the possibility of improper actions and the set of exposed healthcare information, due to the use of attributes and fine-grained access policies.

In summary, our preliminary analysis shows that both ABAC and BLAC have fine-grained features that make them more suitable for preventing both unauthorized and improper access. Due to its use of pseudoroles, as outlined in section 2 and detailed in [14], BLAC is likely to decrease unauthorized and improper access to healthcare information, with better performance and lower costs.

5 Concluding Remarks

Contributions. The main contributions of this paper are (1) the construction of a threat model based on a generic access control mechanism modeled for a healthcare system, and (2) the preliminary assessment of the effectiveness of the RBAC, ABAC and BLAC access control models against the constructed threat model.

Summary. HIPAA requires access control mechanisms to ensure the privacy and security of shared healthcare information. The increased number of data breaches involving patient health information caused by insider attacks in the healthcare domain proves that currently deployed access control models are inadequate. Motivated by the NIST call [32], the BLAC model [14] was recently proposed to address the need for fine-grained access control with performance that combines the best of RBAC and ABAC.

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. We constructed a threat model to address the security objective of minimizing unauthorized and improper use, disclosure, modification, and destruction of patient health information by insiders, based on a set of access policies defined by healthcare organizations and healthcare laws and regulations. This model was used to identify two insider threats: unauthorized access and improper access of health information by healthcare providers. We conducted a preliminary assessment of how well RBAC, ABAC and its variants, and the BLAC model performed against these threats.

Current Status. The BLAC model is currently being implemented in healthcare and finance settings. In a related study, we plan to conduct a further investigation of approaches for measuring and evaluating the performance and complexity of the BLAC model against currently used access control mechanisms.

Acknowledgments

Sushair Alshehri was supported by a Ph.D. fellowship from the Ministry of Higher Education, Saudi Arabia. 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 Ministry of Higher Education, Saudi Arabia.

References

1. PwC's Health Research Institute: Old data learns new tricks: Managing patient security and privacy on a new data-sharing playground (September 2011)

2. iHealthBeat: Hospital workers fired for improper access of football players' ehrs. California HealthCare Foundation (Feb 2011)
3. Chickowski, E.: Healthcare unable to keep up with insider threats. Dark Reading (May 2012)
4. Erin McCann: EHR vendor to report HIPAA breach. Government Health IT (Mar 2013)
5. CBC News: Doctor probed for improper health record access (Dec 2011)
6. Michelle McNickle: 10 of the largest data breaches in 2012. Healthcare IT news (Jun 2012)
7. Roney, K.: Titus regional medical center nurse fired over hipaa violation. Beckers Hospital Review (Jan 2012)
8. Vijayan, J.: Three fired for accessing records of tucson shooting victims. Computerworld (Jan 2011)
9. Marc Winger: HIPAA Increases Financial Penalties For Repeat Violations To Address Increasing Healthcare Data Breaches. Zephyr Networks (Feb 2013)
10. Alawneh, M., Abbadi, I.M.: Defining and analyzing insiders and their threats in organizations. In: Proceedings of the 2011 IEEE 10th International Conference on Trust, Security and Privacy in Computing and Communications. TRUSTCOM '11, Washington, DC, USA, IEEE Computer Society (2011) 785–794
11. Bertino, E., Ghinita, G.: 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, New York, NY, USA, ACM (2011) 10–19
12. Ingalsbe, J., Kunimatsu, L., Baeten, T., Mead, N.: Threat modeling: Diving into the deep end. Software, IEEE **25**(1) (2008) 28–34
13. Swiderski, F., Snyder, W.: Threat Modeling. Microsoft Press, Redmond, WA, USA (2004)
14. Alshehri, S., Raj, R.K.: Secure Access Control for Health Information Sharing Systems. Technical report, Rochester Institute of Technology, Department of Computer Science (May 2013) Available at: `https://ritdml.rit.edu/handle/1850/16463`.
15. Ferraiolo, D., Kuhn, R.: Role-based access control. In: In 15th NIST-NCSC National Computer Security Conference. (1992) 554–563
16. InterNational Committee for Information Technology Standards (formerly NCITS): Role based access control standard (May 2012) INCITS 359-2012.
17. Sandhu, R., Coyne, E., Feinstein, H., Youman, C.: Role-based access control models. Computer **29**(2) (Feb 1996)
18. Healthcare Information and Management Systems Society: Health Information Exchanges: Similarities and Differences HIMSS HIE Common Practices Survey Results White Paper (January 2011) `http://www.himss.org`.
19. Bertino, E., Bonatti, P.A., Ferrari, E.: Trbac: a temporal role-based access control model. In: Proceedings of the fifth ACM workshop on Role-based access control. RBAC '00, New York, NY, USA, ACM (2000) 21–30
20. Damiani, M.L., Bertino, E., Catania, B., Perlasca, P.: Geo-rbac: A spatially aware rbac. ACM Trans. Inf. Syst. Secur. **10**(1) (February 2007)
21. Joshi, J., Bertino, E., Latif, U., Ghafoor, A.: A generalized temporal role-based access control model. Knowledge and Data Engineering, IEEE Transactions on **17**(1) (jan. 2005) 4 – 23
22. Ray, I., Toahchoodee, M.: A spatio-temporal role-based access control model. In: Proceedings of the 21st annual IFIP WG 11.3 working conference on Data and applications security, Berlin, Heidelberg, Springer-Verlag (2007) 211–226
23. Sainan, L.: Task-role-based access control model and its implementation. In: Education Technology and Computer (ICETC), 2010 2nd International Conference on. Volume 3. (june 2010) V3–293 –V3–296
24. United States Department of Health & Human Services: The HIPAA Privacy Rule (2002) `http://www.hhs.gov/ocr/privacy/hipaa/administrative`.
25. Hu, V.C., Ferraiolo, D., Kuhn, R., Friedman, A.R., Lang, A.J., Cogdell, M.M., Schnitzer, A., Sandlin, K., Miller, R., Scarfone, K.: Guide to Attribute Based Access Control (ABAC) Definition and Considerations (Draft) (April 2013) NIST Special Publication 800-162, `http://csrc.nist.gov/publications/drafts/800-162/sp800_162_draft.pdf`.

26. Stepien, B., Matwin, S., Felty, A.: Advantages of a Non-Technical XACML Notation in Role-Based Models. In: 9th Annual International Conference on Privacy, Security, and Trust, IEEE (2011)
27. Alipour, H., Sabbari, M., Nazemi, E.: A policy based access control model for web services. In: Internet Technology and Secured Transactions (ICITST), 2011 International Conference for. (dec. 2011) 472 –477
28. Lang, B., Foster, I., Siebenlist, F., Ananthakrishnan, R., Freeman, T.: A flexible attribute based access control method for grid computing. Journal of Grid Computing **7** (2009) 169–180
29. Shen, H., Hong, F.: An attribute-based access control model for web services. In: Parallel and Distributed Computing, Applications and Technologies, 2006. PDCAT '06. Seventh International Conference on. (dec. 2006) 74 –79
30. Yuan, E., Tong, J.: Attributed Based Access Control (ABAC) for Web Services. In: Proceedings of the IEEE International Conference on Web Services. ICWS '05, Washington, DC, USA, IEEE Computer Society (2005) 561–569
31. Zhu, J., Smari, W.: Attribute based access control and security for collaboration environments. In: Aerospace and Electronics Conference, 2008. NAECON 2008. IEEE National. (july 2008) 31 –35
32. Kuhn, D.R., Coyne, E.J., Weil, T.R.: Adding attributes to role-based access control. Computer **43**(6) (June 2010) 79–81
33. Microsoft Corporation: Threat Modeling (2013)
34. Microsoft Corporation: Threat Modeling Web Applications (2013)
35. The Open Web Application Security Project (OWASP): Application Threat Modeling (2013)
36. Marco Morana and Tony UcedaVelez: Threat modeling of banking malware-based attacks using the P.A.S.T.A. framework (2011)
37. Octotrike: Trike Threat Model (2013)
38. United States Department of Health & Human Services: The HIPAA Security Rule (2003) `http://www.hhs.gov/ocr/privacy/hipaa/administrative`.
39. 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`.
40. 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`.
41. Schneier, B.: Attack Trees. Dr. Dobb's Journal of Software Tools **24**(12) (Dec 1999)