MEDINFO 2010
C. Safran et al. (Eds.)
IOS Press, 2010

doi:10.3233/978-1-60750-588-4-879

Healthcare Chains – Enabling Application and Data Privacy Controls for Healthcare Information Systems

Esraa Omran[a], Tyrone Grandison[b], Shereef Abu Almaati[c]

[a]*Department of Computing, Engineering and Technology, Sunderland University, Sunderland, United Kingdom.*
[b] *Health Informatics, IBM Almaden Research Center, San Jose, CA 95120, USA*
[c] *The American University of Kuwait, Safat 13034, Kuwait.*

Abstract

Healthcare applications that have access control, disclosure management and or privacy enforcement requirements may implement the respective solutions to these issues at the application level or at the database level or in both. Unfortunately, there are technical and non-technical factors that influence what can be done. In this paper we present a flexible, simple and novel approach to seamlessly imbuing current healthcare applications and their supporting infrastructure with security and privacy functionality, while being cognizant of these factors. This approach is called the Chain method. This paper will highlight the smaller design footprint, the increased ease of implementation and use of the Chain method, while demonstrating that it is as powerful and effective as traditional methods.

Keywords:

Privacy, Healthcare

Introduction

Globally, healthcare companies are increasing under pressure to provide technology, which delivers a core functionality, e.g. practice management, X-ray scanning, etc., and protects the sensitive information locked in their systems. Currently, a lot of healthcare vendors have significant investment in their existing product offerings, which tend to be developed with older technical building blocks and does not support the current social, legislative and technical requirements for privacy protection in medical systems.

The first step in enabling current healthcare vendors to augment their systems to address the current market demands, is the provision of non-intrusive technology that allows them to seamlessly handle arbitrary permissions. Unfortunately, these permissions are varied; ranging from storage access rights to "execute" rights for methods of individual classes. The focus of this paper is on the specification and enforcement of access rights to shared data resources, such as Electronic Medical Records (EMRs) and Personal Health Records (PHRs). We introduce novel technology – the Chain method [1] – which enables easy and non-intrusive definition and adherence to easily understood privacy rules.

Methods

Our methodology involved designing the information flow for the healthcare environment of our partner – the Kuwait Hospital, then implementing the Chain method behind a healthcare application at Kuwait Hospital. We then evaluated the technique against comparable technologies in the field.

What is the Chain Method?

The Chain Method [1] (hereafter referred to as *Chains* or *Chain*) is a new and novel paradigm for specifying and enforcing access controls on both applications and data; by formalizing the stakeholders, the acts that they perform and the entities that they act on. The intuition is that in modeling the actors, entities and their interactions (i.e. the private information handling model), one gets a more realistic view of the system's workflow. This model is then refined into a permissions matrix and then sent to an enforcement mechanism.

At the conceptual privacy level, Chains allow for the transformation from purpose-based systems into systems built on chains of limited acts. This is highly desirable because contemporary privacy research has lighted the fact that specifying or deriving purposes in the real world is a difficult problem [2]. A natural consequence of using Chains is that the approach doesn't need a huge number of purposes and doesn't potentially hide important user information from authorized users.

As briefly mentioned previously, the Chain method enables the simplification of the mapping/translation process from user actions to low-level implementation-level mechanisms. Chains leverage concepts from Role-Based Access Control (RBAC) [3, 4], where users have assigned roles and access purpose permissions are granted to roles associated with tasks or functionalities, not directly to individual users. In traditional RBAC, supporting dynamic changes in purposes may involve the assignment of a role to attributes with hierarchy inheritance characteristics, which allows an access purpose to be assigned

to a specific subset of users in the same role. In the Chain method, roles are assigned to acts.

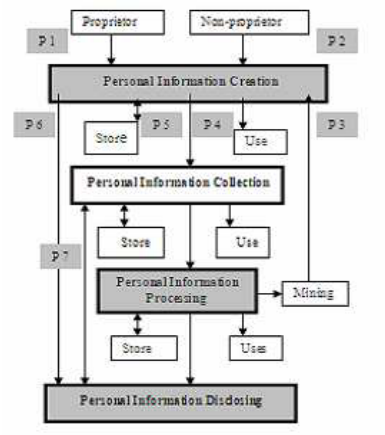

Figure 1-Personal Information Flow Model

The Chain method assumes that each purpose, i.e. conceptual task or function, can be translated to a series of actions (i.e. *chain of acts*) on personal information. The implicit assumption in the Chains is that any piece of personal information does not need more than a limited number of acts to be dealt with, such as creating, storing, processing and disclosing. This limited set can be used to design a lightweight and durable database that could safeguard personal information privacy [1]. The Personal Information Flow Model (PIFM) is a representation of the movement of information and the actions taken. The basic PIFM (Figure 1) consists of informational privacy entities and processes and is divided into a limited set of discrete actions. New personal information may be created at one or more points, e.g. P1, P2, and P3 in Figure 1, by proprietors or non-proprietors. The created information is used either at P4 (e.g. a decision for surgery needs to be made), P5 (e.g. receptionist stores patient information), or P6, where it is immediately disclosed (e.g. a physician explains to the patient his health case). Processing involves analysis and use of the personal information, e.g. mining for adverse drug reactions, longitudinal diagnostic analysis for rare conditions, etc.

Prototype

This work is the first practical instance of the Chain method. Using the concept of the Personal Information Flow Model, we built a system that involves the data owner, i.e. proprietor/patient, and healthcare entities and practitioners, i.e. doctors, nurses, clinics, hospitals (Figure 3).

Each of the labeled arrows represents specific actions that can be taken. In constructing the PFIM, we had to define a healthcare ontology (Figure 3) and integrate it with the (industry-agnostic) Chain method.

The classification that the ontology provides defined the meaning of each act. It also enabled the adjudication of the decision concerning the chain that a particular act should be added to. Finally, we specified the users that could access a particular chain or set of chains. The prototype is shown in Figure 4.

The ontology (Figure 3) was used by distributed semantic agents whose job it is to manage access and protect personal information. The actual interface to other systems and actors wishing to access personal information will use web services (Figure 5).

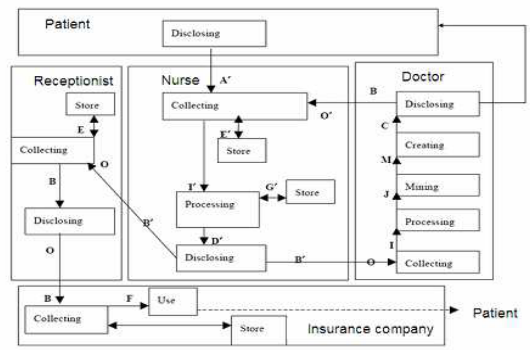

Figure 2- Healthcare PIFM for Kuwait Hospital

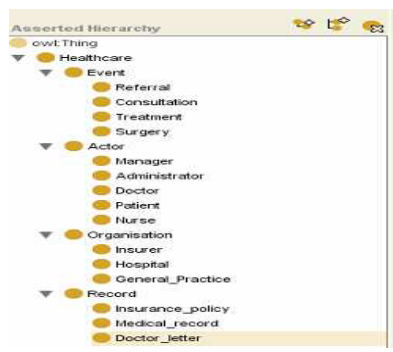

Figure 3- Healthcare Ontology constructed using Protégé OWL environment

These Personal Information Manager (PIM) agents assess information requests, potentially using other agents to verify the request and its circumstance, i.e. current state, using the ontology and analyzing the permitted acts based on the chain approach. This shared ontology is then be mapped onto a local ontology, where appropriate, to map onto the data store, where the personal information is actually stored. The benefit of this architecture is its flexibility to be applied easily to existing systems that store personal information and to manage the access using the shared ontology, while mapping it to the actual data store using a local lower level ontology as described in the ONAR approach [4].

Any request, whether by a user or by a system, in this framework will be dealt with in the following manner:

1. Authenticate the user (locally or through a trusted authentication server)
2. Verify the request to determine legitimacy using the ontology to establish whether the request is reasonable and should be acquiesced to (involves reasoning about the request and checking other systems for verification where needed)
3. Determine the location of information
4. Verify the acts and chains of acts for legitimacy
5. Extract the records
6. Execute the request

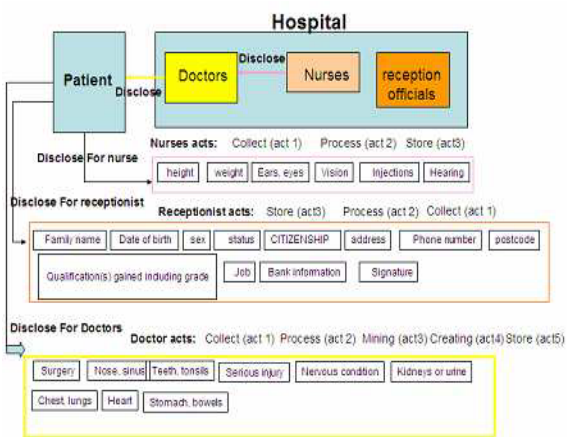

Figure 4- The Prototype

The database was designed based on our findings from a real HealthCare provider in Kuwait. After discussion with physicians, nurses and receptionists to better understand how the work in the hospital is performed, what are the problems of the existing database system they have are, and what are the requirements that should be added to their database system, we have designed the system as shown in Figure 5.

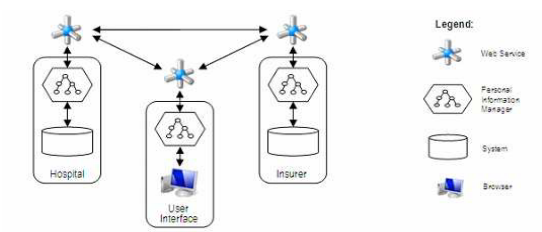

Figure 5- Designed System Architecture

The system in Figure 5 leveraged the chain ontology work and codified the requirements of the current system users.

Experiments

We developed two sets of experiments. The first is a user study, which is still ongoing. This test involves the system users at Kuwait Hospital providing feedback on their experience with Chains. The second is a comparative analysis with other similar technologies in terms of ease of and effort required for enforceability.

The first technology chosen for the comparative analysis was Role-based Access Control (RBAC) [3, 4]. In RBAC, a Role has a list of permissions. By placing a user in a Role, that user is granted access to the systems or resources associated with the Role. Users are assigned to one or more Roles, and each Role is related with a permission or set of permissions to IT resources. Using this method, a Role establishes the relationship between the users and the systems that they are authorized to access and provides a much more efficient way to decide who has access to what resources. An interesting consequence of RBAC is that organizations no longer have to work at the application or system levels, instead they can use roles to assign the appropriate permissions as a group. As a user's job tasks change they are removed from their old roles and placed in new roles based on their job title.

The second technology chosen was Task Based Access Control (TBAC) [6]. This technology is well suited for distributed computing and information processing activities with multiple points of access, control, and decision making such as found in workflow and distributed process and transaction management systems. TBAC varies from traditional access controls and security models in many respects [6]. Instead of having a system-centric view of security, TBAC approaches security modeling and enforcement at the application and enterprise level, which makes it more desirable in real world enterprises.

For our experiments on ease of and effort required for enforceability, we chose real privacy policies[1] found on the websites of the top 100 healthcare companies named in the 2009 Thomson Reuters study [7], as well as the privacy policy from Kuwait Hospital. The results were the same across the board.

For generality, we will walk through the process and show the results for a typical example, the OSF Healthcare System[2]. Given the following HIPAA policy statement[3] from their Web site:

> *"OSF may share your information with a medical care institution or medical professional for the purpose of verifying insurance coverage or benefits, informing you of a medical problem of which you may not be aware, or conducting an operations or services audit."*

We represent and implement it in RBAC, TBAC and Chains. A representational model of the above policy statement in RBAC is shown in Figure 6. We used the standard techniques for mapping from natural language to the each technology [1, 3, 4, 6].

[1] In this context, privacy policy here refers to a virtual combination of the privacy policy and the HIPAA privacy practices notice.
[2] OSF HealthCare is a multi-state corporation operating facilities in Illinois and Michigan.
[3] The full privacy policy is at http://www.osfhealthcare.org/hipaa

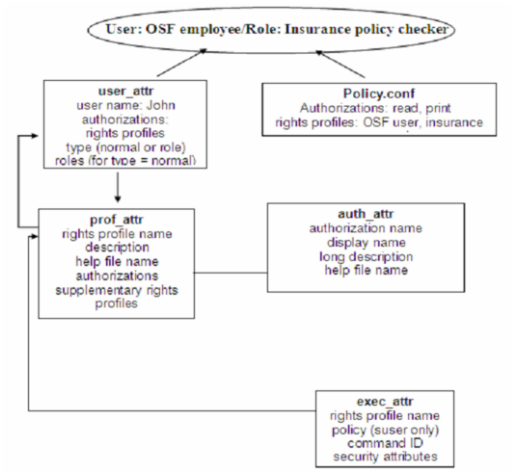

Figure 6- Policy in RBAC

Figure 7 shows the representation of the OSF policy statement in TBAC.

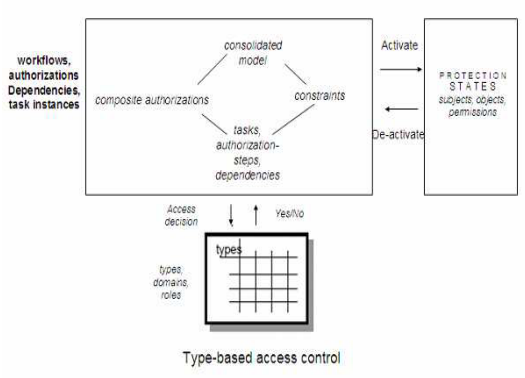

Figure 7-Policy in TBAC

Figure 8 illustrates the Chains representation of the OSF privacy statement.

As these representations will be required to be translated into a standard form, both for healthcare domain reasons and in order to do a fair comparison, each representation was transformed into OWL[4], which has a direct mapping into HL7 [8].

Representing the three models above in the OWL language, we find that the Chain model is the easiest to be translated as it contains less statements and simpler syntax.

[4] OWL (The Web Ontology Language) is a family of knowledge representation languages for authoring ontologies that is endorsed by the World Wide Web Consortium.

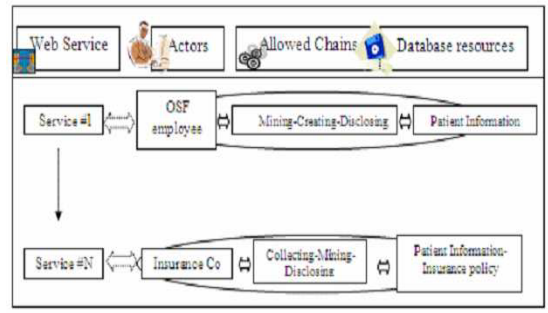

Figure 8- Policy using the Chain Method

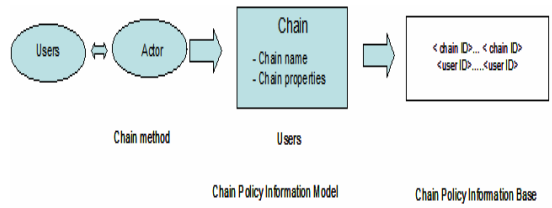

Figure 9- Chain model using OWL

Figure 9 shows that the result of the transformation process would be of the form:

<User ID>.... <User ID>
<Chain ID>.....<Chain ID>
……

Figure 10 shows the refinement process from the RBAC statements to OWL.

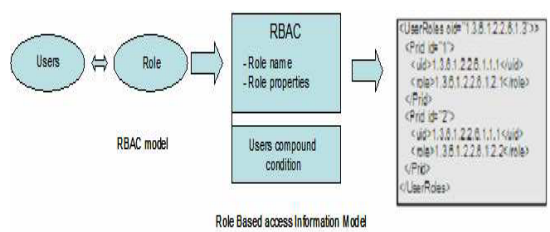

Figure 10- RBAC model using OWL

The resulting policy statement is of the form:

rbac:ssod a owl:SymmetricProperty, owl:TransitiveProperty;
rdfs:domain rbac:Role;
rdfs:range rbac:Role;
.....
<RoleName> rdfs:subclassOf rbac:Role.
<ActiveRoleName> rdfs:subclassOf
rbac:ActiveRole;rdfs:subclassOf <RoleName>.
<RoleName> rbac:activeForm <ActiveRoleName>

As TBAC is based on RBAC, the translation of the TBAC statements is similar to the translation of RBAC statements (Figure 10).

For ease of enforceability, we measure the number of tables that needs to be accessed in the determination of an access or disclosure decision. Table 1 shows that the number of required accessed tables in the Chain method is always the minimum (1), while, for this example (OSF Healthcare), TBAC and RBAC several orders of magnitude more. While the effort required may vary in TBAC and RBAC from policy to policy, the trend is that the effort is always more than Chains (and sometimes the same).

For the effort required in enforcement, this is measured by the work that has to be performed in evaluating the attributes and conditions in a statement (all other things, like low-level enforcement platform details, being equal).

Table 1- Ease of Enforceability

Access Based Method	Chain Method	TBAC	RBAC
Number of tables	1	8	5

Table 2 shows the re-thinking of the underlying representational model in Chains yields benefits in terms of the number of checks that have to be performed during policy enforcement.

Table2- Effort Required in Enforcement

Access Based Method	Chain Method	TBAC	RBAC
Number of attributes/conditions	2	4	5

Discussion

This paper is the first work in literature that has made a comparison between the three privacy preserving methods: Chain, RBAC and TBAC. This comparison was based on scientific criteria that have compared the number of tables, conditions and attributes required to design each of the methods. The chain outstands the two other methods with the small number of required tables and conditions. Also the comparison has shown the complicity of translating the RBAC and TBAC in OWL compared with the simplicity of the OWL sentences in the chain case. We also recognized that EPAL [9], Hippocratic Database (HDB) technology [10] and P3P [11] are related technologies in the field and that it is important to empirically compare them to the Chain method. But we need first to put the HDB and the Chain on same acting level (either to put them both on the application or the data level). This is one of our plans for future work.

Conclusion

In this paper, we have identified that the problem of defining, acquiring, inferring and consistently using purpose-based data disclosure technologies is difficult. We introduced the Chain Method – technology that allows easier specification of (security and privacy) policy at both the application and data level. We have prototyped the Chain method in a real healthcare provider and provided our initial results on the ease and effort involved in enforcement in a Chain-enabled system.

Acknowledgments

Our thanks to Professors Albert Bokma and David Nelson form Sunderland University for their valuable contribution on this paper.

References

[1] Al-Fedaghi S. “Beyond Purpose-Based Privacy Access Control”, The 18th Australasian Database Conf, Ballarat, Australia, Jan 29 - Feb 2, 2007.

[2] Byun J-W, Bertino E and Li N. “Purpose based access control of complex data for privacy protection”, Proceedings of the tenth ACM symposium on Access control models and technologies, June 01-03, 2005, Stockholm, Sweden.

[3] Sandhu RS. "Role-based Access Control". Advances in Computers, Vol 46, Pages 237-286. 1998. Ed: Marvin Zelkowitz. ISBN-13: 978-0120121465.

[4] Ferraiolo DF and Kuhn R, Role-Based Access Control. Proceedings of the 15th NIST-NSA National Computer Security Conf, 554-563, 1992.

[5] Tektonidis D, Bokma A, Oatley G, Salampasis M. “Onar: An Ontologies-Based Service Oriented Application Integration Framework”. In proceedings of the First International Conference on Interoperability of Enterprise Software and Applications. Feb 23-25, Lecture Notes in Computer Science (Interoperability of Enterprise Software and Applications), ISBN: 1-84628-151-2, Geneva, Switzerland, 2005.

[6] Thomas RK and Sandhu RS. " Task-based Authorization Controls(TBAC): A Family of Models for Active and Enterprise-oriented Authorization Management", Proceedings of the IFIP WG11.3 Workshop on Database Security, Lake Tahoe, California, Aug 11-13, 199

[7] Thomson Reuters. "100 Top Hospitals: 2009". March 30, 2009. http://www.modernhealthcare.com/section/lists?djoPage=product_details&djoPid=10537&djoTry=1249923457 Accessed: October 15, 2009.

[8] Health Level Seven Inc., “HL7 Standard”. http://www.hl7.org/ Accessed Oct 15, 2009.

[9] Enterprise Privacy Authorization Language (EPAL 1.2), Nov10 2003. http://www.w3.org/Submission/2003/SUBM-EPAL-20031110/ Accessed Oct 15, 2009

[10] Agrawal R, Kiernan J, Srikant R and Xu Y. “Hippocratic databases”, Proceedings of the 28th international conference on Very Large Data Bases, p.143-154, Aug 20-23, 2002, Hong Kong, China.

[11] World Wide Web Consortium. "The Platform for Privacy Preferences 1.0 Specification". April 16, 2002. http://www.w3.org/TR/P3P/ Accessed: Oct 15, 2009.