

QUT Digital Repository:
<http://eprints.qut.edu.au/>



Alhaqbani, Bandar S. and Fidge, Colin J. (2007) Access Control Requirements for Processing Electronic Health Records. In ter Hofstede, A. and Benatallah, B. and Paik, H.-Y., Eds. *Proceedings Business Process Management 2007 Workshops: First International Workshop on Process-Oriented Information Systems in Healthcare (ProHealth 2007)* 4928, pages pp. 371-382, Brisbane - Australia.

© Copyright 2007 Springer

This is the author-version of the work. Conference proceedings published, by Springer Verlag, will be available via SpringerLink.

<http://www.springer.de/comp/lncs/> Lecture Notes in Computer Science

Access Control Requirements for Processing Electronic Health Records

Bandar Alhaqbani¹ and Colin Fidge²

¹Information Security Institute,

²School of Software Engineering and Data Communications,
Queensland University of Technology, Brisbane, Australia.

b.alhaqbani@isi.qut.edu.au c.fidge@qut.edu.au

Abstract. There is currently a strong focus worldwide on the potential of large-scale Electronic Health Record systems to cut costs and improve patient outcomes through increased efficiency. A number of countries are developing nationwide EHR systems to aggregate services currently provided by isolated Electronic Medical Record databases. However, such aggregation introduces new risks for patient privacy and data security, both by linking previously-separate pieces of information about an individual, and by creating single access points to a wide range of personal data. It is thus essential that new access control policies and mechanisms are devised for federated Electronic Health Record systems, to ensure not only that sensitive patient data is accessible by authorized personnel only, but also that it is available when needed in life-critical situations. Here we review the traditional security models for access control, Discretionary Access Control, Mandatory Access Control and Role-Based Access Control, and use a case study to demonstrate that no single one of them is sufficient in a federated healthcare environment. We then show how the required level of data security can be achieved through a judicious combination of all three mechanisms.

1 Introduction

The healthcare domain—as one of the world’s largest hybrid organizations—stands to gain enormously from increased adoption of Information and Communications Technologies. Electronic Health Record (EHR) systems are the latest evolution of healthcare ICT, and countries such as Australia, the United Kingdom and the USA are working on plans for national EHR systems [9].

An Electronic Health Record is defined by Iakovidis [12] as “digitally stored healthcare information about an individual’s lifetime with the purpose of supporting continuity of care, education and research, and ensuring confidentiality at all times”. It is a mechanism for integrating healthcare information currently collected in both paper files and Electronic Medical Record (EMR) databases by a variety of healthcare providers [16].

Electronic Health Records enable efficient communication of medical information, and thus reduce costs and administrative overheads [5]. Furthermore,

EHRs will help to reduce incidents of medication error—in current healthcare systems, medical data is entered and can be interpreted in inconsistent and possibly ambiguous ways. Moreover, a patient’s health records are currently often dispersed over multiple sites with no single healthcare professional having access to all of this data. Nationwide EHR systems aim to solve these problems.

However, to achieve these potential benefits, the healthcare industry needs to overcome several significant obstacles. Currently, medical information is stored in a variety of proprietary formats using numerous off-the-shelf and custom-built medical information systems. This results in a severe inter-operability problem in the healthcare sector [4].

Also, the security of the patient’s medical data is a major issue [15] which, if not addressed in both a technically-sufficient and transparent way, will lose the patient’s confidence in and trust of the EHR system. In a worst-case scenario, patients may resorting to falsifying information, in an attempt to preserve their privacy, thus affecting the integrity of the stored data and potentially leading to life-threatening situations such as inappropriate medication. Chhanabhai et al. [2] have shown in their EHR usability survey that 73.3% of participants were highly concerned about the security and privacy of their health records. The study indicated that consumers are ready to accept the transition to EHR systems, but only provided they can be assured of the system’s security.

Several solutions are available to overcome the security concerns associated with EHR systems. Cryptographic technology, through the use of Public Key Infrastructure [3], allows confidential information to be transmitted safely via an insecure communications medium such as the Internet. On its own, however, cryptography merely handles the security of data transmission and does not address the issue of what kind of data is transmitted, or solve the problem of who has access to the data at the sending and receiving ends.

To do this we need to consider access control mechanisms that limit who can see Electronic Health Records and how they can manipulate them. Access control mechanisms have been through a lot of developments [14] (in academia and industry) in order to satisfy the needs of healthcare domains. However, developments to date have not been sufficient to meet the security requirements of a federated healthcare environment [8]. Most of the models developed so far have been designed to satisfy healthcare security requirements in a controlled environment, such as the Electronic Medical Record database maintained within a hospital. By contrast, access control mechanisms for EHRs must be safe for use in open networks, such as the Internet, and with peripheral equipment that was not designed for highly-secure operations, such as a patient’s home computer.

Discretionary Access Control (DAC), Mandatory Access Control (MAC) and Role-Based Access Control (RBAC) are well-established access control principles and have been recognized as official standards. Each was designed to overcome limitations found in its predecessor. DAC, the first standard introduced, controls each user’s access to information on the basis of the user’s identity and authorization [17]. MAC, the second standard introduced, governs access on the basis of the security classification of subjects (users) and objects in the system.

RBAC, the third standard introduced, regulates user access to information on the basis of the activities particular types of users may execute in the system.

In this paper, we demonstrate through case studies that none of these three mechanisms in isolation is sufficient for the privacy and security requirements of Electronic Health Record systems. We then explain how a careful combination of all three access control standards can be used to deliver the essential security requirements of a federated EHR system.

2 Related and previous work

An access control mechanism is intended to limit the actions or operations that a legitimate user of a computer system can perform [17]. This research area has witnessed a lot of developments in the last two decades that have resulted in the widespread adoption of three different access control models. In this section we introduce these three models and point out which of their known limitations would apply to the healthcare domain.

2.1 Discretionary access control

Discretionary Access Control is a means of restricting access to objects based on the identity of subjects and/or groups to which they belong [7]. The controls are discretionary in the sense that a user or subject given discretionary access to a resource is capable of passing that capability along to another subject. The identity of the users and objects is the key to discretionary access control. DAC policies tend to be very flexible and are widely used. However, DAC policies are known to be inherently weak for two reasons [10, 7]: granting read access is transitive and DAC policies are vulnerable to “Trojan horse” attacks.

DAC policies are commonly implemented through Access Control Lists (ACLs) and owner/other access control mechanisms, but these mechanisms are difficult to manage because addition and deletion of users or data objects requires discovery and treatment of all dependent entries in the DAC matrix.

In an Electronic Health Record system the access control requirements are more complex than allowed for by Discretionary Access Control because the data in an EHR is nominally ‘owned’ by the patient [9], but is also updated by healthcare professionals, and is stored on infrastructure belonging to healthcare providers and regulators. Indeed, a DAC model could create new security problems due to the patient’s mismanagement of their own records [7].

2.2 Mandatory access control

A Mandatory Access Control policy, which is known to prevent the “Trojan horse” problem [7, 17], means that access control decisions are made by a central authority, not by the individual owner of an object, and the owner cannot change access rights. The need for a MAC mechanism arises when the security policy

of a system dictates that [7] protection decisions must not be decided by the object's owner, and the system must enforce data protection decisions.

Mandatory Access Control typically occurs in military-style security. Usually a security labeling mechanism and a set of interfaces are used to determine access based on the MAC policy. For example, a user who is running a process at the *Secret* classification level should not be allowed to read a file with a label of *Top Secret*. This is known as the “simple security rule”, or “no read up”. By contrast, a user who is running a process with a label of *Secret* should not be allowed to write to a document with a label of *Confidential*. This rule is called the “ \star -property” or “no write down”. Multilevel security models such as the Bell-La Padula Confidentiality [6] and Biba Integrity [1] models are used to formally specify this kind of MAC policy. Nevertheless, unintended information transfer can occur in systems using MAC through covert channels, whereby information of a higher security class is deduced indirectly by intelligently combining information visible to a lower security class [10].

However, using Mandatory Access Control mechanisms in an EHR environment is likely to be very difficult due to the huge number of users who participate in those systems, the wide range of data types, and the desire to give patients ownership and (partial) control over their own medical records. Nevertheless, implementing some form of MAC policy is inevitable in an EHR system, since medical authorities must be ultimately responsible for assigning access rights [5].

2.3 Role-based access control

Role-Based Access Control decisions are based on the roles that individual users have as part of an organization. Users take on assigned roles (e.g. doctor, nurse or receptionist in our case). Access rights (or permissions) are then grouped by role name, and the use of resources is restricted to authorized individuals [10]. Under RBAC, users are granted membership into roles based on their competencies, credentials and responsibilities in the organization. User membership in roles can be revoked easily and new operations established as job assignments dictate. This simplifies the administration and management of permissions since roles can be updated without updating the permissions for every individual user. Moreover, use of role hierarchies provides additional advantages since one role may implicitly include the operations associated with another role. Also, RBAC can satisfy the “least privilege access” requirement [17], which involves granting the minimum set of privileges required for individuals to perform their job functions. Separation of Duty (SoD) is incorporated into the RBAC model [18] to ensure that a user would not be allowed to execute two roles simultaneously as per the organization's policy.

Role-Based Access Control has gained a lot of attention in healthcare security research due to its ability to provide practical fine-grained access policy administration for a large number of users and resources, for being a neutral policy, and for supporting the ‘need-to-know’ security principle.

However, some access request evaluations are complex, due to the need to consider other contextual parameters in the evaluation phase. To overcome this

problem, the Contextual RBAC model adds contextual parameters (e.g. time and location) to the RBAC model [19]. Nevertheless, even Contextual RBAC is insufficient to support the dynamic permission assignments that are needed in the healthcare domain, so Motta et al. [13] extended the model so that permission assignment is based on certain evaluation mechanisms using contextual attributes that are available at access time, and Wilikens et al. [19] used a trust level as a measurement to assign permissions.

Unfortunately, this extended process would add yet more complexity to an EHR system, which requires establishing a connection between the EHR system and the Hospital Information Systems (HISs) that are responsible for handling administrative work within a hospital, in order to collect those contextual attributes which are not immediately visible, for example the requestor's current medical role (e.g. as an emergency department doctor).

3 Healthcare access control requirements

In the previous section we reviewed the capabilities of Discretionary Access Control, Mandatory Access Control and Role-Based Access Control. In order to better understand what kind of access control solution is needed for an Electronic Health Record system, in this section we summarize the specific access control requirements peculiar to EHR systems, illustrated by a small case study, and review the weaknesses of the existing mechanisms in this situation.

A control mechanism for Electronic Health Record access must satisfy all EHR participants' needs, i.e. patients, medical practitioners and medical authorities. Each participant needs to access certain fields of the health record in order to carry out his job. Also, the various participants need the ability to set specific access controls over the record. The following privacy and security requirements have been identified as crucial to healthcare environments:

1. Each healthcare unit should have the freedom to design its own security policy and to enforce it within its domain [15].
2. Healthcare providers (e.g. General Practitioners) should have the flexibility to arbitrarily define the security of a particular document if so required.
3. Patients should have the right to have control over their own health records, including whether or not to grant access to certain medical practitioners [15].
4. Patients should be able to hide specific items of information contained in their health records from selected medical practitioners..
5. Patients should have the ability to delegate control over their health records to someone else under certain conditions (e.g. mental illness).
6. Managing access control policies should be an easy task, in order to ensure that the system is used and to preserve trust in the system.
7. It is important that legitimate uses of health records are not hindered, e.g. overall system availability service levels, and "need-to-know-" data access requirements in emergencies.

Ensuring each patient's privacy and data security is vital for an Electronic Health Record system. Unlike paper-based models, where an exposure or intrusion is confined to a single document or file, a federated EHR system creates the possibility of a patient's entire medical history being compromised by a single action. However, each of the traditional access control models, reviewed in Section 2, can satisfy only some of the above-listed requirements.

To see the access control weaknesses inherent in these previous models, consider the following scenario:

Frank prefers to go to two General Practitioners, Tony and Karen. Frank has two sensitive records in his Electronic Health Record, mental illness and sexual issues. Frank is happy to let Tony have access to his EHR, including his protected data field within his sexual record, but he wants to hide his protected data field within his mental illness record from Tony. On the other hand, Frank will allow Karen to access his EHR, including his protected data field within his mental illness record, but not his protected data field within his sexual record. Apart from these two GPs, Frank won't allow anyone to access his protected data field of his mental or sexual health records. In addition, Frank's father John suffers from Alzheimer's disease, so Frank must manage the access control rights to his father's EHR.

Even this simple and unremarkable scenario creates problems for each of the traditional access control policies, as explained below.

Discretionary Access Control: To use a DAC model we first need to know who owns the Electronic Health Record because DAC assumes that the owner of the data is the one who controls access to it. However, in healthcare, an EHR is partially owned by each of the patient, medical practitioner and medical authority [11], immediately creating an issue with respect to ownership. Furthermore, assuming that Frank has ownership of his EHR, he could nominate and grant access to his trusted/preferred medical practitioners (Tony and Karen), but it would be a difficult task for Frank to identify the specific medical data that is needed by each GP. The 'need-to-know' principle is required here, and in order to have it Frank is required to know the information that is needed for each medical practitioner and then set the access controls accordingly. By granting patients such control over their records, we may hinder the legitimate use of the EHR and, most likely, create another security problem due to the patient's mismanagement of their records. However, delegation of access control can be implemented easily in DAC, as Frank's father owns his EHR, so he can delegate the access control to his son.

Mandatory Access Control: In a MAC model, Frank won't have any sort of control, because the EHR system will be responsible for setting the security labels for users and EHR data objects. Therefore, Frank can't express his access control wishes over his EHR. Also, the 'need-to-know' principle can't be fully achieved here either, even if we apply a security level hierarchy. It's possible that two users might have the same security clearance (e.g. Tony and Karen), but should have different access permissions over a certain data object (e.g. mental

health data). In the MAC case, we can't assign more than one security label for the data object, therefore providing selective access to data objects is difficult. Moreover, there is no existence of delegation of access control due to the fact that the patient has no control over his EHR.

Role-Based Access Control: In an RBAC model, the 'need-to-know' principle can be satisfied by defining the permissions/operations that are required by a specific medical role, and this process could be done by an appropriate medical expert. However, in this situation Frank won't be able to hide his sensitive medical fields as he won't have any control over the permission assignments. In order to allow Frank to express his wishes, the security officer must allow Frank to modify the permissions, roles, user-role and role-permission assignments. Frank would need to create three roles in order to satisfy his needs, which would become an unacceptably time-consuming and complicated task for most patients and is likely to lead to a conflict of access control settings. Delegation of roles in RBAC is permitted if the security officer would allow Frank's father to delegate his roles to his son. Generally, RBAC seems a better choice than DAC and MAC, though it is still not an adequate solution for EHR system security.

In summary, it is clear than none of the existing models is adequate on its own, but that each of them has some feature which is essential to an EHR security model. DAC allows patients to control which data can be seen by particular medical practitioners, MAC allows the medical authority to control access to specific kinds of data, and RBAC allows access rights to be associated with certain medical roles.

4 A combined access control protocol

Although the access control requirement for Electronic Health Records cannot be satisfied by any one access control model alone, we contend that a careful integration of all three existing models is sufficient. Combining existing models, rather than developing an entirely new one for healthcare, allows us to take advantage of the well-understood properties and established implementations for these models.

4.1 Overview of the combined protocol

In the combined model access to a particular Electronic Health Record data item is granted only if it satisfies all three policies. The challenge is to determine where and how each of the access control constraints is introduced.

The basis for our combined protocol is shown in Fig. 1. An Electronic Health Record schema is shown where each EHR field has two MAC-based security labels: one is assigned by the patient and the second is assigned by the medical practitioner. These labels are used to express the sensitivity class of the data field. Also a DAC-style Access Control List (ACL) is maintained by the patient, whereby the patient nominates his/her preferred/trusted medical practitioners and sets the security clearance for each of them. This security clearance allows




Fig. 1. The logical structure of the combined access control protocol

the medical practitioner to access sensitive data that may not be allowed for other medical practitioners. Access to EHR fields is further restricted by overall RBAC-based access control managed by the medical authority.

4.2 Maintenance and enforcement of access control constraints

Each of the participants in the EHR system (patients, medical practitioners and medical authorities) needs to maintain some aspect of the combined access control policy, and is constrained in what information they can view as a result. In this section we describe the sequence of events needed to do this.

We will start with the patients' access control requirements, where the patients want to decide who is authorized to access their Electronic Health Records; and to determine what is the sensitive information in their EHRs, and who is authorized to access it. These requirements are satisfied by executing the following steps using the DAC and MAC interfaces in our combined access control policy:

1. The patients nominate the names of specific practitioners who they trust, and this is done through the use of DAC interface in Fig. 1 to construct the Access Control List (ACL).
2. To categorize data fields as sensitive/protected information, the patient needs to assign security labels to these data fields by using the MAC interface to update the patient's Electronic Health Record schema.

3. To allow specific medical practitioners to gain access to security-classified data fields in the patient's EHR, the patient, via the MAC interface, assigns the same security label of the sensitive data field to the authorized medical practitioners' ACL.

In practice, however, we do not suggest using the “no read up” and “no write down” rules that are introduced in MAC because it would be too complex a task for most patients to keep track of the transitive relationships introduced by a full hierarchy of security levels. Instead patients should just be presented with simple access/no-access settings.

Medical practitioners, as EHR consumers, have certain access control requirements that are important. Medical practitioners need to:

- access all the information that is required to fulfill their medical role in normal scenarios (e.g. a standard consultation with a GP), unless the patient has excluded that practitioner from accessing the particular data field;
- access all the information that is required in emergency cases regardless of the patient's access control settings; and
- hide some medical information from the patient.

Medical practitioners' access control requirements are also satisfied here. The following settings show how these requirements are met:

1. The Medical authority defines roles, permissions and role-permission assignments via the RBAC interface. This process is done by domain experts who know the access requirements for each medical role. Therefore, the ‘need-to-know’ principle is achieved and medical practitioners' access needs will not be limited unless the patient has set some access control restriction through either the DAC or MAC interface.
2. Since RBAC can incorporate contextual attributes into roles assignment, it would be possible for a medical practitioner to have an access role as a GP in a day clinic or as a GP in an emergency department. To allow the GP in an emergency department to access security-classified data records, the RBAC policy would assign a security label to these critical roles to allow them access to secure data. In an emergency case, the DAC constraints are not evaluated, due to the fact that the patient won't be able to know who the attending medical practitioners will be in an emergency.
3. To hide some medical information from the patient, the medical practitioner, through the use of the MAC interface, sets protection labels for these fields which hide the existence of such data in the patient's EHR.

Finally, the medical authority in charge of providing the Electronic Health Record and acts as the ‘security officer’ in the RBAC interface. It controls defining roles and permissions and the assignment of permissions to roles in order to associate specific medical roles with the information needed to fulfill them.

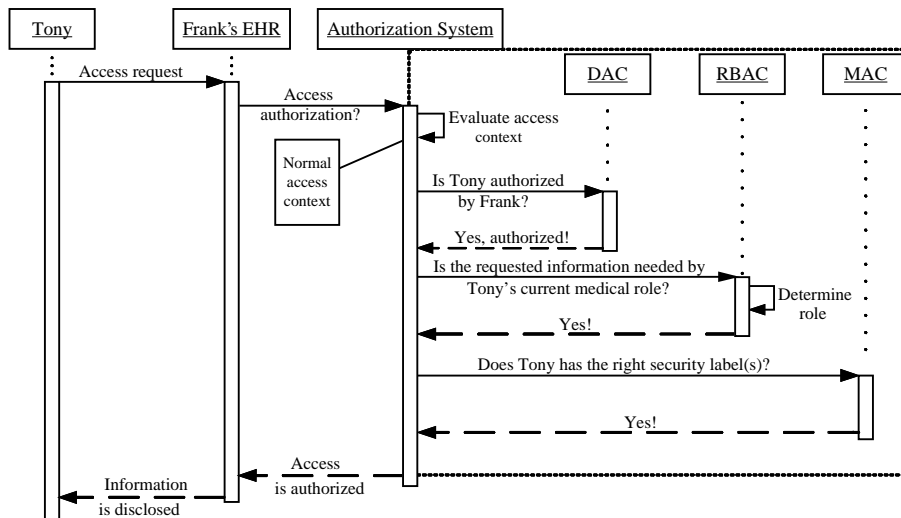


Fig. 2. The authorization evaluation process in the motivational example

4.3 Motivational example revisited

Now, let's revisit the motivational scenario from Section 3 to see how our access control protocol would satisfy Frank's wishes.

1. Frank will classify his mental and sexual data fields as 'sensitive' information by setting security labels for each record S1 and S2 respectively.
2. He will nominate his preferred GPs Karen and Tony to access his EHR. As Frank is happy to allow Karen to access his sensitive mental data field number 2, he will assign the S1 security label to her, which means that she is authorized to access any sensitive information that has a S1 label, in addition to her authorized access as per her medical role. For the same reason, Frank will assign the S2 security label to Tony.
3. When Tony requests access to Frank's Electronic Health Record, to see his sexual history, the following access evaluation occurs (Fig. 2):
 - (a) Evaluate access context, 'normal' or 'emergency'. If it's an emergency go to step 3c, otherwise continue.
 - (b) DAC policy: Does Frank authorize Tony to access his EHR?
 - (c) RBAC policy: Determine Tony's current medical role (e.g. day clinic GP, Emergency doctor) based on current contextual conditions.
 - (d) RBAC policy: Is the requested information needed by Tony's current medical role?
 - (e) MAC policy: Is Tony cleared to access this sensitive record?
 - (f) Tony is granted access to Frank's sexual data only if his access request passes all the steps above.

Also, as Frank needs to take responsibility for his father's Electronic Health Right, the following actions can be performed.

1. Frank's father John needs to delegate the control over his EHR to Frank through the use of the DAC interface.
2. Frank can now set the access rights to his father's EHR.

As well as these static assignments, we also need to consider temporary changes to access requirements. For instance, assume that Tony asks to see Frank's mental health record because he thinks that Frank's sexual issue is affected by some mental illness. This means that Frank must give Tony temporary access to his sensitive mental data field.

1. Frank will grant Tony another security label, S1.
2. Tony now has two security labels S1 and S2 from Frank, which means that he is authorized to access both of Frank's sensitive data fields contained in his sexual and mental health records.
3. After the consultation, Frank can revoke this permission by deleting S1 from Tony's profile.

On the other hand, a medical practitioner may need to change the status of certain fields. For instance, assume that Tony asks Frank to take a blood test which turns out to be positive for HIV. Given Frank's mental state, Tony would prefer to hide the pathology results until Frank's next in-house consultation.

1. Tony assigns a 'hide' flag to the HIV lab result field in Frank's EHR, so that Frank can't see any information contained in that specific field.
2. However, this information can be seen by Frank's authorized medical practitioners, such as the blood bank to which Frank regularly donates.

5 Conclusion and future work

Emerging plans for national Electronic Health Record systems raise new concerns about patient privacy and data security, by merging medical records that were previously kept separate and by making them accessible through single access points. None of the three standard access control models, Discretionary Access Control, Mandatory Access Control and Role-Based Access Control, are adequate for an EHR system in isolation. Nevertheless, we have explained how a careful combination of all three access control models can provide the security functionality needed for an EHR system.

At the time of writing we are assessing the security issues associated with a prototype Service Oriented Architecture for healthcare records. Our goal is to determine whether such an 'application-oriented' networking environment can be used to implement the combined access control protocol described above.

Acknowledgement This research was supported in part by ARC grant LP0776344, *Information Security Evaluation of Embedded Computer Software*.

References

1. K. J. Biba. Integrity Considerations for Secure Computer System. Technical report, Mitre Corporation, 1977.
2. P. Chhanabhai and A. Holt. Consumers are Ready to Accept the Transition to Online and Electronic Records if They can be Assured of the Security Measures. *Medscape General Medicine*, 9(1), 2007.
3. L. Demuyne and B. De Decker. Privacy-Preserving Electronic Health Records. In *Communications and Multimedia Security*, volume 3677 of *Lecture Notes in Computer Science*, pages 150–159. 2005.
4. M. Eichelberg, T. Aden, Riesmeier J., A. Dogac, and G. Laleci. A Survey and Analysis of Electronic Healthcare Record Standards. *ACM Computing Surveys*, 37(4):277–315, 2005.
5. HealthConnect Business Architecture, version 1.0., 2003.
6. D. E. Bell and L. J. LaPadula. Secure Computer Systems: Unified Exposition and Multics Interpretation. Technical report, Mitre Corporation, 1976.
7. D. Ferraiolo, D. Kuhn, and R. Chandramouli. *Role-Based Access Control*. Artech House, 2003.
8. B. Finance, S. Medjdoub, and P. Pucheral. Privacy of Medical Records: From Law Principles to Practice. In *Computer-Based Medical Systems, 2005. Proceedings. 18th IEEE Symposium on*, pages 220–225, 2005.
9. D. Tracy Gunter and P. Nicolas Terry. The Emergence of National Electronic Health Record Architectures in the United States and Australia: Models, Costs, and Questions. *Journal of Medical Internet Research*, 7(1):e3, 2005.
10. V. Hu, D. Ferraiolo, and D. Kuhn. Assessment of Access Control Systems. Technical report, National Institute of Standards and Technology, September 2006.
11. L. Iacovino. Trustworthy Shared Electronic Health Records: Recordkeeping Requirements and HealthConnect. *Journal of Law and Medicine*, 12:40–60, 2004.
12. I. Iakovidis. Towards Personal Health Record: Current Situation, Obstacles and Trends in Implementation of Electronic Healthcare Record in Europe. *International Journal of Medical Informatics*, 52(1–3):105–115, 1998.
13. G. Motta and Sérgio Shiguemi Furuie. A Contextual Role-Based Access Control Authorization Model for Electronic Patient Record. *IEEE Transactions on Information Technology in Biomedicine*, 7(3):202–207, 2003.
14. J. Park and R. Sandhu. Towards Usage Control Models: Beyond Traditional Access Control. In *SACMAT '02: Proceedings of the 7th ACM symposium on Access Control Models and Technologies*, pages 57–64, New York, USA, 2002. ACM Press.
15. P. Ray and J. Wimalasiri. The Need for Technical Solutions for Maintaining the Privacy of EHR. In *Engineering in Medicine and Biology Society, 2006. EMBS '06. 28th Annual International Conference of the IEEE*, pages 4686–4689, 2006.
16. W. Rishel, T. Handler, and J. Edwards. A Clear Definition of the Electronic Health Record. Technical report, Gartner, 2005.
17. R. Sandhu and P. Samarati. Access Control: Principles and Practice. *IEEE Communications Magazine*, 32(9):40–48, 1994.
18. R. Simon and M. Zurko. Separation of Duty in Role-Based Environments. In *IEEE Computer Security Foundations Workshop*, pages 183–194, 1997.
19. M. Wilikens, S. Feriti, A. Sanna, and M. Masera. A Context-Related Authorization and Access Control Method Based on RBAC: A Case Study from the Health Care Domain. In *Proceedings of the 7th ACM symposium on Access control models and technologies*, pages 117–124. ACM Press, 2002.