Modelling Context-Aware Security for Electronic Health Records

Pravin Shetty
Monash University, Australia

Seng Loke
La Trobe University, Australia

INTRODUCTION

The Internet has proven to be the most convenient and demanding facility for various types of businesses and transactions for the past few years. In recent years, business information systems have expanded into networks, encompassing partners, suppliers, and customers. There has been a global availability (Anderson, 2001; BSI Global, 2003) of resources over the Internet to satisfy different needs in various fields. The availability factor has called for various security challenges in fields where information is very valuable and not meant for all. Potential threats to information and system security come from a variety of sources. These threats may result in violations to confidentiality, interruptions in information integrity, and possible disruption in the delivery of services. So it is essential to manage the flow of information over the network with the required level of security. There are many security technologies and models that have been introduced which are capable of realizing the functions and objectives of information system security.

This article first gives a brief overview of what we term basic security policies of an integrated security model. Then it suggests context-based security policies for a health organization scenario using contextual graphs augmented with details about specific security actions, which relate to the security policies enumerated in the integrated security model.

The plan of the article is as follows. We first overview the three concepts in detail and briefly describe the concept of contextual (meta-policy) graphs. We then develop a context-based security meta-policy for securing patient records based on the security policies overviewed and discuss related work, before concluding the paper.

BASIC SECURITY POLICIES

Mobile ambients were first proposed by Cardelli and Gordon (1998a, 1998b) and then further extended by Bugliesi, Castagna, and Crafa (2004); and Braghin et al. (2002) were very efficient in modeling multilevel security issues. These three notions are very effective in modeling a foolproof security solution in a computing scenario by stating various security steps to be taken in the corresponding scenario. On this basis we have five cases that form the basic security policies in this article which we note can be concisely and precisely modeled using the mobile ambients formalism, though we omit such details of the formalism here and only describe the policies in plain language. The article uses them in appropriate scenarios depending on the context. Thus, the combined use of these five policies and a contextual graph representing the contexts of use of these policies provides a context-based security solution for pervasive environments. This section briefly describes the five policies using ambient (representing a boundary of security restrictions) notions.

Policy 1: Authenticate Returning Mobile Agent

When a privileged process (agent or person) leaves the parent ambient (e.g., a host institution) to execute some external independent activities, it relinquishes its local privileges and authority within its bounding parent ambient and ambient community. It exits the parent and might later return to the parent ambient. At this point an *authentication mechanism* is needed to check the authenticity of the returning original process. Cardelli and Gordon (1998a, 1998b, 1999) suggest that these high-level privileges must not be automatically restored to the returning agents/processes without first verifying their identity. This is to preserve the security and integrity of the ambient as well as the services and resources contained within it.

Policy 2: Firewall Access

If any agent/process has to enter an ambient, it has to know the name of the ambient and also possess the capability to enter it. The functionality of firewall is achieved with the help of restriction primitives and with the help of anonymity of the ambient name. Thus without knowing the ambient name, no process or agent can exit or enter the parent ambient. This helps in achieving protection of the resources from unwanted agents. The ambient name could be interpreted as a secret password.

Policy 3: Encryption Using Shared Keys to Secure the Data While Communicating

Cardelli and Gordon (1998a, 1998b, 1999) also put forth the encryption primitives to communicate between two ambients or between an ambient and a remote agent. These primitives helped in maintaining the *confidentiality* of the message or data. Consider a Plaintext message *M*. The encryption of the plaintext message is done with the help of the encryption key *k*. A name can represent a shared key, as long as it is kept secret and shared only by certain parties. A shared key can be reused multiple times, for example, to encrypt a stream of messages. A message encrypted under a key *k* can be represented as a folder that contains the message and whose label is *k* (Cardelli & Gordon, 1998a, 1998b).

Policy 4: Security Across Multiple Levels

In general, an enclosed ambient environment would typically contain numerous subambients as well as active processes, agents, and information resources. These groups of subambients within an ambient may be arbitrarily nested and organized in a hierarchical structure. Ambients and processes at the higher level of the nested structure are responsible for managing resources that are more vital and important than those at a lower level. In such multilevel environments, it is necessary to restrict the access to the flow of information depending upon the need and the security levels. Information can only flow from lower levels of security to higher levels and not conversely. A policy for this assigns levels to users and restricts information flow among the users.

Policy 5: Movement of Data and Entities Through Different Communities

The multilevel security policy mandatory access control security in the boxed ambients provided restricted access to information based on the various security levels in the hierarchical levels. The access is defined by the level at which the agents are which are predetermined based on their needs. But Braghin et al. (2002) were of the view that the implementation of mandatory access control security is complex, as agents and processes may move from one security level to another. The agents themselves may be confidential or may be carrying secure/confidential information. Thus there is no way of ensuring the agents will not be illegally attacked, accessed, or executed by untrustworthy entities at the lower security levels. The *security boundary* concept put forth by Braghin et al. (2002) guarantees absence of information leakage.

According to this concept, every high-level data or process should be encapsulated in a boundary ambient. A boundary ambient can be opened only when it is nested into another pre-specified boundary ambient. A policy for this states that the protected information cannot be read without being contained within some safety boundary (e.g., physically, an item cannot be viewed in the absence of a bodyguard).

CONTEXTUAL (META-POLICY) GRAPHS

Contextual meta-policy graphs are derived from contextual graphs (Braghin et al., 2002; Bugliesi et al., 2001a, 2001b). We replace the security actions in contextual graphs by security policies, which in turn, represent the security actions accordingly. By virtue of this embedding of policies (such as the five mentioned above) into a contextual graph, the graph becomes a meta-policy construct. The contextual meta-policy graphs are very general and can be used to depict security architecture in any scenario. The use of such graphs is to provide a high-level picture of the security framework, thereby avoiding lower details (security actions), which makes the overall architecture less cumbersome. The security actions are triggered according to the policies, which are predefined and programmed. The five security policies depicted in the integrated (in that such policies complement one another and have holistic coverage over security situations) security model above forms the basis of the graphs presented in this article. The contextual meta-policy graphs use various combinations of these five policies to define the various access paths to the patient's records. The next section presents the overall security architecture using such graphs.

SECURING ELECTRONIC HEALTH RECORDS USING CONTEXTUAL META-POLICY GRAPHS

The integrated security model discussed above is used to implement context-based security. The following model can be used in any context-aware situation. This article will explain the use of the model for securing health records. The various access possibilities depending on the context are described with the help of the contextual meta-policy graphs. This section is divided into two subsections. The first gives the details of the security policies used along with the low-level details (what security actions are triggered in each policy). The second subsection gives the actual approach for security purposes, along with an explanation of the individual access possibilities.

POLICY DEFINITION

This section outlines the various types of users associated with the electronic health records. It also specifies the various types of contextual information to be used for implementing the security policies. Finally, the various security actions to be taken depending upon the information received are also stated.

Figure 1. Context graphs showing the various possible scenarios and security actions taken

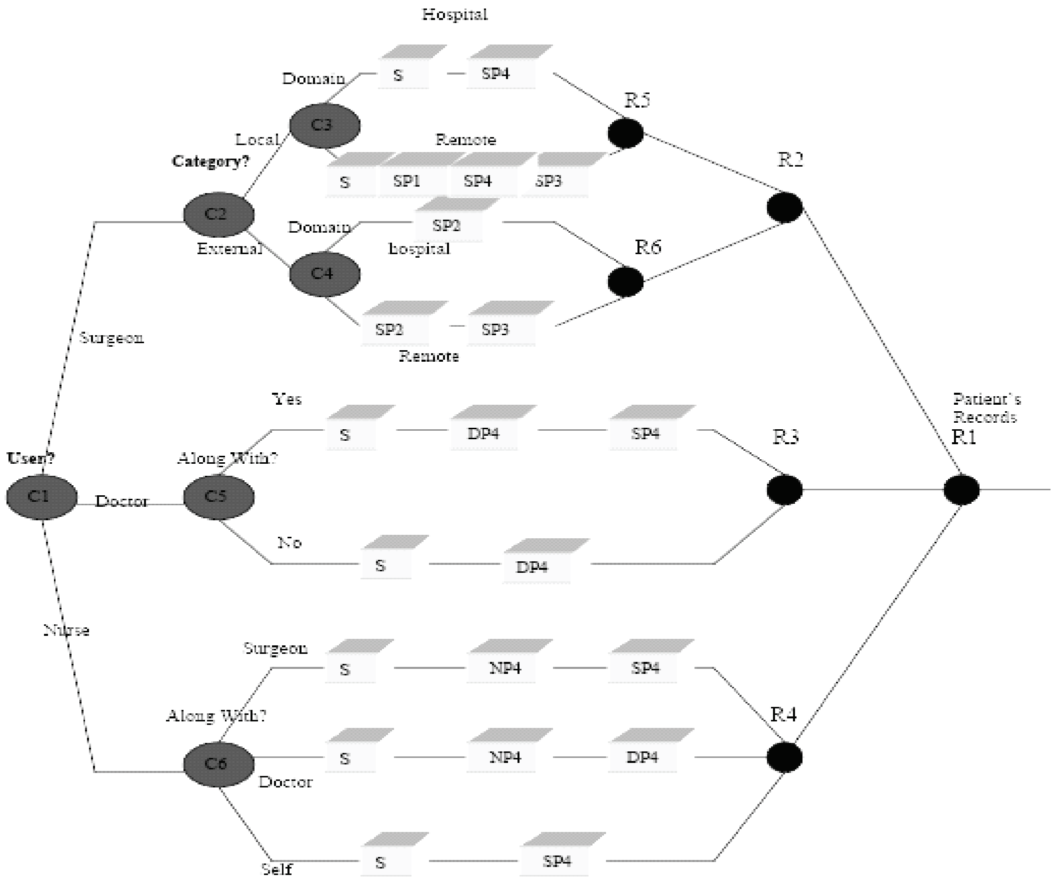

Roles: Surgeon, Doctor, Nurse

Contextual Information Considered: Role, Location, Place

The five policies, collectively labeled the 'integrated security model', can be effectively used for implementing security in a hospital scenario for accessing a patient's records. The five policies can be described in context for the hospital scenario as follows.

The Secret code *S* in the diagram can be considered as a secret name of the hospital network, which has to be known by each and every local person who wants to access the patient's detail. This secret code is similar to the secret name of the ambient in the case of ambient terms. It is not known to foreign entities. The details about the individual usernames and passwords are stored in the systems database along with the corresponding roles.

Policy 1: Authentication

This security policy defines the way in which a valid user can access the hospital network once s/he is out of it. This type of authentication is required for the surgeon and the doctor who have to access the network and hence the patient's health records. When the surgeon goes out of the hospital network, temporarily s/he is given a secret password, which helps her/him to authenticate her/himself when s/he wants to access the network again. Apart from the local surgeon and the local doctor, no one else can access the records from outside the hospital. This security policy triggers the action such as asking the user for the secret password, which will help her/him to authenticate her/himself.

Policy 2: Foreign Agent Authentication

This security policy defines the access method for a foreign entity such as a surgeon from some other medical institution in case of emergency. The security infrastructure provides such foreign agents, which are required in case of an emergency with some special combination of passwords analogous to that of the foreign agents concept in mobile ambients. This security policy triggers actions such as asking the foreign agent for the three sets of passwords given to her/him. The first will help her/him to validate her/himself. In case the access is remote, then it is a password; otherwise, it is done by retina scan or any other biometrics. The second password is used to allow the external surgeon inside the network. The third password will help her/him to access the patient's records in the mode as per her/his role.

Policy 3: Encryption of Data

This policy secures the data from falling into destructive hands by using various encryption techniques. This security policy is required when the user is accessing the network from a remote place. For example, there might be a case where a local surgeon might have to access the patient record from a remote place and do some modifications according to the present condition of the patient. This transfer of data should be safe and confidential and not intercepted by a malicious intruder. Thus, encryption is required which is provided by this policy.

Policy 4: Security Levels

This security policy talks about role-based access control. In the hospital, the three main users considered in this article are the surgeon, the doctor, and the nurse. They have access to the patient's records in different modes and according to the need of their positions. For example the surgeon has the access to the records in all three modes (i.e., read/write/update). A doctor's access is limited to two modes (i.e., read/write), whereas the nurse's access is restricted to just one mode (i.e., read). This type of security hierarchy is analogous to the security level structure policy of the multilevel access model. This helps control access to the valuable resources depending on the role of the user in the medical institution.

Policy 5: Third-Party Authentication

This security policy discusses authentication from the third mediating party in the communication between two entities from different medical institutions. This type of authentication is required to make sure that the entities involved in the communication are authorized. This security policy plays an important role when an outside surgeon needs access to the medical record or communicates with the local network of the medical institution from a remote place. The third party should be a reliable party, giving the authorization of any information passed between the two communicating entities.

The five security policies defined above are based on the concepts of mobile ambients defined in the former sections and depict the various security approaches taken depending on the context. The next subsection elaborates on the various contexts in a hospital scenario and the appropriate security policies to be taken to have secured access to the patient's records using a context meta graph.

Security Paths Defined in a Contextual Meta-Policy Graph

The above model gives the overall security approach used in a hospital environment to access electronic health records of the patients. The behavior of the security infrastructure according to various contexts is described as follows:

1. **User→ Type→ Domain :: Surgeon→Local→Hospital:-** ***Policy 4.*** If the user is a local surgeon and s/he wants to access the patient's records from the hospital, then s/he has to follow security policy 4. When s/he tries to access the records, the system will ask for her/his username and password, which will depict her/his role in the institution. Being a surgeon, s/he enjoys the highest level of rights. S/he can access the records in all the three modes (read/write/update). For access from within the hospital, authentication can be provided by various biometrics.
2. **User→ Type→ Domain :: Surgeon→Local→Remote:-** ***Policy 1 + Policy 4 + Policy 3.*** When the local surgeon, say John, wants to access the information from outside the hospital, which can be in the case of an emergency, then a combination of three security policies is followed. First he has to get back into the hospital network by providing his secret key/password according to policy 1, which will validate him as a local user. Then with the help of his username and password, which is associated with his role (policy 4), he can access the records. As he is in a remote place, care should be taken that the data is not visible to the outside world. This is achieved by encrypting the channel according to policy 3. Thus, using the combination of three policies, the local surgeon is provided access to the records from a remote place.
3. **User→Type→Domain :: Surgeon→External→ Hospital:-** ***Policy 2.*** In some emergency cases, it becomes necessary to invite an external surgeon to the local institution. For such cases, the system stores all the information of such emergency persons. Policy 2 defines the access for such a surgeon using predetermined passwords. The surgeon is provided with three sets of passwords. The first will help him/her to validate him/herself. The second password is used to allow the external surgeon into the network. The third password will help him/her access the patient's records in the mode according to his/her role. Biometrics instruments can provide the required authentication when he/she is accessing from the hospital. If the external surgeon is accessing the records from the hospital, then he/she can use his/her third password to get access according to his/her role. The first and the second passwords are not required in this case.
4. **User→Type→Domain :: Surgeon→External→ Remote:-** ***Policy 2 + Policy 3.*** When the external surgeon has to access the records from a remote place, then the access takes place as determined by the combination of two security policies. Policy 2 is as defined as above. In such a case, he/she must have all three sets of passwords with him/her in order to get into the hospital network and then access the patient's records. Policy 3 is used for encryption of the information to provide safe and confidential communication.
   The surgeon, whether local or external, does not need to be with any other staff, as he/she enjoys the maximum access rights. He/she has to follow the appropriate security procedures depending upon his/her category (i.e., local or external) and his/her location of access.
5. **User→Along with :: Doctor→ Along with Surgeon:-** ***Policy 4 (for doctor) + Policy 4 (for surgeon).*** When a local doctor wants to access the records, then s/he can do so by using a combination of policy 4 applied to two roles. If s/he is with the surgeon, then s/he can access the patient's record in a full mode (read/write/edit) with the same access privileges as the surgeon. But for this, the surgeon has to first specify her/his

username and password according to policy 4 so that her/his role is specified. If the surgeon is not physically present in the hospital, then the doctor cannot access the records in full mode. Also, a doctor is not allowed access to the records from a remote place.

6. **User→Along with :: Doctor→ Alone:-** ***Policy 4 (for doctor).*** If a surgeon is not with the local doctor, then the doctor accesses the patient's information in the restricted mode (read/write). S/he can access the information according to policy 4. S/he has to present her/his username and password so that her/his role will be represented in the system. An important point in this security architecture is that the doctor can never access the patient's records from a remote place.
7. **User→Along with :: Nurse→ Along with Surgeon:-** ***Policy 4 (for nurse) + Policy 4 (for surgeon).*** A nurse, say Jane, is at the lowest security level in the hospital hierarchy. If she is with the surgeon, she is allowed to view the records in full mode as the surgeon has. For that, the surgeon must present his/her role to the system first, and then the nurse can access the information according to her role as in security policy 4. Remote access is not allowed in this case.
8. **User→Along with :: Nurse→ Along with Doctor:-** ***Policy 4 (for nurse) + Policy 4 (for doctor).*** When the nurse, Jane, is with a doctor in the hospital, then she can have access to the patient's information according to the doctor's privileges. The doctor first provides his/her username and password, and then the nurse can provide her role as per security policy 4. Remote access is also not allowed in this case.
9. **User→Along with :: Nurse→ Alone:-** ***Policy 4 (for nurse).*** When nurse Jane is alone, she is only allowed one mode (read). She can read the information but cannot delete or modify it. She can access the information using her username and password according to security policy 4. But she is not allowed to access the records from a remote place.

RELATED WORK

There has been some work on security policies in the field of electronic health records systems in past years. Reid, Cheong, Henricksen, and Smith (2003) presented a model that uses role-based access control to restrict the access to the health records on a need-to-know basis. The prototype described maintains databases consisting of explicit 'allow' and explicit 'denial' lists. The proposed model also permits allow and deny policies to successively qualify each other in a role hierarchy supporting inheritance. Thus, the access control framework exhibits a great flexibility and efficiency in the range of access policies that it can support.

Mostéfaoui and Brézillon (2004) put forth the concept of contextual graph for modeling security in context-aware environments. They present a new model for policy specification based on the new approach. The security policy based on such an approach depends on the contextual information of the user and the environment. Contextual graphs have proved to be very effective in modeling a complex situation. Mostéfaoui and Brézillon (2004) also mention how contextual graphs are used to model security in a context-aware environments. In their paper they gave an example of how context-based security is used in a hospital scenario, but it does not employ our meta-policy scheme in the way we do above.

CONCLUSION

Due to the ubiquitous nature of the today's computing world, security is of utmost important. The traditional static authentication techniques are no longer valid and justified. This situation is due to the lack of consideration for context in existing security systems. Context-based security helps the security policy to, in effect, adapt to the new "threats" as they come. It aims at providing flexible security models for distributed infrastructures, where the user and application environments are continually changing. In this article, we have presented an approach that helps with context-based security in a medical scenario. The type and nature of the authentications that are demanded by the security policy depend on the information that is collected from the environment. Further, the contextual graph approach helps to add/modify secure paths based on the newly detected contexts that need to be utilized for fine-grain security. The model presented is a generalized model that can be used in any context-aware environment or enterprise, from the office to factories.

REFERENCES

Anderson, R. (2001). *Security engineering.* New York: John Wiley & Sons.

Braghin, C., Cortesi, A., & Focardi, R. (2002a). Security boundaries in mobile ambients. *Computer Languages, 28(1),101-127.*

Braghin, C., Cortesi, A., Focardi, A., & van Bakel, S. (2002b). *Boundary inference for enforcing security policies in mobile ambients.* Retrieved from http://www.informatics.sussex.ac.uk/users/vs/myths/reports/papers/boundaries-tcs02.pdf

BSI Global. (2003). *Information security.* Retrieved from http://www.bsi-global.com/ Information+Security/ Overview/Why.xalter

Bugliesi, M., Castagna, G., & Crafa, S. (2001a). Boxed ambients. *Proceedings of TACS* (pp. 38-63). Berlin: Springer-Verlag (LNCS 2215).

Bugliesi, M., Castagna, G., & Crafa, S. (2001b). *Reasoning about security in mobile ambients* (pp. 102-120). Berlin: Springer-Verlag (LNCS 2154).

Bugliesi, M., Castagna, G., & Crafa, S. (2004). Access control for mobile agents: The calculus of boxed ambients. *ACM Transactions on Programming Languages and Systems, 26*(1), 57-124.

Cardelli, L. (1999). *Abstraction for mobile ambients.* Retrieved from http://research.microsoft.com/Users/luca/Papers/Abstractions%20for%20Mobile%20Computation.A4.pdf

Cardelli, L., & Gordon, A.D. (1998a). Mobile ambients. *Proceedings of FOSSACS* (pp. 140-155). Berlin: Springer-Verlag (LNCS 1378).

Cardelli, L. and Gordon, A. D. (1998b). *Mobile ambients.* Retrieved from http://www.cis.upenn.edu/~lee/98cis640/Lectures/fm3.ppt#24

Cardelli, L., & Gordon, A.D. (2004). *Mobile ambients.* Retrieved from http://classes.cec.wustl.edu/~cs673/1

Covington, M.J., Fogla, P., Zhan, Z., & Ahamad, M. (2002). *A context-aware security architecture for emerging applications.* Retrieved from http://www.acsac.org/2002/papers/71.pdf

Mostéfaoui, G., & Brézillon, P. (2004a). Modeling context based security with contextual graphs. *Proceedings of the 2nd IEEE Annual Conference on Pervasive Computing and Communications Workshops.*

Mostéfaoui, G., & Brézillon, P. (2004b). Context-based security policies: A new modeling approach. *Proceedings of the 2nd IEEE Annual Conference on Pervasive Computing and Communications Workshops.*

Reid, J., Cheong, I., Henricksen, M., & Smith, J. (2003). *A novel use of RBAC to protect privacy in distributed health care information systems* (pp. 403-415). Information Security Research Center, Queensland University of Technology, Australia.

Wikipedia. (n.d.). *Ubiquitous computing.* Retrieved from http://en.wikipedia.org/wiki/Ubiquitous_computing

KEY TERMS

Contextual Graph: Graph whose edges represent the values that context information take, and have three types of nodes: branching and recombination nodes, and nodes representing security actions. A path through the graph represents a security action taken in response to particular context information.

Mobile Ambients: A process calculi that emphasizes the notion of boundaries and how processes with such boundaries interact.

Pervasive Computing: Integrates computation into the environment, rather

than having computers which are distinct objects. Other terms for ubiquitous computing include ubiquitous computing, calm technology, things that think, and everyware (Wikipedia, n.d.).

Security Action: Action taken to secure a resource, from authentication to encryption to other informational and physical measures (e.g., putting a man on guard).

Security Policy: A description of security actions to take under different circumstances. Such policies are typically specified as rules in a formal language.