Application of a Multilevel Access Model in the Development of a Security Infrastructure for a Clinical Information System

Steven J. Henkind, M.D., Ph.D.
First Consulting Group
Mamaroneck, New York

Janis M. Orlowski, M.D.
Rush-Presbyterian-St. Luke's Medical Center
Chicago, Illinois

Patricia C. Skarulis
Rush-Presbyterian-St. Luke's Medical Center
Chicago, Illinois

ABSTRACT

A number of security models including the military model, the Institute of Medicine model, and the matrix model have been utilized, or proposed, for protecting clinical information systems. These models have a number of limitations, however, and of particular concern, they focus on security as opposed to access. In this paper we describe a multilevel access model which can overcome some of these limitations. This model is currently being utilized in the development of an improved security infrastructure for a clinical information system.

INTRODUCTION

Until recently, clinical information system (C.I.S.) security issues have received relatively little attention. Although there has been a great deal of research on the topic of computer security, review of the literature reveals only a modest body of work related specifically to the security of clinical information systems; e.g., [6,11,14].

A number of factors, however, are leading to increased recognition of the importance of security issues. For example, the Joint Commission's information management audits will include security [7], and several pieces of pending legislation explicitly refer to the protection of computerized patient data, e.g, [15]. In addition, both published [1,8,12] and unpublished reports of recent C.I.S. security breaches have served to increase awareness.

Although it is essential to have adequate security, it is also true that excessive security can be detrimental (for example, if a physician cannot access lab data on an individual in extremis because "it is not his patient"). There is frequently a tradeoff between access and security and the two must be carefully balanced [5]. There are other tradeoffs as well; for example, security mechanisms can significantly increase system overhead and thereby degrade system performance, and certain security mechanisms may be inordinately expensive.

Conceptually, security mechanisms may be broken into three categories: authentication, authorization, and audit. Authentication involves a user proving their identity (e.g., by means of a password or biometric system) in order to log on to the system. Authorization involves the specification and enforcement of access rights to specific pieces of data. Audit is the process of checking for security breaches.

In this paper we describe a multilevel access model which we developed in order to improve authorization control for the Rush-Presbyterian-St. Luke's C.I.S. A particular emphasis in construction of this model was the need to balance access and security. The model was developed after careful consideration of a number of other models that have

0195-4210/92/$5.00

been used (or proposed) for securing clinical information systems: 1) the military model, 2) the Institute of Medicine (I.O.M.) model, and 3) the matrix model. These other models are described below.

SECURITY MODELS

The basic idea of the military model [3] is to classify documents hierarchically in terms of their sensitivity; e.g., "top secret," "secret," "confidential," and "unclassified." Conceptually, one may think of this model as an onion skin with the most sensitive documents on the inside, and the least sensitive on the outside. Users are allowed to access information only up to the depth which corresponds to their maximum security clearance. These "mandatory" controls can be supplemented with "discretionary" controls that allow a finer level of granularity within a given layer. Note, however, that discretionary controls are overlaid on top of mandatory controls, but do not supersede them. It is important to note that this model is best viewed as a security model as opposed to an access model [11] - although a user's security clearance should be determined on a "need to know" basis, there is very little emphasis on the modeling of users needs, rather the emphasis is on the sensitivity of the data.

Although this model has worked well in the military, it is not clear that it is appropriate for healthcare. For example, there are many situations wherein a healthcare worker may need access to sensitive data, but not to nonsensitive data; e.g., a social worker will need to know if a client is an alcoholic (highly sensitive), but probably does not need access to a patient's blood type (less sensitive). As another example, although an orderly may need to know that a patient has a contagious disease (so that s/he will wear a face mask), it is unlikely that s/he would need to know a patient's marital status. The model has other limitations as well (e.g., some of the limitations of the matrix model).

The Institute of Medicine model [4] proposes that data be classified into three levels: "extremely sensitive," "sensitive," and "least sensitive." This hierarchical model is very similar to the militaries' mandatory access control model, and it suffers from the same limitations.

The basic idea of the matrix model is that each unit of data has an associated set of permissions for each user, where a unit of data can be a file in the context of an operating system [9], a relation or, in some cases, a field in the context of a database [2], etc. Permissions specify the actions that a user can perform (e.g., read, write, execute, etc.). Typically, such a model is implemented by means of a matrix. Note that, for reasons of computational efficiency, users may be classified into groups, as may units of data.

Among the strengths of the matrix model are: 1) It can support a higher level of granularity than the military model. 2) In theory, it is easy to implement the user's access requirements - merely fill in the matrix.

Among the limitations of the model are the following: 1) although it allows for easy implementation of access requirements, it does not assist in the determination of those requirements. 2) The matrix model does not address certain issues at all, for example, cross patient searches (e.g., a given patient's creatinine and address is probably not sensitive, but a list of all patients's creatinines and addresses might be - because this list could be used as market research by an enterprising nephrologist!). 3) The matrix model has theoretical limitations in certain domains; e.g., in a relational database using views, it is easier to achieve finer granularity of control over read access, as opposed to update access [13].

A MULTILEVEL ACCESS MODEL

In an effort to overcome some of the limitations described above, we designed a multilevel access model. This model consists of four levels of principles and issues. The key components of the model are described below. Note that, in the examples that follow, we focus on physicians, but the discussion can be generalized to other healthcare workers as well.

1) General Access (and Security) Principles

A) It is nearly impossible to achieve complete security; e.g., a persistent hacker can penetrate virtually any system if given enough time.

B) Access and security are (generally) inversely related.

2) General Access Principles For Healthcare

A) The need to access data is role specific; a given physician's need to access data depends on the particular role that he or she is playing at a particular point in time. The types of roles may be classified broadly as patient care, administration, and research (note that there is potential overlap here, e.g., Q.A., U.R., education, etc.).

B) The time urgency of access to data depends on which particular role a physician is playing:

High Patient Care

⇓ Administration

Low Research

C) The need for individual patient data depends on which particular role a physician is playing:

High Patient Care

⇓ Administration

Low Research

D) The need for data across patients depends on which particular role a physician is playing:

Low Patient Care

⇑ Administration

High Research

Example for principles B, C, and D: A physician performing patient care activities might (B) immediately need the hematocrit of (C) the individual patient Jones, but would not need (D) a list of all patient's hematocrits.

E) In an emergency, a physician may need access to any data.

3) Access Issues That Will Need to be Resolved in All (or Most) Healthcare Institutions

Note that the particular resolution of these issues will differ from institution to institution, depending upon the underlying environment; e.g., teaching versus community hospital, closed versus open medical staff, etc.

A) Are individual patients "assigned" to individual physicians (i.e., can physicians only access data on their "own" patients)?

B) What data is a physician authorized to access on a given patient (e.g., can physicians access hospital financial data as well as clinical data)?

C) Are cross patient searches allowed?

D) Is the construction of secondary databases allowed?

4) Access Issues That May Need to be Resolved Within Specific Healthcare Institutions

A) Certain types of access within psychiatric hospitals.

B) Certain types of access within substance abuse treatment centers.

UTILIZATION OF THE MODEL: A CASE STUDY

We developed an initial version of the multilevel access model approximately a year and a half ago. Since that time we have refined the model and, in addition, we have used it as a fundamental framework for expanding and improving the security of the Rush-Presbyterian-St. Luke's clinical information system.

Rush-Presbyterian-St. Luke's Medical Center (RPSLMC) is comprised of several hospitals, with over a thousand total beds, medical and other graduate schools, and a number of other healthcare facilities. The RPSLMC clinical information system services the institution by means of a campus-wide network.

Our basic approach has been to use the access model to identify and clarify access requirements, and then use those access requirements to develop security

requirements and mechanisms. The following are some examples (from Rush-Presbyterian-St. Luke's Medical Center) of how the principles and issues specified in the model can be applied to the development of a security infrastructure.

Principles 1A) (it is difficult to defeat a determined hacker), and 1B) (access and security are inversely related): We have made an explicit decision to accept certain types of security vulnerabilities in exchange for encouraging access. In order to minimize risks, however, we are developing a comprehensive set of policies and procedures, audit mechanisms, and penalties, so that transgressors can be detected and penalized.

Principle 2E) (in an emergency, a physician may need any data): In light of this principle we will utilize a "Code Red" facility wherein a physician may, at his or her own discretion, circumvent most security mechanisms in order to access any patients' information. In order to prevent abuses, however, code red accesses will be subject to mandatory, and detailed, investigation.

Model Level 3) (issues requiring resolution within a given institution): In order to resolve specific access issues we have convened several committees with representatives from the medical staff, administration, and other departments. These committees are currently formulating policies, procedures, and detailed profiles of user access rights.

Issue 3A) (assign individual patients to individual physicians?): A number of hospitals have implemented this successfully [10], but in a teaching hospital this may be extremely difficult, if not impossible, to accomplish. Accordingly, we have resolved this issue by electing not to make such an assignment.

Issue 3B) (what data is a physician authorized to access?): We have elected to allow physicians to access most, but not all, data on a given patient (e.g., they will not be able to access certain types of hospital financial data). This will be implemented by means of a matrix. Access rights for non-physicians will probably be more highly restricted. A number of multidisciplinary committees are currently formulating guidelines.

Issue 3C) (cross-patient searches): In general, most cross-patient searches are used for research or administrative purposes (Principle 2D). Therefore, they are generally not time urgent (Principle 2B). Hence, general query language facilities will not be available on the system (arbitrary searches will need to be approved). Since certain cross-patient searches may, in fact, be needed for patient care, e.g., infection control, these searches will be pre-authorized and available by means of a menu.

Issue 3D) (creation of secondary databases): Any secondary database to be created is almost certainly intended for research or administrative use, hence, creation of such a database is generally not time urgent (Principle 2B). Therefore, file transfers will need prior approval (security based upon software and policies/procedures). In addition, workstations, at certain locations, will be diskless (security based upon hardware).

ADVANTAGES OF THE MODEL

The model can be used to overcome some (although not all) of the limitations of other approaches (e.g., it assists in the formulation of access requirements).

The model is at a high enough level that it encourages one to "see the forest through the trees" (e.g., initially concentrating on the creation of a security matrix will likely be extremely time consuming - and it is also likely that important issues will be overlooked).

The model is independent of specific hardware, software, and policies and procedures.

One of the most difficult tasks in implementing a security infrastructure is establishing user "buy in." This model is simple, and focuses on user needs as opposed to system needs. As a result, we have found that the model is easy to present to clinicians and administrators, and, so far, it has met with ready acceptance.

OPEN ISSUES

Although we believe that our model is both powerful and usable, there are a number of issues which it does not address successfully and which would make good grounds for future research:

It can be difficult to protect the contents of lengthy unstructured notes (because they cannot be readily broken into meaningful fields, the level of granularity

of access must be quite coarse - generally all or nothing). As clinicians begin to move towards more structured formats for recording information (e.g., input into encounter forms or templates), this problem will become more tractable.

Access by inference: Certain data elements, although individually non-sensitive, may become sensitive when aggregated together; e.g., certain patterns of liver enzymes are suggestive of alcohol abuse. This issue has been investigated in other domains, however, it is not clear whether solutions from these domains would be effective - or even appropriate - in the context of healthcare.

Field content: Certain fields may become sensitive depending upon their contents - e.g., the diagnosis field may not be sensitive (usually), but certain diagnoses will make it sensitive (e.g., a diagnosis of appendicitis is not sensitive, but a diagnosis of HIV is). Similar issues have been investigated in the context of data base management; for example, some relational database systems can control access depending upon field content (e.g., forbid a SELECT for all salaries > $50,000). These techniques would seem to warrant further investigation in the medical domain.

CONCLUSION

A variety of factors are leading to increased recognition of the need to adequately secure clinical information systems. It is not clear that conventional security models are workable, or even appropriate, in the healthcare domain. In particular, they emphasize security as opposed to access. It is our belief that new models - based upon the access requirements of users - will be needed. We have developed one such model, and have been utilizing it successfully as a framework for improving security.

References

[1] Brophy JT, Tresnowski BR. Workgroup for Electronic Data Interchange: Report to Secretary of U.S. Department of Health and Human Services. July 1992; Appendix 4 pg 32.

[2] Date CJ. An Introduction to Database Systems: Volume II. Reading, Massachusetts: Addison-Wesley, 1983.

[3] Department of Defense, Trusted Computer System Evaluation Criteria. August 1983.

[4] Dick RS, Steen EB, editors. The Computer-Based Patient Record: An Essential Technology for Health Care. Washington D.C.: National Academy Press, 1991.

[5] Gardner E. Computer Dilemma: Clinical Access vs. Confidentiality. Modern Healthcare November 3 1989, 32-42.

[6] Griesser G, Jardel FP, Kenny DF, Sauter K (eds). Data Protection in Health Information Systems - Where do We Stand. Amsterdam, North-Holland, 1983.

[7] Joint Commission on Accreditation of Healthcare Organizations. Draft: Principles For Information Management in Health Care Organization. 1992.

[8] Juni JE, Ponto R. Computer-Virus Infection of a Medical Diagnostic Computer: Letter to the Editor. NEJM 1989;320 #12:811-812.

[9] Kernighan BW, Pike R. The UNIX Programming Environment. Englewood Cliffs, New Jersey: Prentice-Hall, 1984.

[10] Minard B. Full-Time, Real-Time System Security. Computers in Healthcare October 1987;51-57.

[11] Orr GA, Brantley BA. Development of a Model of Information Security Requirements for Enterprise-Wide Medical Information Systems. SCAMC 1992;16:287-291.

[12] Pasternack A. Lawyers Eye On-Line Records: Patient Data, Business Issues Top Concerns. HIMSS News Oct. 1992; vol 3 #8:1-5.

[13] Sandhu RS, Jajodia S. Limitation of Relational Data Base Access Controls. Information Systems Security 1993;vol 2 # 1,57-71.

[14] Shea S, Sengupta S, Croswell A, Clayton PD. Network Information Security in a Phase III Integrated Academic Information Management System (IAIMS). SCAMC 1992;16:283-6.

[15] Stark. H.R. 200. 103d Congress, 1st Session. January 5, 1993.