Development of a Model of Information Security Requirements for Enterprise-Wide Medical Information Systems

Geoffrey A. Orr, Ph.D.
B. Alton Brantley, Jr., M.D., Ph.D.

Center for Information Technology
The Pennsylvania State University College of Medicine
Hershey, Pennsylvania

ABSTRACT

Information security methods developed within the narrow frameworks of operating system design, specific database models, and military security methods all concentrate on representation of the objects of access control, rather than on the information needs of the subjects. This approach does not adequately support the needs of the varied users of medical information systems, who must have access to information in support of multiple organizational roles. A new conceptual approach to access control in medical settings based on user requirements is discussed.

INTRODUCTION

As enterprise-wide information systems emerge as part of the transition to open networked systems [8], issues of how to manage the conflicting access and security needs of the organization become crucial, particularly in the medical center setting. An ongoing review of the literature in computer and information system security reveals a fragmented body of knowledge, with efforts focused primarily on operating system, network [10, 3], and database [2, 4, 5, 9, 13, 15 16] security. These approaches are designed primarily to support Department of Defense requirements and control network intrusion attacks [10], rather than a broad domain-independent/portable view of security. Especially absent in this National Security and *hacker*-oriented discussion is focus on sophisticated security mechanisms in the areas of medical and clinical information systems.

The discussion which follows identifies requirements for a more responsive model of permissions in the healthcare delivery setting, and a preliminary description of such a model. This approach supports the notion that an implementation independent modeling system for information access must be developed and used, so that policy decisions mediating access to data are not colored by technical considerations [9, 14].

OPERATING SYSTEM SECURITY

Much of the discussion of data security is focused on security mechanisms as an aspect of operating system design. The central issue in this discussion is the establishment of access rights or capabilities, which allow subjects (users, processes, etc.) various kinds of access (read, write, execute, etc.) to objects (physical or abstract objects, including i/o devices and files). These relationships form an access control matrix, where the subjects and objects of access control make up rows and columns, and specific sets of permissions make up the cell entries [3].

An example of this approach is the widely implemented Unix operating system access control scheme, where users are granted ownership of the files and processes that they create, and are allowed to control read, write, and execute permissions on these objects for themselves and other users. Access control for non-owners of objects is provided by the assignment of users to groups. The owner of a file may assign a specific group name to that file, and then specify which of read, write, or execute permissions the group members may apply to that file. By creating a rich structure of groups, flexible file permissions may be provided. In this scheme, the subject dimension of the access control matrix contains user names (as owners of objects), group names, and the user cate-

0195-4210/92/$5.00

gory "other" (users other than the owner or group members). The object dimension of the access control matrix contains file and pseudo-device names.

Capabilities-based systems [6] do not rely on the less flexible properties of object ownership by individual subjects and assignment of object access rights to subject groups, but consider only subjects, objects, and access types. Thus, a capability consists of "data presented by a process to gain access to an object" [12], or "an address to some information ... and a finite set of rights" [14]. Access rights may still be associated between subjects and objects using a matrix representation, but the model does not depend *a priori* on the notion of object ownership or group membership. Capability models usually allow the right to copy and transfer a capability from one subject to another, providing for the propagation or withholding of access rights from subject to subject.

This aspect of the capabilities approach (and the concept of granting "group" and "other" privileges in Unix) provides an example of a discretionary control system [7], where subjects are granted the ability to allow or deny access to an object by other subjects. Identification and control of access at the level of abstract subjects and objects allows security models which provide fine grained, highly flexible access control over operating system objects. Discretionary control is one of the least rigorous security mechanisms in the Department of Defense specifications for secure systems [7]. A more stringent requirement is for access control lists, where a single user subject must be explicitly given access to a single object by the central access control authority, without intermediate definition of a subject group.

DATABASE SECURITY

Discussion of the capabilities model in the context of operating system design has focused on the subjects and objects commonly found at the operating system level- users, processes, files, and other addressable objects such as peripheral devices. Database management systems (DBMS) provide facilities that hide the physical aspects of data storage and retrieval by providing higher level logical data models, data definition languages, and query languages. Since DBMS provide their own set of logical abstractions for objects, they have also needed to provide access control mechanisms for those objects.

In SQL (structured query language) relational DBMS, access control has been implemented using the concepts of ownership, access permissions explicitly authorized using the "grant" command, and definition of database views (virtual relations) [2]. In this approach, the fundamental object of access control is the relation, which consists of a set of fields and their data. Relations are owned by the users who create them, and those owners can grant permissions (such as SELECT, INSERT, or UPDATE) to other users. In some cases, the notion of groups of users to whom access is granted has been implemented, using either groups defined at the operating system level, as in SQL/RT [16], or by special designation within the DBMS [11]. The notion of discretionary control has also been extended in some DBMS by means of the "grant with grant" option, as seen in DB2, for example [2]. This option allows the owner of a relation to grant to another user the authority to grant specific permissions to other users.

The ability to manage access at the level of individual fields within a relation, and to further restrict access based on values of the field, is provided by means of the database view. Views are subsets of actual relations that are stored as definitions rather than as actual data. When the view is activated (through an SQL statement), the data are accessed from the base relation using the view definition. The use of views allows a single relation to hold data that supports the access needs of a number of users who could not be allowed access to the base relation for security reasons. Views also allow subsets of data from multiple relations to be brought together to form user-specific sets.

SUBJECTS IN ACCESS CONTROL SYSTEMS

In the case of both operating system and database security, emphasis has been focused on classification and definition of the objects that are accessed, rather than on the subjects that perform the access. What is missing from these discussions is a detailed model of the "user as a subject" and the user's information needs.

This problem is even more serious when one considers the impact of the military classification and access orientation on current thinking in system and database security. The military approach to classification of objects has used an ordinal scale of sensitivity or secrecy, consisting of levels such as Confidential, Secret, and Top Secret. Additional controls or compartments are established based on similarity of content or domain. Access to multilevel objects of this type is granted based on the security clearance of an individual, and governed by the principles of "least privilege" and "need to know" a particular piece or collection of information [7]. Usually, "need to know" is determined by job assignment and tends to be identified statically. Current work in secure databases is focused on the requirements imposed by this approach and preventing accidental release of data to a user of the system who is not cleared for that level of data [5, 13, 15].

With the development of central integrated databases, and integrated access to distributed databases in support of broad institutional information utilization, a narrow approach to user authentication and permissions is inadequate, because it cannot flexibly support the varied levels of information use in the enterprise.

Instead of only focusing on characterization, classification, and manipulation of the objects of an access control system, it is crucial to focus on characterization and classification of the subjects of the access control system, e.g. its users. The basis of this is a model of the information utilization policies and patterns of the institution as it exists, including both computer- and paper-based information processing. Such a model will reflect the specifics of the institutional/professional domain and will help define requirements for the actual access control system. In terms of the military model of access control, this approach emphasizes the detailed description and organization of the "need to know" as it exists in the organization.

Such a model of the information use of an organization is complex, and building it constitutes an empirical task, initially involving more observation and description than abstraction and prescription. Subsets of information must be explicitly associated with identifiable activities in the organization, and thus with groups of individuals. As information utilization is documented, patterns may emerge that will guide more closely focused modeling efforts. Some areas will benefit from more sophisticated modeling, while others will be adequately represented by simple documentation.

Overly broad and unduly restrictive definitions of access can be caused by an inflexible control system that is based on a poor understanding of how information is used in support of specific tasks. Inflexibility in determining what a specific job function's need to know and commensurate access should be is negligent execution of the responsibility to allow or deny access in support of patient care and confidentiality.

A MODEL OF COMPLEX DATA ACCESS

A model of data access requirements can be viewed as a model of data use, if we assume that the ordinary information use of individuals in the organization is the starting point for potential authorized use. It is proposed that data use in the context of an organization can be expressed in terms of one or more roles, each of which consists of one or more tasks. Tasks can be defined using an input-process-output model, so that each unique task is associated with a set of outputs defined on the task, and a set of inputs and processes that are necessary and sufficient to produce the outputs. For purposes of the current discussion, inputs and outputs can be represented by the concept of the data set, which consists of one or more variables (set names), each of which has a domain that is delineated by a set of constraints. The fundamental constraint on the elements of a data set or variable is its data type. Other constraints may be based on allowable values, ranges, etc., such that a data set contains a uniquely identifiable collection of elements.

Individuals are assigned roles based on their membership in the organization; these roles have various tasks associated with them. Each task has associated with it certain information and processing access requirements. For each role identified, the data sets and processes that support the tasks associated with that role must be identified. Permissions are then granted to the individual for that su-

perset of data which is composed of the union of all of the subsidiary (task associated) data sets. For individuals with multiple roles in an organization, this permission set includes data sets associated with tasks defined for each of the multiple roles.

Use of the term "role" as the identifying label for an individual's information processing activities and access needs allows natural expression of several other notions that are crucial to a responsive system of permissions. In practice, roles are not only statically assigned through the management hierarchy of the organization, but are also assumed and delegated dynamically, according to certain rules. The implications of this are considered in the next section.

Role Assumption and Delegation

In organizational settings that demand fast reactions to changing circumstances, e.g., medicine, a static approach to access permissions requires that individuals be granted access to the broadest scope of information they are expected to need in order to deal with the set of situations expected. In the terminology proposed above, individuals must be assigned to roles that have tasks and data sets defined broadly enough to give access to the maximum set of data expected to be needed in the most extreme case. This clearly violates the security principle of least privilege and need to know for most of the routine tasks associated with the work assignment, resulting in wider access to confidential data than is normally needed.

Instead of assigning static role definitions that must allow access to the maximum set of data anticipated, persons who must respond to varying situations will be assigned a base role, which includes permissions that allow the assumption of any of a set of more broadly defined or comprehensive roles. Thus, for normal situations, the user's access is constrained, but in certain well defined instances, the user may (on their own authority) assume one or more pre-defined roles that provide broader access permissions.

For example, in the event of a medical emergency, such as a cardiac arrest, the closest available physician can take charge of the patient and issue orders. In essence, the event in the patient's life effectively invokes a consultation request to the closest physician. Continuing this scenario, such an event also immediately triggers a consultation to the hospital service responsible for managing cardiac arrests, such as the cardiac care unit or the intensive care unit. When the resuscitation team arrives on site, they immediately have all privileges of the attending physician, including the right to review the chart, order medications, and transfer the patient. This sequence of events does not require the intervention of individuals who currently have authorized access to the patient, and is primarily a manifestation of the underlying notion that appropriate care for the patient supercedes all other models of authority over physician functions.

The "assumable role" approach extends the notion of discretionary control, expanding the user's discretion to include assumption of broader access rights. Clearly, role assumption must be restricted and managed so that only roles having been carefully specified are allowed to be assumed by a person who has been assigned a role that is authorized assumption. A system that allows role assumption must also implement a journaling or audit trail system so that role assumption behavior is monitored. This approach allows for implementation of the observed users approach [1] in a context where the importance of confidentiality of data is paramount, except when it conflicts with life-critical medical needs.

Closely related to the idea of role assumption is the notion of role delegation, which allows individuals who have been assigned a given role to delegate that role to another individual. This terminology implements the concept of discretionary control in our model in a very natural way, and provides for hierarchichal organization of access privileges where appropriate.

Integration of Multiple Roles

While some individuals play narrowly defined roles (with perhaps as few as one or two defined tasks), others will be expected to carry out several roles simultaneously. The most obvious example of this situation is that of the clinical faculty member who at once plays the role of attending physician, consulting physician, lecturer, re-

searcher, and perhaps chair/member of one or several committees. A responsive access control system must provide integrated access to data in support of these roles without artificial boundaries.

A major concern of the multilevel secure database literature [4] is "indirect access by inference," where authorized access to facts may result in the inference of data that should be classified at a higher, unauthorized level. In a medical setting, mechanisms that inhibit inference should be used with extreme caution, and should be the exception rather than the rule.

CONCLUSIONS

Advances in networking, operating system, and database technology make enterprise-wide information sharing possible through distribution of data storage, retrieval, and application processing. In the medical setting, protection of sensitive data and optimal use of information in support of patient care are potentially competing goals that pose great challenges for the design and implementation of access control mechanisms as a component of institutional information systems. In the absence of broad approaches in security research that allow for these unique requirements, it is necessary to develop such security models from within the field of medical informatics, with the requirements of medical computing clearly identified as the standard.

References

[1] Clark, D.D. and Wilson, D.R., A comparison of commercial and military computer security policy, *Proceedings 1987 Symposium on Security and Privacy*, IEEE Computing Society, 1987; 184-194.

[2] Date, C.J., *An Introduction to Database Systems, Volume 1*, Addison-Wesley, Reading, Massachusetts, 1990.

[3] Deitel, H.M., *An Introduction to Operating Systems*, Addison-Wesley, Reading, Massachusetts, 1990

[4] Denning, D.E., Commutative Filters for Reducing Inference Threats in Multilevel Database Systems, *Proceedings 1985 Symposium on Security and Privacy*, IEEE Computing Society, 1985.

[5] Denning D.E., Akl, S.G., Heckman, M., Lunt, T.F., Morgenstern M., Neumann P.G., and Schell, R.R., Views for Multilevel Database Security, *IEEE Transactions on Software Engineering*, 1987; SE-13(2): 129-140.

[6] Dennis, J.B. and Van Horn, E.C., Programming semantics for multiprogrammed computations, *Communications of ACM*, 1966; 9: 143-155.

[7] Department of Defense, Trusted Computer System Evaluation Criteria, DoD Document 5200.28-STD, Fort George Meade, MD, 1985.

[8] Digital Equipment Corporation, *Open Systems Handbook: A Guide to Building Open Systems*, Maynard, Mass. 1991.

[9] Fugini, M., and Martella, G., Conceptual Modeling of Authorization in Database Systems, *The Journal of Systems and Software*, 1987; 7: 3-13.

[10] Holbrook, J.P. and Reynolds, J.K., *Site Security Handbook*, Internet RFC 1244.

[11] Ingres Corporation, *INGRES/SQL Reference Manual, Release 6.4*, 1080 Marina Village Parkway, Alameda, CA, 1991.

[12] Leffler, S.J., McKusick, M.K., Karels, M.J., and Quarterman, J.S., *The Design and Implementation of the 4.3BSD UNIX Operating System*, Addison-Wesley, Reading, Massachusetts, 1989.

[13] Null, L.M. and Wong, J., A Unified Approach for Multilevel Database Security Based on Inference Engines, *SIGCSE Bulletin*, 1989; 21(1): 108-111.

[14] Snyder, L., Formal Models of Capability-Based Protection Systems, *IEEE Transactions on Computers*, 1981; C-30(3): 172-181.

[15] Tsai, W.T., Keefe, T.F., Thomsen, D.J., and Thuraisingham, M.B., AI Applications in Multilevel Database Security, *Computer Security Journal*, 1990; 6(1): 63-80.

[16] Ullman, J.D., *Principles of Database and Knowledge-Base Systems, Volume 1*, Computer Science Press, 1988.