Chapter XXXVII
Access Control Models

Romuald Thion
University of Lyon, France

ABSTRACT

Please submit an abstract for your chapter (a brief introduction of the chapter and/or explanation of what topics will be covered in the chapter). Your abstract should be between 150 and 200 words in length.

INTRODUCTION

Information knowledge has been acknowledged for a long time in warfare. For example, Tzu's section III. Attack by Stratagem (1910) describes the importance of knowledge:

If you know the enemy and know yourself, you need not fear the result of a hundred battles. If you know yourself but not the enemy, for every victory gained you will also suffer a defeat. If you know neither the enemy nor yourself, you will succumb in every battle. (p.)

This quotation points out that information knowledge is among the most important factors in winning a war, this quotation is a 2,500 year old introduction to information warfare. Information warfare means a strategy for acquiring an enemy's information, while defending one's own. It is a kind of warfare where information and attacks on information and its system are used as a tool of warfare.

Common mechanisms enhancing security and protecting one's own information are cryptography, *authentification*, or *authorization*. This topic focuses on a particular aspect of security mechanisms: authorization, also known as access control. This concept, in its

broadest sense, came about prior to computer science; chests, locks, fences, and guards have always been used to protect valuable information from foes.

Access control has been used since the very beginning of distributed systems in which multiple users can share common resources. With the increased dependence of defense on computer systems, the U.S. Department of Defense (DoD) investigated the vulnerability of government systems in the late 1960s, leading to the first definitions of access control principles. Researchers also considered the problem. For example, Lampson's (1974) access control matrix is the first formal mathematical description of what access control is. The DoD investigation led to a definition of multilevel access control, relating to classified documents, such as unclassified, confidential, secret, and top-secret, identifying clearly the separation between authorization and authentification. From then on, access control has been abundantly studied, extended, and commercialized to fill the security gap of computer systems, and is a major tool for preventing cyber terrorists from accessing sensitive data.

THE PURPOSE OF ACCESS CONTROL

In computer systems, access control denotes whether a *subject* (e.g., process, computer, human user, etc.) is able to perform an *operation* (e.g., read, write, execute, delete, search, etc.) on an *object* (e.g., a tuple in a database, a table, a file, a service, and, more generally, any resource of the system) according to a *policy*. These concepts are commonly encountered in most access control and computer security literature. The right to carry out an operation on an object is called permission. *Access control policies* define the subjects' permissions in a computer system, in order to enforce the security of an organization. One of the fundamental best practices in security is developing, deploying, reviewing, and enforcing security policies.. These policies are organized according to an access control model. The model may add intermediate concepts between subjects and permission to organize policies. Intermediate concepts are chosen among tasks, groups, roles, or confidentiality labels, for example. They aim at making policies, management, and definition easier, fitting in as best as possible with the internal structure and needs of the protected system (Ferraiolo, Kuhn, & Chandramouli, 2003).

Informally speaking access control means to decide "who can do what." Access control is arguably the most fundamental and most pervasive security mechanism in use in computer systems.

Information security risks are commonly categorized into:

- **Confidentiality:** Information must be kept private; only authorized users can read the information.
- **Integrity:** Information must be protected from being altered; only authorized users can write the information.
- **Availability:** Information must be available for use.

The purpose of access control is to preserve the confidentiality and integrity of information and, to a lesser extent, availability. Access control aims at providing only useful permissions to subjects, thus avoiding improper writing (mainly related to integrity) and reading (mainly related to confidentiality) operations. Access control is not as obviously related to availability, but it has an important role. A cyber terrorist who is granted unauthorized access is likely to bring the system down (Ferraiolo et al., 2003). Moreover, access control provides protection against internal attacks and information disclosure. With an authorization mechanism, a sleeping agent, who is member of an organization, a renegade, or a cyber spy, is not able to access the most valuable information. Information leakage is a major threat for private industry, whose intellectual property, business processes, and methodology are targeted by cyber terrorists.

ACCESS CONTROL MODELS

An access control model defines relationships among *permissions*, *operations*, *objects,* and *subjects.* We distinguish here the difference between *users,* the people who use the computer system, and *subjects,* computer processes acting on behalf of users. Several intermediate concepts have been introduced over the past decades to organize these relationships. This section surveys three widespread access control models: mandatory, discretionary, and role-based.

Lampson's Matrix and Discretionary Access Control

Access control terminology was established in the late 1960s by Lampson (1974), when he introduced the formal notions of subjects, objects, and access control *matrix.* An access control matrix is a simple representation in which each entry [i,j] of the matrix specifies the operations granted to subject *i* on resource *j.* An example from the medical field is shown in Table 1. For example, user *Charly* (more precisely processes invoked by user Charly) is allowed to write and read/access both *administrative* and *medical records* objects and read/access to *prescriptions.*

Such a matrix can be read either:

- **By rows:** Thus the matrix is interpreted as *capabilities list,* defining what is allowed for each user, for example, "David: **read** access on *medical* and *administrative* records";
- **By columns:** Thus the matrix is interpreted as a*ccess control list* (ACLs), defining which permissions are granted to each object, for example, "prescriptions: **read** access by *Alice* and *Charly.*"

Nowadays, an access control matrix tends is rarely used with the increasing number of resources and users; this model is not adequate for large organizations. The main goal of new models (e.g., role-based access control) is to overcome these limitations by proposing organizational grouping of subjects or resources.

Discretionary access control (DAC) (Department of Defense (DoD) National Computer Security Center, 1985) is one of the most widespread access control models. It can be seen as an access control matrix including an *ownership* relation, allowing subjects to settle policies for their own objects. This principle is implemented in the Unix/Linux operating systems to control access to files (e.g., a chown command that changes the owner of a file). This mechanism permits granting and revocations of permissions to the discretion of users, bypassing system administrator control. Even though DAC mechanisms are in widespread commercial use, they suffer from several difficulties among which are the following:

- Users can settle insecure rights, for example, the classical "chmod 777," which allows any permission to anybody in Unix/Linux system
- Transitive read access, for example, if Bob is allowed to read Charly's file, he can copy its

Table 1. A sample access control matrix

	Medical record	*Administrative record*	*Prescriptions*
Alice	**W,R**	**R**	**R**
Bob		**R**	
Charly	**W,R**	**W,R**	**R**
David	**R**	**R**	

content into a new file (of which Bob is the owner) and allow other users to read its content

Thus, safety has been shown to be undecidable (Harrison, Ruzzo, & Ullman, 1976) in access control matrix. It is impossible to prove whether an initial set of access rights that is considered safe would remain safe. The system may grant "unsafe" rights because the system has no control over permissions passed from one user to another. Thus, the use of this access control should be limited to noncritical structures. This model provides security, but as the risks of serious damages or leakage are really high once the system is compromised, it should not be used by potential targets of cyber terrorism (e.g., governmental organizations, large companies, and chemical, biological, or war industries).

Bell-LaPadula, Lattice-Based and Mandatory Access Control

In many organizations, end users do not "own" the information to which they are granted access. Information is the property of organizations, and no user should be able to settle its own permission. To overcome the difficulties of DAC in confidentiality critical environments, mandatory access control (MAC) has been developed (Bell & LaPadula, 1973). MAC was designed to deal with classified documents in computer systems (e.g., military ones). The basic principle of MAC is to control access according to the user's clearance and the object's classification. These classifications are divided into security levels (one can refer to MAC as a multilevel access control); the higher the level is, the more confidential the information is. For example, the common government classifications are *unclassified*, *confidential*, *secret,* and *top-secret*.

The principles shown in Table 2 have been formalized by Bell-LaPadula into a mathematical model suitable for defining and evaluating security in computer systems, making it possible to analyze their properties. We note that security levels are related to an organization's information flow; they represent the hierarchical structure of the organization. Users are able to write to a higher classification in order to transmit documents within their hierarchy.

Fully ordered levels are quite restrictive. The basic MAC principles of Table 2 can be applied to partially ordered levels by combining several classifications; these are called lattice-based access control models (a product of a lattice is a lattice). Figure 1 illustrates these combinations.

This access control model is arguably the most effective in maintaining confidentiality. However, it suffers from a rigidity that commercial companies cannot accept. Thus, this model should be used either by highly hierarchical organizations (e.g., banks or armies) or for critical parts of information systems of

Table 2. MAC control principles

• Only administrators, not data owners, make changes to an object's security label.
• All data is assigned a security level that reflects its relative sensitivity, confidentiality, and protection value.
• All users can read from classifications lower than the one they are granted.
• All users can write to a higher classification.
• All users are given read/write access to objects only of the same classification.
• Access is authorized or restricted to objects, depending on the labeling on the resource and the user's credentials.

Figure 1. A product of lattices

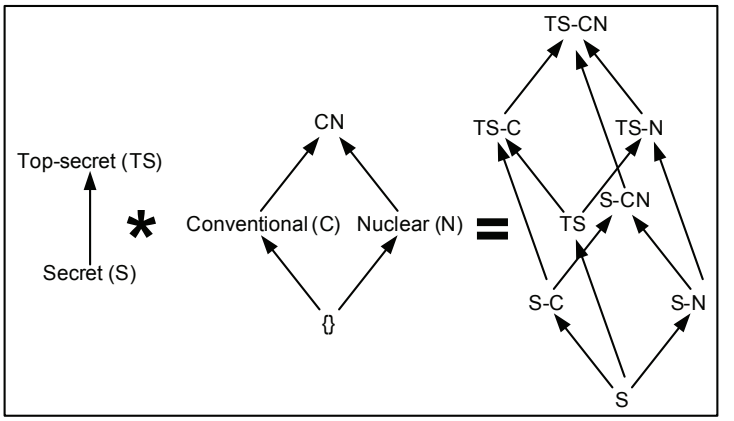

organizations in which several access control models cohabit. This model and the simple, but effective, classification of data it imposes has to be taken into account in any security planning, particularly for organizations threatened by cyber terrorism. Historically, this model originated from investigations on information warfare.

Role-Based Access Control

Role-based access control (RBAC) models constitute a family in which permissions are associated with roles (the intermediate concept of *roles* can be seen as collections of permissions), and users are made members of appropriate roles. Permissions are not directly assigned to users (Figure 2). The definition of role is quoted from Sandhu, Coyne, Feinstein, and Youman (1996): "A role is a job function or job title within the organization with some associated semantics regarding the authority and responsibility conferred on a member of the role" (p. X)

RBAC was developed to overcome administration difficulties encountered in large commercial organizations for which DAC was impracticable and MAC much was too restrictive. As the major part of access control decisions is based on the subjects' function or job, introducing roles greatly simplifies the management of the system. Since roles in an organization are relatively consistent with respect to user turnover and task reassignment, RBAC provides a powerful mechanism for reducing the complexity, cost, and potential for error in assigning permissions to users within the organization (Ferraiolo et al., 2003). RBAC was found to be among the most attractive solutions for providing access control in electronic commerce (e-commerce), electronic government (e-government), or electronic health (e-health) and is also a very active research field.

An important feature of the RBAC model is that roles are hierarchical; roles inherit permissions from their parents. Thus, roles are not flat collections of groups of permissions. Hierarchy aims at increasing system administrator productivity by simplifying distribution, review, and revocation of permissions. A sample role hierarchy is shown in Figure 3. In this example, Physician and Nurse inherit Employee, thus every permission assigned to the role Employee is also assigned to both Physician and Nurse roles. By transitivity, Cardiologist and Surgeon roles inherit all the permissions granted to both Physician and Employee.

RBAC constitute a family of four conceptual models, readers may encounter these specific acronyms:

Figure 2. RBAC model (without constraints)

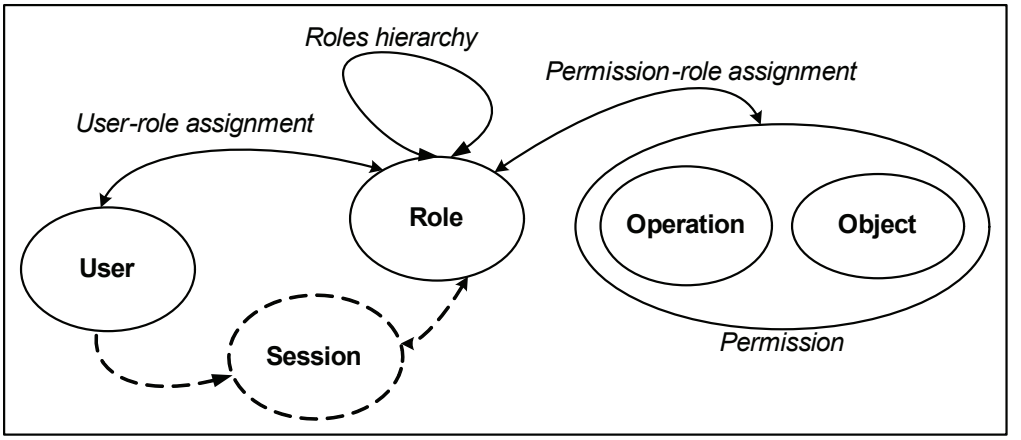

Figure 3. A sample role hierarchy

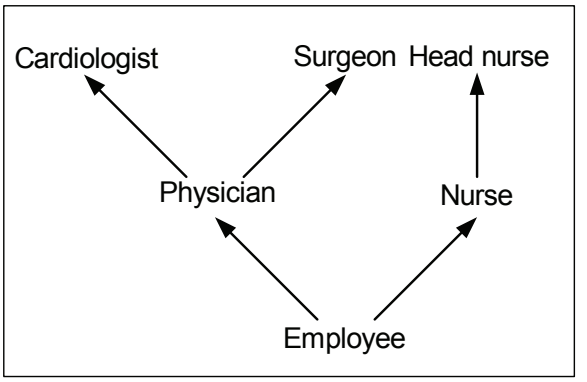

- $RBAC_0$ contains the core concepts of the model;
- $RBAC_1$ adds role hierarchy to $RBAC_{0;}$
- $RBAC_2$ adds static (not related to sessions) and dynamic (related to sessions) constraints between core concepts; and
- $RBAC_3$ includes all aspects of $RBAC_1$ and $RBAC_2$.

Nowadays, an international consensus has been established (National Institute of Standards and Technology (NIST), 2004) It describes the requirements and functionalities of RBAC implementations. RBAC's evolution from concept to commercial implementation (IBM Corporation, 2002) and deployment was quite rapid. For example, the U.S. Health Insurance Portability and Accountability Act of 1996 explicitly defines RBAC requirements. The introduction of RBAC in a large organization like Siemens, for example, was developed in Roeckle, Schimpf, and Weidinger (2000).

Clearly, this access control model is attractive for large organizations, where many different users' profiles are involved. It is arguably the most cost effective model. From a cyber terrorism perspective, it is interesting to point out that this model is considered

as "neutral policy." It can coexist with other policies. Thus, it can be thought of as the main access control model (for day-to-day operations) of an organization. A more restrictive one, a mandatory model, for example, should be reserved for sensitive services or information. Such architecture will protect against insiders (e.g., angry users), but also against external attackers (e.g., cyber spies, cyber terrorists) targeting valuable, sensitive, or critical information.

Other Access Control Models

DAC, MAC, and RBAC are among the most widely used access control models, but several others exist. This subsection surveys Biba's integrity model, the Chinese-Wall policy and Clark-Wilson model.

Biba's (1977) integrity model was introduced as an alternative to the Bell-LaPadula (1973) model to enforce integrity in military-oriented policies, focusing mainly on confidentiality. In Biba's model, security levels are integrity-oriented, for example, the levels are *critical*, *important,* and *ordinary*. The properties of the Biba model are similar to Bell-LaPadula's, except that read and write permissions are reversed. A subject is permitted read (respectively write) access to an object, if the object's (resp. subject's) security level dominates subject's (resp. object's) level.

Clark and Wilson (1987) have compared commercial security policies and military-oriented policies, pointing out their differences. They proposed two general security principles: *separation of duties* (SoD) and *well-formed transaction* to ensure information integrity. The Clark-Wilson model is commercially oriented; it ensures that information is modified only in authorized ways, by trusted people. Whereas military models can be defined in terms of low-level operations, such as read and write, Clark-Wilson's is application-level oriented. It defines a higher abstract notion of transaction.

Chinese-wall policies (Brewer & Nash, 1989) are for business transaction what Bell-LaPadula's policies are to the military. Brewer and Nash identified the notion of *conflict of interest* (COI). The objective of Chinese-wall policies is to avoid such conflicts. The basis of the policies is that subjects are only allowed access to information that does not conflict with any that they already possess (i.e., held on the computer and that has been previously accessed). Informally speaking, "users cannot go through the wall between conflicting classes of interest."

CURRENT ISSUES

Research into access control models aims at providing more expressive models that are able to take into account emerging trends on geographical, temporal, context-aware, and pervasive computer systems. Nowadays, nomadic computing devices and wireless communications force inclusion of geographical and context awareness in access control models. RBAC models have received particular attention, mainly because RBAC is now a de facto standard. For example, the geographical-RBAC model (Bertino, Catania, Damiani, & Perlasca, 2005) is a spatially aware access control model for location-based services and mobile applications. This research tends to be a major concern for the security of wireless information systems. These new proposals may protect against roaming attackers who are looking for wireless access points to target. A practice called war-driving. With the development of ubiquitous mobile computing, its introduction into cyber terrorism, and targeted fields, such as health (e.g., emergency units equipped with PDA and wireless communication devices) or oil companies (e.g., sensor infrastructures with query capabilities that are used by workers' laptops), dealing with geographical and temporal aspects is one of the major trends in access control.

Researchers also have focused on policy administration (Sandhu, Bhamidipati & Munawer, 1999; Ferraiolo, Chandramouli, Ahn & Gavrila, 2003) and common security description in XML (Organization for the Advancement of Structured Information Standards (OASIS), 2005). In fact, policies can be huge in international structures, and can involve thousands of users and hundreds of security administrators. For example, constraints have been introduced to

reflect specificities of organizations, such as mutually exclusive roles or prerequisites. Unfortunately, these constraints may obfuscate the meaning of policies and can lead to inconsistencies; this is especially true in distributed systems where policies from different suborganizations must cohabit. Recent research tries to fill this gap, proposing methods for distributed policies and tools facilitating design and maintenance of access control policies. This aspect of the research is of great importance because most security flaws are due to misconfigurations or administrative mistakes. An organization protecting itself against cyber terrorism must define a security policy and enforce it via access control mechanisms. However, it has to verify the implementation of the policies and be sure that no flaws have been introduced, from either inattention or malevolence.

CONCLUSION

Access control is a fundamental aspect of security and is paramount to protecting private and confidential information from cyber attackers. Understanding the basics of access control is fundamental to understanding how to manage information security. Several models have been developed over the decades to enhance confidentiality, integrity, availability, or administration flexibility. Being sometimes clearly military or commercially oriented, they share common criteria:

- Being built on formal mathematical models (matrix; lattice, entity-relation, etc.)
- Guaranteeing a set of properties (confidentiality of information, integrity of transactions, absence of conflicts of interest, etc.)

However, access control itself is not a panacea; it is a cornerstone of security, but is useless without rigorous security management or if built over insecure authentification mechanisms. It may be a lot easier for cyber criminals to endorse someone else's identity, than to gain unauthorized access inside a computer system that uses access control mechanisms.

The rising threat of cyber terrorism has been taken into account by researchers in access control. The authors Belokosztolszki and Eyers (2003) have highlighted several threats related to cyber terrorism against distributed access control policies. Such aspects of access control have to be investigated more deeply in order for us to protect ourselves against cyber terrorists.

REFERENCES

Bell, D. E., & LaPadula, L. J. (1973). *Secure computer systems: Mathematical foundations and model.* The Mitre Corporation.

Belokosztolszki, A., & Eyers, D. (2003). Shielding the OASIS RBAC infrastructure from cyber-terrorism. *Research Directions in Data and Applications Security,* 3-14.

Bertino, E., Catania, B., Damiani, M. L., & Perlasca, P. (2005). GEO-RBAC: A spatially aware RBAC. *10th Symposium on Access Control Models and Technologies,* 29-37.

Biba, K. J. (1977). *Integrity considerations for secure computer systems.* The Mitre Corporation.

Brewer, D., & Nash, M. (1989). The Chinese wall security policy. *Proceedings of the IEEE Computer Society Symposium on Research in Security and Privacy,* 215-228.

Clark, D. D., & Wilson, D. R. (1987). A comparison of commercial and military computer security policies. *IEEE Symposium of Security and Privacy,* 184-194.

Department of Defense (DoD) National Computer Security Center. (1985). *Department of Defense trusted computer systems evaluation criteria* (DoD 5200.28-STD).

Ferraiolo, D. F., Chandramouli, R., Ahn, G. J., & Gavrila, S. I. (2003). The role control center: Features

and case studies. *8th Symposium on Access Control Models and Technologies,* 12-20.

Ferraiolo, D. F., Kuhn, R., & Chandramouli, R. (2003). *Role-based access controls*. Artech House.

Harrison, M., Ruzzo, W., & Ullman, J. (1976). Protection in operating systems. *Communication of the ACM, 19*(8), 461-471.

IBM Corporation. (2002). *Enterprise security architecture using IBM Tivoli security solutions.*

Lampson, B. (1974). Protection. *ACM Operating System Reviews, 8*(1), 18-24.

National Institute of Standards and Technology (NIST). (2004) *Role-based access control* (NIST Standard 359-2004).Retrieved from http://csrc.nist.gov/rbac

Organization for the Advancement of Structured Information Standards (OASIS). (2005). *eXtensible access control markup language* (XACML 2.0).

Roeckle, H., Schimpf, G., & Weidinger, R. (2000). Process-oriented approach for role-finding to implement role-based security administration in a large industrial organization. *Proceedings of the Fifth ACM Workshop on Role-Based Access Control,* 103-110.

Sandhu, R. S., Bhamidipati, V., & Munawer, Q. (1999). The ARBAC97 model for role-based administration of roles. *ACM Transactions on Information System Security, 2*(1), 105-135.

Sandhu, R. S., Coyne, E. J., Feinstein, H. L., & Youman C. E. (1996). Role-based access control models. *Computer, IEEE Computer Society Press, 29*(2), 38-47.

Tzu, S. (1910). *Sun Tzu on the art of war, the oldest military treatise in the world* (L. Giles, Trans.). Retrieved from http://classics.mit.edu/Tzu/artwar.html

U.S. Health Insurance Portability and Accountability Act (HIPAA). (1996). Retrieved from http://cms.hhs.gov/hipaa

TERMS AND DEFINITIONS

Access Control (or Authorization): The process of determining whether a subject (e.g., process, computer) is able to perform an operation (e.g., read, write) on an object (e.g., a file, a resource in the system).

Access Control Model: This is the underlying model upon which security policies are built. The access control model defines concepts and relations between them to organize access control.

Access Control Policy: This is the set of rules built on an access control model that defines the subjects, objects, permissions, and other concepts within the computer system. Authorization decisions are based upon access control policies settled in the system.

Discretionary Access Control (DAC): This is an access control model in which it is the owner of the object that controls other users' access to the object.

Mandatory Access Control (MAC): This refers to an access control model in which decisions must not be decided upon by the object owner. The system itself must enforce the protection decisions (i.e., the security policy).

Role-Based Access Control (RBAC): This is an access control model in which access decisions are based on the roles that individual users have as part of an organization.