18

Panel Discussion: Role-based Access Control and Next-Generation Security Models

Roshan Thomas (panel chair)
Odyssey Research Associates
301 Dates Drive, Ithaca, NY 14850., rthomas@oracorp.com

Elisa Bertino and Pierangela Samarati, Università di Milano
Hans H. Brüggemann, University of Essen,
Bret Hartman, Odyssey Research Associates
Ravi Sandhu, George Mason University
T. C. Ting, University of Connecticut

Abstract

This purpose in organizing this panel was to promote discussion and to bring to the forefront the many issues related to next generation security models. Each of the sections below discuss the individual contributions of the various panelists.

Keywords

Security models, role-based access control, authorization

1 INTRODUCTION

Roshan Thomas
Odyssey Research Associates
301 Dates Drive, Ithaca, NY 14850
rthomas@oracorp.com

It is now widely recognized that security requirements of systems can be viewed at different levels (stages) of abstraction. The security requirements at the higher stages can be refined and elaborated at lower stages. Given these stages of elaboration, it is possible to formulate security models for each of these stages, as well as classify existing models as to where they belong. In fact, it is possible to derive a related taxonomy of security models for the various stages. At the

D. L. Spooner et al. (eds.), *Database Security IX*

highest level we have models to capture organizational policy and requirements that pertain to security. These requirements are then applied to the interface between the organization and the computer system and captured by computer policy models. Computer policy models in turn are implemented by access control models, which in turn map to implementation models, and so on.

To date most research and development in security models have been primarily aimed at specifying and implementing internal requirements and related rules of operation within computers. Consequently, there is a mature body of literature on access control and implementation models. The Bell-LaPadula model for multilevel security and the typed access matrix model (also called TAM) all fall into this category. However, as observed by Dobson, research into security models for higher stages of elaboration are still in its infancy.

Future research needs to address frameworks, tools, and languages for next generation higher level security models. The emphasis in these models will be on policy-oriented as opposed to mechanism-oriented abstractions and facilities. These models will approach security requirements from an enterprise perspective and be able to capture enterprise-specific policies and security requirements.

Recent research efforts into role-based access control (RBAC) models should form a good starting point for discussion of higher level models. However, in addition to roles, we may require abstractions to express sequences of access control decisions and various dependencies between them, failures and exceptions, and internal controls in organizations such as those based on separation of duties. The model of transaction control expressions (TCE's) proposed by Sandhu and the model of task-based authorizations proposed by Thomas and Sandhu are an attempt to address these issues.

To summarize, we list some of the many open research questions.

1. What are the frameworks, models, and abstractions necessary to approach next generation access control and security models?

2. How can the independent but related security objectives, namely, confidentiality, integrity, availability, and accountability, be uniformly addressed in next generation models?

3. How can security models capture more application logic and be more enterprise and policy oriented?

4. How do we specify exceptions and failures?

5. How can we (and do we need to?) uniformly address multilevel and commercial security requirements?

6. What features of access control models would make them suitable for database systems and object-oriented computing environments?

2 NEXT-GENERATION AUTHORIZATION MODELS FOR DATABASE SYSTEMS

Elisa Bertino and Pierangela Samarati
Dipartimento di Scienze dell'Informazione
Università di Milano
Via Comelico, 39/41, 20135 Milano, Italy
{bertino, samarati}@dsi.unimi.it

Traditional authorization models, proposed for the protection of information in operating systems and database management systems, can be classified in two categories: mandatory and discretionary, according to the type of policy they apply. Mandatory security policies govern the access to the information by the individuals on the basis of the classifications of *subjects* and *objects* in the system. Discretionary protection policies govern the access of users to the information on the basis of the users's identity and rules specifying, for each user and each object in the system, the accesses the user is allowed on the object.

Although both discretionary and mandatory models have been successfully used, modern applications call for richer and more flexible access control models. The need for these "next-generation" authorization models is primarily driven by two factors: development of new data models and the need for representing new protection requirements.

Development of new data models concerns the fact that data management systems are today moving towards richer and more complex modeling paradigms. This shift results in more powerful data models from the point of view of data representation and retrieval. The richer semantics of these new data models makes traditional access control models inadequate for their protection. An example of these new data paradigms is the object-oriented model, which is today one of the most active areas in both academic and industrial worlds. Several object oriented database systems have been developed and a number of relational DBMSs (RDBMSs) are currently being extended with object-oriented capabilities. The rich semantics of the object-oriented data models, providing concepts such as inheritance, versions, and composite objects, introduces new protection requirements which the traditional authorization models do not address. Several authorization models for the protection of object-oriented systems have been proposed and are still under study. However, no model proposed so far addresses the problem satisfactorily. Other paradigms such as deductive and active data models are today receiving increasing attention.

Traditional access control models need to be adapted and extended to take into consideration the data management characteristics proper of such models. Another data modeling paradigm which is receiving increasing attention is the hypertext/hypermedia paradigm. The key point of such an approach is that every piece of information is connected, via links to related pieces of information. Data are often unstructured (consider, for instance, the case of image, voice, and text data) or have very irregular structures (consider, for instance, the case of WWW pages). Traditional authorization models do not "fit" in this flexible framework. Applying traditional authorization models

to these new data models may result in two drawbacks. The first is that violations of the security policy are possible. If the different relationships existing between the data are not taken into consideration, the user may be able to retrieve information for which he is denied direct access, or, vice versa, be unable to access information for which he is authorized. The second is that the application of the access control, which is not able to capture the richness of the relationships among data, may result in imposing rigidity on the use of the system. We must note that, if on the one hand new data models rise new protection requirements, on the other hand, they are more flexible with respect to authorization specification. As a matter of fact, by taking into consideration the semantics of the target data model, more flexible and richer protection requirements can also be specified. As an example, consider the hypertext data model. Authorizations can be specified that allow users to access a node containing information depending on the navigation path along the hypertext followed by the user in arriving at the node.

A second factor pushing for new security models is the need to support a larger variety of authorization policies. Traditional identity-based (discretionary) or label-based (mandatory) models do not allow to express many practical requirements. In this respect, exception handling, temporal authorizations, explicit denials, task-based authorization, roles are all needed facilities. Mandatory policies rise from rigid requirements, like those of the military environment. Discretionary policies rise from cooperation yet autonomy requirements, like those of the academic environment. Neither policy satisfies the need of most of the application environments, such as commercial enterprises.

An alternative protection paradigm currently under study by a number of researchers is represented by the so called role-based policies. Role-based policies regulate the accesses on the basis of the activities the users execute in the system. These policies have several advantages. First, they permit enforcement of the least privilege: a given role can be granted only the authorizations necessary to fulfill the accesses connected with the role's responsibility. Second, it permits enforcement of the separation of duties principle: roles allowed (or simultaneously allowed) to each user may be restricted so that no user will be able to misuse the system. Third, since authorizations to access data are specified for roles instead than for users, the authorization management task is simplified (for example, if a user changes job it is sufficient to assign him a new role without changing all the authorizations to access the data). Morevover, the consideration of roles allows for more powerful administrative policies, for instance certain roles can be given the privilege to administer the authorizations of their subordinate roles. More complex protection requirements can be represented by taking time into consideration. In traditional access control models, an access is either allowed or denied and authorizations are valid from the time they are granted until the time they are revoked. In many real-life situations, access pernissions, or denials, may be limited in time, be granted only temporarily, or depend on other factors, such as temporal relationships between authorizations and between accesses. It is important that the authorization model be able to mirror those situations as well.

The Database Systems Group of the University of Milan, in collaboration with the Center for Secure Information Systems of George Mason University, is actively involved in addressing some of the above issues. In particular, discretionary authorization models and high-assurance access control models have been widely investigated for object-oriented database systems and these

attempts have tried to provide formal definitions and semantics. Research is, however, continuing on issues related to performance, standards (such as CORBA), and formal verification of security properties of object-oriented database schemas. Discretionary authorization models have been also investigated for relational database systems. The goal here is to extend the expressive power of conventional authorization models to provide increased flexibility. Research has so far covered timestamped and non-timestamped authorization models, recursive and non-recursive authorization revocation, negative authorizations and exception handling, temporal authorization models. Research issues currently being investigated include customizable authorization models (able to support a variety of authorization policies), periodic and history-based authorization models, and performance.

An important area which we are also addressing is the development of authorization administration tools. When dealing with authorization models providing a large variety of options (like the above-mentioned models), it is crucial that the authorization administrator be supported by tools able to visualize the authorization base state and to anticipate the consequences of certain choices. A tool has already being developed, for the case of an authorization model with negative authorizations and both recursive and non-recursive revocation. Finally, new applications, whose security requirements have been only partially identified, are being investigated, namely mobile computing, workflow management, digital libraries, computer-supported cooperative work (CSCW) systems.

3 OBJECT-ORIENTED RIGHTS FOR REDUCING COMPLEXITY AND MAKING DESIGN DECISIONS PERSISTENT

Hans H. Brüggemann
Department of Mathematics and Computer Science,
University of Essen, D-45117 Essen, Germany
jimmy@informatik.uni-hildesheim.de

1. My panel contribution primarily focuses on the mechanisms of right administration and access control. For reducing complexity, I suggest to use the basic object-oriented concepts, namely classes and class hierarchies, for modelling the basic units of an action independently: subjects, operations (access types), and granules (i.e. objects to protect).

Thus rights containing class names implicitly give rights to its class members, and moreover the class hierarchy determines how rights are implicitly given to other classes.

Then user groups can simply be modelled by subject classes. Roles can now be expressed by a smaller package of rights, more exactly by a smaller package of capabilities, i.e. (operation, granule)-pairs. The rights system can make use of the application hierarchies (for subject, operation,

and granule classes) and no additional role hierarchy is needed.

2. For making design decisions persistent, there should be a clear distinction between the intended design decisions (i.e., intended permissions and intended prohibitions) and the possibilities for later right updates. Thus rights should be explicitly marked as permissions or as prohibitions. If prohibitions are not explicit, e.g. the later insertion of a permission might destroy an intended, but implicit prohibition.

Moreover priorities for rights should be explicit. Implicit priorities have undesirable drawbacks: If prohibitions dominate permissions, it is not possible to specify an exception which is a permission. If the more specialized right dominates the more general right only one specialization hierarchy can be considered, but usually we have at least two hierarchies (for subjects and granules) and thus incomparable items. If the newer right dominates the older right, the update semantics become very obscure.

3. For further reducing complexity, exceptions are helpful. The general case can be easily expressed using class names (and class hierarchies). The exception can be expressed by using more specialized classes or objects (together with a higher priority). Exceptions can be on several levels, thus the same right can be an exception in a first context and a general case in a second context. Five or four levels of exceptions are not uncommon in everyday examples like rules for take over or right of way.

The impact of priorities can be limited by priority scopes. This is in particular useful for distributed right administration or for further right modularization.

A goal is a common framework for the specification of rights (for confidentiality), duties (for integrity), and abilities (for availability). Duties and its opposite, so-called freedoms, or abilities and inabilities, can be syntactically specified using the same mechanisms as described above. The difficulty is in the semantic interaction of these concepts.

4 SECURITY MODELS AND OBJECT TECHNOLOGY

Bret Hartman,
Odyssey Research Associates
301 Dates Drive, Ithaca, NY 14850
bret@oracorp.com

Object technology (OT) is the solution of choice for fully distributed applications that go well beyond the current client/server paradigm. OT supports flexible interoperability of components that are distributed throughout large-scale heterogeneous networks. The two industrial approaches to

OT are the Component Object Model (COM), developed by Microsoft, and the Common Object Request Broker Architecture (CORBA), developed by the Object Management Group (OMG).

It is widely recognized that the successful deployment of OT hinges on effective security. The commercial marketplace imposes difficult requirements on any proposed security solutions: they must be scalable for very large applications; they must be customizable to specific vertical markets, and they must work across a large set of heterogeneous configurations of software, operating systems, hardware, and networks.

We still have several technical barriers that must be overcome if we are to satisfy these market requirements. No single security solution will work for all markets -- as we have seen from past experience, a single low-level security model enforced solely by the operating system or hardware is inadequate. The multitude of security mechanisms required at each layer to support a more complex model presents a different problem: How can the mechanisms be integrated in a distributed system to ensure that security is tamper-proof and always invoked? Finally, the lack of assurance guidelines for distributed object-based systems means that identifying effective security products will be difficult.

A new generation of security models is needed for distributed object systems. A candidate model would provide a means for describing application-specific security policies and would provide the basis for lower level enforcement mechanisms. RBAC appears to be an excellent candidate that can address the technical barriers listed above. RBAC can be tailored to cover a wide variety of enterprise policies and effectively combines both non-disclosure and integrity constraints. RBAC also naturally meshes with an object model since the fundamental unit of access control, namely the operation, corresponds to a method in OT terms.

Although there is general agreement over the basics of RBAC, there are several choices for the granularity of access control. Typically, an RBAC policy consists of a set of roles, a set of operations on objects, and an access control policy that constrains how subjects acting in a particular role may execute operations on objects. There are three classes of RBAC policies: coarse-grained, medium-grained, and fine-grained.

Coarse-grained RBAC controls access at the granularity of an application rather than an operation. Domain-type enforcement supported by some existing systems is an example of this class of RBAC.

Medium-grained RBAC controls access at the granularity of the operation. Medium-grained RBAC policy is a function of the subject, role, and object.

Fine-grained RBAC further constrains access as a function of the state or arguments to the operation. Examples include:

- permitting a nurse to prescribe a drug only up to a certain maximum dosage -- for greater

doses, a doctor is required;

- preventing a customer from making more than 3 withdrawals/day or more than $1000/day, whichever comes first; - denying administrator access to an audit log if the administrator is simultaneously using email.

These choices for granularity of RBAC all need to be supported to satisfy marketplace requirements for secure OT.

5 ROLE-BASED ACCESS CONTROL AND NEXT-GENERATION ACCESS CONTROL MODELS

Ravi Sandhu
ISSE Department, MS 4A4
George Mason University
Fairfax, VA 22033
sandhu@isse.gmu.edu

In this brief position paper I would like to address several points that were raised during the panel. I began my presentation by observing that traditional access control models are based on the concept of subject and have eliminated the user as an explicit part of the model. Often times a

subject is misdefined as a user. In my opinion the next generation of access control models must deal with the user-subject relationship explicitly and not dismiss it as trivial. The Bell-LaPadula model actually does distinguish between user and subject, since a user can log on at various security levels dominated by the user's clearance. Each session established by the user is a separate subject. If the subject-user distinction is not made we reach strange conclusions where high users are unable to send low email, and Chinese Walls are deemed un implementable in the lattice model. These misinterpretations actually have been made and are not merely hypothetical.

In my mind role-based access control (RBAC) brings out the user-subject distinction in a particularly useful way. Along with colleagues at SETA Corporation, I have developed a framework of models for RBAC that brings out this distinction clearly. A paper describing this framework will appear in a forthcoming issue of IEEE Computer.

RBAC and lattice-based access controls both recognize that a user must log on with different privileges on different occasions. When a top-secret user logs on as an unclassified subject that user can be viewed as assuming an unclassified role in the system. This suggests that lattice-based access controls are a special case of RBAC. This can be formally demonstrated. One approach is given in the paper on Modeling Mandatory Access Control in Role-Based Security Systems by

Nyanchama and Osborn in this proceedings. Other approaches are also possible. The relationship of RBAC to lattice-based access controls needs to be clearly understood. It should not be taken to mean that whenever RBAC is implemented we first need to implement Bell-LaPadula. Rather the fact is that RBAC can be configured to give the same effect as Bell-LaPadula, but it does not have to be configured this way.

During the discussion there were objections that RBAC does not solve all access control problems. I have never claimed that it does. Instead I have always stated that RBAC is not a panacea but is a good approach for a large number of situations. I would be most suspicious of any form of access control that claims to be a panacea.

To summarize my main points are that the user-subject distinction is of fundamental importance to access control and the user-subject association needs to be explicitly modeled, that RBAC is one approach which exploits this distinction in a useful way, that RBAC includes traditional lattice-based controls as a special case, and that role-based access control (or probably anything else for that matter)is not a panacea for all access control problems.

6 ROLE-BASED SECURITY MODELS

T. C. Ting
Department of Computer Science
University of Connecticut
Storrs, CT U.S.A.
tting@nsf.gov

This is a position paper which was presented in a panel at the 1995 IFIP 11.3 Working Conference on Database Security. It summarizes the key points of the presentation and discussion.

Role-Based Security Model is only one of the potential solutions which is particularly suited for DAC. However, both MAC and DAC are important security considerations. No single security solution will work for all concerns. Role-based security model is consistent with object-oriented technology, and therefore, it is particularly significant in object-oriented security-critical information systems. The successful development of object-oriented information systems hinges heavily on effective security approaches.

The importance of a data object varies from applications to applications. Role-based models provide an effective means for providing application dependent data security by tailoring the application's tasks and security requirements. The data access functions and security requirements are expressed and enforced via the user roles. Roles are used to facilitate well-formed data access transactions and to enforce data access control functions based on the least privilege and the need-to-know principles. One of the major security objective is to protect the organizations and the

individuals whose data is stored in the database. The user's data access functions in a given application system are clearly coded in user roles which specify the data access and security actions. Each user within an application is assigned a set of roles which reflect the user's responsibility within the application.

A user must be assigned to a set of consistent roles. The assignment of the roles to users is the responsibility of the application's manager which is responsible for the performance as well as the data protection of the application. The responsible manager must be able to visualize the functions and security requirements of each and every role before assigning them to the individual users. Tools must be available to them for visualizing composite functions and security conflicts when multiple roles are assigned to an individual. Each application database user in an organization has a profile of roles which specifies the user's functional responsibilities and security restrictions. Individual's title suggests a general interpretation about the responsibility and authority. However, the composition of roles is one's actual responsibility and authority. Individuals can be provided with special tasks which may or may not be suggested by his or her title. When develop role-based security models one should focus on the detailed functional and security specifications of each role and don't mix titles with roles.