Proc. First Australasian Conference on Information Security and Privacy, Wollongong, Australia, June 23-26, 1996

Access Control: The Neglected Frontier

(Invited Paper)

Ravi Sandhu

Abstract. Access control is an indispensable security technology. However, it has been relatively neglected by the research community. Over the past ten years, the doctrine of mandatory and discretionary access controls has slowly become discredited but no dominant doctrine has emerged to replace it. There are promising candidates such as role and task-based access controls but these are still in their formative stages and have not gained wide acceptance. This paper gives my personal perspective on these issues and identifies some of the important access control issues that researchers and practitioners should focus on.

1 Introduction

Information and system security is a multi-faceted discipline. Security presents diverse and conflicting objectives. Availability, confidentiality, integrity and privacy have been explicitly recognized in the security literature for some time. Other objectives such as intellectual property protection, copyright, secure electronic transactions, metering systems and information and resource usage for specific purposes are emerging to the forefront.

Security objectives will continue to be refined, expanded and elaborated over the next decade. Security will have a very different meaning in ten years than it does today. Much of classical security thinking has been oriented towards single organizations. In the future we will increasingly have to consider systems encompassing multiple organizations with different and, often conflicting, security objectives. Even within a single organization, as true enterprise-wide computing emerges we will see such conflicts. All this presents interesting policy and technical challenges.

The principal security technologies today are cryptography, access control, authentication, intrusion detection and recovery, risk analysis and assurance. No single technology can solve real security problems under realistic assumptions. Each addresses a piece of the problem. Fortunately security technologies are mutually supportive and there is no fundamental conflict or incompatibility between them. As technologists we must strive to use the most appropriate mix of technologies to achieve overall security objectives.

These security technologies are all important and all deserving of interest from researchers and funding agencies. However, from my perspective it does appear that access control has been relatively neglected in the last decade compared to other technologies, particularly cryptography. There are many reasons for this. It is not my goal here to analyze these in detail. Rather, I would like to draw the attention of researchers and systems developers to access control as an area with tremendous potential for achieving significant results. It is a frontier that has not yet been heavily mined and offers high payoff in terms of achieving practical security.

This paper gives my personal perspective on the neglected frontier of access control. It begins by reviewing classic access control doctrine which is based on the twin pillars of discretionary and mandatory access control. This is followed by a discussion of what is wrong with this doctrine and what alternatives are being pursued.

2 Discretionary Access Control (DAC)

Discretionary access control (DAC) has its genesis in the academic and research setting from which time-sharing systems emerged in the early 1970's. A classic paper by Lampson [Lam71] introduced the basic ideas. DAC is based on the notion that individual users are "owners" of objects and therefore have complete discretion over who should be authorized to access the object and in which mode (e.g., read or write). Ownership is usually acquired as a consequence of creating the object.

In DAC-think if Alice owns an object it is at her pleasure, whim or fancy that she decides to grant Bob access to it. Later on, should she change her fancy she can revoke Bob's access. There are many subtle issues in DAC. A question that arose almost immediately was whether or not Bob can further grant access to Charlie, so the notion of a "grant option" or "copy flag" was invented [GD72]. In turn this led to problems of cascading revoke [GW76, Fag78]. Furthermore if Alice can grant access to a group of users but at the same time withhold access from Bob even if Bob is a member of that group additional subtleties arise [GSF91, Lun88, RBKW91].

All these subtleties of DAC are still being discussed, debated and refined in the literature. Nevertheless the driving principle of DAC is ownership, so much so we should perhaps be calling it owner-based DAC.

I will leave the readers with a DAC conundrum. Suppose Alice grants a permission X to Bob with the grant option. Bob then grants X to Charlie, followed by a grant X from Alice to Charlie. Now Alice revokes X from Charlie. Should Alice's revoke override Bob's grant or should Bob's grant override Alice's revoke? The exercise is to check what System R [GW76, Fag78] would have done in this situation, and to argue the other alternative is equally, if not more, reasonable.

DAC has an inherent weakness that information can be copied from one object to another, so access to a copy is possible even if the owner of the original does not provide access to the original [SS94]. Moreover, such copies can be propagated by Trojan Horse software without explicit cooperation of users who are allowed access to the original.

3 Mandatory Access Control (MAC)

Mandatory access control (DAC) was invented to enforce lattice-based confidentiality policies [BL75, Den76] in face of Trojan Horse attacks. Subsequently it was shown how to apply MAC for integrity and aggregation objectives (such as Chinese Walls) [Bib77, Lip82, San93]. MAC ensures that even in the presence of Trojan Horses information can only flow in one direction in a lattice of security labels (from low confidentiality to high confidentiality, or equivalently from high integrity to low integrity).

MAC enforces one-directional information flow in a lattice assuming there are no covert channels by which information can flow in prohibited ways. Covert channels are expensive to eliminate even if they could all be identified and analyzed. In the late 1980's it became apparent that many low-level hardware performance improvement technologies, such as cache memory, result in very high speed covert channels [KZB+90]. Information can be leaked through these channels at disk and LAN speeds. Moreover the faster the hardware the faster these covert channels get. The covert channel problem remains a major bottleneck for high assurance MAC. MAC also does not solve the inference problem where high information is deduced by assembling and intelligently combining low information.

4 Beyond MAC and DAC

There has been a persistent criticism of the MAC-DAC doctrine over the past decade. The criticism has not been universally accepted but is has been steady and has come from a number of authors. We can roughly divide the critics into two classes as follows.

4.1 Real MAC is more than classical lattice-based MAC

Traditional lattice-based is a very narrow interpretation of the term "mandatory." Lattice-based MAC cannot enforce integrity policies. A more general notion of MAC is needed for integrity. Various authors have suggested trusted pipelines and type enforcement [BK85], well-formed transactions and constrained data items [CW87] and controls based on static and dynamic properties [San90]. Some of these arguments can be extended to confidentiality applications.

In my view inadequacy of lattice-based MAC stems from its reductionist approach of controlling access in terms of read and write operations. Operations such as credit and debit both require read and write access to the account balance and therefore cannot be distinguished for access control purposes in lattice-based MAC.

Several authors have argued that by appropriate construction of lattices it is possible for lattice-based MAC to accommodate policies that do not appear at first sight to be compatible with MAC. For instance, it had been argued that the Chinese Wall policy cannot be implemented using lattice-based MAC [BN89] but it was subsequently shown how to do this [San93]. Attempts to implement the Clark-Wilson integrity model using lattices were described by [Lee88]. Foley shows how various exceptions to information flows in a lattice can be accommodated by modifying the lattice [Fol92]. The question of how far lattice-based MAC can be pushed to support the information flow component of security policies is still not fully resolved.

4.2 Real DAC is more than classical owner-based DAC

Traditional owner-based DAC is but one form of DAC. A general model for propagation of access rights, commonly called HRU, was proposed by Harrison, Russo and Ullman [HRU76]. Unfortunately this model has very weak safety properties so it is difficult to determine the precise consequences of a propagation policy. Pittelli [Pit87] established the connection between lattice-based MAC and HRU by showing how the former can be simulated in the latter.

A variety of models and policies for propagation of rights were developed [LS77, San88a, MS88]. The SPM model of [San88a] was based on the premise that reducing expressive power may facilitate safety analysis. Of course, if we reduce expressive power too much the resulting model will not be very useful. The take-grant model has efficient safety analysis but very limited expressive power [LS77]. It turns out that SPM has strong safety analysis and has considerable expressive power [San92], including the ability to simulate lattice-based MAC. With a slight extension [AS92] it is formally equivalent to monotonic HRU. For monotonic systems (i.e., systems in which only those permissions that are restorable can be revoked) we can have general safety analysis and expressive power simultaneously.

The typed access matrix (TAM) model introduces types into HRU. Augmented TAM (ATAM) further adds the ability to test for absence of rights [SG93]. A systematic analysis of the relative expressive power of different variations of TAM and ATAM was recently completed [Gan96]. These results indicate that even very simple variations of ATAM retain ATAM's full expressive power. From an implementation viewpoint this is an encouraging result. Simple models should have simple implementations and these can provide complete expressive power. From a safety perspective the results are disappointing because the premise of the successful SPM work does not extend to non-monotonic systems.

These results need to be interpreted carefully. General algorithms for safety analysis of non-monotonic systems are rather unlikely to exist. However, it is still possible to effectively analyze and develop safety results for individual systems.

The analogy is to program verification where general verification algorithms do not exist, but individual programs can be verified. The problem with verification technology is the limited size of program that can be verified. In access control systems the policies should not require millions of lines of ATAM specification, but could perhaps be done in hundreds or thousands of lines. So case-by-case safety (and even liveness) analysis of access control policies might be practical.

5 Role Based Access Control

Role-based access control (RBAC) has recently received considerable attention as a promising alternative to traditional discretionary and mandatory access controls (see, for example, [FK92, SCY96, SCFY96]). In RBAC permissions are associated with roles, and users are made members of appropriate roles thereby acquiring the roles' permissions. This greatly simplifies management of permissions. Roles are created for the various job functions in an organization and users are assigned roles based on their responsibilities and qualifications. Users can be easily reassigned from one role to another. Roles can be granted new permissions as new applications and systems are incorporated, and permissions can be revoked from roles as needed.

An important characteristic of RBAC is that by itself it is policy neutral. RBAC is a means for articulating policy rather than embodying a particular security policy (such as one-directional information flow in a lattice). The policy enforced in a particular system is the net result of the precise configuration and interactions of various RBAC components as directed by the system owner. Moreover, the access control policy can evolve incrementally over the system life cycle, and in large systems it is almost certain to do so. The ability to modify policy to meet the changing needs of an organization is an important benefit of RBAC.

There is similarity between the concept of a security label and a role. In particular, the same user cleared to say Secret can on different occasions login to a system at Secret and Unclassified levels. In a sense the user determines what role (Secret or Unclassified) should be activated in a particular session. In [San96] it is shown how traditional lattice-based MAC can be simulated using RBAC96 model of [SCFY96]. This establishes that traditional MAC is just one instance of RBAC thereby relating two distinct access control models that have been developed with different motivations. It is also practically significant, because it implies that the same Trusted Computing Base can be configured to enforce RBAC in general and MAC in particular. This addresses the long held desire of multi-level security practitioners that technology which meets needs of the larger commercial marketplace be applicable to MAC. The classical approach to fulfilling this desire has been to argue that MAC has applications in the commercial sector. So far this argument has not been terribly productive. RBAC, on the other hand, is specifically motivated by needs of the commercial sector. Its customization to MAC might be a more productive approach to dual-use technology.

In large systems the number of roles can be in the hundreds or thousands. Managing these roles and their interrelationships is a formidable task that often is highly centralized and delegated to a small team of security administrators. Because the main advantage of RBAC is to facilitate administration of permissions, it is natural to ask how RBAC itself can be used to manage RBAC. We believe the use of RBAC for managing RBAC will be an important factor in the long-term success of RBAC. Decentralizing the details of RBAC administration without loosing central control over broad policy is a challenging goal for system designers and architects.

Since RBAC has many components, a comprehensive administrative model would be quite complex and difficult to develop in a single step. Fortunately administration of RBAC can be partitioned into several areas for which administrative models can be separately and independently developed to be later integrated. In particular we can separate the issues of assigning users to roles, assigning permissions to roles and defining the role hierarchy. In many cases, these activities would be best done by different administrators. Assigning permissions to roles is typically the province of application administrators. Thus a banking application can be implemented so credit and debit operations are assigned to a teller role, whereas approval of a loan is assigned to a managerial role. Assignment of actual individuals to the teller and managerial roles is a personnel management function. Design of the role hierarchy relates to design of the organizational structure and is the function of a chief security officer under guidance of a chief information officer.

6 Task Based Access Control

The overriding concern of models we have discussed so far (DAC, MAC, HRU, TAM, ATAM, RBAC) has been the fine-grained protection of individual objects and subjects in the system. This approach has served as a reasonable basis for these model, but it lacks the concepts and expressiveness of an information-oriented model that captures the organizational and distributed aspects of information usage.

Increased automation of organizational functions and workflows, and the subsequent need to computerize information systems that often have distributed processing needs. Increased automation always carries it with the risk of weakened controls, especially when human judgment and paper-based checks and balances are taken out of the loop. The emergence of multi-system applications and information-related services that cross departmental and organizational boundaries, call for modeling constructs and integrity mechanisms beyond those existing for centralized systems.

Modern organizations encompass complex webs of activities (tasks) that often span departmental and organizational boundaries. Tasks are authorized and initiated by users in accordance with their roles, responsibilities, and duties (obligations) in the organization. One can view an organization as a system that is required to maintain a certain state (or standard) of integrity. Organizational procedures and internal controls then have to ensure that the tasks carried out in the organization preserve such a state of integrity. Now when we computerize organizational functions, we are faced with the problem of maintaining the required integrity in our computer-based information systems.

These considerations lead to the notion of task-based authorizations (TBA) and access control (TBAC) [TS94]. TBA is concerned with modeling and management of the authorizations of tasks (activities) in information systems. The central objective is preservation of integrity, but confidentiality applications are also possible. In a paper-based system, authorizations manifest as signatures on documents propagating through the organization. The analog to this in a computerized information system would be digital signatures on electronic documents. As such, we believe that task-based authorizations are central to the successful evolution of the concept of the "paper-less office".

A key element of TBA is the fact that authorization is transient and dependent on organizational circumstances. Consider the ability to issue a check. In RBAC we can associate this permission with a role, say, APM (accounts-payable-manager). This association is long-lived. A user who can exercise this role is capable of issuing many checks. In TBA the authority to issue a check is not directly associated with the role. We can say that role APM is a necessary requirement for issuing checks but it is not sufficient. In addition we require that a suitable authorization should have been obtained for the particular check in question. In the paper world this is achieved by obtaining one or more approval signatures on a voucher prior to issuance of the check. Techniques such as transaction control expressions [San88b] can be used to enforce this one-time permission.

7 Conclusion

In this paper I have given a high-level personal perspective on access control models and their future. I do believe this is a neglected frontier where much interesting and practically useful work remains to be done. I have identified some questions which merit particular attention.

There is considerably more literature than I have cited here. The papers I have cited are the ones that have influenced my own thinking most strongly.

Acknowledgment. My research on access control is partly supported by grants from the National Science Foundation and the National Security Agency.

References

AS92. P.E. Ammann and Ravi S. Sandhu. The extended schematic protection model. The Journal Of Computer Security, 1(3&4):335–384, 1992.

Bib77. K.J. Biba. Integrity considerations for secure computer systems. Technical Report TR-3153, The Mitre Corporation, Bedford, MA, April 1977.

BK85. W. Boebert and R. Kain. A practical alternative to hierarchical integrity policies. In *NBS-NCSC National Computer Security Conference*, pages 18–27, 1985.

BL75. D.E. Bell and L.J. LaPadula. Secure computer systems: Unified exposition and Multics interpretation. Technical Report ESD-TR-75-306, The Mitre Corporation, Bedford, MA, March 1975.

BN89. D.F.C. Brewer and M.J. Nash. The chinese wall security policy. In *Proceedings IEEE Computer Society Symposium on Security and Privacy*, pages 215–228, Oakland, CA, May 1989.

CW87. D.D. Clark and D.R. Wilson. A comparison of commercial and military computer security policies. In *Proceedings IEEE Computer Society Symposium on Security and Privacy*, pages 184–194, Oakland, CA, May 1987.

Den76. D.E. Denning. A lattice model of secure information flow. *Communications of the ACM*, 19(5):236–243, 1976.

Fag78. R. Fagin. On an authorization mechanism. *ACM Transactions on Database Systems*, 3(3):310–319, 1978.

FK92. David Ferraiolo and Richard Kuhn. Role-based access controls. In *15th NIST-NCSC National Computer Security Conference*, pages 554–563, Baltimore, MD, October 13-16 1992.

Fol92. Simon Foley. Aggregation and separation as non-interference properties. *The Journal Of Computer Security*, 1(2):159–188, 1992.

Gan96. Srinivas Ganta. *Expressive Power of Access Control Models Based on Propagation of Rights*. PhD Thesis, George Mason University, 1996.

GD72. G.S. Graham and P.J. Denning. Protection – principles and practice. In *AFIPS Spring Joint Computer Conference*, pages 40:417–429, 1972.

GSF91. Ehud Gudes, Haiyan Song, and Eduardo B. Fernandez. Evaluation of negative, predicate, and instance-based authorization in object-oriented databases. In S. Jajodia and C.E. Landwehr, editors, *Database Security IV: Status and Prospects*, pages 85–98. North-Holland, 1991.

GW76. P.P. Griffiths and B.W. Wade. An authorization mechanism for a relational database system. *ACM Transactions on Database Systems*, 1(3):242–255, 1976.

HRU76. M.H. Harrison, W.L. Ruzzo, and J.D. Ullman. Protection in operating systems. *Communications of the ACM*, 19(8):461–471, 1976.

KZB$^+$90. P.A. Karger, M.E. Zurko, D.W. Bonin, A.H. Mason, and C.E. Kahn. A vmm security kernel for the vax architecture. In *Proceedings IEEE Computer Society Symposium on Security and Privacy*, pages 2–19, Oakland, CA, May 1990.

Lam71. B.W. Lampson. Protection. In *5th Princeton Symposium on Information Science and Systems*, pages 437–443, 1971. Reprinted in *ACM Operating Systems Review* 8(1):18–24, 1974.

Lee88. T.M.P. Lee. Using mandatory integrity to enforce "commercial" security. In *Proceedings IEEE Computer Society Symposium on Security and Privacy*, pages 140–146, Oakland, CA, May 1988.

Lip82. S.B. Lipner. Non–discretionary controls for commercial applications. In *Proceedings IEEE Computer Society Symposium on Security and Privacy*, pages 2–10, Oakland, CA, May 1982.

LS77. R.J. Lipton and L. Snyder. A linear time algorithm for deciding subject security. *Journal of the ACM*, 24(3):455–464, 1977.

Lun88. Teresa Lunt. Access control policies: Some unanswered questions. In *IEEE Computer Security Foundations Workshop II*, pages 227–245, Franconia, NH, June 1988.

MS88. J.D. Moffett and M.S. Sloman. The source of authority for commercial access control. *IEEE Computer*, 21(2):59–69, 1988.

Pit87. P. Pittelli. The bell-lapadula computer security model represented as a special case of the harrison-ruzzo-ullman model. In *NBS-NCSC National Computer Security Conference*, 1987.

RBKW91. F. Rabitti, E. Bertino, W. Kim, and D. Woelk. A model of authorization for next-generation database systems. *ACM Transactions on Database Systems*, 16(1), 1991.

San88a. Ravi S. Sandhu. The schematic protection model: Its definition and analysis for acyclic attenuating schemes. *Journal of the ACM*, 35(2):404–432, April 1988.

San88b. Ravi S. Sandhu. Transaction control expressions for separation of duties. In *Fourth Annual Computer Security Application Conference*, pages 282–286, Orlando, FL, December 1988.

San90. Ravi S. Sandhu. Mandatory controls for database integrity. In D.L. Spooner and C.E. Landwehr, editors, *Database Security III: Status and Prospects*, pages 143–150. North-Holland, 1990.

San92. Ravi S. Sandhu. Expressive power of the schematic protection model. *The Journal Of Computer Security*, 1(1):59–98, 1992.

San93. Ravi S. Sandhu. Lattice-based access control models. *IEEE Computer*, 26(11):9–19, November 1993.

San96. Ravi Sandhu. Rationale for the RBAC96 family of access control models. In Ravi Sandhu, Ed Coyne, and Charles Youman, editors, *Proceedings of the 1st ACM Workshop on Role-Based Access Control*. ACM, 1996.

SCFY96. Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, February 1996.

SCY96. Ravi Sandhu, Ed Coyne, and Charles Youman, editors. *Proceedings of the 1st ACM Workshop on Role-Based Access Control*. ACM, 1996.

SG93. Ravi S. Sandhu and S. Ganta. On testing for absence of rights in access control models. In *IEEE Computer Security Foundations Workshop*, Franconia, NH, June 1993. 109–118.

SS94. Ravi Sandhu and Pierangela Samarati. Access control: Principles and practice. *IEEE Communications*, 32(9):40–48, 1994.

TS94. Roshan Thomas and Ravi S. Sandhu. Conceptual foundations for a model of task-based authorizations. In *IEEE Computer Security Foundations Workshop 7*, pages 66–79, Franconia, NH, June 1994.