(1) Workshop Summary

Ravi Sandhu

George Mason University and SETA Corporation
sandhu@isse.gmu.edu

1.0 Overview

The First ACM Workshop on Role-Based Access Control (RBAC) was held at the National Institute of Standards of Technology (NIST) in Gaithersburg, Maryland, on 30 November and 1 December 1995. The workshop was sponsored by the Association for Computing Machinery Special Interest Group in Security, Audit, and Control (ACM SIGSAC), the Washington DC Chapter ACM, and in cooperation with NIST. SETA Corporation of McLean, Virginia hosted the workshop, provided support for organizing the workshop, and supported preparation of the workshop proceedings. Copies of the workshop proceedings were provided to the workshop participants and copies of the proceedings are available from ACM.

1.1 Purpose

The Call for Papers describes the organizers' motivation in creating this series of workshops. Relevant portions are quoted below.

> "In a nutshell, the essence of Role-Based Access Control (RBAC) is that rights and permissions are assigned to roles rather than to individual users. Users acquire these rights and permissions by virtue of being assigned membership in appropriate roles. This simple idea greatly eases the administration of authorizations. The basic concepts behind RBAC have been around since the advent of multi-user computing and information systems in the late 60's and early 70's. There has been a recent resurgence of interest in RBAC. This is in large part due to the user community's expression of interest in RBAC, and disenchantment with traditional mandatory and discretionary access controls.
>
> The ACM Workshop on Role-Based Access Control has been created to bring together users, vendors, and researchers who are interested in fostering and promoting RBAC. The workshop's objectives are to provide a forum for rapid dissemination of new ideas and developments in RBAC, and to cultivate convergence toward a standard framework for RBAC and related access control issues.
>
> This is the first in a series of workshops to be held on a fairly frequent basis. Ideally, we would like these workshops to develop a standard reference model for RBAC. We recognize this cannot be accomplished in a single meeting, but we are seeking progress toward this end at a rapid pace. In the first workshop we seek input from users regarding their access control

Copyright 1996 Association for Computing Machinery. Permission to make digital/hard copy of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage; the copyright notice, the title of the publication, and its date appear; and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

ACM RBAC Workshop, MD, USA
© 1996 ACM 0-89791-759-6/95/0011 $3.50

needs, from vendors regarding plans for products, and from researchers concerning conceptual frameworks from which to approach these issues.

Although there is much agreement on the basic concepts and value of RBAC, there remains a number of issues on which different researchers and vendors are proposing different approaches. The user community is also often doing access control and management in a way that is very similar to RBAC without actually applying that name. At the same time, the scope of RBAC ranges from very simple and straightforward at one end, to very sophisticated and complicated at the other. Much remains to be done to develop a scientific and engineering discipline in this arena. The ACM Workshops on RBAC are primarily intended to support this goal."

The workshop attracted attendees from the U.S., Canada, and various West European countries. Many of the attendees met for the first time at the workshop. We had representation from users, vendors, academia, research laboratories, and standards organizations. Because the need for RBAC is pervasive in computer systems, it was particularly gratifying that we had representation from the database, network, distributed systems, and operating systems communities. Concepts of RBAC have evolved more or less independently in these communities and it is important to have workshops such as this to foster cross-fertilization of ideas.

1.2 Results

The workshop was successful in its modest goal of taking a first step toward a consensus reference model for RBAC. There was substantial agreement among attendees regarding the general outlines of RBAC and its various components. There was considerable discussion about the details, including the reconciliation of different terminology used by different groups working in the field. There was also extensive discussion concerning the priorities and importance of various aspects of RBAC. Nonetheless, in general, there was substantial agreement.

1.3 Issues

From the workshop discussion two issues have emerged as significant ones for further work in similar workshops and study groups. Firstly, any scientific discipline needs an internally consistent and widely used terminology and vocabulary. In the early stages, as the discipline emerges, different people use the same words to mean different things, and sometimes major concepts have not yet been articulated and named. As the discipline matures, a de facto standard terminology emerges. Efforts to impose a standard terminology by committee are rarely successful. These efforts can be premature if major concepts are still emerging in the field. As the field matures, development of a de facto standard can be encouraged and helped by workshops where disagreements about terminology and standards can be articulated and discussed. Discussions at the first workshop indicate that RBAC is at the right point of maturity to merit such efforts.

Secondly, RBAC is an expansive concept. In order to scope the problem the RBAC community needs to clearly articulate what is excluded as being outside the scope of RBAC proper. This relates to the terminology issue because we need to decide what should legitimately be called RBAC. But there is a bigger issue than terminology here. A sound technical discipline draws boundaries around its major concepts for technical reasons. The RBAC community needs to clearly identify where it is useful to draw these boundaries for technical reasons (and not merely for convenience of terminology).

2.0 Session Summaries

The rest of this report describes the sessions in chronological order. This summary is based on my personal impressions appropriately revised after feedback from other attendees on an earlier draft of the report.

2.1 What Is RBAC?

After introductory remarks from representatives of the various sponsoring organizations, the workshop's first session on "What is RBAC?" followed. I presented the first talk. My objective was to present a framework of models for RBAC and the rationale which led to this framework. (This family of RBAC models was recently published in IEEE *Computer*, Feb. 1996 [SAND96a].) The central notion of this framework is that users and permissions are brought together indirectly by roles. A user acquires a permission by virtue of being assigned to a role that has been assigned that permission. The framework begins with a base model to which role hierarchies and constraints are added in the extended models. This incremental approach is motivated by the use of the term RBAC in the literature to include simple as well as sophisticated concepts. As is appropriate in a workshop setting there were many questions and comments during the presentation. These included questions about several design decisions made by our team in constructing this framework of models. Some of these are discussed in our paper in Part II, *Individual Presentations*, of the proceedings.

My talk was followed by a presentation from Virgil Gligor of the University of Maryland concerning a RBAC model that he has developed. Virgil began by noting that if RBAC proponents believe the notion of a role is fundamentally different from that of a group of users we should be able to articulate the essential difference. Since access review is among the basic aspects of access control, he proposed we distinguish between roles and groups by requiring that roles allow us to perform per-subject review of rights. Moreover, per-subject review is needed for separation-of-duty and conflict-of-interest properties generally considered important in RBAC. Virgil observed that per-subject review is difficult in large distributed systems using traditional access control lists (ACLs) because, given a user or group identity, one does not know where to start the review. Potentially, one would have to search the entire set of objects' ACLs to determine the extent of a user's permissions. In discussion, some attendees argued that such searches to determine per-subject review are conducted by several practical systems (such as Novell's NetWare, IBM's RACF, and IBM's OS/400 list). Virgil maintained that none of these systems provide this support with "group" mechanisms alone, and their access control models do not scale up to large distributed systems. Virgil argued that access control policies, such as RBAC, that need both per-object and per-subject review could be derived from the storage of the access matrix, redundantly, on both rows and columns.

The second session continued the theme of "What is RBAC?" Emile Lupu of the Imperial College in London, UK, presented a talk on a policy-based framework for RBAC. His co-author and thesis advisor Morris Sloman has been conducting research in this area from the perspective of distributed systems management. Emile's talk introduced the concepts of obligations and duties going beyond the access control notion of permissions. The authors were careful to point out that obligations and duties are outside the scope of access control. In their view, roles are a larger concept beyond access control and it is important to distinguish and recognize the scope of access control roles in contrast to roles in general.

Chris Sundt presented a talk on ICL's experiences with deploying RBAC in actual products and their experience with real users. Chris noted that RBAC makes it easy to split the administration of the role-user relationship from that of the role-permission relationship. He said users find this very useful. He also introduced the notion of an affiliation whereby a role can be further qualified. Thus a role of branch manager could be qualified by an affiliation to a particular branch thereby conferring branch manager permissions only within that branch. This was another facility that ICL found to be popular with users. Chris also emphasized that practical implementations of RBAC need to cope with distributed multi-vendor environments.

The second session on "What is RBAC?" concluded with three shorter talks. Sylvia Osborn of the University of Western Ontario, Canada, presented an overview of RBAC research by her group. Among other things, this work has developed algorithms for recognizing redundant assignment of permissions to roles, as well as algorithms for deletion and addition of roles in a hierarchy and similar operations. Luigi Guiri of Fondazione Ugo Bordoni in Rome, Italy presented an extension of a model of Baldwin's based on a role as a named protection domain. Luigi argued that a role should be viewed a set of named protection domains (and roles). He presented a role algebra for constructing such sets based on the notion of and-roles (that are simultaneously activated) and or-roles (that are mutually exclusive). For logistical reasons, the third short talk of this session was actually presented later but logically belongs in this session where it had been originally scheduled. The talk was by Fang Chen who is a student of mine at George Mason University. He described some preliminary work in designing a language to express constraints on components of RBAC. Mutually exclusive roles, where one user cannot be assigned both roles, are perhaps the most common example of constraints in RBAC, but there are also many other useful constraints.

2.2 What Are User Needs?

After lunch, the third session on "What are user needs?" had two presentations. The first talk by Yahya Al-Salqan of West Virginia University described an application of RBAC in the health care domain.

In this project RBAC, as provided by the Oracle database management system (DBMS), was being considered as a means to enforce patient privacy and confidentiality requirements. Possible extensions to an inter-organizational environment were also being studied. The second talk by Trent Jaeger looked at the possibility of using RBAC in collaborative systems. RBAC facilitates the use of least privilege because different code can be executed with different roles by the same user. This makes it safer to use agents supplied by other users, since these agents can execute on a user's workstation but with restricted roles. Further, different roles can be assigned to different agents depending upon the trust and requirements of these agents.

3.0 Prioritizing RBAC Features

The fourth session consisted of a group exercise developed by Charles Youman of SETA Corporation. The exercise was conducted in two break-out groups. The objective of the exercise was to rank order different aspects of RBAC with respect to their priority or importance. The results indicated that there were some differences in priorities assigned by individual attendees, but by and large there was considerable agreement on what the more important aspects were. This is an encouraging finding which suggests that there is substantial consensus within the RBAC community upon which a widely accepted reference model can be developed. This exercise is further described by Edward Coyne in Part I of this proceedings. This concluded the first day.

3.1 Available and Emerging Technologies

The second day began with a number of short talks in a session called "Available and Emerging Technologies." LouAnna Notargiacomo of Oracle Corporation described various aspects of RBAC in Oracle 7 and Trusted Oracle 7. Oracle has pioneered the use of roles in relational DBMSs and these features are being incorporated into SQL standards. Jeremy Epstein of Cordant Inc. presented a talk on "NetWare 4 as an Example of Role-Based Access Control." Jeremy's talk focused on identifying how NetWare can implement the concepts of the RBAC models introduced earlier in my talk. NetWare has a built-in concept of Organizational Roles but attaches very little semantics to it beyond that associated with any other NetWare Directory Service object. Jeremy's finding was that some aspects of RBAC do translate quite easily on to NetWare roles but others would be difficult to support. These two talks demonstrate that there is available technology on popular platforms that can be used today to support RBAC (at least to some extent).

The remaining papers in this session addressed emerging technologies. T. C. Ting of the National Science Foundation (on leave from University of Connecticut) described ongoing RBAC research at University of Connecticut concerned with implementing RBAC in object-oriented systems. Their project seeks to develop RBAC notions consistent with object-oriented concepts such as encapsulation, information hiding, and inheritance. They have used a health-care case study in their project. John Barkley of NIST gave a talk on "Implementing Role-Based Access Control using Object Technology." In this project, John used a concept of layered objects to facilitate flexible administration while minimizing the impact of role changes on applications. A prototype demonstration of these concepts is available at NIST's RBAC home page (http://waltz.nist.gov/rbac). Roshan Thomas of Odyssey Research Associates (ORA) presented a talk on "RBAC and Distributed Object-Based Enterprise Computing." Roshan described ongoing work at Odyssey on next-generation security models for workflow processing that involve RBAC as a component. Roshan also mentioned ORA's efforts at including some of these concepts in the Object Management Groups's (OMG's) Common Object Request Broker Architecture (CORBA) initiative.

3.2 Role Engineering and RBAC Transition

The next session consisted of open discussion on two important issues in making RBAC practical. Edward Coyne of SETA Corporation facilitated discussion on role engineering. The definition of roles, assignment of permissions to roles, and definition of other components of an RBAC model, is essentially a requirements engineering process. Attendees agreed that this is an important and complex topic which should be approached with an engineering methodology. Charles Youman of SETA Corporation facilitated discussion on RBAC transition. The question is how to get from here (no RBAC) to there (RBAC). Attendees agreed that RBAC must co-exist with other access control mechanisms. Moreover, RBAC would need to be deployed incrementally in an organization rather than totally replacing legacy access control. Further details of these two discussion sessions are given by Edward Coyne and Charles Youman in Part I of the proceedings.

3.3 Consensus Reached and Remaining Issues

After lunch, I moderated a discussion session on "Consensus Reached and Remaining Issues." A number of open issues were enumerated and are listed elsewhere in Part I of this proceedings. The discussion focussed entirely on the most important which was to define the concept of a role and distinguish it from the familiar concept of a group in access control. There was agreement that the notion of a directory and a file are fairly standard notions in operating systems (OSs) even though the details do vary from one OS to another. Similarly the concept of a group as a collection of users (and possibly other groups) is well-known in access control systems. The discussion attempted to develop a notion of role along the same line. The concept of a role in access control is currently being used in at least two different ways. Some use role to mean a named collection of permissions (and possibly other roles). Others use role to mean a named collection of permissions and a named collection of users (and possibly other roles). The discussion could not reconcile whether one of these was the "correct" use of the term role. Further details and thoughts on this discussion are given elsewhere in Part I of this proceedings. It might seem that a discipline that cannot even agree on the meaning of its key term (role in this case) is in deep trouble. I am, however, much more optimistic and feel there is substantial consensus that can overcome these minor disagreements in terminology. Perhaps the term role can be used in both ways, but we just need to make clear how it is being used in a given context. Or perhaps we need to agree as a community to use two different terms for these two different concepts of a role. It would be inappropriate to discard the entire discipline of RBAC just due to some minor disagreement in basic terminology.

4.0 Future Plans

The final session of the workshop was devoted to discussion of future plans. There was general agreement that we should continue this series of workshops. As the series evolves, the workshop organizers are hopeful that the workshop deliberations will have a positive impact on security practice. Some thoughts on the format and goals of future workshops are given in Part I of this proceedings by David Ferraiolo and Richard Kuhn of NIST.

In conclusion, I reiterate my earlier statement that the first RBAC workshop was successful in its modest goal of taking the important first steps toward a consensus reference model for RBAC. Overall success of this series will depend upon what happens subsequently. Anyone interested in being involved should contact me (E-mail: sandhu@isse.gmu.edu).

This page intentionally left blank.